Distributed Systems and Natural Disasters

BitTorrent as a Global Witness

Zachary S. Bischof, John S. Otto, and Fabián E. Bustamante
Northwestern University
{zbischof, jotto, fabianb}@eecs.northwestern.edu

ABSTRACT

Peer-to-peer (P2P) systems represent some of the largest distributed systems in today's Internet. Among P2P systems, BitTorrent is the most popular, potentially accounting for 20-50% of P2P file-sharing traffic. In this paper, we argue that this popularity can be leveraged to monitor the impact of natural disasters and political unrest on the Internet. We focus our analysis on the 2011 Tohoku earthquake and tsunami and use a view from BitTorrent to show that it is possible to identify specific regions and network links where Internet usage and connectivity were most affected.

1. INTRODUCTION

On Friday, March 11th, 2011 at 2:46 PM local time, a 9.0 magnitude earthquake triggered a massive tsunami off the coast of Tohoku, Japan. Once the tsunami reached land, tsunami runup reached as high as 40 meters above sea level and up to 10 km inland, leaving 561 km^2 of Japan inundated. The Tohoku earthquake is one of the top 5 most powerful earthquakes on record [18]. In addition to the significant loss of life, the event caused extensive damage to infrastructure – ranging from buildings and roads to power plants and submarine communication cables.

In this paper, we argue that globally popular distributed systems provide a unique perspective to monitor the impact of this and similar natural and man-made phenomena on communication networks. We show the potential of this approach by analyzing the impact of the Tohoku event using the view provided to us through BitTorrent, a popular P2P file-sharing system.

To evaluate the impact of this particular event, we focus our analysis on application and network data from peers located in or communicating with others through Japanese networks. We study three months of data from January 1st, 2011 to March 31st, 2011, which covers the event as well as its aftermath. This data is contributed by peers using the popular Vuze BitTorrent client [17] and running the Ono [3] and NEWS [4] extensions we have made available previously.

Beyond application specific data (such as number of connected peers and download/upload rates per torrent), our dataset includes per-connection data such as transfer rates and the results of traceroutes to random subsets of connected peers. The dataset does not include any information that can identify the downloaded content. We combine this unique perspective with several datasets that characterize the severity of the earthquake and tsunami, by geography, to understand the correlation between real-world, application and network-level events.

We first examine the impact of the Tohoku earthquake and tsunami on BitTorrent usage. Our data reveals the expected relationship between application usage and proximity to the event epicenter – with a total BitTorrent usage drop of up to 80%. We then use our collected traceroute measurements to characterize the impact of the event on the underlying network. We identify a subset of popular, high-latency network-level hops as potential submarine cables and analyze their dynamics before, during and after the event. Our analysis corroborates and extends, to the best of our knowledge, the reported set of links that were potentially affected by the earthquake.

After discussing our approach and related work in Sec. 2, we describe the data sets we study in Sec.3 and present our analysis of system and network dynamics that coincide with the Tohoku earthquake and tsunami in Sec. 4. We discuss ongoing work and conclude in Sec. 5.

2. APPROACH AND RELATED WORK

This paper studies the impact of natural disasters as reflected in the state and dynamics of a globally distributed application. We do this by combining the application and network level view of ubiquitous distributed applications with publicly available data related to the event of interest. We have previously applied this approach to analyze BitTor-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ACM SWID 2011, December 6, 2011, Tokyo, Japan.
Copyright 2011 ACM 978-1-4503-1044-4/11/0012 ...$10.00.

rent's view of the anthropogenic Egypt and Libya Internet disconnections from early 2011 [1].

This approach is closely related to previous work on the use of globally distributed applications for early detection and localization of network anomalies [4]. Others have analysed the network impact of real-world events such as natural disasters and political unrest (e.g. [5, 10, 11, 14, 16]) leveraging data from in-network monitoring, targeted measurements and publicly available data. Recently, Schulman and Spring [15] present an interesting analysis of the impact of small weather events on residential host failures.

3. DATA SETS

For our study of the impact of the March 2011 Japanese earthquake and tsunami on the Internet, we integrate measurements from the event[1] with data from BitTorrent users spread around the world. The following paragraphs describe both datasets.

3.1 Earthquake and Tsunami

To quantify the geographic distribution of earthquake severity across Japan, we use a publicly available data set of displacement – how far the land moved – across 1,218 points in Japan [9]. We then map each of these points to one of the 47 prefectures (subnational jurisdictions) in Japan. We are able to map over 95% of these points to a prefecture; the remaining 5% are located underwater and are therefore not considered. The mapped points cover each of the 47 prefectures with a median of 20 measurement points, ranging from 7 to 160 points. We average the data points for each prefecture, and normalize by the maximum prefecture average.

We also map tsunami data as reported by the Japan Meteorological Agency [8] to prefectures. The data set contains data for 148 locations (at city-level granularity), where each point measures the maximum increase in water level. Using a combination of automated and manual techniques, we are able to map 95% of these measurements to a prefecture; the unmapped data points were either islands with small populations or ambiguously named locations with minimal reported tsunami impact. The mapped points cover 37 of the 47 prefectures. We use the same methodology as with the earthquake data to obtain the normalized tsunami impact by prefecture.

3.2 BitTorrent

We consider a 3-month period of data collected by users of the BitTorrent peer-to-peer file-sharing system, spanning from January 1st, 2011 to March 31st, 2011. This period covers the Tohoku event as well as a number of weeks previous to the event; we use the data from before the event to characterize "typical" behavior prior to the earthquake. We also study almost three weeks of data following the earthquake to identify presistent changes in BitTorrent and the network that could have resulted from the event. We divide our data set following the different time periods shown in Tab. 1, and refer to the subset of our data that we use in each of our analyses.

Label	**Start**	**End**
Before	1 Jan., 00:00	10 Mar., 5:00
Day before	10 Mar., 05:00	11 Mar., 5:00
Event	*11 Mar., 05:00*	*12 Mar., 5:00*
Day after	12 Mar., 05:00	13 Mar., 5:00
After	13 Mar., 05:00	31 Mar., 00:00

Table 1: Time periods in our data set. All times are in UTC. The 5:00 cutoff coincides with the earthquake, which occurred at 5:46 UTC.

The BitTorrent data we analyze is contributed by users of the Ono [3] and NEWS [4] plugins for the Vuze BitTorrent client [17], collectively representing over 1.4 million installations. With the user's consent, the plugins anonymously report usage statistics and the results of passive monitoring and active measurements. These data allow us to see system usage patterns as well as the conditions of the underlying network. For the analyses in this paper we use two types of data: BitTorrent population statistics (who our users connect to) and traceroute measurements to a subset of connected peers (the paths that traffic takes through the network).

Specifically, our 3-month data set is comprised of globally diverse reports from 452 K unique peers located in 182 different countries and 3,307 networks. Although only a small fraction of these instrumented peers are actually located in Japan (0.2%), a much larger portion (22.0%) have traffic through Japan and are therefore valuable to our analysis. Further, 66% of the networks and 95% of the countries that our dataset covers have peers that generate traffic relevant to our analysis.

3.2.1 *BitTorrent population data*

For our analysis of BitTorrent usage, we count the number of unique connected BitTorrent peers (by IP address) across all instrumented peer for each hour. Given the median BitTorrent session length of 3.5 hours [12], this sampling rate is sufficient to capture a continuous, representative sample of peers participating in the BitTorrent system while avoiding potential aliasing problems. By mapping IP addresses to geographic regions,[2] we can determine the size of the BitTorrent population in each geographic region over time.

Our dataset includes peers in all 47 of Japan's prefectures in numbers that closely correlate with prefecture population. Figure 1 plots the relationship between the prefecture's population size and the average number of BitTorrent peers seen per hour for that prefecture. There is a strong positive correlation ($r = 0.90$) between prefecture population and

[1] All Japan map data comes from [2] and [7].

[2] GeoLite City `http://www.maxmind.com/app/geolitecity`

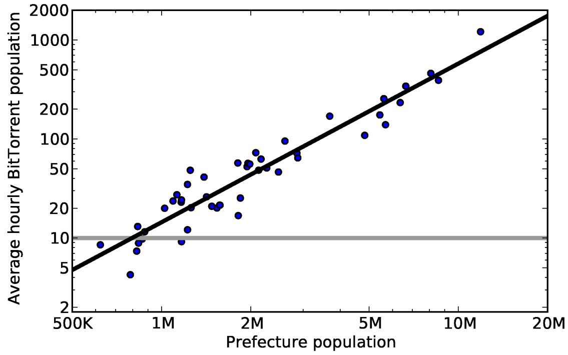

Figure 1: Scatter plot of prefecture population versus the average number of BitTorrent peers seen per hour in that prefecture. There is a strong positive correlation between both metrics ($r = 0.90$). The black line shows the result of linear regression for the two variables. Our dataset provides a representative view of BitTorrent usage throughout Japan. The gray line indicates the minimum BitTorrent population threshold for inclusion in our analyses – 10 peers per hour.

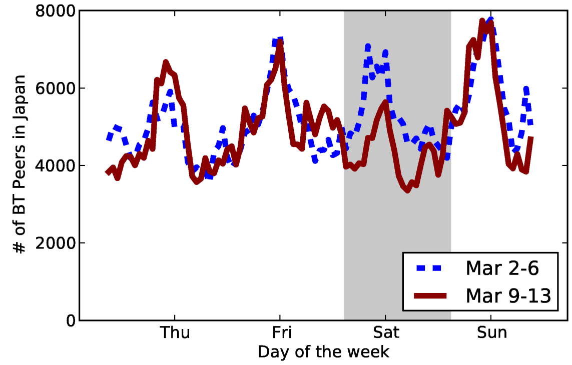

Figure 2: Overall number of peers seen online in Japan, for the week of the event (March 9-13) and before the event (March 2-6). Week day ticks on the x-axis mark the start of that day, and all times are in local time, JST. The gray shaded region starts at the time of the earthquake on Friday and lasts for 24 hours. Diurnal patterns are perturbed for about 24 hours after the earthquake.

the number of BitTorrent peers our dataset finds in that region.

Of the 47 prefectures, we exclude 6 that have fewer than 10 BitTorrent peers seen per hour as these prefectures introduce significant noise into our results. After dropping the smallest prefectures, our BitTorrent population data overlaps with the earthquake data in 41 prefectures (87%) and with the tsunami data in 33 prefectures (70%). In total, these prefectures account for the vast majority of Japan's population (95.9%).

3.2.2 Traceroute data

To see the path that BitTorrent traffic takes through the network, each instrumented peer reports the result of regularly run traceroute measurements to a random subset of the peers to which they connect. Altogether, for the observation period we have nearly 116 M traceroute measurements, of which 2.1 M (1.8%) come from, go to, or traverse a Japanese network. We use this subset of 2.1 M traceroutes in our analysis of the impact of the event on the Internet.

4. ANALYSIS

In this section, we study the impact of the Tohoku earthquake and tsunami first at the application layer in terms of BitTorrent usage, and then on routing dynamics at the network layer.

4.1 Impact on BitTorrent

We begin by analyzing the relationship between BitTorrent usage and the impact of the Tohoku earthquake and tsunami, by Japanese prefecture.

Using our data set, we estimate the number of BitTorrent users online in a geographic region. Basically, each connected peer is an indication of an active BitTorrent user. We compile the set of unique IP addresses seen across all our instrument peers, which gives us a representative sample of online BitTorrent users. We apply this analysis to a sequence of time intervals to build a time series of the number of users online in a region. For this analysis, we only consider BitTorrent peers in Japan, which we identify using IP-to-geo mappings.

Figure 2 shows the number of Japanese peers seen in each hour across two periods: the week before (Mar 2-6) and the week of the earthquake (Mar 9-13). Before the earthquake, we see diurnal patterns of usage where most users are online near midnight. Furthermore, the Thursday and Friday night peaks are approximately equal. In the 24 hours after the earthquake, however, we observe a 25% reduction in online peers relative to the previous week.

We next examine how the impact on BitTorrent (i.e. population decrease on the day of the event) varies across Japanese prefectures. Specifically, we compare each prefecture's average number of BitTorrent users across the "Day before" and "Event" periods (Tab. 1). Figure 3 shows that about 30% of prefectures saw more than a 20% reduction in the number of users online on the day of the earthquake. The rest of the prefectures were relatively unaffected, with at most a 20% decrease in the number of users online. This is consistent with other reports indicating that some regions of Japan were more significantly impacted by the earthquake and tsunami than others [8, 9].

Figure 4 maps the prefectures with the strongest BitTorrent impact (left), and those that were most affected by the earthquake (center) and tsunami (right). Darker colored

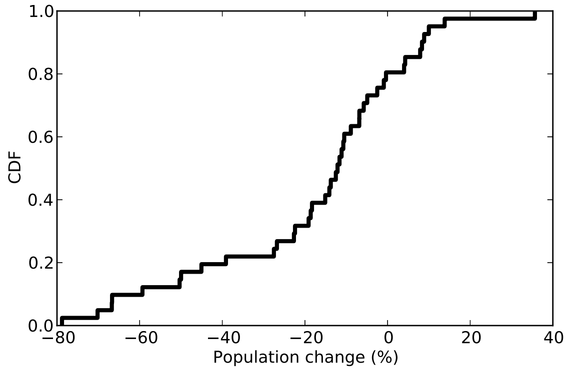

Figure 3: Distribution of percent change in average BitTorrent population for each prefecture, comparing the 24 hours prior to the earthquake and the 24 hours after the earthquake. Most prefectures were relatively unaffected, but about 30% of them had significant declines of over 20% in BitTorrent usage.

prefectures indicate more significant impact. On each map, the "X" marks the epicenter of the earthquake. Overall, the prefectures that had the largest percent decreases in BitTorrent population tended to be closest to the epicenter (Fig. 4a). The map of earthquake intensity shows a similar pattern (Fig 4b). The tsunami predominantly affected prefectures located on the east coast of Japan (Fig. 4c). Overall, the prefectures with the largest declines in BitTorrent population appear to have been affected most significantly by either the earthquake or the tsunami.

We quantify the relationship between impact on BitTorrent usage and the earthquake or tsunami intensities across prefectures by performing linear regression and correlation analyses. Figure 5a shows a scatter plot between tsunami intensity and BitTorrent impact across prefectures. We find a significant correlation ($r = -0.62$), meaning that regions more affected by the tsunami tended to have larger declines in BitTorrent usage. Interestingly, when we apply a similar analysis to earthquake intensity versus BitTorrent impact (Fig. 5b), we find an even stronger correlation ($r = -0.81$).

We believe that one possible reason for the difference in correlation strength between earthquake and tsunami intensity is due to the granularity of our analysis. The earthquake had a relatively uniform intensity across all points in each prefecture – all cities were roughly affected equally. However, in the case of tsunami intensity, the cities on the coast were most likely to be impacted, while inland cities may not have been affected by the tsunami at all. For instance, Fukushima (the prefecture most affected by the tsunami) is a outlier in Fig. 5a at 100% tsunami intensity, but only -45% BitTorrent population change. The relatively small BitTorrent impact would make sense because only a small fraction of Fukushima cities are near the coast; most cities in the prefecture were unaffected by the tsunami.

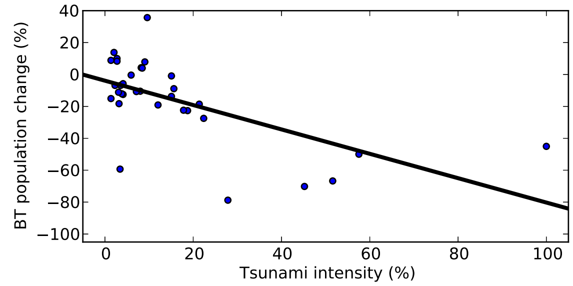

(a) Tsunami intensity vs. BitTorrent impact

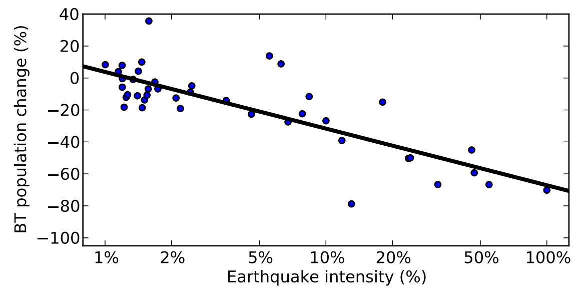

(b) Earthquake intensity vs. BitTorrent impact

Figure 5: Scatter plot of tsunami and earthquake intensity versus BitTorrent impact by prefecture. The black line shows the linear regression between the two variables. For earthquake intensity, this is computed on a log scale. There is a negative correlation between changes in BitTorrent population and tsunami impact ($r = -0.62$); the more severe the tsunami was, the greater the decline in BitTorrent usage. For earthquake intensity, we find an even stronger correlation ($r = -0.81$).

Overall, our application-level analysis of the impact of the earthquake and tsunami on BitTorrent usage revealed the expected relationship: prefectures that were most affected by the disaster had the most significant reductions in BitTorrent usage the day after the event. While we found that earthquake intensity was better correlated with BitTorrent usage impact than the tsunami intensity, a finer-grained analysis at the city level might reveal a stronger correlation between changes in BitTorrent usage and tsunami intensity.

4.2 Network impact

In this section, we examine the impact of the earthquake and tsunami on the flow of traffic at the network level. We identify popular, high-latency IP-level hops in our traceroute dataset as a heuristic for detecting submarine cables, and focus on the dynamics of the popularity of such links during and after the event. We consider the frequency of a link's appearance in traceroutes to be a proxy for the overall number of BitTorrent connection across that link.

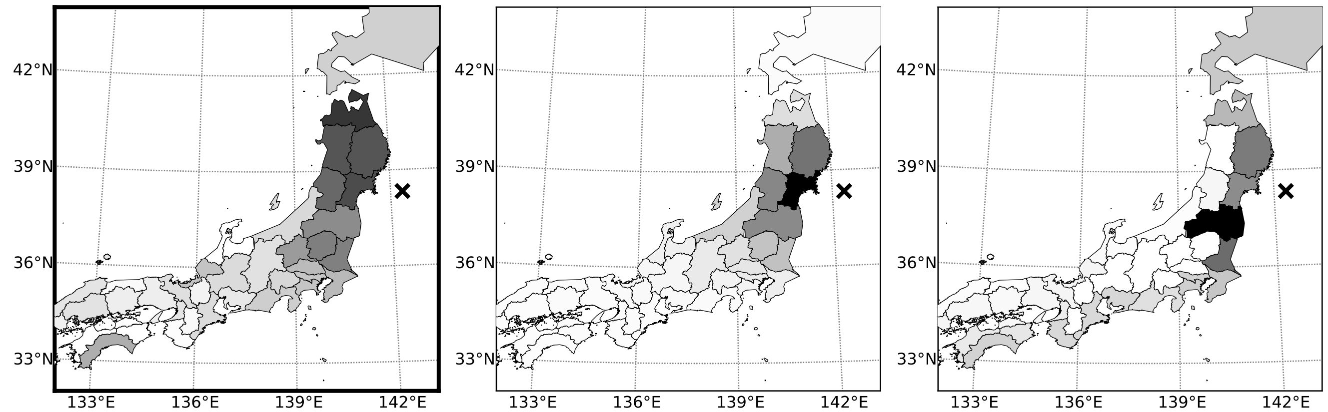

(a) BitTorrent change (b) Earthquake intensity (c) Tsunami intensity

Figure 4: Maps of Japan showing, by prefecture, percentage change in BitTorrent population size (left), earthquake intensity (center) and tsunami intensity (right). The "X" marks the earthquake epicenter. Darker shaded regions were more severely impacted by the earthquake or tsunami, or had a larger reduction in the number of BitTorrent peers seen online in each region.

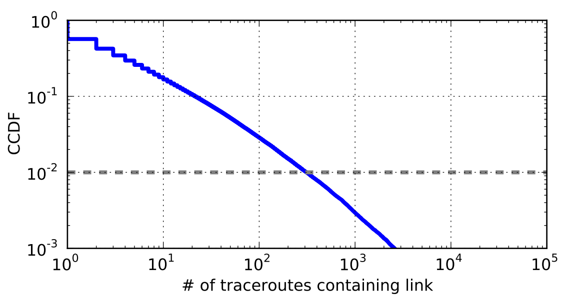

Figure 6: CCDF of the number of times each "link" (pair of consecutive IP addresses in a traceroute) appears. We selected the top 1% most popular links as a threshold (dashed gray line) to identify the set of the most popular links.

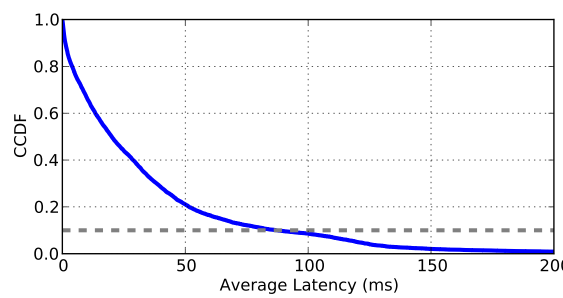

Figure 7: CCDF of the latency across the top 1% of popular links (from Fig. 6). To limit our analysis to long-haul links, we selected the top 10% links (greater than 87 ms latency).

4.2.1 Selecting popular long-haul links

To study the dynamics of the network as a result of the earthquake, we focus on long-haul submarine cables given their potential impact on all communication in and out of Japan. To identify long-haul submarine cables we rely on two heuristics:

- *Popularity:* Traceroute paths converge at submarine cables because there are relatively few alternate routes. Therefore, these links appear more frequently in our dataset.
- *High latency:* Long-haul links (e.g. transcontinental submarine cables) connect distant locations and therefore have relatively high propagation delays. These delays manifest in traceroute latency measurements.

Figure 6 shows the distribution of the number of times a link appears in the set of traceroutes from the "Before" period. Given that there is a relatively small set of submarine cables compared with the total number of links in the Internet, we expect that links over submarine cables will appear more frequently. Of the links in this set of traceroutes, we select the top 1% most popular links. The dotted line in Fig. 6 shows this threshold.

Next, we examine the distribution of link latency of this set of 1% most popular links. We calculate the round trip time over each hop using traceroute data. We select the minimum of the 3 latencies for each hop to reduce the impact of intermittent spikes in latency. Figure 7 shows the distribution of average link latency. From this subset of links, we select the top 10% by highest latency, which corresponds to an average round trip time of at least 87 ms.

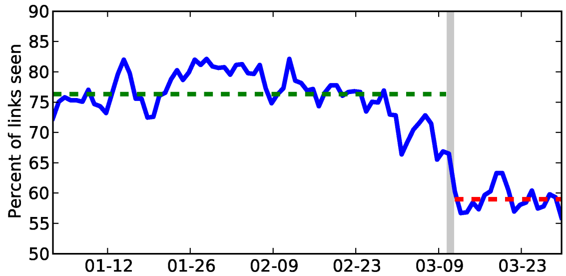

Figure 8: The percentage of popular long-haul links that appear each day. The shaded region represents the day of the earthquake, March 11. The dotted lines before and after the earthquake represent the average percentage of these links that appeared per day during that period. The average drop from 76% to 59% after the disaster indicates that a significant fraction of previously popular links were less frequently used afterwards.

This threshold is comparable to the 80 ms latency of the fastest trans-Pacific cable, PC-1 [13].

4.2.2 Dynamics of popular links

Figure 8 shows the percentage of popular links seen each day, between January and March. Before the earthquake, approximately 76% of these links appear at least once each day. However, from March 11th until the end of March, we see only 59% of these links on a given day. This drop in the percentage of links seen implies that traffic has shifted to other routes in the network.

Figure 9 shows four scatter plots contrasting the popularity of a link between two days. The left plots compare a typical Thursday with a typical Saturday; the plots on the right are from the day before and after the earthquake. Points near the light dotted line ($y = x$) represent links that appear about the same number of times each day, while points at $y = 0$ are links that do not appear at all on the second day. Links that appear at most half as often (or at least twice as often) on the second day are in the shaded regions below $y = \frac{1}{2}x$ (or above $y = 2x$).

Figures 9c and 9d focus on the subset of links near the origin in Figures 9a and 9b, respectively. These figures more clearly show a larger number of links appearing less than half as often (or not at all) after the earthquake, compared with the typical distribution. For instance, on Saturday, March 5th, 10.6% of the links appeared half as often as on March 3rd. However, on March 12th, after the event, 21.3% of links appear half as often as on the 10th – twice more than on a typical day. The white points in Fig. 9c and 9d represent links that did not appear on the Saturday of that week. Note that between the day before and after the earthquake, a significantly larger number of links disappared, compared to the number of links that disappeared the week before.

4.2.3 Shifting link popularity

In the previous section, we focused on links that were popular before the earthquake. Here, we apply the heuristics in Sec. 4.2.1 to the traceroutes in the "After" period of our dataset. We add these links to the set of popular, high-latency links from the "Before" period to also identify links with increasing popularity after the event.

Figure 10 shows the probability density function of the relative popularity of links in the "Before" and "After" periods. The majority of links occur with approximately the same frequency (i.e. near 1/1) before and after the event. However, we observe a cluster of links with a ratio of approximately 100/1 – these are links that seem to disappear after the event. There is also a small cluster of links near 1/100; these links either appear in our traceroutes for the first time or become significantly more popular after the earthquake. Also note that the largest peak in the center is slightly shifted to the left; we believe that this due to the increasing popularity of common links that remain operational after the event.

We aggregate these links by AS number and show information about several of these networks in Tab. 2. For each network, we compute the fraction of links whose popularity increased or decreased by an order of magnitude, or appeared or disappeared from our traceroutes after the disaster. We also corroborate these results with networks reported to be affected by the earthquake via news reports.

KDDI and NTT, two large Japanese telecom networks with the most links in our dataset, both had a significant fraction of their links decline in popularity or disappear from our traceroutes after the earthquake. This coincides with reports in the media that some of their submarine cables were damaged, which could result in re-routing along a different path and therefore a change in the set of popular links in our data set. In addition, for both of these networks we also identify links that increase significantly in popularity; these links may be along backup routes for the damaged cables.

We also found several networks that do not have links with significant changes in popularity due to their location or regional scope. Both Oi (a major Brazilian ISP) and one of NTT's regional networks (AS4713) had no links with significant changes in popularity. This is expected, since Oi does not, to the best of our knowledge, operate trans-Pacific cables, and regional networks are unlikely to communicate over long-distance links.

Our analysis at the network level shows a significant shift in the ASes belonging to KDDI, NTT (AS2914), Tata, IIJ, and China Unicom. In all of these ASes, a large number of popular links became an order of magnitude less popular; all of these networks saw at least 10% of popular links disappear from traffic through Japan, according to our traceroutes. While some links may have been damaged in the disaster, some shifts may have been a result of routing

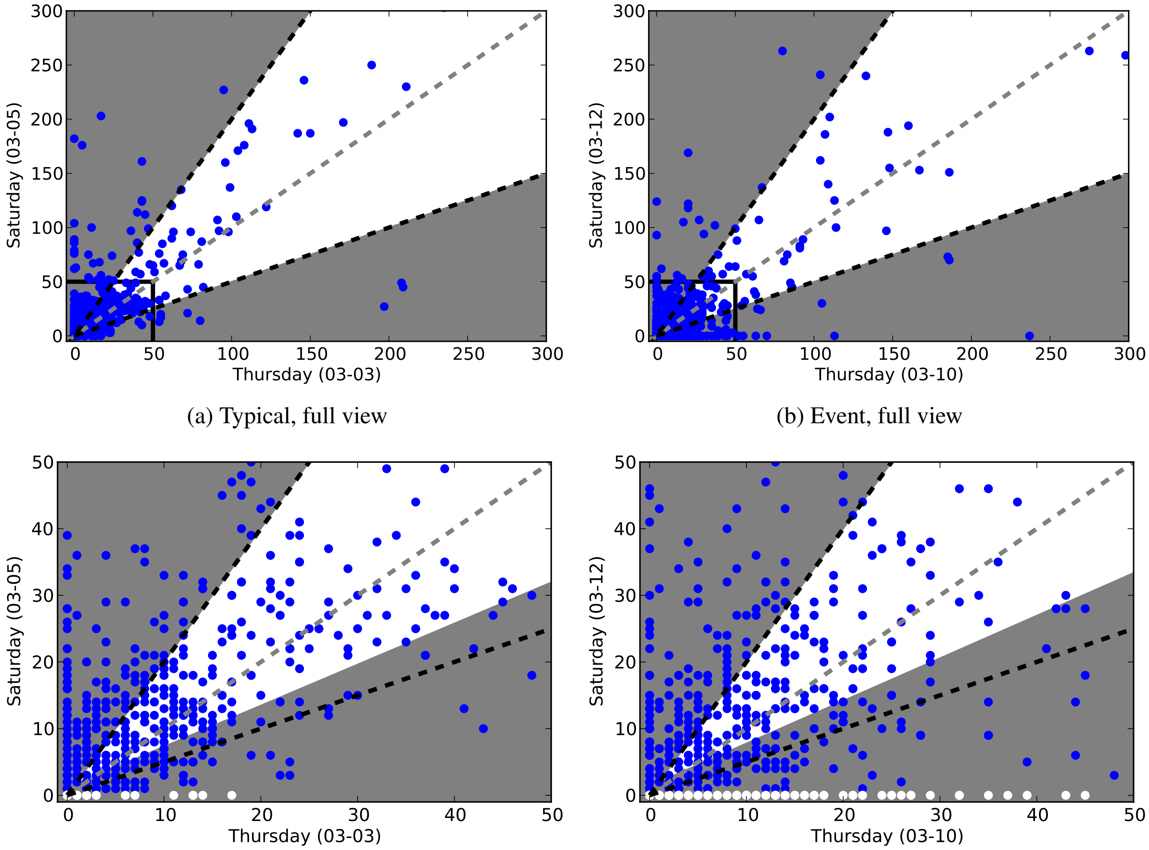

(a) Typical, full view

(b) Event, full view

(c) Typical, zoomed

(d) Event, zoomed

Figure 9: Scatter plots comparing the *number of times a link appears* on two days (x and y axes). The diagonal light gray line represents the consistent distribution, where the link is seen the same number of times each day. Links in the shaded regions appeared at least half as often in one of the days, representing a significant change in popularity. Top Left: comparison of links between a typical Thursday & Saturday. Top Right: comparison of the days before and after the event (Thursday 3/10 and Saturday 3/12). The bottom two figures are zoomed in to focus on links appearing less than 50 times on both days. The white points represent links that did not appear on Saturday of that week.

changes or if some traffic were no longer routed through Japan.

5. CONCLUSION

In this paper, we examined the effects of the Tohoku earthquake and tsunami on BitTorrent usage in Japan and began studying its impact on the underlying network. By leveraging the view of a popular P2P system, we find a geographic correlation between the impact of the disaster and a decrease in BitTorrent usage. Furthermore, low-level measurements performed at end hosts allow us to identify links that disappeared and changes in traffic routing.

Our analysis of the Tohoku event is an example of a more general approach to understand the impact of real-world events on Internet systems and networks from data provided by large-scale distributed systems. We have also applied this approach to identify country-wide Internet outages in Egypt and Libya.[3] As part of ongoing and future work, we are investigating other application- and network-level metrics of potential interest and identifying current world events for study.

6. REFERENCES

[1] AquaLab Blog: A View from the Edge. http://aqualab.cs.northwestern.edu/blog/.

[2] CHINA HISTORICAL GIS PROJECT. Tokugama Japan GIS, Demo Version, Feb 2004. http://www.fas.harvard.edu/~chgis/japan/.

[3] CHOFFNES, D. R., AND BUSTAMANTE, F. E. Taming the torrent: A practical approach to reducing cross-ISP traffic in peer-to-peer systems. In *Proc. of ACM SIGCOMM* (2008).

[4] CHOFFNES, D. R., BUSTAMANTE, F. E., AND GE, Z. Using the crowd to monitor the cloud: Network event detection from edge systems. In *Proc. of ACM SIGCOMM* (2010).

[3] http://aqualab.cs.northwestern.edu/blog/egypt-libya-peers.html

ASN	ISP	Links	Decreased	Disappeared	Increased	Appeared	Reported
2516	KDDI	354	15.8%	15.3%	6.8%	0.6%	[6]
2914	NTT	333	33.0%	26.7%	13.8%	3.9%	[6]
6453	Tata	93	36.6%	26.9%	9.7%	4.3%	–
2497	IIJ	65	18.5%	13.9%	27.7%	23.1%	–
3356	Level3	62	9.7%	6.5%	3.2%	0.0%	–
3549	GBLX	53	1.9%	1.9%	0.0%	0.0%	–
10026	Pacnet	36	11.1%	8.3%	11.1%	8.3%	[6]
7738	Oi (Telemar)	34	0.0%	0.0%	0.0%	0.0%	–
4837	China Unicom	33	36.4%	36.4%	0.0%	0.0%	[6]
701	Verizon Business	26	3.9%	0.0%	0.0%	0.0%	–
4713	NTT	9	0.0%	0.0%	0.0%	0.0%	–

Table 2: Comparison of the changes in the frequency of popular links among ASes. "Links" is the number of links in our set belonging to a particular AS. The "Decreased" and "Increased" columns show the percentage of links that decreased or increased in popularity by a factor of 10 before or after the disaster. "Disappeared" and "Appeared" are, respectively, the percentage of links that appeared in traceroutes only before or only after the day of the event. The "Reported" column lists a reference to reports of links known to be damaged in that network.

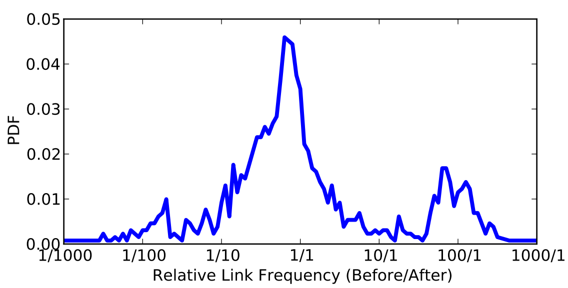

Figure 10: Probability density function of link popularity ratio between "Before" and "After" traceroute datasets. Most links appear with roughly the same frequency in both periods, corresponding to the peak near 1/1. However, we identify clusters of peaks near 100/1 and 1/100, which correspond to links that significantly decreased or increased in popularity between the two periods.

[5] DAINOTTI, A., SQUARCELLA, C., ABEN, E., CLAFFY, K. C., CHIESA, M., RUSSO, M., AND PESCAPÈ, A. Analysis of country-wide internet outages caused by censorship. In *Proc. of IMC* (2011). To appear.

[6] FLETCHER, O., AND OSAWA, J. Rush to fix quake-damaged undersea cables, March 2011. `http://online.wsj.com/article/SB10001424052748704893604576199952421569210.html`.

[7] GADM. Database of Global Administrative Areas. `http://www.gadm.org`.

[8] JAPAN METEOROLOGICAL AGENCY. Tsunami observations. `http://www.jma.go.jp/en/tsunami/observation_04_20110313180559.html`.

[9] JPL, NASA. Japan Quake Displacement, Japan, March 14, 2011. `http://geocommons.com/overlays/100267`.

[10] LABOVITZ, C. Craig Labovitz's Blog. `http://monkey.org/~labovit/blog/`.

[11] LABOVITZ, C. Security to the Core: Middle East Internet Scorecard (February 12-20). `http://asert.arbornetworks.com/2011/02/middle-east-internet-scorecard-february-12-%E2%80%93-20/`.

[12] OTTO, J. S., SÁNCHEZ, M. A., CHOFFNES, D. R., BUSTAMANTE, F. E., AND SIGANOS, G. On Blind Mice and the Elephant: Understanding the network impact of a large distributed system. In *Proc. of ACM SIGCOMM* (2011).

[13] Pacific Crossing - PC-1 Network. `http://www.pc1.com/pc1_network.cfm`.

[14] RENESYS BLOG. Japan Quake, March 2011. `http://www.renesys.com/blog/2011/03/japan-quake.shtml`.

[15] SCHULMAN, A., AND SPRING, N. Pingin' in the rain. In *Proc. of IMC* (2011). To appear.

[16] TeleGeography. `http://blog.telegeography.com/`.

[17] VUZE, INC. Vuze. `http://www.vuze.com`.

[18] WIKIPEDIA. 2011 Tōhoku earthquake and tsunami. `http://en.wikipedia.org/wiki/2011_Tohoku_earthquake_and_tsunami`.