

A Systematic Framework for Unearthing the Missing Links:

Measurements and Impact

Yihua He

Georgos Siganos

Michalis Faloutsos

Srikanth Krishnamurthy

University of California, Riverside

Apr 12, 2007



Motivation

- A complete Internet topology at AS level is important when we try to:
 - Create more accurate models for simulation. (“Realistic bgp traffic for test labs”, Sigcomm 2002)
 - Analyze/design next generation inter-domain routing protocols that take advantage of the topology properties. (“HLP”, Sigcomm 2005)
 - Study the spread of epidemics, such as worms. (“The effect of network topology on the spread of epidemics”. Infocom 2006)



However...

- The topology that people use is not complete
 - Most people only use AS topologies derived from AS_Path attribute in BGP RIBs.
 - And many people only use RIBs from Routeviews
- This is ok in some cases, however...
- The probed topology is probably biased: the number of peer-to-peer links between ASes is probably under-estimated:
 - Current inter-domain routing is policy-based
 - For an AS, routes learned from its peers will not be advertised to its providers.



So what do we want to do?

- We try to find missing links between ASes
 - How to find
 - Where to find
 - What properties do these missing links have
 - How could we possibly find more missing links
- What's the impact of these links? Are these missing links going to change our view of the Internet?
 - Topological impact
 - "Financial" impact



Some of the Highlights

- We find more AS links
 - Find 40% more AS links
 - Find 300% more peer-peer AS links
 - Note all these new found links are either in the BGP RIB snapshots or verified by our traceroute tool (RETRO)
- Locate the missing AS links
 - Most missing AS links are peer-peer AS links
 - Most missing AS links are at IXPs
- Evaluate the impact of the missing links
 - Topologically: let's settle the power-law debate
 - Financially: "middle-class" ASes get richer



Tools built and used

- IXP participants identifying tool
 - Infer IXP participants' AS numbers by IXP IP addresses and web-achieves.
 - Heuristics to “guess” IXP participants from a variety of sources such as traceroute data and whois/dns records.
 - Good accuracy (80%-90%), low false positives (<15%).
 - Used to infer a superset of possible IXP edges.
- RETRO (REverse Trace Route)
 - By using public traceroute servers
 - From 1200+ vantage points, 350 ASes and 55 countries
 - Used to verify if a suspected link really exists or not

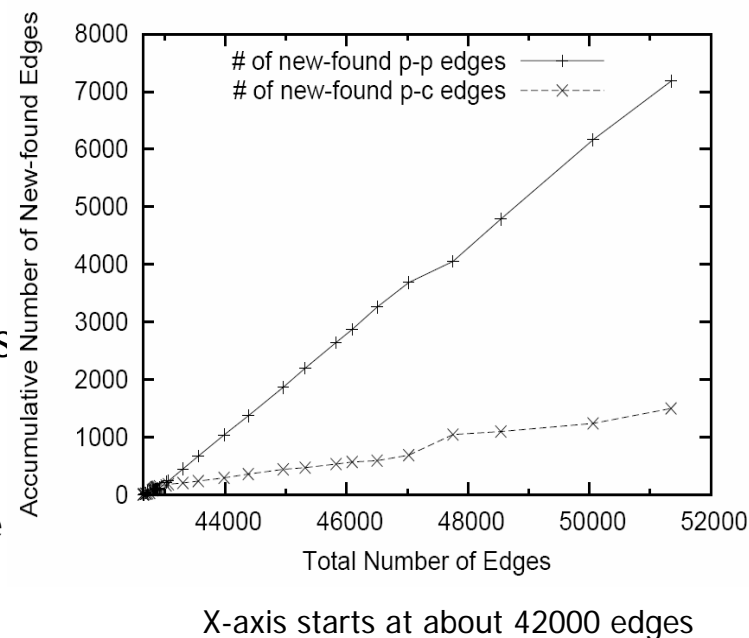


Roadmap

- On finding missing AS links
 - Adding more BGP dumps
 - Exploring IRR
 - IXPs and missing links
 - Validation by RETRO
- Impact on the Internet
 - Topologically
 - Financially

Adding more BGP dumps

- 33 more BGP table dumps, from:
 - Other Oregon route collectors
 - RIPE RIS: www.ripe.net/ris
 - Public route collectors: www.cs.ucr.edu/bgp
 - One snapshot per dump
- Peer-peer edges and provider-customers edges increase with different rate:
 - P-C edges are major in any graphs observed
 - The more complete graph, the more percentage of p-p edges
 - More than 80% new discovered edges are p-p edges
- **Most missing AS links are peer-peer type**



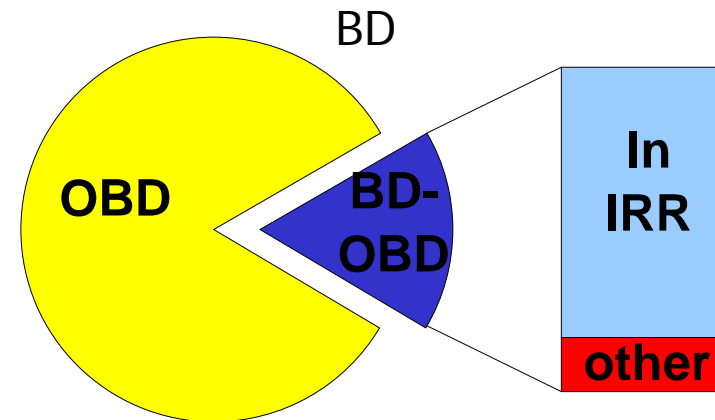


Exploring IRR

- IRR=Internet Routing Registry
- ASes register their routing policies with their neighbors in IRR
- Nemecis (G. Siganos et al. Infocom 2004) is a tool to process IRR, as well as infer AS relationships for the AS links in IRR
- But IRR is known to be not perfect:
 - IRR may not be updated
 - IRR may contain errors/typos
 - IRR is not complete
- To be confident, we will not just combine the links we find from IRR, even for those passed the Nemecis checks.
- Instead we try to get some hints from IRR.....

IRR: Good Hints for Missing Links

- If we only knew OBD, would IRR be a good data set to estimate what are missing from OBD?
 - 83% (7251 out of 8702) of the edges in BD and not in OBD are actually in IRR .
- IRR has many more (50k+) edges that we have not seen yet. Are they real?



OBD: Oregon Routeview BGP Dump

BD: OBD+33 other BGP tables



IXPs and missing links

- Infer IXP participants with our tool
 - Get IXP's IP blocks from pch.net
 - 110 IXPs, 2348 distinct participants
 - Some ASes participate in multiple IXPs
- Potential IXP edges: full mesh
 - This is a superset of real IXP AS edges
 - Total 141,865 distinct AS edges (IXPall)



IXPs and missing links (2)

Name	# of Edges	\cap <i>IXPall</i>	Perc.
<i>BD-OBD</i>	8702	6910	79%
<i>peerBD-OBD</i>	7183	6197	86%
<i>peerIRRnc-BD</i>	39894	23979	60%
<i>peerIRRdual-BD</i>	13905	11477	83%

- Many missing peer-peer links could be at IXPs!
 - 86% of the missed (in BD but not in OBD) peer-peer links are in *IXPall*
 - 83% of the possible missing (in *peerIRRdual* but not in BD) are in *IXPall*
- So are they really at the IXPs? Let's see...



Validation by RETRO

- Now we would like to verify them by traceroute
 - When we want to verify an AS edge [AB], we try to traceroute from AS A to AS B (or AS B to AS A)
 - We say AS edge [AB] does exist if we observe one of the following IP sequences in traceroutes:
 - $[IP_{AS_A}, IP_{AS_B}]$ ---direct link
 - $[IP_{AS_A}, IP_{IXP}, IP_{AS_B}]$ ---via some IXP



Validation by RETRO (2)

Table 7: RETRO verifies peer-peer links in IRR missing from BD

Name	# of Edges	# of RETRO candidates	# of confirmed peering		
			total	via IXP	direct
<i>peerIRRnc-BD</i>	39894	8791	5646	5317	329
<i>peerIRRdual-BD</i>	13905	4487	3529	3351	178

- Many suspects have been verified
 - 5646 out of 8791 (62%) for *peerIRRnc-BD*
 - 3529 out of 4487 (79%) for more reliable *peerIRRdual-BD*
- In confirmed AS edges, most of them are IXP edges!
 - 5317 out of 5646 (94%) for *peerIRRnc-BD*
 - 3351 out of 3529 (95%) for *peerIRRdual-BD*
- The result suggests most missing peer-to-peer links are indeed at IXPs



Summary

- Most missing AS links are peer-peer type
- Many missing AS links are in IRR
- Most missing AS links are at IXPs




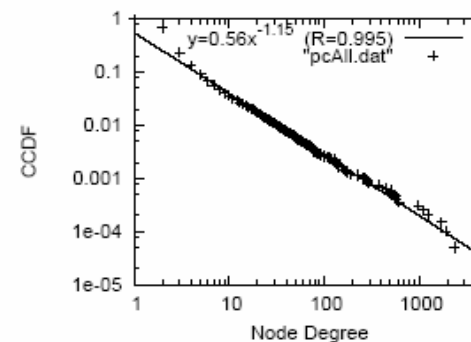
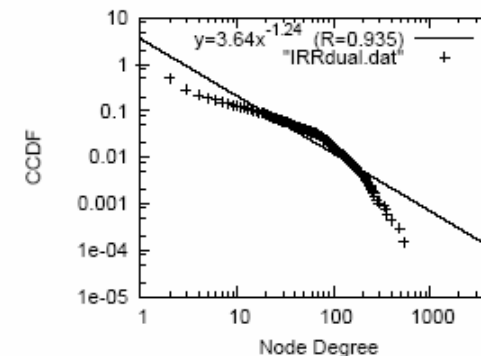



Roadmap

- On finding missing AS links
 - Adding more BGP dumps
 - Exploring IRR
 - IXPs and missing links
 - Validation by RETRO
- **Impact on the Internet**
 - Topologically
 - Financially

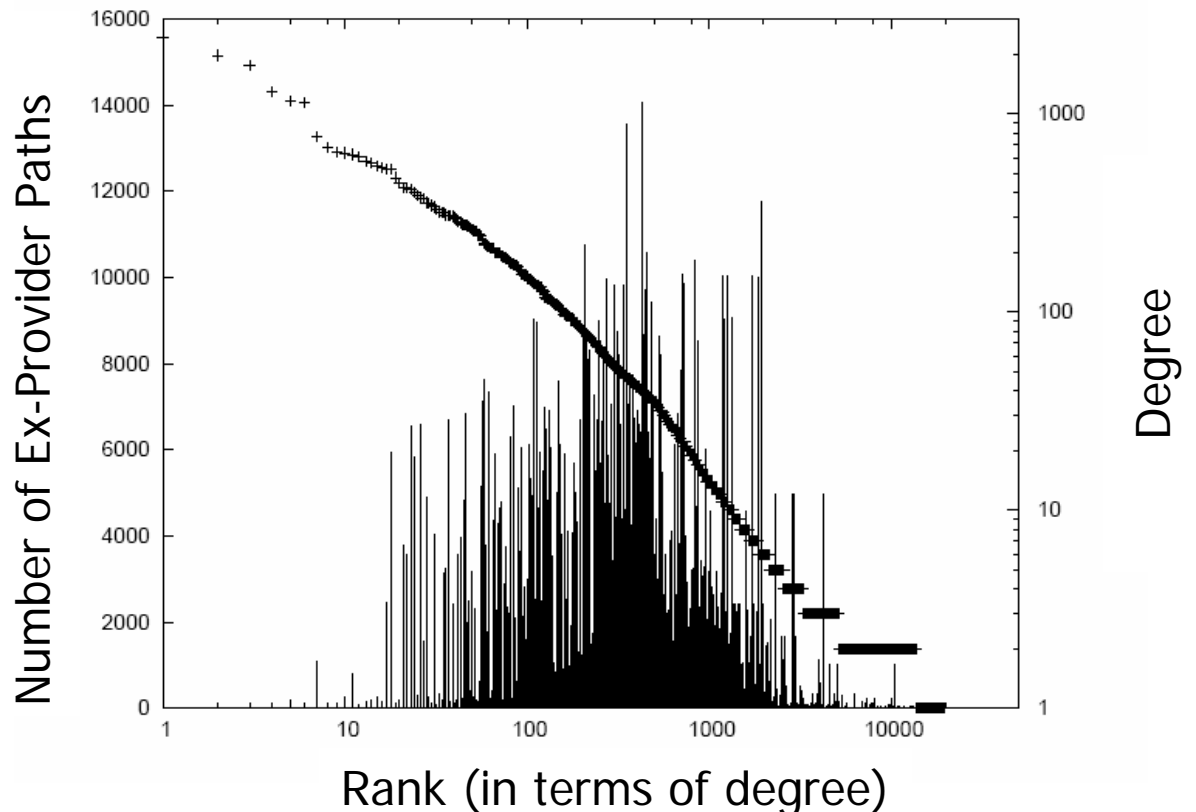
Power-law debate

- The power-law debate
 - By adding sources other than BGP tables, the degree distribution does not look like a strict power-law [infocom 2002]
- We try to settle it by modeling P-P and P-C edges separately
- No matter what the sources are:
 - P-C links follow a power-law distribution
 - P-P links follow a Weibull distribution
 - All with higher than 99% correlation coefficient.



The effect of the new edges on investigating ISP revenue

- Peer-to-peer links make ASes avoid using their providers to reach many destinations
- For many ASes, more than 50% of the paths stop going through their providers.
- By stop going to providers, ISPs lower their costs and increase their revenues.
- The “middle-class” ASes (rank 100 to 1000) benefit the most from these changes



The number of **ex-provider paths** is the number of paths for which an AS stops going to its providers once the new edges are added



Conclusion and future work

- Claim:
 - Synthesize many sources: traceroute, IRR, tables, IXPs, and cross-verify them
 - Significant change in AS model (pp links)
- Usage:
 - Improve the quality of topology generation/reduction
 - Improve AS policy inference
 - Improve understanding routing security and robustness
- Are we done?
 - Exploit the location and patterns we identified
 - Search for hidden links: Backup links, active bgp measurement



Visit our website

- www.cs.ucr.edu/~yhe/LordOfLinks/
 - Data, talks, FAQs
 - Tools will be organized and added shortly
- Since Jan 2007, 100+ visits, 30+ downloads (from different domains)
 - MSR, IBM, UIUC, UMich, UCLA ...
- More tools and data will be available in the future.



Thank you!

Questions?