Understanding E-Commerce Systems under Massive Flash Crowd: Measurement, Analysis, and Implications

Junqiang Ge, *Student Member, IEEE,* Ye Tian, *Member, IEEE,* Lei Liu,
Rongheng Lan, and Xinming Zhang, *Member, IEEE*

Abstract—Leading e-commerce providers have built large and complicated systems to provide countrywide or even worldwide services. However, there have been few substantive studies on e-commerce systems in real world. In this paper, we investigate the systems of Tmall and JD, the top-two most popular e-commerce websites in China, with a measurement approach. By analyzing traffics from campus network, we present a characterization study that covers several features, including usage patterns and shopping behaviors, of the e-commerce workload; in particular, we characterize the massive *flash crowd* in the Double-11 Day, which is the biggest online shopping festival in the world. We also reveal Tmall and JD's e-commerce infrastructures, including *content delivery networks* (*CDNs*) and *clouds*, and evaluate their performances under the flash crowd. We find that Tmall's CDN proactively throttles bandwidths for ensuring low but guaranteed throughputs, while JD still follows the best-effort way, leading to poor and unstable performances; both providers do not have sufficient capacities in their private clouds, resulting in extraordinarily long transaction latencies. Based on the insights obtained from measurement, we discuss the design choices of e-commerce CDNs, and investigate the potential benefits brought by incorporating client-side assistances in offloading massive flash crowd of e-commerce workloads.

Index Terms—E-commerce system, performance measurement, flash crowd, content delivery network (CDN), cloud.

1 Introduction

With the technological advances on Web, cloud computing, and mobile Internet, *electronic commerce* (*e-commerce*) becomes increasingly popular in recent years. It is estimated that in 2015, the retail e-commerce sales worldwide amounted to \$1.67 trillion [1], and in China, 11.1% of the retail sales were on the Internet [2].

For enabling large volume of online transactions and providing countrywide or even worldwide services, leading e-commerce providers such as Amazon and Alibaba have built large and complicated systems. In e-commerce, service availability and system's performance are critical to providers, as it is estimated that for a leading e-commerce Website like Amazon, one second of service latency is worth tens of thousands of US dollars [3]. However, there have been few substantive studies on large-scale e-commerce systems in real world. In this paper, we focus on *Tmall* [4] and *JD* [5], which are the top-two most popular e-commerce websites in China, and investigate their systems with a measurement approach. We analyze the workloads upon Tmall and JD that are collected from our campus network, and investigate behaviors and performances of the e-commerce infrastructures with passive and active measurements. In particular, we characterize the flash crowd [6] in the *Double-11 Shopping Day* [7], which is the biggest online shopping festival in the world, and evaluate the e-commerce systems' performances under the flash crowd.

As far as we know, this work is the first measurement study on large-scale e-commerce systems in real world. We find that the Double-11 Shopping Day impose a massive workload on e-commerce systems; it is challenging for the systems to accommodate the service requests during the flash crowd; and users have poor service experiences. Based on the observations from measurement, we have insightful discussions on improving large-scale e-commerce systems under massive flash crowd. The main contributions of this paper are summarized as follows.

- We introduce effective methodologies that combine passive and active measurements to study Tmall and JD's large-scale e-commerce systems, including the *content delivery networks* (*CDNs*) and the *clouds*. In particular, we have developed methodologies for identifying and correlating network traffics with various e-commerce workloads (e.g., product browsing, shopping cart operations, checkouts, etc.). To overcome shortages of PlanetLab nodes and open recursive local DNS (LDNS) servers in China, which are widely used in previous studies [8] [9], we propose to employ VPN servers as measurement vantage points, and have developed a suite of techniques that enable us to accurately evaluate performances of e-commerce CDNs and clouds.
- We present the first characterization study on the e-commerce workload, in particular, the massive flash crowd in the Double-11 Shopping Day. Our study covers several features, including usage patterns and

- *The authors are with the Anhui Key Laboratory on High Performance Computing and Applications, School of Computer Science and Technology, University of Science and Technology of China, Hefei, Anhui 230026, P.R. China. Phone:+86-551-63602445, Fax: +86-551-63601013, E-mail: gjq1993@mail.ustc.edu.cn, yetian@ustc.edu.cn, {lliu0610, lan0729}@mail.ustc.edu.cn, xinming@ustc.edu.cn.*
- *Corresponding author: Ye Tian (yetian@ustc.edu.cn).*

shopping behaviors. For example, we find that there exists a Zipf-like popularity among the products, most of the shopping cart operations are read-only queries, the Double-11 Shopping Day attracts many users, who are more willing to buy than usual, and the rush buying behaviors at the very beginning of the shopping festival impose a massive flash crowd on the e-commerce systems.

- We investigate the behaviors and performances of the e-commerce infrastructures, including the CDNs and the clouds. We find that both Tmall and JD have decent CDN throughputs, but the throughputs degrade significantly under the Double-11 flash crowd, despite that several efforts, such as expanding CDN footprint and capacity-based client mapping, have been made to accommodate the massive content requests. We observe that Tmall's e-commerce CDN adopts a proactive bandwidth throttling to provide low but guaranteed throughputs under the Double-11 flash crowd, while JD still follows the best-effort way. As for the e-commerce clouds, both Tmall and JD do not have sufficient capacities during the busy hours, and users suffer extraordinarily long latencies during their e-commerce transactions.
- Finally, we discuss the implications of the observations from our measurement study. We argue that for an e-commerce CDN to serve a massive flash crowd as the one in the Double-11 Day, the design choices of utilizing massive data centers and the DNS-based client redirection are preferred. We also show that by incorporating assistances from clients, considerable flash crowd workload on both e-commerce CDN and cloud can be offloaded.

The remainder part of this paper is organized as the following: Section 2 discusses the related work. Section 3 presents an overview on Tmall and JD. We describe our measurement methodology in Section 4. Section 5 characterizes the e-commerce workload. We evaluate performances of the e-commerce infrastructures in Section 6. Section 7 discusses the implications for improving e-commerce systems. Finally, we conclude this paper and discuss the future work in Section 8.

2 Related Work

E-commerce has dramatically changed people's daily life. There are many studies on the behaviors/mis-behaviors of buyers and sellers in the cyber marketplace (e.g., [10] and [11]), however, few substantive works focus on e-commerce infrastructures. In this paper, we present a measurement study on the large-scale e-commerce systems of Tmall and JD, the top-two most popular e-commerce websites in China.

There have been a rich literature on measuring and evaluating commercial CDNs. Huang et al. [8] evaluate two representative commercial CDNs: Akamai and Limelight, and show that their different design philosophies lead to different performances. Triukose et al. [12] carry out a measurement study on Akamai, and estimate the performance of a more consolidated CDN deployment. Wendell et al. [6] analyze the sudden spikes of CDN traffics on an open CDN and discuss the implications. Adhikari et al. [13] unveil the architecture of YouTube's video CDN in details. He and Tian [14] [15] reveal the tradeoff between energy and traffic costs of a video CDN, and propose a capacity provisioning algorithm for optimizing the overall cost. Adhikari et al. [9] study the CDN selection problem in Netflix and Hulu, and present a CDN selection strategy that can significantly increase users' available bandwidth while still conforming to the business constraints. Our work differs from the previous works in two aspects: First, for overcoming shortages of PlanetLab nodes and open recursive local DNS (LDNS) servers, which are widely used in previous works, we propose to use VPN servers as vantage points in our CDN measurement, and have developed novel techniques for evaluating CDN performance with VPN server vantage points. Second, we focus on a special event, namely the Double-11 Online Shopping Festival, study and evaluate the CDNs under the event's massive flash crowd. As far as we know, such a massive flash crowd has not been thoroughly investigated in previous works. We also point out the preferred design choices for an e-commerce CDN.

For evaluating cloud services, Li et al. [16] compare a number of public cloud service providers, and develop a tool called `CloudCmp` to benchmark their performances. Nasiriani et al. [17] further investigate the capacity dynamism of the Amazon EC2 instances. Bermudez et al. [18] study Amazon's Web Service provided by EC2 and S3 with a passive measurement approach, and show that there is room for improvement on Amazon's server selection strategy. Chen et al. [19] investigate the impact of the front-end server placement strategies on the performance of a dynamic content cloud service. Ren et al. [20] investigate a cloud file system, and propose a benchmarking suite for evaluating cloud file services. Unlike these works, we develop novel measurement-based techniques that enable us to evaluate performances of a real-world cloud database service persistently and in large scale. With the methodology, we are the first to systemically investigate e-commerce cloud database services under the extreme circumstance of the Double-11 massive flash crowd.

For improving cloud services, Wang et al. [21] study a customer-provided cloud platform and propose an instance recommendation mechanism. Niu et al. [22] design an optimization algorithm to distribute service requests between private and public clouds. Singh et al. [23] propose a framework for a cloud to allocate resources for meeting the service level agreement (SLA). Amiri et al. [24] survey the workload prediction methodologies, which are critical in cloud resource management. However, few works focus on cloud system under massive flash crowd. With the insights and data obtained from our measurement, in this paper we discuss preferred CDN design choices, and analyze the potential benefits brought by incorporating client-side assistances for an e-commerce system.

3 Overview

3.1 Background

`Tmall.com` [4] and `JD.com` [5] are the top-two most popular Chinese business-to-consumer (B2C) online retail websites. Both websites currently host tens of thousands

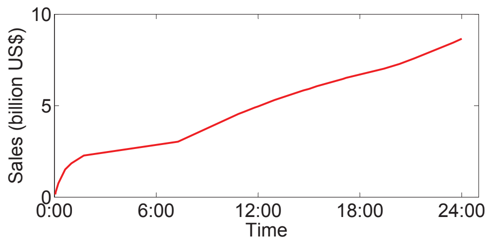

Fig. 1. Tmall's sales growth in 2015's Double-11 Day, data from [25].

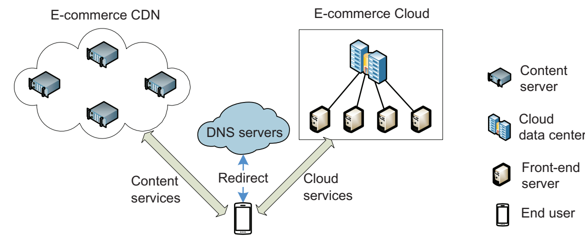

Fig. 2. E-commerce system architecture.

of online stores, and serve hundreds of millions of buyers. According to Alexa, as of Feb. 2016, Tmall and JD were the 8^{th} and 13^{th} most visited websites in China respectively.

Tmall and JD have very similar website organization. Each online store on Tmall or JD has a catalog Web page that lists the items for sale, and each item also has a Web page, which displays texts, high-resolution images and photos provided by the seller for describing/advertising the item, and some item pages even have videos embedded.

Each registered user on Tmall or JD has a virtual shopping cart. When browsing goods online, a user can add an item into his shopping cart by clicking the "add to shopping cart" button on the item page. On the shopping cart page, a user can sort out the items he has previously added, and select some of them to remove from the shopping cart; the user can also select some items he intends to buy, click the "check out" button for the bill, and pay for them online.

Both Tmall and JD hold several cyber shopping festivals, among them, the *Double-11 Online Shopping Day* [7], which is on Nov. 11 each year, is the biggest and most well-known one. It is reported that in 2015's Double-11 Day, only Tmall alone had sales at US$ 8.7 billion [25], which is greater than the online sales in 2015's Thanksgiving Day, Black Friday, and Cyber Monday combined [26]. Fig. 1 shows how the Tmall sales grew in 2015's Double-11 Day. Note that the fastest growth happened in the first hour (i.e., 0:00$\sim$1:00). As we will see in this paper, this is because a lot of buyers rush to buy at the very beginning of the day, in case that the products they are interested in are sold out very soon.

The large volume of online transactions in the Double-11 Day impose a massive workload on the e-commerce infrastructures, and in particular, the rush buying behaviors at the very beginning of the day incur a flash crowd of service requests. Clearly, under the explosive growth of the workload, it is very challenging for Tmall and JD to ensure their services' availabilities and guarantee the users' experiences. In fact, in each year's Double-11 Day, there are complaints on the accessibilities of the e-commerce services. In this paper, we will systematically investigate Tmall and JD's e-commerce systems under the Double-11 flash crowd.

3.2 Architecture Overview

Tmall and JD provide most of their services on Web. Roughly speaking, the e-commerce services can be categorized into two kinds: the *content service* and the *cloud service*. As shown in Fig. 2, the content service is provided with CDN. Both Tmall and JD build their dedicated CDNs for distributing e-commerce contents, such as static Web pages, javascripts, and high-resolution images. The e-commerce CDN consists of many *content servers* that are deployed at various locations and ISP networks.

Both Tmall and JD run their private e-commerce clouds, which provide services such as search engine, recommendation, shopping cart, billing, etc. As shown in Fig. 2, an e-commerce cloud is composed of at least one *cloud data center* and many *front-end Web servers*. The cloud data center maintains elastic computing/database/storage capacities for all back-end jobs, such as handling users' database reads/writes/queries regarding their shopping carts, maintaining sales and inventory databases, executing ranking and recommendation algorithms, etc. The front-end servers proxy between end users and the cloud data center. More specifically, a front-end server receives a user's service request in HTTP or JSON, processes and forwards it to the cloud data center; when a response is returned from the data center, the front-end server generates a dynamic Web page containing the service response and sends it back to the user.

To provide nationwide services in China, both Tmall and JD employ DNS redirection [8] to assign content and front-end servers to users. For example, when a user from our campus network wants to access his shopping cart on Tmall, his DNS query for the name "cart.tmall.com" is firstly resolved to the Canonical Name (CNAME) "cart.tmall.com.danuoyi.tbcache.com", where the suffix "danuoyi.tbcache.com" indicates that the name is managed by Alibaba's DNS system, then by resolving the CNAME, a front-end server at the address "121.194.7.253" is assigned to handle the user's operation requests on his shopping cart.

4 Measurement Methodology

In this section, we describe our methodologies in the measurement study on Tmall and JD's e-commerce systems.

4.1 Passive Measurement Methodology

For the passive measurements on Tmall and JD, we collect traffics of the two e-commerce websites at the gateway of our university campus network, which connects tens of thousands of computers from offices, laboratories, student dormitories, etc.

We employ a high-performance network traffic analyzer named *iProbe* [27] to collect the e-commerce traffics. For each HTTP flow, iProbe keeps a record in the log file that contains the fields such as HTTP method and URL, source/destination addresses and ports, flow size in terms of Bytes and packets in both directions, etc. With iProbe, we have collected two datasets:

TABLE 1
URLs and methods for requesting some e-commerce services on JD from PC and smartphone.

	PC Web browser	Smartphone app.
Login	passport.jd.com/loginservice.aspx?string (get HTML)	jpns.m.jd.com/client.action (post JSON)
Search	search.jd.com/Search?string (get HTML)	search.m.jd.com/client.action (post JSON)
Shopping cart	**Add**: cart.jd.com/gate.action?string (get HTML) **Remove**: cart.jd.com/removeSkuFromCart.action (post JSON) **Query**: cart.jd.com/cart (get HTML)	cart.m.jd.com (post JSON)
Checkout	trade.jd.com/shopping/order/submitOrder.action (post JSON)	order.m.jd.com/client.action (post JSON)

- One dataset contains all the HTTP flows associated with Tmall and JD in a week between 00:00 Mar. 1, 2016 and 23:59 Mar. 7, 2016. The dataset covers five weekdays and two weekends on Mar. 5 and Mar. 6. In the remainder part of this paper, we refer to this dataset as *WEEK*.
- We also collect in 2015's Double-11 Online Shopping Day from 00:00 to 23:59 on Nov. 11, 2015. We refer to this dataset as *D11*.

For each HTTP flow recorded by iProbe, we look up its domain name to decide which e-commerce service the flow is about. For example, Table 1 lists the methods and URLs that are used for requesting some JD e-commerce services from different user devices.

Our datasets only contain the HTTP flows between campus users and e-commerce servers, thus we do not capture all the Tmall's e-commerce traffics, as Tmall provides services on both HTTP and HTTPS simultaneously. But fortunately for JD, we have captured all its e-commerce traffics, since JD provides all its services over non-encrypted HTTP connections, although it encrypts some data pieces such as user passwords in JSON.

4.2 Active Measurement Methodology

4.2.1 VPN Server Vantage Point

In addition to the passive measurements, we also actively probe Tmall and JD's e-commerce infrastructures from a number of vantage points. Previous studies employ PlanetLab nodes or open recursive local DNS (LDNS) servers as the measurement vantage points [8] [9], however, as there are few PlanetLab nodes on the Chinese Internet [28] and the LDNS servers are usually not open in China, we employ a commercial virtual private network (VPN) service named "517VPN"[1], and use its numerous VPN servers, which are widely distributed in China, as our vantage points.

In our active measurement, after establishing a Layer 2 Tunneling Protocol (L2TP) connection between our measurement computer in the campus network and a remote VPN server, we run our probing programs to request various e-commerce services through the tunnel. On behalf of the measurement computer, the VPN server resolves the services' domain names from its local DNS server, and sends the service requests to the content or front-end servers that are assigned to it by the e-commerce DNS system.

1. http://www.517vpn.cn/

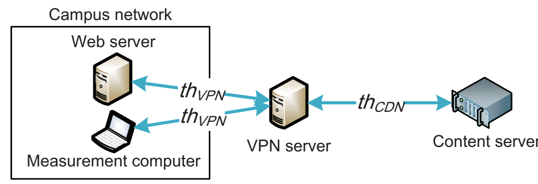

Fig. 3. Methodology for measuring the CDN throughput.

4.2.2 Probing E-Commerce CDN

We have developed a suite of techniques for evaluating e-commerce CDN and cloud from the VPN server vantage points. For evaluating the e-commerce CDNs, we focus on the image distribution service, as images are the most requested e-commerce content. In our measurement, we maintain URLs of a collection of over 100 high-resolution images, whose sizes are between 500 and 600 kB. The images are on Tmall and JD's item pages. After connecting to a VPN server, we employ the tool `cURL` [29] to request and download the images from the content server that is assigned to our vantage point, and record the download throughput as $th_{cURL} = \frac{s_{img}}{t_{dl}}$, where s_{img} is the image size, and t_{dl} is the download time.

To bypass any transparent Web caches, we apply the methodology used in [12] by appending a random string to the image URL. For example, for downloading an image from Tmall's CDN at "http://img.alicdn.com/$\cdots$/xx.png", we actually request an URL like "http://img.alicdn.com/$\cdots$/xx.png?rand_str" with cURL, where "rand_str" is a random string that changes each time.

We want to point out that th_{cURL} reported by cURL is indeed the minimum of the throughput th_{VPN} between our measurement computer and the VPN server it connects to, and the throughput th_{CDN} between the VPN server and the content server, i.e.,

$$th_{cURL} = \min\{th_{VPN}, th_{CDN}\}$$

To obtain th_{CDN}, we need to identify the cases that th_{cURL} is constraint by th_{CDN} by estimating th_{VPN} and comparing it with th_{cURL}.

For estimating th_{VPN}, we set up a Web server in our campus network, which hosts the same high-resolution image files whose URLs on Tmall and JD are collected for probing their content servers. As demonstrated in Fig. 3, we run cURL on our measurement computer, which is also in the campus network, download the image files from the campus server via the remote VPN server, and measure the download throughput as th'_{cURL}. Obviously in this case

$$th'_{cURL} = \min\{th_{VPN}, th_{VPN}\} = th_{VPN}$$

In our measurement, each time we probe Tmall or JD's content server and obtain the throughput th_{cURL}, we also estimate the throughput th'_{cURL} by downloading the same image files from the campus server at the same time. We compare th_{cURL} with th'_{cURL}: If they are close to each other, we simply ignore the result, as in this case, th_{cURL} is constraint by th_{VPN}; but when th_{cURL} is much smaller than th'_{cURL}, we can infer that the measured throughput th_{cURL} is actually constraint by the throughput th_{CDN} between the VPN server vantage point and the content server, i.e.,

$$th_{CDN} = th_{cURL}, \quad \text{if } th_{cURL} \ll th'_{cURL} \tag{1}$$

Fortunately on most VPN servers, we have $th_{cURL} \ll th_{VPN}$. This is reasonable as we will see in Section 6.2, the throughput bottleneck is at the content server, on which our image downloading request is competing the server's bandwidth with many other clients; moreover, the e-commerce CDN may proactively throttle the per-flow bandwidth, so as to ensure the service availability to as many users as possible.

4.2.3 Probing E-Commerce Cloud

We consider the shopping cart service, which is essentially a cloud database service, as an example to evaluate Tmall and JD's e-commerce clouds. We select the shopping cart service for two reasons: First, it is the most requested e-commerce cloud service by users; Second, the service demands timely responses, that is, the e-commerce system will always seek to respond a user's shopping cart operation request immediately, and if extraordinarily long latencies have been detected, it's a good indication that the e-commerce cloud is overloaded.

Motivated by the above observation, we have developed a probing program, which mimics a user's shopping cart operations by constructing the corresponding URLs and JSON data as demonstrated in Table 1 to add items, remove items, and query shopping carts on Tmall and JD. Fifteen buyer accounts on each website from volunteers are employed in our study. In each probe, after connecting to a VPN server, the probing program first clears the shopping cart, then it adds a number of random items into the cart, after the clearing and adding, the probing program queries the shopping cart and measures the latency between the time t_i of issuing the query request and the time t_r that the response is received, as

$$l_{query} = t_r - t_i$$

Note that l_{query} includes the round trip time (RTT) rtt_1 between our measurement computer and the VPN server, the RTT rtt_2 between the VPN server and the front-end server, and the time l_{cloud} for the e-commerce cloud to process the query request and return the response. Therefore, after each probe, we measure $(rtt_1 + rtt_2)$ by pinging the front-end server through the VPN server, and calculate the cloud's query response latency l_{cloud} as

$$l_{cloud} = l_{query} - (rtt_1 + rtt_2) \tag{2}$$

5 CHARACTERIZING E-COMMERCE WORKLOAD

In this section, we analyze and characterize the e-commerce workload from out collected datasets.

5.1 General Usage Pattern

We first study the overall e-commerce traffics. Fig. 4(a) presents the Tmall and JD flows in the WEEK dataset from Mar. 1 to Mar. 7, 2016 on an hourly basis, and in Fig. 4(b), we present the e-commerce flows in every 10 minutes on a normal weekday of Mar. 3. From the figures we can see that the e-commerce workloads exhibit a strong diurnal pattern, and there are more flows observed during weekdays than weekends. By examining Fig. 4(b), which presents the usages of the two e-commerce websites on a weekday, we observe that more e-commerce workloads happen during the working hours, suggesting that unlike online social network and video streaming websites, which are visited by users more frequently in their spare time (e.g., evenings, weekends) [30] [31], campus users are more likely to visit the e-commerce websites during breaks in their working hours.

In Fig. 4(c), we present the e-commerce flows of Tmall and JD in the 24 hours of 2015's Double-11 Shopping Day (Nov. 11, 2015). By comparing Fig. 4(c) with Fig. 4(b), one can see that for both e-commerce websites, there are much more traffics in the Double-11 Day than in a normal weekday. This is reasonable as many users visit Tmall and JD in the Double-11 Day because of the promotions. Fig. 4(c) also shows that the traffic peaks for both websites appear at 0:00, then drop quickly within the next a few hours, which conforms to the sales growth in Fig. 1. The observation can be explained with the fact that many campus users stay up late on the night of Nov. 10, so that they can rush to buy as soon as the discounted prices became effective at the very beginning of the shopping festival.

From Fig. 4, one can see that Tmall and JD's workloads fluctuate very similar to each other. In fact, the Pearson correlations between Tmall and JD's workloads in Fig. 4(a)-(c) are as high as 0.9871, 0.9792, and 0.9498 respectively. The high correlations suggest that the two websites provide very similar services, and are accessed by campus users in almost a same way.

We also compare the campus traffic data with the public available Double-11 sales data in [32], and find that the four sales peaks, which happen in 0:00∼1:00, 8:00∼11:00, 15:00∼17:00, and 20:00∼22:00, all have their corresponding traffic peaks in Fig. 4. By comparing the campus traffic with the general sales data, we can see that our captured traffic can present the general e-commerce behaviors very well.

5.2 PC and Smartphone Usage Patterns

In addition to web browsers, both Tmall and JD have developed smartphone applications to enable users to shop with their smartphones. We are interested to know whether users with different devices access e-commerce services differently.

The WEEK and D11 datasets captured by iProbe do not contain any user device information. However, as demonstrated in Table 1, when requesting some services, PC and smartphone users employ different methods and URLs,

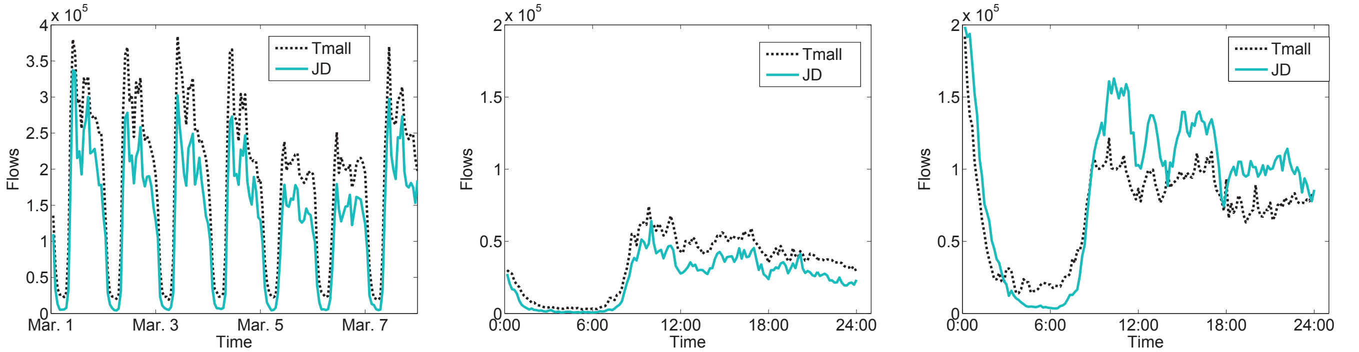

(a) Flows per hour in a week (b) Flows per 10 min. in a weekday (c) Flows per 10 min. in Double-11 Day

Fig. 4. Tmall and JD's traffics (a) in a week, (b) in a normal weekday, and (c) in the Double-11 Day.

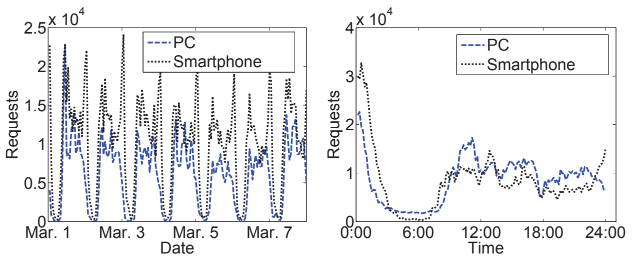

(a) Flows per hour in a week (b) Flows per 10 min. in Double 11 Day

Fig. 5. Three types of JD service requests initialized from PC and smartphone (a) in a week and (b) in the Double 11 Day.

which enable us to identify their devices. In our analysis, we select three JD e-commerce cloud services, namely the login service, the shopping cart service, and the checkout service, and differentiate the requests initialized from PC and smartphone respectively.

In Fig. 5(a), we present the PC and smartphone requests in the WEEK dataset. We can see that in general, the campus users prefer smartphone to PC, which conforms to the fact that smartphone replaces PC to become the major Internet access device in recent years. It is also observed that smartphones are heavily used in the evening and night, as the requests from smartphones keep increasing and reach to a peak around mid-night, while the PC usage declines since early evening.

We also compare the PC and smartphone usages on the Double 11 Day in Fig. 5(b). From the figure we can see that there are more PC usages during the regular hours (i.e., morning, afternoon, and evening) than smartphone, as it is convenient for users to use PC at their workplaces and dorms to access the e-commerce websites. But during the other time, smartphones are used more heavily. The most obvious discrepancy between PC and smartphone usages happens at late evening, as we can see that PC usage starts to decline after 22:00, while usage from smartphones begins to rise and reaches to the second highest peak of the day, suggesting that many users switch to smartphones to continue to access e-commerce websites at that time.

TABLE 2
Comparison of service requests in WEEK and D11 datasets.

	WEEK	**D11**
Browse an item	15,631	16,854
Shopping cart query/adding/removing	22,633	38,826
Checkout	1,214	2,015
Add an item to shopping cart	5,368	8,754
Items added / items browsed	0.343	0.519

5.3 Shopping Behavior

As in a conventional supermarket, there are a few steps for a user to complete a purchase on Tmall or JD. First, a user learns details of the products from their item pages; and if interested, he clicks the "add to shopping cart" button on the item page to add the item into his virtual shopping cart; the user can examine and organize his shopping cart by adding and removing items; finally, if the user decides to buy some items, he can select them from the shopping cart, click the "check out" button for the bill, and pay for them online.

In this subsection, we investigate the campus users' shopping behaviors, in particular, we focus on the item browsing, shopping cart query/adding/removing, and checkouts that correspond to the major online shopping steps. Since iProbe doesn't record the JSON data exchanged within a flow, we limit our scope to the requests issued from PC Web browsers on JD, as the requests can be identified with their URLs (Table 1).

In Table 2, we list the item browsing, shopping cart query/adding/removing, and checkout requests in the WEEK and D11 datasets. From the table one can see that for all the services, there are more requests in a single day of the Double-11 Day than in a week. As a user can only browse one item and add one item to his shopping cart per request, we also list the numbers of the shopping cart adding requests, and compute the ratios between the items added to shopping carts and the items browsed in the two datasets. We find that in the WEEK dataset, the ratio is 0.343, but in the D11 dataset, it is increased to 0.519, suggesting that users are more willing to buy in the Double-11 Shopping Day than usual.

We then study the item browsing activity. Fig. 6 presents the frequencies of the items browsed by campus users in the WEEK and D11 datasets. We find that in both datasets, the item popularity follows a Zipf-like distribution with

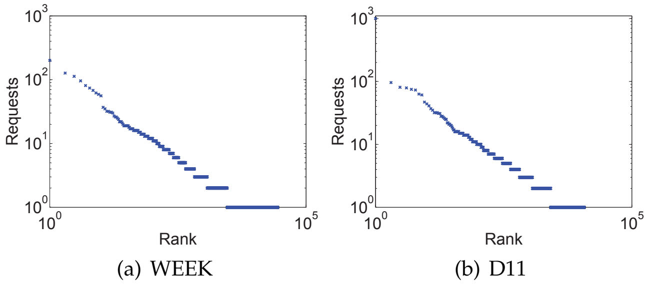

Fig. 6. Frequencies of items being browsed in (a) WEEK and (b) D11 datasets.

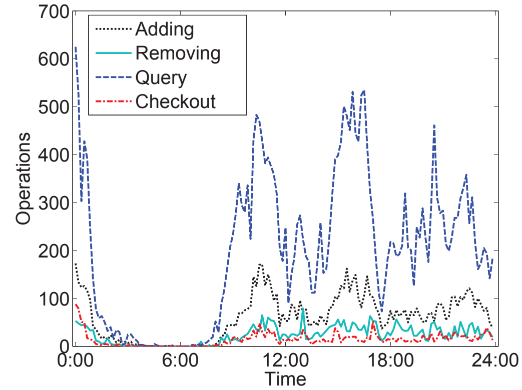

Fig. 7. Shopping cart adding, removing, query and checkout requests on JD per 10 min. in D11 dataset.

exponent about 0.6. We then examine the top-10 most browsed items, and find that in the WEEK dataset, seven of the top-10 most popular items are smartphones, and the other three are food/fruit, which fit campus users' consumption habits. On the other hand, the top-10 most popular items on the Double-11 Shopping Day are more diverse: besides smartphones, the 3^{rd} to 6^{th} most browsed items are high-end liquors, while such items are rarely browsed in the WEEK dataset. Further investigation shows that these liquors have big price reductions in the Shopping Festival. The discrepancy implies that the promotions of the Double-11 Shopping Day significantly influence the campus users' shopping behaviors.

In Fig. 7, we decompose the shopping cart operations, namely, adding, removing, query, and checkout, and present numbers of each type of the requests in every 10 minutes in the D11 dataset. From the figure we find that as many as 66.9% of the shopping cart operations are queries, while the checkouts, which are critical to both users and e-commerce providers, account for only 4.7% of the total requests. We explain the frequent shopping cart queries with two reasons: First, e-commerce users frequently compare new items they are browsing with the ones they have added in their shopping carts; Second, they frequently check whether the items in their shopping carts are still available to buy, as in both Tmall and JD, when an item is sold out, it will be marked as "unavailable" in the cart.

By further examining Fig. 7, we find that 8.8% and 14.9% of the shopping cart adding and checkout operations happen within the first hour (i.e., 0:00$\sim$1:00) of the Double-11 Day. The observation suggests that a lot of users make their purchase decisions at the very beginning of the shopping festival, which also explains the rapid sales growth in Fig. 1. Clearly, such a rush buying behavior will impose a flash crowd of service requests on the e-commerce systems.

5.4 Summary

We summarize our results on characterizing the e-commerce workload as the following:

- *Workload characteristic*: Campus users are more likely to visit the e-commerce websites in their work breaks; PC and smartphone users access e-commerce services differently; and there exists a Zipf-like popularity among the shopping items.
- *Massive flash crowd*: Promotions of the Double-11 Shopping Day attracts many users, who have more diverse interests and are more willing to buy than usual; most of the shopping cart operations are read-only queries; and the rush buying behavior at the very beginning of the shopping festival impose a massive workload on the e-commerce systems.

6 Evaluating E-Commerce Infrastructure

In this section, we examine Tmall and JD's e-commerce infrastructures, including the CDNs and the clouds, and evaluate their performances under various circumstances, in particular, the massive flash crowd in the Double-11 Shopping Day.

6.1 E-Commerce Infrastructure Discovery

We first discover Tmall and JD's e-commerce infrastructures with active measurements. In our study, we employ 241 VPN servers from "517VPN", which are distributed in 152 cities of 27 provinces and 40 autonomous systems (ASes) in China, as our vantage points. For discovering the content and front-end servers that are assigned to each vantage point, we connect to the remote VPN servers, resolve the domain names corresponding to various e-commerce services, and collect the IP addresses returned from the vantage points' LDNS servers.

We perform the server discovery on Mar. 3, 2016, and have collected 101 and 24 image content server IPv4 addresses for Tmall and JD respectively. Although JD has fewer content server addresses, we find that all its addresses are gateway addresses (i.e., the addresses in the form of `a.b.c.1`), suggesting that these addresses may bind to reverse proxies, each representing a cluster of real content servers.

We then geo-locate the content servers with the methodology proposed in [28], and looks up the ASes they belong to [33]. We cluster the content servers that are in a same AS and a same city as a *CDN node*. As a result, for Tmall we have clustered 52 CDN nodes, which are in 40 cities and 21 ASes; while for JD, the content servers are clustered into 21 nodes distributed in 16 cities and 14 ASes.

We also collect and geo-locate the front-end servers. For JD, we find that all the front-end servers are located in Beijing. While for Tmall's shopping cart and checkout services,

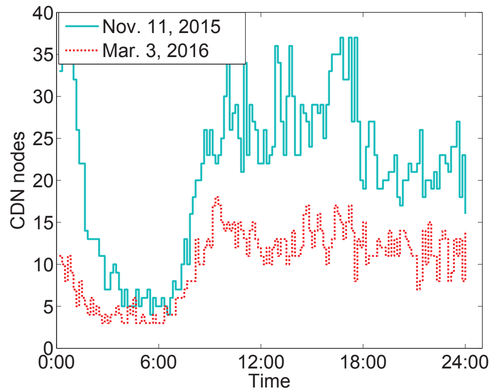

Fig. 8. JD CDN nodes encountered by campus users in every 10 minutes in 2015's Double-11 Day and a normal weekday of Mar. 3, 2016.

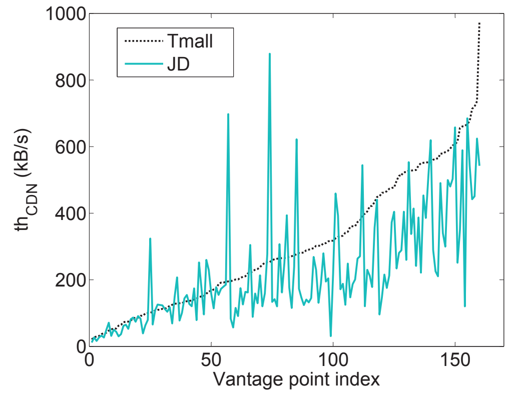

Fig. 10. Tmall and JD's CDN throughputs th_{CDN} collected from 160 vantage points in a normal weekday.

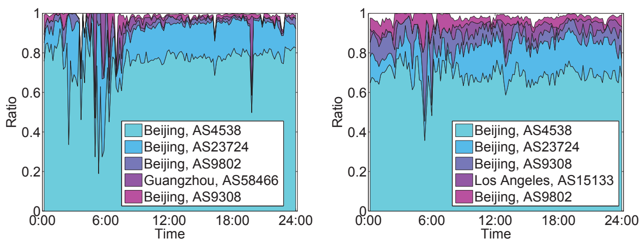

(a) Normal weekday (b) Double-11 Day

Fig. 9. Percentages of the image traffics from the top-5 CDN nodes in (a) a normal weekday and (b) the Double-11 Day.

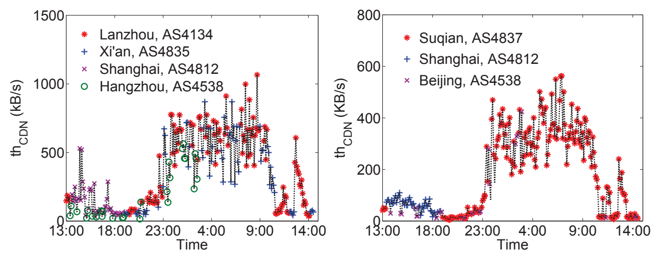

(a) Tmall (b) JD

Fig. 11. (a) Tmall and (b) JD's CDN throughputs th_{CDN} collected from Lanzhou, AS4134 in a normal weekday for 24 hours.

the front-end servers are co-locate with the content servers, but for the other services, all the front-end servers are in AS37963, the AS in Hangzhou operated by Alibaba. The geo-locations of the front-end servers suggest that, Tmall's e-commerce cloud data center is at Hangzhou, while JD's data center is located in Beijing.

Note that we do not claim to have discovered all the infrastructures for Tmall and JD, but consider our findings as a *snapshot* of their e-commerce systems. In fact, e-commerce infrastructure evolves over time, as we will see in the following subsection, for accommodating the massive service requests in the Double-11 Day, JD had greatly expanded its e-commerce CDN footprint, while Tmall announces that it operated two cloud data centers instead of one during 2015's Double-11 Shopping Day [34].

6.2 Evaluating E-Commerce CDN

6.2.1 Passive Measurement

We employ the WEEK and D11 datasets to identify all the CDN nodes of JD that satisfy campus users' image requests. Fig. 8 compares the CDN nodes that we have encountered in every 10 minutes, in 2015's Double-11 Shopping Day and an ordinary weekday of Mar. 3, 2016, for 24 hours. We can see that more CDN nodes are employed in the Double-11 Day than usual. Moreover, the CDN nodes we have encountered in the D11 dataset greatly exceed the total number of the nodes that we have discovered in Section 6.1, suggesting that JD has greatly expanded its e-commerce CDN footprint, for accommodating the explosive growth of the content requests in the shopping festival.

We then identify the top-5 CDN nodes that contribute most image traffics in Nov. 11, 2015 and Mar. 3, 2016, and present the ratio of the traffics from each node in 10-minute intervals in Fig. 9 respectively. From the figures we can see that in both days, about 70% of the traffics are from the CDN node at Beijing, AS4538, probably because the node is in a same AS as our campus network. However, it is interesting to find that during the Double-11 Day, 4.2% of the image traffics are from the node Los Angeles, AS15133, and the AS belongs to EdgeCast, a commercial CDN in North America. The observation implies that unlike the distance and latency-based mappings that are observed in previous studies [35] [36] [37], during the Double-11 Day, JD's e-commerce CDN adopts a *capacity-based mapping* by redirecting clients to the CDN nodes that have spare capacities. Under such policy, spare oversea servers are scheduled to satisfy the domestic content requests.

6.2.2 Active Measurement

CDN performance in usual time

We employ the methodology as described in Section 4.2 to actively probe Tmall and JD's e-commerce CDNs. From each VPN server vantage point, we randomly select five high-resolution images from our collection, and use Equation (1) to estimate the throughput th_{CDN} between the VPN server and the CDN content server. The measurement was performed in the afternoon of Mar. 3, 2016.

In Fig. 10, we present Tmall and JD's CDN throughputs collected from 160 VPN server vantage points. The VPN servers are distributed in 104 cities and 32 ASes in China, and from each VPN server, we have $th_{cURL} \ll th'_{cURL}$

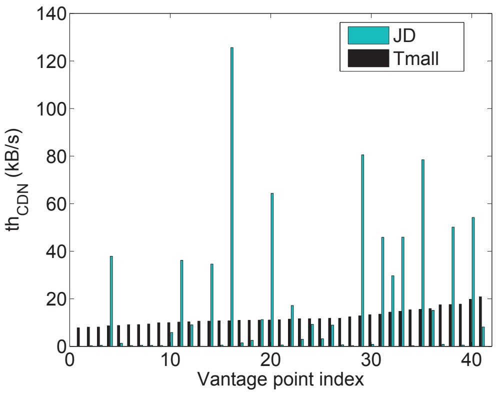

Fig. 12. Tmall and JD's CDN throughputs th_{CDN} collected from 41 vantage points under the Double-11 flash crowd.

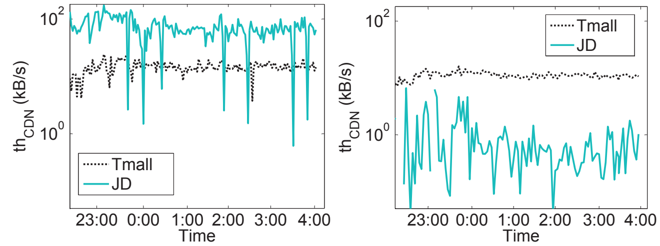

(a) Shanghai, AS4134 (b) Mianyang, AS38283

Fig. 13. Tmall and JD's CDN throughputs th_{CDN} collected from (a) Shanghai, AS4134, and (b) Mianyang, AS38283, during the Double-11 flash crowd.

for both e-commerce websites, ensuring that the collected throughput is indeed the throughput between the VPN server and the content server, i.e., $th_{CDN} = th_{cURL}$. From the figure we can see that both CDNs have decent performances: Tmall achieves an averaged CDN throughput of $302.14\,\mathrm{kB/s}$ and JD has $223.87\,\mathrm{kB/s}$. We also find that on many vantage points, Tmall has a higher throughput than JD, probably because it has more CDN nodes that are more proximate to our vantage points, as observed in Section 6.1.

We then select one VPN server, which is located in the city of Lanzhou, AS4134, and continuously probe the e-commerce CDNs for 24 hours starting from 13:00, Mar. 3, 2016. The results are presented in Fig. 11. Note that in the figures we differentiate the CDN nodes that are assigned to our vantage point, and mark them with different labels. From the figures we can see that the CDN throughputs th_{CDN} vary significantly within the 24 hours, and exhibit a strong diurnal pattern that is opposite to the workload pattern in Fig. 4. In addition, although the CDNs assign different content servers to our vantage point in different hours, however, we believe that the bandwidth competition at the content server is the major reason of the throughput variances, as one can see from the figure, the throughputs from a same CDN node also vary significantly in the peak and off-peak hours.

CDN performance under Double-11 flash crowd

We probe Tmall and JD's e-commerce CDNs since 0:00, Nov. 11, 2015, the very beginning of 2015's Double-11 Shopping Day. We employ 41 VPN servers, which are distributed in 37 cities and 13 ASes in China, as our vantage points. Fig. 12 presents the throughputs th_{CDN} of the two CDNs that we have collected from different vantage points. Comparing with Fig. 10, we can see that both CDNs have much lower throughputs, suggesting that the content servers were overwhelmed by the massive content requests. Moreover, we find that the performance of Tmall's CDN is much more stable than JD, as the throughputs from different Tmall CDN nodes range between $8{\sim}20\,\mathrm{kB/s}$, which is quite stable; on the other hand, the JD CDN's throughputs vary greatly from a few $\mathrm{B/s}$ to hundreds of $\mathrm{kB/s}$, and from 16 of the 41 vantage points, the measured throughputs are lower than $1\,\mathrm{kB/s}$. The observation suggests that during the peak hours of the Double-11 Shopping Day, Tmall enforces a *proactive per-flow bandwidth throttling* on its content servers, so as to guarantee a low but stable throughout for ensuring the service availability; on the other hand, JD still follows the best-effort way to provide content service under the Double-11 flash crowd.

To better justify our point, we select two vantage points, one is in Shanghai, AS4134 and the other is in the city of Mianyang, AS38283. From the two vantage points we continuously probe Tmall and JD's e-commerce CDNs every 30 seconds, from 22:00, Nov. 10 to 4:00, Nov. 11, 2015, which covers the Double-11 flash crowd. We present the results in Fig. 13. Note that the y-axle is in log scale. From the figures we can see that on the two vantage points, the throughputs th_{CDN} of Tmall's e-commerce CDN are quite stable between $10{\sim}20\,\mathrm{kB/s}$, confirming that Tmall has enforced a proactive bandwidth throttling to provide low but guaranteed throughput. On the other hand, the JD CDN's throughputs th_{CDN} observed from the two vantage points differ greatly: from the Shanghai vantage point, JD has a much higher throughput than Tmall, as there are a few JD CDN nodes located nearby; while from Mianyang, which is far away from most of JD's CDN nodes, the throughput we have observed barely exceeds $1\,\mathrm{kB/s}$ after 0:00. The observation confirms our analysis that under the Double-11 flash crowd, Tmall enforces a proactive bandwidth throttling on its CDN to provide a guaranteed service, while JD still provides the service in a best-effort way.

6.3 Evaluating E-Commerce Cloud

In this section, we evaluate Tmall and JD's e-commerce clouds. We focus on the shopping cart service, which is essentially a cloud database service, and apply Equation (2) to estimate the latency l_{cloud} between the time that the e-commerce cloud receives a shopping cart query request and it returns the response through the front-end server. Since the shopping cart service is a real-time service, when extraordinarily long latencies have been detected, it is a good indication that the e-commerce cloud is overloaded.

Cloud performance in usual time

In our first measurement study, we connect to two VPN servers, which are located in Shanghai, AS4812 and Wuhan, AS4134, as our vantage points. We continuously probe Tmall and JD's shopping cart services every 5 minutes for 24 hours from Mar. 10 to Mar. 11, 2016, and plot the query response latencies l_{cloud} in Fig. 14. From the figures we can

TABLE 3
Summaries of the shopping cart query response latencies.

	Tmall		JD	
Vantage point	Shanghai, AS4812	Wuhan, AS4134	Shanghai, AS4812	Wuhan, AS4134
Baseline latency l_{base}(ms)	178.46	174.21	108.11	94.48
Mean/std. latencies of l_{cloud} in Double-11 Day (ms)	484.96/482.16	440.67/342.45	600.38/510.85	611.63/440.37

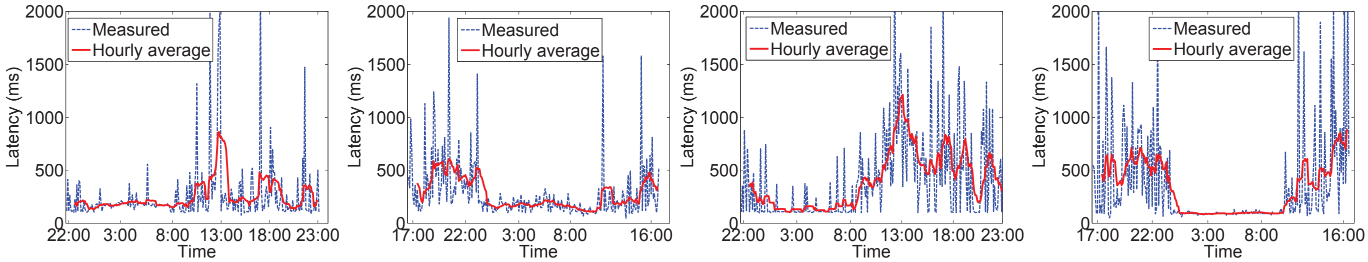

(a) Tmall, Shanghai (b) Tmall, Wuhan (c) JD, Shanghai (d) JD, Wuhan

Fig. 14. Shopping cart query response latencies l_{cloud} on Tmall and JD's e-commerce clouds and their hourly averages collected from Shanghai, AS4812 and Wuhan, AS4134 in a normal weekday.

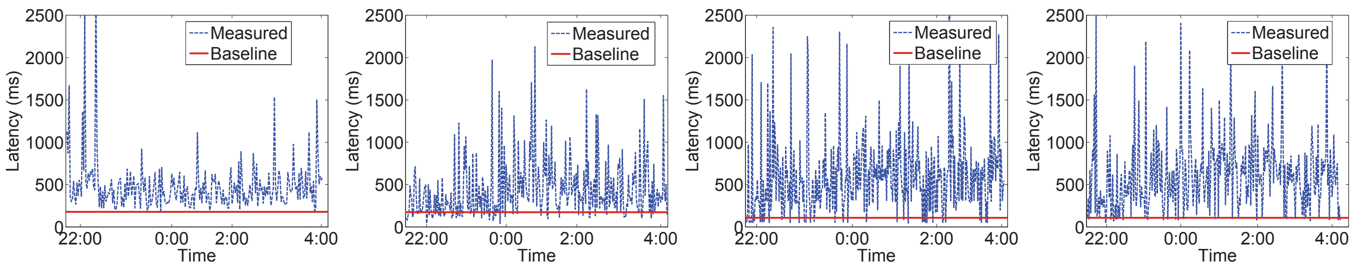

(a) Tmall, Shanghai (b) Tmall, Wuhan (c) JD, Shanghai (d) JD, Wuhan

Fig. 15. Shopping cart query response latencies l_{cloud} on Tmall and JD's e-commerce clouds collected from Shanghai, AS4812 and Wuhan, AS4134 during the Double-11 flash crowd.

see that in the night and early morning hours, both clouds have small and stable latencies, but during the daytime and evening hours, a lot of long response latencies are detected. We also plot the hourly averages of l_{cloud} in the figures for better demonstrating the differences. Our observation suggests that, both Tmall and JD's e-commerce clouds are overloaded during the busy hours, and the cloud database systems queue our shopping cart queries without processing them immediately, resulting in the long latencies; on the other hand, in the off-peak hours, the cloud system can process most of our probing queries without delay, so we have short and stable latencies as in Fig. 14. In the following, we refer to the latency required for the e-commerce cloud to respond the query when it is not overloaded as a *baseline latency* of the shopping cart service, denoted as l_{base}.

To estimate l_{base}, we employ the measurement results collected from the two vantage points between 0:00 and 7:00, filter out the outliers by removing a few samples that are over twice of the median, and calculate the mean as the baseline latency. Table 3 lists the baseline latencies for Tmall and JD from the two vantage points, we can see that for each e-commerce website, the estimations from different vantage points are very close, suggesting that they are estimated correctly.

Cloud performance under Double-11 flash crowd

We repeat the measurement experiments on Tmall and JD's shopping cart services between 22:00 Nov. 10 and 4:00 Nov. 11, 2015, which covers the Double-11 flash crowd. Fig. 15 presents the latencies l_{cloud} collected from the Shanghai and Wuhan vantage points, and compares them with the baseline ones. From the figures we can see that, under the Double-11 flash crowd, a lot of extraordinarily long latencies have been detected, and the latencies vary significantly. Table 3 lists the mean shopping cart query latencies and its standard variances, which are of several times longer than l_{base}.

Our observation suggests that both Tmall and JD have difficulties in guaranteeing the quality in their e-commerce cloud services, as it is very common for users to suffer extraordinarily long transaction latencies. We believe that the problem arises because the private e-commerce clouds of Tmall and JD do not have sufficient capacities, especially during the peak hours of the Double-11 Day.

6.4 Summary

We summarize our major results on evaluating Tmall and JD's e-commerce infrastructures as the following:

- *E-commerce CDN*: Both Tmall and JD's e-commerce CDNs have decent throughputs, but the performances degrade significantly under the Double-11 flash crowd, despite that the providers have expanded the e-commerce CDN footprint and employed a capacity-based client mapping policy; Tmall's e-commerce CDN adopts a proactive per-flow bandwidth throttling to provide low but guaranteed throughput, while JD still provides content service in a best-effort way.

- *E-commerce cloud*: Both Tmall and JD's private e-commerce clouds do not have sufficient capacities to guarantee the service quality in busy hours, and users suffer extraordinarily long latencies in their e-commerce transactions.

7 Implications

In this section, we discuss the implications from observations of our measurement study.We analyze the e-commerce CDN design choices, and show that utilizing massive data centers and employing DNS-based redirection are preferred over their counterparts. We also investigate the benefits of incorporating client-side assistances in offloading workloads from the e-commerce CDN and cloud under flash crowd.

7.1 E-Commerce CDN Design Choices

7.1.1 Massive data centers vs. numerous small clusters

There have been a long-time debate on the CDN architecture design [15]. One CDN design philosophy is to enter deep into the ISPs, with Akamai as a representative example [8] [38]. In such a CDN, content servers are deployed at point-of-presences (PoPs) in as many ISPs as possible. By placing content server clusters deep into the ISP networks, CDN can deliver contents from the servers that are proximate to users, thus provides good content delivery performances regarding latency and throughput. The other design philosophy, with Limelight [8] and Google [39] as examples, is to bring ISPs to home. That is, a CDN build limited number of massive data centers at a few key locations that are close to many transit and eyeball ISPs' PoPs, and connects to them at these locations. The CDN design philosophy is motivated from the recent observations that majority of inter-domain traffics on today's Internet flow directly between large content providers and consumer networks, and the global Internet becomes flatter with ISP networks peering with each other much more densely, due to the reasons such as the emerging Internet exchange points (IXPs) [40].

Both CDN designs rely on a mapping system to assign clients to CDN nodes for serving their content requests. As we have seen in Section 6.2, an e-commerce CDN under massive flash crowd employs a capacity-based mapping policy, which assigns clients to CDN nodes that have spare capacities. Such a mapping policy requires that, the global load-balancer of the CDN must accurately estimate real-time capacity of each CDN node, so as to make the "right" assignments. Clearly, when there are errors in the estimations, the CDN will either overload some CDN nodes, or leave some other nodes under-utilized.

With the above consideration, we argue that for an e-commerce CDN under flash crowd, the design choice with massive data centers is preferred. Note that the benefit brought by the CDN design with numerous clusters, that is, to serve clients from proximate servers, no longer exists under the capacity-based mapping, as under such a policy, it is common to assign clients to distant servers with spare capacities. More importantly, with a small number of massive data centers, the influence of capacity estimation errors can be considerably reduced, comparing with estimating capacities of numerous small clusters.

We demonstrate the benefit of massive data centers with a simple analysis. Suppose a CDN has n clusters of content servers deployed in a geographical region, with the i^{th} cluster having a spare capacity of $c_i(t)$ at time t. The CDN makes an capacity estimation as $c'_i(t) = c_i(t) \pm \Delta c_i(t)$, where $\pm\Delta c_i(t)$ is the estimation error. For all the n clusters, the sum of the estimation errors should be

$$\Delta C(t) = \sum_i^n \Delta c_i(t)$$

However, when the n small clusters are "merged" into one massive data center, given the same estimation errors for the clusters as $\pm\Delta c_i(t)$, the aggregated error for the merged data enter becomes

$$\Delta C'(t) = \sqrt{\sum_i^n (\Delta c_i(t))^2}$$

Mathematically, $\Delta C'(t)$ is smaller than $\Delta C(t)$ as $n > 1$. For example, if $\Delta c_i(t) = \Delta c(t)$ for all the clusters, we have

$$\Delta C'(t) = \frac{\Delta C(t)}{\sqrt{n}}$$

In other words, the error for estimating n small clusters is $\sqrt{n}$ times of the error for estimating one massive data center of a same overall capacity. From the analysis, we can see that comparing with a CDN of numerous small clusters, a CDN that employs a few massive data centers is more robust against capacity estimation errors, thus can assign clients to nodes more accurately according to their capacities.

7.1.2 DNS-based redirection vs. anycast redirection

Another CDN design choice is how to redirect clients to content servers. CDNs such as Akamai employ the DNS-based redirection by resolving CNAMEs to IP addresses of content servers based on its mapping policy (e.g., latency-based or capacity-based mapping) [41]. Some other CDNs such as Bing apply the anycast redirection [37]. Anycast is a routing strategy where same IP address is announced from many locations throughout the Internet. Then BGP routes a client to one location with the best BGP path to the client.

Obviously for an e-commerce CDN that adopts the capacity-based mapping policy, it must use the DNS-based redirection, so as to dynamically change its mapping decisions based on the real-time capacities of the CDN nodes. If the CDN uses the anycast redirection, clients will always be redirected to the nodes of the best BGP path, regardless the nodes' capacities. For example, with anycast redirection, a CDN will never redirect domestic content requests to an oversea CDN node, as we have observed on JD's CDN in Section 6.2, since the oversea node has an obvious longer BGP path.

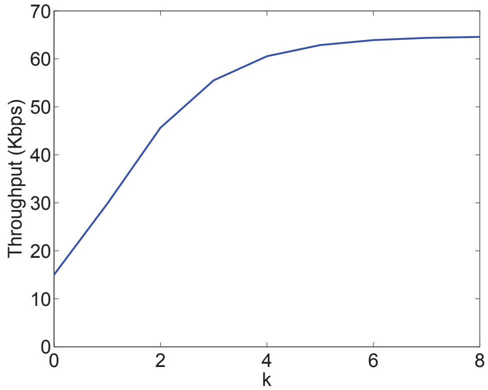

Fig. 16. CDN throughput per request achieved with peer-assisted e-commerce content distribution, under various list length of k.

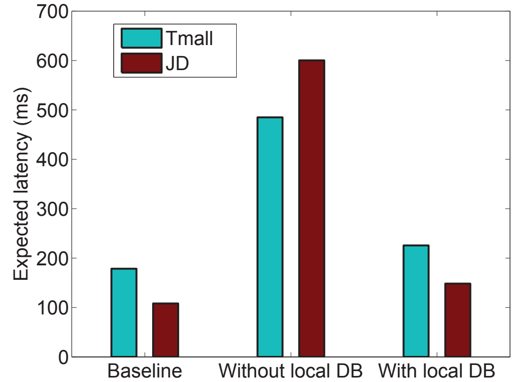

Fig. 17. Comparison of the baseline latency and the response latencies with and without local database for shopping cart cloud service.

7.2 Client-side Assistance

As observed in our measurement study, both the e-commerce CDN and cloud do not have sufficient capacities for serving the massive workloads in the Double-11 Day, leading to poor service performances. In this subsection, we investigate the benefits brought by incorporating assistances from clients. More specifically, we consider to use peers to help to distribute the static e-commerce contents, and use local database to handle the read-only shopping cart queries. We analytically show that by involving the client-side assistances, considerable workloads on the e-commerce CDN and cloud can be offloaded.

7.2.1 Offloading e-commerce CDN workload

We suggest that during massive flash crowd, each e-commerce client caches large-sized e-commerce contents such as high-resolution images that it has ever downloaded in its local storage, and uses the cached replicas to serve subsequent requests from other clients. An e-commerce CDN's content server works as a tracker in the peer-assisted content distribution by keeping a list L of up to k recent clients that have downloaded and cached the content. When a new request for the same content arrives at the server, the server returns L, and the requesting client contacts the peers in the list for the content.

With simple analysis, we can derive that with the peer assistances, the fraction of the content requests that are served by peers can be expressed as

$$r_{CDN} = \sum_{i=1}^{M} p_i \times (1 - (1-q)^{n_i}) \tag{3}$$

In the above equation, M is the total number of contents, p_i is the probability of the i^{th} popular content being requested as shown in Fig. 6, q is the probability that a client is unable to serve for reasons such as NAT or becoming offline, and n_i is the expected number of the peers that can serve the i^{th} popular content, calculated as

$$n_i = \min\{k, N \times (1 - (1-p_i)^s\}$$

where N is the number of the clients covered by the tracker server and s is the local cache size.

We show the effectiveness of the peer-assist e-commerce content distribution with numerical results. We assume that without peer assistance, each content request has a $15\,\mathrm{kB/s}$ throughput, as we have observed on Tmall's CDN in Section 6.2.1. With a faction r_{CDN} of content requests as in Equation (3) being offloaded, the content server's bandwidth is shared among the remaining $(1 - r_{CDN})$ fraction of the requests. In Fig. 16, we show the throughput per request under various list lengths k. With $M = 10,000$, $N = 1,000$, $s = 20$, and $q = 0.5$. From the figure we can see that by tracing no more than 5 recent requesting clients, the e-commerce CDN's throughput for each request can be considerably improved to over $60\,\mathrm{kB/s}$.

7.2.2 Offloading e-commerce cloud workload

We then consider improving the e-commerce cloud services, in particular, the shopping cart service. From Fig. 7 one can see that among all the shopping cart service requests, 66.9% are read-only queries, which do not incur any update on users' shopping cart data, or the provider's inventory and sales database in the e-commerce cloud.

Motivated by the observation, we suggest that e-commerce client maintains a local database for the shopping cart service: when a user issues a shopping cart adding, removing, or checkout service request, the request is handled by the cloud database, and local database synchronizes with the cloud database immediately; however, local database directly handles a user's shopping cart query requests without involving the cloud. The e-commerce cloud also proactively updates a client's local database when an item in the user's shopping cart is sold out, or its price has been changed. Note that comparing with shopping cart queries, such incidents happen rarely.

As in previous work [22], we model an e-commerce cloud as an $M/M/1$ queueing system, whose expected service response latency can be expressed as

$$latency = \frac{1}{\mu - \lambda} \tag{4}$$

where μ is the service rate of the e-commerce cloud, and λ is the arrival rate of the service requests.

We employ our measurement results in Section 6.3 to estimate the $M/M/1$ parameters. For estimating the service rate μ, we use the baseline latency in Table 3 as the time that $M/M/1$ serves a request without queueing, i.e., $\frac{1}{\mu} = l_{base}$. For estimating λ, we use the mean latency measured under the Double-11 flash crowd in Table 3 (denoted as l_{d11}), and let $\frac{1}{\lambda+\mu} = l_{d11}$. Note that λ is the rate of all types of shopping cart operations, including adding, removing, query, and checkout. After applying the local database, since the shopping cart queries are handled locally, the new service request arrival rate becomes $\lambda' \approx 0.331 \times \lambda$, and the cloud's service response latency is reduced to

$$latency' = \frac{1}{\mu - \lambda'} \tag{5}$$

We estimate the $M/M/1$ parameters with the measurement results collected from Shanghai, AS4812, and compare the service response latencies for Tmall and JD, with and without applying the local database strategy respectively in Fig. 17. We also present the baseline latencies for comparison. From the figure we can see that by employing a local database to handle the read-only queries, service response latency can be significantly reduced.

We recognize that both the peer-assisted content distribution and the local database introduce additional complexities and management overheads. For example, for providing content service with peer-cached replicas, the e-commerce system must ensure the integrity of the cached replica against malicious content pollution, and preserve users' privacies. While for applying the local-database strategy, the e-commerce cloud must trace user's device for accessing the e-commerce cloud service, and proactively synchronizes local databases with the cloud database when user switches his access device, or uses multiple devices simultaneously. However, given the poor service qualities as we have observed during the Double-11 Shopping Day, which potentially causes huge income loss for e-commerce provider [3], we believe that such complexities and overheads are worthwhile.

8 Conclusion

In this paper, we present a measurement study on the e-commerce systems of Tmall and JD, the top-two most popular e-commerce websites in China. We analyze the e-commerce workloads that are collected from our campus network, and investigate the behaviors and performances of the e-commerce infrastructures, including the e-commerce CDNs and clouds, with passive and active measurements. In particular, we analyze the massive flash crowd of e-commerce workload in 2015's Double-11 Shopping Day, which is the biggest online shopping festival in the world, and evaluate the e-commerce systems' performances during the Double-11 Day's peak hours.

We discuss implications of the observations from our measurement study. We analyze that the design choices of utilizing massive data centers and DNS-based redirection are preferred for e-commerce CDNs, and by involving client-side assistances, considerable e-commerce workloads on both CDN and cloud can be potentially offloaded. Our study provides insights on the e-commerce workload and infrastructure, and is valuable in the design and management of large-scale e-commerce systems.

For the future work, we seek to collect e-commerce traffics from more locations by deploying iProbe at a few ISP point-of-presences (PoPs), and analyze usage patterns and shopping behaviors of residential users. We also seek to investigate logs of content and front-end servers from a commerce CDN that carries e-commerce traffics, infer the e-commerce service qualities, and identify the bottlenecks.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China under Grants 61672486, 61379130, and 61672485, in part by the Key Project of the New Generation Broadband Mobile Wireless Communication Network under Grant 2017ZX03001019-004, and in part by the Anhui Provincial Natural Science Foundation under Grant 1608085MF126.

The authors would like to thank Mr. Huanjie Zhang from the USTC Network Information Center for the assistance in the network traffic measurement.

References

[1] Statista, "Global retail e-commerce sales 2014-2019," http://www.statista.com/statistics/379046/worldwide-retail-e-commerce-sales/.

[2] H. Ma, "E-commerce sales reaches to 11% of the total retail sales (in Chinese)," May 2016, http://ep.ycwb.com/epaper/ywdf/html/2016-05/17/content_40900.htm.

[3] CDNetworks, "How much is each second worth in ecommerce? check out these 15 states!" Aug. 2013, https://www.cdnetworks.com/blog/how-much-is-each-second-worth-in-ecommerce-check-out-these-15-stats/.

[4] Tmall.com, https://www.tmall.com/.

[5] JD.com, http://www.jd.com/.

[6] P. Wendell and M. J. Freedman, "Going viral: Flash crowds in an open CDN," in *Proc. of Internet Measurement Conference (IMC'11)*, Berlin, Germany, Nov. 2011.

[7] "Singles' Day," https://en.wikipedia.org/wiki/Singles'_Day.

[8] C. Huang, Y. A. W. J. Li, and K. W. Ross, "Measuring and evaluating large-scale CDNs," Microsoft Research, Tech. Rep. MSR-TR-2008-106, 2008.

[9] V. K. Adhikari, Y. Guo, F. Hao, V. Hilt, and Z.-L. Zhang, "Measurement study of Netflix, Hulu, and a tale of three CDNs," *IEEE/ACM Trans. on Networking*, vol. 23, no. 6, pp. 1984–1997, 2015.

[10] C. Lo, D. Frankowski, and J. Leskovec, "Understanding behaviors that lead to purchasing: A case study of Pinterest," in *Proc. of KDD'16*, San Francisco, CA, USA, Aug. 2016.

[11] H. Xu, D. Liu, H. Wang, and A. Stavrou, "E-commerce reputation manipulation: The emergence of reputation-escalation-as-a-service," in *Proc. of WWW'15*, Florence, Italy, May 2015.

[12] S. Triukose, Z. Wen, and M. Rabinovich, "Measuring a commercial content delivery network," in *Proc. of WWW'11*, Hyderabad, India, Mar. 2011.

[13] V. K. Adhikari, S. Jain, Y. Chen, and Z.-L. Zhang, "Vivisecting YouTube: An active measurement study," in *Proc. of IEEE INFOCOM'12*, Orlando, FL, USA, Mar. 2012.

[14] H. He, Y. Zhao, J. Wu, and Y. Tian, "Cost-aware capacity provisioning for internet video streaming CDNs," *Computer Journal*, vol. 58, no. 12, pp. 3255–3270, 2015.

[15] Y. Tian, M. Zhao, and X. Zhang, *Internet Video Data Streaming: Energy-saving and Cost-aware Methods*. Singapore: Springer, 2017.

[16] A. Li, X. Yang, S. Kandula, and M. Zhang, "CloudCmp: Comparing public cloud providers," in *Proc. of Internet Measurement Conference (IMC'10)*, Melbourne, Australia, Nov. 2010.

[17] N. Nasiriani, C. Wang, G. Kesidis, and B. Urgaonkar, "A measurement-based study of effective capacity dynamism on Amazon EC2," Pennsylvania State University, Tech. Rep. CSE-16-005, 2016.

[18] I. Bermudez, S. Traverso, M. Mellia, and M. Munafo, "Exploring the cloud from passive measurements: the Amazon AWS case," in *Proc. of IEEE INFOCOM'13*, Apr. 2013.

[19] Y. Chen, S. Jain, V. K. Adhikari, and Z.-L. Zhang, "Characterizing roles of front-end servers in end-to-end performance of dynamic content distribution," in *Proc. of Internet Measurement Conference (IMC'11)*, Berlin, Germany, Nov. 2011.
[20] Z. Ren, W. Shi, J. Wan, F. Cao, and J. Lin, "Realistic and scalable benchmarking cloud file systems: Practices and lessons from AliCloud," *IEEE Trans. on Parallel and Distributed Systems*, 2017, accepted.
[21] H. Wang, F. Wang, J. Liu, and J. Groen, "Measurement and utilization of customer-provided resources for cloud computing," in *Proc. of IEEE INFOCOM'12*, Orlando, FL, USA, Mar. 2012.
[22] Y. Niu, B. Luo, F. Liu, J. Liu, and B. Li, "When hybrid cloud meets flash crowd: Towards cost-effective service provisioning," in *Proc. of IEEE INFOCOM'15*, Hong Kong, China, Apr. 2015.
[23] S. Singh, I. Chana, and R. Buyya, "STAR: SLA-aware autonomic management of cloud resources," *IEEE Trans. on Cloud Computing*, 2017, accepted.
[24] M. Amiri and L. Mohammad-Khanli, "Survey on prediction models of applications for resources provisioning in cloud," *J. of Network and Computer Applications*, vol. 82, pp. 93–113, 2017.
[25] R. Fu, "Tmall sales over $1 billion in 7 mins on Double 11 2015," Nov. 2015, http://www.chinainternetwatch. com/15497/double-11-2015/.
[26] M. Kaplan, "Sales report: 2015 Thanksgiving Day, Black Friday, Cyber Monday," Dec. 2015, http://www.practicalecommerce.com/articles/94777-Sales-Report-2015-Thanksgiving-Day-Black-Friday-Cyber-Monday/.
[27] "iProbe white paper," http://www.wxjuyun.com/ service/ziyuanzhongxin /baipishu/2015/1006/25.html.
[28] Y. Tian, R. Dey, Y. Liu, and K. W. Ross, "Topology mapping and geolocating for China's Internet," *IEEE Trans. on Parallel and Distributed Systems*, vol. 24, no. 9, pp. 1908–1917, 2013.
[29] "cURL," https://curl.haxx.se/.
[30] M. Kihl, R. Larsson, N. Unnervik, and et al., "Analysis of FaceBook content demand patterns," in *Proc. of International Conference on Smart Communications in Network Technologies (SaCoNeT'14)*, Vilanova i la Geltru, Spain, Jun. 2014.
[31] M. Zink, K. Suh, Y. Gu, and J. Kurose, "Characteristics of YouTube network traffic at a campus network – measurements, models, and implications," *Computer Networks*, vol. 53, no. 4, pp. 501–514, 2009.
[32] Caixin, "Big data in double 11 (in Chinese)," Nov. 2015, http://m.companies.caixin.com/m/2015-11-13/100873929.html.
[33] Cymru, "Team Cymru IP to ASN lookup v1.0," https://asn.cymru.com/.
[34] Xinhua Finance, "Alibaba launches Hangzhou + Beijing double centers strategy," Sep. 2015, http://en.xinfinance.com/html/Companies/2015/146145.shtml.
[35] R. Krishnan, H. V. Madhyastha, S. Srinivasan, S. Jain, A. Krishnamurthy, T. Anderson, and J. Gao, "Moving beyond end-to-end path information to optimize CDN performance," in *Proc. of Internet Measurement Conference (IMC'09)*, Chicago, IL, USA, Nov. 2009.
[36] R. Torres, A. Finamore, J. R. Kim, M. Mellia, M. M. Munafo, and S. Rao, "Dissecting video server selection strategies in the YouTube CDN," in *Proc. of ICDCS'11*, Minneapolis, MN, USA, Jun. 2011.
[37] M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye, "Analyzing the performance of an anycast CDN," in *Proc. of Internet Measurement Conference (IMC'15)*, Tokyo, Japan, Nov. 2015.
[38] E. Nygren, R. K. Sitaraman, and J. Sun, "The Akamai network: A platform for highperformance internet applications," *ACM SIGOPS Operating Systems Review*, vol. 44, no. 3, pp. 2–19, 2010.
[39] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan, "Mapping the expansion of Google's serving infrastructure," in *Proc. of Internet Measurement Conference (IMC'10)*, Barcelona, Spain, Nov. 2010.
[40] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet interdomain traffic," in *Proc. ACM SIGCOMM'10*, New Delhi, India, Aug. 2010.
[41] F. Chen, R. K. Sitaraman, and M. Torres, "End-user mapping: Next generation request routing for content delivery," in *Proc. ACM SIGCOMM'15*, London, UK, Aug. 2015.

Junqiang Ge received the bachelor's degree in computer science from the Central South University, Changsha, China, in 2015. He is currently working toward the master's degree with the School of Computer Science and Technology, University of Science and Technology of China, Hefei, China. His research interests include Internet measurement and performance evaluation. He is a student member of the IEEE.

Ye Tian received the bachelor's degree in electronic engineering and the master's degree in computer science from the University of Science and Technology of China (USTC), Hefei, China, in 2001 and 2004, respectively, and the Ph.D. degree from the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, in 2007. He joined USTC in 2008 and is currently an Associate Professor with the School of Computer Science and Technology, USTC. His research interests include Internet measurement, performance evaluation, future Internet architectures, and software-defined networking. He has published over 50 papers and co-authored a research monograph published by Springer. He is the winner of the Wilkes Best Paper Award of Oxford The Computer Journal in 2016. He is currently serving as an Associate Editor of the Springer Frontiers of Computer Science Journal. He is a member of the IEEE.

Lei Liu received the bachelor's degree in software engineering from the Xidian University, Xi'an, China, in 2016. He is currently working toward the master's degree with the School of Computer Science and Technology, University of Science and Technology of China, Hefei, China. His research interests include the measurement studies on Internet and cloud systems.

Rongheng Lan received the bachelor's degree in statistics from the University of Science and Technology of China, Hefei, China, in 2016. He is currently working toward the master's degree with the School of Computer Science and Technology, University of Science and Technology of China, Hefei, China. His research interests include performance modeling and evaluation of Internet systems.

Xinming Zhang received the BE and ME degrees in electrical engineering from the China University of Mining and Technology, Xuzhou, China, in 1985 and 1988, respectively, and the PhD degree in computer science and technology from the University of Science and Technology of China, Hefei, China, in 2001. Since 2002, he has been with the faculty of the University of Science and Technology of China, where he is currently a professor with the School of Computer Science and Technology. He is also with the National Mobile Communications Research Laboratory, Southeast University, Nanjing, China. From September 2005 to August 2006, he was a visiting professor with the Department of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology, Daejeon, Korea. His research interest includes wireless networks. He has published more than 70 papers in wireless ad hoc and sensor networks. He is a member of the IEEE.