



Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

ACAS: An anomaly-based cause aware auto-scaling framework for clouds

Sara Kardani Moghaddam*, Rajkumar Buyya, Kotagiri Ramamohanarao

Cloud Computing and Distributed Systems (CLOUDS) Laboratory, School of Computing and Information Systems, The University of Melbourne, Australia



HIGHLIGHTS

- Proposing an anomaly prediction module to detect future performance problems.
- Applying unsupervised Isolation-Trees based approach with a new strategy for model updating.
- Investigating types of performance anomalies including CPU and memory bottlenecks.
- Proposing a two-level auto-scaler by categorizing anomalies into two levels local and global.
- Extending CloudSim simulator by adding characteristics of the elastic VMs.

ARTICLE INFO

Article history:

Received 14 September 2017
 Received in revised form 9 April 2018
 Accepted 12 December 2018
 Available online 26 December 2018

Keywords:

Cloud computing
 Anomaly detection
 Performance management
 Isolation-trees
 Vertical scaling

ABSTRACT

Cloud computing as a model to deliver distributed resource and services on the pay-as-you-go policy has become increasingly popular for all organizations including industry. However, the inherent dynamicity in this environment makes it prone to various types of performance problems which introduce many challenges in the area of distributed resource management. Advances in the big data learning approaches can bring the opportunity for a data aware dynamic management of resources in the cloud. The collected data from the performance indicators of the system can be a valuable source of information to identify unusual behaviors in the resource consumptions or application performance. Different types of problems can cause the performance degradations at VM or system level. System administrators are overwhelmed with the huge amount of data to be analyzed to find the problems and overall health of the system. In this paper, we argue that a better selection of dynamic resource scaling policies can be employed for better performance by predicting the anomalies in the system and narrowing down the possible cause of the anomaly to one of the attributes of the system. Therefore, we propose a 2-level cause aware auto-scaling framework which leverages two types of resource management solutions, horizontal and vertical, as the corrective actions when the performance is degraded. We show the effectiveness of vertical scaling strategy as a quick solution for cases that a VM is exposed to some type of the local anomaly, while the horizontal scaling solutions can be used for system wide anomaly to add new VMs in the system. Moreover, our data analysis module can predict anomalies to give sufficient time to the scaling system to make an effective scaling decision. The proposed unsupervised anomaly detection module leverages a new updating strategy for renewing the models which considers the changes in the state of the system to reduce the overhead of recurrent model trainings. We have performed a comparison of the proposed framework with an approach which is used by several popular cloud providers to show the advantage of mixing the multi-level auto-scaling with the knowledge of anomaly detection analysis in resolving performance problems in the cloud.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

Cloud computing as a model to deliver computing infrastructure and services is bringing both opportunities and challenges in the area of distributed resources management. A wide range of

customers from individuals to commercial companies are leveraging the power of offered services in an on demand pay-as-you-go model. On the other hand, public providers such as Amazon and Google aim at satisfying the demand of their customers as well as meeting the service level agreements (SLA) knowing that the violation of SLAs can cost them money and their reputation.

Elastic resource management offers the providers the ability to dynamically adjust the resources based on the type and number of requests from the users. For example, existing auto-scalers

* Corresponding author.

E-mail address: skardani@student.unimelb.edu.au (S.K. Moghaddam).

such as Amazon elastic auto-scaler [1] enable the system to dynamically add and remove Virtual Machine (VM) instances as a response to the observed performance degradations in the system. Workload fluctuations targeting hosted applications are one of the main underlying reasons for these performance problems. This is a highly important observation, especially for the large scale web application systems where the interaction between users and web servers can change frequently, affecting the pattern of workloads and resource requirements. On the other hand, there is a variety of problems that can happen locally in one VM such as a bug in the application code, resource bottlenecks or hardware faults. This type of problems can affect the local performance of the VM adversely. Distinguishing these faults from system wide problems can help auto-scalers to make more informed decisions by focusing on the solutions that target directly the root cause of the anomalies. To achieve this goal, we can divide the data-aware resource management problem into two main subproblems, data analysis and resource management by auto-scaling, that can be dealt with separately.

First, it should be mentioned that different types of performance problems in the VMs usually leave distinctive signs in the performance indicators of the machine. Therefore, continuously monitoring the behavior of resources by collecting the values of important attributes provides system administrators a valuable source of the data that can be analyzed to have timely information about the performance of the system. Big data learning approaches offer the necessary concepts and tools to analyze the collected data and find interesting patterns of unexpected behaviors or anomalies encountered by the system.

The second part of the problem focuses on the auto-scaling solutions to be triggered when a performance problem is identified by analyzing the collected data from the system. There is a variety of resource management solutions including horizontal scaling, elastic VM management, migrations, resource contention management, etc. for alleviating the performance degradations. However, when the scaling actions should be triggered and which type of the scaling action selected are different challenges which are investigated in this paper.

With regard to the aforementioned challenges, we propose an Anomaly and Cause Aware auto-Scaling (ACAS) framework consisting of three main modules, monitoring, data analyzer, and resource auto-scaler which exploits two types of the resource adjustment policies, horizontal and vertical scaling. In this work, we focus on the local anomalies such as CPU and memory bottlenecks as well as system wide load problems that can affect the performance of the applications. Accordingly, the **contributions** of this paper are as follows:

- Proposed a proactive, unsupervised anomaly detection on predicted performance data of the VMs to identify future anomalies of the system. We have also developed a strategy for deciding when the anomaly detection model need to be updated to reduce the recurrent model training overheads
- Proposed a two-level auto-scaler by categorizing anomaly problems into two levels of local and global anomalies. Local anomalies are targeted by elastic VM solutions while system-wide anomalies are resolved by the global scaler
- Achieved better scalability by breaking down the problem of performance management to local and global anomaly detection
- Extended CloudSim simulator functionality by adding characteristics of the elastic VMs
- Performed extensive experimental evaluation of the proposed system under various loads and demonstrated ACAS ability to maintain a better quality of performance compared to the existing resource management solutions

An extensive set of experiments are performed targeting both types of local anomalies and global load problems. The experiments show that distinguishing between VM specific anomalies and system wide load problems help auto-scaler to take advantage of fast vertical scaling policies to increase bottleneck resources for one VM while the proactive anomaly detection helps to trigger early system wide horizontal scaling actions to reduce the number of SLA violations.

The rest of this paper is organized as follows: Section 2 introduces some of the existing works in the field of data-aware resource management. Section 3 presents the motivation, an overview of the approach and a brief introduction to Isolation-Trees based anomaly detection. Section 4 presents the details of learning algorithms and explains communication among the modules. Section 5 presents the experiments and the results and finally, the paper is concluded with the future directions in Section 6.

2. Related work

The idea of utilizing data learning techniques for the performance analysis in the cloud has been of great interest to the researchers in recent years. The work presented in [3] investigates the feasibility of Isolation-Trees based anomaly analysis to detect anomalies in data from IaaS data centers, focusing on the behavior of the algorithm to the presence of seasonality/trends in their dataset. ACAS also leverages the same concept of Isolation-Trees in the anomaly detection part of the problem. However, ACAS is a complete framework that covers the problems of online learning and model updating, root cause analysis and resource management modules. [9] proposes a method for long-term load prediction in Google data centers, considering load as the main factor involved in resource management solutions. Another work presented in [5] considers a single attribute, number of required processors at a certain time, for resource utilization estimation. [22] presents a regression based workload prediction framework to improve the utilization of the resources while reducing the cost. To achieve this goal, they use the knowledge from workload prediction to decide the time and amount of resources to be changed in the system, considering both types of vertical and horizontal scaling. [23] combines workload prediction and reinforcement learning to find the best configuration for VM resources. The feedbacks from application performance and resource utilizations are used to calculate the reward and update the resource configuration strategy for better selection of future actions. Compared to our framework, the aforementioned works address the problem of resource management by focusing on the workloads as the only influential factor for performance analysis and ignore other sources of performance problems in the system. [10] follows a more systematic approach to the problem of VM management in the cloud by modeling the problem as a feedback-based control approach. The Proportional-Integral-Derivative (PID) based controller is designed to manage the number of VMs in the system, aiming at keeping the service quality in accordance with the agreement levels. [2] designs a reinforcement learning approach to gradually learn from the environment and decide on the VM level scaling of the system to alleviate the performance problems occurring due to the load fluctuations in the system. Different to our model, these works consider the management of resources only at VM level by changing the number of VMs in the system. [13] presents an automatic anomaly identification technique for adaptively detecting performance anomalies such as Disk and Memory related failures. Proposed method investigates the idea that a subset of the principal components of metrics can be highly correlated to specific failures in the system. BARCA, proposed in [6], is another framework for online identification of anomalies in distributed applications which divides anomaly detection process into two

steps. First, a one class classifier is employed to distinguish normal behavior from unexpected ones. Second, a multi-class classifier is used to separate different types of anomalies from detected abnormal behaviors. [12] investigates proactive anomaly detection in data stream processing systems. The proposed solution includes a phase of predicting resource utilization and then applying an anomaly identification algorithm on the predicted values. The target anomalies are injected and the training is done on a labeled dataset of different anomaly occurrences in the past data. Although these works focus on the same problem of anomaly detection, how this information can be used for resource management is not investigated. Alternatively, [21] addresses the performance problem by integrating a 2-dependent Markov model as the predictor and tree-augmented Bayesian networks (TAN) for anomaly detection. Based on the knowledge from learning algorithm, they apply some type of the vertical scaling or migration to minimize the performance degradation. [14] focuses on the cost effectiveness of vertical scaling approaches and proposes a threshold based scaling strategy to combine different scaling approaches including self-healing, fine-grained resource scaling and VM level scaling to meet QoS while reducing cloud providers' costs. [18] addresses the problem of shared memory management among multiple VM with the over-subscription approach and elastic VM technique. A threshold based strategy on the value of memory related metrics is utilized to trigger the memory adjustment actions while live migration is used to avoid the SLA violations when the total memory demands of the VMs exceed the available memory of the physical machine. In contrast, our proposed work ACAS focuses on the effectiveness of horizontal and vertical scaling policies by leveraging the capabilities of unsupervised learning approaches for situations that system is exposed to the local and load related anomalies. Another study by [8] investigates unsupervised behavior learning problem for proactive anomaly detection. The proposed framework leverages Self-Organizing Maps (SOM) to map a high dimensional input space (performance metrics) to a lower dimensional map without losing the structural information of original instances. In contrast, we show that resource management process can make use of the knowledge from proactive anomaly detection and root cause identification to address the specific anomalies occurring in each VM.

Table 1 compares the above mentioned works by highlighting the main components and characteristics of the proposed solutions for the problem of performance analysis and management in large distributed systems.

3. Preliminary

In this section, we explain the motivation and an overview of our approach. Then, a short introduction to the concept of isolation-based anomaly detection is presented.

3.1. Motivation and approach overview

Virtualization technologies are the core concept in the functionality of cloud models. The possibility of running many VMs and applications on one physical host brings new opportunities, as well as adds more complexity to the design of these environments.

Resource management concept is much more challenging due to the inherent dynamic nature of cloud environment, where we can host different range of applications with a variety of demands and workload types. This is especially important for the large scale web applications, in which the pattern of the incoming requests from the users can change quickly creating a dynamic environment where the configurations of resources should frequently be adjusted to satisfy the demanded SLAs. Elastic resource management, as a solution for this problem, leverages the VM based scaling of the

system known as horizontal scaling. Public cloud providers offer customized policies of horizontal scaling to satisfy the resource demands of the applications based on the load in the system. Even though the VM based scaling policy is a common approach to manage performance problems in the cloud environment, it may not be a proper solution for a different category of performance problems caused by the faults in one VM. For example, consider a situation where a memory-intensive process is started in the same VM hosting a web based application which is consuming all the available memory, ignoring the demands of the web application. Therefore, the lack of the free memory can cause performance degeneration such as longer than usual response times from the web server. The conventional scaling approaches add new VMs into the system even though the problem is not caused by the load growth from a higher number of user requests. Having the same load in the system, newly added instances incur extra costs including both resource and license costs as well as higher resource wastage due to added resources which are not utilized. Considering this scenario, there are other types of the problem that can create similar effects on the utilization of the resources. For example, it is shown that the web applications are prone to many of the performance problems which involves CPU and memory resources [20].

On the other hand, existing auto-scaling solutions such as Amazon elastic auto-scaler are designed in a way that are more suitable to track the changes at the system level. They do not consider VM based problems particularly when the changes in one VM does not have an immediate impact on the average performance of the system. One solution to address this category of problems is to have a resource management solution at fine-grained levels of control. Elastic VM architecture enables on-the-fly tuning of VM resources without turning off the VM which avoids the delays of rebooting the system. Given the above explanations, we formulate the problem as the selection of proper resource scaling policy to satisfy the quality of service (QoS) by analyzing the state of the system to distinguish resource level bottlenecks from system wide load problems. To be more explicit about the system state, we use the following definition:

System State: State or behavior of the system at each time is an abstract representation of operational attributes and performance indicators of the system which can be recognized in normal or abnormal/anomalous condition. The main indicators of an abnormal state are the presence of unexpected patterns or values in the load and resource level measurements of VMs and applications.

The proposed framework addresses the resource management problem at the service provider level who has access to the VMs hosting the application to monitor system and application level metrics. In this paper, we target a category of performance anomalies known as resource bottlenecks and particularly two problems insufficient CPU and memory in one VM. Therefore, by tracking resource level metrics of VMs, one can utilize vertical scaling functionality to increase the amount of RAM capacity or the number of CPU cores of one VM to quickly respond to the performance degradations of the system. When there is a system level degradation, the framework employs horizontal scaling to add new VMs into the system.

In the next section, the components of ACAS framework for cause aware auto-scaling in the cloud are explained in more detail. However, before that, we present a brief introduction to IForest algorithm which is the key function used in the proposed anomaly identification approach in the data analyzer module. It should be highlighted that several other algorithms can also be used for the problem of anomaly detection [7,13,15]. However, as we explain more in the following section, IForest is extremely efficient and have been shown to work with a wide variety of distributions of data and do not require prior knowledge of these distributions which makes it a notable alternative to be utilized for real-time tracking of anomalous behavior.

Table 1
Related works on cloud performance management.

Work	Data analysis method	Resource management	Proactive	Unsupervised ^a	Vertical scaling
[3]	IForest (AD) ^b	X	X	✓	-
[9]	Bayes Model (Workload Analysis)	✓	✓	-	X
[10]	Control Theory	✓	X	-	X
[2]	Reinforcement Learning	✓	✓	-	X
[6]	SVM (AD)	X	✓	✓	-
[12]	Markov models, Bayes Classifier (AD)	X	✓	X	-
[21]	Markov models, Bayes Classifier (AD)	✓	✓	X	✓
[14]	Threshold-Based Rules	✓	X	-	✓
[18]	Threshold-Based Rules	✓	X	-	✓
[23]	Reinforcement Learning	✓	✓	-	✓
[8]	Self-Organizing Maps (AD)	X	✓	✓	-
Proposed work (ACAS)	Isolation-based Trees (AD)	✓	✓	✓	✓

^aThis column is applicable for the works with the focus on anomaly detection.
^bAnomaly Detection.




Fig. 1. A simple Isolation Tree for two attributes CPU and Memory.

3.2. A brief summary of isolation-trees

The concept of isolation based anomaly detection which is used in Isolation-Forest (IForest) algorithm [16] is based on the idea that an anomaly instance is isolated in attribute space faster than a normal instance. IForest leverages this definition to create an ensemble of random binary trees on the attribute space of the problem with the expectation that an anomaly instance is isolated quickly close to the root node of the tree.

In order to generate a random tree, IForest first randomly selects a sample of ψ instances and starts by selecting a random attribute with a random value based on the measurements in the selected sample. This step corresponds to creating one node in the tree. Then, sample instances are divided into two categories based on their values for the corresponding attribute. Each category is assigned to a new child node and this process continues for newly created nodes until there is only one instance left or all the instances have the same values or the length of the tree has reached the predefined maximum value. Fig. 1 shows a simple Isolation Tree generated on a sample space of 7 instances, each representing the utilization values of CPU and memory of one VM. Let $X = (x_1, x_2, \dots, x_7)$ be the set of the input observations for the IForest algorithm. To create the root node, the value of c_1 is selected based on the observed utilization values of CPU attribute in X . As Fig. 1 shows, the instance x_1 is the only instance with a utilization less than c_1 which creates a leaf node on the left side of the root node, while the remaining instances create the right child node. The right child node includes more than one instance with different utilization values. Therefore, another attribute is selected to divide the remaining instances, creating two new sub nodes.

This process continues until the terminating conditions are met. In this example, instance x_1 can be a possible anomaly as it seems to be in a different range of CPU utilization values compared to the other instances.

In order to build a model, IForest generates t Isolation trees from the training data. An anomaly score can be calculated for each new observation which represents the degree of anomalousness of the record compared to the past observations used for the training of the models. Each instance traverses all the randomly generated trees, starting from the root of the tree until it reaches a leaf node. Then, the score is calculated by averaging the length of the traversed paths on all the trees and using a formula presented in [16]. The anomalous score is inversely proportional to the average length of the branches. Shorter the length, more likely the instance is anomalous. The authors demonstrate that the algorithm has a low linear-time complexity with small memory requirements. The algorithm’s worst time and space complexity for training is $O(t\psi^2)$ and $O(t\psi)$ respectively, where t is the number of trees in the ensemble. Since t and ψ are small values, both training and testing can be performed quickly. The extensive comparisons of IForest done in [16] with other state-of-the-art anomaly detectors such as Random Forests or Local Outlier Factor (LOF) demonstrate the superiority and robustness of IForest in terms of the accuracy as well as the processing time. Moreover, the authors show that the IForest performs near optimal (largely insensitive to the values of t and ψ) with the default parameter settings which is consistent with our observations during ACAS experiments. This is a very important feature, especially for highly dynamic environments, where the frequent updating of the models is required. For these reasons, we have chosen to employ IForest for our anomaly detection for each VM.

4. System design

Fig. 2 depicts an overview of the proposed framework and how the components work together. The framework is modeled based on a web based application with the application and database servers hosted on the cloud VMs. These applications are known for the exposure of many performance degradations caused by the changes in the workload or CPU and memory related faults. However, the definitions are generic and can be applied to any distributed application. The components of the application can be distributed on different VMs, while each VM has its own monitoring component, data analyzer and local scaler modules installed. The data analyzer box in Fig. 2 shows the details of the local analyzer module on each VM. The scaling decisions are performed at two levels, local and global. The local scaler is responsible for the vertical scaling decisions at one VM, while the global scaler performs horizontal scaling decision in the system. The global scaler and the load balancer are parts of a separate master node which plays as

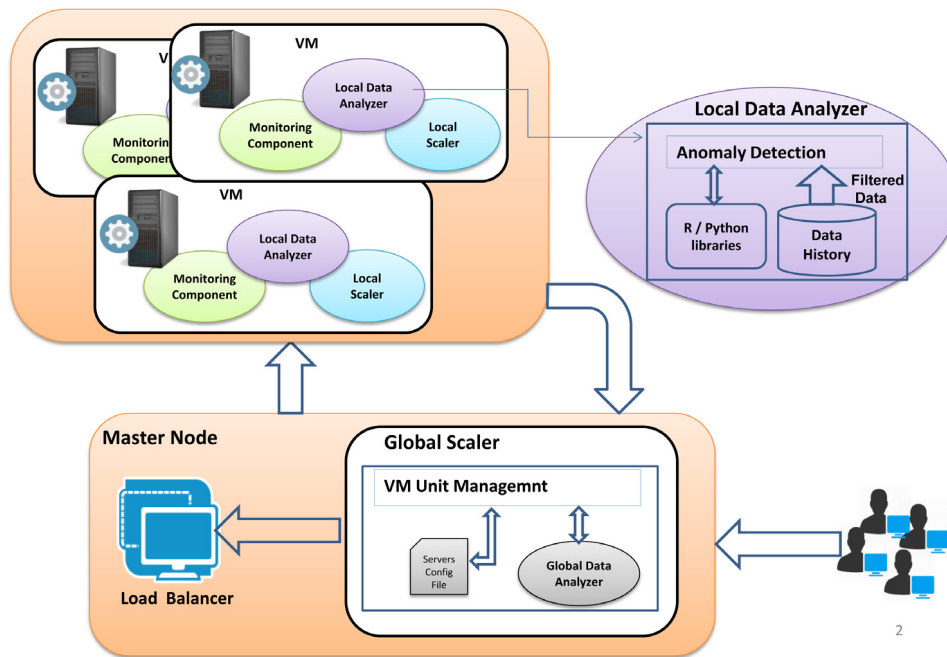


Fig. 2. A high level system model.

the central broker for the whole system. Therefore, the incoming requests are distributed among existing VMs (application servers) based on the load balancer configuration and registered servers at the master.

Each VM monitors the performance of its own resources and collects a variety of attributes such as CPU and memory utilization, and disk I/O rates which can model the state of the system. During regular intervals, collected data are sent to the local data analyzer to be processed for the possible signs of performance problems occurring in the near future. Therefore, at the first step, future values of each metric are predicted. There is a wide range of algorithms that can be used for the prediction and modeling of time series data. We have tested two algorithms ARIMA and feed-forward Neural Networks (NN) for this step and finally selected NN due to the observed stability of its predictions in the presence of the noise in our dataset. NN is utilized to generate a separate model for each metric and predicts the future values based on the learned models from the past observations of the system. Upon receiving the newly predicted values, the anomaly detection algorithm calculates an anomaly score for each observation and sends a new alert if the new score exceeds the threshold θ . Before proceeding, we should remark that anomaly detection module considers every deviation in the values of the attributes from the past state of the system as an anomaly which is reflected in the calculated anomaly scores. However, from the service providers perspective, a performance anomaly is important when it shows a possible breach of the SLA objectives; otherwise, it can be ignored. Therefore, to be clear about an anomaly event which is considered by the resource management module for taking the corrective actions we pursue the following definition in the next sections:

Anomaly Event: A continuous change in the behavior of the system which is reflected as unexpected trends in the values of the monitored attributes of the VM while at least one of the metrics shows the possibility of breaching the threshold for the maximum accepted utilization.

The first part of the aforementioned definition is handled by anomaly detection module to detect the attributes that show a transition in their state based on the details provided in Section 4.1. The second part of the definition confines the performance anomalies to the anomaly events that are breaching the performance

thresholds. This part is considered by resource management module as described in Section 4.2.

During anomaly detection phase and at the time of observing anomalous behavior, the system asks the cause detection module to analyze the state of different observed metrics and find a possible cause for detected anomalies. The suggested causes of the problem from this module are used as additional knowledge in the auto-scaler components to help them make more informed decisions regarding the scaling policies.

The results from anomaly detection module are sent to the local and global auto-scaler components. The local scaler is responsible for resource configurations at VM level also known as vertical scaling policies. In contrast, the global scaler is aware of the state of the whole system and is responsible for changing the number of VMs in the system known as horizontal scaling policies. Algorithm 1 shows a summary of the main steps of ACAS framework at the local and global level. The details of these steps and the priority of different scaling policies are explained in the following algorithms and subsection. A list of all the notations used in following sections are listed in Table 2.

4.1. Anomaly prediction based on isolation-trees models

Given one VM measurements, the goal is to find if the collected values show a different pattern compared to the past behavior (lines 3–7 Algorithm 1). Therefore, having a sequence of past observations from one VM, an ensemble of Isolation-Trees is generated using the IForest algorithm. After the training is done and each VM has the initial models of its performance, the anomaly detection process starts to analyze the new measurements collected from the VM. Algorithm 2 shows the sequence of required steps for the process of anomaly prediction in ACAS. This process is called regularly to check the recent performance of the VM.

In order to give the system enough time to trigger auto-scaling actions, we need to detect anomalies in the future data. Therefore, the first step is to predict the future values of each metric for the VM (lines 1–2). NN algorithm is exploited as the prediction function (f_p) to forecast the w values of each metric based on the recent measurements from the system. Predicted values are fed as

Algorithm 1: Cause Aware Resource Scaling in ACAS

```

input      :  $V = (VM_1, VM_2, \dots, VM_M)$ : A list of all
              registered VMs in the system
1 while The system is running and in the beginning of
  performance-check interval do
2   for  $VM_i \in V$  do
   /* This part of the code is executed
   locally in each VM */
3   if  $VM_i$  has not initialized the IForest models and
   there are enough data collected for training then
4     | Initialize IForest models for  $VM_i$ ;
5   end
6   Collect the recent monitored values for different
   metrics of  $VM_i$ 
7   Call Algorithm 2 on the collected observations to
   predict future data and find the possible
   performance anomalies and suggested causes;
8   call Algorithm 3 to check if  $VM_i$  requires a new
   vertical scaling to be done by local scaler; If
   scaling is done,  $VM_i$  goes into a locked state for a
   predefined time.
9   end
   /* This part of the code is executed in
   the master node */
10  Initialize all indicators in  $f_i$  to 0;
11  for  $VM_i \in V$  do
12    | if  $VM_i$  is not in a locked state and is moving to
    | critical condition based on the Algorithm 4 then
13    | |  $f_i \leftarrow 1$ 
14    | end
15  end
16  Decide on a new horizontal scaling action based on
   the information provided by  $f_i$ ;
17 end

```

Table 2
Description for notations.

Notation	Description
K	Minimum number of alerts before an anomaly record is created
w	prediction window size
lw	Number of observations in prediction learning window
tw	Number of observations in training window
L	Minimum number of violations before the threshold based approach (baseline) starts an action
r	Number of attributes for each observation
Ll	Log Time for monitoring system to record a new observation
θ	Threshold for anomaly score. Values greater than θ will be considered as anomaly
X^m	A record of monitored metrics (attributes) from environment
th_i	Usage threshold for attribute i
f_i	An indicator of anomalouslyness of attribute i
S	Anomaly scores
ψ	Number of randomly selected samples from input instances as the input of IForest algorithm

the inputs to the trained models which calculate an anomaly score for each predicted record. The anomaly scores show the degree of abnormality of the observations compared to the data used in training phase (line 3).

It should be highlighted here that we expect to encounter cases where ACAS may miss some of the anomalies due to the wrong measurements or wrong predictions resulting from the dynamic nature of the target environment. Therefore, ACAS also considers more reactive mechanisms which try to adjust the scores of anomaly points when a violation in the system is detected. To make this point clear, let S_i be the score for the prediction P_i . ACAS checks

if S_i actually reflects the violation observed at time t_i and if it does not (meaning that $S_i < \theta$ and $P_i \geq th_i$), it deliberately increases the score S_i to a higher value so other components of the framework handle situation as a new anomaly state.

4.1.1. Model updating

One question to be answered is how the system decides to update the anomaly detector models. The inherent dynamicity in cloud workloads and the possibility of different types of failures highlight the importance of updating models so they can show the most recent state of the system. In this regard, three different states of the system are distinguished as follows:

- **Transition State:** The system is recognized as in transition if it meets two main conditions. First, newly observed values differ from the past training data in the patterns and/or values. Therefore, we expect to see higher anomaly scores calculated to show the abnormality of recent behavior compared to the historical records. Second, the system has not reached a stable state, meaning that a continuous change of the variables is still observable. The focus of this paper is on the transitions which cause the average values of the attributes to change with the assumption that the patterns remain unaffected. For example, consider a situation that an incremental trend is continuously impacting the values of one of the attributes in the system.
- **Changed State:** The system has reached the changed state when the new observations show deviations compared to the recorded data used for the training. However, the system has reached a stable condition meaning that no significant changes in the average values of the attributes are detected. In terms of the conditions mentioned for the transitions state, a system at changed state satisfies the first condition only.
- **Normal State:** The system is at the normal state when none of the above conditions is satisfied, meaning that the average values of the attributes for recent observations do not show significant changes compared to the training data. As a result, the calculated anomaly scores do not indicate any abnormal behavior demonstrating a stable environment.

The anomaly detection module decides to update the model if it finds the corresponding VM is at the *changed* state (lines 6–12). The reason is that, at this state, the high number of anomaly alerts shows the previously trained models are not representing the current state of the system. Moreover, the system has reached a new stable environment and new models are required to enable the anomaly detection module to perform in accordance with the changes. The updating procedure continues until the new models correctly reflect the new state or another transition in the system starts. It should be mentioned that ACAS does not consider *transition* state a proper time for updating the models as some of the attributes are showing significant changes in their values and new models quickly become obsolete, resulting in many unnecessary updates.

4.1.2. Cause identification

The cause detection procedure tries to provide some knowledge about the possible resource level root causes of the performance problem to help the scaling modules make a more informed decision about the proper scaling policies. Therefore, if the output scores from anomaly detection module show a possible anomaly is occurring in the VM, the next step is to identify the underlying reason for the problem. The category of changes addressed in this paper are the ones that impact the average values of the attributes with an increasing or decreasing trend. Therefore, to find an attribute with a trend in the values, we follow an approach which

Algorithm 2: Anomaly Detection

```

input      :  $D = (X_1^m, X_2^m, \dots, X_{lw}^m), X_i^m \in R^{1 \times r}$ : A matrix
              of  $lw$  records, each record including
              measurements for  $r$  features
Parameter:  $w$ : Prediction Window
               $\theta$ : Anomaly Score Threshold
output    : (Anomaly Alert, Cause of Anomaly)
1  $c \leftarrow -1$ 
   /* Prediction function  $f_p$  is used to
   predict future values of data.  $X^m$ 
   corresponds to the Measured data and  $X^p$ 
   presents Predicted data. */
2  $(X_{lw+1}^p, X_{lw+2}^p, \dots, X_{lw+w}^p) = f_p(X_1^m, X_2^m, \dots, X_{lw}^m)$ 
3  $S_i = \text{AnomalyScore}(X_{lw+i}^p), 0 < i \leq w$ : Find the anomaly
   scores with IForest algorithm. Then, check if these
   scores should be adjusted based on a reactive approach
   if some violation is already happening in the system.
4  $\text{anomalyDetected} \leftarrow (\text{Count}(S > \theta) > \text{Length}(S)/2)$ 
5 if  $\text{anomalyDetected}$  then
6   Initialize all indicators in  $f_i$  to 0
7   for feature  $i \in D$  do
8     if system is in changing state on dimension  $i$  then
9        $f_i \leftarrow 1$ 
10    end
11  end
12  Decide about updating the models based on the
   information provided by  $f_i$ .
13  Identify the cause of abnormality and assign it to  $c$ .
14 end
15 return ( $\text{anomalyDetected}, c$ )

```

fits a regression line on the data and calculates the slope of the line as a measure of the existing trends in the data.

One point worth noting here is how to distinguish load problems from other local anomalies. One observation to be followed is that when the performance of the system is impacted due to the changes in the incoming workload, we expect to see more than one attribute affected and changes their state. Accordingly, ACAS checks whether most of the attributes in the system are recognized at the *transition* state simultaneously and then flags the anomaly as a load problem.

4.2. Resource management module

The management of resources in a continuously changing environment requires the integration of resource configuration policies at different layers of granularity. Depending on the type of the problem and identified root causes, some policies may work better at meeting the SLA objectives such as time or cost of the solution. In this paper, two policies horizontal and vertical scaling of the resources are considered. Horizontal policies address resource configuration strategies which change the number of active VMs in the system. In contrast, vertical policies are defined at finer grains of control (Elastic VMs) and adjust the amount of allocated resources based on the new demands of the VM. Since the scaling happens online and there is no need to reboot the instance, vertical scaling is much faster and does not add extra costs for a software license or wasted resources.

Upon receiving an anomaly alert from anomaly prediction module, the framework should create a new record to flag the beginning of a new anomaly event in the target VM. However, we need to consider the transient changes in the system that may cause false alarms. As a result, a new anomaly event is recorded at time t if the

current observation is showing an anomaly alert as well as all the past observations in the window $\{t - K, t - K - 1, \dots, t - 1\}$. In other words, the system ignores the first K alarms for one VM until there will be at least $K+1$ consecutive alerts notifying an anomalous behavior. A proper value for K can be selected considering the trade-off between computation overheads, the stability of the environment and the performance degradation tolerance. Small values of K may cause the system to perform unnecessary checks of the performance or decide on preventive actions for many false alarms, while large values of K increase the time it takes for the system to start a scaling action in response to the performance problems.

In the proposed framework, some conditions should be met before resource manager decides on a new scaling action for the system. The following subsections and Algorithm 3 explain these conditions.

Algorithm 3: Vertical Scaling Policy

```

input      :  $\text{counter}$ : Number of recent alerts for the VM
input      :  $\text{anomalyDetected}$ : True if recent anomaly
              score exceeds threshold
input      :  $\text{cause}$ : The root cause detected for the
              current anomaly
Parameter:  $K$ : Minimum Number of Alerts to Record an
              Anomaly
               $\theta$ : Anomaly Score Threshold
   /* reset the counter when the system is in
   normal state. */
1 if  $\text{anomalyDetected}$  then
2    $\text{counter} \leftarrow \text{counter} + 1$ 
3 else
4    $\text{counter} \leftarrow 0$ 
5 end
6 if  $\text{counter} > K$  then
7   if system is not in cooling period &&  $\text{cause} \neq \text{Load}$ 
   then
8     If system is moving toward critical condition
     based on Algorithm 4, start a vertical scaling
     action.
9   end
10 end

```

4.3. Per-VM vertical scaling policies

After receiving a confirmed anomaly event for one VM, the VM starts to check if some type of the resource adjustment is required. ACAS considers scaling strategy only when a performance degradation or SLA violation is observed. In this case, we consider the breach of the resource utilization thresholds as a sign of the violation of SLA objectives. Let th_i be the threshold for resource i . If the utilization of this resource at time t is more than th_i , system records a violation of SLA starting from time t . Therefore, no corrective action is triggered if there are enough spare resources to fulfill the requests during next time intervals. One question to be answered here is that what is the best time interval to predict the future utilization of resources. Since the online resource adjustments in elastic VMs become effective almost immediately, we take one time interval away from the recent observation as the prediction interval. Therefore, the framework sends back the list of all metrics that are predicted to violate their respective thresholds at the next time interval.

Since vertical scaling is a response to local anomalies happening in a VM, the load problem is ignored at this step and local resource adjustment is triggered if the detected problem is related to one of the resource level metrics of the VM. Depending on the metric

Algorithm 4: Identification of System Criticality

```

input      :  $D = (X_1^m, X_2^m, \dots, X_{lw}^m)$ ,  $X_i^m \in R^{1 \times r}$ ; A matrix
              of  $lw$  records, each record including
              measurements for  $r$  features
input      : cause: Root cause detected for current
              anomaly
Parameter: LI: Log Interval
1 delay  $\leftarrow 0$ 
2 if cause  $\neq$  Load then
3   | delay  $\leftarrow$  VerticalScalingDelay
4 else
5   | delay  $\leftarrow$  HorizontalScalingDelay
6 end
7 windowLength  $\leftarrow$  delay/LI
8  $P = (X_{i_{w+windowLength}}^p) = f_p(X_1^m, X_2^m, \dots, X_{lw}^m)$ 
9 for feature  $i \in P$  do
10  | if  $P_i$  exceeds  $th_i$  then
11  |   |  $f_i \leftarrow 1$ 
12  | end
13 end
14 Decide about the criticality of system based on the
    information provided by  $f_i$ 

```

detected as the root cause of the problem, the system decides about changing the number of CPU cores or the amount of memory capacity of the VM to prevent performance degradations in the application. After starting an auto-scaling process, the VM will enter in a locked state which means that during this time no other scaling action is performed. The reason is that it takes some time for the system to adapt to the changes of the resources, so the first few anomaly alerts are ignored to give the system enough time to reach a stable state.

4.4. Horizontal scaling policies

A horizontal scaling policy is performed if there are no VM in the locked state, meaning that there has not been any vertical scaling in the recent intervals that can affect the state of the system. First, the state of all VMs is checked and the number of VMs which are moving toward critical condition is recorded. One VM is recognized in critical condition if at least one of the main attributes is predicted to breach the threshold in the near future. Similar to vertical scaling procedure, we consider a rough estimate of the time it takes to boot a new VM in the system as prediction interval. In other words, ACAS asks for enough time to add a new VM before the system enters the anomaly state. If all the active VMs are found moving toward the violation state, an alert to add a new VM is issued. Afterwards, the system starts a cooling period when no scaling will take place. This waiting time is required so the load balancer can detect new VM and start sending new requests to that.

5. Performance evaluation

The proposed framework incorporates multiple components from resource monitoring, resource configuration and data analysis. The framework is general and should be applicable to different types of applications and workloads. However, in order to demonstrate the effectiveness of ACAS, we select web applications which are shown to be prone to many performance problems involving CPU and memory resources [20]. The main focus of this work is the performance of the application layer which can be easily affected by the behavior of users, buggy codes or other malfunctioning applications.

To validate the framework, we use CloudSim discrete event simulator [4] which is a framework for modeling and simulation of cloud computing infrastructures. CloudSim has been used extensively for validation of cloud services and applications that can be hard to be validated in real implementation as we need a controlled environment where one could perform analysis of the system with and without data analysis or auto-scaling methods including elastic VMs. An extension of the CloudSim is leveraged that implements an analytical performance model of 3-tier applications in the cloud and multi-cloud environments [11]. CloudSim offers both flexibility as well as the extensively validated models of reference workloads that helped us to create a near real environment.

5.1. Experimental settings

The experimental environment is simulated as one cloud data center hosting the application and database servers. The application servers are modeled with the initial configuration of one virtual core, 3.75 GB of RAM and Linux operating system. The VM start-up times are modeled based on the performance study done by [17].

The following experiments are based on an extension of CloudSim which models the workloads on Rice University Bidding System (RUBiS) benchmarking environment [19]. RUBiS is a benchmark that implements the core functionality of an auction site including browsing, bidding and selling modeled based on eBay.com. RUBiS follows a 3-tier web based framework consisting of the client, application and database servers. Sessions are the unit of works defined in the RUBiS and represent a sequence of requests from one customer interacting with the application. The resource usage of each session is monitored and modeled in the CloudSim based on the work done by [11]. In total, there are 4 attributes CPU, memory, I/O usages as well as the number of sessions which are collected during each experiment for data analysis part. For the details of how the workload is modeled and validation of the extracted models you can refer to the work [11]. To implement the prediction step, we utilize *forecast* package implemented in R¹ which models a feed-forward neural network with lagged inputs for forecasting univariate time series. The final prediction is an average of the results from 20 trained networks; each network is trained on lag-1 of all input values. Therefore, each network has one input (with a bias node), one hidden layer with one node (with a bias node) and one final output node which is analogous to AR model but with a non-linear function. The averaging on the all networks helps the prediction result to be more robust in the presence of noise.

In order to demonstrate the functionality of ACAS in resolving local performance problems with the help of fine grained resource scaling, we have also extended CloudSim framework to enable the on-the-fly changes of the resource configurations without turning the VM off. Two main resource types CPU and RAM are considered in this implementation. However, the codes are general and can easily be extended for other types of the resources. The amount of changes in each scaling action can be configured to be a percentage of the original capacity of the resource. For the following experiments, the capacity of CPU resources increases by one core (100% of initial configuration) while the RAM storage is increased by 20% for each scaling action. Moreover, the anomaly detection models in ACAS are generated using IsolationForest package implemented in R environment.² In order to connect the anomaly detection module to the simulation environment which is developed in JAVA, we

¹ <https://www.rdocumentation.org/packages/forecast/versions/8.1/topics/nnetar>

² <https://sourceforge.net/projects/iforest/>.

Table 3
Experiment configurations.

Variable	Description	Value
K	Minimum number of alerts before an anomaly record is created	6
lw	Number of observations in learning window	60
tw	Number of observations in training window	300
L	Minimum number of violations before baseline approach starts an action	2
LI	Monitoring Interval (Log Time)	60
θ	Anomaly Score Threshold	0.55

utilize *ReqEngine* interface which supports calling R implemented functions from Java environment.

Each experiment has a duration of 18 h with sessions arrival time modeled as a Poisson distribution with a frequency that is defined as a function of time [11].

In order to evaluate different aspects of the proposed framework, four cases of experiments have been run. In two cases, the behavior of ACAS is tested in the presence of local anomalies in the VMs. In two other cases, the system is exposed to workload increases and the functionality of the framework is analyzed. Two types of the resource level bottlenecks, insufficient memory and CPU, are simulated. In both cases, we focus on the impact of increasing trends on the corresponding attribute. In order to simulate memory problems in CloudSim, a predefined percentage of the memory storage is removed from the available memory at consecutive interval times which creates an incremental trend in the used memory of the VM. For insufficient CPU, a predefined percentage of the available CPU capacity is flagged as reserved assuming a different CPU-intensive application starts running as a background process along with the target application.

The idea of performance anomaly detection has been widely investigated in the research area. However, most of them follow supervised approaches or are designed for specific scenarios or focus on data analysis part of the problem without providing the details of an integrated framework for the purpose of resource management. On the other hand, many of popular public cloud providers such as Amazon [1] use a threshold based auto-scaling approach for dynamic scaling of their resources. In the threshold based approaches, system continuously tracks the state of the resources in the system and an anomaly alert is triggered if the utilization of monitored metrics exceeds a predefined threshold. For example, a new machine is added to the system if the CPU utilization is more than 80% for five continuous sampling intervals. Therefore, for the comparison purpose, we have implemented the same threshold method as our baseline approach. To have a comparable experiment, the thresholds for the baseline auto-scaler are the same as the triggering thresholds of ACAS framework. In the all experiments, this value is equal to 70% and is similar for both CPU and memory. A cooling period of 15 min is considered for the baseline simulation. Therefore, no two auto-scalings are performed in a time interval less than the cooling period. Table 3 shows the values of parameters used in the experiments.

5.2. Experiments and results

In the first experiment, we investigate the behavior of ACAS based on a sample workload similar to Fig. 3. The experiment starts by sending requests to a load balancer which distributes the load among application servers on a round robin basis. In order to start training the models, we follow the observations from [16] which suggests that 2^8 generally is enough to consider as the sample size (ψ) for training phase. Considering this and based on the nature of the dataset and empirical experiments to apply IForest as an online anomaly detection algorithm, 300 is selected as the

Table 4

Number of times that resource utilization exceeds the threshold before the first auto-scaling action is triggered. NA means no scaling is performed.

Anomaly Type	Algorithm	
	ACAS	Threshold Method
CPU	1	NA (> 100)
Memory	10	NA (> 100)
System Load	5	8

training window size (tw) to be considered for the sampling and training purpose. Therefore, the anomaly module waits for the first 300 observations to pass and then initializes the first anomaly detection models by training IForest algorithm with the last 300 records as it is shown in Fig. 3. The first 30 records are ignored for the system to stabilize.

After the first initialization, anomaly detection module starts to regularly check the performance of the system by applying the generated anomaly detection models on the recent collected observations at the configured time intervals (presented as a while loop in Algorithm 1). However, depending on the state of the system and based on the definitions discussed in Section 4.1, models may need to be updated occasionally to represent the new state of the system. As we can see in Fig. 3, after the first model initialization, a low rate increase of the incoming load is started which corresponds to a transition state based on our definitions. Therefore, the first update of the models recorded for this experiment is occurring around 435th observation when the system is identified at the end of the transition and entering a new normal state. Similarly, other updates occur occasionally during the experiment due to the fluctuations in the utilization data. However, there are also several gaps that no update has occurred during that times. These gaps are consistent with our observations of the stability of average utilization data and the functionality of ACAS which has not detected any *transition* that requires new model trainings. For example, there is no update between observations 570 to 640 or there are only 7 updates between observations 645 to 760. The reduction in the number of updates helps the system to decrease the overhead of recurrent trainings to create new models. The same procedure with similar reasoning is applicable for the next load increase, starting around observation 800, that changes the state of the system from *normal* to *transition* and also triggers an auto-scaling action which adds a new VM to the system.

The next experiments are designed to test the presence of the local anomalies in VMs. The initial configuration is done by adding 3 application and 2 database servers in the system. Then, one VM is randomly selected as an anomalous VM. For both experiments of CPU and memory anomaly, we wait for a minimum of 5 h and then, at a random time, the injection of the anomaly in the VM is started. Table 4 shows the number of recorded observations that the attribute corresponding to the detected root cause exceeds the threshold. NA in the table means that there was no auto-scaling action in the response to the injected fault in the target VM which is equal to a 100% violation of the SLA. The exact number of the violations depends on the duration of the corresponding anomaly in the system which may last for hours.

For more clarification of the way these policies are reacting in the presence of local performance problems, Figs. 4 and 5 present the utilization of the corresponding attribute for each fault and for both ACAS and threshold methods. Regarding CPU, the ACAS has increased the number of cores by one as soon as it predicts the criticality of the CPU utilization measurements. One violation is observed in this case which is a result of the fast changes in the attribute values which the prediction function has not caught. In contrast, the threshold approach is monitoring the average state of the whole system, missing the local faults occurring at the

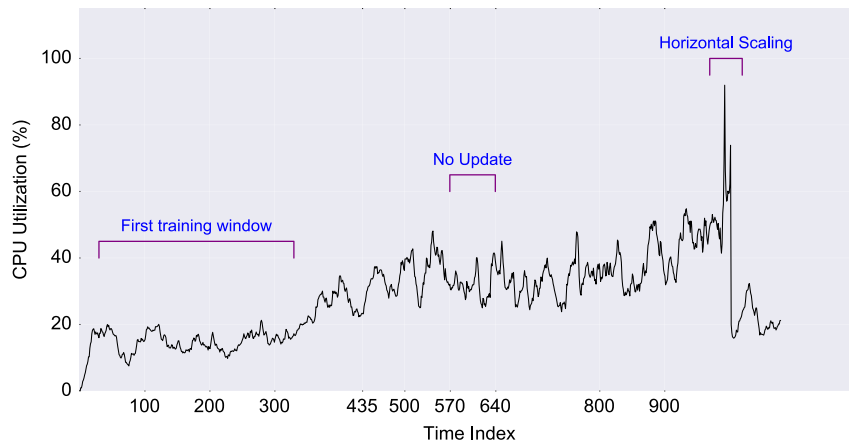
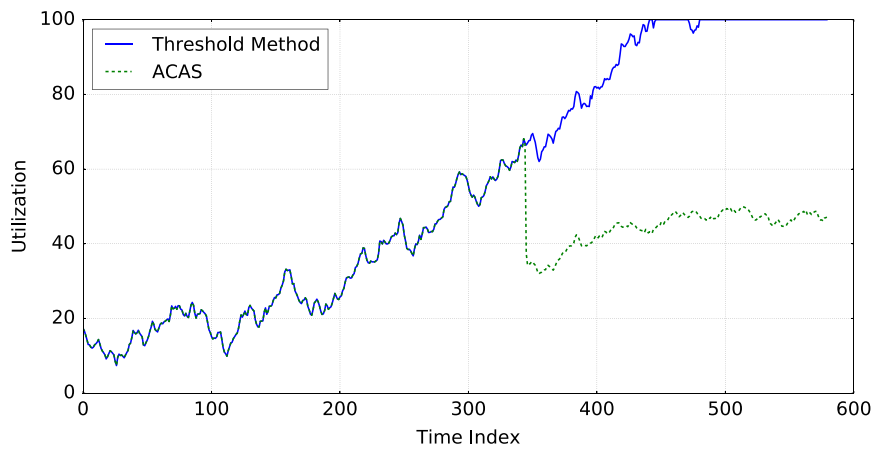



Fig. 3. The process of ACAS on a sample workload including the first training window and one horizontal scaling action. One part of the data that is analyzed with the same models (no model update occurred during this time) is also annotated.



(a) CPU Utilization




(b) Response Time (Log scale)

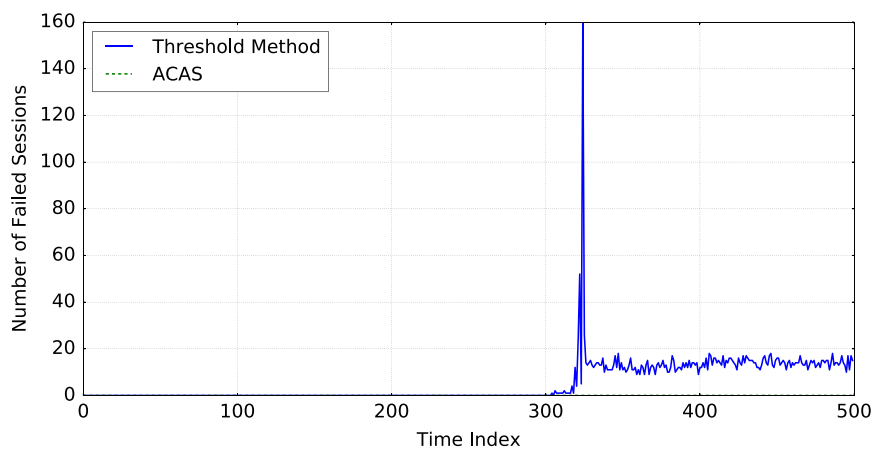
Fig. 4. Vertical auto-scaling for CPU bottleneck. ACAS avoids high response times by timely reaction to the predicted performance problem.

anomalous VM. It is worth mentioning that even having a per-VM monitoring mechanism for the threshold approach can only help to trigger a horizontal scaling with the condition that the monitored values show a minimum of L violations before auto-scaler starts triggering an action. The L value should be chosen reasonably to avoid unnecessary scalings in the presence of temporal changes in the system. In our simulations, L has a small value equal to 2.

However, depending on the application instability, this value can be higher which leads to even more violations. This situation is a result of the lack of the knowledge about preceding trends to the anomaly state. ACAS solves this problem by keeping the track of the patterns in the data and performing the scaling when the conditions of being in a continuous anomaly state and the violation of the threshold values are met.



(a) Memory Utilization



(b) Failed Sessions

Fig. 5. Vertical auto-scaling for Memory bottleneck. ACAS avoids failed sessions by timely reaction to predicted performance problem (ACAS line for the failed sessions is zero for duration of the experiment).

The above reasoning is also applicable for memory bottlenecks. One point to mention is that the simulated RUBIS application shows a CPU intensive behavior. Therefore, memory usage has fewer fluctuations and shows more clear change points which can be detected with higher accuracy. Fig. 5 shows two sequential vertical scalings of memory which adds 20% of the initial capacity each time. The first scaling happens before any violation is observed which shows the prediction part of ACAS helps the scaler to perform a proactive action to predict the future anomaly events and start a corrective action. The results show that the memory usage drops down by 20%. However, the utilization continues to increase which causes the start of the second scaling action. This time, however, a few numbers of violations of the memory usage are observed. The reason is that for a small duration of time after the first scaling, the system is recognized in a new *changed* state which is followed by an update of the models. Therefore, the initial increases in the memory do not trigger anomaly alerts which cause the system to start the second action after some delays. In this case, the reactive part of the approach helps the system to detect the anomaly state when the violations are observed.

Fig. 5 also shows the number of failed sessions for both policies. A session is flagged as failed if the VM does not have enough memory to process its requests. As the figure shows, in the experiments with the threshold method, the number of failed sessions has increased as a result of ignoring the local fault in the VM. In contrast, ACAS has properly adjusted the configurations corresponding to

the bottleneck resource which avoids the unusual increases in the failed sessions.

As demonstrated by aforementioned experiments, the proactive vertical scaling helps to quickly target the bottleneck resource and reduce the number of violations by adjusting the amount of resources accordingly. This process also helps to reduce the cost as well as the energy consumptions compared to the conventional way of adding new VM machines in the system. It is also worth noticing that the local execution of anomaly detection reduces the complexity of training the anomaly detection models. As it is explained in Section 3 the time and space complexity of IForest algorithm is constant when the same number of training observations is used for model generation.

The next set of the experiments analyzes the behavior of the system when the input workload of the machines suddenly increases. Two types of the problem have been considered. The first experiment simulates an environment where one VM is exposed to an increasing workload while other VMs in the system stay in their normal state. Therefore, one VM is randomly selected and the number of the requests sent to this VM is increased. This scenario can happen in different cases such as a result of a misconfigured balancer service which assigns a higher weight to one VM. Fig. 6 shows the impact of the load increase on the response time of the target VM for both policies. As we expect, the threshold approach is not successful at detecting the local performance problem and many violations of the response time are observed. In contrast,

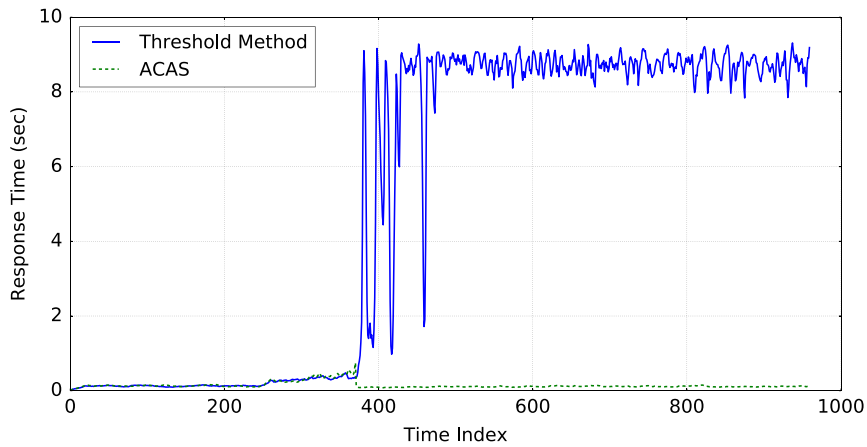
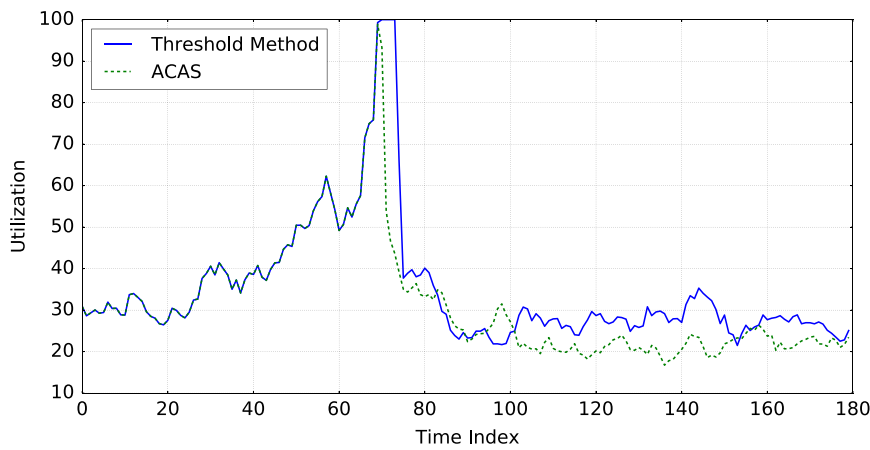



Fig. 6. Response time of one application server when the machine is overloaded.



(a) CPU Utilization



(b) Response Time

Fig. 7. CPU Utilization and Response Time of one application server when the system is overloaded. ACAS is able to proactively trigger a horizontal scaling action compared to reactive response of the threshold method which causes more SLA violations.

the local anomaly detection approach utilized by ACAS helps to identify the problem as soon as the metrics show an increasing trend followed by exceeding the thresholds.

The second experiment for the load problem simulates an overloaded system where the number of incoming requests to the balancer is increased, resulting in the increase in the resource usage

of every machine at the same time. This scenario is a common case in the web applications known as flash crowds when sudden surges in the traffic to a web site causes high delays in the response time making it virtually unreachable for the users. As Fig. 7 shows, both policies make similar decisions and add a new VM after the problem is recognized. However, ACAS approach is able to react to




Fig. 8. CPU utilization of one application server when the machine is overloaded. The marked points are the records detected as anomaly.




Fig. 9. Detected anomaly points and the model update times for the duration of the experiment. Red points show the observations that detected as an anomaly. Blue points show the times that a model update occurred in the system.

the problem immediately at the same time that the first breach of the threshold is detected which causes the system to return back to the normal state after 5 observations of the violation of CPU and memory metrics. In contrast, the threshold approach does not have a knowledge of the past behavior of the system and therefore delays the triggering of the auto-scaling action for L observations. In our experiments, this value is set equal to 2 which results in about 8 violations before the system goes back to the normal state. Larger values for L , more SLA violations in the system.

Finally, a set of the plots presenting the relation between anomaly scores and model updates are shown for a sample experiment in ACAS framework. Fig. 8 shows the CPU utilization of one application server. The marked points are the observations recognized as anomalies, meaning that the corresponding anomaly scores are higher than 0.55. Fig. 9 presents a combined view of the anomaly detection process for the same workload, including detected anomaly points along with the anomaly update times. Each point at the top line shows that the observation at the corresponding time was detected as an anomaly, while the gaps between these points reflect normal or transition states of the system. The first 330 points are ignored as they are used during the training phase and detection process was not activated at that time. Similarly, each point at the bottom line shows that a model update happened at the time of the corresponding observation. As we can see, at the times that the system is recognized in the normal state, no update is occurring, meaning that the models are reflecting the current state of the system. Another observation

from these figures indicates that the updates are delayed when an anomaly event is started while the system is recognized as being in the transition state. An example of this condition can be seen between observation 900 to 1100 which is reflected by the gaps among the points at the bottom line.

6. Conclusions and future work

Elastic VMs with the accompanying knowledge from analyzing performance data can bring new opportunities to offer better resource management solutions in the distributed environment. In this work, we show how fine grained resource configurations can help to improve the auto-scaling solutions for a category of local anomalies occurring in one VM. The proposed ACAS framework utilizes a low overhead anomaly detection solution based on the Isolation-Trees and combines it with a cause identification procedure to enable appropriate auto-scaling solution taking into consideration the nature of the anomaly.

For the future work, we plan to improve the cause detection method with the aim of detecting more complex anomaly problems, the ones that can affect the pattern of the data in long run or influence multiple attributes. As an example of more complex types of anomalies, distributed attacks on the cloud resources are a challenging problem, especially considering the other sources of data such as packet information which should be included along with the previously discussed resource level measurements. In this case, we may need to add a profiling step to record a signature of

old and new anomalies, categorizing them based on the underlying reason for the problem. Moreover, we can work to further improve the framework by adjusting the thresholds for both anomaly scores and attribute dependent violations automatically based on the type of the anomaly, the intensity of changes in the data, etc. Finally, we plan to extend our work to consider alternative virtualization methods such as containers in which the same Operating System (OS) is shared among multiple container-hosted applications. While the proposed ACAS framework is applicable to these environments and we expect similar performance of our framework when applied to containers at the resource level, the nature of containers may introduce some additional requirements in terms of measuring resource utilization with respect to each container. Therefore, the isolation of performance may become more challenging which requires finer levels of data analysis such as monitoring system call traces and interactions among containers to alleviate corresponding performance degradations.

References

- [1] Amazon, Amazon 2018, URL <https://aws.amazon.com/>.
- [2] J.V. Bibal Benifa, D. Dejeu, RLPAS: Reinforcement learning-based proactive auto-scaler for resource provisioning in cloud environment, *Mobile Netw. Appl.* (2018).
- [3] R.N. Calheiros, K. Ramamohanarao, R. Buyya, C. Leckie, S. Versteeg, On the effectiveness of isolation-based anomaly detection in cloud data centers, *Concurr. Comput.: Pract. Exper.* (2017) e4169.
- [4] R.N. Calheiros, R. Ranjan, A. Beloglazov, C.A.F. De Rose, R. Buyya, CloudSim: A toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms, *Softw. Pract. Exper.* 41 (1) (2011) 23–50.
- [5] K. Cetinski, M.B. Juric, AME-WPC: Advanced model for efficient workload prediction in the cloud, *J. Netw. Comput. Appl.* 55 (2015) 191–201.
- [6] J.A. Cid-Fuentes, C. Szabo, K. Falkner, Online behavior identification in distributed systems, 34th IEEE Symposium on Reliable Distributed Systems, Montreal, Quebec, Canada, 2015, pp. 202–211.
- [7] C.A. Cunha, L. Moura e Silva, Separating performance anomalies from workload-explained failures in streaming servers, in: Proceedings of the 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (Ccgird 2012), CCGRID'12, IEEE Computer Society, Washington, DC, USA, 2012, pp. 292–299.
- [8] D.J. Dean, H. Nguyen, X. Gu, UBL: Unsupervised behavior learning for predicting performance anomalies in virtualized cloud systems, in: Proceedings of the 9th International Conference on Autonomic Computing, ACM, New York, NY, USA, 2012, pp. 191–200.
- [9] S. Di, D. Kondo, W. Cirne, Google hostload prediction based on bayesian model with optimized feature combination, *J. Parallel Distrib. Comput.* 74 (1) (2014) 1820–1832.
- [10] D. Grimaldi, V. Persico, A. Pescape, A. Salvi, S. Santini, A feedback-control approach for resource management in public clouds, in: IEEE Global Communications Conference, 2015, pp. 1–7.
- [11] N. Grozev, R. Buyya, Performance modelling and simulation of three-tier applications in cloud and multi-cloud environments, *Comput. J.* 58 (1) (2013) 1–22.
- [12] X. Gu, H. Wang, Online anomaly prediction for robust cluster systems, in: Proceedings of the 25th IEEE International Conference on Data Engineering, 2009, IEEE, Shanghai, China, 2009, pp. 1000–1011.
- [13] Q. Guan, S. Fu, Adaptive anomaly identification by exploring metric subspace in cloud computing infrastructures, in: Proceedings of the 32nd IEEE International Symposium on Reliable Distributed Systems, SRDS'13, IEEE Computer Society, Braga, Portugal, 2013, pp. 205–214.
- [14] R. Han, L. Guo, M.M. Ghanem, Y. Guo, Lightweight resource scaling for cloud applications, in: Proceedings of the 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (Ccgird 2012), CCGRID'12, IEEE Computer Society, Washington, DC, USA, 2012, pp. 644–651.
- [15] B. Hong, F. Peng, B. Deng, Y. Hu, D. Wang, DAC-Hmm: Detecting anomaly in cloud systems with hidden Markov models, *Concurr. Comput.: Pract. Exper.* 27 (18) (2015) 5749–5764.
- [16] F.T. Liu, K.M. Ting, Z.H. Zhou, Isolation-based anomaly detection, *ACM Trans. Knowl. Discov. Data* 6 (1) (2012) 3:1–3:39.
- [17] M. Mao, M. Humphrey, A performance study on the VM startup time in the cloud, in: Proceedings of the 2012 IEEE Fifth International Conference on Cloud Computing, CLOUD'12, IEEE Computer Society, Washington, DC, USA, 2012, pp. 423–430.
- [18] G. Moltó, M. Caballer, C. de Alfonso, Automatic memory-based vertical elasticity and oversubscription on cloud platforms, *Future Gener. Comput. Syst.* 56 (C) (2016) 1–10.
- [19] RUBIS, Rice University Bidding System, URL <http://rubis.ow2.org/>.
- [20] B. Subraya, *Integrated Approach to Web Performance Testing: A Practitioner's Guide*, IGI Global, 2006.
- [21] Y. Tan, H. Nguyen, Z. Shen, X. Gu, C. Venkatramani, D. Rajan, PREPARE: Predictive performance anomaly prevention for virtualized cloud systems, Proceedings of the 32nd IEEE International Conference on Distributed Computing Systems, Macau, China, 2012, pp. 285–294.
- [22] J. Yang, C. Liu, Y. Shang, Z. Mao, J. Chen, Workload predicting-based automatic scaling in service clouds, in: Proceedings of the 2013 IEEE Sixth International Conference on Cloud Computing, CLOUD'13, IEEE Computer Society, Washington, DC, USA, 2013, pp. 810–815.
- [23] L. Yazdanov, C. Fetzer, VScaler: autonomic virtual machine scaling, in: Proceedings of the 2013 IEEE Sixth International Conference on Cloud Computing, CLOUD'13, IEEE Computer Society, Washington, DC, USA, 2013, pp. 212–219.



Sara Kardani Moghaddam received her Bachelor's degree from the Shiraz University of Technology and completed a master degree in Information Technology at Sharif University of Technology. She is currently a Ph.D. student, at the University of Melbourne, Australia. Her major research interests are large scale distributed systems, performance management, anomaly detection and time-series analysis.



Dr. Rajkumar Buyya is a Redmond Barry Distinguished Professor and Director of the Cloud Computing and Distributed Systems (CLOUDS) Laboratory at the University of Melbourne, Australia. He is also serving as the founding CEO of Manjrasoft, a spin-off company of the University, commercializing its innovations in Cloud Computing. He served as a Future Fellow of the Australian Research Council during 2012–2016. He has authored over 625 publications and seven text books including "Mastering Cloud Computing" published by McGraw Hill, China Machine Press, and Morgan Kaufmann for Indian, Chinese and international markets respectively. He also edited several books including "Cloud Computing: Principles and Paradigms" (Wiley Press, USA, Feb 2011). He is one of the highly cited authors in computer science and software engineering worldwide (h-index=123, g-index=271, 78,000+ citations). Microsoft Academic Search Index ranked Dr. Buyya as #1 author in the world (2005–2016) for both field rating and citations evaluations in the area of Distributed and Parallel Computing. "A Scientometric Analysis of Cloud Computing Literature" by German scientists ranked Dr. Buyya as the World's Top-Cited (#1) Author and the World's Most-Productive (#1) Author in Cloud Computing. Recently, Dr. Buyya is recognized as a "Web of Science Highly Cited Researcher" in 2016, 2017, and 2018 by Thomson Reuters, a Fellow of IEEE, and Scopus Researcher of the Year 2017 with Excellence in Innovative Research Award by Elsevier for his outstanding contributions to Cloud computing. He is currently serving as Co-Editor-in-Chief of Journal of Software: Practice and Experience, which was established over 45 years ago. For further information on Dr. Buyya, please visit his cyberhome: www.buyya.com.



Professor Ramamohanarao (Rao) Kotagiri received PhD from Monash University. He was awarded the Alexander von Humboldt Fellowship in 1983. He has been at the University Melbourne since 1980 and was appointed as a professor in computer science in 1989. Rao held several senior positions including Head of Computer Science and Software Engineering, Head of the School of Electrical Engineering and Computer Science at the University of Melbourne and Research Director for the Cooperative Research Center for Intelligent Decision Systems. He served or serving on the Editorial Boards of the Computer Journal, Universal Computer Science, IEEE TKDE, VLDB Journal and International Journal on Data Privacy. Rao is a Fellow of the Institute of Engineers Australia, a Fellow of Australian Academy Technological Sciences and Engineering and a Fellow of Australian Academy of Science.