

Estimating the Capacities of Function-as-a-Service Functions

*This is the preprint version of the accepted paper at CloudAM'21 workshop (UCC)

Anshul Jindal
Technical University of Munich
Garching (near Munich), Bavaria, Germany
anshul.jindal@tum.de

Shajulin Benedict
Indian Institute of Information Technology
Kottayam, Kerala, India
shajulin@iiitkottayam.ac.in

Mohak Chadha
Technical University of Munich
Garching (near Munich), Bavaria, Germany
mohak.chadha@tum.de

Michael Gerndt
Technical University of Munich
Garching (near Munich), Bavaria, Germany
gerndt@in.tum.de

ABSTRACT

Serverless computing is a cloud computing paradigm that allows developers to focus exclusively on business logic as cloud service providers manage resource management tasks. Serverless applications follow this model, where the application is decomposed into a set of fine-grained Function-as-a-Service (FaaS) functions. However, the obscurities of the underlying system infrastructure and dependencies between FaaS functions within the application pose a challenge for estimating the performance of FaaS functions. To characterize the performance of a FaaS function that is relevant for the user, we define Function Capacity (FC) as the maximal number of concurrent invocations the function can serve in a time without violating the Service-Level Objective (SLO).

The paper addresses the challenge of quantifying the FC individually for each FaaS function within a serverless application. This challenge is addressed by sandboxing a FaaS function and building its performance model. To this end, we develop *FnCapacitor* – an end-to-end automated Function Capacity estimation tool. We demonstrate the functioning of our tool on Google Cloud Functions (GCF) and AWS Lambda. *FnCapacitor* estimates the FCs on different deployment configurations (allocated memory & maximum function instances) by conducting time-framed load tests and building various models using statistical: linear, ridge, and polynomial regression, and Deep Neural Network (DNN) methods on the acquired performance data. Our evaluation of different FaaS functions shows relatively accurate predictions with an accuracy greater than 75% using DNN for both cloud providers.

CCS CONCEPTS

• **Computer systems organization** → *Cloud computing*.

KEYWORDS

Serverless Computing, Function-as-a-Service, Function Capacity

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference '17, July 2017, Washington, DC, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

ACM Reference Format:

Anshul Jindal, Mohak Chadha, Shajulin Benedict, and Michael Gerndt. 2021. Estimating the Capacities of Function-as-a-Service Functions: *This is the preprint version of the accepted paper at CloudAM'21 workshop (UCC). In *Proceedings of ACM Conference (Conference '17)*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

With the advent of Amazon Web Services (AWS) Lambda in 2014, serverless computing has gained popularity and more adoption in different application domains such as machine learning [6, 8], linear algebra computation [10, 29], and map/reduce-style jobs [22]. Function-as-a-Service (FaaS), a key enabler of serverless computing, allows a traditional application to be decomposed into fine-grained functions that are executed in response to event triggers or HTTP requests on a FaaS platform [7]. The FaaS platform is responsible for deploying and providing resources to the FaaS functions.

Since serverless computing environments abstract the underlying system infrastructure configurations away from the users, most of the public cloud providers in their FaaS offerings allow users to configure only a small set of configuration parameters: memory allocation and the maximum number of function instances, also called as concurrency [1–3]. Moreover, cloud providers speedup function execution when a higher memory is configured [25]. Furthermore, the heterogeneity in the underlying nodes can lead to variations in the execution time of the FaaS functions [9, 18]. Therefore, estimating the maximum number of requests the deployed functions can handle such that the response times adhere to the SLOs constraints poses a set of challenges. Towards this, we define the maximal number of concurrent invocations that a FaaS function can serve within a time interval without violating the SLOs when deployed with particular memory configuration and fixed maximum function instances as the *Function Capacity (FC)*. In this paper, we consider the 95th percentile execution time of a FaaS function as the SLO.

To highlight the effects of various parameters configuration on the performance and the Function Capacities of FaaS functions, we deployed a compute-intensive (calculates prime numbers till 10000000) serverless function written in Python on AWS Lambda and Google Cloud Function (GCF) [3]. We fixed the 95th percentile execution time of the function to 20 seconds. Figure 1a shows how the Function Capacity (i.e., the maximum number of requests per second the function can handle) first increases with varying memory




Figure 1: AWS lambda and Google Cloud Function (GCF) FCs variation with memory configurations, and concurrency.

sizes upto a certain point (2048MB). After that, it becomes constant when the function concurrency is fixed to 100 for both the cloud providers. The same Figure 1a also shows the variation in FC with cloud providers. The variation in the system resources causes differences in performance between the identical function deployments for the same FaaS platforms. Figure 1b shows the execution time of corresponding runs, and one can see that it decreases with the increase in memory, and after a point (2048MB), it also becomes constant. Lastly, Figure 1c shows the linear increase in FC with the increase in function concurrency, keeping the memory fixed.

The examples above highlight some factors that can affect the performance and the FCs. However, there are many other factors such as cold starts, I/O and network conditions, type of container runtimes, and co-location with other functions affecting the performance and FCs which the users are not aware of [12, 30]. Additionally, the dependencies between the functions within a serverless application can also affect the FCs. To this end, we develop **FnCapacitor**, a tool that can estimate the FCs of the functions adhering to the given SLOs, the specified memory configurations, and function concurrency. Our key contributions are as follows:

- We develop and present a novel python-based tool called *FnCapacitor* for automatic estimation of Function Capacities of FaaS functions within a serverless application (§3).
- As part of *FnCapacitor*, we present a functions sandboxing method which can be used to sandbox individual functions from the overall serverless application (§3.1).
- We showcase the effect of different deployment configurations : memory allocation (§5.1) and function concurrency (§5.3), on the Function Capacity for different FaaS functions within a sample application. To the best of our knowledge, none of the previous works involving FaaS [5, 11, 20, 24, 28, 30] considered function concurrency parameter.
- Although our approach is generic and *FnCapacitor* can be easily extended to support other commercial and open-source FaaS platforms, we demonstrate the functioning of *FnCapacitor* with Google Cloud Functions (GCF) and AWS Lambda (§5) on a sample serverless application consisting of 8 functions. We further present the performance results of the formulated FC estimation models on each of them (§5.4).




Figure 2: Typical FaaS function invocation procedure.

- We open source the collected data and developed tool for further research¹.

Paper Organization Section 2 introduces the basic FaaS function invocation procedure and Function Capacity concept. In Section 3 the developed tool *FnCapacitor* is described. Section 4 describes the experimental setup and the hyper-parameters of the different algorithms. In Section 5, our evaluation results are presented. Section 6 describes some of the prior works in this domain. Finally, Section 7 concludes the paper.

2 TAXONOMY

2.1 FaaS function invocation procedure

Function-as-a-Service based *function* is a piece of code containing a *handler method* responsible for processing the *events* that are passed to the function when invoked and these are executed within a *FaaS platform*. FaaS based functions can be invoked by a user’s HTTP request or another type of event created within the FaaS platform or the cloud infrastructure. These include changes to data in a database, files added to a storage system, or a new virtual machine instance is created. The FaaS platform is responsible for providing resources for function invocations and performs automatic scaling. This is done by creating an *execution environment* which provides a secure and isolated runtime environment for the function. The functions can be written using various languages, and a language-specific environment called *runtime* is created in the execution environment. The runtime relays invocation events, context information, and responses between the FaaS platform and the function.

¹https://github.com/ansjin/faas_capacitor

The first time the function is invoked, the FaaS platform creates an *instance* of the function (execution environment) and runs its *handler method* in it to process the event. When the handler exits or returns a response, it stays active and becomes available to handle other events. If the function is invoked again while the first event is being processed, the FaaS platform creates another *instance*, and the two events are processed *concurrently*. As more events come in, the FaaS platform routes them to available instances and creates new instances as needed. When the number of requests decreases, the FaaS platform stops unused instances to have free scaling capacity for other functions. FaaS platforms usually have an upper limit on how many maximum concurrent instances called *function concurrency* can be created, such as 1000 for AWS lambda and 3000 for GCF. Figure 2 summarizes the overview of the typical FaaS function request procedure.

2.2 Function Capacity (FC)

We define the *Function Capacity (FC)* as the maximal number of concurrent invocations that a FaaS function when deployed on a FaaS platform with a certain memory configuration and fixed maximum function instances, can serve within a time interval without violating the SLOs when deployed. Ideally, if an instance i serves n_i^t number of invocations within a time interval t and C is the function concurrency for the FaaS platform, then the FCf , where f is any function, can be calculated by the equation (1).

$$FCf = n_i^t \times C \quad (1)$$

However, in practice, there are many other factors such as cold starts, I/O and network conditions, type of container runtimes, and co-location with other functions affecting the performance and FCs of FaaS functions [30]. Therefore, we follow a modeling approach to estimate the capacities of the FaaS functions. In this paper, we consider the 95th percentile execution time of a FaaS function as the SLO.

3 FNCAPACITOR

To estimate the capacities of the FaaS functions within a serverless application, we have developed *FnCapacitor*, a python-based automated end-to-end function capacity estimation tool. Given the SLO requirements, *FnCapacitor* is responsible for estimating the FCs of FaaS functions at different deployment configurations (memory allocation and function concurrency). Figure 3 provides an overview of the high-level architecture of *FnCapacitor* and the interaction between its components.

FnCapacitor takes a YAML file as input that specifies the initial FaaS platform configuration parameters (minimum memory allocation and functions timeout) for the function deployment, serverless application to be deployed, and configuration parameters for the load generator and the modeling (step ①). Since a serverless application consists of multiple functions and the performance of one function could affect the others depending upon it, therefore in the next step, the individual functions from the given application are segregated (step ②). These sandboxed functions are then deployed on the FaaS platform with some initial configurations (step ③). After the deployment, *FnCapacitor* repeatedly changes the functions configurations (steps ④ - ⑤) and generates a user workload to the function's API




Figure 3: High-level architecture of the *FnCapacitor*.

endpoint (step ⑥) for collecting various monitoring metrics data (§4.2). The collected metrics data is stored in InfluxDB, which is used for creating the function performance models (step ⑦). The created function models are then used for estimating the Function Capacities for different deployment configurations (step ⑧).

We now describe the sub-components of *FnCapacitor*:

3.1 Functions Sandboxing

Usually, a serverless application consists of multiple functions, and the performance of one function could affect the others depending upon it. Therefore to measure the pure performance of the functions i.e., where their performance is not affected by others, we build this component for sandboxing individual functions through a mockup of their neighbors. It isolates each function and substitutes its direct neighbors with dummy functions accepting the requests and sending the responses in the same format, but without any additional processing allowing to measure the pure performance of only that function and build models using this data.

Firstly, this component for each function within the serverless application replaces the calls to other functions with calls to a *proxy-function*. This *proxy-function* serves as an intermediary between the sandboxed function and other functions and takes the originally called function names from the sandboxed function and the input payload to them as the input. This allows every invocation to other functions to go through this *proxy-function* and this dummy *proxy-function* then will invoke the following functions based on the input received. At the same time, copies of these requests and responses are stored in the *FnCapacitor*'s MongoDB database for creating function mockups. It is to be noted that, Backend-as-a-Service (BaaS) services such as database, storage, queues, etc., are out of the scope of this work for sandboxing as it is assumed that these BaaS services provides high scalability and serve the user requests within the defined SLOs. Following this, each function receives its own sandboxed deployment where mockup functions replace the direct neighbors. These mockup functions will respond with the response stored in MongoDB. As a result, the time taken by the dependent functions to respond becomes negligible and therefore allows to build a pure performance model of the functions.

3.2 Function Deployer, Destroyer and Updater

As the name suggests, it is responsible for deploying, destroying, and updating the functions to be modeled on the GCF and AWS Lambda using the serverless framework [27]. For deployment and updating, different configuration parameters: memory configurations and function concurrency are taken into account. The deployed functions are scaled automatically by adding multiple function instances by the FaaS platform. This component automatically deletes the functions when the test is finished to free up the resources.

3.3 Load Generator

The Load Generator is implemented using a load testing tool - *k6*. It uses a script for running the tests where the function endpoint and request parameters are specified. As part of each test, the number of requests per second (RPS) generated by *k6* is varied and depends on the number of Virtual Users (VUs) and the time taken by each request to complete. VUs are the entities in *k6* that execute the test and make HTTP(s) requests. The load generation metrics from *k6* and FaaS platform monitoring metrics data from the cloud providers are exported and stored inside an InfluxDB instance.

3.4 Performance Modeler & Capacity Estimation

This is the main component of *FnCapacitor* and is responsible for analyzing the correlation between the different monitoring metrics (§4.2) and the deployment configurations. It uses the collected data stored in InfluxDB to create models of the functions and estimate their Function Capacities. Modeling approaches used in this work are categorized under two categories:

- **Statistical Approaches:** We consider linear, polynomial, ridge, and random forest regression for modeling the relationship.
- **Deep Neural Networks (DNNs):** DNNs are designed to solve complex problems by building relationships among multiple dependent and independent variables. As a result, we use them for formulating the models.

The collected data from InfluxDB is pre-processed by removing outliers and dividing the data into training and test set. Following this, different models, i.e., statistical and DNNs, are trained on the training data set. Due to the training data being sparse, k-fold cross-validation (in our case k=6) is followed for training the model [31].

Lastly, Function Capacities are estimated on the new or test dataset using the trained models. The prediction accuracy is computed using the R^2 score value from the actual and predicted Function Capacities for the test dataset.

4 EXPERIMENTAL CONFIGURATION

In this work, we have fixed the total duration of a test to 30 minutes for the deployed serverless application. A test consists of the memory allocation configurations: <256MiB, 512MiB, 1GiB, 2GiB and 4GiB> and function concurrency: <10, 20, 30, 40 and 50>. AWS functions are deployed in the europe-central1 region, and GCF functions are deployed in the eu-west3 region. The number of VUs during the load generation are varied from 5 to 500 depending on the number of requests the functions can serve. As a result, for each function, 5 (total memory configurations) \times 5 (function concurrency configurations) = 25 tests were conducted.




Figure 4: High level workflow of the evaluated application.

We partition the collected data into training and test set (33% of the total data). We used a part of the training data set as a validation set for selecting the hyperparameters of the different models. We select the hyperparameters through an exhaustive grid search. For the DNN model, we use a 12-layer fully connected neural network architecture, with each layer having 64 units. We use the Rectified Linear Unit (ReLU) as activation function [23]. We use mean absolute error as the loss function for training the DNN model with Adam as the optimizer.

4.1 Application used for evaluation

To investigate the performance of each deployment configuration, we used a subset of the benchmarks provided with the FaaSProfiler [28]. We created an application shown in Figure 4 for our use case. The functions used as part of this application are summarized in Table 1.

The application flow starts with the *nodeinfo* function, which exposes an HTTP endpoint and provides the user with basic information about the system such as hostname, underlying architecture, number of CPUs, etc. This then invokes the compute-intensive *primes-python* function to judge the built models' abilities on compute-intensive applications, and it further invokes *linpack* and *dd* asynchronously and waits for their response to come back. *dd* invokes *gzip-compression*. It further invokes *lr-prediction* in a sequence. *lr-prediction* queries the model and data from the google cloud storage (created in the google compute platform in the europe-west3 region, AWS lambda functions also use this storage bucket) and then performs prediction. Once the response is available to *primes-python* function from both the invocations, it sends back the response to *nodeinfo* function, which in turn invokes *nodejs-endpoint* and *nodejs-endpoint* further invokes *sentiment-analysis* function. We use these eight heterogeneous functions to test the accuracy of the modeling on different types of functions.

4.2 Monitoring Metrics

We extracted following monitoring metrics from the GCF and AWS lambda² with the data sampling rate as one minute:

- *concurrent_instances*: Number of concurrent function instances.
- *invocations*: Number of times the function is executed.
- *execution_duration*: Amount of time function code spends in processing an event.

²<https://docs.aws.amazon.com/lambda/latest/dg/monitoring-metrics.html>

Table 1: Functions comprising the evaluated application.

Functions	Description	Runtime
nodeinfo	Gives basic characteristics of node like CPU count, architecture, uptime.	Node.js 14
primes-python	Calculates prime numbers till 100000.	Python3
linpack	It solves a dense linear system of equations in double precision and returns the results in GFlops. Problem size (number of equations) is fixed to 100.	Python3
dd	It is based on Unix dd command-line utility for converting and copying files. 128bytes as a block size and 5 times conversion is used as parameters.	Python3
gzip-compression	It first creates a file filled with random numbers of size 1MB and then compresses it using gzip compression scheme.	Python3
lr-prediction	It first downloads a linear regression model trained on user reviews data from the storage bucket along with the test data and performs prediction on it.	Python3
nodejs-endpoint	It is a simple REST endpoint which returns the current time along with the time zone.	Node.js 14
sentiment-analysis	Analyzes the sentiment of a provided string using the Python TextBlob library	Python3




Figure 5: execution_durations of the sandboxed functions when executed with a load of 50 RPS with five different memory configurations and no limit on the function_concurrency.

- *memory_usage*: Function’s maximum memory usage.
- *allocated_memory*: Memory allocated to the function.
- *function_concurrency*: Maximum number of concurrent instances allowed for processing events.

5 EXPERIMENTAL RESULTS

In this section, we first describe the impact of heterogeneity in the memory allocations on function’s *execution_duration* and *concurrent_instances*, and then the effect of *function_concurrency* on the Function Capacities for both the cloud providers and on different functions. Following this, we present results of FCs estimation using the different modeling approaches.

5.1 Memory effect on function execution duration

Figure 5 shows the *execution_durations* of the load testing of 50 invocations per second on the application when the functions are sandboxed and deployed with five different memory configurations on GCF and AWS lambda platforms. We observe the following:

Decrease in execution_duration with the increase in memory and becoming constant: From Figure 5, we can see that for most of the functions and across two FaaS platforms, there is a general trend of decrease in *execution_duration* with the increase in memory, and after a certain point (2048MB memory configuration), either it is becoming constant (for all functions running on AWS, and *lr-regression*, *sentiment*, and *linpack* on GCF) or increasing (for all




Figure 6: concurrent_instances of the sandboxed functions for handling the load of 50 RPS with five different memory configurations and no limit on the function_concurrency.

other functions on GCF). This can be attributed to an increase (2x) in the number of allocated clock cycles for a memory configuration of 4096MB as compared to 2048MB [15].

In general, AWS lambda has a lower execution duration for most of the functions at all memory configurations as compared to GCF: We can observe from the Figure 5 that, for most of the functions except the three compute-intensive functions at a lower memory configuration (256MB and 512MB), AWS lambda process function events faster than the GCF. For example, for *dd* microbenchmark at 256MB configuration, AWS lambda takes 5.2x times less than the GCF at the same memory allocation, and even it can process faster than GCF allocated with 4GB. *nodeinfo*, *sentiment*, *linpack*, and *web-endpoint* took almost the same amount across different memory configurations and FaaS platforms.

5.2 Memory effect on function’s concurrent instances

Figure 6 shows the *concurrent_instances* per function in the serverless application when it is load tested with a load of 50 invocations per second on GCF and AWS lambda platform for five different memory configurations. We observe the following:

AWS lambda creates more concurrent_instances as compared to GCF: From Figure 6, we can see that, for most of the functions and across different memory configurations, AWS creates a higher




Figure 7: Function Capacity of the functions when deployed on the two FaaS platforms for different `function_concurrency` with memory configuration fixed to 256MB.

number of `concurrent_instances` as compared to GCF for handling the same amount of load.

Decrease in number of `concurrent_instances` with the increase in memory configuration: As the memory is increased for each function, the number of functions `concurrent_instances` for both the FaaS platforms either remain constant or has decreased. This trend can be attributed to the fact that a higher resource instance can serve the requests faster and hence can process more requests per unit time. Therefore, fewer instances are required to handle the same amount of load when allocated with lower memory configurations.

Slow-processing functions are scaled to higher number of `concurrent_instances` to match up with the fast-processing functions : From Figure 5, we can observe that, `lr-regression` and `gzip` have the highest function `execution_time` as compared to the other functions and when observing the number of `concurrent_instances` for those two functions in Figure 6, we can see that they are the highest. This concludes that the compute-intensive (slow-processing) functions require higher scaling to match up the other fast-processing functions to handle the same amount of load. Such visualization can also be used to understand the bottleneck function in the serverless application; for example, in our case, it is `lr-regression`.

5.3 Effect of function concurrency on the FC

Figure 7 shows the actual capacity measurements for the two FaaS platforms for different `function_concurrency` configurations with memory configuration fixed to 256MB for all the functions. The capacities depicted are the average of the five runs for both FaaS platforms. In general, it can be inferred that for most of the functions, FCs vary linearly with the `function_concurrency` for both the FaaS platforms. Also, a single instance of AWS lambda can process a higher number of requests than the single instance on GCF. However, for the two compute-intensive functions namely:

`gzip-compression`, `ml-lr-prediction` we see a similar FCs for both the platforms at different `function_concurrency`. In case of GCF, for simple web-based functions : `sentiment-analysis`, `nodeinfo`, and `nodejs-endpoint` the linear increasing slope is not constant. From Figure 7, one can see that, for the three FaaS functions the linear slope changed after `function_concurrency` of 6. This means that, after the `function_concurrency` of 6, each instance can process more number of requests as compared to the instance used when the `function_concurrency` is less than 6. In general, the trend is linear for all other functions and both FaaS platforms. However, they are not exactly following the ideal lines. Therefore, one needs modeling approaches for the estimation of FCs for both the FaaS platforms.

5.4 Function Capacity estimation

On analyzing the impact of varying memory configurations on the performance of the different FaaS functions, we use the metrics `<concurrent_instances, execution_duration, allocated_memory, memory_usage, function_concurrency>` obtained from the collected load test data as input parameters for the different models (§3.4). For a given set of input parameters, all models predict `<function_invocations>` which is equivalent to the FC.

Table 2 shows the comparison of the accuracy results for estimated FCs for the different modeling approaches (§3.4) with the best ones underlined for both the FaaS platforms. For determining the accuracy of the formulated models, we use the R^2 score [4].

In general, it was found that the accuracy measurements for Function Capacity estimation for AWS lambda are higher than the GCF for most of the FaaS functions since AWS lambda exhibits more linear behavior as compared to GCF (Figure 7). **Linear Regression (LR)** leads to best results when the parameters are linearly correlated to the FC. For most FaaS functions, the parameters are linearly correlated to FC, leading to an accuracy value greater than 80%.

Table 2: Comparison of accuracy results (R^2 score) for estimated FCs for the different modeling approaches.

Function	LR		PLR		RR		RFR		DNN	
	GCF	AWS	GCF	AWS	GCF	AWS	GCF	AWS	GCF	AWS
dd	81.9	98.1	88.27	97.7	82.2	98.0	88.5	98.1	<u>91.1</u>	<u>98.2</u>
gzip	83.5	91.4	89.8	94.8	83.6	91.4	93.0	94.6	<u>93.6</u>	<u>94.9</u>
primes	75.8	95.4	78.6	95.1	76.4	95.6	83.1	<u>96.7</u>	<u>85.0</u>	96.5
linpack	86.3	58.7	87.7	75.9	86.4	76.3	88.5	87.2	<u>88.8</u>	<u>92.4</u>
sentiment	65.6	33.6	72.9	92.9	52.2	24.9	<u>76.0</u>	<u>97.4</u>	74.4	96.2
lr-pred.	90.9	99.4	93.0	99.5	90.7	99.4	<u>95.4</u>	98.7	94.7	<u>99.5</u>
nodeinfo	80.6	87.5	88.4	87.6	79.6	87.9	89.6	88.2	<u>90.2</u>	87.6
endpoint	77.2	36.5	80.8	67.9	76.3	35.7	<u>82.8</u>	<u>81.0</u>	77.8	80.2




Figure 8: Box plot showing the prediction accuracy on the test data across k-folds using DNN model for both FaaS platforms.

For both the FaaS platforms, the accuracy for the *nodejs-endpoint*, and *sentiment-analysis* FaaS functions is comparatively less as compared to other FaaS functions since most parameters in them are non-linearly correlated with FC. **Polynomial Linear Regression (PLR)** leads to highly accurate results for most of the function types due to its ability to model non-linear relations among the parameters. With PLR, the estimated FC accuracies are higher than the linear regression, which is attributed to most parameters being non-linearly correlated. **Ridge Regression (RR)** produced approximately the same results as that of linear regression and worked well for certain function types. On the other hand, **Random Forest Regression (RFR)** can provide the best results among the statistical approaches.

The **Deep Neural Network (DNN)** method outperformed all the statistical approaches for most of the FaaS functions since it is capable of modeling the linear and non-linear correlation between the parameters. For most function types, the FC estimation accuracy is approximately above 75%. In Figure 8, we show the prediction accuracy percentage for k-folds in the case of DNN on the test data using the box plot for both FaaS platforms, and all the FaaS functions.

6 RELATED WORK

FaaS is priced as per the pay-as-you-go model, where they charge the number of function invocations and memory. The majority of the prior work is based on optimizing the scalability, cost, execution time, integration support, and the constraints associated with FaaS services such as AWS Lambda, GCF, and Azure Functions (AF) provided by public cloud providers [13, 16, 17, 19, 20]. However, there is no prior research for quantifying and estimating the FCs

based on the different deployment configurations using GCF and AWS Lambda. We present prior work from two aspects:

6.1 Performance impact of the underlying system

FaaSProfiler [28] is the first to take a bottom-up approach in analyzing the architectural implication to unwrap the server-level overheads in the FaaS model. They analyzed the difference between native and in-FaaS function execution and calculated the additional server-level overheads like computational overheads, memory consumption, bandwidth usage, and management overheads like orchestration, queuing, scheduling, and power consumed. Furthermore, Lee et al. [24] compared the performance of various serverless computing environments offered by public cloud providers by showcasing the results of throughput, network bandwidth, file I/O and compute performance on concurrent function invocations. Wang et al. [30] performed an in-depth study of resource management and performance isolation with three popular serverless computing providers: AWS Lambda, Azure Functions, and GCF. Their analysis demonstrates a reasonable difference in performance between the FaaS platforms. Moreover, the deployed function instances are associated with VMs having a variety of configurations with different underlying host micro-architectures. In this work, we showcase a strong correlation between the function’s performance and the resources allocated to it based on the different configurations.

Additionally, Figiela et al. [14] developed a cloud function benchmarking framework. Compute-intensive functions were deployed in major cloud providers’ FaaS platforms. The authors observe variation in response time duration based on the different underlying hardware and runtime environment. These observations encouraged us to proceed with estimating the FCs of functions when deployed with different deployment configurations.

6.2 Performance modeling of the FaaS application

Pawlik et al. [26] state that to assess the feasibility of running an application on the FaaS platform, we have to determine the SLO of the application. This can be achieved by constructing a reliable performance model capable of analyzing a function performance, which requires knowledge about the performance of the infrastructure. Cloud service providers abstract details such as the number of cores, memory available, and network I/O capacity in the underlying hardware, usually limiting the available information to function time limit, maximum memory. The allocated memory also affects the provisioned CPU quota [25]. In our previous work [21], we developed

a tool for estimating the capacity of a microservice application when it is sandboxed. We followed the similar approach in this work for the capacities estimation of the FaaS functions.

7 CONCLUSION AND FUTURE WORK

In this paper, we demonstrated the impact of various configuration parameters on the Function Capacity of the two FaaS platforms (AWS Lambda and GCF). The methodology and the tool *FnCapacitor* introduced in this paper aim to solve the problem of estimating the Function Capacity at a specific deployment configuration. In summary, *FnCapacitor* can be used in two ways:

- By application developers in an offline manner for estimating the FCs of FaaS functions within their application for different deployment configurations. Based on the estimated FCs and the requirements, the developer can deploy the functions with the right configurations.
- By application developers in an online manner, the tool collects the monitoring data of the already existing FaaS functions. It automatically builds the models in the background without additional load testing. The built models can then be used to update the deployment configurations of the functions depending on the required SLOs.

In the future, we plan to extend *FnCapacitor* with other public serverless compute providers and to open source FaaS platforms. Creating a function scheduling framework based on the estimated FCs, the functions can be created with the suitable deployment configurations is another future direction.

ACKNOWLEDGMENTS

This work was supported by the funding of the German Federal Ministry of Education and Research (BMBF) in the scope of the Software Campus program. Google Cloud credits in this work were provided by the *Google Cloud Research Credits* program with the award number NH93G06K20KDXH9U.

REFERENCES

- [1] [n. d.]. AWS Lambda, <https://docs.aws.amazon.com/lambda/latest/dg/invocation-scaling.html>. Accessed on 09/24/2020.
- [2] [n. d.]. Azure Functions hosting options, <https://docs.microsoft.com/en-us/azure/azure-functions/functions-scale>. Accessed on 02/18/2021.
- [3] [n. d.]. Cloud Functions Overview. <https://cloud.google.com/functions/docs/concepts/overview>. (Accessed on 08/22/2020).
- [4] [n. d.]. The regression score, https://scikit-learn.org/stable/modules/generated/sklearn.metrics.r2_score.html. Accessed on 12/17/2020.
- [5] Nabeel Akhtar, Ali Raza, Vatche Ishakian, and Ibrahim Matta. 2020. COSE: Configuring Serverless Functions using Statistical Learning. In *IEEE INFOCOM 2020 - IEEE Conference on Computer Communications*. 129–138. <https://doi.org/10.1109/INFOCOM41043.2020.9155363>
- [6] Joao Carreira, Pedro Fonseca, Alexey Tumanov, Andrew Zhang, and Randy Katz. 2019. Cirrus: A serverless framework for end-to-end ml workflows. In *Proceedings of the ACM Symposium on Cloud Computing*. 13–24.
- [7] Paul Castro, Vatche Ishakian, Vinod Muthusamy, and Aleksander Slominski. 2019. The Rise of Serverless Computing. *Commun. ACM* 62, 12 (Nov. 2019), 44–54. <https://doi.org/10.1145/3368454>
- [8] Mohak Chadha, Anshul Jindal, and Michael Gerndt. 2020. Towards Federated Learning Using FaaS Fabric. In *Proceedings of the 2020 Sixth International Workshop on Serverless Computing (Delft, Netherlands) (WoSC'20)*. Association for Computing Machinery, New York, NY, USA, 49–54. <https://doi.org/10.1145/3429880.3430100>
- [9] Mohak Chadha, Anshul Jindal, and Michael Gerndt. 2021. Architecture-Specific Performance Optimization of Compute-Intensive FaaS Functions. *CoRR* abs/2107.10008 (2021). [arXiv:2107.10008](https://arxiv.org/abs/2107.10008) <https://arxiv.org/abs/2107.10008>
- [10] Ryan Chard, Tyler J. Skluzacek, Zhuozhao Li, Yadu Babuji, Anna Woodard, Ben Blaiszik, Steven Tuecke, Ian Foster, and Kyle Chard. 2019. Serverless Supercomputing: High Performance Function as a Service for Science. [arXiv:1908.04907](https://arxiv.org/abs/1908.04907) [cs.DC]
- [11] Simon Eismann, Long Bui, Johannes Grohmann, Cristina L. Abad, Nikolas Herbst, and Samuel Kounev. 2021. Sizeless: Predicting the optimal size of serverless functions. [arXiv:2010.15162](https://arxiv.org/abs/2010.15162) [cs.DC]
- [12] Lennart Espe., Anshul Jindal., Vladimir Podolskiy., and Michael Gerndt. 2020. Performance Evaluation of Container Runtimes. In *Proceedings of the 10th International Conference on Cloud Computing and Services Science - CLOSER.* INSTICC, SciTePress, 273–281. <https://doi.org/10.5220/0009340402730281>
- [13] Chen-Fu Fan., Anshul Jindal., and Michael Gerndt. 2020. Microservices vs Serverless: A Performance Comparison on a Cloud-native Web Application. In *Proceedings of the 10th International Conference on Cloud Computing and Services Science - CLOSER.* INSTICC, SciTePress, 204–215. <https://doi.org/10.5220/0009792702040215>
- [14] Kamil Figiela, Adam Gajek, Adam Zima, Beata Obrok, and Maciej Malawski. 2018. Performance evaluation of heterogeneous cloud functions. *Concurrency and Computation: Practice and Experience* 30, 23 (2018).
- [15] GoogleCloud. [n. d.]. Cloud Functions Pricing. <https://cloud.google.com/functions/pricing>. (Accessed on 08/22/2020).
- [16] Jake Grogan, Connor Mulready, James McDermott, Martynas Urbanavicius, Murat Yilmaz, Yalemisew Abgaz, Andrew McCarren, Silvana MacMahon, Vahid Garousi, Pooyan Jamshidi, et al. [n. d.]. An analysis of Function-as-a-Service (FaaS): vendors, challenges and implications for software developers. ([n. d.]).
- [17] Anshul Jindal, Mohak Chadha, Michael Gerndt, Julian Frielinghaus, Vladimir Podolskiy, and Pengfei Chen. 2021. Poster: Function Delivery Network: Extending Serverless to Heterogeneous Computing. In *2021 IEEE 41st International Conference on Distributed Computing Systems (ICDCS)*. 1128–1129. <https://doi.org/10.1109/ICDCS51616.2021.00120>
- [18] Anshul Jindal, Julian Frielinghaus, Mohak Chadha, and Michael Gerndt. 2021. Courier: Delivering Serverless Functions Within Heterogeneous FaaS Deployments. In *2021 IEEE/ACM 14th International Conference on Utility and Cloud Computing (UCC'21) (Leicester, United Kingdom) (UCC '21)*. Association for Computing Machinery, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3468737.3494097>
- [19] Anshul Jindal and Michael Gerndt. 2021. From DevOps to NoOps: Is It Worth It?. In *Cloud Computing and Services Science*, Donald Ferguson, Claus Pahl, and Markus Helfert (Eds.). Springer International Publishing, Cham, 178–202.
- [20] Anshul Jindal, Michael Gerndt, Mohak Chadha, Vladimir Podolskiy, and Pengfei Chen. 2021. Function delivery network: Extending serverless computing for heterogeneous platforms. *Software: Practice and Experience* 51, 9 (2021), 1936–1963. <https://doi.org/10.1002/spe.2966> <https://onlinelibrary.wiley.com/doi/pdf/10.1002/spe.2966>
- [21] Anshul Jindal, Vladimir Podolskiy, and Michael Gerndt. 2019. Performance modeling for cloud microservice applications. In *Proceedings of the 2019 ACM/SPEC International Conference on Performance Engineering*. 25–32.
- [22] Eric Jonas, Qifan Pu, Shivaram Venkataraman, Ion Stoica, and Benjamin Recht. 2017. Occupy the cloud: Distributed computing for the 99%. In *Proceedings of the 2017 Symposium on Cloud Computing*. 445–451.
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2017. ImageNet Classification with Deep Convolutional Neural Networks. *Commun. ACM* 60, 6 (May 2017), 84–90. <https://doi.org/10.1145/3065386>
- [24] Hyungro Lee, Kumar Satyam, and Geoffrey Fox. 2018. Evaluation of production serverless computing environments. In *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*. IEEE, 442–450.
- [25] W. Lloyd, S. Ramesh, S. Chinthalapati, L. Ly, and S. Pallickara. 2018. Serverless Computing: An Investigation of Factors Influencing Microservice Performance. In *2018 IEEE International Conference on Cloud Engineering (IC2E)*. 159–169. <https://doi.org/10.1109/IC2E.2018.00039>
- [26] Maciej Pawlik, Kamil Figiela, and Maciej Malawski. 2018. Performance evaluation of parallel cloud functions. *Poster Presented at ICPP* (2018).
- [27] Serverless. 2020. Documentation. <https://serverless.com/framework/docs/>. [Online; Accessed: 4-February-2020].
- [28] Mohammad Shahrad, Jonathan Balkind, and David Wentzlaff. 2019. Architectural implications of function-as-a-service computing. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*. 1063–1075.
- [29] Vaishaal Shankar, Karl Krauth, Qifan Pu, Eric Jonas, Shivaram Venkataraman, Ion Stoica, Benjamin Recht, and Jonathan Ragan-Kelley. 2018. Numpywren: Serverless linear algebra. *arXiv preprint arXiv:1810.09679* (2018).
- [30] Liang Wang, Mengyuan Li, Yinqian Zhang, Thomas Ristenpart, and Michael Swift. 2018. Peeking behind the curtains of serverless platforms. In *2018 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)*. 133–146.
- [31] Sanjay Yadav and Sanyam Shukla. 2016. Analysis of k-fold cross-validation over hold-out validation on colossal datasets for quality classification. In *2016 IEEE 6th International conference on advanced computing (IACC)*. IEEE, 78–83.