# **INDICES:** Applying DDDAS Principles for Performance Interference-aware Cloud-to-Fog Application Migration

Shashank Shekhar and Ajay Dev Chhokra and Anirban Bhattacharjee and Yogesh Barve and Shweta Khare and Guillaume Pallez and Hongyang Sun and Aniruddha Gokhale and Gabor Karsai

**Abstract** An increasing number of interactive applications and services, such as online gaming and cognitive assistance, are being hosted in the cloud because of the elastic properties and cost benefits of distributed data centers. Despite these benefits, the longer and often unpredictable end-to-end network latencies between the end user and the cloud can be detrimental to time-critical response to the applications. Although technology enablers, such as Cloudlets or Micro Data Centers (MDCs), are increasingly being leveraged by cloud infrastructure providers to address the network latency concerns, existing efforts in re-provisioning services from the cloud to the MDCs seldom focus on ensuring that the performance properties of the migrated services are met. This chapter demonstrates the application of Dynamic Data Driven Applications Systems (DDDAS) principles in the systems software layer, to address these limitations by: (a) determining when to re-provision; (b) identifying the appropriate MDC and a suitable host within that MDC that meets the performance considerations of the applications; and (c) ensuring that the cloud service provider

Shashank Shekhar

Siemens Corporate Technology, Princeton, NJ, USA, e-mail: shashankshekhar@siemens.com and Ajay Dev Chhokra

Vanderbilt University, Nashville, TN, USA e-mail: ajay.d.chhokra@vanderbilt.edu and Anirban Bhattacharjee

Vanderbilt University, Nashville, TN, USA e-mail: anirban.bhattacharjee@vanderbilt.edu and Yogesh Barve

Vanderbilt University, Nashville, TN, USA e-mail: yogesh.d.barve@vanderbilt.edu and Shweta Khare

Vanderbilt University, Nashville, TN, USA e-mail: shweta.p.khare@vanderbilt.edu and Guillaume Pallez

INRIA, Bordeaux, France e-mail: guillaume.pallez@inria.fr

and Hongyang Sun

Vanderbilt University, Nashville, TN, USA e-mail: hongyang.sun@vanderbilt.edu and Aniruddha Gokhale

Vanderbilt University, Nashville, TN, USA e-mail: a.gokhale@vanderbilt.edu and Gabor Karsai

Vanderbilt University, Nashville, TN, USA e-mail: gabor.karsai@vanderbilt.edu

continues to meet customer service-level objectives while keeping its operational and energy costs low. Empirical evaluations using a setup comprising a cloud data center and multiple MDCs composed of heterogeneous hardware are presented to validate the capabilities of the INDICES (Intelligent Deployment for ubiquitous Cloud and Edge Services) framework to process DDDAS methods. It should also be noted that the capabilitis created through INDICES are aimed to satisfy a broad set of applications requiring real-time data deliver and thus also satisfy the support requirements of environments enabling DDDAS-based applications.

# **1** Introduction

The cloud has become an attractive hosting platform for a variety of interactive and soft real-time applications, such as scientific research, cognitive assistance, health monitoring systems and collaborative learning, due to its elastic properties and cost benefits. Despite these substantial advantages, the response time considerations of the users mandate lower latencies for the applications, especially for the Dynamic Data Driven Applications Systems (DDDAS), such as stochastic estimation of mobile platforms, cyber security and analytics, and urban situation awareness. For example, prior works [33, 38] have shown that in highly interactive applications, latencies exceeding 100 milliseconds (ms) may be too high for acceptable user experience. However, real-world experiments have shown that the latencies experienced by geographically distributed users of an interactive service may tend to be on the order of several hundreds of milliseconds [57]. Consequently, there is a need to bound the resulting response times within acceptable limits.

For any cloud-hosted interactive application, the key factors that affect the round trip latencies are the network delay between the client and the cloud, particularly the roundtrip delay between the nearest access point of the client and the cloud, and the time it takes to serve the client request in the cloud. Other factors are important but can be negligible, such as the time taken by the thin client (e.g., lightweight computer, such as an edge-node or Internet of Things IoT) to communicate with the nearest access point [42] or time for the load balancer at the cloud front-end [32]. Thus, any improvement in response times must focus on reducing the network delays and the server processing time.

In recent years, edge computing, cloudlets [52] or Micro Data Centers (MDCs) [3] have emerged as one of the key mechanisms to manage and bound the transit latency by supporting cloud-based services closer to the clients. MDCs can be viewed as "data center in a box," which act as the middle tier in the emerging "mobile device–MDC–cloud" hierarchy [52]. In the rest of the chapter, the term MDC is used to represent all emerging mechanisms, such as Cloudlets, Micro Datacenters (MDCs), Locavore infrastructures, etc. MDCs possess key attributes of soft states, sufficient compute power and connectivity, proximity to clients, and conformance to standard cloud technologies.

Recent efforts [14,15,39,66] have leveraged the cloud, MDCs and mobile ad-hoc networks by focusing primarily on cyber-foraging [4,52], where tasks are offloaded from mobile devices to the cloud/MDCs for faster execution and conserving resources on the mobile client endpoints. However, only recently has there been an increasing interest in moving tasks from the central clouds to the MDCs. Those that do, however, have seldom considered the resulting application performance because these efforts tend to overlook the fact that servers within the MDC may themselves get overloaded, thereby worsening the user experience as compared to that of a traditional cloud-hosted interactive service. On the other hand, efforts that consider performance of MDCs make very simplistic assumptions regarding their performance models.

In this chapter, focus is placed on the performance of MDCs, specifically on the key factors contributing to performance degradation of applications running in MDCs and data centers in general. One fundamental system property that is often overlooked in prior works is performance interference, which is caused by co-located applications in virtualized infrastructures [11, 18, 35, 37]. *Performance interference,* being an inherent property of any virtualized system, manifests itself in MDCs also and therefore must be factored in by any approach that is performance-aware. Thus, the methods discussed here consider a "just-in-time" and "performance-aware" service migration approach for migrating cloud-based interactive services hosted in a centralized cloud data center to an MDC.

To support such as vision, a number of challenges manifest themselves as described below - that must be addressed by any just-in-time and performance-aware cloud-to-fog application migration solution:

- *Hardware heterogeneity:* Differences in hardware configurations of the servers in a traditional data center and in an MDC will provide different performance profiles, and hence should be accounted for in the analyses.
- Performance Interference: Noisy neighbors [12] cause performance issues, which
  must be considered in both traditional data centers and MDCs. However, since
  an MDC is orders of magnitude smaller than a traditional data center, performance interference may be more pronounced and manifests more rapidly than in
  traditional data centers.
- Network performance measurements: Accurate latency and bandwidth measurements are required to reliably work over Wide Area Networks (WANs). Measurements are important, since accurately estimating the value for the transit latency is critical in the problem formulation and solution.
- System performance measurements: Reliable application and server performance measurement and logging techniques are required to precisely measure the service execution (i.e., service execution time on the hardware of the data centers or MDCs).

The cloud-fog resource spectrum is a highly dynamic system and hence no *a priori*, statically-defined solution is going to address these multitude of challenges all at once. Instead, solutions that rely on dynamic data-driven techniques such as that envisioned by the Dynamic Data Driven Applications Systems (DDDAS) paradigm [7, 16] offer the most promise. DDDAS proposes an approach where

data-driven models of a system are learned through dynamic instrumentation of the system, the models are simulated to conduct what-if analysis, and in turn these insights are used in a feedback loop to steer the system along the desired trajectory. This chapter describes the DDDAS approach at the systems software layer to address these aforementioned challenges and present a just-in-time and performance-aware cloud-to-fog application migration approach that involves the following contributions:

- A data-driven modeling technique to estimate the performance of a cloud application on different hardware platforms subject to performance interference stemming from various co-located applications.
- Formulate server selection as an optimization problem that finds an apt server among multiple micro data centers to migrate an application to, so that its performance needs can be met while minimizing the deployment cost of the service provider.
- Describe INDICES (INtelligent Deployment for ubIquitous Cloud and Edge Services), which is a framework that implements DDDAS-based algorithms for online performance monitoring, performance prediction, network performance measurements, server selection and application migration.
- Show experimental results to validate the INDICES framework and evaluate the efficacy of system.

The rest of the chapter is organized as follows: Section 2 presents the INDICES system model and assumptions; Section 3 describes the problem formulation addressed in this research effort; Section 4 delves into details of the solution provided through the methods discussed, including the design and implementation; Section 5 presents empirical results that validate the capabilities sought through the methods presented here; Section 6 compares related work; and finally Section 7 presents concluding remarks referring to lessons learned and future work.

# 2 System Model and Assumptions

This section formally defines and describes the system model used in the present study, and also introduces the assumptions made for formulating the problem.

# 2.1 Components of Application Response Times

Typically, the total end-to-end latency  $t_{total}$  or response time experienced by cloudhosted interactive applications consists of several parts as shown in Equation (1) and which forms the motivation for the work presented here:

4

#### where

- *t<sub>client</sub>* is the processing delay at the client endpoint;
- *t<sub>access</sub>* is the sum of inbound and outbound message transmission delays between the client and its nearest network access point;
- *t<sub>transit</sub>* is the sum of inbound and outbound communication delays between the network access point and the cloud data center;
- *t*<sub>datacenter</sub> is the communication delay from the data center front end (e.g., a web server and load balancer) to the target server in the data center that actually handles the request in both directions;
- *t<sub>server</sub>* is the processing delay at the target server.

Among these costs,  $t_{client}$  and  $t_{access}$  cannot be controlled and managed by the cloud service provider. Also, since  $t_{datacenter}$  is usually less than 1 ms [13], it is practically negligible. On the other hand, a cloud provider can control and manage  $t_{transit}$  and  $t_{server}$ , both of which are key factors in meeting the response time requirements of interactive applications. Note that  $t_{transit}$  is governed by the number of hops incurred by the application messages to traverse the wide area network (WAN) to reach the cloud data center and for the responses to traverse back to the user.

# 2.2 Architectural Model

The work presented in this chapter is geared towards platform-as-a-service (PaaS) cloud providers who seek to meet service-level objectives (SLOs) of soft real-time applications, such as data analytics, online gaming, augmented reality, or virtual desktop, by improving application response times via the exploitation of micro data centers (MDCs); of which such applications have stringent real-time response requirements which put to utmost test of the DDDAS-based systems software capabilities sought, and thus will also satisfy the support requirements of environments needed to support DDDAS-based applications. In doing so, cost considerations and energy savings for the PaaS provider in operating and managing the resources beyond the traditional data centers are also critical issues while ensuring that such an approach provides an additional source of revenue to the PaaS provider. Revenue generation issues are, however, beyond the scope of this chapter.

Figure 1 depicts the INDICES architectural model that consists of a Centralized Data Center (CDC), owned by a PaaS cloud provider. The CDC is connected to a group of Micro Data Centers (MDCs), denoted by  $M = \{m_1, m_2, \ldots, m_n\}$ . These MDCs are deployed at the edge, and are either owned by the CDC provider or leased from an edge-based third party MDC provider. A leased MDC is assumed to be exclusively under the control of the CDC provider. The sharing of MDCs across different CDC providers is not addressed in this work and forms a dimension of the future work. Once an MDC is leased, all its resources are considered to be part of

the CDC provider, and hence customers of the CDC can be transparently diverted to the MDC using their CDC-based security credentials.



Fig. 1 Architectural Model of INDICES (Intelligent Deployment for ubiquitous Cloud and Edge Services)

The CDC comprises a set of compute servers,  $H_{cdc}$ , that can be used to execute applications. The CDC also contains a global manager gm, which is responsible for detecting and mitigating global SLO violations. It is assumed that, for each MDC  $m \in M$ , there exist links to the CDC with a backhaul bandwidth of  $b_m$ . Each MDC m also comprises a set of compute servers,  $H_m$ , that can be allocated to the CDC for its operations at a specified cost. One of the hosts from  $H_m$  or a specially designated MDC host acts as the local manager  $(lm_m)$  for that MDC and is responsible for the DDDAS functionality of data collection, performance estimation, latency measurements and MDC-level decision making, which is the DDDAS feedback loop. This decision-making logic is deployed at the MDC by the CDC provider. For convenience, let  $H_{mdc} = \bigcup_{m \in M} H_m$  denote the set of all servers in the MDCs, and let  $H_{total} = H_{mdc} \bigcup H_{cdc}$  denote the set of all servers from both the CDC and the MDCs. Table 1 lists the key notations in the architecture model.

6

Table 1 Key notations in the architectural model

Notation	Description		
CDC	Centralized Data Center, located at the cloud for execution of applications		
М	Set of Micro Data Centers, located at the edge for execution of applications		
H <sub>cdc</sub>	Set of compute servers (hosts) located at the CDC		
$H_m$	Set of compute servers (hosts) located at an MDC $m \in M$		
H <sub>total</sub>	Set of all compute servers (hosts) located at both the CDC and the MDCs		
H <sub>mdc</sub>	Set of all compute servers (hosts) located at the MDCs		
$b_m$	Backhaul bandwidth to transfer data from the CDC to hosts of $H_m$		

# 2.3 Application Model

For the system model, consider a set of latency-sensitive applications, denoted as Apps, that can be collaborative or single user and interactive or streaming in nature. Each application  $a \in Apps$  is initially deployed in the CDC, with a set  $U_a$  of users, and is assumed to be containerized inside a virtual machine (VM). It is assumed that for a collaborative application a, its users are located in proximity of each other where they incur similar round trip latencies. These scenarios are common when one considers for example collaborative educational applications such as [10] where the users are a group of students working from a school library or a coffee shop, computer vision-based applications in museum and stadium settings [20], or a single user system such as augmented reality-based assisted industrial troubleshooting [29] where image processing operations are performed in the cloud. Table 2 lists the key notations used in the application model.

#### 2.3.1 Application Performance

Each application  $a \in Apps$  can be hosted on any active host in CDC or MDC, i.e.,  $\eta \in H_{total}$  that provides virtualization using a hypervisor or virtual machine monitor (VMM), such as KVM [36] and Xen [2]. Let  $eed_a$  represent the expected execution duration for which the application will be used by the end-user clients. An interactive or streaming application comprises multiple individual interactions between the user and the application that are designated as *streaming steps*.

Each interactive or streaming step of application a consists of both a latency and an execution time. A step takes an estimated execution time  $eet_{a,\eta}$  on host  $\eta$ ; for collaborative applications, which indicates the time needed for all users to have completed that step. Section 4.2 discusses in detail a systematic way of estimating these per-step execution times. In addition, for each user  $u \in U_a$  of the application, let  $el_{a,\eta,u}$  represent the estimated round-trip network latency. Hence,  $ess_{a,\eta}$  denotes the total expected response time of a streaming step, it ensues:

$$ess_{a,\eta} = eet_{a,\eta} + \max_{u \in U_a} el_{a,\eta,u}$$
(2)

For an application a,  $\phi_a$  is defined to be a bound on the acceptable response time for each interactive step of the application. Thus, formally, the SLO for application a hosted on host  $\eta$  should satisfy:

$$ess_{a,\eta} \le \phi_a$$
 (3)

Table 2 Key notations in the application model

Notation	Description				
Apps	Set of applications initially deployed on the CDC				
PA	Subset of applications in <i>Apps</i> that suffer from performance degradation				
$U_a$	Set of users executing application $a \in Apps$				
$eed_a$	Expected execution duration of application $a \in Apps$				
$eet_{a,\eta}$	Estimated execution time of a streaming step of application $a \in Apps$ executed				
-	on host $\eta \in H_{total}$				
$el_{a,\eta,u}$	Estimated round-trip latency experienced by user $u \in U_a$ from a streaming				
	step of application $a \in Apps$ executed on host $\eta \in H_{total}$				
$ess_{a,\eta}$	Total expected response time of a streaming step of application $a \in Apps$				
	executed on host $\eta \in H_{total}$				
$\phi_a$	Bound on acceptable response time for each streaming step of application				
	$a \in Apps$				
$s_a$	Size of the snapshot of application $a \in PA$ to be migrated from CDC to MDC				
$ci_{a,\eta}$	Initialization cost of migrating application $a \in PA$ to host $\eta$				
$transfer_{a,\eta}$	Transfer time for migrating application $a \in PA$ from CDC to host $\eta$ in MDC				
	and for initializing it				
$\delta_a$	Bound on acceptable transfer time for migrating and initializing application				
	$a \in PA$				

#### 2.3.2 Migration Cost

During normal execution, some applications may suffer from performance degradation. *PA* denotes the subset of *Apps* running on the CDC that suffer from performance degradation. These impacted *PA* applications can be identified reactively by the end-user client, which notices missed deadlines using special instrumentation features supplied in the client-side "app" installed as part of the PaaS platform and notifies the CDC service. Alternatively, they can be identified proactively via a predictive decision based on the existing user profiles, where the system predicts that the users are likely to experience SLO violations if they had connected from their profiled location during a certain time period. The goal of the methods presented here is to minimize the number of SLO violations.

To minimize the SLO violations, the system can decide to migrate those applications from the CDC to the hosts in the MDCs, which is a decision made by the DDDAS-based algorithm. Migrating (or transferring) an application incurs two costs: the cost to transfer a snapshot of the application on the new host, and the initialization cost to start the application on the new host.

For an application  $a \in PA$  transferred to host  $\eta \in H_m$ ,  $s_a$  denotes the size of the snapshot of the application (independent of the host). The snapshot is transferred to host  $\eta$  using the corresponding MDC's backhaul bandwidth  $b_m$  and,  $ci_{a,\eta}$  denotes the initialization cost of migrating application a to host  $\eta$  before the application can start processing requests on the MDC host. Once the user-specific state has been transferred, there is minimal interaction between the CDC-based server and the MDC-based server for the remainder of the functioning of application a. In this chapter, the work discussed does not consider further consolidation of resources where applications migrate back to the CDC. The transfer time *transfera*, $\eta$  incurred while migrating application a from the CDC to host  $\eta$  of an MDC is therefore defined as follows:

$$transfer_{a,\eta} = \frac{s_a}{b_m} + ci_{a,\eta} \tag{4}$$

In general, to migrate the application to the new machine, one needs the transfer duration to be small compared to the application's remaining expected execution duration  $eed_a$ . The small duration is a necessary condition to motivate the use of MDC resources and for the ensuing solution to be relevant. To ensure this, the methods developed here do not require sending entire images of the VM or the container from the CDC to MDC. Instead, a layered file system architecture is used at the MDC that is pre-populated with base images used at the CDC as described in Section 4.4. The images assumption is realistic because it is expected that an MDC is either owned entirely or leased exclusively by the CDC provider. The methods developed here are aimed to ensure that the transfer duration is within a threshold  $\delta_a$  defined by the application users before they start to observe improved response time. These requirements are concretized with the following constraint:

$$transfer_{a,\eta} \le \delta_a \ll eed_a \tag{5}$$

Finally, another critical issue that must be accounted for is that any migration of a new application from the CDC to an MDC should not violate the SLOs of the existing applications in that MDC. To model the violation constraint, let  $Apps_{\eta}$ represent the set of all applications currently running on an MDC host  $\eta$ . Then, for each application  $b \in Apps_{\eta}$ , one must verify that its response time bound remains satisfied, i.e.,  $ess_{b,\eta} \le \phi_b$ , after the migration of the new application.

# **3** Problem Statement and its Formulation

This section formally presents the problem statement. Recall that the objective of the methods is to improve response times for cloud-hosted interactive applications that are experiencing performance degradation by migrating them to MDCs. To that end, two key problems must be addressed. First, there must be have a systematic approach for understanding the causes of performance degradation and for determining if an application is impacted. Second, there must be an effective approach by which an

application can be migrated from the CDC to an MDC without impacting existing MDC-based applications while minimizing the cost incurred by the cloud provider. Both these problems are addressed using the DDDAS approach as explained in the remainder of this section.

# 3.1 Performance Estimation Challenges

The performance of an application depends on several factors, including the workload, the hardware hosting platform, and co-located applications that cause performance interference [11, 18, 35]. It is thus important for any solution to account for all of these factors in order to accurately estimate an application's performance in both the CDC and the MDCs. The roles of these factors and the challenges in performance estimation are described next.

## 3.1.1 Workload Estimation

For the cloud-hosted interactive applications of interest here, it is assumed that the workload variation is not significant within a single user session of the service. However, different sessions may have different workloads. For example, in an image processing application, the quality and hence the size of the captured and relayed image may vary depending on different clients' mobile devices. Thus, each workload is considered as a different application setting, which is reflected by the application-specific response time (as described in Section 2.3.1).

## 3.1.2 Hardware Heterogeneity

The CDC and MDCs may consist of heterogeneous hardware platforms and hence each application's performance can vary significantly from one hardware platform to another [18]. Therefore, an accurate benchmark of performance is needed for each hardware platform. DDDAS enables meeting this objective through effective instrumentation, data collection and model building.

# 3.1.3 Performance Interference

Server virtualization platforms such as KVM [36] and Xen [2] provide high degree of security, as well as fault and environment isolations for applications running in virtualized containers, i.e., virtual machines (VMs). However, the level of isolation is inadequate when it comes to performance isolation even though the cloud providers have well-defined resource sharing mechanisms due to two primary reasons:

- Presence of non-partitionable shared resources: VMs can provide isolation guarantees by applying strict CPU reservations and static partitioning of disk and memory spaces. There are also solutions available to limit the network bandwidth. Yet, on-chip resources including cache spaces, cache and DRAM bandwidths, and interconnect networks are difficult to partition [27]. Recently, Intel has introduced Cache Allocation Technology [9] to partition the last level cache (LLC). However, it is still not widely used and cannot be applied to older generation servers. In addition, the presence of shared storage disk is a leading cause of performance degradation [53]. The load imposed on these shared resources by one application is detrimental to all the cache-, memory- and I/O-sensitive applications [46].
- ٠ Hypervisor overhead: The virtual machine monitor or hypervisor has its own overhead. In traditional hypervisors such as KVM and Xen, each virtual machine runs its own operating system which leads to overhead. In addition, the virtual CPUs (vCPUs) can be de-scheduled and virtual RAM can be swapped out without notification. This leads to performance anomalies. In recent years, Linux Container-based virtualization techniques, such as Docker and LXC, have grained transaction due to its low overhead. These resource sharing mechanisms divide the kernel space from the host machine, which alleviates some of the performance concerns but do not provide the same level of performance and security isolations as the VMs [49]. In addition, to maximize the server utilization in shared clusters, cloud providers tend to overbook resources such as CPU cores. This precludes strict CPU reservations and leads to even the lower level caches (L1 and L2) getting shared. In addition, if the overbooked workload goes beyond the server capacity, contention takes place and the applications suffer from performance issues.

## **3.2 Optimization Problem Formulation**

The objective of INDICES' DDDAS-based framework is to assure the SLOs for all the identified applications in *PA* by migrating them to the MDC hosts, while minimizing the overall deployment cost.

#### 3.2.1 Objective Function

To formalize the optimization problem, the following binary variables are defined to indicate the decision taken by INDICES' DDDAS actuation logic for deploying the applications in the set *PA* onto the hosts in  $H_{mdc}$ .

$$x_{a,\eta} = \begin{cases} 1 & \text{if } a \in PA \text{ is deployed on } \eta \in H_{mdo} \\ 0 & \text{otherwise} \end{cases}$$

Thus, the total number of decision variables is  $|PA| \times |H_{mdc}|$ .

The total cost of deploying applications consists of two parts as shown in Equation (6), where the first part indicates a fixed cost incurred due to extending the lease of a host  $\eta$  (if the host is selected), and the second part indicates the cost of deploying an application *a* to host  $\eta$ , which includes both transfer and initialization costs.

$$C^{total} = \sum_{\eta \in H_{mdc}} C_{\eta}^{fixed} + \sum_{\eta \in H_{mdc}} \sum_{a \in PA} x_{a,\eta} \cdot C_{a,\eta}^{deploy}$$
(6)

Each MDC host  $\eta$  involves a monetary allocation cost as it is either leased or could be leased to other providers if owned by the centralized cloud. In addition, running servers involve operational costs, such as the need for power and cooling. Thus, the provider wants to use an MDC server for the shortest duration possible and hence the deployment cost depends on the duration for which the MDC server is used. The cost can be modeled as the extra duration for which the server has to be turned on due to the deployment of the applications in *PA* and it can be represented by a nonlinear function as shown in Equation (7), where the constant  $\alpha_{\eta}$  denotes the cost gradient for powering on host  $\eta$ .

$$C_{\eta}^{fixed} = \alpha_{\eta} \cdot \max\left(0, \max_{a \in PA}(x_{a,\eta} \cdot eed_{a}) - \max_{b \in Apps_{\eta}}eed_{b}\right)$$
(7)

The other objective for the cloud provider is to select hosts which bear minimum transfer and initialization costs. The cloud provider cost is represented as a linear function of transfer and initialization time, as shown in Equation (8), where the constant  $\beta_{a,\eta}$  represents the cost gradient for transferring application *a* to host  $\eta$ .

$$C_{a,\eta}^{deploy} = \beta_{a,\eta} \cdot transfer_{a,\eta} \tag{8}$$

## 3.2.2 Optimization and Constraints

Using the decision variables and the objective function, one can now formulate the optimization problem as an Integer Non-Linear Program (INLP) as follows:

minimize C<sup>total</sup>

subject to 
$$\sum_{\eta \in H_{mdc}} x_{a,\eta} = 1, \forall a \in PA$$
 (9)

$$\sum_{\eta \in H_{mdc}} x_{a,\eta} \cdot ess_{a,\eta} \le \phi_a, \forall a \in PA$$
(10)

$$\sum_{\eta \in H_{mdc}} x_{a,\eta} \cdot transfer_{a,\eta} \le \delta_a, \forall a \in PA$$
(11)

$$ess_{b,\eta} \le \phi_b, \forall b \in Apps_\eta, \forall \eta \in H_{mdc}$$
 (12)

$$x_{a,\eta} \in \{0,1\}, \forall a \in Apps_{\eta}, \forall \eta \in H_{mdc}$$
(13)

The following explains the constraints in the above formulation based on the architectural and application models:

- Constraint (9) restricts each application *a* ∈ *PA* to be deployed on only one MDC host.
- Constraint (10) enforces the response time constraint for each application *a* ∈ *PA* (i.e., Inequality (3)).
- Constraint (11) enforces the deployment constraint for each application *a* ∈ *PA* (i.e., Inequality (5)).
- Constraint (12) ensures the response time constraint for each existing application  $b \in App_n$  in each MDC host  $\eta \in H_{mdc}$ .
- Constraint (13) restricts the decision variables to be binary.

Due to the NP-hardness of the above INLP formulation, the optimization relies on a greedy heuristic to solve it. Section 4.5 describes the proposed heuristic for server selection.

# 4 Design of the DDDAS-based INDICES Framework

This section details the INDICES framework, which solves the optimization problem from Section 3.2. To that end, the solution depends on accurately and reliably estimating: (a) the execution time of the impacted application and network latencies suffered by its clients; (b) similar parameters for the already running applications on different hosts of different MDCs, which are then used in selecting the appropriate host to migrate an impacted application; and (c) the transfer time for migrating the state of the impacted application. This section describes the framework architecture and the detailed techniques for solving the optimization problem.

# 4.1 INDICES Architecture and Implementation

Before delving into the detailed techniques used to solve the optimization problem, the high-level architecture of INDICES is described. Given the scale of the system, a centralized approach to performance prediction and cost estimation for every application hosted in the CDC/MDC and its clients is infeasible. Thus, a hierarchical DDDAS approach is taken, where individual MDCs with their local managers and the global manager of the CDC participate in a two-level decision making process as shown in Figure 1.

Figure 2 shows the local decision-making component of INDICES. Each MDC is composed of a management node and several servers on which the applications execute within virtual machines. Each individual host in the system has a performance monitoring component that logs the data at the local manager  $lm_m$ . The local manager

consists of various DDDAS components: a data collector, a latency estimator, a performance predictor and a cost estimator.



Fig. 2 Local Decision

The performance monitor, which executes the DDDAS instrumentation functionality, instruments the host, and collects system level metrics such as CPU, memory and network utilizations, as well as micro architectural metrics such as instructions per second (IPS) who are retired (out-of-order instructions completed) and cache misses. This information is periodically logged to the local manager for processing. The performance monitoring framework is based on the collectd [24] system performance statistics collection tool. To collect micro architectural performance metrics, the INDICES system uses a Python plugin for collectd using Linux perf. The perf plugin detects if the hardware platform is known, and accordingly executes code that collects hardware specific performance counter statistics. The information is then forwarded to the  $lm_m$  using AMQP [56] message queuing protocol. The  $lm_m$  runs a server developed in the Go programming language, which persists the data in the InfluxDB database, designed specifically for time-series data.

#### 4.2 Execution Time Estimation via DDDAS Model Execution

The constraints in the optimization problem require an accurate understanding of the predicted execution time duration of an application if it were to execute at an MDC, as well as the execution times of the existing applications executing on the hosts of the MDCs. Hence, an application's expected performance profile and in turn its interference profile [31] when co-located with other applications on different hardware platforms given the hardware heterogeneity across the CDC and MDCs is built by defining a DDDAS model using the instrumented dynamic data. Although prior efforts [46, 62, 64] have used retired instructions per cycle (IPC) or last-level cache (LLC) miss rate as the performance indicators, Lo et. al [44] have shown the limitations of these metrics for latency-sensitive applications. Thus, in the present work the *execution time* is considered as the primary indicator of performance.

The *interference profile* of an application [40, 46, 61, 65] is a property that identifies: (a) the degree to which the application will degrade the performance of other running applications on the host – known as *pressure* – and (b) how much the application's own performance suffers due to interference from other applications – known as *sensitivity*. The performance degradation of an application depends, to varying degrees, on different system components and architectures, and other co-located applications. Several prior efforts have used pairwise application execution to estimate their sensitivity and pressure [40, 46, 61, 65]. However, these solutions are not viable for a data center given the significantly large number of hosted applications. Some other efforts [62] pause non-critical applications to measure pressure and sensitivity of live applications, which may not be a realistic solution.

Thus, the generated DDDAS model for a given application *a*, captures the application's performance on a host with hardware configuration *w* and is modeled by Equation (14), where *Y* is the execution time, *X* is a vector of system-level metrics that quantify the state of the host, and the function  $f_a^1$  models the relation between the state of the host machine and the performance of the application *a*. Moreover, the system-level information needed is obtained through Equation (15), which depicts through function  $f_a^2$  the change in the state of the host with hardware configuration *w* if application *a* were to be hosted on it. Equation (15) provides an indirect measure of performance interference, since its output can be used to calculate the change in execution time of an already running application by plugging the new state vector  $X_w^{new}$  into Equation (14) and solving it for each running application.

$$Y = f_a^1(X_w) \tag{14}$$

$$X_w^{new} = f_a^2(X_w^{old}) \tag{15}$$

Another required step is to determine the appropriate DDDAS instrumentation and data collection strategy, which in this case is to identify the right systemlevel metrics to use. Previous works [17, 18, 31] have identified several sources of interference including caches, prefetchers, memory, network, disk, translation lookaside buffers (TLBs), and integer and floating point processing units. Both Intel and AMD architectures provide hardware counters to monitor the performance of

Samor	Hardware	Sockets/Cores/	L1/L2/L3	Mem Type-	Mamory Bandwidth	Count	
Server	Model	Threads/GHz	Cache (KB)	/MHz/GB	Memory Bandwidth	Count	
					64 *(UNC_IMC_NORMAL_READS.ANY		
А	i7 870	1/4/2/2.93	32/256/8192	DDR3/1333/16	+ UNC_IMC_WRITES.FULL.ANY) / time	2	
					in sec		
	Voon				64 *(UNC_IMC_NORMAL_READS.ANY		
В	W3530	1/4/2/2.8	32/256/8192	DDR3/1333/6	+ UNC_IMC_WRITES.FULL.ANY) / time	1	
					in sec		
	Coro2Duo				64 * BUS_TRANS_MEM.ALL_AGENTS		
C	00550	1/4/1/2.83	32/6144/-	DDR2/800/8	*1e9 * CPUFrequency /	1	
	Q9550	550			CPU_CLK_UNHALTED.CORE		
	Ontoron				64 * SamplingPeriod *		
D		2/6/1/2.1	64/512/5118	DDR3/1333/32	DRAM_ACCESSES_PAGE.ALL / time in	9	
	4170HE				sec		

 Table 3
 Server Architectures

micro-architectural components. However, not all the sub-components can always be monitored. Moreover, the list of available counters is significantly smaller for older generation servers. Due to these constraints and driven by the need to support a broadly applicable solution, the following host metrics for performance monitoring are:

- **System Metrics:** CPU utilization, memory utilization, network I/O, disk I/O, context switches, page faults.
- **Hardware Counters:** Retired instructions per second (IPS), cache utilization, cache misses, last-level cache (LLC) bandwidth and memory bandwidth. The bandwidth metrics are not directly available and the counters can vary from one hardware to other. In the analysis, the LLC bandwidth and memory bandwidth were found to be highly correlated and hence the memory bandwidth and not LLC bandwidth is selected due to its easier availability on different architectures and versions. Table 3 lists the hardware counter-based equations for calculating memory bandwidth, which are derived from [1,21].
- Hypervisor Metrics: Scheduler wait time, scheduler I/O wait time, scheduler VM exits. These metrics are the summation for all the executing virtual machines for the KVM hypervisor.

By applying standard supervised machine learning techniques on the collected metrics, Equations (14) and (15) are estimated using the following sequence of steps:

- 1. **Feature Selection:** Feature selection is the process of finding relevant features in order to shorten the training times and reduce errors due to over-fitting. The present work adopted the Recursive Feature Elimination (RFE) approach using Gradient Boosted Regression Trees [22] and applied RFE in a cross-validation loop to find the optimal number of features that minimizes a loss function (e.g. mean squared error).
- 2. Correlation Analysis: To remove the linearly dependent features, correlation analysis is required. This step further reduces the training time by decreasing the dimensions of the feature vector. The Pearson Correlation Coefficient [43] eliminates highly dependent metrics with a threshold of  $\pm 0.8$ .

3. **Regression Analysis:** In this step, curve fitting is performed using ensemble methods with a standard off-the-shelf Gradient Tree Boosting (GTB) method, which is widely used in the areas of web page ranking and ecology. The primary advantage of GTB lies in its ability to handle heterogeneous features and its robustness to outliers.

The performance estimation of the applications, which is effectively conducting what-if-analysis by simulating the developed DDDAS model, consists of two phases: (1) Offline Phase, and (2) Online Phase. The offline phase occurs at the CDC to find estimators, while the online phase is performed by the local manager  $(lm_m)$  of the MDCs to estimate the performance of the target application and also to estimate the performance degradation of the running applications. The two phases are described next.

## 4.2.1 Offline Phase

Whenever the data center receives a request for migrating an application that has not been profiled, it is benchmarked on a single host with a given hardware configuration and then co-located with other applications to develop its interference profile. However, since the number of profiling configurations can be huge, a uniformly distributed subset of possible co-location combinations for profiling is selected. The estimators can be found either by following the three above-listed steps or by choosing an existing estimator of some application based on similarity between the projected performance and the actual performance. A hybrid approach is used here, which first predicts the performance of the new application and its interference profile using estimators of an existing application for the same hardware specifications. If the difference between the measured performance and the estimated performance are within a pre-defined threshold, then the new application is considered to be similar in performance to the existing application. Among all such similar applications, the estimator of the application with the least error is selected for all MDC hardware configurations, which saves profiling time and cost. However, if there is no match, the application profile is developed, i.e., the existing DDDAS model is refined, by performing feature pruning followed by model fitting on each unique hardware platform maintained by the data center.

## 4.2.2 Online Phase

The learned models are then exported and forwarded to the MDC local manager  $lm_m$  for the available hardware platforms in the MDC to estimate the performance of any application to be deployed in the MDC. Since each MDC is small in size and typically illustrates limited heterogeneity in the supported hardware, the number of estimation models will be small. On receiving a request from the global manager gm, the local manager  $lm_m$  estimates the performance of an application by feeding the estimator with presently logged data set using estimation function (14). The

pressure on the existing set of applications  $App_{\eta}$  on host  $\eta$  is calculated by first applying Equation (15) on the target application and then Equation (14) for the existing applications.

## 4.3 Network Latency Estimation

The constraints of the optimization problem also require an accurate understanding of the network latencies incurred by the clients, specifically the worst among all the clients of each application. The latency information is needed for identifying the appropriate host in the appropriate MDC to which an impacted application can be migrated such that it satisfies the SLOs for the worst suffering client while not unduly affecting the existing applications of the MDC hosts.

Thus, estimating the latency to different MDC servers is another key component for achieving the targeted SLOs. To that end, the clients who suffer SLO violations from Equation (3) must be determined. In each client, the instrumented "app" that is installed by the user as part of the client application periodically reports to gmthe application response time it is observing. To not overwhelm the gm, such data logging needs not occur directly on the gm; instead, it can be logged on an ensemble of servers that then report to the gm, or the application server can itself gather data and forward the information when SLO violations occur.

Since there could be multiple MDC choices to migrate an impacted application to, the first step in the algorithm developed in the present work requires reducing the target set of MDCs for latency estimation to decrease the load and amount of time for server selection. To that end, the logged performance data from the clients is used to extract its IP address in order to determine the closest MDCs to that client. The extracted client IP address may not be accurate since often internet users have private addresses and the reported external address is that of the network router or one from the pool of network provider's addresses in case the connection is via a cellular network. However, the extracted IP address information is sufficient, as the client location is used to reduce the set of MDCs that need to be queried. The client's geo-location and consequently its region is derived from the IP address.

The next step is measuring the latencies to the nearby MDCs. To obtain a reliable latency estimate, HTTP-based and TCP socket-based latency measurement techniques are used for HTTP-based and plain TCP-based cloud applications, respectively. Also, additional protocols can be added to this list based on the protocol used. Subject to the collected information, the *gm* forwards to the client app a list of "nearby" MDC gateway servers that are also the local managers  $lm_m$ , which each host a server for the purpose of latency measurement. The client then posts *n* requests to each  $lm_m$  with a file that it typically posts to the cloud for processing (e.g., an image for image processing application) and also the average size of the response it receives from the application. The server responds with a response for the same number of bytes. For each of the *n* interactions, the client records the elapsed time and thus measures  $t_{client} + t_{access} + t_{transit}$ . The client app selects the SLO latency (e.g.,  $95^{th}$  percentile) from the *n* latencies for each  $lm_m$  and reports it to *gm*. The latency assessment also accounts for the delay due to bandwidth size as the actual request data is transferred instead of a ping. It is similar to the *speed test* for measuring the download/upload speeds of an internet provider.

## 4.4 State Transfer Estimation

The final constraint of the optimization problem requires estimating the cost of the state transfer. The local managers calculate the state transfer cost using Equation (4) and use it in local decision making. Once the *gm* selects a host for migrating an application, the application's state has to be transferred before the clients can be switched to the new server location. In this regard, there exist several solutions available for WAN-scale virtual machine migration [8, 52, 55, 59]. The present work leverages the cloud virtual disk format such as qcow2 features for WAN migration. The VM disk is composed of a base image and can contain several overlays on top of it for change sets. The VM overlay when combined with the base image constructs the VM that needs to run for serving the clients.

This base image can contain just an operating system such as Ubuntu (www.ubuntu.com) or an entire software stack such as OpenCV (www.opencv.org) for image processing. The base image is assumed to be present on MDC hosts to save on migration costs and can be shared by multiple VMs. For the target application, overlays are created using external snapshots. The VM overlay is the state that gets transferred to the host and is synthesized with the base image for execution.

Once the application starts running, it informs the gm and all the application clients are redirected to the new application URL. The URL direction happens for custom clients by forwarding the new location to the clients which can then use the new URL for processing. However, for browser-based clients, the communication with the gm occurs via application server due to cross-domain restriction and the existing application issues HTTP-redirect to the new location. In the future, the solution provided through the present work will be enhanced to support live migration of VMs using solutions such as Elijah cloudlet [28] or the recently introduced Docker Linux container's live migration feature [48].

# 4.5 Solving the Optimization Problem at Runtime

The final piece of the puzzle is to solve the optimization problem described in Section 3.2. The problem cannot be solved offline due to the changing dynamics of the system. Moreover, due to the non-linearity and NP-hardeness of the problem, the work here employs a heuristics-based DDDAS decision making algorithm as described in Algorithm 1 (Deployment Server Selection) to solve it efficiently in

an online setting. The algorithm selects aptly a suited server in an MDC while minimizing the overall deployment cost for the entire system.

## Algorithm 1 Deployment Server Selection

1:	<b>Input</b> : set <i>Apps</i> of all applications running on the CDC				
2:	: <b>Output</b> : a server on an MDC for migrating each application $a \in PA \subseteq Apps$ that experiences				
	performance degradation				
3:	for all $a \in Apps$ do				
4:	$ess_{a,cdc} \leftarrow EstTotalExecTime(a, CDC)$				
5:	if $ess_{a,cdc} > \phi_a$ then				
6:	PA.insert(a)				
7:	if $PA = \emptyset$ then return	▹ Do nothing			
8:	for all $a \in PA$ do				
9:	$eed_a \leftarrow GetExpectedExecutionDuration(a)$				
10:	$clientLoc \leftarrow GetLocation(U_a)$				
11:	$nearbyMDCs \leftarrow FindNearbyMDCs(clientLoc)$				
12:	for all $m \in nearbyMDCs$ do				
13:	$H_m \leftarrow GetServerList(m)$				
14:	for all $\eta \in H_m$ do				
15:	$transfer_{a,\eta} \leftarrow EstTransDur(\eta, a)$				
16:	if $transf er_{a,\eta} > \delta_a$ then				
17:	skip $\eta$	Constraint violated			
18:	$perf_{a,\eta} \leftarrow PredictPerfInterf(\eta, a)$				
19:	for all $b \in App_{\eta}$ do				
20:	$ess_{b,\eta} \leftarrow EstTotalExecTime(b, \eta, perf_{a,\eta})$				
21:	if $ess_{b,\eta} > \phi_b$ then				
22:	skip η	Constraint violated			
23:	$ess_{a,\eta} \leftarrow EstTotalExecTime(a, \eta, perf_{a,\eta})$				
24:	if $ess_{a,\eta} > \phi_a$ then				
25:	skip $\eta$	Constraint violated			
26:	$C_{a,\eta}^{total} \leftarrow EstTotalCost(transfer_{a,\eta}, eed_a)$				
27:	$[C_{min}^{total}, \eta_{min}] \leftarrow \min_{n} \{C_{a,n}^{total}\}$				
28:	migrate application a to server $\eta_{min}$				

There are two phases in the algorithm. In the first phase, the set *PA* of applications suffering from SLO violations (Lines 3–6) is identified. In the second phase, a suitable server is selected for migrating each of these applications (Lines 8–28), unless the set *PA* is empty, in which case the algorithm simply returns (Line 7). Otherwise, for each application  $a \in PA$ , the locations of the clients of the application are identified (Line 10), which are used to perform a lookup for the nearby MDCs (Line 11). Then the server within the nearby MDCs that provides the best performance is chosen which is carried out in parallel across all the nearby MDCs (for-loop starting from Line 12).

For each such server  $\eta$ , it must first check the deployment constraint as shown by Constraint (11) (Lines 15–17) and then predict the performance interference of application *a* were it to execute on the server (Line 18). The resulting interference is used to check the response times (Constraint (12)) for the existing set  $App_{\eta}$  of appli-

20

cations on that server (Lines 19–22), as well as the response time (Constraint (10)) for application a itself (Lines 23–25). Finally, the total cost is calculated according to Equation (6) on whether all of these constraints can be met (Line 26). The server with the minimum cost is then selected across all identified MDCs (Line 27), and the application is migrated to the selected server and clients are redirected to the migrated application.

# **5** Experimental Validation

In this section, experimental results for validating the DDDAS-based INDICES framework are presented in the context of a latency sensitive application use case.

# 5.1 Experimental Setup

Table 3 illustrates the hardware platforms and their counts used in the experiments considered in the present study. The CDC uses Openstack (www.openstack.org) cloud OS version 12.0.2 where the guests receive their own public IP addresses. The MDC servers are managed directly by libvirt virtualization APIs (www.libvirt.org) and the guests communicate via port forwarding on the host. Each machine has Ubuntu 14.04.03 64-bit OS, QEMU-KVM hypervisor version 2.3.0 and libvirt version 1.2.16. Guests are configured with 2 GB memory, 10 GB disk, Ubuntu 14.04.03 64-bit OS and either 1 or 2 vCPUs. Since VM migration within a CDC is of not present, the CDC heterogeneity is not a factor.

As described in Section 4.2, to preclude profiling every new application on all the hardware, initial training data are needed. PARSEC, Splash-2 [6] and Stream [47] benchmarks are used to generate the training data to train the INDICES' DDDAS model. PARSEC and Splash-2 target those chip-multiprocessors used typically in modern-day data centers, and provide a rich set of applications with different instruction mix, cache and memory utilization, needed for stressing different system subcomponents. For the present study, 20 tests were selected from these benchmarks for data generation and validation. Also used is the Stream benchmark that specifically targets cache/memory bandwidth, which is one of the key sources of performance interference.

Due to lack of access to servers in different geographical regions, this study used the network emulation tool, *netem*, and hierarchical token bucket-based traffic control, *tc-htb*, for emulating the desired network latencies from the client to the CDC and different MDCs.

## 5.2 DDDAS Target Application Use Case

To validate the efficacy of the framework developed here, an image processing application was used. The application performs feature detection, which is a critical and expensive part of any of the computer vision problem such as object detection, facial recognition [25], etc. Also used the well-known Scale Invariant Feature Transform (SIFT) [45] to find the scale and rotation independent features. The experiment considers an augmented reality application as the use case, where the client device is capturing video frames at a continuous rate. The application augments information to the captured frame by processing the frame. A Minnowboard Turbot (www.minnowboard.org) is used as the client device connected to a webcam that continuously captures frames at a rate of 5 frames per second (fps). The frame resolution is 640X360 pixels and average frame size is 56 KB. The server comprises a Python-based application that receives frames over a TCP socket, processes it, and responds with the identified features along with the processing time. The client expects to receive a response before capturing the next frame, implying that 200 ms is the deadline for the application. Although the use case considers the performance for a single client connected to the cloud-hosted application, it can easily be extended to multiple clients residing in a similar latency region.

When the image processing application is submitted for hosting in the cloud, it is executed on different hardware platforms in isolation to find its base execution times. For hardware platforms A, B, C, D shown in Table 3, the base execution times are measured to be 86, 91, 146, 157 ms, respectively. Table 4 displays the emulated ping latency from this client to CDC or different MDCs in the same region as the client. The table also lists their server composition, and the measured 95<sup>th</sup> percentile network latency while sending TCP/IP and HTTP post requests of 56 KB size and receiving a response of size less than 1 KB. The expected duration for which the client needs to perform the image processing,  $eed_a$ , is set as 1 hour and the SLO is set to 95%.

MDC	Distance	Ping latency	TCP la-	HTTP la-	Servers	
		(±20%)(ms)	tency(ms)	tency(ms)		
$H_1$	1 hop	<1	2	6	1C + 1D	
$H_2$	2 hops	5	14	28	1A + 2D	
$H_3$	Multi hops	20	54	96	1B + 2D	
$H_4$	Multi hops	30	76	142	1A + 3D	
$H_5$	Central	50	127	220	1D	

 Table 4 CDC and MDC setup for application use case

## 5.3 Evaluating the Performance Estimation Model

The use-case application is first benchmarked on hardware platform D in order to develop its performance estimators. The threshold to discern applications with similar interference performance profile, as described in Section 4.2.1, is set to 10% error. However, as illustrated in Figure 3, none of the existing applications meets the criteria. Thus, it was decided not to use any of the existing estimators for the use-case application and benchmark the application on all hardware configurations to develop their estimators. With this approach, it was found that the mean estimation error for the use-case application is less than 4% on all the platforms with low standard deviations as depicted in Figure 4. One can account for this estimation error in the response time constraint (Equation (3)) for stricter SLO adherence.



Fig. 3 Estimation of SIFT Profile Similarity with PARSEC Benchmark



Fig. 4 SIFT Application Performance Estimation Error

## 5.4 Evaluating the Server Selection Algorithm

The server selection algorithm is compared against two other approaches based on minimum number of hops and least loaded server (among reachable MDCs), respectively. Table 4 highlights that the minimum hop is 1. There are 2 servers in the minimum hop MDC  $H_1$  with hardware configuration types C and D. While interference load on both servers, it's ensured that the total load on the server does not exceed its capacity in terms of memory and vCPUs to eliminate unrealistic performance deteriorations. For the least-loaded server algorithm, the server with least existing allocated resources is considered, i.e., containing only a single VM. Servers with no existing load are not considered as it entails acquiring a new server and thus causing additional cost to the service provider. In this case, the server of hardware type D in MDC  $H_4$  is found to be least loaded. Applying SLO from Equation (3) and the server selection algorithm, INDICES finds 2 servers of type A and D from MDC  $H_2$  and one server of type B from MDC  $H_3$ .

Figure 5 shows the response time comparison of each of the suitable servers found by INDICES against the least loaded server for the expected duration  $eed_a$  (one hour) of the application. It's observed that, in this scenario, the least loaded server has 100% SLO violation because of the network latency. However, the servers found by INDICES meet the SLO 100%, 99.38% and 98.94%, respectively, which are all well over the target SLO of 95%. Figure 6 shows the corresponding response time for the minimum hop algorithm. The 2 servers found by the algorithm meet the SLO only 66.64% and 60.64% of times due to performance interference.



Fig. 5 INDICES vs Least Loaded Server Selection

Applying Algorithm 1 further, INDICES finds the server of type B from MDC  $H_3$  to be most suitable, since the objective is to select a server with the minimum cost to the service provider if it can meet the SLO. Thus, it prefers a server which already has an application that is going to run longer and has better bandwidth from the CDC server for migration. Figure 7 compares 3 migration scenarios: (a)



Fig. 6 INDICES vs Minimum Hop Server Selection



Fig. 7 Application Switch-Over Performed by INDICES under 3 Different Scenarios

an overlay with the software stack already present on the target server and the bandwidth is 10 Mbps; (b) same as the previous scenario but with a bandwidth of 1 Mbps; (c) overlay is not present on the target server and the compressed file of size 938 MB has to transferred over 10 Mbps bandwidth. In all the scenarios, the application overlay and configuration files have to be transferred and the application has to be initialized. It is observed that the server selection takes  $\approx 1$  second, but the migration and initialization take 32 seconds, 56 seconds and 190 seconds, respectively, for scenarios (a), (b) and (c). Thus, the overlay-based image transfer should be the preferred methodology whenever applicable.

# 6 Related Work

In this section, the work presented in this Chapter is compared and contrasted with related work along three dimensions: network latency-based server selection, performance interference-based server selection and performance-aware edge computing. Unlike the analysis presented here, the comparative survey conducted has found that existing works seldom consider all these three dimensions holistically. Further,

seldom do any of these prior efforts utilize dynamic instrumentation, data-driven model building and simulation, and actuation in a feedback loop in their entirety as prescribed by the DDDAS paradigm and implemented in INDICES.

#### 6.1 Network Latency-based Server Selection

DONAR [58] addresses the global replica selection problem using a decentralized, selection algorithm where the underlying protocol solves an optimization problem that takes into account client performance and server load. CloudGPS [19] is a server selection scheme that considers network performance, inter-domain transit traffic and server workload for decision making. This work also reduces the network distance measurement costs. Dealer [30] targets geo-distributed, multi-tier and interactive applications to meet their stringent deadline constraints by monitoring individual component replicas and their communication latencies, and selects the combination that provides the best performance. Kwon et al. [41] applied network latency profiling and redundancy for cloud server selection while suggesting using cloudlets. The analysis discussed here contends that these efforts consider simplistic models of server workload and their impact on performance, and do not cater to edge resource management.

#### 6.2 Performance Interference-aware Server Selection

Paragon [18] identified the sources of interference that impact application performance and developed micro benchmarks for heterogeneous hardware. The system benchmarks applications and classifies them to find collocation patterns for scheduling. Sherlock [34] developed approach to estimate the performance interference effect due to last level cache contention in containerized environment. The system profiles application and defines a metric *IScore*, that measures the degree of performance degradation of the application. It does not rely on the hardware counter of the system, and does not take into account the heterogeneity in the hardware platform.

MEDEA [26] allows the users to define anti-affinity application constraints, which captures resource interference characteristics of the applications. These constraints are utilized by the server scheduler for optimizing placement decisions of the services to be deployed in the cluster. SMiTe [65] designed rulers for estimating sensitivity and degree of contention between applications when they are collocated. Bubble-Flux [62] assures QoS for latency-sensitive applications by dynamic interference profiling of shared hardware resources and collocating latency-sensitive applications with batch applications. These works, however, do not apply to virtualized data centers where the hypervisor places its own overhead on the resources and impacts performance. Moreover, INDICES requires virtualized environments to support migration of applications on heterogeneous platforms. DeepDive [51]

first identifies an abnormal behavior using a warning system and employs an interference analyzer by cloning the target VM and running synthetic benchmarks. Such an approach can be a costly runtime operation. Wu, et al. [60], developed a DDDAS-based containerized method for latency-aware collection and processing of wide-are motion imagery, which in addition to the imagery, had to schedule the VMs in parallel to process sections of the extremely large images.

Prior work [11] conducted by the authors designed a performance interferenceaware resource management framework that benchmarks applications residing in virtual machines and applies a neural network-based regression mechanism that estimates a server's performance interference level. However, hardware heterogeneity and per application performance were not considered. Heracles [44] mitigates performance interference issues for latency-sensitive applications by partitioning different shared resources. However, partitioning for resources, such as memory bandwidth is still not available, and moreover, cache partitioning is only available on newer hardware which cannot be applied to existing hardware.

# 6.3 Performance-aware Edge Computing

Zhou et al. [66] described a multi attribute decision analysis algorithm to offload tasks amongst mobile ad-hoc network, cloudlet and public cloud. Their work performs cost estimation considering execution time, power consumption, bandwidth and channel conjunction level which is utilized by the decision making algorithm. The approach utilizes ThinkAir [39] for offloading the tasks. However, they target only Java-based tasks and the solution is not catered to latency-sensitive applications such as those targeted by the present work and its methods.

Fesehaye et al. [23] described a design to select between cloudlets and central cloud server for interactive mobile cloud applications based on the number of hops, mobility and latency. SEGUE [63] is an edge cloud migration decision system that applies state-based Markov Decision Process (MDP) model incorporating network and server states. Both the approaches have not been evaluated on real systems and the results are only simulation-based.

SmartRank. [54] is a tool for offloading facial recognition from mobiles to cloudlets, and the scheduling is performed based on the round trip time and CPU utilization. In the approaches presented here, the cost to the cloud provider is optimized while maintaining the end user SLOs. Recent efforts build on the container system, cloud-fog-edge developments towards a distributed latency-aware imagery distribution system [50].

# 7 Conclusions

This chapter presents INDICES, a framework based on the DDDAS principles for dynamic cloud resource management that exploits the available edge/fog resources in the form of micro data centers, which are used to migrate cloud-hosted applications closer to the clients so that their response times can be improved. In doing so, the INDICES framework ensures that existing edge-deployed services are not unduly impacted in terms of their performance nor are the operational and management costs for the cloud provider overly affected. These objectives are met using an online optimization problem, which is solved using a two-level cooperative and online process between system-level artifacts developed under the work presented here and deployed at both the micro data centers and centralized cloud data center. The experimental results obtained support the capabilities sought to be achieved through the methods presented here by showing the efficacy of the INDICES framework using a realistic application use-case.

This work has opened up many new challenges and directions, which form opportunities for future work. These challenges are presented below:

- Lack of benchmarks: There is a general lack of open source and effective benchmarking suites that researchers can use to conduct edge/fog computing studies. The DDDAS paradigm can be used here to define a collection of models that serve as parameterized benchmarks that can be configured and customized for different application use-cases.
- Collecting metrics under hardware heterogeneity: The plethora of deployed hardware configurations with different architectures and versions makes it hard to collect various performance metrics. Modern architectures are making it easier to collect more fine-grained performance metrics, but much more research is needed in identifying effective approaches to control the hardware and to derive the best performance out of them. It will be interesting to investigate how the dynamic instrumentation capabilities envisioned in DDDAS can be applied in this context.
- Workload consolidation and migration across MDCs: In this work, once an application is migrated to an MDC, it will complete its operation until termination. Future work planned by the authors will consider dynamic server consolidation across MDCs and CDCs. This dimension of future work concerns refining the existing DDDAS decision-making algorithm in the INDICES framework.
- Reconciling application state: In the methods developed here, it was assumed that once the application state is transferred to the MDC, there is no additional state that accumulates at the CDC. However, for a broader set of applications, not all application states may be transferrable to the MDC and may have to be reconciled periodically with the CDC, which gives rise to interesting consistency versus availability tradeoffs. Addressing this requirement will entail potentially using multiple DDDAS loops that interact with each other: one at the CDC and multiple loops one each at the MDC level.
- Distributed user base: Also, the work presented assumed that all distributed users of an interactive applications are located in close proximity to each other.

However, for applications such as online gaming, this assumption may not hold for which additional research will be necessary.

- Energy savings and revenue generation: This work did not address revenue generation issues stemming from the use of edge resources. Moreover, energy savings is only indirectly referred to through the experimental results presented . Addressing these limitations shapes directions for future work. This dimension of future work also involves refining the existing DDDAS decision making algorithm to incorporate additional constraints and refining the objective.
- Shared micro data centers: Also, the methods derived here assumed that an MDC is exclusively controlled by a CDC provider. In the future, it is likely that MDC providers may lease their resources to multiple different CDCs. Additional research is needed to address situations where MDCs are shared including those that address security and isolation guarantees. This scenario has the potential to incorporate dimensions of trust and security into multiple, interoperating DDDAS loops, which will need significant new research.

All scripts, source code, and experimental results for INDICES are available for download from a project called FECBench [5], which provides additional enhancements to INDICES, from https://github.com/doc-vu/fecbench.

# Acknowledgments

This work is supported in part by the AFOSR DDDAS FA9550-13-1-0227 and FA9550-18-1-0126, and NSF US Ignite CNS 1531079. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of AFOSR and NSF.

## References

- 1. Detecting memory bandwidth saturation in threaded applications. URL https://software.intel.com/en-us/articles/ detecting-memory-bandwidth-saturation-in-threaded-applications
- Abels, T., Dhawan, P., Chandrasekaran, B.: An overview of xen virtualization. Dell Power Solutions 8, 109–111 (2005)
- Bahl, V.: Cloud 2020: Emergence of micro data centers (cloudlets) for latency sensitive computing (keynote). Middleware 2015 (2015)
- Balan, R., Flinn, J., Satyanarayanan, M., Sinnamohideen, S., Yang, H.I.: The Case for Cyber Foraging. In: Proceedings of the 10th Workshop on ACM SIGOPS European Workshop, EW 10, pp. 87–92. ACM, New York, NY, USA (2002). DOI 10.1145/1133373.1133390. URL http://doi.acm.org/10.1145/1133373.1133390
- Barve, Y., Shekhar, S., Khare, S., Bhattacharjee, A., Kang, Z., Sun, H., Gokhale, A.: FECBench: A Lightweight Interference-aware Approach for Application Performance Modeling. In: To Appear in the IEEE International Conference on Cloud Engineering (IC2E), p. 10. Prague, Czech Republic (2019)

- Bienia, C., Kumar, S., Singh, J.P., Li, K.: The parsec benchmark suite: characterization and architectural implications. In: 17th international conference on Parallel architectures and compilation techniques, pp. 72–81. ACM (2008)
- 7. Blasch, E., Ravela, S., Aved, A.: Handbook of dynamic data driven applications systems. Springer (2018)
- Bradford, R., Kotsovinos, E., Feldmann, A., Schiöberg, H.: Live wide-area migration of virtual machines including local persistent state. In: 3rd international conference on Virtual execution environments, pp. 169–179. ACM (2007)
- Cache allocation technology improves real-time performance. http://www. intel.com/content/dam/www/public/us/en/documents/white-papers/ cache-allocation-technology-white-paper.pdf
- Caglar, F., Shekhar, S., Gokhale, A., Basu, S., Rafi, T., Kinnebrew, J., Biswas, G.: Cloud-hosted Simulation-as-a-Service for High School STEM Education. Elsevier Simulation Modelling Practice and Theory: Special Issue on Cloud Simulation 58(2), 255–273 (2015). URL http: //dx.doi.org/10.1016/j.simpat.2015.06.006
- Caglar, F., Shekhar, S., Gokhale, A., Koutsoukos, X.: An Intelligent, Performance Interferenceaware Resource Management Scheme for IoT Cloud Backends. In: 1st IEEE International Conference on Internet-of-Things: Design and Implementation, pp. 95–105. IEEE, Berlin, Germany (2016)
- Caglar, F., Shekhar, S., Gokhale, A., Koutsoukos, X.: Intelligent, performance interferenceaware resource management for iot cloud backends. In: 2016 IEEE First International Conference on Internet-of-Things Design and Implementation (IoTDI), pp. 95–105. IEEE (2016)
- Choy, S., Wong, B., Simon, G., Rosenberg, C.: The brewing storm in cloud gaming: A measurement study on cloud to end-user latency. In: Proceedings of the 11th annual workshop on network and systems support for games, p. 2. IEEE Press (2012)
- Chun, B.G., Ihm, S., Maniatis, P., Naik, M., Patti, A.: Clonecloud: elastic execution between mobile device and cloud. In: Proceedings of the sixth conference on Computer systems, pp. 301–314. ACM (2011)
- Cuervo, E., Balasubramanian, A., Cho, D.k., Wolman, A., Saroiu, S., Chandra, R., Bahl, P.: Maui: making smartphones last longer with code offload. In: Proceedings of the 8th international conference on Mobile systems, applications, and services, pp. 49–62. ACM (2010)
- Darema, F.: Dynamic data driven applications systems: A new paradigm for application simulations and measurements. In: International Conference on Computational Science, pp. 662–669. Springer (2004)
- Delimitrou, C., Kozyrakis, C.: ibench: Quantifying interference for datacenter applications. In: Workload Characterization (IISWC), 2013 IEEE International Symposium on, pp. 23–33. IEEE (2013)
- Delimitrou, C., Kozyrakis, C.: Paragon: QoS-aware Scheduling for Heterogeneous Datacenters. In: ACM SIGPLAN Notices, vol. 48, pp. 77–88. ACM (2013)
- Ding, C., Chen, Y., Xu, T., Fu, X.: Cloudgps: a scalable and isp-friendly server selection scheme in cloud computing environments. In: IEEE 20th International Workshop on Quality of Service, p. 5. IEEE Press (2012)
- Drolia, U., Guo, K., Gandhi, R., Narasimhan, P.: Edge-caches for vision applications. In: Edge Computing (SEC), IEEE/ACM Symposium on, pp. 91–92. IEEE (2016)
- 21. Drongowski, P.J., Center, B.D.: Basic performance measurements for amd athlon<sup>™</sup> 64, amd opteron<sup>™</sup> and amd phenom<sup>™</sup> processors. AMD whitepaper **25** (2008)
- Elith, J., Leathwick, J.R., Hastie, T.: A working guide to boosted regression trees. Journal of Animal Ecology 77(4), 802–813 (2008)
- Fesehaye, D., Gao, Y., Nahrstedt, K., Wang, G.: Impact of cloudlets on interactive mobile cloud applications. In: Enterprise Distributed Object Computing Conference (EDOC), 2012 IEEE 16th International, pp. 123–132. IEEE (2012)
- Forster, F.: Collectd The System Statistics Collection Daemon. http://collectd.org (2017)

30

- Forsyth, D.A., Ponce, J.: Computer vision: a modern approach. Prentice Hall Professional Technical Reference (2002)
- Garefalakis, P., Karanasos, K., Pietzuch, P., Suresh, A., Rao, S.: Medea: Scheduling of long running applications in shared production clusters. In: Proceedings of the Thirteenth EuroSys Conference, EuroSys, vol. 18, p. 4 (2018)
- Govindan, S., Liu, J., Kansal, A., Sivasubramaniam, A.: Cuanta: quantifying effects of shared on-chip resource interference for consolidated virtual machines. In: Proceedings of the 2nd ACM Symposium on Cloud Computing, p. 22. ACM (2011)
- Ha, K., Abe, Y., Chen, Z., Hu, W., Amos, B., Pillai, P., Satyanarayanan, M.: Adaptive vm handoff across cloudlets. Tech. rep., Technical Report CMU-CS-15-113, CMU School of Computer Science (2015)
- Ha, K., Pillai, P., Lewis, G., Simanta, S., Clinch, S., Davies, N., Satyanarayanan, M.: The impact of mobile multimedia applications on data center consolidation. In: Cloud Engineering (IC2E), 2013 IEEE International Conference on, pp. 166–176. IEEE (2013)
- Hajjat, M., Maltz, D., Rao, S., Sripanidkulchai, K., et al.: Dealer: application-aware request splitting for interactive cloud applications. In: 8th international conference on Emerging networking experiments and technologies, pp. 157–168. ACM (2012)
- Islam, M.S., Gibson, M., Muzahid, A.: Fast and qos-aware heterogeneous data center scheduling using locality sensitive hashing. In: 2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom), pp. 74–81. IEEE (2015)
- Iyengar, A., Challenger, J., Dias, D., Dantzig, P.: High performance web site design techniques. IEEE Internet Computing 4(2), 17–26 (2000)
- Jarschel, M., Schlosser, D., Scheuring, S., Hoßfeld, T.: Gaming in the clouds: Qoe and the users' perspective. Mathematical and Computer Modelling 57(11), 2883–2894 (2013)
- 34. Joshi, K., Raj, A., Janakiram, D.: Sherlock: Lightweight detection of performance interference in containerized cloud services. In: 2017 IEEE 19th International Conference on High Performance Computing and Communications; IEEE 15th International Conference on Smart City; IEEE 3rd International Conference on Data Science and Systems (HPCC/SmartCity/DSS), pp. 522–530 (2017). DOI 10.1109/HPCC-SmartCity-DSS.2017.68
- Kambadur, M., Moseley, T., Hank, R., Kim, M.A.: Measuring interference between live datacenter applications. In: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, p. 51. IEEE Computer Society Press (2012)
- Kivity, A., Kamay, Y., Laor, D., Lublin, U., Liguori, A.: kvm: the Linux virtual machine monitor. In: Proceedings of the Linux Symposium, vol. 1, pp. 225–230 (2007)
- Koh, Y., Knauerhase, R., Brett, P., Bowman, M., Wen, Z., Pu, C.: An analysis of performance interference effects in virtual environments. In: Performance Analysis of Systems & Software, 2007. ISPASS 2007. IEEE International Symposium on, pp. 200–209. IEEE (2007)
- Kohavi, R., Henne, R.M., Sommerfield, D.: Practical guide to controlled experiments on the web: listen to your customers not to the hippo. In: 13th ACM SIGKDD international conference on Knowledge discovery and data mining, pp. 959–967. ACM (2007)
- Kosta, S., Aucinas, A., Hui, P., Mortier, R., Zhang, X.: ThinkAir: Dynamic Resource Allocation and Parallel Execution in the Cloud for Mobile Code Offloading. In: INFOCOM, 2012 Proceedings IEEE, pp. 945–953 (2012). DOI 10.1109/INFCOM.2012.6195845
- Kuang, W., Brown, L.E., Wang, Z.: Modeling Cross-Architecture Co-Tenancy Performance Interference. In: 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), pp. 231–240. IEEE (2015)
- Kwon, M., Dou, Z., Heinzelman, W., Soyata, T., Ba, H., Shi, J.: Use of network latency profiling and redundancy for cloud server selection. In: 2014 IEEE 7th International Conference on Cloud Computing, pp. 826–832. IEEE (2014)
- Lai, A.M., Nieh, J.: On the performance of wide-area thin-client computing. ACM Transactions on Computer Systems (TOCS) 24(2), 175–209 (2006)
- Lee Rodgers, J., Nicewander, W.A.: Thirteen ways to look at the correlation coefficient. The American Statistician 42(1), 59–66 (1988)

- Lo, D., Cheng, L., Govindaraju, R., Ranganathan, P., Kozyrakis, C.: Improving resource efficiency at scale with heracles. ACM Transactions on Computer Systems (TOCS) 34(2), 6 (2016)
- Lowe, D.G.: Distinctive image features from scale-invariant keypoints. International journal of computer vision 60(2), 91–110 (2004)
- Mars, J., Tang, L., Hundt, R., Skadron, K., Soffa, M.L.: Bubble-up: Increasing utilization in modern warehouse scale computers via sensible co-locations. In: 44th annual IEEE/ACM International Symposium on Microarchitecture, pp. 248–259. ACM (2011)
- McCalpin, J.D.: A survey of memory bandwidth and machine balance in current high performance computers. IEEE TCCA Newsletter 19, 25 (1995)
- Merkel, D.: Docker: lightweight linux containers for consistent development and deployment. Linux Journal 2014(239), 2 (2014)
- Morabito, R., Kjällman, J., Komu, M.: Hypervisors vs. lightweight virtualization: a performance comparison. In: Cloud Engineering (IC2E), 2015 IEEE International Conference on, pp. 386–393. IEEE (2015)
- Nikouei, S.Y., Chen, Y., Aved, A., Blasch, E., Faughnan, T.R.: I-safe: Instant suspicious activity identification at the edge using fuzzy decision making. In: Proceedings of the 4th ACM/IEEE Symposium on Edge Computing, pp. 101–112 (2019)
- Novaković, D., Vasić, N., Novaković, S., Kostić, D., Bianchini, R.: DeepDive: Transparently Identifying and Managing Performance Interference in Virtualized Environments. In: USENIX Conference on Annual Technical Conference, USENIX ATC'13, pp. 219–230. USENIX Association, Berkeley, CA, USA (2013). URL http://dl.acm.org/citation.cfm?id= 2535461.2535489
- Satyanarayanan, M., Bahl, P., Caceres, R., Davies, N.: The Case for VM-Based Cloudlets in Mobile Computing. Pervasive Computing, IEEE 8(4), 14–23 (2009). DOI 10.1109/MPRV. 2009.82
- Shafer, J.: I/o virtualization bottlenecks in cloud computing today. In: Proceedings of the 2nd conference on I/O virtualization, pp. 5–5. USENIX Association (2010)
- Silva, F.A., Maciel, P., Matos, R.: Smartrank: a smart scheduling tool for mobile cloud computing. The Journal of Supercomputing pp. 1–24 (2015)
- 55. Travostino, F., Daspit, P., Gommans, L., Jog, C., De Laat, C., Mambretti, J., Monga, I., Van Oudenaarde, B., Raghunath, S., Wang, P.Y.: Seamless live migration of virtual machines over the man/wan. Future Generation Computer Systems 22(8), 901–907 (2006)
- Vinoski, S.: Advanced message queuing protocol. IEEE Internet Computing 10(6), 87–89 (2006). DOI doi.ieeecomputersociety.org/10.1109/MIC.2006.116
- Wang, Y.A., Huang, C., Li, J., Ross, K.W.: Estimating the performance of hypothetical cloud service deployments: A measurement-based approach. In: INFOCOM, 2011 Proceedings IEEE, pp. 2372–2380. IEEE (2011)
- Wendell, P., Jiang, J.W., Freedman, M.J., Rexford, J.: Donar: decentralized server selection for cloud services. ACM SIGCOMM Computer Communication Review 40(4), 231–242 (2010)
- Wood, T., Ramakrishnan, K., Shenoy, P., Van der Merwe, J.: Cloudnet: dynamic pooling of cloud resources by live wan migration of virtual machines. In: ACM Sigplan Notices, vol. 46, pp. 121–132. ACM (2011)
- Wu, R., Liu, B., Chen, Y., Blasch, E., Ling, H., Chen, G.: A container-based elastic cloud architecture for pseudo real-time exploitation of wide area motion imagery (wami) stream. Journal of Signal Processing Systems 88(2), 219–231 (2017)
- Xu, C., Chen, X., Dick, R.P., Mao, Z.M.: Cache contention and application performance prediction for multi-core systems. In: Performance Analysis of Systems & Software (ISPASS), 2010 IEEE International Symposium on, pp. 76–86. IEEE (2010)
- Yang, H., Breslow, A., Mars, J., Tang, L.: Bubble-Flux: Precise Online QoS Management for Increased Utilization in Warehouse Scale Computers. In: ACM SIGARCH Computer Architecture News, vol. 41, pp. 607–618. ACM (2013)
- Zhang, W., Hu, Y., Zhang, Y., Raychaudhuri, D.: Segue: Quality of service aware edge cloud service migration. In: 8th IEEE International Conference on Cloud Computing Technology and Science (CloudCom). IEEE (2016)

- Zhang, X., Tune, E., Hagmann, R., Jnagal, R., Gokhale, V., Wilkes, J.: Cpi2: Cpu performance isolation for shared compute clusters. In: Proceedings of the 8th ACM European Conference on Computer Systems, EuroSys '13, pp. 379–391. ACM, New York, NY, USA (2013)
- Zhang, Y., Laurenzano, M.A., Mars, J., Tang, L.: Smite: Precise qos prediction on realsystem smt processors to improve utilization in warehouse scale computers. In: 47th Annual IEEE/ACM International Symposium on Microarchitecture, pp. 406–418. IEEE Computer Society (2014)
- Zhou, B., Dastjerdi, A.V., Calheiros, R.N., Srirama, S.N., Buyya, R.: A context sensitive offloading scheme for mobile cloud computing service. In: Cloud Computing (CLOUD), 2015 IEEE 8th International Conference on, pp. 869–876. IEEE (2015)