IIR

Internet
Infrastructure
Review

Feb.2019
Vol.41

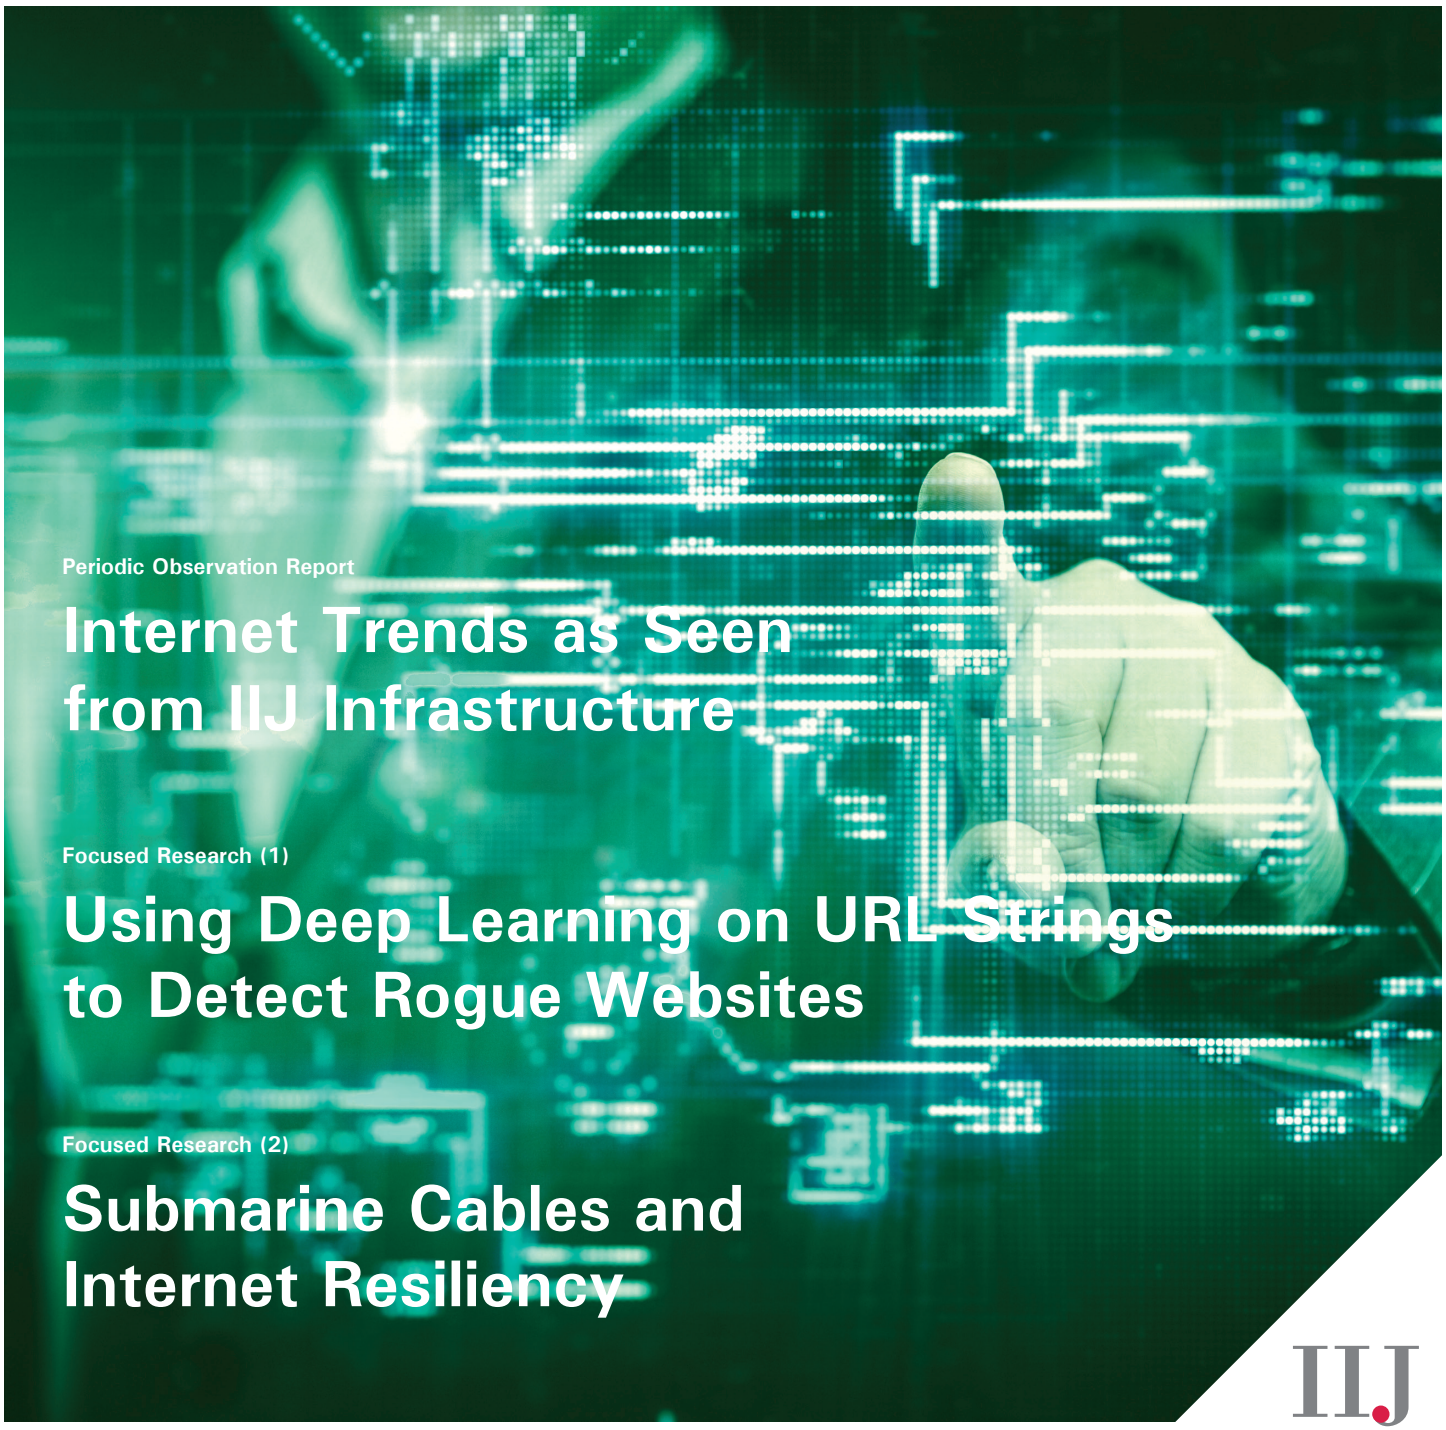

Periodic Observation Report

Internet Trends as Seen from IIJ Infrastructure

Focused Research (1)

Using Deep Learning on URL Strings to Detect Rogue Websites

Focused Research (2)

Submarine Cables and Internet Resiliency

IIJ
Internet Initiative Japan

Internet Infrastructure Review

February 2019 Vol.41

Executive Summary

News about 5G, the fifth generation of mobile telecommunications technology, is on the rise. In the US, Verizon began launching 5G services in some cities starting October 2018. The service, called Verizon 5G Home, offers fixed-wireless access (FWA) to the home, rather than mobile connectivity.. News coming out of South Korea, meanwhile, said that three mobile carriers are launching 5G services there in December 2018. The initial aim, apparently, will be to provide solutions to industry.

In Japan, 5G services are set to roll out in full from 2020, but the country's three big mobile carriers plan to launch pre-commercial 5G services during 2019. Information about those pre-commercial services revealed by the carriers at a Ministry of Internal Affairs and Communications (MIC) hearing in October 2018 indicates that they will cover diverse ground, including services related to regional revitalization as well as sports stadium VR experiences.

Frequency bands for 5G services are slated to be allocated by the end of March 2019, ahead of the pre-commercial services, and the MIC announced the guidelines for the process this past November. A total of 10 frequency ranges spanning 1,800 MHz are to be allocated for nationwide use from the 3.7GHz, 4.5GHz, and 28GHz bands, and additionally, a range of 200MHz in the 4.5GHz band and 900MHz in the 28GHz band are being left open, with the idea being that the MIC will consider allocating frequency ranges for privately operated applications and the like. How these frequency ranges will be allocated depends on the course of debate ahead, but perhaps we can look forward to seeing some distinctive services being offered by operators other than the incumbent nationwide carriers.

The IIR introduces the wide range of technology that IIJ researches and develops, comprising periodic observation reports that provide an outline of various data IIJ obtains through the daily operation of services, as well as focused research examining specific areas of technology.

In our periodic observation report in Chapter 1, we look at Internet trends as viewed from IIJ infrastructure. Of the various types of data we observe via IIJ network infrastructure, in this issue we analyze BGP route numbers, DNS queries, IPv6 traffic, mobile traffic, and backbone traffic. The analysis of BGP route numbers yielded some interesting results. For example, with IPv4 addresses being exhausted, we are seeing a rise in /22, /23, and /24 routes, possibly because address blocks are being split up for the purpose of transfer, and we also find that 32-bit only ASNs, the number of which continues to rise, appear to be operated only in the IPv4 space. We also note that the data on BGP route numbers, DNS queries, and the various types of traffic we look at indicate that use of IPv6 continues to progress.

Chapter 2 is our first focused research report, summarizing our attempt to apply deep learning to the task of identifying rogue URLs. The WWW is without doubt responsible for driving the explosive spread of the Internet and having a major impact on the way our society operates. Malicious uses of that WWW, which now pervades our society, include fraudulent sites masquerading as legitimate online services and banking facilities. Protecting users from accessing such rogue sites is a crucial task for people whose job it is to provide network services. Many approaches to identifying rogue sites have been proposed. Here, we evaluate an approach that uses a simple neural network and report high accuracy in the classification of URLs.

Chapter 3 presents our second focused research report, which looks at submarine cables. With 99% of all international data being carried by submarine cables, these cables are utterly crucial to the global Internet. Here, we describe the growth and state of the submarine cable network based on publicly available information and put forward an approach for examining the impact of submarine cable disruptions on the Internet based on various observational data on actual disruptions.

Through activities such as these, IIJ strives to improve and develop its services on a daily basis while maintaining the stability of the Internet. We will continue to provide a variety of services and solutions that our customers can take full advantage of as infrastructure for their corporate activities.

Junichi Shimagami

Mr. Shimagami is a Senior Executive Officer and the CTO of IIJ. His interest in the Internet led to him joining IIJ in September 1996. After engaging in the design and construction of the A-Bone Asia region network spearheaded by IIJ, as well as IIJ's backbone network, he was put in charge of IIJ network services. Since 2015, he has been responsible for network, cloud, and security technology across the board as CTO. In April 2017, he became chairman of the Telecom Services Association of Japan MVNO Council.

Internet Trends as Seen from IIJ Infrastructure

To provide Internet services, IIJ operates some of the largest network and server infrastructure in Japan. Here, we examine and discuss current Internet trends based on information obtained through the operation of this infrastructure.

We cover the topics of network routing information and DNS query information, as well as IPv6 usage and mobile connectivity services. We also report on the current state of the backbone network that supports the bulk of IIJ's traffic.

Topic 1

BGP / Number of Routes

Following on from last year's IIR Vol. 37 (https://www.iij.ad.jp/en/dev/iir/037.html), we start by looking at IPv4 full-route information advertised by our network to other organizations (Table 1, Figure 1). During the past year, RIPE NCC finished allocating/assigning its last /8 block (making it the second organization to do so, after ARIN). The size of allocations from the IANA Recovered IPv4 Pool to RIRs has also fallen to /22 (1,024 addresses). Increasingly, therefore, the acquisition of IPv4 addresses is becoming reliant on address transfers.

The last eight years has seen the largest increase in the total number of routes, which now exceeds 700,000. The growth rates for the /22 and /23 prefixes are above 10%, and taken together, the /22, /23, and /24 prefixes combined saw 89% growth in the number of routes to now account for 79% of all routes. As address blocks are increasingly split up for the purpose of transfer, it will be worth keeping an eye on the extent to which the proportion of routes accounted for by these prefixes grows.

Next we take a look at IPv6 full-route data (Table 2). The biggest increase here has also come in the past eight years.

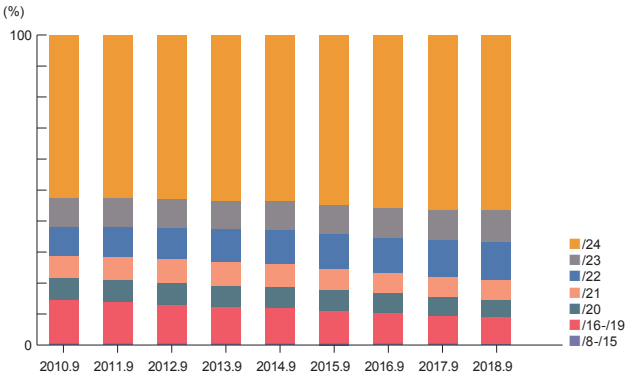

Figure 1: Percentage Breakdown of Number of Routes by Prefix Length for Full IPv4 Routes

Table 1: Number of Routes by Prefix Length for Full IPv4 Routes

	/8	/9	/10	/11	/12	/13	/14	/15	/16	/17	/18	/19	/20	/21	/22	/23	/24	total
Sep. 2010	20	10	25	67	198	409	718	1308	11225	5389	9225	18532	23267	23380	30451	29811	170701	324736
Sep. 2011	19	12	27	81	233	457	794	1407	11909	5907	9885	19515	26476	26588	35515	34061	190276	363162
Sep. 2012	19	14	29	84	236	471	838	1526	12334	6349	10710	20927	30049	31793	42007	39517	219343	416246
Sep. 2013	16	11	30	93	250	480	903	1613	12748	6652	10971	22588	32202	34900	48915	42440	244822	459634
Sep. 2014	16	12	30	90	261	500	983	1702	13009	7013	11659	24527	35175	37560	54065	47372	268660	502634
Sep. 2015	18	13	36	96	261	500	999	1731	12863	7190	12317	25485	35904	38572	60900	52904	301381	551170
Sep. 2016	16	13	36	101	267	515	1050	1767	13106	7782	12917	25229	38459	40066	67270	58965	335884	603443
Sep. 2017	15	13	36	104	284	552	1047	1861	13391	7619	13385	24672	38704	41630	78779	64549	367474	654115
Sep. 2018	14	11	36	99	292	567	1094	1891	13325	7906	13771	25307	39408	45578	88476	72030	400488	710293
Growth rate*	-1	-2	0	-5	8	15	47	30	-66	287	386	635	704	3948	9697	7481	33014	56178

*Since September 2017

That said, prefixes /33 through /48 account for 74% of all routes, and of those routes, we calculate that over 76% correspond to route advertisements for blocks that have been split into smaller fragments since being allocated (assigned). The number of routes is one gauge of the spread of IPv6, so an increase in this metric is desirable, but with fragments accounting for over half of the total, we seem to be somewhat removed from the original IPv6 ideal (?) of being able to limit growth of the routing table through consolidated route advertising, and this is a little disappointing.

Lastly, let's also take a look at IPv4/IPv6 full-route Origin AS figures (Table 3). With IANA's 16-bit Autonomous System Number (ASN) Pool having been exhausted in July 2016, the number of 16-bit Origin ASNs turned downward from 2016. The number of 32-bit only ASNs (allocation started in January 2007), meanwhile, has continued to rise steadily, but the vast majority appear to be operated only in the IPv4 space. This seems to indicate that even new organizations have not given much thought to using IPv6 despite probably having acquired ASNs and started running BGP during a time when the stockpile of IPv4 addresses had been exhausted, and this tells us that the IPv6 rollout still has a long way to go. That said, IPv4 addresses will certainly be harder and harder to obtain ahead, so whether this trends persists or not is also something we will be keeping close tabs on.

Table 2: Number of Routes by Prefix Length for Full IPv6 Routes

	/16-/28	/29	/30-/31	/32	/33-/39	/40	/41-/43	/44	/45-/47	/48	total
Sep. 2010	38	3	10	2023	33	2	9	4	17	436	2575
Sep. 2011	68	13	22	3530	406	248	45	87	95	2356	6870
Sep. 2012	102	45	34	4448	757	445	103	246	168	3706	10054
Sep. 2013	117	256	92	5249	1067	660	119	474	266	5442	13742
Sep. 2014	134	481	133	6025	1447	825	248	709	592	7949	18543
Sep. 2015	142	771	168	6846	1808	1150	386	990	648	10570	23479
Sep. 2016	153	1294	216	8110	3092	1445	371	1492	1006	14291	31470
Sep. 2017	158	1757	256	9089	3588	2117	580	1999	1983	18347	39874
Sep. 2018	168	2279	328	10897	4828	2940	906	4015	2270	24616	53247
Growth rate*	10	522	72	1808	1240	823	326	2016	287	6269	13373

*Since September 2017

Table 3: IPv4/IPv6 Full-Route Origin AS Numbers

ASN	16-bit (1–64495)					32-bit only (131072–4199999999)				
Advertised route	IPv4+IPv6	IPv4 only	IPv6 only	total	(IPv6 -enabled)	IPv4+IPv6	IPv4 only	IPv6 only	total	(IPv6 -enabled)
Sep. 2010	2083	32399	67	34549	(6.2%)	17	478	3	498	(4.0%)
Sep. 2011	4258	32756	115	37129	(11.8%)	90	1278	13	1381	(7.5%)
Sep. 2012	5467	33434	125	39026	(14.3%)	264	2565	17	2846	(9.9%)
Sep. 2013	6579	34108	131	40818	(16.4%)	496	3390	28	3914	(13.4%)
Sep. 2014	7405	34555	128	42088	(17.9%)	868	4749	55	5672	(16.3%)
Sep. 2015	8228	34544	137	42909	(19.5%)	1424	6801	78	8303	(18.1%)
Sep. 2016	9116	33555	158	42829	(21.7%)	2406	9391	146	11943	(21.4%)
Sep. 2017	9603	32731	181	42515	(23.0%)	3214	12379	207	15800	(21.7%)
Sep. 2018	10199	31960	176	42335	(24.5%)	4379	14874	308	19561	(24.0%)

Topic 2

DNS

IIJ provides a full resolver to enable DNS name resolution for its users. In this section, we discuss the state of name resolution, and analyze and reflect upon data from servers provided mainly for consumer services, based on a day's worth of full resolver observational data obtained on May 7, 2018.

ISPs notify users of the IP address of full resolvers via various protocols, including PPP, DHCP, RA, and PCO, depending on the connection type, and they enable users to automatically configure which full resolver to use for name resolution on their devices. ISPs can notify users of multiple full resolvers, and users can specify which full resolver to use, and add full resolvers, by altering settings in their OS, browser, or elsewhere. When more than one full resolver is configured on a device, which ends up being used depends on the device's implementation or the application, so any given full resolver is not aware of how many queries a user is sending in total. When running full resolvers, therefore, this means that you need to keep track of query trends and always keep some processing power in reserve.

Observational data on the full resolver provided by IIJ show fluctuations in user query volume throughout the day, with volume hitting a daily trough of about 0.05 queries/sec per source IP address at around 4:00 a.m., and a peak of about 0.22 queries/sec per source IP address at around 1:00 p.m. Broken down by protocol (IPv4 and IPv6), the trends in query volume are virtually the same (no major differences) during daytime hours, whereas IPv6 queries per IP address show a tendency to rise after 8:00 p.m. This suggests that the computing environment needed to allow use of IPv6 in the home is coming into place.

Recent years have seen a tendency for queries to rise briefly at certain round-number times, such as hour marks. The number of query sources also increases, which tells us that the increase is possibly due to tasks being scheduled on user devices and automated network access that occurs when a device is activated by, for example, an alarm clock function. Diving a little deeper, we note an increase in queries 14 seconds before every hour mark. The increase in queries that occurs on the hour tapers of gradually, but with the spike that occurs 14 seconds before the hour, query volume immediately returns to about where it was. Hence, because a large number of devices are sending queries in almost perfect sync, we can infer that some sort of lightweight, quickly completed tasks are being executed. Some implementations, for example, may have a mechanism for completing basic tasks, such as connectivity tests or time synchronization, before bringing the device fully out of sleep mode, and the queries used for these tasks could be behind the spike.

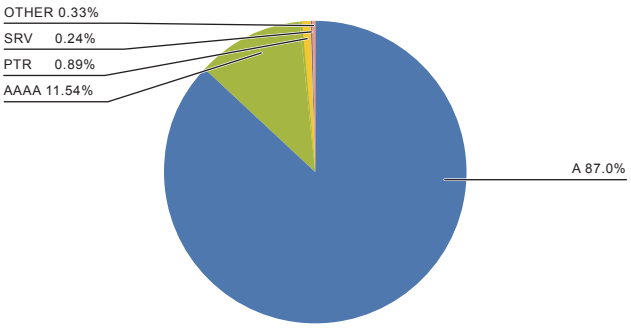

Figure 2: IPv4-based Queries from Clients

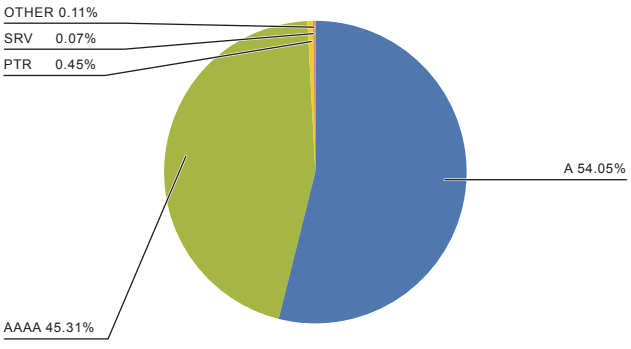

Figure 3: IPv6-based Queries from Clients

Looking at the query record types, most are A records that query the IPv4 address corresponding to the host name and AAAA records that query IPv6 addresses. ANY queries have decreased relative to last year.

With ANY queries being widely abused for reflection attacks and the IETF continuing to discuss relevant countermeasures, this type of query is gradually falling out of use, which would explain the decline this year. Turning to trends broken down by query IP protocol, the number of query source IPs and the number of actual queries are both higher for IPv6-based queries than for IPv4 queries. The trends in A and AAAA queries differ by IP protocol, with more AAAA record queries being seen for IPv6-based queries. Of IPv4-based queries, around 87% are A record queries and 11% AAAA record queries (Figure 2). With IPv6-based queries, meanwhile, AAAA record queries account for a higher share of the total, with around 54% being A record and 45% being AAAA record queries (Figure 3).

Topic 3

IPv6

Around a year has passed since we last looked at the state of IPv6 in Internet Infrastructure Review Vol. 37. In this issue, we look at what volume of overall traffic on the IIJ backbone is IPv6 and what protocols are mainly being used. We also look specifically at the state of and factors behind mobile services traffic, an area where IPv6 traffic is on the rise.

Traffic

As before, we again present IPv4 and IPv6 traffic measured using IIJ backbone routers at core POPs (points of presence—Tokyo, Osaka, Nagoya), shown in Figure 4. The data span the year from October 1, 2017 to September 30, 2018. Over the year, IPv4 traffic increased by around 20% while IPv6 traffic rose by around 80%. IPv6 accounts for around 6% of overall traffic, an increase from around 4% last year. Figure 5 plots the data for the same period on a

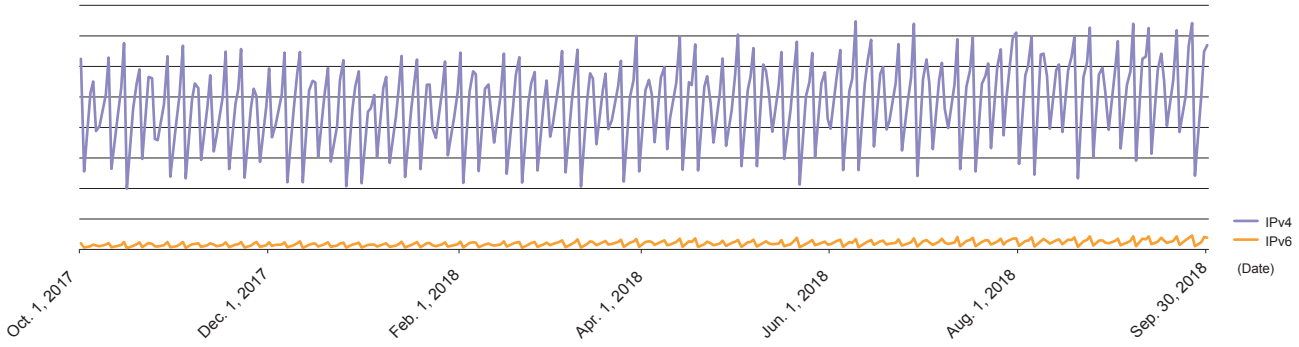

Figure 4: IPv4/IPv6 Traffic Measured via IIJ Backbone Routers at Core Points of Presence (Tokyo, Osaka, and Nagoya)

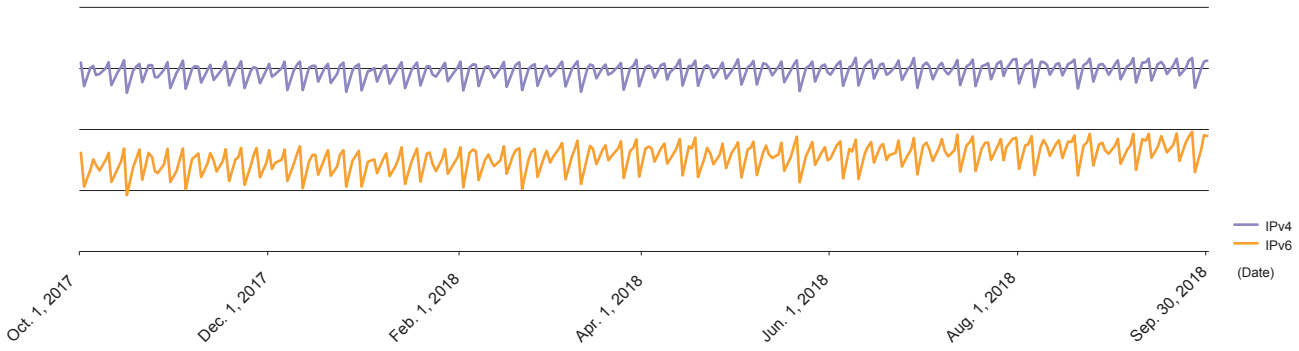

Figure 5: IPv4/IPv6 Traffic Measured via IIJ Backbone Routers at Core Points of Presence (Tokyo, Osaka, and Nagoya)—Log Scale

log scale. Although the absolute volume of traffic accounted for by IPv6 is less than a tenth that for IPv4, the rate of growth for IPv6 is clearly higher than for IPv4.

Next, Figures 6 and 7 show the top annual average IPv6 and IPv4 traffic source organizations (BGP AS Number) for the year from October 2017 through September 2018.

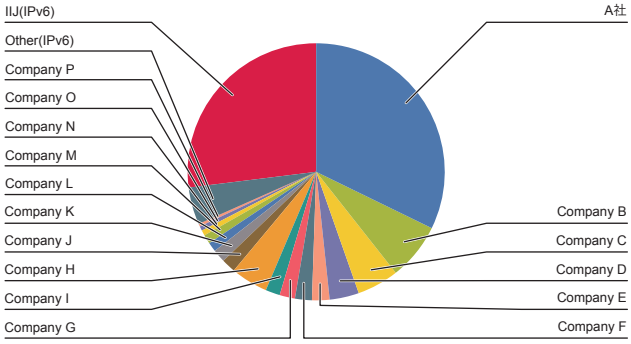

Figure 6: Top Annual Average IPv6 Traffic Source Organizations (BGP AS Number) from October 2017 to September 2018

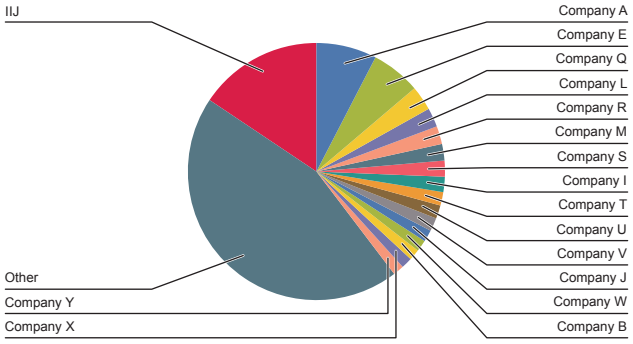

Figure 7: Top Annual Average IPv4 Traffic Source Organizations (BGP AS Number) from October 2017 to September 2018

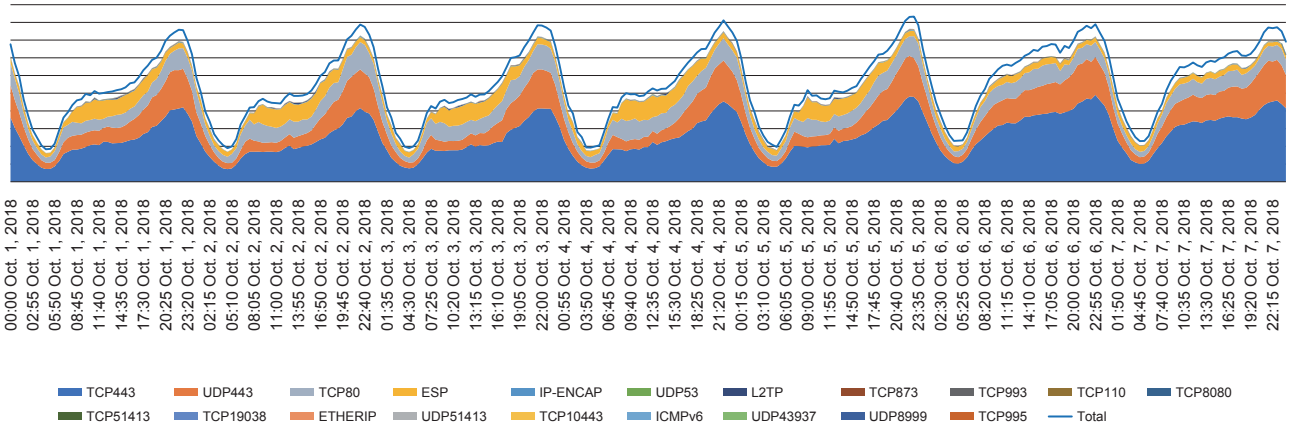

Figure 8: Breakdown of IPv6 Traffic by Protocol Number (Next Header) and Source Port Number

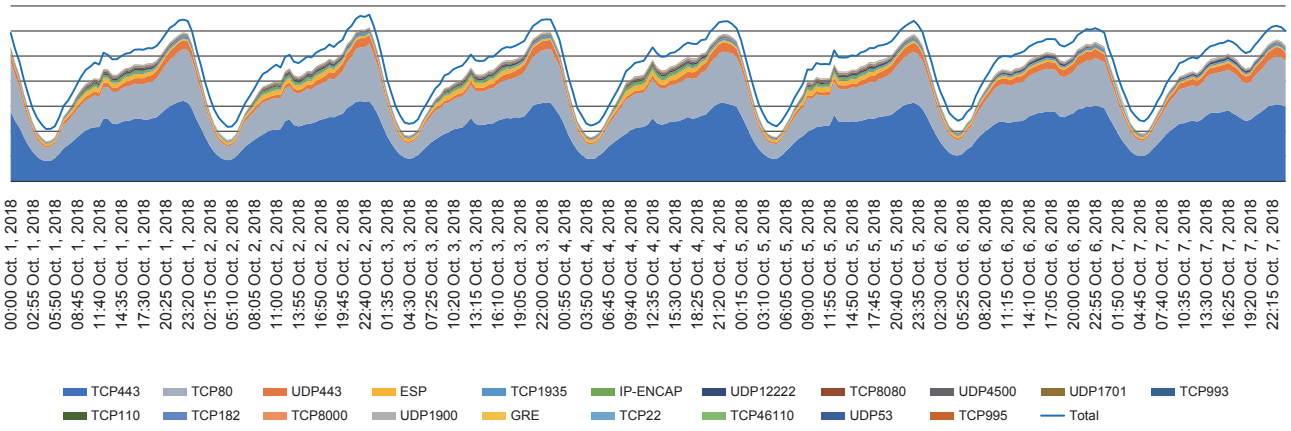

Figure 9: Breakdown of IPv4 Traffic by Protocol Number and Source Port Number

Last year's top ranking organization is still in the No. 1 spot but with its slice of the pie now only about half what it was due to a narrowing of the traffic volume gap versus No. 2 down. The data indicate that use of IPv6 is progressing at the No. 2 level on downwards as well. Providers of IPv6 IPoE access services via FLET'S Hikari Next come in at No. 4 and No. 6, which may tell us that the spread of IPv6 IPoE is leading to an increase in the use of IPv6.

Protocols Used

Figure 8 plots IPv6 traffic according to protocol number (Next Header) and source port number, and Figure 9 plots IPv4 traffic according to protocol number and source port number (for the week starting October 1, 2018).

In a trend similar to last year, TCP/UDP 443 and TCP 80 now account for an even greater share of the total, with Web-based applications accounting for most of the traffic. This goes for not only IPv6 but IPv4 as well.

In addition, IP-ENCAP (Protocol Number 4) rose from No. 6 last year to No. 5 in the rankings this year. Although the series is too thin to be visible in the plots, the traffic numbers have more than doubled since last year, which we surmise indicates an increase in traffic using IPv4-over-IPv6 technologies such as DS-Lite (RFC6333).

Mobile Services IPv6 Traffic

In a new addition to this periodic report, we now look at IPv6 traffic on mobile services.

Figure 10 plots traffic on IIJ mobile services over a two-year period from October 1, 2016. IPv6 traffic starts to surge right around the middle of the plot, corresponding to late September 2017.

This coincides with the release of iOS version 11, the operating system on US-company Apple's iPhones and iPads. In iOS 11, the MNVO APN profile (config file for mobile

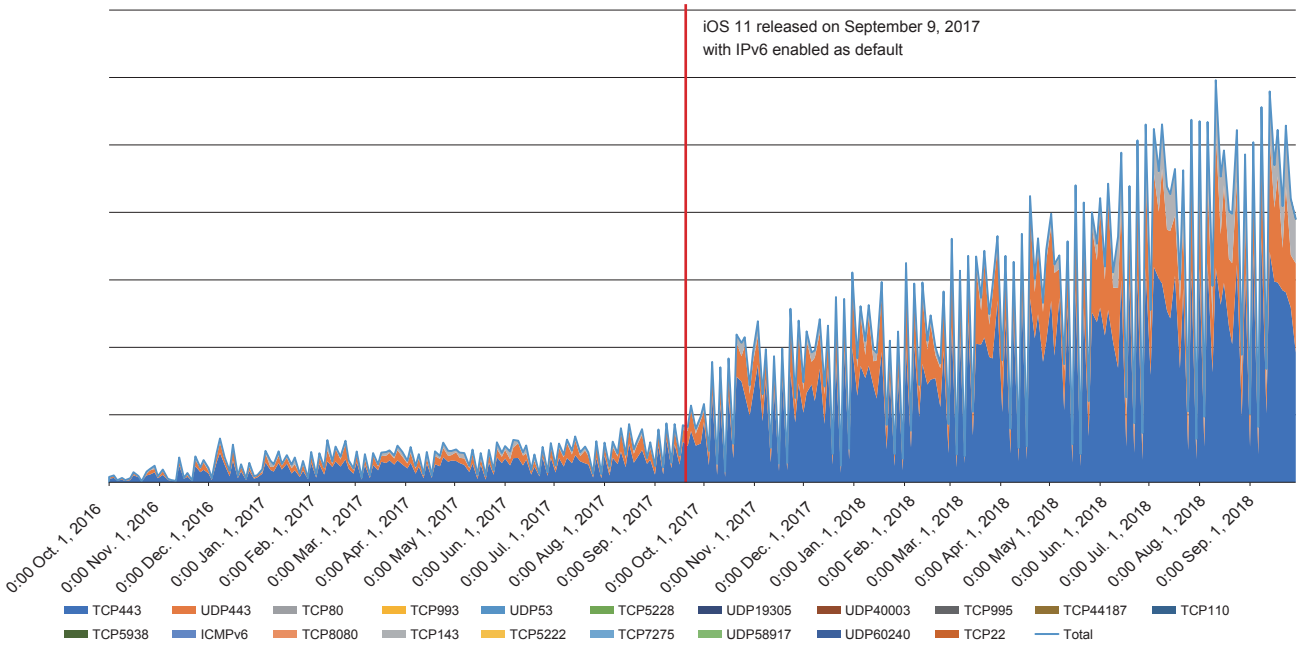

Figure 10: Mobile Services Traffic over Two Years from October 1, 2016

network connections) defaults to IPv6, and this may be why a lot of user devices started using IPv6.

Web-based application traffic accounts for almost all (over 98%) of mobile IPv6 traffic.

Summary

In this issue, we examined IPv6 traffic volume, protocols used, and IPv6 traffic for mobile services as a separate category. Overall, IPv6 traffic is on the rise, with the rate of growth outpacing IPv4. One reason for this seems to be the advance of IPv6 on everyday user devices due to factors such as the spread of IPv6 IPoE services on FLET's Hikari Next and the release of Apple iOS 11, as well as the range of IPv6 -based service providers growing increasingly diverse. The last regional registry with IPv4 addresses remaining, AFRINIC (RIR for the African region), is expected to exhaust its IPv4 pool around the middle of 2019, so the use of IPv6 looks set to rise further and further.

Topic 4

Mobile and Broadband

We now analyze mobile and broadband traffic. Note that broadband in this section does not include FLET'S IPoE.

Figure 11 plots download and upload (from the user's perspective) traffic volume (bps) for both mobile and broadband normalized based on the peak values for each series. The plot shows that the peak in mobile traffic comes around noon and in broadband around 10:00 p.m. Mobile connections are often used when people are out of the house, so traffic is higher during the day. Increases can also be observed around times when people are commuting to and from work or school, so the data correlate strongly with people's daily movements. Broadband, meanwhile, is generally used at home once people return for the day, so traffic is heavier at night.

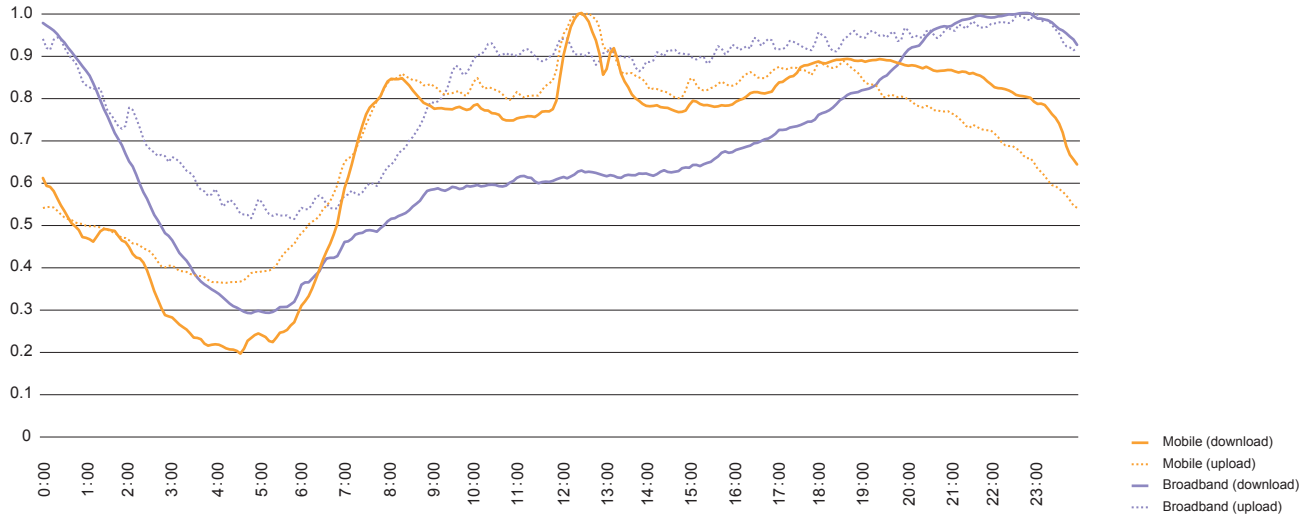

Figure 11: Traffic Indexed to Peak Levels

The fluctuations over the course of a day are larger for both mobile and broadband. Broadband traffic tends to rise gradually over the course of the day, from morning through to the nighttime peak. Mobile traffic, on the other hand, rises sharply in the morning and hovers at high levels up until right before midnight.

For reference, information on the traffic of Japan's five mobile carriers published in a Ministry of Internal Affairs and Communications document titled "Current State of Mobile Communications Traffic in Japan (June 2018)" [in Japanese] also indicates that traffic peaks at nighttime, as it does for IIJ broadband. Compared with MNO services, IIJ's mobile services (MVNO) currently attract more customers in the early adopter segment. We surmise that these customers also have broadband connections in the home and that they offload a high proportion of their nighttime traffic to broadband. As MVNO services spread further and attract more customers from the majority customer segments, the traffic peak is likely to shift into the nighttime, like that for MNO services.

Next, we compare download-to-upload ratios. Figure 12 shows this ratio, found by dividing download traffic (bps) by upload traffic, for both mobile and broadband.

The plot shows that the download ratio is higher for broadband than it is for mobile. Technological advances continue to push mobile communication speeds higher, but broadband still generally offers a more stable, high-speed connection. Plus, while broadband data is basically unlimited, mobile services are often subject to various forms of data transfer caps. This is probably why broadband sees higher capacity downloads.

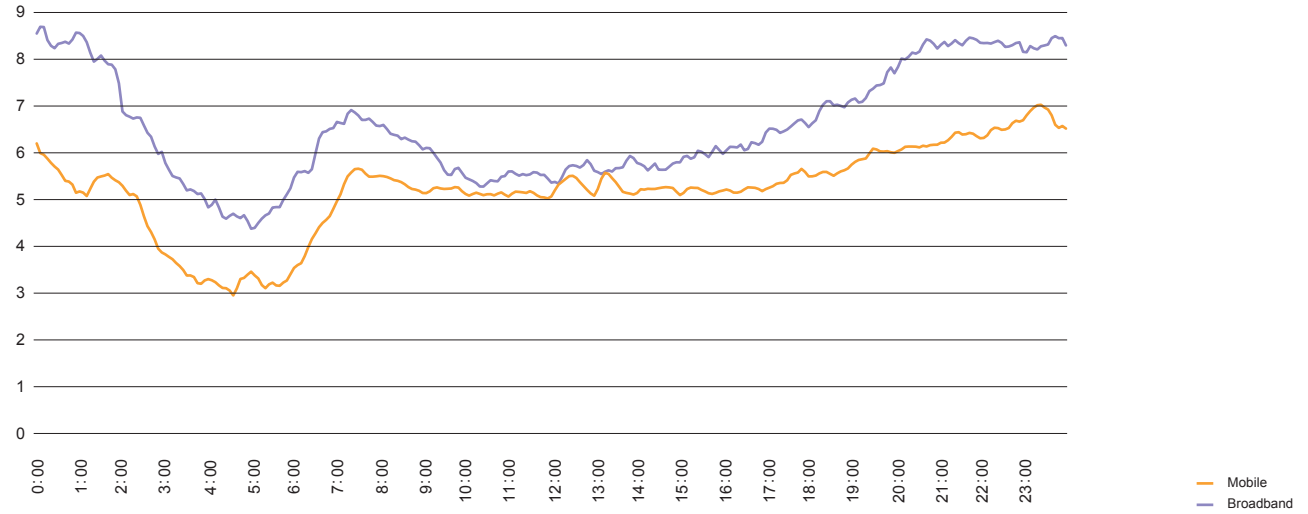

Figure 12: Download Relative to Upload Traffic

Next, we compare protocols. Figures 13 and 14 show a percentage breakdown (protocol, source port) of download traffic volume (bps) for mobile and broadband, respectively.

In both mobile and broadband, HTTP protocols (443/tcp and 80/tcp) account for about three quarters of the total. Also prominent is another relatively new protocol called QUIC, which uses 443/udp. The QUIC destinations are noticeably biased toward specific Internet service providers. Interestingly, 443/tcp (i.e., HTTPS) accounts for a greater share in mobile than in broadband. Most mobile users are quite possibly using smartphones, but rather than simply surfing Web pages via a stock-standard browser, it is possible that they frequently have occasion to use various applications designed for specific purposes and that the HTTPS protocol is commonly used by such applications.

Next, we look at IPv6 usage rates. Table 4 gives a percentage breakdown by connection type for mobile and broadband. The table shows that IPv6 usage is higher for mobile, albeit only slightly. With most new smartphones sold today being IPv6 capable, the groundwork for using IPv6 is naturally falling into place, even if users are unaware of it. With NTT's FLET'S service, users who have a compatible home gateway can use IPv6 without consciously having to do anything, but those users who provide their own broadband router need to configure it to use IPv6 themselves. FLET'S offers two connection types, PPPoE and IPoE, with the number of users connecting via IPoE, which faces fewer speed bottlenecks, rising of late. IPv6 is standard for IPoE connections, and to use IPv4 on IPoE connections, users need an environment that provides protocols to enable IPv4 over IPv6, such as DS-Lite.

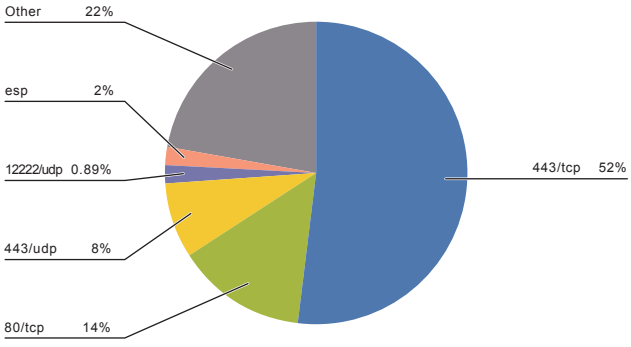

Figure 13: Percentage Breakdown of Download Traffic Volume (bps) by Protocol (Mobile)

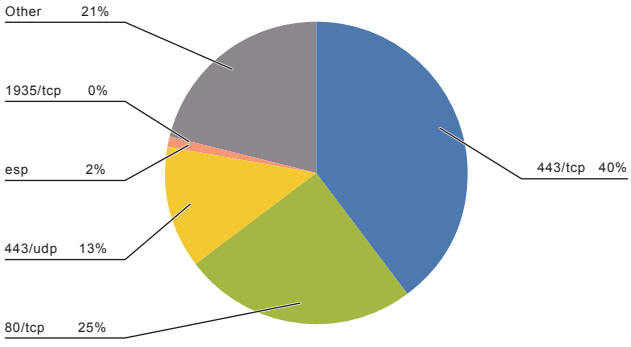

Figure 14: Percentage Breakdown of Download Traffic Volume (bps) by Protocol (Broadband)

A certain number of mobile users are exclusively IPv6. The number of services offering equivalent content via both IPv6 and IPv4 is rising, so it may be possible to meet demand in some cases with IPv6 alone, but it is surprising to see that there are so many of these users (when configured to use IPv4/IPv6 simultaneously, some devices make separate IPv4 and IPv6 connections for some reason, and this could be behind what we are observing).

Finally, let's break down mobile use by device. Mobile user communications use a protocol called GTP, and as part of the GTP connection process, devices report their IMEI (International Mobile Equipment Identity) to the network. IMEIs are composed of information that includes the device manufacturer and product name, so going through this data can give us a decent idea of what devices users are using.

Figure 15 shows a percentage breakdown by manufacturer. Even in a global context, Apple devices are noted as being highly prevalent in Japan, and they account for almost 40% of devices on IIJ's personal mobile services. This figure is astonishing given that, although IIJ does sell smartphones, it does not carry Apple products. And this shows just how appealing users find Apple devices to be.

The mobile user experience is very heavily influenced by the device (i.e., smartphone) used, and the mobile carriers put a lot of effort into smartphone sales for this reason. MVNO users, in some cases, continue to use the device they had when signing up for an MNO service, and so understanding what sort of smartphones are being used, including what devices are being brought over from MNO services, is crucial to the services strategy of MVNOs like IIJ.

Table 4: Breakdown of Connection Types for Mobile and Broadband

	IPv4	IPv6	IPv4v6
Mobile	70.43%	0.02%	29.55%
Broadband	75.48%	24.52%	NA

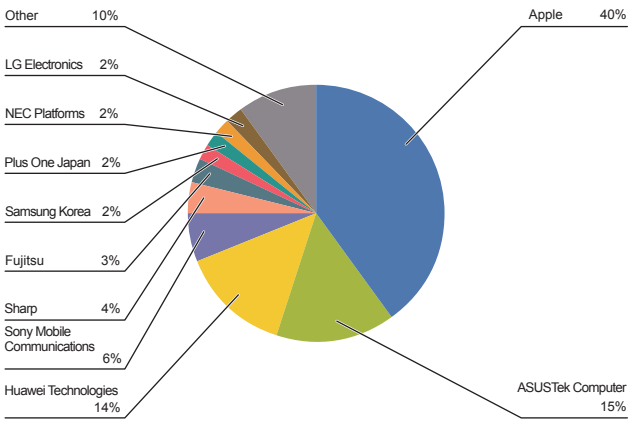

Figure 15: Device Breakdown

Topic 5

IIJ Infrastructure (Backbone)

IIJ monitors network status from a variety of angles to enable it to operate the IIJ network properly. In this issue, we take a look at one of the key metrics used: total traffic.

IIJ is an ISP, and one metric used to gauge an ISP's size is total traffic. Yet we cannot find many instances of what total traffic actually indicates being clearly explained. Here, we define total traffic to be the total bandwidth of transmissions into and out of IIJ's backbone. Backbone here means our routers collectively. It does not include transmission source/destination hosts. Although some transmissions are directed at the routers themselves, the volume is negligible.

With these definitions in place, let's look at IIJ backbone inflows/outflows. We can break these into three broad categories.

1. IIJ connection services customers
   - Connection services including Internet access services and datacenter access services
   - IIJ GIO (cloud service) Internet access service
   - Broadband (NTT East and West's FLET'S etc.) access services
   - Mobile connectivity services
2. IIJ service hosts
   - Email- and Web-related services
   - Content delivery
3. Interconnection business
   - Interconnections with other ISPs (peer)
   - Interconnections with cloud/content businesses

Based on this, Figure 16 shows plots of total traffic of the past 10 years. The data series are stacked. The outbound data are observations made at entry points, and the inbound data are observations made at exit points. Some traffic is eliminated within the backbone, such as that involved in

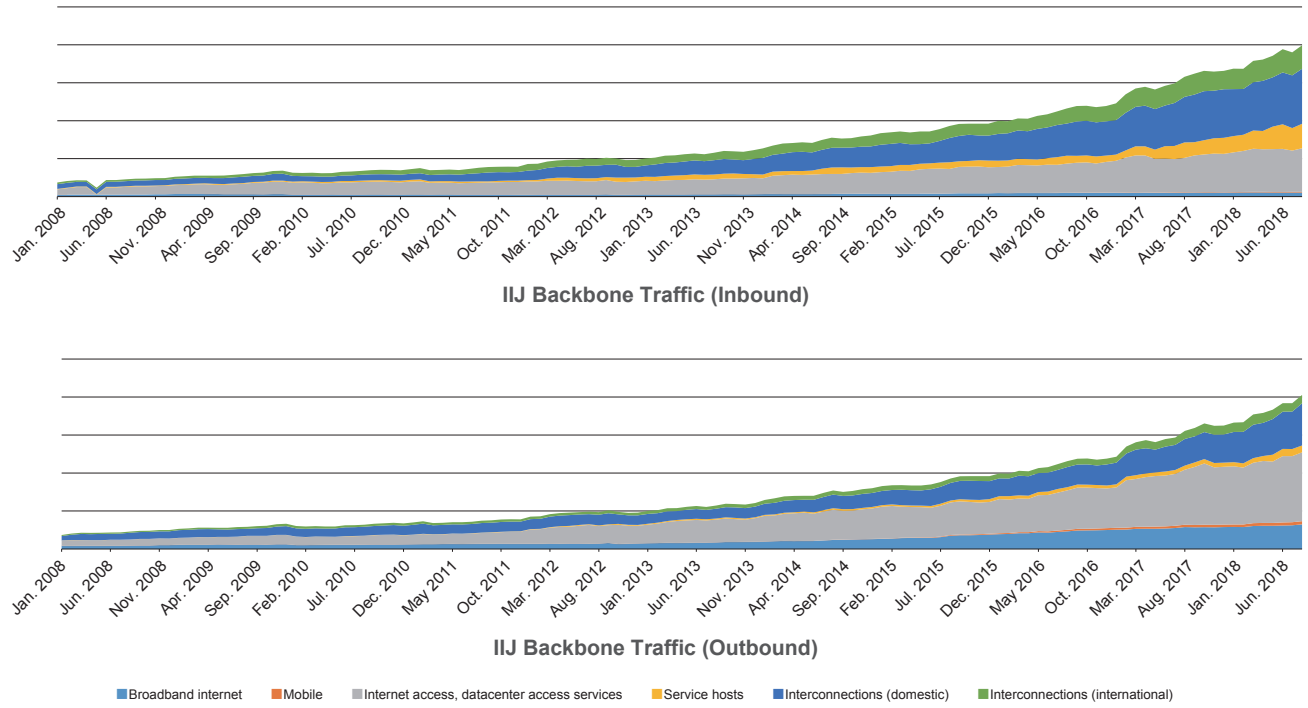

Figure 16: Total Traffic over 2008–2018

attacks, but generally all traffic that comes into the backbone also exits at some point, so the totals are almost the same.

The first thing you notice is that traffic has grown more than 10-fold over the past 10 years, with that growth accelerating. It shows no signs of easing. In the outbound plot, the three bottom series represent traffic that goes out to customers. Mobile traffic has been rising steadily since about three years ago but still accounts for only a small proportion of the total. Roughly speaking, the bulk of broadband and mobile traffic is accounted for by personal-use customers. Although this is rising steadily, the rate of growth in traffic to business customers (includes business customers to which IIJ provides personal services), shown in gray, has been higher over the past 10 years. Among IIJ customers, business traffic is growing more than personal-use traffic.

Let's look at inbound traffic. There tends to be less broadband and mobile traffic here, which means that not much information is sent out. Broadband traffic is growing, but the rate of growth is small, so it is accounting for a smaller and smaller proportion of the total. The share of traffic accounted for by service hosts is rising. This can be ascribed to the growth of content delivery and Web-based services.

Turning to a comparison of outbound and inbound traffic, we see that for broadband, although inbound traffic has only grown about twofold, outbound has grown close to ninefold. Here, too, it is evident that individual content offerings are becoming larger and larger. On the interconnection front, domestic inbound and outbound traffic look fairly well matched, whereas inbound is clearly greater for international interconnections. This seems to indicate that IIJ's content is not so much attractive when it comes to the international interconnection business.

This section has walked through a number of plots of IIJ's total traffic. Different insights emerge even from within the same traffic dataset depending on how you look at the data or what you focus on. Aside from traffic, we also record other completely different metrics such as latency or errors within the backbone. Going forward, we will continue to monitor IIJ's network and periodically report on changes in our observations.

1.BGP / Number of Route
Tomohiko Kurahashi
Infrastructure Planning Department, Service Infrastructure Division, IIJ.

2.DNS
Yoshinobu Matsuzaki
Infrastructure Planning Department, Service Infrastructure Division, IIJ

3.IPv6
Taisuke Sasaki
Deputy General Manager, Network Technology Department, Service Infrastructure Division, IIJ

4.Mobile
Takafusa Hori
Manager, Network Technology Department, Service Infrastructure Division, IIJ

5.IIJ Infrastructure (Backbone)
Takanori Sasai
Manager, Backbone Technology Section, Network Technology Department, Service Infrastructure Division, IIJ

Using Deep Learning on URL Strings to Detect Rogue Websites

While the Internet is now used to provide a range of useful services, it is also increasingly being used maliciously. As it is difficult to keep track of everything on huge, complicated systems manually, a range of automation strategies are employed. Security applications of deep learning have attracted attention in recent years. If deep learning can be used to assist humans in fields where experience and knowledge are crucial, this should enable a greater number of people to engage in higher-level tasks and, as a result, make it possible to provide safe services. In this issue, we present our attempt to use deep learning to prevent cyberattacks.

2.1 Advent of the Web and the Battle against Malicious Sites

Some 30 years have passed since Tim Berners-Lee, then a fellow at CERN (the European Organization for Nuclear Research), released CERN httpd, the first ever World Wide Web (WWW) server software. It constituted a means of using hypertext on the Internet, and combined with HTTP (Hypertext Transfer Protocol) and URLs (Uniform Resource Locators), it provided the technology to connect the world's information resources in a blink of an eye. It is no coincidence that the 1980s and 90s were also a time that saw the explosive spread of TCP/IP-equipped BSD UNIX, particularly among educational institutions, laying the groundwork for connecting the world's computers to one another. The release of Mosaic, a GUI-based Web browser, by the United States' National Center for Supercomputing Applications (NCSA) also made it possible for non-computer-experts to easily access the world's information. Web technologies continue to evolve even now, with new services popping up all over the place daily.

Something common to all technologies is that those technologies with the potential to make the world a better place can also make it worse. As all sorts of services become available online, so too emerge attempts to deceive and defraud via the Web. A common example involves setting up a fake version of a well-known service or banking website and sending out fake emails or other communications to steer users into the site, which is then used to steal their personal information, passwords, and the like. Fraud and deception existed before the advent of the Web, of course, but just as with email spam, digitization has lowered the cost and made it possible to target a larger number of people. The information age has brought benefits for both legitimate society and its underbelly alike.

Protecting users from accessing malicious sites has been a key issue for network services operators in recent years. On their end, ISPs commonly provide services that block rogue sites for their users. If you are responsible for your organization's information systems, perhaps your activities involve having some sort of security software installed for your organization's users. The technology commonly used to defend against rogue sites at present basically uses blacklists. But, as you can probably imagine, it is not really feasible to cover the entirety of the vast space that is the Web using blacklists alone. Researchers have been looking for more efficient ways of recognizing malicious sites. One attempt, for example, has involved using the domain names of known rogue sites to mechanically derive similar strings, thereby producing a large number of new potential malicious domain names from a short blacklist[*1]. Another has involved going beyond the idea of mere lists to look at features such as when the domain was registered and its Google search ranking, with domain names that have only recently been registered, have a low search ranking, and so on being treated as less trustworthy[*2]. Other techniques that have been proposed include actually retrieving and analyzing the content of web pages via a transparent proxy to detect whether a site is malicious or not[*3]. And as deep learning has advanced in recent years, it has also increasingly been employed in security applications.

*1 P. Prakash, M. Kumar, R. R. Kompella, and M. Gupta, "PhishNet: Predictive blacklisting to detect phishing attacks," in 2010 Proceedings IEEE INFOCOM, ser. INFOCOM 2010, pp.1–5.

*2 S. Garera, N. Provos, M. Chew, and A. D. Rubin, "A framework for detection and measurement of phishing attacks," in Proceedings of the 2007 ACM Workshop on Recurring Malcode, ser. WORM '07. New York, NY, USA: ACM, November 2007, pp.1–8.

*3 Y. Zhang, J. I. Hong, and L. F. Cranor, "CANTINA: A content-based approach to detecting phishing web sites," in Proceedings of the 16th international conference on World Wide Web, ser. WWW '07. ACM, May 2007, pp.639–648.

2.2 Meaning Hidden in URL Strings

The battle against rogue sites is a never-ending one. As soon as a technique for defending against such sites is devised, a mechanism for avoiding it appears. Even so, it is still important to consider new defense techniques if we are to make the Internet safer.

The proxy approach of actually retrieving page content to determine whether a site is rogue has the advantage in terms of detection rates. The act of actually accessing a site, however, can be dangerous in some cases. Given processing load, privacy, and other issues, methods that do not involve retrieving any actual content have also been proposed. The simplest of these is to look solely at the URL itself. The issue here is whether the strings that make up the URL contain any information that can indicate whether a site is rogue or not.

No one has a precise answer to this question. But a look at past research shows that some people have thought there may be meaning to be found. One well-known idea, for example, is to look at whether the domain name is a pronounceable string. Domain names often have something to do with actual goods or services, so URL strings are often based on natural language words and names and thus inevitably turn out to be strings that humans can pronounce. Some malware uses mechanically generated domain names (from a domain generation algorithm, or DGA), and in many cases these names consist of strings that cannot be pronounced. The idea is that if a distinction can be made here, it may be possible to distinguish between ordinary and suspicious Internet access.

Another idea is that an unusually large number of subdomains (host names with lots of dots in them) and an unusually deep path (URLs with lots of slashes in them) often indicate malicious intent. Under this approach, it is common practice to consider various criteria based on empirical rules, combining those criteria to determine whether something is malicious or not.

And at the forefront of techniques in this area, researchers are looking into the use of deep learning to assess URLs.

2.3 Resurgence of Deep Learning

Deep learning began to rise in the popular mindset around five or six years ago. Neural networks themselves, which are used in deep learning, have actually been around for long enough to be labeled as classical. However, the approach of deep learning, which uses multilayer neural networks, was long thought of as being hamstrung by practical impediments given the amount of calculation involved and the technical difficulties in getting models to learn properly. Flash forward to the 2010s, though, and the development of techniques that produced remarkable results in the area of image recognition flung the field into the spotlight. Opinion varies, but the most common account seems to be that the deep learning-based image recognition system[*4] demonstrated by Alex Krizhevsky and colleagues at the ImageNet Large Scale Visual Recognition Challenge (ILSVRC) in 2012 marked the start of the modern wave of deep learning that has propagated through to today. The model achieved a sharp reduction of 10 percentage points in the error rate from the previous mark of around 25%, demonstrating that deep learning can be applied to real-world scenarios. Use of deep learning subsequently became widespread, mainly in image and voice recognition, and it has also been applied to natural language translation, document classification, and even the strategy game Go.

In networking as well, researchers continue to put forward deep learning-based techniques, mainly in the area of security. In this article, we describe our proposal[*5] for detecting malicious sites based on URL strings, but we note that this is naturally not the first proposed approach of its type in the

*4 A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," Advances in neural information processing systems, 2012.

*5 K. Shima, D. Miyamoto, H. Abe, T. Ishihara, K. Okada, and Y. Sekiya, "Classification of URL bitstreams using Bag of Bytes," in Proceedings of First International Workshop on Network Intelligence (NI 2018), 2018.

world and that we expect many researchers and engineers to put forward even better strategies going forward. Creating something that will work indefinitely is not easy in network environments, particularly in distributed autonomous environments such as the Internet. Systems and data change with the times, and it is impossible to know all of the information therein because we can only ever see a portion of the world at once, and the information we can see becomes outdated almost as quickly as you can blink.

Deep learning is not an all-powerful approach. We do not yet know whether it will produce an intelligence that exceeds human capabilities, an idea that is often bandied about, but we do know what it is currently capable of.

Deep learning is a subset of machine learning in which a set of operations are performed on a given input vector to produce a separate output vector. It is used in classification problems and identification problems. An example of an image recognition application is a system that accepts an image of a cat (converted into a vector representation) and gives either a 0 or a 1 as output to indicate whether or not the image is a cat. Deep learning requires a large amount of data to determine what set of operations to perform. In the cat example, this would be a large quantity of cat images as well as images of objects other than cats. This is called the training data. When the data are labeled so that the answers are known, this is called supervised learning, and when this is not the case, it is called unsupervised learning (semi-supervised techniques that fall between these two also exist). Deep learning has produced great results with these sorts of classification problems.

2.4 Vectorizin URLs

Now let's move on to our URL classification problem. Our objective is to determine whether a given URL points to an ordinary, unproblematic site or to a rogue site. To use deep learning methods, we first need to convert the URLs into a vector representation that a deep learning model can take as input.

Before the rise of deep learning, the task of defining these vectors (feature engineering) was crucial to machine learning. This is because how you define what information is necessary and sufficient for differentiating your data ahead of time greatly influences performance. As mentioned previously, in URL classification, a variety of factors have been studied and validated as features that can be used to distinguish URLs, including whether the strings are pronounceable, the number of dots and slashes, the ratios of alphabet, symbol, and number characters, the position of characters, the frequency of n-gram strings, and so on. Haphazardly increasing the number of features can affect how long it takes to crunch the numbers, so with conventional machine learning methods of the past, experts with a deep knowledge of the dataset in question had to carefully select features likely to be useful through painstaking analysis of existing data.

It is said that deep learning, in contrast, can discover features for itself when trained using a large amount of data. In reality, it is not that simple, as careful preprocessing of the data often influences the final results, but it is also true that quantity of data can mitigate the effort needed for feature engineering to an extent.

To classify URLs in our system, we will not use existing features. Instead, we define a simple transformation to convert URL strings into fixed-length vectors. Past wisdom certainly can be used to classify URLs, but it is not necessarily possible to define features that will always be useful when working with other datasets in the future. There is also the tentative prospect of perhaps being able to use the same strategy on other datasets if we discover that it is possible to distinguish URLs using simple preprocessing and a large quantity of training data.

The vectorization procedure we adopted is as follows.

1. Split the URLs into individual characters
2. Convert the characters to hexadecimal ASCII codes
3. Extract byte values beginning at the start of the host part and the path part separately, shifting 4 bits at a time
4. Count how many times each value (from 0x00 to 0xFF) appears in the host part and path part, respectively, to form 256-dimensional vectors
5. Combine the 256-dimensional vectors created from the host part and the path part to form a 512-dimensional vector
6. Normalize the vector

Figure 1 shows steps 1 through 3. In step 4, we then count the values. The counts (in parentheses) for the sequence shown in Figure 1 are as follows.

0x16 (1), 0x2E (3), 0x42 (1), 0x61 (1), 0x64 (1), 0x69 (2), 0x6A (2),0x70 (1), 0x72 (1), 0x77 (5), 0x96 (2), 0xA2 (1), 0xA7 (1), 0xE6 (3)

Denoting the host vector as ***V*** and the *i*th element as v_i (where *i* is the extracted value), $v_{0x16} = 1$, $v_{0x2E} = 3$, $v_{0x42} = 1$, ... , $v_{0xE6} = 3$. Positions corresponding to unextracted values will contain a 0. Any original URL of any length can be converted into a 512-dimensional vector in this manner. But the longer the URL, the larger the size of the vector, so we normalize the vectors in step 6. We define the resulting 512-dimensional vector to be the "URL feature vector".

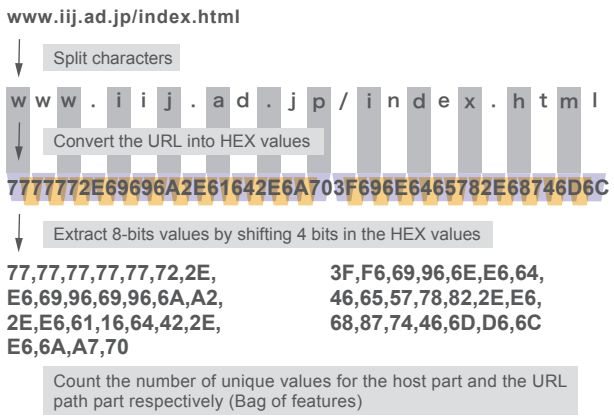

Figure 1: URL Vectorization

2.5 Designing the Neural Network

We are now ready to convert the URLs to fixed-length vectors. Next we have to decide how to train on that URL feature vector. In our attempt here, we used a simple 3-layer neural network. Although this is a rather shallow network for deep learning, it should be sufficient to determine whether this sort of method is effective or not.

Figure 2 shows the topology of the neural network we used. Some readers may have seen this topology somewhere before. This three-layer, fully connected topology appears as a sample in the Chainer open source deep learning library (https://chainer.org/) developed by Preferred Networks, and is used to recognize handwritten numerals from the MNIST dataset (http://yann.lecun .com/exdb/mnist/). Our work is based on this, with the following two changes.

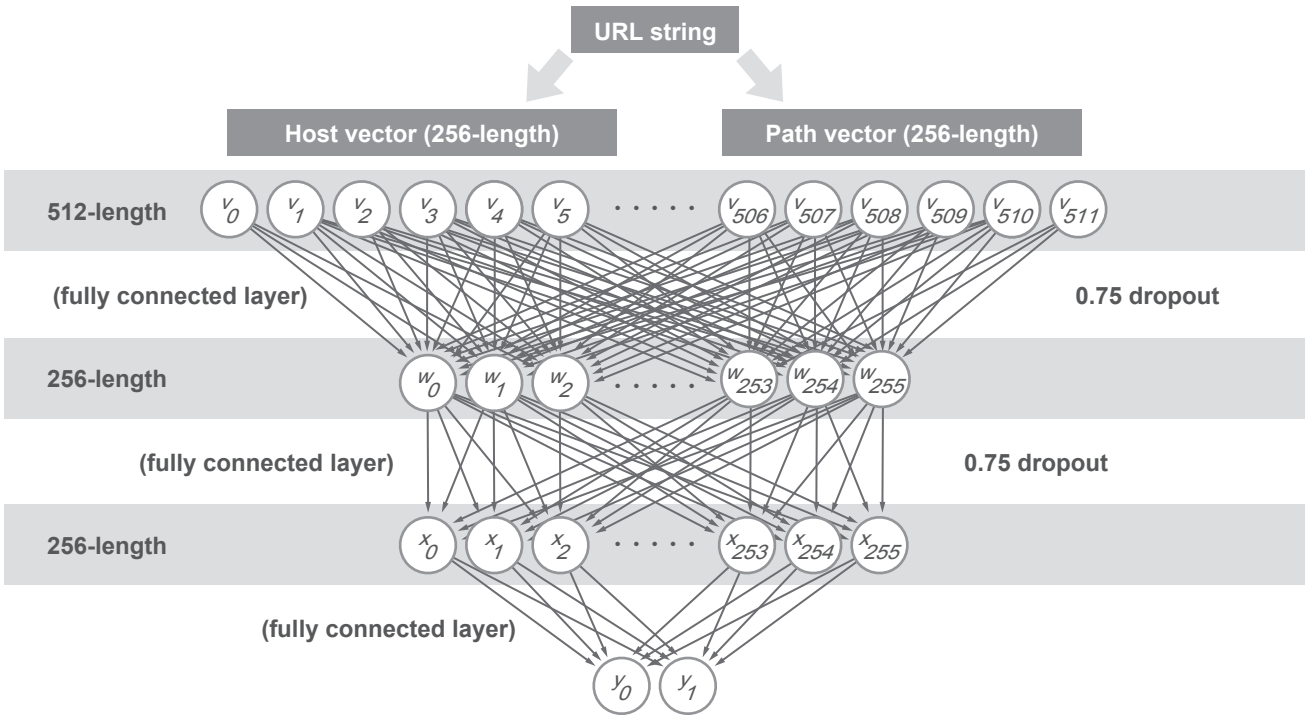

Figure 2: Neural Network Topology

1. Number of input/output dimensions: With the MNIST sample, input images are 28x28 pixels, yielding 784 input dimensions. And the output is 10-dimensional because the output values are the digits 0 through 9. Our input is the 512-dimensional URL feature vector, and our output is 2-dimensional, being the value 0 or 1, indicating whether a site is rogue or not.
2. Dropout rate: The MNIST sample does not use dropout, a method for preventing overfitting, but we observed serious overfitting with our data and thus set a fairly high dropout rate.

We used Chainer to verify our proposed approach. The neural network model we built in Chainer is shown in Table 1.

2.6 Data Source Selectionk

With our data structures and neural network model in place, we can now use actual data to verify our approach. Two major issues present themselves in dynamic environments like the Internet.

1. Data accuracy: With datasets like MNIST, the data are already fully validated and properly labeled (in the case of MNIST, this means the handwritten digits and the values they represent). Accurately labelling data observed/collected via the Internet, meanwhile, can be problematic. If the model is trained on incorrect information, it will naturally end up predicting incorrect answers.
2. Completeness: It is not possible to show whether the data used to train the model is a true representation of the general picture. If the model is trained on biased data, it will be unable to cope with different patterns when they appear. In the case of handwritten digit recognition, the problem space is somewhat limited since it only involves the digits 0 through 9, whereas URL strings on the Internet represent a virtually unlimited space, so the scope of applicability will naturally differ.

Under the provision that such problems exist, it is important to prepare data that is as accurate and complete as possible. The data we use comprises active phishing sites listed on PhishTank (https://www.phishtank.com/). PhishTank by no means provides an exhaustive list of all the rogue sites out there, so although completeness is not guaranteed, the data offer a degree of credibility since the process of determining whether a site is a phish or not involves human verification via a voting system. Compiling data on ordinary (non-rogue) sites is more difficult. For verification, we take sites appearing in a particular research institution's access logs and exclude those listed on PhishTank, defining such sites to be non-rogue sites. To enhance completeness, however, we would need to perform additional tests based on different types of access logs.

```
from chainer import Chain
import chainer.functions as F
import chainer.links as L
class Model(Chain):
    def __init__(self):
        super(Model, self).__init__()
        with self.init_scope():
            self.l1 = L.Linear(None, 256)
            self.l2 = L.Linear(None, 256)
            self.l3 = L.Linear(None, 2)
    def __call__(self, x):
        h1 = F.dropout(F.relu(self.l1(x)),
                       ratio=0.75)
        h2 = F.dropout(F.relu(self.l2(h1)),
                       ratio=0.75)
        y = self.l3(h2)
        return y
```

Table 1: Neural Network Model Built in Chainer

2.7 Applicability of Deep Learning

We randomly extract around 26,000 URLs each from the two types of data sources prepared per the previous section. Of this, 80% is used for training. When we used the remaining 20% to assess accuracy, we found that we were able to correctly distinguish between rogue and non-rogue URLs for 94% of the data. Opinion may vary on whether this is a good result or not. Some classification methods put forward in the past have achieved better outcomes than this figure. In some cases, those methods also used information other than just the strings (e.g., Whois information, Google search ranking), so simple comparisons with our approach are not possible. Another attempt employed deep learning on URL strings alone[*6] in a manner similar to our approach, yielding better classification accuracy than we achieved. Yet when we independently implemented the neural network proposed in that study and tested it on our dataset, we were unable to replicate the high accuracy reported in that paper. The takeaway here is that even with the same neural network model, accuracy can vary significantly depending on the dataset used for training.

Since it is impossible to obtain all the world's data, trained models will inevitably carry some bias. The term big data seems to have fallen by the wayside a bit lately, but we think it is likely that organizations with large stores of wide-ranging data will continue to occupy an advantageous position in the deep-learning world as well; indeed, that advantage may even widen.

2.8 Conclusion

We have described our attempt to apply deep learning to the task of identifying rogue URLs. Despite only using a simple neural network to test our approach, we were able to classify URLs with 94% accuracy. This exercise also reaffirmed the difficulties in collecting data and the advantages that having data imparts.

We can expect deep learning to increasingly be applied to network data ahead. With computing power increasing, it has become relatively easy to use deep learning. We hope to incorporate new technologies into our approach as we work toward making the Internet even safer going forward.

Acknowledgments
This research was supported by JST, CREST, JPMJCR 1783.

Keiichi Shima
Deputy Director, Research Laboratory, IIJ Innovation Institute

*6 J. Saxe and K. Berlin, "eXpose: A character-level convolutional neural network with embeddings for detecting malicious URLs, file paths and registry keys," CoRR, vol. abs/1702.08568, February 2017.

Submarine Cables and Internet Resiliency

The IIJ Innovation Institute is engaged in research that helps to improve the resiliency of the global Internet through measurement and analysis. In this chapter, we report on research that aims to improve resiliency by understanding the submarine cables that underpin the Internet. This is a summary of a paper presented at ACM HotNets 2018[*1].

3.1 Introduction

Ninety-nine percent of all international data is carried by submarine cables[*2]. Deployments of the submarine network date back to the mid-19th century, and total capacity of this undersea infrastructure is now growing at an exponential rate. Today, a complex mesh of hundreds of cables stretching over one million kilometers[*3] connects nearly every region in the world (Figure 1). It comprises both the operation backbone of major corporations' global services and cables that ensure connectivity to regions with limited terrestrial connectivity[*4*5].

Yet, despite the impressive scale and criticality of the submarine cable network, past studies have either treated it as a black box or focused on specific events and their impact on particular links, and its role in the global Internet is not well understood.

Here, we describe the growth and state of the submarine cable network based on publicly available information and put forward an approach for examining the impact of submarine cable disruptions on the global network based on observational data.

3.2 Background to the Submarine Cable Networks

The first commercial submarine cable was laid across the English Channel in 1850. Early cables were made of stranded copper wires and used for telegraphy. Fiber-optic cables were developed in the 1980s and the first fiber-optic transatlantic cable (TAT-8) was put into operation in 1988. Today nearly all cables are fiber-optic cables. In modern cables, the core optical fibers are protected by multiple layers, depending on the cable depth, including a copper tube, an aluminum water barrier, stranded steel wires, and a polyethylene shield (Figure 2). Cables vary in thickness from 10cm in diameter, weighing around 40t/km for shore-end cable, to 2cm in diameter, weighing about 1.5t/km, for deep-sea cable.

Most submarine cables have been constructed and are managed by consortia, and shared by multiple network operators. TAT-8, for instance, had 35 participants including most major international carriers at the time (including

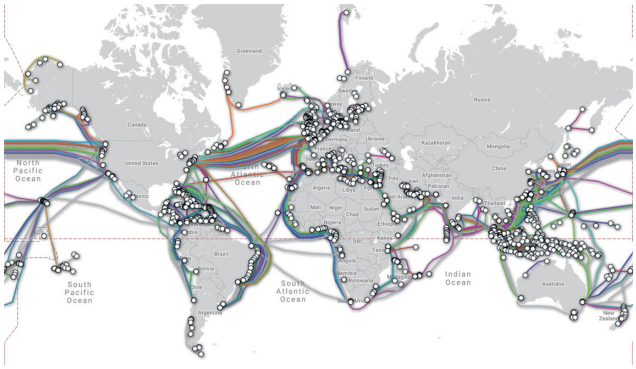

Figure 1: TeleGeography's Submarine Cable Map (June 2018)[*6]

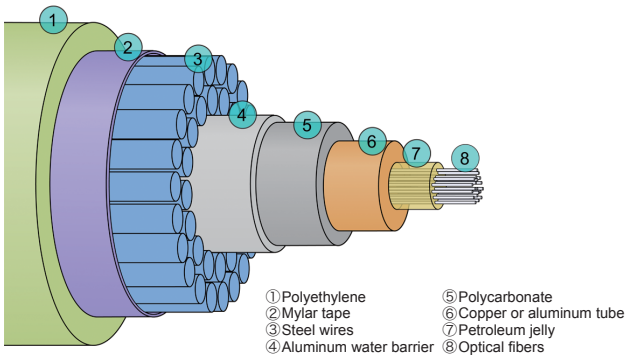

Figure 2: Cross Section of Submarine Cable with Multilayer Protection

*1 Zachary S. Bischof, Romain Fontugne, Fábian E. Bustamante. Untangling the world-wide mesh of undersea cables. In Proc. of HotNets, November 2018.
*2 P. Edwards. A map of all the underwater cables that connect the Internet, 2015 (https://bit.ly/2Ep19i4).
*3 The various threats to subsea cables. Ultramap (https://bit.ly/2Ld9LKW).
*4 NEC begins construction of submarine cable links to the islands of Palau,Yap and Chuuk. NEC, May 2017 (https://bit.ly/2JqQQaE).
*5 Z. S. Bischof, J. P. Rula, and F. E. Bustamante. In and out of Cuba: Characterizing Cuba's connectivity. In Proc. Of IMC, October 2015.
*6 TeleGeography. Submarine cable map (https://www.submarinecablemap.com/).

AT&T, British Telecom, and France Telecom)[*7]. The latest construction boom, however, seems to be driven by content providers, such Google, Facebook, Microsoft, and Amazon. According to TeleGeography, the amount of capacity deployed by content providers rose 10-fold between 2013 and 2017, outpacing all other customers[*8].

3.2.1 Problems Related to Submarine Cables

As the total length of submarine cables continues to expand rapidly, so too does the chance of network disruptions due to cable problems. The natural environment poses a number of risks for starters, from large-scale disasters like earthquakes and tsunamis, to undersea landslides and ocean currents that can scrape cables across the rocky surfaces and shark attacks on some cables.

Even more than natural forces, human actions—intentional or not—are the biggest threat to cables, with approximately 70% of disruptions being caused by fishing trawlers and ship anchors[*3], as well as growing concern over intentional attacks on vulnerable cables. For instance, US Navy officials have stated concern upon observing Russian submarines and spy ships operating near important submarine cables[*9*10].

While the high degree of connectivity available in certain areas may limit the consequence of cable disruptions, other regions appear to be particularly vulnerable[*11*12*13]. The Asia America Gateway cable (AAG), notorious for frequent breakdowns, connects Southeast Asia and the US, handling over 60% of Vietnam's international Internet traffic. In 2017 alone, the AAG suffered at least five technical errors[*14].

In another incident, divers off the coast of Egypt were arrested for cutting the SE-WE-ME-4 submarine cable, leading to a 60% drop in Internet speeds[*11*15]. Other incidents have resulted in entire countries being taken offline due to a single submarine cable cut, such as Mauritania in April 2018[*12].

To understand these risks, it is necessary to clarify the role of the submarine network as a component of the global network. Routes that appear to be distinct paths at the network layer may rely on the same cable at the physical layer.

For particularly critical routes (e.g., transpacific or transatlantic), large network operators often utilize multiple cables. Yet even with full details on the Layer 3 topology, the lack of visibility as to which routes and submarine cables networks are connected by makes it difficult for third parties to quantify the dependence of Internet connections on particular submarine cables.

3.2.2 The World's Submarine Cables

Data on submarine cables are publicly available on a number of websites. Here, we use data collected from two sites—TeleGeography's Submarine Cable Map[*6] and Greg (Mahlknecth)'s Cable Map[*16]—to describe the growth and current state of submarine network infrastructure in terms of the number and capacity of the cables. Both sites present a global map of hundreds of submarine cables with details on each cable. While there is a large overlap between them, we find significantly more cables in TeleGeogaphy's Map than in Greg's.

A caveat is that both resources only list details on publicly announced cables[*17]. TeleGeography estimates that by early 2018 there were approximately 448 submarine cables in service globally[*18], 90% of which were publicly announced. Most of the remaining privately owned and unannounced cables belong to content provider networks—such as Facebook and Google—who have made significant

*7 N. Starosielski. The Undersea Network. Duke University Press.
*8 A. Mauldin. A complete list of content providers' submarine cable holdings.
*9 M. Birnbaum. Russian submarines are prowling around vital undersea cables. It's making NATO nervous. The Washington Post, December 2017 (https://wapo.st/2NW71QP).
*10 D. E. Sanger and E. Schmitt. Russian ships near data cables are too close for US comfort. The New York Times, October 2015 (https://nyti.ms/2uqCnXh).
*11 C. Arthur. Undersea internet cables off Egypt disrupted as navy arrests three. The Guardian, March 2013 (https://bit.ly/2mlluZK).
*12 C. Baynes. Entire country taken offline for two days after undersea Internet cable cut. Independent, April 2018 (https://ind.pn/2L0zIOn).
*13 R. Noordally, X. Nicolay, P. Anelli, R. Lorion, and P. U. Tournoux. Analysis of Internet latency: The Reunion Island case. In Proc. Of AINTEC, 2016.
*14 B. Anh. Vietnam Internet returns to normal after AAG repairs. Submarine Telecom Forum, June 2018.
*15 A. Chang. Why undersea Internet cables are more vulnerable than you think. Wired, April 2013 (https://bit.ly/2KYFP5Y).
*16 G. Mahlknecht. Greg's cable map (https://www.cablemap.info/).
*17 A. Mauldin. A complete list of content providers' submarine cable holdings. Telegeography blog, November 2017 (https://bit.ly/2Lw7DLm).
*18 Telegeography. Submarine Cable 101 (https://bit.ly/2qcGSTc).

investments in undersea cables as part of their inter-datacenter networks[*17]. Although we focus here on those cables that are part of the public Internet, understanding the relation between the public and private submarine cable network is an open research question.

Each site lists the name of the cable, a list of its landing points, an approximate cable length, a ready-for-service date, and for some cables, links to external websites. Figure 3 shows an example of the data made available by TeleGeography, including cable length, owners, and landing points.

3.2.3 Growth and State of the Network

The submarine network has seen consistent linear growth in number of cables since the late 1980s. Using the data collected from the TeleGeography site, Figure 4 plots the number of cables currently in use based on ready-for-service dates (includes cables slated to go into operation by the end of 2020). As Figure 4 (left axis) shows, over the last thirty years there has been, on average, a new cable activation per month. Note that this data set is missing cables that were decommissioned. For example, TAT-8 (constructed in 1988) was the first fiber-optic cable in the Transatlantic Telephone (TAT) series of cables, but it was decommissioned in 2002 and is not part of TeleGeography's current dataset. The currently active TAT-14 cable began operating in 2001. Thus, the graph shows a lower bound on the total number of cables active each year.

The submarine network has grown not just in number of cables but also in the length of these cables. Figure 4 also plots the total length of currently active cables per year (right axis). By 2018, the total length of currently active cables had grown to over 1.2 million km.

The graph shows an interesting spike in lengths starting around 2015. The only period with faster growth corresponds with the dot-com boom (1997–2001).

Today, the global submarine infrastructure is capable of transferring over 1 Pbps of traffic, with total capacity growing multiple orders of magnitude in the last few decades. Using the bandwidth capacities from Greg's Cable Map, we plotted the total global bandwidth for currently active submarine cables, shown in Figure 5. A comparison with Figure 4 indicates that recently constructed cables are responsible for carrying a large portion of Internet traffic. Figure 6 shows the average bandwidth capacity of new cables from Figure 5 and Figure 4. Despite some noise in the early 1990s, we see that the average bandwidth capacity of cables grew by 2–3 orders of magnitude through around 2015. While average cable capacity remained relatively consistent between 1995 and 2010, capacity has spiked again in recent years.

These data represent conservative estimate as the sources do not include decommissioned cables and are restricted to publicly announced cables.

Asia-America Gateway (AAG) Cable System

Email link

RFS: November 2009
Cable Length: 20,000 km
Owners: Telekom Malaysia, AT&T, Starhub, PLDT, CAT Telecom Public Company Limited, Airtel (Bharti), Telstra, Telkom Indonesia, BT, Eastern Telecom, PT Indonesia Satellite Corp., Spark New Zealand, Viettel Corporation, Saigon Postel Corporation, Vietnam Telecom International, Brunei International Gateway, BayanTel, Ezecom
URL: http://www.asia-america-gateway.com

Landing Points

Changi North, Singapore
Keawaula, Hawaii, United States
La Union, Philippines
Lantau Island, Hong Kong, China
Mersing, Malaysia
San Luis Obispo, California, United States
Sri Racha, Thailand
Tanguisson Point, Guam
Tungku, Brunei
Vung Tau, Vietnam

Figure 3: Example of TeleGeography's Data

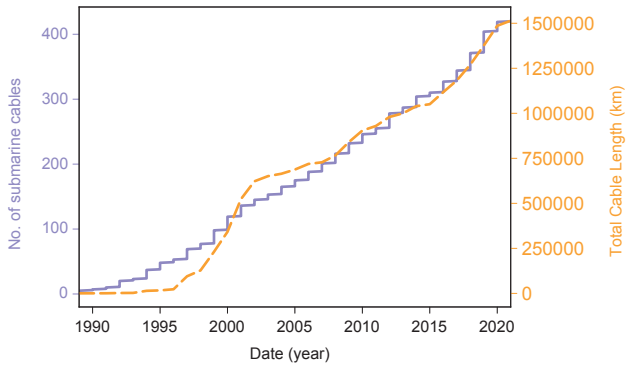

Figure 4: Number of Cables and Total Cable Length

3.3 Submarine cables and the Internet

We discuss the relationship between these submarine cables and the Internet. We set out three high-level tasks here: (1) creating an abstract graph of the submarine cable network, characterizing connectivity and identifying regions that are particularly susceptible to cable disconnections; (2) inferring the relationship between network-level resources and specific submarine cables in order to connect observations at the physical and network layers, and (3) exploring the consequences of submarine cable failures for Internet users.

3.3.1 Graphing Submarine Cable Connections

The first task is to derive an abstract graph of the submarine network. While seemingly simple, mapping cables, each with multiple landing points in different countries and land masses, on a single plot is no easy task.

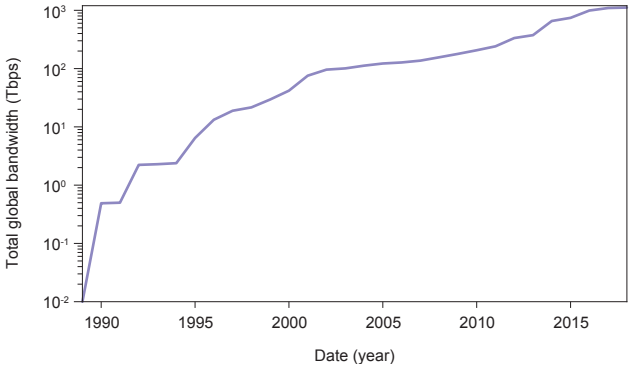

Figure 5: Time Series of Total Bandwidth of Currently Active Cables (per Greg's Cable Map)

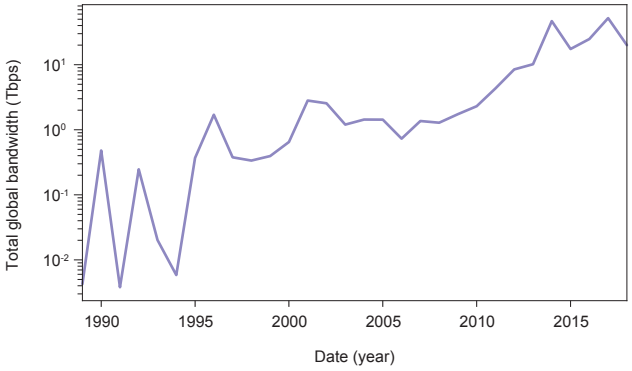

Figure 6: Time Series of Average Bandwidth of New Cables (per Greg's Cable Map)

In a first approximation, one could group cities connected by terrestrial network infrastructure into edges on the graph, using the submarine links between them as vertices. Take, for example, the Greenland Connect cable, shown in Figure 7, which connects Canada with two landing points in Greenland and one in Iceland. This approach will group the two Greenland points as having a land-based connection, with submarine connections between Canada and Greenland and between Greenland and Iceland[*19]. But a continuous landmass does not necessarily imply a terrestrial network connection. For example, although Panama and Colombia are contiguous neighbors, the lack of any transit infrastructure across the Darién Gap means that for connectivity purposes, these are essentially separate regions. We are currently using map data from Google Maps and Open Street Map to aid in identifying these disconnected regions.

A more difficult problem appears when landing points are close by. Consider the ACE (Africa Coast to Europe) cable, shown in Figure 8, and the Jasuka cable from Telkom Indonesia, in Figure 9. Unlike the Greenland example, ACE has 22 landing points connecting tens of countries on the west coast of Africa to two locations in continental Europe (Portugal and France). Even if one could imagine grouping the European points into a single vertex, it is unclear how to best group the west Africa points. The Jasuka cable, connecting 11 points around the island of Sumatra, further complicates matters—here, the exact definition of landing points is not even clear.

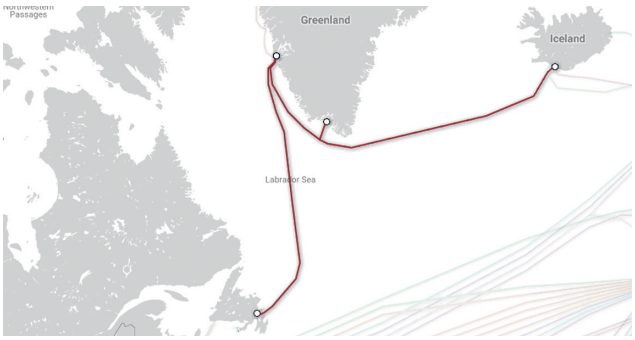

Figure 7: Greenland Connect (per TeleGeography Cable Map)

*19 In reality, even this "simple" example is not so straightforward; despite being on the same landmass, we need to treat the landing points in Greenland as separate due to the lack of infrastructure connecting the cities.

We plan to apply a variation of our basic approach, using other publicly available records, while building a common repository for the inferred view. Using this abstraction of the submarine cable network will help us to study the dependability of geographical areas to physical cables and identify high-risk links from a connectivity perspective.

3.3.2 Mapping onto the Internet

Most studies on Internet topology rely solely on measurements at the network layer. Inferring network reliability from such analysis has limits, as traffic that appears to be traveling via separate network paths could potentially be relying on the same physical resource. Besides shared infrastructure such as datacenters, submarine cables are commonly co-owned or leased by multiple network operators (e.g., TAT-14 is co-owned by over 30 network operators).

Understanding the relationship between network-level measurements and the underlying cables is key to accurately assessing the resiliency of the Internet[*20] Toward that understanding, we envision a service that, given a traceroute, can annotate the appropriate hops with the submarine physical links traversed.

We have started to explore this possibility using the RIPE Atlas[*21] topology data to identify submarine cable hops. RIPE Atlas is an Internet measurement project run by RIPE NCC, connecting traceroutes and other data from users around the globe. Using over 500 million traceroutes collected by the RIPE Atlas project between January and April 2018, we estimated the latency between routers at each hop using a method that we developed for estimating RTTs[*22]. Here, we use pairs of router IP addresses that appeared adjacently in traceroutes with summary statistics of their differential RTT. There is a large disparity among separate RTT data sources, but statistical processing of large quantities of data can lead to greater precision.

We then use RIPE's geolocation service[*23] to get an approximate location for each router IP address. For each IP pair for which we were able to geolocate both IPs, we then compared the geographical distance and differential RTT between them to determine whether it is possible for the path between them to traverse any of the submarine cables. Specifically, we assume the path traverses a particular

Figure 8: ACE (Africa Coast to Europe) (per TeleGeography Cable Map)

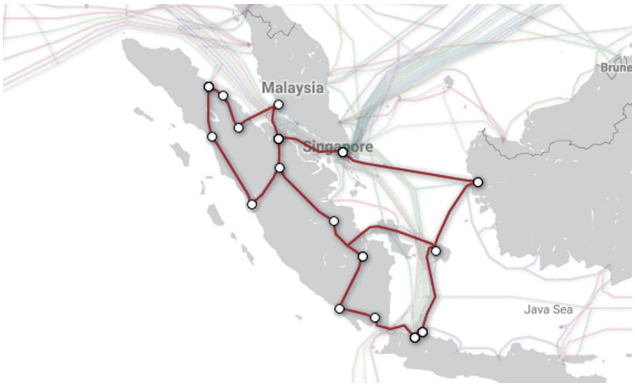

Figure 9: Telkom Indonesia's Juska (per TeleGeography Cable Map)

*20 R. Durairajan, P. Barford, J. Sommers, and W. Willinger. Intertubes: A study of the US long-haul fiber-optic infrastructure. In Proc. of ACM SIGCOMM, August 2015.
*21 RIPE NCC. RIPE Atlas (http://atlas.ripe.net).
*22 R. Fontugne, C. Pelsser, E. Aben, and R. Bush. Pinpointing delay and forwarding anomalies using large-scale traceroute measurements. In Proc. of IMC, November 2017.
*23 M. Candela. Multi-approach infrastructure geolocation. Presentation at RIPE 75, October 2017.

submarine cable and calculate the distance between the IPs through a pair of the cable's landing points. We compare this value with the distance found using the differential RTT and the speed of light in a fiber cable to assess the possibility of that particular submarine cable having been used.

After running this analysis for each pair of IPs in our RIPE Atlas dataset, we identified 3,429 unique IP pairs that could have possibly traversed a submarine cable.

While promising as a starting point, we face a number of challenges with this approach. For starters, we are unable to obtain a location for some of these routers (e.g., because data needed to get an accurate location estimate do not exist). Also, 90% of IP pairs mapped to two or more possible cables. This is not surprising given that multiple cables share similar landing points and co-location facilities, and that limits on accuracy are inherent in RTT-based analysis.

We are working on adding other methods to improve accuracy. For example, using information about which operators use each cable should help to narrow down the set of cables that could possibly be used by the AS to which IP addresses belong.

Another approach we are investigating is the use of cable outage information for cable identification. Submarine cable outages, due to maintenance or faults, are frequent. Such service outages are often reported by the news or by individuals or research groups on Twitter. Relationships can be inferred from the correlation between service outages and RTTs.

A report by Palmer-Felgate and Booi[*24] used data on over 1,000 submarine cable faults between 2008 and 2014 to create a model of cable outages and repairs. The results indicated that cables had at most two nines of availability, with the majority having outages for 9 or more days per year. By viewing historical traceroute data and comparing with reports of cable outages, we can identify IP pairs that disappear in sync with cable faults.

3.3.3 Quantifying Cable Failure Aftermaths

Mapping router IP addresses to specific physical cables will allow us to study the impact of submarine cable outages on Internet users.

Using traceroutes from RIPE Atlas, we studied the impact of a number of cable cuts in recent months. While collecting reports of submarine cable damage, we observed a number of recent outages and repairs in Southeast Asia. While these issues did not result in any major network outages, we did notice a significant impact on latency.

One of these events is damage to the SEA-ME-WE-3 cable on May 10, 2018. SEA-ME-WE-3 is one of the longest cables in the world, reaching from western Australia to western Europe via the Middle East. Once this cable was

*24 A. Palmer-Felgate and P. Booi. How resilient is the global submarine cable network? SubOptic, 2016 (https://bit.ly/2L5JHST).

damaged, certain traffic had to be rerouted via longer alternative routes, resulting in increased latency. Figure 10 shows latency measurements between Australia and Singapore before and after the cut. We see that RTTs more than tripled, from 97ms to over 320 ms. This latency spike continued for days after the cable break, as repairs to submarine cables can take weeks.

Another possible source of performance degradations is submarine network misconfiguration or maintenance.

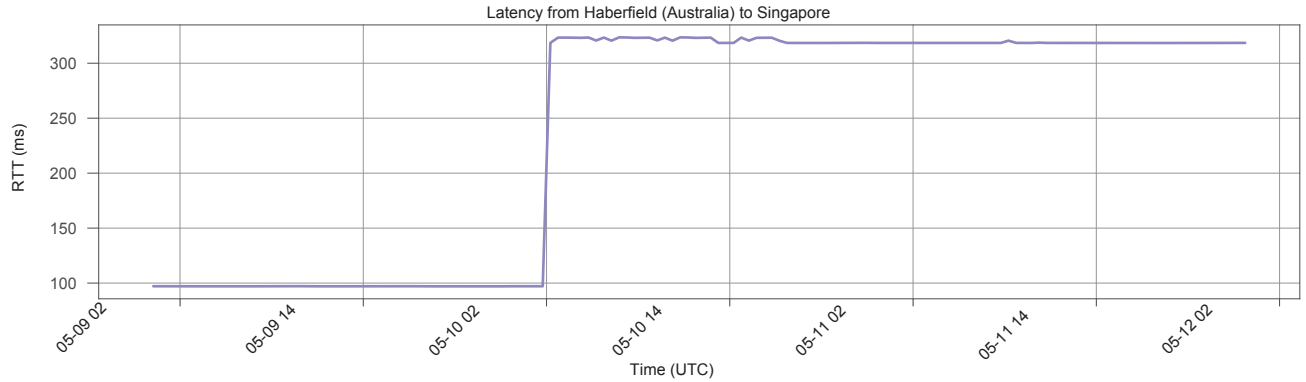

Figure 10: South-East Asia - Middle East - Western Europe 3 (SEA-ME-WE-3) undersea cable break and latency between Australia and Singapore (May 10, 2018)

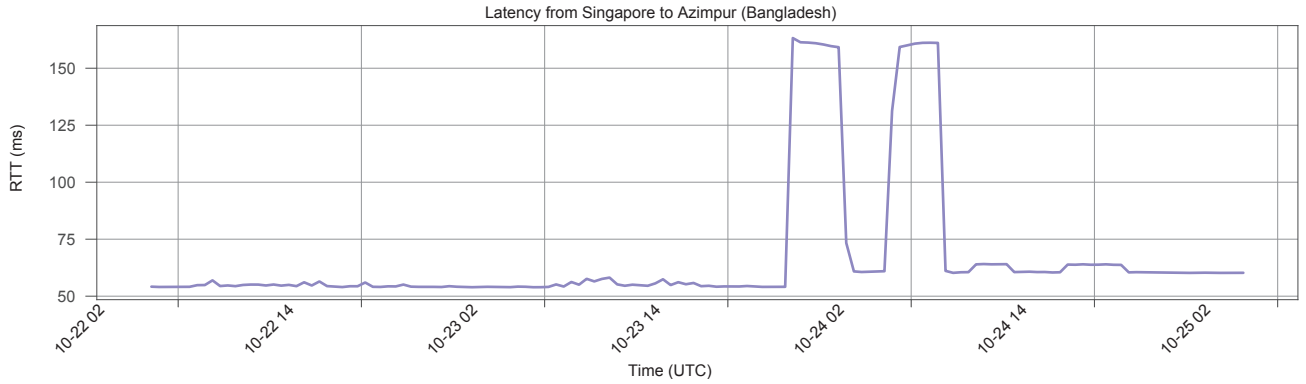

Figure 11: South-East Asia - Middle East - Western Europe 4 (SEA-ME-WE 4) cable reconfiguration (October 2017)

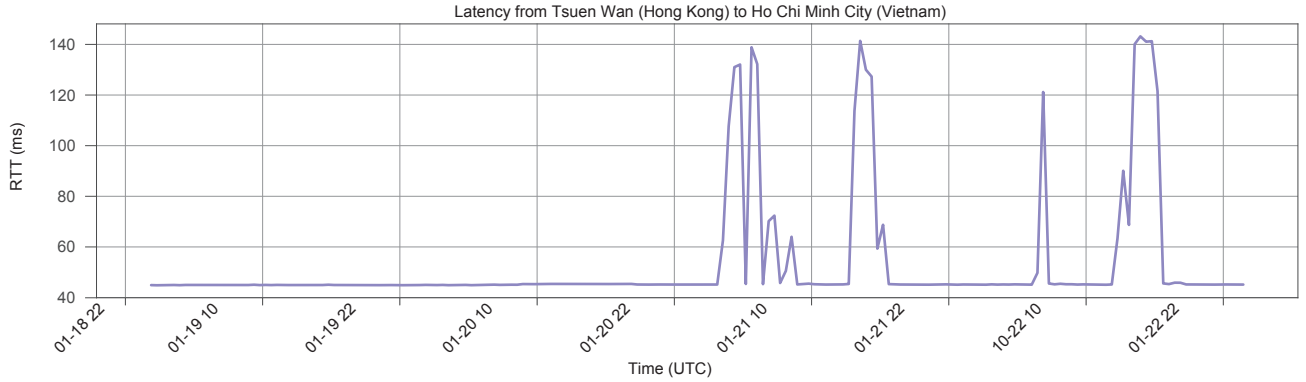

Figure 12: Latency between Hong Kong and Vietnam during Asia-America Gateway (AAG) cable reconfiguration (January 2018)

Figure 11 shows a latency increase due to reconfigurations on the SEA-ME-WE 4 submarine cable[*25]. We observed an almost tripling of latency between Singapore and Bangladesh over a period of about 12 hours.

Similarly, we observed an increase in latency between Hong Kong and Vietnam, coinciding with the reconfiguration of the Asia-America Gateway (AAG) cable starting on January 21, 2018, as shown in Figure 12.

Annotating intercontinental traceroutes with the submarine cables traversed along the path will help in diagnosing the cause of spikes such as these. Cables disappearing from traceroutes could signify a cable cut or change in routing behavior. Correlating this information will aid in understanding the underlying cause of performance anomalies.

The IP paths to submarine cables mapping can also assist network operators in understanding the dependence on a network to submarine cables. This information is important for planning future expansions of network infrastructure. For example, an operator looking to add a new upstream ISP to improve resiliency could select an ISP that uses different submarine cables from its existing providers.

Furthermore, tracking cables that appear in traceroutes would also help identify cables that are heavily utilized in a given region. Such cables could have a significant impact on performance and routing if damaged. Durairajan et. al conducted a similar study of the terrestrial long-haul fiber-optic infrastructure in the US,*20 identifying high-risk links and making suggestions for deploying new links in specific regions to reduce both risk and latency. We plan to conduct a similar analysis on the submarine network.

3.4 Conclusion

As we continue to invest on the defense of the virtual network, our limited understanding of the physical network that enables it will become its most serious vulnerability. We have put forward an approach for combining information on cables and measurements on the network layer to explore the state of the submarine cable network using publicly available data. Taking connectivity risks to physical routes into account, we believe this approach can be used to assess Internet redundancy and resiliency.

Acknowledgments
This work was funded in part by the JSPS fellowship program and NSF CNS award 1619317.

Zachary Bischof
Visiting Researcher, IIJ Innovation Institute
Zachary conducts experimental research on networks and large-scale distributed systems. He aims to characterize broadband networks through DNS and traffic analysis.

Romain Fontugne
Senior Researcher, IIJ Innovation Institute
Fabián E. Bustamante
Professor at Northwestern University, United States

*25 T. D. Star. Internet to be slow for next 4 days (https://bit.ly/2LmlNSn).

February 2019 Vol.41

Internet Initiative Japan

About Internet Initiative Japan Inc. (IIJ)

IIJ was established in 1992, mainly by a group of engineers who had been involved in research and development activities related to the Internet, under the concept of promoting the widespread use of the Internet in Japan.
IIJ currently operates one of the largest Internet backbones in Japan, manages Internet infrastructures, and provides comprehensive high-quality system environments (including Internet access, systems integration, and outsourcing services, etc.) to high-end business users including the government and other public offices and financial institutions.
In addition, IIJ actively shares knowledge accumulated through service development and Internet backbone operation, and is making efforts to expand the Internet used as a social infrastructure.

IIJ-MKTG020-0039

Internet Initiative Japan Inc.

Address: Iidabashi Grand Bloom, 2-10-2 Fujimi, Chiyoda-ku,
Tokyo 102-0071, Japan
Email: info@iij.ad.jp URL: https://www.iij.ad.jp/en/