

# FUNNEL: Assessing Software Changes in Web-Based Services

Shenglin Zhang, Ying Liu, *Member, IEEE*, Dan Pei, *Senior Member, IEEE*, Yu Chen, Xianping Qu, Shimin Tao, Zhi Zang, Xiaowei Jing, and Mei Feng

**Abstract**—The detection of performance changes in software change roll-outs in Internet-based services is crucial for an operations team, because it allows timely roll-back of a software change when performance degrades unexpectedly. However, it is infeasible to manually investigate millions of performance measurements of many roll-outs. In this paper, we present an automated tool, FUNNEL, for rapid and robust impact assessment of software changes in large Internet-based services. FUNNEL automatically collects the related performance measurements for each software change. To detect significant performance behavior changes, FUNNEL adopts singular spectrum transform (SST) algorithm as the core algorithm, uses various techniques to improve its robustness and reduce its computational cost, and applies a difference-in-difference (DiD) method to differentiate the true causality from the random correlations between the performance change and the software change. Evaluation through historical data in real-world services shows that FUNNEL achieves accuracy of more than 99.7 percent. Compared with previous methods, FUNNEL’s detection delay is 38.02 to 64.99 percent shorter, and its computation speed is 4.59-7,098 times faster. In real deployment, FUNNEL achieves a 98.21 percent precision, high robustness, fast detection speed, and shows its capability in detecting unexpected behavior changes.

**Index Terms**—Software change, performance change, singular spectrum transform, difference in difference

## 1 INTRODUCTION

IN large web-based services such as search engine, online shopping, and social networking, the operations team needs to frequently conduct software changes, i.e., software upgrades and configuration changes, in order to deploy new features, fix bugs, and improve service performance. Although each software change is extensively tested on test-beds before deployment, errors and bugs may still occur in the operational environment because of diverse hardware/software systems, complex interactions, and the large scale of devices [1], [2]. Therefore, the operations team typically deploys software changes using a “Dark Launching” [3] approach. Instead of rolling out a software change to all servers at one time, the operations team deploys the software change on a subset of servers at the beginning and continuously monitors a predefined list of **Key Performance Indicators** (KPIs) to determine the impact of the software change. The KPIs cover a wide range of performances, including user-perceived issues (e.g., Web page response delay),

service performance (e.g., advertisement click count), hardware health (e.g., server memory utilization), and so on.

If the KPIs on the server subset perform as expected, the software change will be rolled out to all servers. Otherwise, the software change should be rolled back as soon as possible. Generally, service performance degradation may induce poor user experience [1], [2] or revenue drop [4], [5].

Thus it is important to detect significant KPI changes rapidly, whether positive or negative, to allow a timely roll-back. However, the operations team generally assesses the impact of software changes manually in Web-based services, which has been demonstrated to be error-prone, cumbersome, and almost impossible to scale to a larger size. As a result, some critical issues caused by software changes may fly under the operations team’s radar, hurting the applications performance and user experience [1], [2].

In this study, we focus on software changes and their impact on KPIs. Our objective is to build an *automated* tool that detects behavior changes *rapidly* and *accurately* in a broad range of KPIs after a software change, and *accurately* determines whether the behavior changes are caused by the software change. Usually, KPI changes caused by software changes include level shifts, e.g., a sudden increase in memory utilization, or ramp-ups/ramp-downs, e.g., a deteriorating condition. Similar to [6], we focus on level shifts and ramp up/downs induced by software changes in this paper. How to reduce the detection delay, i.e., the delay between the occurrence of a KPI change and its detection is a real challenge for Web-based services for both *scalability* and *robustness* reasons. First, the impact of a software change can be observed in any of the huge number of KPIs of the services and servers that share the same spatial scope with the software change. In our studied scenarios, there are hundreds, even thousands, of KPIs that should be monitored after each

- S. Zhang and Y. Liu are with the Institute for Network Sciences and Cyberspace, and the Tsinghua National Laboratory for Information Science and Technology, Tsinghua University, Beijing 100084, China. E-mail: slzhangsd@gmail.com, liuying@cernet.edu.cn.
- D. Pei is with the Department of Computer Science and Technology, and the Tsinghua National Laboratory for Information Science and Technology, Tsinghua University, Beijing 100084, China. E-mail: peidan@tsinghua.edu.cn.
- Y. Chen, X. Qu, S. Tao, and Z. Zang are with the Baidu, Inc., Beijing 100085, China. E-mail: chenyu034@hotmail.com, {quxianping, taoshimin, zangzhi}@baidu.com.
- X. Jing and M. Feng are with the PetroChina, Inc., Beijing 100007, China. E-mail: {jxw, fm}@petrochina.com.cn.

Manuscript received 3 Nov. 2015; revised 26 Jan. 2016; accepted 3 Mar. 2016. Date of publication 9 Mar. 2016; date of current version 2 Feb. 2018. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TSC.2016.2539945

single software change that has occurred, while there are tens of thousands of software changes occurring every day in our scenario. This forces the operations team to monitor several million KPIs every day. The large scale of the problem calls for an algorithm with low computation overhead. Second, with the requirement of short detection delay, we do not have the luxury of using data smoothing and aggregation to achieve robustness any more, and the behavior change detection method should be quite robust [7].

To the best of our knowledge, in the literature there exists no rapid and robust method for impact assessment of software changes in large Web-based services. There exists several studies of impact assessment in the area of network infrastructure [6], [7]. However, the CUmulative SUM (CUSUM) used in [6] suffers from long detection delay [7], because the cumulative sum may take a long time before it exceeds the threshold. Multiscale Robust Local Subspace (MRLS) method applied in [7] achieves low detection delay, but the iteration of Singular Value Decomposition (SVD) used in subspace computation with  $l_1$ -norm exhibits high computational complexity [8]. While it works in [7] for thousands of time-series metrics in backbone networks, MRLS would spend too much computational resources for millions of KPIs in Web-based services, and is hence not feasible in our scenario. The iteration of SVD is essential to MRLS for improving robustness, and it is hardly possible to reduce the computation overhead of MRLS.

Singular Spectrum Transform (SST) [9] has emerged as a popular performance change detection method recently. SST has been demonstrated to be accurate with short detection delay [10], [11]. However, based on SVD, SST still suffers from high computational cost [12], and its accuracy degrades fast in the face of noises [13]. KPIs in Web-based services are quite diverse intrinsically, exhibiting various characteristics including strong seasonality (e.g., Web page view count), high variability (e.g., server CPU context switch count), and stationarity (e.g., server memory utilization). In addition, the baseline used to compare the performance after software changes may be contaminated by the impact of previous software changes and/or other factors, as is the case in large infrastructure networks [7].

In this paper, we designed and implemented FUNNEL, an automated tool for assessing the impact of software changes rapidly and robustly in large Web-based services. For a given software change, FUNNEL automatically determines the correct spatial scope of the impact, collects a broad range of KPIs, detects KPI changes, and determines the KPI changes caused by software changes. FUNNEL adopts matrix compression and implicit inner product calculation to reduce SST's computation overhead. The short detection delay and the low computational cost of the improved SST make timely mitigation possible when FUNNEL is deployed online. Furthermore, FUNNEL improves the robustness for SST, making it work quite well across diverse types of KPIs, especially variable KPIs.

In addition to software changes, other factors including seasonality, network hardware breakdowns, malicious attacks, etc., can also give rise to KPI changes. The impact of other factors can over-shadow the assessment of software changes [14]. To achieve the accurate inference of the impact of software changes, it is non-trivial for the operations team

to exclude KPI changes caused by other factors. However, neither CUSUM and MRLS, nor the improved SST can exclude the KPI changes induced by other factors. FUNNEL uses a classic method, difference in difference (DiD) [15], [16], [17], to solve the problem by comparing the relative performance between the treated group and the control group. DiD compares the KPIs of service processes and servers in which the software changes have been conducted (hereafter, collectively referred to as *t*servers/*t*instances, where “t” stands for “treated”) with the KPIs of service processes and servers of the same service (a server is usually dedicated to a specific service in our context) without the software change (hereafter, collectively referred to as *c*servers/*c*instances, where “c” stands for “control”). If the operations team conducts the software change without using the Dark Launching method, and there is no server/process in *c*servers/*c*instances, FUNNEL compares the impacted KPIs related to the software change with historical measurements of KPIs to exclude seasonality. FUNNEL adopts a relatively long baseline to address baseline contamination (in our prototype implementation, KPIs of 30 days before the day of software change are used to construct the baseline) [7]. Moreover, the large number of KPIs of *c*servers/*c*instances also alleviates the impact of baseline contamination.

Our main contributions can be summarized as follows:

- 1) We identify the problem of rapid, robust, and automated impact assessment of software changes in large Web-based services, and its research challenges in terms of scalability, robustness, detection delay, and computation cost.
- 2) We propose FUNNEL, the first approach in the literature that addresses the above challenges. The core idea of FUNNEL is to use SST, a rapid change point detection algorithm, as our algorithm basis. To reduce SST's computation overhead, FUNNEL adopts matrix compression and implicit inner product calculation. To improve the robustness for SST against noises, FUNNEL uses a DiD method to compare with various control groups.
- 3) Evaluation through historical data shows that FUNNEL performs significantly better than CUSUM [6] and MRLS [7], and the operational experiences of real deployment show FUNNEL's good performance and value. We conducted extensive evaluations using manually labeled KPIs collected from 144 different software changes in real-world Web-based services. Specifically, FUNNEL achieved an accuracy of more than 99.7 percent, a detection delay that is 58.82 percent of that of MRLS and 33.76 percent of that of CUSUM. In addition, FUNNEL is more than 7,000 times faster than MRLS in computational speed, and over four times faster than CUSUM. Furthermore, we have deployed the prototype of FUNNEL to dozens of real-world Web-based services. FUNNEL achieved a 98.21 percent precision based on one week's observation. A few representative cases were presented to show FUNNEL's robustness, detection speed, and capability to detect unexpected KPI changes. In one specific case, FUNNEL reduced the detection delay of one important

incident to less than 10 minutes, far less than the 1.5 hour used in manual assessment.

The rest of the paper is organized as follows. We provide an introduction to software change and KPI in Section 2, and describe the design of FUNNEL in Section 3. The evaluation of FUNNEL is presented in Section 4, followed by a description of the deployment of FUNNEL and case studies in Section 5. Finally, we review related works in Section 6 and conclude our paper in Section 7.

## 2 SOFTWARE CHANGES AND KPIS

In this section, we provide a brief introduction of software changes and KPIS.

### 2.1 Scope of Studied Software Changes

In this paper, we focus on two types of software changes on servers in large web-based services, *software upgrades* and *configuration changes*, for the following three reasons. (1) The operations team typically care about the unexpected consequences that are potentially due to these planned changes; (2) These changes are *controllable* by the operations team via command line interfaces and *observable* in logs; (3) We have observed that these two types constitute more than 90 percent of the tens of thousands of software changes in our data.

*Software upgrades.* With the current rapid evolution of the Internet, new features are continuously being deployed with software upgrades. The operations team also conducts software upgrades to fix bugs or improve service performance. In a large service, it is often the case that one software upgrade implements multiple features or bug fixes, and FUNNEL considers such a software upgrade as a whole. FUNNEL decides whether the whole software upgrade introduces any KPI change but does not attempt to distinguish which individual feature or bug fix introduces KPI changes. In addition, the web services in our scenario usually trigger a software upgrade when operators try to patch the software, and thus we classify the patches as software upgrades.

*Configuration changes.* Using command line interfaces, the operations team can change the configurations by using specific commands. The configuration change can be in the operating system (OS) or infrastructure software (e.g., a configuration change in Apache), service configuration (e.g., an increase in the number of threads in a service process *by command lines*), deployment scale (e.g., an increase in the number of servers where a service is deployed), or data source (e.g., an update to the strategy that calculates the valid page view counts).

With above focuses, the following perspectives are out of scope for this paper.

- 1) We do not consider the software changes on the network devices such as routers and switches, which have been already studied in depth in [6], [7], [14];
- 2) We do not consider software changes that were external to the company, e.g., a change in a peer company, since these changes might be invisible to the studied company's operations team.
- 3) We do not explicitly consider the interactions across multiple concurrent or consecutive software changes on a same server, which can be considered as one combined change as a straw man approach.

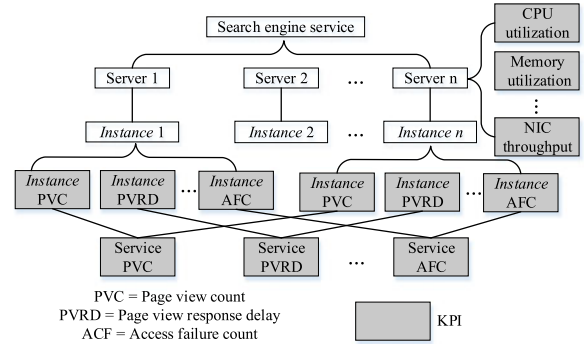


Fig. 1. The relationship among service, server, instance and KPI.

More detailed studies along the last three directions are left as future work.

### 2.2 KPI

In the studied Web-based services, there are hundreds of thousands of servers providing various types of services. Each service (e.g., search, web mail, social networking) runs on one or more servers with a specific process on each server. An *instance* denotes a *process* of a specific service on a specific server. A KPI is a performance metric of a *given server/service/process*. There are three types of KPIS that need to be monitored for software changes assessment: server KPIS, *instance* KPIS and service KPIS. The operations team deploys an agent on each server to monitor the status of each *instance* and collect the KPIS of all *instances* continuously. For example, immediately after the process serves a customer with some web page view, the page view count is incremented and a new page view response delay is recorded. In addition, by analyzing server log files that record the system status, the agent is able to periodically collect server KPIS, such as CPU utilization, memory utilization, and NIC throughput. A service KPI is an aggregation of all *instance* KPIS in the service. Fig. 1 shows an example of the relationship among service (search engine service), server (server 1, server 2, ..., server  $n$ ), *instance* (instance 1, instance 2, ..., instance  $n$ ) and KPI (page view count, page view response delay, access failure count, CPU utilization, memory utilization, and NIC throughput).

After collecting the measurements of KPIS of servers and *instances*, the agent on each server delivers the measurements *via datacenter networks* to a centralized Hadoop-based database, which also stores the service KPIS aggregated based on the KPIS of the *instances*. The database also provides a subscription tool for other systems, such as FUNNEL, to periodically receive the subscribed measurements based on the server, *instance*, and service. The data collection interval at the servers is typically 1 minute, and thus the time granularity of the input KPI measurements of FUNNEL is 1 minute. If we set the time granularity of the KPI measurements larger, the detection delay of FUNNEL will be larger. Although the accuracy of FUNNEL may be impacted by the time granularity of the measurements [18], the accuracy of FUNNEL is relatively good (see Section 4.2) with the time granularity of 1 minute, we set the time granularity of the KPI measurements as 1 minute. Within one second, the measurements subscribed by FUNNEL are pushed to FUNNEL *via datacenter networks*.

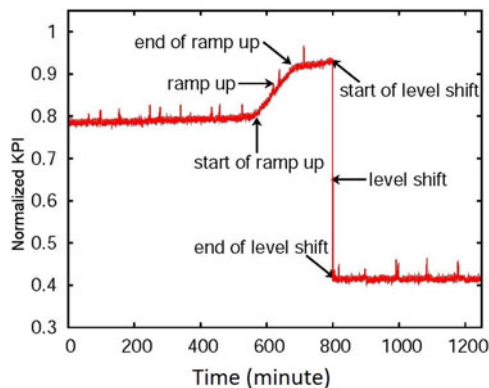


Fig. 2. Examples of level shift and ramp up/down.

In large services, there might exist some KPIs of dubious quality. To the best of our knowledge, there is no previous work on eliminating low-quality KPIs in web-based services. In this paper, we do not focus on eliminating low-quality KPIs either. FUNNEL detects all KPI changes in the impact set regardless of the quality of the KPI, and delivers the results to the operations team. The operations team will then determine whether the performance changes in the low-quality KPIs are induced by the software change or not.

As described in [19], even the operations team does not know exactly what is a “good” threshold for a specific KPI, e.g., due to seasonal variation. Therefore, it is difficult to obtain the expected values from the operations team for FUNNEL, and we do not give a threshold for any KPI in this paper.

### 2.3 KPI Changes

A KPI change is defined as a non-transient change (e.g., lasting more than 7 minutes<sup>1</sup>) in a KPI that is introduced by a software change. In this paper, we focus on behavior changes evidenced by individual KPI time series. As Fig. 2 shows, the changes can be either level-shifts immediately after software changes, or ramp ups or downs that ensue gradually over time after the software changes.

Behavior changes in KPIs can validate the expected impacts, e.g., a decrease in CPU utilization after a configuration change aimed to increase efficiency, or show the unexpected impact occurrences, e.g., a sudden increase in page view response delay after a software upgrade.

The assessment of any *single* software change should be in a low-computational way, for the following two reasons. (1) Hundreds to thousands of KPIs should be detected after a single software change; (2) The operations team should assess tens of thousands of *software changes* every day.

## 3 FUNNEL DESIGN

Recall that our aim is to provide an automatic tool for rapid and robust impact assessment of software changes. To achieve this goal, we designed FUNNEL, which is composed of two main components, as shown in Fig. 3.

1. The threshold for transient changes is subjective in reality. It may be related to the tolerance level or the sensitivity of the operations team, or the nature of the web-based services. In this paper, choosing 7 minutes as the threshold for non-transient KPI changes is based on the experience of the operations team.

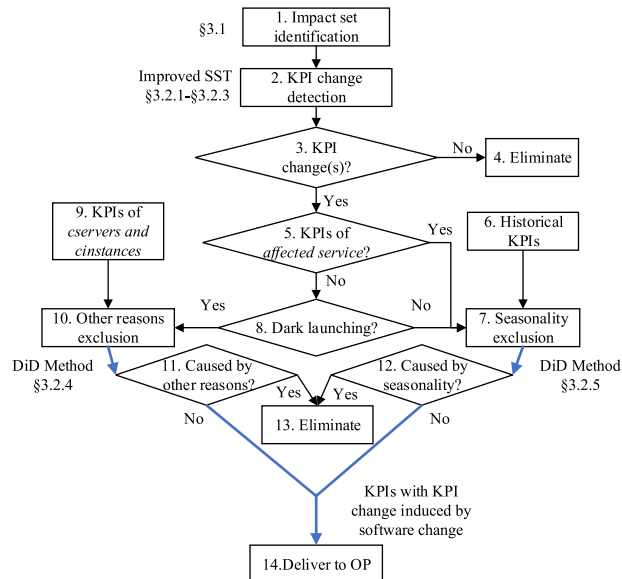


Fig. 3. FUNNEL design.

**Impact set identification.** In the studied search engine company, the operations team names the services based on the service hierarchy. We believe this practice is not uncommon in other companies. FUNNEL derives the relationship among services using the naming rules. Inspired by [7], FUNNEL automatically identifies the impact set, that is, the set of servers, instances, and services that may be impacted (Section 3.1) based on the change deployment logs and the relationships among services.

**Performance change detection and determination.** FUNNEL uses an SST based performance change detection method. It improves the robustness of SST (Section 3.2.2) and introduces matrix compression and implicit inner product calculation to solve the high computational cost of SST (Section 3.2.3). The improved SST can detect KPI changes rapidly and robustly (step 2 in Fig. 3). FUNNEL then determines whether the changes are caused by software changes. If the KPI is not the KPI of *affected services* (services that are related to the service where the software change is deployed) (step 5 in Fig. 3) and if the operations team rolls out the software change using Dark Launching (step 8 in Fig. 3), FUNNEL excludes the impact of other factors (step 10 in Fig. 3) by applying a DiD method (Section 3.2.4) using KPIs of *cservers* and *cinstances* (step 9 in Fig. 3). Otherwise, based on the DiD method, FUNNEL uses historical measurements of KPIs (step 6 in Fig. 3) to exclude the impact of seasonality (Section 3.2.5), i.e., the time of day and the day of week effects (step 7 in Fig. 3).

### 3.1 Impact Set Identification

The effect scopes of different types of software changes are different: some are local, e.g., a configuration change aimed at balancing traffic only influences the performance of the servers where the change is conducted, while others are global, e.g., an upgrade in an advertising system can have an impact on almost all types of web-based services. A false impact scope of a software change may give rise to delayed detection after a performance degradation or a large number of false alarms.

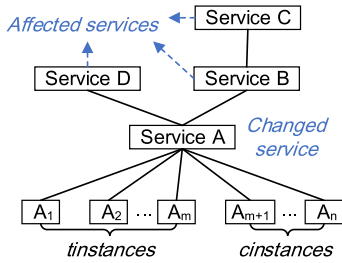


Fig. 4. Example of service relationship.

It is straightforward to identify the impact scope of a software change on servers because only the performance of the servers where the software change is conducted can be directly affected. Therefore, the impact set of servers consists of *tserver*s, i.e., the servers on which the software change is deployed. The set of *tserver*s is directly obtained from the software change logs.

As Fig. 4 shows, for a software change deployed on Service A, suppose that Service A has multiple *instances* ( $A_1, A_2, \dots, A_n$ ), Service A is related to Service B and Service D (i.e., Service A send requests and responses to Service B and Service D, and the relationships among services are available to the operations team), and Service B is related to Service C. We refer to service A as the *changed service*, and service B, service C and service D as the *affected services*. For a specific software change deployed on ( $A_1, A_2, \dots, A_m$ ), ( $A_1, A_2, \dots, A_m$ ) are the *tinstances*, and  $A_{m+1}, \dots, A_n$  are the *cinstances*, which are the service A's *instances* running on the servers on which the change is not deployed yet. The impact set consists of the *tinstances*, the *changed service*, and the *affected services*. We do not include any *instances* of the *affected services* in the impact set because it is unlikely that any *instance* KPI of the *affected services* is individually affected for load balancing reason. Instead, they are more likely to be affected by the same impact. Thus, including their aggregation, i.e., the *affected service* KPI in the impact set, is sufficient for studying the impact of software changes on the *affected services*.

Except for the *affected services*, the elements in the impact set all belong to one service. The operations team usually do not deploy two software changes in one service at the same time based on common practice (to avoid complications). Therefore, except for the *affected services* all the elements in the impact set of a specific software change are not likely to be impacted by other software changes. For the *affected services* in the impact set, if they are the *affected services* or the *changed services* of other software changes, the operations team can manually determine the cause of the behavior changes in the *affected services* when the results are delivered by FUNNEL to them.

In addition, it is possible that two services share the same network infrastructure, such as TOR, aggregation switch, access router. However, based on intuition and the operations team's experience, even if some two services share the same network infrastructure, the KPIs of one service is little impacted by the other service. Therefore, we do not consider the interference of services that share the same network infrastructure.

FUNNEL investigates *all* the KPIs in the impact set automatically. A time-series is constructed for each KPI by

dividing the original event series into equal time-bins. One min is used as the time-bin in FUNNEL since our goal is to achieve a rapid impact assessment of software changes.

### 3.2 Change Detection and Determination

For a given time series of KPIs, our mission is to determine whether behavior changes (level shifts, ramp up/downs) exist and whether they are caused by software changes. Our objective is to provide a rapid and robust KPI change detection and determination tool, which goes beyond traditional change detection methods that compare means, medians, distributions, etc., before and after software changes. For example, the CUSUM approach used in [6] suffers from low accuracy in the face of KPIs with strong seasonality. In addition, its long detection delay also makes the CUSUM method unsuitable for our scenario. What's more, when applying a simple moving average method to detect the KPI changes after a software change, we find that if we set the threshold to be small, there will be too many false positives, while if we set the threshold to be large, the detection delay will be too large, and it is hardly possible that we can find a balance.

SST, which is based on SVD, has been shown to be accurate and rapid in performance change detection in other research fields [10], [11]. SST projects training data into a normal subspace and finds the difference between the normal subspace and the data that needs to be tested. However, SST suffers from low accuracy in the case of normal subspace contamination. Since the training data may contain outliers as a result of previous software changes, network device breakdowns, network attacks, etc., baseline contamination often occurs in practice. In addition, due to the complexity of SVD, SST suffers from high computational cost.

Our core idea is to improve the robustness and reduce the computational cost of SST, and then combine the improved SST with DiD to determine the impact of software changes. We first provide a brief description of SST's detection of performance changes.

#### 3.2.1 SST: Singular Spectrum Transform

SST [11] is based on SVD of the Hankel matrix [20], and its main idea is to find the difference before and after a point  $x(i)$  at time  $t$ . Specifically, the algorithm compares a representation of the dynamics of a few points before  $x(i)$  and a few points after  $x(i)$ . The difference is normalized by  $x_s(i)$ .

The dynamics of the points before  $x(i)$  can be denoted by a Hankel matrix:

$$B(t) = [q(t - \delta), \dots, q(t - 1)], \quad (1)$$

where  $q(t) = [x(t - w + 1), \dots, x(t)]^T$ , and  $w$  and  $\delta$  are the size and the number of the overlapping windows respectively.

To find the singular values and vectors of the Hankel Matrix, SVD is used:

$$B(t) = U(t)S(t)V(t)^T, \quad (2)$$

where  $S(t)$  is the singular value matrix and  $S(i - 1, i - 1) \leq S(i, i) \leq S(i + 1, i + 1)$ , and  $U(t)$  and  $V(t)^T$  are unitary matrices. The columns of  $U(t)$  are the eigenvectors of  $B(t)B(t)^T$ . The first  $\eta$  eigenvectors of  $U(t)$  ( $U_\eta(t)$ ) are used to denote the past change pattern.

Then, for the future points of  $x(i)$ , a similar procedure is used to find the largest change in the dynamics by concatenating  $\gamma$  overlapping windows of size  $\omega$ , starting at  $\rho$  points after  $x(i)$ :

$$A(t) = [r(t + \rho), \dots, r(t + \rho + \gamma - 1)], \quad (3)$$

where  $r(t + \rho) = x(t + \rho), \dots, x(t + \rho + \omega - 1)^T$ .

The eigenvector  $\beta(t)$ , which represents the direction of the maximum change in the future of the time-series, is:

$$A(t)A(t)^T u^\rho = \mu u^\rho, \quad (4)$$

$$\beta(t) = u_\gamma^\rho, \quad (5)$$

where  $\gamma = \arg \min_i(\mu_i)$ .

Then, the projection of  $\beta(t)$  onto  $U_\eta$  is used to quantify the discordance between  $\beta(t)$  and  $U_\eta$ :

$$\alpha(t) = \frac{U_\eta(t)^T \beta(t)}{\|U_\eta(t)^T \beta(t)\|}. \quad (6)$$

The change score is calculated as the cosine of the angle between  $\alpha(t)$  and  $\beta(t)$  as:

$$x_s(t) = 1 - \alpha(t)^T \beta(t). \quad (7)$$

The intuition behind the algorithm is that  $\beta(t)$  is in or very near to the direction of the maximum change in the past denoted by  $U_\eta$  when there is no change in the time-series.

SST suffers from three problems as follows.

- 1) Five different parameters ( $\rho, \gamma, \eta, \delta, \omega$ ) must be specified in SST. Domain knowledge can help find proper values for  $\delta$  and  $\omega$ , but fails in choosing the other parameters;
- 2) SST degrades fast in terms of accuracy when the input time-series includes significant noises [13];
- 3) SST is based on SVD which suffers from high computational cost, and thus SST is not computationally efficient and not suitable for KPI change detection when the number of KPIs is large [12].

### 3.2.2 Improving the Robustness of SST

In this section, we describe our approach to addressing the first two aforementioned problems. Our approach is motivated by the study reported in [13].

As suggested in [12], [13], we set  $\rho = 0, \gamma = \delta$ , and since a value of 3 or 4 is suitable for  $\eta$  even when  $\omega$  is on the order of 100 empirically, we set  $\eta = 3$ .

To estimate the change score around every point and alleviate the impact of noises in the time series, we attempt to use more information from the future matrix  $A(t)$  than SST does by utilizing the  $\eta$  eigenvectors of  $A(t)A(t)^T$  with the smallest corresponding eigenvalues ( $\lambda_{1:\eta}$ ), rather than using only the first one.

Based on Eq. (4), the eigenvector  $\beta_i(t)$  is calculated as:

$$\beta_i(t) = u_i^g, \quad (8)$$

where  $i \leq \eta$ , and  $\lambda_{j-1} \leq \lambda_j \leq \lambda_{j+1}$  for  $1 \leq j \leq \omega$ .

The change score of the point  $x(i)$  at time  $t$  is calculated:

$$\hat{x}(t) = \frac{\sum_{i=1}^{\eta} \lambda_i \times \varphi_i(t)}{\sum_{i=1}^{\eta} \lambda_i}, \quad (9)$$

where

$$\varphi_i(t) = 1 - \sum_{j=1}^{\eta} (\beta_i(t) u_j^T)^2, \quad (10)$$

where  $u_j$  is the column of  $U_\eta$ . As described in [12], Eq. (10) is consistent with Eq. (7).

To alleviate the effect of noise like spikes induced by *bustly workloads* [21] on the final change score, we then filter the sections where the median and the median absolute deviation (MAD) of  $\hat{x}(t)$  remain nearly constant. The intuition behind the filtering step is that the accuracy of SST is reduced mainly when the noise takes over the original signal in the time-series. Since the mean and standard deviation for Gaussian distribution are not very robust in the presence of large changes or outliers, we use the median and MAD rather than the mean and standard deviation [13]. The combination of median and MAD has been proved to be a more robust approach even when outliers occur [7]. The change score  $\hat{x}(t)$  is then updated:

$$\tilde{x}(t) = \hat{x}(t) \times \frac{|median_a(t) - median_b(t)|}{|\sqrt{MAD_a(t)} - \sqrt{MAD_b(t)}|}, \quad (11)$$

where  $median_a(t)$  and  $median_b(t)$  are the medians of  $x(t)$  in the time-series of length  $(2\omega - 1)$  before and after the point  $x(i)$  at time  $t$ , respectively, and

$$MAD_a(t) = median(x_{ia}(t) - median_a(t)), \quad (12)$$

where  $x_{ia}(t)$  is the time-series of length  $(2\omega - 1)$  before  $x(i)$  at time  $t$ . It is similar to  $MAD_b(t)$ .

### 3.2.3 Reducing the Computational Cost of SST

In addition to the robustness problem, SST also suffers from high computation cost because of the SVD procedure and thus its direct deployment in our scenario is not feasible. This section summarizes the work in [12]. First, we applied the Implicit Krylov Approximation (IKA) algorithm to reduce the computation cost for SST. The essence of the IKA algorithm is matrix compression and implicit inner product calculation, and the efficiency of the algorithm was demonstrated in [12].

First, let  $C = B(t)B(t)^T$ , and  $\eta < k < \omega, k \in N$ . Since it is empirically true that the change score is not very sensitive to  $\delta$  [12], we set  $\delta = \omega$  as the IKA algorithm requires, and then, we get  $\gamma = \delta = \omega$ . In addition, suppose that  $T_k$  is a  $k$ -dimensional tridiagonal matrix;  $a_1, \dots, a_k$  and  $b_1, \dots, b_{k-1}$  are the diagonal and subdiagonal elements of  $T_k$ . At each time  $t$ , we first compute  $\beta_i(t)$ , and then run Lanczos ( $C, \beta_i(t), k$ ) [22] to obtain  $T_k$  as Algorithm 1 describes. Using the QL iteration [23], the eigenvectors of the tridiagonal matrix  $T_k$  can be calculated extremely fast. Based on the  $\eta$  top eigenvectors  $x_1, \dots, x_\eta$  of  $T_k$ , we can obtain  $\varphi_i(t)$  as:

$$\varphi_i(t) \simeq 1 - \sum_{j=1}^{\eta} x_j^2. \quad (13)$$

Then the change score can be calculated by Eqs. (9) and (11).

---

**Algorithm 1.** Lanczos( $C, \beta_i(t), k$ )

---

**Input:**

$$C \in \mathbb{R}^{\omega \times \omega}$$

$$\beta_i(t) \in \mathbb{R}^{\omega}$$

$$\eta < k < \omega, k \in \mathbb{N}$$

1:  $r_0 = \beta_i(t), b_0 = 1, q_0 = 0, s = 0.$

2: **while**  $s < k$  **do**

3:  $q_{s+1} = r_s / b_s;$

4:  $s \rightarrow s + 1;$

5:  $a_s = q_s^T C q_s;$

6:  $r_s = C q_s - a_s q_s - b_{s-1} q_{s-1};$

7:  $b_s = \|r_s\|;$

8: **end while**

9: **return**  $\{a_1, \dots, a_k\}$  and  $\{b_1, \dots, b_{k-1}\}$

**Output:**

$$T_k$$


---

Based on [12], the dimension of the Krylov subspace  $k$  can be set as:

$$k = \begin{cases} 2\eta, \eta \in \text{even} \\ 2\eta - 1, \eta \in \text{odd}. \end{cases} \quad (14)$$

For a service that needs quick mitigation on false software changes,  $\omega$  can be set to a small value such as 5. For a service that needs more precise assessment of software changes,  $\omega$  can be set to a larger value such as 15.

### 3.2.4 Excluding Other Reasons for Non-Affected-Service KPIs in Dark Launching

In addition to software changes, the KPI changes can also be induced by other factors including seasonality, network attacks, etc. The impact of these factors can over-shadow the impact assessment of software changes.

To solve this problem, motivated by [14], [24], [25], we apply a split testing method to compare the relative performance of the treated group (KPIs of servers/services/instances in the impact set) and the control group (KPIs of cservers/cinstances). The intuition behind the comparison is that other factors except software changes influence both the treated group and the control group.

*Treated group and control group.* In Dark Launching, the operations team first deploys the software change on a subset of servers and instances, and then, rolls it out to all servers and instances that belong to the same service. It is straightforward to identify the control group and the treated group of servers based on the software change logs. The KPIs of *tserver*s constitute the treated group of servers, while the KPIs of *cservers*, i.e., servers that belong to the same service as *tserver*s but without software changes, constitute the control group of servers. Since the KPI of the *changed service* is an aggregation of the KPIs of the *tinstances*, determining the relative performance of the *tinstances* is sufficient: if no performance changes in *tinstances* are caused by software changes, it is not necessary to study the impact on the *changed service*. Therefore, the treated group of service

consists of the KPIs of *tinstances*, while the control group of services is constituted of KPIs of *cinstances*, i.e., instances that belong to the same service as *tinstances* but without software changes. This is based on the four following observations about web-based services.

- 1) *Instances* and servers that belong to the same service exhibit high-level spatial correlation or statistic dependency because of the similarity in traffic load, memory usage, etc., thanks to load balancing;
- 2) It is very likely that any non-software change factors will introduce similar performance impact on all servers and instances of the same service;
- 3) A performance change induced by the software change introduces a relative difference in performance between the treated group and the control group;
- 4) As studied in [26], [27], only a small fraction (less than 3 percent) of edge links are hotspots in datacenter networks. In other words, a small fraction of servers are hotspots. Even though there are KPIs of hotspot servers in the control group, most of the KPIs in the control group are not of hotspot servers. We use the *average* of all of the KPIs in the control group to eliminate performance changes caused by other factors, and the large number of KPIs in the control group can alleviate the impact of hotspots.

The operation team's common practice is not to deploy two software changes in a specific service at the same time. Furthermore, the KPIs of a specific control group all belong to a specific service. Therefore, KPIs in the control group are very unlikely to be impacted by some other software changes.

For each *affected service*, if it is influenced by software changes, all of its *instances* will be affected. Therefore, there are no *cinstances* for *affected services*. We exclude the "Full Launchin" and other reasons for *affected services* in Section 3.2.5.

*DiD method.* DiD is one of the most popular tools in econometrics [16] and health care [17] for evaluating the effects of interventions that are instituted at a particular point in time. In this study, DiD is used to compare changes over time in the treated group with those over time in the control group, and to attribute the difference-in-differences to the effects of software changes. DiD is based on the assumption that, in the absence of software changes, the difference between the average KPIs for the treated group and those for the control group remains stable over time. Although [14] argued that DiD suffered from low accuracy in impact assessment of changes in cellular networks, it turns out to perform quite well in our scenario because the treated group and the control group exactly follow the above assumption.

The basic DiD framework can be described as follows. Let  $Y(i, t)$  be the KPI  $i$  at time  $t$ . The performance is observed in a pre-software-change period  $t = 0$ , with length  $\omega$  as described in Section 3.2.1, and in a post-software-change period  $t = 1$  with length  $\omega$ . Some fraction of KPIs are exposed to the software change between these two periods. If KPI  $i$  has been exposed to the software change prior to period  $t$ , then  $D(i, t) = 1$ , and  $D(i, t) = 0$  otherwise. In other words, those KPIs with  $D(i, 1) = 1$  are in the treated group, and those KPIs with  $D(i, 1) = 0$  are in the control

group. Obviously,  $D(i, 0) = 0$  for all  $i$ , because the KPIs are exposed to the software change only after the first period.

To obtain the standard errors and significance levels for the DiD estimator, a linear parametric model is used [15]:

$$Y(i, t) = \theta(t) + \alpha \cdot D(i, t) + \xi(i) + v(i, t), \quad (15)$$

where  $\theta(t)$  is a time-specific parameter,  $\alpha$  denotes the impact of software changes, and  $\xi(i)$  is an KPI-specific parameter.  $v(i, t)$  is a transient shock at each period, i.e.,  $t = 0, 1$ , with mean zero. A sufficient condition for determination using DiD is  $P(D(i, 1) = 1 | v(i, t)) = P(D(i, 1) = 1)$ .

Then, the impact estimator of software change,  $\alpha$ , is:

$$\alpha = \{E[Y(i, 1) | D(i, 1) = 1] - E[Y(i, 1) | D(i, 1) = 0]\} - \{E[Y(i, 0) | D(i, 1) = 1] - E[Y(i, 0) | D(i, 1) = 0]\}, \quad (16)$$

where  $E(\cdot)$  denotes the expectation.

If the KPI changes are caused by factors excluding software changes, then there is no change in the relative performance between the treated group and the control group, thus the DiD impact estimator,  $\alpha$ , should be near zero. Therefore, if  $\alpha \approx 0$ , we consider that the performance changes are not induced by software changes. If  $\alpha \gg 0$  or  $\alpha \ll 0$ , then we consider that there is a relative increase or decrease in the treated group as compared to the control group, and the likelihood that the performance changes are caused by a software change is high.

Empirically, for a service which is sensitive to KPI change, such as advertisement, online shopping, the threshold of  $\alpha$  can be set to a small value like 0.5. Otherwise, the threshold can be set larger.

### 3.2.5 Excluding Other Reasons for Related Services and Full Launching Manner

For *affected services*, there are no *cservers* or *instances*. Moreover, if the operations team deploys software changes on all servers at one time (i.e., Full Launching), the control group is also empty. In addition to software changes, seasonality may also give rise to KPI changes [7], [14], [28]. Thus, we need to exclude KPI changes induced by seasonality. For a specific KPI of servers/services/instances in the impact set, FUNNEL compares the measurements of KPI around the software change and the KPI measurements in the same period of day but on historical days, since seasonality impacts the KPI measurements around the software change and the historical KPI measurements similarly.

*Treated group and control group.* For a given KPI of a server/service/instance in the impact set, the treated group consists of the measurements of KPIs around the software change, while the control group consists of the historical measurements of KPIs. Specifically, to exclude the performance changes due to the time-of-day or day-of-week pattern and exclude the influence of baseline contamination, we use the measurements of KPIs of 30 days before the day of the software change to construct the control group. This is based on the observation that there is almost no relative performance change between the control group and the treated group if the performance change is induced by seasonality.

FUNNEL also applies DiD method to compare the relative performance between the control group and the treated group. Specifically,  $t = 0$  for a pre-software-change period of length  $\omega$  as described in Section 3.2.1 in the treated group, and for the same period of day but on historical days in the control group. Similarly,  $t = 1$  denotes a post-software-change period, of length  $\omega$ , in the treated group, and the same period of day but on historical days in the control group.

## 4 FUNNEL EVALUATION

In this section, we evaluate FUNNEL's performance. We compare FUNNEL with other software change assessment methods including CUSUM [6] and MRLS [7] that have previously been deployed for upgrade assessment in the network infrastructure. We implemented FUNNEL, CUSUM, and MRLS with C++. We use the software change and KPI data from a few real-world Web-based services offered by a top global search engine. The evaluation using real-world data is challenging because of the lack of a ground truth [14]. We used manual assessment results by the operations team as our ground truth for evaluation.

The focus of FUNNEL is to detect KPI behavior changes in a *rapid* and *robust* way, which is achieved by the improved SST. The introduction of DiD method helps to determine whether the KPI changes are caused by software changes, by comparing the performance between the control/treated group. We decided not to compare FUNNEL with Litmus because the robust regression approach in Litmus [14] is mainly used for inferring the relative performance change between the treated and the control group, and the comparison adds no additional value.

Our results in this section show FUNNEL performs significantly better than CUSUM and MRLS in accuracy, detection delay, and computational cost.

### 4.1 Data Sets

In cooperation with the operations team, we randomly picked 19 moderate-sized services (the manual assessment efforts for large services would be prohibitive) over a 2-day period, then we got 6,277 software changes in total. We ran the algorithm in Section 3.1 to identify the impact set, and collected the KPI measurements of *tservers/tinstances/changed services/affected services* and the KPI measurements of *cservers/cinstances*. Note that running this algorithm is equally beneficial to FUNNEL, CUSUM and MRLS, and is not biased towards FUNNEL. The details of manual inspection are as follows. For a given software change, first we aggregated the KPIs of *tinstances/tservers* by calculating the average measurements of *tinstances/tservers*. The operations team then manually inspected whether behavior changes occurred in the KPIs of the *changed service*, the *affected services*, and the *aggregation KPIs* of the *tinstances* and the *tservers* around the time of the software change. The operations team found that in total 83 software changes had behavior changes in the KPIs of *changed service/affected service* or *aggregation KPIs* of *tservers/tinstances* shortly after each of these software changes. For each behavior change, the operations team then manually inspected the control group and determined whether it is actually caused by software change. Eventually, the operations team selected 72 out of



6,277 software changes that induced KPI changes. We admit that behavior changes could occur in one or more *individual* KPIs of the *tserver*/*tinstance*s in the 6,194 (6,277 - 83) software changes. However, manually inspecting all the KPIs of the 6,194 software changes, i.e., about 600 thousands KPIs, is a huge amount of work. Alternatively, we *randomly* selected 72 out of the remaining 6,194 software changes. For all selected 72 software changes, the operations team manually carried out the investigation, and found that there were no KPI changes induced by the software changes. We use the 72 software changes that induced changes in KPIs and the 72 selected software changes that did not for the evaluation, and the manual assessment results of the 144 (72 + 72) software changes served as the ground truth.

The data used in our evaluation exhibited different characteristics including seasonality, variability, stationarity, and different levels of baseline contaminations, which provided a relatively exhaustive validation.

Generally, the CPU context switch count of servers varies frequently, while the memory utilization remains stationary. In addition, both the CPU context switch count and the memory utilization indicate the health status of servers. Specifically, the CPU context switch count indicates the computational efficiency and the number of threads after software changes, while the memory utilization indicates whether a software change introduces memory leaking. Thus we used the CPU context switch count and the memory utilization as the KPIs of all the servers in the evaluation. The KPIs of a given service/*instance* are defined by the operations team and they differ from one service/*instance* to another service/*instance*.

We collectively refer to the combination  $(S_i, c_i, k_i)$  as an *item*, where  $S_i$  denotes a software change,  $c_i$  is a server, service, or *instance* in the impact set, and  $k_i$  is a KPI of  $c_i$ .

A total of 9,982 *items* were included in the evaluation, to which 144 software changes (described earlier in this section), and 931 servers were related. More specifically, the  $k_i$  of 931, 931, 8,120 *items* were CPU context switch count, memory utilization, KPIs of services/*instances*, respectively. In addition, based on the operations team's assessment, there were 968 *items* that had been labelled as having performance changes introduced by software changes in the 72 software changes that induced behavior changes.

To show that the software changes we selected were representative, we collected all software changes *on all services* in the studied company in the same two days period, i.e., 23,873 software changes, and analyzed whether the distribution of the *item* number of the 23,873 software changes were consistent with that of the *item* number of the studied 144 software changes used in the evaluation. Similarly, we apply the CPU context switch count and the memory utilization as the KPI of each server, and the KPI of each *instance* is predefined by the operations team. Fig. 5 shows the cumulative distribution functions (CDFs) of the number of *items* of each software change in the 23,873 software changes and each of the 144 software changes used in the evaluation. We also present the CDFs of the number of servers/*instances* of the two sets of software changes. We observe that the set of the software changes for the evaluation is consistent with the overall data set for the two day period.

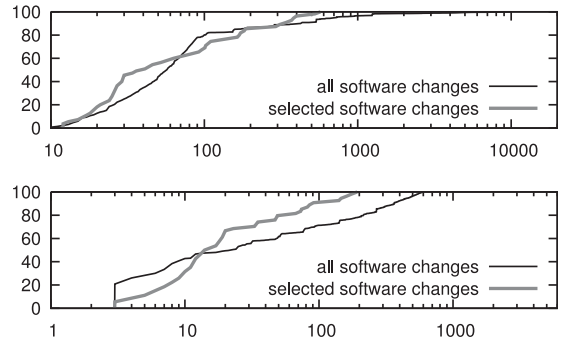


Fig. 5. CDFs of the number of servers/*instances* and *items* of each software change.

Based on empirical experience, if a software change in an web-based service has a negative impact, the KPIs usually change shortly after the software change. The operations team think that 1 hour is enough for software change assessment. Among the 144 software changes, 108 were deployed with Dark Launching, while the remaining 26 were not. If a software change was conducted with Dark Launching, and the KPIs were not of *affected services*, we will construct the treated group using the KPIs of *tserver*/*tinstance*s 1 h before and after the software change, and construct the control group using the KPIs of *cserver*/*cinstance*s in the same period. Otherwise, the treated group will consist of the measurements of KPIs 1 h before and after the software change, and the control group will consist of the measurements of KPIs in the same period but on historical days (30 days).

We constructed a time-series,  $x(1), x(2), \dots, x(n)$ , for each *item* by dividing the original measurements of KPIs into equal time-bins of 1 min. Each method took a *time window* of  $x(i), x(i+1), \dots, x(i+W)$  as its input to construct a matrix (FUNNEL, MRLS) or calculate a cumulative sum (CUSUM). For a fair comparison of the accuracy among FUNNEL, CUSUM and MRLS, the length of the sliding input *time window*  $W$  for each method was set as the one that achieved the best accuracy ( $W_{\text{FUNNEL}} = 34$  (i.e.,  $\omega$  in Section 3.2.3 is set 9),  $W_{\text{MRLS}} = 32$ ,  $W_{\text{CUSUM}} = 60$  in our scenario). The time window moves forward every minute. For example, FUNNEL first detects and determines performance changes for  $x(1), x(2), \dots, x(34)$ , and then for  $x(2), x(3), \dots, x(35)$ , etc. Note that the values of other parameters of CUSUM, MRLS and FUNNEL ( $\alpha$ ) are also set to the best for the corresponding algorithm's accuracy. We believe the above method draws the same conclusion as the method that changing the value of the parameters, calculating the accuracies and plotting the receiver operating characteristic (ROC) curves.

Empirically we set a threshold of 7 minutes in FUNNEL to declare a change in a time series as a level-shift or ramp-up/down rather than a one-off event.

## 4.2 Comparison of Accuracy

We now compare the accuracy of FUNNEL, the improved SST without DiD, CUSUM, and MRLS. For each *item*, based on the ground truth provided by the operations team, we knew the outcome - either having the KPI changes induced by software changes or not. For each method, we labelled its outcome as true positive (TP), true negative (TN), false positive (FP), and false negative (FN). True positives were

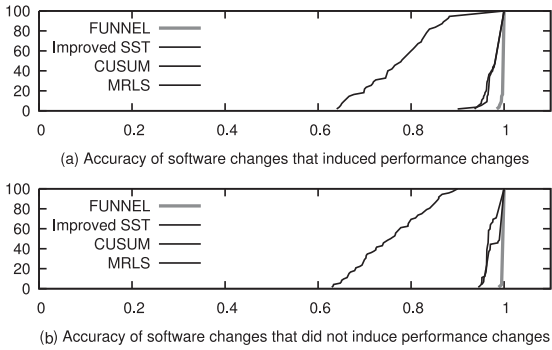


Fig. 6. CDFs of the precision, recall, TNR and accuracy of each software change.

*items* with KPI changes caused by software changes that were accurately determined as such by the method, and true negatives were *items* that were accurately determined as having no KPI changes induced by software changes. If the method determined a KPI change caused by a software change while there was no KPI change or the KPI change was not induced by a software change, we then labelled the *item* as a false positive. False negatives were KPI changes induced by a software change that were incorrectly missed by the method. We calculated Precision, Recall, true negative rate (TNR), and Accuracy as: Precision =  $\frac{TP}{TP+FP}$ , Recall =  $\frac{TP}{TP+FN}$ , TNR =  $\frac{TN}{TN+FP}$ , Accuracy =  $\frac{TP+TN}{TP+TN+FP+FN}$  [14].

#### 4.2.1 Comparison Results

For each software change and each method, we calculated the Accuracy, and Fig. 6 shows the CDFs of the Accuracy across different software changes of the 72 software changes that induced performance changes as well as the 72 ones that did not. The improved SST only improved robustness for SST, and even though it had a small number of false negatives in the software changes assessment, just like CUSUM, it failed to exclude the KPI changes induced by other reasons, i.e., seasonality, thus the improved SST had relatively low Accuracy. With the introduction of DiD method to eliminate KPI changes caused by other factors, FUNNEL achieved fewer false positives and higher Accuracy. Because MRLS was designed to detect KPI changes that included spikes that we did not care about in our

scenario, and variable KPIs occupied a large portion of all KPIs, MRLS performed poorly in Accuracy.

Based on the characteristics of KPIs, the 9,982 *items* were divided into three types: seasonal, stationary, and variable [7]. It is clear that the *items* with CPU context switch count were variable, and the *items* with memory utilization were stationary. If the KPI of a given *item* has strong seasonality, then the *item* is seasonal. Similarly, the *item* with KPI of strong variability is variable, and the rest are stationary. For *items* of service, 705 *items* were seasonal, 2,702 *items* were stationary, and 4,713 *items* were variable.

To examine the performance of FUNNEL in handling seasonal and variable KPIs, we reorganized the above KPI items based on the KPI types, and then compared different methods. For each method, we multiplied the TPs, TNs, FPs, and FNs of the 72 software changes that did not introduce behavior changes with 86 (6,194/72), and added the result to the TPs, TNs, FPs, and FNs of the 72 software changes which induced KPI changes. We believe that this provides a reasonable approximation of real TPs, TNs, FPs, and FNs. Eventually, we calculated the accuracy, recall, precision and TNR based on the *synthetic* true positives, true negatives, false positives, and false negatives. Table 1 shows the aggregated results for each method and each KPI type. FUNNEL performed the best across all three kinds of KPIs. The improved SST, CUSUM and MRLS performed well in the face of stationary KPIs. However, because CUSUM and the improved SST failed to exclude the impact of seasonality, they detected performance changes induced by software changes with low accuracy when the KPIs had strong seasonality. MRLS was sensitive to spikes, and it was hardly feasible to modify MRLS to detect level shifts or ramp up/downs only. Thus a large portion of the *items* with variable KPIs were incorrectly determined as having KPI changes.

#### 4.3 Comparison of Computational Cost

Since millions of KPIs should be monitored to determine whether they are impacted by software changes, it is important that the computational cost of the impact assessment be relatively low for the sake of scalability and deployability. In this section, we compare the computational cost of FUNNEL, CUSUM, and MRLS.

TABLE 1  
The Precision, Recall, TNR and Accuracy of Seasonal, Stationary and Variable Data for FUNNEL, Improved SST, CUSUM and MRLS

| Algorithm    | Type       | Total   | Precision | Recall  | TNR     | Accuracy |
|--------------|------------|---------|-----------|---------|---------|----------|
| FUNNEL       | Seasonal   | 28,500  | 98.28%    | 100.00% | 100.00% | 100.00%  |
|              | Stationary | 129,943 | 100.00%   | 100.00% | 100.00% | 100.00%  |
|              | Variable   | 215,339 | 68.47%    | 99.48%  | 99.88%  | 99.88%   |
| Improved SST | Seasonal   | 28,500  | 1.10%     | 100.00% | 81.93%  | 81.96%   |
|              | Stationary | 129,943 | 14.28%    | 100.00% | 98.44%  | 98.44%   |
|              | Variable   | 215,339 | 15.04%    | 99.48%  | 98.50%  | 98.50%   |
| CUSUM        | Seasonal   | 28,500  | 0.76%     | 84.21%  | 77.97%  | 77.98%   |
|              | Stationary | 129,943 | 10.34%    | 98.52%  | 97.78%  | 97.78%   |
|              | Variable   | 215,339 | 17.92%    | 96.34%  | 98.82%  | 98.81%   |
| MRLS         | Seasonal   | 28,500  | 100.00%   | 87.72%  | 100.00% | 99.98%   |
|              | Stationary | 129,943 | 9.23%     | 97.33%  | 97.51%  | 97.51%   |
|              | Variable   | 215,339 | 0.61%     | 97.04%  | 57.85%  | 57.95%   |

TABLE 2  
Comparison of Computational Time

| Method                          | FUNNEL        | CUSUM    | MRLS    |
|---------------------------------|---------------|----------|---------|
| Run time per <i>time window</i> | 401.8 $\mu$ s | 1.846 ms | 2.852 s |
| # Cores for one million KPIs    | 7             | 31       | 47,526  |

Motivated by the evaluation method presented in [29], all three methods were deployed on the same server (CPU information: 12 Intel(R) Xeon(R) CPU E5645 @ 2.40 GHz) with a single thread. The CPU utilization remained 100 percent while the process of each method was running so that we could use the total time to evaluate complexity [29].

Table 2 shows the average computational time taken by FUNNEL, CUSUM and MRLS to detect the KPI changes in a single *time window*. With the introduction of matrix compression and implicit inner product calculation, FUNNEL was highly computational efficient. Specifically, compared with MRLS, FUNNEL was more than 7,000 times faster in computational speed. Moreover, FUNNEL reduced 77.42 percent of computational cost as compared to CUSUM. Suppose that one million KPIs need to be monitored and determined for the impact assessment of software changes, the KPIs are collected and detected every minute, and the implementation runs on the same types of CPU. As the last row in Table 2 shows, if we apply MRLS as the software change assessment method, we need at least 47,526 cores, i.e., 3,960 servers if each server has 12 cores. CUSUM needs 31 cores and 3 servers. However, with FUNNEL, one server is fully capable of detecting and determining all one million KPIs. FUNNEL is thus quite competent to assess the impact of software changes online in large web-based services.

#### 4.4 Comparison of Detection Delay

When detecting level shifts, and ramp up/downs, all detection methods need some data points (thus time) to finish. However, rapid detection of performance changes and determination of software change impacts is quite necessary for timely damage mitigation for web-based services. In this section, we compare the detection delay of different methods and show that FUNNEL has a much lower detection delay, thanks to the use of SST.

For each *item*, the operations team labelled the start of KPI changes (if any, as shown in Fig. 2), which served as the ground truth information for comparing different methods. We defined the detection delay as the time between the start of a KPI change and its detection by a method. Suppose that a method correctly detects and determines a KPI change first when the input *time window* is  $x(i+1), x(i+2), \dots, x(c), \dots, x(i+w)$ , and the KPI change starts at time  $c$ , then the detection delay is  $(w-c)$  minute. Note that the detection delay defined above does not include the delay due to the computational cost (previously evaluated separately in Section 4.3). During evaluation, each method is given sufficient CPU power to finish processing one-window's worth of data within one time window.

Fig. 7 shows the Complementary Cumulative Distribution Functions (CCDFs) for the detection delay of FUNNEL, CUSUM, and MRLS. The median delays were  $M_{MRLS} = 21.3 \text{ min}$ ,  $M_{FUNNEL} = 13.2 \text{ min}$ , and  $M_{CUSUM} = 37.7 \text{ min}$ .




Fig. 7. CCDFs of detection delay.

FUNNEL reduced 38.02 percent of detection delay compared with MRLS, and 64.99 percent for CUSUM.

As aforementioned, in FUNNEL we set a threshold of 7 minutes to declare a change in a time series as a level-shift or ramp-up/down rather than a one-off event. Occasionally, MRLS can detect a level shift within 7 minutes, at the cost of much more false positives. In these cases, FUNNEL is slower than MRLS as shown in the top left corner of Fig. 7

The distribution of the detection delay for FUNNEL was more concentrated and the longest detection delay of FUNNEL was much shorter than that of CUSUM and MRLS. This gives FUNNEL a very significant advantage, because a software change may introduce a great loss due to impairment if an unexpected impact of a software change causes poor user experience or degraded advertising income and is detected only after a long delay. Thus, in terms of detection delay, overall FUNNEL is more suitable than CUSUM and MRLS for online impact assessment of software changes in large web-based services.

## 5 OPERATIONAL EXPERIENCE

We deploy the multi-threaded FUNNEL prototype on one server with a 12-core Intel(R) Xeon(R) CPU E5645 @ 2.40 GHz. FUNNEL assesses the software changes of a few dozens of services offered by the search engine company. Table 3 shows some daily statistics for a specific one-week period which we studied in details in this section.

Ideally, we would like to measure the Accuracy, Precision, Recall and TNR as in Section 4.2 for this one-week period. However, it is prohibitive for the operations team to label more than 2 million KPIs of *all* services daily. Therefore, we made a compromise, and only asked operations team to verify the KPI changes detected by FUNNEL (10 thousands per day, fourth column in Table 3). This allows us to calculate Precision =  $\frac{TP}{TP+FP}$ , which is shown in the last column in Table 3.

Based on the operational experience of FUNNEL deployment, we show four representative cases. We applied the FUNNEL's prototype to validate expected performance changes that were direct outcomes of configuration changes in the first two cases, and to detect unexpected performance changes induced by software upgrades in the last two. The

TABLE 3  
Statistics About the Implementation of FUNNEL

| #software changes | # software changes that have impact | #KPIs     | #KPI changes | Precision |
|-------------------|-------------------------------------|-----------|--------------|-----------|
| 24,119            | 268                                 | 2,256,390 | 10,249       | 98.21%    |

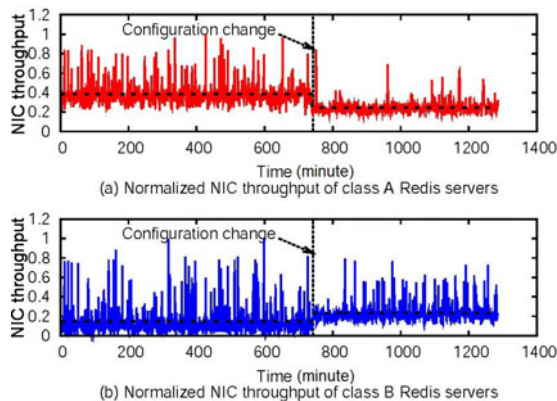


Fig. 8. KPI changes induced by a configuration change in the Redis query service.

cases show that FUNNEL assesses software changes rapidly (Sections 5.2 and 5.3) and accurately even in the face of variable (Sections 5.1) and seasonal (Sections 5.4) KPIs.

### 5.1 Robustness to Variable KPIs

Redis is an advanced data structure store which is used as database, cache and message broker [30]. After a configuration change in the Redis query service, FUNNEL determined that 16 out of 118 KPIs in the impact set had behavior changes which were caused by this configuration change. More specifically, FUNNEL found that immediately after the configuration change, some Redis servers (class A Redis servers) witnessed a negative level shift in NIC throughput as Fig. 8a shows, while in the NIC throughput of other Redis servers (class B Redis servers) a positive level shift occurred, as shown in Fig. 8b.

The operations team verified that the configuration change was aimed at load balancing and achieved the expected result. Specifically, the Redis query tasks were assigned to class A Redis servers first. Unless the class A Redis servers reached the limit of NIC bandwidth capacity, the query tasks would not be allocated to class B Redis servers. This led to a situation where the NICs of class A Redis servers were always busy, while the utilization of the NICs of class B servers was low, which degraded the performance of the Redis query service and shortened the life of NICs in class A Redis servers. The operations team launched a configuration change to balance the traffic between class A and class B Redis servers. As Fig. 8 shows, the configuration change successfully balanced the traffic and had the expected effect.

This case study shows that, although the NIC throughput had strong variability by its nature, FUNNEL still successfully detected and determined the KPI changes induced by the configuration change, demonstrating FUNNEL's capability to handle variable KPIs.

### 5.2 Rapid Detection of Expected KPI Changes

In the second case, we present the performance of FUNNEL in assessing a configuration change in the operational metadata acquisition service. During the deployment of the FUNNEL's prototype, there occurred a configuration change in the operational meta-data acquisition service, after which FUNNEL identified that there were 44 KPIs in the impact set, and determined that the configuration change induced behavior changes in 14 KPIs.

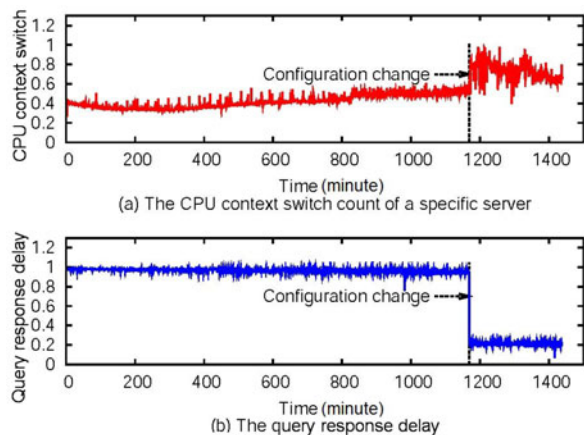


Fig. 9. KPI change induced by a configuration change in server/service status acquisition service.

More specifically, FUNNEL detected KPI changes in the CPU context switch count on servers where the configuration change was deployed, as Fig. 9a shows, as well as the KPI changes in the query response delay, as Fig. 9b shows.

We delivered the assessment result to the operations team, and they verified that the result was as expected. Specifically, the operational metadata acquisition service provides a query tool that enables the employees in the company to query the status of any server or service. The operations team carefully monitors the performance of the service, and if a performance degradation occurs in the service, the operators will mitigate it as soon as possible. The query tool used to deliver the parameters to the operational metadata database by using a single thread. But this caused frequent timeouts after the deployment of the query tool thus the operations team found that the majority of delays were induced by this single thread configuration. Therefore, the operations team launched a new configuration change that would make the query tool multi-threaded, which naturally caused the symptoms in Fig. 9.

The KPI changes were detected and determined 7 minutes after the configuration change, showing that FUNNEL is capable of online impact assessment of software changes.

### 5.3 Rapid Detection of Unexpected KPI Changes

To see how FUNNEL speeds up the detections of unexpected behavior changes, we randomly picked a small fraction of software changes for which FUNNEL does not directly deliver the detection results to the operations team. Instead, the operations team independently assess software changes without the help of FUNNEL. In the below case, we compare the speed of FUNNEL and manual inspection, and demonstrate the performance of FUNNEL in detecting unexpected performance changes

During the deployment of FUNNEL, the operations teams made a software upgrade aimed at improving the performance of the advertising system. Advertising system is a very large and complex system [4], and in fact 36,752 KPIs are included in the impact set in the software upgrade according to FUNNEL. Manually investigating this upgrade's impact is infeasible. 10 minutes after the software upgrade, FUNNEL detected 1,141 KPI changes induced by the software upgrade. More specifically, as Fig. 10 shows, the normalized ratio of clicks on advertisements considered as cheats, i.e., clicks on




Fig. 10. The normalized ratio of clicks on advertisements considered as cheats.

advertisements that are not considered as human behavior by the anti-cheating system, increased dramatically immediately after the software upgrade was conducted.

Without using FUNNEL, the operations team independently found the performance change after 1.5 h and then quickly fixed the issue. The operations team confirmed that the occurrence of the unexpected behavior changes due to the software upgrade was a real significant incident. It was later carefully investigated. Specifically, the anti-cheating service injects a check program, which is implemented with JSON, into every advertisement to determine whether a click on the advertisement is performed by a human or is a cheat executed by an automated program. However, the software upgrade failed to load the JSON program on the iPhone browsers, i.e., the check program was not effective for iPhone users. Therefore, all clicks on the advertisements by iPhone users were considered as cheats, which resulted in the decline in advertising revenue. The number of cheating clicks monitored by the anti-cheating service returned to the normal level after the operations team had remedied the situation, and thus, a negative level shift occurred 1.5 h after the software upgrade.

Had the operations team used FUNNEL to assess the impact of the erroneous software upgrade, they could have much more timely mitigated the loss caused by the software upgrade.

#### 5.4 Robustness to Seasonal KPIs

In the last case, we show how FUNNEL performs in the face of seasonal KPIs. With rapid development of the Internet, online shopping has witnessed an explosive growth. The stability of the online shopping service is of vital importance to guarantee user experience and service income. When the service is upgraded, the operations team carefully inspects the KPIs that maybe indicative of performance impairment or software bugs. Such KPIs should be detected rapidly for timely mitigation. However, given the large number of KPIs in the online shopping service (tens of thousands of KPIs should be monitored), it becomes incredibly challenging to identify the impact of software upgrades manually. We applied FUNNEL to detect and determine KPI changes after a software upgrade in the online shopping service.

As is in Section 5.3, FUNNEL did not deliver the results to the operations team at first, and the operations team assessed the software upgrade independently. FUNNEL detected several KPI changes and determined that these KPI changes were caused by the software upgrade about 8 minutes after the upgrade. Concretely, the software upgrade tried to merge multiple mget requests [30] (a command used to query data from Redis server) into one




Fig. 11. The normalized number of successful orders.

mget request, which however, degraded the performance of Redis proxy in handling mget requests. The computational resources of Redis servers were then exhausted, and a large number of web page view timeouts and failures occurred. As a consequence, the shopping online service witnessed a degradation in the number of successful orders and in the income. The operations team detected the unexpected performance changes and correlated them to the software upgrade about half an hour after the software upgrade.

Fig. 11 shows the normalized number of successful orders of the shopping service, and FUNNEL detected a negative level shift in the KPI, and determined that the behavior changes were caused by the software upgrade. The KPIs had strong seasonality, but FUNNEL still accurately assessed the impact of the software upgrade, which demonstrated that FUNNEL performed extremely well even in the face of seasonal KPIs.

## 6 RELATED WORK

The impact of software changes has attracted considerable attention in recent years [6], [7], [14], [31], [32], [33]. Mahimkar et al. [6] developed MERCURY to detect the performance impact of upgrades in large operational networks. MERCURY uses the CUSUM method to detect behavior changes in KPIs, and applies statistical rule mining and network configuration to identify commonality across the behavior changes. PRISM [7] is developed to reduce the detection delay between the behavior change and the detection. The MRLS method was developed in PRISM to rapidly and robustly detect maintenance-induced behavior changes. However, the MRLS method suffers from high computational cost, and thus is not appropriate in our scenario. Litmus was developed to address the assessment of changes in cellular networks where external factors may over-shadow the assessment [14]. Litmus applies a spatial regression algorithm for the comparison of the treated and the control group.

The detection of behavior changes has a very rich literature. In [34], Principal Component Analysis (PCA) was applied for anomaly detection in network. Yamada et al. proposed a change point detection method, additive Hilbert-Schmidt Independence Criterion (aHSIC), which was based on supervised learning, and used the weighted sum of the SIC scores for incorporating feature selection [29]. To detect changes in seasonal time-series,

Chef et al. applied a time series decomposition based method, week-over-week [28].

Split testing, also known as the control/treatment test, has been widely studied in the fields of econometrics [15], [16], health care [17], and web domains [24], [25]. Web users were exposed to either the treated group or the control group in the study reported in [24], [25]. To analyze the impact of the treatment, the authors recorded the interactions of users and compared the metrics of the treated and the control group. Stuart et al. [17] applied propensity scores to address the disadvantage of the DiD algorithm when the compositions of the control group and the treated group change over time. However, in our scenario, the compositions of the control and the treated group remain unchanged over time. That is, it is not necessary to introduce propensity scores in FUNNEL.

## 7 CONCLUSION

We designed and implemented a new tool, FUNNEL, for rapidly and robustly assessing the impact of software changes in large web-based services. For each software change, FUNNEL analyzes all the services, processes, and servers that may be influenced and automatically constructs the impact set. FUNNEL then detects performance changes in the impact set rapidly and robustly by improving the robustness and reducing the computational cost of SST, and determines performance changes induced by software changes using a DiD method. We evaluated FUNNEL by comparing it with CUSUM and MRLS using 144 software changes and showed that FUNNEL achieved high accuracy with short detection delay and low computational cost. Operational experiences of FUNNEL deployment show that FUNNEL can assess the impact of software changes rapidly and robustly in the face of seasonary and variable data. Compared to manual efforts, in one specific case FUNNEL shortens the assessment from 1.5 hours to 10 minutes, saving both time and money.

## ACKNOWLEDGMENTS

The authors are grateful to Bo Wang, Bolin Yang, Wenhui Xie, Ming Hua, Wei Huang, Xiaoguang Miao, and Jiantao Chi, for their invaluable help at various stages of this project. They strongly appreciate Yosef Azzabi, Dapeng Liu, Juexing Liao, Beibei Miao, Yongqian Sun, and Xiaohui Nie for their helpful suggestions and elaborative proofreading. The work was supported by National Natural Science Foundation of China (NSFC) under Grant Nos. 61472214 and 61402257, the National Key Basic Research Program of China (973 program) under Grant No. 2013CB329105, the Tsinghua University Self-determined Project under Grant No. 2014z21051, the State Key Program of National Science of China under Grant No. 61233007, the National Natural Science Foundation of China under Grant No. 61472210, the National High Technology Development Program of China (863 program) under Grant No. 2013AA013302.


## REFERENCES

- [1] (2014). Back up and running [Online]. Available: <https://blogs.dropbox.com/dropbox/2014/01/back-up-and-running/>
- [2] (2012). Google apps incident report, gmail partial outage [Online]. Available: <http://www.google.com/appsstatus/ir/plibxfjh8whr44h.pdf>
- [3] (2008). Facebook chat [Online]. Available: [https://www.facebook.com/note.php?note\\_id=14218138919](https://www.facebook.com/note.php?note_id=14218138919)
- [4] R. Bhagwan, R. Kumar, R. Ramjee, G. Varghese, S. Mohapatra, H. Manoharan, and P. Shah, "Adtributor: Revenue debugging in advertising systems," in *Proc. 11th USENIX Conf. Netw. Syst. Des. Implementation*, Apr. 2014, pp. 43–55.
- [5] C. Luo, J.-G. Lou, Q. Lin, Q. Fu, R. Ding, D. Zhang, and Z. Wang, "Correlating events with time series for incident diagnosis," in *Proc. 20th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, New York, NY, USA, Aug. 2014, pp. 1583–1592.
- [6] A. Mahimkar, H. H. Song, Z. Ge, A. Shaikh, J. Wang, J. Yates, Y. Zhang, and J. Emmons, "Detecting the performance impact of upgrades in large operational networks," in *Proc. SIGCOMM*, New Delhi, India, Aug. 2010, pp. 303–314.
- [7] A. Mahimkar, Z. Ge, J. Wang, J. Yates, Y. Zhang, J. Emmons, B. Huntley, and M. Stockert, "Rapid detection of maintenance induced changes in service performance," in *Proc. 7th Conf. Emerging Netw. Exp. Technol.*, Tokyo, Japan, Dec. 2011, no. 13, pp. 1–12.
- [8] Z. Lin, M. Chen, and Y. Ma, "The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices," *arXiv Preprint arXiv:1009.5055*, 2010.
- [9] T. Idé and K. Inoue, "Knowledge discovery from heterogeneous dynamic systems using change-point correlations," in *Proc. 5th SIAM Int. Conf. Data Mining*, Newport Beach, CA, USA, Apr. 2005, pp. 571–575.
- [10] H. Hassani, S. Heravi, and A. Zhigljavsky, "Forecasting european industrial production with singular spectrum analysis," *Int. J. Forecasting*, vol. 25, no. 1, pp. 103–118, 2009.
- [11] V. Moskvina and A. Zhigljavsky, "Change-point detection algorithm based on the singular-spectrum analysis," *Commun. Statist.: Simul. Comput.*, vol. 32, no. 2, pp. 319–352, 2003.
- [12] T. Idé and K. Tsuda, "Change-point detection using Krylov subspace learning," in *Proc. 7th SIAM Int. Conf. Data Mining*, Apr. 2007, pp. 515–520.
- [13] Y. Mohammad and T. Nishida, "Robust singular spectrum transform," in *Proc. 22nd Int. Conf. Indus., Eng. Other Appl. Appl. Intell. Syst. Next-Gener. Appl. Intell.*, Jun. 2009, pp. 123–132.
- [14] A. Mahimkar, Z. Ge, J. Yates, C. Hristov, V. Cordaro, S. Smith, J. Xu, and M. Stockert, "Robust assessment of changes in cellular networks," in *Proc. 9th ACM Conf. Emerging Netw. Exp. Technol.*, Santa Barbara, California, USA, Dec. 2013, pp. 175–186.
- [15] O. C. Ashenfelter and D. Card, "Using the longitudinal structure of earnings to estimate the effect of training programs," *Rev. Econ. Statist.*, vol. 67, no. 4, pp. 648–660, Aug. 1985.
- [16] W. R. Shadish, T. D. Cook, and D. T. Campbell, *Experimental and Quasi-Experimental Designs For Generalized Causal Inference*. Belmont, CA, USA: Wadsworth, 2002.
- [17] E. A. Stuart, H. A. Huskamp, K. Duckworth, J. Simmons, Z. Song, M. E. Chernew, and C. L. Barry, "Using propensity scores in difference-in-differences models to estimate the effects of a policy change," *Health Serv. Outcomes Res. Methodol.*, vol. 14, no. 4, pp. 166–182, Aug. 2014.
- [18] Q. Wang, Y. Kanemasa, M. Kawaba, and C. Pu, "When average is not average: Large response time fluctuations in n-tier systems," in *Proc. 9th Int. Conf. Autonom. Comput.*, 2012, pp. 33–42.
- [19] D. Liu, Y. Zhao, H. Xu, Y. Sun, D. Pei, J. Luo, X. Jing, and M. Feng, "Opprentice: Towards practical and automatic anomaly detection through machine learning," in *Proc. ACM Conf. Internet Meas. Conf.*, Oct. 2015, pp. 211–224.
- [20] N. Aoki, *State Space Modeling of Time Series*. Cambridge, U.K.: Cambridge Univ. Press, 1990.
- [21] N. Mi, G. Casale, L. Cherkasova, and E. Smirni, "Burstiness in multi-tier applications: Symptoms, causes, and new models," in *Proc. 9th ACM/IFIP/USENIX Int. Conf. Middleware*, 2008, pp. 265–286.
- [22] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD, USA: The Johns Hopkins Univ. Press, 2012, vol. 3.
- [23] W. H. Press, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*. Cambridge, U.K.: Cambridge Univ., 2007.
- [24] R. Kohavi, A. Deng, B. Frasca, R. Longbotham, T. Walker, and Y. Xu, "Trustworthy online controlled experiments: Five puzzling outcomes explained," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Aug. 2012, pp. 786–794.
- [25] R. Kohavi, R. M. Henne, and D. Sommerfield, "Practical guide to controlled experiments on the web: Listen to your customers not to the hippo," in *Proc. 13th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, San Jose, CA, USA, Aug. 2007, pp. 959–967.

- [26] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proc. 10th ACM SIGCOMM Conf. Internet Meas.*, Melbourne, Australia, Nov. 2010, pp. 267–280.
- [27] S. Kandula, S. Sengupta, A. Greenberg, P. Patel, and R. Chaiken, "The nature of data center traffic: Measurements & analysis," in *Proc. 9th ACM SIGCOMM Conf. Internet Meas. Conf.*, Chicago, Illinois, USA, Nov. 2009, pp. 202–208.
- [28] Y. Chen, R. Mahajan, B. Sridharan, and Z.-L. Zhang, "A provider-side view of web search response time," in *Proc. SIGCOMM*, Aug. 2013, pp. 243–254.
- [29] M. Yamada, A. Kimura, F. Naya, and H. Sawada, "Change-point detection with feature selection in high-dimensional time-series data," in *Proc. 23rd Int. Joint Conf. Artif. Intell.*, Beijing, China, Aug. 2013, pp. 1827–1833.
- [30] (2015). Redis [Online]. Available: <http://redis.io/>
- [31] M. Canini, D. Venzano, P. Peresini, D. Kostic, and J. Rexford, "A nice way to test openflow applications," in *Proc. 9th USENIX Conf. Netw. Syst. Des. Implementation*, San Jose, CA, USA, Apr. 2012, p. 10.
- [32] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker, "Abstractions for network update," in *Proc. SIGCOMM*, Aug. 2012, pp. 323–334.
- [33] R. R. Sambasivan, A. X. Zheng, M. De Rosa, E. Krevat, S. Whitman, M. Stroucken, W. Wang, L. Xu, and G. R. Ganger, "Diagnosing performance changes by comparing request flows," in *Proc. 8th USENIX Conf. Netw. Syst. Des. Implementation*, Boston, MA, USA, Apr. 2011, pp. 43–56.
- [34] A. Lakhina, M. Crovella, and C. Diot, "Mining anomalies using traffic feature distributions," in *Proc. SIGCOMM*, Aug. 2005, pp. 217–228.




**Shenglin Zhang** received the BE degree in network engineering from the School of Computer Science and Technology, Xidian University in 2012. He is currently working toward the PhD degree in the Institute of Network Sciences and Cyberspace, Tsinghua University. His research interests include failure detection and prediction in datacenter networks, BGP AS path behavior, and P2P optimization.




design and implementation. She is a member of the IEEE.

**Dan Pei** received the BE and MS degree in computer science from the Department of Computer Science and Technology, Tsinghua University in 1997 and 2000, respectively, and the PhD degree in computer science from the Computer Science Department, University of California, Los Angeles (UCLA) in 2005. He is currently an associate professor in the Department of Computer Science and Technology, Tsinghua University. His research interests include network and service management in general. He is a senior member

of the IEEE and ACM.



**Yu Chen** received the BS and MS degree in computer science from the Department of Computer Science and Technology, Peking University in 1998 and 2001, respectively. He used to be a senior engineer in Baidu, Inc, and he is currently working in a startup company on education.



**Xianping Qu** received the BS degree in information and computing science from the School of Mathematical Sciences, Fudan University in 2009. He is currently a senior engineer in Baidu, Inc. His research interests include monitoring, anomaly detection, root cause analysis, and automatic operations of web-based services.



**Shimin Tao** received the BS degree in public administration management from the College of Politics & Law and Public Administration, Hubei University in 2003. He is currently a technical manager in Baidu, Inc. His research interests include monitoring, anomaly detection, root cause analysis, and automatic operations of web-based services.



**Zhi Zang** received the BS degree in computer science from the School of Computer Science and Technology, Nanjing University in 2005 and the ME degree in computer software and theory from the Institute of Software, Chinese Academy of Sciences in 2008. He is currently a senior engineer in Baidu, Inc. His research interests include monitoring, anomaly detection, root cause analysis, and automatic operations of web-based services.



**Xiaowei Jing** received the BE degree in computer science and application from the Harbin University of Science and Technology in 1995 and the MS degree in computer science from the Harbin Engineering University in 2002. He is currently a senior engineer in the IT Management Department, Petrochina, Inc.



**Mei Feng** received the PhD degree from the Chemical Process Engineering Department, Newcastle University in 2001. She is currently a professor in the Computer Application Department, Research Institute of Petroleum Explore and Development (RIPED). Her research interests include computer network construction and information security.