Internet Initiative Japan

Broadband measurement in Japan

Broadband Performance and QoE: Measurement Technologies Improving the Marketplace

Kenjiro Cho, IIJ

2018-06-23

Kenjiro Cho

- Research Director at IIJ
- a board member, mawi-wg chair of WIDE Project
- involved in measurement research for 20 years
  - data collection and analysis at ISP/academic networks, especially on traffic
  - large-scale content distribution: traffic management
  - data-center networking and cloud services

broadband traffic in Japan

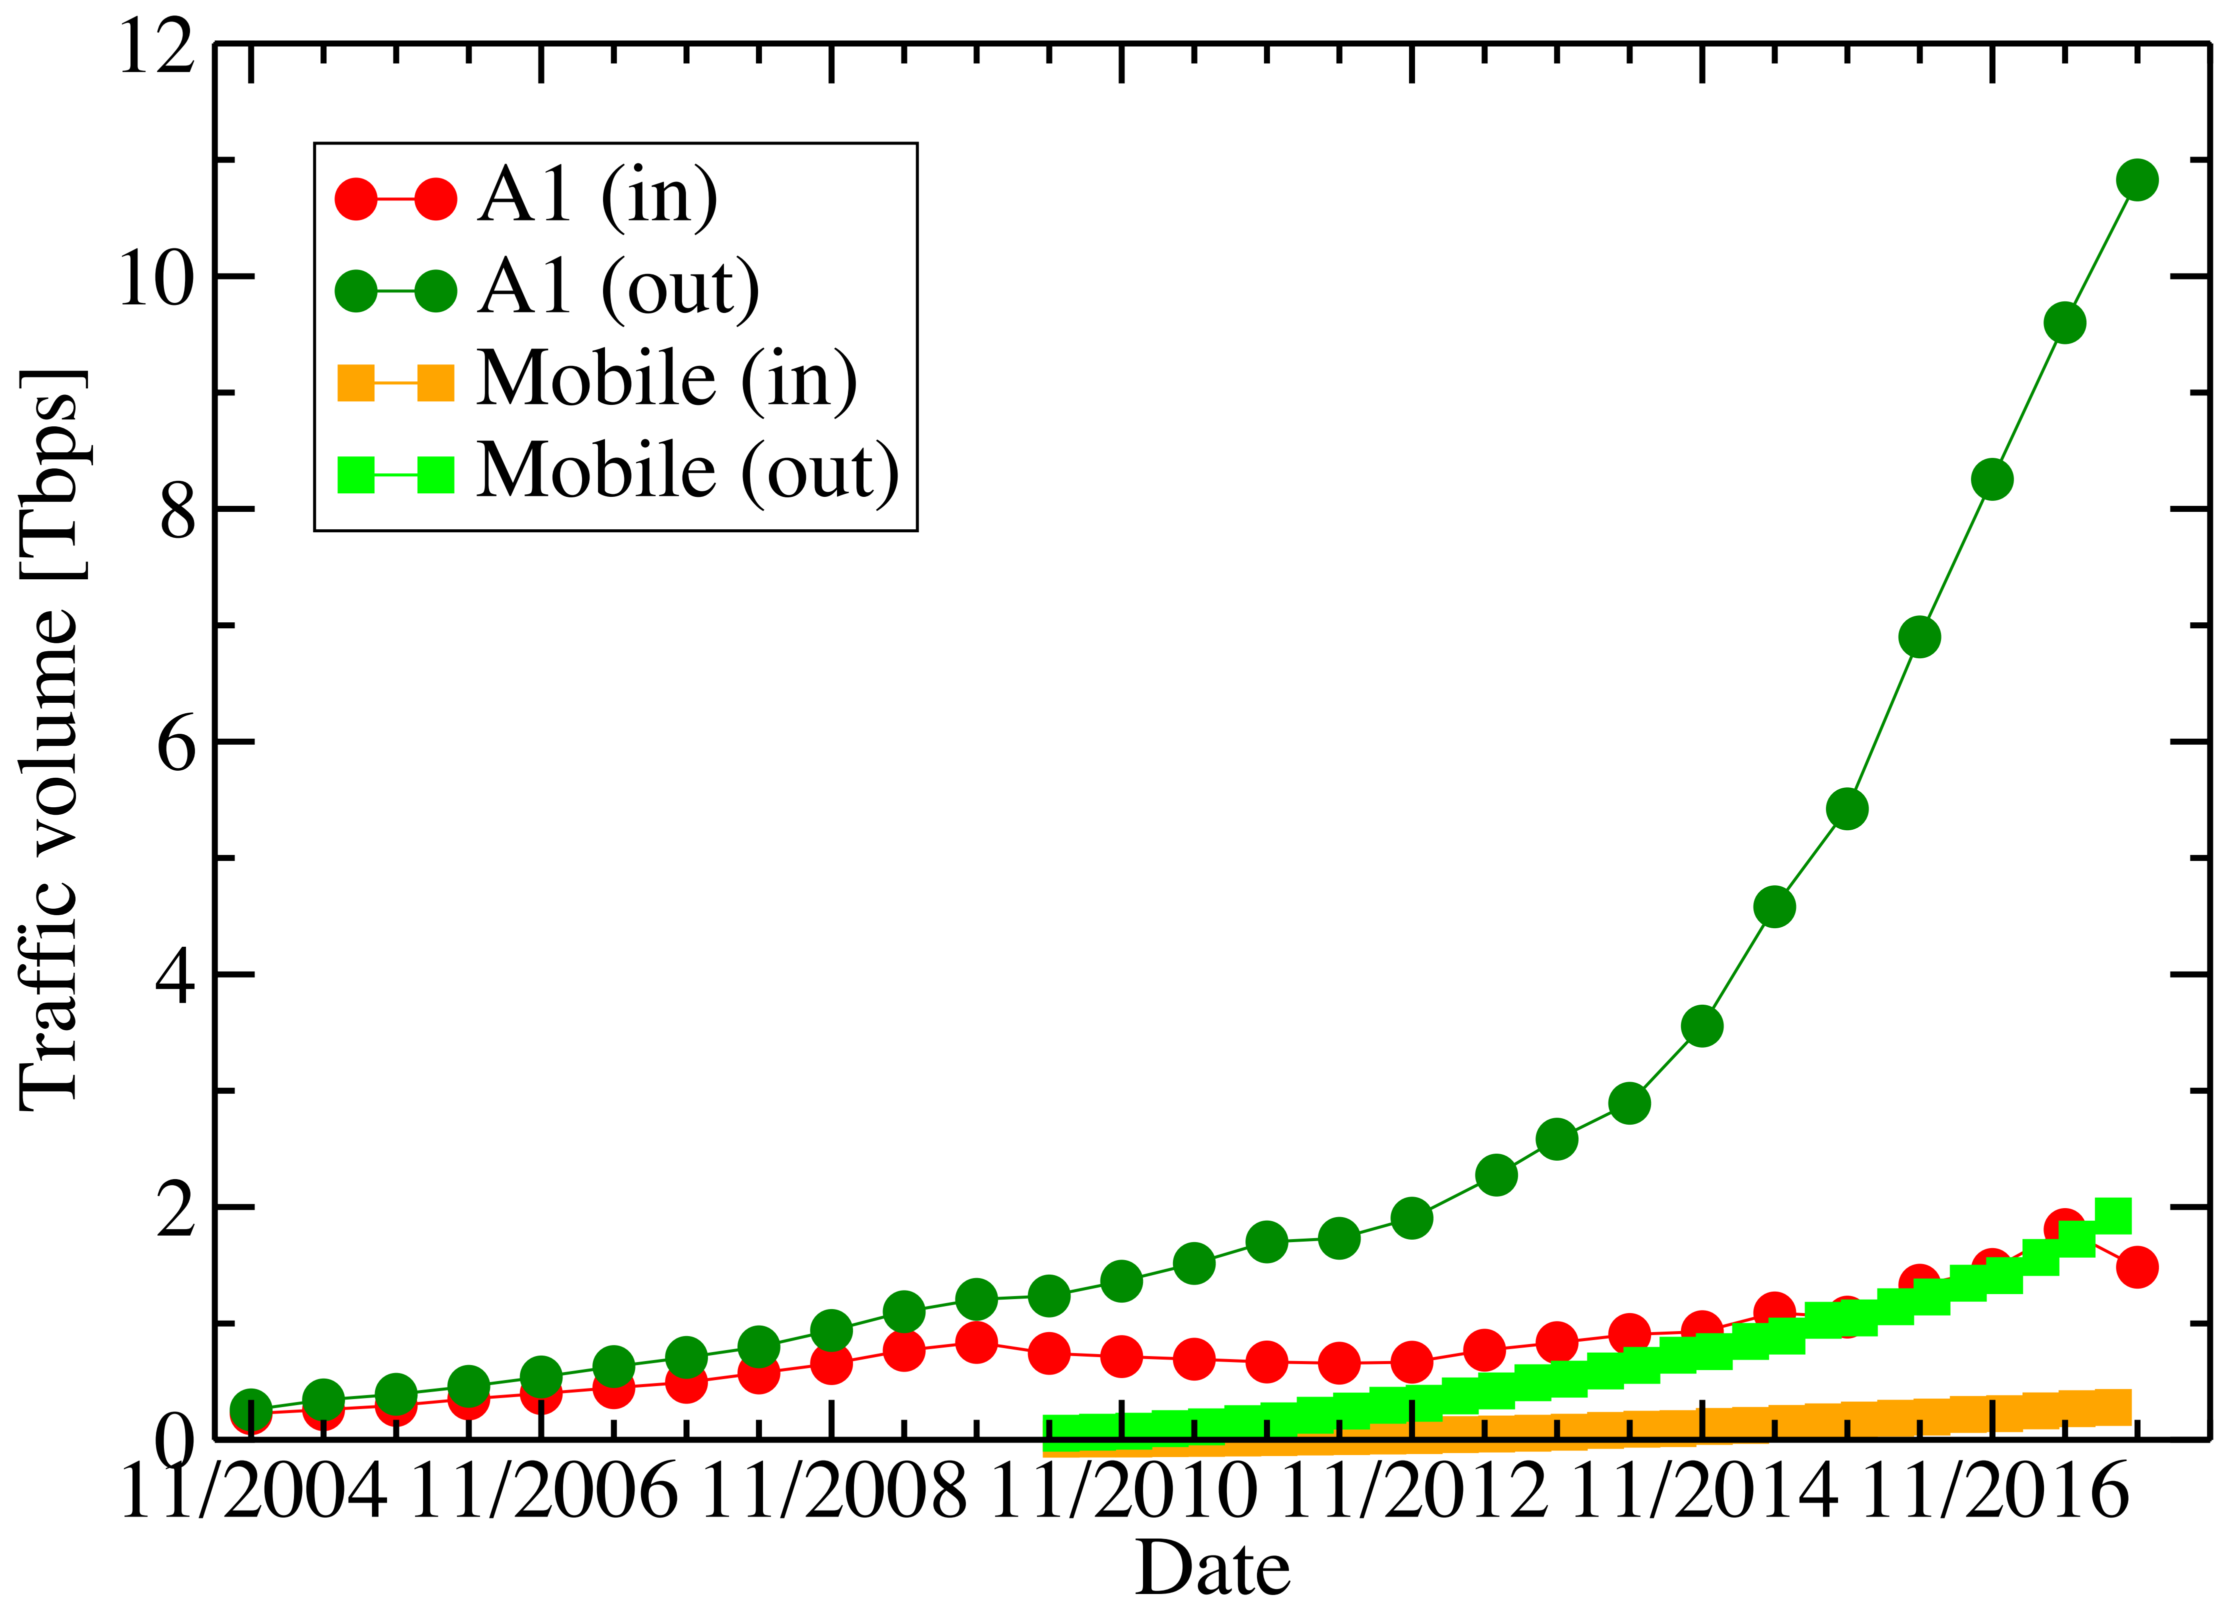

estimated total broadband traffic in Japan

http://www.hongo.wide.ad.jp/InternetTraffic/data/

overall traffic trends

- steady increase driven by different major factors
  - web (late 90's), broadband deployment/P2P-file sharing (early 2000's), short-videos (late 2000's), video subscriptions (mid 2010's)
- recent trends
  - video streaming by subscriptions
  - software updates: become larger and more frequent
  - cloud uploads: SNS, multiple-devices, auto-backups
  - faster access: NGN, 1G/10G access, LTE+
  - closer content: CDN, cache, peering
  - advertisement and analytics

possible bottlenecks in broadband

- many potential bottlenecks: hard to distinguish!

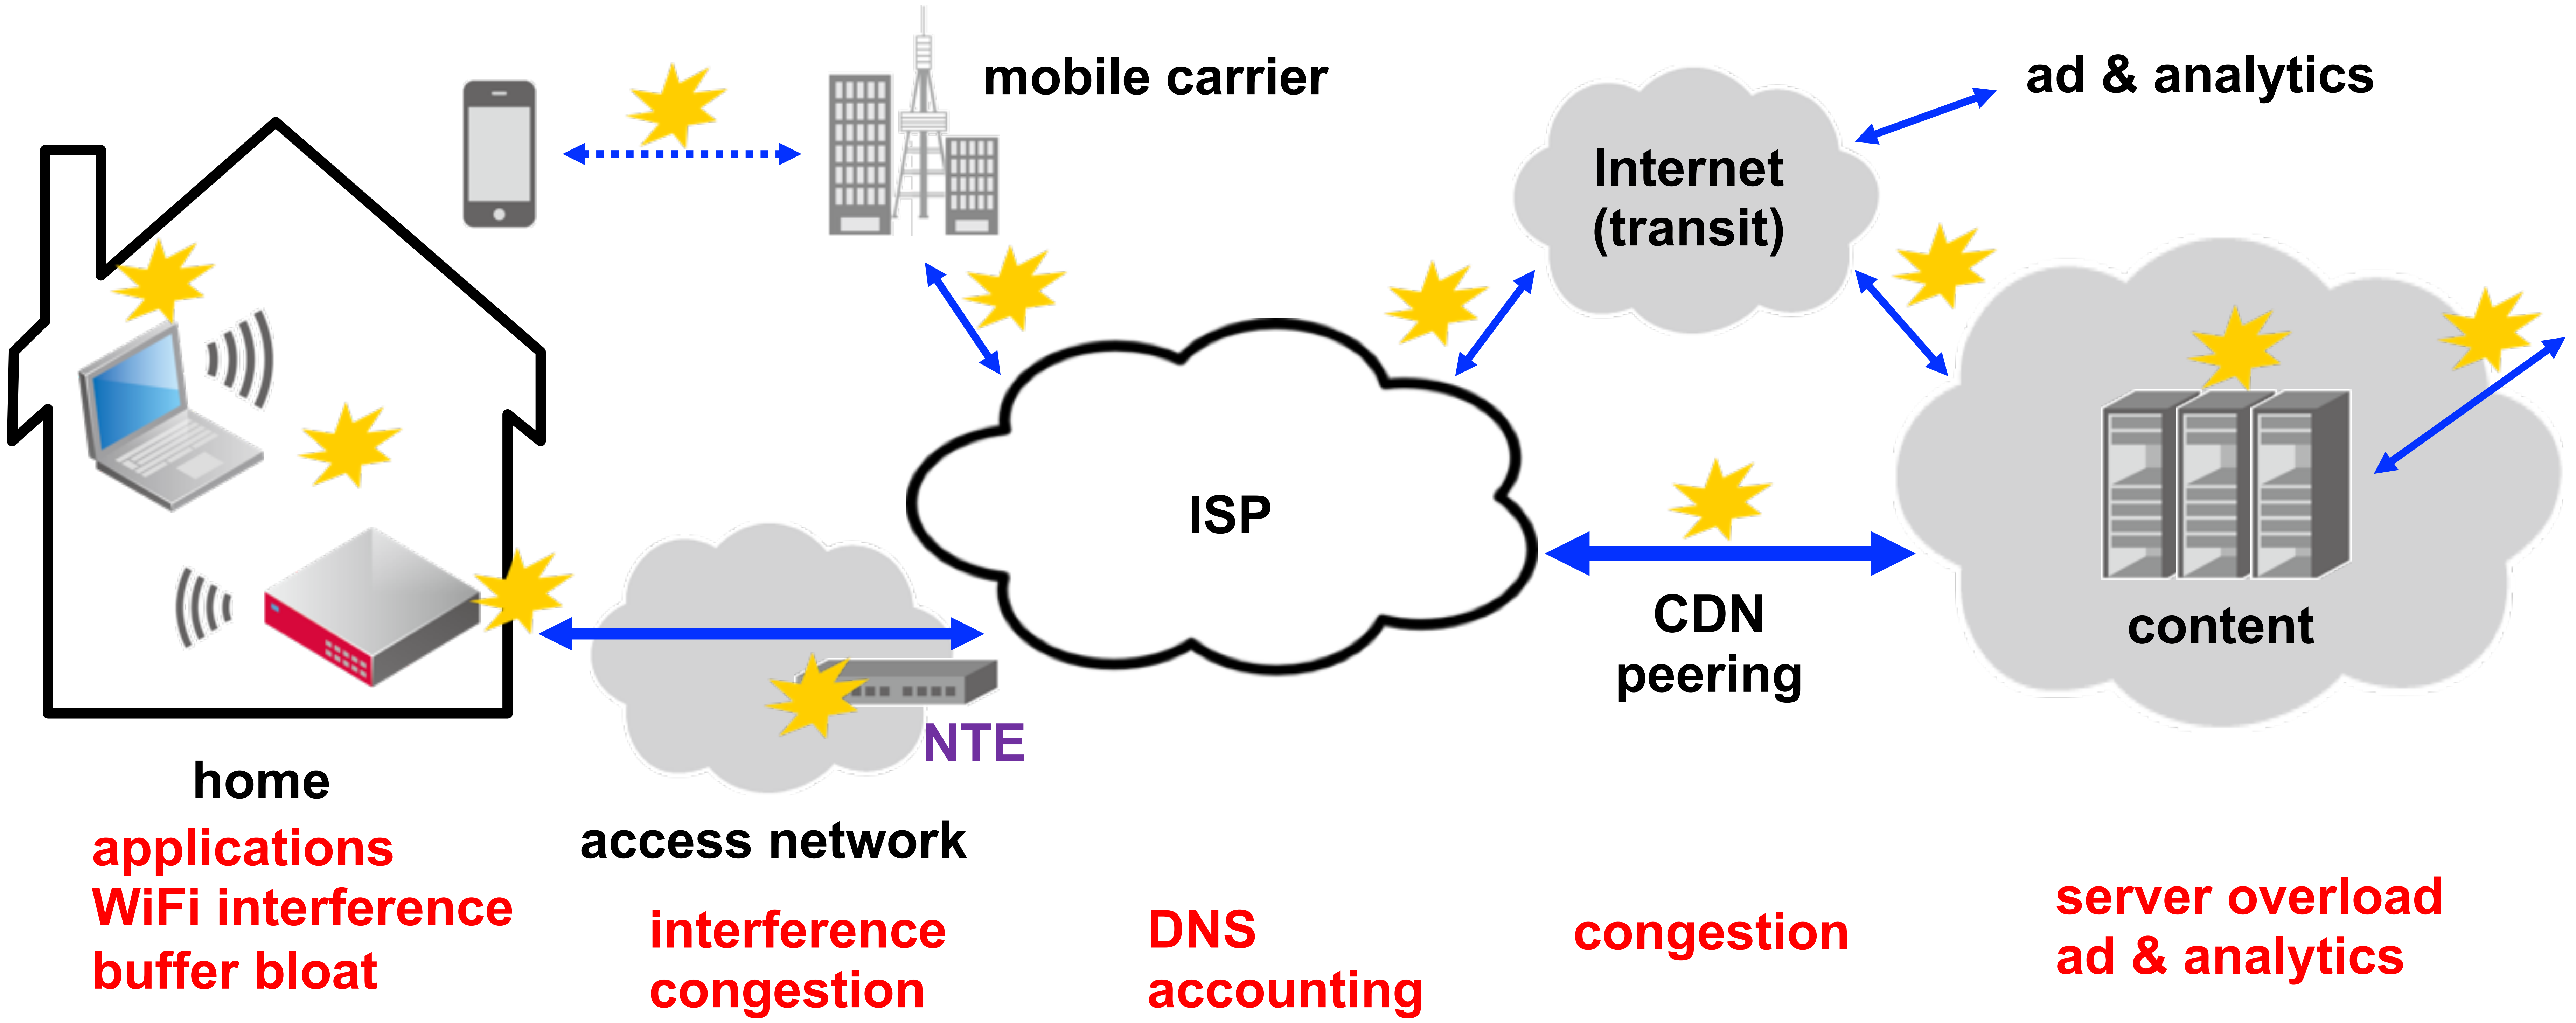

various measurements

- ISP-side
- client-side
- content-provider side

ISP-side measurements

- interface counters
- flow exports
- DNS resolvers

traffic between ISPs

- transit services: metered billing (e.g., by 95th-percentile over 5-min volumes)
- peering: bilateral settlement-free customer-traffic exchange

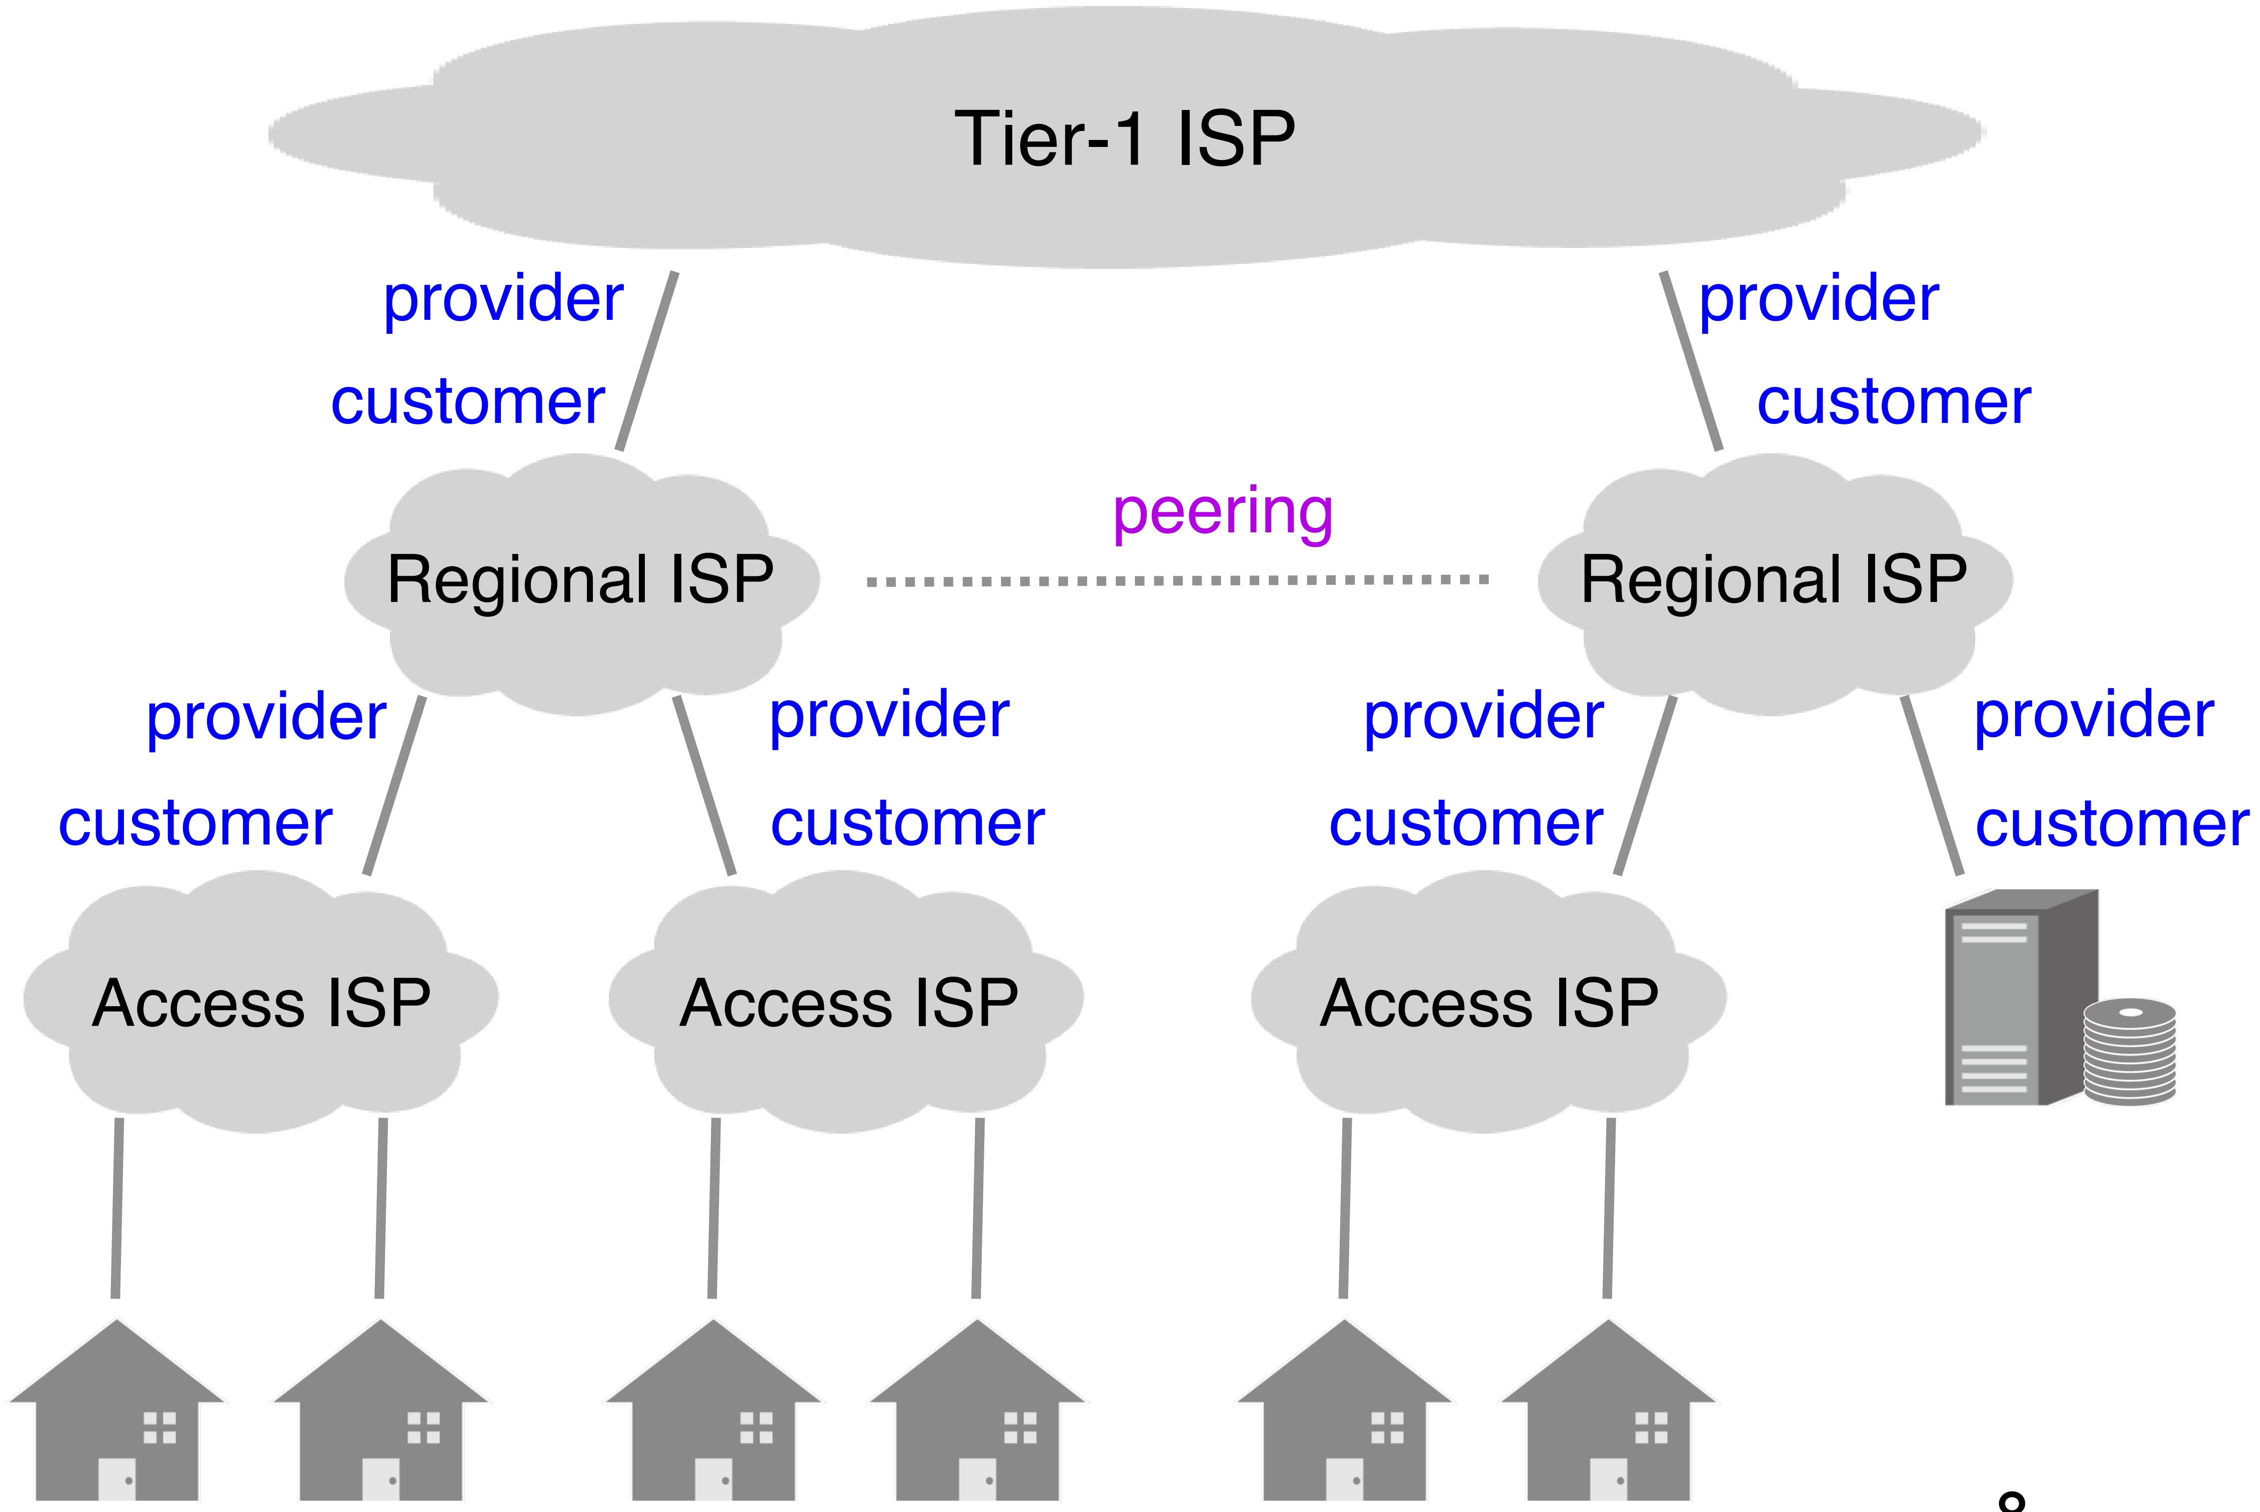

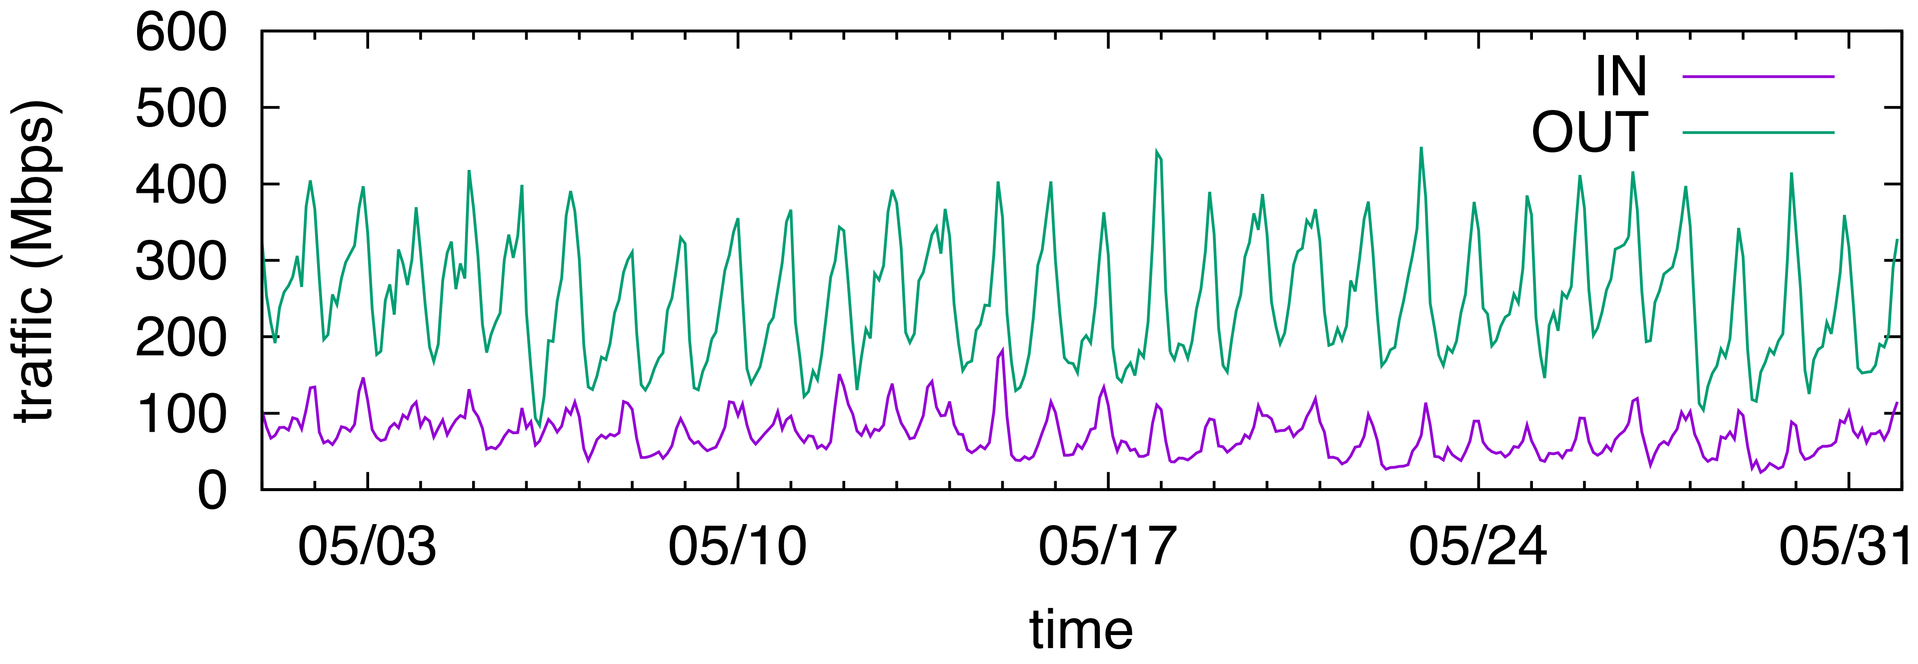

diurnal broadband traffic

2011.11

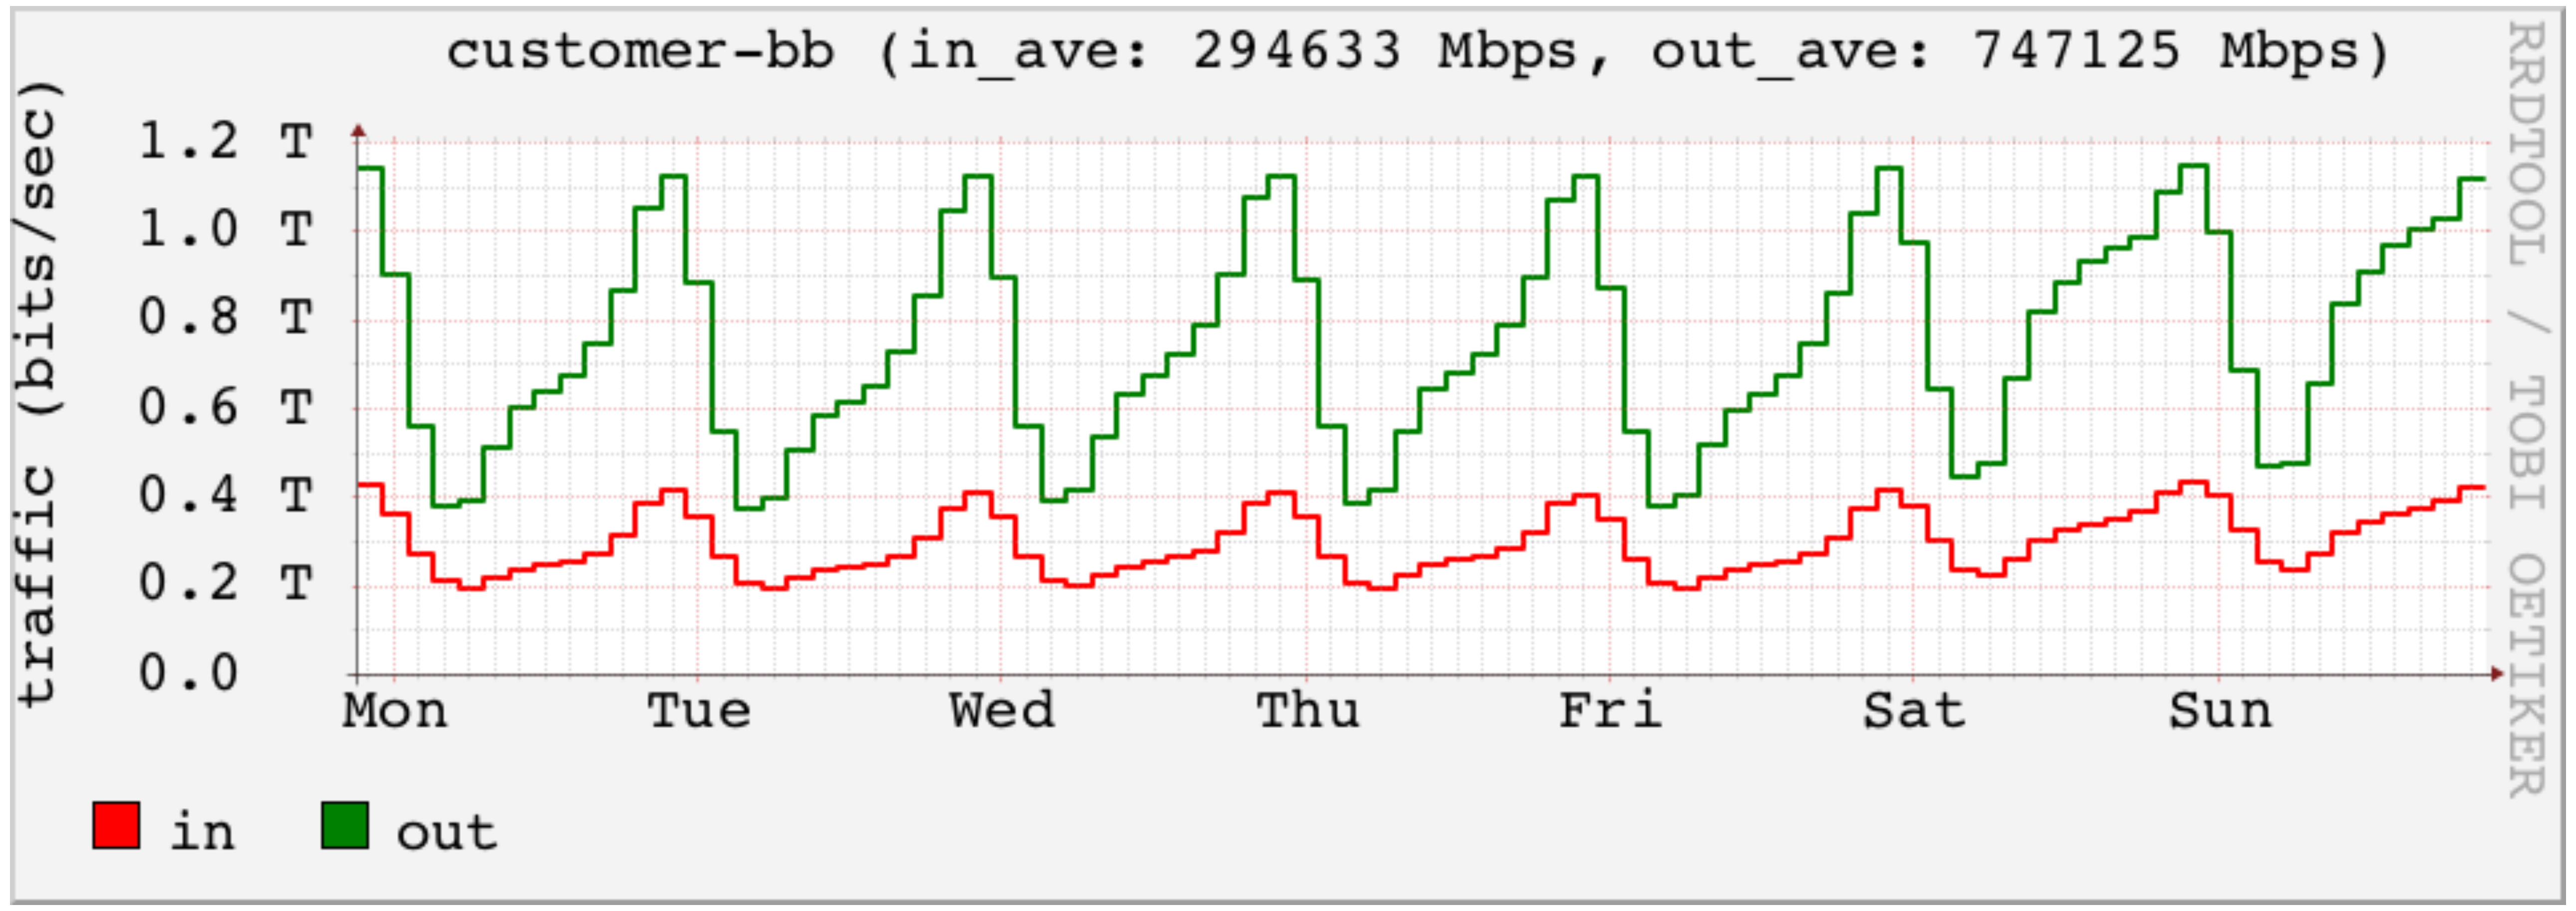

2017.11

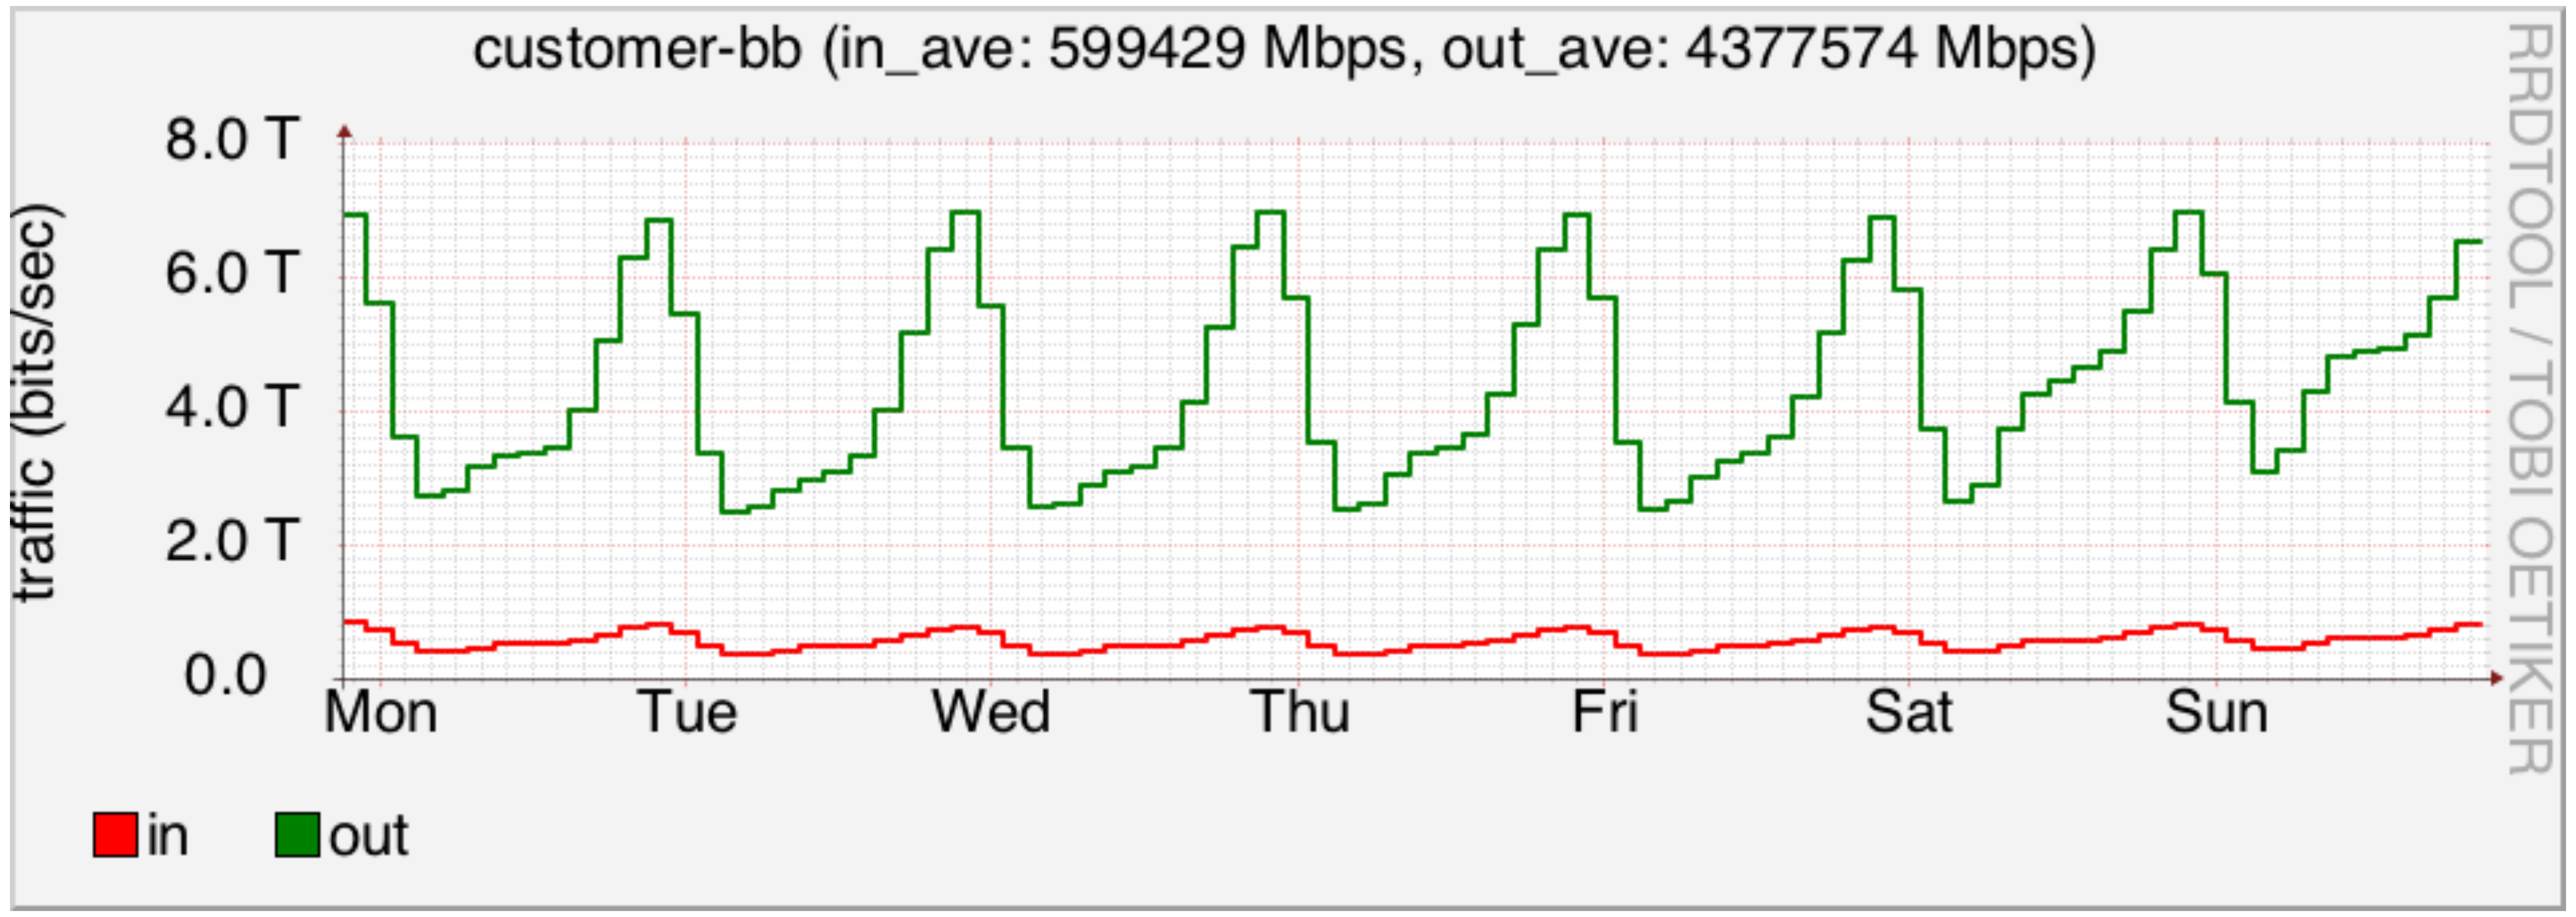

traffic usage per customer

- daily traffic usage distribution
  - mean of a week as weekdays and weekends differ

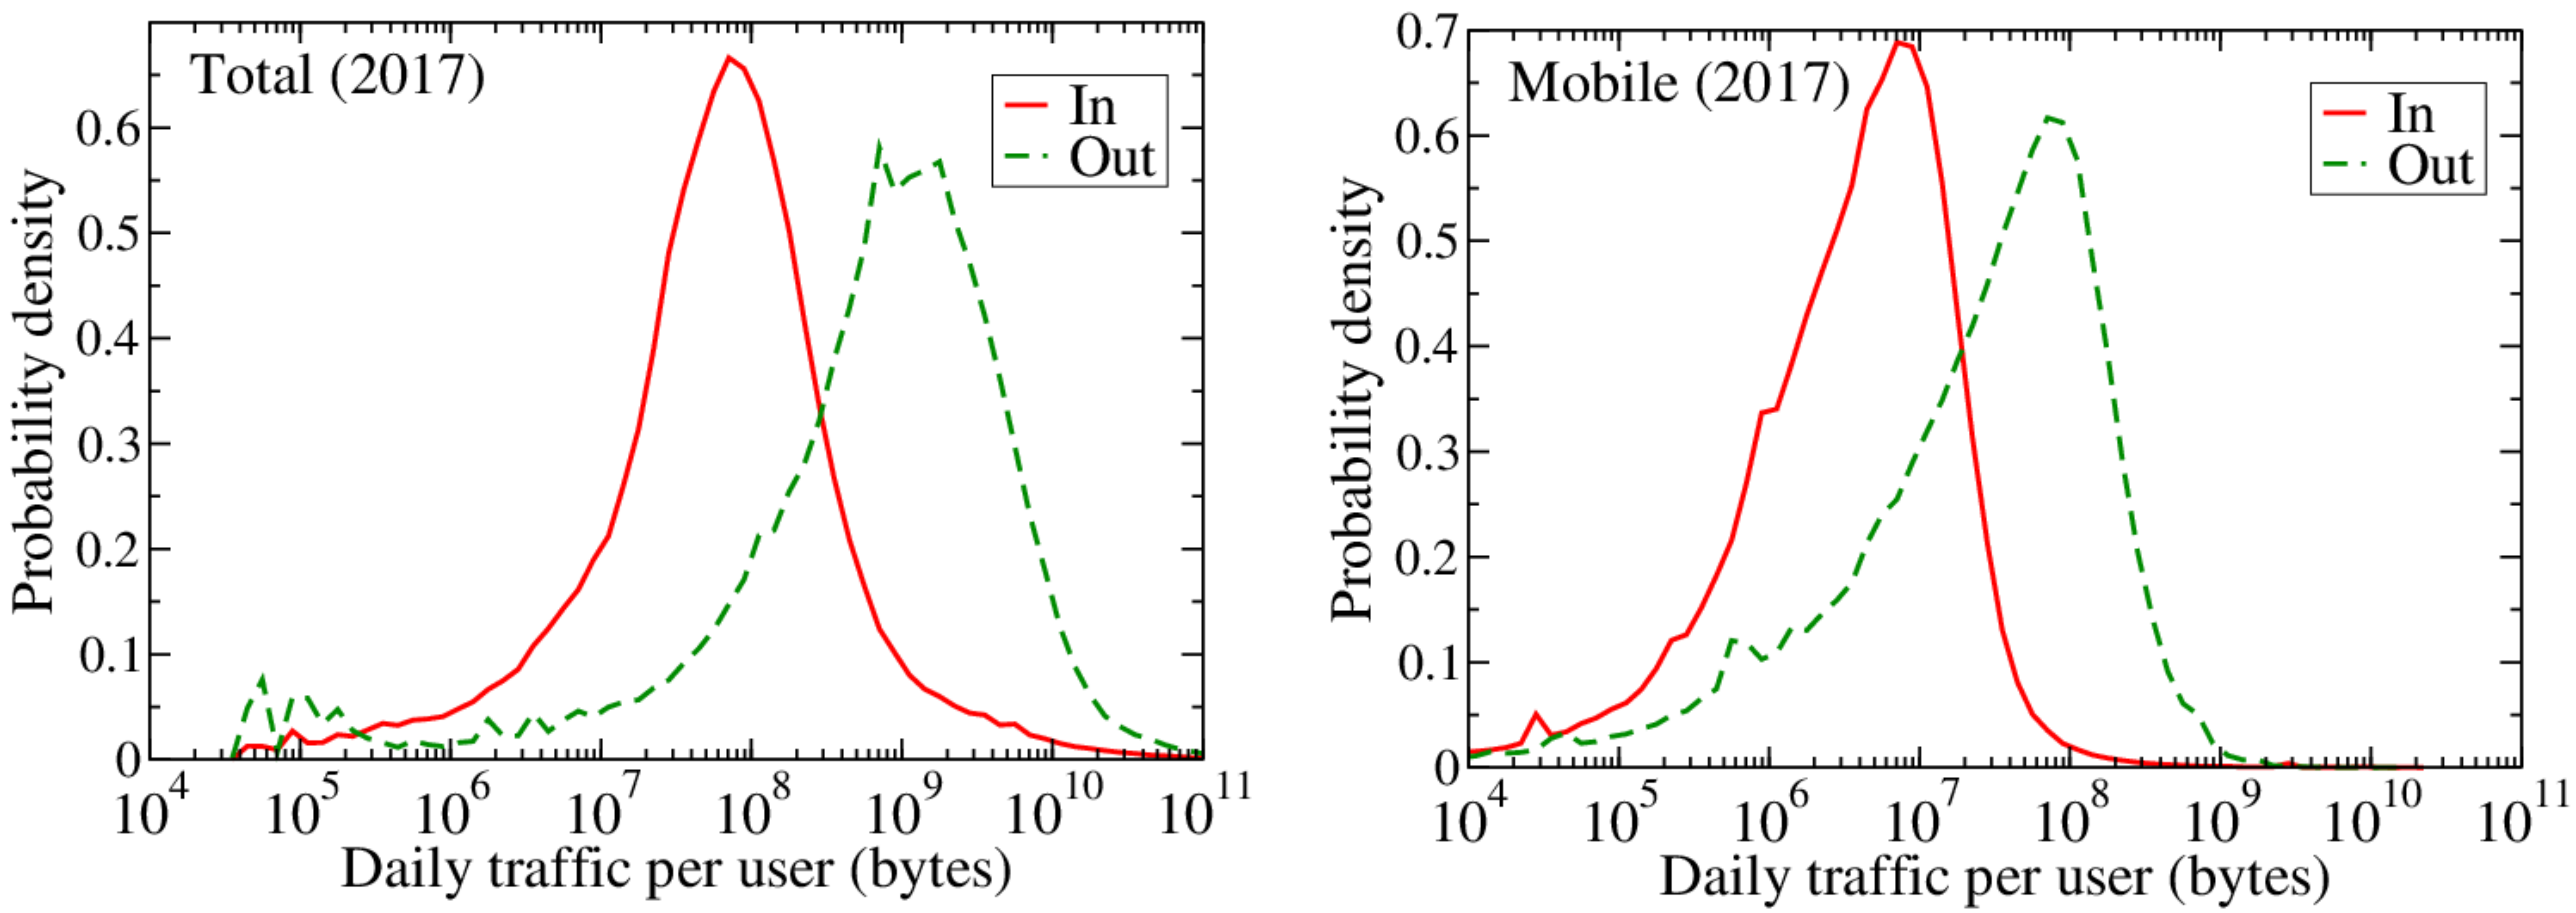

daily traffic usage per user: broadband (left) and mobile (right)

per-user traffic over 12 years

- broadband downloads over 12 years
  - mean: x5.9, mode: x39
- heavy user traffic hasn't increased much
- mode of mobile traffic: about broadband 10 years ago

broadband

	IN (MB/day)			OUT (MB/day)		
year	mean	median	mode	mean	median	mode
2005	430	3	3.5	447	30	32
2007	433	5	4	712	58	66
2008	483	6	5	797	73	94
2009	556	7	6	971	88	114
2010	469	8	7	910	108	145
2011	432	9	8.5	1,001	142	223
2012	410	12	14	1,026	173	282
2013	397	14	18	1,038	203	355
2014	437	22	28	1,287	301	447
2015	467	33	40	1,621	430	708
2016	475	48	56	2,081	697	1,000
2017	520	63	79	2,624	835	1,260

mobile

	IN (MB/day)			OUT (MB/day)		
year	mean	median	mode	mean	median	mode
2015	6.0	2.7	5.5	46.6	19	40
2016	7.8	3.6	7	63.0	27	63
2017	12.0	4.3	7	77.4	35	79

broadband 2005 vs. 2017

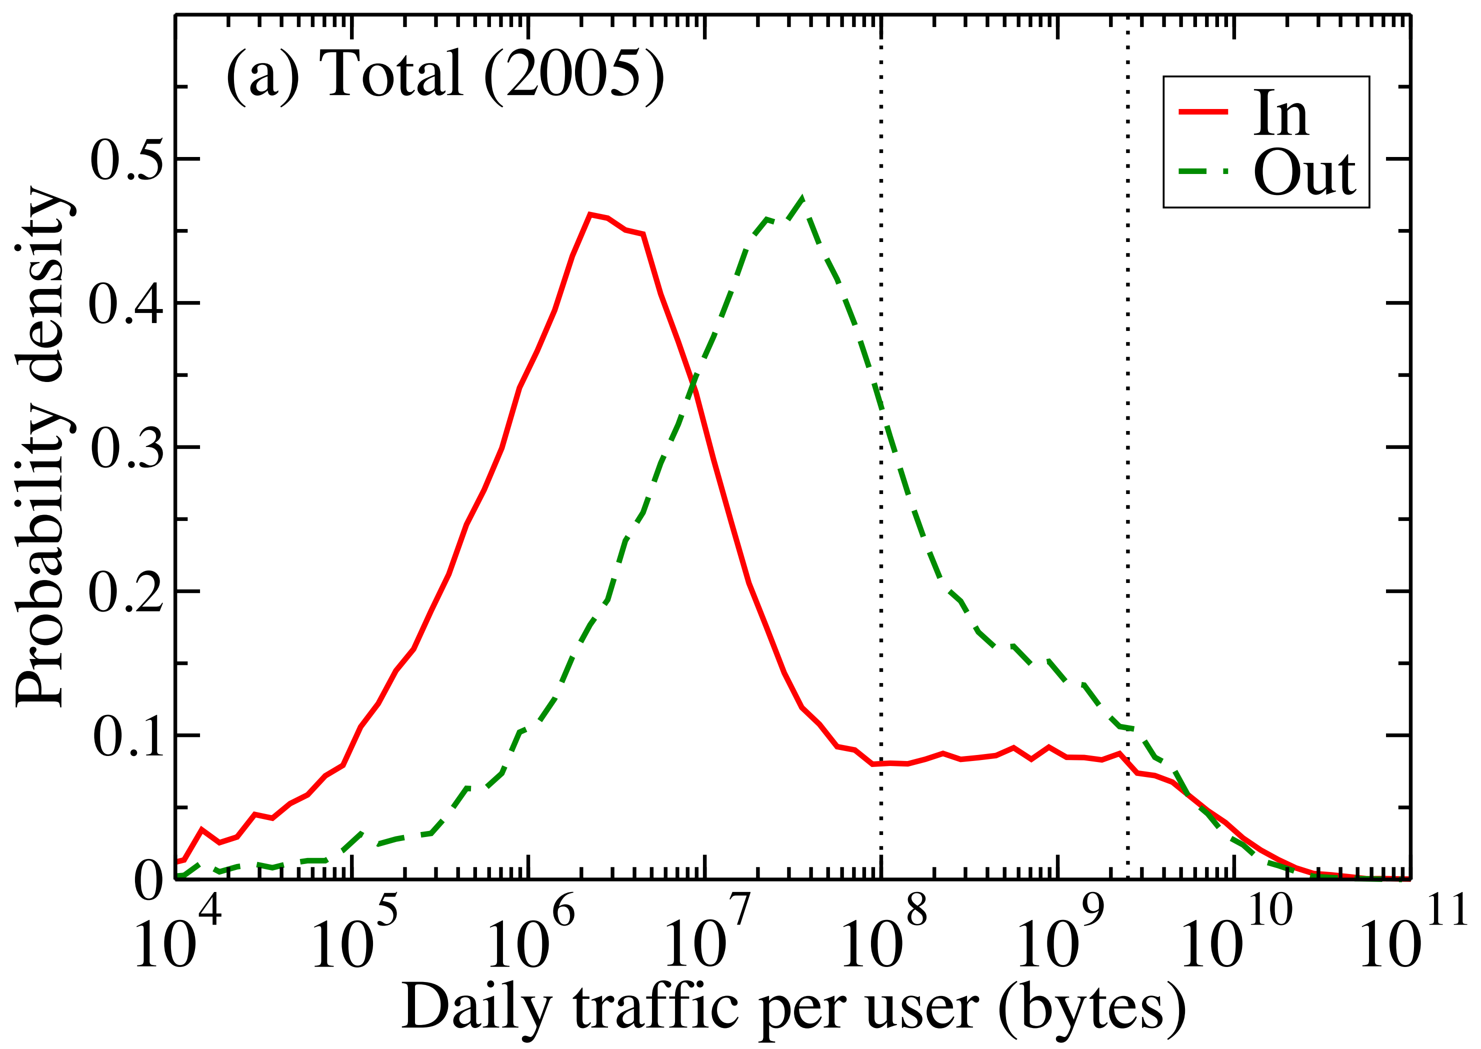

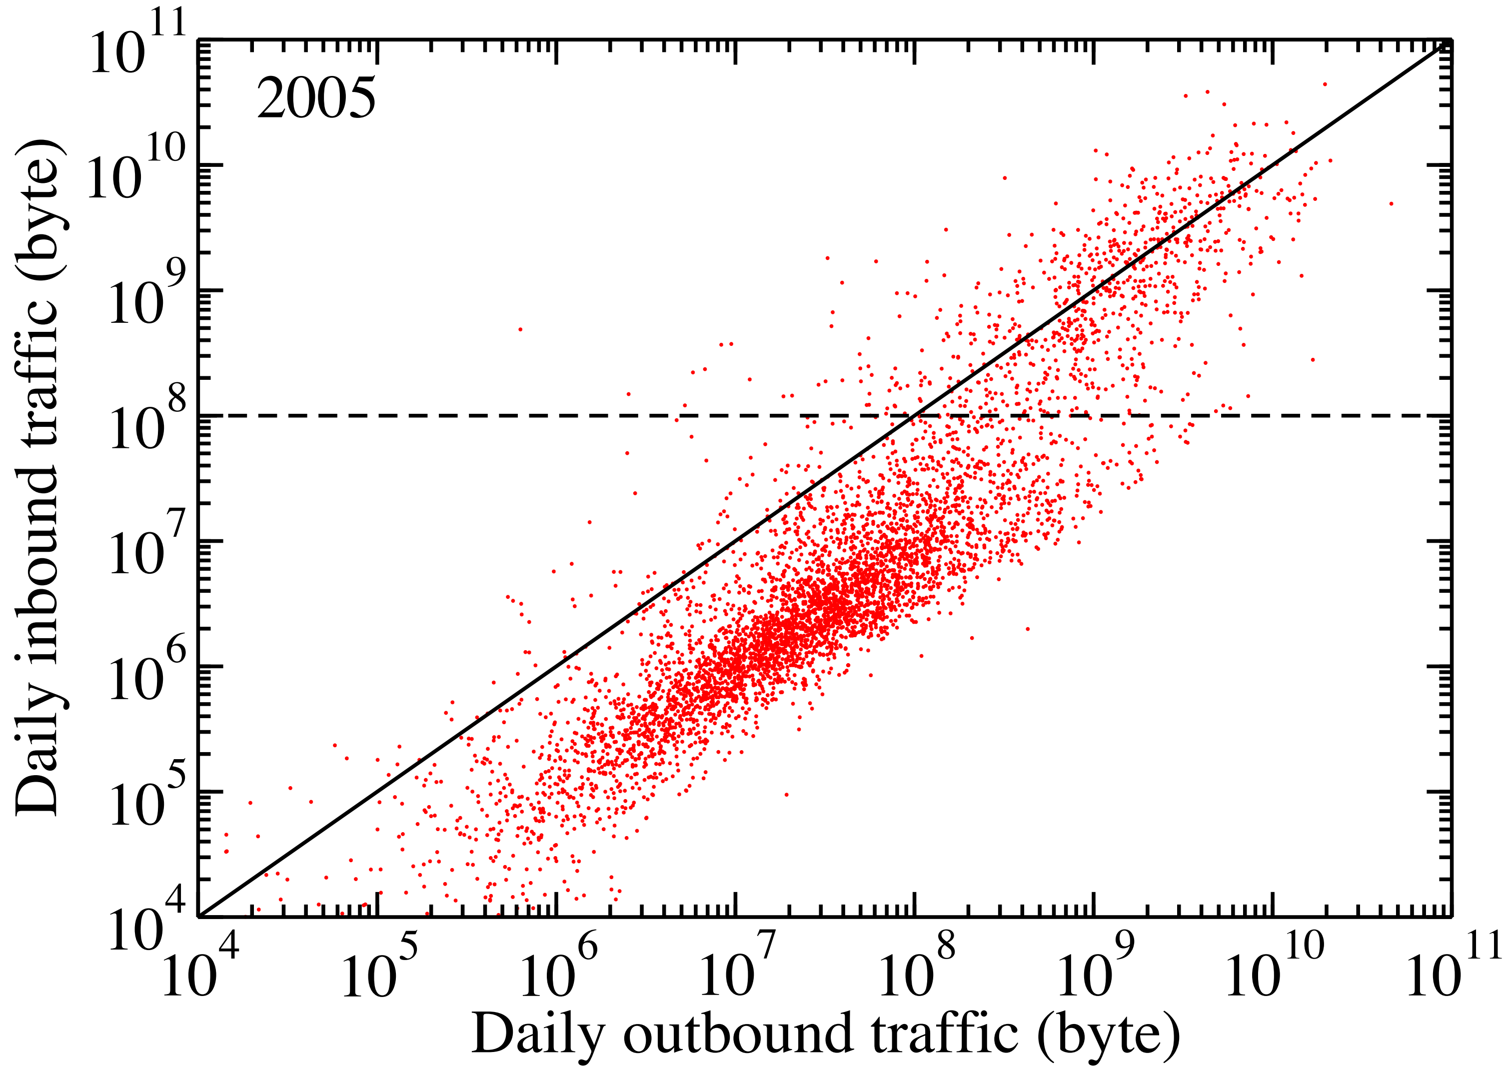

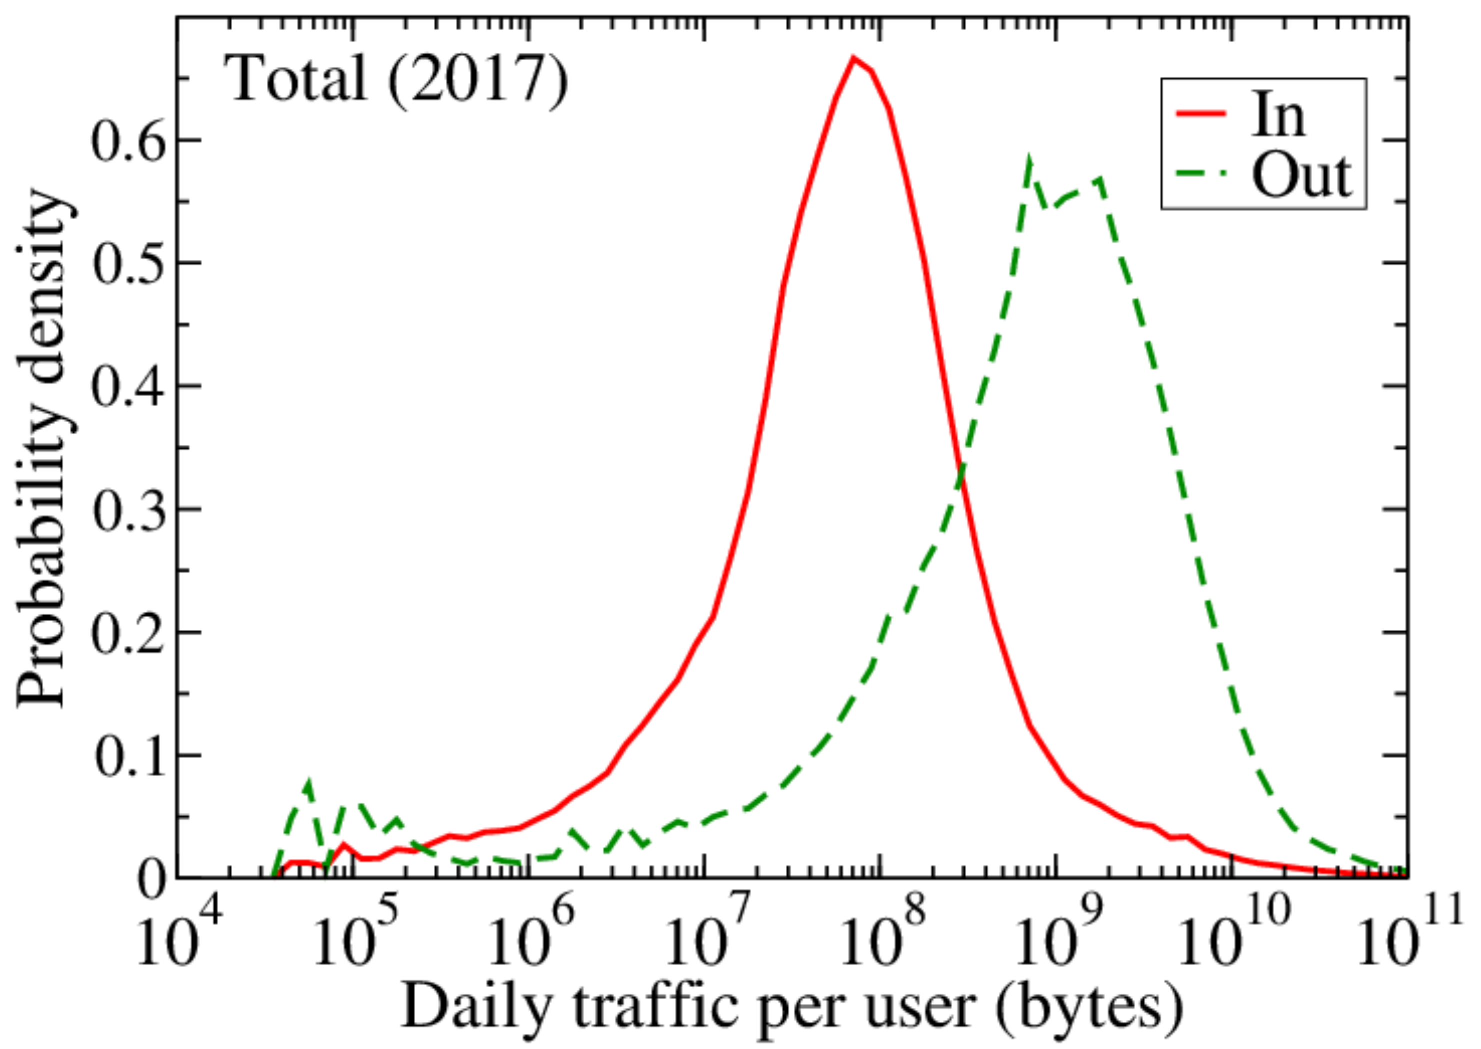

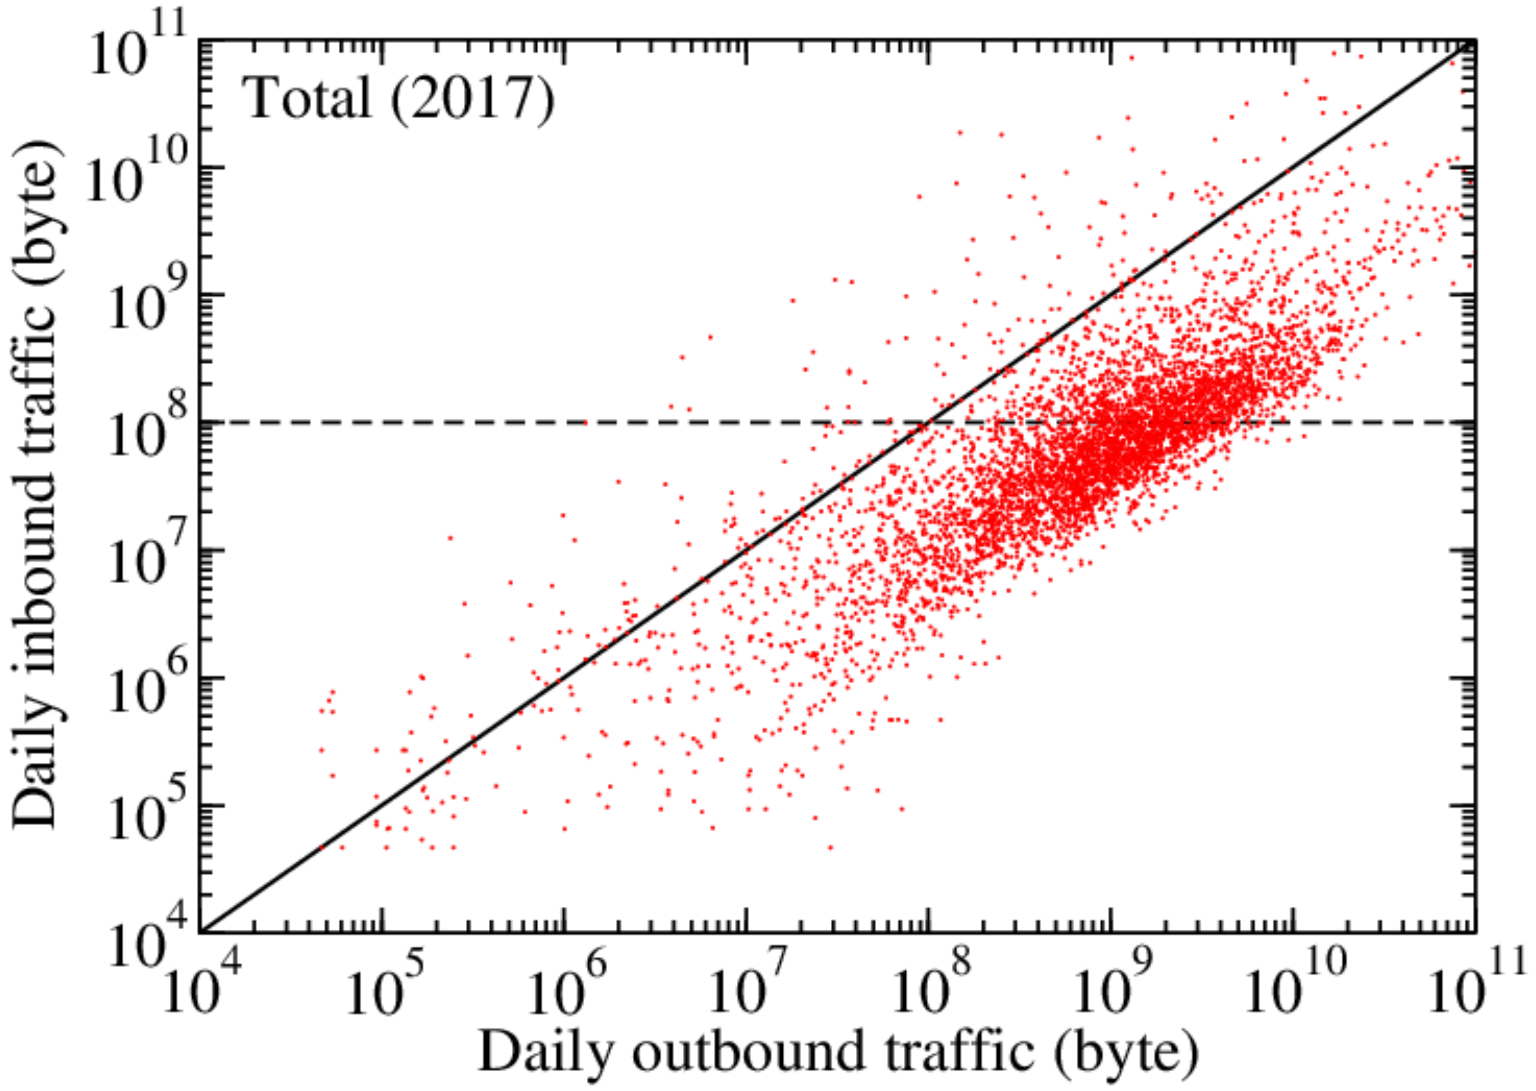

traffic usage skew among users

- heavy users: statistically distributed, not outliers
  - broadband: top 10% users account for 60% of OUT, 85% of IN
  - mobile: top 10% users account for 48% of OUT, 62% of IN

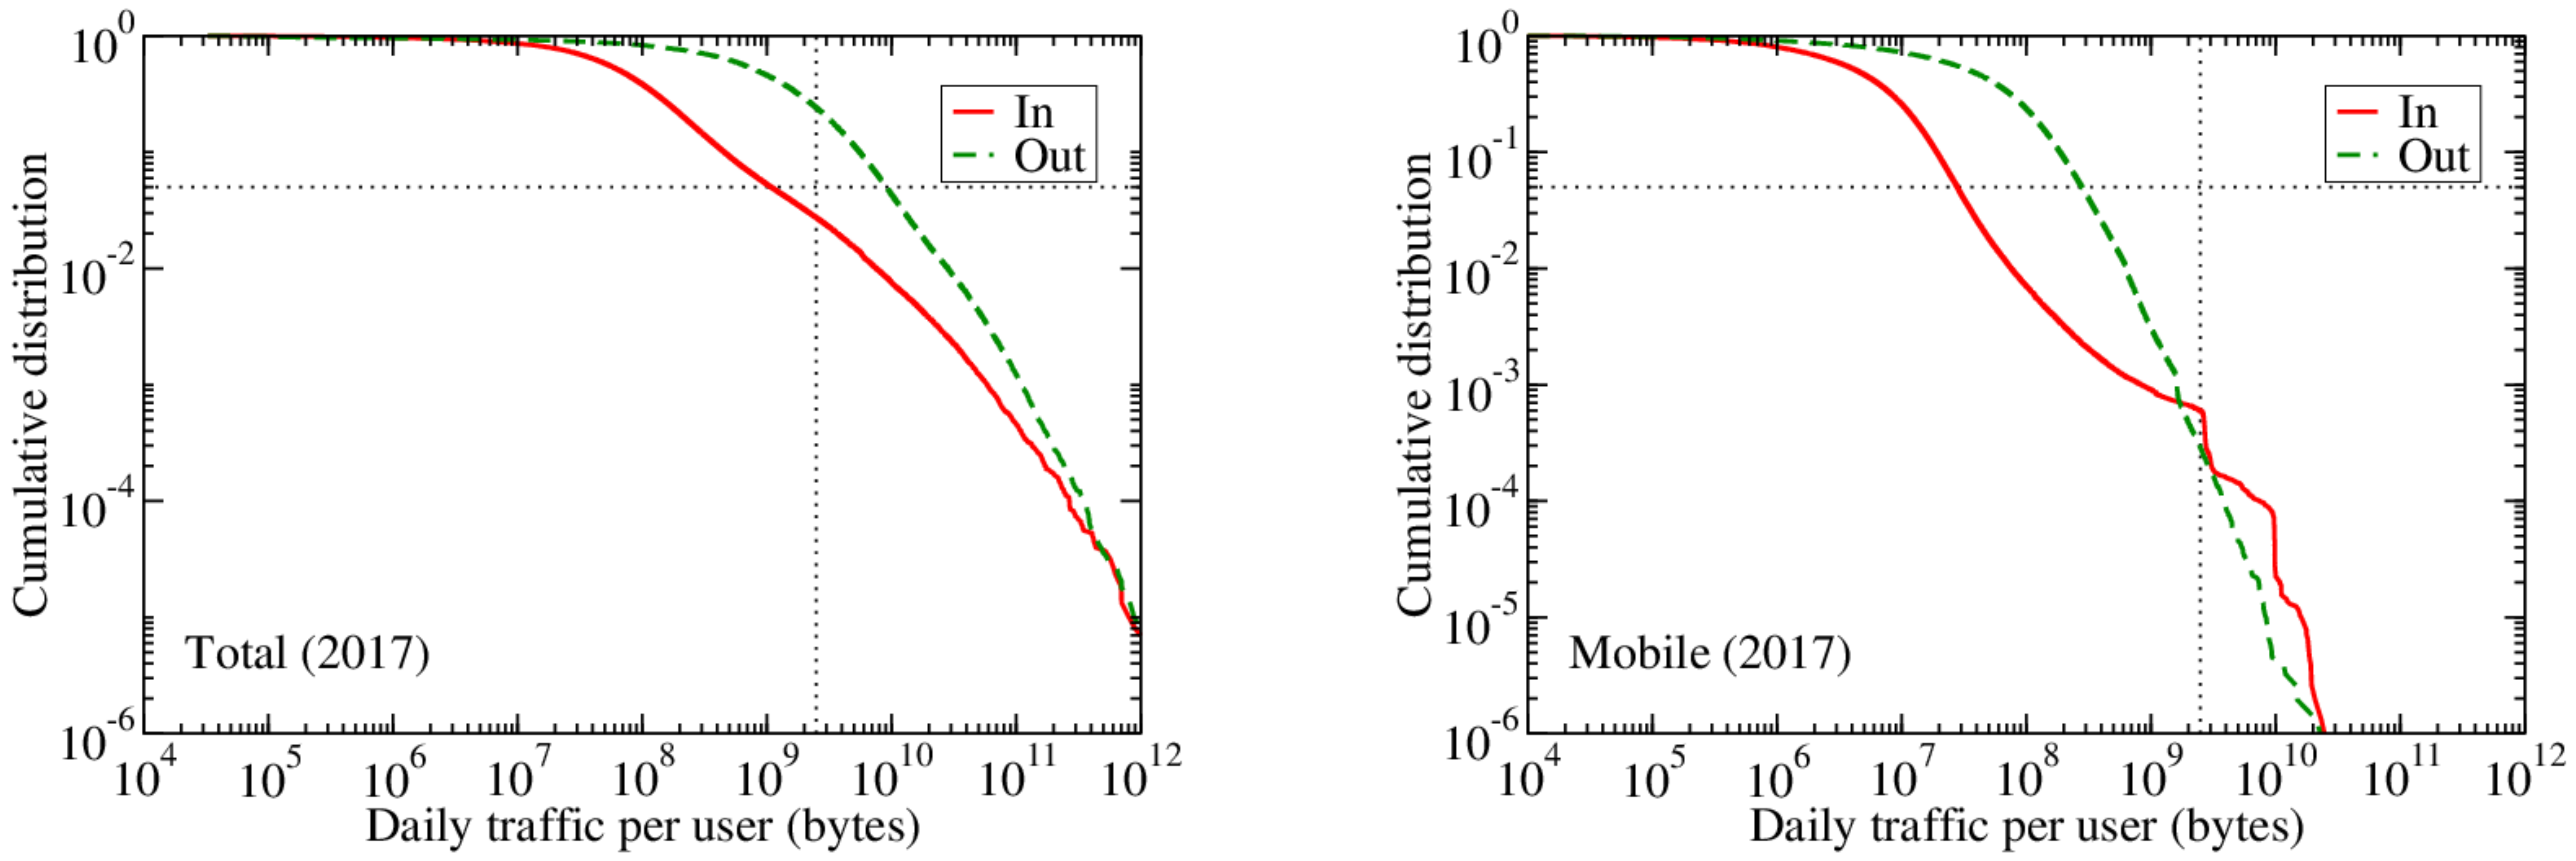

CCDF of daily traffic usage per user: broadband (left) and mobile (right)

port usage

- port no: min(sport, dport)
  - increasing TCP 443, UDP 443

broadband

protocol port	2016 total (%)	2016 light users	2017 total (%)	2017 light users
TCP	**82.8**	**93.3**	**83.9**	**92.3**
(< 1024)	63.3	89.9	72.9	88.6
443(https)	30.5	39.6	43.3	52.5
80(http)	37.1	49.2	28.4	35.2
182	0.3	0.2	0.3	0.3
81	0.4	0.7	0.2	0.2
993(imaps)	0.1	0.1	0.2	0.1
22(ssh)	0.2	0.0	0.1	0.0
110(pop3)	0.1	0.1	0.1	0.1
(>= 1024)	13.7	3.2	11.0	3.7
1935(rtmp)	1.5	1.7	1.1	1.2
8080	0.2	0.1	0.3	0.1
UDP	**11.1**	**4.0**	**10.5**	**4.9**
443(https)	2.4	2.8	3.8	3.7
4500(nat-t)	0.2	0.1	0.2	0.1
ESP	**5.8**	**2.6**	**5.1**	**2.7**
IP-ENCAP	**0.2**	**0.0**	**0.3**	**0.0**
GRE	**0.1**	**0.0**	**0.1**	**0.0**
ICMP	**0.0**	**0.0**	**0.0**	**0.0**

mobile

protocol port	2016 total (%)	2017 total (%)
TCP	**94.4**	**84.4**
443(https)	43.7	53.0
80(http)	46.8	27.0
31000	0.2	1.8
993(imaps)	0.5	0.4
1935(rtmp)	0.3	0.2
81	0.5	0.1
UDP	**5.0**	**11.4**
443(https)	1.5	7.5
12222	0.1	0.1
4500(nat-t)	0.2	0.2
53(dns)	0.2	0.1
ESP	**0.4**	**0.4**
GRE	**0.1**	**0.1**
ICMP	**0.0**	**0.0**

port usage over a week

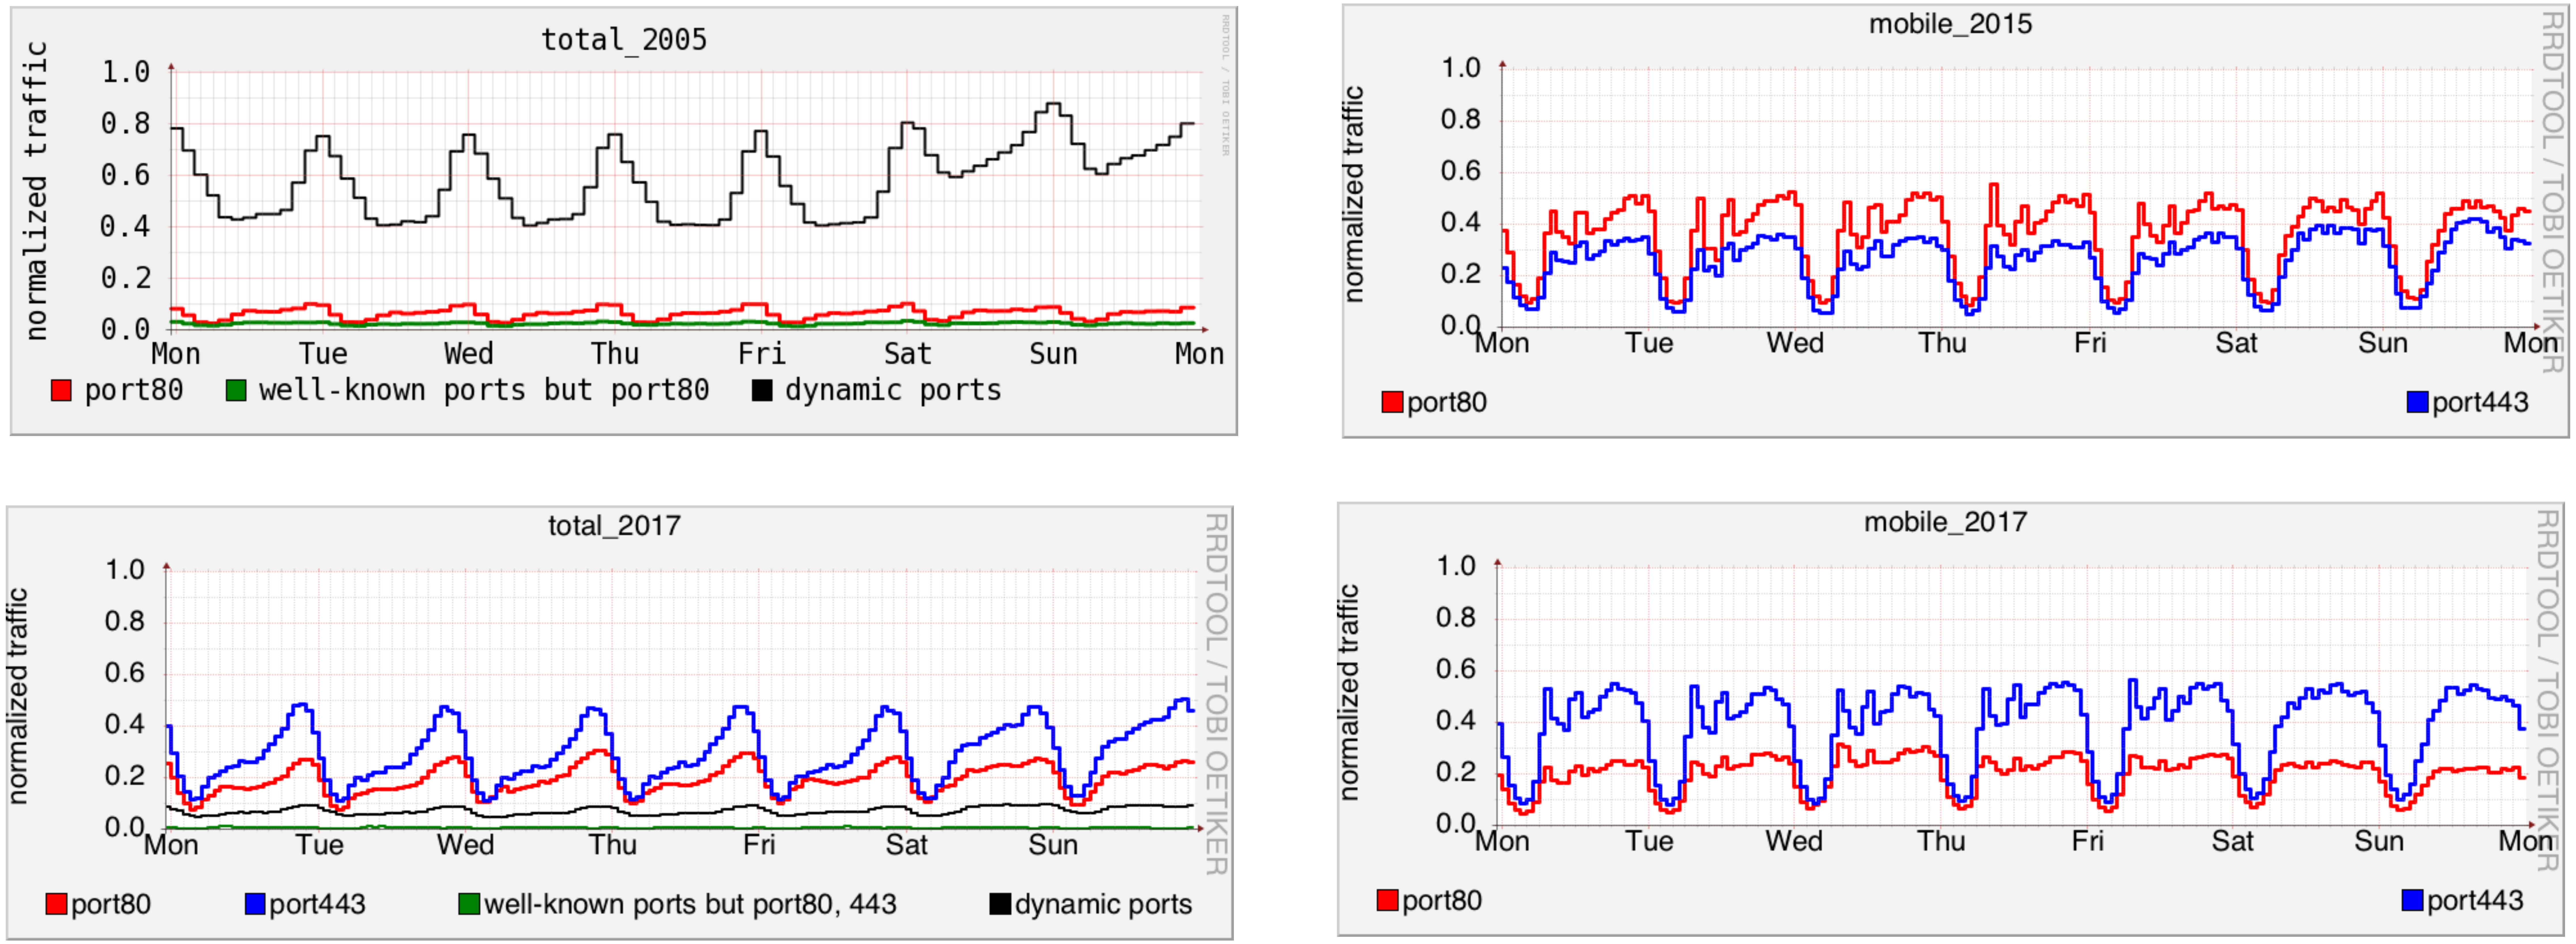

TCP port usage over a week: broadband (left) and mobile (right)

redundancy matters

- example: Japan earthquake in 2011

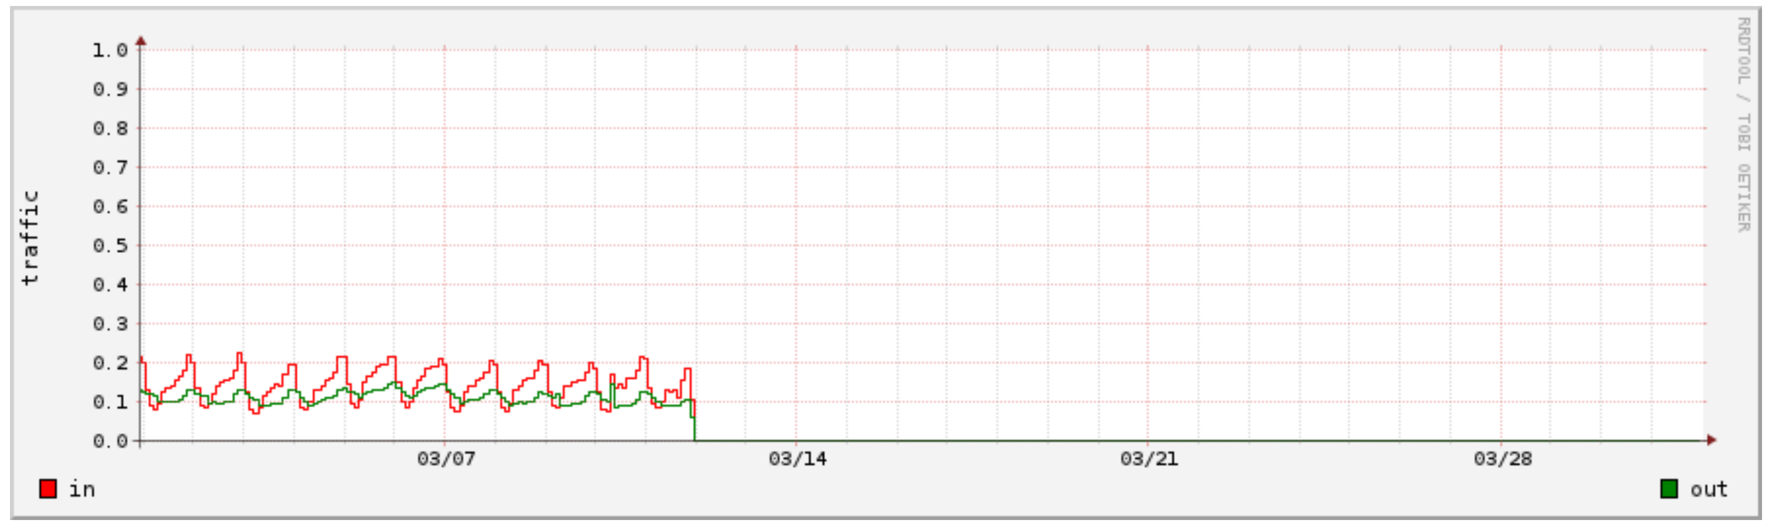

a failed link:
one of the US-JP submarine cables damaged

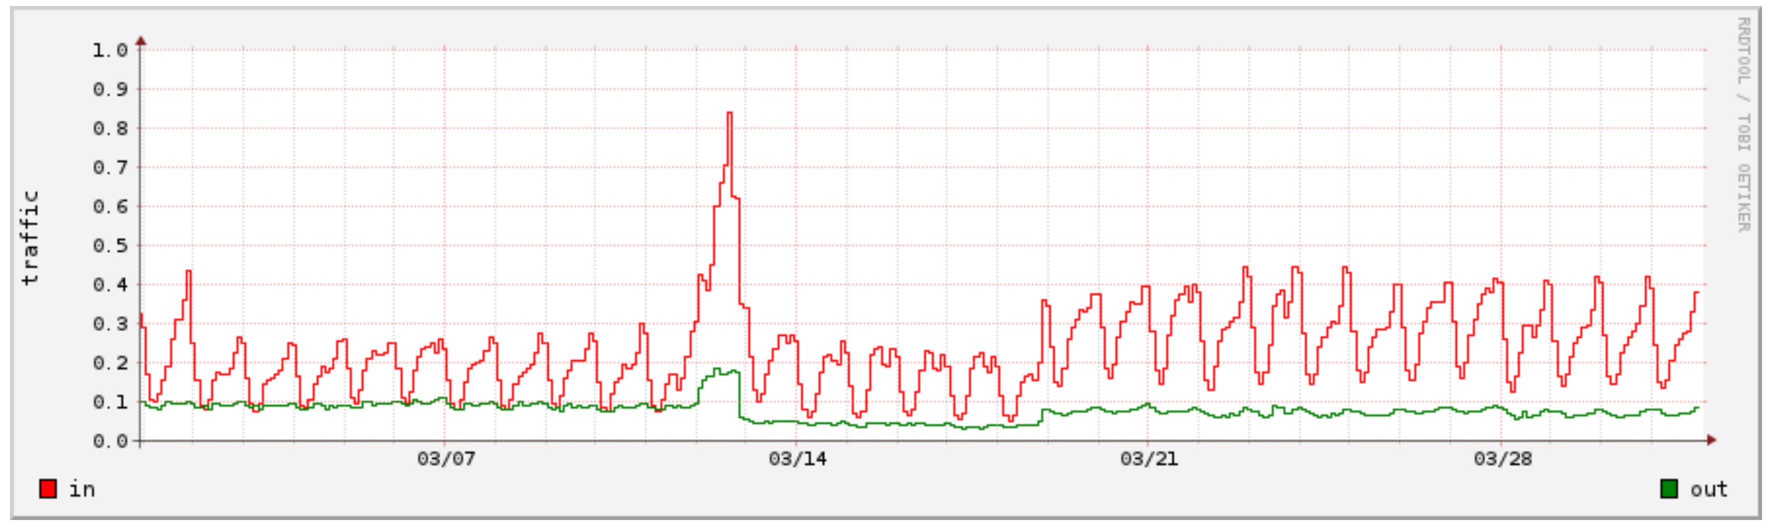

a survived link:
immediately, traffic rerouted to available links

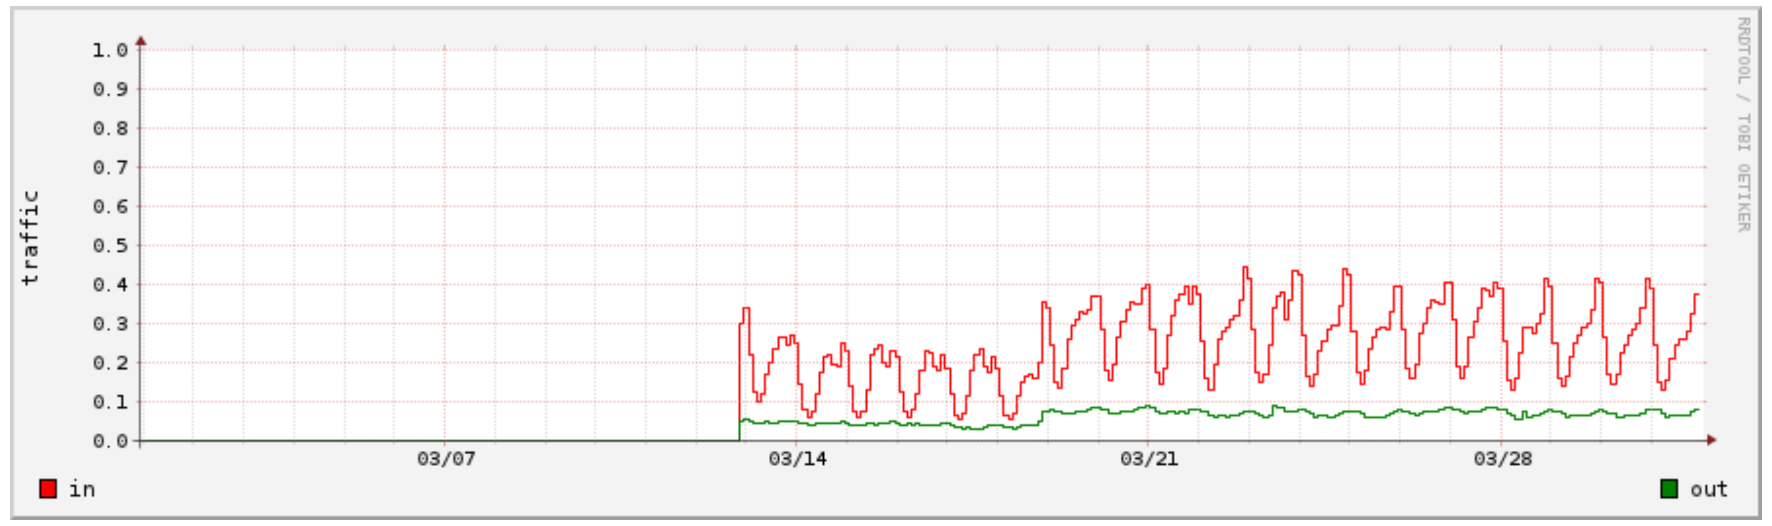

an alternate link:
one-day after, a new link brought up

IPv6 in Japan (by APNIC Labs)

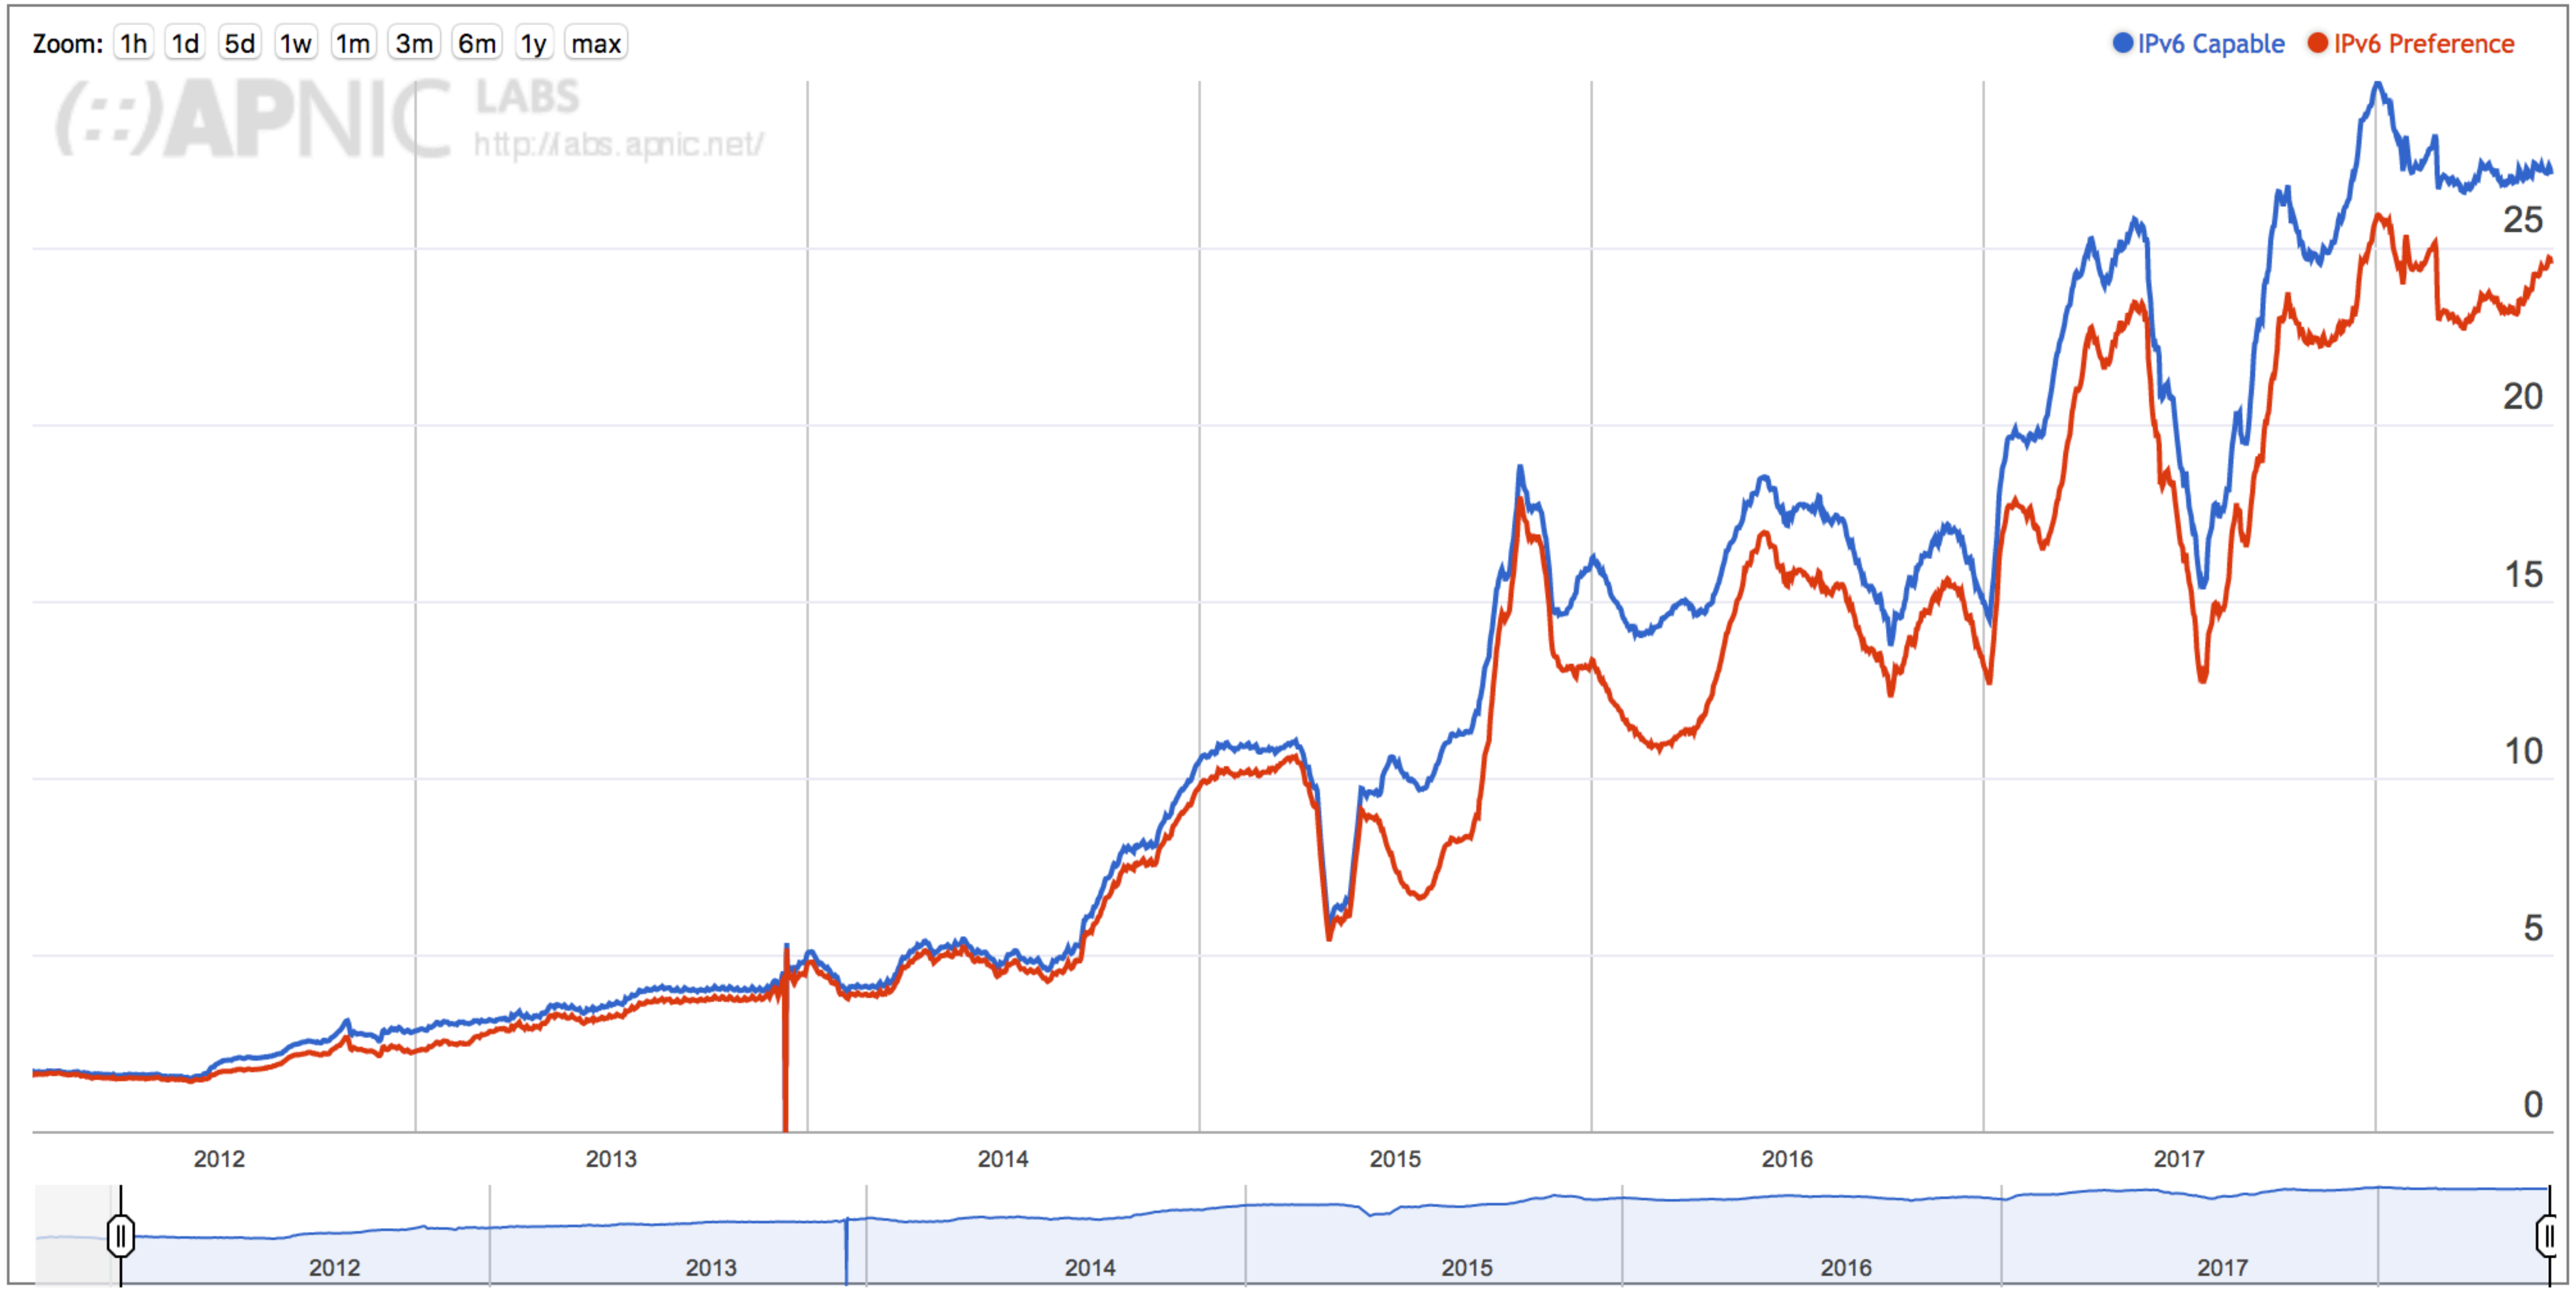

https://stats.labs.apnic.net/ipv6/JP

hyper-giants observed from IIJ's broadband traffic

- by mapping TCP 80 and 443 src addresses to ASes

ASes	traffic volume (%)
google	30.1
akamai	7.1
apple	6.6
limelight	4.8
amazon	4.5
yahoo japan	3.4
microsoft	3.1
netflix	2.4
facebook	2.3
edgecast	1.9

Upstream		Downstream		Aggregate	
BitTorrent	18.37%	Netflix	35.15%	Netflix	32.72%
YouTube	13.13%	YouTube	17.53%	YouTube	17.31%
Netflix	10.33%	Amazon Video	4.26%	HTTP - OTHER	4.14%
SSL - OTHER	8.55%	HTTP - OTHER	4.19%	Amazon Video	3.96%
Google Cloud	6.98%	iTunes	2.91%	SSL - OTHER	3.12%
iCloud	5.98%	Hulu	2.68%	BitTorrent	2.85%
HTTP - OTHER	3.70%	SSL - OTHER	2.53%	iTunes	2.67%
Facebook	3.04%	Xbox One Games Download	2.18%	Hulu	2.47%
FaceTime	2.50%	Facebook	1.89%	Xbox One Games Download	2.15%
Skype	1.75%	BitTorrent	1.73%	Facebook	2.01%
	69.32%		74.33%		72.72%

sandvine

Table 1 - Top 10 Peak Period Applications - North America, Fixed Access

peak traffic in North America
Sandvine Global Internet Phenomena Report 2016

client-side measurements

- more accurate views from users
- probe deployment and sampling issues
- crowd-sourcing: promising for scalability
  - e.g., RIPE Altas, CAIDA ark

RIPE Atlas

- started in 2010, 10k+ probes worldwide, 300+ in Japan
- periodic ping and traceroute + custom measurements
- abundant but noisy data, requiring statistical filtering

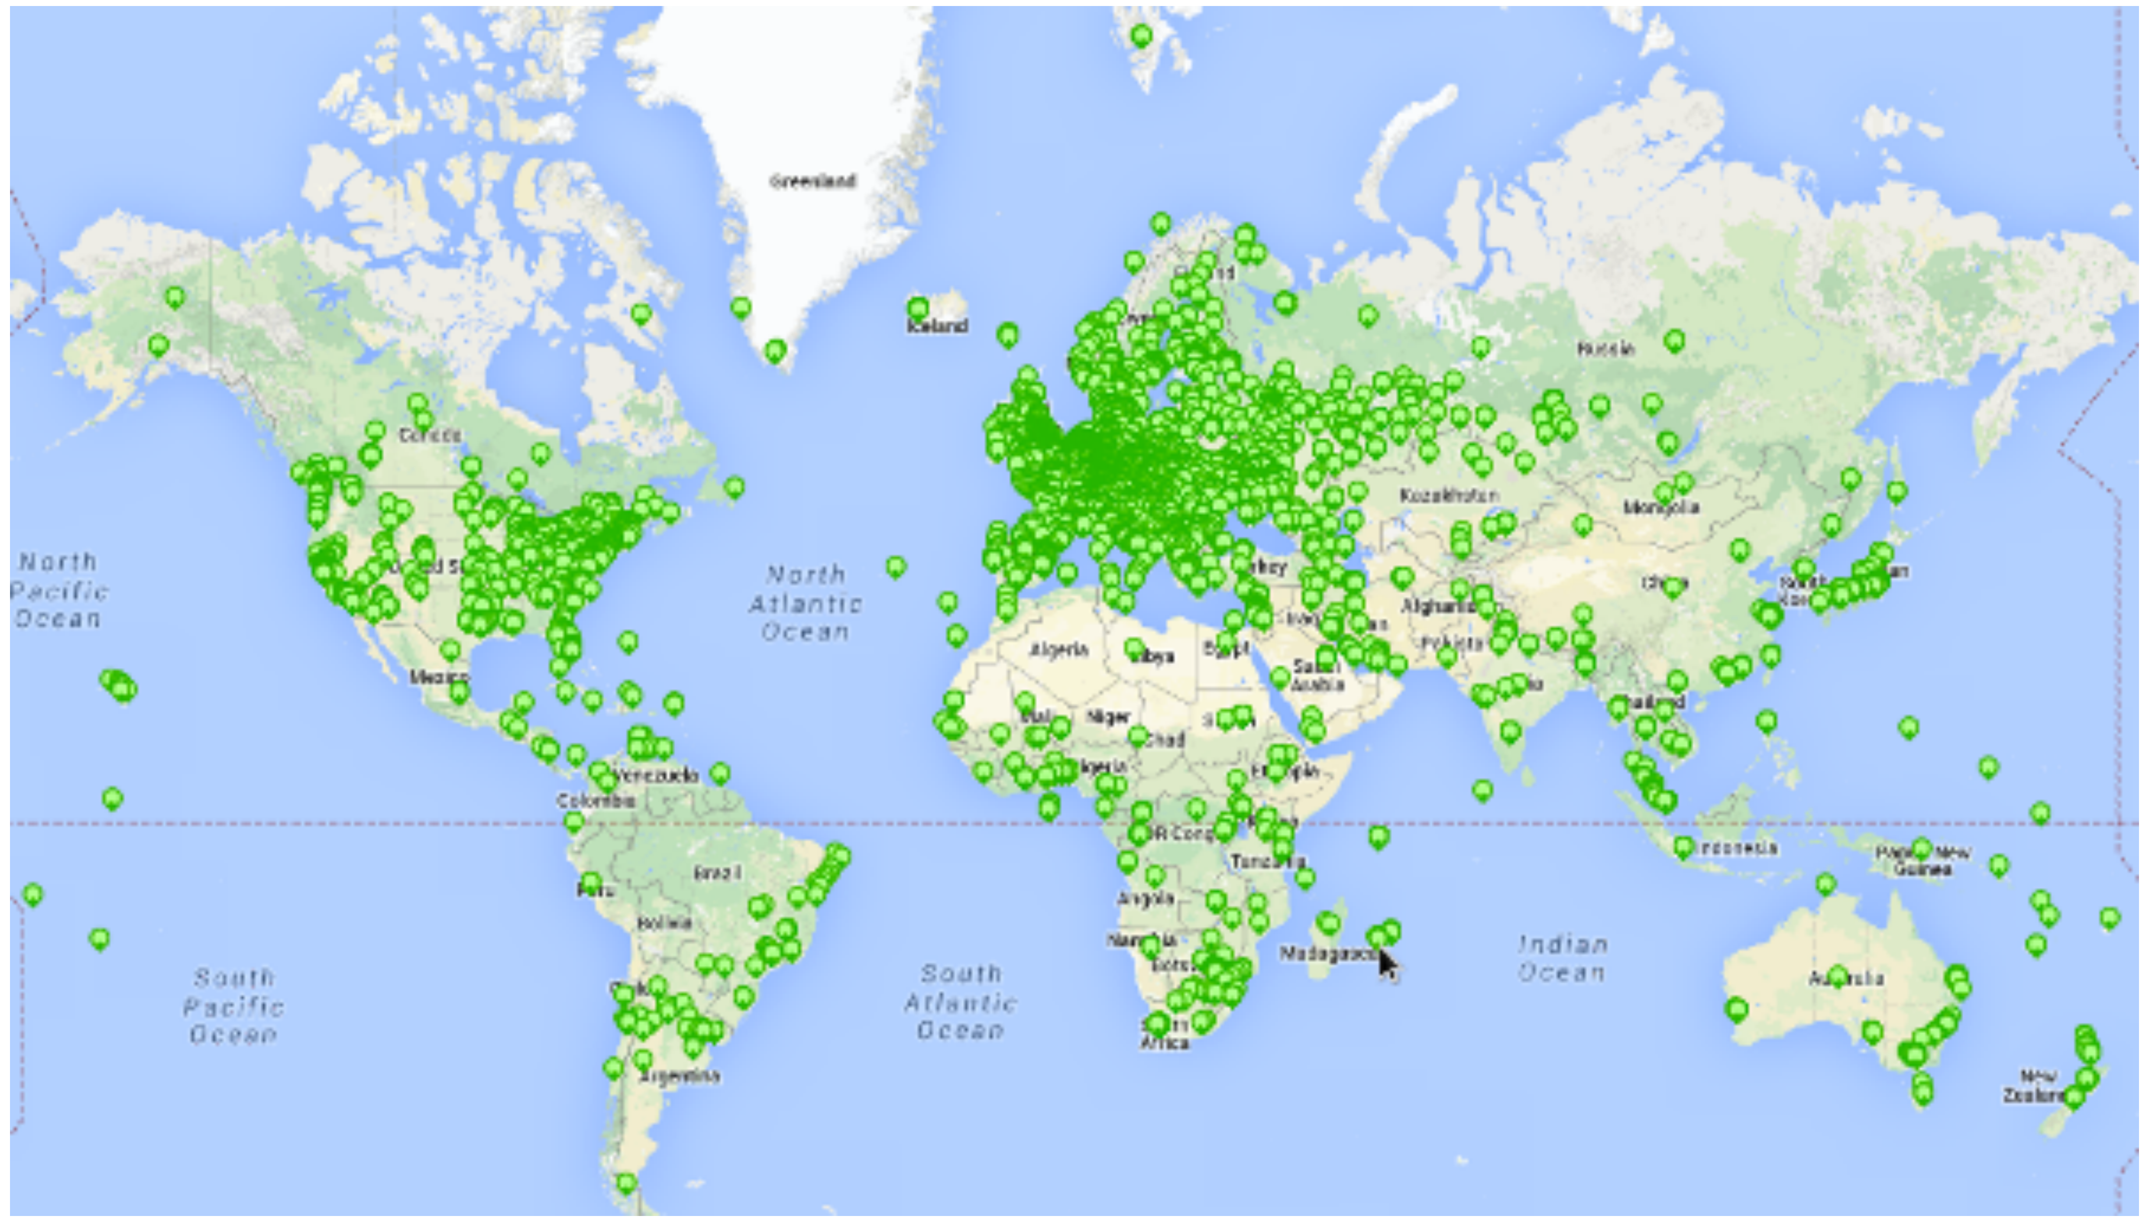

inferring access delays from RIPE Atlas traceroutes: preliminary results

- mean differential-RTT of Japanese access links over 5 weeks
  - between home routers and ISP's first hop
  - aggregated by the first hop address
- diurnal congestion observed in evening hours
  - no such congestion observed in non-FLETS access

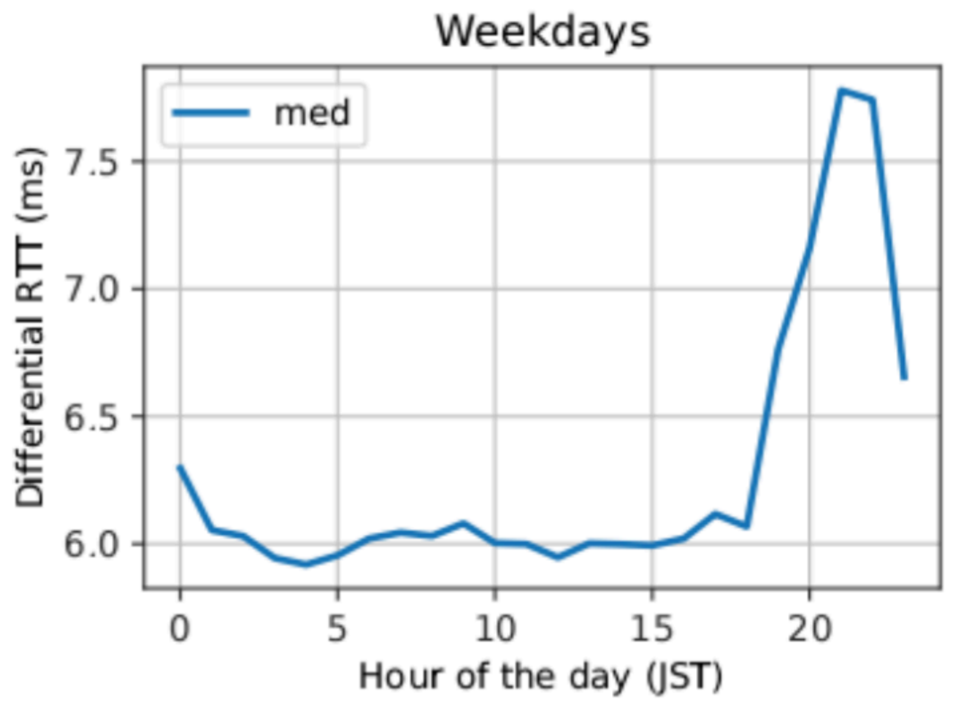

118.23.22.236
OCN (Tokyo)
6-8ms

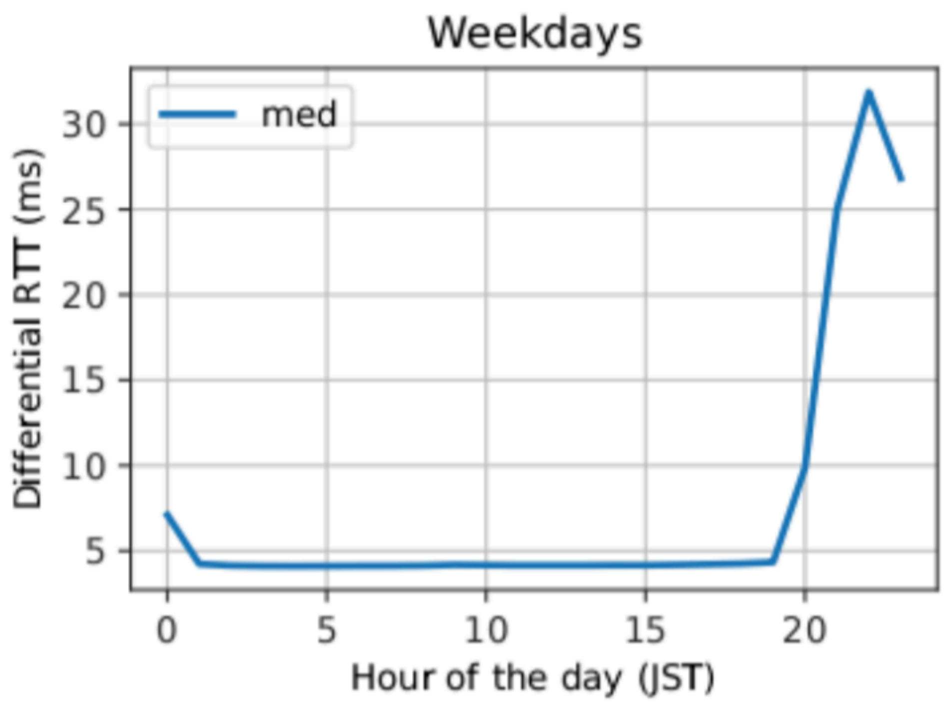

203.180.8.11
IIJ (Tokyo)
5-32ms

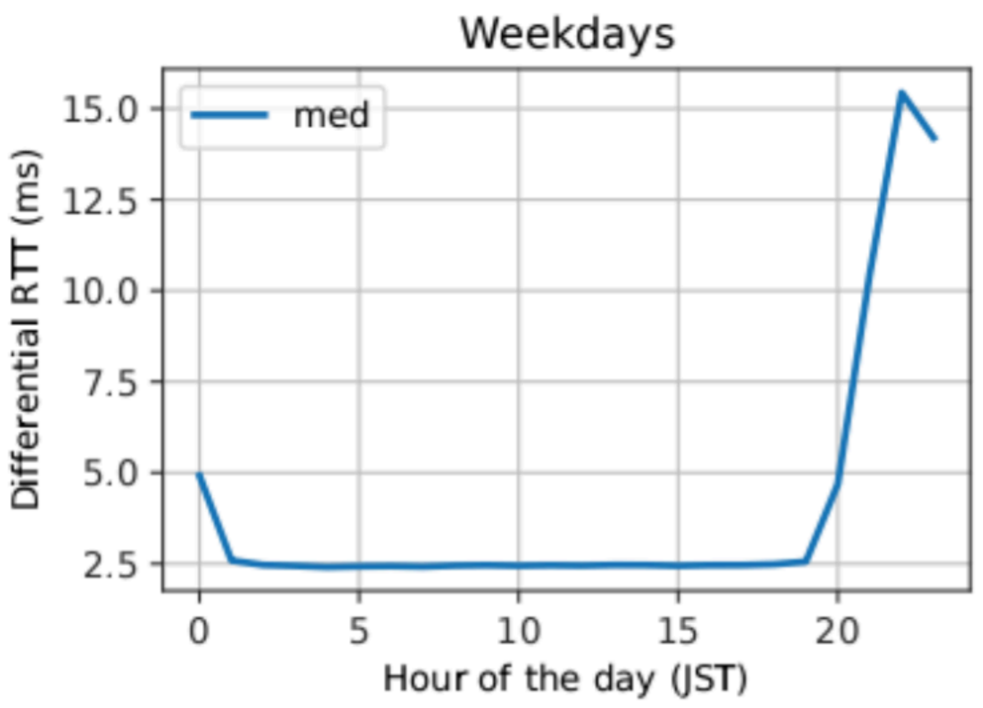

221.110.198.225
Softbank (Tokyo)
2.5-15ms

content-provider measurements

- http video-streaming
  - video as series of segment files
  - video replay under client's control
    - adaptive bit-rate: selected by client
  - pros: cost and scalability as web technologies are used (but traffic behavior is different)
  - cons: extra-delay, diverse clients' behaviors

Netflix ISP Speed Index (Japan)

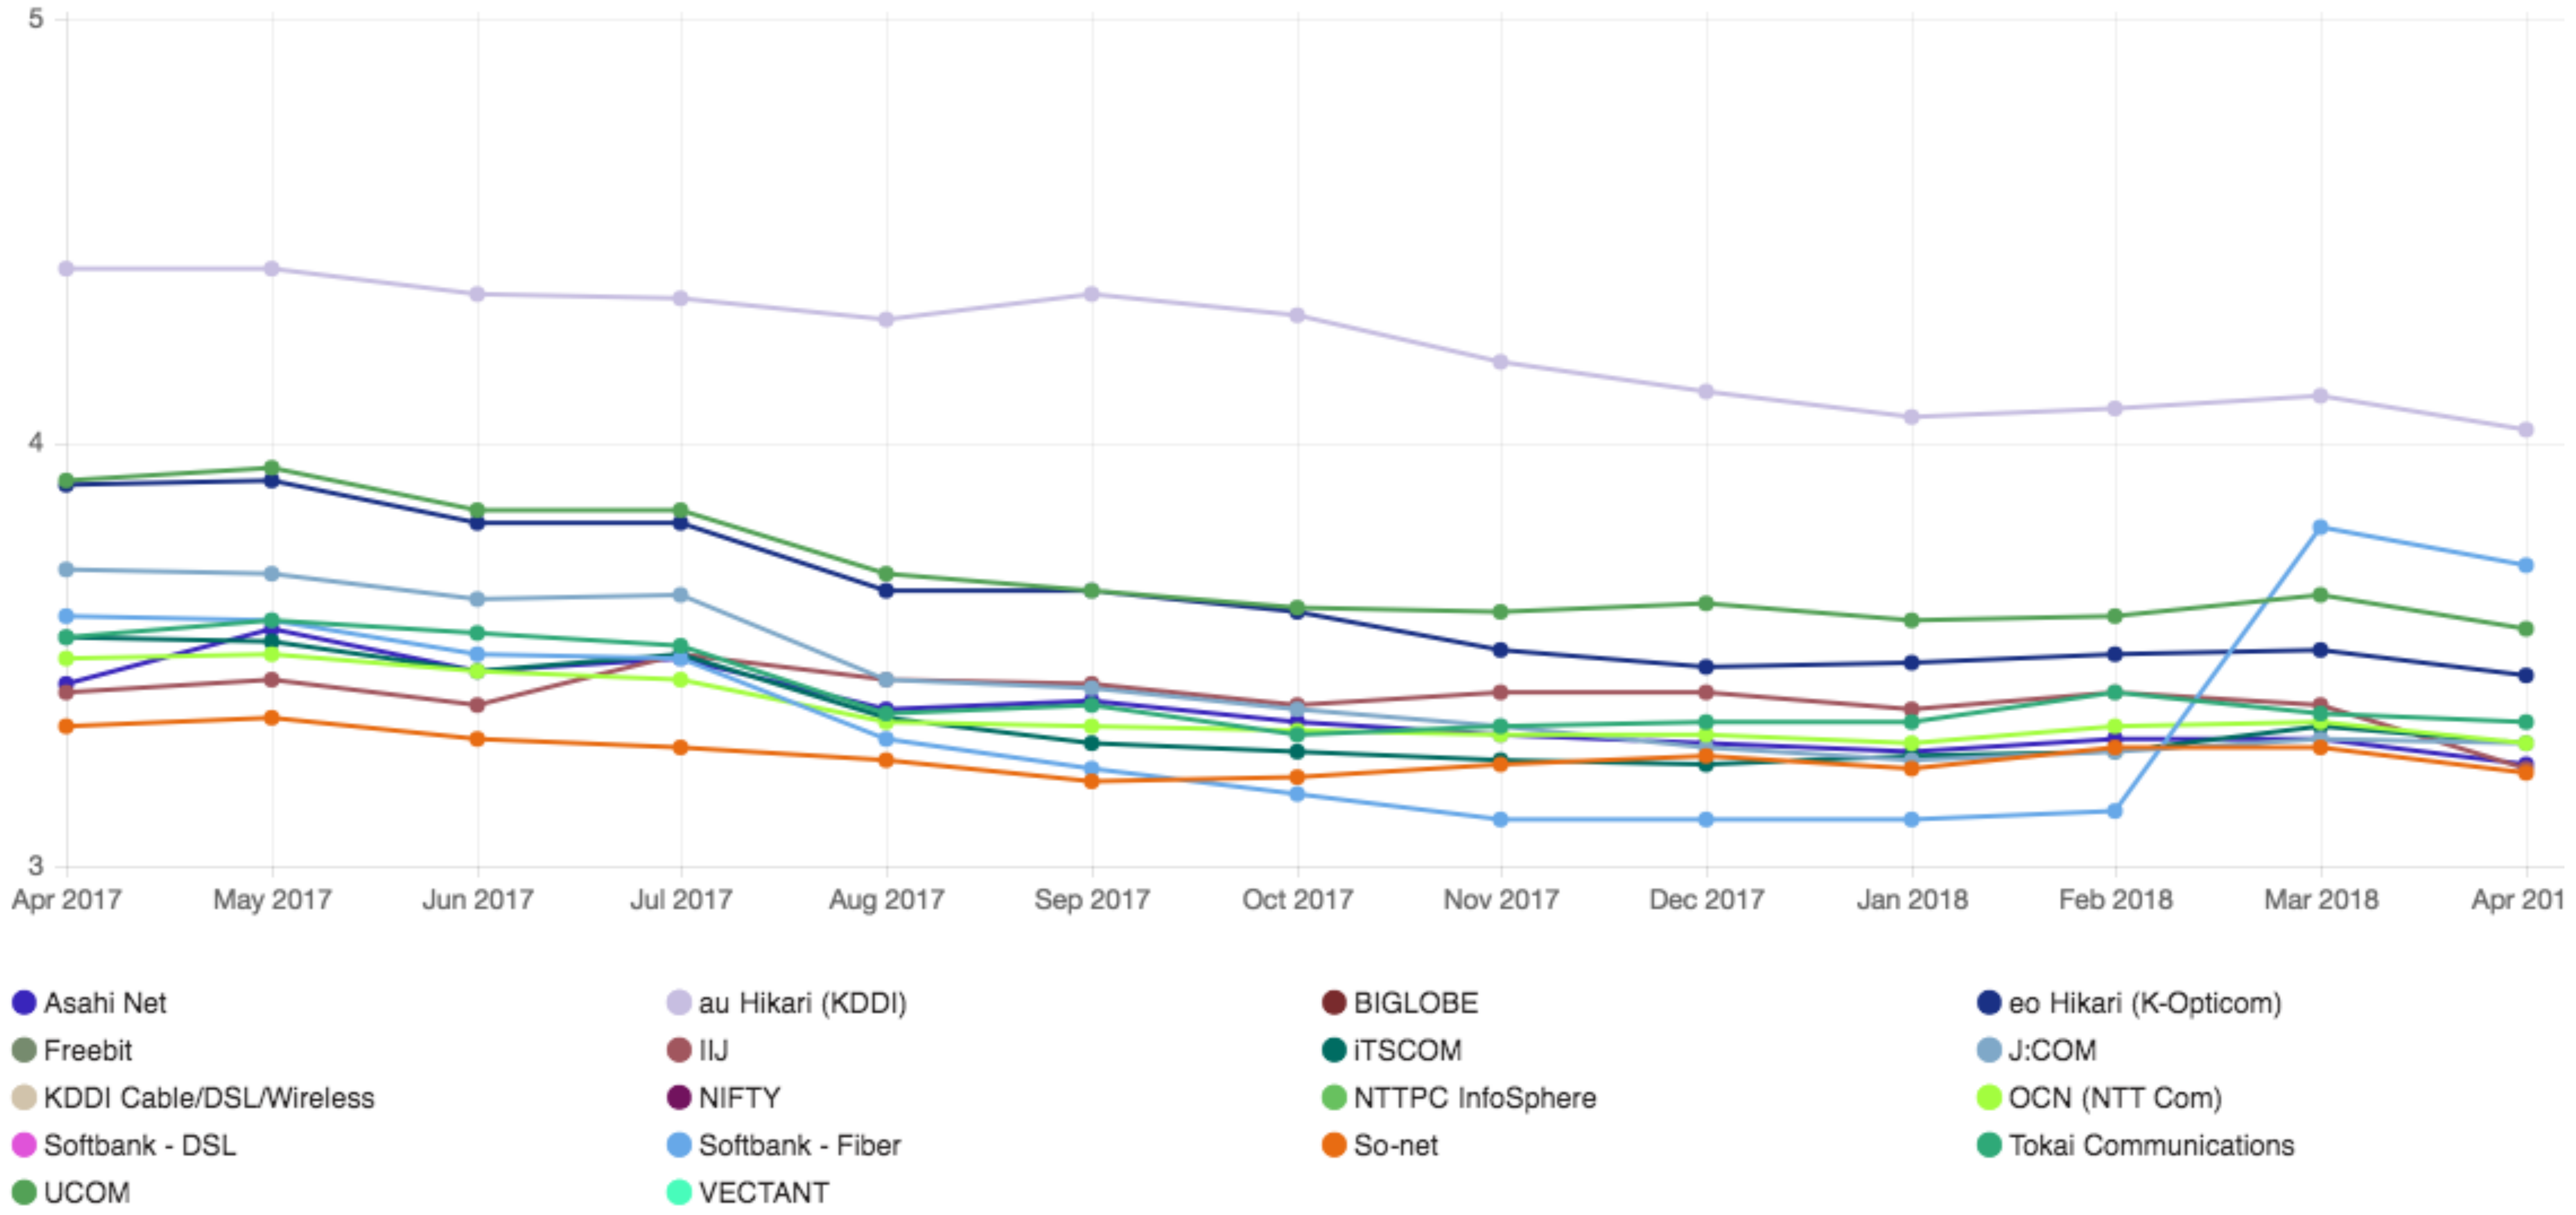

source: https://ispspeedindex.netflix.com/country/japan/

Netflix ISP Speed Index (US)

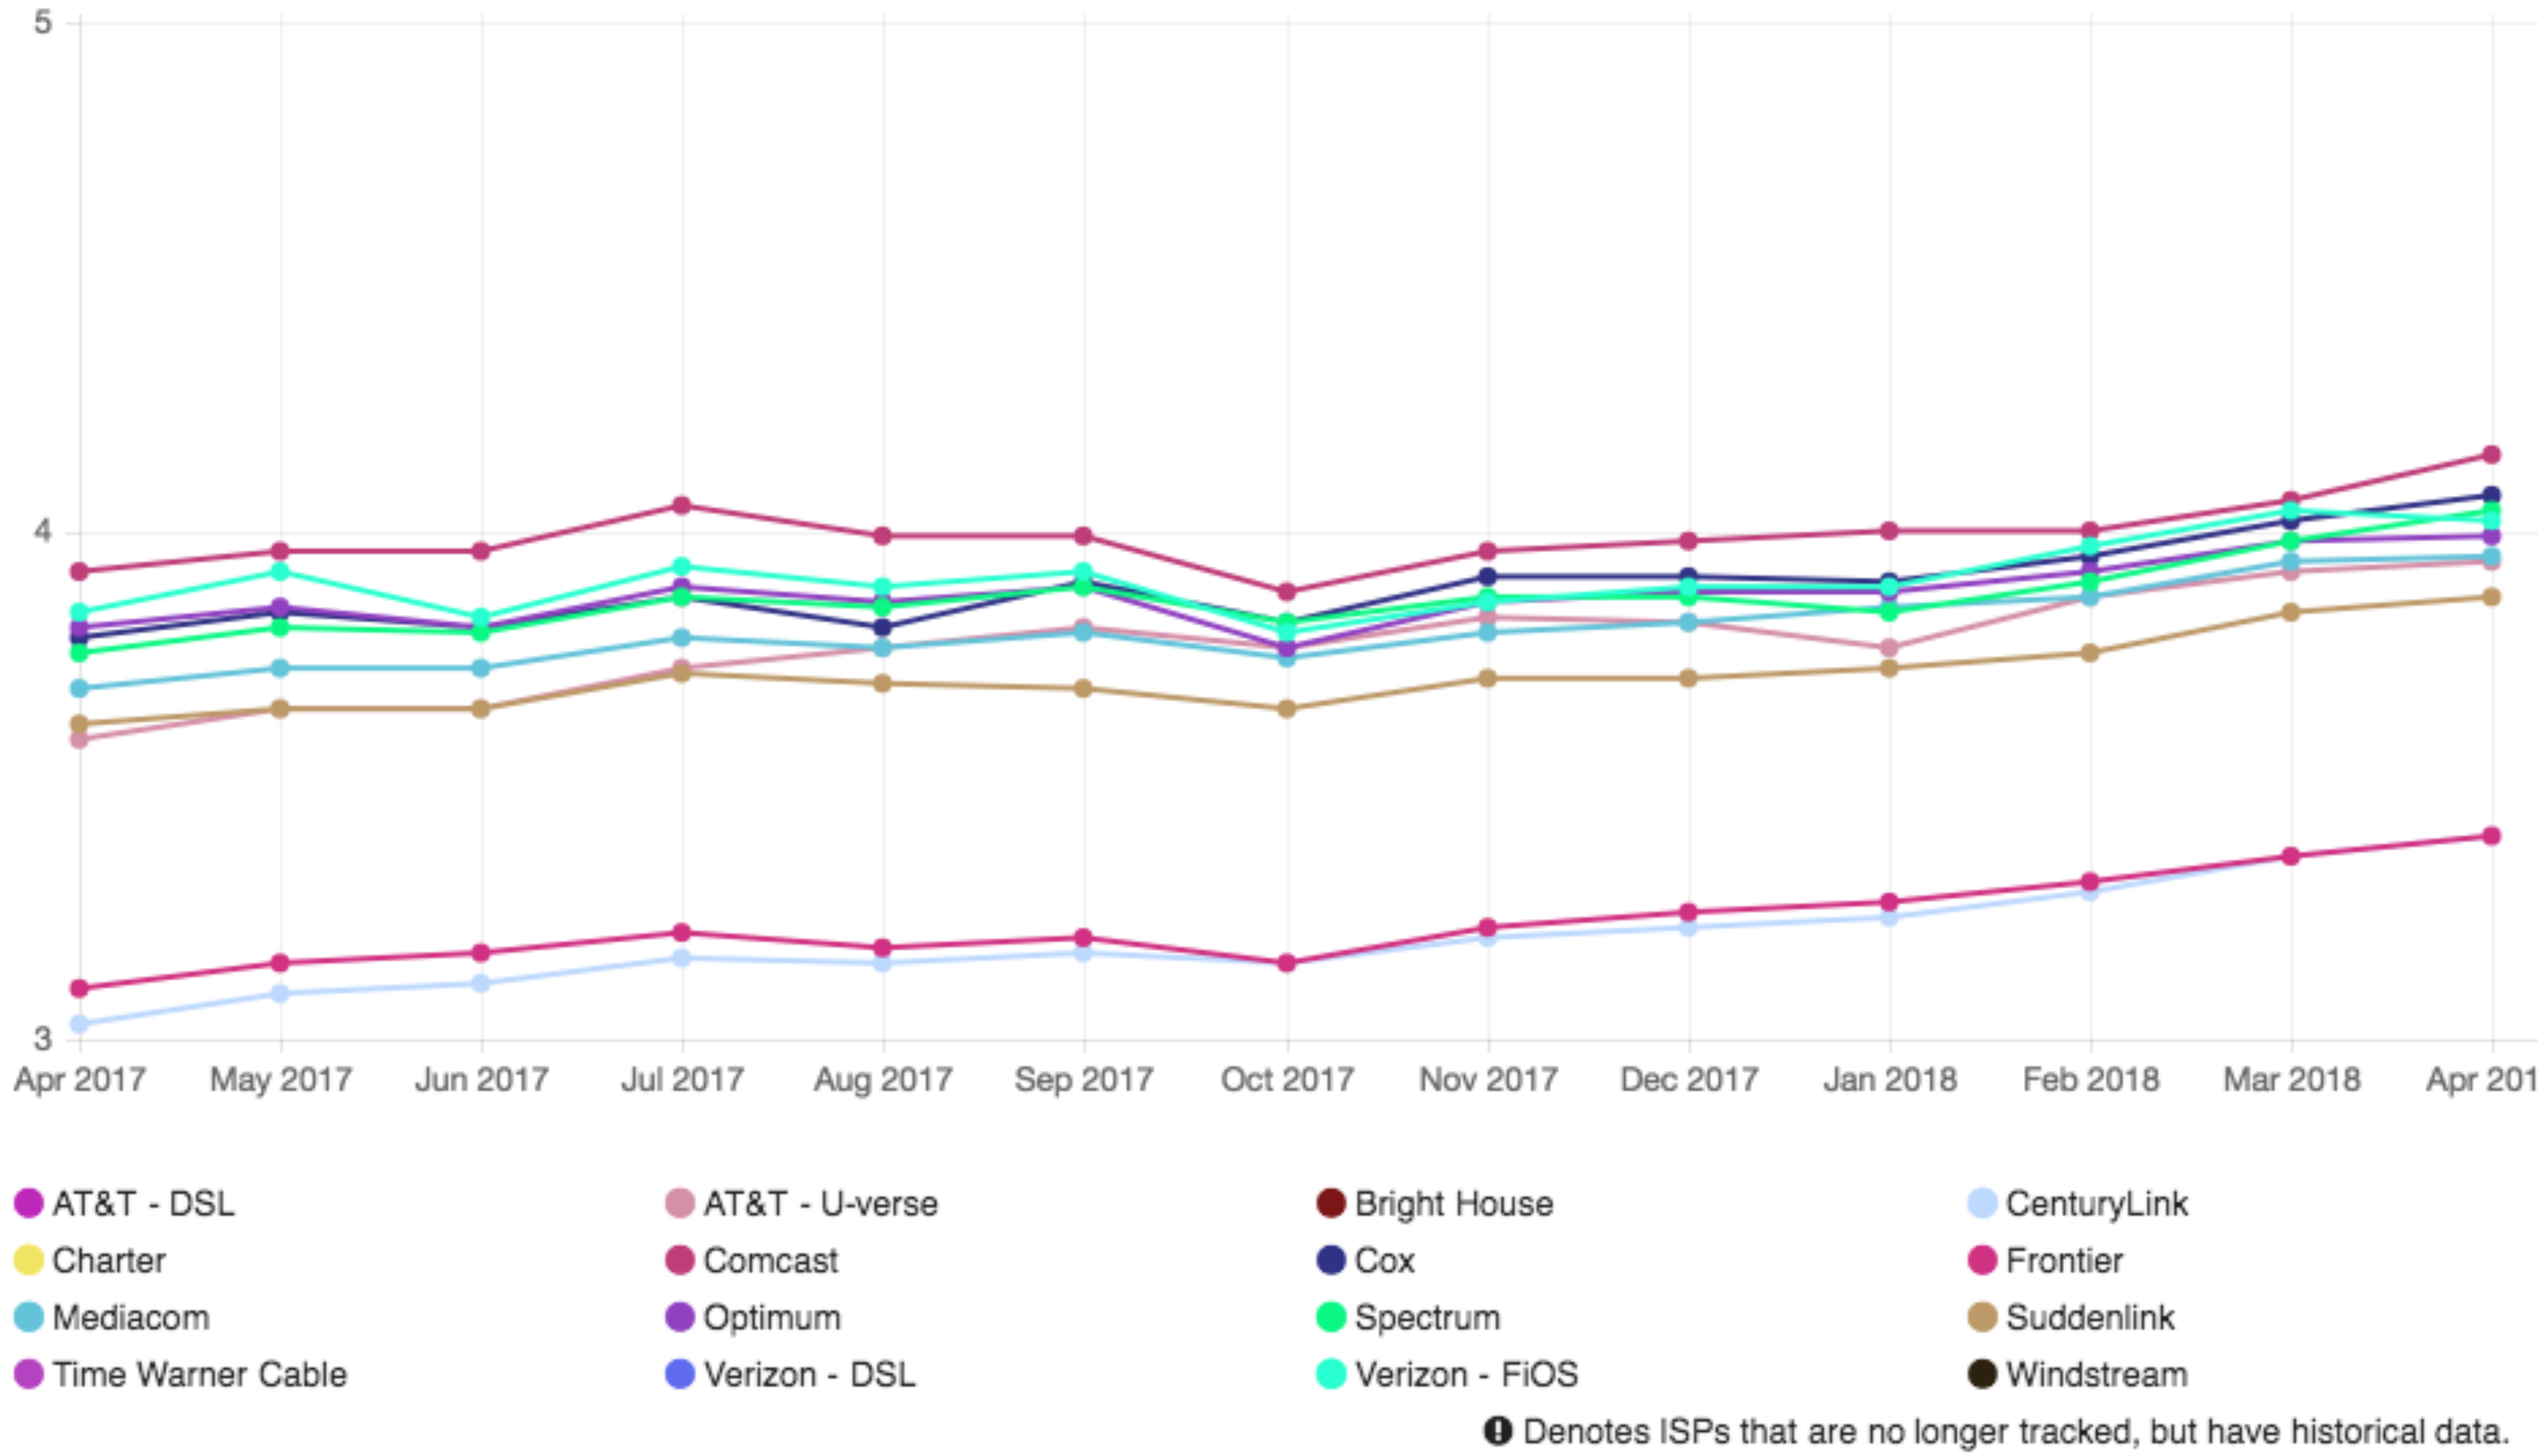

source: https://ispspeedindex.netflix.com/country/us/

on-demand vs. live-streaming

- challenges for live-streaming
  - client: to shorten delays (fewer buffering) still avoiding rebuffering
  - server: large-scale bursty requests

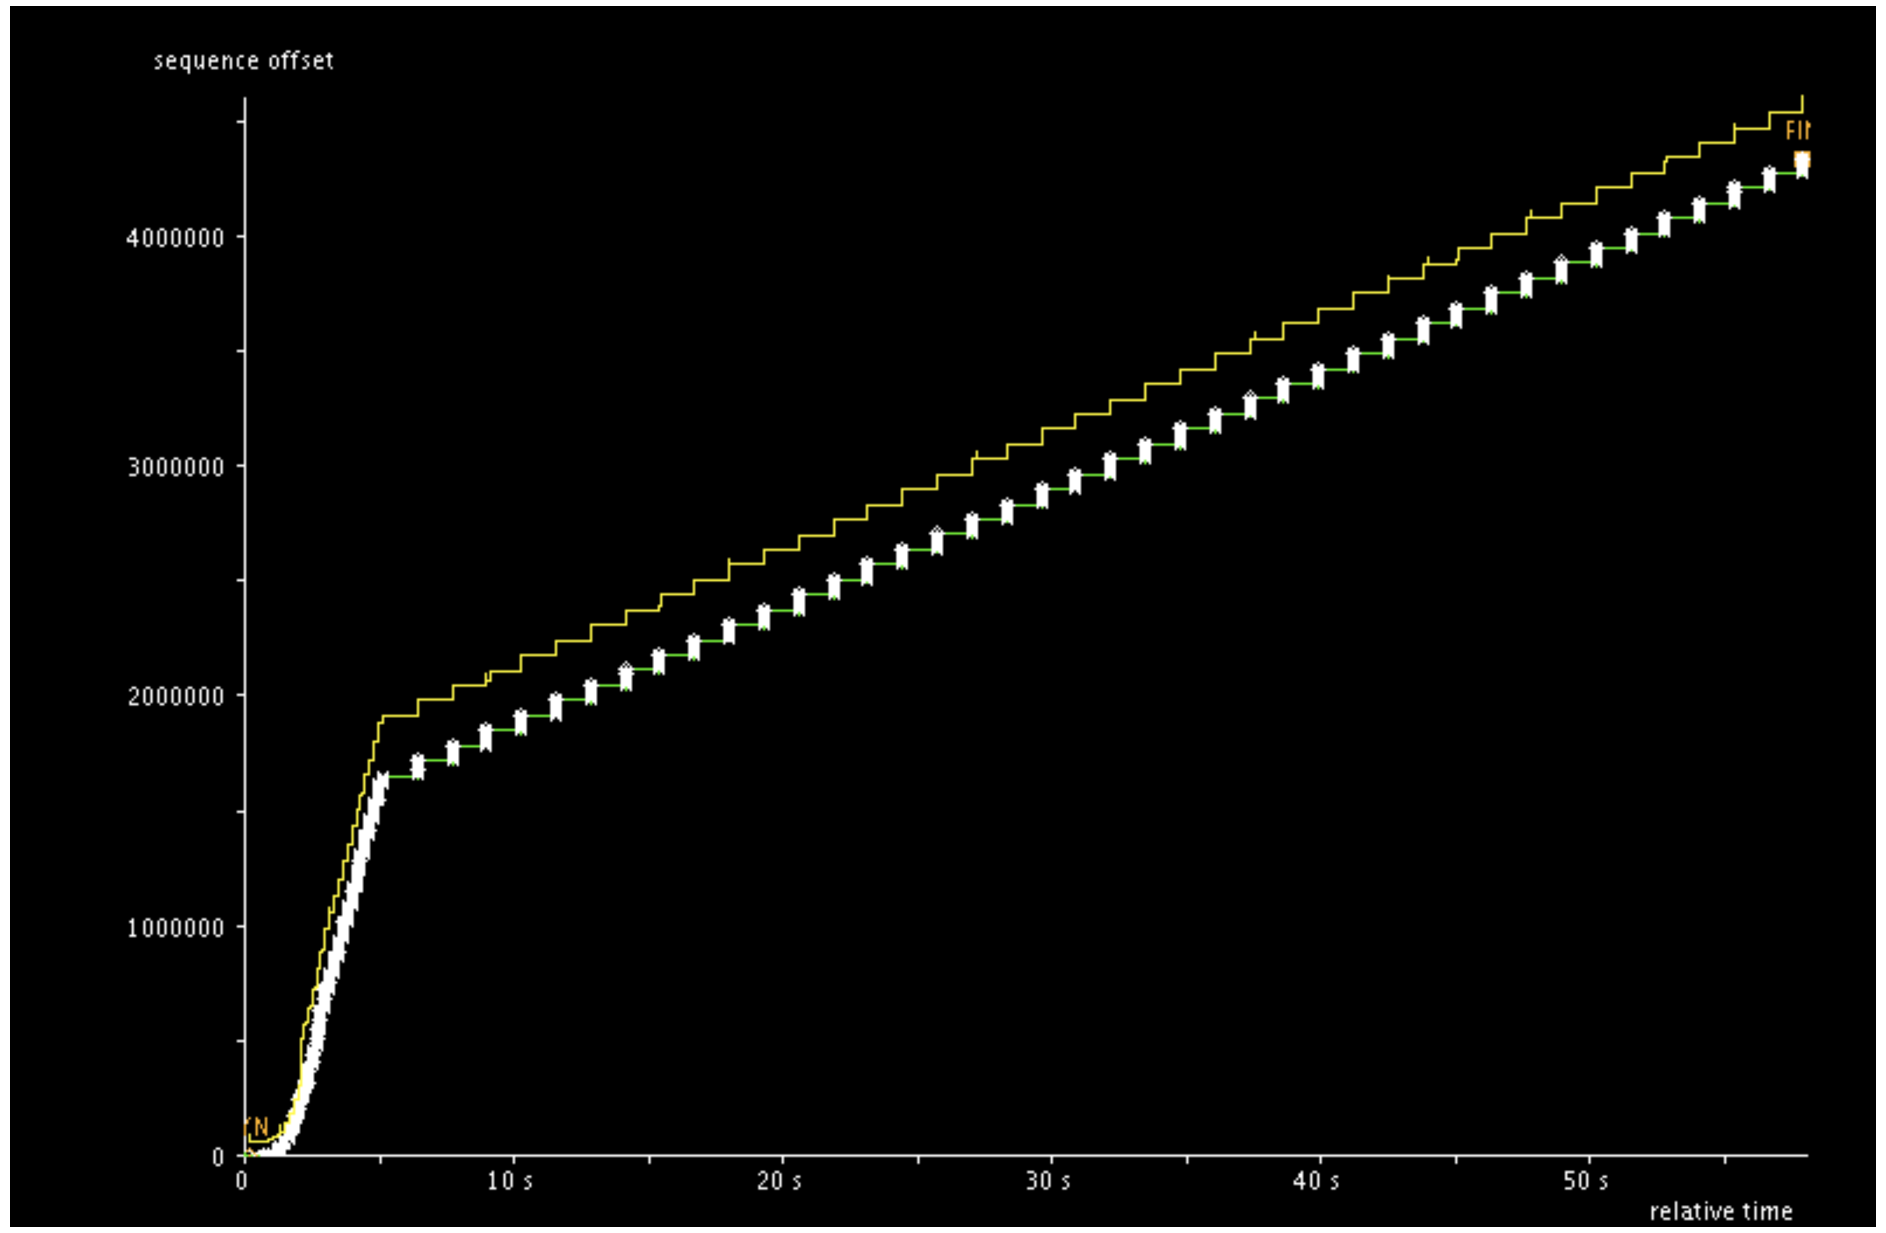

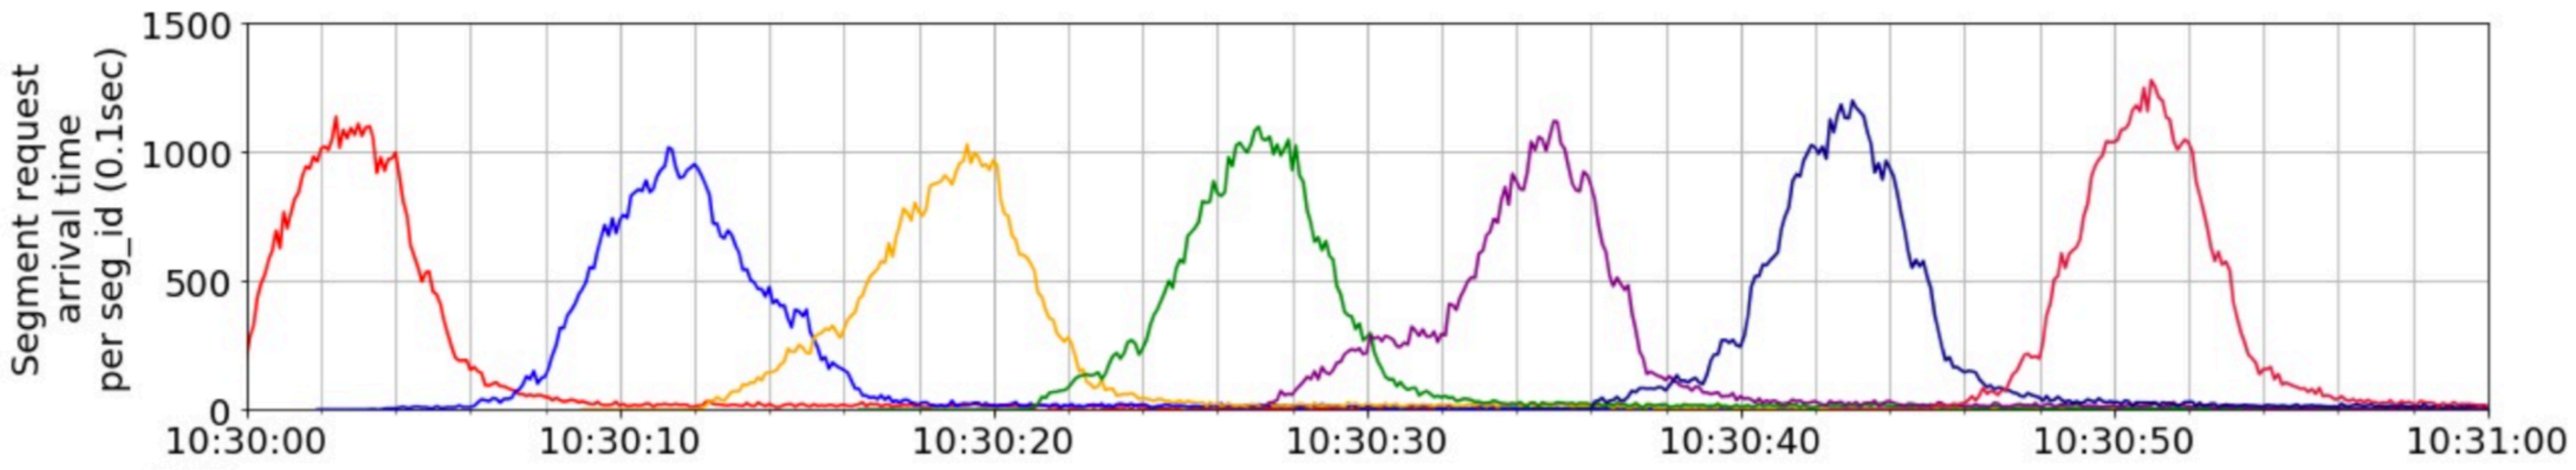

video segment requests arriving at a server

S. Alcock, R. Nelson. Analysis of YouTube Application Flow Control. 2011.
http://www.wand.net.nz/~salcock/youtube/

summary

- steady increase in traffic
  - driven by different technical and non-technical factors
  - dominated by videos (short and long)
- technical trends
  - rise of https, QUIC, IPv6
- QoE
  - better access networks, closer content, better algorithms