

# PerfSim: A Performance Simulator for Cloud Native Microservice Chains

Michel Gokan Khan, *Member, IEEE*, Javid Taheri, *Senior Member, IEEE*,  
Auday Al-Dulaimy, *Member, IEEE*, and Andreas Kassler, *Senior Member, IEEE*

**Abstract**—Cloud native computing paradigm allows microservice-based applications to take advantage of cloud infrastructure in a scalable, reusable, and interoperable way. However, in a cloud native system, the vast number of configuration parameters and highly granular resource allocation policies can significantly impact the performance and deployment cost. For understanding and analyzing these implications in an easy, quick, and cost-effective way, we present PerfSim, a discrete-event simulator for approximating and predicting the performance of cloud native service chains in user-defined scenarios. To this end, we proposed a systematic approach for modeling the performance of microservices endpoint functions by collecting and analyzing their performance and network traces. With a combination of the extracted models and user-defined scenarios, PerfSim can then simulate the performance behavior of all services over a given period and provide an approximation for system KPIs, such as requests' average response time. Using the processing power of a single laptop, we evaluated both simulation accuracy and speed of PerfSim in 104 prevalent scenarios and compared the simulation results with the identical deployment in a real Kubernetes cluster. We achieved ~81-99% simulation accuracy in approximating the average response time of incoming requests and ~16-1200 times speed-up factor for the simulation.

**Index Terms**—performance simulator, performance modeling, cloud native computing, service chains, simulation platform

## 1 INTRODUCTION

Cloud Native Computing is an emerging paradigm of distributed computing that “empower organizations to build and run scalable applications in modern, dynamic environments such as public, private, and hybrid cloud” [1]. One of the key purposes of introducing this paradigm was to answer the increasing need for mitigating the efforts of application-level clustering in the cloud and inline with the emergence of microservice architecture that promotes decoupling components of a software system into multiple independently manageable services, known as a microservice [2]. Since Google first introduced Kubernetes during the Google Developer Forum in 2014, as an approach for “decoupling of application containers from the details of the systems on which they run” [3], it becomes the de-facto enabler for utilizing microservice architecture based on technologies such as OS-level virtualization, better known as containers [4].

Amongst the main advantages of cloud native computing is the possibility for allocating highly granular resources to large-scale chains of services in a cluster. This additional granularity, while facilitating the scalability of service chaining, imposes complexity in resource allocation, traffic shaping and placement of containers in a cluster. Therefore, as microservices networks grow larger, the need for tools and techniques that shed light on service chaining implications upon system performance becomes critical. Furthermore, by the rising trend of cloud native computing in containerized cluster environments, many researchers are nurturing new

methods and schemes to implement new deployment [5] and performance optimization techniques [6] at various levels: starting from modern container scheduling [7]–[10] methods to predictive request load-balancing [11] and resource auto-scaling algorithms [12], [13].

Analyzing the performance behavior of a service chain in a real testbed gives the most reliable results. However, evaluating in real clusters is not always possible. In many cases, it might be too costly, notably time-consuming, and sometimes various skill sets are required to configure, run and manage the testbed. Moreover, in most performance optimization techniques, various scenarios need to be evaluated in a timely manner to eventually minimize a cost function. Performing such evaluations in a real testbed, while providing accurate results, imposes a dramatic burden for achieving a scalable and efficient optimization method.

Moreover, modern cloud native distributed systems have important performance affecting properties that earlier generation software systems such as monoliths (single-tiered software systems consisting of multiple tightly-coupled components) didn't have to cope with as much. Properties such as agile horizontal/vertical scaling, highly granular resource allocation, context-awareness, contention with other services, service chaining, and dynamic load-balancing between replicas.

To mitigate the aforementioned implications of using real testbeds for performance evaluation of cloud native microservice chains, we proposed *PerfSim*, a simulation platform that aims to approximately predict the performance of service chains under various placement, resource allocation, and traffic scenarios using profoundly limited resources of a laptop. We also proposed a systematic performance modeling approach to model the time-predictable endpoint functions of microservices using performance traces gen-

- M. Gokan Khan, J. Taheri, and A. Kassler are with the Department of Mathematics and Computer Science, Karlstad University, Sweden. E-mails: {michel.gokan,javid.taheri,andreas.kassler}@kau.se.
- A. Al-Dulaimy is with the School of Innovation, Design and Engineering, Mälardalen University, Sweden. E-mail: auday.aldulaimy@mdh.se

erated by the profiling and tracing tools such as *perf* and *ebpf* as well as distributed network tracing programs such as *jaeger* and *zipkin*. Using these models and a user-defined scenario, PerfSim can then simulate performance behavior of all service chains under a desired placement, resource allocation policy, network topology and traffic scenario.

Using profoundly limited resources of a laptop, we evaluated the simulation accuracy and speed of PerfSim under 104 prevalent scenarios by deploying and running them on a real Kubernetes cluster and comparing measured KPIs with the simulation results (i.e., average requests latency). We used *sfc-stress*, a synthetic service chain generation toolkit, for generating various service chains and microservice-based workloads. In our evaluation, we achieved ~81-99% accuracy in predicting the average latency of incoming requests and ~16-1200 times speed-up factor between the simulation time and actual execution time on a real cluster. With the same laptop, we also simulated a large service chain consisting of 100 microservices interconnected with 200 links over 100 hosts and showed that PerfSim can be effectively used for large-scale simulations.

To summarize, with PerfSim we contributed to the relatively new, but rich area of cloud native computing by enabling a fast, accurate and easy way for evaluating various user-defined policies in microservice-based applications.

## 2 RELATED WORKS

In this section, we briefly review existing works in the areas of (1) simulation platforms, (2) emulation tools and (3) analytical performance modeling approaches. Table 1 presents a comparison between the key properties of the most popular frameworks with PerfSim. In this table, besides common features, we also compared additional challenges imposed by simulating performance of microservices in cloud native environments.

### 2.1 Simulation tools

Cloud native systems have intricate provisioning and deployment requirements. Evaluating and predicting the performance of such services, studying the impact of provisioning policies, and correlating workload models with achievable performance are not straightforward due to the diversity and complexity of system interactions. Even though studying these interactions on real testbeds provides the most accurate results, several researchers and companies argue that computer simulation can be a powerful method to test multiple scenarios and evaluate various policies before enforcing them at different levels [14]–[17]. In recent years, several works aim to predict the services' performance using computer simulations to develop adequate resource policies and other decisions at different levels to meet the required *Quality of Service* (QoS) in a dynamic manner.

CloudSim [15] is one of the most popular simulators in this category that is designed to simulate cloud computing infrastructures. It allows modeling of various data centers, workload scheduling, and allocation policies. Using CloudSim in many scenarios can boost the development of innovative methods and algorithms in the cloud computing paradigm without the need for deploying applications in

the production environments. In recent years, several tools and modules had been designed based on CloudSim. For example, the iFogSim toolkit [18] inherits the features of CloudSim, and extends them by the ability of modeling IoT and Edge/Fog environments. The authors in [19] also built upon CloudSim to simulate the specifications of edge/fog computing and support the required functionalities. There exist several other CloudSim-based simulators for simulating various use cases in fog or edge environments [20], [21].

CloudSim and modules/plugins/tools based on CloudSim provide a sophisticated and straightforward way for researchers to simulate cloud/edge/fog computing infrastructure for modeling service brokers, data-centers, and scheduling policies. However, they are not designed with the purpose of simulating rigorous performance testings of microservices under extremely granular resource allocation and placement policies that exists in containerized services in cloud native applications. Moreover, to effectively simulate a set of service chains, we need to precisely specify the links and connections between each microservice within all chains, and then route the requests based on the user-defined traffic scenarios and network topology model. Something that the CloudSim category of simulators has not been designed to address.

In addition to CloudSim, there are other simulators in the context of cloud simulation. *Yet Another Fog Simulator* (YAFS) [22] is a discrete-event simulator based on *Simpy* that allows to simulate the impact of applications' deployment in edge/fog computing environments through customizable strategies. YAFS allows to model the relationships between applications, infrastructure configurations, and network topologies. It uses those relationships to predict network throughput and latency in dynamic and customized scenarios, such as path routing and service scheduling. Even though YAFS offers a novel approach towards simulating performance of large network topologies, it's not designed to simulate the microservices' performance within the context of complex service chains over a set of hosts in a cluster.

Apart from cloud/edge/fog simulators, there exist other simulators that focus on specific aspects of cloud computing paradigm, such as energy efficiency or network modeling and optimization. For example, GreenCloud [14], is a packet-level simulator for capturing energy footprint of data center components with the aim towards providing an environment for researchers to design more energy-aware data centers [23], [24]. Other examples are NS-3 [17], OMNet++ [25] and NetSim [26] that are primarily designed to simulate various types of networks and topologies. Although these simulators can smoothly simulate specific aspects of cloud systems that they are designed for (e.g., networking and energy efficiency), they cannot be used to simulate cloud native applications' performance; or at least can only be used to simulate specific aspects of service chains, such as their network performance or placement efficiency.

### 2.2 Performance emulators

Another popular experimentation approach in evaluating performance of cloud systems is emulation. Using emulation, users can analyze the system performance behavior supported by the available hardware in a more realistic way

CloudSim category [15]	Type	Target Environment	Granular Resource Allocation	Vertical Scaling	Horizontal Scaling	CPU Scheduling	Multi-threaded	Contention Endpoints	Fast Performance Modeling	PTE Independent/Low Cost	Service Chains Support	Multi-chain Support	Dynamic Control	Load-balancing	Congestion Control	Advanced Routing
CloudSim category [15]	S	Cloud		✓		✓		✓	✓		✓	✓	✓	✓		
YAFS [22]	S	Fog/Edge						✓	✓	✓	✓			✓	✓	
GreenCloud [14]	S	Cloud						✓	✓		✓	✓				
NS-3 [17]	S	Any						✓	✓	✓	✓	✓	✓	✓	✓	✓
OMNet++ [25]	S	Any						✓	✓	✓	✓	✓	✓	✓	✓	✓
NetSim [26]	S	Any						✓	✓	✓	✓	✓	✓	✓	✓	✓
Emulators [29] [33] [34]	E	Any	✓	✓	✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓
PerfSim	S	Cloud/Edge	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

S=Simulator,E=Emulator

Resources      Threads      Services      Network

TABLE 1: Comparison of performance characteristics in popular simulation and emulation frameworks with PerfSim

than using simulation [27]. Mininet [28], [29] is one of the widely used network emulators that can form a network of virtual switches, routers, controllers, links, and hosts. The hosts in Mininet are able to run Linux network software, while switches offer flexible custom routing and SDN. The bright reputation of Mininet inspired many researchers to extend Mininet in various ways. For example, MaxiNet [30] extended MiniNet to support spanning emulated networks over multiple hosts. Then EmuFog [31] extended MaxiNet to support fog computing deployment models. Another MaxiNet-based emulator, called Fogbed [32] enables the emulation of fog nodes as software containers under various network configurations to mimic edge/fog environments.

After Mininet introduced the idea of using containers for network and process emulation, other works also started to adopt a similar concept. For example, Dockemu [33] adopted both Docker for emulating network nodes and NS-3 for simulating the network traffic. NEaaS [34], a cloud-based network emulation platform, utilizes both Docker and virtual machines to emulate various networking scenarios.

However, the widespread popularity of such emulators stands in stark contrast to the rather trembling fact that they are bounded to the available computation power and bandwidth of the underlying hardware they are deployed upon, and consequently, they cannot be efficiently used to predict the resource utilization aspects of large-scale and complex cloud native applications. Moreover, emulation cannot dramatically improve the speed of evaluating various resource allocation or placement policies, and therefore can't be used for exhaustive policy testing.

### 2.3 Analytical performance modeling approaches

There is a third category of approaches towards modeling the performance of cloud native applications which based upon an analytical framework. Contrary to the simulation/emulation techniques, which employ a bottom-up approach in modeling and simulating various performance aspects of cloud applications, analytical methods adopt a top-down approach for that purpose by collecting and analyzing application KPIs via running various stress tests on the system or using historical performance measurements.

As an example, the work proposed in [35] aims to model the response time of microservices based on stress testing

and concurrently collecting performance traces for predefined intervals to learn a predictive auto-scaling model, such that the response time requirements are satisfied. Those performance traces are used to learn the resource provisioning policy model using a regression analysis approach. In [36], microservices are being tested individually using service-based sandboxing to construct the corresponding model, the analyzed data and performance model are presented to the user. Considering modeling throughput and latency, in our previous work [37], we used a combination of stress-testing and regression modeling approaches to understand the impact of microservices' resource configurations, model the correlation between the KPIs and resource configurations in a smaller *Performance Testing Environment* (PTE) to predict the performance in a larger *Production Environment* (PE).

Although the aforementioned methods benefit from accurate performance measurements on real systems and can be considered as the most widely used techniques in performance engineering in production environments, they suffer from three major limitations. Their first limitation is the high cost of test environment preparation that can accurately provide performance measurements. As for gathering accurate performance insights, the tests need to be performed on either a production environment or a PTE that mimics it; that imposes a high cost of deployment in either case. The second limitation is the slow stress-testing procedure as for getting accurate results, several hours of performance testing is required. Last but not least of their limitations is applicability in utilizing advanced optimization techniques, such as deep learning or meta-heuristics based optimization heuristics, which requires numerous re-deployments to learn, train, or find an optimal policy.

## 3 PERFSIM DESIGN AND IMPLEMENTATION

In this section, we introduce the system architecture and implementation details of PerfSim and mathematical notation used in this paper (summarized in Table 2).

### 3.1 Modeling elements of cloud native systems

PerfSim has models for various elements of cloud native systems and their underlying infrastructure.

As described in the previous sections, the joint problem of placement and resource allocation of a service chain

Hosts	$H = \{h_k\}_{k=1}^{ H }$	Set of hosts	$W^{r^H}$	Weights of resources	$h_k^{\text{cores}}, h_k^{\text{clock}}$	Cores count and CPU clock (hz) in $h_k$
	$r^H \in R^H$	Set of resource names	$\hat{V}_H(r^H, h_k)$	Initial $r^H$ capacity of $h_k$	$V_H(r^H, h_k)$	Current $r^H$ capacity of $h_k$
Network Topology	$P = \{\rho_z\}_{z=1}^{ P }$	Set of all routers	$L = {}^H L \cup {}^\tau L$	Set of all network links	$G(\tau) = (P, L)$	Graph represent. of network topology $\tau$
	${}^H L = \{{}^H l_o\}_{o=1}^{ {}^H L }$	Set of host $\leftrightarrow$ router links	$\rho_z^{\text{in bw}}, \rho_z^{\text{out bw}}$	In/out bw of $\rho_z$ (Bps)	$G(\tau)_{h_a, h_b}$	Network path between $h_a$ and $h_b$
	${}^\tau L = \{\tau l_o\}_{o=1}^{ {}^\tau L }$	Set of router $\leftrightarrow$ router links	$l_o^{\text{lat}}, \rho_z^{\text{lat}}$	Latencies of $l_o$ & $\rho_z$ (ns)	$\rho^{h_k} \in P$	The router connected to $h_k$
Services and Endpoints	$S = \{S_i\}_{i=1}^{ S }$	Set of all services	$r^S \in R^S$	Set of res. ctrl. params	$\Pi = (\rho_{i,j} \in \{0, 1\})$	Service replica placement matrix
	$S_i = \{\hat{s}_j^i\}_{j=1}^{ S_i }$	Set of replicas of $S_i$	$\hat{V}_S(r^S, S_i), V_S(r^S, \hat{s}_j^i)$	Initial/current $r^S$ cap. of $\hat{s}_j^i$	$\Pi(\hat{s}_j^i) \in H$	The host that $\hat{s}_j^i$ is currently placed
	$F_i = \{f_n^i\}_{n=1}^{ F_i }$	Set of endpnt. funcs of $S_i$	$A_{S_i}, \tilde{A}_{S_i}$	Affinities/anti-affinities	$H^{\hat{s}_j^i}, \psi_{\hat{s}_j^i}^{h_k}$	Eligible hosts and score of $h_k$ for $\hat{s}_j^i$
Service Chains and Traffic	$C = \{C_l\}_{l=1}^{ C }$	Set of service chains	${}^l F = \{f_n^l\}_{n=1}^{ {}^l F }$	Set of endpnt. funcs in $C_l$	${}^l U = \{u_o^l\}_{o=1}^{ {}^l U }$	Set of user requests in $C_l$
	${}^l S = \{S_i\}_{i=1}^{ {}^l S } \subseteq S$	Set of services in $C_l$	${}^l F_i = \{f_n^i\}_{n=1}^{ {}^l F_i } \subseteq {}^l F$	Set of endpnt. funcs of ${}^l S_i$	$C_l^{\text{rate}}, C_l^{\text{duration}}, C_l^{\text{batch}}$	Arrival rate/duration & batch size of $C_l$
	${}^l E = \{e_x\}_{x=1}^{ {}^l E }$	Ordered set of links in $C_l$	$G(C_l) = ({}^l F, {}^l E)$	Graph represent. of $C_l$	${}^l u_o^{\text{in time}}, {}^l u_o^{\text{exe time}}$	Req. arrival & exe time of ${}^l u_o$
Threads	${}^l f_n = \{t_m\}_{m=1}^{ {}^l f_n }$	Set of threads in ${}^l f_n$	$t_m^{\text{maccess}}$	Mem. accesses of $t_m$	$t_m^{\text{blk rw}}$	blkio R/W of $t_m$ (in bytes)
	${}^l f_n^i = \{t_m\}_{m \in {}^l f_n^i}$	Set of threads in ${}^l f_n^i \subseteq {}^l f_n$	$t_m^{\text{crefs}}, t_m^{\text{cmisss}}$	Cache refs & misses of $t_m$	$t_m^{\text{idle}}$	Total idle time of $t_m$
	$t_m^{\text{inst}}, t_m^{\text{CPI}}$	Inst. count & CPI of $t_m$	$t_m^{\text{cpenalty}}$	Avg. lost cycles per miss	$\hat{f}_{h_k}$	Set of running threads on $h_k$

TABLE 2: Table of the base notations used in this paper

over a cluster is complex as multiple dimensions impact the achievable performance and latency. However, several parameters have only marginal impact on the prediction quality. Consequently, we aim at providing a simplified model that is tractable but powerful enough in order to predict the performance of typical cloud native service chain architectures when deployed on several well-known container orchestration platforms such as Kubernetes.

### 3.1.1 Hosts

An  $h_k \in H = \{h_k\}_{k=1}^{|H|}$  is a model of Linux-based physical machine that has  $h_k^{\text{cores}}$  number of CPU cores with  $h_k^{\text{clock}}$  clock speed (in Hertz). It has also a set of consumable resources, such as a memory capacity (in bytes), a NIC with a limited ingress and egress network bandwidth (in bytes per second), a local storage with a limited storage read/write bandwidth and storage capacity. The CPU resource is measured in millicores. Each host  $h_k$  introspects the OS to determine the  $h_k^{\text{cores}}$  and then multiplies it by 1000 to denote its total capacity (Eq. 1).

$$h_k^{\text{millicores}} = h_k^{\text{cores}} \times 1000 \quad (1)$$

To facilitate our formulation, we denote hosts consumable resources as  $R^H = \{\text{millicores}, \text{mem}, \text{in bw}, \text{out bw}, \text{blkio bw}, \text{blkio size}\}$ . For each resource  $r^H \in R^H$ , a host  $h_k$  has an initial resource capacity denoted as  $\hat{V}_H(r^H, h_k)$  as well as a current capacity denoted as  $V_H(r^H, h_k)$ .

### 3.1.2 Network topology

A network topology, denoted as  $\tau$ , has a directed acyclic graph representation of  $G(\tau) = (P, {}^\tau L)$  with routers  $\rho_z \in P$  as nodes, and directed links  $\tau l_o = (\rho_i, \rho_j) \in {}^\tau L$  between them as edges. To separately model egress and ingress bandwidth of interconnection between nodes, we hypothetically assumed there are always 2 links between each pair of routers with opposite directions. We denote the maximum ingress/egress bandwidth of a router as  $\rho_z^{\text{in bw}}$  and  $\rho_z^{\text{out bw}}$ .

Each host  $h_k$  is connected to a router  $\rho^{h_k} \in P = \{\rho_z\}_{z=1}^{|P|}$  using a hypothetically separated egress/ingress directed links, denoted as  ${}^H l_o \in {}^H L$  that have a pair of ordered host $\rightarrow$ router ( $h_k, \rho_i$ ) and router $\rightarrow$ host ( $\rho_i, h_k$ ) links.

Each router  $\rho_z$  and link  $l_o \in L$ , may respectively add extra latency of  $\rho_z^{\text{lat}}$  and  $l_o^{\text{lat}}$  nanoseconds to each request. These additional latencies differs from the delay caused by network congestion which we modelled separately.

### 3.1.3 Services

We assume that each chain of services is composed of a set of  $|S|$  containerized services  $S = \{S_i\}_{i=1}^{|S|}$  such that each  $S_i \in S$  may have  $|S_i|$  number of single-process/multi-threaded replicas  $\{\hat{s}_j^i\}_{j=1}^{|S_i|}$  load balanced among  $|H|$  hosts and connected through  $|P|$  routers. A given service  $S_i$  has a set of endpoint functions  $F_i = \{f_n^i\}_{n=1}^{|F_i|}$  that can be executed based on the type of incoming request. Each endpoint function  $f_n^i$  may spawns a set of threads  $t_m \in f_n^i$ . We will explain properties of these threads in section 3.1.10.

### 3.1.4 CPU scheduler and resource controller

A service  $S_i$  may optionally have resource constraints on its containerized replicas. We denote this set of resources as  $R^S = \{\text{CPU requests}, \text{mem requests}, \text{in bw}, \text{out bw}, \text{blkio bw}, \text{blkio size}\}$ , which corresponds to resources  $R^H$  in each host ( $R : R^S \rightarrow R^H$ ). By default, there is no reservation or usage limitation on any resource.

In Linux kernel, a service replica is in a form of a container and is being managed by *Control Groups* (cgroups) [38] that is responsible for auditing and restricting a set of processes. The limits associated with cgroups isolate the resource usage of a collection of processes.

We partially modeled the behaviour of cgroups in PerfSim by allowing a service  $S_i$  to initially reserve  $\hat{V}_S(r^S, S_i)$  resource capacities for each of its replicas. When reserving CPU requests  $\hat{V}_S(\text{CPU requests}, S_i)$ , PerfSim allows



a replica to use as much as CPU units available, even more than its reserved quota. We also partially modeled the behaviour of CPU bandwidth controller [39] of Linux by allowing to define  $S_i^{\text{CPU limits}}$  (in millicores) as an upper bound for replicas' CPU consumption. In kernel level, CPU request and limit is translated to cgroups `cpu.shares`, `cpu.cfs_quota_us` (in microseconds) and `cpu.cfs_period_us` (in microseconds) control parameters. We denote them as  $S_i^{\text{CPU share}}$ ,  $S_i^{\text{CPU quota}}$ , and  $S_i^{\text{CPU period}}$ . However, for simplicity, we assume a fixed value for  $S_i^{\text{CPU period}}=100\text{ms}$  in our calculations and only considered CPU share and quota when estimating available resources for each replica. A usage example for this model is simulating Kubernetes *Best Effort* and *Guaranteed QoS* classes for CPU resources which we will cover in section 5.2.

The  $\hat{V}_S(\text{mem requests}, S_i)$  specifies the initial memory allocation of replicas of  $S_i$  (in bytes). We modelled the `memory.limit_in_bytes` control parameter in cgroups in a way that it may only affect the placement of replicas.

To control the traffic shaping and the maximum network bandwidth of replicas in  $S_i$ , the  $\hat{V}_S(\text{in bw}, S_i)$ , and  $\hat{V}_S(\text{out bw}, S_i)$  control parameters can be used (in Bps).

The storage size of replicas in PerfSim is controlled by a model of `blkio` cgroups controller and denoted as  $\hat{V}_S(\text{blkio size}, S_i)$ . We also denote the available bandwidth for blkio reads and write as  $\hat{V}_S(\text{blkio bw}, S_i)$ . We intended to model and simulate the disk bandwidth throttling feature of cgroups `blkio.throttle.write_bps_device`.

### 3.1.5 Affinity controller

A service may optionally have a set of in-service affinity/anti-affinity rules for its placement. We define two binary decision vectors  $A_{S_i}$  and  $\tilde{A}_{S_i}$  that are respectively holding all the affinity and anti-affinity rules related to service  $S_i$ . Each vector has a dimension of  $|S| \times 1$  where each row represents one of the services. An entry  $A_{S_i}[S_j]$  takes on the value of '1' if replicas of services  $S_i$  and  $S_j$  must be placed at the same host and '0' otherwise; and similarly an entry  $\tilde{A}_{S_i}[S_j]$  takes on a value of '1' if any replicas of both services must not be placed at the same host and '0' otherwise. It's also possible to define host affinities, but due to the page limitation we omitted the formal definition.

### 3.1.6 Placement algorithm

The placement of each  $\hat{s}_j^i$  among all hosts, takes place using a placement algorithm. By default, we implemented the *Least Allocated* bin packing algorithm which partially simulates the container scheduling mechanism of Kubernetes. This algorithm, favor hosts with fewer resource requests and is controlled with weights for each resource  $r^H \in R^H$ , denoted as  $W^{r^H}$  to be used for scoring nodes. We described the details of this algorithm in the next section.

### 3.1.7 Container scheduler

We use a binary matrix  $\Pi$  of size  $(\sum_{i=1}^{|S|} |S_i|) \times |H|$  to keep track of service replicas placement. In this matrix,  $p_{j,k} \in \Pi = 1$  if the replica  $\hat{s}_j \in \{S_i\}_{i \in S}$  is placed in host  $h_k$  and  $p_{j,k}=0$  if otherwise. Since each replica can only be placed on one host,  $\Pi(\hat{s}_j^i) \in H$  indicates the host that replica  $\hat{s}_j^i$  is placed. After placing each replica, the amount of available resources in each host  $V_H(r^H, h_k)$  are recalculated for all resources.

### 3.1.8 Service chains

The goal of services is to serve requesters through a set of service chains  $C = \{C_l\}_{l=1}^{|C|}$ . These service chains represent the connection and traffic flow between endpoint functions of services and are specified as a directed flow network  $G(C_l)=(^lF, ^lE)$  where  $^lF$  is the subset of endpoint functions which have a role in providing service chain  $C_l$ , and  $^lE$  is the subset of all virtual links belonging to the service chain  $C_l$ . In other words, all endpoint functions within  $C_l$  are connected through  $^lE$  number of ordered directed virtual links  $^lE=(^le_1=(^lf_a, ^lf_a'), \dots, ^le_{|^lE|}=(^lf_c, ^lf_c'))$ . Each  $^le_v \in ^lE$  has a request payload  $^le_v^{\text{payload}}$  in bytes, which indicates the request size flowing from edge's first vertex to the second.

The first node in a  $C_l$  is its only request entry point, known as the *source* node, that receives requests from potential request senders and the cluster load-balancer directs the traffic to a replica  $\hat{s}_j^i$  based on a desired load-balancing algorithm (default is *Round-robin*). A service chain has also one or more *sink* nodes with no immediate outgoing edges.

### 3.1.9 Traffic control

Assuming a service chain  $C_l$  is active for  $C_l^{\text{duration}}$  seconds. The traffic rate of each  $C_l$  is being controlled with 2 control parameter  $C_l^{\text{rate}}$  and  $C_l^{\text{batch}}$ ; every  $\frac{1}{C_l^{\text{rate}}}$  seconds, a batch of  $C_l^{\text{batch}}$  requests will arrive to service chain  $C_l$ . In other words, each service chain  $C_l$  may be requested by a plurality of  $C_l^{\text{batch}}$  users with an average arrival rate of  $C_l^{\text{rate}} \frac{\text{req}}{\text{s}}$ . Therefore, the number of requests  $^lU$  for service chain  $C_l$  over a period of time  $C_l^{\text{duration}}$  will be equal to  $C_l^{\text{duration}} \times C_l^{\text{rate}} \times C_l^{\text{batch}}$ . We denote set of all user requests of  $C_l$  as  $^lU = \{^lu_o\}_{o=1}^{|^lU|}$  and their arrival time and execution time are denoted by  $^lu_o^{\text{in time}}$  and  $^lu_o^{\text{exe time}}$ , respectively.

### 3.1.10 Service endpoint functions and associated threads

During the life cycle of the  $C_l$ , each endpoint function  $^lf_n \in ^lF$  spawns  $^lf_n$  number of threads  $\{t_m\}_{m=1}^{|^lf_n|}$  on one of its replicas. Since each replica is placed on a host, we denote a subset of all threads running on a host  $h_k$  as  $\hat{t}_{h_k}$ .

Each thread  $t_m \in ^lf_n$  has a set of properties that specified in the original model and extracted based on measurements of the performance traces and monitoring tools during the modeling phase which we will explain in section 3.2.

A  $t_m$  executes  $t_m^{\text{inst}}$  CPU instructions in which  $t_m^{\text{maccs}}$  number of them are due to memory accesses with  $t_m^{\text{CPI}}$  average CPU Cycle per Instruction (CPI). Moreover,  $t_m$  has total of  $t_m^{\text{crefs}}$  cache references in which, when  $\hat{s}_j^i$  is deployed on an isolated and single core reference machine,  $t_m^{\text{cmis}}$  number of those references will be cache misses with average miss penalty of  $t_m^{\text{cpenalty}}$  cycles. During the simulation, PerfSim dynamically re-calculates cache misses based on various factors such as co-located active threads, number of cores, CPU requests, and cache size.

A thread may also read/write a total of  $t_m^{\text{blk rw}}$  bytes of data from/to the storage device. Moreover, a thread might be idle for some time during its active life-cycle and we denote  $t_m$  accumulated period of the idle time as  $t_m^{\text{idle}}$ .

On a real Linux machine, threads are load balanced among all CPU runqueues in the host using the Linux's

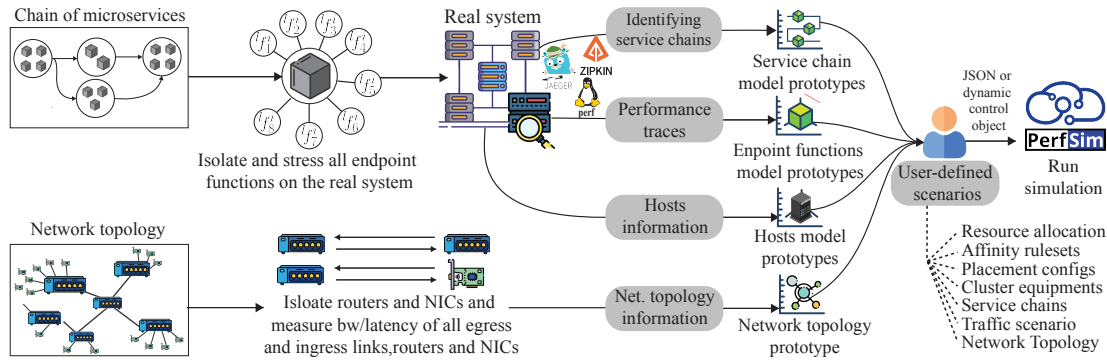


Fig. 1: An overview of our proposed systematic method for modeling various service chains and microservices with their corresponding endpoint functions as well as underlying network topologies encapsulated in user-defined scenarios

*Completely Fair Scheduler (CFS)*. We implemented a simplified version of CFS in PerfSim to imitate the impact of thread scheduling mechanism on the performance. We described the details of our implementation in section 5.

### 3.2 Modeling process

In order to accurately simulate the performance of service chains in a given scenario, PerfSim requires to have performance models of each and every active endpoint function within running services in the chain. To introduce such performance models in PerfSim, the user can either provide a pre-defined model file, or run a new modelling procedure. In the latter case, we proposed a systematic performance testing and modeling approach to extract the KPIs of endpoint functions, construct aforementioned modeling elements, and identify each equipment’s resource capacities in the simulation. Once such performance models being extracted for an endpoint function, the models can be re-used for simulating various types of user-defined scenarios without any need for a new profiling phase.

As presented in Figure 1, the modeling process starts by placing microservices of all service chains on a set of reference hosts in the *Performance Testing Environment (PTE)*. Then, one by one the microservices get isolated on a single host to start stressing each of its endpoint functions by flowing a single request through the function, and in parallel, an automated script captures various performance traces and resource utilization measurements and stores the correlated traces it in a separate database.

While running the network traffic, various network tracing tools would be used to identify the connections between microservices and identify service chains within the cloud native system. For example, a *service mesh* layer based on *istio* [40], enables distributed tracing capabilities through *Envoy* [41] which “allows developers to obtain visualizations of call flows in large service oriented architectures” by taking advantage of tracing tools such as *LightStep* [42], *Zipkin* [43], and *Jaeger* [44]. Leveraging a service mesh layer allows PerfSim to extract service chain flow models which facilitates the modeling procedure. However, these models can also be provided manually via PerfSim’s JSON model files or dynamic control objects.

After completing the profiling phase for all endpoint functions, a similar procedure can optionally be repeated

with all the target network equipments, such as routers and NICs, using network performance measurement tools such as *iperf*. The intention is to separately measure available egress/ingress bandwidths of each network equipment and links latency to effectively simulate the network congestion.

Finally, based on these extracted models, various user-defined scenarios can be specified by the user. These scenarios includes resource allocation policies, affinity rulesets, placement configurations and algorithm, equipments used in a cluster (i.e., active hosts and router), service chains, incoming traffic scenarios, and underlying network topology. To start simulating the given scenario, all the aforementioned models and simulation parameters can be fed into PerfSim using a JSON model file or a Python script (i.e. dynamic control object). We will cover the details of defining a scenario in section 3.3.1.

### 3.3 System Architecture

Figure 2 presents the layered architecture of PerfSim. We employed a multi-tier and object-oriented approach to architect PerfSim: (1) The presentation layer (client), (2) business layer (PerfSim’s core), and (3) data layer (results db).

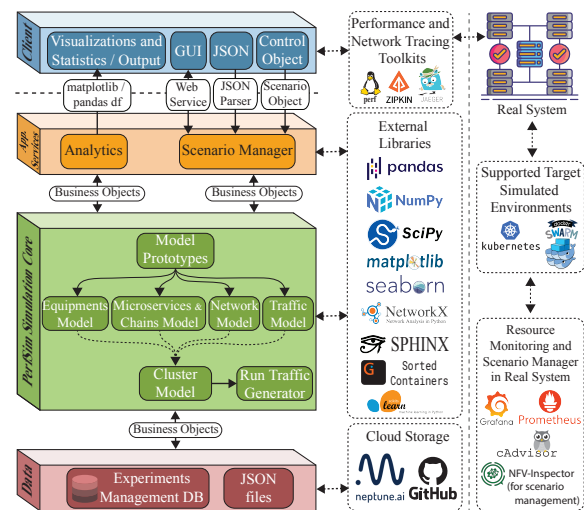


Fig. 2: Software architecture of PerfSim

### 3.3.1 Presentation layer: defining scenarios

To retrieve simulated performance insights using PerfSim, a scenario should be specified for the simulator.

In the presentation layer, PerfSim user defines the simulation scenarios and/or provides the control object. The user can provide the scenario in 3 general ways:

- Using a JSON file for static simulations
- Using a `perfsim.Cluster` object for dynamic simulations
- Using the GUI for quick experimentation

A scenario's structure in PerfSim is as follows:

- 1) **Prototypes** define the model of all microservices, end-point function, threads, hosts, routers, links and traffics.

```
"prototypes": {
  "microservices": { {∀Ci ∈ C ∨ Si ∈ S ∨ fni ∈ Fi →
    "fni": {∀tm ∈ fni →
      {tminst, tmcpu, tmmacs, tmcrefs, tmcmis, tmcpenalty, tmblk rw}}},
    "hosts": {∀ host types →
      "{type_name}": {hcores, hclock, {VH(rH, *)|rH ∈ RH}}},
    "routers": {∀ router types →
      "{type_name}": {ρlat, ρin bw, ρout bw}},
    "links": {∀ link types →
      "{type_name}": {llat}},
    "traffics": {∀ traffic types →
      "{type_name}": {Crate, Cduration, Cbatch}}
```

Listing 1: Prototypes JSON object in the scenario file

- 2) **Equipments** define the hosts and routers that are going to be used in the cluster scenarios. These equipments should be created based on the defined "prototypes".

```
"equipments": {
  "hosts": {∀hk ∈ H → "hk": "{host_type_name}"},
  "routers": {∀ρz ∈ P → "ρz": "{router_type_name}"}
```

Listing 2: The JSON object defining hosts and routers

- 3) **Topologies** define the network topologies that are going to be used in the cluster scenarios.

```
"topologies": {∀ topologies and links →
  "{topology_name}": {"nodes": {P ∪ H}, "edges": {L}}}
```

Listing 3: Defining network topologies and hosts links

- 4) **Service chains** define the chain of microservices as a graph (nodes and edges) including all link payloads.

```
"sfcs": {∀Ci ∈ C → "Ci": {"nodes": {F}, "edges": {E}}}
```

Listing 4: Defining service chains graph  $G(C_i)$

- 5) **Resource allocation scenarios** consists of resource policy templates that later on can be assigned to each replica in a cluster scenario.

```
"res_alloc_scenarios": {∀ res. alloc. scenario →
  "{res_scenario_name}": {∀rS ∈ RS → "rS": {VS}(rS, *)}}
```

Listing 5: Defining resource allocation scenarios

- 6) **Placement scenarios** define the algorithm, weights and configuration of the container scheduler. By default, PerfSim uses the *Least Allocated* placement algorithm.

```
"placement_scenarios": {∀ placement scenario →
  "{placement_scenario_name}": {
    "algorithm": "{algorithm_name}",
    "options": {e.g., ∀rH ∈ RH → "rH": WrH}}
```

Listing 6: Placement algorithms and related options

- 7) **Affinity rule-sets** define the affinity/anti-affinity rule-set scenarios that can be used in the cluster scenario.

```
"affinity_rulesets": {∀ affinity ruleset →
  "{affinity_set_name}": {
    "affinity": [∀Si ∈ S → ASi],
    "anti-affinity": [∀Si ∈ S → A̅Si]}}
```

Listing 7: Rulesets containing affinity/anti-affinity rules

- 8) **Cluster scenario** defines ultimate cluster scenarios and consists of a combination of all aforementioned sections, as well as additional configuration to define number of replicas for each service and network timeouts.

```
"cluster_scenarios": {∀ cluster scenario →
  "{scenario_name}": {
    "service_chains": {∀Ci ∈ C →
      "Ci": {
        "traffic_type": "{traffic_type_name}",
        "nodes_settings": {∀Si ∈ S →
          "Si": {
            "replica_count": |Si|,
            "res_scenario": "{res_scenario_name}"}}}},
    "placement_scenario": "{placement_scenario_name}",
    "topology": "{topology_name}"}}
```

Listing 8: The ultimate simulation scenarios object

Similar to static JSON-based scenarios, a user can define the ScenarioManager object in pure Python and provide the control object to PerfSim to start the simulation. As shown in Listing 1, since the JSON-based scenario is human-readable, a PerfSim user can easily modify scenarios and/or add new ones as needed. Moreover, a user can define and simulate various types of scenarios without any need for performance trace logs or profiling data. The aforementioned scenarios can be provided based on either extracted performance models mentioned in the previous section or based on user-defined models. A PerfSim user can then retrieve simulated insights of the final performance of the system in the given scenario.

### 3.3.2 Business and data layers

In the business layer, PerfSim creates all necessary objects to perform the simulation. The details of modeling and simulation in PerfSim is covered in the next section.

After performing the simulation, the extracting performance simulation results will be saved in a database. The user can choose between saving the result in (1) a JSON file or (2) store it in a database (e.g., neptune.ai or MySQL).

## 4 PLACEMENT, CHAINING AND ROUTING OF SERVICES OVER THE CLUSTER

In PerfSim, the first step towards approximating the performance of service chains in a user-defined scenario is to place replicas on a cluster of hosts. After the placement, PerfSim simulates the traffic scenario by first processing the defined service chain flow graphs, identifying parallel subchains, and extracting the exact route between service replicas based on the given network topology  $\tau$ .

### 4.1 Placement of service replicas on the cluster

By default, we implemented a simplified version of the Least Allocated bin packing algorithm for the placement of

**Algorithm 1:** Simplified placement of service replicas among several hosts in a cluster

```

1 function filter_hosts( $\hat{s}_j^i$ )
2    $A_S \leftarrow x$  in  $A_{S_i}[x]$  if  $A_{S_i}[x]=1$ 
3    $\tilde{A}_S \leftarrow x$  in  $\tilde{A}_{S_i}[x]$  if  $\tilde{A}_{S_i}[x]=1$ 
4    $H^{\hat{s}_j^i} \leftarrow \emptyset$  if  $A_S = \emptyset$  else  $H$ 
5   foreach  $S' \in A_S$ 
6     foreach  $s' \in S'$ 
7        $H^{\hat{s}_j^i} \cdot \text{push}(\Pi(s'))$ 
8   foreach  $\bar{S}' \in \tilde{A}_S$ 
9     foreach  $\bar{s}' \in \bar{S}'$ 
10       $H^{\hat{s}_j^i} \cdot \text{pop}(\Pi(\bar{s}'))$ 
11  foreach  $h_k \in H'$ 
12    foreach  $(r^S, r^H) \in R : R^S \rightarrow R^H$ 
13      if  $V_S(r^S, \hat{s}_j^i) > V_H(r^H, h_k)$ 
14         $H^{\hat{s}_j^i} \cdot \text{pop}(h_k)$ 
15  return  $H^{\hat{s}_j^i}$ 
16 function score( $\hat{s}_j^i, h_k$ )
17    $\psi^{\min} \leftarrow 0, \psi^{\max} \leftarrow 100, \psi_{\hat{s}_j^i}^{h_k} \leftarrow$  based on Eq. 4
18  return  $\psi_{\hat{s}_j^i}^{h_k}$ 
19 function place( $\hat{s}_j^i, h_k$ )
20  update_values( $\Pi, V_H, V_S$ )
21 function schedule()
22  foreach  $\hat{s}_j^i \in Q$ 
23     $H^{\hat{s}_j^i} \leftarrow \text{filter\_hosts}(\hat{s}_j^i)$ 
24    if  $H^{\hat{s}_j^i} = \emptyset$ 
25       $Q \cdot \text{pop}(\hat{s}_j^i) \cdot \text{push}(\hat{s}_j^i)$ 
26     $h^* \leftarrow \text{first } h \in H^{\hat{s}_j^i}, \bar{\psi}_{\hat{s}_j^i}^{h_k} \leftarrow \text{score}(\hat{s}_j^i, h^*)$ 
27    foreach  $h_k \in H^{\hat{s}_j^i}$ 
28      if  $\bar{\psi}_{\hat{s}_j^i}^{h_k} > \text{score}(\hat{s}_j^i, h_k)$ 
29         $h_{\hat{s}_j^i}^* \leftarrow 1^{\text{st}} h \in H^{\hat{s}_j^i}, \bar{\psi}_{\hat{s}_j^i}^{h_k} \leftarrow \text{score}(\hat{s}_j^i, h_{\hat{s}_j^i}^*)$ 
30    place( $\hat{s}_j^i, h_{\hat{s}_j^i}^*$ )
31     $Q \cdot \text{pop}(S_j^i)$ 

```

service replicas, partially imitating the Kubernetes scheduler's placement strategy. The key idea behind this strategy is to ensure that replicas are placed on hosts with adequate resources and balance out the resource utilization of hosts. It consists of four steps: (1) enqueueing replicas, (2) filtering hosts, (3) scoring hosts, and (4) placement of the replica on the host. All these steps are managed by a *schedule* procedure. Algorithm 1 represents each of these steps and Figure 3 represents an example of such final placement.

**4.1.1 Enqueueing replicas**

Replicas are being placed one after the other. When a service requests for scheduling a replica on the cluster, scheduler adds that request to a queue (denoted as  $Q$  in Algorithm 1) and then attempts to find a suitable host for placing it.

$$Q = \{\hat{s}_j^i | \hat{s}_j^i \in S_i \in S\} \quad (2)$$

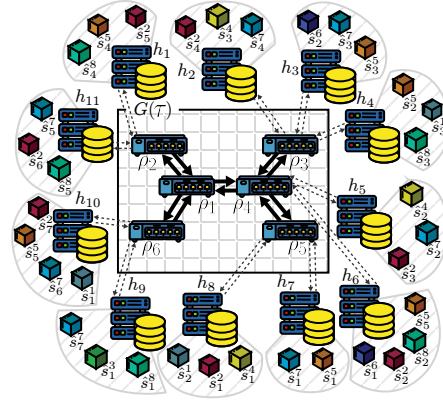


Fig. 3: An example placement of replicas over multiple hosts connected to a network with topology  $\tau$

**4.1.2 Filtering hosts**

Considering all affinity/anti-affinity rules as well as resource constraints of the service replica, the scheduler needs to filter out all the hosts that matches the specified limitations by the service. The filtered nodes, denoted as  $H^{\hat{s}_j^i}$ , are potentially eligible to host the replica  $\hat{s}_j^i \in S_i$ .

$$H^{\hat{s}_j^i} = \{h_k | \text{filter\_hosts}(\hat{s}_j^i)\} \quad (3)$$

**4.1.3 Scoring hosts**

In the Least Allocated bin packing strategy, in order to fairly distribute a replica  $\hat{s}_j^i$  among several hosts in a cluster, all eligible hosts in the  $H^{\hat{s}_j^i}$  are scored based on the request to capacity ratio of primary resources  $r^S \in R^S$  and  $r^H \in R^H$  considering the weight of each resource denoted as  $W^{r^H}$ . The score  $\psi_{\hat{s}_j^i}^{h_k}$  for each host  $h_k$  related to a replica  $\hat{s}_j^i$  are between  $\psi^{\min}$  and  $\psi^{\max}$  are calculated as follows:

$$\psi_{\hat{s}_j^i}^{h_k} = \frac{\sum_{(r^S, r^H) \in R} \left(1 - \left(\frac{V_H(r^H, h_k) - V_S(r^S, \hat{s}_j^i)}{\bar{V}_H(r^H, h_k)}\right)\right) W^{r^H}}{\sum_{r^H \in R^H} W^{r^H}} \quad (4)$$

**4.1.4 Placement of replicas**

After calculating the score of all eligible nodes for hosting replica  $\hat{s}_j^i$ , the host  $h_{\hat{s}_j^i}^*$  with the lowest score will be selected:

$$h_{\hat{s}_j^i}^* = \left\{ h^* \in H^{\hat{s}_j^i} | \psi_{\hat{s}_j^i}^{h^*} = \max\{\psi_{\hat{s}_j^i}^{h_k} | \forall h_k \in H^{\hat{s}_j^i}\} \right\} \quad (5)$$

**4.2 Service chain flow graphs**

A network flow graph  $G(C_l) = (S, E)$  of service chain  $C_l$  is a directed ordered graph that can have any form, from a simple tree-like service chain to a complex multi-degree cyclic multigraph. To efficiently simulate these service chains, we should be able to estimate the processing time of each service replica  $l_s^i \in l_{S_i} \in l_S \subseteq S$ , and then considering the outgoing request payload and network congestion, estimate the network transfer time between two services  $(l_{S_i}, l_{S_i'}) \in l_E$ . To accomplish these tasks, we should first identify the exact execution order of service requests.



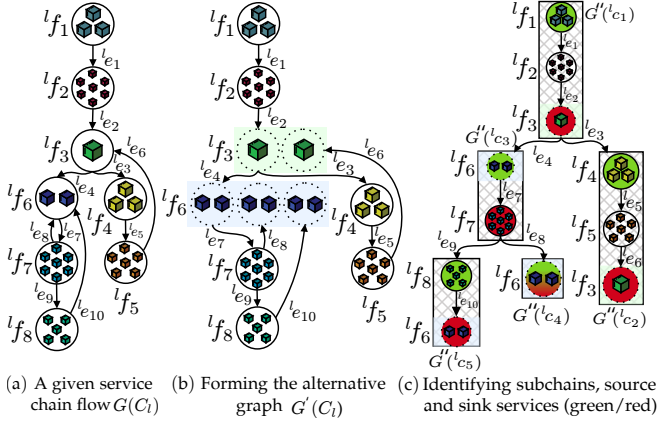


Fig. 4: Forming alternative graph  $G'$  and flow graph  $G''$

In any given service chain  $C_l$ , there might be subchains that run in parallel and their execution order are entirely depends on their execution time; this in turn depends on a vast range of parameters, starting from the processing power of the host to allocated resources to each service.

To identify all subchains within a service chain flow graph  $G(C_l)$ , we first form an alternative graph  $G'(C_l) = (S', E', F)$  from the original graph  $G(C_l)$  in which every service  $l_s^i \in S'$  is visited only once. We form  $G'(C_l)$  by duplicating services in  $S_i \in S$  that are visited more than once; in other words, duplicate service nodes where has indegree of  $\delta^-(l_s^i) \geq 2$ , and rerouting connecting edges to newly generated nodes. Then, based on newly formed  $G'(C_l)$ , we now identify all  $n^{G'(C_l)}$  subchains in  $C_l$ , denoted as  $\{l_{c_x}\}_{x=1}^{n^{G'(C_l)}}$ . Each subchain  $l_{c_x}$  has a flow subgraph  $G''(l_{c_x}) = (S'', E'', F'')$  that has a source service, which is the first service of the subgraph, and a sink service, which is the last service of the subgraph. In case a  $G''(l_{c_x})$  has only one service, that single service will be both source and sink node. We form each subgraph  $G''(l_{c_x})$  of the subchain  $l_{c_x}$  by first identifying its source service. Starting from the the very first source service  $l_s^i \in S'$  that is initiating first subchain  $l_{c_1}$ , when a service  $l_s^i \in S'$  has an outdegree of  $\delta^+(l_s^i) \geq 2$ , which means it has more than one outgoing edges  $\{e_y^i = (l_s^i, l_s^x)\}_{i \in E'}$ , then all immediate services  $l_s^x$  will be marked as source services and will initiate a new subchain. A subchain ends where we meet a service node with an outdegree of  $\delta^+(l_s^y) \neq 1$ . A request  $l_{u_o}$  ends when all sink service nodes conclude their executions.

As an example, Figure 4 demonstrates all aforementioned steps. In subfigure (a), we see an example flow graph  $G(C_l)$  of service chain  $C_l$  with all its services  $l_s$  (circles) and their containing service replicas (cubes) together with all edges  $l_e$  connecting services together (arrows) with the source service node  $S_1^l$ . In this graph, we see node  $S_3^l$  initiate two parallel subchains and has a indegree  $\delta^-(S_3^l) = 2$ . Also,  $S_6^l$  forms a cycle with  $S_7^l$  with indegree  $\delta^-(S_6^l) = 3$ . In the second subgraph (b), we form the alternative graph  $G'(C_l)$  by duplicating  $S_3^l$  and  $S_6^l$  and rerouting connected edges  $e_{1,6}$ ,  $e_{1,8}$ , and  $e_{1,10}$ . In subfigure (c), we identify all  $n^{G'(C_l)} = 5$  subchains  $\{l_{c_x}\}_{x=1}^5$  and form  $\{G''(l_{c_x})\}_{x=1}^5$ .

### 4.3 Routing requests

When a subchain  $l_{c_x}$  of a service chain  $C_l$  receives a user request  $l_{u_o} \in lU$ , based on the  $G''(l_{c_x})$ , it'll be routed to one of the service replicas of the source service  $l_s^{i''x} \in lS''^x \subseteq lS \subseteq S$  that is chosen based on a round-robin load balancing algorithm (neglecting the session affinity possibility). If the chosen replica  $s_j^1 \in l_s^{i''x}$  is either placed in the same host as the requester's service replica or if  $l_s^{i''x}$  is the source service replica  $S_1^l \in lS$  of the service chain  $C_l$ , then, assuming  $l_s^{i''x}$  is not a sink service, after the request is being process in the  $s_j^{i''1}$ , it will simply get routed to the next replica inside the same host. But in case the next replica is not in the same host, then it needs to get routed to the destination replica based on the network topology graph  $G(\tau)$ .




Fig. 5: Routing a request among replicas over a network

Considering  $G(\tau)$  is a directed acyclic graph and there should be one-and-only-one active path carrying information from a source to destination in a network, we extract the path between all pairs of hosts and form a set of  $\binom{H}{2}$  sub-graphs denoted as  $\{G(\tau)_{h_a, h_b} = (P_{h_a, h_b}, L_{h_a, h_b})\}_{h_a, h_b \in H}$  where  $P_{h_a, h_b}$  is a set of routers between  $h_a$  and  $h_b$  that are connected using  $L_{h_a, h_b}$  links. As an example in Figure 3, we have eleven hosts  $H = \{h_1, \dots, h_{11}\}$  that are connected together using six routers  $P = \{\rho_1, \dots, \rho_6\}$  with a tree topology  $\tau$  and network topology graph  $G(\tau)$ . Figure 5 represents the path from  $s_1^1$  to  $s_1^2$  based on the network topology  $\tau$  represented in Figure 3.

## 5 APPROXIMATING EXECUTION TIME OF THREADS IN A MULTI-CORE HOST

The next step after identifying and deploying service chains and network routers is to estimate the execution time of each endpoint function  $l_f_n^i \in lF_i \subseteq lF$  when a request arrives alongside other running processes in a host. In each cluster scenario, there exist a set of service replicas that each of them are capable of running a set of endpoint functions. When a request arrives, an endpoint function may propagate a few threads during its execution and each of its threads, when load-balances on a core's runqueue, may perform a few different tasks at a time. To be precise, a thread may perform one of the following tasks at any given time:

- Execute a CPU intensive task on the processor
- Read/write bytes of data from/to memory/cache
- Read/Write a file from/to storage
- Send/receive packets of information over the network
- Be in the *idle* mode

Additionally, other running threads inside a host have direct influence on the CPU time of an application because of the CPU scheduling mechanism in the operating system. For instance, Linux's CFS load-balance threads over dozens of CPU runqueues based on their load, CPU share, CPU quota and other parameters. Moreover, a typical multi-threaded service have a main thread and a number of

worker threads that coordinate with each other through synchronization primitives, that results additional overhead on execution time [45].

Even though predicting the exact execution time of a thread in a host can be very complex due to the aforementioned complexities, and considering that cycle-level simulation of multi-threaded services can be a very time-consuming procedure, we can predict a rough approximation of its run time by properly categorize type of running tasks in a host and approximate the execution time of each task by considering each thread's parent cgroups' resource constraints as well as other parallel active threads in the host. Consequently, this approximation may neglects a few performance factors such as microarchitecture dependent variable (i.e., CPU instruction sets and CPU architecture). However, during the modeling phase in a reference PTE, these effects will be measured and considered in the model, and therefore they contribute to the final approximation.

To approximate the total execution time of a thread, we categorize its tasks into 3 main types and accumulate measured values:

- 1) Accumulated CPU instructions to be executed
- 2) Accumulated stall cycles due to cache misses, memory access or block I/O
- 3) Payload to be sent to the next service over the network

## 5.1 Load-balancing of Threads over CPU runqueues

The first thing before being able to approximate the execution time of threads on a multi-core hardware, we need to estimate threads placement in a set of active CPU cores. This is crucial for the estimation, as the execution time is directly being affected by the other running threads in a core's runqueue. Therefore, we implemented a simplified version of Linux's CFS load-balancing algorithm [46] in PerfSim. We presented the details of our implementation in Algorithm 2. For simplicity, we assume there are only one NUMA node (one CPU socket with multiple cores/runqueues in it).

The CFS balances cores runqueues based on their *load* which is a measure that is a combination of threads accumulated weights (CPU shares) and average CPU utilization. To estimated weight of a thread  $t_m \in f_n^i$  on a core, we first divide CPU shares of its parent service  $S_i \in \mathcal{S}$  with total number of its running threads, and then divide the result with sum of all CPU shares currently running on the core:

$$t_m^{\text{share}} = \frac{S_i^{\text{CPU share}}}{|f_n^i|} \quad (6)$$

Assuming  $\hat{f}_{h_k}$  is the set of running threads placed on  $h_k$ :

$$\forall t_m \in \hat{f}_{h_k} h_k \rightarrow t_m^{\text{weight}} = \frac{t_m^{\text{share}}}{\sum_{t_j \in \hat{f}_{h_k}} t_j^{\text{share}}} \quad (7)$$

To calculate load of a thread:

$$t_m^{\text{load}} = \frac{t_m^{\text{runnable sum}} \times t_m^{\text{weight}}}{t_m^{\text{runnable period}}} \quad (8)$$

In which,  $t_m^{\text{runnable sum}}$  is the amount of time that the thread was runnable and  $t_m^{\text{runnable period}}$  is the total time that the thread could potentially be running.

## Algorithm 2: Simplified CFS load-balancing algorithm in a host $h_k$

```

1 sched_domains = (NUMA, pairs)
2 sched_groups[NUMA][0] = {core}_{core ∈ h_k}
3 sched_groups[pairs] =
  { {(1st core in pair, 2nd core in pair)}_{core ∈ core pairs} }_{core pairs ∈ h_k}
4 function load_balance_threads(h_k)
5   foreach current_core ∈ {1...h_kcores}
6     foreach sd ∈ {sched_domains}
7       first_core ← first_idle_core(sd)
8       if first_core = ∅
9         first_core ← first_core(sd)
10      if first_core = current_core
11        continue
12      foreach sched_group ∈ sd
13        sched_group.load =
14          avg(t_mload ∈ sched_group)
15      current_sg =
16        current_sched_group(current_core)
17      bussiest_sg = bussiest_sched_group()
18      if bussiest_sg.load > current_sg.load
19        do
20          bussiest_core =
21            bussiest_sg.next_bussiest_core()
22          load_balance(current_core, bussiest_core)
23        while load balance was not successful

```

Additionally, the load-balancing procedure takes place by considering cache locality and its hierarchical levels called *scheduling domains*. Within each scheduling domain, load-balancing occurs between a set of cores, called *scheduling groups*. Since we assumed there is only one CPU socket in each host, there exist only two scheduling domains: NUMA node level and core-pair level. Algorithm 2 represents our simplified version of the CFS algorithm used in PerfSim.

## 5.2 Approximating CPU time

Here we explain PerfSim's approach towards approximating CPU time of threads when co-located in a host.

### 5.2.1 Calculating auxiliary CPU share of a thread

When a thread co-locates with other threads on a multi-core machine, its execution time may be affected by the threads on the runqueue. Since we do not intend to simulate CPU time in a cycle-by-cycle basis, we approximate the effect of CPU bandwidth and CPU quota of threads by first defining an auxiliary CPU share for each thread, denoted as  $t_m^{\text{share ratio}}$ . For simplicity, we assume threads run in either (1) the *Best Effort* mode, which implies there is no CPU quota and CPU bandwidth defined, or (2) in the *Guaranteed* mode where CPU shares are fixed and guaranteed. For every threads in a runqueue, we calculate  $t_m^{\text{share ratio}} \in (0, 1024]$  as follows:

$$\forall t_m \in \hat{f}_{h_k} \text{ s.t. } t_m \in f_n^i \rightarrow t_m^{\text{share ratio}} = \begin{cases} \frac{t_m^{\text{share}} \times 1024}{\sum_{t_j \in \hat{f}_{h_k}} t_j^{\text{share}}} & \text{if } S_i^{\text{CPU quota}} = 0 \\ t_m^{\text{share}} & \text{elif } t_m^{\text{share}} \leq 1024 \\ 1024 & \text{elif } t_m^{\text{share}} > 1024 \end{cases} \quad (9)$$

### 5.2.2 Approximating cache miss rate based on cache stores and CPU size

The cache miss rate of a single thread can get affected by various hardware dependent factors such as:

- CPU L1, L2, and Last Level Cache (LLC) sizes
- Support of CPU Cache Allocation Technology (CAT) allowing software control over LLC allocation per process
- microarchitecture-level details

The cache miss rate may also get affected by following significant software dependant factors:

- Allocated CPU size of process
- Memory access rate of co-located threads in runqueue

Since the focus of PerfSim was to model the performance of microservices in merely software-oriented scenarios, such as container's placement, container/host affinity/anti-affinity, and resource allocation policies, we assumed a fairly similar microarchitecture and cache implementation details between the CPUs of PTE and the PE.

Hence, during the performance modeling phase, the effect of both (1) CPU size and (2) memory access rate (store/loads) of co-located threads is measured and a logarithmic regression model is being trained to predict the excessive cache miss rate of each thread. We denote the first effect on  $t_m \in f_n^i$  as  $t_m^{\text{CMC}}$  and the second one as  $t_m^{\text{CMT}}$ .

$$t_m^{\text{CMC}} = a \ln \left( \frac{S_i^{\text{CPU limits}}}{|f_n^i|} \right) + b \quad (10)$$

$$t_m^{\text{CMT}} = a \ln \left( \sum_{t_j \in \text{runqueue of } t_m} t_j^{\text{maccs}} \right) + b \quad (11)$$

We then apply these excessive penalties to the measured miss rate on the reference PTE as follows:

$$t_m^{\text{miss rate}} = \frac{t_m^{\text{cmisss}}}{t_m^{\text{crefs}}} \times (t_m^{\text{CMT}} + 1) \times (t_m^{\text{CMC}} + 1) \quad (12)$$

### 5.2.3 CPU time approximation

We used the *CPU Performance Equation* promoted in [47] to approximate the CPU time of an isolated thread  $t_m \in f_n^i$  on an idle core (in the Best Effort mode):

Isolated CPU Time( $t_m$ ) =

$$t_m^{\text{inst}} (t_m^{\text{CPI}} + \frac{t_m^{\text{maccs}}}{t_m^{\text{inst}}} \times t_m^{\text{miss rate}} \times t_m^{\text{cpenalty}}) \times \frac{1}{h_k^{\text{clock}}} \quad (13)$$

Considering that  $t_m^{\text{inst}}$  is the accumulated number of instructions that thread  $t_m$  executes during its lifetime, it also includes additional instructions that thread performs due to mem accesses, cache misses, blkio r/w and net I/O.

$$t_m^{\text{cycle penalty}} = \frac{t_m^{\text{maccs}}}{t_m^{\text{inst}}} \times t_m^{\text{miss rate}} \times t_m^{\text{cpenalty}} \quad (14)$$

$$\hat{t}_m^{\text{relative share}} = \frac{t_m^{\text{CPI}} \times \hat{t}_m^{\text{share ratio}}}{t_m^{\text{CPI}} + t_m^{\text{cycle penalty}}} \quad (15)$$

Assuming  $\hat{t}_m^{\text{inst}}$  is the remaining instructions at any given time  $\hat{t}$ , we approximated the CPU time of a thread on a CPU runqueue at any given time as follows:

$$\hat{t}_m^{\text{CPU time}} = \frac{\hat{t}_m^{\text{inst}} \times t_m^{\text{CPI}} \times \frac{1}{h_k^{\text{clock}}}}{\hat{t}_m^{\text{relative share}}} \quad (16)$$

Therefore, given a time period  $\Delta T \geq \hat{t}_m^{\text{CPU time}}$  (nanoseconds), we can estimate the executed instructions in that period as follows:

$$\hat{t} + \Delta T t_m^{\text{inst}} = \hat{t}_m^{\text{inst}} - \frac{\Delta T \times \hat{t}_m^{\text{relative share}}}{t_m^{\text{CPI}} \times \frac{1}{h_k^{\text{clock}}}} \quad (17)$$

### 5.3 Approximating the storage I/O time

Assuming a thread  $t_m$  is placed on host  $h_k$  with accumulated reads/writes of  $t_m^{\text{blk rw}}$ , we can approximate its storage I/O time as Equation 18.

$$t_m^{\text{blkio time}} = \frac{t_m^{\text{blk rw}}}{h_k^{\text{blkio bw}}} \quad (18)$$

At any given time  $\hat{t}$ , the remaining blkio time is denoted by  $\hat{t}_m^{\text{blkio time}}$ .

### 5.4 Approximating execution time of a thread

As described in the previous section, the execution time of a thread consists of the CPU time, blkio rw time and idle time. Therefore, given a time  $\hat{t}$ , we can estimate the execution time of thread on a core's runqueue as follows:

$$\hat{t}_m^{\text{exe time}} = \hat{t}_m^{\text{CPU time}} + \hat{t}_m^{\text{blkio time}} + \hat{t}_m^{\text{idle time}} \quad (19)$$

### 5.5 Approximating network transfer time

When a service replica  $l_s^i \in lS_i \in lS$  attempts to send  $l_e^v$  payload bytes of data to another service replica  $l_s^j \in lS_j \in lS$ , they are either placed on the same host and the network transfer time is equal to zero, or they placed in two different hosts  $h_a = \Pi(l_s^i)$  and  $h_b = \Pi(l_s^j)$ . In the latter case, a request needs to traverse an ordered subgraph  $G(\tau)_{h_a, h_b} = (N_{h_a, h_b}, L_{h_a, h_b})$  where  $N_{h_a, h_b}$  consist of both hosts  $h_a$  and  $h_b$  and all routers in between them. Also,  $L_{h_a, h_b}$  is all the network links between two hosts, including host links and topology links.

Depending on bandwidth usage of active requests flowing over a network link, the available bandwidth of a link changes over time. We therefore, denote the maximum available bandwidth of each link  $l_o \in L_{h_a, h_b}$  at any give time  $\hat{t}$  as  $\theta(l_o)_\hat{t}$  and calculate the maximum available bandwidth between all links in  $L_{h_a, h_b}$  as follows:

$$\Theta(L_{h_a, h_b})_\hat{t} = \min(\{\theta(l_o)\}_{l_o \in L_{h_a, h_b}}) \quad (20)$$

We then for any given time  $\hat{t}$ , calculate the request bandwidth between  $l_s^i$  and  $l_s^j$ , denoted as  $\theta(l_s^i, l_s^j)_\hat{t}$ , as follows:

$$\theta(l_s^i, l_s^j)_\hat{t} = \min \left( \Theta(L_{h_a, h_b})_\hat{t}, V_S(\text{out bw}, l_s^i), V_S(\text{in bw}, l_s^j) \right) \quad (21)$$

To calculate the network time between 2 service replica  $l_s^i$  and  $l_s^j$ , denoted as  $\omega(l_s^i, l_s^j)_\hat{t}$ , we calculate the transfer time by dividing the payload size with available bandwidth, and then sum it up with routers and links latency as follows:

$$\omega(l_s^i, l_s^j)_\hat{t} = \frac{\hat{t}_m^{\text{payload}}}{\theta(l_s^i, l_s^j)_\hat{t}} + \sum_{\rho_z \in P_{h_a, h_b}} \rho_z^{\text{lat}} + \sum_{l_o \in L_{h_a, h_b}} l_o^{\text{lat}} \quad (22)$$



## 5.6 Putting all together: The simulation

During each iteration, PerfSim predicts the next event to simulate. We categorized events into five main categories:

- **Request Generation** simulates the incoming traffic to each service chain by generating requests.
- **Threads Generation** checks for new queued threads and generates them.
- **Threads Execution Time Estimation** estimates the execution time of threads and checks whether their execution is going to end before the next transmission completes, or vice versa.
- **Threads Execution** estimates consumed instructions in all threads and repeat load-balancing among available cores in each host.
- **Network Transmissions:** Estimates the remaining payload in each active transmission in the network.

## 6 SIMULATION ACCURACY EVALUATION

In this section, we address the followings to highlight the accuracy, speed, applicability and significance of PerfSim.

- 1) **Potential threats to evaluation validity** (Section 6.2) to explain our approach for ensuring evaluation validity.
- 2) **Performance modeling accuracy** (Section 6.3) to reflect and report the PerfSim's simulation error using a comprehensive set of prevalent scenarios (Table 4).
- 3) **Execution time of PerfSim prototype** (Section 6.4) to report the amount of time required to simulate each scenario using an early-stage prototype of PerfSim.
- 4) **Simulating large scale service chains** (Section 6.5) to demonstrate the applicability of PerfSim for simulating large-scale service chains.
- 5) **Challenges and limitations** (Section 6.6) to highlight PerfSim's practical limitations as well as key challenges in simulating performance of computer systems.

### 6.1 Experimental setup

**Cluster setup.** To evaluate the simulation error of service chains' execution time between PerfSim and the real setup, we deployed a Kubernetes cluster using four physical machines (as compute nodes), and connected them through one or more Netgear 10Gbps routers (depending on the topology). All our servers are based on Intel Core i7 microarchitecture and their hardware details specified in Table 3. The latency of routers for processing a packet is 7.3e5ns, and links latencies are 4.2e5ns.

**Workload.** To evaluate various scenarios described in this section, we used *sfc-stress* [48] that is a customizable synthetic service chain benchmarking suit capable of generating different types of service chains with various types of CPU/memory/blkio/storage intensive as well as user-defined workloads. In *sfc-stress*, workloads are comparable to benchmark suites used for microservices, such as DeathStarBench [49]. They are similar to real-world cloud native microservices in various ways, such as capability of invoking any requested number of threads when their endpoint function being called, REST API based, written in Node.js, cloud native, deployable on both Kubernetes and Docker, fully parametric, and easily customizable. Using *sfc-stress* provided us with the flexibility of easily changing

workloads type, size, threads count, execution duration, resource allocation settings per service (both vertically and horizontally), and most importantly design custom service chains and generate periodic automated requested based on given arrival rate (request per second).

**Simulation parameters.** We used the pre-release prototype version (`alpha-0.1`) of PerfSim for this evaluation. To drive the simulations, we first extracted performance traces of various endpoint functions in *sfc-stress* and fed it to PerfSim as the simulation's initial parameters (Table 3). We only set  $W^{\text{millicores}}$  and  $W^{\text{mem}}$  to  $\underline{1}$  and set other weights to  $\underline{0}$  (similar to Kubernetes settings).

**Placement of service replicas.** We used Kubernetes default kube-scheduler with weights specified in Table 3 to place and govern service replicas.

### 6.2 Potential threats to evaluation validity

As with any evaluation, potential threats exist to reduce its validity. We, therefore, identified several of these threats before starting our evaluation process and paid great attention to address them during our evaluation.

#### 6.2.1 Evaluation consistency

During our initial experiments, we realized using only one experiment might not be representative of the evaluation scenarios performance behaviour. Thus, we repeated each experiment four times and calculated the average for each data point. In the end, we calculated moving average for both simulated and actual results as appropriate.

#### 6.2.2 Procedural rigor

Collocating hosts in our testbed with other nodes in a cluster may introduce additional noise to the system. Moreover, we realized connecting various hosts to the same router may affect the performance of the network intensive scenarios. To mitigate this problem and prevent any possible noise affecting the experiments results, we isolated the entire testbed by locating hosts in a separate rack and making sure no other hosts are connected to the routers.

We also noted that keeping energy saving options in the BIOS may slightly affect the CPU performance by altering its clock frequency. To mitigate this effect, we disabled energy saving options in the BIOS and optimized it for production.

#### 6.2.3 Comprehensiveness of evaluation scenarios

Another threat to our evaluation was to neglect common scenarios used in today's systems. Nowadays a real distributed system may consist of different type of services, each with various resource demands, workload sizes, interconnections, and expected average traffic. Thus, to ensure the comprehensiveness of our evaluation scenarios when studying the accuracy of PerfSim, we designed 6 category of scenarios, each focusing on a distinct aspect of PerfSim. These categories, as presented in Table 4, includes CPU intensive, memory intensive, network intensive, scenarios with multiple replicas, and scenarios with multiple endpoint functions per service.

### 6.3 Latency prediction accuracy

In this section, we evaluated the accuracy of PerfSim when simulating 104 different scenarios in 6 distinct category.

### 6.3.1 Evaluating distinctive classes of scenarios

We perform our evaluation by executing various workloads on our real Kubernetes cluster for  $\Delta T = 60$  seconds and compared average execution time of requests ( $u_o^{\text{exe time}}$ ) with the ones computed using PerfSim. We summarized the simulation error of all scenarios in Table 4 and illustrated the latency over requests line graphs in Figure 8. We also illustrated the average latency ( $u_o^{\text{exe time}}$ ) of each scenario as a bar graph in the same figure. We repeated each experiment four times to avoid any unaccounted artifact.

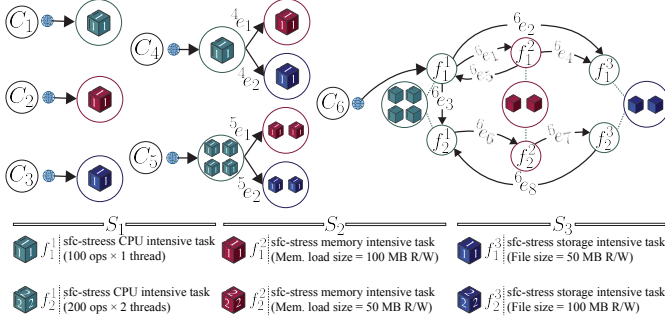


Fig. 6: Service chain scenarios ( $e_v^{\text{payload}}=50\text{MB}$  for all edges)

Each scenario has (1) a service chain as illustrated in Figure 6, (2) a network topology as illustrated in Figure 7, (3) an arrival rate  $C_l^{\text{rate req}}$ , and (4) a resource allocation setting.

In scenarios 1-20, our goal was to evaluate PerfSim simulation accuracy when assigning different CPU sizes to a single CPU-intensive container. For example, in scenario #1 we assigned 100 millicores to the only replica available in  $S_1$  and generated a traffic at a rate of  $1 \frac{\text{req}}{\text{s}}$  on  $C_1$  and measured requests average execution time ( $u_o^{\text{exe time}}$ ). We repeat the exact same scenario using PerfSim and calculate the simulation percentage error by comparing extracted  $u_o^{\text{exe time}}$  in PerfSim with the one measured in our real testbed.

In scenarios #2-19, we increase the CPU size by 100 millicores in each scenario and repeat the same procedure; in scenario #20, we also evaluated the case where all resources run in the best effort mode. In scenarios #21-40, we repeated all aforementioned procedures but for the memory intensive workloads (also implemented in the sfc-stress toolkit) to measure the accuracy of PerfSim for simulating memory-intensive workloads. Similarly, in scenarios #41-60, we focused on the disk-intensive workload.

To evaluate simulations involving networks and service chains, we designed scenarios #61-80. In the first half of this category (scenarios #61-70), we aimed to evaluate the accuracy of PerfSim when assigning egress bandwidth to an outgoing service in service chain  $C_4$  (Figure 6), and in the second half (scenarios #71-80), our goal was to measure the




Fig. 7: Network topologies used in the evaluation

Hosts	H	$h_k^{\text{lock}}$	$h_k^{\text{cores}}$	$h_k^{\text{rH}}$					
		$h_k^{\text{millicores}}$	$h_k^{\text{mem}}$	$h_k^{\text{in bw}}$	$h_k^{\text{out bw}}$	$h_k^{\text{blkio bw}}$	$h_k^{\text{blkio size}}$		
$h_1$	1.59								
$\vdots$	GHz	4	4000	16	1	1	657	500	
$h_4$				GB	Gbps	Gbps	MBps	GB	

Services	S	F	Endpoint Functions						
			T	Threads					
				$t_m^{\text{inst}}$	$t_m^{\text{cpu}}$	$t_m^{\text{macs}}$	$t_m^{\text{crefs}}$	$t_m^{\text{cmss}}$	$t_m^{\text{miss penalty}}$
$S_1$	$f_1^1$	$t_{1,1,1}$	1.4e9	0.7432	3.1e8	1.0e6	1.0e5	4	$\epsilon$
		$t_{1,2,1}$	3.1e9	0.750	7.2e8	1.2e6	1.3e5	4	$\epsilon$
	$f_2^1$	$t_{1,2,2}$	3.1e9	0.715	6.6e8	1.7e6	2.2e5	3	$\epsilon$
$S_2$	$f_2^1$	$t_{2,1,1}$	1.7e9	0.520	3.4e8	2.9e6	2.0e6	5	$\epsilon$
		$t_{2,2,1}$	1.0e8	0.4912	7.4e8	5.5e6	4.1e6	5	$\epsilon$
$S_3$	$f_1^3$	$t_{3,1,1}$	2.1e8	0.6660	4.3e7	1.5e6	5.7e5	5	5.1e7
		$t_{3,2,1}$	5.1e8	0.7199	2.2e7	4.3e6	2.3e6	5	1.0e8
	$f_2^3$	$t_{3,2,1}$	5.1e8	0.7199	2.2e7	4.3e6	2.3e6	5	1.0e8

TABLE 3: Constant control parameters of hosts, routers, links and workloads used in the evaluation

accuracy when assigning ingress bandwidth to the incoming node  $S_3$ . Scenarios #81-100 were designed to measure the same effect but when the system is deployed in slightly more complex network topology  $\tau_2$  (Figure 7).

Scenario Type	Experiment Number	$C_l^{\text{rate}} (\frac{\text{req}}{\text{s}})$	Target Resource Capacity	$\tau$		Avg. Error
				$u_o^{\text{exe time}}$	PerfSim Actual	
CPU intensive	1-19	$C_1^{\text{rate}}=1$	$S_1^{\text{CPU requests}}=\{100k\}_{k=1}^{19}$	$\tau_1$		
	20	$C_1^{\text{rate}}=1$	Best Effort	$\tau_1$	38.23	1.237%
Memory intensive	21-39	$C_2^{\text{rate}}=1$	$S_1^{\text{CPU requests}}=\{100k\}_{k=1}^{19}$	$\tau_1$		
	40	$C_2^{\text{rate}}=1$	Best Effort	$\tau_1$	9.654	6.662%
Storage intensive	41-59	$C_3^{\text{rate}}=1$	$S_3^{\text{CPU requests}}=\{100k\}_{k=1}^{19}$	$\tau_1$		
	60	$C_3^{\text{rate}}=1$	Best Effort	$\tau_1$	1.877	19.27%
Network intensive	61-70	$C_4^{\text{rate}}=1$	$S_4^{\text{out bw}}=\{100k\}_{k=1}^{10}$	$\tau_1$	99.60	5.086%
	71-80	$C_4^{\text{rate}}=1$	$S_3^{\text{in bw}}=\{100k\}_{k=1}^{10}$	$\tau_1$	34.45	12.77%
	81-90	$C_4^{\text{rate}}=1$	$S_1^{\text{in bw}}=\{100k\}_{k=1}^{10}$	$\tau_2$	99.62	4.306%
	91-100	$C_4^{\text{rate}}=1$	$S_3^{\text{in bw}}=\{100k\}_{k=1}^{10}$	$\tau_2$	34.45	9.772%
	Multi-replica	101-102	$C_5^{\text{rate}}=\{1, 3\}$	Best Effort	$\tau_1$	14.57
Multi-endpoint	103-104	$C_6^{\text{rate}}=\{\frac{1}{2}, \frac{1}{3}\}$	Best Effort	$\tau_1$	40.98	10.10%

Avg. error=9.203%

TABLE 4: Evaluation scenarios and associated errors

To evaluate scenarios involving multiple replicas, we designed scenarios #101-102. In these scenarios, we assigned 4 replicas to  $S_1$ , 2 replicas to  $S_2$ , and 2 replicas to  $S_3$ . We tested 2 different arrival rates on  $C_5$ : (a)  $1 \frac{\text{req}}{\text{s}}$  in scenario #101, and (b)  $3 \frac{\text{req}}{\text{s}}$  in scenario #102. In scenarios 103-104, we aimed to evaluate deployment that involved multiple endpoint functions, as well as, cases where an endpoint function spawns multiple threads (e.g.,  $f_2^1$ ). Another goal of these scenarios was to test the performance of PerfSim in scenarios with complex service chains (i.e.,  $C_6$ ) with multiple loops and nested diamond-shaped connections.

### 6.3.2 Reflections on simulation accuracy

In the vast majority of our experiments, as shown in Figure 8, PerfSim predicted the latency trend with a high accuracy (with average error rate of 9% error). However, in all scenarios, the actual service latency has visible fluctuations across time (the orange line) while PerfSim's predictions (blue line) did not capture those slight variations. This is because PerfSim is a discrete-event simulator and captures changes in the system state based on events such as request gener-

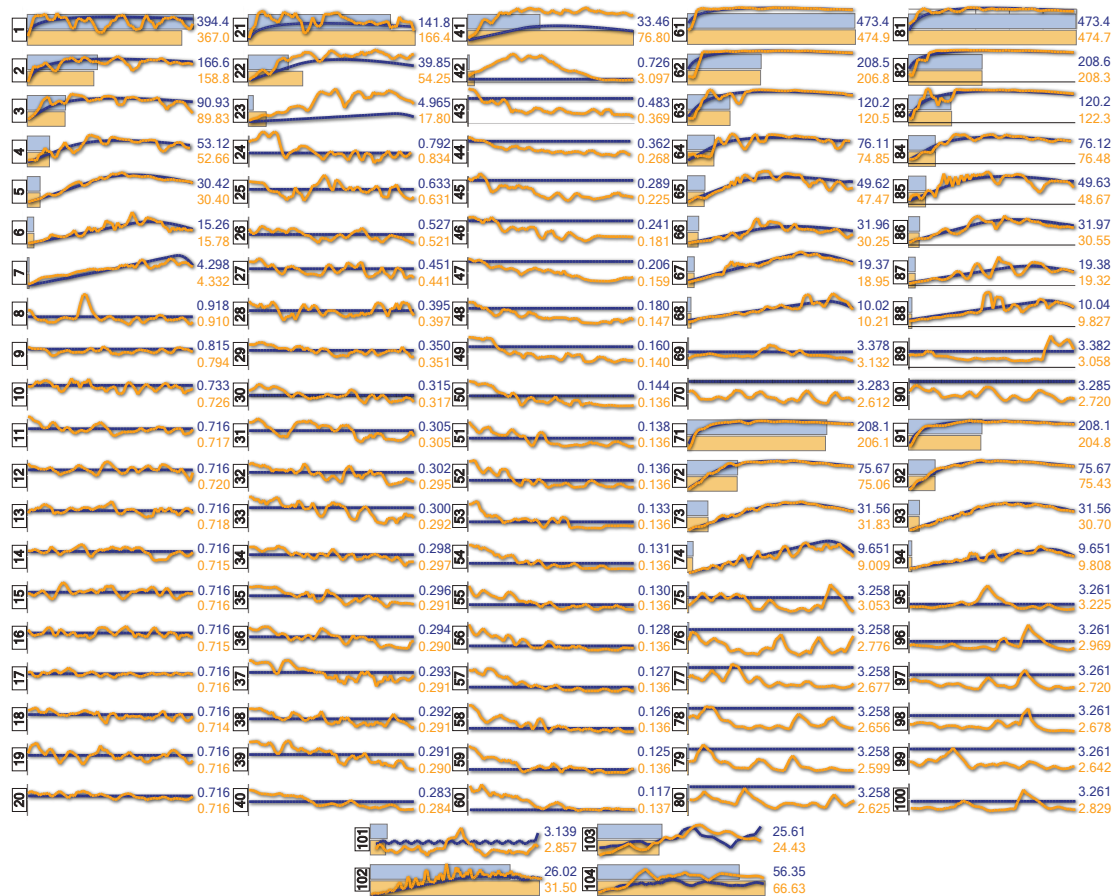


Fig. 8: Evaluation results representing latency vs requests line graphs of PerfSim (blue) and actual (orange) as well as the bar graph of their average latency ( $u_o^{exe\ time}$ )

ation/conclusion, thread spawn/kill, network transmission start/end, queue start/end, task scheduling start/end, etc. Therefore, as PerfSim events are not defined based on the CPU clock resolution (i.e., kernel clock) to maintain fast simulation speed, it doesn't consider the CPU noise-factor to simulate slight fluctuations (e.g., slight changes in CPU frequency or clock speed).

Another observation captured from the experiments results is the clear harmony between almost all simulation results and the reality. Nonetheless, in some intense scenarios (e.g., scenarios 21-23 or 41-42) where very low CPU resources has been assigned to a memory intensive or storage intensive services (100-300 millicores), even though PerfSim accurately simulated the trend of performance, we observe a gap between PerfSim's approximated latency and the actual obtained latency in reality. The nature of this variation is due to the use of a static model for each service in our experiments. In highly overloaded or intense scenarios, specially in a memory-intensive task, cache-miss rate may significantly affected and undergo substantial changes, which even though has already been considered in PerfSim's core, combining it with ever increasing rate of context switches and CPU migrations in such intense scenarios, creates a high competition between service threads from one hand and system/user threads on the other hand which puts the entire system in a chaotic situation. This problem can be addressed by utilizing more dynamic performance models per service

and considering system level threads in the simulation.

#### 6.4 Execution time of PerfSim

One of the key ideas behind designing PerfSim is to eliminate the barrier for utilizing advanced machine learning techniques (such as deep reinforcement learning) for optimizing the performance of large-scale service chains that require fast efficiency in estimating various resource allocation and placement scenarios. In this section, we evaluated the execution time of all 104 scenarios and compared them with the simulation time in PerfSim.

Due to the fact that the main users of PerfSim will be researchers and performance engineers with limited access to computation resources, and to highlight the lightweight nature of PerfSim, a single personal laptop has been used to measure the simulation speed. We ran PerfSim in single-threaded mode using a MacBook Pro with a 2.6 GHz Intel Core i7 CPU and 16GB of RAM.

It represents the detailed speed comparison in Figure 9. As shown in the graph, a single-thread execution of PerfSim is 16-1200 times faster to evaluate the performance of a scenario when compared to running/evaluating the same deployment on a real cluster.

#### 6.5 Simulating large scale service chains

When designing PerfSim, we initially aimed to use it for optimizing large-scale service chains. Therefore, one of the




Fig. 9: Execution time (in seconds) of all scenarios in both PerfSim and real deployment (note the difference in scales).

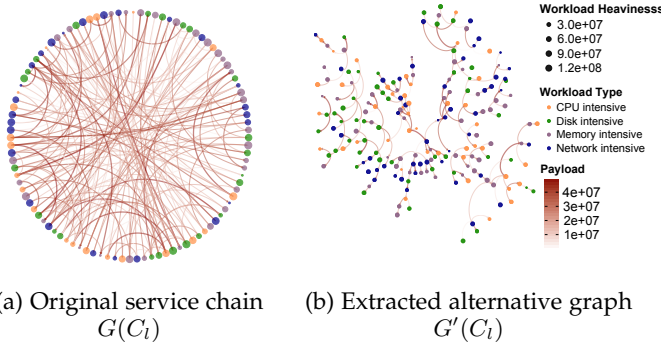


Fig. 10: Simulating a large service chain with different types of workload, payload sizes and heaviness (e.g., large number of instruction, cache r/w)

aspects of this evaluation has been focused on the ability of PerfSim to simulate large-scaled service chains. For this purpose, we simulated the performance of a large randomized service chain deployed over a cluster of 100 host. We illustrated both our generated service chain  $G(C_i)$  and extracted alternative graph  $G'(C_i)$  in Figure 10. This service chain consists of 100 nodes and 200 edges with randomized payloads, workload heaviness and workload type.

Similar to the previous scenarios, we used a single-threaded version of PerfSim for the entire experiment. We presents both simulation parameters and results in Table 5. As shown, PerfSim was able to simulate the performance of this large-scale service chain for up to 10 times faster than targeted execution time  $\Delta\hat{T}$  while using reasonable CPU and memory resources.

Simulation parameters		Simulation results	
Parameter	Value	Key	Value
$C_i^{rate}$	$0.1 \frac{req}{s}$	$\bar{i}_{u_o}^{exe\ time}$	9,744.95ms
$\Delta\hat{T}$	600s	Sim. time	68.80s
$ S $	100	$ S' $	201
$ E $	200	$n^{G'(C_i)}$	161
$ H $	100	Exec. threads	12060
$\tau$	$\tau_1$	CPU usage	99.9%
Resources	Best effort	Mem. usage	1.64GB

TABLE 5: Simulation parameters and summary of results in the large service chain scenario using PerfSim prototype

## 6.6 Discussion on challenges and limitations

In this section, we discuss main challenges of performance simulation, as well as, key limitations for using PerfSim.

### 6.6.1 Comments on time-predictability of tasks

Time-predictability of a task in a computer software mainly refers to the properties of the phenomenon execution time, including the execution pattern of instructions or spectrum of occurrence events related to a job’s execution time [50], [51]. These properties directly affect the process of modeling a job’s execution time and, in many cases, can make it impossible to extract any reusable pattern [52]. Hence, it is a vital property to hold in real-time and embedded systems due to their time and mission critical nature.

To shed light on the surface of the problem, we provide an example in Listing 2, representing 2 functions ( $f_1$  and  $f_2$ ), both calculating  $n$  number of MD5 hashes. Function  $f_1$  receives the value of  $n$  as an argument whereas  $f_2$  initiates  $n$  with a randomly generated number between 0 and  $10^{10}$ . For these functions, modeling  $f_1$  is relatively straightforward because its execution time is directly related to the input argument  $n$ ; for  $f_2$  on the other hand, it is extremely hard to approximate the execution time because it depends on the randomly generated variable  $n$ .

```

//high time-predictability
function f1(n){
  for($i=0; $i<$n; $i++){
    createHash('md5');
  }
}

//low time-predictability
function f2(){
  let n = randInt(0,10**10);
  for($i=0; $i<$n; $i++){
    createHash('md5');
  }
}
    
```

Listing 9: An example of time-predictability in functions

With the growing convergence of High-Performance Computing (HPC) systems and cloud computing paradigm, modern design patterns tend to progressively consider the properties of time-predictable computing in their approaches to be able to plan for the underlying resource allocation policies, infrastructure costs, as well as, to avoid Service Level Agreement (SLA) penalties. However, many organizations are still not ready for such a transition, and this imposes immense challenges in modeling and simulating the performance of their designed systems.



### 6.6.2 Modeling the software architecture

The software architecture, as the key foundation of any software system, can significantly impact the performance of chain of services. As described in the previous sections, PerfSim can adequately model architecture of a software system by defining the connection between microservices and hosts, and translating each microservice into a set of endpoint functions while describing their properties.

However, one of the critical challenges in simulating large-scale systems is to model components that may change their behaviour based on the context. For example, some architectures foresee an HTTP cache server in the front-line of their main webserver to reduce the latency of requests or API calls. When a request arrives, these cache servers check their temporary storage for the requested content and only send the request to the upstream server when they have not found the requested content in the cache or when it is outdated. Modeling such services that can change their behavior based on the situation, although possible in PerfSim, requires designing content-aware modules for the simulator. For this, interested users can extend the corresponding classes to override the methods responsible for consuming resources, as well as the pattern of executing endpoint functions and their corresponding threads.

### 6.6.3 Model portability

We rigorously evaluated PerfSim in various scenarios and demonstrated its accuracy in different settings. However, we find out that even though the extracted models can be used in different hosts and settings, they cannot be generalised/used to simulate the performance in other microarchitecture settings. For example, a model extracted for an Intel CPU architecture cannot be used to accurately simulate the performance of the ARM-compatible version of the software in an ARM-based architecture (e.g., Apple's M-series chips - even with dynamic binary translation enabled).

### 6.6.4 Simulating highly overloaded scenarios

Two paramount features of PerfSim are (1) its ability to simulate contention for various resources when multiple containers are packed in a host, and (2) simulating network congestion when microservices communicate over a network topology. However, when a process enters a highly overloaded state (i.e., when the rate of incoming requests is more than the processing capacity of serving processes), the request queue will grow without a bound, and in the meanwhile, the system aggressively throttles CPU and network resources. This, in turn, will increase the rate of context switches, CPU migrations, and cache misses.

In large-scale systems, performance engineers generally define a rigorous set of constraints and resource management policies throughout a software's design and deployment lifecycle to prevent the system from entering such unpredictable situations [53]. Hence, accurate simulation of such chaotic situations are extremely challenging.

### 6.6.5 Simulating containers on system virtual machines

One of the widely used schemes for deploying container orchestrators (e.g., Kubernetes) on a cluster involves the use of hypervisor-based virtualization platforms, such as

OpenStack or VMWare. Such deployment scheme becomes popular as system virtual machines (1) allows the use of various operating systems in one machine, (2) allows OS-level customization for each service, and (3) provide better data and resource isolation between neighbouring services. However, study shows that virtualization suffers from noticeable performance overhead due to the additional OS layer [54], [55].

One of the main intentions of designing PerfSim is to simulate thread-level performance of services and to enable accurate prediction of endpoint functions performance when various threads compete over resources in a host. Thus, the additional complexity of OS-level threads in the aforementioned schemes, if not appropriately modelled, may affect the accuracy of simulation in scenarios where the virtualization overhead is considerably high.

### 6.6.6 Challenges of performance modeling

As described earlier, the process of performance modeling occurs on a PTE that mimics a PE to (1) obtain a complete isolation for services to collect accurate measurements for each endpoint function and (2) to prevent interference with the experience of real cloud users. However, preparing a PTE may impose both technical and financial challenges that needs to be addressed in the design of DevOps process flows. Automating the entire process of performance testing after each release may dramatically reduce the cost of modeling and dramatically increase the accuracy of simulations.

## 7 CONCLUSION AND FUTURE WORKS

In this work, we presented PerfSim as a systematic method and simulation platform for modeling and simulating the performance of large-scale service chains in the context of cloud native computing. Using performance tracing and monitoring tools, PerfSim allows performance modeling of various microservices and their corresponding service chains in cloud native orchestration platforms (such as Kubernetes that we used in this article) and enables the possibility to simulate the effectiveness of different resource management scenarios and placement policies. We evaluated the accuracy of PerfSim in a set of prevalent scenarios and obtained ~81-99% prediction accuracy as well as ~16-1200 times speed-up factor in comparison to testing on a real system (excluding the time needed for setting up the system and configuring the cluster). We evaluated the capability of PerfSim in simulating large-scale service chains and showed that using a single-core of a laptop, it can achieve a 10 fold speed-up factor to simulate a scenario with complex service chains (composed of 100 microservices interconnected with 200 links deployed over a cluster with 100 hosts).

Our future work will cover three main categories: (1) improving speed, (2) improving accuracy, and (3) expanding use cases of PerfSim. To increase its speed-up factor, we plan to refactor it to harness the parallel processing capabilities of GPUs. To improve its accuracy, we will add more layers/features, such as the support of defining cache-size per container (to simulate various scenarios based on Intel's CAT) or adding the simulation of poll mode in NICs (e.g., scenarios involving DPDK). Because advanced performance optimization methods can harness PerfSim's fast scenario

assessment capability to quickly perform numerous policy trials and errors, we will work towards proposing novel performance optimization methods. We will also use PerfSim to train deep reinforcement learning agents that can autonomously control different policies in the cluster. We will also implement accounting features in PerfSim that would allow detailed report on cost of deployment (e.g., when deploying on a public cloud) based on pricing/cost information provided within the scenario config files.

## ACKNOWLEDGMENTS

Parts of this work has been supported by the Knowledge Foundation of Sweden (KKS) through the Project AIDA (20200067). The contribution of Auday Al-Dulaimy to this work has been performed with the support of the KKS under the SACSys project.

## REFERENCES

- [1] "Cncf cloud native definition v1.0," Cloud Native Computing Foundation, 06 2018. [Online]. Available: <https://github.com/cncf/toc/blob/master/DEFINITION.md>
- [2] T. Cerny, M. J. Donahoo, and M. Trnka, "Contextual understanding of microservice architecture: Current and future directions," *SIGAPP Appl. Comput. Rev.*, vol. 17, no. 4, p. 29–45, Jan. 2018.
- [3] "An update on container support on google cloud platform," Google Cloud Platform Blog, June 2014. [Online]. Available: <https://cloudplatform.googleblog.com/2014/06/an-update-on-container-support-on-google-cloud-platform.html>
- [4] E. Casalicchio, *Container Orchestration: A Survey*. Springer International Publishing, 2019, pp. 221–235.
- [5] B. Yang, A. Sailer, and A. Mohindra, "Survey and evaluation of blue-green deployment techniques in cloud native environments," in *Service-Oriented Computing – ICSOC 2019 Workshops*, S. Yangui, A. Bouguettaya, X. Xue, N. Fati, W. Gaaloul, Q. Yu, Z. Zhou, N. Hernandez, and E. Y. Nakagawa, Eds. Springer International Publishing, 2020, pp. 69–81.
- [6] Y. Mao *et al.*, "Resource management schemes for cloud-native platforms with computing containers of docker and kubernetes," *arXiv preprint arXiv:2010.10350*, 2020.
- [7] Y. Hu *et al.*, "Concurrent container scheduling on heterogeneous clusters with multi-resource constraints," *Future Generation Computer Systems*, vol. 102, pp. 562 – 573, 2020.
- [8] B. Liu *et al.*, "A new container scheduling algorithm based on multi-objective optimization," *Soft Computing*, vol. 22, pp. 7741–7752, 2018.
- [9] T. Menouer, "Kcss: Kubernetes container scheduling strategy," *The Journal of Supercomputing*, pp. 1–27, 2020.
- [10] Y. Mao *et al.*, "Speculative container scheduling for deep learning applications in a kubernetes cluster," *ArXiv*, vol. abs/2010.11307, 2020.
- [11] S. Burroughs, "Towards predictive runtime modelling of kubernetes microservices," Ph.D. dissertation, Hamilton, New Zealand, 2021, masters.
- [12] E. Casalicchio, "A study on performance measures for auto-scaling cpu-intensive containerized applications," *Cluster Computing*, vol. 22, no. 3, pp. 995–1006, Sep 2019.
- [13] W.-S. Zheng and L.-H. Yen, "Auto-scaling in kubernetes-based fog computing platform," in *New Trends in Computer Technologies and Applications*, C.-Y. Chang, C.-C. Lin, and H.-H. Lin, Eds. Singapore: Springer Singapore, 2019, pp. 338–345.
- [14] D. Kliazovich *et al.*, "Greencloud: A packet-level simulator of energy-aware cloud computing data centers," in *2010 IEEE Global Telecommunications Conference GLOBECOM 2010*, 2010, pp. 1–5.
- [15] R. N. Calheiros *et al.*, "Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms," *Software: Practice and experience*, vol. 41, no. 1, pp. 23–50, 2011.
- [16] C. Puliafito *et al.*, "Mobfogsim: Simulation of mobility and migration for fog computing," *Simulation Modelling Practice and Theory*, vol. 101, p. 102062, 2020, modeling and Simulation of Fog Computing.
- [17] G. F. Riley and T. R. Henderson, *The ns-3 Network Simulator*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 15–34.
- [18] H. Gupta *et al.*, "ifogsim: A toolkit for modeling and simulation of resource management techniques in the internet of things, edge and fog computing environments," *Software: Practice and Experience*, vol. 47, no. 9, pp. 1275–1296, 2017.
- [19] C. Sonmez, A. Ozgovde, and C. Ersoy, "Edgecloudsim: An environment for performance evaluation of edge computing systems," *Transactions on Emerging Telecommunications Technologies*, vol. 29, no. 11, p. e3493, 2018.
- [20] M. C. Silva Filho *et al.*, "Cloudsim plus: A cloud computing simulation framework pursuing software engineering principles for improved modularity, extensibility and correctness," in *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 2017, pp. 400–406.
- [21] J. Byrne *et al.*, "Recap simulator: Simulation of cloud/edge/fog computing scenarios," in *2017 Winter Simulation Conference (WSC)*. IEEE, 2017, pp. 4568–4569.
- [22] I. Lera, C. Guerrero, and C. Juiz, "Yafs: A simulator for iot scenarios in fog computing," *IEEE Access*, vol. 7, pp. 91 745–91 758, 2019.
- [23] C. Fiandrino *et al.*, "Performance and energy efficiency metrics for communication systems of cloud computing data centers," *IEEE Transactions on Cloud Computing*, vol. 5, no. 4, pp. 738–750, 2017.
- [24] D. Boru *et al.*, "Energy-efficient data replication in cloud computing datacenters," in *2013 IEEE Globecom Workshops*, 2013, pp. 446–451.
- [25] A. Varga, *A Practical Introduction to the OMNeT++ Simulation Framework*. Springer International Publishing, 2019, pp. 3–51.
- [26] A. Wahid-Ul-Ashraf, M. Budka, and K. Musial, "Netsim – the framework for complex network generator," *Procedia Computer Science*, vol. 126, pp. 547–556, 2018, knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 22nd International Conference, KES-2018, Belgrade, Serbia.
- [27] R. Beuran, *Introduction to network emulation*. CRC Press, 2012.
- [28] *MiniNet*. [Online]. Available: <http://mininet.org/>
- [29] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: Rapid prototyping for software-defined networks," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, ser. Hotnets-IX. New York, NY, USA: ACM, 2010.
- [30] P. Wette *et al.*, "Maxinet: Distributed emulation of software-defined networks," in *2014 IFIP Networking Conference*, 2014, pp. 1–9.
- [31] R. Mayer *et al.*, "Emufog: Extensible and scalable emulation of large-scale fog computing infrastructures," in *2017 IEEE Fog World Congress (FWC)*. IEEE, 2017, pp. 1–6.
- [32] A. Coutinho *et al.*, "Fogbed: A rapid-prototyping emulation environment for fog computing," in *2018 IEEE International Conference on Communications (ICC)*. IEEE, 2018, pp. 1–7.
- [33] M. A. To, M. Cano, and P. Biba, "Dockemu—a network emulation tool," in *2015 IEEE 29th international conference on advanced information networking and applications workshops*. IEEE, 2015, pp. 593–598.
- [34] J. Lai *et al.*, "Network emulation as a service (neaaS): Towards a cloud-based network emulation platform," *Mobile Networks and Applications*, pp. 1–15, 2020.
- [35] M. Abdullah *et al.*, "Learning predictive autoscaling policies for cloud-hosted microservices using trace-driven modeling," in *2019 IEEE International Conference on Cloud Computing Technology and Science*, 2019, pp. 119–126.
- [36] A. Jindal, V. Podolskiy, and M. Gerndt, "Performance modeling for cloud microservice applications," in *Proceedings of the 2019 ACM/SPEC International Conference on Performance Engineering*, ser. ICPE '19. New York, NY, USA: ACM, 2019, p. 25–32.
- [37] M. Gokan Khan *et al.*, "A performance modelling approach for sla-aware resource recommendation in cloud native network functions," in *2020 6th IEEE Conference on Network Softwarization (Net-Soft)*, 2020, pp. 292–300.
- [38] *cggroups*. [Online]. Available: <https://www.kernel.org/doc/Documentation/cgroup-v2.txt>
- [39] P. Turner, B. B. Rao, and N. Rao, "Cpu bandwidth control for cfs," in *Proceedings of the Linux Symposium*, 2010, pp. 245–254.
- [40] istio project authors, "istio: An open platform to connect, manage, and secure microservices." Accessed: 2021. [Online]. Available: <http://istio.io>
- [41] envoy project authors, "Envoy: A cloud-native high-performance edge/middle/service proxy," Accessed: 2021. [Online]. Available: <http://envoyproxy.io>

- [42] lightstep project authors, "Lightstep: A distributed tracing library," Accessed: 2021. [Online]. Available: <http://lightstep.com>
- [43] zipkin Project Authors, "Zipkin: A distributed tracing system," Accessed: 2021. [Online]. Available: <http://zipkin.io>
- [44] jaeger project authors, "Jaeger: A distributed tracing system," Accessed: 2021. [Online]. Available: <http://jaegertracing.io>
- [45] S. De Pestel *et al.*, "Rppm: Rapid performance prediction of multithreaded applications on multicore hardware," *IEEE Computer Architecture Letters*, vol. 17, no. 2, pp. 183–186, 2018.
- [46] J.-P. Lozi *et al.*, "The linux scheduler: A decade of wasted cores," in *Proceedings of the Eleventh European Conference on Computer Systems*, ser. EuroSys '16. New York, NY, USA: ACM, 2016.
- [47] D. Patterson, *Computer architecture : a quantitative approach*. San Mateo, Calif: Morgan Kaufman Publishers, 1990.
- [48] M. Gokan Khan, "sfc-stress: A customizable synthetic service chain," <https://github.com/michelgokan/sfc-stress>, 2021.
- [49] Y. Gan *et al.*, "An open-source benchmark suite for microservices and their hardware-software implications for cloud & edge systems," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 3–18.
- [50] R. Kirner and P. Puschner, "Time-predictable computing," in *Software Technologies for Embedded and Ubiquitous Systems*, S. L. Min, R. Pettit, P. Puschner, and T. Ungerer, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 23–34.
- [51] V. Nelis *et al.*, "The challenge of time-predictability in modern many-core architectures," *OpenAccess Series in Informatics*, vol. 39, pp. 63–72, 01 2014.
- [52] M. Schoeberl, "Is time predictability quantifiable?" in *2012 International Conference on Embedded Computer Systems (SAMOS)*, 2012, pp. 333–338.
- [53] M. K. Molloy, *Fundamentals of Performance Modeling*, 1st ed. USA: Prentice Hall PTR, 1988.
- [54] S. Shirinbab, L. Lundberg, and E. Casalicchio, "Performance evaluation of container and virtual machine running cassandra workload," in *2017 3rd International Conference of Cloud Computing Technologies and Applications (CloudTech)*, 2017, pp. 1–8.
- [55] S. Shirinbab, "Performance implications of virtualization," 2019.



**Auday Al-Dulaimy** is a postdoctoral researcher at the School of Innovation, Design and Engineering, Mälardalen University, Sweden. He received PhD degree in Computer Science from Beirut Arab University, Lebanon in 2017; and B.Sc and M.Sc degrees in Computer Science from Al-Nahrain University, Iraq in 2000 and 2003 respectively. His major areas of research interest includes Distributed Systems, Cloud Computing, Internet of Things, and Edge Computing.



**Michel Gokan Khan (M'13)** received the M.S. degree in software engineering from Iran University of Science and Technology, Tehran, Iran, in 2016. He is currently pursuing the Ph.D. degree in computer science at Karlstad University, Sweden. His main research interest includes network function virtualization, cloud native applications, mathematical and machine learning based performance optimization for telecom clouds, cloud computing, and applications of AI in distributed systems.



**Andreas J. Kassler (SM'12)** received his M.Sc. degree in mathematics/computer science from Augsburg University, Germany in 1995 and his Ph.D. degree in computer science from University of Ulm, Germany in 2002. He is currently Professor with the Department of Computer Science, Karlstad University, Sweden. He teaches wireless networks and advanced topics in computer networking. His main research interests include Software Defined Networks, Future Internet, and Network Function Virtualization.



**Javid Taheri** received the Ph.D. degree in mobile computing from the School of Information Technologies, The University of Sydney, Australia. He is currently a Professor with the Department of Computer Science, Karlstad University, Sweden. His major areas of interest are profiling, modeling, and optimization techniques for private and public cloud infrastructures, profiling, modeling, and optimization techniques for software-defined networks, and network-aware scheduling algorithms for cloud computing.