Linking Internet traffic dynamics with the BGP topology : the view of two stub ISPs

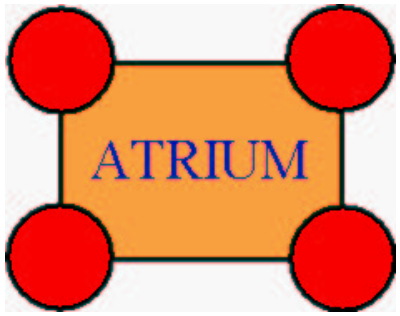

Steve Uhlig and Olivier Bonaventure

Infonet group

University of Namur, Belgium

Email : suh@info.fundp.ac.be

URL : http://www.infonet.fundp.ac.be/

The Internet Today

A large network :

- more than 100,000,000 hosts
- divided into more than 11,500 ASs
- more than 100,000 routes covering about 25 % of the IPv4 address space

What do we know about the distribution of the traffic inside the Internet ?

- routes, application mix, packet sizes,...

What is the topological distribution of the traffic ? Does it vary with time ? How stable is the traffic distribution ?

Existing studies : routing

Routing table studies

- CAIDA, Geoff Huston, RIPE,...

Principle

- Frequently Collect BGP tables at key Internet routers and track their evolution

Results

- Table size is (was) growing quickly
- Speed growth has decreased since we speak about it
- Provides some information about available routes and types of AS (connectivity)
- But no correlation between the actual traffic and the collected routing tables

Existing studies : traffic

Packet traces analysis

- CAIDA, MCI, Sprint, ATT and many others...

Principle

- Collect raw packet traces on key Internet links
- Packet trace often anonymized before analysis

Results

- Packet size distribution
- Application/protocol mix
- Microflows durations and distributions
- Self-similar stuff...
- But usually no topology-related information

Existing studies : both traffic and routing

Tier-1 ISPs

- ATT, Sprint, MCI,...

Principle

- Collect data (SNMP, Netflow, ...) on most routers or key routers to obtain traffic matrix

Results

- Mainly methods and sometimes tools to collect this data
- Usually, the statistics are confidential

Our study

Goal

- Obtain a better understanding of the (temporal and topological) distribution of the Internet traffic

Questions

- What fraction of the Internet sends traffic to a typical ISP ?
- What is the topological distribution of Internet traffic ?
- What is the variability of interdomain traffic ?

Our study (2)

Collection of the traffic statistics

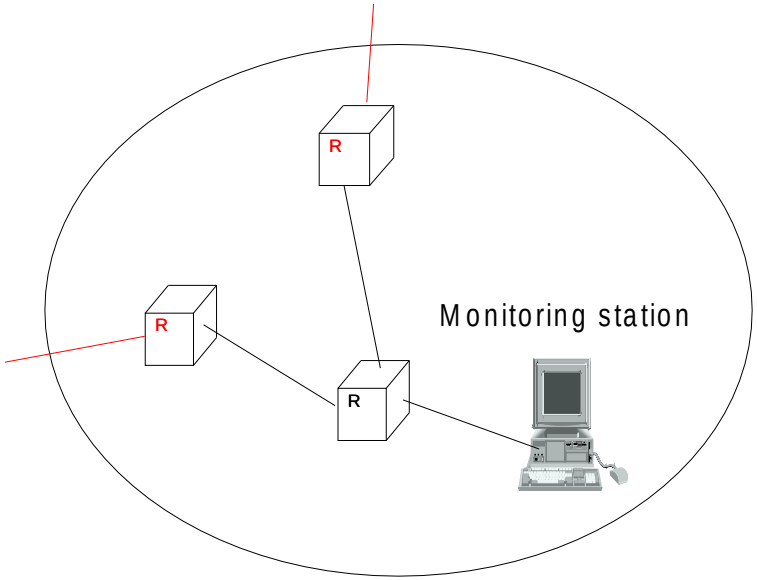

- Enable Netflow on external interfaces of border routers
- Collect Netflow data on monitoring station
- Download BGP routing table of border router

Our study (3)

Information provided by Netflow for each “layer-4 flow” seen by router :

- Timestamp for flow start and flow end
- Source and destination IP addresses
- Number of bytes/packets
- Protocol, TOS
- Input and output interface
- Source and destination ports, TCP flags
- Source and destination ASs (if router used its BGP table to route)
- Source and destination netmasks (if router used its BGP table to route)

Collected data was processed by own tools

The first studied ISP

BELNET

- Belgian ISP for research, education/government
- Typical user : student or researcher on campus network (Ethernet) attached to 34 Mbps backbone
- Interdomain links :
  - Two transit providers with E3/T3
  - Direct connection to TEN-155
  - Router on the Belgian and Dutch Interconnection points (BNIX and AMS-IX)
  - 63 direct peers
- One week Netflow trace collected with cflowd in Dec. 1999 covering all interdomain traffic
- Granularity of the trace : one minute, no AS info

The research ISP

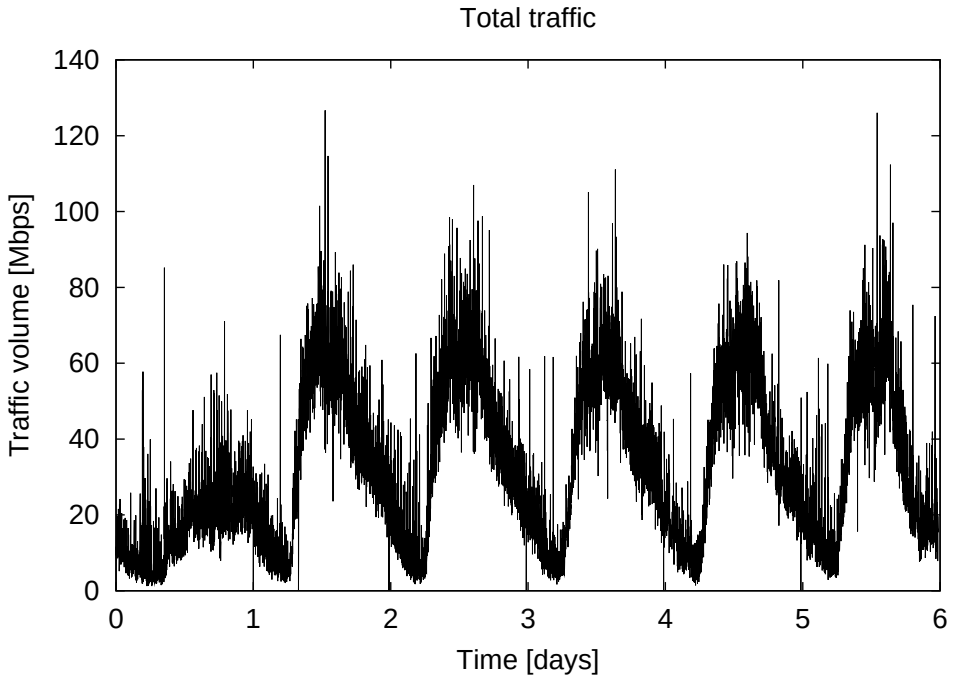

- Total volume : 2.1 terabytes in 6 days
- Business hours are busiests hours
- Average incoming traffic : 32 Mbps
- Peak incoming traffic during 1 minute : 126 Mbps

The research ISP (2)

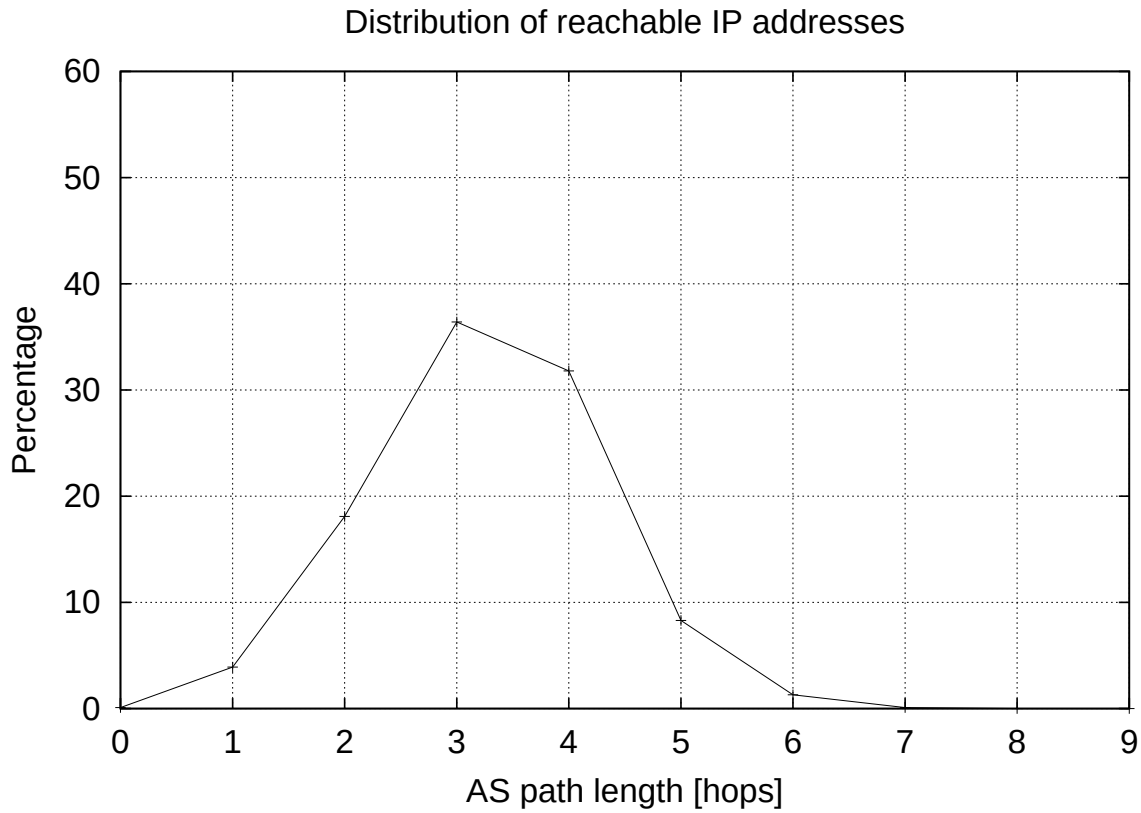

- BGP routing table : 68909 prefixes, 6298 AS
- 24 % of the IPv4 addressing space
- Average AS path length : 4.5 AS hops ($\neq$ from figure above)

The second studied ISP

YUCOM

- ISP offering free Internet access to dialup users
- Typical end-user : dialup modem
- Interdomain links :
  - Two transit ISPs with E3/T3
  - One router on Belgian Interconnection Point (BNIX)
  - 15 direct peers
- Netflow trace :
  - Covers all interdomain links for five consecutive days
  - Collected with flow-tools in May 2001
- Granularity of the trace : msec, no AS information

The dialup ISP

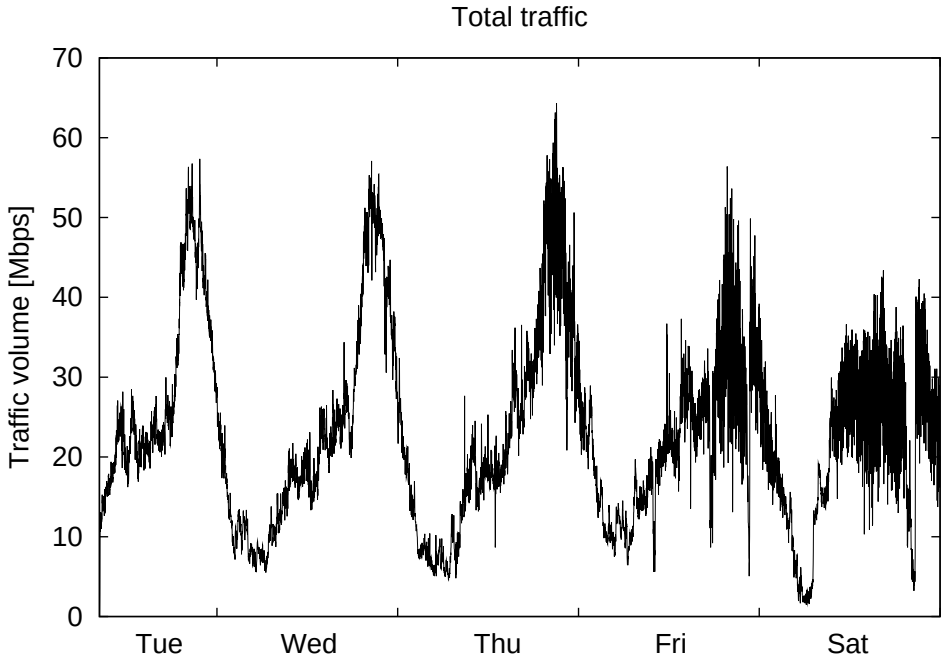

- Total volume : 1.1 terabytes in 5 days
- Evening hours are busiest hours
- Average incoming traffic : 23 Mbps
- Peak incoming traffic during 1 minute : 64 Mbps

The dialup ISP (2)

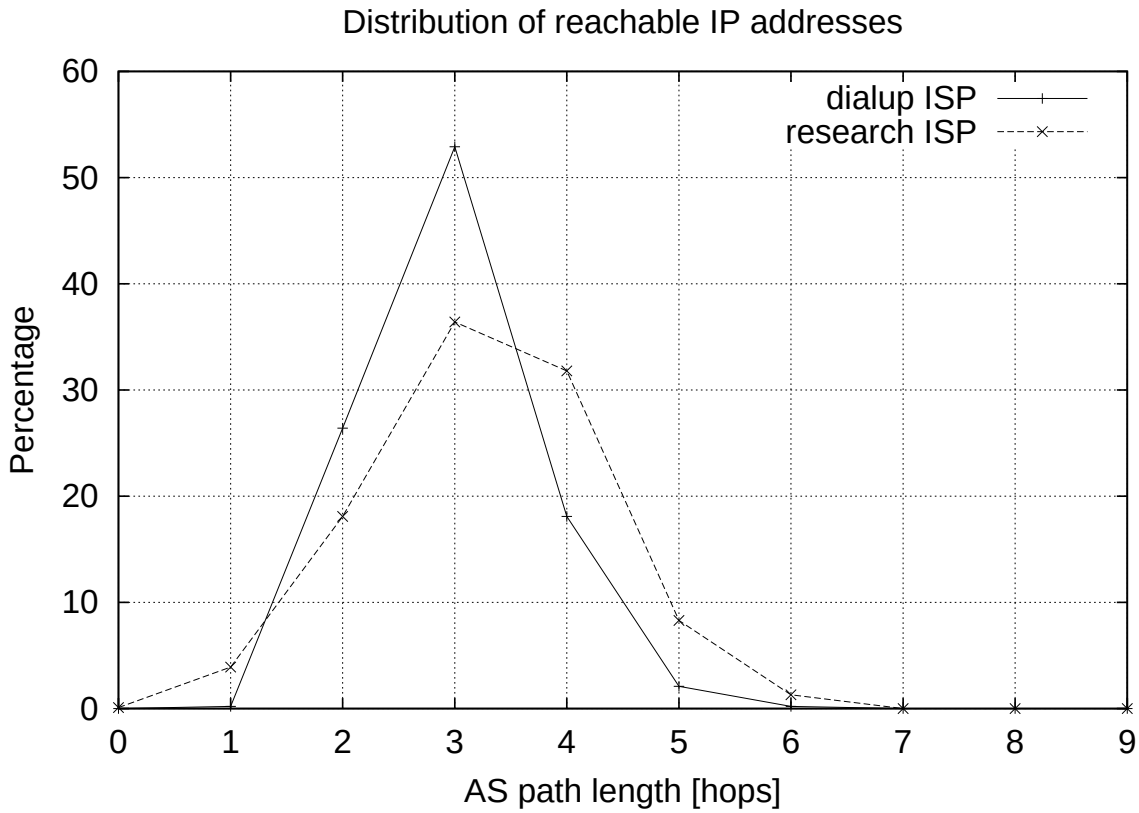

- BGP routing table : 102345 prefixes, 10560 AS
- 26 % of the IPv4 addressing space
- Average AS path length : 4.2 AS hops ($\neq$ from figure above)

Our study

Goal

- Obtain a better understanding of the (temporal and topological) distribution of the Internet traffic

Questions

- What fraction of the Internet sends traffic to a typical ISP ?
- What is the topological distribution of Internet traffic ?
- What is the variability of interdomain traffic ?

Topological traffic distribution
Research ISP

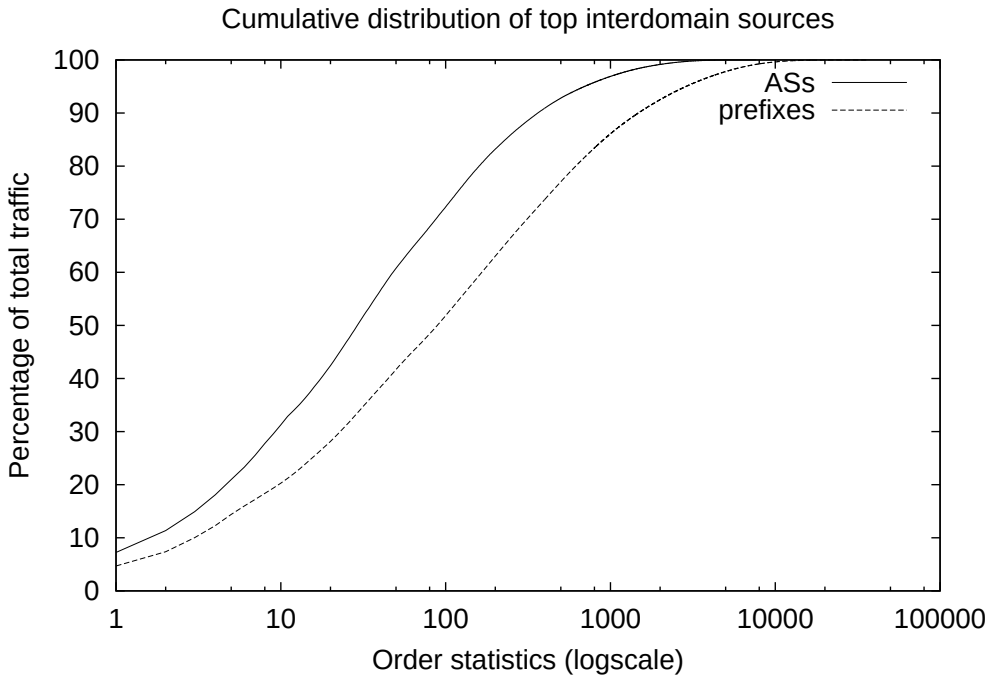

- Top 100 ASs (prefixes) capture 60 % (40 %) of total traffic
- 9.8 % of the ASs (4.5 % of the prefixes) capture 90 % of total traffic
- 5606 AS (out of 6298, $\sim$ 89 %) sent IP packets
- 719 AS sent less than 1 Mbytes of data

Topological traffic distribution
Dialup ISP

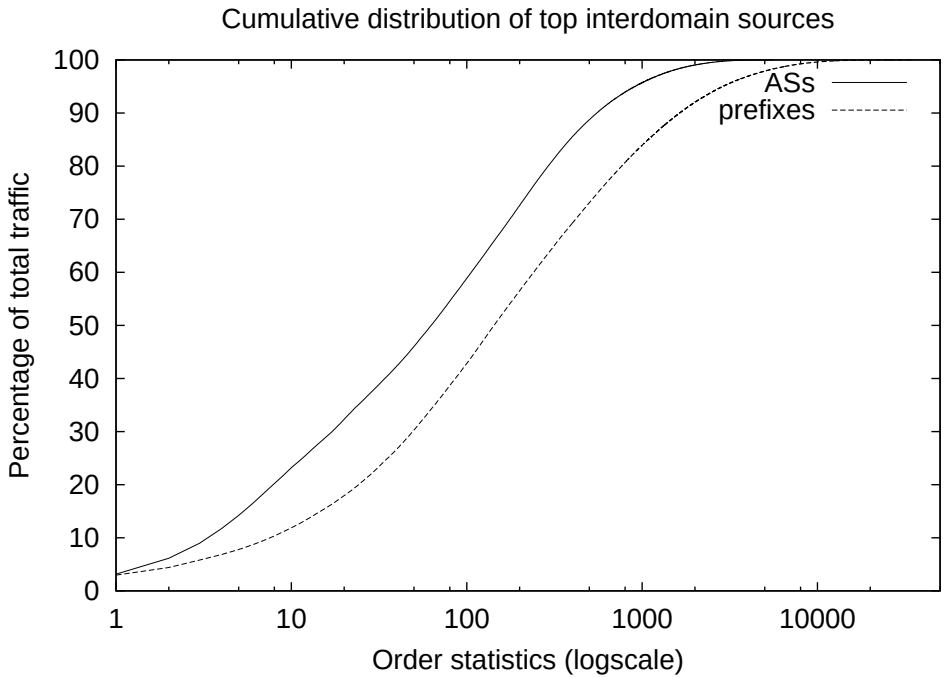

- Top 100 ASs (prefixes) capture 72 % (52 %) of total traffic
- 4.7 % of the ASs (4.1 % of the prefixes) capture 90 % of total traffic
- 7668 AS (out of 10560, $\sim$ 72 %) sent IP packets
- 4000 AS sent less than 1 Mbytes of data

Day-to-day traffic stability
Research ISP

- Do top ASs represent a stable concept in terms of traffic capture ?

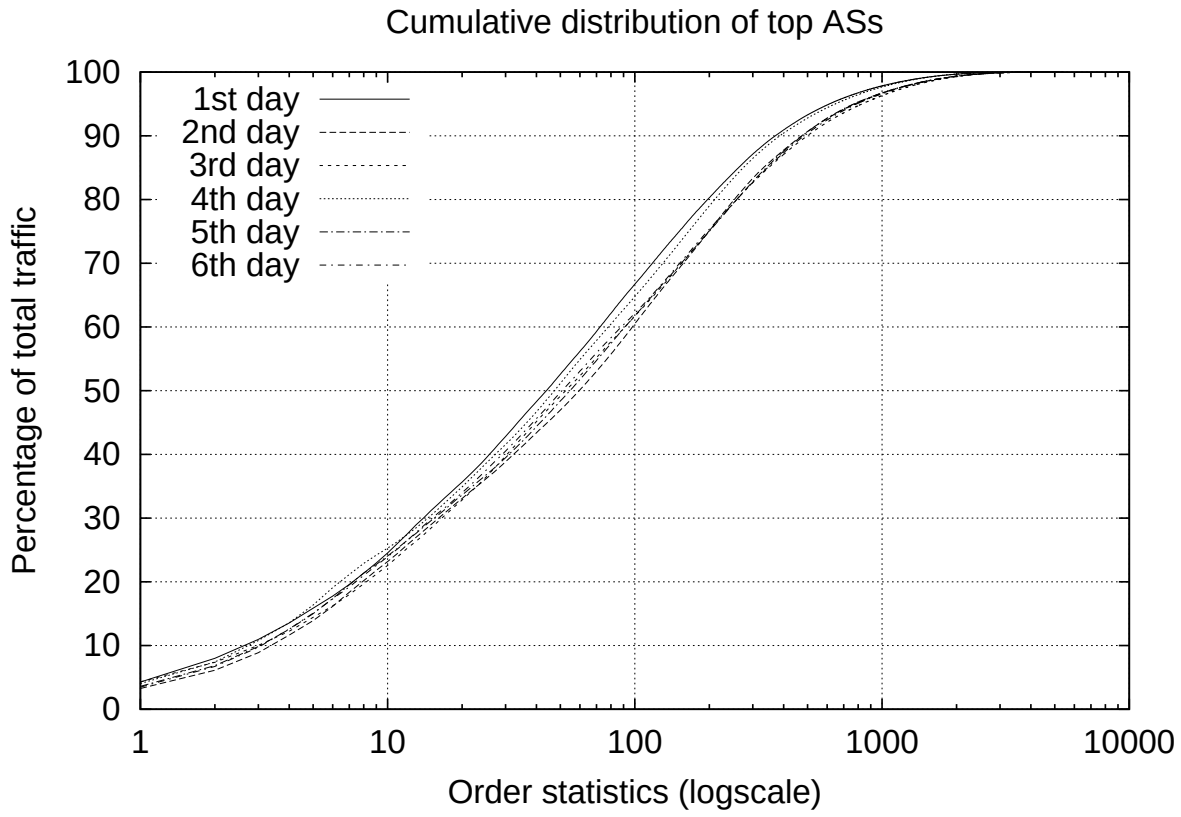

Topological traffic distribution does not change significantly from day-to-day (traffic capture is stable)

Day-to-day traffic stability (2)
Research ISP

- How well do yesterday's top ASs predict today's top ASs ?

Day	Top 10	Top 100	Top 1000
$1 \rightarrow 2$	5	70	773
$2 \rightarrow 3$	9	79	819
$3 \rightarrow 4$	8	64	723
$4 \rightarrow 5$	7	58	713
$5 \rightarrow 6$	9	74	821

Top ASs do not change significantly from day-to-day (important ASs are stable)

Our study

Goal

- Obtain a better understanding of the (temporal and topological) distribution of the Internet traffic

Questions

- What fraction of the Internet sends traffic to a typical ISP ?
- What is the topological distribution of Internet traffic ?
- What is the variability of interdomain traffic ?

AS hop distance vs. traffic
Research ISP

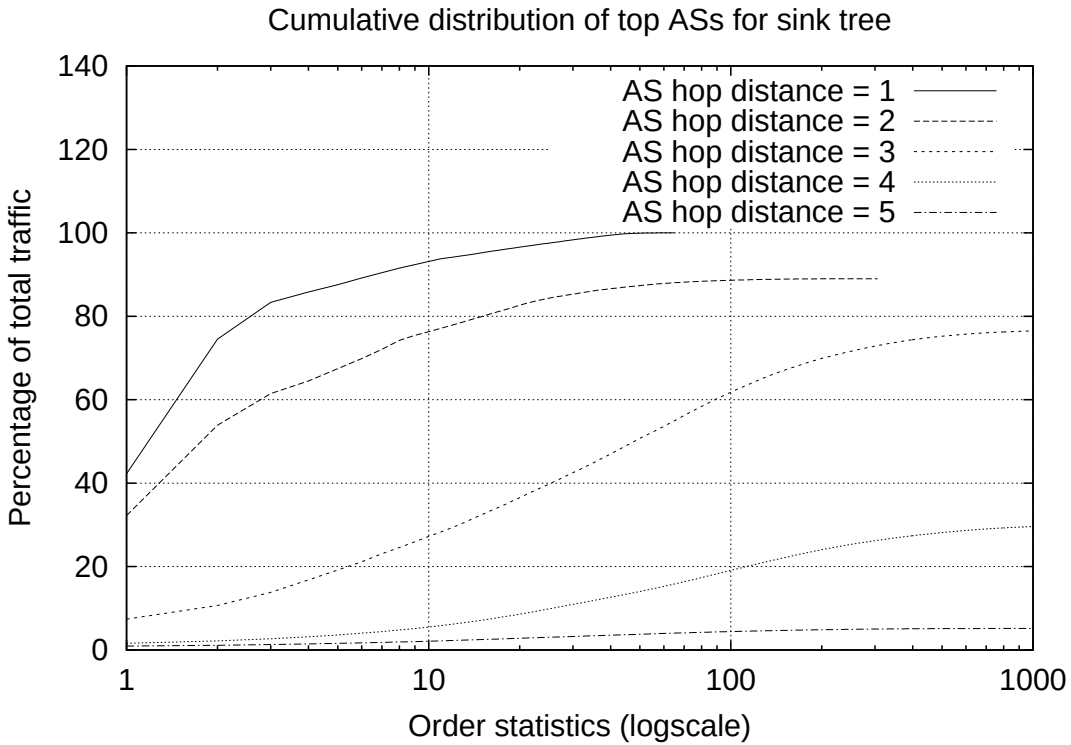

- curves tell cumulative traffic percentage seen by top ASs for each distance (both local and transit traffic)
- traffic coming from level i = difference between curve $i+1$ and i

AS hop distance vs. traffic (2)
Research ISP

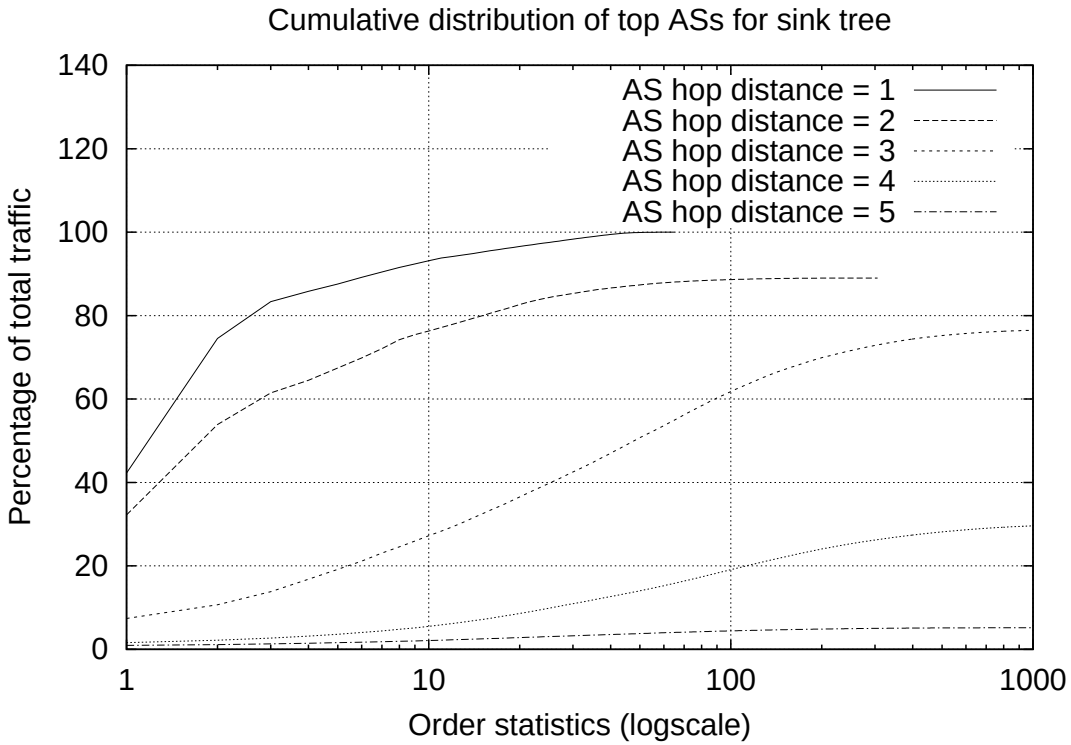

- Direct peers send only $\sim$ 11 % of the traffic
- A few ASs see a large fraction of the traffic (transit + local) at a distance of 1-2 hops
- Most traffic produced at AS hop distance of 3

AS hop distance vs. traffic
Dialup ISP

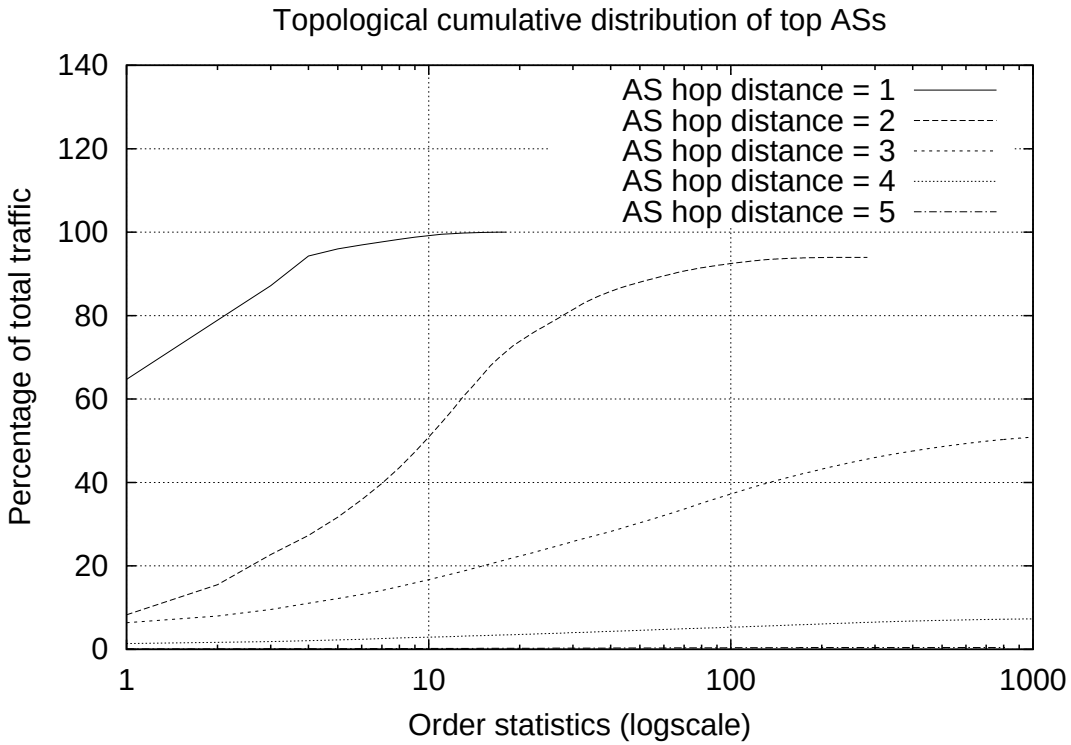

- Direct peers send only $\sim$ 6.1 % of the traffic
- Very few ASs see a large fraction of the traffic (transit + local) at a distance of 1 hop (more ASs at 2 hops)
- Most traffic produced at AS hop distance of 2-3

Our study

Goal

- Obtain a better understanding of the (temporal and topological) distribution of the Internet traffic

Questions

- What fraction of the Internet sends traffic to a typical ISP ?
- What is the topological distribution of Internet traffic ?
- What is the variability of interdomain traffic ?

Activity of the AS sources : Research ISP

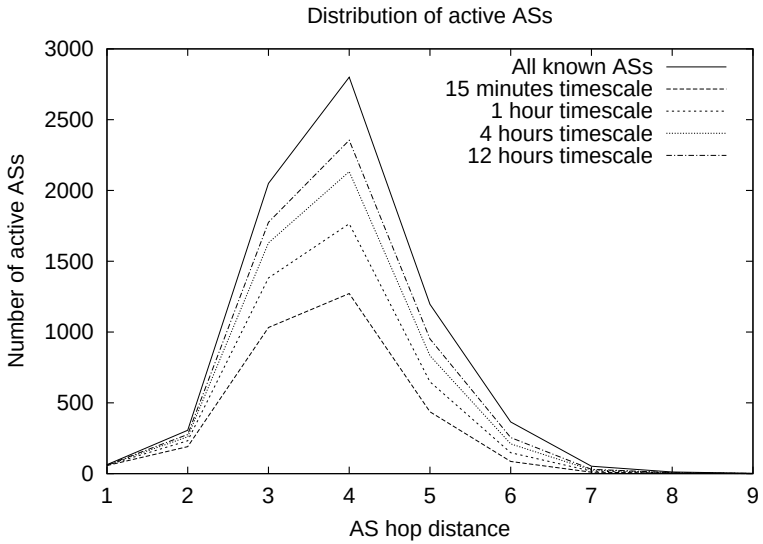

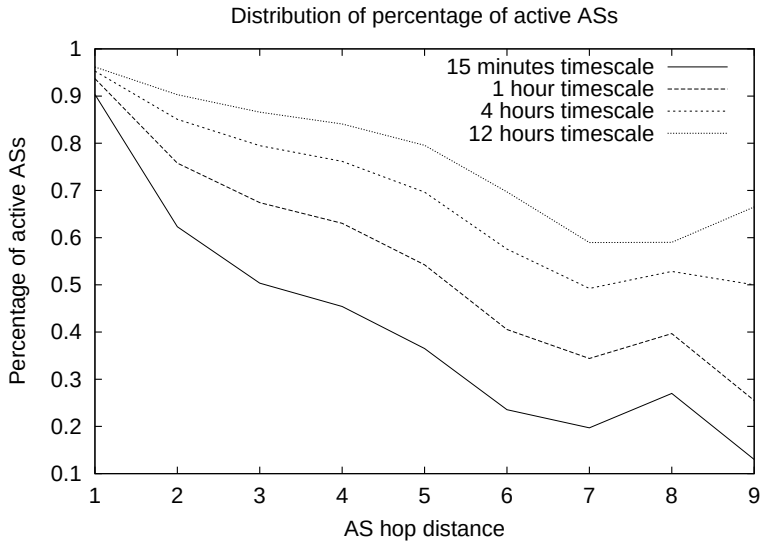

- Compute the average number of ASs generating at least one packet
- Direct peers are almost always active
- for 12 hours blocks, 5711 active ASs on average ($\sim$ 90 % of all ASs)
- for 15 minutes blocks, 3088 active ASs on average ($\sim$ 49 % of all ASs)

Variability of the activity of the AS sources could be a concern.

Activity of the AS sources : Dialup ISP

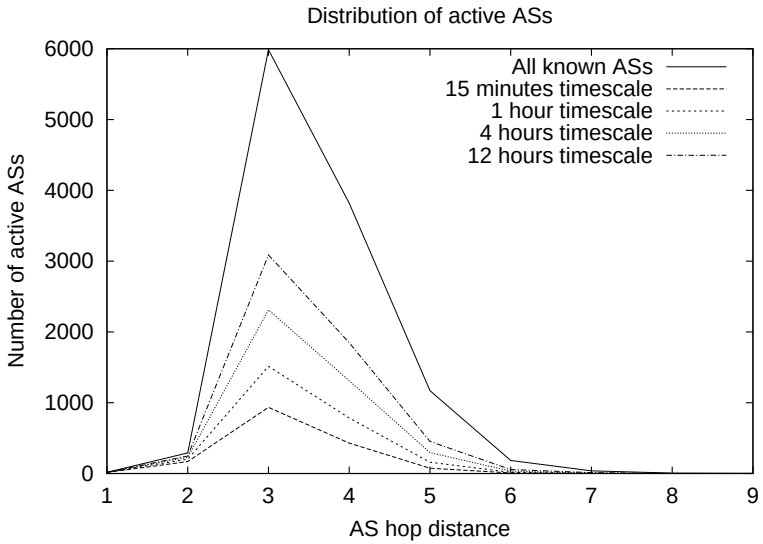

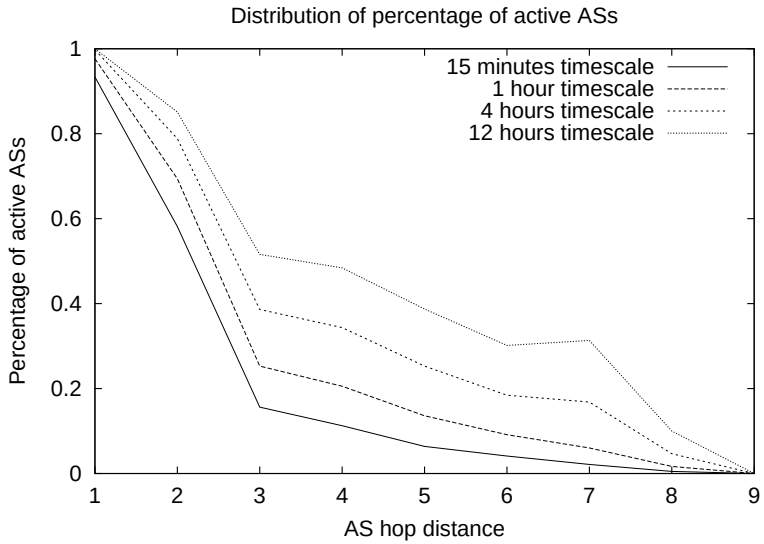

- Direct peers are almost always active
- for 12 hours blocks, 5721 active ASs on average ($\sim$ 54 % of all ASs)
- for 15 minutes blocks, 1632 active ASs on average ($\sim$ 15 % of all ASs)

Variability of the activity of the AS sources could be a concern.

Activity of the high-bandwidth AS sources : Research ISP

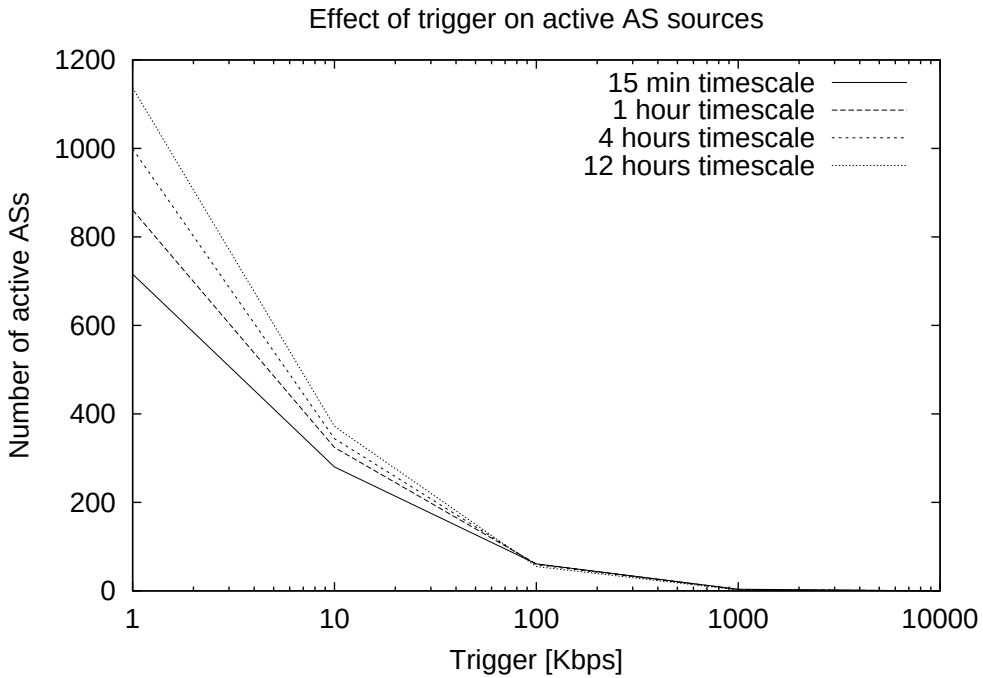

- 720 1 Kbps ASs are active during each 15 minutes period on average (out of 1633)
- Only 50 ASs are active on average with a bandwidth of 100 Kbps
- Average number of high-bandwidth sources is stable with time

The number of high bandwidth sources appears reasonable.

Activity of the high-bandwidth AS sources : Dialup ISP

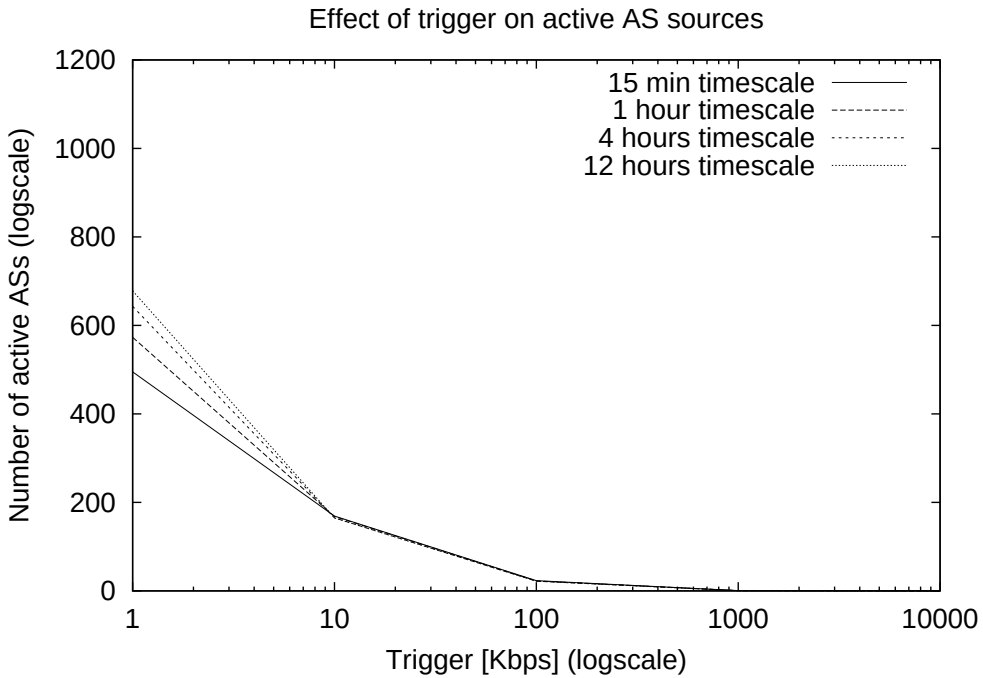

- 500 1 Kbps ASs are active during each 15 minutes period on average (out of 3088)
- Only 25 ASs are active on average with a bandwidth of 100 Kbps
- Average number of high-bandwidth sources is stable with time

The number of high bandwidth sources appears reasonable.

Conclusion

This study reveals several important issues to consider for interdomain traffic engineering :

- The Internet is becoming more and more flat, both when looking at BGP table and at traffic sources
- Even a medium ISP receives packets from a large fraction of the Internet : research ISP receives from 3088 ASs during 15 min, dialup ISP receives from 1632 AS during 15 min.
- Even with prefix/AS level aggregation, traffic variability in volume is a very important concern. (we have shown these flows are self-similar)

For more information

- Implications of interdomain traffic characteristics on traffic engineering. Steve Uhlig and Olivier Bonaventure. European Transactions on Telecommunications Journal, Special Issue on Traffic Engineering, January-February 2002.
- The macroscopic behaviour of Internet traffic : a comparative study. Steve Uhlig and Olivier Bonaventure Infonet-TR-2001-10, July 2001.
- All reports and papers available from `http://www.infonet.fundp.ac.be/doc/tr`