

Could Ash Cloud or Deep-Sea Current Overwhelm the Internet?

Rocky K. C. Chang, Edmond W. W. Chan, Weichao Li, Waiting W. T. Fok, and Xiapu Luo

Department of Computing, The Hong Kong Polytechnic University

Hunghom, Hong Kong, SAR China

Abstract

In this paper, we are initially set out to evaluate how the ash cloud from the Eyjafjallajökull volcano impacted the Internet traffic. Based on our path measurement for European websites, we have observed significant congestion whose onset seems to coincide with the period of disrupted air traffic. However, after expanding the scope of investigation, the path congestion was in fact caused *indirectly* by a submarine cable fault which received far less attention than the ash-cloud news. The paths under our monitoring were overloaded by taking on additional traffic diverted from the faulty cable. To probe further into the path congestion, we use loss pairs to measure the correlation between packet loss events and delay.

1 Introduction

In recent years, there were quite a few natural phenomena that rocked the Internet. Notably, earthquakes and typhoons were responsible for a few Internet blackouts and cable failures. In August 2009 alone, three submarine cables were damaged by an undersea mudslide caused by the typhoon Morakot, and another was severed by an Earthquake near Hualien [5]. Submarine cables could also be damaged by deep-sea current, climate change, ship anchors, and fishing activities [6].

The eruption of ashes from the Eyjafjallajökull volcano in Iceland [1] seriously affected the air traffic to and from the UK and many European countries mainly during 15-23 April. While the media attention was drawn to the impact on air travelers, there were also reports on its cascading effects on the Internet. For example, it has been reported that the traffic from Citrix' GoToMeeting web-conferencing service was doubled in the first week after the airspace was closed [3]. According to Akamai's report, the overall web traffic in the northern Europe was well above normal levels [4]. But France Telecom and Deutsche Telekom did not notice significant increases in mobile network usage [4].

In this paper, we are initially set out to measure the impact of the ash cloud on the Internet performance. Measuring the actual impact is very difficult, because there is a lack of effective monitoring infrastructure. Assessing relevant information and statistics from various network service providers is also a formidable task, as they are usually considered business secrets. In the previous incidents involving cable failures (and network attacks as well), the Border Gateway Protocol (BGP) routing information [7] and flow-based monitoring systems [8] were used to report the extent of the impact. However, the route for a congested path may still remain stable. The flow-based monitoring, on the other hand, can only be used for measuring paths whose traffic passes through the monitoring system, but it cannot be used to measure arbitrary paths.

In section 2, we first report the results of our active measurement of some paths to European websites, which has been conducted routinely. Our measurement "accidentally" captured severe congestion on paths to some commercial sites during the period of air traffic disruption. A more careful investigation, however, leads us to discover a submarine cable fault that was responsible for the network congestion. In section 3, we use loss pairs to further characterize the loss behavior of the congested paths. In section 4, we conclude this paper with a summary of lessons learnt from the journey of this investigation.

2 Unplanned measurement

As part of our network measurement project, we have been monitoring Internet paths from eight locations at Hong Kong to websites in Asia, Europe, Australia, and the US. Our measurement uses OneProbe [10] for the continuous monitoring and Tcptraceroute [13] for analyzing the forward paths. We do not use ping which is not accurate for measuring the path quality experienced by data packets and cannot support a high sampling rate. OneProbe is designed for measuring the quality of paths

to any TCP-based service. Since it uses TCP data packets for probing, the measurement is able to capture the performance experienced by data packets. Moreover, the probes are generally not affected by firewalls and other middleboxes. Each probe measurement reports several path metrics, including the round-trip time (RTT), packet loss rates, packet reordering rates, and capacity. The last three metrics can further be distinguished for forward and reverse paths.

2.1 The overall results

Each of the eight measurement probers, referred to as UA, ..., UH, measures eight websites located in the UK and Europe. Four of them are commercial sites and the other four academic sites. There are therefore a total of 64 paths. For each path, OneProbe measures continuously for one minute and then repeats the measurement after idling for nine minutes. During the measurement, OneProbe dispatches a sequence of Poisson-modulated probes with a mean rate of 5 Hz and a packet size of 576 bytes. This section mainly reports the median RTT and average loss rates computed from a one-minute measurement. The packet reordering events were scarce, therefore not included for further analysis. In the next section, we also report loss-pair statistics.

We first summarize the overall measurement using heat-map time series for the period of 8 March-21 May 2010. Heat maps are very useful for visualizing contrasts which are the differences in the path metrics for our measurement. For the heat-map time series, we divide the measurement period into one-hour bins and calculate the median value for each bin for the six samples obtained during the hour. After computing the bin values, we divide the range from 0 to the maximum value into nine levels linearly and assign each level with a different gray level. A darker color refers to a higher value for the metric. We have plotted for RTT (in seconds), forward-path loss rate (in %), and reverse-path loss rate (in %).

Figure 1 shows the time series for paths to two commercial sites in the UK and Finland, referred to as ENG and NOK, respectively. Note that the RTTs surged on around 14 April for both paths and lasted till 3 May. The forward-path losses became very significant also around 14 April. Before that, the paths did not see significant loss rates. The patterns for the reverse-path losses, however, are very different. There were heavy losses during 23-28 April for both paths. By comparing the two sets of graphs, it is not difficult to conclude that the two paths suffered from RTT surge and heavy losses around the same time. Therefore, it is highly likely that the contributing sources were the same for both paths.

Figure 2 shows the time series for paths to another commercial site and an academic site, both in the UK, referred to as BBC and TEIN2, respectively. Unlike the

ENG and NOK paths, no RTT surge and heavy forward-path losses are observed for these two paths. The lighter color from UF is the result of the prober's unsuccessful measurement. Similar to TEIN2, the other three academic paths did not see any RTT surge and heavy losses. However, BBC also suffered from heavy reverse-path loss for the period of 23-28 April. We do not include the results for the other commercial website at the UK, because the measurement was not entirely successful.

2.2 A root-cause analysis

Since the path performance for both NOK and ENG is similar, we choose only one path to NOK for further analysis. Figure 3(a) shows the time series of RTT and loss rate for UB→NOK. We also show the time series for UB→BBC in Figure 3(b). Although the BBC paths did not suffer the same heavy forward-path loss as the NOK and ENG paths, Figure 3(b) shows that there are actually two loss peaks starting on 14 April. A traceroute analysis reveals that at the onset of the path congestion observed on around 14 Apr 2010 07:18:00 GMT, the forward paths to ENG, NOK, and BBC went through the same provider FLAG (AS15412). However, on 16 Apr 2010 07:39:00 GMT, the BBC paths changed from FLAG to GLOBEINTERNET (AS6453). Both FLAG's and GLOBEINTERNET's next AS hop were London IX. After the change, the BBC path saw very stable RTT and insignificant packet losses. Another point is that the reverse-path loss suffered by the BBC path was more severe than the NOK and ENG paths.

At the first glance, the onset of the path congestion matches quite well with the period of air traffic disruption started on 15 April. However, a careful analysis of their correlation reveals several problems:

1. The onsets of the path congestion and air traffic disruption do not entirely match. A travel warning due to the ash problem was announced on 14 April, but the announcement of shutting down the British airspace was first made in the morning of 15 April.
2. Some of the peak loss rate and RTT occurred on weekends, and these network traffic could not possibly be introduced by an increased usage of Internet for conducting business meetings.
3. Path congestion can still be observed at the end of the measurement period, although the RTT was much less than before. However, the air traffic basically returned to normal at the end of April. Although there were some intermittent airspace closures in May, their impact is much less than the prolonged closure in April.

Prompted by the three problems above, we have decided to expand the scope of our investigation. We first analyze where the forward-path congestion occurred by comparing the traceroutes of the 64 paths. Since only

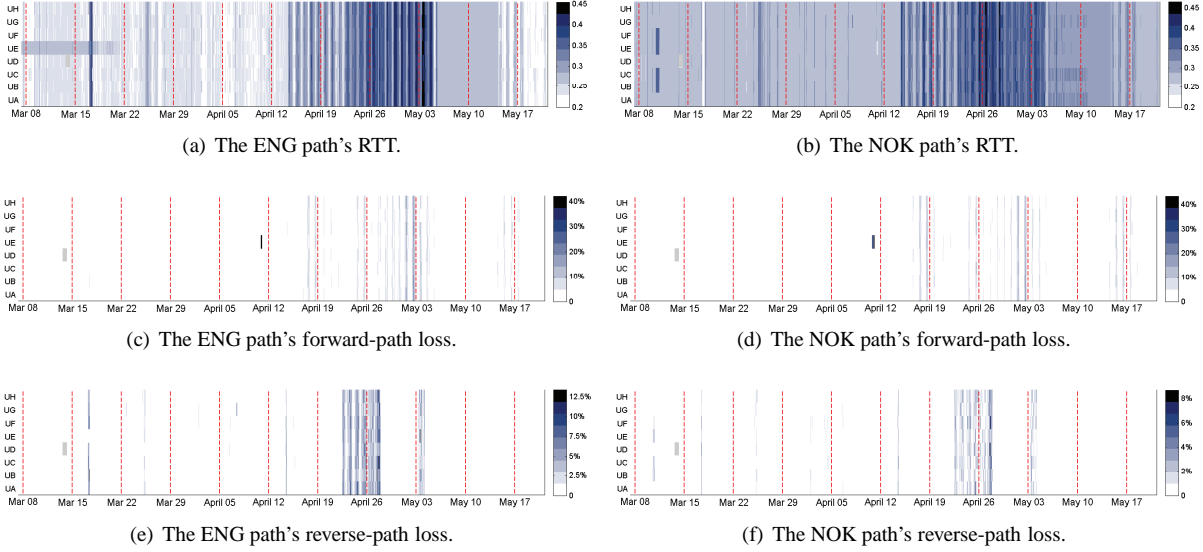


Figure 1: Heat-map time series for the ENG path's and NOK path's RTT and packet loss rates.

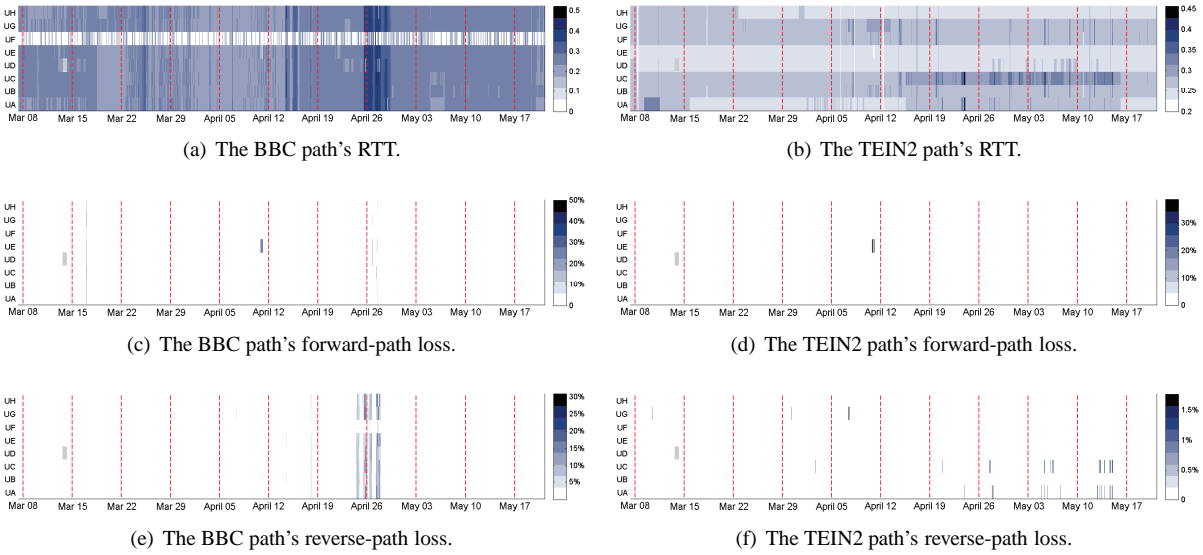


Figure 2: Heat-map time series for the BBC path's and TEIN2 path's RTT and packet loss rates.

the ENG and NOK paths experienced the congestion, we perform “intersection” operations on the two sets of ASs on the paths. The only AS that is common to both sets of paths is the FLAG network which, however, does not appear in other paths (except for the BBC paths before 16 Apr 2010 07:39:00 GMT). After making the inference, our subsequent investigation discovers that the SeaMeWe-4 cable, connecting Asia and Middle, was damaged near Italy on 14 April (the exact time is usually not reported) [11]. This incident, which did not receive too much attention, actually caused limited Internet connectivity to and from Middle east and Asia, including Pakistan and India [12], and the cable was restored at the

end of April [2].

The Internet traffic affected by the SeaMeWe-4 cable fault was shifted to satellite, land based networks or two other submarine cables to Europe: the SeaMeWe-3 and FLAG Europe-Asia cables [12]. Therefore, a more plausible explanation for the path congestion for the ENG and NOK paths is the result of taking on additional traffic diverted from the SeaMeWe-4 cable¹. Moreover, on 05 May 2010 12:03:00 GMT, our traceroute analysis shows that the ENG and NOK paths were changed from

¹In private communication, our explanation presented to FLAG was not rejected.

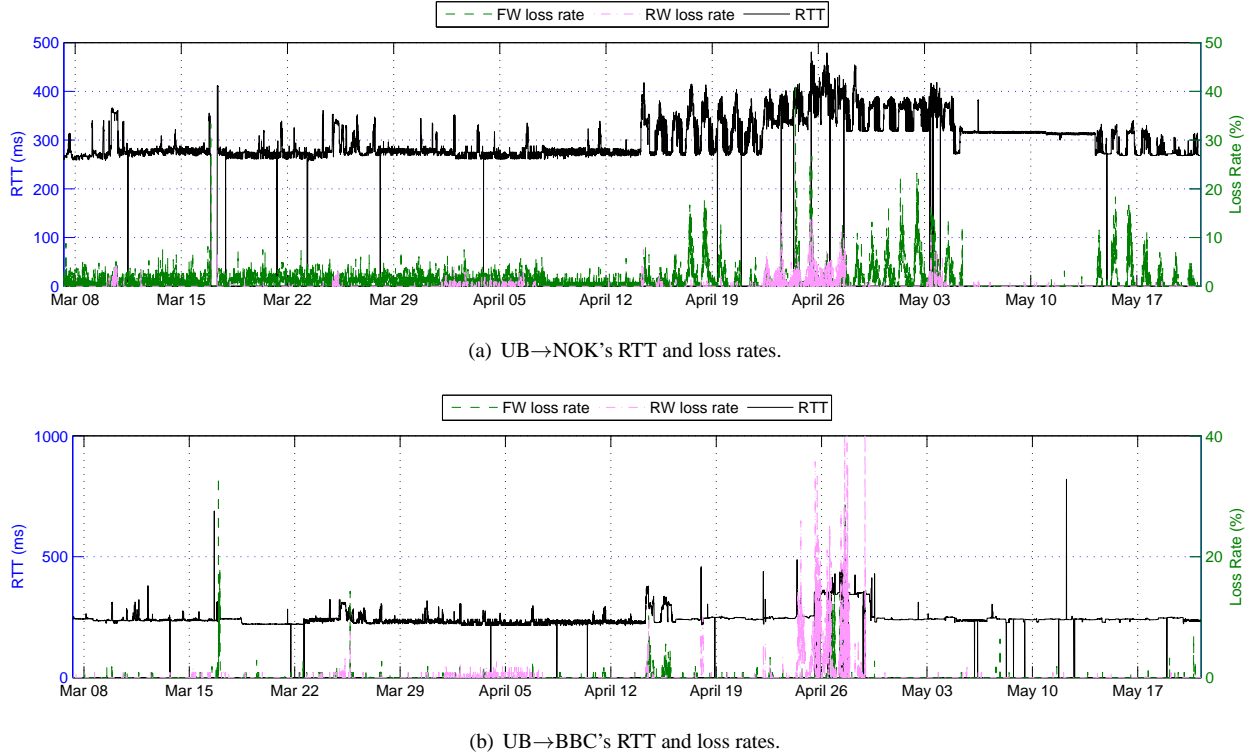


Figure 3: Time series for UB→NOK's and UB→BBC's RTT and loss rates.

a westbound direction to an eastbound direction (relative to Hong Kong). The new path to the London IX was through the West Coast of the US². The new route enjoyed very stable RTT performance and low losses. However, on 14 May 2010 05:53:00, the paths were changed back to the old routes. Due to the lack of time, we have not further analyzed the cause for the reverse-path losses for the ENG, NOK, and BBC paths.


3 Loss-pair measurement of the congested paths

Besides measuring the loss and delay individually, it is also useful to correlate a packet loss event and the delay that would have been experienced by the lost packet. We measure the delay-loss correlation using loss pairs. A packet pair is referred to as a *loss pair* [9] if exactly one packet (the first or second) in the pair is lost. If the two packets traverse the path close to each other, then the residual packet's delay could be used to infer the lost packet's delay. OneProbe measures two types of loss pairs (first or second packet being the lost packet). We use 1P-LP (2P-LP) to refer to the loss pairs for which the first (second) probe packet is lost on the forward path and use 1R-LP and 2R-LP similarly for the reverse-path loss pairs.


²In private communication, some major backbone outages of FLAG network occurred at Alexandria in late April.

Figure 4 plots the residual packets' RTT (i.e., loss-pair RTTs) of the four types of loss pairs for UA→NOK. We also superimpose the residual packets' RTT with the RTT time series obtained from all packet pairs' first packet (labeled by 'All') to identify which parts of the RTT time series the loss-pair RTTs were located. As shown, the loss-pair RTTs provide different path signatures for the forward and reverse paths. While almost all the loss-pair were clustered on the RTT peaks, a significant number of them also occurred in the forward path during the period of 8 March-13 April.

We divide the entire measurement period into four subperiods based on the RTT characteristics shown in Figure 4: (i) 8 March-13 April, (ii) 14 April-21 April, (iii) 22 April-4 May, and (iv) 14 May-21 May; and plot the distributions of the path queueing delays for 2P-LP and 2R-LP in Figures 5(a) and 5(b), respectively. The *path queueing delay* of a loss pair is defined as the residual packet's delay subtracted by the minimum observable delay. We divide the entire measurement period into one-day bins and obtain the minimum observable RTT in each bin to compute the path queueing delay for all the loss pairs found in that bin. Figure 5 shows that almost all the path queueing delays obtained during subperiods (ii) and (iii) for both forward and reverse paths are greater than (i). Therefore, the packet loss events during the subperiods (ii) and (iii) were the result of more intense con-




(a) Forward path.




(b) Reverse path.

Figure 4: Time series of loss-pair RTTs for UA→NOK.

gestion in the path. For subperiod (iv), the forward-path queuing delay was still significantly higher than (i) but was much reduced compared with (ii) and (iii).



(a) 2P-LP.



(b) 2R-LP.

Figure 5: Distribution of the path queuing delay for UA→NOK.

4 Conclusions and lessons learnt

Going back to the question posted in the title, our measurement results did not observe any correlation between the air traffic disruption due to the ash cloud and the performance of the paths under our monitoring³. However, our results showed significant network congestion which was the result of carrying additional traffic diverted from a faulty submarine cable. The impact of this kind cannot be observed by routing information, because the path under measurement did not use the faulty path.

This paper is part of our efforts of monitoring critical network infrastructure due to natural phenomena and intentional attacks. Through this experience of detect-

³In private communication, we learned that a major tier-one provider did not experience any congestion.

ing and analyzing the path-quality degradation, we have learnt the following lessons.

1. Measuring end-to-end network paths actively is necessary for monitoring critical network infrastructure. The BGP information, though useful for detecting link outages, generally cannot reveal degradation of path performance. On the other hand, it is very difficult to diagnose the network problems in this paper using passive measurement.
2. Correlating the measurement of multiple paths for the same destination is very useful for locating the problematic AS. Although we have measured only eight websites in the UK and Europe, we were able to identify the AS in which the main congestion occurred. Using multiple probes located close to one another on the AS topology could facilitate the cross-path correlation analysis.
3. Correlating IP routes with end-to-end path measurement is also necessary for diagnosing path problems. However, extra caution must be exercised in drawing sound conclusions from the analysis. In our case, for example, we observed some IP route changes inside the FLAG network close to the onset of path congestion. Despite this temporal correlation, it turns out that these route changes were not responsible for the path congestion.

Acknowledgments

This work is partially supported by a grant (ref. no. ITS/355/09) from the Innovation Technology Fund in Hong Kong and a grant (ref. no. H-ZL17) from the Joint Universities Computer Centre of Hong Kong.

References

- [1] 2010 eruptions of Eyjafjallajökull. Available from http://en.wikipedia.org/wiki/2010_eruptions_of_Eyjafjallajökull.
- [2] UPDATE Repairs to SEA-ME-WE 4 cable successfully completed.
- [3] B. Bulik. Video conferencing players see volcanic cloud's silver lining. Available from http://adage.com/digital/article?article_id=143394.
- [4] M. Campbell. Web traffic, video meetings surge as flights grounded. Available from <http://www.bloomberg.com/apps/news?pid=20601109&sid=aLIB2jOAJF2A>.
- [5] R. Clark. Subsea cable breaks 'as bad as 2006 quake'. Available from <http://www.telecomasia.net/content/subsea-cable-breaks-bad-2006-quake>.
- [6] M. Green, S. Drew, L. Carter, and D. Burnett. Submarine cable network security. Presented at APEC, April 2009.
- [7] G. Huston. Cable quakes. Presented at RIPE 54, May 2007.
- [8] Y. Kitamura, Y. Lee, R. Sakiyama, and K. Okamura. Impact of Taiwan earthquake, December 26th, 2006. In *Proc. APAN Network Research Workshop*, August 2007.
- [9] J. Liu and M. Crovella. Using loss pairs to discover network properties. In *Proc. ACM Internet Measurement Workshop*, 2001.

- [10] 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 Annual Tech. Conf.*, 2009.
- [11] A. Mir. SEAMEWE-4 under repair. Available from <http://telecompk.net/2010/04/19/seamewe-4-under-repair/>.
- [12] Nabeel. Underwater Internet cables crippled near Mediterranean. Available from <http://nabtron.com/under-water-internet-cables-crippled-near-mediterranean/1802/>.
- [13] M. Toren. tcptraceroute. <http://michael.toren.net/code/tcptraceroute/>.