

The Impact of Residential Broadband Traffic on Japanese ISP Backbones

Kensuke Fukuda
NTT/WIDE
fukuda@t.ecl.net

Kenjiro Cho
IIJ/WIDE
kjc@ijlab.net

Hiroshi Esaki
U.Tokyo/WIDE
hiroshi@wide.ad.jp

ABSTRACT

This paper investigates the effects of the rapidly-growing residential broadband traffic on commercial ISP backbone networks. We collected month-long aggregated traffic logs for different traffic groups from seven major ISPs in Japan in order to analyze the macro-level impact of residential broadband traffic. These traffic groups are carefully selected to be summable, and not to count the same traffic multiple times.

Our results show that (1) the aggregated residential broadband customer traffic in our data exceeds 100Gbps on average. Our data is considered to cover 41% of the total customer traffic in Japan, thus we can estimate that the total residential broadband traffic in Japan is currently about 250Gbps in total. (2) About 70% of the residential broadband traffic is constant all the time. The rest of the traffic has a daily fluctuation pattern with the peak in the evening hours. The behavior of residential broadband traffic deviates considerably from academic or office traffic. (3) The total traffic volume of the residential users is much higher than that of office users, so backbone traffic is dominated by the behavior of the residential user traffic. (4) The traffic volume exchanged through domestic private peering is comparable with the volume exchanged through the major IXes. (5) Within external traffic of ISPs, international traffic is about 23% for inbound and about 17% for outbound. (6) The distribution of the regional broadband traffic is roughly proportional to the regional population.

We expect other countries will experience similar traffic patterns as residential broadband access becomes widespread.

1. INTRODUCTION

The availability of residential broadband access has made tremendous advances over the past few years, especially in Korea and Japan where both the penetration rate and the average line speed are much higher than other countries. A government survey shows that there are 14.5 million broadband subscribers in Japan as of February 2004; 11 million DSL subscribers, 2.5 million CATV Internet subscribers,

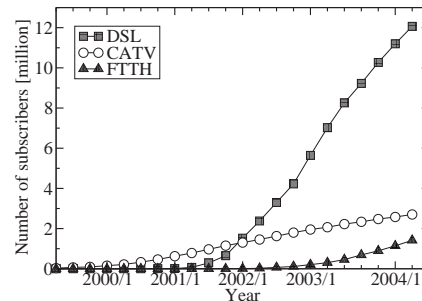


Figure 1: Increase of residential broadband subscribers in Japan

and 1 million FTTH subscribers [16]. The number of broadband access subscribers is still increasing as shown in Figure 1 [16]. At the same time, broadband access technologies are shifting to higher speed such as 50Mbps DSL and 100Mbps FTTH.




Figure 2: Traffic growth at the major Japanese IXes

As residential broadband access becomes widespread, we are observing an unprecedented traffic increase on commercial backbone networks. Figure 2 shows the aggregated peak traffic at major IXes (JPNAP[8], JPIX[7], and NSPIX[10]) in Japan, and illustrates the growth in backbone traffic [16]. The impact of residential broadband traffic is not only in volume but also in usage patterns. The peak hours have shifted from office hours to evening hours, and emerging file sharing or other peer-to-peer communications with audio/video contents exhibit behavior considerably different from traditional world wide web [14, 6]. There are striking differences in traffic patterns from earlier observations [13, 4, 9, 3, 5].

Although a drastic change in backbone traffic has already been observed, it is difficult to plan for the future because residential broadband traffic is undergoing a transformation; new innovations in access networking technologies continue to be developed, and new applications as well as their usage are emerging to take advantage of low-cost high-speed connectivity.

There is a strong concern that, if this trend continues, Internet backbone technologies will not be able to keep up with the rapidly growing residential traffic. Moreover, commercial ISPs will not be able to invest in backbone networks simply for low-profit residential traffic.

It is critical to ISPs and policy makers to understand the effects of growing residential broadband traffic but it is difficult both technically and politically to obtain traffic data from commercial ISPs. Most ISPs are collecting traffic information for their internal use but such data contain sensitive information and are seldom made available to others. In addition, measurement methods and policies differ from ISP to ISP so that it is in general not possible to compare a data set with another set obtained from a different ISP.

In order to seek out a practical way to investigate the impact of residential broadband traffic on commercial backbone networks, we have formed an unofficial study group with specialists including members from seven major commercial ISPs in Japan.

Our goal is to identify the macro-level impact of residential broadband traffic on ISP backbones. More specifically, we are trying to obtain a clearer grasp of the ratio of residential broadband traffic to other traffic, changes in traffic patterns, and regional differences across different ISPs. As the first step, we have collected aggregated bandwidth usage logs for different traffic groups. Such statistics will provide reference points for further detailed analysis, most likely by sampling methods. In this paper, we report findings in our data sets that residential broadband traffic presents a significant impact on ISP backbones.

2. METHODOLOGY

There are several requirements in order to solicit ISPs to provide traffic information. We need to find a common data set which all the participating ISPs are able to provide. The required workload and investment for ISPs to provide the data set should not be high. The data set should be coarse enough not to reveal sensitive information about the ISP but be meaningful enough so that the behavior of residential broadband traffic can be analyzed. It is also desirable to be able to cross-check the consistency of the results with other data sets. The data sets should be summable in order to aggregate them with those provided by other ISPs.

We found that most ISPs collect interface counter values of almost all routers in their service networks via SNMP, and archive per-interface traffic logs using MRTG [12] or RRDtool [11]. Thus, it is possible for the ISPs to provide aggregated traffic information if they can classify router interfaces into a common set.

Our focus is on traffic crossing ISP boundaries which can be roughly divided into customer traffic, and external traffic such as peering and transit. For practical purposes, we selected the 5 traffic groups shown in Figure 3 for data collection. The descriptions of the groups are in Table 1. It is impossible to draw a strict line for grouping (e.g. residential/business and domestic/international) on the global

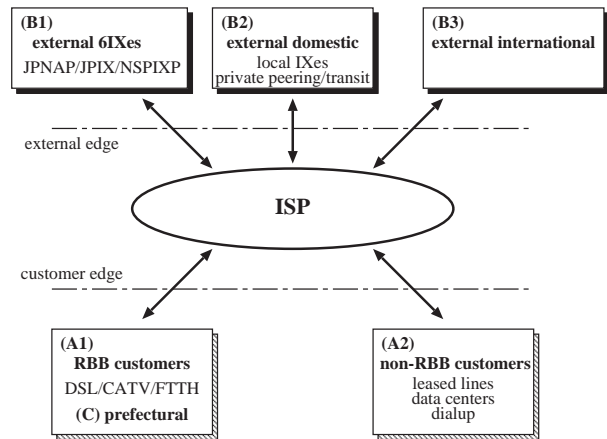


Figure 3: 5 traffic groups at ISP boundary for data collection

Internet so that these groups are chosen by the existing operational practice of the participating ISPs. We re-aggregate each ISP's aggregated logs, and only the resulting aggregated traffic is used in our study so as to not reveal a share of each ISP.

Our main focus is on (A1) RBB (Residential Broadband) customers but other items are used to understand the relative volume of (A1) with respect to other types of traffic as well as to cross-check the correctness of the results. (A2) non-RBB customers is used to obtain the ratio of residential broadband traffic to total customer traffic. The total customer traffic (A) is $A = (A1) + (A2)$. (B1) external 6IXes and (B2) external domestic are used to estimate the coverage of the collected data sets. (B3) external international is used to compare domestic traffic with international traffic. The total external traffic (B) is $B = (B1) + (B2) + (B3)$. (C) prefectural is to observe regional differences. This group covers only 2 major residential broadband carriers who provide aggregated links per prefecture to ISPs; other carriers' links are not based on prefectures. This group is a subset of (A1).

In general, it is meaningless to simply sum up traffic values from multiple ISPs since a packet could cross ISP boundaries multiple times. Customer traffic is, however, summable because a packet crosses customer edges only once in each direction, when entering the source ISP and exiting the destination ISP. The numbers for external traffic are overestimated since a packet could be counted multiple times if it travels across more than 2 ISPs. However, the error should be relatively small in this particular result since these ISPs are peering with each other.

We collected month-long traffic data that was sampled every two hours from the participating ISPs because a 2-hour resolution is the highest common factor for month-long data. This is because both MRTG and RRDtool aggregate old records into coarser records in order to bound the database size. In MRTG, 2-hour resolution records are maintained for 31 days in order to draw monthly graphs. RRDtool does not have fixed aggregation intervals but most operators configure RRDtool to maintain 1-hour or 2-hour resolution records for a period longer than needed for monthly graphs.

We developed a perl script to read a list of MRTG and

Table 1: Descriptions of traffic groups

traffic group	description	notes
(A1) RBB customers	residential broadband customer lines	includes small business customers using RBB
(A2) non-RBB customers	includes leased lines, data centers, dialup lines	may include RBB customers behind leased lines
(B1) external 6IXes	links for 6 major IXes (JPNAP/JPIX/NSPIX in Tokyo/Osaka)	
(B2) external domestic	external domestic links other than the 6IXes (regional IXes, private peering, transit)	domestic: both link-ends in Japan. includes domestic peering with global ASes
(B3) external international	external international links	
(C) prefectural	RBB links divided into 47 prefectures in Japan	prefectural links from 2 RBB carriers

RRDtool log files, and aggregate traffic measurements for a give period with a given resolution. It outputs “timestamp, in-rate, out-rate” for each time step. Another script produces a graph using RRDtool. We provided the tools to the ISPs so that each ISP can create aggregated logs by themselves. It allows ISPs not to disclose the internal structure of their network or unneeded details of its traffic.

The biggest workload for the ISPs is to classify the large number of per-interface traffic logs and create a log list for each group. For large ISPs, the total number of the existing per-interface traffic logs exceeds 100,000. To reduce the workload, ISPs are allowed to use the internal interface of a border router instead of a set of external (edge) interfaces if the traffic on the internal interface is an approximation of the sum of the external interfaces. In this case, we instruct the tool to swap “in” and “out” records since the notation in the per-interface logs depicts the perspective of the routers but inbound/outbound records in our data sets signify the ISPs’ point of view.

3. RESULTS

We analyzed traffic logs for September and October in 2004 from seven major ISPs in Japan. Each ISP provided traffic logs with 2-hour resolution for those two months. The results were obtained by aggregating all the traffic logs provided by the seven ISPs. 2-hour boundaries were computed in UTC by MRTG and RRDtool so that they fell on odd hours in Japanese Standard Time that is nine hours ahead of UTC.

For weekly data analysis, we took the averages of the same weekdays in the month. We excluded two holidays in September and one holiday in October from the weekly analysis since their traffic pattern is closer to that of weekends. We also excluded another two days in October from the weekly analysis as one ISP failed to record traffic logs during this period.

3.1 Customer Traffic

Figure 4 shows the weekly traffic of RBB customers, consisting of DSL/FTTH/CATV residential users (A1). This group also includes small business customers using residential broadband access. Note that the plot is the mean rate and not the peak rate, even though the peak rate is often used for operational purposes. The residential broadband customer traffic has already exceeded 100Gbps in total. The inbound and outbound traffic are almost equal, and about 70Gbps is constant for both directions, probably due to peer-to-peer applications which generate traffic




Figure 4: Aggregated RBB customer weekly traffic in September 2004. Darker vertical dotted lines indicate the start of the day (0:00 am in local-time).

independent of daily user activities. The diurnal pattern indicates that home user traffic is dominant, i.e., the traffic increases in the evening, and the peak hours are from 21:00 to 23:00. Weekends can be identified by larger daytime traffic although the peak rates are close to weekdays. The outbound traffic to customers is slightly larger than the inbound, even though it is often assumed that home users’ downstream traffic is much larger than upstream. We believe that peer-to-peer applications contribute significantly to the upstream traffic.




Figure 5: Aggregated non-RBB customer weekly traffic in September 2004

Figure 5 shows the weekly traffic of non-RBB customers (A2). This group contains leased lines, data centers, and other customers (e.g., dialup customers). It also includes leased lines used to accommodate residential broadband access within the customer networks (e.g., second or third level ISPs) since ISPs do not distinguish them from other leased lines. As a result, the traffic pattern still appears to be dominated by residential traffic, which is indicated by the peak hours and the differences between weekdays and weekends. However, we also observe office hour traffic (from 8:00 to 18:00) in the daytime on weekdays but traditional office commercial traffic appears to be smaller than residential

customer traffic. Note that we cannot directly compare the traffic volume of (A2) with that of (A1) because (A2) was provided by only four of the seven ISPs.




Figure 6: Aggregated total traffic from ABILENE in October 2004. Time is in CDT.

The traffic patterns common to Figure 4 and 5 are quite different from well-known academic or business usage patterns. For example, Figure 6 shows the weekly traffic of ABILENE [1], an Internet2 backbone network for universities and research labs. From Figure 6, it is clear that office hour traffic is dominant; traffic peaks occur around noon, and there is less user activity on weekends.

3.2 External Traffic




Figure 7: Weekly external traffic to/from the 6 major IXes in September 2004

The external traffic groups are used to understand the total traffic volume in Japan. Figure 7 shows traffic to and from the six major IXes (B1). It is apparent that the traffic behavior is strongly affected by residential traffic.




Figure 8: Weekly other domestic external traffic in September 2004

Figure 8 shows the external domestic traffic (B2) including regional IXes, private peering and transit but not including traffic for the six major IXes. The traffic pattern is very similar to Figure 7.

Figure 9 shows international traffic (B3). The inbound traffic is much larger than the outbound, and the traffic pattern is clearly different from the domestic traffic. The peak hours are still in the evening, but outbound traffic volume fluctuates less than inbound traffic, suggesting that




Figure 9: Weekly international external traffic in September 2004

the traditional behavior of content downloading to Japan still dominates international traffic.

3.3 Prefectural Traffic

In order to investigate regional differences (i.e., between metropolitan and rural areas), we collected regional traffic rates of the 47 prefectures. Figure 10 illustrates aggre-

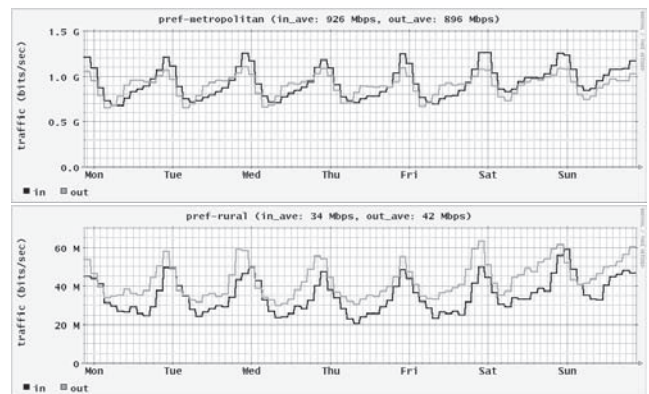


Figure 10: Example prefectural traffic: a metropolitan prefecture (top) and rural prefecture (bottom)

gated traffic of one metropolitan prefecture (top graph) and of one rural prefecture (bottom graph). Both graphs exhibit similar temporal patterns such as peak positions and weekday/weekend behavior. In addition, about 70% of the average traffic is constant regardless of the traffic volume. These characteristics are common to other prefectures. One noticeable difference found is that metropolitan prefectures experience larger volumes of office hour traffic, probably due to heavy business usage.

Figure 11 is a scatter plot of traffic and population for the 47 prefectures. We found that a prefecture's traffic is roughly proportional to the population of the prefecture. We obtained similar results when the number of Internet users found in [15] is used instead of the population. The result indicates that there is no clear regional concentration of heavy hitters of the Internet. That is, the probability of finding a heavy hitter in a given population is constant.

In order to analyze the scaling property of traffic volume — to find a typical size of prefectural traffic volume, we show the (complementary) cumulative distribution of prefectural traffic on a log-log scale in Figure 12. The plot conforms to a power law distribution with a cutoff point at 700Mbps, meaning that there is no typical size of prefectural traffic volume. In other words, most prefectures generate a small




Figure 11: Relationship between population and traffic for prefectures

amount of traffic, still prefectures with high traffic volume are observable with a certain probability. It is also observed that the plots for the top 5 largest prefectures deviate from the power law. To investigate this power law decay, we show the cumulative distribution of prefectural populations in the sub-panel. The plots reveal that the power law appearing in traffic volume is derived from the power law decay of prefectural populations, as can be inferred from the linear relationship between traffic and populations in Figure 11. Thus, we can conclude that the probability of finding a heavy hitter in a given population is constant and the distribution of aggregated traffic volume directly depends on the pc

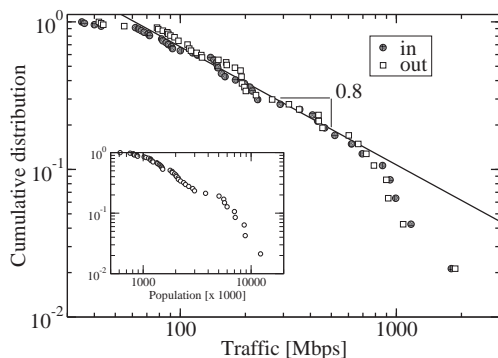


Figure 12: Cumulative distribution of prefectural traffic. Sub-panel indicates the cumulative distribution of populations for comparison.

3.4 Summary of Traffic

The monthly average rates in bits/second of the traffic groups are shown in Tables 2 through 5.

Table 2 is the average rates of aggregated customer traffic. As explained before, the non-RBB customer traffic was obtained only from the four ISPs so that it is difficult to directly compare (A1) with (A2). Thus, we estimated the ratio of the RBB customer traffic (A1) to the total customer traffic (A) from only four ISPs' data with both (A1) and (A2). The estimated ratio $(A1)/(A1+A2)$ is 65% for inbound and 67% for outbound.

Table 2: Average rates of aggregated customer traffic

	(A1)customer-RBB (7 ISPs)		(A2)customer-non-RBB (4 ISPs)	
	inbound	outbound	inbound	outbound
Sep	98.1G	111.8G	14.0G	13.6G
Oct	108.3G	124.9G	15.0G	14.9G

Table 3: Average rates of aggregated external traffic

	(B1)ext-six (7 ISPs)		(B2)ext-dom (7 ISPs)		(B3)ext-intl (7 ISPs)	
	in	out	in	out	in	out
Sep	35.9G	30.9G	48.2G	37.8G	25.3G	14.1G
Oct	36.3G	31.8G	53.1G	41.6G	27.7G	15.4G

Table 3 summarizes the average rates of aggregated external traffic. We observe that the total volume of external domestic traffic (B2), mainly private peering, exceeds the volume for the six major IXes (B1). From this result, it can be concluded that simply relying on data from IXes to estimate and understand nation-wide traffic may be misleading, because a considerable amount of traffic is exchanged by private peering. At the same time, it is possible that the volume of private peering is larger in our measurement than the rest of the Japanese ISPs because private peering is usually exercised only between large ISPs. The ratio of international traffic to the total external traffic is 23% for inbound and 17% for outbound.

Table 4: Average rates of total customer traffic and total external traffic

	(A)customer(A1+A2)		(B)external(B1+B2+B3)	
	inbound	outbound	inbound	outbound
Sep	112.1G	125.4G	109.4G	82.8G
Oct	123.3G	139.8G	117.1G	88.8G

There is a relationship between the total customer traffic (A1 + A2) and the total external traffic (B1 + B2 + B3) in Table 4. If we assume all inbound traffic from other ISPs is destined to customers, the inbound traffic volume for the total external traffic (B) should be close to the outbound traffic volume for the total customer traffic (A). Similarly, the outbound traffic volume of (B) should be close to the inbound traffic volume of (A). However, the non-RBB customer data is provided by only 4 ISPs. If we interpolate the missing ISPs in the non-RBB customer traffic using the ratio from the four reporting ISPs, the total inbound customer traffic is estimated to be 152.1Gbps, and that outbound to be 167.8Gbps. Though these volumes are higher than those for the total external traffic, this is probably because the total customer traffic contains traffic whose source and destination belong to the same ISP.

Lastly, we examined the relationship between our IX traffic data (B1) and the total input rate of the six major IXes, as obtained directly from these IXes [16]. In comparison with the published total incoming traffic of these IXes, our data represent 41% of the total traffic as shown in Table 5. If we assume this ratio to be the traffic share of the seven ISPs, the total amount of residential broadband traffic in Japan is roughly estimated to be 250Gbps.

To check consistency, we collected the September results




Figure 13: Cumulative distribution of daily traffic per user: all prefectures (left), a metropolitan prefecture (middle) and a rural prefecture (right)

Table 5: IX traffic observed from ISPs and from IXes

	(B1)ext-6ix outbound	traffic observed by IXes inbound
Sep	30.9G	74.5G
Oct	31.8G	77.1G

and the October results separately in October and November respectively. These results are consistent so that we are fairly confident about their accuracy. However, the traffic increase from September to October was higher than our projection; the traffic of the six IXes increased by only about 3% but the other groups increased by about 10%. We suspect that some links missing in the September measurements could have been added later for the October measurements, and we expect further measurement will shed light on this issue.

3.5 Distribution of per-customer traffic

In order to verify our assumption that the distribution of heavy hitters is similar across different regions, we obtained per-customer traffic information for October 2004 from one of the participating ISPs. The inbound/outbound traffic volumes of residential broadband customers for each prefecture were collected by means of sampled NetFlow [2] and matching customer IDs with the assigned IP addresses. Although this data set is from only one ISP, the results appear to be consistent with the aggregated results. The results are also consistent with earlier measurements on peer-to-peer traffic by Sen and Wang [14]; peer-to-peer traffic is extremely variable and highly skewed among participating nodes.

Figure 13 shows the (complementary) cumulative distribution of daily traffic per customer on a log-log scale, and compares all the prefectures (left) with one metropolitan prefecture (middle) and one rural prefecture (right). The daily traffic volume is the average of the month, and the distribution is computed independently for inbound and outbound traffic. It is common to the three plots that about 4% of the customers use more than 2.5GB/day (or 230kbits/sec) and, beyond this point, the slope of the distribution changes. Thus, heavy hitters can be statistically identified as customers using more than 2.5GB/day. The distribution also shows that outbound traffic is dominant for most customers but it does not hold for heavy hitters. These trends are con-

sistent across different prefectures, and the differences are only in the tail length affected by the number of customers, which confirms that the distribution of heavy hitters is similar across

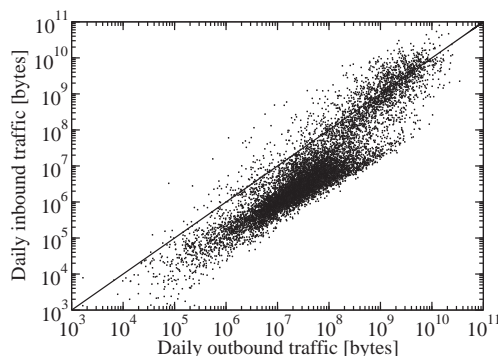


Figure 14: Correlation of inbound and outbound traffic volumes in one metropolitan prefecture

Figure 14 is a log-log scatter plot to show the correlation between inbound and outbound traffic volumes for each customer. There is a positive correlation as expected, and the highest density cluster is below and parallel to the unity line where the volume of outbound (down-streaming for customers) is about ten times larger than that of inbound. This is probably due not only to application characteristics but also to the restriction of asymmetric access lines. In a higher volume region, a different cluster appears to exist around the unity line. The slope of the cluster seems to be slightly larger than 1, which explains the inversion of inbound and outbound traffic volumes in Figure 13. A plausible interpretation of excess upstream traffic of heavy hitters is that symmetric high bandwidth of FTTH access lines complements the shortage of upstream bandwidth of DSL lines. It can be also observed that, across the entire traffic volume range, the inbound/outbound traffic ratio varies greatly, up to 4 orders of a magnitude. This plot is taken from a metropolitan prefecture but the correlation is common to all prefectures.

Figure 15 shows the cumulative distribution of traffic volume of all of the prefectures with heavy hitters in decreasing order of volume. Again, the distribution is computed independently for inbound and outbound traffic. The graph reveals skewed traffic distribution among customers; the top




Figure 15: Cumulative distribution of traffic volume with heavy hitters in decreasing order of volume

$N\%$ of heavy hitters use $X\%$ of the total traffic. For example, the top 4% of customers at the knee point of the distribution in Figure 13 use 75% of the total inbound traffic, and 60% of the outbound.

4. CONCLUSION

The widespread deployment of residential broadband access has tremendous implications to our lives. Although its effects to the Internet infrastructure are difficult to predict, it is essential for ISPs to prepare for the future to accommodate innovations brought by empowered end-users.

Residential broadband traffic has already contributed to a significant increase in commercial backbone traffic. In our study, residential broadband traffic accounts for two thirds of the ISP backbone traffic, which should have a significant impact on the pricing and cost structures of the ISP business.

The properties of residential broadband traffic differ considerably from those of academic or office traffic often seen in literature. The constant portion of daily traffic fluctuations is about 70%, much larger than ones found in earlier reports [3, 5]. Research results obtained from campus or other academic networks may no longer apply to commercial traffic. More research efforts should be directed to measurement and analysis of residential broadband traffic.

The inbound/outbound rates are roughly equal throughout our data sets. Many access technologies employ asymmetric line speed for inbound and outbound based on the assumption that content-downloading is dominant for normal users. However, this assumption does not hold in our measurements.

Our measurements also suggest that a large amount of traffic is exchanged by private peering so that data from IXes may not be an appropriate index of nation-wide traffic volume.

The prefectural results show that traffic volume is roughly proportional to regional population. It indicates a unique characteristic of the cyber-world in which activities are not bound by time and place. If this is the case, it would affect the design of capacity planning for the future Internet.

For future work, we will continue collecting aggregated traffic logs from ISPs. We are also planning to do more detailed analysis of residential broadband traffic by selecting a few sampling points.

Acknowledgments

We are grateful to the following ISPs for their support and assistance with our data collection; IIJ, Japan Telecom, K-Opticom, KDDI, NTT Communications, POWEREDCOM, and SOFTBANK BB. We would like to thank the Ministry of Internal Affairs and Communications of Japan for their support in coordinating our study, and for providing the statistics of broadband subscribers. We would also like to thank Akira Kato of the University of Tokyo, Atsushi Shionozaki of Sony CSL, Randy Bush of IIJ America, and anonymous reviewers for their valuable input.

5. REFERENCES

- [1] ABILENE aggregate traffic statistics. <http://abilene.internet2.edu/observatory/summary-view.html>.
- [2] Cisco Sampled NetFlow. http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120limit/120s/120s11/12s_sanf.htm.
- [3] A. Feldmann, A. G. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True. Deriving traffic demands for operational IP networks: methodology and experience. In *SIGCOMM*, pages 257–270, Stockholm, Sweden, 2000.
- [4] M. Fomenkov, K. Keys, D. Moore, and k. claffy. Longitudinal study of Internet traffic in 1998-2003. In *WISICT*, Cancun, Mexico, Jan. 2004.
- [5] C. Fraleigh, S. Moon, B. Lyles, C. Cotton, M. Khan, D. Moll, R. Rockell, T. Seely, and C. Diot. Packet-level traffic measurements from the sprint IP backbone. *IEEE Network*, pages 6–16, 2003.
- [6] K. P. Gummadi, R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan. Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In *SOSP-19*, pages 314–329, Bolton Landing, NY, Oct. 2003.
- [7] Japan Internet Exchange Co., Ltd. (JPIX). <http://www.jpix.co.jp>.
- [8] Multifeed JPNAP service. <http://www.jpnap.net>.
- [9] S. McCreary and k. claffy. Trends in wide area IP traffic patterns. In *ITC Specialist Seminar*, Monterey, CA, Sept. 2000.
- [10] NSPIX: Network Service Provider Internet eXchange Project. <http://nspixp.wide.ad.jp>.
- [11] T. Oetiker. RRDtool: Round Robin Database Tool. <http://ee-staff.ethz.ch/~oetiker/webtools/rrdtool/>.
- [12] T. Oetiker. MRTG: The multi router traffic grapher. In *USENIX LISA*, pages 141–147, Boston, MA, Dec. 1998.
- [13] S. Saroiu, K. P. Gummadi, R. Dunn, S. D. Gribble, and H. M. Levy. An analysis of Internet content delivery systems. In *OSDI 2002*, pages 315–328, Boston, MA, Dec. 2002.
- [14] S. Sen and J. Wang. Analyzing peer-to-peer traffic across large networks. In *Second ACM SIGCOMM Internet Measurement Workshop*, pages 137–150, Marseille, France, Nov. 2002.
- [15] White book for the Internet 2004 (in Japanese). The Internet Association Japan, July 2004.
- [16] A study group report on the next generation IP-based infrastructure (in Japanese). The Ministry of Internal Affairs and Communications of Japan, June 2004.

