# Traffic Measurement Activities of the WIDE project

Kenjiro Cho
IIJ Research Lab

# characteristics of WIDE traffic

- WIDE
  - internet research through live network
- WIDE has its own backbone operated by members
  - backbone includes
    - international links
    - IXes
    - root name server
    - various link types up to 10GbE
  - carrying both commodity traffic and experiments
    - commodity: university traffic, WIDE members
    - experiments: new products, our technologies under development
      - IPv6 everywhere
    - events (including firedrills)
  - not a typical internet but a showcase

# traffic measurement and analysis in WIDE

- measurement activities across research groups
- broad perspectives
  - tracking long-term trends
  - analysis (with wide range of granularity)
  - operational tools (trouble-detection/shooting)
  - evaluation of new technologies
- emphasis on
  - wide-area
  - multi-point
    - measurement on backbone
  - long-term
    - continuation by group effort

# international collaboration

- CAIDA (the Cooperative Association for Internet Data Analysis)
  - collaboration since 2003 on DNS, topology, routing measurement
  - WIDE/CAIDA measurement workshops were held 5 times
- University of Waikato
  - development of the scamper tool for topology measurement
- RSSAC (ICANN Root Server Systems Advisory Committee)
  - root name server measurement
  - WIDE, CAIDA, ISC OARC, USC/ISI
- other collaboration
  - routeviews, RIPE, INRIA, AIT

# recent activities

- IPv6 AS core map
  - collaboration with CAIDA
- residential broadband traffic analysis
  - with 7 major Japanese ISPs and government
- dual-stack path analysis
  - identify IPv6 network problems by comparative path analysis
- DNS measurement
  - root-server measurement
  - tracking trends (e.g., EDNS0, DNSSEC, AAAA)

# IPv6 AS CORE MAP, collaboration with CAIDA

- visualize the outdegree of ASes, locations are mapped to longitude

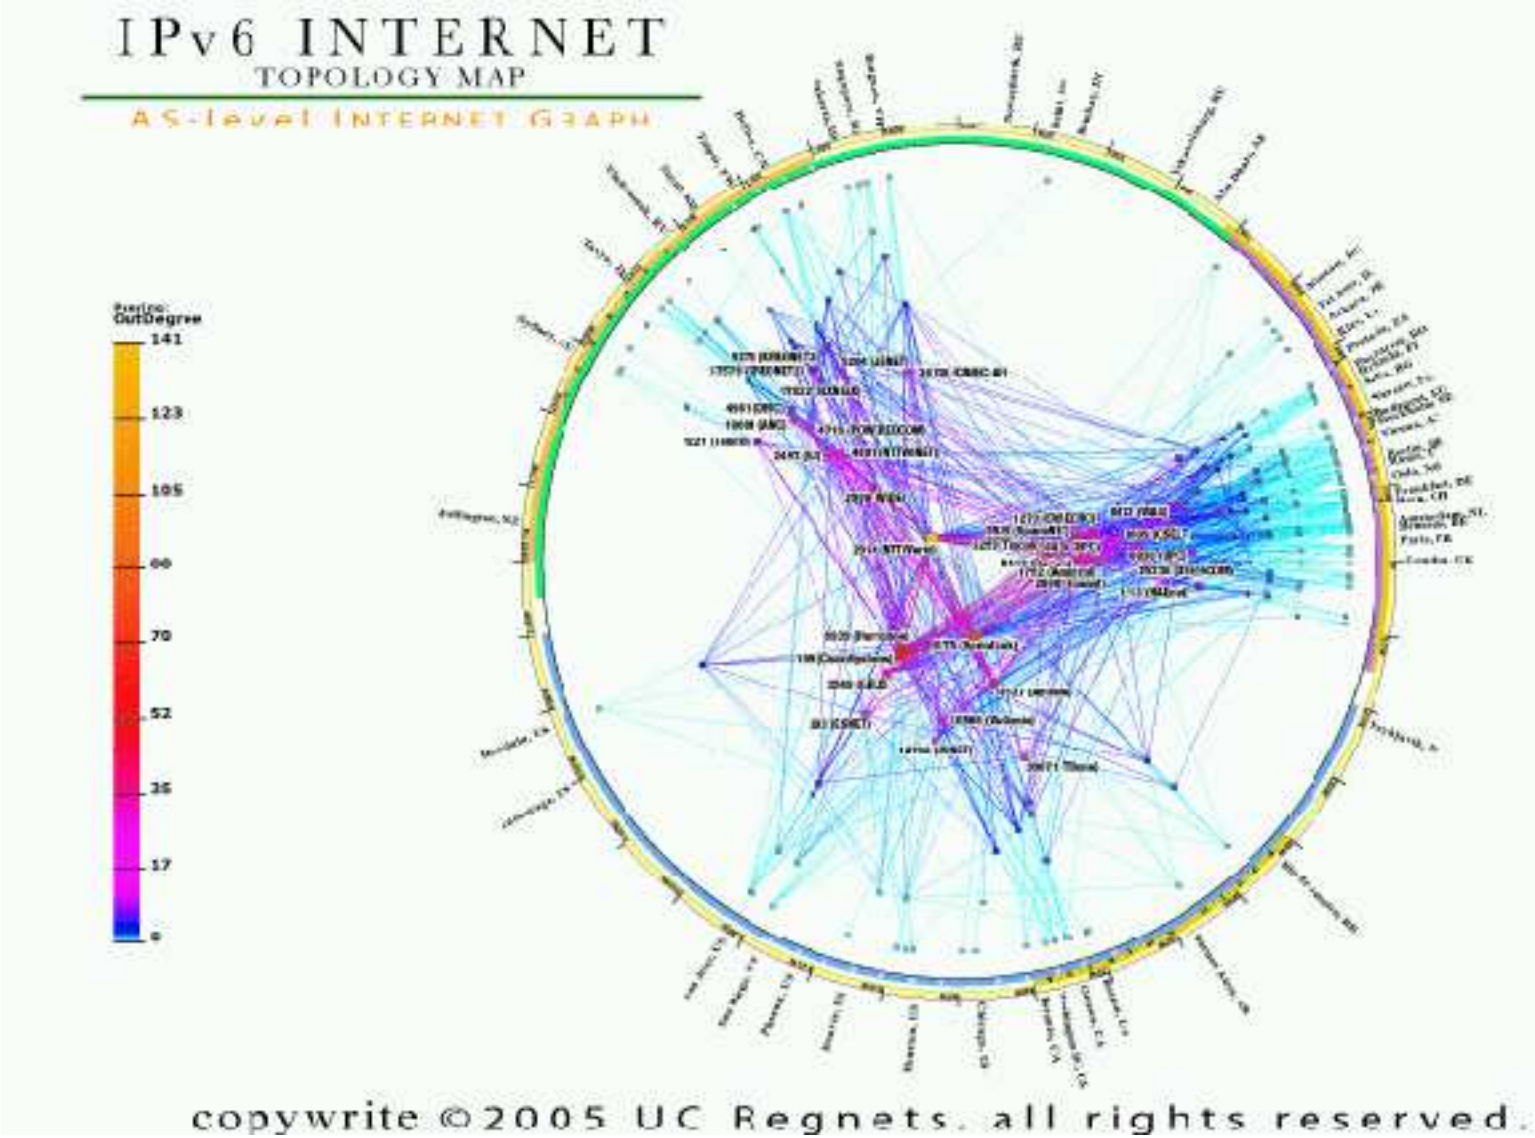

# IPv4 vs IPv6

- IPv6 AS graph is much sparser, less US-centric

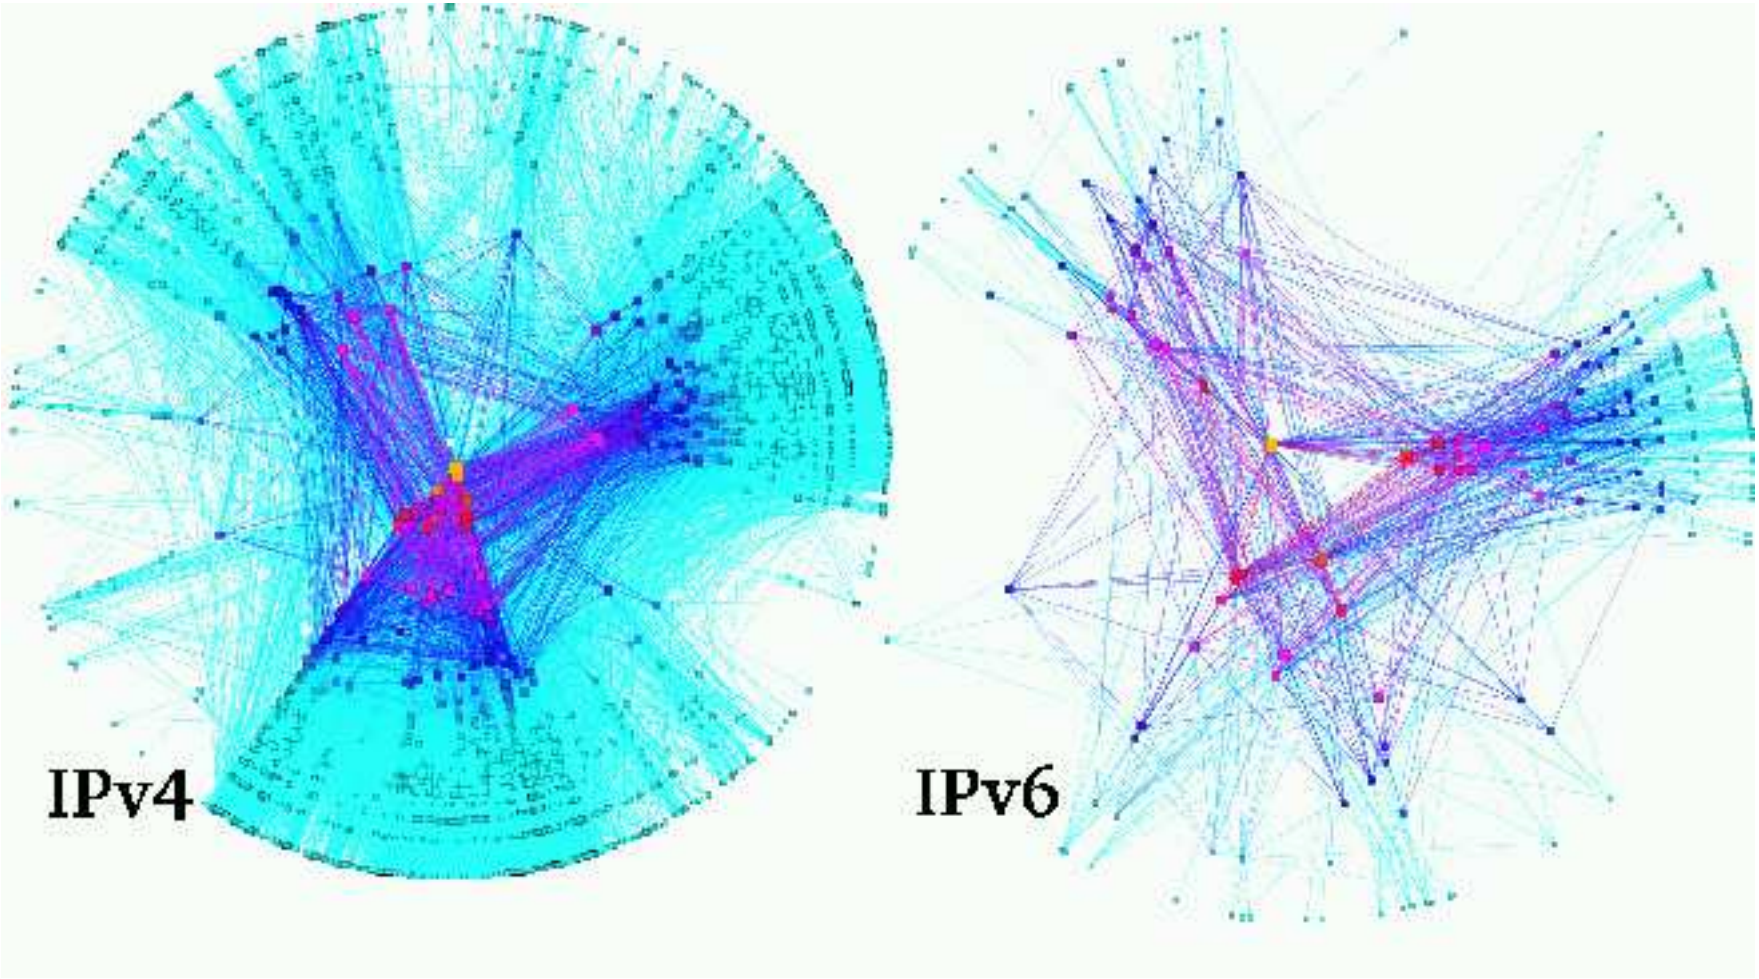

# residential broadband traffic analysis

- 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

- to identify the macro-level impact of RBB traffic on ISP backbones
  - a study group was formed with specialists
    - members from 7 major Japanese ISPs and government
- goals: traffic measurement across multiple ISPs, to identify
  - ratio of RBB traffic to other traffic
  - changes in traffic patterns
  - regional differences

# major findings in residential broadband traffic

- our data is considered to cover 41% of total Japanese traffic
  - total RBB traffic in Japan is estimated to be 470Gbps (2005/05)
- 70% of RBB traffic is constant, peak in the evening hours
- RBB traffic is much larger than office traffic, so backbone traffic is dominated by RBB traffic
- traffic volume exchanged via private peering is comparable with volume exchanged via major IXes
- regional RBB traffic is roughly proportional to regional population

# data collection across major 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

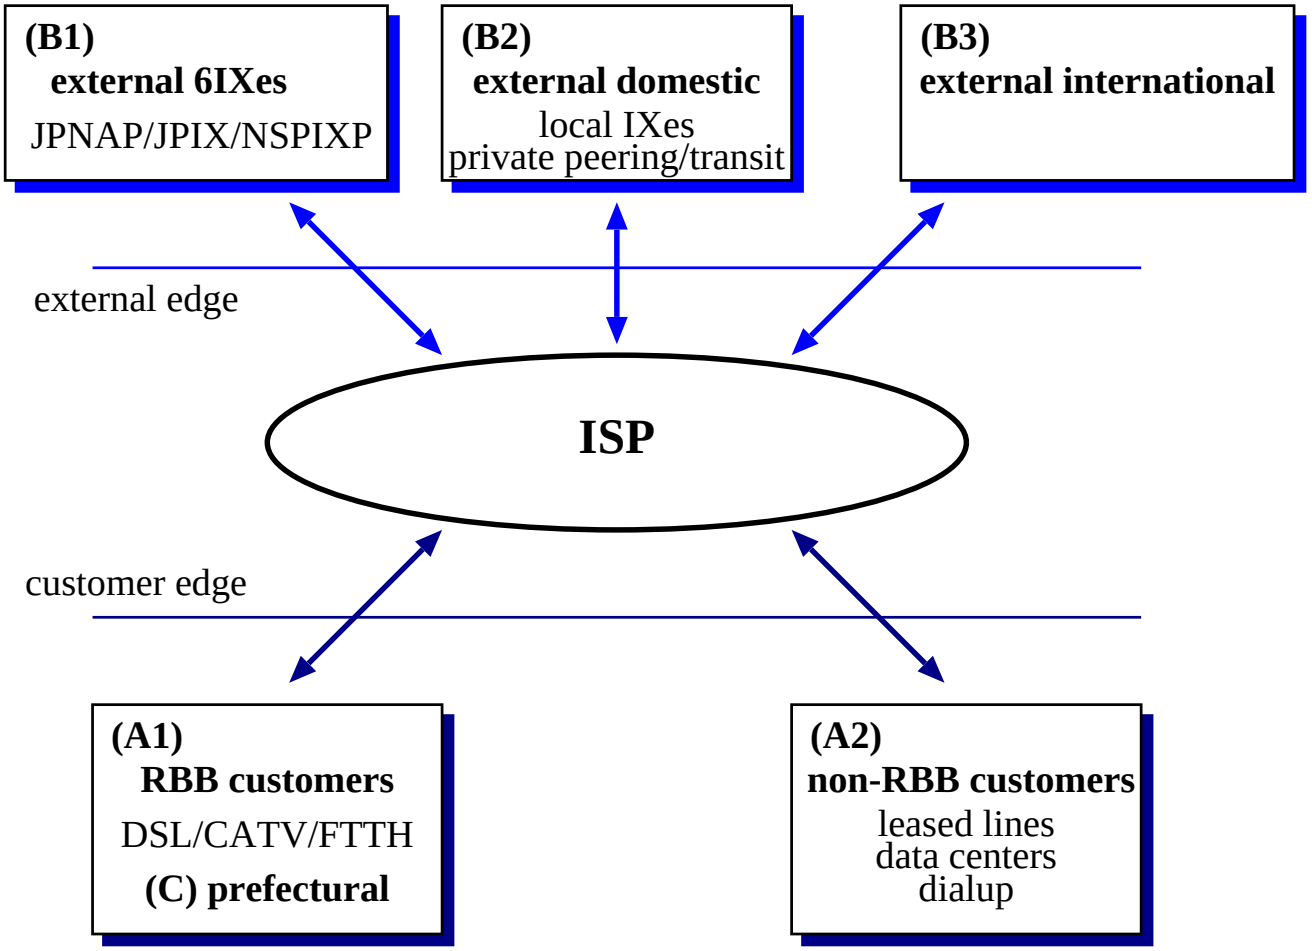

# RBB customer weekly traffic
## in May 2005

- DSL/CATV/FTTH customer traffic of the 7 ISPs
  - inbound and outbound are almost equal
  - almost 200Gbps on average!
  - 150Gbps is constant, probably due to p2p applications
  - daily fluctuations: peak from 21:00 to 23:00

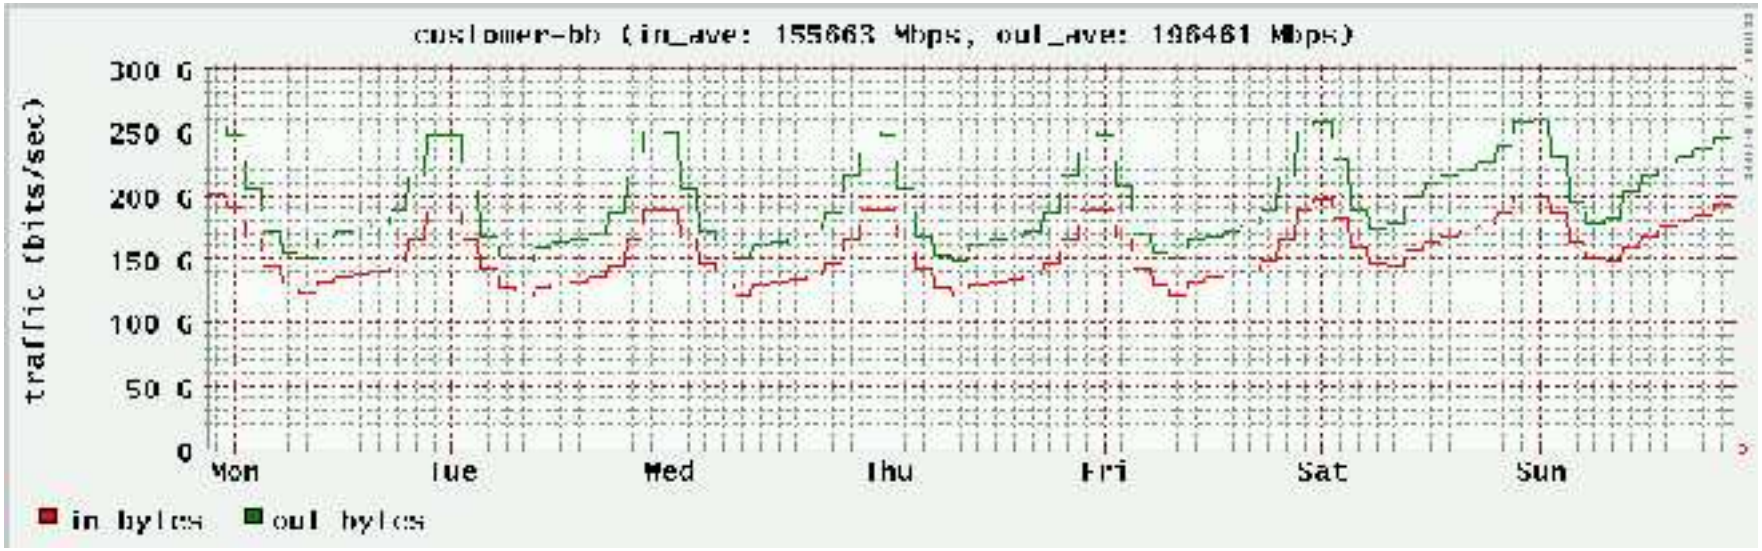

# non-RBB customer weekly traffic
## in May 2005

- leased lines/data center/other customers
  - home user traffic is still dominant (by peak hours)
    - because leased lines include 2nd/3rd level ISPs
  - larger office hour traffic than RBB customer traffic
- only 4 ISPs provided data for this group
  - some ISPs have too many routers, historycally mixed up settings

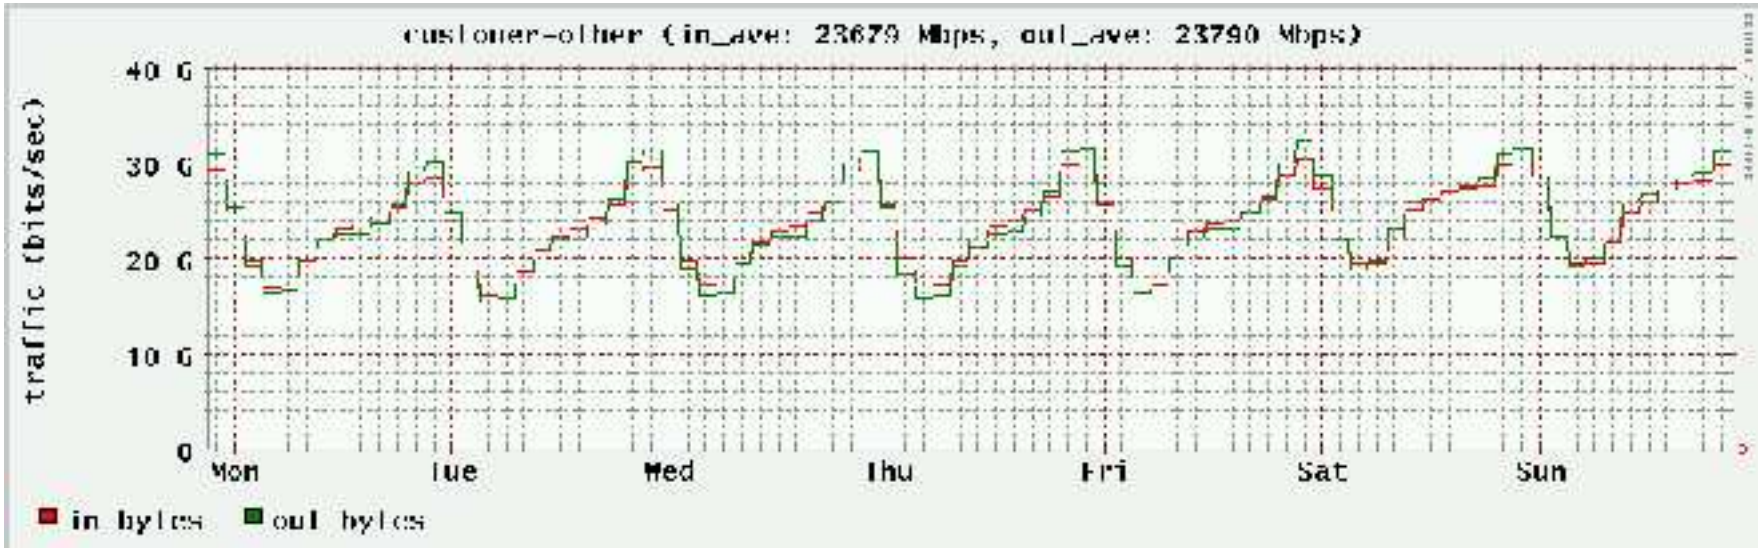

# weekly external traffic to/from 6 major IXes
## in May 2005

- IX traffic is also strongly affected by residential traffic

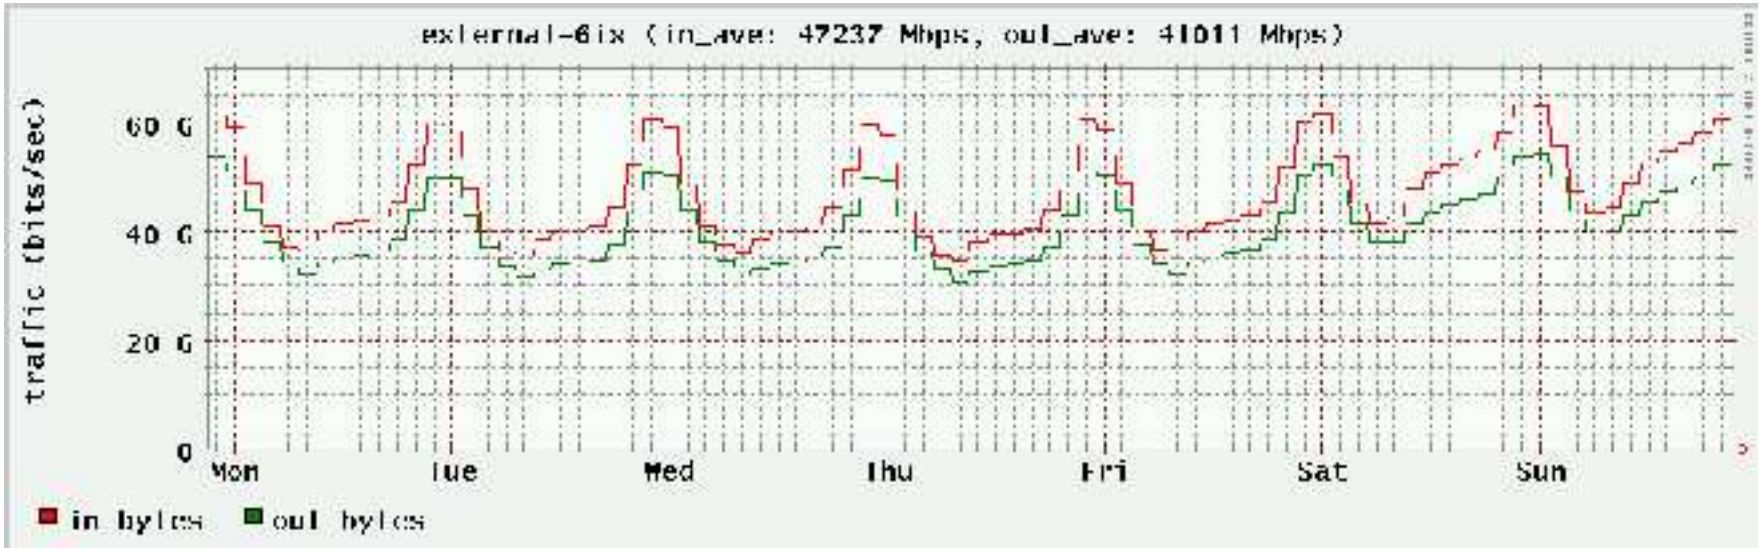

# weekly other domestic external traffic
## in May 2005

- private peering/transit, regional IXes (mainly private peering)
  - traffic volume and pattern are similar to IX traffic

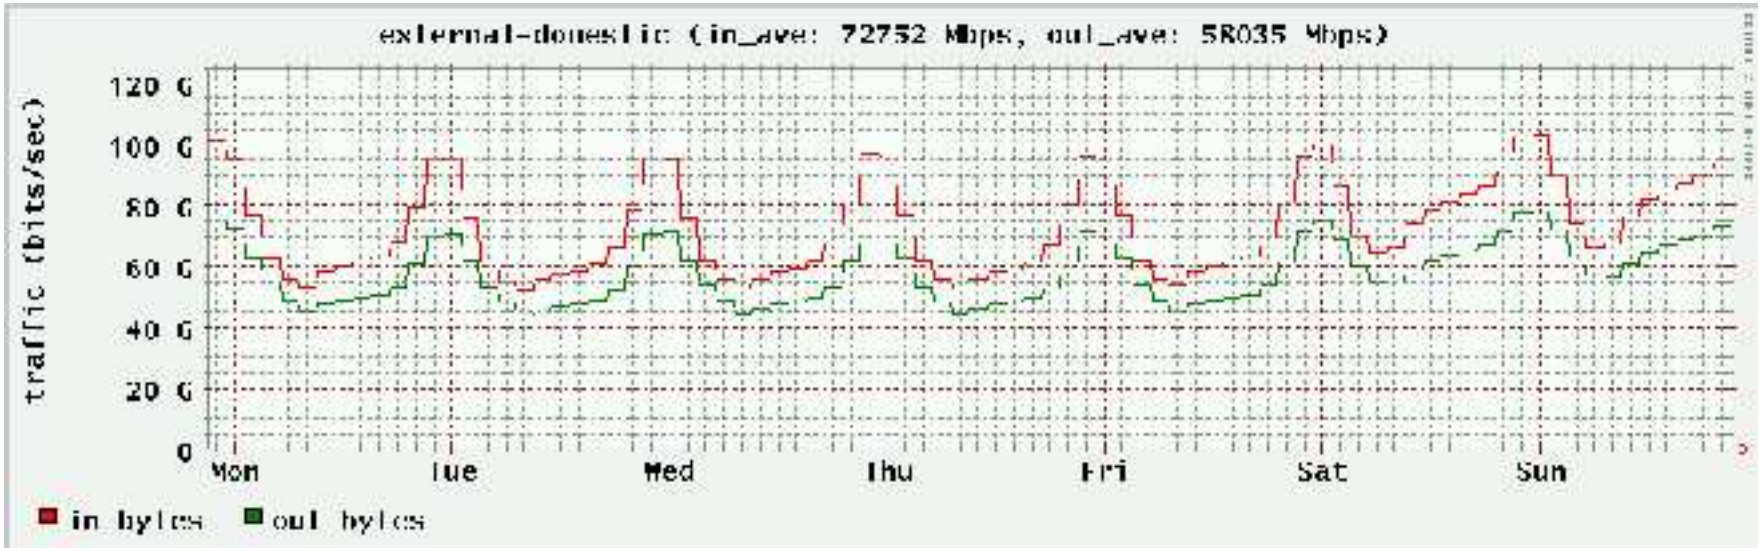

# weekly international external traffic
## in May 2005

- international traffic
  - inbound much larger than outbound
    - traditional content downloading seems still dominant

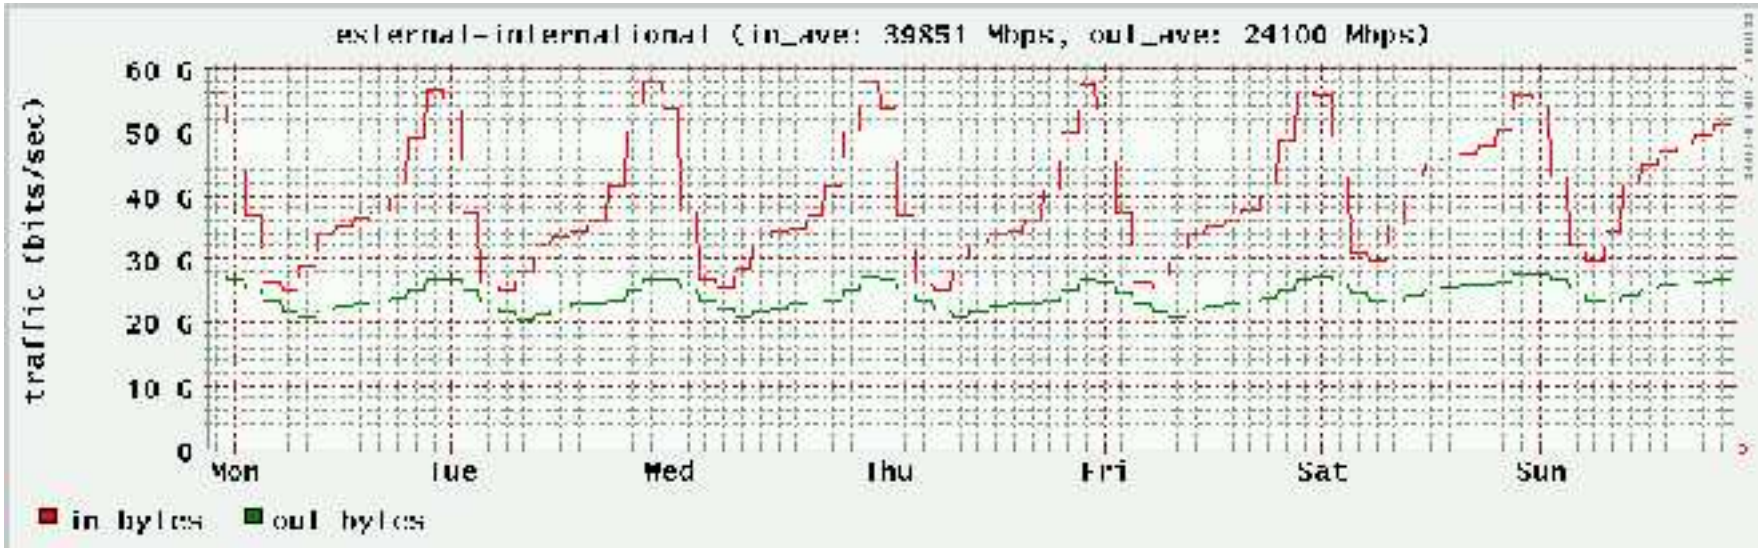

# traffic growth

- 25-50% increase in 6 months!

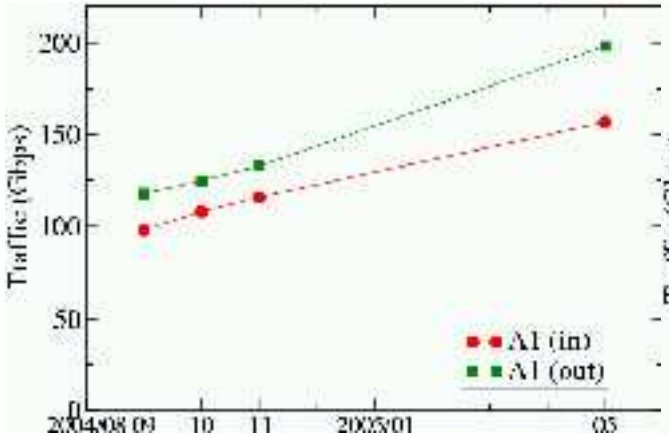


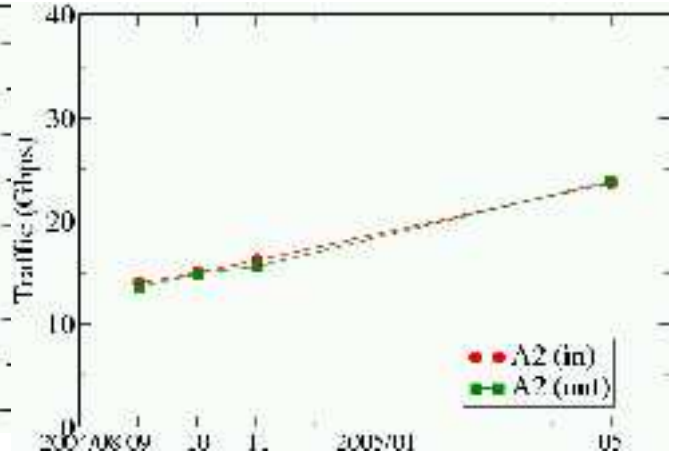


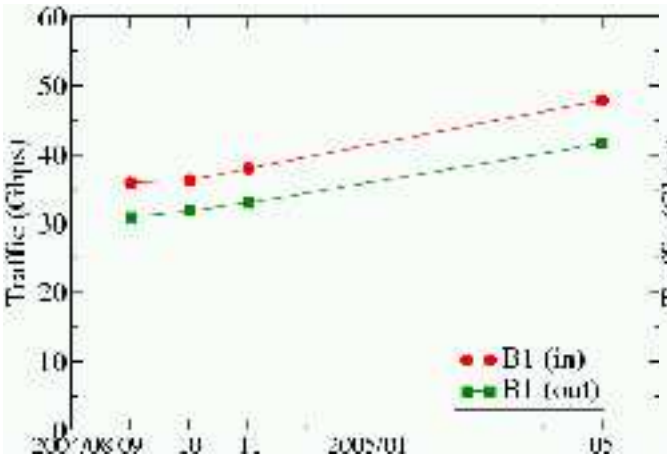


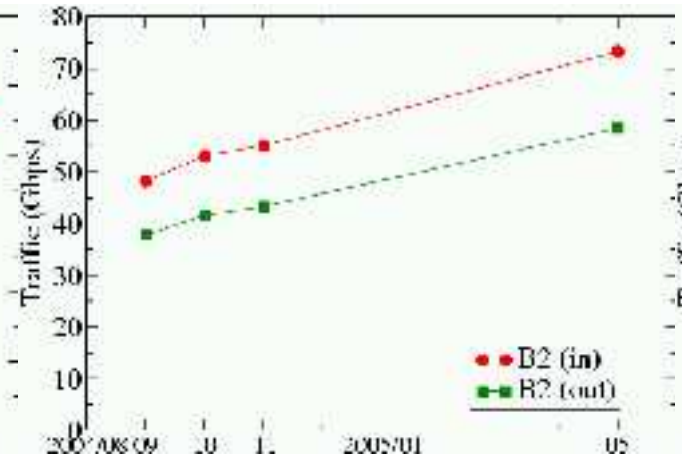


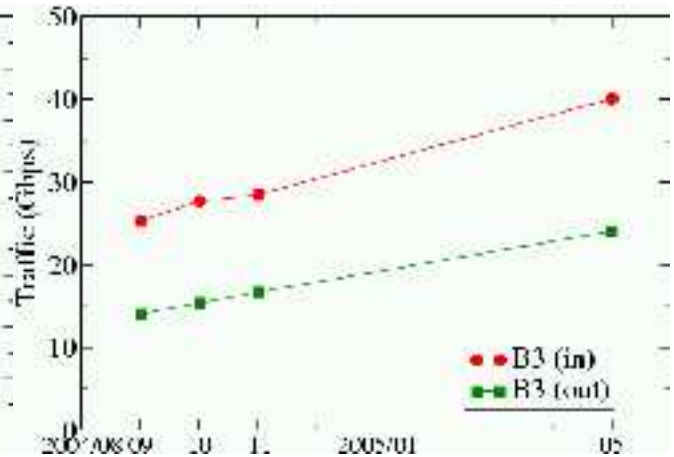

# prefectural population and traffic

- a scatter plot of population and traffic volume
  - traffic is roughly linear to population!
  - similar result with the number of Internet users

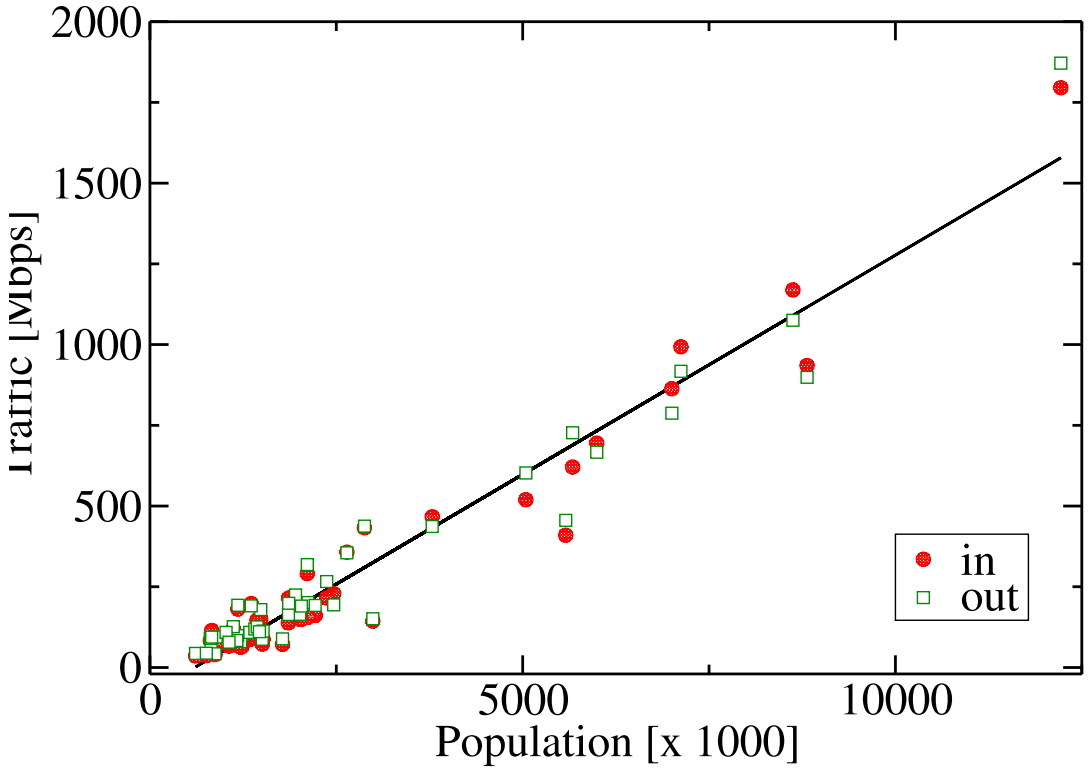

# distribution of per-customer traffic in one ISP

- one of the ISPs provided per-customer traffic info for October 2004
  - by sampled NetFlow and matching customer ID with assigned IP addresses
- we used average daily traffic volume per customer for analysis
- results are consistent with the aggregated traffic

# cumulative distribution of daily traffic per user

- complementary cumulative distribution on a log-log scale
  - distribution similar in all prefectures, differences only in tail length
  - knee point: 4% of customers use more than 2.5GB/day (230kb/s)
  - outbound is dominant for most customers but not for heavy hitters
- even heavyhitters follow the distribution, are not exceptional anymore

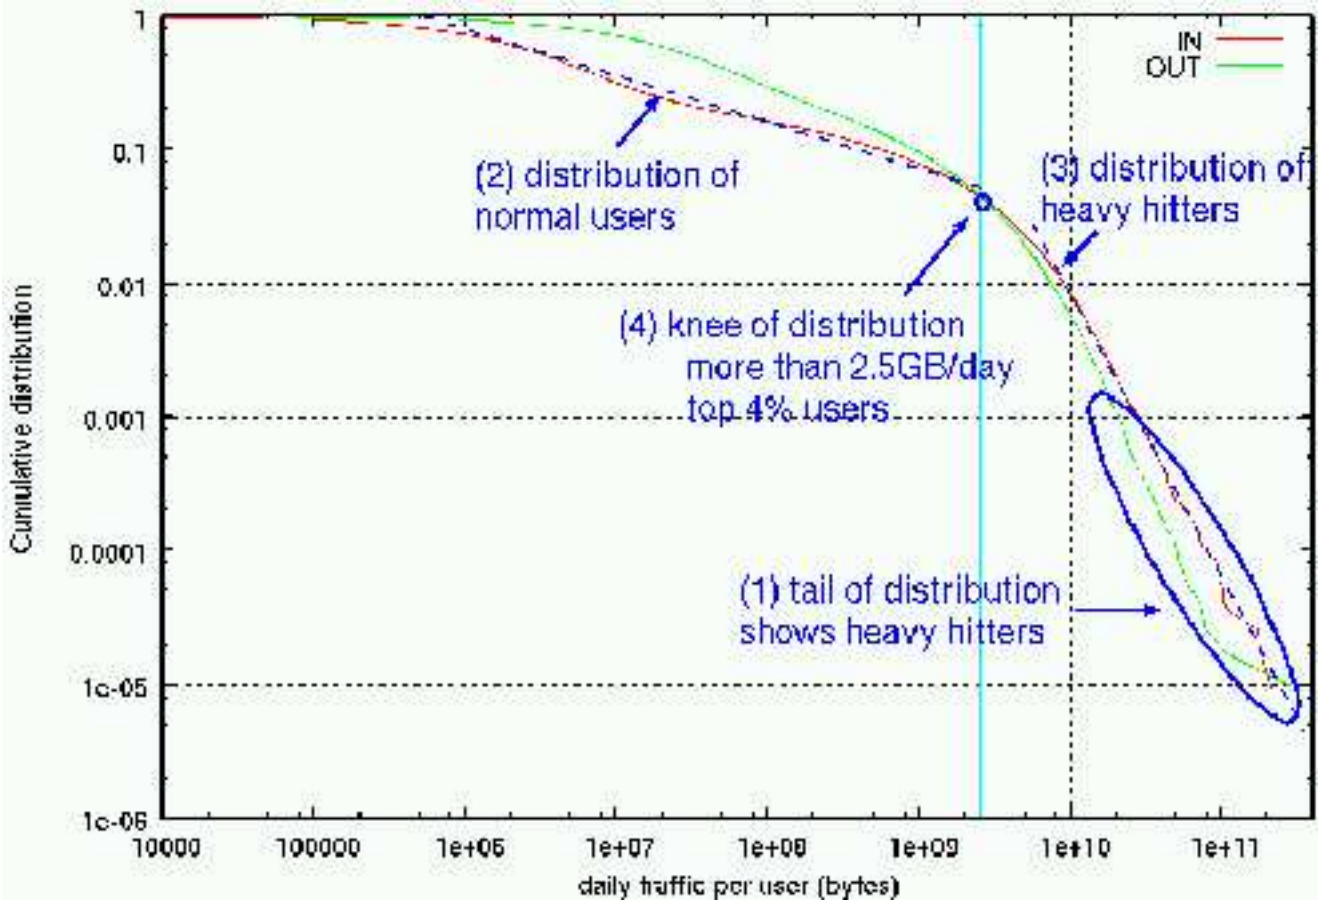

# correlation of inbound and outbound per customer

- high density cluster below and parallel to the unity line
  - outbound is 10 times larger than inbound
- in higher volume region, another cluster around the unity line
  - seems like p2p file-sharing

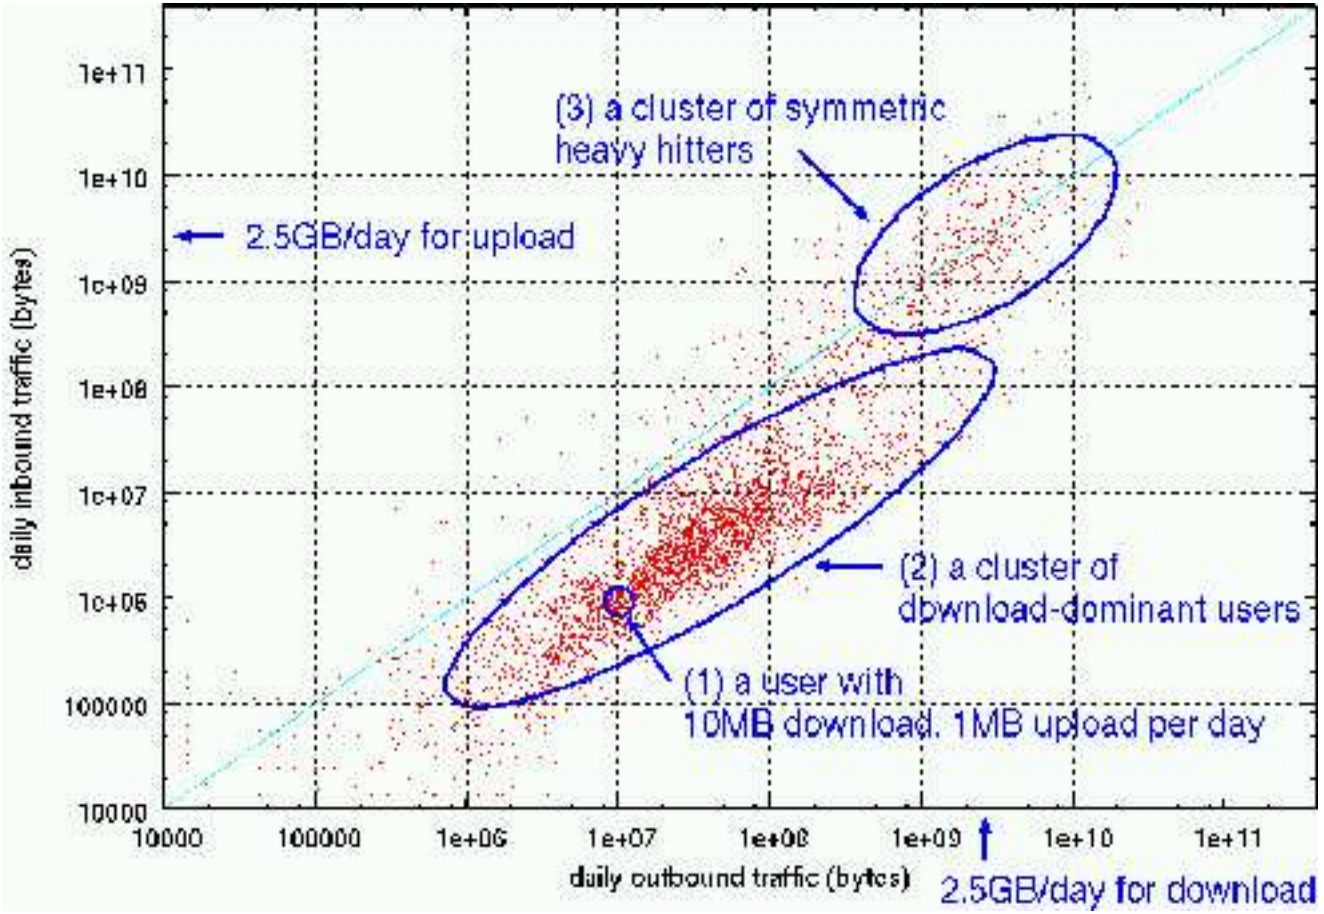

# prefectural traffic distribution

- using geo-IP database from Cyber Area Research Inc
  - only covers ISP users (not including data-centers, leased lines)
  - src on Y-axis, dst on X-axis
    - ALL, 47 prefectures, other DOMestic, INTernational, INValid
- prefectures with larger population have larger share
- user-to-user (p2p) is 63%, 90% is domestic

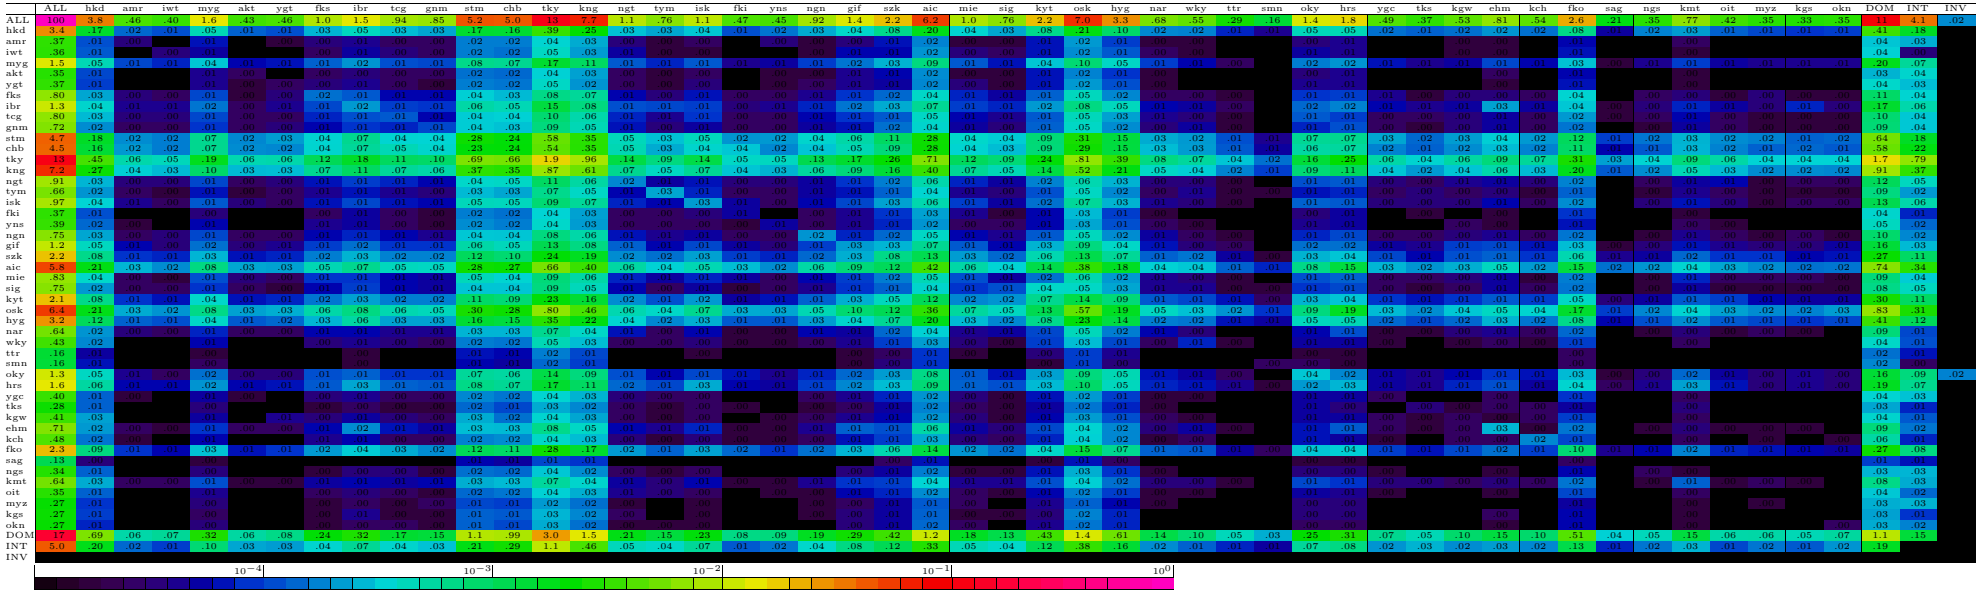

# prefectural traffic distribution normalized to src/dst

- similar distribution, only small locality (1-3%) is found

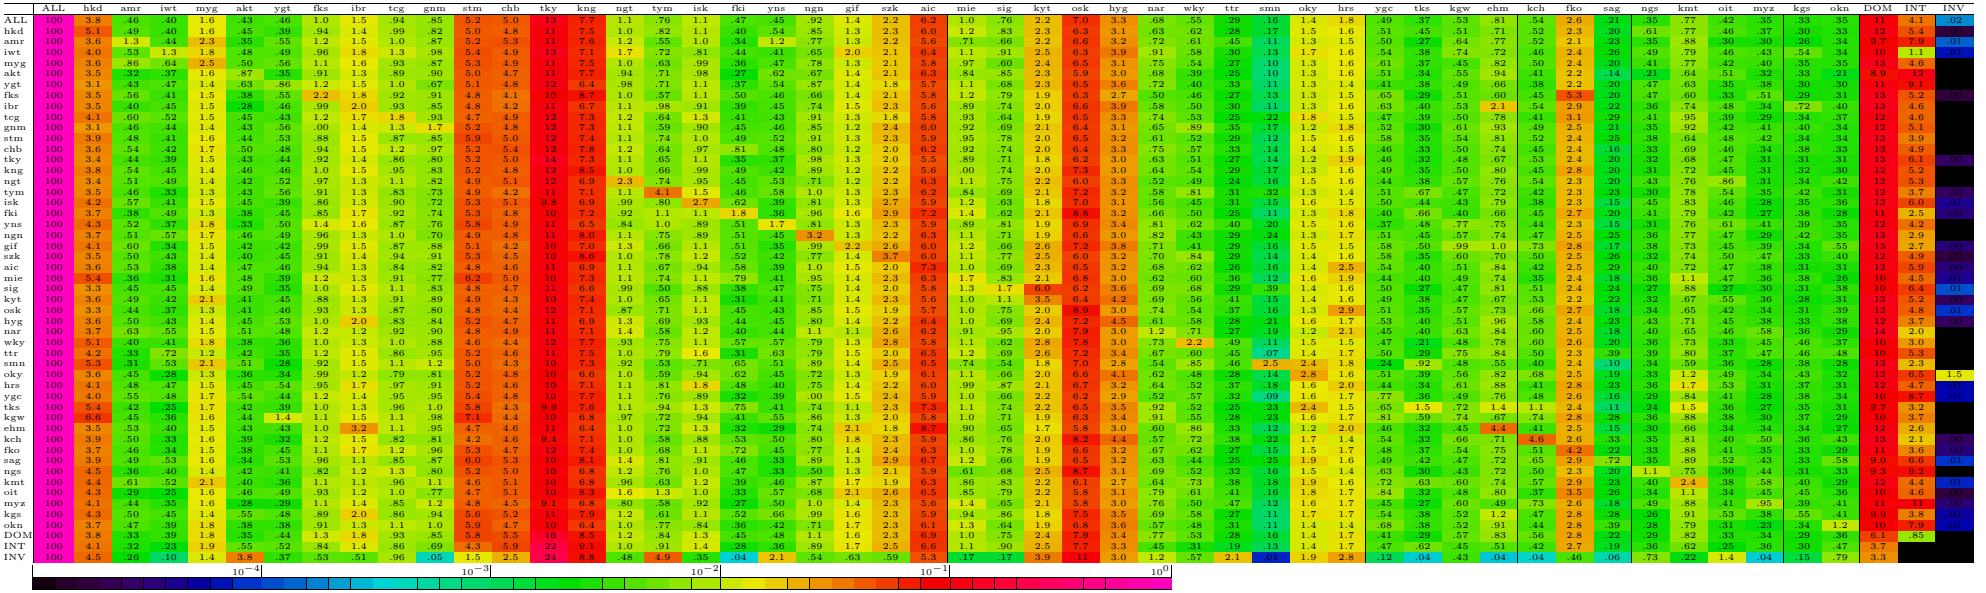

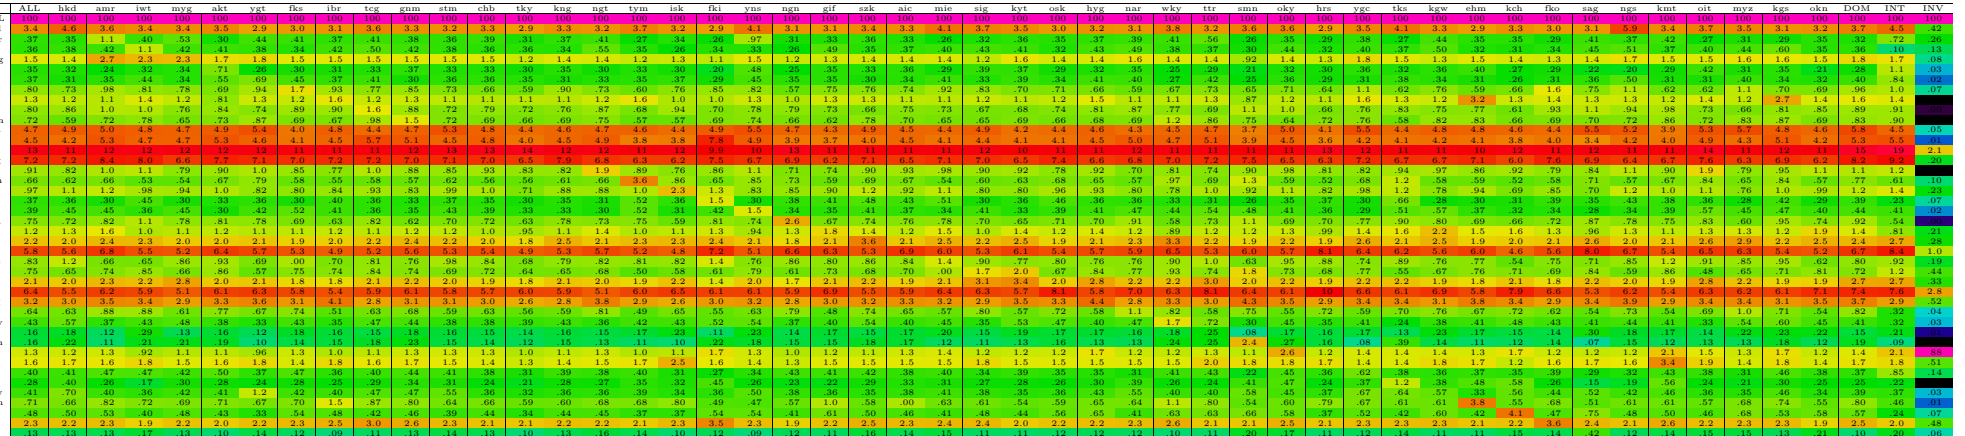

# protocols/ports ranking

- 80% is TCP dynamic ports

| protocol | ratio(%) | port # | name | ratio(%) |
|---|---|---|---|---|
| TCP | 97.43 | | | |
| ($port < 1024$ | 13.99) | 80 | http | 9.32 |
| | | 20 | ftp-data | 0.93 |
| | | 554 | rtsp | 0.38 |
| | | 443 | https | 0.30 |
| | | 110 | pop3 | 0.17 |
| | | 81 | - | 0.15 |
| | | 25 | smtp | 0.14 |
| | | 119 | nntp | 0.13 |
| | | 21 | ftp | 0.11 |
| | | 22 | ssh | 0.09 |
| | | - | other | 2.27 |
| ($port >= 1024$ | 83.44) | 6699 | winmx | 1.40 |
| | | 6346 | gnutella | 0.92 |
| | | 7743 | winny | 0.48 |
| | | 6881 | bittorrent | 0.25 |
| | | 6348 | gnutella | 0.21 |
| | | 1935 | macromedia-fsc | 0.20 |
| | | 1755 | ms-streaming | 0.20 |
| | | 2265 | - | 0.13 |
| | | 1234 | - | 0.12 |
| | | 4662 | edonkey | 0.12 |
| | | 8080 | http-proxy | 0.11 |
| | | - | other | 79.30 |
| UDP | 1.38 | 6346 | gnutella | 0.39 |
| | | 6257 | winmx- | 0.06 |
| | | - | other | 0.93 |
| ESP | 1.09 | | | |
| GRE | 0.07 | | | |
| ICMP | 0.01 | | | |
| OTHER | 0.02 | | | |

# dual-stack path analysis

- ideas
  - develop techniques specifically designed for dual-stack network
    - take measurements for both IPv4 and IPv6 at the same time
    - compare IPv6 results to IPv4 results
    - extract problems which exist only in IPv6
- methodology
  - dual-stack node discovery
    - create a dual-stack node list by monitoring DNS replies
  - dual-stack ping
    - run ping/ping6 to target dual-stack nodes
    - select a few representative nodes per site (/48) by rtt ratios
  - dual-stack traceroute
    - run traceroute/traceroute6 to the selected nodes
    - visualize results for comparative path analysis

# distribution of IPv6/IPv4 RTTs from 3 locations

- individual nodes far above the unity line: leaf issues
- clusters above the unity line: backbone issues

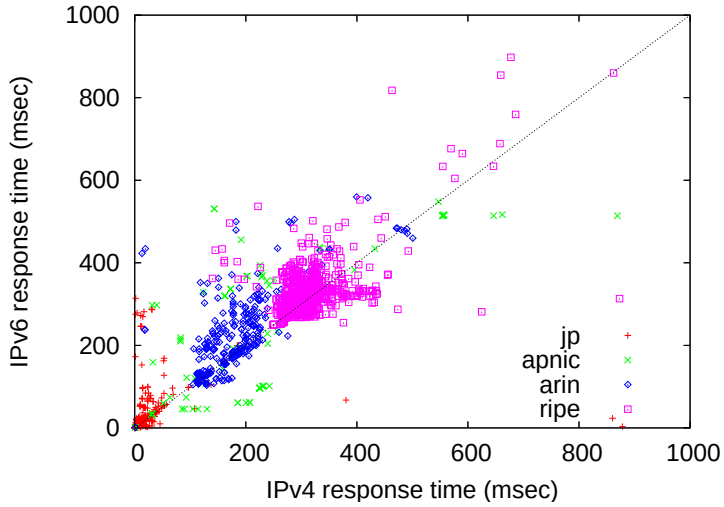


WIDE

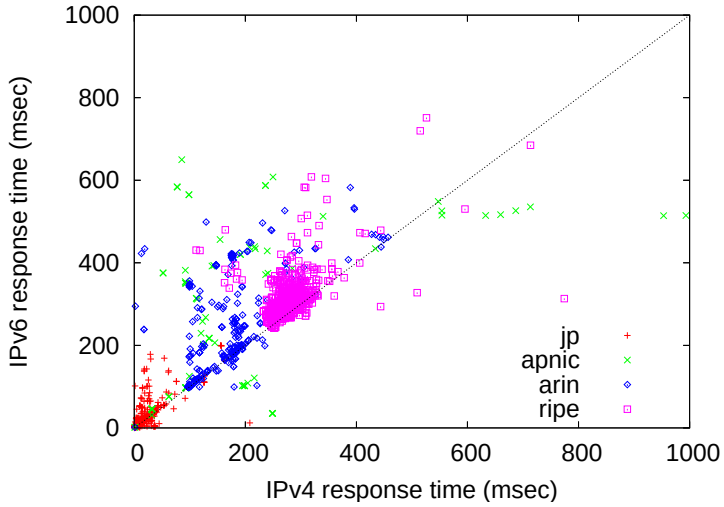


IIJ

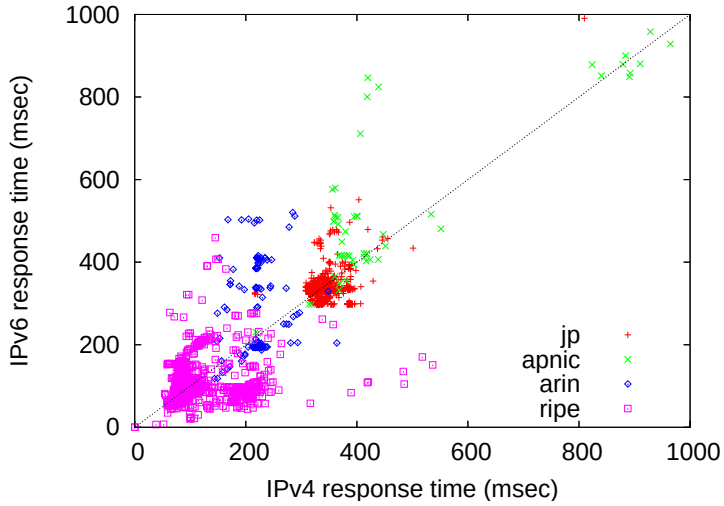


ES

# dual-stack traceroute to ABILENE from WIDE

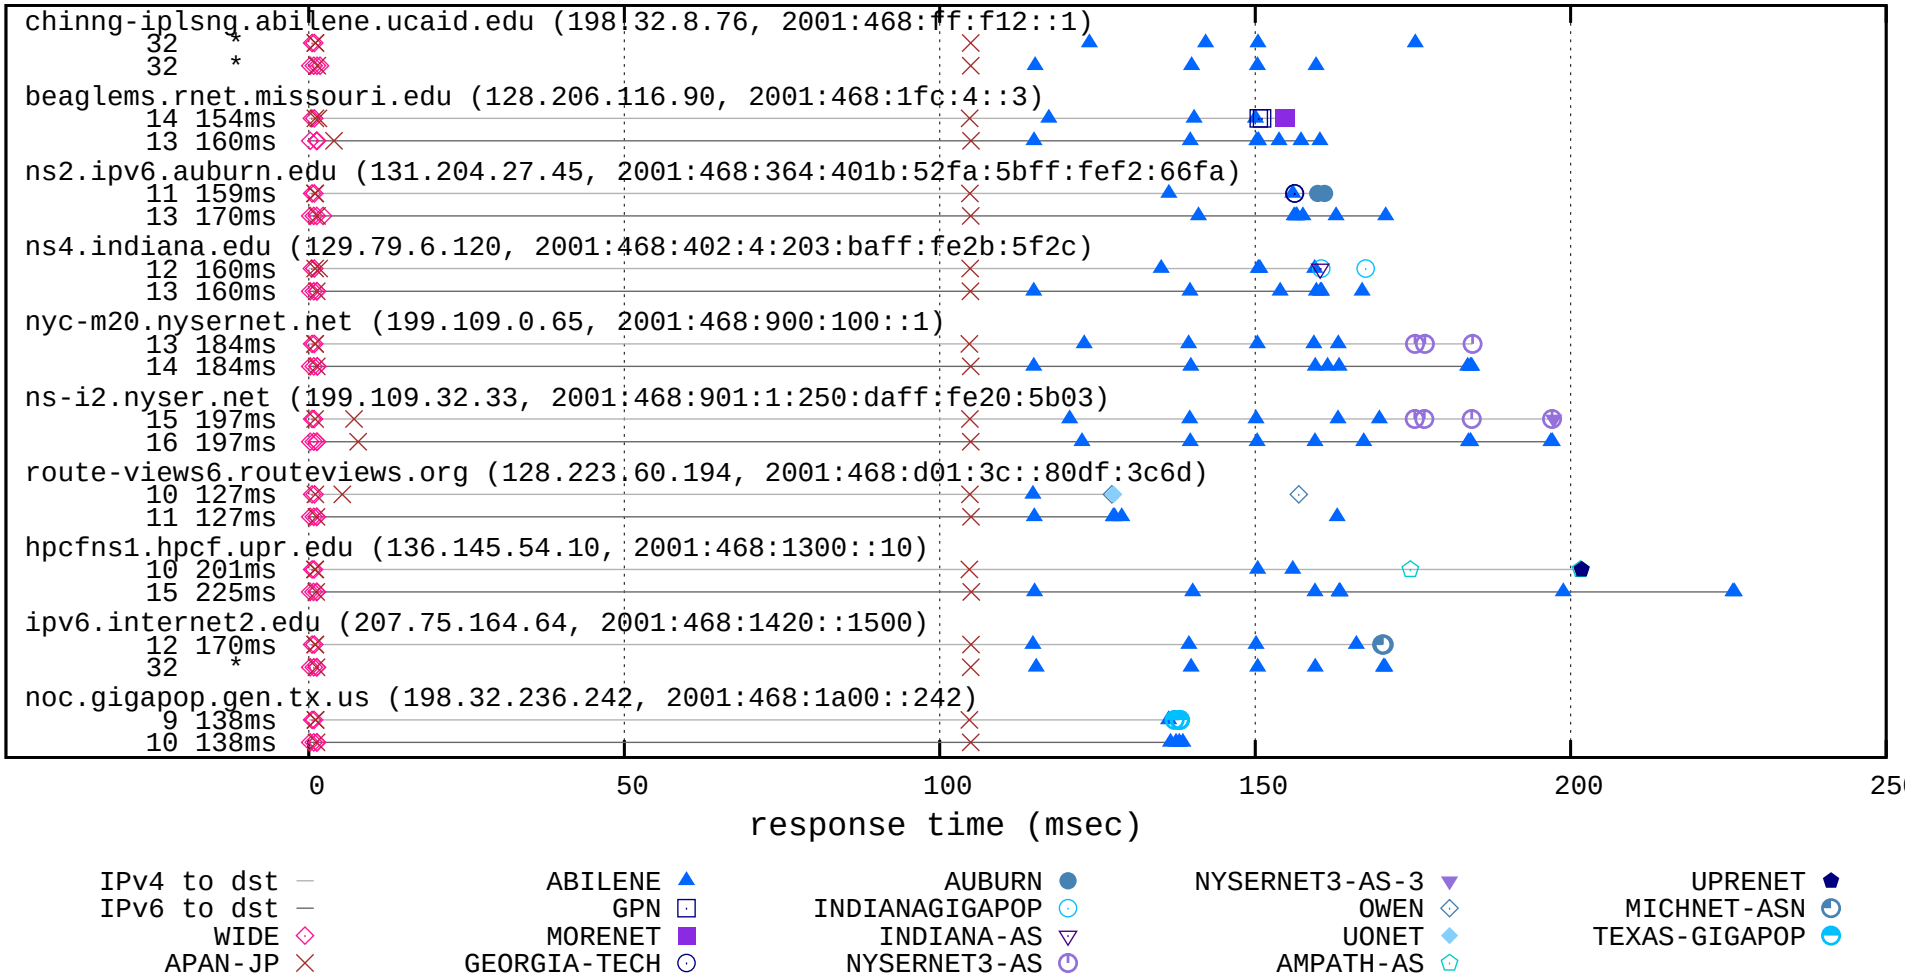


- similar RTTs and hops for IPv4 and IPv6
  - naitive dual-stack paths

# dual-stack traceroute to ABILENE from IIJ

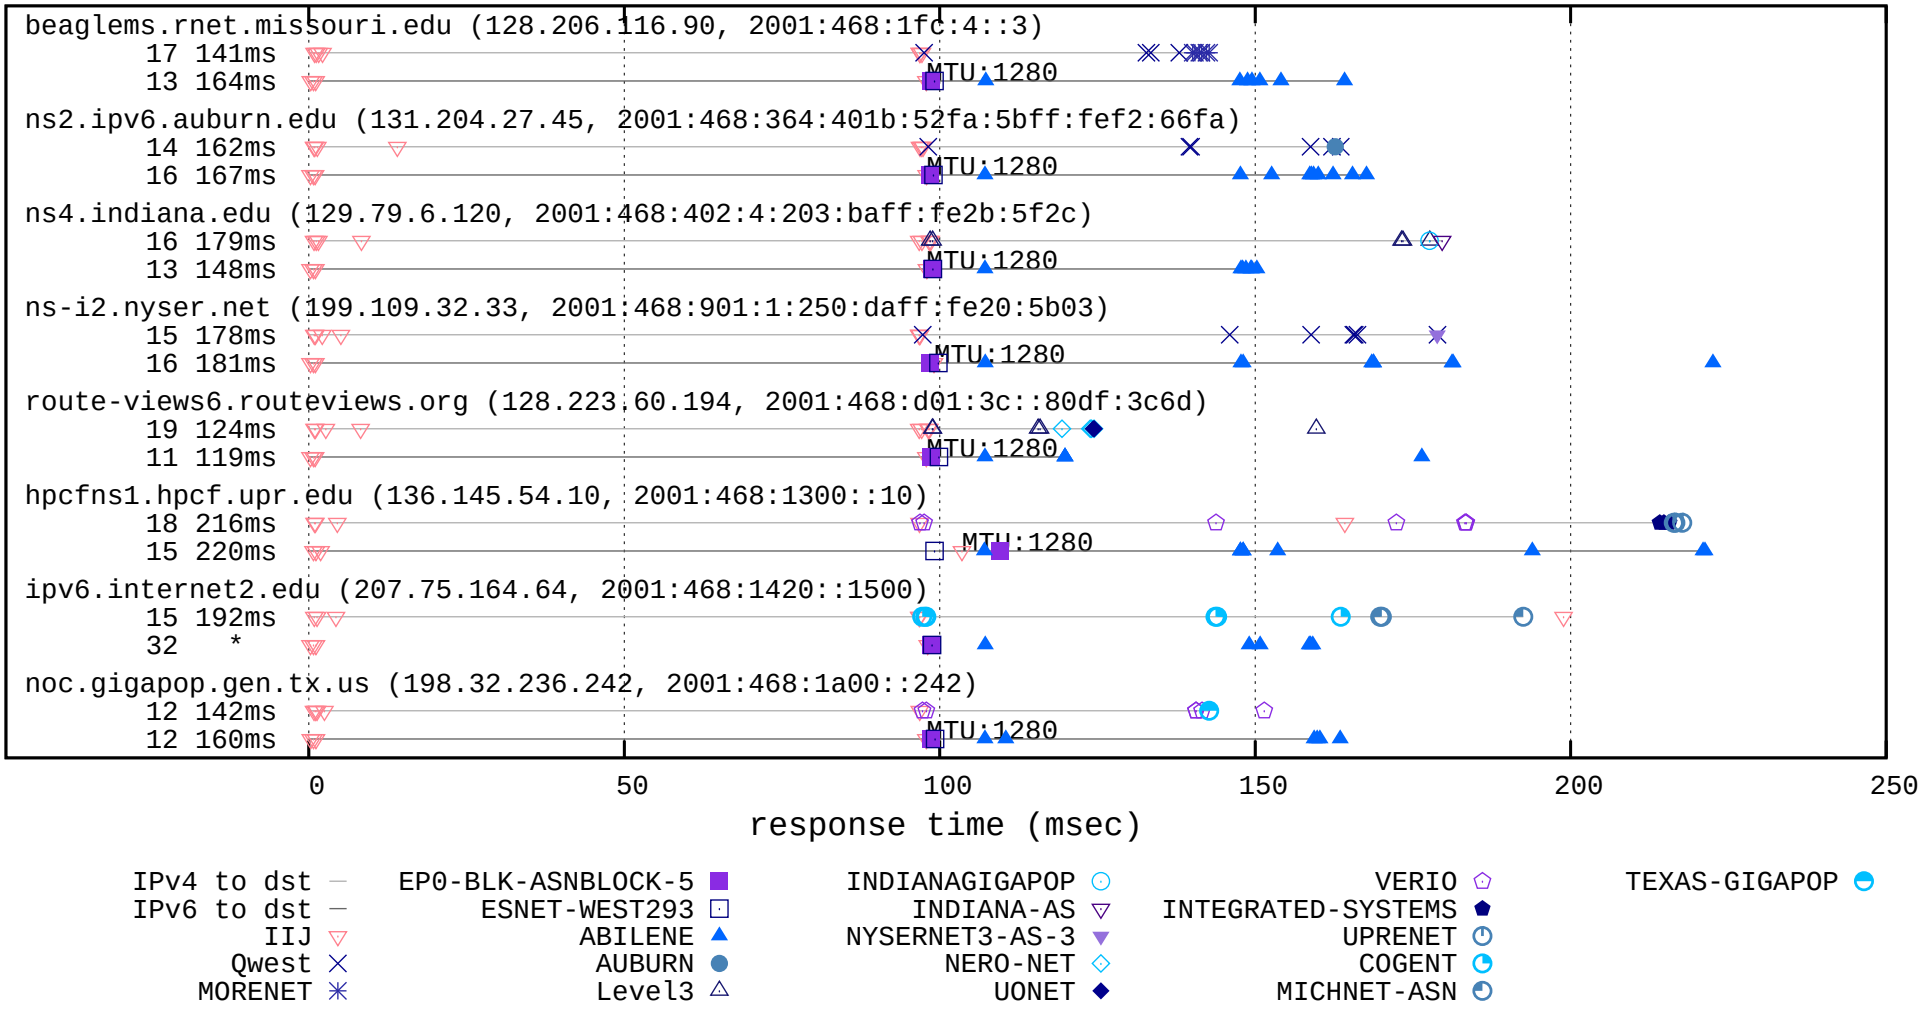


- similar RTTs but different paths for IPv4 and IPv6
  - more common in the current Internet

# dual-stack traceroute to ABILENE from ES

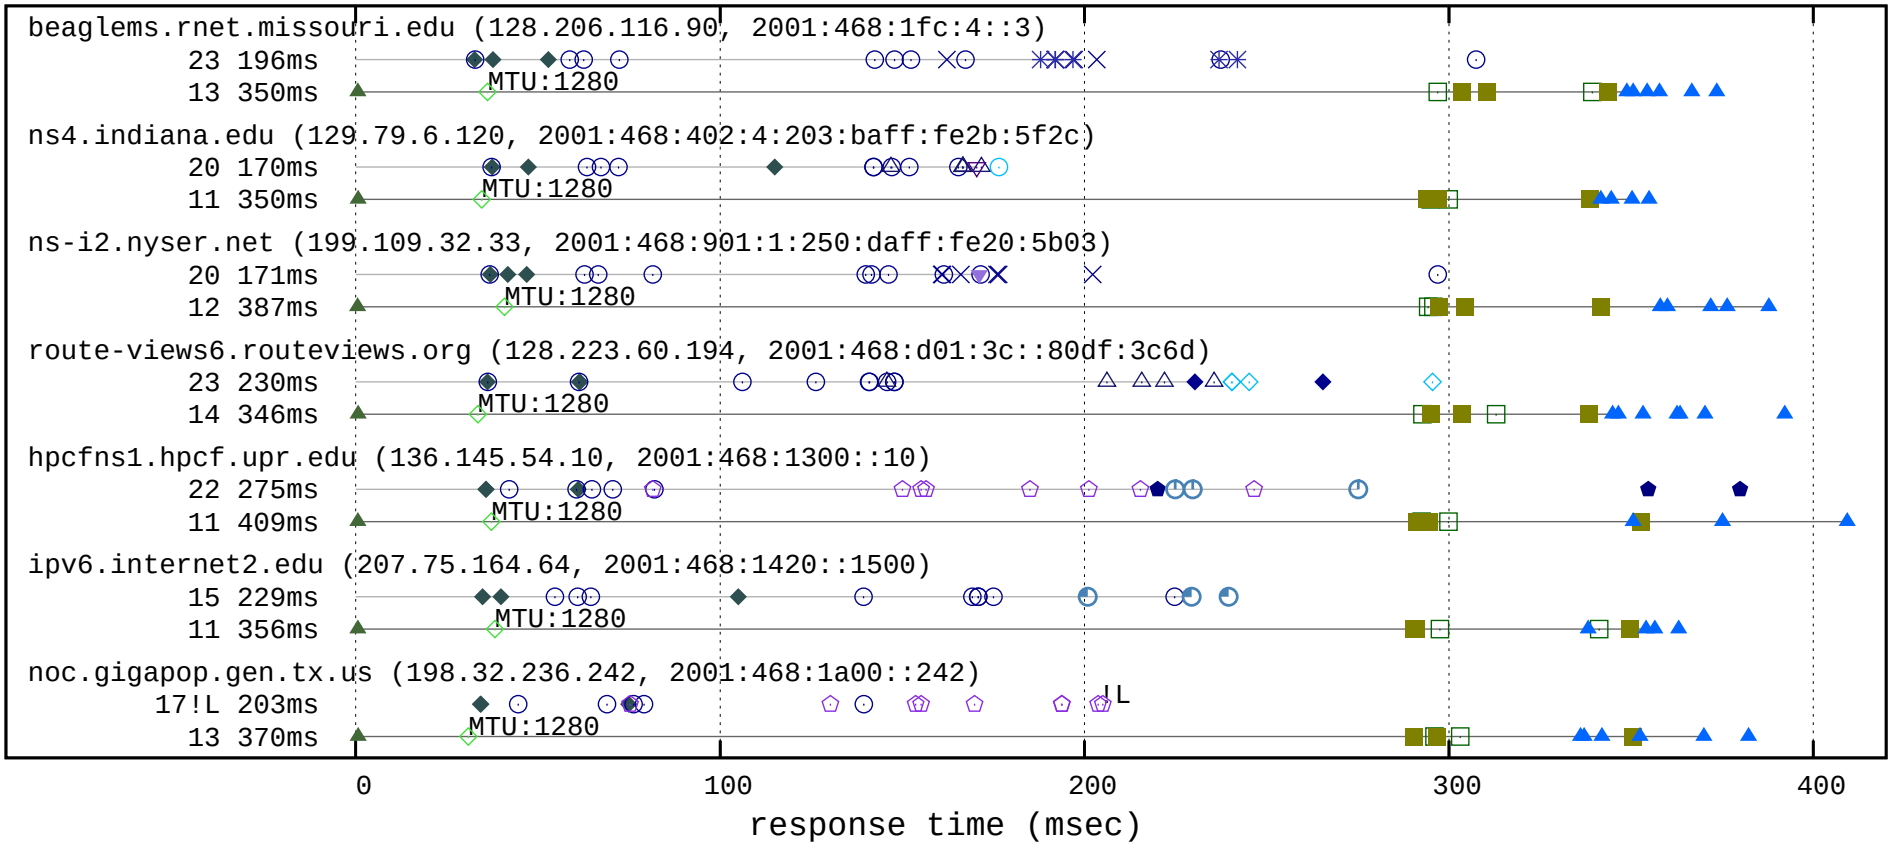


- IPv6 RTT larger than IPv4
  - roundabout tunnel
- note: ES has better paths to other US sites

# summary

- traffic measurement and analysis in WIDE
  - emphasis on wide-area, multi-point, long-term measurement
  - on our own backbone operated by us

- recent activities
  - IPv6 AS core map
  - residential broadband traffic analysis
  - dual-stack path analysis