Discovering Interdomain Prefix Propagation using Active Probing

Lorenzo Colitti

lorenzo@ripe.net - colitti@dia.uniroma3.it

The problem

Prefix propagation

- Focus: an AS *Z*
- *Z* would like to know how the rest of the Internet treats its prefixes

- Motivations:
  - Predict the effect of network faults
  - Perform effective traffic engineering
  - Develop peering strategies
  - Evaluate QoS provided by upstreams
  - ...

- For simplicity, we consider an arbitrary prefix *p* in *Z*

Per-prefix discovery

- Easy!
- Get paths from RIS and ORV, merge into a graph

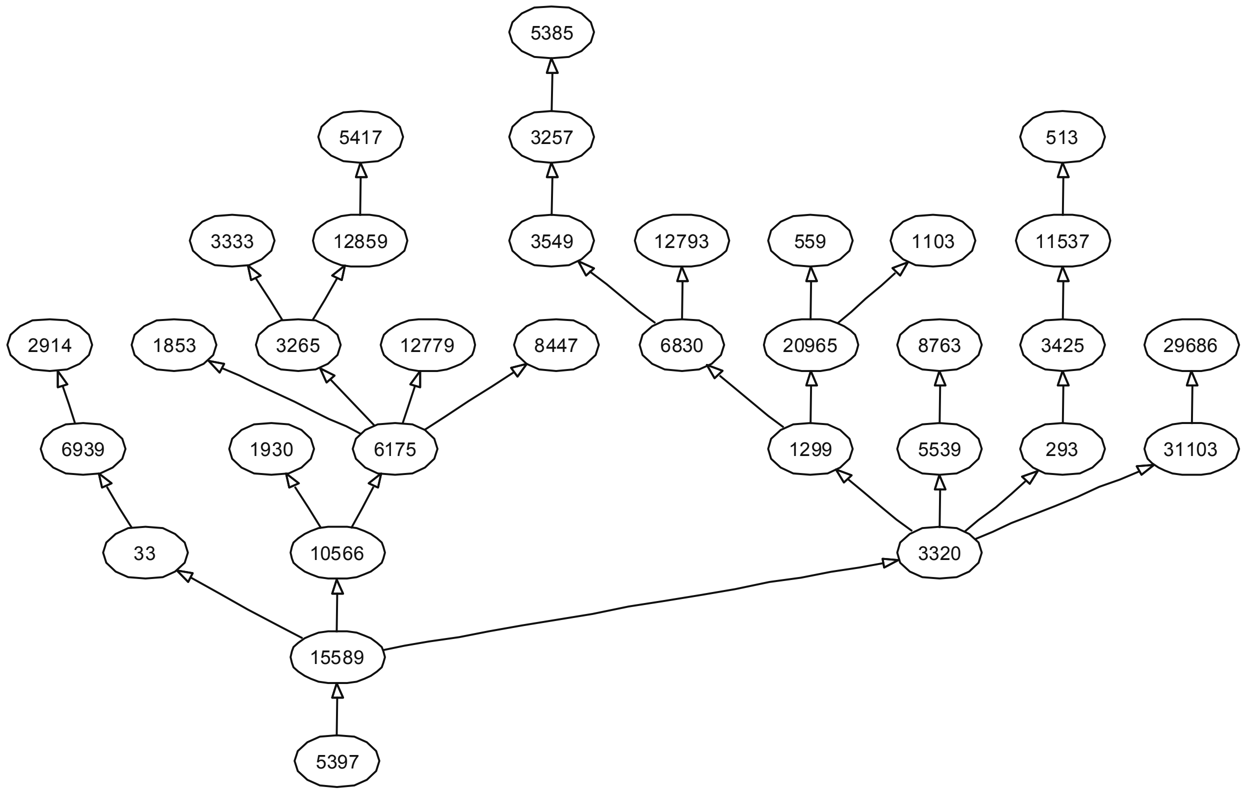

What about alternate paths?

- These are the paths used by the network to reach *p* when the discovery is performed
- But which paths **could** be used in other conditions?
  - Obviously many more
- For example, what would happen:
  - In case of network faults?
  - If we did inbound traffic engineering?
  - ...

Discovering AS adjacencies

- Another approach:
  - Take all paths seen by RIS and ORV for all Internet prefixes
  - Merge into an interdomain graph
  - Examine portion of graph around Z

- We obtain a subset of the interdomain topology around *Z* at the time of exploration

But what about policies?

- BGP is not about routing, it's about policy
- We have the interconnections between AS
- But without policies, we know nothing!
  - Like having a city map without one-way streets
- The topology itself tells us nothing on which paths can be used to reach Z

Idea

- By manipulating BGP announcements for *p* we can force the network to use alternate paths
- We can then use per-prefix discovery methods
  - We reveal alternate paths
    - ... but only those that can be used to reach *p*
- This is what we wanted

Methodology

Feasibility

- An AS-path is **feasible** for *p* if “the routing policies of the Internet allow the announcement of *p* with that path ”
  - Active paths
  - Backup paths
  - Alternate paths

- A peering between two ASes is feasible for *p* if it’s part of an AS-path that is feasible per *p*
  - i.e. if it’s possible, in some state of the Internet, that traffic to *p* could flow through that peering

Feasibility graph

- Directed graph, nodes = ASes, arcs = feasible peerings

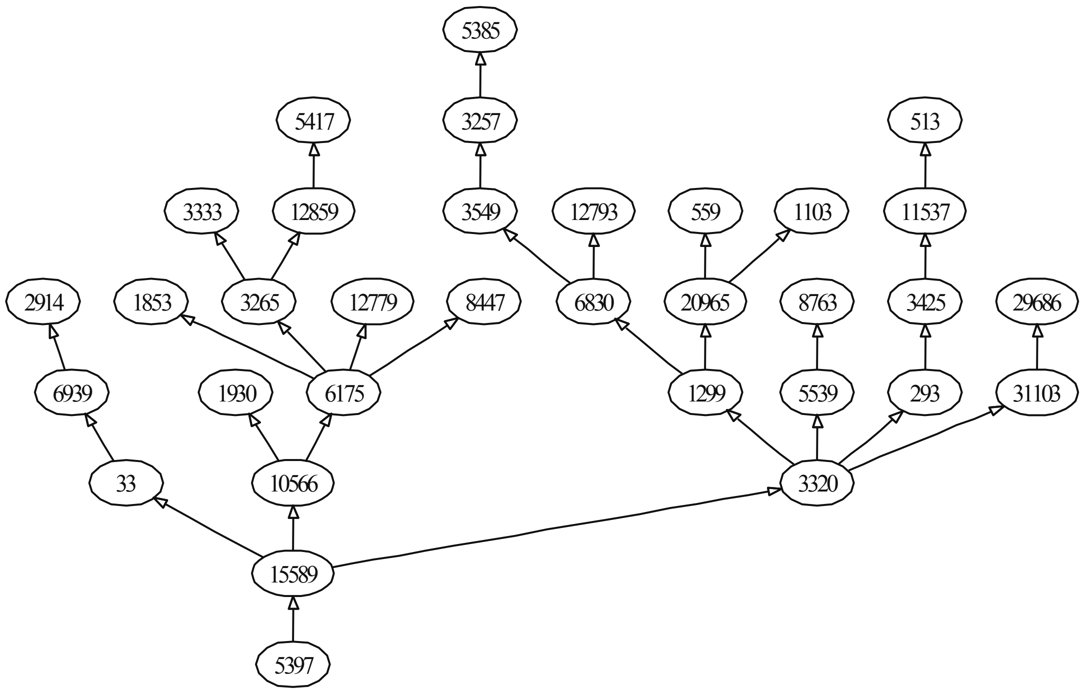

- Shows us [a subset of] the portion of interdomain topology involved by traffic flows to *p*

Methodology

- Basic idea:
  - Send BGP announcements
  - Observe the results using RIS, ORV, looking glasses, etc.

- Two primitives
  - "Withdrawal observation"
    - Sends a withdrawal and observes convergence
      - BGP explores alternate paths
  - "AS-set stuffing"
    - Prohibits an announcement from being propagated by certain ASes by putting them in the path
      - Forces BGP to choose alternate paths

Withdrawal Observation

- When a prefix is withdrawn, BGP explores alternate paths before concluding that it is unreachable
- So:
  - Withdraw *p*
  - Observe convergence process
  - Record all alternate paths chosen by BGP
  - Merge all paths into a feasibility graph
- Rapidly obtains a rich feasibility graph

AS-set

- Normally used in route aggregation
- Indicates that information on who exactly originated a certain announcement was lost
- e.g.:
  - L'AS 701 has customers AS1, AS2, AS3
  - The three customers have contiguous address space
  - AS 701 can aggregate the three announcements in one

$$\left.\begin{array}{l}701\ 1\\701\ 2\\701\ 3\end{array}\right\}\Rightarrow 701\ \{1,2,3\}$$

AS-set stuffing

- If an AS receives an announcement with its own number in the path, it discards it to prevent routing loops
- We can stop an announcement from traversing a given AS by putting that AS in the path
  - If we use an AS-set, path length does not change
  - The announced paths end in ... $Z \{A_1, ..., A_n\}$
- As far as *p* is concerned, it's as if the ASes A_i had been eliminated from the topology
  - We name the ASes A_i "prohibited"

Applications

Applications

- Withdrawal observation allows:
  - Topology discovery
    - Faster than AS-set stuffing
- AS-set stuffing allows:
  - Topology discovery
  - Path feasibility determination
  - Path preference comparison
  - Measuring performance in alternate routing states
    - What-if studies on Internet routing

Topology discovery

- Objective: discover ASes and peerings not ordinarily visible
  - Simple algorithm using AS-set stuffing: "level-by-level exploration"
    - We name ℓ ("level") of an AS the topological distance from Z
    - Start from Z with increasing values of ℓ
      - Prohibit all ASes at distance ℓ
      - Merge all paths discovered into feasibility graph
    - If new ASes at distance ℓ have appeared, prohibit them
    - Otherwise, proceed with level $\ell + 1$
  - Or just use withdrawal observation

Example: prohibit level 2

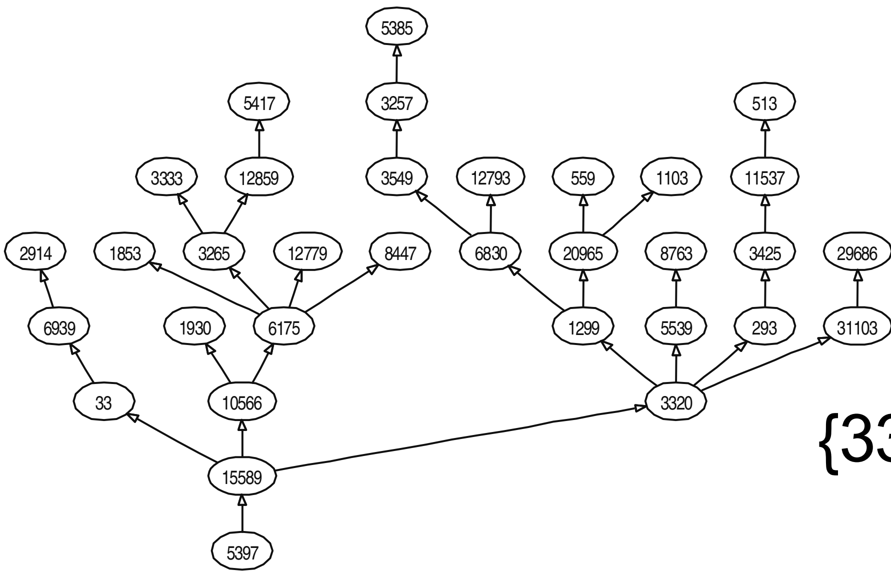

{33,3320,10566}

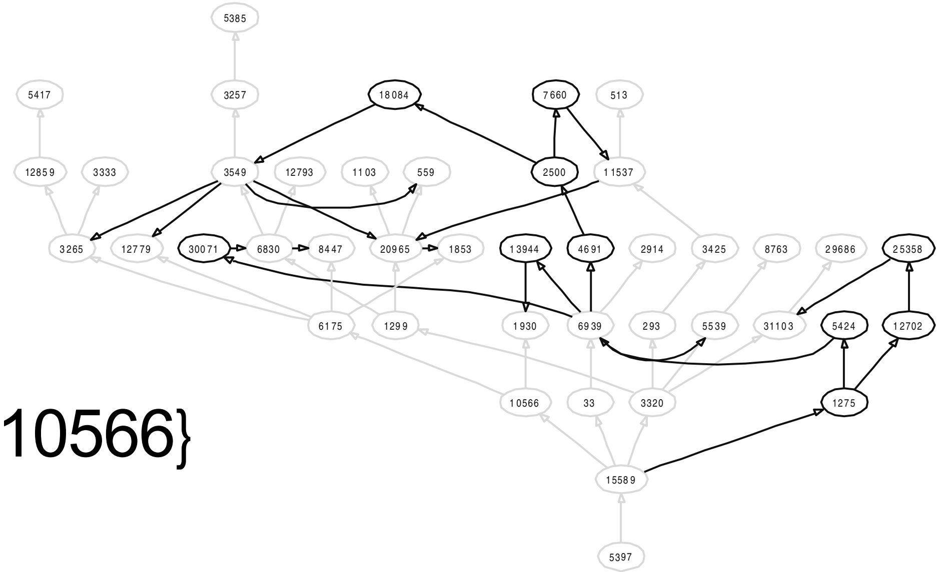

32 ASes
33 peerings

42 ASes
57 peerings

After four levels

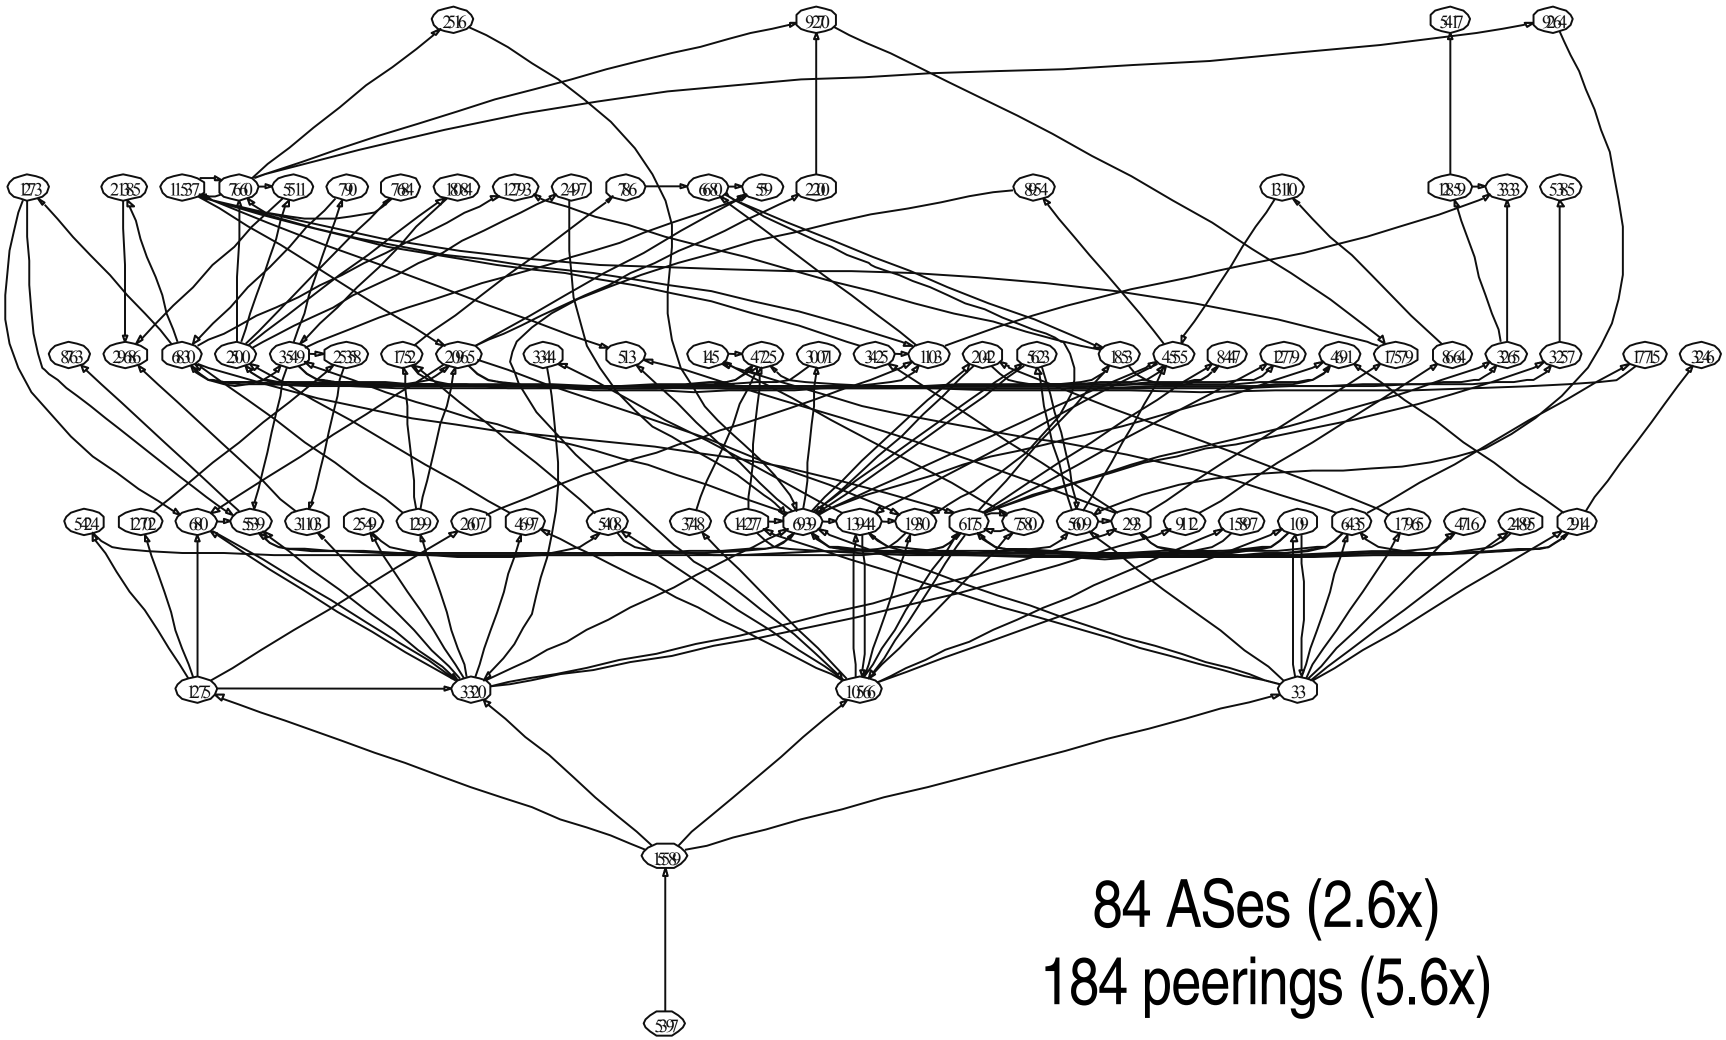

84 ASes (2.6x)
184 peerings (5.6x)

Path feasibility determination

- Route collector *C* sees path ***ZFGC***
- Is path ***ZADC*** feasible?
- Z announces {*B*,*F*,*G*}

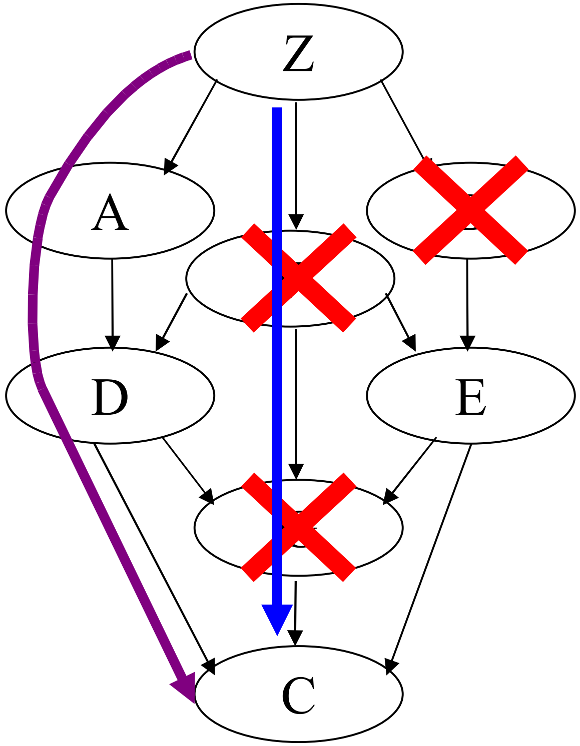

- If *C* sees ***ZADC***, ***ZADC*** is feasible
- If *C* does not see any path, ***ZADC*** is not feasible

Path preference determination

- *C* sees ***ZFGC***
- ***ZADC*** and ***ZBEC*** are also feasible
- Which does C prefer?

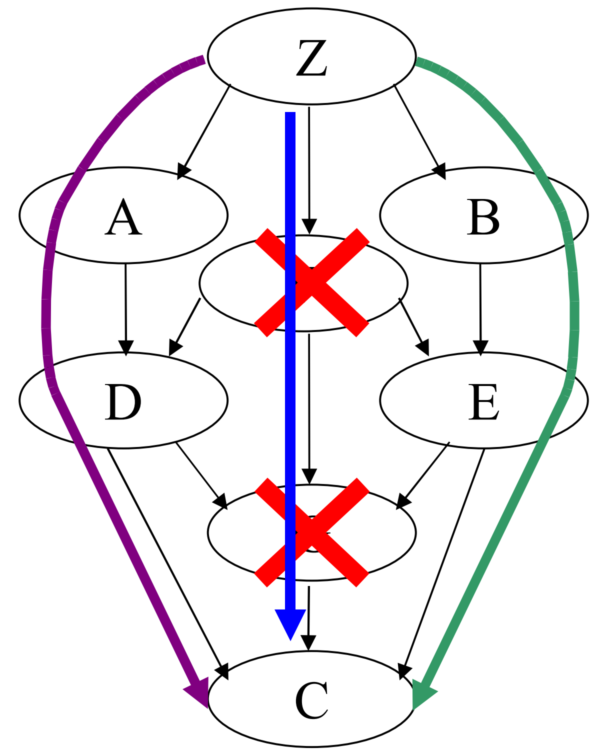

- Z announces {*F*,*G*}

- C's best path is the path it prefers

Measurements in altered routing states

- Routing changes made with AS-set stuffing are steady-state
- This enables “what-if” analysis of performance
- “How would performance change if we used ISP A instead of B?”
  - Use AS-set stuffing to change the topology
  - Then measure performance
    - Even ping could suffice
      - Outgoing path stays the same

Results

Testing and evaluation

- We first tested our techniques on the IPv6 Internet
  - IPv6: Nov 2004 – Feb 2005 (CASPUR)
  - IPv4: Jun 2004 – Jul 2004 (RIPE NCC)
- Lessons learned:
  - Interdomain routing is a sensitive topic
  - Wear a flame-proof suit

Topology discovery: results

- Both methods are significantly better than stable state routing
- The topologies produced by AS-set stuffing are slightly richer

Method	IPv6		IPv4	
	AS	Peering	AS	Peering
Stable state	32	31	24	23
Withdrawal	94 (2.9x)	211 (6.8x)	28 (1.2x)	49 (2.1x)
AS-set	97 (3.0x)	222 (7.2x)	29 (1.2x)	55 (2.4x)

New ASes and peerings found in the discovery process

Topology discovery: results by level

IPv6

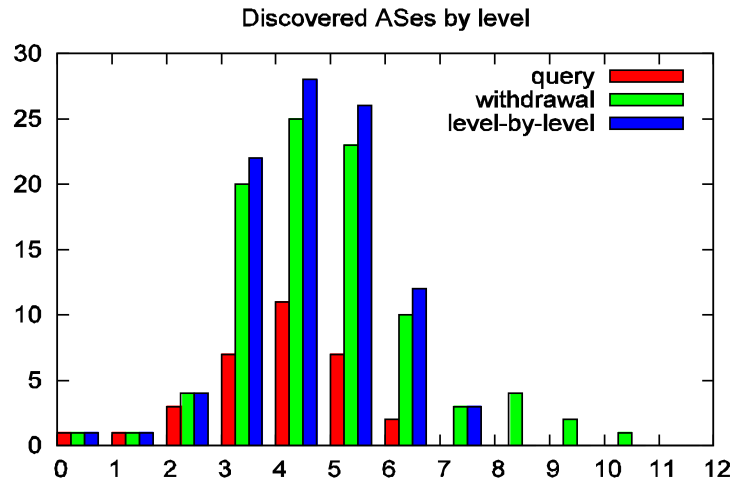

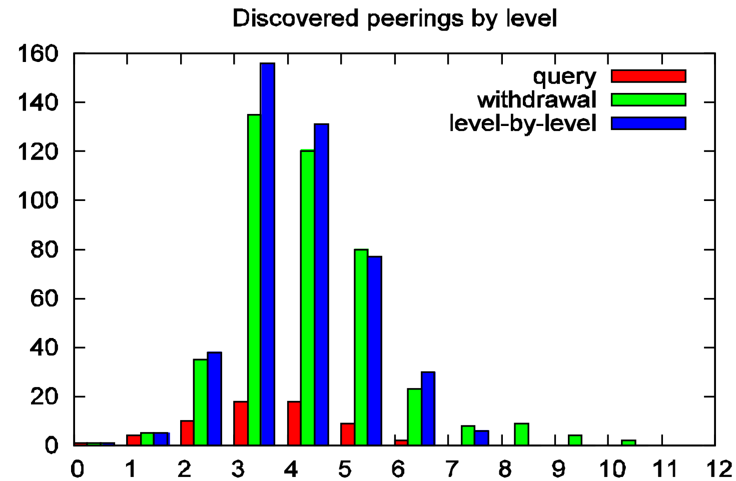

IPv4

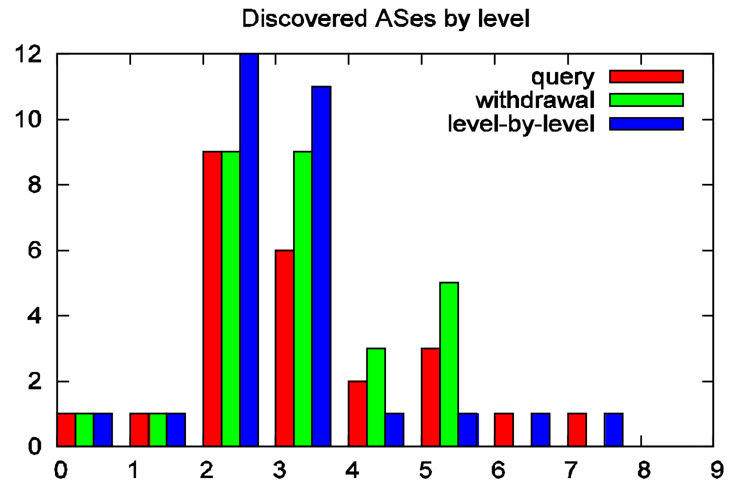

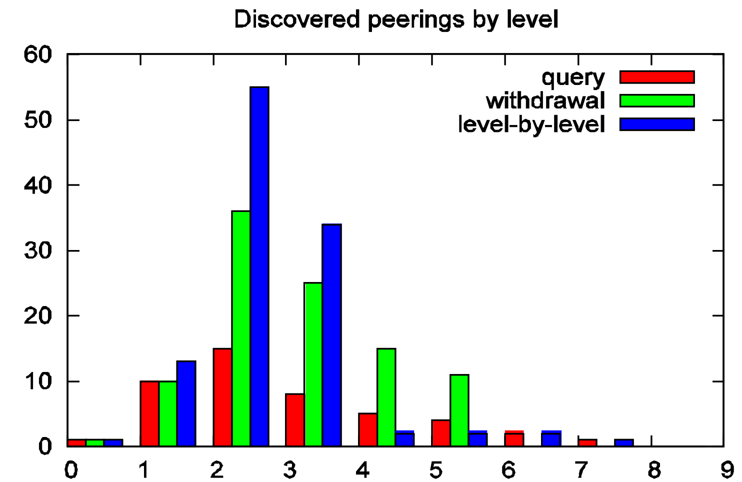

Comparison with existing methods

- We compared our per-prefix discovery methods to per-path methods
  - *W* is a feasibility graph obtained using withdrawal observation
  - *C* obtained by fusing all AS-paths from ORV
    - Seen at time withdrawal was made
  - *I* is the graph induced on C by the nodes in *W*
  - Compare number of arcs seen

Comparison with existing methods: results

- *W* has ~ 50% (IPv6) or 25% (IPv4) of the arcs in *I*
- The topology of the complete graph is much richer
  - This is not a weakness, it is a strength!
    - We only discover feasible peerings
    - *Z* has little interest in peerings that are not feasible
- We also discovered a few arcs not seen in C
  - Probably backup paths only seen during convergence

Comparison with existing methods: details

IPv6

I	W	I only	W only
312	158 (51%)	175	21 (13%)
334	168 (50%)	189	23 (14%)
302	154 (51%)	174	26 (17%)

IPv4

I	W	I only	W only
241	61 (25%)	181	1 (2%)

Conclusions

Conclusions

- Existing methods do not permit the discovery of alternate paths that could be used in case of faults or routing changes
- Our methods allow an ISP to:
  - Discover alternate paths
  - Partially deduce other ASes' routing policies
  - Measure performance in alternate routing states
- Testing on the IPv4 and IPv6 Internet shows they are effective

Bibliography

- L. Colitti, G. Di Battista, M. Patrignani, M. Pizzonia, M. Rimondini, *Discovering Interdomain Prefix Propagation Using Active Probing*, to appear in Proc. IEEE ISCC 2006, June 2006
- L. Colitti, G. Di Battista, M. Patrignani, M. Pizzonia, M. Rimondini, *Active BGP Probing*, Technical Report RT-DIA-102-2005, November 2005

http://www.dia.uniroma3.it/~compunet/bgp-probing/

compunet@dia.uniroma3.it

Questions?