

MonoScope: Automating Network Faults Diagnosis Based on Active Measurements

Waiting W. T. Fok[§], Xiapu Luo^{§‡}, Ricky Mok[§], Weichao Li[§], Yujing Liu[‡], Edmond W. W. Chan[†],
and Rocky K. C. Chang[§]

Department of Computing[§]
Shenzhen Research Institute[‡]
The Hong Kong Polytechnic University
{cswtfok|csxluo|cskpmok|csweicli|csrchang}
@comp.polyu.edu.hk

School of Computer[‡]
National University of
Defense Technology
China
liuyujing@nudt.edu.cn

Noah's Ark Lab[†]
Huawei Technologies
China
edmond.chan@huawei.com

Abstract—Network faults, such as router failures, cable outages, and configuration errors, could seriously affect network performance. In this paper, we use network faults very loosely, including those that will yield unfavorable end-to-end network performance, such as packet reordering and suboptimal routes. Diagnosing network faults on end-to-end paths is a very challenging problem, because it generally involves other domains. Even if it can be done, the process is very time consuming, because multiple sources of data, which are scattered in different places, are needed for such diagnosis.

In this paper, we consider the problem of making the fault diagnosis as automatic as possible. Based on coordinated active measurement from a set of end systems, we propose a procedure of detecting network faults and identifying their locations. Although this procedure cannot be fully automated for the time being, we show that some of the components could be automated, and we are automating them in a preliminary system called MonoScope. We demonstrate the efficacy of this procedure through several real-world cases that we have encountered in our four years of network monitoring experience.

I. INTRODUCTION

Network failures or faults are not uncommon in the Internet and data center networks today [10], [14], [22]. They could be resulted from software and hardware bugs, misconfigurations, accidents, natural disasters, and even deliberate acts. In this paper, we use network faults very loosely, including those that will yield unfavorable end-to-end network performance, such as packet reordering and suboptimal routes. For example, a load balancer generates a high packet reordering rate, and an ISP discriminates a customer by assigning a suboptimal path to its routes, thus resulting in a poorer performance as compared with other “similar” customers.

The possible impact of a network fault includes disruption in network connectivity and performance degradation. In the former, diagnosis can be performed based on routes [14]. In the latter the network connectivity may not be disrupted, but the end-to-end performance is affected. Since disruption in network connectivity also affects network performance, in this paper we consider fault diagnosis based on end-to-end path performance and route tracing. This diagnosis problem is very challenging, because it generally involves other domains. Even if it can be done, the process is very time consuming, because multiple sources of data, which are scattered in different places, are needed for such diagnosis.

Nowadays only large-scale network faults leading to huge financial losses or blackout in major cities are reported in mass media. Most events remain out-of-sight while affecting end-system network services. The systems available for the edge to diagnose performance problems, such as Scriptroute [21], M-Lab [2], and Speedtest [3], restrict the diagnosis to some pre-selected remote servers. Therefore, they cannot be used for diagnosing arbitrary paths. Individual installation of active measurement tools, such as perfSONAR [4] and HTTP/OneProbe [17], also does not help a user to effectively diagnose network faults for end-to-end paths.

Based on our four years of experience on coordinated active measurement from a set of end systems, we propose in this paper a procedure of detecting network faults and identifying their locations. The key idea is to synchronize the end systems' path measurement to the same remote destination, so that a network fault in these paths could be localized by detecting its impact on their path performance. The procedure thus consists of three consecutive steps. The first is to detect performance changes in individual paths. The detected changes will then be clustered in the time domain. Those path changes occurring in a small time window are suspected to come from the same network fault. In the last step, the routes for those paths identified in the second step will be analyzed to localize the fault.

Although this procedure cannot be fully automated for the time being, we show in this paper that some of the components could be automated. Based on our preliminary system called MonoScope, the first two steps could be automated. But there are still a number of challenges to overcome to optimize their performance in terms of accuracy and detection time. However, due to the diversity of network faults, the last step will be extremely difficult to automate for a general-purpose system. This difficulty will be illustrated through the case of diagnosing a suboptimal path and a submarine cable fault.

The main contribution of this paper is a three-step procedure to diagnose faults for end-to-end paths based on coordinated active measurement from a set of end systems (in §II). By analyzing each step, we also discuss the feasibility of automating it in an operational system. We then demonstrate the efficacy of this procedure through several real-world cases that we have encountered in our four years of network monitoring experience (in §III). After discussing the related works in §IV,

we conclude the paper in §V.

II. METHODOLOGY

A. The scope of network monitoring

We have been conducting network monitoring for the Hong Kong Academic and Research Network (HARNET) for over four years. One of the most common activities is to diagnose network performance problems revealed by the measurement results. The HARNET connects mainly eight local universities, including our own university, through an optical ring, and a local ISP (which is called HARNET ISP hereafter) provides the Internet connectivity service. Besides connecting to HARNET, each university also subscribes additional service from other ISP. Being an academic network, their primary concern is the performance of reaching other academic networks, such as CERNET and Internet2.

As shown in Figure 1, the HARNET monitoring system consists of eight measurement systems, each deployed at each university. They are installed with HTTP/OneProbe (OneProbe@university) [17] and Tcptraceroute for measuring end-to-end network performance and tracing routes for over 50 web servers in Hong Kong, Asia, Oceania, Europe, North America, and Africa. The measurement results are connected and analyzed by Planetopus [16], a network measurement management system, and the results are then made available to the HARNET users. The web servers are located at academic networks and commercial sites. The academic sites were chosen by the HARNET users, and others were added later on to monitor network performance for specific regions, such as Africa during the 2010 FIFA World Cup in South Africa. Therefore, the network paths under the monitoring cover a wide geographical area and different types of networks.

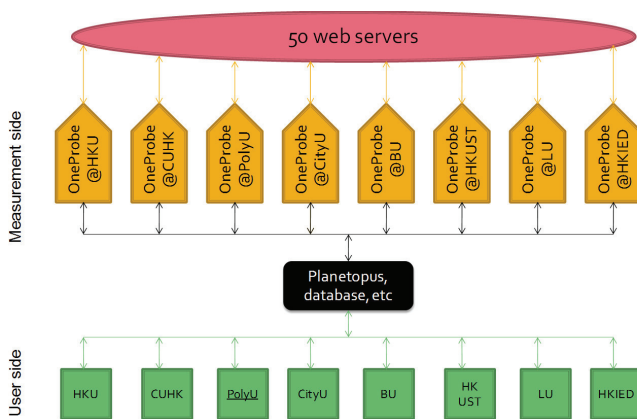


Fig. 1. The HARNET monitoring system.

B. Data collection

We diagnose network faults based on the end-to-end path performance data and routes collected by us and other publicly available network data, such as BGP routes. Moreover, an accurate diagnosis occasionally needs information from news feeds, such as reports on submarine cable outages [8], [7].

1) *Path performance*: To facilitate the diagnosis, it is highly desirable to obtain as much performance metrics as possible. Since we cannot deploy measurement agents at the remote nodes of the network paths under monitoring, our system is not able to obtain one-way network delay. However, we can obtain other metrics for one-way paths using HTTP/OneProbe [17] which dispatches a sequence of back-to-back packet pairs for measurement:

- 1) **ROUND-TRIP TIME (RTT)** Each probe consisting of two data TCP packets carrying an HTTP request will induce a response that consists of one or two TCP data packets carrying a (partial) HTTP response from the web server. The RTT is therefore measured by the arrival time of response packet subtracted by the departure time of the corresponding probe packet. Moreover, we use only the first packet's RTT, because the second packet's RTT may be biased by the first packet. We take the median from around 100 observations for a time-series analysis.
- 2) **PACKET LOSS RATE** We obtain forward-path and reverse-path loss rates based on the HTTP/OneProbe measurement. The response packets contain information for determining whether any of the packets in the probe and response is lost. We compute a packet loss rate for each direction by dividing the number of probes that suffer loss by the total number of probes.
- 3) **END-TO-END PACKET REORDERING RATE** We obtain forward-path and reverse-path packet reordering rates based on the HTTP/OneProbe measurement. The response packets contain information for determining whether the packet pairs in the probe are reordered or that in the response are reordered if two packets can be induced from the server. We compute a packet reordering rate for each direction by dividing the number of probes that experience reordering by the total number of probes.

2) *Coordinated traceroutes*: Our measurement systems also perform regular traceroutes for the forward paths. However, unlike other traceroute measurement, our system performs coordinated traceroute where the eight measurement systems perform traceroute to the same destination at the same time. The individual traceroutes, after resolving the IP aliases for routers, are fused into a single traceroute diagram as illustrated in Figure 2 where the eight measurement systems (UA-UH) trace the paths to a web server in `www.jp.apan.net`.

Moreover, previous works, such as [18], show that IP address, BGP prefix, and RTT can be used to determine the geographic location of the hops. They are especially useful for popular ASes, and we use the information to identify common forward paths in the next section. We also characterize route change from the continuous traceroute data obtained for each path using Jaccard distance [7].

3) *BGP data for reverse-path diagnosis*: Public sources, such as [5] and [6], offer BGP data for analysis. The BGP information for the measurement systems' ASes shows possible reverse paths. Moreover, BGP updates can be used to study the impact of network disruption events [11].

C. A fault diagnosis procedure

1) *Coordinated measurement*: The major mechanism that facilitates effective fault diagnosis is to conduct the path per-

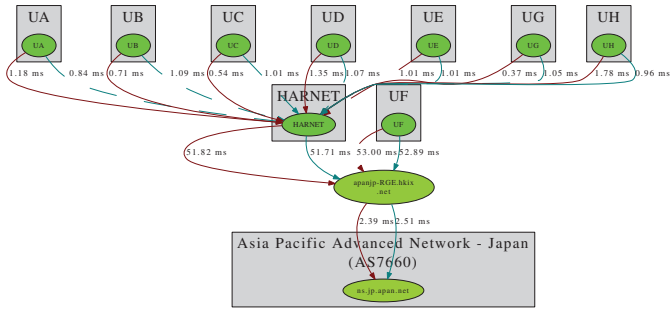


Fig. 2. A combined traceroute from the eight measurement systems to a web server in `www.jp.apan.net`.

formance and traceroutes in a coordinated manner, so that these measurements could be correlated for localizing the faults. Figure 3 depicts two scenarios of coordinated measurement. In Figure 3(a), each measurement node measures the paths to multiple destinations around the same time. Thus, if the performance of all the path deteriorates around the same time, the fault location should be at the segments close to the measurement node. In Figure 3(b), all measurement nodes measure the paths to a given destination around the same time. Thus, if the performance of the network segments close to the destination deteriorates, the performance degradation will show up in all the paths.

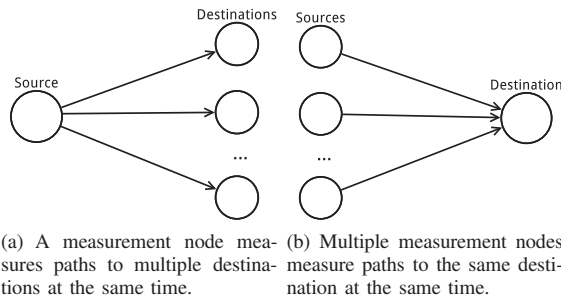


Fig. 3. Two scenarios of coordinated network measurement.

Our monitoring system setup follows Figure 3(b), because there is generally route diversity for the paths from the eight measurement nodes to a given destination. This route diversity could be used to diagnose problems at the destination network and at some common hops between the two ends of a path. At the same time, problems at HARNET and HARNET ISP can be diagnosed by path measurement launched by several measurement nodes at the same time. The measurement nodes are therefore time synchronized, and they are scheduled to measure the paths to the same web server around the same time.

2) *A three-step procedure:* Based on the measurement results obtained from the coordinated measurement, we propose three major steps to detecting and locating network faults:

- 1) **CHANGE-POINT DETECTION** The first step is to detect change points in the time series of network path performance under monitoring. The change points may reveal the onset of performance degradation through one or more

performance metrics. Besides, the forward-path routes are also analyzed for changes. This step therefore is applied to each network path independently.

- 2) **CHANGE-POINT CORRELATION** This step is to identify a similar change in multiple paths. If a change is detected for only a single path, we cannot possibly locate the fault. However, if a change is observed in multiple paths around the same time, these paths will be further analyzed in the next step for fault localization. The step therefore will provide such set of candidate paths, whenever identified, to the next step.
- 3) **FAULT LOCALIZATION** Given the set of candidate paths from the last step and corresponding traceroutes, this step tries to localize the faults. If the change detected in step 1 indicates a forward-path fault (e.g., a high forward-path loss rate), the common hops appearing in the set of forward paths are the candidate locations. However, if the network fault occurs on the reverse path, an AS-path analysis will be performed to identify the fault locations.

a) Change-point detection: Detecting changes in a time-series of RTT, packet loss rate, or packet reordering rate could be framed as a change-point detection problem which discovers time points at which properties of time-series data change. This problem is also known as anomaly detection and event detection. The change-point detection problem has been studied extensively for different types of problems, such as fraud detection in cellular systems and intrusion detection in computer networks. Common approaches to this problem include Sequential Probability Ratio Test, Cumulative Sum, and Exponentially Weighted Moving Average.

For detecting changes in forward-path routes, we quantify the degree of forward-route change using the Jaccard distance of AS-path and IP-path as we have done in [7]. In the time series of Jaccard distance values, a spike represents a significant route change. To analyze the reverse-path changes, we use the publicly available BGP data. The number of BGP updates for a specific AS can be used to quantify the routing dynamics of the network. Similar to the traceroute analysis, a spike in the time series represents a significant reverse-path route change of the AS.

b) Change-point correlation: This step attempts to cluster the changes detected on different paths from the first step. The rationale is that a network fault can induce performance changes on multiple paths if they traverse the location of the network fault. Therefore, we set a time window to collect changes occurred within the window and identify them to be the result of the same network fault. This set of candidate paths will be passed to the next step for fault localization. In our measurement setting, we use a window of 20 minutes for change-point clustering, because each path is measured every ten minutes, and 20 minutes is enough for obtaining at least one measurement.

c) Fault localization: This last step is to localize the network fault that induces performance change as precise as possible based on the set of candidate paths identified from the second step and the coordinated traceroute results. If there is only a single fault¹, it could be located on the

¹The multiple-fault case is much more challenging, and it is not considered in this paper.

forward path or reverse path. To diagnose a forward-path fault, we use the coordinated traceroute results to identify the hops/nodes that are common to the set of candidate paths. If such hops/nodes exist, the network fault is very likely located on them. Otherwise, it is possible that the fault is on the reverse path. Diagnosing a reverse-path fault is much more challenging, because performing reverse traceroute is still a daunting task even when a reverse traceroute facility [13] is available. Instead, we use the AS paths obtained from public BGP data. Similar to the coordinated traceroute, we obtain the ASes that are common on the candidate paths.

The network faults detected in our four-year experience with HARNET monitoring are commonly found in HARNET ISP and critical links in Internet, such as submarine cables. Moreover, we have detected network faults on the web server's side and CDN services, such as YouTube. In the next section, we will illustrate the three-step procedure using three real network faults discovered by the HARNET monitoring.

D. What can be automated?

In our current HARNET monitoring system, we have automated the path performance measurement and traceroute, but have not automated the fault diagnosis process. For the three-step procedure, our preliminary experience shows that the first two steps can be automated. However, the challenge is to decrease the false positives and false negatives, and to decrease the detection delay. Moreover, there are two types of time-series data: RTT, and the packet loss and reordering rates. In our experience, the bursty nature of the loss and reordering events makes it more challenging to detect them.

The third step will be the most difficult to automate. As will be illustrated in the next section, identifying a submarine cable fault requires manual analysis of the performance time-series and traceroutes, and ground truth about the submarine cable layout and the traffic going through them. This observation, however, is not surprising, because our network monitoring system is not designed specifically for monitoring submarine cables.

III. THREE REAL-WORLD CASES


We present in this section three cases where path performance degradation was observed and analyzed according to the methodology proposed in §II. In §III-A, we diagnose an ISP-scale packet reordering problem, which was introduced by a per-packet load balancer on a downstream link. In §III-B, the HARNET ISP performed reverse-path traffic engineering that affected some, but not all, universities' path performance. In §III-C, our system captured a network fault caused by Japan's earthquake in 2011 which damaged five submarine cables landed at Japan. It is important to note that the IP routes were unchanged in all three cases. Therefore, route analysis based on IP and BGP cannot diagnose these problems.

A. Anomalous reverse-path reordering rate


Step 1 - Change-point detection: The reverse-path packet reordering rate (RVR) of all paths were negligible on or before 17 October 2010. On 18 and 24 October, however, two RVR bursts that lasted for less than two days were found in the UF-Citibank path (where Citibank is a web server). For the

UF-Citibank path and some other paths, the diurnal patterns of high RVR became permanent since 29 October. The RVR reached as much as 45%.

Figure 4(a) shows the average RVR for the UF-Citibank path in October and November 2010. We plot the mean RVR using raw data and four aggregation levels. As expected, a higher aggregation will yield a smoother curve, thus facilitating a more accurate detection for changes. However, the downside is a longer detection time. The standard deviation of the average rates in Figure 4(b) further shows that the 2-hour aggregation is sufficient for smoothing the RVR. After smoothing the RVR, it is not difficult to see that the change in the RVR could be detected automatically by a time-series anomaly detection algorithm.



(a) The average rates.



(b) The standard deviation.

Fig. 4. Time series for the average and standard deviation of reverse-path reordering rates for the UF-Citibank path.

Step 2 - Change-point correlation: The surges in RVR were found not only in the UF-Citibank path, but also in the paths from other measurement nodes to 17 overseas destinations. As an example, Figure 5 shows the four-week RVR time series around 30 Oct 2010 for the 17 affected paths originated from UF and the paths from all eight measurement nodes (UA to UH) to NISSAN, another web server selected for path measurement. The RVR surges for these 25 paths occurred around the same time. Therefore, a window of reasonable size will be able to automatically cluster them together for further analysis in the next step.

Step 3 - Fault localization: Figure 6 illustrates the network connections of HARNET. We have obtained this information from an authoritative source and have also confirmed it by our traceroute. As shown, HARNET is peered with overseas academic and research networks, and connects to local networks via the Hong Kong Internet eXchange (HKIX).

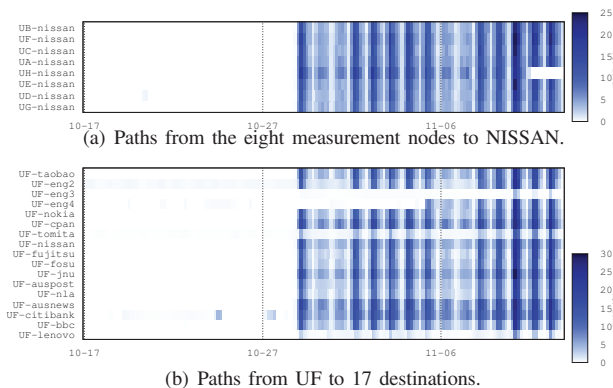


Fig. 5. Heat-map diagram of reverse-path reordering between 17 October 2010 and 13 November 2010.

Moreover, HARNET provides connectivity to other overseas network through the HARNET ISP.

Since we have the ground truth about the IP addresses of each router in HARNET, we can easily map IP paths to actual paths within HARNET. Packets routed through a specific router reflect the default forwarding path in use to other networks. We note that the paths that do not suffer from the high RVR problem all went through the links to HKIX and overseas academic and research networks (i.e., the dotted lines). On the other hand, the rest that were affected went through the link to HARNET ISP, and their common IP hops are within HARNET ISP. Therefore, the network fault responsible for the high RVR must be located at the link between HARNET and HARNET ISP or within HARNET ISP.




Fig. 6. A high-level view of HARNET's connectivity.

Impact of our diagnosis: We reported our finding to the HARNET users, and the information was subsequently passed to the operators in HARNET ISP who admitted that the high RVR rates were possibly generated by a per-packet load-balancing function at a router on the link between HARNET and HARNET ISP. On 13 May 2012, as promised by HARNET ISP, the router was replaced, and the RVR returned to normal ever since.

B. Poor latency performance from ISP's traffic engineering

Step 1 - Change-point detection: On 24 January 2011, we observed a 70 ms increase of RTT from two measurement nodes UA and UD to Citibank's in the US (www.citibank.com). Figure 7 shows the RTT, forward-path loss rate (FWL)

and reverse-path loss rate (RVL) for the UA-Citibank path in Jan 2011. The RTT before 24 January was stable but was increased sharply after that. This abrupt increase could be easily picked up by an abrupt change detection algorithm.




Fig. 7. RTT and loss rates of the UA-Citibank path in January 2011.

Step 2 - Change-point correlation: This step is to correlate the RTT surges found in the UA-Citibank and UD-Citibank paths with other change points detected from other paths. Figures 8(a) and 8(b) show the RTT of the paths from all eight measurement nodes to Citibank, and the paths from UA to the 17 overseas destinations. In Figure 8(a), besides the 70 ms increase for the UA-Citibank and UD-Citibank paths, other paths experienced only 20 ms increase in their RTTs. Moreover, Figure 8(b) shows the RTTs for the paths from UA to the 17 destinations. However, no change points from non-UA paths coincide with the change point in the UA path.

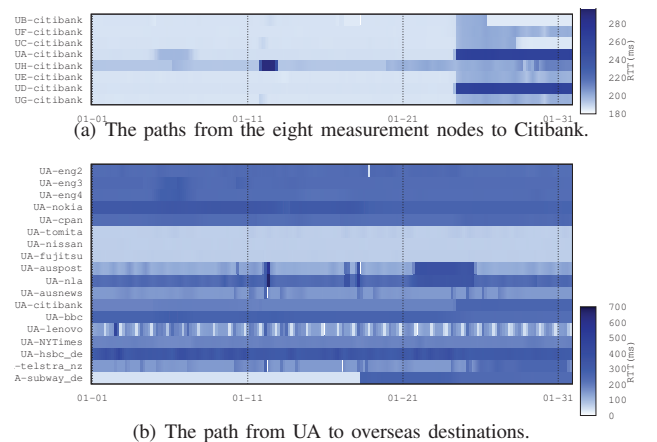


Fig. 8. A heat-map diagram of RTT in January 2011.

Step 3 - Fault localization: Since the two change points found for the UA-Citibank and UD-Citibank paths are not clustered with others, we investigate the underlying network fault by collecting information concerning the remote destination. We first quantify the change in the IP forward paths by the Jaccard distance of the traceroute results. Figure 9 shows the result for the UA-Citibank path. The Jaccard distance for the AS path is zero, therefore not shown here. Similar patterns were found for all eight paths from the measurement nodes to Citibank. Since the Jaccard distance is small around 24 January 2011, we rule out that the forward route change is responsible for the RTT surge.

We used undns [20] to resolve the geolocation of the hops obtained by traceroute. The forward path is denoted as FR in Figure 10. Except for five low-latency subpaths in the local




Fig. 9. Jaccard distance of the UA-Citibank IP-path from 10/1/2011 to 6/2/2011.

NEWT network (AS9381) which is HARNET ISP, all the eight AS paths are the same.

From [1], we found that the destination network CITIGROUP (AS32287) has two major upstream providers, AT&T (AS7018) and LEVEL3 (AS3356). We then collected reverse-path information through looking glasses (LG) of both ASes. A single LG at AT&T (AS7018) provides traceroute function. We used the function to identify a common reverse AS-path RR1 for all eight universities.

On the other hand, the BGP routes collected from LEVEL3 (AS3356) show interesting results. First of all, in several LEVEL3's LGs, the route to CITIGROUP (AS32287) directly went through LEVEL3's facilities in Dallas (ipcolo1.Dallas1). The two ASes seemed to be connected at the site. Using LEVEL3's LG in Dallas, two sets of AS paths labeled by RR_2 & RR_3 for the eight local universities were identified. RR_2 was assigned to UA and UD, and RR_3 to other universities.

We then performed traceroute at the LEVEL3's LG. We used undns again to infer the hops' geolocations. The underlying physical paths of RR_2 and RR_3 , as shown in Figure 10, went through very different geographic regions. While RR_3 passed through a trans-Pacific link (US West-coast-Japan-HK), RR_1 and RR_2 traversed a trans-Atlantic link (US East-coast-Europe) and an Eurasia link (Europe-Southeast Asia-HK). RR_2 was longer than RR_3 in terms of geographical distance, thus the propagation delay also.

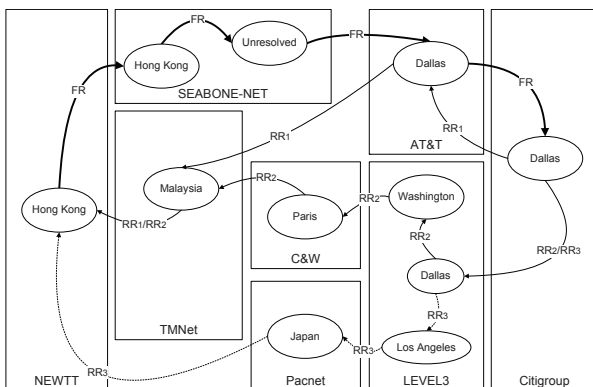


Fig. 10. Forward and reverse routes between UA-UH and Citibank.

Our path measurement matches LEVEL3's reverse paths RR_2 and RR_3 but not AT&T's RR_1 . It is because if the UA-Citibank and UD-Citibank paths went through the longer RR_2 , the RTT difference of 60 ms is reasonable. If all round-trip paths went through FR and RR_1 , then the RTT difference cannot be explained.

Impact of our diagnosis: We reported our findings to the HARNET users, and they were then forwarded to

HARNET ISP. At around 2:00am (HKT) of 23 July 2011, a few hours after the UA user filed a complaint to HARNET ISP, the RTTs for the UA and UD paths decreased significantly from 300 ms to 200 ms. A subsequent lookup at LEVEL3's Dallas LG revealed that RR_3 was assigned to all eight local universities.

C. Performance degradation due to submarine cable damage

Step 1 - Change-point detection: On 11 March 2011, a 9.0-magnitude earthquake occurred near Japan. Numerous aftershocks were recorded days after the main earthquakes. From the UA's path measurement, the reverse-path loss rate rose up to 5% in ten overseas paths for a five-day period immediately after the earthquake. A change of RTT pattern can also be found on 11 March for the UA-nla path. The RTT surged from 320 ms to 420 ms and was much more fluctuated than before. Figure 11 shows a heat map diagram, depicting the RTT and reverse loss rate for 10 out of 19 overseas paths measured from UA. The y-axis shows, from top to bottom, the measured paths to 3 hosts in Europe, 5 hosts in the US, 1 host in Australia, and 2 hosts in Japan. The time that 9.0 earthquake occurred is marked with a vertical arrow. These increases in RVL and RTT could be detected by a change detection algorithm.

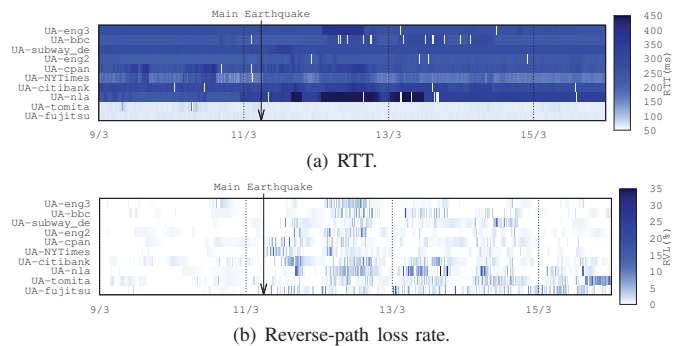


Fig. 11. A heat-map diagram of RTT and reverse-path loss rate from 09 March 2011 to 16 March 2011 for ten paths originated from UA.

Step 2 - Change-point correlation: The performance degradation detected in step 1 could be grouped into five clusters. They were in the form of bursts in RVL lasting for less than a day, once in each of the five consecutive days immediately after the earthquake. The first cluster on 11 March contained 5 paths, and four subsequent clusters from 12 March to 15 March include additional five paths, thus a total of 10 paths. These ten paths that observed the RVL surges were clustered together to diagnose the underlying network fault.

Step 3 - Fault localization: Tracerouting the ten paths showed that the traffic were forwarded to five different carriers after leaving HARNET ISP, and the destination hosts were located in ten different ASes. But around 40 paths to Hong Kong or overseas destinations did not observe the high RVL. Therefore, it is not caused by a network fault in the local network, HARNET ISP, or remote networks. Due to the scale of affected paths in different destination, the "discriminative" performance experienced by the universities in the last case is not exhibited here.

To diagnose this network fault with impact on a larger region, we first find the possible geolocation of the paths according to the available submarine cable systems. The rationale for this approach is that submarine cable systems are now the major media for carrying inter-continental data traffic. Figure 12 illustrates a simplified version of the global submarine cables systems. To make it simple, we intentionally miss out cables connecting islands, such as Iceland, Greenland, and Papua New Guinea. We have observed that the cable systems are highly redundant and concentrated in coastal cities of developed areas, such as North America, Europe, and Asia. For example, network traffic between Asia and Europe can go through either:

- 1) The submarine cables running between Asia and Africa through Middle East, and Mediterranean cables to Europe;
- 2) The transpacific cables between Asia (Japan inclusive or not) and the North America, continental cable networks through West and East coast of the North America, and transatlantic cables to Europe;
- 3) Same as above but through Australia before reaching the North America.




Fig. 12. A simplified map of worldwide submarine cables.

We then shortlist possible geographic paths by comparing end-to-end RTT with the estimated delay of each geographic path. The propagation delay component of an RTT path serves the baseline of the RTT. On top of that other components include processing delays in network devices and end hosts, and queuing delays in routers during peak hours.

The actual geographical path and the length of the submarine cable links are not public information. We can therefore only estimate the shortest path and minimum latency. Instead of conducting a full study on network delay of the Internet, we propose a workaround. We first estimate the intra-region delays as the propagation delays of the most direct path, including land and submarine cables, of the two far-most points within a region. We inflate the end-to-end geographical delay by a sum of the intra-region delay of intermediate regions on the path and half of the delays in source and destination regions.

Table I summarizes the minimum RTTs of inter-regional links (white) and possible intra-regional links (grey). If no link is available between two regions, the cell is blacked out. For example, the RTT of the UA-BBC path is 220 ms. For the aforementioned three possible paths running between Asia and Europe, their estimated RTTs are:

- 1) Asia-Middle East-Europe: $30 + 200 + 30 = 260$ ms;
- 2) Asia-North America-Europe: $30 + 150 + 60 + 65 + 30 = 335$ ms;
- 3) Asia-Australia-North America-Europe: $30 + 65 + 40 + 130 + 60 + 65 + 30 = 420$ ms.

Therefore, the first path is likely the actual geographical path for the UA-BBC path.

TABLE I. MINIMUM RTTs (IN MS) OF INTER-REGIONAL LINKS (WHITE) AND AVERAGE RTTs (IN MS) OF INTRA-REGIONAL LINKS (GREY).

Region	EU	Afr	Asia	JP	Aus	NA	SA
EU	60	90	200			65	
	Afr	150	200				65
		Asia	60	45	65	150	
			JP	20	90	90	
				Aus	40	130	
					NA	60	65
						SA	150

Finally, we infer which regional paths were the origin of the performance degradation by analyzing the shortlisted geographic paths of the ten paths identified in step 2. Table II presents the shortlisted geographic paths that carry the traffic sourced from Hong Kong and ended at remote networks located in Japan, Australia, Europe, and the US. Each possible path is numbered in an ascending order of the number of links used, e.g., path(2) for HK-Europe consisting of two links: Asia-N.America and N.America-Europe. Although we cannot conclude which regional links were affected, the links that were used by more than two affected paths are good candidates. They include Asia-Japan, Japan-N.America, and Asia-N.America.

TABLE II. POSSIBLE GEOGRAPHIC PATHS BETWEEN NODES IN HONG KONG AND OVERSEAS NETWORKS.

Links	Japan	Australia	Europe	US
Asia-Japan	◇1	◇24	◇3	◇2
Japan-N.America	×	◇4	◇3	◇2
Asia-Australia	×	◇1	×	×
Japan-Australia	×	◇2	×	×
Asia-N.America	×	◇3	◇2	◇1
Australia-N.America	×	◇34	×	×
Asia-Africa	×	×	◇1	×
Africa-Europe	×	×	◇1	×
N.America-Europe	×	×	◇23	×

◇ = Link possibly used by at least one path,
 × = None used the link

Impact of our diagnosis: Besides the analysis in step 3, we have also analyzed BGP data from RIPE [5] and RouteViews [6]. We have observed that among the paths from 128 different ASes to UA's network, 31.54% of them passed through FLAG (AS15412). It was a crucial provider for transmitting data traffic from the Internet to UA. After the earthquake, at 12:00 am on 11 March 2011, the RIB of RouteViews showed that there was a sharp increase in the number of AS paths going through FLAG. About 22,000 AS paths were re-routed to FLAG for transiting the traffic to other users. We can infer that this BGP rerouting behavior worsened the congestion of FLAG. The UA paths passing through FLAG were therefore expected to experience high packet loss and long queuing delays.

IV. RELATED WORK

Turner et al. [22] reported the statistics about the failure events in a large regional network consisting of over 200 routers. Their results show that the causes include from hardware, software, power, configuration, and so on. They also included the impact of the causes among which software and power have a higher impact. In this paper, our measurement

covers submarine cable outages which are not covered in [22]. Moreover, we consider a broader meaning of network faults, including ISP's network configurations that cause very high packet reordering rates and suboptimal routes.

Another related problem is to detect network outages in the Internet. Both Qian et al. [19] and the Hubble system [12] use a similar approach of actively sending probes to the IPv4 address space. Their objective is to determine how many addresses are unreachable due to outages and cluster them to events. The objective of our system, on the other hand, is to detect and localize network faults that will affect the path performance which generally will not deteriorate to the point of unreachable. Therefore, the network faults considered in this paper cannot be detected by the methods in [19] and [12].

Detecting and avoiding route failures is another frequently addressed problem in the literature, e.g., [19], [12], [23], [9], [14]. In particular, Zhang et al. [23] diagnosed routing disruptions associated with any large networks using a collaborative probing launched from end systems. Our three-step procedure is in fact similar to theirs on the block level, and coordinated traceroute (called collaborated traceroute) is also used for their diagnosis. However, our goal is to diagnose network faults responsible for performance degradation, and route disruption is just one of the possible causes for performance change.

Instead of launching active probing, some systems rely on passive analysis to detect network faults. However, such systems usually require much information that is not accessible to end users. The system proposed by Kompella et al. requires SNMP data along with the information about Shared Risk Link Groups and their relationships [15]. Based on log information including configuration files, syslog messages, administrator notices, and BGP data, Turner et al. investigate the causes and the impact of network faults in the CENIC network [22].

V. CONCLUSIONS

In this paper, we proposed a three-step procedure to diagnose network faults, which encompass all types of anomalies, based on coordinated active path measurement. We have devised the three steps based on our experience on fault diagnosis in HARNET. The first step, where change points are detected, could be automated with existing change-detection algorithms. Step 2 could also be made automatic by clustering the change-points detected on multiple paths using a moving time window. However, the most complicated analysis for fault localization in step 3 cannot be automated easily. As illustrated in the three case studies, step 3 could involve analyzing reverse paths based on traceroute and BGP routes, and RTT of inter-regional links. We are currently in the process of automating the first two steps in our preliminary system called MonoScope.

ACKNOWLEDGEMENTS

This work is partially supported by grants (ref. no. H-ZD91, H-ZL17) from the Joint Universities Computer Centre of Hong Kong, and a grant (ref. no. G-YK26) from The Hong Kong Polytechnic University and a Tier-2 project (ref. no. GHP/027/11) from the Innovation Technology Fund in Hong Kong. We also thank the reviewers for their careful reading of the manuscript and comments.

REFERENCES

- [1] The CAIDA AS Relationships Dataset, 2011-01-16. <http://www.caida.org/data/active/as-relationships/>.
- [2] Measurement Lab. <http://www.measurementlab.net>.
- [3] Ookla. <http://speedtest.net/>.
- [4] perfSONAR. <http://www.perfsonar.net/>.
- [5] RIPE. <http://www.ripe.net>.
- [6] University of Oregon RouteViews project. <http://www.routeviews.org/>.
- [7] E. Chan, X. Luo, R. Chang, W. Fok, and W. Li. Non-cooperative diagnosis of submarine cable faults. In *Proc. PAM*, 2011.
- [8] R. Chang, E. Chan, W. Li, W. Fok, and X. Luo. Could ash cloud or deep-sea current overwhelm the Internet? In *Proc. USENIX HotDep*, 2010.
- [9] N. Feamster, D. Andersen, H. Balakrishnan, and F. Kaashoek. Measuring the effects of Internet path faults on reactive routing. In *Proc. ACM SIGMETRICS*, 2003.
- [10] P. Gill, N. Jain, and N. Nagappan. Understanding network failures in data centers: measurement, analysis, and implications. In *Proc. ACM SIGCOMM*, pages 350–361, 2011.
- [11] Y. Huang, N. Feamster, A. Lakhina, and J. Xu. Diagnosing network disruptions with network-wide analysis. In *Proc. ACM SIGMETRICS*, 2007.
- [12] E. Katz-Bassett, H. Madhyastha, J. John, A. Krishnamurthy, D. Wetherall, and T. Anderson. Studying black holes in the Internet with Hubble. In *Proc. NSDI*, 2008.
- [13] E. Katz-Bassett, H. V. Madhyastha, V. K. Adhikari, C. Scott, J. Sherry, P. Van Wesep, T. Anderson, and A. Krishnamurthy. Reverse traceroute. In *Proc. USENIX NSDI*, 2010.
- [14] E. Katz-Bassett, C. Scott, D. Choffnes, I. Cunha, V. Valancius, N. Feamster, H. Madhyastha, T. Anderson, and A. Krishnamurthy. LIFE-GUARD: Practical repair of persistent route failures. In *Proc. ACM SIGCOMM*, 2012.
- [15] R. Kompella, J. Yates, A. Greenberg, and A. Snoeren. IP fault localization via risk modeling. In *Proc. USENIX NSDI*, 2005.
- [16] W. Li, W. Fok, E. Chan, X. Luo, and R. Chang. Planetopus: A system for facilitating collaborative network monitoring. In *Proc. IFIP/IEEE Intl. Symp. Integrated Network Management*, pages 911–924, 2011.
- [17] X. Luo, E. Chan, and R. Chang. Design and implementation of TCP data probes for reliable and metric-rich network path monitoring. In *Proc. USENIX ATC*, 2009.
- [18] V. Padmanabhan and L. Subramanian. An investigation of geographic mapping techniques for Internet hosts. In *Proc. ACM SIGCOMM*, volume 31, pages 173–185, 2001.
- [19] L. Quan, J. Heidemann, and Y. Pradkin. Detecting Internet outages with precise active probing. Technical Report ISI-TR-2012-678b, USC/Information Sciences Institute, 2012.
- [20] N. Spring, R. Mahajan, and D. Wetherall. Measuring isp topologies with rocketfuel. In *Proc. ACM SIGCOMM*, pages 133–145. ACM, 2002.
- [21] N. Spring, D. Wetherall, and T. Anderson. Scriptroute: a public Internet measurement facility. In *Proc. USITS*, 2003.
- [22] D. Turner, K. Levchenko, A. Snoeren, and S. Savage. California fault lines: understanding the causes and impact of network failures. In *Proc. ACM SIGCOMM*, 2010.
- [23] Y. Zhang, Z. Mao, and M. Zhang. Effective diagnosis of routing disruptions from end systems. In *Proc. USENIX NSDI*, 2008.