

Empirical Evaluation of Workload Forecasting Techniques for Predictive Cloud Resource Scaling

In Kee Kim, Wei Wang, Yanjun Qi, and Marty Humphrey

Department of Computer Science

University of Virginia

{ik2sb, wwang}@virginia.edu, {yq2h, humphrey}@cs.virginia.edu

Abstract— Many predictive resource scaling approaches have been proposed to overcome the limitations of the conventional reactive approaches most often used in clouds today. In general, due to the complexity of clouds, these reactive approaches were often forced to make significant limiting assumptions in either the operating conditions/requirements or expected workload patterns. As such, it is extremely difficult for cloud users to know which – if any – existing workload predictor will work best for their particular cloud activity, especially when considering highly-variable workload patterns, non-trivial billing models, variety of resources to add/subtract, etc. To solve this problem, we conduct comprehensive evaluations for a variety of workload predictors under real-world cloud configurations. The workload predictors cover four classes of 21 predictors: naive, regression, temporal, and non-temporal methods. We simulate a cloud application under four realistic workload patterns, two different cloud billing models, and three different styles of predictive scaling. Our evaluation confirms that no workload predictor is universally best for all workload patterns, and shows that Predictive Scaling-out + Predictive Scaling-in has the best cost efficiency and the lowest job deadline miss rate in cloud resource management, on average providing 30% better cost efficiency and 80% less job deadline miss rate compared to other styles of predictive scaling.

I. INTRODUCTION

When and how to scale-out and scale-in a cloud application can be one of the most difficult challenges for a cloud application designer/deployer. Fundamentally, basing this decision on the current state of resources (e.g., VMs) for a given cloud application is usually simple and can be effective but is limited due to its *reactive* nature. For example, a standard use of Amazon Web Services’ auto-scaling mechanism is to configure it to add more “worker VMs” when the average current CPU utilization stays above a certain threshold for a short period time. However, fluctuating or generally unpredictable workload patterns can quickly eliminate any potential performance or monetary improvements of such a policy/mechanism. In other words, intuitively any changes to the cloud application system state/configuration assume that the near-past workload will continue for the near-future.

Next-generation scaling approaches attempt to move beyond the limitations of such reactive systems by instead attempting to *predict* the near-future workloads. In such systems, generally, there are two components. The *workload predictor* functionality, and the *resource scaling* functionality, which allocates/deallocates cloud resources and maps work requests to specific resources. Over the past years, many studies have

proposed predictive scaling approaches [1–19]. In general, due to the complexity of cloud environments, these approaches were often forced to make significant limiting assumptions in either the operating conditions/requirements or expected workload patterns. As such, it is extremely difficult for cloud users to know which – if any – existing workload predictor will work best for their particular cloud activity, especially when considering highly-variable workload patterns, non-trivial billing models, variety of resources to add/subtract, etc.

The goal of this work is to comprehensively evaluate existing cloud workload predictors, holistically, and often in a broader context than in the authors’ evaluation methodology. Because the most common metric to evaluate workload predictors is accuracy for the future job arrivals, the first question we seek to answer is: *Which existing workload predictor has the highest accuracy for job arrival rate prediction, when applied for different workload patterns?*

After considering the workload predictor in isolation, we must now evaluate it in combination with the resource scaling component. Therefore, the second question we seek to answer is: *Which existing workload predictor has the best cost efficiency and deadline miss rate (which represents performance and SLA requirements), when applied for different workload patterns, different scaling styles and different pricing models?*

Furthermore, some previous work employed predictive scaling-out, some employed predictive scaling-in, and some employed both. Given these choices of applying predictive scaling, the third question we seek to answer is: *Which style of predictive scaling (predictive scaling-out, predictive scaling-in, or both) achieves the best cost and performance benefit for particular cloud configurations (e.g. billing model, job deadline)? And how much benefit can be achieved?*

To answer these questions, we conducted comprehensive evaluations with a wide range of configurations of predictive cloud resource scaling using 15 existing workload predictors [1–19]. We also included 6 well-known machine learning predictors that have not been used for the predictive scaling before. In total, we examined 21 predictors, covering naive, regression, temporal (time series) and non-temporal methods. We used 24 randomly generated workloads covering four common types of job arrival patterns [20, 21], which are growing, on/off, bursty and random. We also examined scaling operations including **RR** (Scaling-out: **Reactive** + Scaling-in: **Reactive**), **PR** (Scaling-out: **Predictive** + Scaling-in: **Reactive**), **RP** (Scaling-out: **Reactive** + Scaling-in: **Predictive**) and **PP** (Scaling-out: **Predictive** + Scaling-in: **Predictive**). We also considered both hourly and

minutely pricing models. In our experiments, each configuration then covered one workload predictor, one workload, one scaling operation and one billing model. As a result, more than 4K (21x24x4x2) configurations were examined. We run each configuration using the PICS [22] simulators to evaluate each configuration’s cost and deadline miss rate. We chose PICS because it can accurately simulate real public IaaS clouds in short amount of time. Without PICS, it is both timely and financially infeasible to conduct such extensive experiments on real IaaS clouds.

Based on the experimental results, we successfully answer the three questions posed previously. Here we summarize our findings and answers to each question:

To find the best workload predictor in terms of the accuracy for diverse workload patterns (Question #1): The accuracies of different workload predictors vary considerably. The best workload predictors in terms of statistical accuracy are usually orders of magnitudes more accurate than the worst ones. However, no workload predictor is universally the best for all workload patterns. Each workload pattern has its own best workload predictor. We show the best workload predictors for each workload pattern in Section IV-A.

To find the best workload predictor in terms of cloud metrics such as cost efficiency and SLA satisfaction (Question #2): The workload predictor with the highest accuracy is not necessarily the best in terms of cost and deadline miss rate. Additionally, no workload predictor is universally the best for any workload pattern and billing model. However, we observe that the best predictor (in terms of cost and deadline miss rate) for a particular workload pattern is always one of the top 3 most accurate predictors (in terms of statistical accuracy) of that workload pattern. We also show the best workload predictors for each workload pattern in terms of cloud metrics in Section IV-B.

To find the best style of predictive resource scaling in terms of providing the best cost and performance benefits (Question #3): Both predictive scaling-out (**PR**) and predictive scaling-in/out (**PP**) significantly reduces cost and deadline miss rate over purely reactive scaling (**RR**). However, predictive scaling-in (**RP**) performs similarly to **RR**. Overall, **PP** always provides the lowest cost and deadline miss rate for all workload patterns and billing models. On average, **PP** provides 30% less cost and 80% less deadline miss rate compared to **RR** or **RP**, and **PP** offers 14% less cost and 39% less deadline miss rate as compared to **PR**.

A key finding is that users who want to design a new algorithm for predictive resource scaling should consider the top three workload predictors depending on workload patterns and use **PP** for their scaling operations in order to archive better cost efficiency and deadline miss rate. The fact that **PP** is the best makes us further examine the benefit of using different workload predictors for scaling-out and scaling-in operations in **PP**. Our results show that mixing workload predictors has lower cost and deadline miss rate than using a single predictor in many cases. These results further suggest that mixture/ensemble of top workload predictors may be the best for predictive scaling.

The contributions of this paper include:

- An extensive evaluation of 21 workload predictors for their accuracy to predict the future job arrival time.

- A comprehensive evaluation of workload predictors in terms of cost and deadline miss rates. This evaluation considers various workload patterns, three styles of scaling operations, and two different billing models.
- Complete answers for the questions regarding what is the best workload predictor and what is the best style for predictive resource scaling operations.

The rest of this paper is organized as follows: Section II describes workload predictors used in this work. Section III contains the experimental design of this work. Section IV provides evaluation results for all predictors. Section V provides related work and Section VI concludes this paper.

II. WORKLOAD PREDICTORS

We collect a total of 21 workload predictors via an extensive survey of previous research on predictive resource scaling. Each predictor is categorized into one of the following types: 1) naive, 2) regression, 3) temporal, and 4) non-temporal predictor.

A. Naïve Predictors

There exist two types of naive workload predictors that are mean and recent mean-based (*k*NN) methods. The **mean job arrival time-based predictor** forecasts a next job arrival time based on a mean arrival time of all previous jobs. For the scaling-out operation, the cloud application prepares cloud resources as if the next job will be arrived at the predicted result based on mean. For the scaling-in operation, the cloud application waits until the predicted next job’s arrival time when a VM running by the cloud application is idle in order to increase a possibility of reuse of this VM. The **recent mean job arrival time-based predictor (*k*NN)** is a similar approach with mean-based predictor, but this uses the arrival rates of recent *k* jobs and predicts the next job’s arrival time based on a mean arrival time of those recent jobs.

B. Regression-based Workload Predictors

Regression-based approaches can be split into category of global and local regressions. Each category can include linear (1-degree) or polynomial (2 or more degrees) models. In total, we use 6 regression-based predictors, which are global and local regression with linear, quadratic, and cubic models.

Global regression predictors – linear and polynomial regression models forecast the next job arrival rate by creating a linear or polynomial regression model [23] using features including all previous job arrival rate. Here we only consider job arrival rate as the main variable. Therefore, these approaches are a single variable regression models.

Local regression predictors – local Linear and polynomial regression models use locally weighted regression models to estimate the next job arrival time. These approaches consist of two steps: 1) applying a kernel function to select job arrival samples and 2) creating linear or polynomial regression model based on the samples. In this work, we use *k*NN (*k* Nearest Neighbor) as the kernel function for the local regression models to select proper samples. *k*NN calculates a distance between a target object (e.g. next job arrival rate) and all possible samples (e.g. past job arrival rate) by using absolute or Euclidean distance function. *k*NN then selects proper local samples (e.g. *k* recent jobs). Based on the selected samples from the *k*NN, a linear or polynomial regression model is created, and

predicts the next job arrival rate. The major difference between the global and local regression is the size and similarity of job arrival samples used for creating a regression model. Local regression uses smaller number of samples that are more similar to the prediction target. The local regression models often have less overhead for model creation and workload prediction.

C. Temporal (Time-Series)-based Workload Predictors

There exist various temporal (time-series) methods for the future workload prediction because these predictors [5–15, 18, 19, 24] are widely used to analyze workload patterns for cloud computing as well as other domains of computer systems research. We use four categories of temporal models: 1) **ES** (Exponential Smoothing), 2) **AR** (Autoregressive), 3) **ARMA** (Autoregressive and Moving Average), and 4) **ARIMA** (Autoregressive Integrated Moving Average).

- **Exponential Smoothing (ES)** includes four methods: WMA (Weighted Moving Average), EMA (Exponential Moving Average), Holt-Winters DES (Double Exponential Smoothing), and Brown’s DES.

WMA is based on [3]. WMA the next job arrival time by equation-(1). WMA is a weighted sum of observed dataset (e.g. past job arrival information) and sum of weight for each observed data is 1.

$$WMA = \sum_{n=1}^k w_n x_{t+1-n}, s.t. \sum_{n=1}^k w_n = 1 \quad (1)$$

EMA is a similar approach as WMA, but it gives more weight to the most recent observation of job arrivals. We modeled EMA predictor based on [19]. EMA predicts the future job arrival time by $s_t = \alpha x_t + (1-\alpha)s_{t-1}$. α is a smoothing factor ($0 < \alpha < 1$). If α is close to 1, EMA has less smoothing effect and gives more weight to the recent data, and vice versa.

DES is used when the observed data show clearer seasonal trend. We use two types of DES: Holt-Winters and Brown’s DES.

Holt-Winters DES predicts the next job arrival time via equations-(2, 3). Equation-(2) captures a smoothing value at time t . Equation-(3) is the best estimate of the trend at time t . x_0 α is the first observation of raw data, is a smoothing factor ($0 < \alpha < 1$) and β is a trend smoothing factor ($0 < \beta < 1$). Then, the next job arrival time is calculated by $s_t + b_t$.

$$s_t = \alpha x_t + (1 - \alpha)(s_{t-1} + b_{t-1}), \text{where } s_1 = x_1 \quad (2)$$

$$b_t = \beta(s_t - s_{t-1}) + (1 - \beta)b_{t-1}, \text{where } b_1 = x_1 - x_0 \quad (3)$$

Brown’s DES predicts the next job arrival time by equation-(4). s'_t is the first order exponential smoothing model and is expressed by $s'_t = \alpha x_t + (1-\alpha)s'_{t-1}$. x_t is current job arrival and α is a smoothing factor ($0 < \alpha < 1$). s''_t is double-smoothed statistics and is expressed by $s''_t = \alpha s'_t + (1-\alpha)s''_{t-1}$.

$$T_{NextJobArrival} = \left(2 + \frac{\alpha}{1-\alpha}\right) s'_t - \left(1 + \frac{\alpha}{1-\alpha}\right) s''_t \quad (4)$$

- **AR** is a weighted sum of previous data in particular amount of time. AR is desirable to predict next job arrival time because it is a linear combination of previous data of the target object (e.g. job arrival time). [1] **AR(p)** for predictive resource scaling is based on [8] and **AR(p)** model is expressed in equation-(5). p is the order of AR model, φ_i is the set of parameters of the model, c is constant, and ε_t is white noise.

$$X_t = c + \sum_{i=1}^p \varphi_i X_{t-i} + \varepsilon_t \quad (5)$$

- **ARMA** is a combined model of AR and MA (Moving Average) and **ARMA(p, q)** is expressed in equation-(6). The first term is **AR(p)** model with the order of p . The second term is **MA(q)** model with the order of q . We modeled ARMA for workload prediction using [9–13].

$$X_t = \sum_{i=1}^p \varphi_i X_{t-i} + \sum_{i=1}^q \theta_i \varepsilon_{t-i} + c + \varepsilon_t \quad (6)$$

- **ARIMA** is a generalization of ARMA and provides a reliable prediction for non-stationary time-series data by integrating AR and MA models. ARIMA is expressed as **ARIMA(p, d, q)**, where p is the order of AR Model, q is the order of MA, and d is the order of differencing model. We modeled ARIMA for workload prediction based on [14, 15].

D. Non-temporal Workload Predictors

Non-temporal workload predictors have not been applied to the cloud resource scaling before. These predictors, however, have provided accurate prediction results within a deterministic amount of time. We consider several non-temporal approaches to predict the next job arrival time and select three categories of non-temporal machine learning approaches: SVM (Support Vector Machine), decision tree, and ensemble methods.

SVM is an optimal margin-based classifier that tries to find a small number of support vectors (data points) that separate all data points of two classes with a hyperplane in a high-dimensional space [23]. With kernel tricks, it can be extended as a nonlinear classifier to fit more complex classification problems. SVM can be applied to the case of regression as well which contains all the main features that characterize the maximum margin based algorithm. At testing time, the (positive or negative) distance of a data point to the hyper-plane is output as the prediction result. We consider both linear and non-linear SVM. We use linear kernel for Linear-SVM and Gaussian kernel for non-linear SVM (Gaussian-SVM). Linear-SVM is to focus on the workloads that have relatively clear trend factor and Gaussian-SVM is to predict the workloads that have non-linear characteristics.

Decision tree is a non-parametric learning algorithm and it has also been used for both classification and regression problems [23]. Decision tree creates a classification or regression model by applying decision rules derived from features of dataset. Decision tree is known as its simple (tree-based) structure and fast execution time for numerical data.

Ensemble prediction methods employ multiple numbers of predictors to obtain better generalizability and increase performance. Ensemble methods use bagging or boosting approaches to reduce variance (bagging) or bias (boosting) on prediction results. We select three ensemble methods including gradient boosting, extra-trees, and random forest [23].

III. EXPERIMENTAL DESIGN

A. Design of Cloud Resource Management System

To evaluate all workload predictors, we designed a cloud management system for resource scaling as shown in Figure 1. This system consists of three modules: job portal, resource management module (RMM) and predictive scaling module

(PSM). The job portal is an entry for the workloads (jobs from end-users). A job’s arrival triggers two other modules. A newly arrived job is passed to the RMM. The RMM selects a proper VM based on the job’s duration and deadline. More specifically, the RMM creates a list of VMs that meets the deadline of the job by comparing the performance of different VM types with the job’s duration and deadline. the RMM then selects the most cost efficient VM (i.e., cost/performance-ratio [25]) from the candidates. Note that the algorithm, used in this cloud resource management, focuses primarily on improving the job deadline satisfaction than reducing the cloud cost. Once a proper VM is selected, the RMM schedules the job to the selected VM via “Earliest Deadline First” scheduling. The selected VM is then used to execute the job. A new job’s arrival activates the PSM as well. The job’s arrival information is stored in the workload repository. The workload information from this repository is used by both predictive scaling-out and scaling-in.

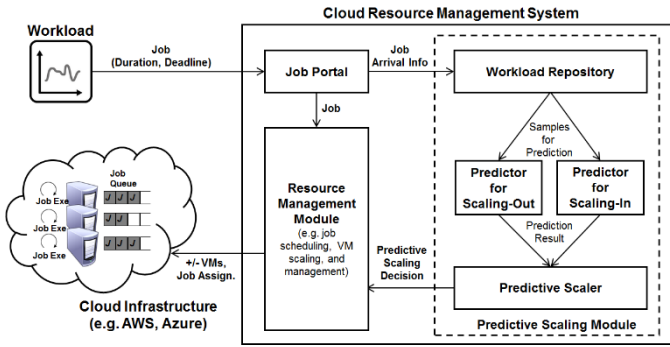


Fig. 1. Cloud Resource Management System with Predictive Scaling.

Predictive Scaling-out Operation (Algorithm 1) is triggered when a job arrives. A prediction obtains proper amount of job arrival samples for prediction (ln 1) and forecasts the next job’s arrival time in the future (ln 2). Based on the prediction result, this operation chooses a proper type of VM for the future job (ln 3) as explained in the previous paragraph (ln Algorithm 1, we assume that the duration and deadline of the future job are known). This operation selects a list of currently running VMs to execute the next job (ln 4). If the list is empty (ln 6), then a new VM will be created for the next job (ln 7).

Algorithm #1 – Predictive Scaling-Out

```

Require: A new job arrivals
1:  $samples \leftarrow get\_sample\_for\_prediction ()$ 
2:  $next\_job \leftarrow predict\_next\_job\_arrival (samples)$ 
3:  $vm\_type \leftarrow select\_proper\_vm\_type (next\_job)$ 
4:  $vm\_list \leftarrow current\_running\_vms (vm\_type)$ 
5:
6: if  $vm\_list$  is empty then
7:    $create\_vm (vm\_type, time\_to\_start)$ 
8: else
9:    $do\_nothing ()$ 
10: end if

```

Predictive Scaling-in Operation (Algorithm 2) is triggered when a VM is idle – no jobs in both processor and work queue. The workload predictor for scaling-in operation estimates the next job arrival time to the idle VM (ln 1–2). Then we compare the predicted job arrival time with maximum VM startup delay [26]. If the job arrival time is smaller than the max

startup delay, we choose to keep this VM for at least max startup delay time; otherwise, we choose to terminate this VM (ln 4–8). The rationale behind this choice is explained as follows. For any new job, starting a new VM for it takes (startup time + job duration) to execute it. However, if we use existing idle VM, it takes (job arrival time + job duration) to execute it. Therefore, if job arrival time is smaller than startup time, it is faster and cheaper to use the existing VM; otherwise, it is cheaper to use new VM. Additionally, we choose to terminate a VM only at nearest billing boundary, because we have already paid for this billing cycle. If the next job arrives within the predicted time, the job is assigned/processed by this VM (ln 12). If there is no more jobs until the predicted time, this VM is terminated (ln 16).

Algorithm #2 – Predictive Scaling-In

```

Require:  $vm$  is idle
1:  $samples \leftarrow get\_sample\_for\_prediction (vm)$ 
2:  $next\_job \leftarrow predict\_next\_job\_arrival (samples)$ 
3:
4: if  $next\_job$  arrival  $\leq max\_startup\_delay$  then
5:    $scale\_in\_time \leftarrow next$  billing boundary after  $next\_job\_arrival$ 
6: else
7:    $scale\_in\_time \leftarrow this$  billing boundary
8: end if
9:
10: repeat
11:   if  $next\_job$  arrives then
12:     go to  $new\_job\_processing\_routine()$ 
13:   end if
14: until  $scale\_in\_time$ 
15:
16: terminate ( $vm$ )

```

TABLE I. WORKLOAD DATASET INFORMATION. (WL: WORKLOAD, JOB AR: JOB ARRIVAL, JOB DR: JOB DURATION, JOB DL: JOB DEADLINE)

WL	# Jobs	Mean Job AR	Std. Dev. Job AR	Job DR	Job DL
Growing	10,000	25s	60.5	Average: 450s, Std. Dev.: 270	Average: 500s, Std. Dev.: 250
On/Off			375.5		
Bursty			35.3		
Random			270		




Fig. 2. Cloud Workload Patterns with Predictive Scaling.

B. Cloud Workload Patterns

We generate synthetic workloads based on the four canonical cloud workload patterns (Fig. 2). We create 6 workloads for each workload pattern with different mean and standard deviation of job arrival time/duration to reflect various and realistic cloud usage scenarios. The detail of each dataset is described in Table I. For the growing workload pattern, we first generate a quadratic base function and then apply Poisson

distribution to randomize the arrival time of a particular job. The on/off workload pattern has four active periods and three inactive periods. For the active periods of the on/off workload pattern, we use growing and declining quadratic functions. The bursty workload pattern has 6-7 peak spikes periods and other stable periods. To generate the random workload pattern, we use Poisson distribution for the random job arrivals.

C. Implementation Details

For the evaluation of all predictive scaling configurations, we implemented our cloud resource management system on top of PICS [26], a Public IaaS Cloud Simulator. In addition to the simulation model, we implement all predictors and scaling-in/out mechanisms (described in Section II) using numpy version 1.8, Pandas (Python Data Analysis Library), and scikit-learn machine learning packages.

IV. EVALUATION

A. Evaluation of Accuracy for Workload Predictors

As the first evaluation of this work, we measure all 21 workload predictors' accuracy for predicting the future job arrival time under four different workload patterns. We employ MAPE (Mean Absolute Percentage Error) [23] to statistically measure the prediction accuracy.

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{T_{actual,i} - T_{predicted,i}}{T_{actual,i}} \right| \quad (7)$$

Figure 3 shows MAPE results of all 21 predictors. Average of the MAPE for all 21 predictors is 0.6360. Overall, two SVM approaches have the best MAPE results. Other three best predictors are WMA, ARMA, and Brown's DES. The MAPE value of Gaussian-SVM predictor (the best predictor) is 0.3665, which is 42% less result than average of all predictors. However, the best predictors in overall do not necessarily mean the best predictor for each workload pattern, so we also present the performance of all predictors for each workload pattern.




Fig. 3. MAPE Results of All 21 Predictors.

Table II shows the MAPE results for each workload pattern. Due to page limitation, Table II only contains the best 3 predictors, average results, and the worst predictor. As shown in Table II, different workload patterns have different best predictors: **Linear-SVM** (growing), **Gaussian-SVM** (on/off), **ARIMA** (bursty), and **Gaussian-SVM** (random). Table II also shows that the MAPE results of the top three predictors are very similar to each other for growing and bursty workloads. The MAPE results of random workload prediction are lower compared to other workloads, indicating random workload is the most difficult to predict. This difficulty is primarily caused by

the fact that job arrival times in random workloads do not have clear trend pattern for predictors to discover.

TABLE II. THE MAPE RESULTS OF WORKLOAD PREDICTORS UNDER FOUR DIFFERENT WORKLOAD PATTERNS. (WL: WORKLOAD, GR: GROWING, OO: ON/OFF, BR: BURSTY, RN: RANDOM)

WL	Rank	Predictor	MAPE	WL	Rank	Predictor	MAPE
GR	1	L-SVM	0.28	OO	1	G-SVM	0.22
	2	AR	0.29		2	ARMA	0.30
	3	ARMA	0.30		3	L-SVM	0.44
	Avg.	—	0.51		Avg.	—	0.69
	Worst	Qua. Reg.	2.75		Worst	Loc. Cub. Reg.	1.25
BR	1	ARIMA	0.38	RN	1	G-SVM	0.45
	2	BRDES	0.41		2	Lin.Reg	0.46
	3	L-SVM	0.43		3	L-SVM	0.46
	Avg.	—	0.75		Avg.	—	0.52
	Worst	mean	3.35		Worst	Dec. Tree	0.62




Fig. 4. Prediction Overhead (Log-scaled Sum of Total Prediction Time) for All 21 Workload Predictors for the Future Job Arrival Time Prediction.

Moreover, obtaining the prediction results in a deterministic amount of time is a critical issue for the predictive resource scaling. We also measured computation overhead of predictors as shown in Figure 4. kNN, WMA, and EMA are the fastest predictors. The overall prediction time for 10K jobs of these three predictors are 0.48 (kNN), 0.96 (WMA), and 1.86 seconds (EMA). However, some temporal approaches (AR, ARMA, and ARIMA) and ensemble methods (extra trees, gradient boosting, and random forest) have longer prediction time. The highest overhead predictor is ARMA, which takes 6031.52 seconds for 10K jobs.

B. Performance of Different Styles of Predictive Scaling

We measure the performance of four different styles of scaling operations for cloud resource management:

- **Baseline: RR** (Scale-Out: **R**eactive + Scale-In: **R**eactive)
- **PR** (Scale-Out: **P**redictive + Scale-In: **R**eactive)
- **RP** (Scale-Out: **R**eactive + Scale-In: **P**redictive)
- **PP** (Scale-Out: **P**redictive + Scale-In: **P**redictive)

PR is the most common style of predictive scaling for cloud application. **PR** predictively scales out cloud resources and reactively scales in cloud resources. **RP** is another way of predictive scaling, and it uses a reactive way for scaling-out and a predictive approach for scaling-in. **PP** is a combination of **PR** and **RP**, and it leverages a workload predictor for both scaling-out and scaling-in operations. For this evaluation, we use **RR** (Scaling-Out: **R**eactive + Scaling-In: **R**eactive) as a baseline.

RR adds a new VM (scaling-out) when a new job needs an extra VM, and terminates an idle VM (scaling-in) at its billing boundary. (i.e., hourly bound or minutely bound).

Due to page limitation, we only show the results of the predictive scaling operations with the most accurate 8 workload predictors from Section IV-A. These predictors are Linear-SVM, Gaussian-SVM, ARMA, AR, WMA, ARIMA, Brown’s DES, Linear regression, which cover the overall best 5 predictors and the best 3 predictors for each workload pattern.

To evaluate the predictive scaling operations, we use two common cloud metrics (cost and job deadline miss rate) and two different billing models (hourly and minutely pricing model). Cost is to evaluate each scaling operations’ cost efficiency, and job deadline miss rate represents the SLA-satisfaction requirement. We use two different billing models for cloud infrastructure, because major commercial IaaS clouds employ either hourly (e.g. AWS) or minutely (e.g. MS Azure) pricing model. We also use four different workload patterns (growing, on/off, bursty, and random workload).

The goals of this evaluation are:

- Measuring the actual benefits from predictive scaling.
- Determining the best style of predictive scaling.
- Finding the best workload predictor for each workload pattern in terms of cloud metrics.

TABLE III. CASE #1 – NORMALIZED COST AND JOB DEADLINE MISS RATE OF **PR**.

(a) Hourly Pricing Model								
	Growing		On/Off		Bursty		Random	
	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL
G-SVM	50%	22%	44%	8%	53%	15%	114%	47%
L-SVM	38%	14%	44%	9%	49%	18%	114%	49%
ARMA	38%	16%	43%	12%	53%	24%	90%	49%
WMA	50%	16%	43%	14%	52%	17%	114%	48%
AR	44%	19%	43%	12%	52%	20%	114%	48%
ARIMA	38%	14%	43%	14%	52%	15%	114%	49%
BRDES	38%	17%	43%	16%	53%	21%	114%	55%
Lin.Reg.	38%	23%	45%	12%	61%	24%	114%	55%
Average	42%	18%	44%	12%	53%	19%	111%	50%

(b) Minutely Pricing Model								
	Growing		On/Off		Bursty		Random	
	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL
G-SVM	104%	43%	99%	13%	97%	43%	106%	59%
L-SVM	105%	43%	99%	13%	96%	43%	106%	59%
ARMA	103%	42%	100%	11%	98%	45%	105%	59%
WMA	103%	42%	99%	15%	98%	43%	107%	59%
AR	104%	41%	99%	14%	98%	45%	107%	59%
ARIMA	104%	42%	98%	14%	98%	44%	107%	59%
BRDES	104%	43%	99%	14%	98%	43%	106%	59%
Lin.Reg.	108%	41%	100%	12%	97%	44%	107%	59%
Average	104%	42%	99%	13%	98%	44%	106%	59%

Case #1 – PR (Scale-Out: Predictive + Scale-In: Reactive): Table III(a) shows a normalized cost and job deadline miss rate (all results are normalized to **RR**) of **PR** for hourly pricing model. The results show that **PR** can improve 47%–58% of cost efficiency for growing, on/off, and bursty workloads. However, for random workload, **PR** has 11% of worse cost efficiency over the baseline. In terms of job deadline miss rate, **PR** has 50%–88% of less job deadline miss rate over the **RR** (baseline). For the hourly pricing model, **PR** shows relatively poor performance for random workload in both cost efficiency and job deadline

miss rate. This is because random workload is harder to predict than the other workload patterns. We rank the workload predictors for the **PR** based on the deadline miss rate, because it is more important to ensure that jobs meet their deadlines. Only after the deadline requirements are met, our cloud resource manager in Section III-A optimizes for cost efficiency. The best workload predictors for **PR** are: **Linear-SVM** (13.7%) for growing, **Gaussian-SVM** (8.2%) for on/off, **ARIMA** (15.2%) for bursty, and **Gaussian-SVM** (47.3%) for random workload.

Table III(b) shows evaluation results of **PR** for minutely pricing model under four workload patterns. **PR** has similar cost efficiency with **RR**, but it has 41%–87% of less job deadline miss rate than the baseline. Thus, **PR** provides better job deadline satisfaction without dramatically increasing cost. The reason that **PR** has similar cost efficiency with **RR** is that the minutely pricing model is designed to provide better cost efficiency than hourly model to the user. So it is very hard to improve the cost efficiency for the minutely pricing model even though we have a good predictor. The best workload predictors for **RP** with minutely pricing model are: **AR** (41.3%) for growing, **ARMA** (11.4%) for on/off, **WMA** (42.6%) for bursty, and **Gaussian-SVM** (58.6%) for random workload.

TABLE IV. CASE #2 – NORMALIZED COST AND JOB DEADLINE MISS RATE OF **RP**.

(a) Hourly Pricing Model								
	Growing		On/Off		Bursty		Random	
	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL
G-SVM	101%	99%	71%	100%	134%	100%	78%	99%
L-SVM	100%	99%	77%	99%	138%	99%	78%	99%
ARMA	101%	99%	78%	98%	137%	98%	78%	111%
WMA	101%	100%	71%	99%	134%	100%	80%	99%
AR	101%	100%	83%	99%	150%	101%	79%	100%
ARIMA	101%	99%	77%	99%	138%	97%	78%	99%
BRDES	101%	99%	71%	99%	138%	98%	78%	100%
Lin.Reg.	100%	99%	75%	100%	134%	100%	78%	100%
Average	101%	99%	75%	99%	138%	99%	78%	101%

(b) Minutely Pricing Model								
	Growing		On/Off		Bursty		Random	
	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL
G-SVM	110%	101%	121%	99%	109%	99%	86%	99%
L-SVM	108%	100%	121%	99%	107%	99%	86%	99%
ARMA	105%	99%	120%	100%	108%	100%	89%	100%
WMA	107%	99%	122%	99%	109%	99%	86%	99%
AR	108%	99%	123%	100%	110%	99%	90%	100%
ARIMA	107%	99%	123%	99%	109%	99%	85%	99%
BRDES	105%	99%	121%	99%	108%	101%	86%	99%
Lin.Reg.	104%	99%	122%	99%	109%	99%	85%	99%
Avg	107%	99%	122%	99%	109%	99%	87%	99%

Case #2 – RP (Scale-Out: Reactive + Scale-In: Predictive): Table IV(a) and IV(b) show the evaluation results of **RP** for both pricing models under four workload patterns. The results indicate that **RP**’s benefit to the cloud system is not as much as the benefits from **PR**. The only benefit from the **RP** is the improved cost efficiency (on/off and random workloads for hourly pricing model, random workload for minutely pricing model). The improvement of cost efficiency is 12%–25%, but it has no benefits for job deadline miss rate.

Case #3 – PP (Scale-Out: Predictive + Scale-In: Predictive): Table V(a) shows normalized costs and job deadline miss rates of **PP** for hourly pricing model under four workload patterns. The results show that **PP** improves 48%–69% of cost efficiency

for the four workloads over **RR**, and has 78%–93% of less job deadline miss rate than **RR**. The best workload predictors for the **PP** with hourly pricing model are: **Linear-SVM** (5.5%) for growing, **Gaussian-SVM** (7.8%) for on/off, **Brown’s DES** (5.1%) for bursty, and **Linear-SVM** (18.1%) for random workload.

Table V(b) shows evaluation results of **PP** for minutely pricing model under four workload patterns. The results show that **PP** slightly improves cost efficiency (< 11%) of the cloud system and has huge improvement (60%–87%) for the job deadline miss rate. The best workload predictors of **PP** with minutely pricing model for each workload pattern are: **AR** (27.5%) for growing, **ARMA** (11.4%) for on/off, **Brown’s DES** (19.7%) for bursty, and **Gaussian-SVM** (39.5%) for random workload.

TABLE V. CASE #3 – NORMALIZED COST AND JOB DEADLINE MISS RATE OF **PP**.

(a) Hourly Pricing Model								
	Growing		On/Off		Bursty		Random	
	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL
G-SVM	30%	9%	34%	8%	50%	9%	48%	22%
L-SVM	30%	6%	34%	9%	47%	6%	48%	18%
ARMA	33%	6%	34%	8%	52%	12%	48%	20%
WMA	32%	7%	33%	12%	56%	9%	48%	22%
AR	30%	4%	39%	12%	56%	9%	48%	24%
ARIMA	30%	6%	38%	12%	52%	12%	48%	27%
BRDES	32%	7%	35%	14%	45%	5%	48%	20%
Lin.Reg.	33%	13%	34%	9%	56%	7%	48%	20%
Average	31%	7%	35%	11%	52%	9%	48%	22%

(b) Minutely Pricing Model								
	Growing		On/Off		Bursty		Random	
	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL	Nor. Cost	Nor. DL
G-SVM	98%	29%	99%	13%	94%	31%	90%	40%
L-SVM	98%	28%	99%	13%	94%	27%	89%	40%
ARMA	98%	29%	100%	11%	97%	37%	89%	41%
WMA	98%	28%	99%	15%	94%	31%	86%	40%
AR	97%	28%	99%	13%	96%	33%	94%	40%
ARIMA	99%	28%	99%	14%	96%	33%	89%	40%
BRDES	97%	28%	99%	14%	95%	20%	89%	40%
Lin.Reg.	97%	31%	100%	12%	92%	30%	86%	40%
Average	98%	29%	99%	13%	95%	30%	89%	40%

Comparison of Three Predictive Scaling Styles: So far we have separately evaluated three predictive scaling styles of the cloud resource management. In the following paragraphs, we present the overall benefit of predictive scaling. And we compare the results of these three predictive scaling styles to determine the best one for cloud resource management. Figure 5 shows that the comparison of average results of normalized cost and job deadline miss rate for **PR**, **RP**, and **PP** in both hourly and minutely pricing models.

For the hourly pricing model (Figure 5(a)), we found that **PP** is the best style for cloud resource management in terms of better cost efficiency and less job deadline miss rate. **PP**’s cost efficiency is 20% (compared to **PR**) and 56% (compared to **RP**) better than other two approaches. Moreover, **PP**’s job deadline miss rate is 13% (compared to **PR**) and 88% (compared to **RP**) lower than others. This result is interesting, because although predictive scaling-in does not improve cost and deadline miss rate by itself (as shown in Case #2), it provides considerable improvement for both metrics when combined with predictive

scaling-out. These results show that predictive scaling in/out approach (**PP**) (with a good workload predictor) help to improve the performance of the cloud resource management.

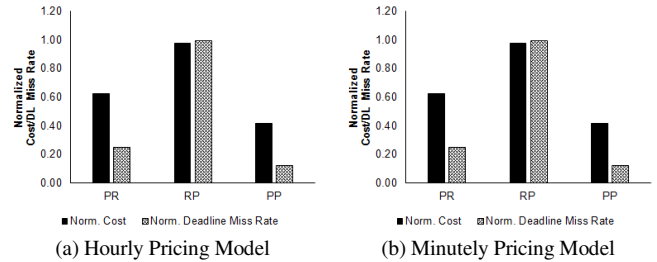


Fig. 5. Comparison of Average Results of Normalized Cost and Job Deadline Miss Rate of Three Scaling Styles (**PR**, **RP**, and **PP**).

For the minutely pricing model (Figure 5(b)), the job deadline miss rate of **PP** outperforms other two styles of predictive scaling operations. **PP** has 12% (compared to **PR**) and 72% (compared to **RP**) of less job deadline miss rate. **PP** also improves cost efficiency of over **PR** and **RP**. These results suggest that **PP** can significantly reduce deadline miss rate without any cost overhead.

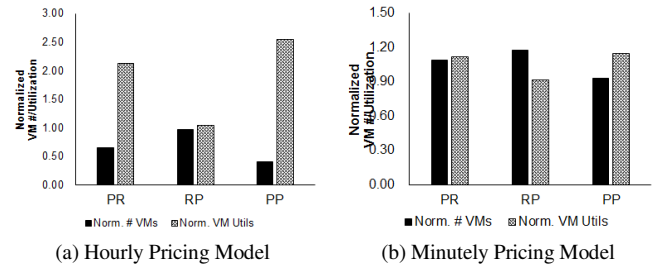


Fig. 6. Comparison of VM Numbers and Utilization of Three Scaling Styles (**PR**, **RP**, and **PP**).

To understand the reasons of 1) **PP** significantly improves cost efficiency (for hourly pricing model) and deadline miss rate (for both pricing model) and 2) **RP** does not improve the performance by itself, we analyze the number of created VMs and VM utilization of three styles. Figure 6 represents the VM numbers and utilization of three predictive scaling styles for both pricing models. For the both pricing models, **PP** creates the lower number of VMs and has higher utilization than other approaches. The reasons that **PP** has high VM utilization and lower number of created VMs are:

- Predictive scaling-out of **PP** uses more currently running VMs for the (near) future jobs, and creates less VMs for the (near) future jobs.
- Predictive scaling-in of **PP** keeps VMs alive for (further) future jobs, which further reduces the new VM creations, and increases the utilizations of existing VMs.

Moreover, the reason that **RP** cannot improve the cloud metrics is related to reactive scaling-out of **RP**. Reactive scaling-out operation creates VMs when jobs actually arrive, so **RP** has to create better performance VMs (more expensive VMs) in order to meet the jobs’ deadline. This is because **RP** has no advance preparation for eliminating the overhead of the VM creation (e.g. startup delay). Also most of VMs should be terminated after processing the current job because they are not used for future jobs. So, predictive scaling-in of **RP** does not help in this case because most of VMs should be destroyed.

V. RELATED WORK

A large body of work has been conducted for the predictive cloud resource management for dynamic workload patterns. There are two major branches in predictive resource management in the clouds. First branch is to focus on predicting the future resource usages (e.g. CPU, memory, and I/O) based on past resource usage history [1–4, 27]. Second branch is a workload predictive approach for cloud resource management. More specifically, this branch focuses on predicting the future job arrival rate by using various workload predictors. Our work is closer to the second branch because we have evaluated prediction techniques for the future job arrivals to the clouds. In order to predict the next job arrival time, previous works employ a variety of predictors. They applied regression [16, 17, 28–30], time-series methods (e.g. ES [5–7, 18, 19], AR [8, 24], ARMA [9–13] and ARIMA [14, 15]) to their cloud systems to effectively manage the cloud resources. Herbst *et al* [31] proposed WCF (Workload Classification and Forecasting) framework, which is a self-adaptive approach for proactive resource management. However, the purpose of our work is different from the goal of previous works (proposing a new predictive scaling mechanism). We are more geared toward evaluating the performance of all existing prediction techniques for cloud resource management (e.g. scaling) by employing various/realistic workload patterns [20, 21], diverse cloud configurations (e.g. billing model), and common performance metrics. Furthermore, we have concentrated on the evaluation of different styles of predictive scale operations: predictive scaling-out only, predictive scaling-in only, and both predictive scaling-in/out. Therefore, our goal is to determine the best workload predictor for the cloud users' specific workload pattern and cloud configurations as well as the best scaling style of the predictive scaling.

VI. CONCLUSION

In order to help the cloud users select the best workload predictor, we comprehensively evaluated 21 workload predictors using statistical and cloud metrics. We found that to design a new predictive cloud resource scaling, the cloud users should consider top 3 workload predictors depending on their workload patterns, and use **PP** approach for their scaling operations in order to maximize the scaling performance. Moreover, we examined the benefit of using different predictors for scaling-out and scaling-in in **PP**. Our results show that mixing workload predictors has lower cost and deadline miss rate than using a single predictor. These results further suggest that ensemble of top workload predictors may be the best for predictive scaling.

REFERENCES

- [1] P.Dinda and D.O'Hallaron. Host Load Prediction using Linear Models. *Cluster Computing*, 3(4):265–280, 2000.
- [2] Z.Gong, X.Gu, and J.Wilkes. PRESS: PRedictive Elastic ReSource Scaling for cloud systems. In Proc. CNSM, 2010.
- [3] Z.Xiao, W.Song, and Q.Chen. Dynamic Resource Allocation Using Virtual Machines for Cloud Computing Environment. *IEEE Trans on Parallel and Distributed Systems*, 24(6), 2013.
- [4] A.Bankole and S.Ajila. Cloud Client Prediction Models for Cloud Resource Provisioning in a Multitier Web Application Environment. In Proc IEEE SOSE2013, 2013.
- [5] S.Niu, J.Zhai, X.Ma, X.Tang, and W.Chen. Cost-effective Cloud HPC Resource Provisioning by Building Semi-Elastic Virtual Clusters. In Proc SC13, 2013.
- [6] C.Mark, D.Niyato, and T.Chen-Khong. Evolutionary Optimal Virtual Machine Placement and Demand Forecaster for Cloud Computing. In Proc. AINA, 2011.
- [7] H. Mi *et al*. Online Self-reconfiguration with Performance Guarantee for Energyefficient Large-scale Cloud Computing Data Centers. In Proc. SCC, 2010.
- [8] N.Bobroff, A.Kochut, and K.Beaty. Dynamic Placement of Virtual Machines for Managing SLA Violations. In Proc IFIP/IEEE Symposium of Integrated Management (IM), 2007.
- [9] N.Roy, A.Dubey, and A.Gokhale. Efficient Autoscaling in the Cloud using Predictive Models for Workload Forecasting. In Proc CLOUD 2011.
- [10] J.Tirado, D.Higuero, F.Isaila, and J.Carretero. Predictive Data Grouping and Placement for Cloud-based Elastic Server Infrastructures. In Proc. CCGRID, 2011.
- [11] M.Dhingra *et al*. Elastic Resources Framework in IaaS, preserving performance SLAs. In Proc. CLOUD 2013.
- [12] U.Sharma, P.Shenoy, and S.Sahu. A Flexible Elastic Control Plane for Private Clouds. In Proc. CAC, 2013.
- [13] S-P.Li and M-H.Wong. Data Allocation in Scalable Distributed Database Systems Based on Time Series Forecasting. In Proc. BigData 2013.
- [14] R.Calheiros *et al*. Workload Prediction Using ARIMA Model and Its Impact on Cloud Applications' QoS. *IEEE Trans on Cloud Computing*, 2015.
- [15] H.Xu Di Niu, B.Li, and S.Zhao. Quality-Assured Cloud Bandwidth Auto-Scaling for Video-on-Demand Applications. In Proc. INFOCOM, Orlando, FL, March, 2012.
- [16] P.Bodik *et al*. Statistical Machine Learning Makes Automatic Control Practical for Internet Datacenters. In Proceedings HotCloud., 2009.
- [17] J.Yang *et al*. A cost-aware auto-scaling approach using the workload prediction in service clouds. *Information Systems Frontiers*, 16(1), 2014.
- [18] S.Koyano *et al*. A High-grained Traffic Prediction for Microseconds Power Control in Energy-aware Routers. In Proc. UCC 2012..
- [19] P.Saripalli *et al*. Load Prediction and Hot Spot Detection Models for Autonomic Cloud Computing. In Proc. UCC 2011.
- [20] C.Fehling *et al*. Cloud Computing Patterns: Fundamentals to Design, Build, and Manage Cloud Applications. 2014.
- [21] Workload Patterns for Cloud Computing. <http://watdenkt.veenhof.nu/2010/07/13/workload-patterns-for-cloud-computing/>.
- [22] I.Kim, W.Wang, and M.Humphrey. PICS: A Public IaaS Cloud Simulator. In CLOUD, 2015.
- [23] T.Hastie, R.Tibshirani, and J. Friedman. *The Element of Statistical Learning: Data Mining, Inference, and Prediction*. 2011.
- [24] A.Chandra, W.Gong, and P.Shenoy. Dynamic Resource Allocation for Shared Data Centers Using Online Measurements. In IWQoS, 2003.
- [25] Q.Zhu and G.Agrawal. Resource Provisioning with Budget Constraints for Adaptive Applications in Cloud Environments. *IEEE Transactions on Parallel and Distributed Systems*, 54(4), 2012.
- [26] M.Mao and M. Humphrey. A Performance Study on the VM Startup Time in the Cloud. In Proc. CLOUD, 2012.
- [27] Z.Shen *et al*. Cloud-Scale: Elastic Resource Scaling for Multi-Tenant Cloud Systems. In Proc. SoCC, 2011.
- [28] W.Iqbal *et al*. Adaptive resource provisioning for read intensive multi-tier applications in the cloud. *Future Generation Computer Systems*, 27(6), 2011.
- [29] S.Muppala, X.Zhou, and L.Zhang. Regression Based Multi-tier Resource Provisioning for Session Slowdown Guarantees. In IPCC, 2010.
- [30] S.Dutta *et al*. SmartScale: Automatic Application Scaling in Enterprise Clouds. In CLOUD, 2012.
- [31] N.Herbst *et al*. Self-Adaptive Workload Classification and Forecasting for Proactive Resource Provisioning. In ICPE, 2013