On the Topological Stability of Interdomain Traffic

Steve Uhlig (UCL)

Computer Science and Engineering Dept.

Université Catholique de Louvain

E-mail : suh@info.ucl.ac.be

joint work with Vincent Magnin (EPFL), Olivier Bonaventure (UCL), Chris Rapier (PSC) and Luca Deri (UniPisa)

Schedule

- Existing literature
- The question
- Measurements
- Analysis
- Lessons

Some studies on interdomain traffic

- Kleinrock et al. (1974) : inter-site traffic distribution in the ARPANET
- Claffy et al. (1993) : packet-level and SNMP traffic stats on NSFNet backbone
- Fang & Peterson (1999) : packet-level traffic stats on MCI backbone (several one hour traces)
- Uhlig & Bonaventure (2002-) : flow-level traffic stats on stub ASes (several one week traces)

Main lesson : few sites get/produce a large fraction of the traffic

Topological properties of interdomain traffic

How does traffic gets distributed on the AS-level graph ?

- transit AS view : no publically available data known (would be interesting...)
- stubs AS view: most traffic within 2-4 AS hops, tree-like distribution (temporary website at `http://bgp.mon-net.org/index.php`)

Who is this traffic ?

- 90^{ies} : mostly Web
- now : P2P, CDN and Akamai-like stuff

$\Rightarrow$ Interdomain traffic distribution could change a lot on AS-level topology

The Question

Existing literature shortcomings:

- existing studies consider a fixed “representative” timespan
- “representativeness” highly depends on study’s objectives (capacity planning, intradomain TE, POP-POP matrix inference,...)
- often traffic data size issues $\Rightarrow$ focus on relatively short time periods compared to expected shifts in interdomain traffic patterns, just enough for particular study

Q: How does traffic distribution on AS-level graph change with time when considering a large time period ?

Measurements

Measured stubs :

- UCL : university with multi-homed provider (BELNET), 1 month of flow-level traffic stats (avg 25 Mbps) + BGP session with BELNET (8.4 TBytes of IP traffic)
- PSC : regional POP in Pennsylvania, 1 day of flow-level commodity traffic stats (avg 165 Mbps) + BGP updates (1.7 TBytes of IP traffic)

Two different multi-homing contexts :

- BELNET peers : several IX points (AMS-IX, SFINX, BNIX) + 2 Tier-1 ISPs + GEANT
- PSC peers : 2 Tier-1 ISPs (+ Internet2)

Measurements (2)

What we measured :

- for each outgoing flow :
  1. find BGP best matching prefix *pref* for flow's destination IP
  2. atttribute to *pref*'s AS path flow traffic
- draw the graph based on all AS paths seen during some given time interval

Trafficed AS path graph

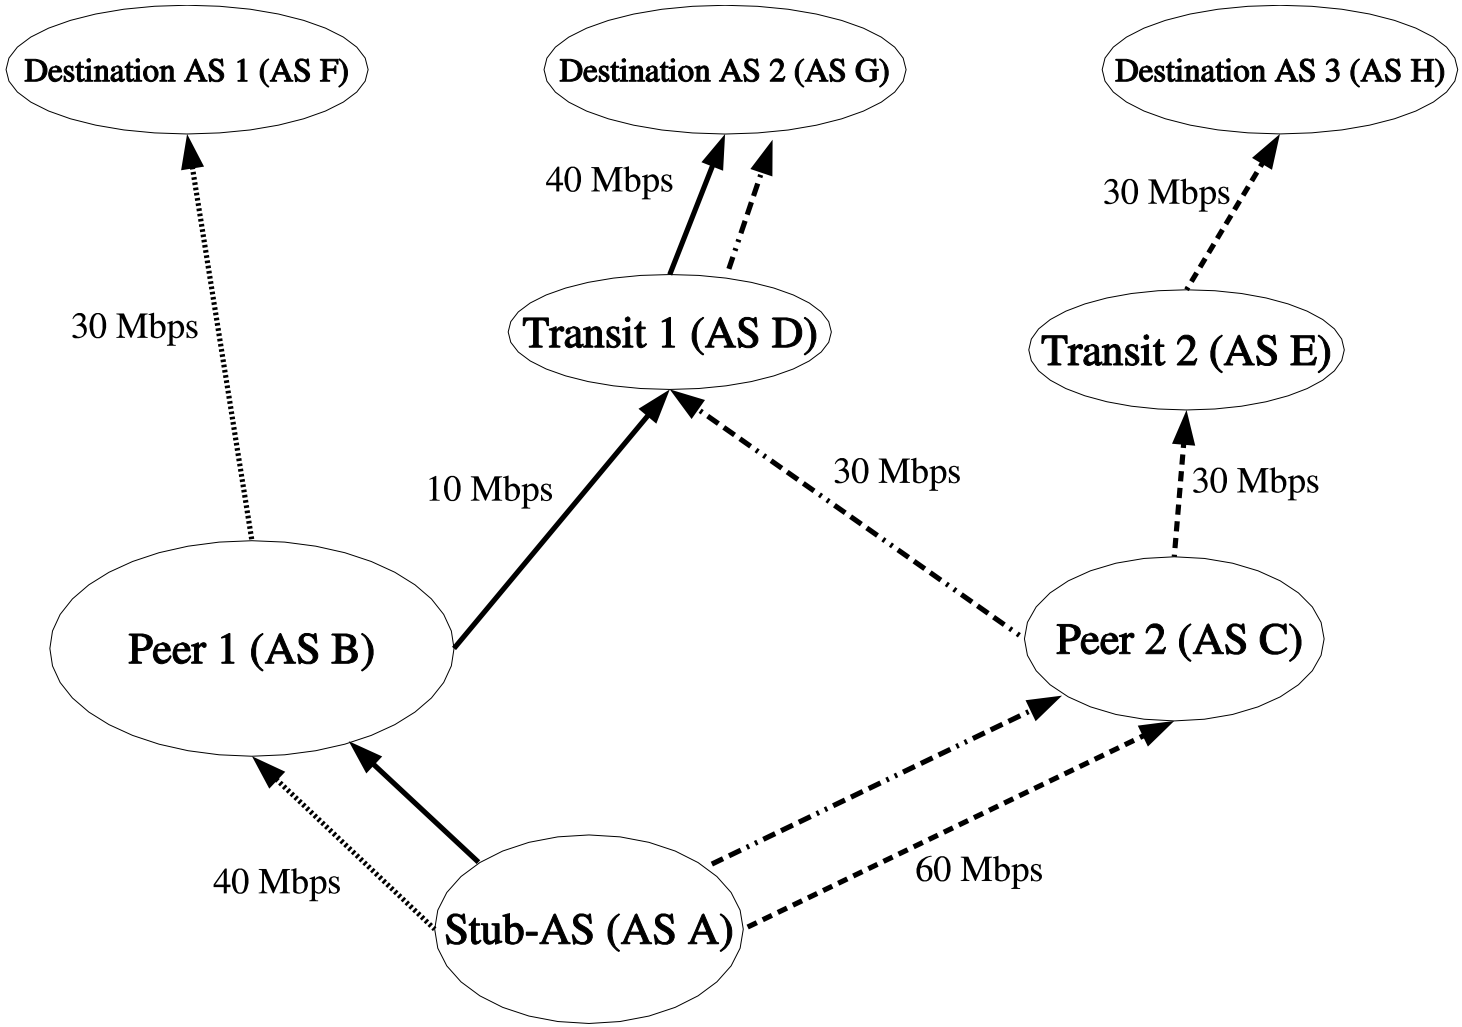

AS paths presence

How much do distinct AS Paths last over long time periods ?

- 103,853 distinct AS paths
- 31,151 carried any traffic at all
- 50 % were present less than 9 minutes
- 42 % of trafficed AS paths were present more than 99 % of the month

How much do **trafficed** AS Paths last over long time periods ?

- 95 % of traffic carried by AS paths that lasted more than 99 % of the time

$\Rightarrow$ most traffic sees stable AS paths

Traffic percentage for AS paths

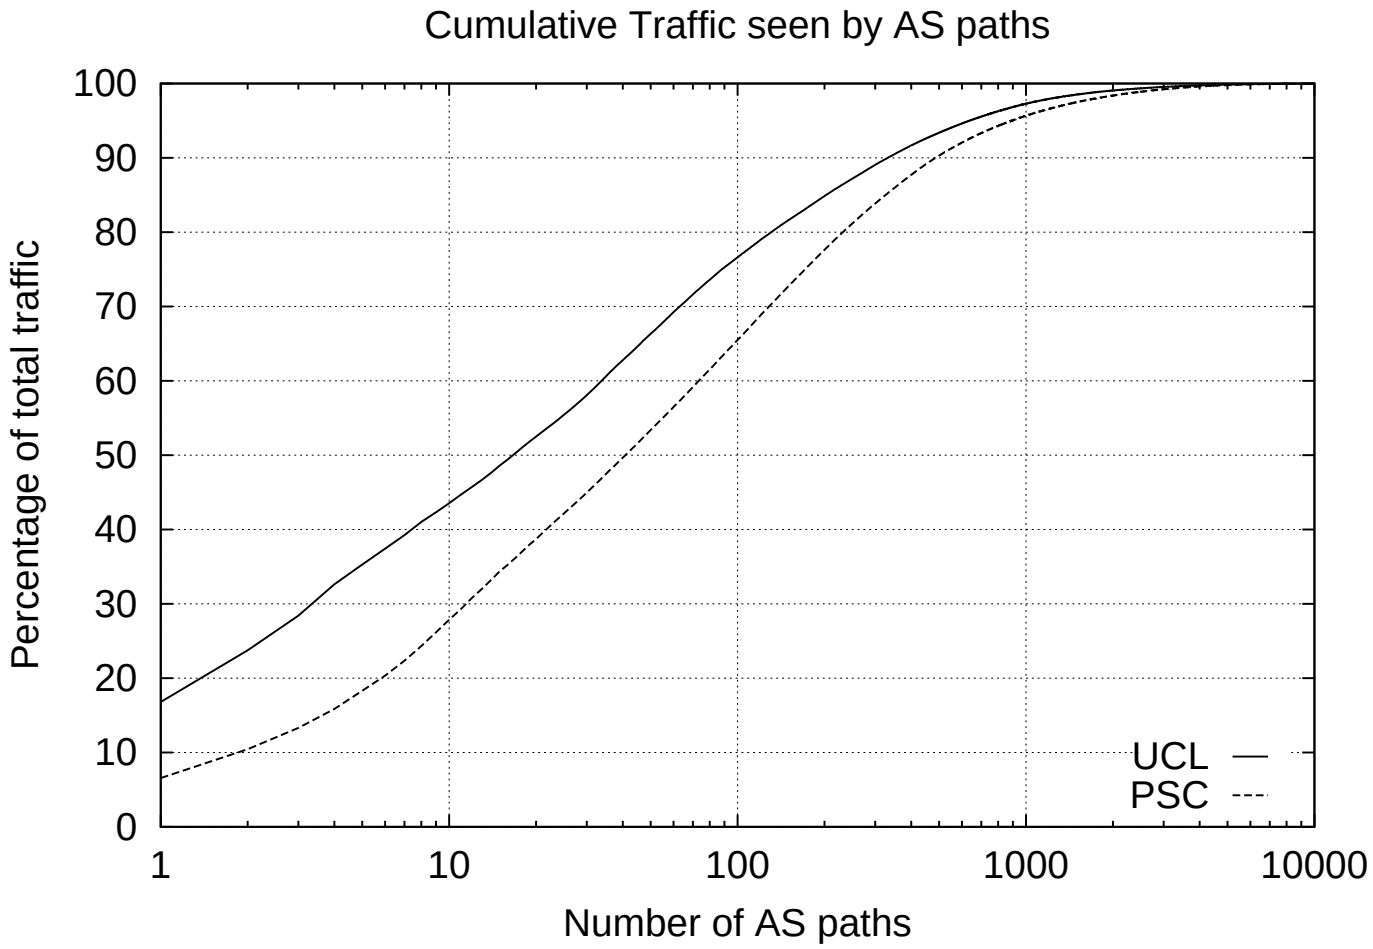

Traffic aggregation for AS paths

- add traffic for each consecutive ASes of each AS path
- remove AS paths and AS pairs having less than *x* % of total traffic
- compute how many AS paths left :

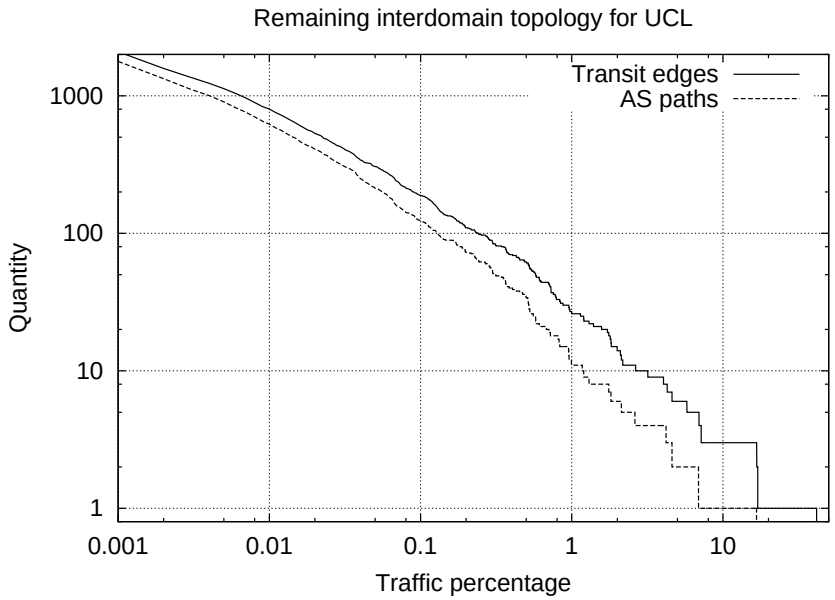

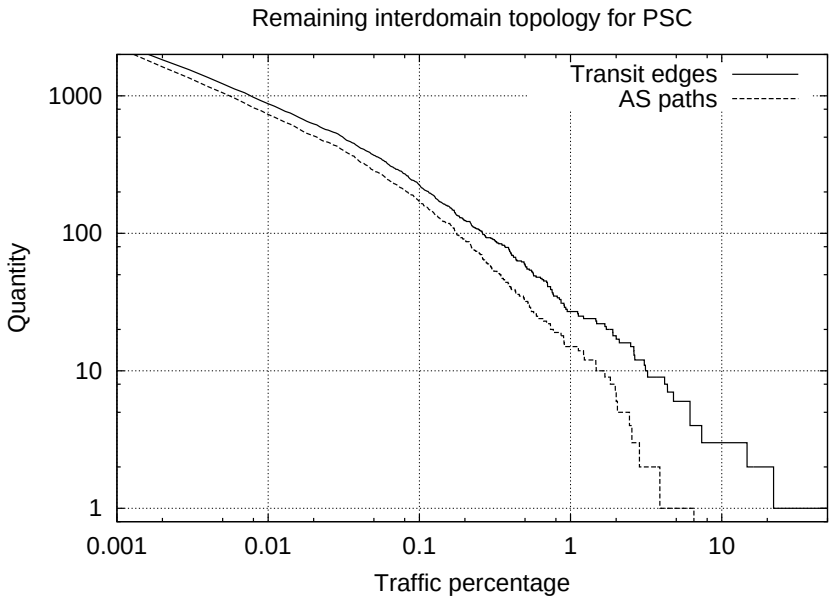

Top AS paths dynamics

For each hour, what is the minimal number of AS paths to capture *x* % of the traffic ?

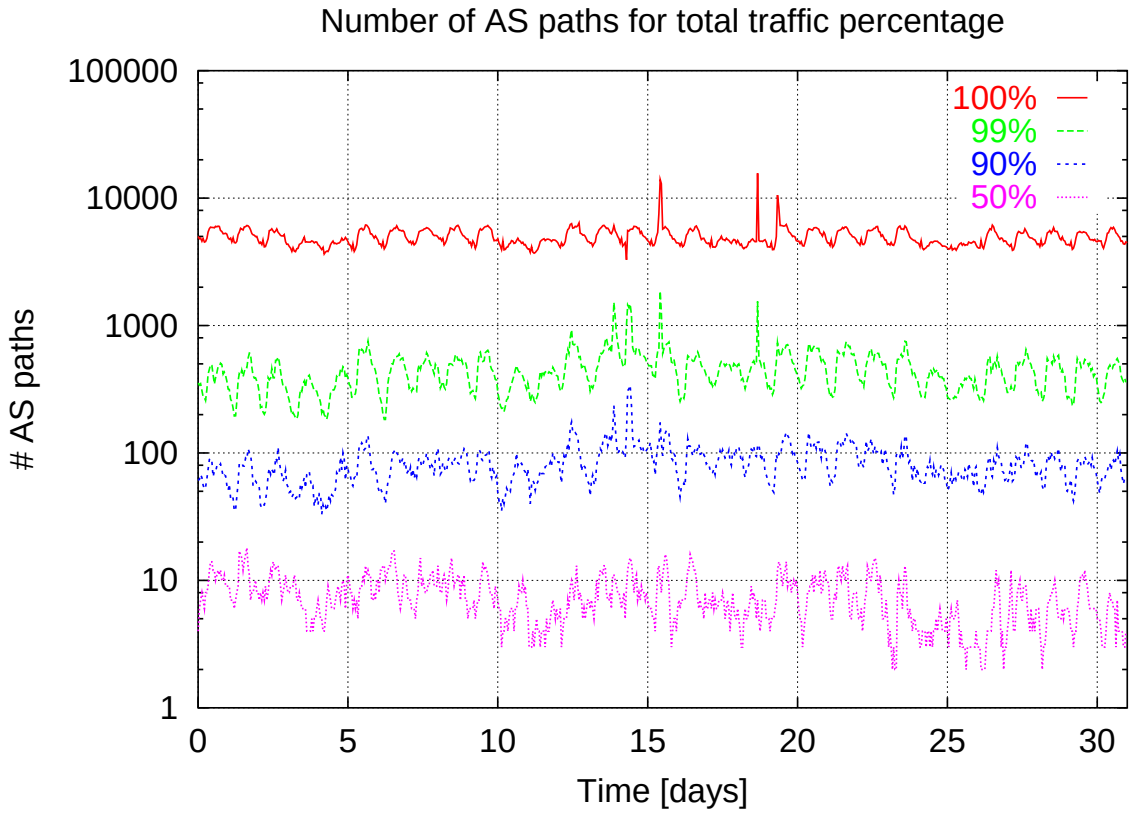

Presence of Top AS paths

1. for each hour, compute the largest traffic AS paths capturing 90 % of the traffic
2. for each AS path found in step 1, mark how many times it is seen among the top 90 %
3. plot presence of these AS paths :

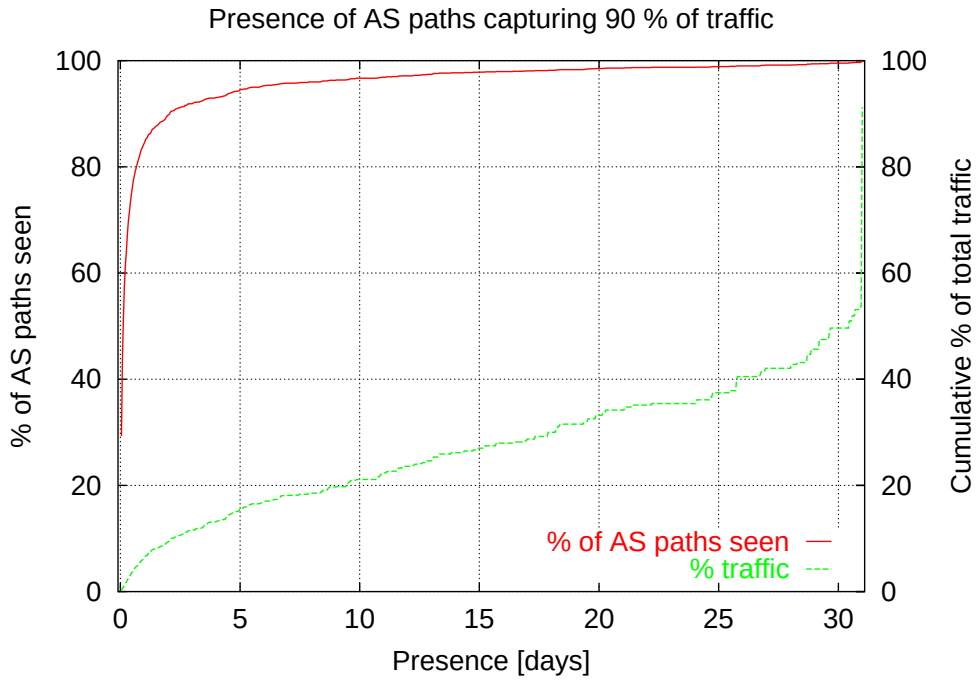

Predicting the future ?

- up to now we assumed that we know top AS paths : what if we do not know them ?
- so compute top AS paths over a time window and compute how much traffic they capture :

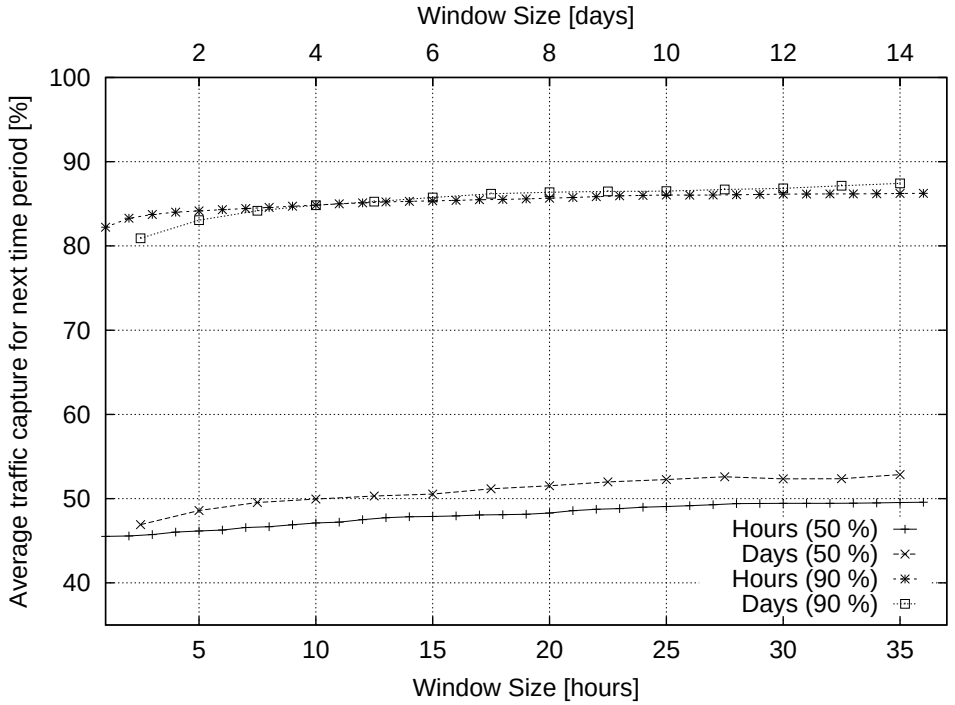

Lessons learned

- AS paths are stable for the largest fraction of the traffic
- part of largest trafficed AS paths are stable, part are unstable
- influencing more traffic requires influencing more AS paths and a larger topology
- stubs see a “stable” transit core where aggregation occurs
- leafs of the trafficed interdomain graph (for stubs) are unstable