



**UNIVERSIDADE FEDERAL DE CAMPINA GRANDE  
CENTRO DE ENGENHARIA ELÉTRICA E INFORMÁTICA  
CURSO DE BACHARELADO EM CIÊNCIA DA COMPUTAÇÃO**

**DAVID FERREIRA QUARESMA**

**VALIDATION OF A SIMULATION MODEL FOR FAAS  
PERFORMANCE BENCHMARKING USING PREDICTIVE  
VALIDATION**

**CAMPINA GRANDE - PB**

**2021**

**DAVID FERREIRA QUARESMA**

**VALIDATION OF A SIMULATION MODEL FOR FAAS  
PERFORMANCE BENCHMARKING USING PREDICTIVE  
VALIDATION**

**Trabalho de Conclusão Curso  
apresentado ao Curso Bacharelado em  
Ciência da Computação do Centro de  
Engenharia Elétrica e Informática da  
Universidade Federal de Campina  
Grande, como requisito parcial para  
obtenção do título de Bacharel em  
Ciência da Computação.**

**Orientador: Professor Dr. Thiago Emmanuel Pereira da Cunha Silva.**

**CAMPINA GRANDE - PB**

**2021**



Q18v Quaresma, David Ferreira.  
Validation of a simulation model for FaaS performance benchmarking using predictive validation. / David Ferreira Quaresma. - 2021.

11 f.

Orientador: Professor Dr. Thiago Emmanuel Pereira da Cunha Silva.

Trabalho de Conclusão de Curso - Artigo (Curso de Bacharelado em Ciência da Computação) - Universidade Federal de Campina Grande; Centro de Engenharia Elétrica e Informática.

1. Function as a Service - FaaS. 2. Experimentos de simulação. 3. Validação de modelo de simulação. 4. Modelo de simulação. 5. Modelo de validação. I. Silva, Thiago Emmanuel Pereira da Cunha. II. Título.

CDU:004(045)

**Elaboração da Ficha Catalográfica:**

Johnny Rodrigues Barbosa  
Bibliotecário-Documentalista  
CRB-15/626

**DAVID FERREIRA QUARESMA**

**VALIDATION OF A SIMULATION MODEL FOR FAAS  
PERFORMANCE BENCHMARKING USING PREDICTIVE  
VALIDATION**

**Trabalho de Conclusão Curso  
apresentado ao Curso Bacharelado em  
Ciência da Computação do Centro de  
Engenharia Elétrica e Informática da  
Universidade Federal de Campina  
Grande, como requisito parcial para  
obtenção do título de Bacharel em  
Ciência da Computação.**

**BANCA EXAMINADORA:**

**Professor Dr. Thiago Emmanuel Pereira da Cunha Silva**

**Orientador – UASC/CEEI/UFCG**

**Professora Dra. Melina Mongiovi Cunha Lima Sabino**

**Examinador – UASC/CEEI/UFCG**

**Professor Tiago Lima Massoni**

**Professor da Disciplina TCC – UASC/CEEI/UFCG**

**Trabalho aprovado em: 25 de Maio de 2021.**

**CAMPINA GRANDE - PB**

## RESUMO (ABSTRACT)<sup>1</sup>

No artigo Controlling Garbage Collection and Request Admission to Improve Performance of FaaS Applications, nós verificamos e avaliamos o impacto dos mecanismos de gerência de memória das linguagens de programação no contexto de Funções como Serviço (Function as a Service, FaaS) através de experimentos de simulação. Os resultados desse estudo apontaram um impacto de 11.68% no tempo de resposta das requisições quando os procedimentos de coleta de lixo foram executados durante a execução de uma função CPU intensiva. Como trabalhos futuros, nós listamos algumas ameaças à validade dos resultados obtidos, e entre eles, nós citamos a validação do modelo de simulação usado. A validação do modelo é importante, pois ela também valida os resultados gerados nos experimentos de simulação, o que garante resultados realistas. Neste trabalho, nós propusemos e executamos uma validação do modelo de simulação usado no trabalho anterior. Para atingir este objetivo, nós executamos experimentos de medição em uma plataforma FaaS pública e experimentos de simulação usando o mesmo simulador do artigo anterior. Sendo assim, nós validamos o simulador comparando os resultados obtidos em ambos os experimentos para garantir que o resultado da simulação e o de medição são equivalentes.

---

<sup>1</sup>

# Validation of a simulation model for FaaS performance benchmarking using predictive validation

David Ferreira Quaresma\*  
Federal University of Campina Grande

Campina Grande, Paraíba, Brazil  
david.quaresma@ccc.ufcg.edu.br

Thiago Emmanuel Pereira  
Federal University of Campina Grande

Campina Grande, Paraíba, Brazil  
temmanuel@computacao.ufcg.edu.br

Daniel Fireman  
Federal Institute of Alagoas  
Arapiraca, Alagoas, Brazil  
daniel.fireman@ifal.edu.br

## ABSTRACT

In the paper Controlling Garbage Collection and Request Admission to Improve Performance of FaaS Applications[30], we verified and evaluated the impact of memory management mechanics of programming languages in the context of Functions as a Service (FaaS) via simulation experiments. The results of this study pointed to an impact of up to 11.68% on the response time of requests when a garbage collector procedure was executed during the execution of a CPU-bound function. As future work, we listed a few threats to the validity of the results attained, and among them, we cited the validation of the simulation model used. The validation of the model is important because it validates the results generated in the simulation experiments, which ensures realistic results. In this work, we proposed and executed a validation to the simulation model used in the previous work. To do so, we run measurement experiments in a public FaaS platform and simulation experiments of the same scenarios using the same simulator of the previous paper. Then, we validate the simulator by comparing the results obtained in both experiments to ensure that the simulation result and the measurement one are equivalent.

### ACM Reference Format:

David Ferreira Quaresma, Thiago Emmanuel Pereira, and Daniel Fireman. 2021. Validation of a simulation model for FaaS performance benchmarking using predictive validation. In *Proceedings of Universidade Federal de Campina Grande (UFCG)*. UFCG, Campina Grande, PB, Brazil, 8 pages.

## 1 INTRODUCTION

Function as a Service (FaaS) is a recent offer of computing as a service and a model of Serverless computing [20]. It is an event-driven computing model that is based on the execution of ephemeral stateless functions in containers [32]. A function, in this model, is triggered by events predefined, and all the burden of resource management is outsourced to the platform. An event, for instance, might be a function invocation request from a user client or another function instance. In this way, FaaS promises to reduce the time to deliver software, as well as providing flexible scalability (on function level) and a more precise billing system based on the time in which each function has been processing events [9, 22].

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 UFCG 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 david.quaresma@ccc.ufcg.edu.br

UFCG, May 2021, Campina Grande, PB, Brazil  
© 2021 Universidade Federal de Campina Grande.

Although popular and promising [19], FaaS is still a new Cloud Computing offer. There are studies categorizing it [10], evaluating its performance [23, 29], and proposing new approaches to deal with already known problems [37]. Still, there are open problems in this field and among them, we emphasize the performance degradation of FaaS platforms (and applications) due to the runtime background activity [30].

Using simulation, we estimated the consequences of using Garbage Collector (GC) for FaaS scenarios [30]. We observed that when a GC routine was triggered during function execution, the response time of this function call was increased by up to 11.68% [30]. We also proposed an approach to mitigate this impact by controlling garbage collection and request admission using the Garbage Collector Control Interceptor (GCI), a technique used to avoid the impact of GCs [12, 13]. The results attained of using the GCI in the FaaS context were the reduction of the microservice response time of up to 10.86% and saving of resources costs of 7% [30].

The previous study mentioned is still a work in progress. There are some threats to the validity of results that we made sure to point out, such as the lack of variety of used functions or a more representative set of FaaS functions, the usage of a synthetic workload upon a discrete probability distribution (Poisson), and the absence of validation for the simulator used. Although the simulator was properly tested and verified, no further validation was done. The validation process is important because it gives us credibility about its results. Thus, the main objective of this work is to validate the simulator used in the previous work and evaluate if the model and production environment are similar. To do so, we executed experiments using the simulation model and the scenarios of the previous study and compared them under statistical analysis.

The remainder of this paper is organized as follows. In Section 2, we introduce the basics of our previous work [30], such as the results and conclusions attained, and then we talk about main concepts of validation, such as validation terminology and validation approach to be used in this work. Then, in Section 3, we present our approach to validation, describing the simulator model and its components, the validation process used, and the measurements experiments. In Section 4, we present and discuss the results attained. Finally, in Section 5, we present our final remarks and possible future work.

## 2 BACKGROUND

In this section, we describe the previous study we are going to continue as well as the basic concept behind this work, which is the validation process. In Section 2.1 we introduce our previous study [30] discussing what it is, what results were accomplished, and briefly how it was attained. Following, in Section 2.2 we present

the basics of the validation process we are going to use in this work, such as definition and general usage.

## 2.1 Previous Work

As mentioned in the Introduction, in the previous work [30], we evaluated the drawback of letting FaaS functions be vulnerable to Garbage Collector (GC) activity. The paper “Controlling Garbage Collection and Request Admission to Improve Performance of FaaS Applications” [30] not only showed that a GC might increase the response time by up to 11.68%, but also provided a solution to mitigate this impact that reduced the same metric of up to 10.86% and reduced the resource costs by 7%. As said, although the results attained are very promising, some threats to its validity require further work.

The previous work [30] is a study that evaluated the relation of performance (in terms of service time) degradation of FaaS functions and the memory management procedures of the programming language in which the function was coded with. The main two scenarios studied were, initially, a function execution that was impacted by garbage collections and another function execution free of any GC intervention. In both scenarios, the function used was the image resizer, a popular FaaS function in the literature [2, 27].

Once the GC impact was properly evaluated with a workload much simpler, simulation experiments were performed to evaluate it under a more complex workload configuration. The complex workload configuration used was a series of function invocations in which the time between each invocation was defined following a probability distribution - the Poisson distribution to be precise. The FaaS platform behavior and function environment considered to the experiments were based on the AWS Lambda. The same process was executed to compare the feasibility of the solution proposed.

The results attained in this study, and its conclusions, were obtained using simulation experiments. The simulation model used is described further in Section 3.1. For the scope of the previous work, the precision of the simulated distribution shape of the response time was essential since the strategy behind the solution suggested is highly impacted by the shape of the distribution.

## 2.2 Model Validation

A key concept to understand this work is validation. While verification often stands for ensuring that a conceptual model and its implementation are correct [11, 33], model validation stands for a satisfactory accuracy of the computerized model with the intended application of the model within its domain of applicability [1, 33]. In other words, validation is a process where the used model is tested to ensure that it does represent what it is proposed to represent, considering the scenarios that are important to it. It is an important process because it validates the model tested and also any information acquired from using this model in some simulation experiment.

There are a variety of methods used to validate simulation models and assert its validity, ranging from approaches such as comparisons between models to the usage of data generated by an actual system [26]. In this study, to validate the simulator, we use a Predictive Validation technique. The Predictive Validation approach

consists of using a simulation model to predict, or forecast, the behavior of the target system, and then compare the results attained to determine if they are the similar [34]. We decided to use the Predictive Validation because it was the technique that best feat in our conditions, but there were two other interesting approaches as well, such as the Face Validity and the Comparison to Other Models.

## 3 METHODOLOGY

The main goal of this work is to validate whether the simulator used in the previous work [30] is able to produce realistic results from various scenarios. Thus, in this work, we perform a validation of the simulator mentioned using the predictive validation approach, as highlighted in Section 2.2. We evaluate if the results attained with the simulator are realistic by comparing them to measurement experiment results, which are, by default, realistic since they are collected from a real function in production.

To avoid misleading understandings, we emphasize that since the input of the simulator is generated by measurement experiments, there are two sorts of measurement experiments. One of them (described in Section 3.3.1) was used as simulation input, while the other one (described in Section 3.3.2) was used for comparison purposes. Thus, we will call the measurement experiments used for the simulation input just as input experiments.

The remainder of this section is described as follows: In Section 3.1, we discuss the simulation model used in the previous [30] and current work and briefly describe each main component of the Simulator. Then, in Section 3.2, we discuss the validation process, metrics, and the scope of this validation. Finally, in Section 3.3 we describe the measurement experiments performed to provide input to the simulations as well as realistic results for comparison and validation purposes.

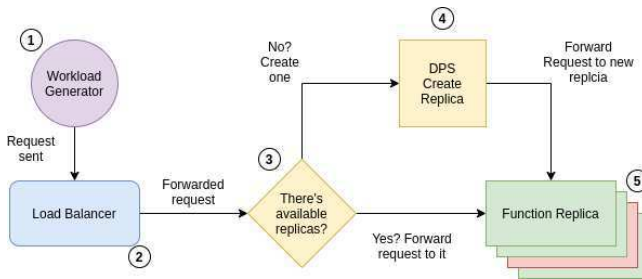
### 3.1 Simulation Model

Even though lots of infrastructures of public FaaS platforms have many aspects in common, such as a simple deployment, a dynamic resource provisioning, and a pay-per-use billing, they are not equals. In fact, they vary a lot in many other aspects. For instance, GCP Cloud Functions [8] and AWS Lambda [4] differ in terms of available events that triggers the execution of functions, resource specs, and supported execution environment. Not only the platforms differ from each other, but applications (the functions) might as well, and their workload too (which varies from applications to applications).

So, during the modeling of the simulator used, some challenges showed up. First, since FaaS is fairly new, how to choose a proper function, platform, and workload to representatively evaluate it? and, for instance, even taking a public FaaS provider as a reference, how to properly simulate its environment if its internal settings are typically not public? To address these issues, in [30], we decided to use AWS Lambda platform as a reference and executed a couple of functions on it to collect as much internal information as possible. Regarding a representative function, we solely choose to use as reference a FaaS function popular in literature and run it using the settings we were able to retrieve. Finally, regarding to the workload, to make things simple, we choose a synthetic workload using the Poisson distribution to generate the inter-arrival of each

function invocation during simulations. Thus, the model to validate is purposed to simulate the AWS Lambda environment, by a chosen function, and a chosen workload. It is a simplified version of the platform having only the main abstractions of it.

The decision of use the AWS Lambda was made based on the fact that AWS Lambda platform is a very popular public offering in this field [31, 38]. Among its characteristics, we highly considered and applied in this simulation model the serial execution of requests by function replicas [5], the termination of function replicas after an idle period of time [17], and the reuse of replicas most used recently. The internal settings collected were used to understand the platform and also to setup the measurement experiment (read Section 3.3 for further details about the measurement experiments) since in our study we used the same internal settings of AWS Lambda to run experiments and measure the response time to be used as blueprint by the simulations. Lastly, the chosen function, as introduced in [2], was a image resizer function popular in FaaS literature [2, 27].



**Figure 1: Request's flow through the simulated model components. The Workload Generator (WG) periodically sends requests to the Load Balancer (LB), which redirects it to the Dynamic Resources Provisioning System (DRPS), the one that make sure to redirect the request to a available function, even if it have to create another function replica.**

This simulation model is able to simulate a FaaS platform environment by only receiving a specific input - a list of response time of each function invocation measured from a FaaS function deployed. This list must be collected through measurement experiments using a sequential workload (as presented later in 3.3.1) to avoid functions invocations of running concurrently, and is required one to each function replica to be simulated. Thus, the simulator reproduce, to each function replica, the respective input file and simulate how the provider scale up and/or scale down based on the demand. Whenever a new invocation arrives and there is no available replica, in this simulation model a new function replica is created and the response time is impacted by a cold start.

The simulator model [15] is composed of four main components [30]: 1) the Workload Generator (WG) - responsible to generate events processed by the function replicas according to a given statistical distribution; 2) the Load Balancer (LB) - responsible to balance load among replicas in a given algorithm fashion; 3) a Dynamic Resources Provisioning System (DRPS) - responsible for scaling the pool of replicas to match the demand; and 4) the Function Replicas - the one responsible to handle requests. Figure 1 summarizes the request processing flow and the relationship between the components of the model. We detail these components following below.

**3.1.1 Workload Generator.** The first entity in this model is the Workload Generator (WG). The primary goal of WG is to send requests to the load balancer according to a statistical distribution. More precisely, the WG defines the time between the sending of each request based on a discrete probability distribution. So, after a given number of simulated time, the WG sends to the load balancer a function invocation. In this work, we modeled the inter-arrival time according to a Poisson distribution. In Figure 1, it is the first step of the requests flow in the simulation environment.

**3.1.2 Load Balancer.** The second entity of the simulated model is the Load Balancer (LB). We chose this model based on observations of how AWS Lambda works. Since we know AWS Lambda scale down its replicas based on inactivity time, we judge that its LB cannot distribute the load using popular policies such as Round-Robin, once it would uniformly restart the idle time counter. So, our LB receives all generated requests from the WG and schedules them to an available replica. Regarding the scheduling, the LB chooses the replica which has most recently become available. Lastly, the requests are processed according to the first-in-first-out (FIFO) policy. In Figure 1, it is the second component an invocation hits after being generated by the WG.

**3.1.3 Dynamic Resources Provisioning System.** The third component of this simulator model is the Dynamic Resources Provisioning System (DRPS). It is responsible to start new function replicas or terminating the idle ones whenever is needed. It starts new replicas when a new request arrives, and there is no available replica to process it. This is the case when all replicas are busy processing other requests or when there is no replica available. The DRPS automatically terminates replicas when they remain idle for some time. This idleness period until replica removal is configurable, and the default value is 5 minutes. In Figure 1, it is composed by the steps 3 and 4, and is the last component an invocation hits before being received by a function replica.

**3.1.4 Function Replica.** The last component of this model is the Function Replica. Our model of the Function Replica dictates its performance, as response times, when processing requests. In addition to the response time, the modeled response time also accounts for runtime environment interference, for instance, the impact of runtime startup, automatic memory management, and cold starts. This last one, cold starts, is a problem of virtualization technique that FaaS has inherited due to the fact that functions are executed in containers [24]. The cold start is a delay in the function execution due to the fact that there is no ready resource, or container, to host the execution [37].

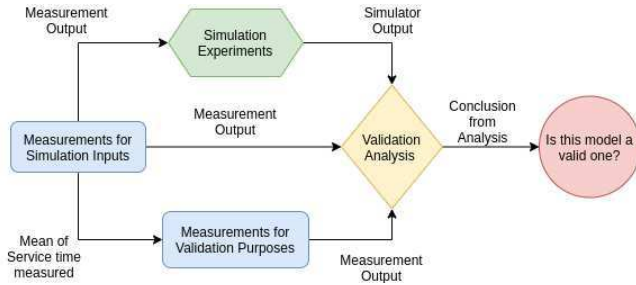
The model of each function replica is represented as a sequence of (*duration, status code*) tuples. Each tuple of this sequence represents the processing of a request. Thus, the model simulates the replica behavior by reproducing the duration and the status code of each tuple in that sequence. For instance, upon receiving the first request, the simulated replica returns the first tuple of the set associated with it. The procedure repeats to all the following replicas and requests.

In Figure 1, when the generated function invocation hits the step 5, the Function Replica assign to this invocation the duration and status code of the current tuple, simulating the behavior of this



invocation to be the same of one of the invocations measured before. In the Figure, function might green or red to illustrates that it might be available or busy. To better represent the state of practice, in our simulations, the model of the function replicas was instantiated based on experimental measurements (explained in Section 3.3), and each simulated replica is associated with a result of a different measurement experiment.

### 3.2 Validation Process



**Figure 2: Flow diagram of the validation process. The first step was the measurement experiment to collect the simulation input and statistics from the function deployed. Then a simulation and measurement experiment were executed based on this data. Finally, the whole data was analyzed in order to validate the simulation model**

The process of validation used in this work was predictive validation. It consists of executing measurement and simulation experiments, both for the same scenarios, and comparing the results obtained under statistical analysis to assert if the simulation model was able to forecast the system behavior. This section describes in detail the validation process designed and each step of its execution as shown in Figure 2. Later, we discuss the simulation experiments used for this process.

To perform the simulation experiment we need to generate the proper input to the simulator, which we get from the input experiments pointed in Section 3.1.4 and described further in Section 3.3.1. Briefly, the input experiments is a simple measurement experiments, an experiment using a sequential workload in order to avoid concurrent execution of functions. After collecting the input required, we execute simulation and measurement experiments using a more complex workload for the same scenarios. Once the measurement and simulation experiment are done, we compare, as said, its results under a confidence interval to ensure that the simulation results can be representative and relevant. We want to validate the results of the simulator and check whether its results are realistic for the scenarios studied previously [30].

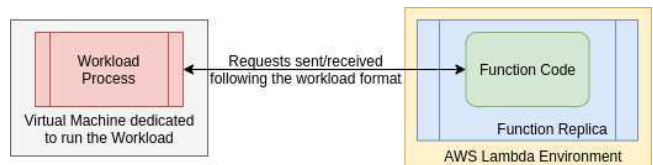
Figure 2 illustrates the whole process of validation in a flow diagram. The first step in the validation process is the input experiments aimed to collect the simulation input and statistical measures from the function behavior in the production environment. More details about this experiments in 3.3.1. Next, we execute two sorts of experiments: a simulation experiment (described in Section 3.4) and a measurement experiment for validation purposes (described

in Section 3.3.2). Both these measurement and simulation experiments focus on collect mainly the service time, i.e. the time took to execute function code.

After that, as in the Figure 2, all the output from the simulation and measurement experiments (including the input experiments) are used in the analysis step, the one in which our validation conclusion comes from. Next, both measurement and input experiments are described in detail later in Section 3.3, and Section 3.4 focused solely in describe the simulation experiment.

### 3.3 Measurement Experiments

As introduced in Section 3, our study contains two types of measurement experiments, the input experiments and the measurement experiments. The input experiments are the measurement experiments executed aiming to generate inputs to the simulation and the measurement experiments are the ones executed in order to generate results for comparison purposes. The input experiments focused on generating the simulator inputs and are essential to the understanding of the function behavior in the AWS Lambda environment. It also was used to generate the workload used in the measurement experiments, since the workload settings required the mean of the function response time. These experiments were carefully design to mitigate interference from external sources to focus only on the runtime environment. Thus, we argue that the observed response time is mainly due to the function service time, i.e., the function code plus runtime procedures, and the cold start. For both two measurement experiments, the illustration in Figure 3 emphasizes the environment of these measurements. The following sections describes in details the measurement experiments executed for simulation input and validation comparison purposes.



**Figure 3: Experiment environment for both measurement experiments. The Workload Generator (WG) running in a different virtual machine sends requests to a function replica in the AWS Lambda environment and collect the service time (time took to process the function code). The WG sends the requests following a predefined distribution (sequential or Poisson, depending on the type of the measurement experiment).**

**3.3.1 Simulation Inputs.** In order to collect the simulation inputs, our experiments measured the performance of a resize function [14], which is a popular application used in performance evaluation[2, 27]. Each function invocation reduces a 560KB sized image to 10% of its original size. To reduce the effect of I/O transfers and network fluctuations in experiment results, the original image is loaded and stored in memory during the function startup, and the resized images are not sent back to the client, only the service time (time to execute the function code) to resize the image. Each

measurement experiment observes the behavior of a single instance of the resize function application running on the AWS Lambda functions environment.

For each run, a process running in a different virtual machine sends a sequence (which means that a request was sent only after the receiving of the response of the previous one) of 5.000 requests, stores its response times and status codes. Our analysis ignored the response time of the 250 (5%) first requests of each run to soften the effects of the warmup [7]. We executed a total of 32 runs. And between each run, we waited one hour to make sure a new instance is created and the effects of cold start properly accounted for each measurement experiment.

The measurement experiments for simulation inputs were also useful to define the Poisson distribution used in the 3.3.2. Thus, we used the mean of the response time of the function invocations measured for the simulation inputs as the lambda value of the Poisson distribution used in the measurements experiments for validation purposes. This way, the mean of the inter-arrival of function invocations is set to be equal to the mean of the response time of the function. Additionally, these measurement experiments were also used for sanity and any additional analysis that require further understanding of the response time of the function behavior in the AWS Lambda environment.

**3.3.2 Measurements for Validation.** The measurement experiments for comparison and validation purposes consist of an execution of a workload in a different virtual machine using a Poisson distribution to define the time between each function invocation. Differently from the measurements in 3.3.1, in this workload we surely have concurrency occurring and triggering new function instances, having the concurrency level depending on the distribution intensity. As well as in 3.3.1, this measurement experiment also used the resizer function [14], in which each function invocation resizes a 560KB image to 10% of its original size, and all procedures to avoid I/O transfers and network fluctuations in the experiments was also used here.

For each run, a workload sent 20.000 requests to the function replica in AWS lambda environment with the inter-arrival following the Poisson distribution, then stores the responses and status codes. In order to avoid warmup [7] influence, we ignored the response time of the 1.000 (5%) first requests of each run. Still, as well as in 3.3.1, between each run we waited for enough time (more than one hour) to make sure that the function replica was not ready to run at the beginning of each experiment. A total of 8 simulations, 4 for each lambda value, was executed.

The Poisson distribution used, as pointed in 3.3.1, was defined with the lambda value equal to the mean of the response time of the function invocations measured during the measurement experiments for the simulation inputs. By defining the lambda value to that, we ensure that the Poisson distribution periodically generates a function invocation before the previous request is finished, and therefore, the system have to trigger a new function replica, increasing the concurrency level.

### 3.4 Simulation experiments

In order to validate the simulation model used in the previous work [30] and described in Section 3.1, we executed a series of

simulation experiments. The main objective of this experiment is to simulate the measurement experiments described in Section 3.3.2 to later compare the results and validate the simulator. This experiment simulated the sending of 20.000 function invocations following the Poisson distribution, using the same settings of the respective measurement experiments. In order to avoid warmup [7] influence, for each run, we removed the first 1.000 (5%) service times simulated. A total of 4 simulations run was executed.

Regarding the input of the simulator, to each possible function instance the simulation might have simulate, a respective input file previously measured is required. These input files were generated through the measurement experiments for the simulation inputs, as described in 3.3.1. A total of 32 input files was used in all simulation experiments to be reproduced among all function replicas created during simulation. Each of this files contained a total of 5.000 tuples as described in Section 3.1.4. Among the several outputs generated from the simulator, in our analysis, we focused mainly on the service time to the validation, but we also used instances and cold start statistics for comprehension concerns.

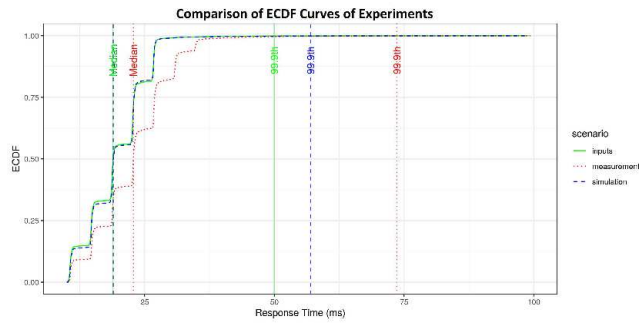
There are two limitations in our model. One is that the simulation might trigger more function instances than the input files passed, and the other is that the file might not have enough entries for the whole simulation. In the first case, if a new function instance is created and there is no unused input file, the simulator will reuse the one that was used less recently. In the latter case, if the number of entries within the input file is not enough for the whole simulation, the simulation will reset the file iteration to just after the cold start entry (to avoid multiples cold starts in the same function instance).

## 4 RESULTS

In this section, we describe the results attained in the validation process described in section 3.2. Briefly, by comparing the results of the measurement experiments with the ones from the simulation experiments, we observed that the simulated model was effective in simulating the shape of the response time distribution of the AWS Lambda environment. We also noticed an increase in the response time mean, median, and percentiles. However, this positive shift of value in the response time does not impact the validity of the conclusions in the study executed in the previous work [30] since the study relied only on the distribution shape for the conclusions, as talked in Section 2.1.

The remaining of this section presents and discuss the results attained in detail. We start our analysis by evaluating the Empirical Cumulative Distribution Function (ECDF) of the simulation experiments and the measurements experiments for input and validation purposes. Then, we studied the shape of the response time distribution generated in simulation and measurement experiments results. Finally, we compared statistics such as mean, median, and percentiles for all experiments, and shared the sanity checks made in order to testify the results observed.

Figure 4 presents the ECDF of the response time for the results attained in the simulation and measurements experiments. The vertical lines in the image indicate the median and percentile 99.9<sup>th</sup> of each distribution of each experiment results. Solid green line represents the results from input experiments, blue dashed lines represents the results from simulation experiments and red dotted

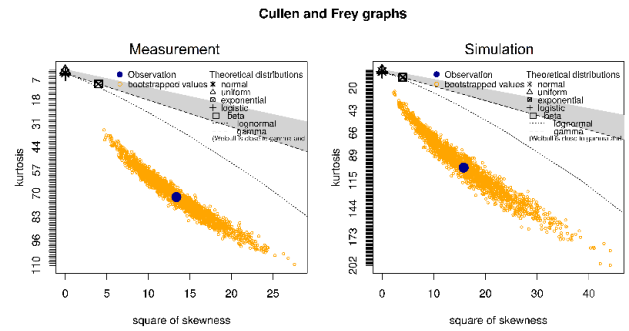


**Figure 4: The Empirical Cumulative Distribution Function (ECDF) curves of all the simulation and measurement experiments. While the ECDF of the simulation and input measurements experiments are quite identical, the measurement for comparison purposes experiments shares the same shape but has higher values of response time.**

lines means results from the measurement experiments. As we can see, in Figure 4 while simulation and input experiments share similar distribution shape and values, being likely identical curves, the measurement ECDF curve has higher response time values, even with a similar distribution shape. In the median, we see an overlap between the input and simulation experiments, indicating that both have the same median. Differently, in the percentile 99.9th, we have a minimal difference between the simulation and input experiments, with the simulation response time being slightly higher over 5ms. Regarding the median and percentiles of the measurement experiments, both the median of the response time is bit higher if compared to the ones obtained in the simulation and inputs experiments results. The measurement response time median is slightly greater (less than 5ms) than the ones from simulation and inputs, while the percentile 99.9th is considerably higher (more than 20ms).

In Figure 5, we have a Cullen and Grey Graph, specifically the skewness-kurtosis plot, For the simulation and measurement experiments. While Skewness is basically a measure of the relative size of the two tails of the distribution, the Kurtosis is the measure of the combined sizes of both tails [25]. Regarding skewness, as symmetric the distribution tails are, the closer to 0 the Skewness is. On the other side, the kurtosis measure is a measures of how peaked a distribution is. It is often measured comparing to the Kurtosis of the normal distribution, which is 3. So, as closer to 3 the Kurtosis of a distribution is, the closer it is to a normal distribution. If you need to evaluate a normal distribution under the skewness-kurtosis graphic, is expected skewness values closer to 0 and kurtosis values closer to 3.

As we can see in the Figure 5, both simulation and measurement experiments have very similar values for skewness and kurtosis, which indicates that the distribution of response time of both experiments is the same. Both distributions are highly tailored distributions, with a fairly large range in peak variation. We ensured that the workload applied in both scenarios was the same by studying how the workload concurrency triggered function instances in both experiments. The triggering of function instances might heavily impact the distribution shape due to cold starts, so peak



**Figure 5: Cullen and Frey graphs for both simulation and measurement experiments results. The skewness of the distribution indicates how symmetrical the distribution is, being skewness of 0 a perfect symmetrical distribution [25]. Kurtosis stands for a measure of whether the distribution is peaked, i.e., the distribution has most of the values in the center) [18]. In the graph, both simulation and measurement experiments shares similarly values for Skewness and Kurtosis.**

of high concurrency level along the workload would trigger more cold start.

We concluded, by analyzing the simulator and CloudWatch [3] logs that both experiment had the same peaks of concurrency and all cold start happened in the begging of the benchmarking. Since this studied focused in validade whether the simulation model is valid for the designed scarios of the previous work [30], we argue that this does not impact the validity of the conclusion of the previous study. However, we assert that for a more generalist usage of this simulation model, a more realistic workload would be required.

**Table 1: Presents a summary table of the percentiles under the confidence interval of 95% calculated for the simulation and measurement experiments. The table shows that the response time of the distributions studied is statistically different and that the measurement experiments has higher response time values.**

Percentile.	Measurement (ms)	Simulation (ms)
50 <sup>th</sup>	[22.83, 22.84]	[18.93, 18.97]
95 <sup>th</sup>	[34.76, 34.79]	[26.89, 26.91]
99 <sup>th</sup>	[39.12, 39.84]	[29.03, 30.68]
99.9 <sup>th</sup>	[69.14, 79.70]	[53.29, 60.28]

Table 1 shows the percentiles values for the measurement and simulation experiments, ranging from 50th to 99.9th. First, we noticed that for all percentiles and for the confidence interval of 95%, all percentiles studied for simulation and measurement experiments are different. Also, the higher the percentile, the higher is the response time of the measurement experiment, just as indicated in the Figure 4 in the median and percentile indicators. The mean of the response time of the simulation input is the same as the ones from input experiments, around 19ms, while the mean of response

time from measurement experiments is about 22ms. According to the range of the confidence interval, the difference between the mean of the simulation and measurement experiments goes from 3.86ms to 3.91ms. Respecting to the percentile of 99.9th, the difference between the simulation and measurement experiments varies from 8.86ms to 26.41ms. We argue that these differences in low enough to be ignored since it these small variations do not critically change the distribution shape as shown in Figure 4 and explained in Figure 5.

To the understanding of the results attained, a fundamental process in the analysis phase is the sanity check. In this process, we executed additional test cases, experimental scenarios, and studied logs from the AWS CloudWatch. We did so in order to understand why the increase of concurrency was also increasing the service time of function execution in the AWS Lambda environment. We noticed that in our results following, as higher our concurrency level was, the higher was the impact in the service time, but not proportionally. A workload with doubled concurrency level implied only plus 3 to 4 milliseconds in the mean of service time. We checked this statistics in the CloudWatch dashboard, and the CloudWatch logs indicated the same as ours. This was consistent and happened across all our experiment repetitions. We argue that the workload increased concurrency probably caused an overhead in the AWS Environment and slightly impacted the function execution, even running in isolated containers, due to the multi-tenancy characteristic of the AWS Lambda platform [35].

## 5 CONCLUSION

In the previous work [30], as introduced in Section 3.1, during the process of modeling the simulator, we defined an assumption in order to simplify the simulation model. We assumed that the input files used in the simulator would have a similar response time distribution shape. The discussion of the Figure 4 and Table 1 showed that the measurement experiments presented slightly higher values for some percentiles, mean, and median. This difference in the response time measured might be justified by the overhead introduced in the provider environment due to the increased concurrency generated by the workload. Still, as shown in the discussion related to Figure 5, the simulation model is highly powerful in simulating the response time distribution shape, proving the assumption used, which was essential to the scope of the previous work [30].

By using the predictive validation technique, we concluded that the simulation model used in the previous work [30], and the conclusions from this study, was valid for the scope of the work. Although we noticed differences in the response time values of percentiles, mean and median, these differences were not impactful enough to invalidate the assumption used and, therefore, can not invalidate the simulation model for the scoped it was used. Furthermore, the validation revealed that the simulation model is highly accurate in simulating the response time distribution shape. However, since this work focused only on the validation of the simulation model used in the Controlling Garbage Collection and Request Admission to Improve Performance of FaaS Applications [30], there are still open problems to tackle as future work to further validation of the simulation model on a generalist usage scope.

As listed in the previous work, we argue that for generic use of the simulation model we need to validate the model using a greater variety of representative functions. We also argue that a generic validation of the model requires a greater variety of platform platforms as reference, including public vendor such as Google Cloud Function [8], Microsoft Azure Function [6] and the already used AWS Lambda Function [4], and also include open-source platforms such as OpenFaaS [28], Fission [16], and Kubeless [21]. Finally, the complete validation of the simulation model requires a realistic workload, which was already provided in the Serverless in the Wild characterization study [36].

## ACKNOWLEDGMENTS

To Thiago Emmanuel and Daniel Fireman, the two who have strongly contributed to my academic education and the development of this work. To Universidade Federal de Campina Grande (UFCC) for providing the foundations and resources required for my education throughout the graduation period. To Laboratório de Sistemas Distribuídos (LSD), the laboratory that gave me the opportunity to work in the academic area and introduce myself in the research field. To VTEX, for providing a laptop and knowledge, through the VTEX Lab program, that allowed me to orchestrate all experiments and execute the analysis upon the results. Lastly, Amazon Web Services (AWS), which, by providing the AWS Free Tier program, makes easier and cost safe the execution of the measurement experiments used in this work.

## REFERENCES

- [1] 1979. Terminology for model credibility. 32 (3 1979), 103–104. <https://doi.org/10.1177/003754977903200304>
- [2] Istemi Ekin Akkus, Ruichuan Chen, Ivica Rimac, Manuel Stein, Klaus Satzke, Andre Beck, Paarijaat Aditya, and Volker Hilt. 2018. SAND: Towards High-Performance Serverless Computing. In *2018 USENIX Annual Technical Conference, USENIX ATC 2018, Boston, MA, USA, July 11-13, 2018*, Haryadi S. Gunawi and Benjamin Reed (Eds.). USENIX Association, 923–935. <https://www.usenix.org/conference/atc18/presentation/akkus>
- [3] Amazon Web Services (AWS). [n. d.]. Amazon CloudWatch: Observability of your AWS resources and applications on AWS and on-premises. <https://aws.amazon.com/cloudwatch/>. Online; Accessed: 2021-15-03.
- [4] Amazon Web Services (AWS). [n. d.]. AWS Lambda: Run code without thinking about servers or clusters. Only pay for what you use. <https://aws.amazon.com/lambda/>. Online; Accessed: 2021-15-03.
- [5] Amazon Web Services (AWS). [n. d.]. Managing concurrency for a Lambda function. <https://docs.aws.amazon.com/lambda/latest/dg/configuration-concurrency.html>. Online; Accessed: 2021-05-09.
- [6] Microsoft Azure. [n. d.]. Azure Functions: More than just event-driven serverless compute. <https://azure.microsoft.com/en-us/services/functions/>. Online; Accessed: 2021-15-03.
- [7] Stephen M. Blackburn, Kathryn S. McKinley, Robin Garner, Chris Hoffmann, Asjad M. Khan, Rotem Bentzur, Amer Diwan, Daniel Feinberg, Daniel Frampton, Samuel Z. Guyer, Martin Hirzel, Antony L. Hosking, Maria Jump, Han Lee, J. Eliot B. Moss, Aashish Phansalkar, Darko Stefanovic, Thomas VanDrunen, Daniel von Dincklage, and Ben Wiedermann. 2008. Wake Up and Smell the Coffee: Evaluation Methodology for the 21st Century. *Commun. ACM* 51, 8 (Aug. 2008), 83–89. <https://doi.org/10.1145/1378704.1378723>
- [8] Google Cloud. [n. d.]. Cloud Functions. <https://cloud.google.com/functions>. Online; Accessed: 2021-15-03.
- [9] Mathheus Nicolas da Silva and Marcus Carvalho. 2019. Análise de Mecanismos de Serverless Computing em Ambientes de Nuvens Computacionais. In *Anais Estendidos do XXXVII Simpósio Brasileiro de Redes de Computadores e Sistemas Distribuídos*. SBC, Porto Alegre, RS, Brasil, 225–232. [https://doi.org/10.5753/sbrc\\_estendido.2019.7791](https://doi.org/10.5753/sbrc_estendido.2019.7791)
- [10] Erwin Eyk, Alexandru Iosup, Johannes Grohmann, Simon Eismann, André Bauer, Laurens Versluis, Lucian Toader, Norbert Schmitt, Nikolas Herbst, and Cristina Abad. 2019. The SPEC-RG Reference Architecture for FaaS: From Microservices and Containers to Serverless Platforms. *IEEE Internet Computing PP* (11 2019), 1–1. <https://doi.org/10.1109/MIC.2019.2952061>

- [11] Phillip Farrington, Harriet Nembhard, D. Sturrock, G. Evans, and Robert Sargent. 2000. Validation And Verification Of Simulation Models. (02 2000).
- [12] D. Fireman, J. Brunet, R. Lopes, D. Quaresma, and T. E. Pereira. 2018. Improving Tail Latency of Stateful Cloud Services via GC Control and Load Shedding. In *2018 IEEE International Conference on Cloud Computing Technology and Science (Cloud-Com)*. IEEE Computer Society, 121–128. <https://doi.org/10.1109/CloudCom2018.2018.00034>
- [13] Daniel Fireman, Raquel Lopes, and João Arthur Brunet Monteiro. 2017. Using Load Shedding to Fight Tail-Latency on Runtime-Based Services. In *Brazilian Symposium on Computer Networks and Distributed Systems*.
- [14] Daniel Fireman and David Quaresma. 2020. GCI FaaS Sim, Handler. <https://github.com/dfquaresma/gci-faas-sim/blob/master/runtime/thumbfunc/src/main/java/com/openfaas/function/Handler.java>.
- [15] Daniel Fireman and David Quaresma. 2020. GCI Simulator. <https://github.com/gcinterceptor/gci-simulator/tree/master/serverlessgo>.
- [16] Fission. [n. d.]. Fission: Open source, Kubernetes-native Serverless Framework. <https://fission.io/>. Online; Accessed: 2021-15-03.
- [17] A Cloud Guru. [n. d.]. How long does AWS Lambda keep your idle functions around before a cold start? <https://acloudguru.com/blog/engineering/how-long-does-aws-lambda-keep-your-idle-functions-around-before-a-cold-start>. Online; Accessed: 2021-05-09.
- [18] Joe Hair, G. Tomas M. Hult, Christian Ringle, and Marko Sarstedt. 2016. *A Primer on Partial Least Squares Structural Equation Modeling (PLS-SEM)*, 2nd edition.
- [19] Vivian Hu. [n. d.]. Serverless Days 2020 Looks at Future of Serverless Architecture. <https://www.infoq.com/news/2020/07/future-serverless-architecture/>. Online; Accessed: 2021-05-12.
- [20] K. Kritikos and P. Skrzypek. 2018. A Review of Serverless Frameworks. In *2018 IEEE/ACM International Conference on Utility and Cloud Computing Companion (UCC Companion)*. 161–168.
- [21] Kubeless. [n. d.]. Kubeless: The Kubernetes Native Serverless Framework. <https://kubeless.io/>. Online; Accessed: 2021-15-03.
- [22] Manoj Kumar. 2019. Serverless Architectures Review, Future Trend and the Solutions to Open Problems. 6 (03 2019), 8. <https://doi.org/10.12691/ajse-6-1-1>
- [23] Hyungro Lee, Kumar Satyam, and Geoffrey Fox. 2018. Evaluation of Production Serverless Computing Environments. 442–450. <https://doi.org/10.1109/CLOUD.2018.00062>
- [24] Johannes Manner, Martin Endreß, Tobias Heckel, and Guido Wirtz. 2018. Cold Start Influencing Factors in Function as a Service. <https://doi.org/10.1109/UCC-Companion.2018.00054>
- [25] Dr. Bill McNeese. [n. d.]. Are the Skewness and Kurtosis Useful Statistics? <https://www.spforexcel.com/knowledge/basic-statistics/are-skewness-and-kurtosis-useful-statistics>. Online; Accessed: 2021-05-12.
- [26] MITRE. 2020. Verification and Validation of Simulation Models. <https://www.mitre.org/publications/systems-engineering-guide/se-lifecycle-building-blocks/other-se-lifecycle-building-blocks-articles/verification-and-validation-of-simulation-models>
- [27] Edward Oakes, Leon Yang, Dennis Zhou, Kevin Houck, Tyler Harter, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. 2018. SOCK: Rapid Task Provisioning with Serverless-Optimized Containers. In *2018 USENIX Annual Technical Conference, USENIX ATC 2018, Boston, MA, USA, July 11-13, 2018*, Haryadi S. Gunawi and Benjamin Reed (Eds.). USENIX Association, 57–70. <https://www.usenix.org/conference/atc18/presentation/oakes>
- [28] OpenFaaS. [n. d.]. OpenFaaS: Serverless Functions, Made Simple. <https://www.openfaas.com/>. Online; Accessed: 2021-15-03.
- [29] Andrei Palade, Aqeel Kazmi, and Siobhán Clarke. 2019. An Evaluation of Open Source Serverless Computing Frameworks Support at the Edge. <https://doi.org/10.1109/SERVICES.2019.00057>
- [30] David Quaresma, Daniel Fireman, and Thiago Emmanuel Pereira. 2020. Controlling Garbage Collection and Request Admission to Improve Performance of FaaS Applications. In *2020 IEEE 32nd International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*. 175–182. <https://doi.org/10.1109/SBAC-PAD49847.2020.00033>
- [31] Gladys Rama. [n. d.]. Report: AWS Lambda Popular Among Enterprises, Container Users. <https://awsinsider.net/articles/2020/02/04/aws-lambda-usage-profile.aspx>. Online; Accessed: 2021-05-12.
- [32] RedHat. [n. d.]. Cloud-Native Applications:What Is Function-As-A-Service (FaaS)? <https://www.redhat.com/en/topics/cloud-native-apps/what-is-faas>. Online; Accessed: 2021-05-07.
- [33] Stewart Robinson. 2000. Simulation Model Verification And Validation: Increasing The Users' Confidence. *Winter Simulation Conference Proceedings* (05 2000). <https://doi.org/10.1145/268437.268448>
- [34] Robert Sargent. 2009. Verification and validation of simulation models. *Journal of Simulation* 7, 162–176. <https://doi.org/10.1109/WSC.2009.5429327>
- [35] Narrative Science. 2020. How AWS Lambda Changed the Game of Multi-tenancy. <https://narrativescience.com/resource/blog/how-aws-lambda-changed-the-game-of-multi-tenancy/>
- [36] Mohammad Shahrad, Rodrigo Fonseca, Ínigo Gouri, Gohar Chaudhry, Paul Batum, Jason Cooke, Eduardo Laureano, Colby Tresness, Mark Russinovich, and Ricardo Bianchini. 2020. Serverless in the Wild: Characterizing and Optimizing the Serverless Workload at a Large Cloud Provider.
- [37] Paulo Silva, Daniel Fireman, and Thiago Silva. 2020. Prebaking Functions to Warm the Serverless Cold Start. <https://doi.org/10.1145/3423211.3425682>
- [38] Andy Yeung. [n. d.]. Top 4 Serverless Computing Platforms in 2021. <https://www.loginradius.com/blog/async/serverless-overview/>. Online; Accessed: 2021-05-08.