Trends in Japanese residential traffic

Kenjiro Cho (IIJ), Kensuke Fukuda (NII),
Hiroshi Esaki (Univ. of Tokyo)

August 2012

trends in residential broadband traffic

global trends

- 2002-2005: p2p file-sharing dominated traffic
- 2006-2010: decline of p2p, rise of hypergiants
  - decline of p2p
    - throttling and volume caps by ISPs
    - campaigning, robbying, court actions by RIAA and others
    - rise of alternative services (e.g., iTunes, Netflix)
    - rise of direct download sites (e.g., MegaUpload)
  - everything on TCP port 80
  - visible players: content providers, CDNs, hosting sites
- 2010-: rise of video streaming services (e.g., Netflix)

differences in Japan

- much higher fiber penetration rate
  - traffic has been stable with modest growth rate (20%/year)
- p2p file-sharing
  - winny's success and failuer, bittorrent not so popular
- slower decline of p2p
  - higher skew in users' traffic usage due to higher bandwidth
- slower takeoff of video streaming services

explosive traffic growth by video content?

many media reports on explosive traffic growth by video content

OPINION | JANUARY 20, 2007

The Coming Exaflood

By BRET SWANSON

Article

Email Printer Friendly Share: Yahoo Buzz

Today there is much praise for YouTube, MySpace, blogs and all the c
technologies that are allowing you and me to transform media and com
Internet applications are at risk, thanks to the regulatory implications of
Proponents of this concept -- including Democratic Reps. John Dingell

Forbes.com

Search: Forbes.com Quotes Video

U.S. EUROPE ASIA HOME PAGE FOR THE WORLD'S BUSINESS LE

HOME BUSINESS TECH MARKETS ENTREPRENEURS LEADERSHIP

E-Mail | Print | Comments | Request Reprints | E-Mail Newsletters | RSS

Commentary

Information Super Traffic Jam

Phil Kerpen 01.31.07, 6:00 AM ET

Phil Kerpen

WASHINGTON, D.C. -A new assessment from Deloitte & Touche predicts that global traffic will exceed the Internet's capacity as soon as this year. Why? The rapid growth in the number of global Internet users, combined with the rise of online video services and the lack of investment in new infrastructure. If Deloitte's predictions are accurate, the traffic on many Internet backbones could slow to a crawl this year absent substantial new infrastructure investments and deployment.

Uncertainty over potential network neutrality requirements is one of the major factors delaying necessary network upgrades. The proponents of such regulations are back on the offensive, heartened by sympathetic new Democratic majorities and the concession made by **AT&T** (nyse: T - news - people) in its **BellSouth** (nyse: BLS - news - people) acquisition. The Google/MoveOn.org coalition fighting for network neutrality mandates calls itself "Save the Internet." But the Internet doesn't need to be saved--it needs to be improved, expanded and bulked up. An attempt to "save" the Internet in its current state would be something akin to saving the telegraph from the telephone.

USA TODAY

Home News Travel Money Sports Life Tech

Money Markets Economy Companies & Execs Media Cars Personal Finance Real E

GET A QUOTE: Enter symbol(s) or Keywords GO DJIA 8,149.09 -679.95 NASDAQ 1,3

Video, interactivity could nab Web users by '10

Updated 11/20/2007 12:29 AM | Comments 52 | Recommend 14 E-mail | Save | Print | Reprints & Permissions | RSS

By David Lieberman, USA TODAY

NEW YORK — Enjoy your speedy broadband Web access while you can.

The Web will start to seem pokey as early as 2010, as use of interactive and video-intensive services overwhelms local cable, phone and wireless Internet providers, a study by business technology analysts Nemertes Research has found.

"Users will experience a slow, subtle degradation, so it's back to the bad old days of dial-up," says Nemertes President Johna Till Johnson. "The cool stuff that you'll want to do will be such a pain in the rear that you won't do it."

Mixx it

Other ways to sha

Digg

del.icio.us

Newsvine

Reddit

Facebook

What's this?

modest traffic growth?

but technical sources report only modest traffic growth worldwide

- MINTS: 40-50% in U.S. and worldwide
- Cisco visual networking index: worldwide growth of 29% per year for 2011-2016 (down from 32% in the 2011 report)

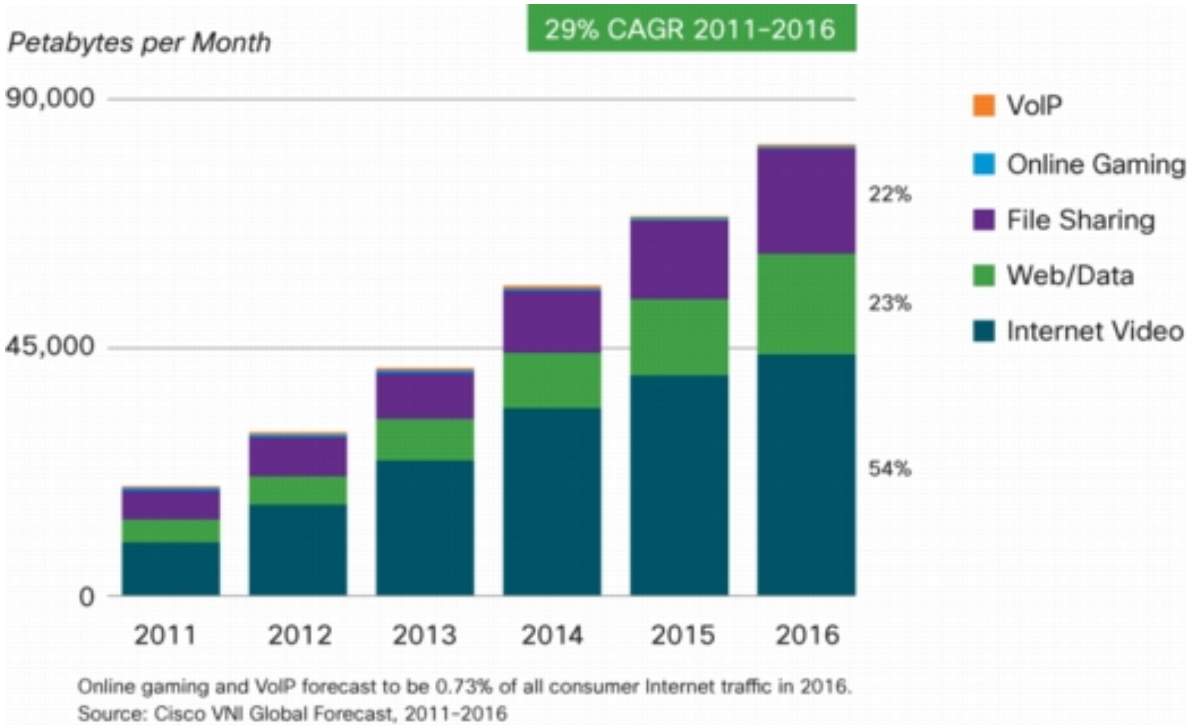

source: Entering the Zettabyte Era (Cisco 2012/5)

- TeleGeography: network capacity also grows by 50% per year

impact of Netflix in US

- Netflix is reported to be responsible for 30% of download traffic in US

	Upstream Traffic		Downstream Traffic		Total Traffic	
Rank	Application	Share	Application	Share	Application	Share
1	BitTorrent	52.01%	Netflix	29.70%	Netflix	24.71%
2	HTTP	8.31%	HTTP	18.36%	BitTorrent	17.23%
3	Skype	3.81%	YouTube	11.04%	HTTP	17.18%
4	Netflix	3.59%	BitTorrent	10.37%	YouTube	9.85%
5	PPStream	2.92%	Flash Video	4.88%	Flash Video	3.62%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS

sandvine

Table 1 - North America - Top Applications by Bytes (Peak Period, Fixed Access)

source: Global Internet Phenomena Spotlight. Sandvine. May 2011.

Netflix on CDN Market

- Netflix starting to deploy its own CDN

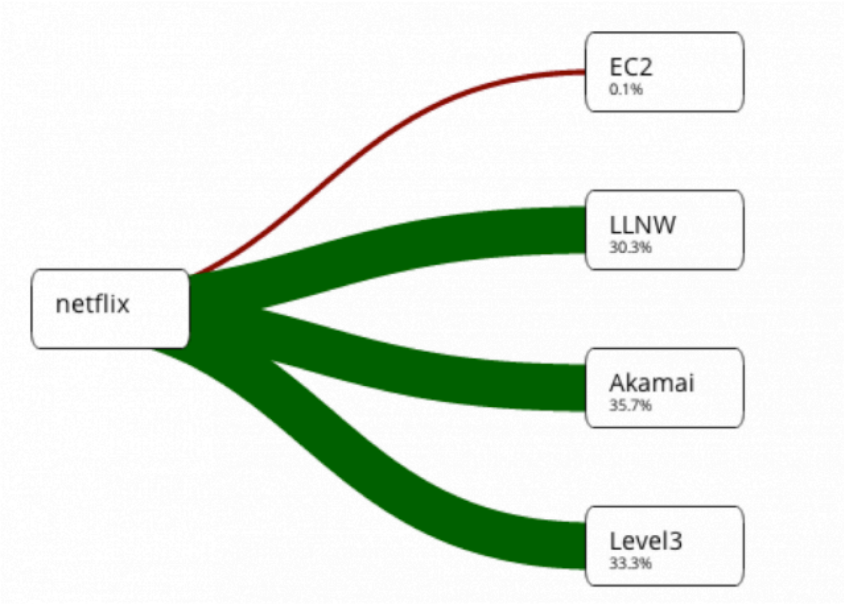

Average Netflix streaming video distribution in several North American networks at the start of 2012.

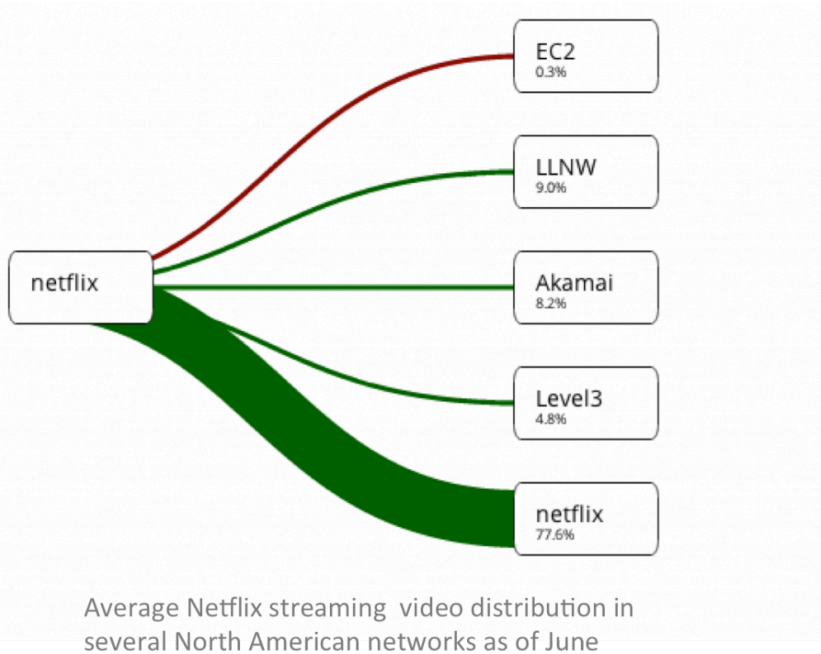

Average Netflix streaming video distribution in several North American networks as of June 2012.

source: First Data on Changing Netflix and CDN Market Share
Craig Labovits. DeepField Networks blog. June 2012.

Filesharing Market

- HyperGiants: responsigle for 50% of Internet traffic
- the rest of 50%: looks diverse, but actually not
  - there exist common infrastructures for file-sharing, adult, hosted p2p sites

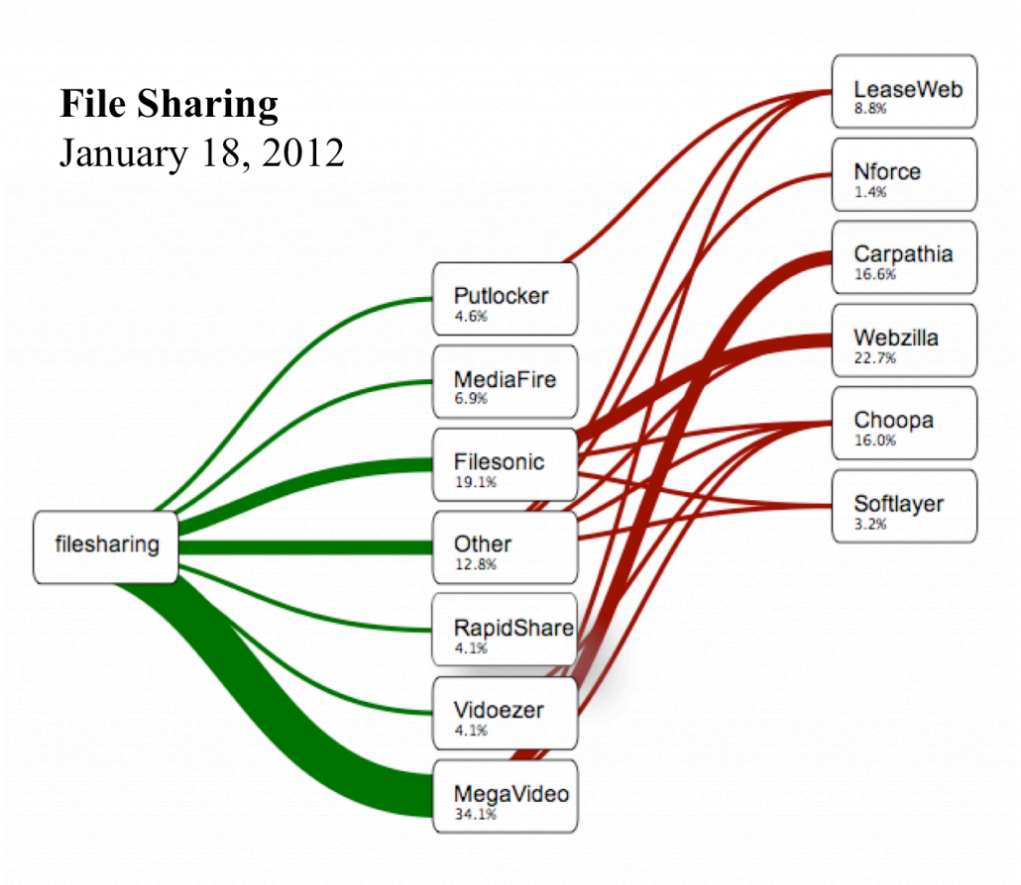

source: The Other 50% of Internet Traffic. Craig Labovits. NANOG54. Feb 2012.

this talk

key question: what is the macro level impact of video and other rich media content on Japanse residential traffic?

- measurements: 2 data sets
  - aggregated SNMP data from 6 ISPs covering 42% of Japanese traffic
  - Sampled NetFlow data from 1 ISP

motivation

why is traffic growth important?

- one of the key factors driving research, development and investiment in technologies and infrastructures
  - with annual growth of 100%, it grows 1000-fold in 10 years
  - with annual growth of 50%, it grows 58-fold in 10 years
- crucial is the balance between demand and supply
  - balanced growth makes both users and ISPs happy
  - traffic surged in 2003-2004 by p2p file sharing
  - situation is chaning: might need to worry about oversupply in the future?

residential broadband subscribers in Japan

38.6 million broadband subscribers as of December 2011

- reached 74% of households, increased by 3% in 2011
- FTTH:21.9 million, DSL:7.1 million, CATV:5.9 million

shift from DSL to FTTH: FTTH has exceeded DSL

- 100Mbps bi-directional fiber access costs 60USD/month
  - effects of sales promotion for VoIP and IPTV?
- significant impact to backbones

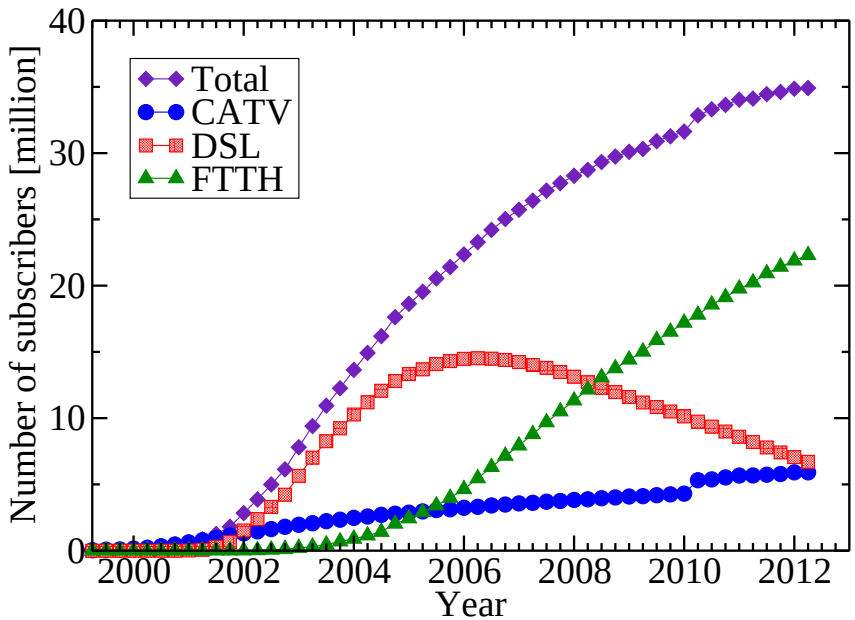

residential broadband subscribers in Japan

traffic growth in backbone

Japanese residential broadband access

- one of the best services (low-cost high-speed) in the world
- the highest in Fiber-To-The-Home (FTTH)

traffic growth of the peak rate at major Japanese IXes

- modest growth since 2005
- drop in 2010, probably due to the campaign for the amended Copyright Act
- no visible impact of the Japan Earthquake at macro level

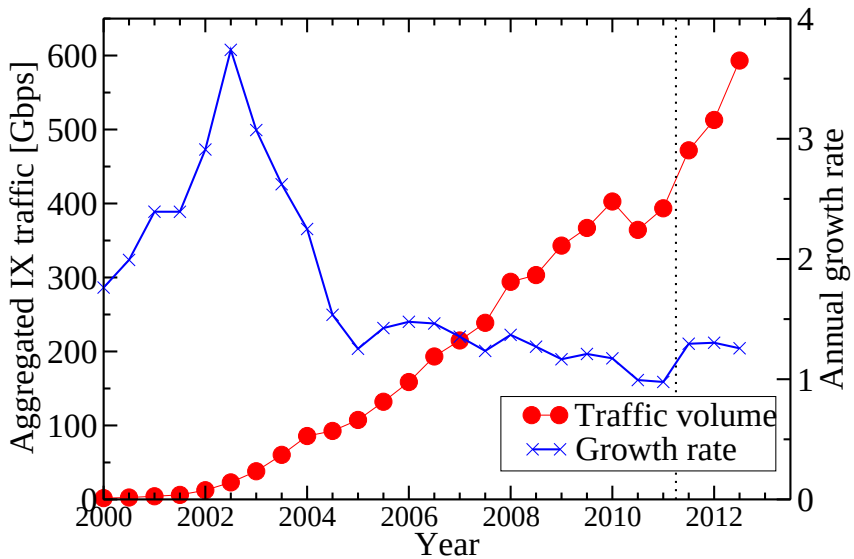

traffic growth of the average rate at major Japanese IXes

motivation: aggregated traffic study by 6 ISPs

concerns about rapid growth of RBB traffic

- backbone technologies will not keep up with RBB traffic
- ISPs cannot invest in backbone simply for low-profit RBB

ISPs and policy makers need to understand the effects of RBB
although most ISPs internally measure their traffic

- data are seldom made available to others
- measurement methods and policies differ from ISP to ISP

lots of IT policy discussions which would affect ISPs

- e.g., net neutrality, content-control, broadband pricing, local IXes
- but mostly based on conjectures or skewed measurements

ISPs' concerns are often not shared by other parties because no data is availabe
ISPs need to speak up for healthy Internet by showing facts

history

2000-2005 e-Japan strategy (by IT strategic headquarter)

- successful in broadband deployment

2004/06 next generation IP intrastructure report by MIC

- an output of one of governmental study groups
- identified issues in backbone
- emphasized importance of long-term measurements for policy making
- suggested cooperative measurement by IPSs, academia, government

2004/07 a study group is formed

- 7 ISPs, 4 researchers from academia, MIC as secretariat
- concensus making
  - first, technical discussions among ops/research people
  - then, talked to top management
- not an official governmental activity but ISPs' voluntary actions
  - ISPs were concerned about government intervention
  - no funding from government

SNMP data collection from 6 ISPs

focus on traffic crossing ISP boundaries (customer and external)

- tools were developed to aggregate MRTG/RRDtool traffic logs

only aggregated results published not to disclose individual ISP share

challenges: mostly political or social, not technical

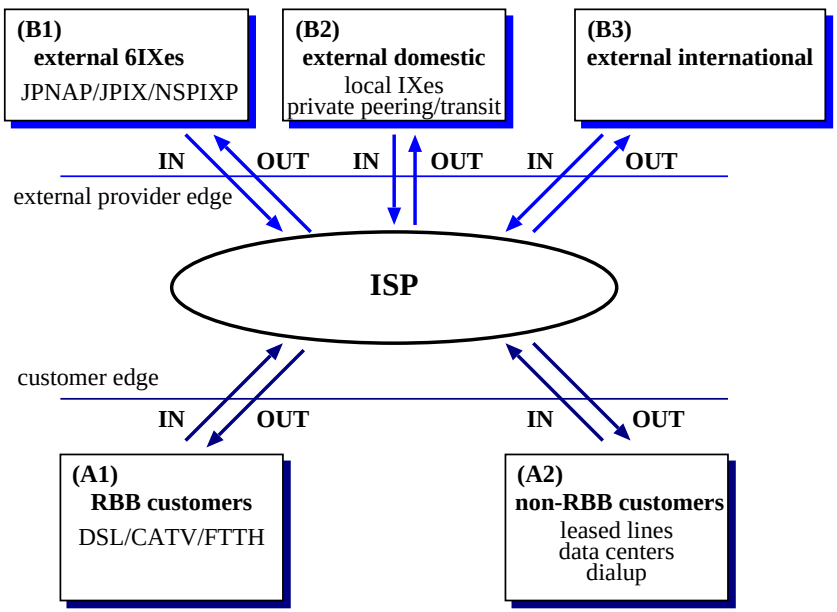

5 traffic groups at ISP cusomer and external boundaries

IN/OUT from ISPs' view

methodology for aggregated traffic analysis

month-long traffic logs for the 5 traffic groups with 2-hour resolution

- each ISP creates log lists and makes aggreagated logs by themselves without disclosing details

biggest workload for ISP

- creating lists by classifying large number of per-interface logs
  - some ISPs have more than 100,000 logs!
- maintaining the lists
  - frequent planned and unplanned configuration changes

data sets

- 2-hour resolution interface counter logs
  - for May and Nov each year
  - by re-aggregating logs provided by 6 ISPs
- our data used to cover 44% of broadband contracts
  - originally we used traffic sahre of the major IXes but the share started to decrease since 2009, propably due to increasing share of content providers
  - so, switched to contract share since 2010/05

traffic growth

- a sharp increase in international inbound due to popular video and other web2.0 services

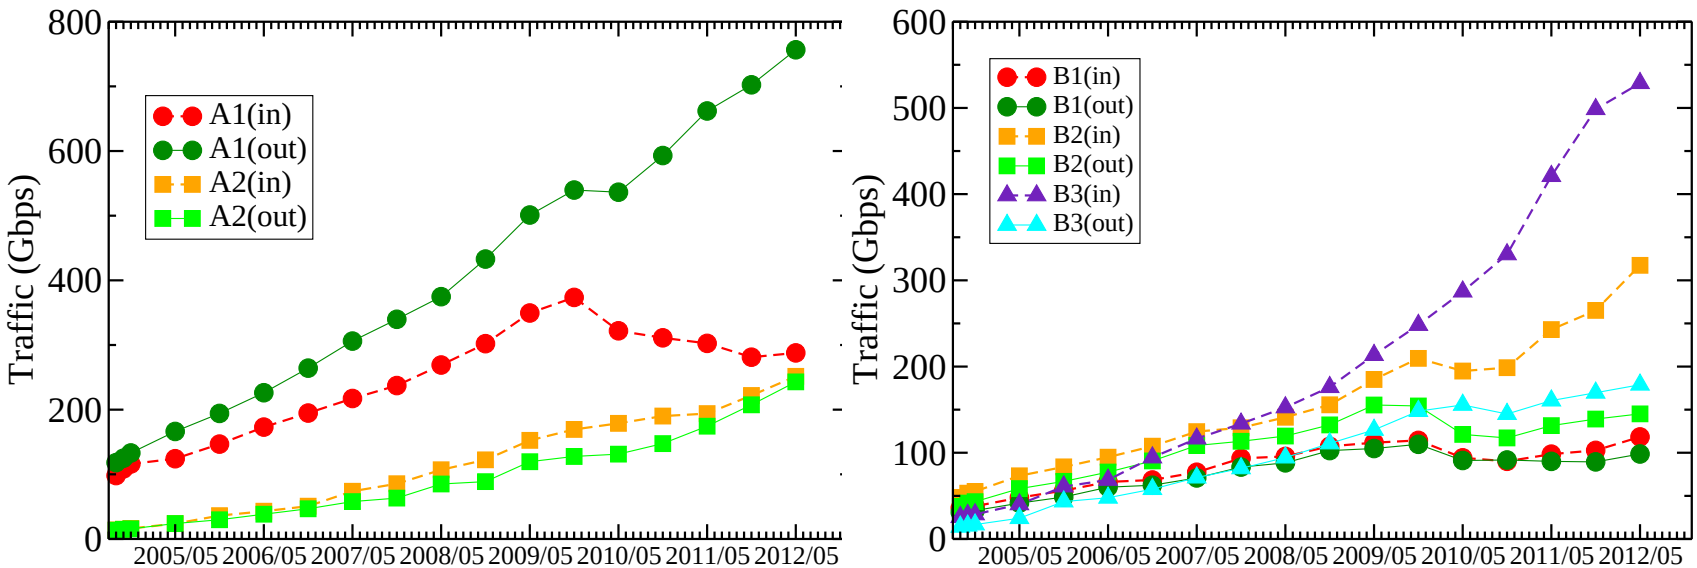

measured traffic growth: customer traffic(left) external ISPs(right)

annual growth rate

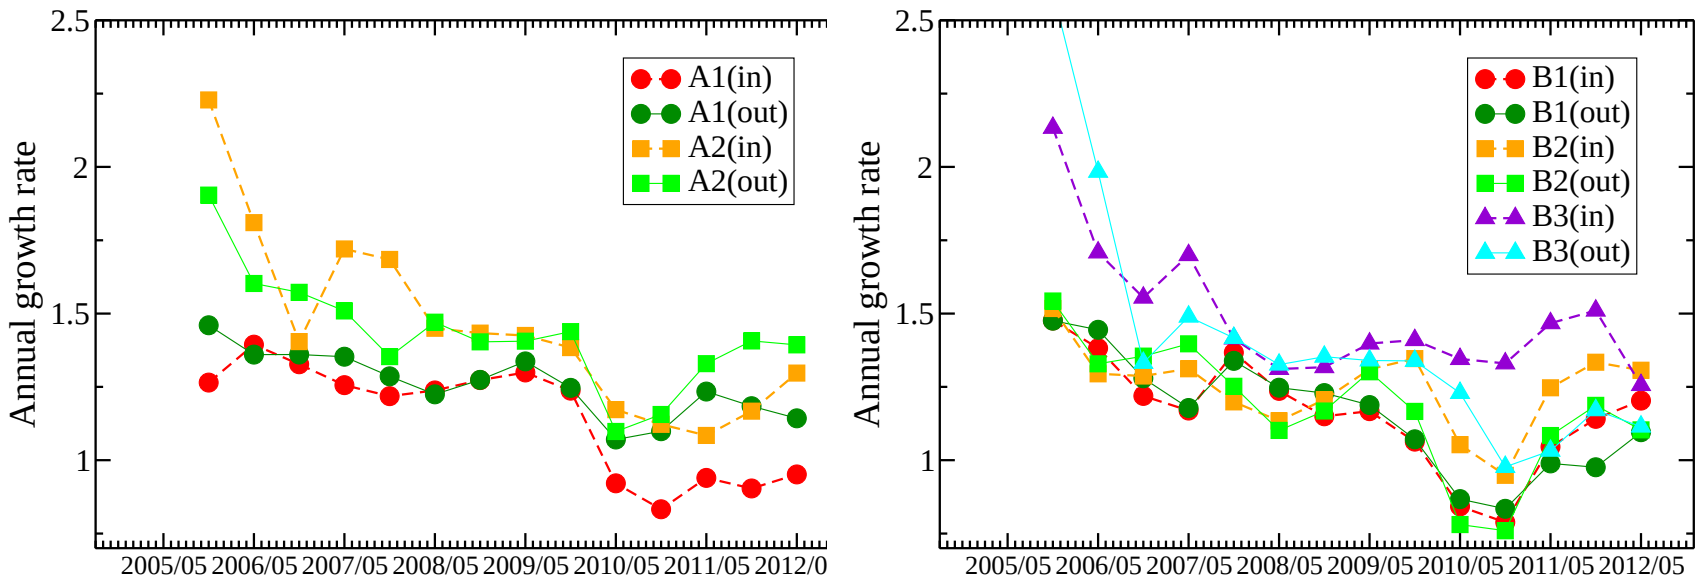

annual growth rates: customer traffic(left) external ISPs(right)

changes in RBB weekly traffic

- traffic patterns by home users (peak at 21:00-23:00)
- 2005: in/out were almost equal (dominated by p2p)
- 2012: outbound (downloading to users) became much larger
  - both constant portion and daily fluctuations grew

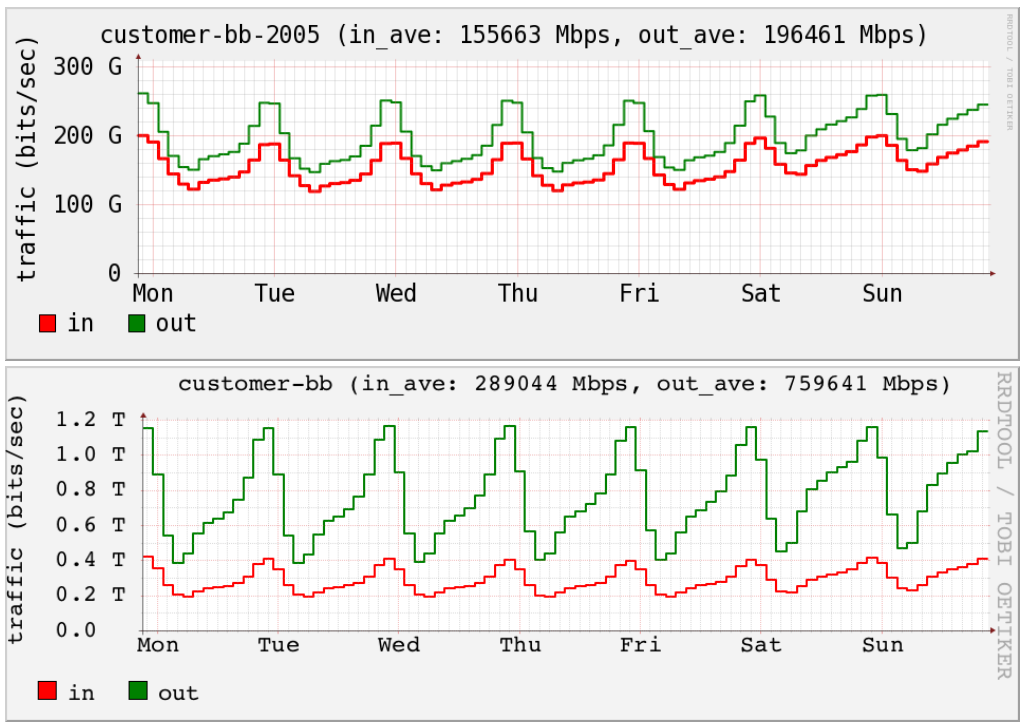

weekly RBB traffic: 2005(top) 2012(bottom)

weekly external traffic in May 2012

external traffic is also strongly affected by RBB traffic

- other-domestic (top right): mainly private peering (also transit, regional IXes)
  - larger than traffic via majior IXes (top left)
- international (bottom): inbound much larger than outbound
  - increasing presence: Google, limelight, Akamai, Amazon, etc

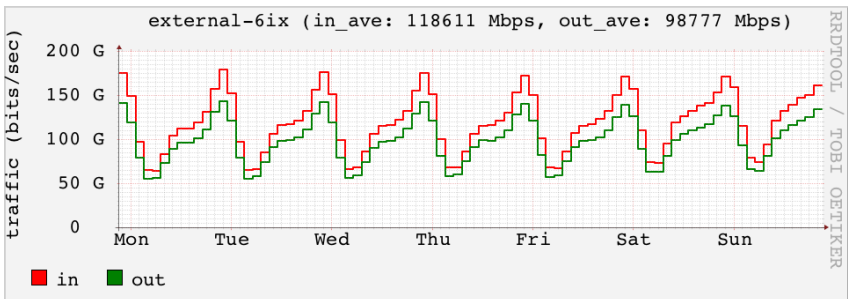

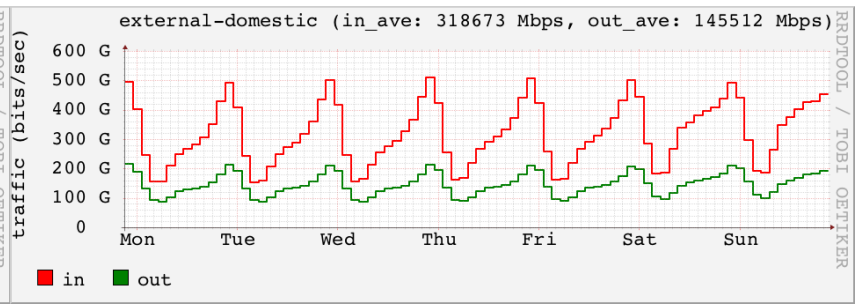

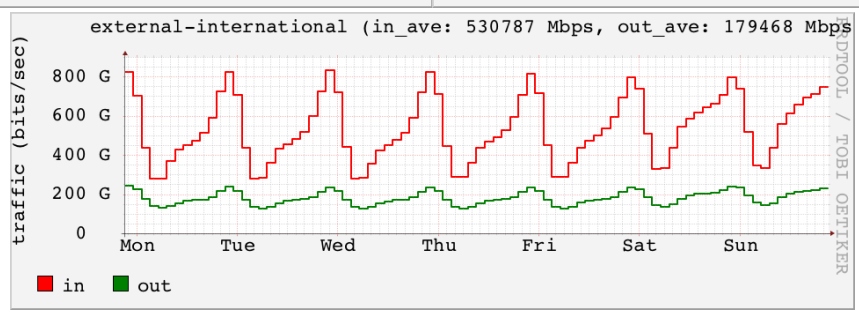

decreasing traffic share of major ISPs

Table: contract share vs. IX traffic share of 6 ISPs

		contract share	IX traffic share
2004	Sep	51.8%	41.5%
	Oct	51.8%	41.9%
	Nov	51.7%	41.6%
2005	May	51.9%	42.0%
	Nov	49.7%	41.4%
2006	May	49.3%	43.1%
	Nov	48.9%	41.5%
2007	May	48.6%	42.4%
	Nov	48.0%	41.8%
2008	May	46.9%	42.6%
	Nov	46.1%	43.8%
2009	May	45.5%	40.6%
	Nov	44.7%	39.6%
2010	May	43.4%	36.9%
	Nov	43.5%	34.7%
2011	May	43.7%	28.6%
	Nov	43.9%	26.4%
2012	May	43.7%	25.8%

aggregated traffic summary

in 2012, we observed

- larger download volume, larger evening-hour volume in RBB
- RBB traffic decreased share in customer traffic
- larger growth of international inbound
- decreasing traffic share of major ISPs

implies a shift from p2p to video and other web2.0 services

analysis of per-customer traffic by IIJ

IIJ's per-customer traffic data (RBB traffic only)

- Sampled NetFlow data
  - from edge routers accommodating fiber/DSL RBB customers
- week-long data since 2004
  - latest news: comparison of 2011 and 2012
  - also 2005: before YouTube

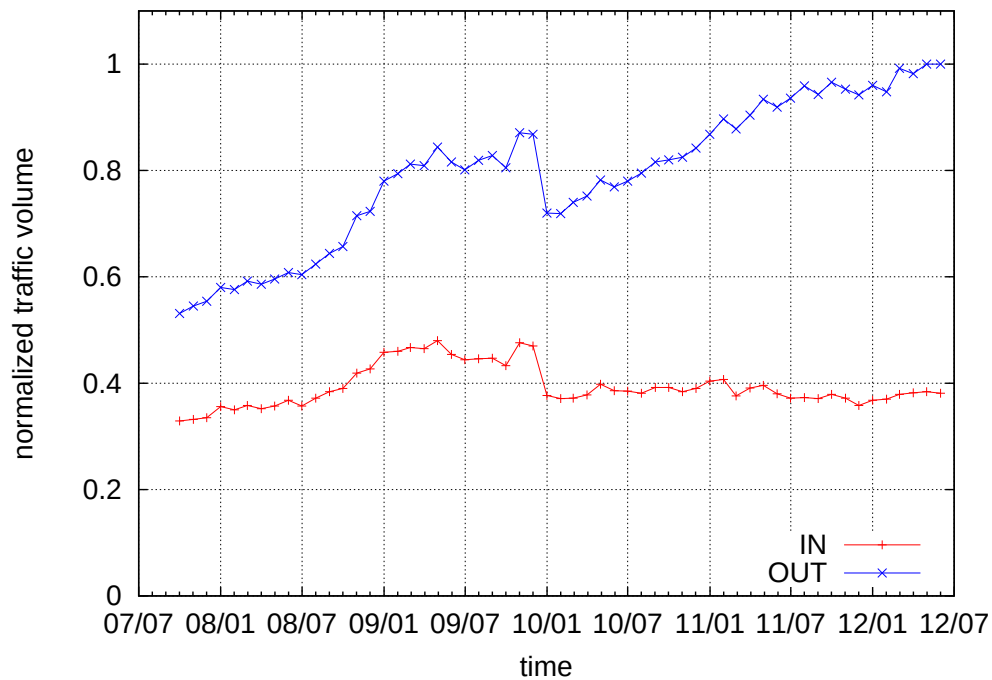

residential broadband per-month traffic

ratio of fiber/DSL active users and total traffic volumes

- in 2012, 91% of active users are fiber users, consuming 95% of traffic
  - active user: unique customer IDs observed in the data set

		active users (%)	total volume (%)
2005	fiber	46	79
	DSL	54	21
2008	fiber	79	87
	DSL	21	13
2011	fiber	88	93
	DSL	12	7
2012	fiber	91	95
	DSL	9	5

PDF of daily traffic per user (2005)

each distribution consists of 2 roughly lognormal distributions

- client-type: asymmetric (majority)
- peer-type: symmetric high-volume

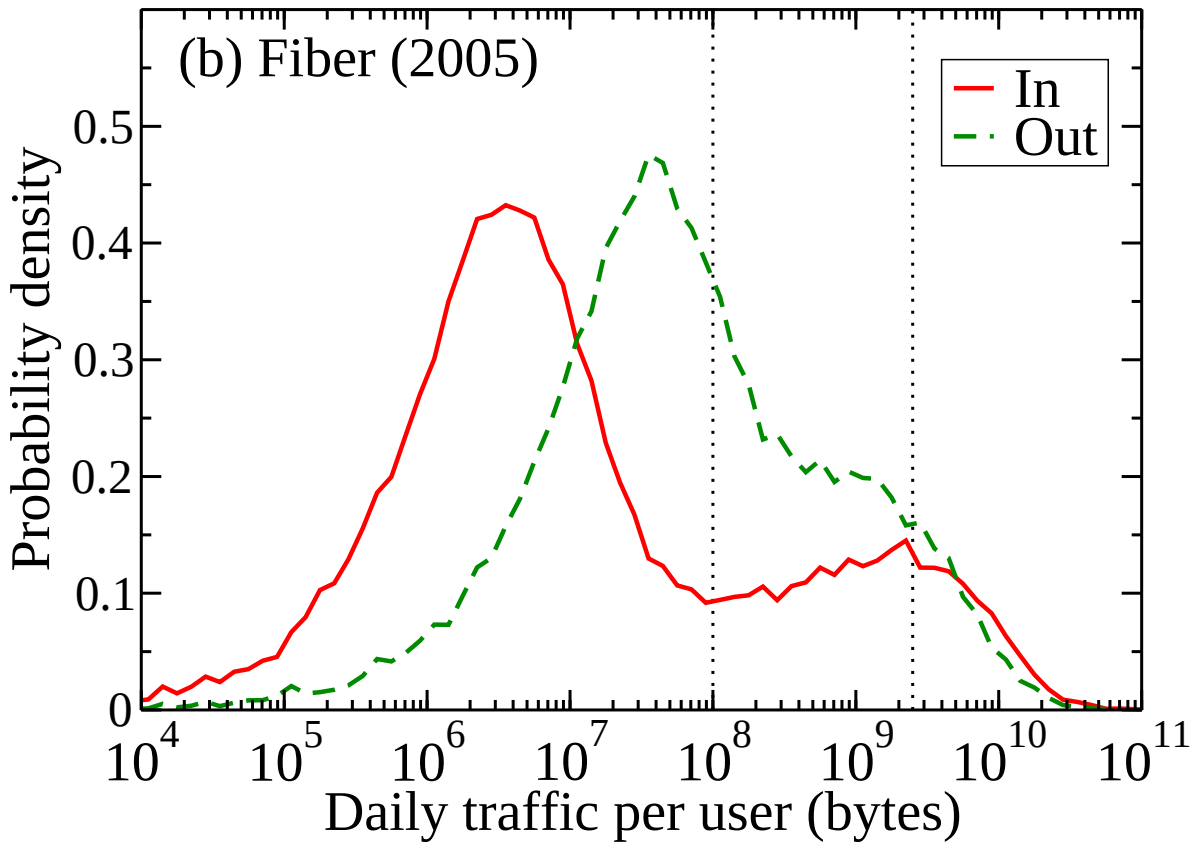

comparing total: 2005, 2011 and 2012

	IN (MB/day)		OUT (MB/day)	
	mean	mode	mean	mode
2005	430	3.5	447	32
2007	433	4	712	66
2008	483	5	797	94
2009	556	6	971	114
2010	469	7	910	145
2011	432	8.5	1001	223
2012	410	14	1026	282

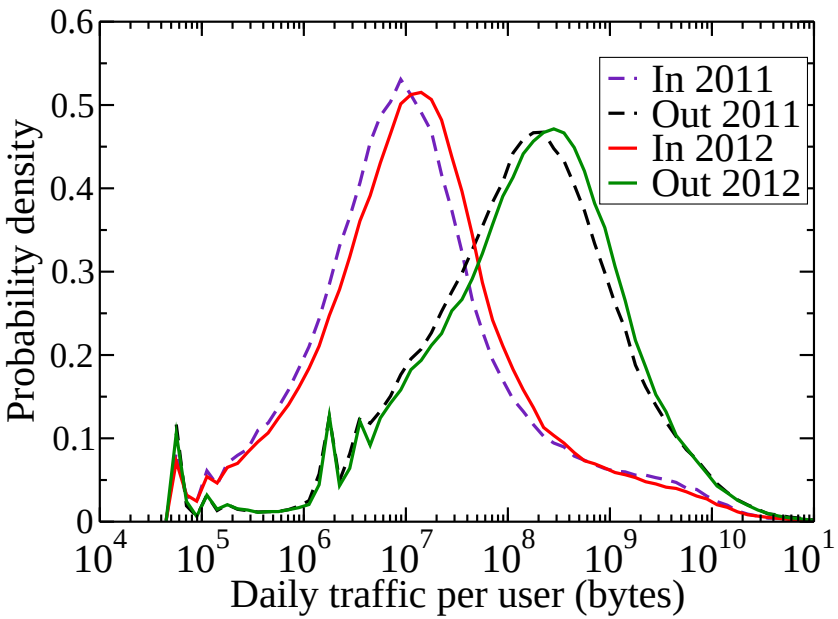

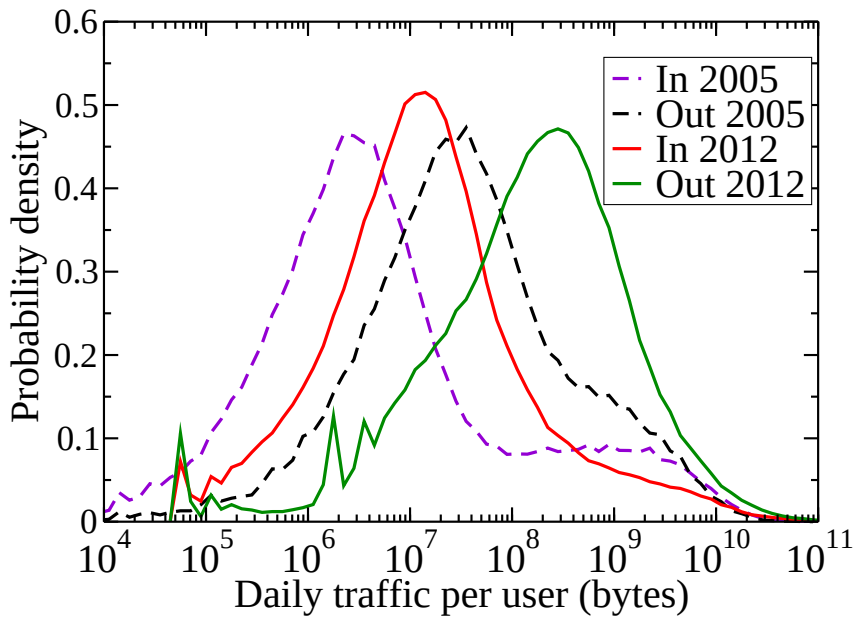

changes in daily traffic per user (2011 vs. 2012) (2005 vs. 2012)

CCDF of daily traffic per user

- heavy-tailed distribution
  - the tail exceeds 800GB/day
- the tail becomes longer and symmetric (no longer need to compensate upstream shortage of DSL)

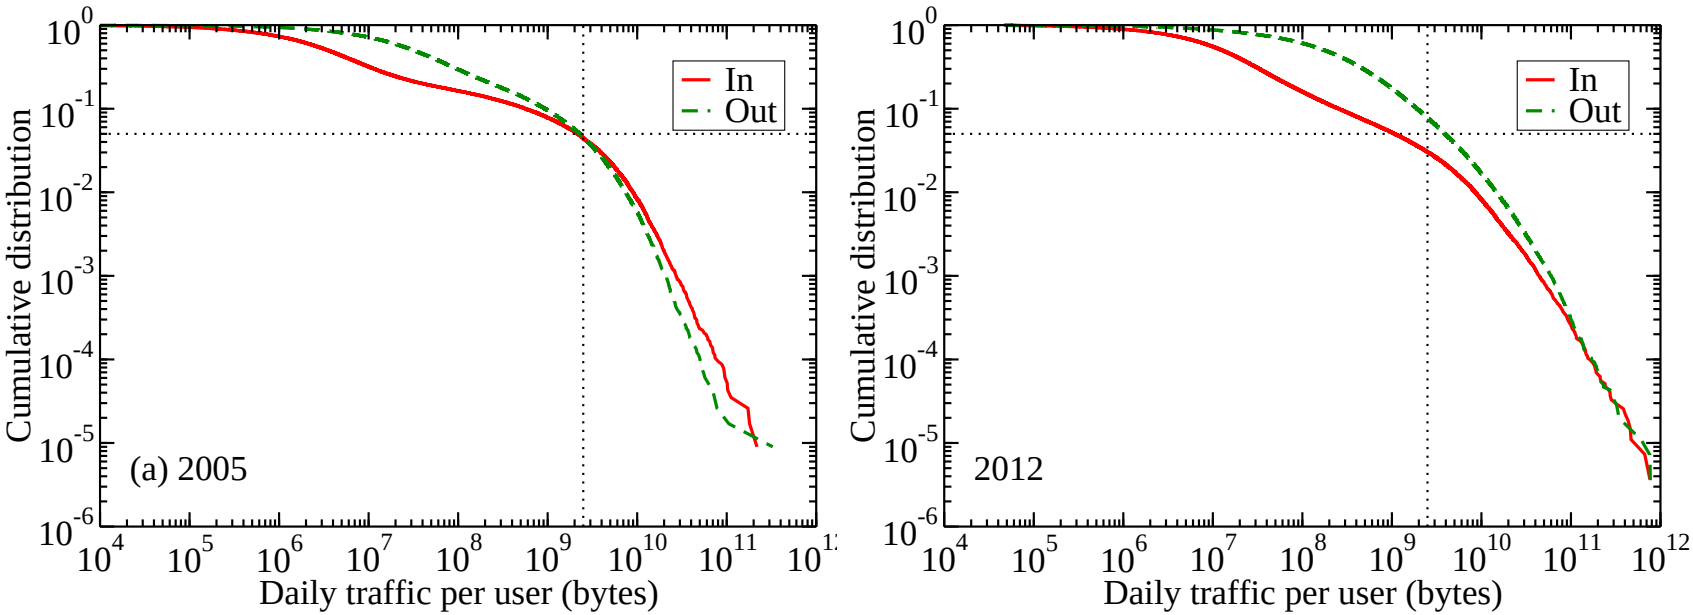

CCDF of daily traffic per user in 2005 and 2012

skewed traffic usage among users

- highly skewed distribution in traffic usage
  - top 10% users consume 73% of download, 95% of upload volumes in 2012
- no noticeable change since 2005
  - long-tailed distribution (common to other Internet data)
  - looks similar even if p2p traffic is removed

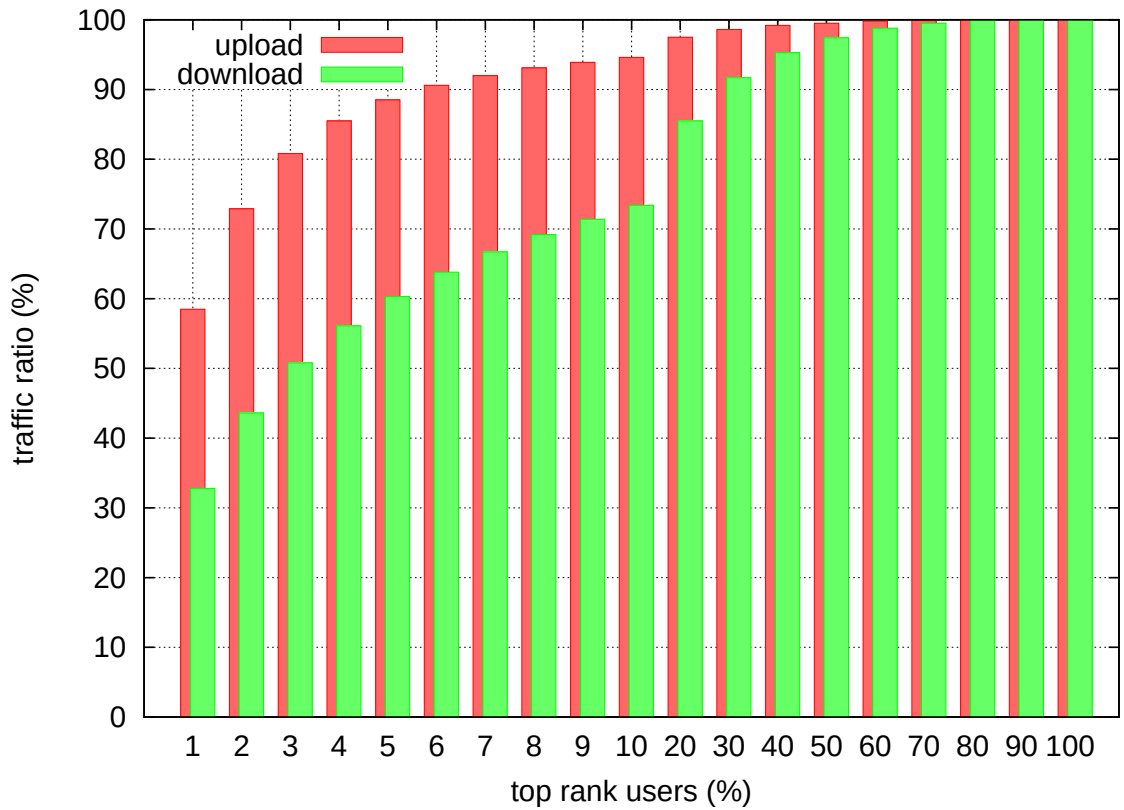

traffic share of top-ranking heavy-hitters

correlation of inbound/outbound volumes per user

2 clusters: client-type users and peer-type heavy-hitters

- the peer-type cluster is disappearing
- difference between fiber and DSL: only heavy-hitter population
- no clear boundary: heavy-hitters/others, client-type/peer-type
- actual individual users have different traffic mix

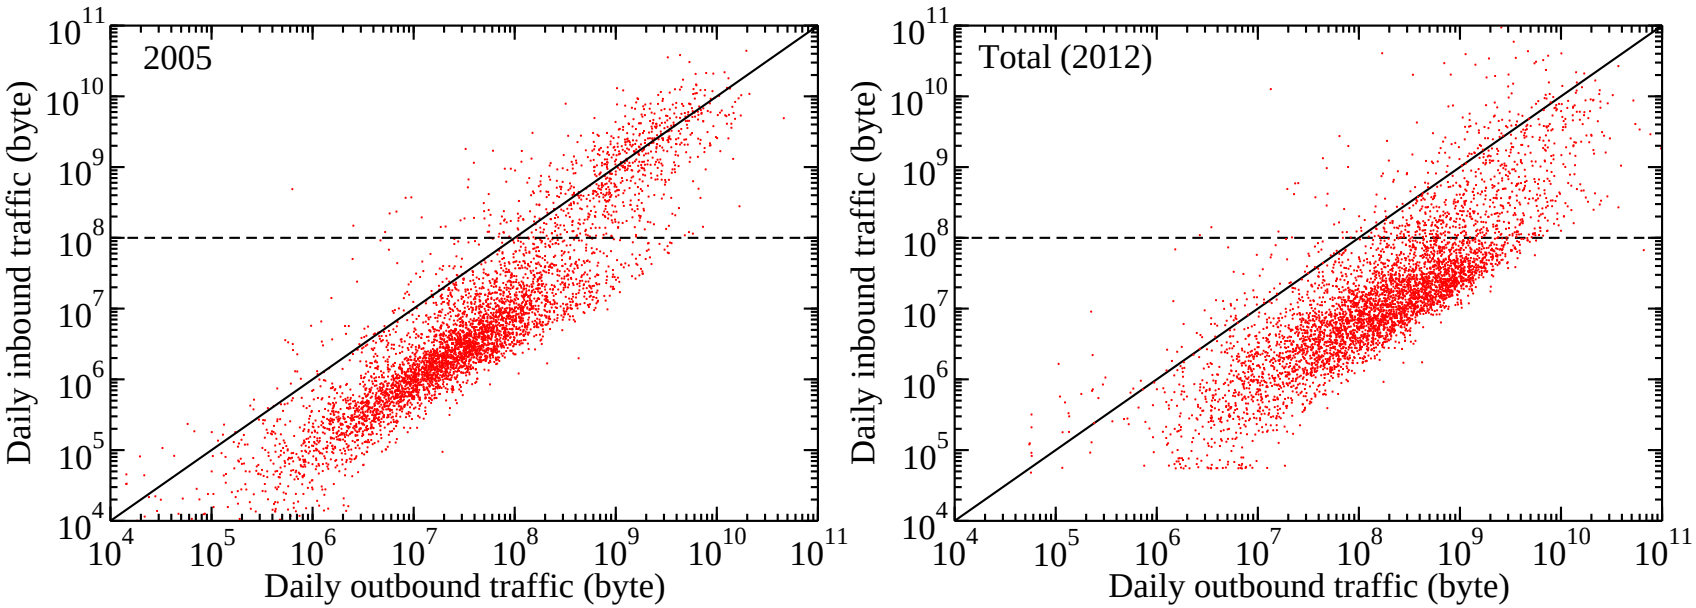

in/out volumes per user: 2005 vs. 2012

protocols/ports ranking

classify client-type/peer-type with threshold: 100MB/day upload

- to observe differences in protocol/port usage
- port number: min(sport, dport)

observations

- dominated by TCP dynamic ports (but each port is tiny)
- TCP port 80 is increasing (again)

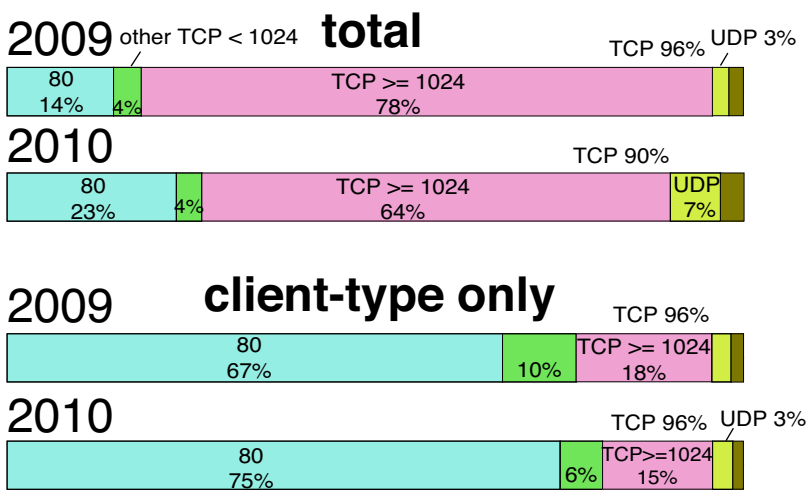

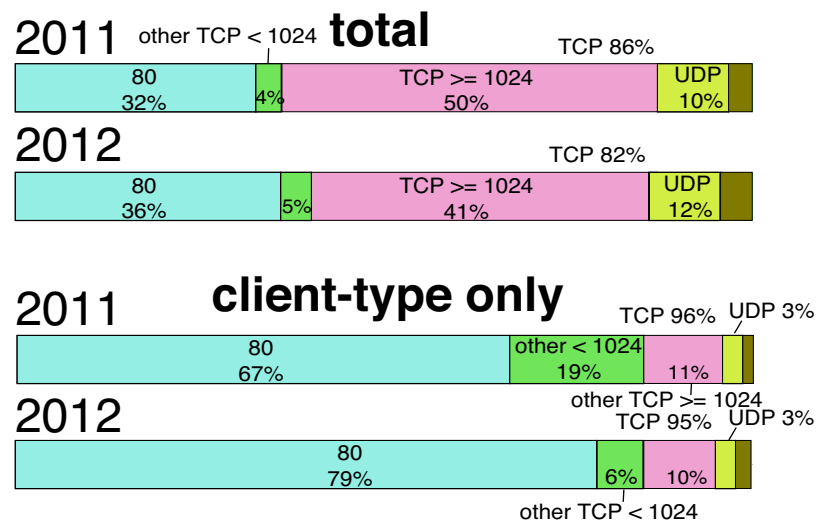

protocols/ports ranking data

protocol	port	2011 total (%)	2011 client type	2012 **total** (%)	2012 client type
TCP	*	**85.95**	**96.28**	**81.86**	**95.09**
	(< 1024)	*36.24*	*85.69*	*41.23*	*85.25*
	80 (http)	32.10	67.30	36.22	79.39
	443 (https)	1.33	1.91	2.45	3.43
	554 (rtsp)	1.33	6.89	0.77	1.01
	22 (ssh)	0.27	0.17	0.22	0.06
	(>= 1024)	*49.71*	*10.59*	*40.63*	*9.84*
	1935 (rtmp)	1.58	1.51	2.12	3.91
	7144 (peercast)	0.38	0.00	0.44	0.04
	6346 (gnutella)	0.68	0.60	0.37	0.09
	8080	0.26	0.14	0.30	0.17
UDP		**10.01**	**2.61**	**12.38**	**2.94**
ESP		**3.56**	**1.02**	**5.29**	**1.79**
GRE		**0.15**	**0.05**	**0.16**	**0.14**
L2TP		**0.13**	**0.00**	**0.14**	**0.00**
IP-ENCAP		**0.10**	**0.01**	**0.09**	**0.01**

temporal behavior of TCP port usage

3 types: port 80, well-kown port but 80, dynamic ports

- total traffic heavily affected by peer-type traffic
- shift from dynamic ports to port 80 for client-type users
- daily fluctuations also observed in dynamic ports
  - slow decay of dynamic port traffic over night

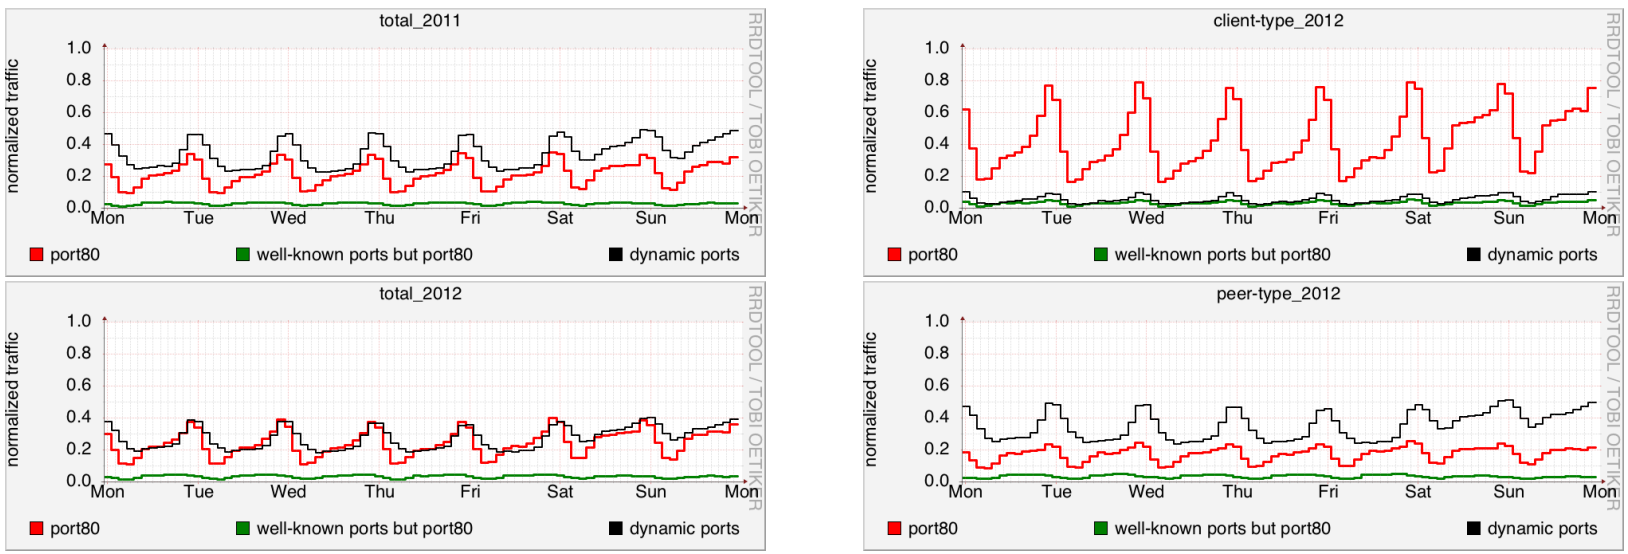

TCP usage: total 2011(top-left) 2012(bottom-left) client-type(top-right) peer-type(bottom-right)

Top 10 ASes in TCP port 80 traffic volume

content providers, CDNs, and hosting sites

Total

2011			2012		
ASN	Org	share(%)	ASN	Org	share(%)
15169	Google	18.7	15169	Google	20.9
2497	IIJ	7.5	2497	IIJ	10.6
20473	Choopa	7.1	22822	LimeLight	7.6
22822	Limelight	6.6	38634	DWANGO	4.3
38634	DWANGO	4.7	40263	FC2	3.7
40263	FC2	3.4	2914	NTT-Com	3.3
40824	WZ Com	3.2	40824	WZ Com	2.8
35415	WebaZilla	3.1	16276	OVH Systems	2.8
24572	Yahoo! JP	2.4	35415	WebaZilla	2.7
30212	Hypermedia	2.3	16265	LeaseWeb	2.7

Client-type

2011			2012		
ASN	Org	share(%)	ASN	Org	share(%)
15169	Google	24.2	15169	Google	26.1
22822	Limelight	8.0	22822	Limelight	9.5
2497	IIJ	5.0	2497	IIJ	9.2
38634	DWANGO	4.9	38634	DWANGO	4.4
40263	FC2	4.4	40263	FC2	4.3
20473	Choopa	4.1	2914	NTT-Com	3.5
24572	Yahoo! JP	3.1	23816	Yahoo! JP	2.7
23816	Yahoo! JP	2.7	24572	Yahoo! JP	2.5
30212	Hypermedia	2.4	30212	Hypermedia	2.0
17506	UCOM	2.0	35415	WebaZilla	1.7

summary of per-customer traffic analysis

- traffic has been stable with modest growth rate (20%/year)
- p2p traffic decreased in population share and volume share
  - still, traffic distribution among users is highly skewed
- visible players: content providers, CDNs, hosting sites

new amended Copyright Act

- on June 20 2012, new amended copyright act has been passed
  - criminization of illeagal downloads
  - it becomes effective on October 1 2012
- 2 views on a traffic drop by the amended copyright act 2010
  - law enforcement is needed to decreace illegal content on the Net
  - it was just a trigger to acceralate the already existing shift from p2p file-sharing to web services
- can be evaluated by the impact of the new amended copyright act

conclusion

- apparent slow traffic growth in Japan
  - due to decline of p2p traffic
  - high penetration of fiber broadband
- steady increase in normal users' volume
- it is difficult to predict future traffic
  - infra is ready for video streaming services
  - users shifting to mobile devices (wifi off-loading?)
  - possible impact by the new amended copyrught act?

acknowledgments

- IIJ, SoftBank Telecom, K-Opticom, KDDI, NTT Communications, SoftBank BB for data collection support
- ministry of internal affairs and communications for coordination