

Scheduling data analytics work with performance guarantees: queuing and machine learning models in synergy

Ji Xue¹ · Feng Yan¹ · Alma Riska² · Evgenia Smirni¹

Received: 17 January 2016 / Revised: 23 March 2016 / Accepted: 30 March 2016 / Published online: 23 April 2016
© Springer Science+Business Media New York 2016

Abstract In today's scaled out systems, co-scheduling data analytics work with high priority user workloads is common as it utilizes better the vast hardware availability. User workloads are dominated by periodic patterns, with alternating periods of high and low utilization, creating promising conditions to schedule data analytics work during low activity periods. To this end, we show the effectiveness of machine learning models in accurately predicting user workload intensities, essentially by suggesting the most opportune time to co-schedule data analytics work. Yet, machine learning models cannot predict the effects of performance interference when co-scheduling is employed, as this constitutes a “new” observation. Specifically, in tiered storage systems, their hierarchical design makes performance interference even more complex, thus accurate performance prediction is more challenging. Here, we quantify the unknown performance effects of workload co-scheduling by enhancing machine learning models with queuing theory ones to develop a hybrid approach that can accurately predict performance and guide scheduling decisions in a tiered storage system. Using traces from commercial systems we illustrate that queuing theory and machine learning models can

be used in synergy to surpass their respective weaknesses and deliver robust co-scheduling solutions that achieve high performance.

Keywords Tiered storage system · Priority scheduling · Workload prediction · Performance guarantees · Hybrid approach · Data analytics work

1 Introduction

High utilization of resources is a design target of scaled-out systems. Often in such systems, data analytics workloads co-exist with high priority user workloads that operate within strict service level objectives (SLOs). Data analytics workloads, e.g., personalized advertising, sentiment analysis, product recommendation, database replication, dominate many systems today. Different than traditional internal work (e.g., garbage collection, snapshots, upgrades), data analytics work requires faster reaction in order to provide timely information [1–3], e.g., a delayed advertisement event update could cause reduced income or a product recommendation should occur before the user leaves the site. Since data analytics to enhance user experience and regular user traffic share the same hardware, their effective resource management can greatly enhance business value and user satisfaction.

Scheduling user traffic and data analytics work in the same system is a challenging task. Scheduling data analytics too aggressively may cause user traffic to suffer from SLO violations. If scheduled too conservatively, data analytics work could not finish in time, thus could lose its value. With user workload traffic that is repeatable and periodic across different time scales [3–5], it is natural to interleave data analytics work with user workload at periods of low user demands.

✉ Feng Yan
fyan@cs.wm.edu

Ji Xue
xuejimic@cs.wm.edu

Alma Riska
alma.riska@netapp.com

Evgenia Smirni
esmirni@cs.wm.edu

¹ Computer Science Department, College of William and Mary, Williamsburg, VA, USA

² NetApp, Waltham, MA, USA

Key to the effective deployment of any policy that interleaves data analytics with SLO-bound user level work, is the prediction of fluctuations in the user workload and especially identifying a priori heavy or spiky loads. Here, we propose *NeuQ*, a neural network/queueing hybrid solution: the queueing model that is the basis of co-scheduling decisions is significantly enhanced by neural networks to improve its accuracy. The queueing models that we use predict the magnitude of the potential performance interference of co-scheduling user and data analytics workloads. An important parameter of the queueing model that greatly affects its prediction accuracy is a priori knowledge of the upcoming arrival intensity, this is successfully provided by the neural network.

To illustrate the effectiveness of *NeuQ*, we consider tiered storage systems as a use case. In storage systems flash-based technologies (e.g., DRAM, SSD) are widely integrated into data centers and scaled-out cloud storage systems [6–8]. Despite their performance advantages over the traditional disk- or tape-based storage technologies (e.g., HDD), their much higher cost per byte prevents flash-based storage technologies to completely replace traditional disk or tape based storage devices. Hybrid, tiered architectures that integrate flash with various disk technologies are common alternatives.

Tiered storage systems adopt a hierarchical design: fast tiers using flash storage devices aiming at boosting performance and slow tiers using HDD devices for the purpose of balancing capacity and cost, as well as for improving reliability [9, 10]. Their efficiency is based on moving data to the right tier based on statistics of data access frequencies. Data moving, i.e., *what* portion of data ought to be transferred and *when* this should take place, affect greatly performance as access times and capacities across different tiers differ by orders of magnitude. This greatly depends on how the upcoming workload intensity is known in advance and is vital for *NeuQ*'s success in offering co-scheduling decisions, as the effects of tier warming greatly depend on timely information on this measure.

NeuQ guarantees user SLOs while maximizing data analytics throughput. To this end, we do not treat data analytics simply as a best effort workload. Instead, our aim is to schedule it as aggressively as the system allows without violating user SLOs. We stress that the above performance targets are by no means a limitation of the proposed hybrid model. Incorporating different targets (e.g., deadlines) for completion of data-analytics workload are also easily handled, as we show here.

The proposed approach is fully automatic and robust. We validate its correctness and efficiency via trace-driven simulation using two case studies: (1) enterprise traces from Wikipedia [4] and (2) arrival traces in a scaled-out storage back-end of a mid-size web service provider that we have

also used in prior work.¹ Our extensive experimental evaluation shows that the prediction of the user traffic arrival intensity and data analytics completion time is remarkably accurate. More importantly, compared to other state-of-the-art scheduling approaches, the proposed solution strictly meets user SLOs while serving aggressively data analytics workloads. In addition, our approach also supports different scheduling objectives.

This paper is organized as follows. Section 2 presents experimental numbers of a multi-tiered storage system that motivate this work. Section 3 presents the hybrid model: the machine learning model for traffic intensity prediction and the queueing model. Section 4 presents extensive experimental evaluation via trace driven simulations to verify the correctness and effectiveness of the proposed methodology. Section 5 discusses related work. We conclude in Sect. 6.

2 Motivation

We showcase the challenges of performance interference by presenting some results from an experimental testbed that consists of a server with a disk enclosure attached to it, providing data services to a host. Its memory is 12 GB and the disk enclosure has 12 SATA 7200RPM HDDs of 3TB each. In our experiments the system memory emulates the fast tier and the disk enclosure is used as the slow tier where the bulk of the data resides. The workload is generated and measured at the host machine. We use *fio* [11] as the IO workload generator and generate two types of IO workloads, user and data analytics, which differ on the active working set size rather than on their access patterns. The working set size for the user workload is 1 GB such that it always fits into the memory of the server that emulates the fast tier. The data analytics work has an active working set of 24 GB, i.e., it does not fit fully into the fast tier and the large slow tier is accessed to retrieve the data. The access pattern for both user and data analytics workload is 100 % small random reads to emulate common enterprise workloads that benefit from prefetching (warm-up) only if the working set can fully (or almost) fit in the high performing tier (i.e., the SSD). We also assume that the system is provisioned in such a way that the fast tier can fit the entire (or the majority of) the active user working set. The fast tier is warmed up via a sequential read of the user working set.

Using *fio*, we generate a random reads workload accessing data stored in our server. The intensity of user IO requests is shown in Fig. 1 and it emulates very closely the load pattern of user requests of the Wikipedia and the storage workloads that we use to drive our experiments in Sect. 3. Note that

¹ The Wikipedia traces are publicly available [4]. Due of confidentiality agreements, the storage system trace or provider details can not be made publicly available.




Fig. 1 User and data analytics IOPS (throughput) and user response time over time

without any data analytics work, the response time of user IO requests remains in the same range of about 150ms. All IOs are served from the fast tier. The user throughput however does increase by one order of magnitude as the arrival intensity increases. This confirms that the storage system does not suffer from queuing delays and it has the capacity to sustain more user load.

We add on the same experiment some data analytics work. Since the user workload pattern has a clear periodicity, we contend that this can be easily captured by a machine learning model (indeed, we show the effectiveness of machine learning models for predicting such periodic patterns in Sect. 4.) Having observed the time pattern of the first 60 min, we co-schedule a data analytics workload at around the 65th min, i.e., when the user traffic subsides significantly. While the data analytics workload immediately gains a lot of throughput, the performance effect of co-scheduling on the user work is huge: response times increase by two orders of magnitude. Data analytics work is stopped around the 70th min, but we observe that poor user RTs are sustained for about 6 min. This is because the data analytics work has evicted part of the user working set from the fast tier/cache.

What could have made the difference in user performance is the timely fast-tier warm up. Figure 2 captures this effect. In this experiment, we use a small data set of 1 GB. We perform two experiments: one with explicit cache warm up (via a sequential read to the user working set) and the other one that brings the workload to cache when needed (no explicit




Fig. 2 Response time with warm up and without warm up across the experiment time

warm up). While it takes only about 2 min to bring 1 GB of data into our fast tier by reading it sequentially, it takes about 10 min to fully warm up the cache by the user workload alone. This number is longer than what we observe in Fig. 1, this is because in Fig. 1 the cache is not completely cold. These experiments demonstrate the importance of incorporating the effects of workload interference and tier warming into any scheduling methodology. The point here is that naively scheduling data analytics during the low period (whose timing can be predicted) does not reach optimized performance. What we need here is a methodology that incorporates tier-warming and accurately predicts when this should start. We address these issues in the following sections.

3 Methodology

In this section, we start with the description of the user workloads used in this work. Then we illustrate how to use a neural network to build a traffic prediction model. Finally, we introduce in details of the *NeuQ* scheduler that is powered by a queueing model.

3.1 Workload

Data center workloads often follow periodic patterns across time [5, 12–14]. In Fig. 3, we demonstrate the arrival intensity of the storage system workload during 1 month period [5]. In Fig. 4, we present the arrival intensity of requests to Wikipedia during October 2007 [4].

Intuitively, the workload of Fig. 4 shows a distinctive day/night pattern as well as weekday/weekend pattern. To capture this, we carry out some statistical analysis by calculating the autocorrelation of the time series of the arrival process. Autocorrelation is the cross-correlation of a signal with itself [15]. Intuitively, it captures the similarity between observations as a function of the lag between them. Autocorrelation values are in the $[-1, 1]$ range, the closer the lag autocorrelation to 1, the closer the two observations. Zero values indicate no relationship among the range of values of the two observations is diametrically different.

The autocorrelation of user arrival intensity for the Wikipedia workload at varying time lags is shown in Fig. 5a, b, that report on autocorrelations at the minute lag (10-min granularity) and day lag, respectively. The figures verify clear daily and weekly patterns, as also observed in Fig. 4. Autocorrelation reaches a maximum value per day across all lags illustrating a clear daily pattern, ditto for Fig. 5b that illustrates a clear weekly pattern. Similar autocorrelation patterns are also observed for the storage workload. In the following section, we use these properties to train a neural network that can model user arrival intensity.




Fig. 3 User request arrival intensity of storage workload over 1 month




Fig. 4 User request arrival intensity to Wikipedia over 1 month in 2007

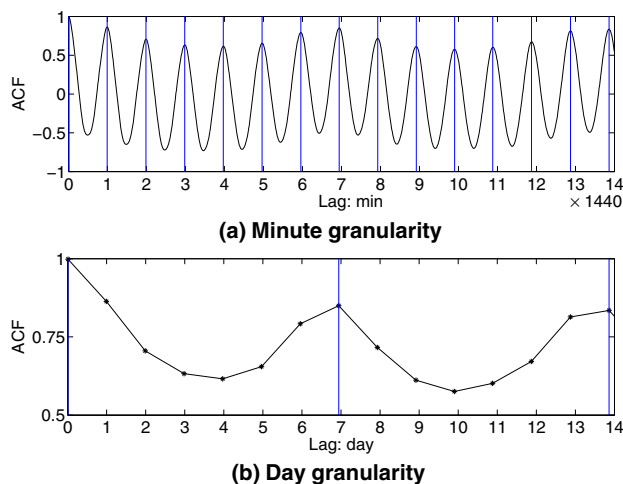


Fig. 5 Autocorrelation of arrivals for different granularities. Note for **a**, the points are at 10-min granularity, the ticks in x-axis is multiplied by 1440, e.g., 2 represents 2×1440 min or 2 days, so the lag is up to 14 days

3.2 Neural network model

Neural networks are a class of machine learning models that can capture periodicity within different time scales. Traditional time series models such as ARMA/ARIMA [16] and Holt-Winters exponential smoothing [17] are usually limited by the linear basis function. Neural networks can instead model non-linear functions of the input, which makes them

effective for time series modeling [18]. A time series analysis consists of two steps: first building a model that represents a time series, and then using the model to forecast future values. If a time series has a regular pattern, then a value of the series should be a function of previous values. If X is the target variable for prediction, and X_t is the value of X at time t , then the goal is to create a model of the form:

$$X_t = f(X_{t-1}, X_{t-2}, X_{t-3}, \dots, X_{t-n}) + \epsilon_t, \quad (1)$$

where ϵ_t is an error term. An artificial neural network consists of multiple nodes, or *neurons*, which are interconnected to mimic a biological neural network [19]. These neurons have adaptive weights, tuned by a learning algorithm, which enables neural networks to approximate non-linear functions of their inputs. The adaptive weights describe the connection strengths between neurons, activated during training and prediction.

To build a neural network model for a time series, selecting the most *relevant* input that can describe the *trend*, *season*, and *noise* is utterly important. To take care of *trend* and *season*, or in other words, to capture the *long-term* pattern, we make use of the correlogram in Fig. 5b. The figure shows that when the lag equals to 7 days, the user traffic is highly and positively correlated. Therefore, as input to our traffic model, we choose the user arrival intensities of exactly one week ago. To capture the temporal change or *noise*, which can be seen as a *short-term* pattern, we look into Fig. 5a and see that for lag

= 10 min the arrival intensities have the highest correlation value. This suggests to consider the user arrival intensities of 10 min ago as another input to the model. With the above inputs as features and current observations, we can train a neural network. When training a neural network, the data are divided into three distinct subsets [20]: the training set, that is used to train the weights (model parameters) of neural network; the validation set, which is used to minimize overfitting thus ensure a generalized solution; and the test set, which is used for evaluating the accuracy of the trained neural network model. Here we use the neural network toolbox provided by MATLAB [21] to train our traffic prediction models.

When training a neural network, even with the same data, the trained model can have different accuracy depending on the configured parameters, e.g., different initial weight values, different divisions of data used for training, validation, and test samples. Our primary attempt to avoid using badly-behaved neural networks (that result in poor prediction), was to train several neural networks on the same data set and select the one with the smallest mean squared error. However, since the future largely remains unknown, it is possible that the “best” neural network fails to do a good prediction in the future. To mitigate this potential problem, we use an ensemble method [22], which averages predictions from different neural networks. Even though the ensemble method may not always provide the optimal prediction, it consistently produces accurate predictions without requiring manual checking.

Last but not least, the computational complexity of the neural network training is not significant, as the prediction model can forecast upcoming traffic for up to a week ahead, suggesting that it is sufficient to update the neural network model as often as once per week for the data in hand. The frequency of training can be adjusted based on different needs.

3.3 Co-scheduling

As Fig. 4 shows, user traffic demonstrates peaks and valleys, suggesting that one could aggressively co-schedule data analytics work during the “low” user periods. Our intention is to quantify how much additional work one could co-schedule such that overall system utilization increases while user SLOs are respected. Because we are aiming to provide a scheduling approach that is easy to deploy and integrate with other management tools, we refrain from making scheduling decisions for user requests too frequently, as this could result in significant overhead. Therefore, our framework divides the time into small windows t_w and makes scheduling decisions only at the beginning of each window²

² We assume $t_w = 1$ min in our experimental evaluation, but this could be adjusted according to the specific system requirement.

3.3.1 Performance model for user traffic in a tiered storage system

To simplify presentation, we assume that the time window that the user SLO is computed is the same as the scheduling window size t_w .

Arrival process: Although in large time windows the arrival process shows a periodical pattern, within each small time window t_w , the arrival process can be viewed as a Poisson process [15].

Service process: In a 2-tiered storage system³, if the working set of the coming IO request is in the fast-tier (e.g., SSD), the request is served in fast tier. Otherwise, the coming IO request is served in the slow-tier (e.g., disk). Here we assume the service process for each tier follows an exponential distribution. The service process for the 2-tiered system can be described by a hyper-exponential model [23] with mean service time:

$$E[s] = h \times E[s_1] + (1 - h) \times E[s_2], \quad (2)$$

where $E[s_1]$ and $E[s_2]$ are the mean service times for the fast tier (tier 1) and slow tier (tier 2), respectively. h is the fast tier hit rate defined as the probability that a request is served from the fast tier. The maximum value of the fast tier hit rate h_{max} is determined by the workload characteristics and the capacity of the fast tier, e.g., if the entire working set can be loaded into the fast tier, then the maximum hit rate is 1, in which case all the requests are served in the fast tier. Based on Eq. 2, the expectation of the squared service time can be computed as follows [23]:

$$E[s^2] = 2! \times (h \times E[s_1]^2 + (1 - h) \times E[s_2]^2), \quad (3)$$

where $E[s_1]^2$ and $E[s_2]^2$ are the square of mean service times for the fast tier (tier 1) and slow tier (tier 2), respectively.

Queuing model: Based on the above assumptions, the 2-tiered storage system can be modeled by a $M/H_2/1$ queue for each small time window. We use the Pollaczek–Khinchine formula [24] to compute the average response time of an $M/H_2/1$ queue:

$$RT = E[s] + \frac{\lambda \times E[s^2]}{2(1 - \rho)}, \quad (4)$$

where $E[s]$ is the mean service time, $E[s^2]$ is the mean of squared service time, λ is average arrival intensity, and ρ is the system utilization

³ For presentation reasons, we use a 2-tiered storage system to explain our methodology, but it could be easily extended to storage systems with more tiers. This discussion also applies to caching.

$$\rho = \lambda \times E[s]. \quad (5)$$

Combining Eqs. 2, 3, and 5 into Eq. 4, we have:

$$RT = h \times E[s_1] + (1 - h) \times E[s_2] + \frac{\lambda \times (h \times E[s_1]^2 + (1 - h) \times E[s_2]^2)}{1 - \lambda \times (h \times E[s_1] + (1 - h) \times E[s_2])}. \quad (6)$$

From Eq. 6, it is clear that within each time window t_w , the average user response time RT is a function of the average arrival intensity of user requests λ , the mean service time for each tier ($E[s_1]$ and $E[s_2]$), and the fast tier hit rate h .

3.3.2 NeuQ scheduler

Co-scheduling data analytics work may have a performance impact on the user traffic. From our performance model above, the first three parameters (λ , $E[s_1]$, and $E[s_2]$) only depend on the characteristics of user workload and performance of hardware devices, so co-scheduling has no impact on these parameters. However, co-scheduling data analytics work does impact the fast tier hit rate h because it evicts the user working set from the fast tier and reduces the probability of serving user requests in the fast tier. Therefore, in order to meet the user SLO (RT_{target}), the fast tier hit rate of user traffic needs to be maintained above a threshold that is computed from the user SLO and arrival intensity. We compute the threshold (h_{target}) based on Eq. 6 by representing the fast tier hit rate as a function of the user response time target and arrival intensity:

$$h_{target} = \frac{\lambda \times (RT_{target} + E[s_1] - E[s_2]) - \sqrt{P}}{2\lambda \times (E[s_1] - E[s_2])}, \quad (7)$$

where:

$$P = \lambda^2 \times (RT_{target}^2 + 2RT_{target} \times E[s_1] + 2RT_{target} \times E[s_2] + (E[s_1] - E[s_2])^2) - 2\lambda \times (RT_{target} - E[s_1] - E[s_2]) + 1. \quad (8)$$

Note that the fast tier hit rate depends on both user response time target and arrival intensity because the latency is composed of service time and queuing waiting time, e.g., during periods of high arrival intensity, a higher fast tier hit rate is needed to achieve the same response time target than during low arrival intensity periods.

Now the question becomes how to make scheduling decisions such that the data analytics work can be completed as fast as possible without affecting the user SLO. Because the data analytics work is usually measured by the average completion time of submitted jobs or the throughput, it is important to quantify and maximize the (cumulative) time

slot allocated to the analytics work within each time window t_w . Figure 2 and its corresponding analysis in Sect. 2 shows the benefit of explicit warm up compared to no explicit warm up. When the system is not as busy with user requests, there are 2 choices: (i) scheduling the data analytics work for time t_{DA} , which reduces the fast tier hit rate or (ii) explicit warming up the fast tier, which recovers the fast tier hit rate.

In order to maximize the time slot allocated to data analytics work but maintain the fast tier hit rate above the threshold h_{target} , different scheduling choices need to be made based on the arrival intensity in the near future. Recall that the future arrival intensity can be predicted using the neural network model introduced in Sect. 3.2. Therefore, we can estimate the fast tier hit rate h_{DA} after scheduling for t_{DA} time units data analytics work:

$$h_{DA} = h_{begin} - \frac{t_{DA}}{t_{evict}}, \quad (9)$$

where h_{begin} is the fast tier hit rate at the beginning of a time window and t_{evict} is the time to evict the entire user working set from a fully warmed up fast tier. Since the data analytics work is very intensive, t_{evict} can be approximated by the time to explicitly warm up the fast tier from completely cold to fully warmed up, which we define as $t_{warmup1}$, this can be easily obtained by a quick profiling experiment such as the one shown in Fig. 2. Assuming that the remaining of idle time is used for explicit warm up of the fast tier, the fast tier hit rate h_{warmup} is:

$$h_{warmup} = \min \left\{ h_{begin} + \frac{(1 - \lambda \times E[s]) \times t_w - t_{DA}}{t_{warmup1}}, h_{max} \right\}, \quad (10)$$

recall that h_{max} is the maximum value of the fast tier hit rate. In addition, serving user requests also changes the fast tier hit rate (no explicit warm up):

$$h_{user} = h_{begin} + \frac{\lambda \times E[s] \times t_w}{t_{warmup2}}, \quad (11)$$

where $t_{warmup2}$ is the time to non-explicitly warm up the fast tier from completely cold to fully warmed up, which can be easily obtained by a quick profiling experiment. Combining Eqs. 9, 10, and 11, we have the fast tier hit rate at the end of the time window h_{end} :

$$h_{end} = \min \left\{ h_{begin} - \frac{t_{DA}}{t_{warmup1}} + \frac{(1 - \lambda \times E[s]) \times t_w - t_{DA}}{t_{warmup1}} + \frac{\lambda \times E[s] \times t_w}{t_{warmup2}}, h_{max} \right\}. \quad (12)$$

Since each time window is very small, we assume that the fast tier hit rate only changes at the end of the window, but stays the same within the window. To meet the SLO, the fast tier hit rate needs to be maintained at or above the threshold i.e., $h_{end} \geq h_{target}$, therefore

$$t_{DA} \leq ((h_{begin} - h_{target}) \times t_{warmup1} \times t_{warmup2} + (1 - \lambda \times E[s]) \times t_w \times t_{warmup2} + \lambda \times E[s] \times t_w \times t_{warmup1}) \times \frac{1}{2 \times t_{warmup2}}. \quad (13)$$

From the above inequality, we can quantify the maximum amount of time to be allocated to data analytics work without violating user SLOs. The co-scheduling decisions are based on Eq. 13. We stress that this inequality is *critical* for the success of workload co-scheduling as it regulates the amount of data analytics work that the system can sustain in order to not violate the user SLO.

Because the targeted fast tier hit rate changes with the arrival intensity, scheduling decisions need to be evaluated well in advance. Here we determine how early a scheduling decision needs to be evaluated such that there is enough time to fully warm up the fast tier during this period. We define $t_{advance}$ as:

$$t_{advance} = (h_{max} - h_{current}) \times t_{warmup1}, \quad (14)$$

where $h_{current}$ is the current fast tier hit rate. Based on the predicted arrival intensity after $t_{advance}$, a correct scheduling decision (the time allocated to data analytics work) for the current time window can be made.

In addition, based on the time slot allocated to the data analytics work (Eq. 13), we can estimate the throughput of data analytics work $Throughput_{DA}$ as follows:

$$Throughput_{DA} = \frac{t_{DA}}{s_{DA}} \quad (15)$$

where s_{DA} is the average service time of data analytics work.

If the scheduling target is meeting deadlines for data analytics work, then the throughput of data analytics work needs to meet the above requirements. Assume that the throughput requirement is $Throughput_{target}$, then we have the following:

$$Throughput_{DA} = \frac{t_{DA}}{s_{DA}} \geq Throughput_{target}, \quad (16)$$

therefore,

$$t_{DA} \geq Throughput_{target} * s_{DA}. \quad (17)$$

The above inequality indicates the minimum amount of time to be allocated to the data analytics work in order to meet the deadlines.

3.3.3 NeuQ scheduler with capacity planning

The interesting question is whether it is possible to meet both user response time SLO RT_{target} and data analytics work throughput SLO $Throughput_{target}$ as in practice, each workload usually has its own performance target. Such scheduling target is achievable through capacity planning, thanks to the elastic storage techniques [25]. When the size of fast-tier increases, serving data analytics work has less impact on evicting user working set that resides in fast-tier, so increasing the capacity of fast-tier allows scheduling more aggressively data analytics work during high user traffic intensity periods. Let's denote v_{evict} as the speed of user working set being evicted from the fast tier. v_{evict} can be computed as: $v_{evict} = \frac{\delta h}{t_f}$, where δh is the change of user work hit rate during the data analytics work serving time period t_f . Recall t_{evict} is the time to evict the entire user working set from a fully warmed up fast-tier (with hit rate of $h_{upper} = 1$) to completely cold (with hit rate of $h_{lower} = 0$), so we have: $v_{evict} = \frac{h_{upper} - h_{lower}}{t_{evict}} = \frac{1}{t_{evict}}$. Let's denote fast-tier capacity as C , then the time to evict all data from fast-tier t_{all} can be computed as: $t_{all} = \frac{C}{\mu_c}$, where μ_c is the service rate of the fast tier. Since the time to evict all data is equal to the time to evict user working set from fully cached to empty, we have:

$$t_{evict} = \frac{C}{\mu_c}. \quad (18)$$

For each time window:

$$h_{end} = h_{begin} + \frac{\rho * t_w}{t_{warmup1}} + \frac{(1 - \rho) * t_w - t_{DA}}{t_{warmup2}} - \frac{t_{DA}}{t_{evict}}. \quad (19)$$

Together with Eqs. 13, 17, 18, and 19, it is possible to compute the minimum fast-tier capacity to achieve the SLOs of both user work and data analytics work.

4 Performance evaluation

An integral part of the effectiveness of workload co-scheduling is the ability to predict accurately the upcoming user workload. We first evaluate the accuracy of the prediction model and then compare the performance of the proposed approach with other methods in the literature. We also show several scenarios that illustrate the effectiveness of NeuQ. Finally, we demonstrate NeuQ can support different scheduling targets.

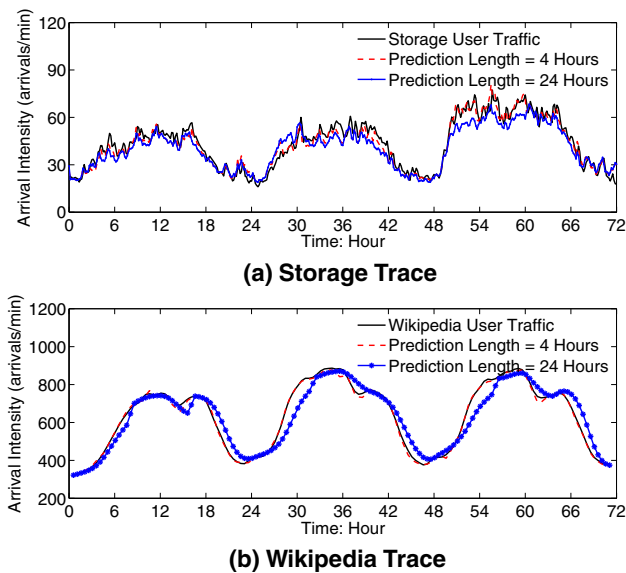


Fig. 6 Traffic predictions from neural networks with different prediction lengths.

4.1 Traffic prediction

We drive our simulations using the storage workload in [5] and the Wikipedia trace that is shown in Fig. 4 (both workloads with granularity of 1 min). We trained neural networks for the two workloads and used the trained models to predict incoming traffic for the time period of the next three upcoming days. We also illustrate the model with two different prediction lengths, i.e., 4 and 24 h. Here, if the prediction length equals to K hours, the neural network directly predicts the arrival intensities in the next K hours. In the proposed *NeuQ* scheduler, the traffic intensity information is usually needed only a few hours ahead, so such prediction length can fully satisfy the scheduler's needs. Consistent with the workload analysis of Sect. 3.1, for each observation during training, we select the observations from the most recent time window and the one from one week ago as features. In Fig. 6, we show the traffic predictions with two prediction lengths, as well as the actual traffic. The figure illustrates that for both of the storage and Wikipedia traces, the shorter the prediction length, the more accurate the predictions. Yet, even predicting 24 h ahead, the overall prediction accuracy is still good.

To quantify better the quality of the two prediction lengths, we use the *absolute percentage error* (APE) metric which is defined as:

$$APE = \frac{|Prediction - Actual|}{Actual} \times 100\%. \quad (20)$$

When APE is close to 0, it suggests an accurate prediction. Figure 7 illustrates the CDFs of the absolute percentage error for the two traffic predictions. For the storage trace, Fig. 7a

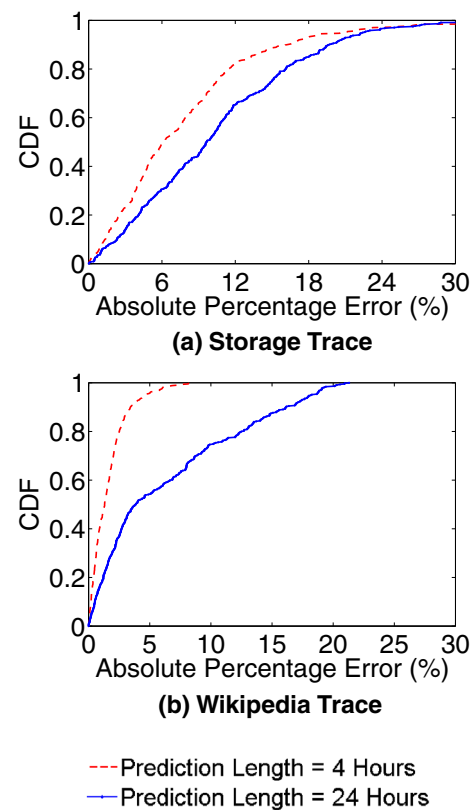


Fig. 7 CDF of absolute percentage error of the neural network predictions

shows that the 80 percentile of the APE for the 4 and 24-h predictions, is less than 12 and 18 % respectively. For the Wikipedia trace, Fig. 7b shows that for the 4-h prediction length, almost all the APEs are no more than 10 %, i.e., predictions are very accurate. For the 24-h case, over 70 % of the APEs are less than 10 %, and nearly 100 % of them are no more than 20 %.

4.2 Scheduler comparisons

Here, we compare our *NeuQ* scheduler with two other state-of-the-art scheduling approaches in the literature. We also compare to a scheduling approach that only uses neural network traffic prediction without any help from queueing models.

– On-line Feedback

The on-line feedback method [26] collects measurements during the most recent time window to check whether user performance requirements are met or violated. If met and provided that the system is in a low utilization state in the current window, then data analytics work is added. In the interest of showing the maximum benefit of this scheduler, we assume that we know all

future upcoming workload, i.e., we are certain about the length of low/high utilization times due to user workload *a priori*.

– State Prediction

The state change prediction based method [5] divides the traffic into high/low states, based on the average arrival intensity. It makes use of a Markovian Modulated Poisson Process (MMPP) model to predict when the high state starts and ends. Based on the state prediction, this method only schedules the data analytics work in the low utilization state, while pro-actively warming up the fast tier with the active user working set right before the arrival of a high utilization state.

– NeuralNet Scheduler

The NeuralNet scheduler decides the amount of data analytics work to schedule based on the user arrival intensity prediction. Here the amount of data analytics work to schedule in each time window increases/decreases linearly with the predicted user arrival intensity in that time window. For example, if for the current window w_1 , λ_1 is the average user arrival intensity, and the amount of scheduled data analytics work is n_1 , then for the incoming time window w_2 with predicted average user arrival intensity equal to λ_2 , the NeuralNet scheduler schedules $\frac{\lambda_1}{\lambda_2} \times n_1$ data analytics work in w_2 . Intuitively, if the user arrival intensity is predicted to increase, then the NeuralNet scheduler schedules less data analytics work in the incoming time window. Otherwise, more data analytics work is scheduled.

– NeuQ Scheduler

The NeuQ scheduler makes scheduling decisions (how much and when to schedule data analytics work) based on the hybrid model developed in Sect. 3.

We conduct trace-driven simulations on the storage workload in [5] and the Wikipedia user traffic (days 26, 27, and 28 in Fig. 4). Because the NeuQ scheduler and the state prediction based scheduler depend on the accuracy of their workload model, we illustrate in Fig. 8 how well the MMPP model predicts changes in system state as well as the neural network prediction, both methods are quite effective.

We start by comparing how fast data analytics work each method can complete, given an SLO for user performance. We assume that the service times for the fast and slow tiers are exponential distributed, and that the average service time between the two tiers differs by two orders of magnitude. For the data analytics work, we assume that the average time to finish one unit of work equals to the mean slow tier service time and we use throughput to measure how fast it can be scheduled. Figure 9 illustrates the user RTs and data analytics throughputs for the four policies. In each graph, the horizontal lines represent the pre-defined user SLOs. Figure 9a

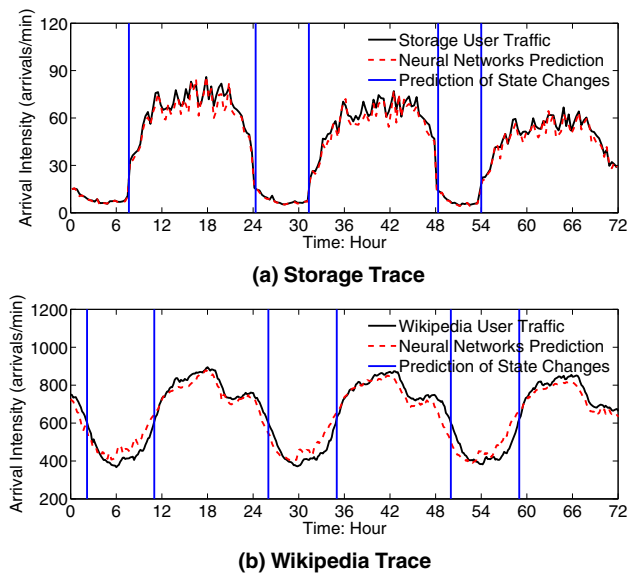


Fig. 8 Storage system and Wikipedia user traffic for simulation, with the state changes and traffic predictions

shows that during heavy user traffic periods (e.g., the 6th, 30th and 54th h), user SLOs are consistently violated due to the aggressive scheduling of data analytics. Similar phenomena are observed in Fig. 9e. The system recovers after the violation is detected. The MMPP-based state prediction method, see Fig. 9b, f, is more conservative and refrains from scheduling data analytics work while the system is in high utilization. Since this scheduler pro-actively warms up the fast tier, it contains SLO violations. However, the stochastic nature of the underlying Markovian-based model results in unavoidable inaccuracies for the *exact* time that the system changes, which results in SLO violations when the high state approaches. In Fig. 9c, g, the NeuralNet scheduler is evaluated. The figures show that although less data analytics work is scheduled in heavy user traffic periods than during hours of low traffic, the user SLO is still violated. Figs. 9d and h illustrate the performance of NeuQ. With NeuQ, there are no SLO violations while data analytics work is served aggressively, and with even higher throughput as with on-line feedback. We also list the percentage of user SLO violations, throughput of data analytics, and its coefficient of variation (CV) for the above four policies in Table 1. NeuQ achieves $2 \times$ to $3 \times$ higher data analytics throughput without any user SLO violation, while for the other policies, the percentage of user SLO violations is much higher and still less efficient on data analytics throughput.

To further illustrate the advantages of our proposed method, we present the CCDFs that illustrate the tails of average user response times of each time window in Fig. 10. NeuQ manages to strictly respect the user SLO (1000 ms for the storage trace and 75 ms for the Wikipedia trace).

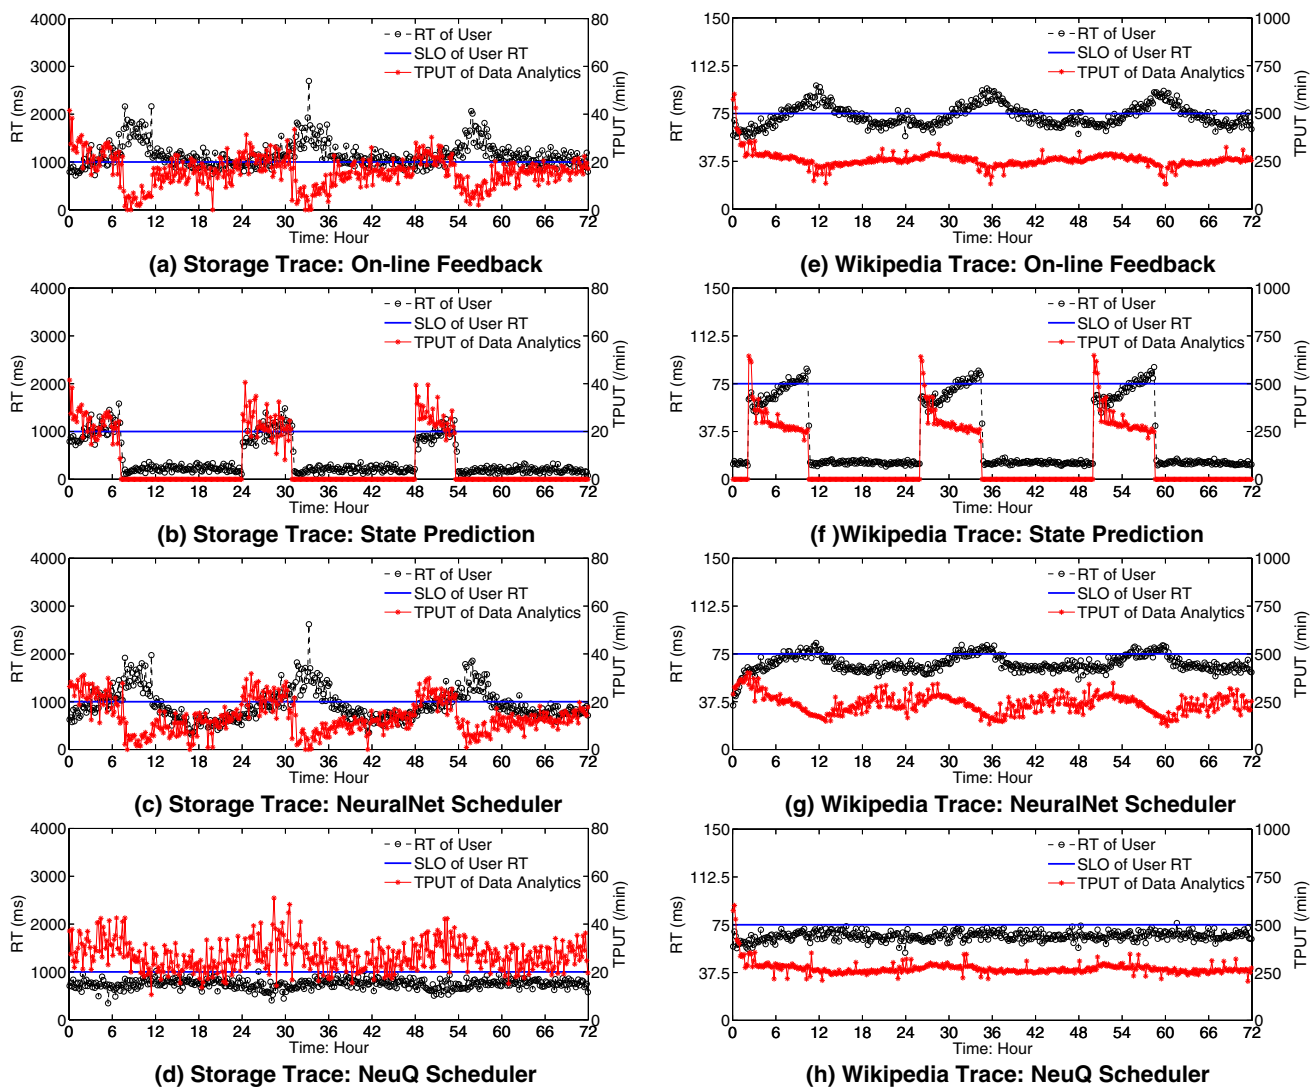


Fig. 9 Performance comparisons via simulation

4.3 Model effectiveness

The performance prediction model that is developed in Sect. 3 is the core of the proposed *NeuQ* scheduler. Recall that the model allows us to regulate the amount of data analytics work to be co-scheduled such that a certain SLO is met for the user work. Similarly, if one needs to increase the throughput of data analytics, this would result in affecting the user RT as well.

Figure 11 illustrates this relationship between user RT (x-axis) and data analytics throughput (y-axis) for the storage and Wikipedia workloads. The figure plots the relationship of these measures as obtained both by simulation and by using the prediction model of Sect. 3 (Eq. 15). Both model and simulation numbers are in good agreement, well-capturing the relationship trends of both measures. Further, by using Fig. 11 one could estimate the maximum data analytics

throughput to be achieved given a certain SLO or conversely the sustained average user RT if a certain throughput for data analytics work is expected.

4.4 Different scheduling targets


NeuQ can be used to support different scheduling targets. We first demonstrate the scheduling target scenario of finishing the data analytics work by the pre-defined deadlines while preserving the performance of user workload as much as possible. We use Storage trace for evaluation and show two different scheduling targets:

- *Scheduling Target 2*: the deadline of data analytics work is 15 min for 500 units of work. The SLO of user workload is SLO 1350 ms.

Table 1 Average performance analysis for simulations

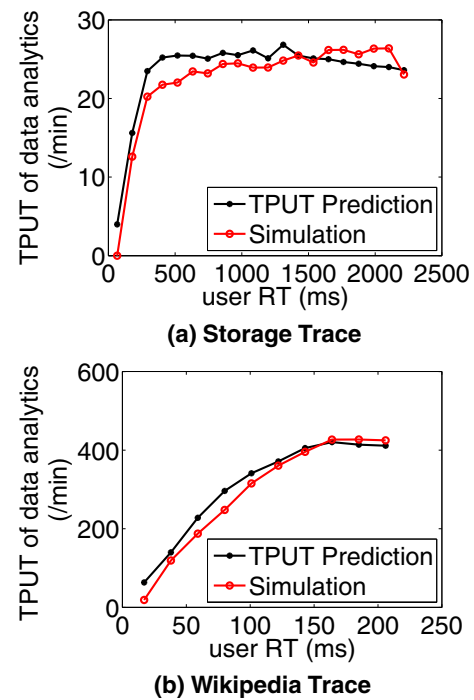
	Storage trace		
	% User SLO Violations	$TPUT_{DA}$ (/min)	CV of $TPUT_{DA}$
On-line feedback	68.29	18.24	0.42
State prediction	11.11	7.58	1.71
NeuralNet scheduler	32.41	15.53	0.52
NeuQ scheduler	0.00	26.45	0.23

	Wikipedia trace		
	% User SLO Violations	$TPUT_{DA}$ (/min)	CV of $TPUT_{DA}$
On-line feedback	40.05	257.38	0.16
State prediction	11.81	110.99	1.45
NeuralNet scheduler	15.51	240.71	0.21
NeuQ scheduler	0.00	271.69	0.13

**Fig. 10** CCDF of average user response time

- *Scheduling Target 2*: the deadline of data analytics work is 1.5 h for 1500 units of work. The SLO of user workload is SLO 850 ms.

Scheduling Target 1 has a tighter deadline for data analytics work while relatively loose SLO for the user workload, and Scheduling Target 2 is on the contrary. Both scheduling sce-

**Fig. 11** Throughput of data analytics versus user response time: model (Eq. 15) and simulation (*NeuQ* scheduler)

narios consider meeting the deadline of data analytics work as first priority and meeting SLO of user workload as a secondary target. The data analytics work that is not finished by the deadline is dropped so that its impact is not propagated to the future analytics work.

The results are presented in Fig. 12. The left y-axis is the user response time measured in ms and right y-axis is the percentage of the finished data analytics work. The x-axis represents the elapsed time (3-day period). For Target 1 (left column of graphs), the results suggest that only *NeuQ* can meet the deadlines of data analytics work (there are a few exceptions, but very few and all above 80 %) and all other methods fail. Meanwhile, *NeuQ* also consistently archives the SLO of user workload, this is not the case for other methods. Target 2 (right column of graphs) shares the same requirement for meeting deadlines of data analytics work. Given the stricter SLO, none of the methods can meet the deadlines of data analytics work while also achieving the SLO for user workload. However, it is clear that *NeuQ* results in smallest violation of the SLO of user workload among these methods. The above experiments suggest that *NeuQ* has great use potential in reaching other targets. For the experiments presented in Fig. 12, especially for Scheduling Target 1, it is necessary to tolerate high user SLOs to meet the 15 min deadline.

Then we evaluate the scheduling target scenario of meeting both user response time target and data analytics work throughput target through capacity planning. We

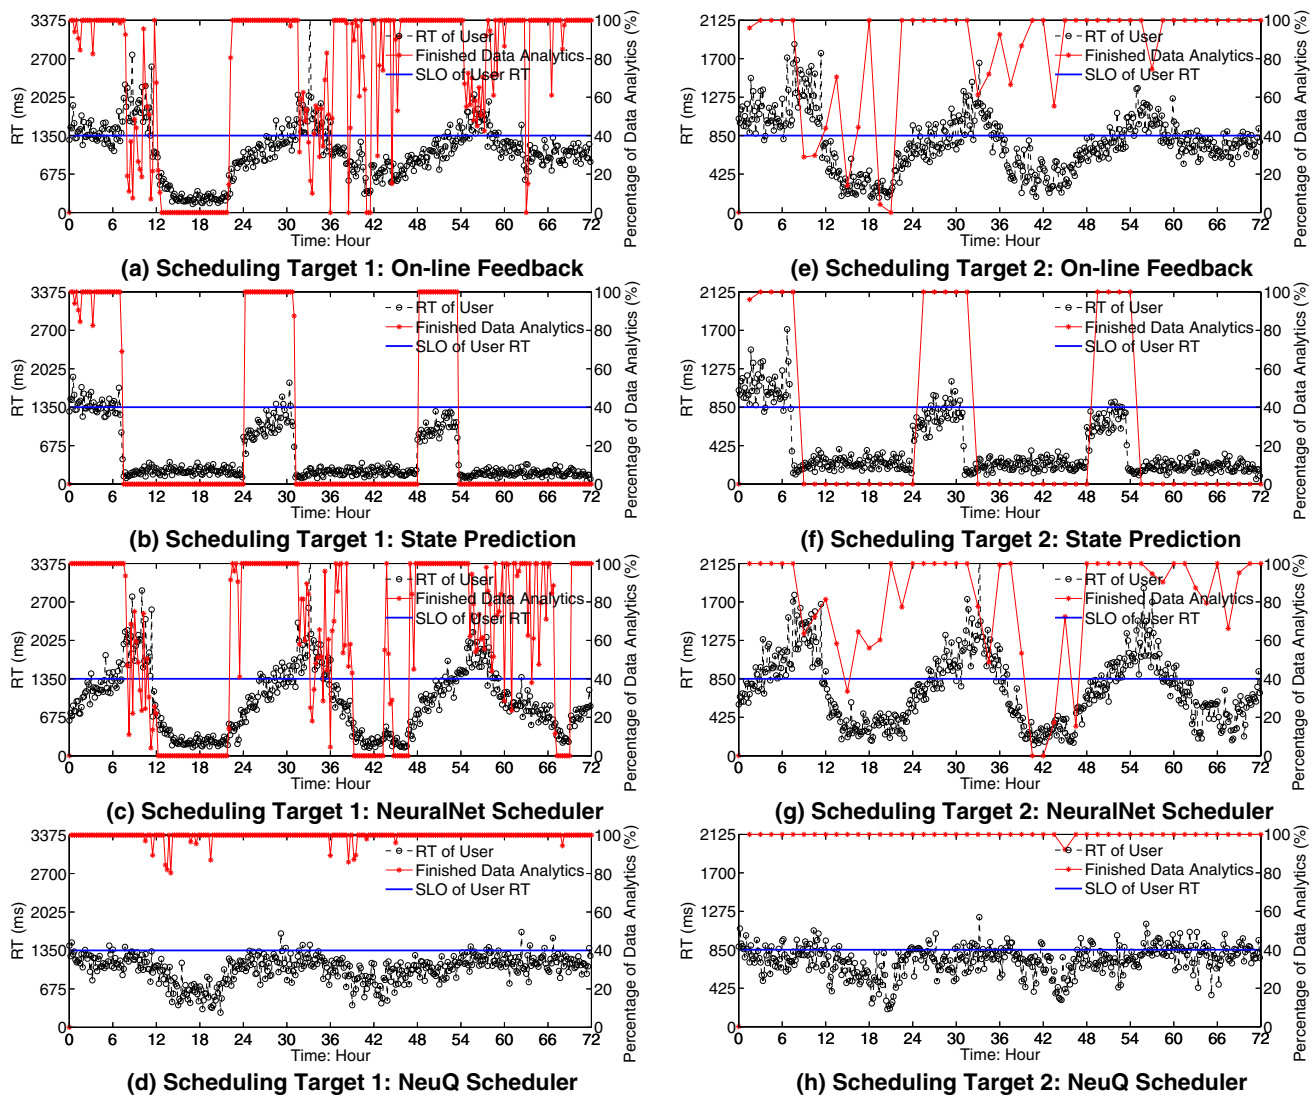


Fig. 12 Performance comparisons via simulation. Scheduling Target 1 (left column of graph): the deadline of data analytics work is 15 min for 500 units of work and the user workload SLO is 1350 ms. Scheduling

Target 2 (right column of graph): the deadline of data analytics work is 1.5 h for 1500 units of work and the user workload SLO is 850 ms

use Wikipedia trace for evaluation and show six different scheduling targets:

- *Scheduling Target 1*: the deadline of data analytics work is 10 min for 3200 units of work. The SLO of user workload is SLO 75 ms.
- *Scheduling Target 2*: the deadline of data analytics work is 10 min for 3600 units of work. The SLO of user workload is SLO 75 ms.
- *Scheduling Target 3*: the deadline of data analytics work is 10 min for 4000 units of work. The SLO of user workload is SLO 75 ms.
- *Scheduling Target 4*: the deadline of data analytics work is 10 min for 3200 units of work. The SLO of user workload is SLO 105 ms.

- *Scheduling Target 5*: the deadline of data analytics work is 10 min for 3600 units of work. The SLO of user workload is SLO 105 ms.
- *Scheduling Target 6*: the deadline of data analytics work is 10 min for 4000 units of work. The SLO of user workload is SLO 105 ms.

These scheduling targets represent a variety selection of different user workload SLO and data analytics work SLO combinations. The results are presented in Fig. 13. The left y-axis is the user response time measured in ms and right y-axis is the percentage of the finished data analytics work. The x-axis represents the elapsed time in 3-day period. Due to the interest of space, the results of other methods are not shown here. Because other methods do not take into consid-

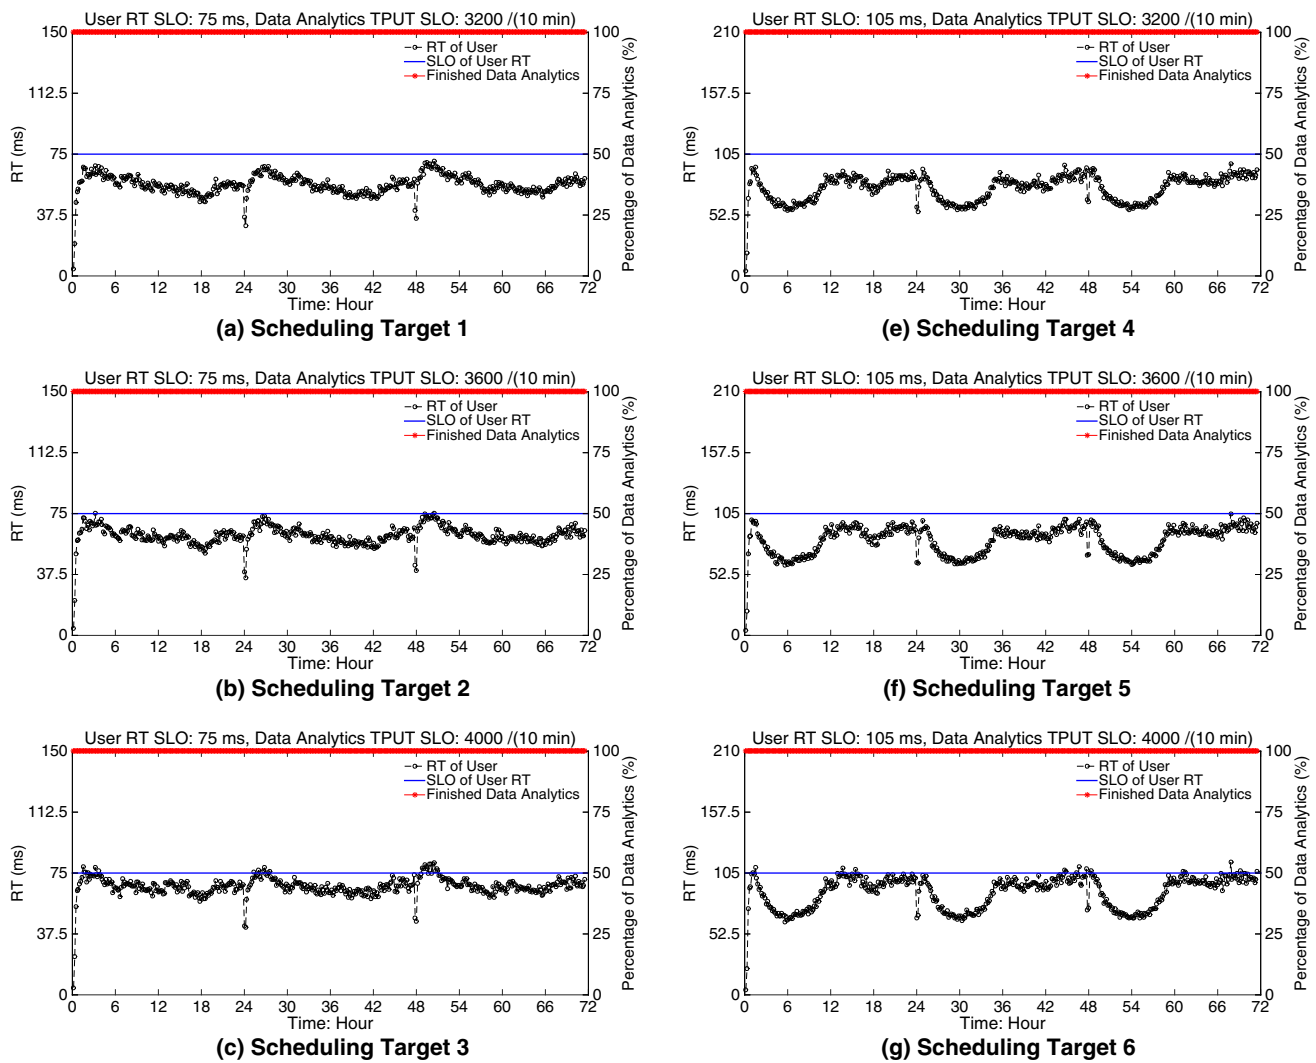


Fig. 13 User response time and finished data analytics work under various scheduling targets

eration of both user workload SLO and data analytics work SLO, as expected, the SLOs could not be met at the same time. For *NeuQ*, it consistently archives the SLOs of both user workload and data analytics work in different scenarios, see Fig. 13. We also plot the user working set evicting speed from fast-tier and fast-tier capacity demands overtime in Figs. 14 and 15 respectively. These plots show how *NeuQ* correctly estimates the user working set evicting speed and fast-tier capacity demands under different system load so that to plan in advance the fast-tier capacity to accommodate both user workload and data analytics work in each time window.

5 Related work

Analytical and simulation models have been widely used to quantify the impact of workload changes to application and/or system performance, see [5,9,13,27–30] and ref-

erences therein. [5] uses a probabilistic model to define “workload states” via hierarchical clustering. After state discovery, the observed workload is used to parameterize a Markov Modulated Poisson Process that can accurately predict the duration of each state as well as transitions from state to state. ARMA/ARIMA [16] have been adopted in [3] to predict the user traffic overtime in order to achieve cost-efficient capacity planning. However, this prediction method is limited to the linear basis function.

Machine learning techniques are used to overcome the limitations of the linear basis function of ARMA/ARIMA models, and are used for effective characterization of TCP/IP [31] and web server views [32]. Machine learning techniques [33] have been also used for performance prediction of total order broadcast, a key building block for fault-tolerant replicated systems. Ensembles of time series models have been used to project disk utilization trends in a cloud setting [22].

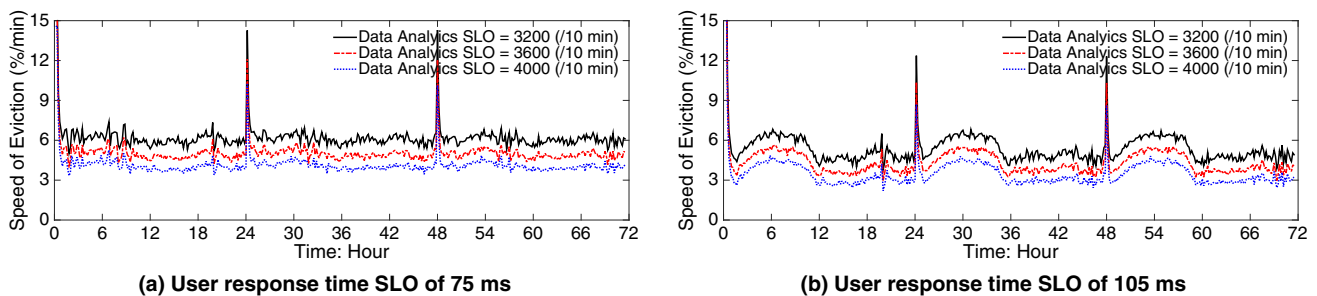


Fig. 14 User working set evicting speed from fast-tier overtime with different scheduling targets

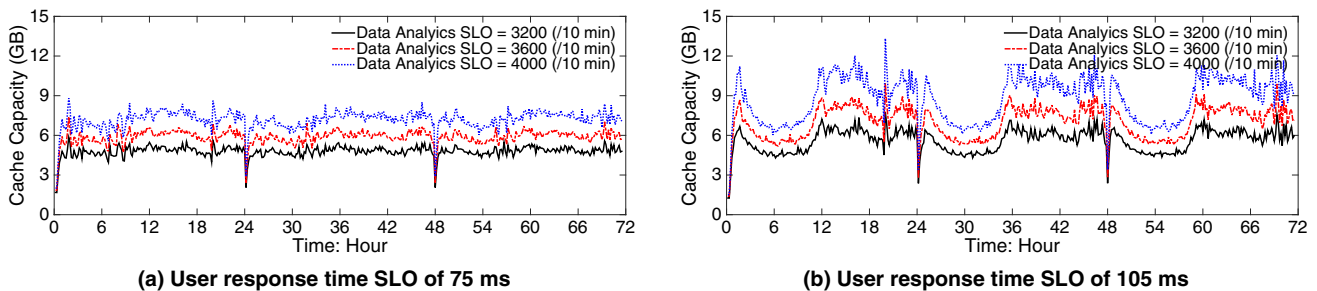


Fig. 15 Fast-tier capacity demands overtime with different scheduling targets

In general, analytical models are restricted by their simplified assumptions while machine learning models are effective in predicting performance for scenarios that have already been observed in the past and fail when new patterns are observed. A gray-box performance model that combines analytical modeling with machine learning has been proposed [34]. The authors advocate the use of analytical models to lower the initial training time of machine-learning based predictors or enhance the accuracy of the analytic model by adjusting its error with the help of machine learning. In contrast to this work, what we propose here is the usage of machine learning to accurately predict specific inputs of a queueing model, which in turn we use to derive scheduling decisions.

6 Conclusion

Co-scheduling data analytics workloads with user workloads in multi-tiered storage systems is a challenging problem. In this paper, we propose *NeuQ* scheduler, a hybrid co-scheduling approach using machine learning and queueing models, that applies neural networks to predict user workload intensities and then appropriately adjusts the input to a queueing model in order to consistently meet user SLOs. Trace-driven simulations show that *NeuQ* can effectively reach performance targets under different user workloads and different performance/scheduling targets from commercial systems.

Acknowledgments This work is supported by NSF Grant CCF-1218758.


References

1. Herodotou, H., Lim, H., Luo, G., Borisov, N., Dong, L., Cetin, F.B., Babu, S.: Starfish: a self-tuning system for big data analytics. *CIDR* **11**, 261–272 (2011)
2. Cohen, J., Dolan, B., Dunlap, M., Hellerstein, J.M., Welton, C.: Mad skills: new analysis practices for big data. *Proc. VLDB Endow.* **2**(2), 1481–1492 (2009)
3. Zhuang, Z., Ramachandra, H., Tran, C., Subramaniam, S., Botev, C., Xiong, C., Sridharan, B.: Capacity planning and headroom analysis for taming database replication latency: experiences with linkedin internet traffic. In: *Proceedings of the 6th ACM/SPEC ICPE*, pp. 39–50 (2015)
4. Urdaneta, G., Pierre, G., van Steen, M.: Wikipedia workload analysis for decentralized hosting. *Elsevier Comput. Netw.* **53**(11), 1830–1845 (2009)
5. Xue, J., Yan, F., Riska, A., Smirni, E.: Storage workload isolation via tier warming: how models can help. In: *Proceedings of the 11th ICAC*, pp. 1–11 (2014)
6. Peters, M.: 3PAR: optimizing I/O service levels. ESG White Paper (2010)
7. Laliberte, B.: Automate and optimize a tiered storage environment—FAST! ESG White Paper (2009)
8. Amazon ElastiCache. <http://aws.amazon.com/elasticache>. Accessed 11 Mar 2015
9. Guerra, J., Pucha, H., Glider, J.S., Belluomini, W., Rangaswami, R.: Cost effective storage using extent based dynamic tiering. In: *FAST*, pp. 273–286 (2011)
10. Oh, Y., Choi, J., Lee, D., Noh, S.H.: Caching less for better performance: balancing cache size and update cost of flash memory cache in hybrid storage systems. In: *FAST*, pp. 313–326 (2012)

11. FIO Benchmark. <http://www.freecode.com/projects/fio>. Accessed 11 Mar 2015
12. Björkqvist, M., Chen, L.Y., Binder, W.: Opportunistic service provisioning in the cloud. In: 5th IEEE CLOUD, pp. 237–244 (2012)
13. Ansaloni, D., Chen, L.Y., Smirni, E., Binder, W.: Model-driven consolidation of java workloads on multicores. In: 42nd IEEE/IFIP DSN, pp. 1–12 (2012)
14. Birke, R., Björkqvist, M., Chen, L.Y., Smirni, E., Engbersen, T.: (Big)data in a virtualized world: volume, velocity, and variety in cloud datacenters. In: *FAST*, pp. 177–189 (2014)
15. Leemis, L.M., Park, S.K.: Discrete-Event Simulation: A First Course. Pearson Prentice Hall, Upper Saddle River (2006)
16. George, B.: Time Series Analysis: Forecasting & Control, 3rd edn. Pearson Education India, Gurgaon (1994)
17. Goodwin, P.: The holt-winters approach to exponential smoothing: 50 years old and going strong. In: *Foresight*, pp. 30–34 (2010)
18. Frank, R.J., Davey, N., Hunt, S.P.: Time series prediction and neural networks. *J. Intell. Robot. Syst.* **31**(1–3), 91–103 (2001)
19. Hassoun, M.H.: Fundamentals of Artificial Neural Networks, 1st edn. MIT Press, Cambridge (1995)
20. Hill, T., O'Connor, M., Remus, W.: Neural network models for time series forecasts. *Manag. Sci.* **42**(7), 1082–1092 (1996)
21. Demuth, H., Beale, M., Hagan, M.: Neural network toolboxTM 6, User Guide
22. Stokely, M., Mehrabian, A., Albrecht, C., Labelle, F., Merchant, A.: Projecting disk usage based on historical trends in a cloud environment. In: Proceedings of the 3rd Workshop on ScienceCloud, pp. 63–70 (2012)
23. Ross, S.M.: Introduction to Probability and Statistics for Engineers and Scientists. Academic Press, Cambridge (2009)
24. Tijms, H.C.: A first course in stochastic models. Wiley, New York (2003)
25. Lim, H.C., Babu, S., Chase, J.S.: Automated control for elastic storage. In: Proceedings of the 7th ICAC. ACM, pp. 1–10 (2010)
26. Cucinotta, T., Checconi, F., Abeni, L., Palopoli, L.: Self-tuning schedulers for legacy real-time applications. In: Proceedings of the 5th EuroSys, pp. 55–68 (2010)
27. Ferrer, A.J., Hernández, F., Tordsson, J., Elmroth, E., Ali-Eldin, A., Zsigri, C., Sirvent, R., Guitart, J., Badia, R.M., Djemame, K., et al.: Optimis: a holistic approach to cloud service provisioning. *Futur. Gener. Comput. Syst.* **28**, 66–77 (2012)
28. Singh, R., Shenoy, P., Natu, M., Sadaphal, V., Vin, H.: Analytical modeling for what-if analysis in complex cloud computing applications. *ACM SIGMETRICS Perform. Eval. Rev.* **40**(4), 53–62 (2013)
29. Zhang, Q., Cherkasova, L., Smirni, E.: A regression-based analytic model for dynamic resource provisioning of multi-tier applications. In: Proceedings of the 4th ICAC, pp. 27–36 (2007)
30. Yan, F., Riska, A., Smirni, E.: Busy bee: how to use traffic information for better scheduling of background tasks. In: Proceedings of the 3rd ACM/SPEC International Conference on Performance Engineering, pp. 145–156 (2012)
31. Cortez, P., Rio, M., Rocha, M., Sousa, P.: Multi-scale internet traffic forecasting using neural networks and time series methods. *Expert Syst.* **29**(2), 143–155 (2012)
32. Li, J., Moore, A.W.: Forecasting web page views: methods and observations. *J. Mach. Learn. Res.* **9**(10), 2217–2250 (2008)
33. Couceiro, M., Romano, P., Rodrigues, L.: A machine learning approach to performance prediction of total order broadcast protocols. In: 4th IEEE SASO, pp. 184–193 (2010)
34. Didona, D., Quaglia, F., Romano, P., Torre, E.: Enhancing performance prediction robustness by combining analytical modeling and machine learning. In: Proceedings of the 6th ACM/SPEC ICPE, pp. 145–156 (2015)



Ji Xue received his B.S. degree in Computer Science from Beihang University (BUAA) in 2012. Currently he is a Ph.D. candidate of Computer Science at the College of William and Mary. His current research interests are system performance evaluation, cloud computing, machine learning, and data mining, especially in time series. He is a member of IEEE.



Feng Yan received his B.S. degree in Computer Science from Northeastern University in 2008 and M.S. degree in Computer Science from The College of William and Mary in 2011. Currently he is a Ph.D. candidate of Computer Science at the College of William and Mary. His current research interests are big data, cloud computing, resource management, scheduling, performance modeling and evaluation, and storage systems. He is a member of IEEE and ACM.



Alma Riska received her Ph.D. in Computer Science from the College of William and Mary, in Williamsburg, VA, in 2002. She was a Research Staff Member at Seagate Research in Pittsburgh, Pennsylvania and a Consultant Software Engineer at EMC in Cambridge, Massachusetts. She is currently Principal Software Engineer at NetApp in Waltham, Massachusetts. Her research interests are on performance and reliability modeling of computer systems, in general, and storage systems, in particular. The emphasis of her work is on applying analytic techniques and detailed workload characterization in designing more reliable and better performing storage systems that can adapt their operating into the dynamically changing operational environment. She is a member of IEEE and ACM.



Evgenia Smirni is the Sidney P. Chockley professor of computer science at the College of William and Mary, Williamsburg, VA. She holds a diploma degree in computer science and informatics from the University of Patras, Greece (1987) and a Ph.D. in computer science from Vanderbilt University (1995). Her research interests include queuing networks, stochastic modeling, resource allocation, storage systems, cloud computing, workload character-

ization, and modeling of distributed systems and applications. She is an ACM Distinguished scientist, and a member of the IEEE and the Technical Chamber of Greece.