Characterizing the Global Impact of P2P Overlays on the AS-level Underlay

- Amir H. Rasti, Reza Rejaie (Univ. of Oregon)
- Walter Willinger (AT&T Labs – Research)

PAM 2010, ETH Zurich, April 2010

Motivation

- Growing interest in P2P/overlay-based applications.
- Increasing concerns among edge ISPs (AS) due to associated external traffic.
- Researchers mainly focused on techniques to reduce external traffic on edge ISPs.
- Global effect of a P2P overlay on individual ASes (AS-level underlay) is challenging and not well understood.
- Assessing this global effect requires
  - Understanding the design & characterization of overlay-based applications
  - Understanding the characterization of AS-level topology & BGP routing
  - Dealing with inaccurate, missing, or ambiguous information about AS-level topology

This paper presents

- A methodology that incorporates a set of best existing practices to assess the global impact of a P2P overlay on AS-level underlay

- Characterization of the global impact of a few real-world P2P overlay on the AS-level underlay using different measures.

Methodology

- Goal: assessing the aggregate imposed load of overlay connections on and between individual ASes

- Methodology: An Overview

1) Capturing the overlay topology

2) Estimating the load on individual overlay connections

3) **Inferring the AS-path associated with individual overlay connections**

4) Determining the aggregate load on and between individual ASes

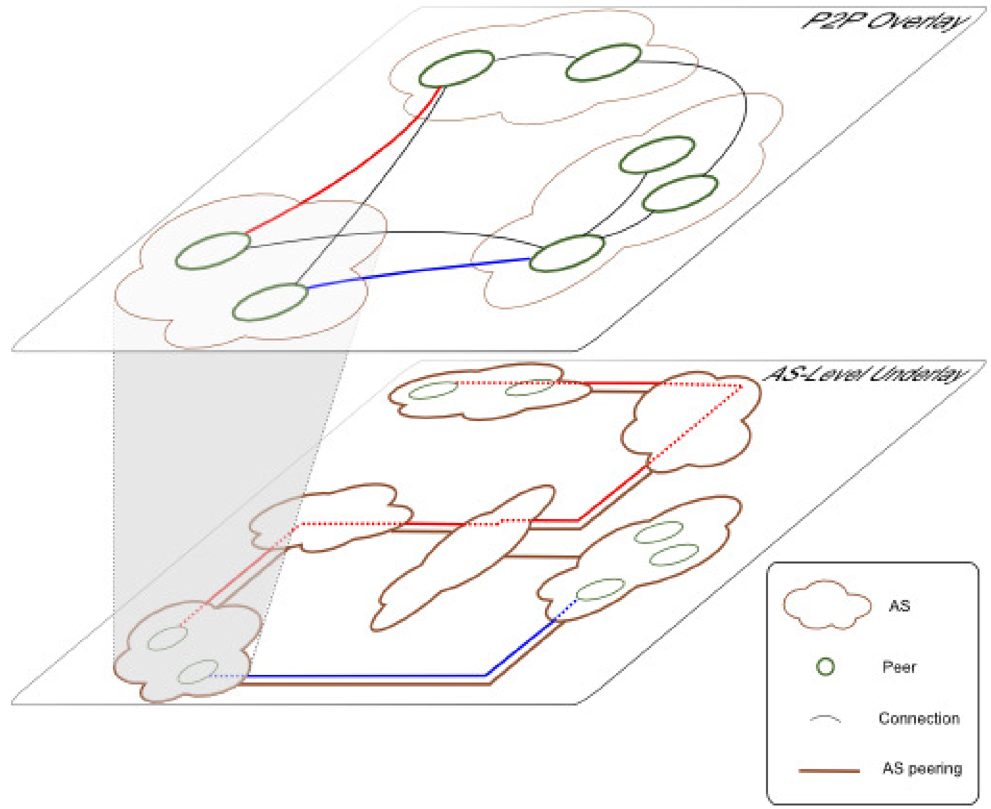

1) Capturing Overlay Snapshot

- The ability to capture overlay snapshots depends on the features offered by the target P2P application
- We captured top-level overlay of Gnutella using a fast crawler mainly due to feasibility
  - These overlay snapshots are realistic, large and accurate.
- More bandwidth-intensive applications (e.g. BitTorrent) offer better candidates but
  - Capturing accurate snapshots may not be feasible
- Our methodology can be used with any given overlay topology

2) Estimating load on Overlay Connections

- Load on individual overlay connections depends on several factors
  - Number of sources, their rate and pattern of traffic generation, and their relative location in the overlay
  - Overlay topology
  - Relaying/Routing strategy at individual peers
- Overlay traffic model or measurement can be used to assign load to individual connections.
- Simplifying assumption: All connections observe the same average load (in each direction)
  - Load can be measured as the number of overlay connections
- More accurate load values (or models) can be incorporated

3) Inferring AS-paths (1)

- Determining the AS path associated with each overlay connection

1) Mapping peers to ASes
   - Using BGP snapshots (from RouteViews) to translate the IP address of each peer to the AS number for its corresponding edge AS.

2) Capturing AS-level topology & Inter-AS Relationship
   - Using annotated AS-level topologies provided by CAIDA
   - These topologies include inferred relationship between connected ASes
     - 1) Customer-Provider, 2) Peer-Peer, 3) Sibling-Sibling
   - It is shown that these topologies are missing a significant fraction of peering links between lower-tier ASes
   - More accurate AS-level topologies can be incorporated.

3) Inferring AS-paths (2)

3) Inferring AS-Paths by simulating BGP
   - Simulating BGP over the annotated AS-level topology using **C-BGP**
   - in a predefined network with configured routers
   - C-BGP represents each AS as a single node, maps inter-AS relationships into intuitive BGP routing policies (see paper for details).
   - Once BGP converges, we obtain AS-level path between any pair of ASes (see tech report for list of problems)

- This approach has two shortcomings
  - It cannot consider possible multiple paths between a pair of ASes
  - Possible error due to unknown & diverse BGP policies between ASes with different relationships

- Carefully-designed measurement may provide a viable alternative to simulation!
  - but has its own challenges and limitations

- Assessing AS Tiers
  - Using TierClassify [Gao] tool to classify ASes into tiers

4) Determining Aggregate Load

- Calculating the aggregate load imposed on each transit AS by summing up the traffic carried by each AS-path on all ***transit (or core)*** ASes and links along the path

Datasets

Snapshot	Gnutella Snapshots		BGP Snapshots		AS-Paths	
	#Peers	#Conn.	#Prefixes	#ASes	#Unique	%Important
G-04	177k	1.5M	165k	19k	192k	2.0
G-05	681k	5.8M	185k	21k	384k	2.9
G-06	1.0M	8.6M	210k	23k	605k	2.8
G-07	1.2M	9.8M	229k	25k	684k	2.7

- Both overlay and underlay have grown significantly over 4 years
- Only 2%-3% of paths carry +100 overlay connections

Characterizing Impact on Underlay

- What is a useful way to present the global impact of an overlay on the AS-level underlay?
- We used the following measures:
  - Diversity & load on individual AS-paths
  - Load on individual transit ASes
  - Identity and evolution of top transit ASes
  - AS path length
  - Propagation of traffic through AS-level hierarchy

Load Distribution (1)

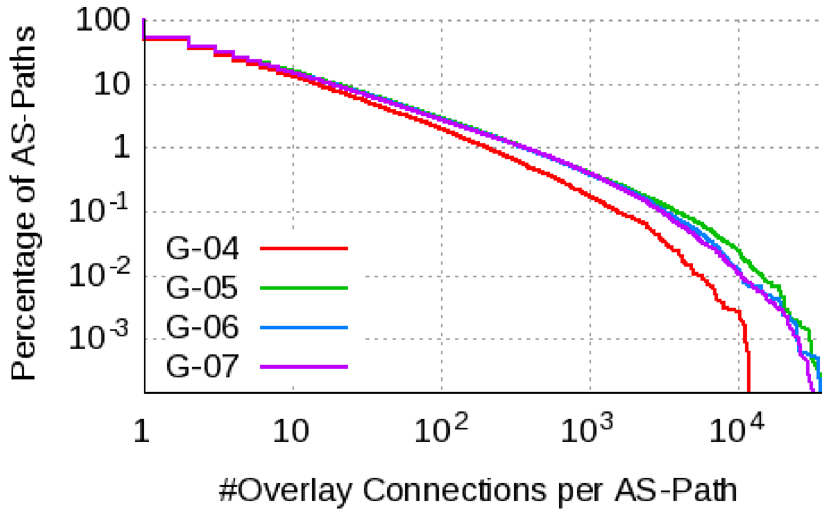

- CCDF of number of overlay connections per AS-Path
  - Skewed for all snapshots
  - Small number of AS-Paths carry a large fraction of load
  - *e.g.* about 10% of AS-Paths each carry 10 connections or less while 1% of AS-Paths carry 200 connections or more

Load Distribution (2)

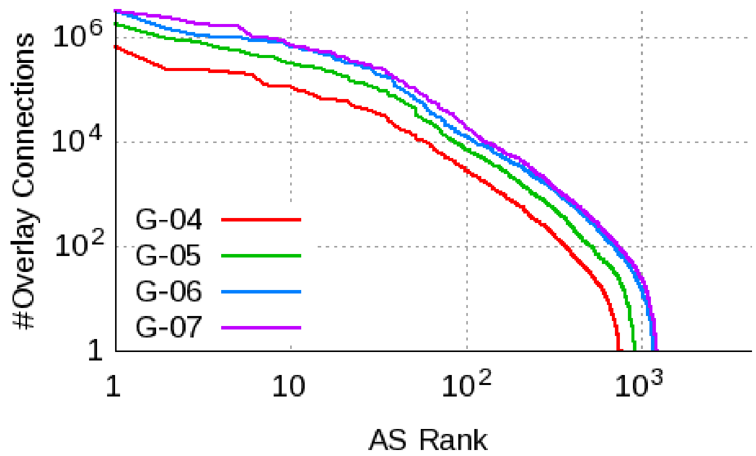

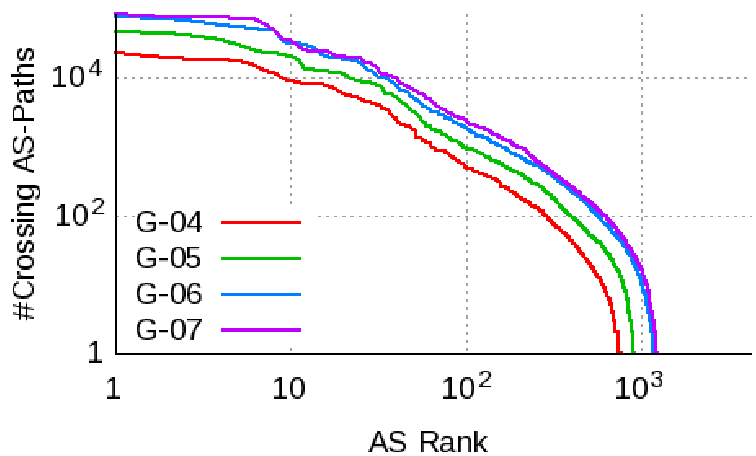

- Sorted histogram of the number of *overlay connections* passing through each transit AS
  - Similarly skewed for all snapshots
  - For G-07, each of the top-10 ASes carry more than 1M connections while each of the top-100 ASes carry more than 10k connections
  - Similar distributions despite changes in overlay and underlay
- Sorted histogram of the number of unique *AS-Paths* passing through each transit AS
  - Similar to previous graph

Load Distribution (3)

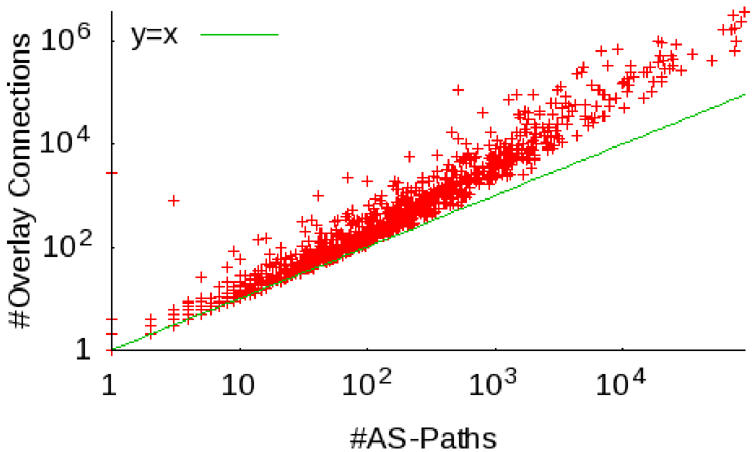

- Why the distribution of load does not change despite significant changes in over- and underlay?
  - The stability of main edge Ases, or
  - The constraint of valley-free routing through the hierarchical structure
- Scatter plot for number of overlay connections vs the number of AS-paths crossing through each AS
- Suggests that the load on each transit AS mostly depends on the number of crossing AS path
  - Its location in the AS hierarchy

Top transit ASes: Identity & Evolution

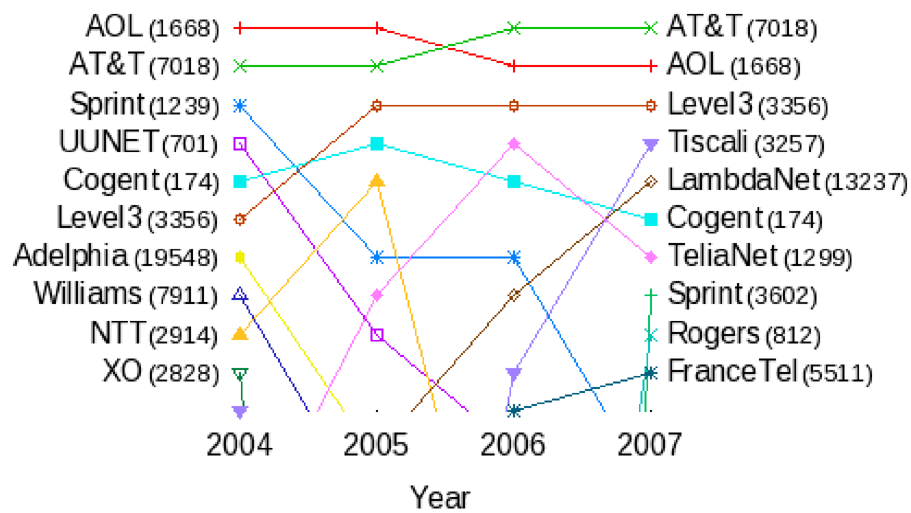

- Top-10 transit ASes carrying the largest number of overlay connections during 2004-2007
  - 4 ASes remain in top-10 all the time
  - Other ASes change rank due to changes in the following
    - Peer location & connectivity
    - AS-level topology & routing policy
  - Exact root causes for changes in AS ranking is hard to determine

AS-Path length

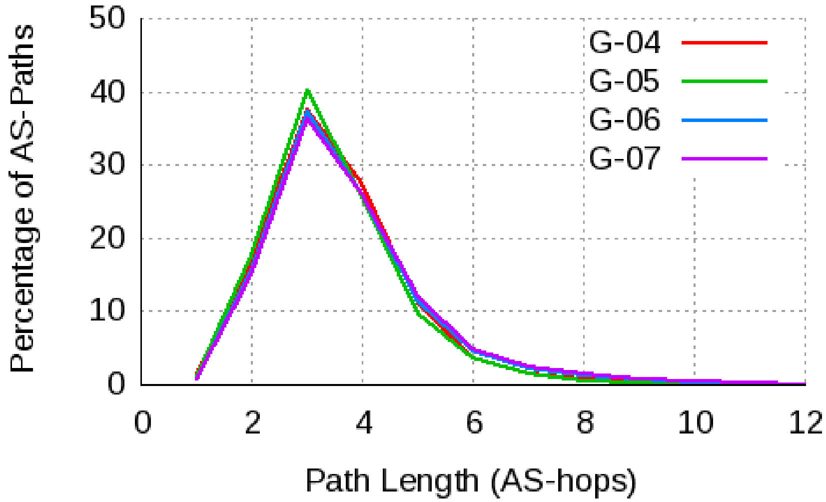

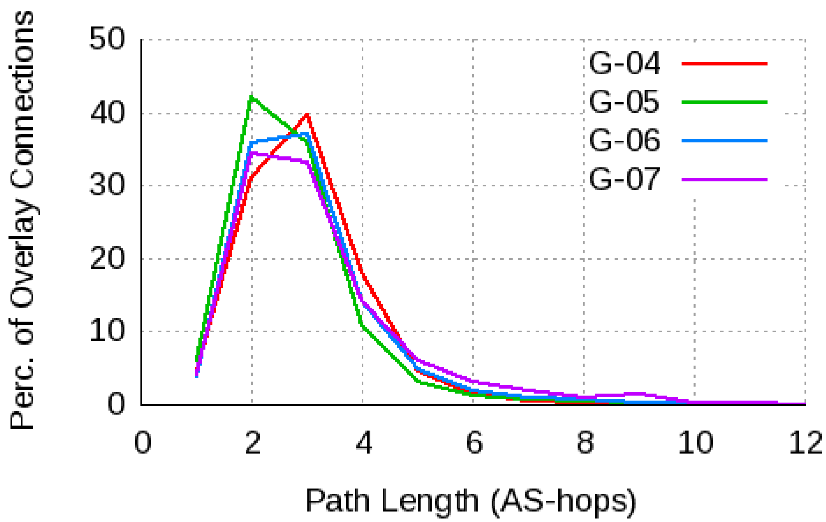

- Empirical density of the length of AS-Paths between all pairs of edge ASes
  - No significant change
  - 40% of paths are 3 hops long
  - 80% of paths are 4 hops long or shorter
- Empirical density of the AS-Path length across all *overlay connections*
  - Weighted version of the top graph
  - Similar pattern with shorter average path length
  - For G-07 mean un-weighted path length is 3.7 while mean weighted path length is 3.2
  - **A higher fraction of overlay connections are connected via shorter paths.**

Load Propagation over AS hierarchy

Snap.	Tier-1		Tier-2		Tier-3	
	Path	Conn	Path	Conn	Path	Conn
G-04	51	84	46	16	2.4	0.0
G-05	59	73	38	27	3.0	0.0
G-06	52	64	38	36	10	0.0
G-07	55	63	41	37	3.6	0.1

- Percentage of AS-Paths/overlay connections reaching tier-N.
  - ~50% of the AS-Paths reach tier-1 and ~40% reach tier-2.
  - Compared to paths, a larger fraction (84%-63%) of overlay connections reach tier-1; more connections are mapped to tier-1 paths.
  - A smaller fraction (16% - 37%) of overlay connections reach tier-2
  - The percentage of connections reaching tier-1 ASes is decreasing and perc. of connections reaching tier-2 is increasing over time.
    - Most likely due to the increased peering relationships at lower tiers

Conclusion & Future Works

- Leveraging a set of best practice, we presented a methodology to assess the global load imposed by an overlay on the underlay
  - The methodology can incorporate more accurate techniques & datasets
- Using this methodology, we characterized the impact of a real-world P2P overlay on the underlay.
- Presented results provide a better understanding of how overlay traffic is mapped on the underlay
  - Traffic is getting dispersed from the core due to increasing ISP peering.

Future works:

- Detailed sensitivity analysis for main components
  - Overlay connectivity and load
  - Accuracy of underlay topology, AS relationship & routing policies
- Examining the effect of geographical footprint of the overlay.
- Incorporating inter-ISP pricing model and perform cost analysis.