Where is the Traffic Going? A Comparative Study of Clouds following Different Designs

Qinkai Wang, *Student Member, IEEE,* Ye Tian, *Member, IEEE,* Xin Yu, Lan Ding, and Xinming Zhang, *Senior Member, IEEE*

Abstract—Cloud computing is critical for today's information society. In this paper, we shed light on two radically different cloud design philosophies: the *DC-cloud* built around massive data centers, and the *ISP-cloud* built upon a large ISP. With extensive measurements on Alibaba, Tencent, and CTYun, we find that both designs have strengths and weaknesses: the ISP-cloud of CTYun has less inflated paths to users within the same ISP, but its paths to external users are more inflated comparing with the DC-clouds of Alibaba and Tencent. By analyzing the clouds' routing policies, we reveal the reasons behind the path inflations: Alibaba and Tencent adopt an *early-exit* policy to use more inflated public Internet paths as early as possible; while CTYun follows a *global and location-agnostic* policy to detour traffic to remote PoPs, leading to highly inflated paths. Based on the insights, we suggest alternative policies and averagely reduce 11.0% latency to 30.5% destinations for Alibaba, 9.8% latency to 18.6% destinations for Tencent, and 54.1% latency to external destinations for CTYun. The results suggest that both cloud designs have rooms for improvement, and an ISP-cloud has the potential to achieve a superior performance, thanks to its inherited advantages from the ISP infrastructure.

Index Terms—Cloud, routing, Internet measurement, performance

1 Introduction

Cloud computing has become an important foundation for today's information society in the past decade. According to Cisco, cloud traffic had reached $14.1\,\mathrm{ZB}$, and constituted over 92% of the global IP traffic in the year 2020 [1]. A cloud provider usually runs a number of inter-connected data centers, and provides services to users on the public Internet. Different providers have different ways to deploy their data centers and interplay with the Internet, and a provider's design choices can greatly impact the network performance of the cloud.

The current design of a cloud follows two radically different philosophies. One philosophy is to *build cloud around massive data centers*. Clouds following this philosophy are usually built by so-called hypergiants, such as Amazon, Google, Alibaba, Tencent, etc., which have already operate a number of massive data centers, and built network infrastructures to inter-connect them. To reach to users on the Internet, hypergiants make use of the facilities such as Internet eXchange Points (IXPs) [2] [3], and peer with as many eyeball Internet Service Providers (ISPs) as possible. We refer to a cloud that follows this design philosophy as a *DC-cloud*.

The other philosophy is to *build cloud upon a large ISP*. Recently, a few tier-1 / tier-2 ISPs start to build their own clouds by deploying data centers on their network infrastructures. By residing in a large ISP, such a cloud does not need to connect to other ISPs explicitly, but relies on the host ISP's inter-domain routing (e.g., Border Gateway Protocol (BGP)) to reach to the rest of the Internet. We refer to a cloud following this philosophy as an *ISP-cloud*. Examples include Verizon, Deutsche Telekom, and the major Chinese ISPs, namely China Telecom, China Unicom, and China Mobile, all build their clouds following this design.

Most previous studies focus on the mainstream DC-clouds like Amazon, Google, Microsoft, etc., while little attention is paid to the other design of the ISP-cloud, and there lacks a comprehensive comparison between the two different designs. In this work, we carry out, to our knowledge, the first comparative study on clouds following the two design philosophies. With a measurement-based approach, we aim to answer the following questions:

1) *Which design, the DC-cloud or the ISP-cloud, provides better paths to users on the Internet?*
2) *What are the reasons behind the paths with poor performances for the clouds of the two designs?*
3) *How the two cloud designs can be improved? and which cloud design has the potential to provide paths to users with a superior performance?*

Measuring, analyzing, and comparing clouds following different designs are challenging. The first challenge is that since different clouds focus on markets of different regions, to enable a fair comparison, the clouds representing different design philosophies should have an overlapped market, and should have comparable shares in this market. The second challenge is that clouds generally do not reveal their networks' footprints and routing policies to the public, making a third-party analysis difficult.

To overcome the first challenge, we select *Alibaba Cloud*[1] and *Tencent Cloud*[2] to represent the design of the DC-

• *The authors are with the Anhui Key Laboratory on High-Performance Computing, School of Computer Science and Technology, University of Science and Technology of China, Hefei, Anhui, 230026, China.*
E-mail: {qinkai, yuxin5836, landing}@mail.ustc.edu.cn, {yetian, xinming}@ustc.edu.cn
• *Corresponding author: Ye Tian (yetian@ustc.edu.cn).*

1. https://www.aliyun.com/
2. https://cloud.tencent.com/

cloud; and we choose *CTYun*[3], the cloud operated by China Telecom, which is the largest ISP in China, as an ISP-cloud representative. Moreover, we restrict our study within mainland China, as it is the domestic and the most important market for the three clouds. According to a recent report [4], Alibaba, Tencent, and CTYun are ranked 1^{st}, 3^{rd}, and 4^{th} respectively in China's public cloud market in 2020.

To overcome the second challenge, we perform extensive measurements on the three clouds. Moreover, we propose a novel IP address geolocating method, discover all the Point of Presences (PoPs) where the clouds exchange traffics with the rest of the Internet, and analyze the clouds' routing policies with extensive traceroute probes.

We make following contributions in this paper. First, we assess and compare performances of the clouds following different design philosophies. In addition to latency, we particularly focus on the metric of *path inflation ratio*, which better captures how well an end-end path is provisioned on the Internet. We find that comparing with the DC-clouds of Alibaba and Tencent, the ISP-cloud of CTYun has the least inflated paths to users within the same ISP, but its paths to external users outside the ISP network are most inflated. We also confirm that paths within a cloud's private wide area network (WAN) is considerably less inflated than paths on the public Internet.

Second, we analyze the reasons behind the path inflations of the two cloud designs. We propose a novel method to geolocate a cloud's border IP addresses, discover the PoPs where the cloud exchange traffic with the Internet, and analyze its routing policy. We find that the DC-clouds of Alibaba and Tencent adopt an *early-exit policy* to use more inflated paths on the public Internet as early as possible; while the ISP-cloud of CTYun follows a *global and location-agnostic policy* inherited from Telecom to detour traffic to remote PoPs.

Third, we explore "what-if" scenarios and suggest alternative routing policies to improve performances for both cloud designs. We suggest that cloud traffic should be carried within a cloud's private network as much as possible, and for an ISP-cloud, PoPs should be selected in a location-aware way to avoid traffic detouring. Evaluating with the measurement data suggests that the alternative policies can reduce 11.0% latency to 30.5% destinations for Alibaba, 9.8% latency to 18.6% destinations for Tencent, and as much as 54.1% latency to non-Telecom destinations for CTYun.

Our study elaborates that both cloud designs have strengths and weaknesses, and have different reasons behind their path inflations. Morevoer, the design of the ISP-cloud has the potential to achieve a superior performance comparing with its DC-cloud counterpart, thanks to its inherited advantages in many data centers and the large-scale ISP network infrastructure.

The remainder part of this paper is organized as follows. We introduce the background and related works in Sec. 2; Sec. 3 describes our measurement methodology; We assess and compare the clouds' performances in Sec. 4; Sec. 5 analyzes the clouds' routing policies; We explore "what-if" scenarios in Sec. 6; Sec. 7 discusses the insights, implications, and limitations; We conclude this paper and discuss the future work in Sec. 8.

3. https://www.ctyun.cn/

2 Background and Related Work

2.1 Background

Both Alibaba and Tencent operate clouds initially for their own businesses, i.e., e-commerce and online social networking, and later provides a wide range of cloud services to the public. A major business application on the Alibaba cloud is Internet finance, which includes data services for banks, securities, insurances, and fund investments. The workload from online trade and bank operations is non-trivial, for example, in 2019's double 11 shopping festival, Alibaba processed 1.29 billion transactions in 24 hours, with a peak of 0.54 million transactions per second [5].

Tencent is the largest and most profitable game company in the world, and the Tencent cloud hosts hundreds of online games that attract billions of online players. For example, "Honor of Kings", the world-wide most popular mobile game run by Tencent, had a record of 100 million daily active players in 2021 [6].

Currently, Alibaba provides Infrastructure-as-a-Service (IaaS) services from 11 data centers, while Tencent provides IaaS services from 6 locations in China. All the 6 Tencent data centers are in metropolises, while 4 of the 11 Alibaba data centers are placed in small cities for reducing the housing and power costs.

CTYun is built and operated by China Telecom, which is the largest ISP in China. As a combination of transit and access networks containing dozens of pubic and private ASes, Telecom's network infrastructure is composed of a national backbone, dozens of provincial networks, and hundreds of metropolitan networks in all the provinces and cities in China [7]. Residing in Telecom's network infrastructure, CTYun provides IaaS services from over 30 data centers, nearly one in each province.

As a state-owned provider, CTYun provides cloud services for many public affairs such as e-government, intelligent transportation, healthcare, education, etc. In particular, CTYun plays a critical role in China's battle against CVOID-19, as many health surveillance systems, which trace potential infection cases with big data, run on CTYun. For example, during the epidemic surge, the Xi'an city's health surveillance system on CTYun handled over $40,000$ health QR-code scan queries per second in peak hours.

2.2 Related Work

Server placement and selection. Clouds follow certain strategies to place servers. For example, Akamai places its servers in numerous ISP PoPs worldwidely, while Limelight has only 20 data centers deployed at key locations [8] [9]. When providing services to users, a cloud needs to select the "right" server for individual users, and the server selection strategy is balanced between various objectives, including user's QoE [10], bandwidth price [11], energy consumption [12], and the handling of flash crowds [13].

Interplay between cloud and ISP. Clouds have various ways to interplay with the rest part of the Internet. For example, YouTube employed a location-agnostic strategy to exchange traffic with a tier-1 ISP, and the traffic volume at each PoP is proportional to the PoP's capacity [14]; Netflix takes advantage of worldwide IXPs to deliver significant amounts of its traffic [3]; Amazon has as many as six different types of peering relationships to connect to other networks [15]. A recent study reveals that the interplays between clouds and ISPs can impact various aspects of user perceived QoE, especially for mobile users [16].

Flattening of the Internet structure. The trend of the flattening [17] in the past decade has fundamentally altered the Internet's hierarchical structure, and makes major cloud providers, which set up peering links with many networks, more and more important. For example, Google and other clouds are observed only one Autonomous System (AS) hop away from most of their users [18], and it is possible for major clouds to bypass Tier-1 ISPs and other transit networks, while are still able to reach 76% of the Internet [19]. A recent study confirms that the widespread cloud data centers are capable to support several latency-critical applications to the global Internet, thanks to the flattening Internet structure [20].

Cloud and Internet paths. End-end paths on the public Internet are notorious for inflations [21], and one major reason is the performance-agnostic inter-domain policy routing (e.g., BGP) widely used on the Internet. Overlay routing and multihoming, which can be easily realized by clouds, are proved to be effective in overcoming the inefficiency [22]. For example, a measurement study [23] shows that paths within a cloud's private WAN are better provisioned and managed, and have lower loss rates and jitters than public Internet paths; a recent study finds that under a multi-cloud scenario, using a cloud's private connectivity can provide an overall superior performance comparing with the best-effort public Internet [24]; and private WANs are found critical in reducing latency imbalances among the cloud paths [25]. Recently, Google and Amazon allow users to choose between their private WANs and the public Internet for carrying the important business traffics [26]. Moreover, cloud overlays have been proposed as a basic network service for improving global Internet data delivery [27].

Although there is a rich literature on network performances of clouds, however, most previous works focus on the mainstream design of the DC-cloud, while very little attention has been paid to the other design of the ISP-cloud. To our best knowledge, this work presents the first comprehensive comparison between the two cloud design philosophies with a measurement approach.

3 Methodology

For assessing a cloud's performance, we collect and analyze a large volume of diverse measurement data. Our basic idea is to send traffic from each cloud data center towards a set of addresses representing the entire Chinese Internet. To this end, we rent a virtual machine (VM) at each data center of a cloud. More specifically, we rent 11 ECS instances on Alibaba, 6 CVM instances on Tencent, and 26 CT-ECS instances on CTYun[4]. A total number of 43 VMs at 29 different cities are employed in this work.

4. CTYun lists a total number of 32 locations on its website, but only 26 are available when we rent VMs.

3.1 Sampling Chinese Internet Address Space

We restrict our study within mainland China, and select IP addresses that can represent the entire Chinese Internet as the measurement targets. By referring to CAIDA's AS-to-organization mapping [28] and BGP routing tables from RIPE RIS [29] and Oregon RouteViews [30], we have collected all the $29,664$ prefixes from the 585 ASes that belong to China. For sampling the Chinese IP address space, we divide the prefixes into a total number of $1,188,516$ /24 Classless Inter-Domain Routing (CIDR) blocks. From each block, we consider the gateway address (i.e., address in the form of `a.b.c.1`) as the destination for traceroute probes. With a commercial geoIP service (i.e., `ipip.net`), we find that the $1,188,516$ gateway addresses are distributed in 352 cities, which are indeed all the cities in China.

We also identify $15,211$ addresses that answer ICMP pings from the $1,188,516$ gateway addresses. These pingable addresses are distributed in 217 ASes and 347 cities.

3.2 Measurement Tool and Performance Metric

We run `scamper` [31] [32] to traceroute from each of the 43 VMs to all the $1,188,516$ gateway addresses sampled from the Chinese Internet. For soliciting accurate paths efficiently, we carefully construct the probe packets and design the give up conditions. More specifically, on Alibaba and Tencent VMs, we configure scamper to send UDP-paris packets to avoid inaccurate path information introduced by multipath load balancing, which is widely employed in cloud environments [33]. We send TCP-paris probe packets on CTYun VMs, as Telecom routers have strict rate restrictions on UDP. For the give-up condition, we give up a probe after encountering 7 consecutive anonymous routers (i.e., 7 continuous '*'s on the path) on Alibaba and Tencent; while on CTYun, we give up after receiving 13 consecutive '*'s, as many Telecom routers do not return the ICMP Time Exceeded message.

We use latency as our primary metric in evaluating a cloud's performance. From the discussion in Sec. 2.1, one can see that for the applications of the Internet finance, online gaming, and health surveillance on the three clouds, latency is the most critical performance metric. Moreover, a cloud is able to reduce its latencies to Internet users by better planning the paths. We do not consider throughput as a metric for two reasons: First, within a cloud, a VM's allocated bandwidth depends on its type and price. For example, Alibaba and Tencent announce that they can provide bandwidths up to $35\,\mathrm{Gbps}$ and $30\,\mathrm{Gbps}$ per VM respectively, and such a high bandwidth is provisioned with software-defined networking [34] and elastic optical network [35] [36]. Second, for a user on the Internet, his achievable throughput is indeed constraint by the access network. For example, a recent study shows that WiFi, LTE, or 5G uplink limits a user's throughput to a cloud server within $100\,\mathrm{Mbps}$, regardless of the server's allocated bandwidth within the cloud [37]. In other words, a cloud can not improve the throughputs to users by itself alone.

For measuring latency, we run scamper to ping from each VM to all the $15,211$ pingable addresses, and from each VM to each address, we consecutively execute 15

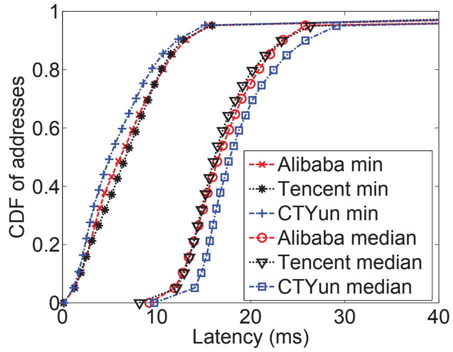

Fig. 1. Minimum and median latencies from VMs in Alibaba, Tencent, and CTYun to $15,211$ pingable addresses.

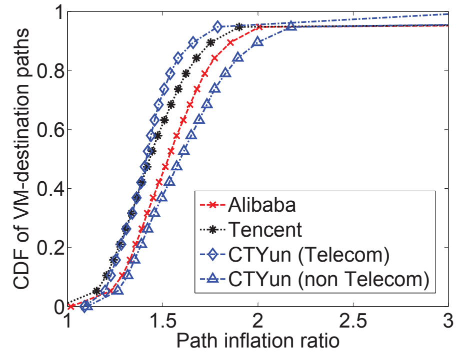

Fig. 2. Inflation ratios of paths from VMs in Alibaba, Tencent, and CTYun to pingable addresses, and for CTYun we differentiate pingable addresses into Telecom addresses and non-Telecom addresses.

pings. To minimize the impact of router buffering delay, the measurements are conducted at 2:00 am in the mid-night, when there is little traffic on the Internet and router buffers are hardly occupied. We complete pinging all the addresses from all the VMs within 60 minutes.

In both the traceroute and the ping measurements, we restrict the packet rate per VM within $300\,\mathrm{pps}$ to avoid triggering any rate control mechanisms on the clouds or ISP networks. A total of $309\,\mathrm{GB}$ measurement data in JSON has been collected in our study.

4 Evaluating Cloud Performance

In this section, we consider the question: which design, the DC-cloud or the ISP-cloud, provides better paths to users on the Internet?

4.1 Latency

To assess cloud latency, we have performed pings from the 43 VMs on the three clouds to the $15,211$ pingable addresses in three nights. For each VM and address pair, 45 RTTs were collected, and we use half of the median as the VM-address latency.

For each address, we compute the minimum and median latencies from VMs of a cloud to it. Note that the minimum latency is a cloud's best-case performance in serving this address, while the median latency indicates the average case.

Fig. 1 presents the minimum and median latencies from the three clouds to the $15,211$ pingable addresses. One can see that the gaps between the two latencies are significant, with the mean differences as $10.63\,\mathrm{ms}$, $10.10\,\mathrm{ms}$, and $13.01\,\mathrm{ms}$ for Alibaba, Tencent, and CTYun respectively. The large gaps suggest that providing service from a data center that is proximate to user is important. Another observation is that CTYun has the lowest minimum latencies among the three clouds. This is because it has much more data centers distributed in a wider geographical area than Alibaba and Tencent, therefore has VMs closer to most pingable addresses. For example, the CTYun VM at Urumqi, the capital city of the Xinjiang province, is much closer to the addresses in Xinjiang, while all the Alibaba and Tencent data centers are over $1000\,\mathrm{km}$ away from this west frontier province. We believe that such an advantage is shared by other ISP-clouds, as generally, a large ISP's infrastructure has already extended over a wide geographical area.

To our surprise, Fig. 1 shows that CTYun has the highest median latency among the three clouds, despite that it resides in the largest ISP network in China. We investigate the reason behind in Sec. 5.

4.2 Path Inflation Ratio

Previous analysis suggests that a cloud's latency is impacted by various factors such as endpoint locations and geographical distances. To filter out these factors, we examine a path's *inflation ratio* instead of its raw latency. More specifically, for an end-end path, we define its inflation ratio as

$$r = \frac{t}{d/v} \tag{1}$$

where t is the measured end-end latency, d is the geographical distance between the two ends, and v is the signal propagation speed in wired medium, which is estimated as $2/3$ of the light speed [38] [39]. Note that an inflation ratio is larger than 1, and the larger the ratio the further the path is deviated away from a hypothetical link that directly connects the two ends. In other words, an Internet path's inflation ratio describes how "straight" the path is, and since it is a ratio, it is irrelevant to the concrete latency and distance values.

Fig. 2 presents the inflation ratios for all the paths from cloud VMs to the pingable addresses. For CTYun, we differentiate the $15,211$ pingable addresses into $4,150$ Telecom addresses and $10,701$ non-Telecom ones. Note that paths from CTYun VMs to Telecom addresses are within the ISP's network, while paths to non-Telecom addresses cross the ISP border.

Fig. 2 shows that for the DC-cloud of Alibaba and Tencent, paths are moderately inflated. Alibaba paths are more inflated than Tencent, as all the Tencent data centers are in metropolises, and have richer ISP connectivity than some of the Alibaba's data centers in small cities. For CTYun, there is a big difference between the paths to Telecom and to non-Telecom addresses: paths to Telecom addresses are least inflated, while paths to non-Telecom addresses are most inflated among the three clouds. The observation suggests that on one hand, an ISP-cloud has well provisioned paths

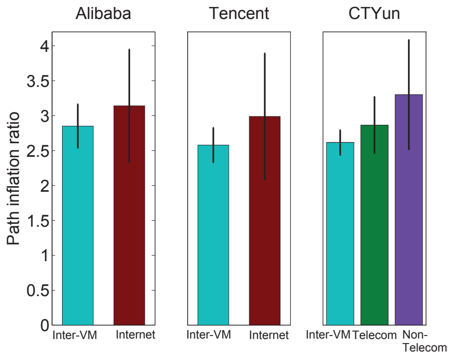

Fig. 3. Comparisons between inflation ratios of a cloud's inter-VM paths and paths from cloud VMs to pingable addresses for Alibaba, Tencent, and CTYun.

to users within the same ISP network, but on the other hand, its paths to external users outside the ISP is highly inflated, probably due to the ISP's inter-domain routing policy, which we analyze in Sec. 5.

4.3 Cloud Private WAN vs. Public Internet

Recent studies suggest that paths within a cloud's private WAN are better provisioned than public Internet paths [23] [27]. In this section, we assess the clouds' private WANs. For Alibaba or Tencent, we examine all the traceroute paths between any pairs of VMs, and find that all the hops on the paths belong to Alibaba or Tencent. We conclude that both clouds have their private WANs that inter-connect their data centers. For CTYun, since it resides in Telecom, we consider the Telecom backbone network that inter-connects its data centers as the cloud's private WAN.

To evaluate a cloud's private network, we measure the latencies between any pair of the cloud's VMs using the method as described in Sec. 3. Fig. 3 presents the inflation ratios of the inter-VM paths for Alibaba, Tencent, and CTYun, and for each cloud, we compare the inter-VM paths with the paths from the cloud to the public Internet. We can see that a cloud's inter-VM paths have much lower inflation ratios, for all the clouds under study.

Moreover, Fig. 3 shows that for CTYun, the inter-VM paths have even lower inflation ratios than the paths from the VMs to the Telecom addresses, although both paths are within the ISP network. We explain the difference with the fact that most of the CTYun data centers directly connect to Telecom's national backbone network, while paths to ordinary Telecom addresses are deep into the provincial and metropolitan networks.

In summary, we find that both cloud designs have strengths and weaknesses: comparing with the DC-clouds of Alibaba and Tencent, the ISP-cloud of CTYun has the least inflated paths to users within the same ISP, but it also has the most inflated paths to external users outside of the ISP network, despite that it has the largest network infrastructure among the three clouds under study. We also confirm that the inter-VM paths within a cloud's private WAN are considerably less inflated than the paths from the clouds to the public Internet.

TABLE 1
IP_c and IP_u addresses of three clouds.

	IP_u	IP_c case I	IP_c case II	IP_c case III
Alibaba	$2,837$	603	0	0
Tencent	$15,346$	0	564	0
CTYun	$4,590$	$3,503$	0	190

5 Analyzing Cloud Routing Policy

In this section, we aim to understand the reasons behind the path inflations by analyzing the routing policies of the clouds following the two design philosophies.

5.1 Finding and Geolocating Cloud-ISP Borders

In order to analyze a cloud's routing policy, we first need to identify and geolocate the borders between the cloud and the Internet. As described in Sec. 3, we perform traceroutes from the 43 VMs in the three clouds to all the $1,188,516$ addresses sampled from the entire Chinese Internet. In each traceroute path, we seek to find where the path leaves the cloud's network. More specifically, we look for an "IP_c, IP_u" tuple on the path, where IP_u is the first IP address not belonging to the cloud's organization (i.e., Alibaba / Tencent / China Telecom), and IP_c is the non-anonymous address immediately before IP_u. After examining all the traceroute paths, we find that the IP_c addresses fall into three cases:

- Case I: IP_c is the last address on the path that belongs to the cloud. This is the ideal case we expect.
- Case II: IP_c is a private address, i.e., address in $10.0.0.0/8$ or $192.168.0.0/16$.
- Case III: IP_c is an unroutable address, that is, address does not belong to any prefix announced by BGP routers on the global Internet, and for such an address, we can not tell its belonging AS and organization.

After examining the traceroute paths, we find all the the IP_u and IP_c addresses that fall into each case for the three clouds, and list their numbers in Table 1. In particular, we find 603 IP_c addresses on Alibaba, which are the cloud's border router addresses. However for Tencent, all the IP_c addresses are private addresses in $10.0.0.0/8$. We believe that these addresses belong to Tencent, which universally uses private addresses to connect to other networks. For CTYun, we find $3,503$ IP_c addresses on border routers, but there are 190 unroutable addresses in $202.97.0.0/12$. Fortunately, our previous study [40] indicates that they belong to Telecom (i.e., they used to be announced by Telecom's BGP routers).

We employ `ipip.net` to geolocate the IP_c and IP_u addresses. However, like many commercial geoIP services, `ipip.net` can not provide city-level location for many router addresses. Moreover, neither private nor unroutable addresses can be geolocated by the service.

To accurately geolocate as many IP_c addresses as possible, we propose an iterative heuristic. We first geolocate IP_c and IP_u addresses with `ipip.net`, then for each ungeolocated IP_c address, we examine all the IP_u addresses that it connects to in the traceroute paths. If it connects to no less than five IP_u addresses with city-level locations, and over

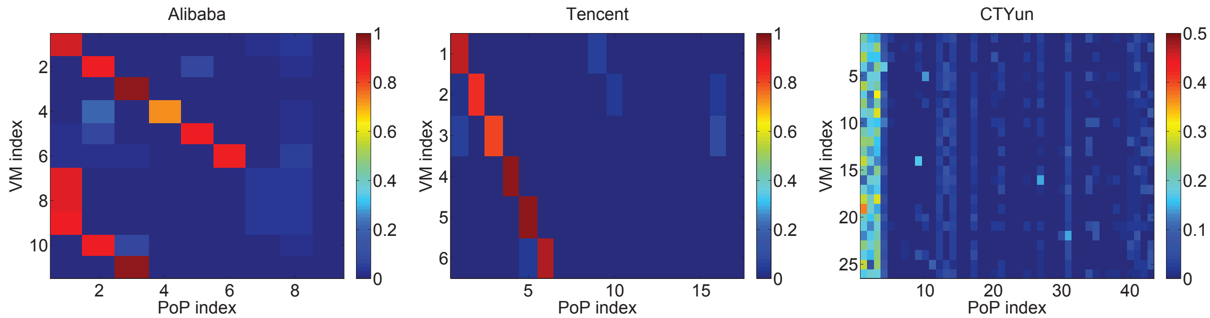

(a) Alibaba (b) Tencent (c) CTYun

Fig. 4. Heatmaps of PoP preferences of VMs in (a) Alibaba, (b) Tencent, and (c) CTYun.

80% of the IP_u addresses are in same city, we geolocate the IP_c address to that city. Similarly, for each IP_u address not geolocated, we examine all the IP_c addresses that connect to it in the traceroute paths. If it is connected by no less than five IP_c addresses with city-level locations, and over 80% of the IP_c addresses are in same city, we assign the IP_u address to that city. We iteratively repeat the two steps until no more IP_c and IP_u addresses can be further geolocated.

With the heuristic, we have successfully geolocated all the 603 IP_c addresses for Alibaba, 542 of the 564 IP_c addresses for Tencent, and $3,660$ of the $3,693$ IP_c addresses on Telecom for CTYun.

5.2 PoPs and Routing Policy

For each cloud, we cluster the IP_c addresses located in same city into a cluster, and refer to such a cluster as a *point of presence* (*PoP*). Typically, a PoP represents a facility where a cloud's network connect to other networks. We have clustered 9 PoPs for Alibaba, 17 PoPs for Tencent, and 43 PoPs for CTYun. Since CTYun resides in Telecom, its 43 PoPs are indeed the PoPs of the ISP to connect to the rest of the Internet.

Our first observation is that most cloud data centers have PoPs co-located in same cities. For example, all the 6 Tencent data centers and all the 26 CTYun data centers have co-located PoPs, and 6 of the 11 Alibaba data centers have PoPs co-located with them.

Another observation is that a traceroute path originating from a VM may leave the cloud's network at any PoP. For a VM, say VM_i, we compute $r_{i,j}$, which is the ratio of the paths from VM_i leaving the cloud's network at PoP_j, defined as

$$r_{i,j} = \frac{\text{num. of paths from } VM_i \text{ leaving the cloud at } PoP_j}{\text{num. of paths from } VM_i} \quad (2)$$

For each cloud, we compute $r_{i,j}$ between all its VMs and PoPs, and present the ratios in heatmaps in Fig. 4. We can see that on the DC-clouds of Alibaba and Tencent, most of a VM's traceroute paths leave the cloud's network at a PoP that is close to the VM's data center. For example, $83-98\%$ of the traceroute paths from Tencent VMs leave the cloud at their co-located PoPs; while on Alibaba, for the 6 data centers with co-located PoPs, $73-97\%$ paths leave at the co-located PoPs, and for the other 5 data centers without co-located PoPs, $87-97\%$ paths leave at the PoPs that are geographically closest to the data centers. From Fig. 4(a) and (b), we conclude that the DC-clouds employ an *early-exit policy* to deliver traffic to the Internet at the PoPs that are as close to their data centers as possible.

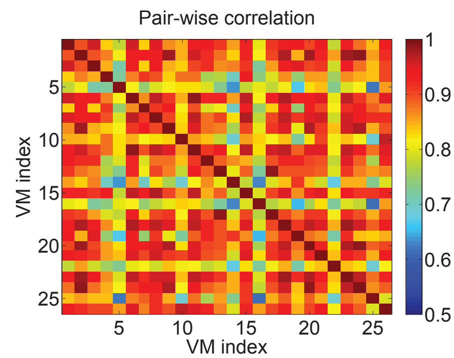

Fig. 5. Pair-wise correlations of PoP preference vectors among all CTYun VMs.

For the ISP-cloud of CTYun, although each data center has a co-located PoP, however, traceroute paths from a VM hardly leave Telecom at the co-located PoP. For example, among the 26 VMs, 20 have fewer than 5% paths exiting Telecom via their co-located PoPs. On the other hand, Fig. 4(c) shows that for all the VMs at different locations, $44-67\%$ of their paths leave Telecom via the three PoPs at Beijing, Shanghai, and Guangzhou, which correspond to the first three columns in the heatmap.

Fig. 4(c) suggests that CTYun VMs have similar preferences to the 43 PoPs. To confirm such similarity, for each CTYun VM, say VM_i, we denote its preference to the 43 PoPs as a vector $[r_{i,1}, r_{i,2}, \cdots, r_{i,43}]$, where $r_{i,j}$ is the ratio of the paths from VM_i exiting Telecom at PoP_j as defined in (2). For any two VMs, we compute the correlation between their PoP preference vectors, and present the pair-wise correlations among all the VMs in Fig. 5. We find that most of the correlations are above 0.9, with an average as high as 0.871. The high correlations confirm that CTYun or Telecom employs a *global and location-agnostic policy* to route traffic from different data centers to external destinations. In such a policy, data centers at different locations have similar preferences to the PoPs, and among them, Beijing, Shanghai, and Guangzhou, which have the largest PoP capacities in terms of number of IP_c addresses, are most preferred.

To further understand the two different routing policies

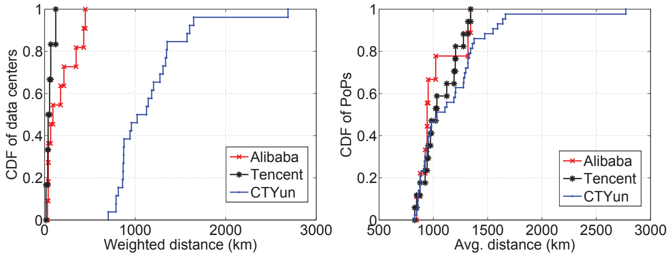

(a) VM-to-PoP distance (b) PoP-to-address distance

Fig. 6. Comparisons of (a) weighted mean VM-to-PoP distances and (b) averaged PoP-to-address distances for Alibaba, Tencent, and CTYun.

adopted by the clouds of different designs, for each cloud, we compute the weighted mean distance between a VM and the PoPs where paths originating from the VM leave the cloud's network, i.e.,

$$dist(VM_i, PoP\_set) = \sum_{PoP_j \in PoP\_set} dist(VM_i, PoP_j) \times r_{i,j} \tag{3}$$

where PoP_set is the cloud's PoP set. We also compute the averaged distance between a PoP and the traceroute destination addresses that paths to these addresses leave the cloud's network via the PoP as

$$dist(PoP_j, addr\_set) = \frac{\sum_{addr_k \in addr\_set} dist(PoP_j, addr_k)}{|addr\_set|} \tag{4}$$

where $addr\_set$ is the set of the destination addresses. Note that for CTYun, $addr\_set$ contains only the non-Telecom addresses.

Fig. 6 presents and compares the two distances among the three clouds. We can see that under the early-exit policy adopted by Alibaba and Tencent, VMs or data centers are geographically close to the PoPs, but PoPs are far away from the traceroute destination addresses. While for the ISP-cloud of CTYun, data centers are far away from the PoPs, and PoPs are also far away from the destinations, or in other words, under the cloud's global and location-agnostic routing policy, traffics from CTYun VMs *detour* to PoPs that are far away from both ends of the path. Clearly, the traffic detouring explains CTYun's high latency and highly inflated paths to external addresses as in Fig. 1 and Fig. 2 respectively.

We recognize that the early-exit policy adopted by Alibaba and Tencent are widely observed in providers, while the global and location-agnostic policy adopted by CTYun also has precedent, for example, YouTube used to adopt such a policy before it is integrated into Google's network [14].

Finally, we conclude that the different routing policies adopted by the two cloud designs impact cloud performances in different ways:

- For the DC-clouds of Alibaba and Tencent, the early-exit policy forces traffics of cloud services to travel on public Internet paths as early as possible, however, Sec. 4.3 indicates that public Internet paths are considerably more inflated than paths within a cloud's private WAN.

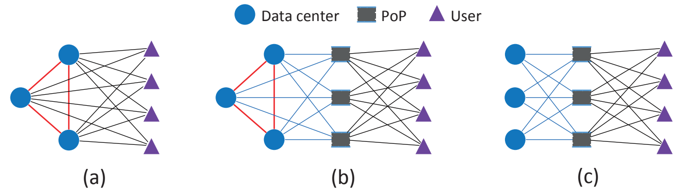

Fig. 7. (a) A graph model for DC-cloud, (b) a graph model for ISP-cloud, (c) another graph model for ISP-cloud.

- For the ISP-cloud of CTYun, the ISP's global and location-agnostic routing policy forces cloud traffic to detour to PoPs that are far away from both the data center and the user, leading to highly inflated paths.

6 Improving Cloud Performance

Our analysis reveals that clouds of different designs adopt different routing policies, which lead to path inflations. In this section, we consider "what-if" scenarios by suggesting alternative policies for improving the clouds' performances.

6.1 Improving DC-cloud

For improving a DC-cloud like Alibaba and Tencent, we suggest that paths from cloud data centers to Internet users should traverse the cloud's private WAN as much as possible.

More specifically, we model a DC-cloud as a combination of a *complete graph* and a *complete bipartite graph* as demonstrated in Fig. 7(a). The complete graph $G = (V, E)$ represents the cloud's private WAN that inter-connects its data centers, where V is the set of the data centers (or VMs) and E is the set of the weighted edges with inter-VM latencies as the weights; the complete bipartite graph $G' = (V, A, E')$ represents the paths from the data centers to Internet users, where V is the VM set in G, A is the set of the Internet users, and E' is the set of the weighted edges with latencies between VMs and users as the weights. Here we use the $15,211$ pingable addresses to represent the users.

Since we have measured the inter-VM latencies and the latencies from each VM to each pingable address, we can construct the combined graph $G \cup G'$ with the measurement data. By applying Dijkstra's algorithm, we can find the shortest path between any VM in V and any pingable address in A. As a path may traverse additional data centers in V, we refer to the corresponding routing policy as the *inter-DC policy*.

Fig. 8(a) and Fig. 9(a) present and compare the VM-address latencies under the inter-DC policy and the early-exit policy currently adopted by Alibaba and Tencent respectively. By traversing a cloud's private WAN, the inter-DC policy reduces the latency for 30.5% paths on Alibaba, and the mean latency reduction is $2.57\,\text{ms}$, which is 11.0% of the latency under the early-exit policy. For Tencent, the inter-DC policy reduces the latency for 18.6% paths, and the mean reduction is $2.43\,\text{ms}$, constituting 9.8% of the latency under the early-exit policy.

Fig. 8(b) and Fig. 9(b) present the ratios of the inter-DC routing paths that traverse zero, one, or more additional

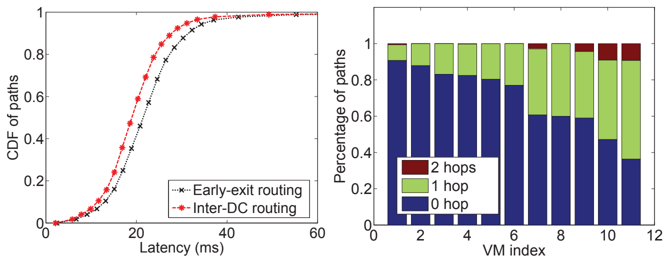

(a) (b)

Fig. 8. (a) Comparison of latencies from Alibaba VMs to pingable addresses under early-exit policy and inter-DC policy, (b) ratios of inter-DC paths for Alibaba that contain various numbers of hops.

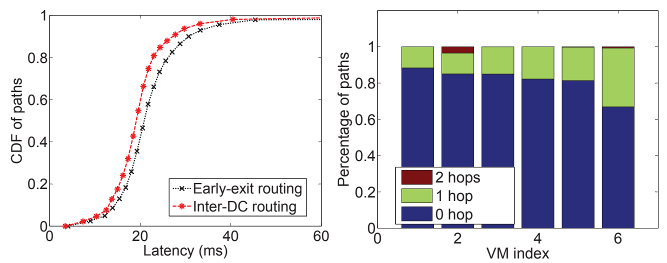

(a) (b)

Fig. 9. (a) Comparison of latencies from Tencent VMs to pingable addresses under early-exit policy and inter-DC policy, (b) ratios of inter-DC paths for Tencent that contain various numbers of hops.

data centers as hops on Alibaba and Tencent respectively. Note that when a path traverses no additional hop, it is indeed the path decided by the early-exit policy. On average, 30.5% and 18.6% of the paths traverse one or more additional data centers on Alibaba and Tencent, suggesting that there are considerable rooms for the clouds to improve their performances by carrying cloud traffics on their private WANs. On the other hand, in most cases, traversing one hop is enough, and there is no need to plan paths passing over two data centers.

Fig. 8 and Fig. 9 show that the performance improvement on Alibaba is more pronounced than on Tencent. This is reasonable as all the Tencent data centers are in metropolises with rich ISP connectivity and co-located PoPs. While on Alibaba, 4 of the 11 data centers are in small cities, and 5 data centers do not have co-located PoPs. So the inter-DC routing policy can better improve Alibaba than Tencent.

6.2 Improving ISP-cloud

The major reason behind path inflations of an ISP-cloud like CTYun is that the cloud inherits the ISP's global and location-agnostic routing policy to detour cloud traffic to remote PoPs. To avoid such detouring, we suggest that the routing policy should be location aware, in addition, as on Alibaba and Tencent, the cloud's private WAN should carry the cloud traffic as much as possible.

Unfortunately, we are unable to construct new paths from VMs to pingable addresses by simply stitching the inter-VM paths with the paths from VMs to pingable addresses, like we have done for Alibaba and Tencent. This is because under the current policy, traffic from any VM always detours to remote PoPs to leave Telecom, regardless of the data center's location. To overcome this problem, we explicitly include PoPs into the cloud model to allow Dijkstra's algorithm to explicitly add PoPs to the alternative paths.

More specifically, we model an ISP-cloud as in Fig. 7(b), which is a combination of a complete graph $G = (V, E)$, a complete bipartite graph $G' = (V, P, E')$, and another complete bipartite graph $G'' = (P, A, E'')$. Here G represents the cloud's private WAN that inter-connects its data centers; G' represents the paths from the data centers to the PoPs, where P is the cloud's PoP set, and E' is the set of the weighted edges between the VMs and the PoPs with the VM-PoP latencies as the weights; G'' represents the paths from the PoPs to external users, where A is the user set, and E'' is the set of the weighted edges with the PoP-user latencies as the weights. Here we use the $10,701$ non-Telecom pingable addresses to represent the external users outside the ISP.

Before applying Dijkstra's algorithm, we need to decide the weights for the edges in E' and E'', which are the VM-PoP latencies and PoP-address latencies. Unfortunately, neither of them can be directly measured, so we estimate the latencies as

$$t = \frac{r \times d}{v} \tag{5}$$

where d is the geographical distance between the two ends, v is $2/3$ of light speed, and r is path inflation ratio. For estimating a VM-PoP latency, we set $r = 2.813$, which is the mean inflation ratio for paths from VMs to Telecom addresses, and we let $r = 3.065$, which is the mean inflation ratio for paths outside Telecom, to estimate the PoP-address latencies. Note that both values are derived from our measurement data. Since the shortest path from any VM in V to any address in A on the combined graph $G \cup G' \cup G''$ may traverse multiple data centers and exit Telecom from a PoP close to the last hop, we refer to the corresponding routing policy as the *inter-DC and location-aware policy*.

For comparison, we also consider another cloud model as shown in Fig. 7(c), in which we do not model the cloud's private network as a complete graph, but simply combine the two complete bipartite graphs of G' and G''. Obviously, the shortest paths on $G' \cup G''$ can avoid traffic detouring, but they do not make use of the cloud's private WAN. We refer to the corresponding routing policy as the *location-aware policy*.

Fig. 10(a) compares the latencies between the CTYun VMs and the non-Telecom pingable addresses under different routing policies. We can see that by avoiding traffic detouring, the location-aware policy reduces the latency $14.79\,\text{ms}$ on average, constituting 42.6% of the latency under the global and location-agnostic policy currently adopted by CTYun. Furthermore, by carrying traffic on the cloud's private WAN with the inter-DC and location-aware policy, we further reduce the latency $18.79\,\text{ms}$, which is 54.1% of the original latency. Note that our improvement on CTYun is much more significant than the ones on Alibaba and Tencent, this is because unlike Alibaba and Tencent, CTYun's global and location-agnostic policy causes expensive traffic detouring.

Fig. 10(b) presents the ratios of the paths under the inter-DC and location-aware policy that traverse zero, one,

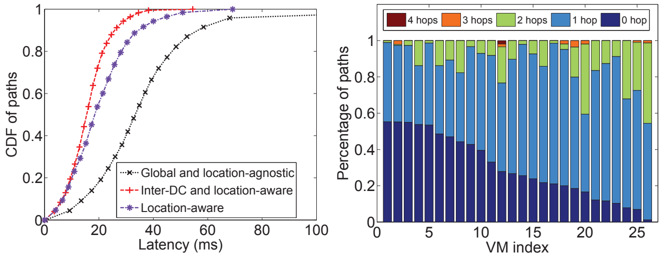

(a) (b)

Fig. 10. (a) Comparison of latencies from CTYun VMs to non-Telecom addresses under global and location-agnostic policy, location-aware policy, and inter-DC and location-aware policy, (b) ratios of paths under inter-DC and location-aware policy for CTYun that contain various numbers of hops.

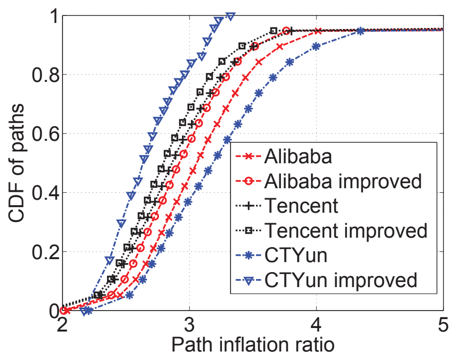

Fig. 11. Comparison of path inflation ratios with and without improvements for Alibaba, Tencent, and CTYun.

or more additional data centers as hops. One can see that although 30.0% of the paths contain no additional hop (i.e., they are indeed the paths decided by the location-aware policy), however, 56.6% paths contain one additional hop, and 13.4% paths contain two or more additional hops. Moreover, the longest paths contain as many as 4 additional data centers within CTYun's private WAN. Comparing Fig. 10(b) with Fig. 8(b) and Fig. 9(b), we can see that the alternative paths traverse more data centers on CTYun than on Alibaba and Tencent, this is because CTYun has much more data centers, and its private WAN, which is indeed Telecom's backbone network, is far more geographically extensive than the ones of Alibaba and Tencent. In other words, the inter-DC policy brings more benefit on an ISP-cloud than on a DC-cloud.

Finally, we compute the inflation ratios for paths from cloud VMs to pingable addresses for Alibaba, Tencent, and CTYun with and without our improvements, and compare them in Fig. 11 and Table 2. We can see that by applying the inter-DC policy, both Alibaba and Tencent have considerably reduced their path inflation ratios; and by

TABLE 2
Comparison of averaged path inflation ratios of various network paths on the three clouds with and without the improvements.

Avg. path inflation ratio	Alibaba	Tencent	CTYun
Inter-VM	2.850	2.579	2.615
VM-user	3.141	2.988	3.301
VM-user (improved)	2.990	2.897	2.653

applying the inter-DC and location-aware policy, CTYun has the least inflated paths among the three clouds. Note that the observation is in contrast with Fig. 2, in which CTYun has the most inflated paths. We conclude that after eliminating the causes of path inflations in both designs, the ISP-cloud of CTYun is capable to provide better paths than the DC-clouds of Alibaba and Tencent, thanks to its inherited advantages in many data centers and the large-scale ISP network infrastructure.

7 Discussion

7.1 Implication

As discussed in Sec. 2.1, Internet finance and online gaming are two major applications on Alibaba and Tencent respectively, and CTYun runs health surveillance applications for many cites and provinces in China. Reducing cloud-user latencies definitely benefit these applications. For example, it is estimated that for high-frequency trading, a millisecond delay decrease can boost a firm's earnings by as much as \$100 million per year [41]. Our suggested inter-DC policy, which averagely reduces $2.57\,\mathrm{ms}$ for 30.5% Alibaba's VM-user paths, can help its clients to make more money from the financial market.

Our suggested inter-DC policy averagely reduces $2.43\,\mathrm{ms}$ for 9.8% VM-user paths for Tencent, and such a reduction benefits online game applications. For example, a latency exceeding $20\,\mathrm{ms}$ triggers motion sickness in immersive applications, such as AR/VR, $360°$ streaming, etc; and a latency over $100\,\mathrm{ms}$ can severely influence a player in cloud games [42]. Considering that the delays caused by computations such as video transcoding and game logic executions are difficult to avoid [37], each millisecond reduction in data transportation is highly valuable.

For CTYun, our suggested inter-DC and location-aware policy, which achieves an averaged latency reduction of $18.79\,\mathrm{ms}$ to non-Telecom addresses, also benefits public affair applications. For example, as nowadays in China, scanning the health QR-code is mandatory in many daily occasions such as taking bus/subway, entering into school/workplace, etc., saving dozens of milliseconds in each query can accumulatively save considerable time for billions of people.

7.2 Depending factor and practical issue

Our suggested alternative routing policies ignore some factors such as cloud-ISP business relation, bandwidth price, PoP capacity, etc. In the following, we discuss how these factors impact the choices of routing policies for clouds following different designs in practice.

For a DC-cloud like Alibaba or Tencent, its private WAN is largely built upon ISP infrastructures. For example, most of the lines connecting Alibaba or Tencent's data centers are leased from ISPs, and typically, ISPs charge the clouds based on bandwidth consumptions. Obviously under such a business relation, when a DC-cloud applies the inter-DC policy by carrying cloud traffic on its own private WAN, its bandwidth expense will certainly increase. On the other hand, a DC-cloud can make use of local facilities such as IXPs to deliver traffic to ISPs at almost zero cost. For

TABLE 3
Key differences in design philosophy, current routing policy, and suggested routing policy between DC-cloud and ISP-cloud.

	DC-cloud	ISP-cloud
Data centers	A few massive data centers, mostly in metropolises.	Dozens of data centers in many different cities.
Data center inter-connection	Inter-connect data centers with private WAN.	Inter-connect data centers with the ISP's backbone network.
Routing policy	Early-exit: Deliver cloud traffic to the Internet via the closest PoP as early as possible.	Global and location agnostic: Deliver cloud traffic to the Internet via PoPs according to a fixed preference based on PoP capacity.
Suggested routing policy	Inter-DC: Use the private WAN to transfer cloud traffic as much as possible.	Inter-DC and location aware: Use the ISP's backbone network to transfer cloud traffic as much as possible, and deliver cloud traffic to the Internet via the PoP close to the destination.

this reason, a DC-cloud has little incentive to apply the suggested inter-DC policy but will stay with the early-exit policy. Nevertheless, for some performance-sensitive and high-valued traffic, a DC-cloud may willing to apply the inter-DC policy.

For the ISP-cloud of CTYun, the advantages brought by its many data centers and extensive network infrastructure are largely offset by the global and location-agnostic routing policy inherited from the ISP. We conjecture that such a policy is not tailored for CTYun, but is applied on all traffics, including the traffics of home and enterprise subscribers, traffics of peering and customer ISPs, as well as the cloud traffic. Note that although inefficient, such a policy has its merit in load balancing among the PoPs [14]. Our study shows that in order to compete with the DC-clouds, an ISP should differentiate its cloud traffic from other traffics, and judiciously apply a different routing policy, so as to fully exploit the ISP-cloud's inherited advantages regarding number of data centers and footprints of the ISP's network infrastructure. Our study shows that after the optimization, an ISP-cloud has the potential to provide a superior performance comparing with its DC-cloud counterpart.

7.3 Limitation

Our work has limitations. First, although we rent 43 VMs at 29 different cities to make traceroutes towards over one million destinations and measure latencies from 43 VMs to more than $15,000$ pingable addresses, however, all the probes are initiated from the clouds. We are unable to make pings and traceroute probes from a large number of widely distributed user clients. The reason is that unlike the Internet in North America and Europe, there are very limited number of vantage points on the Chinese Internet for us to exploit. For example, studies report that there are less than 50 vantage points from China in 6 months on Speedchecker [42], and less than 50 vantage points from China in 12 months on RIPE Atlas [20].

Second, as discussed in Sec. 6.2, since we are unable to directly measure the latencies for the VM-PoP and PoP-address path segments, we apply (5) with the derived path inflation ratios to infer the latencies. The drawback is that although the inferred latencies are statistically accurate, however, an individual path's inferred latency may be quite different from its actual latency in real world.

7.4 Generalized outside China?

Our study is focused on the three major providers of Alibaba, Tencent, and CTYun in mainland China. However, the tools and methods we use in this work, such as `scamper` and the geolocating heuristic, can be applied to other part of the Internet.

Our insights on the DC-cloud are also applicable to other clouds that follow the same design philosophy, and in fact, Google and Amazon recently allow users to choose between their private WANs (with higher prices) and the public Internet for transferring their important business data [26].

Large ISPs outside China also build their clouds, for example, Verizon Business operates 50+ data centers and facilities including Network Access Points (NAPs) and IXPs in United States, Europe, and Latin America to provide an integrated cloud service. From the architectural perspective, Verizon follows the ISP-cloud design philosophy, but it is unknown how it makes use of its global-scale network infrastructure to route cloud traffic to users outside the ISP network. We expect to see more studies on ISP-clouds in future.

8 Conclusion and Future Work

In this paper, we presented, to our best knowledge, the first head-to-head comparison between different cloud design philosophies. Using Alibaba, Tencent, and CTYun as examples, we conducted extensive and thorough measurements that accurately characterize the performances of the clouds. In addition to latency, we particularly focused on path inflation ratio, which better reflects how a network path is planned and provisioned. By applying a novel IP address geolocating method, we accurately identified and located the PoPs where the clouds exchange traffics with the rest of the Internet. We found that the DC-clouds of Alibaba and Tencent adopt an early-exit routing policy, which prefers inflated paths on the public Internet, while the ISP-cloud of CTYun adopts a global and location-agnostic policy to detour cloud traffic to remote PoPs, leading to highly inflated paths.

Based on the insights from the measurement and analysis, we explored "what-if" scenarios and suggested alternative routing policies to improve performances for both cloud designs. By carrying cloud traffics on clouds' private WANs and selecting PoPs in a location-aware way, our suggested alternative policies averagely reduce 11.0% latency to 30.5% destinations for Alibaba, 9.8% latency to 18.6% destinations for Tencent, and 54.1% latency to non-Telecom destinations for CTYun. We also demonstrated that the design of the ISP-cloud has the potential to achieve a superior performance comparing with its DC-cloud counterpart, thanks to its inherited advantages in many data centers and large-scale ISP network infrastructure. Table 3 summarizes the key insights we have made on the two cloud design philosophies in this paper. We believe that our work will provide useful insights for both cloud practitioners and researchers.

For the future work, we seek to deploy measurement tools on some emerging platforms such as OpenNetLab [43], and initiate probes from the network edge (e.g., university / home / cellular networks). Given the flattening trend of the Internet, we also plan to conduct a long-term measurement study to observe how clouds' routing policies evolve over time.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grants 61672486 and 62072425, and in part by the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant XDC02070300.

References

[1] "Cisco annual Internet report (2018-2023) white paper," Mar. 2020. [Online]. Available: https://www.cisco.com/c/en/us/solutions/executive-perspectives/annual-internet-report/index.html

[2] B. Augustin, B. Krishnamurthy, and W. Willinger, "IXPs: Mapped?" in *Proc. IMC'09*, Chicago, IL, USA, Nov. 2009.

[3] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig, "Open connect everywhere: A glimpse at the internet ecosystem through the lens of the Netflix CDN," *ACM SIGCOMM CCR*, vol. 48, no. 1, pp. 28–34, Jan. 2018.

[4] "China's public cloud computing market growth leads the world (in Chinese)," Apr. 2021. [Online]. Available: https://www.idc.com/getdoc.jsp?containerId=prCHC47631621

[5] "This year's Double 11 turnover is 268.4 billion Yuan, Alibaba Cloud holds the world's largest traffic peak (in Chinese)," Nov. 2019. [Online]. Available: https://cloud.ofweek.com/news/2019-11/ART-178803-8110-30416861.html

[6] R. Liao, "Tencent claims record 100m daily users on mobile game honor of kings," Nov. 2020. [Online]. Available: https://techcrunch.com/2020/11/01/tencent-claims-record-100m-daily-users-on-mobile-game-honor-of-kings/

[7] Y. Tian, R. Dey, Y. Liu, and K. W. Ross, "China's Internet: Topology mapping and geolocating," in *Proc. INFOCOM'12 mini-conference*, Orlando, FL, USA, Mar. 2012.

[8] Y. A. Wang, C. Huang, J. Li, and K. Ross, "Estimating the performance of hypothetical cloud service deployments: A measurement-based approach," in *Proc. INFOCOM'11*, Shanghai, China, Jun. 2011.

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

[10] R. Torres, A. Finamore, J. R. Kim, M. Mellia, M. M. Munafò, and S. Rao, "Dissecting video server selection strategies in the YouTube CDN," in *Proc. ICDCS'11*, Minneapolis, MN, USA, Jun. 2011.

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

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

[13] J. Ge, Y. Tian, L. Liu, R. Lan, and X. Zhang, "Understanding e-commerce systems under massive flash crowd: Measurement, analysis, and implications," *IEEE Trans. Services Comput.*, vol. 13, no. 6, pp. 1180–1193, Dec. 2020.

[14] V. K. Adhikari, S. Jain, and Z.-L. Zhang, "YouTube traffic dynamics and its interplay with a tier-1 ISP: An ISP perspective," in *Proc. IMC'10*, Melbourne, Australia, Nov. 2010.

[15] B. Yeganeh, R. Durairajan, R. Rejaie, and W. Willinger, "How cloud traffic goes hiding: A study of amazon's peering fabric," in *Proc. IMC'19*, Amsterdam, Netherlands, Oct. 2019.

[16] F. Michelinakis, H. Doroud, A. Razaghpanah, A. Lutu, N. Vallina-Rodriguez, P. Gill, and J. Widmer, "The cloud that runs the mobile Internet: A measurement study of mobile cloud services," in *Proc. INFOCOM'18*, Honolulu, HI, USA, Apr. 2018.

[17] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *Proc. SIGCOMM'10*, New Delhi, India, Aug. 2010.

[18] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan, "Are we one hop away from a better Internet?" in *Proc. IMC'15*, Tokyo, Japan, Oct. 2015.

[19] T. Arnold, J. He, W. Jiang, M. Calder, I. Cunha, V. Giotsas, and E. Katz-Bassett, "Cloud provider connectivity in the flat Internet," in *Proc. IMC'20*, Pittsburgh, PA, USA, Oct. 2020.

[20] L. Corneo, M. Eder, N. Mohan, A. Zavodovski, S. Bayhan, W. Wong, P. Gunningberg, J. Kangasharju, and J. Ott, "Surrounded by the clouds: A comprehensive cloud reachability study," in *Proc. WWW'21*, Ljubljana, Slovenia, Apr. 2021.

[21] N. Spring, R. Mahajan, and T. Anderson, "Quantifying the causes of path inflation," in *Proc. SIGCOMM'03*, Karlsruhe Germany, Aug. 2003.

[22] A. Akella, J. Pang, B. Maggs, S. Seshan, and A. Shaikh, "A comparison of overlay routing and multihoming route control," in *Proc. SIGCOMM'04*, Portland, OR, USA, Aug. 2004.

[23] O. Haq, M. Raja, and F. R. Dogar, "Measuring and improving the reliability of wide-area cloud paths," in *Proc. WWW'17*, Perth, Australia, Apr. 2017.

[24] B. Yeganeh, R. Durairajan, R. Rejaie, and W. Willinger, "A first comparative characterization of multi-cloud connectivity in today's Internet," in *Proc. International Conference on Passive and Active Network Measurement (PAM'20), LNCS 12048*, Eugene, OR, USA, Mar. 2020.

[25] Y. Pi, S. Jamin, P. Danzig, and F. Qian, "Latency imbalance among internet load-balanced paths: A cloud-centric view," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 4, no. 2, Jun. 2020.

[26] T. Arnold, E. Gürmeriçliler, G. Essig, A. Gupta, M. Calderz, V. Giotsas, and E. Katz-Bassett, "(How much) does a private WAN improve cloud performance?" in *Proc. INFOCOM'20*, Toronto, ON, Canada, Jul. 2020.

[27] O. Haq, C. Doucette, J. W. Byers, and F. R. Dogar, "Judicious QoS using cloud overlays," in *Proc. CoNEXT'20*, Barcelona, Spain, Dec. 2020.

[28] "Inferred AS to organization mapping dataset," accessed on Apr. 12, 2021. [Online]. Available: https://www.caida.org/catalog/datasets/as-organizations/

[29] "Routing information service (RIS)," accessed on Apr. 12, 2021. [Online]. Available: https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

[30] "University of Oregon route views project," accessed on Apr. 12, 2021. [Online]. Available: http://www.routeviews.org/routeviews/

[31] M. Luckie, "Scamper: a scalable and extensible packet prober for active measurement of the Internet," in *Proc. IMC'10*, Melbourne, Australia, Nov. 2010.

[32] "Scamper," accessed on Apr. 12, 2021. [Online]. Available: https://www.caida.org/catalog/software/scamper/

[33] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding traceroute anomalies with paris traceroute," in *Proc. IMC'06*, Rio de Janeiro, Brazil, Oct. 2006.

[34] S. Jain, A. Kumar, S. Mandal, J. Ong, and et al., "B4: Experience with a globally-deployed software defined WAN," in *Proc. SIGCOMM'13*, Hong Kong, China, Aug. 2013.

[35] Z. Zhu, W. Lu, L. Zhang, and N. Ansari, "Dynamic service provisioning in elastic optical networks with hybrid single-/multi-path routing," *J. Lightw. Technol.*, vol. 31, pp. 15–22, Jan. 2013.

[36] Y. Yin, H. Zhang, M. Zhang, M. Xia, and et al., "Spectral and spatial 2D fragmentation-aware routing and spectrum assignment algorithms in elastic optical networks," *J. Opt. Commun. Netw.*, vol. 5, pp. A100–A106, Oct. 2013.

[37] M. Xu, Z. Fu, X. Ma, L. Zhang, Y. Li, F. Qian, S. Wang, K. Li, J. Yang, and X. Liu, "From cloud to edge: A first look at public edge platforms," in *Proc. IMC'21*, Virtual Event, Nov. 2021.
[38] R. Percacci and A. Vespignani, "Scale-free behavior of the Internet global performance," *Eur. Phys. J. B*, vol. 32, pp. 411–414, 2004.
[39] Y. Wang, D. Burgener, M. Flores, A. Kuzmanovic, and C. Huang, "Towards street-level client-independent IP geolocation," in *Proc. NSDI'11*, Boston, MA, USA, Mar. 2011.
[40] Y. Tian, R. Dey, Y. Liu, and K. W. Ross, "Topology mapping and geolocating for China's Internet," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 9, pp. 1908–1917, 2013.
[41] J. Fawcett, "The new finance era: From hft to big data," Jul. 2013. [Online]. Available: https://www.wallstreetandtech.com/electronic-trading/the-new-finance-era-from-hft-to-big-data/a/d-id/1268323.html
[42] T. K. Dang, N. Mohan, L. Corneo, A. Zavodovski, J. Ott, and J. Kangasharju, "Cloudy with a chance of short rtts: Analyzing cloud connectivity in the Internet," in *Proc. IMC'21*, Virtual Event, Nov. 2021.
[43] "OpenNetLab," accessed on Apr. 4, 2022. [Online]. Available: https://opennetlab.org/

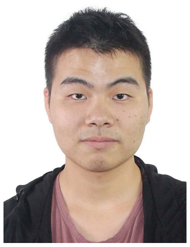

Qinkai Wang received the bachelor's degree in mechanical engineering from Zhejiang Sci-Tech University, Hangzhou, China, in 2020. He is currently working toward the master's degree with the School of Computer Science and Technology, University of Science and Technology of China, Hefei, China. His research interest focuses on Internet measurement. 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 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, Internet architecture, and network softwarization. He has published over 70 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 a Young Associate Editor of Springer Frontiers of Computer Science Journal. He is a member of the IEEE.

Lan Ding received the bachelor's degree in computer science from Anhui University, Hefei, China, in 2019. 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 5G and cloud systems.

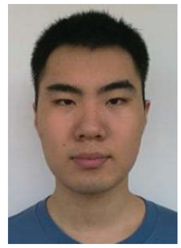

Xin Yu received the bachelor's degree in computer science from University of Science and Technology of China (USTC), Hefei, China, in 2021. He is currently working toward the master's degree with the School of Computer Science and Technology, USTC. His research interest focuses on performance evaluation of Internet systems.

Xinming Zhang received the BE and ME degrees in electrical engineering from 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 (USTC), Hefei, China, in 2001. Since 2002, he has been with the faculty of USTC, where he is currently a Professor with the School of Computer Science and Technology. 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, deep learning, and intelligent transportation. He has published more than 100 papers. He won the second prize of Science and Technology Award of Anhui Province of China in Natural Sciences in 2017. He won the awards of Top reviewers (1%) in Computer Science & Cross Field by Publons in 2019. He is a senior member of the IEEE.