

On the Completeness of Observed AS-level Internet Topology

Presented by:

Hyunseok Chang

EECS Department

University of Michigan

hschang@eecs.umich.edu

<http://topology.eecs.umich.edu>

Inferring AS-level Internet Topology (or AS graph)

Internet modeled as a graph:

- Node: Autonomous System (AS)
- Edge: Pair-wise AS peering relationship & Direct inter-AS physical connection.

Inferred from a set of AS paths found in BGP routing tables.

12 public route servers including Oregon route server provide public access to their BGP routing tables.

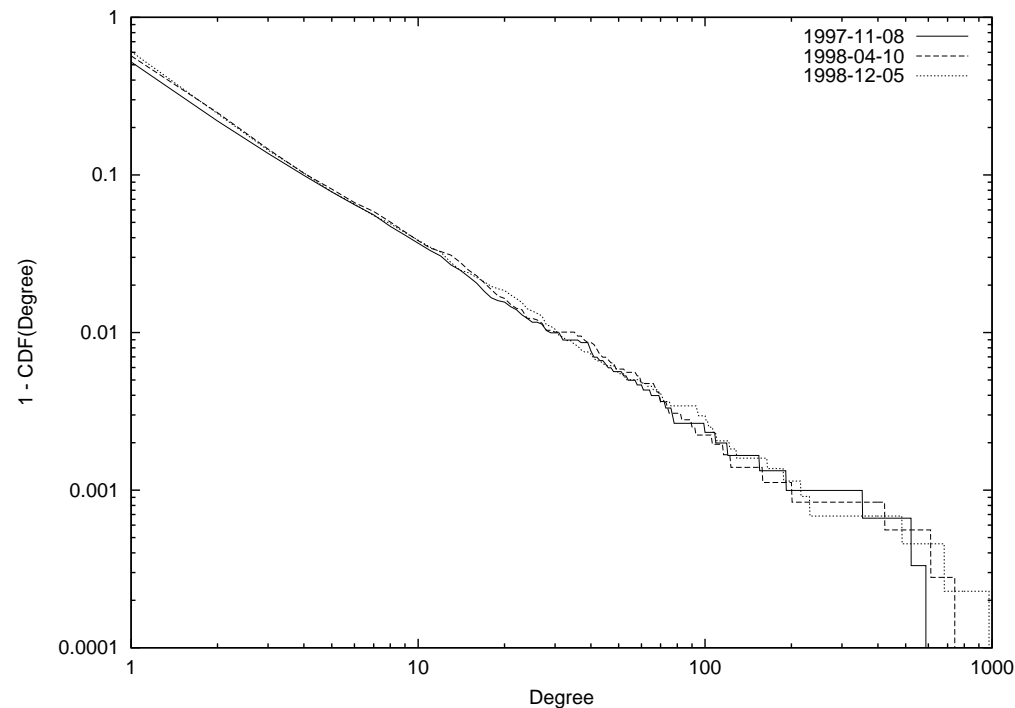
Inferred AS graph

Source: BGP routing tables collected by Oregon route server since Nov. 1997.

- Originate from *geographically-dispersed* and *heterogeneous* ASs.

Exhibits Highly-variable node degree distribution expressed as:

$$f_d \propto d^{-\alpha} \quad (d=1,2,\dots; \alpha: \text{const}).$$



How complete is the *inferred* AS graph?

Problem: We don't have the *complete* AS graph to compare the *inferred* one against.

Approach: instead, look at how well the local connectivity of *individual* ASs is observed by any *other* ASs.

Local connectivity of individual ASs obtained from:

- their BGP routing tables,
- their Looking Glass data,
- Routing Registry database.

Assumption: Now we know their *complete* local connectivity info.

Local vs. Non-local BGP view from BGP routing tables

Let's define a *BGP view* as an instance of AS graph constructed from a BGP routing table.

Then, let's define *local* BGP view and *non-local* BGP view with respect to a given AS.

Now, for each AS, look at what portion of its connectivity revealed by its *local* view is being observed in its *non-local* views.

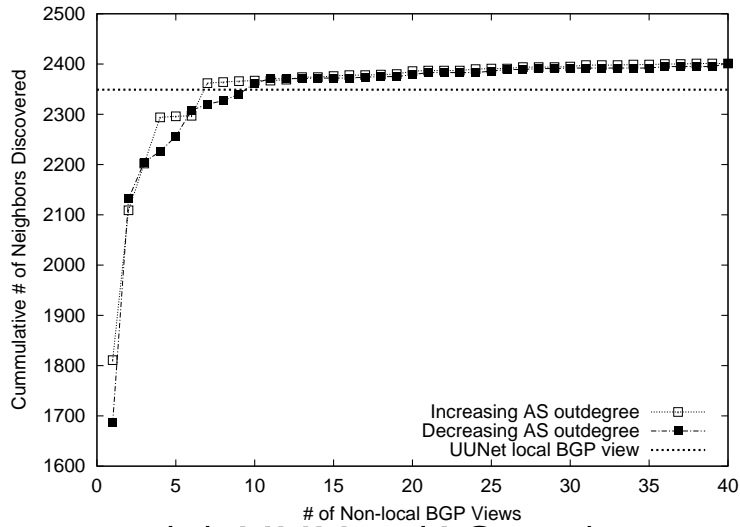
We collected BGP routing tables of 41 different ASs from public route servers:

- North America: 20, Europe: 19, Asia/Pacific: 2
- Backbone: 14, ISP: 24, Etc: 3

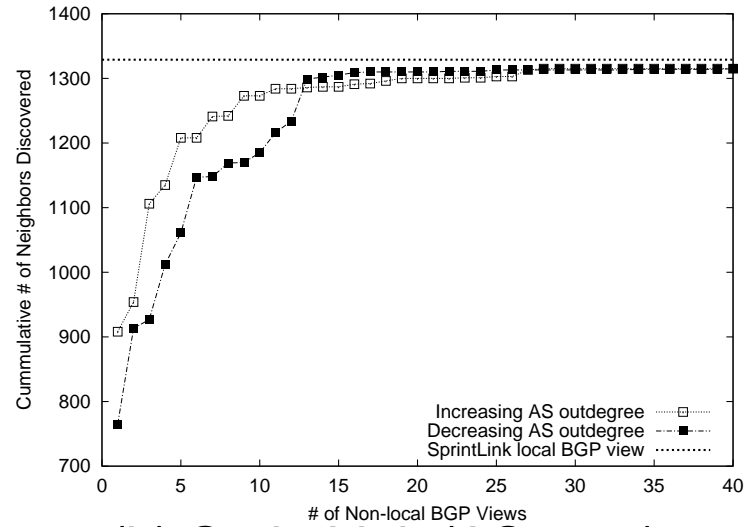
Observations from local/non-local view analysis

- The connectivity of Tier-1 ASs more easily observed than that of non-Tier-1 ASs.
- The connectivity of an AS better observed by its *customer* ASs than by its *peering* ASs.
- The connections between an AS to its upstream ASs easily observed.
- BGP routing tables as a means to infer Internet topology should be collected from end-customer ASs locating in the lowest hierarchy.

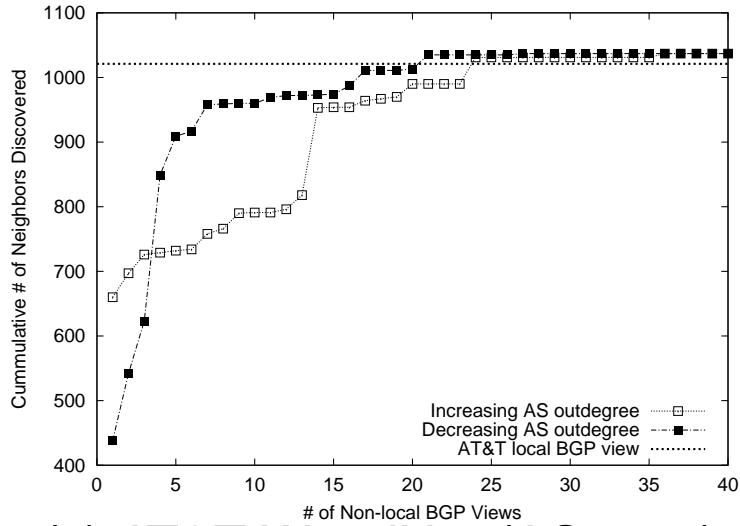
Tier-1 ASs



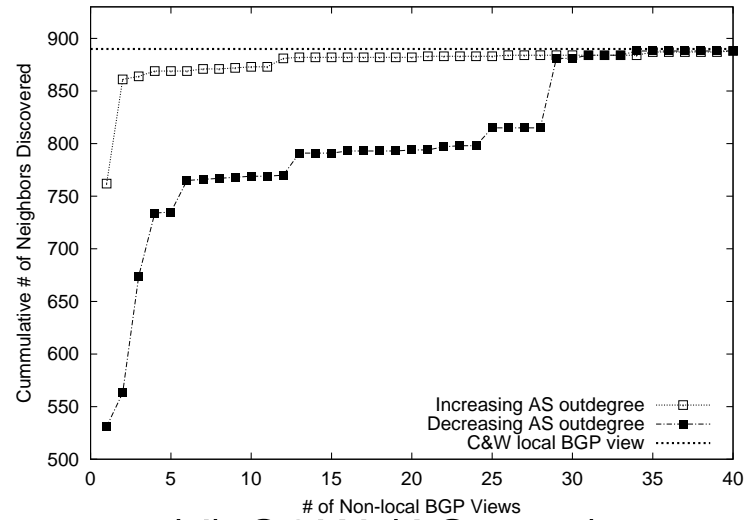
(a) UUNet (AS701)



(b) SprintLink (AS1239)

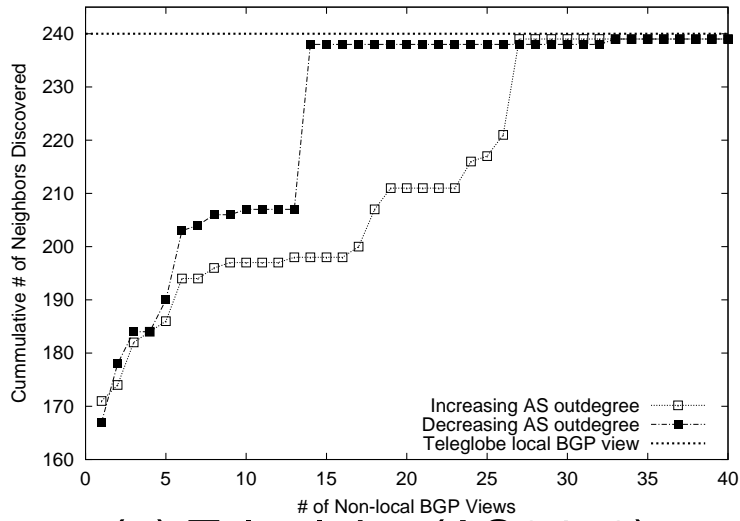


(c) AT&T WorldNet (AS7018)

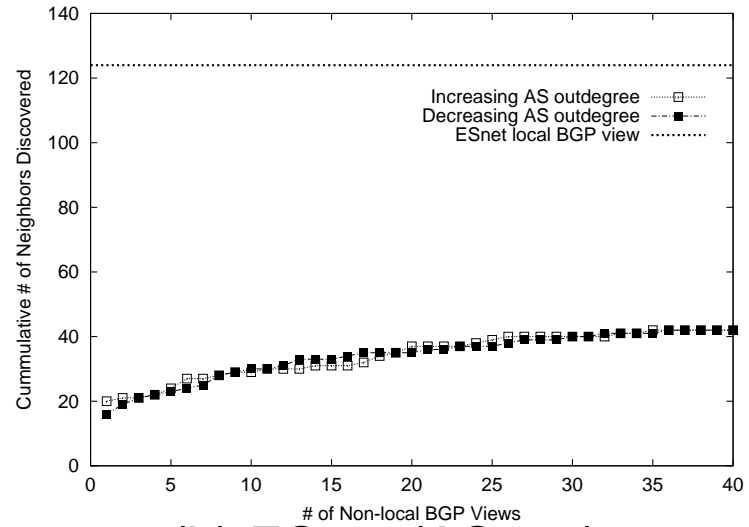


(d) C&W (AS3561)

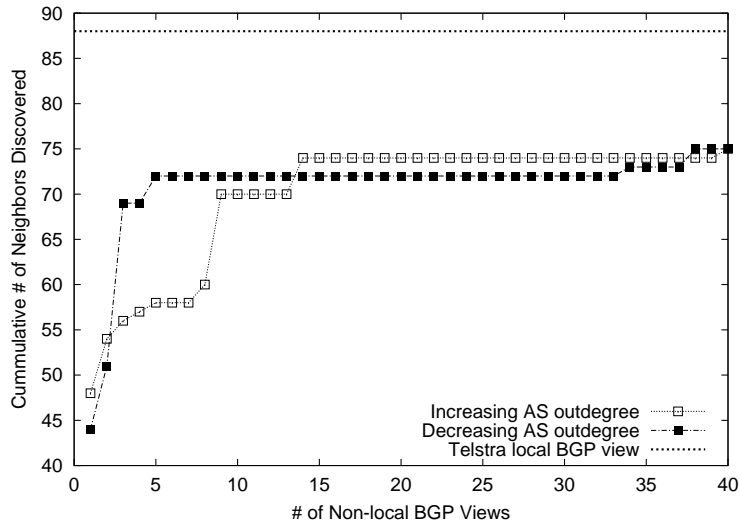
Tier-2 ASs



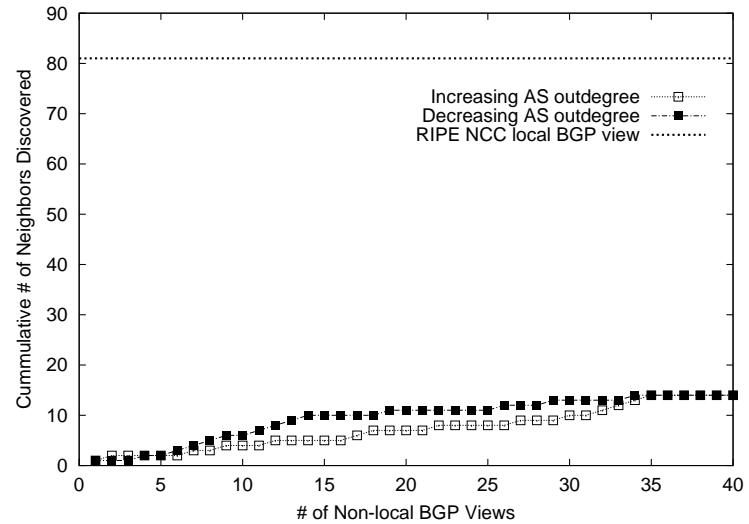
(a) Teleglobe (AS6453)



(b) ESnet (AS293)

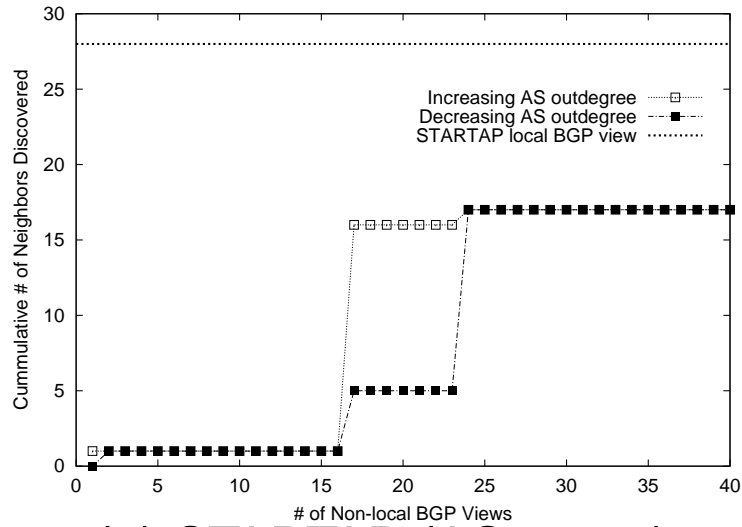


(c) Telstra (AS1221)

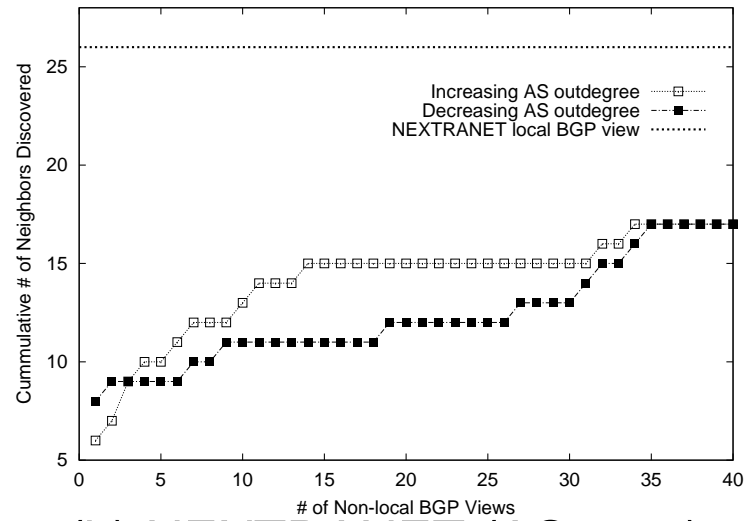


(d) RIPE NCC (AS3333)

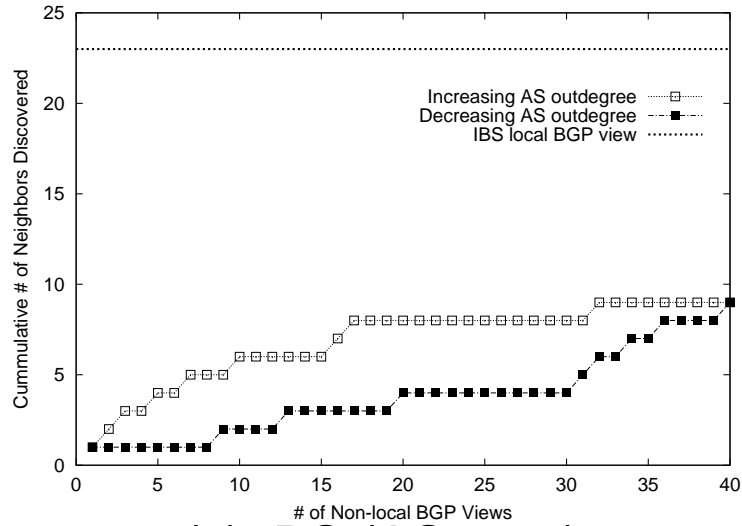
Tier-3 ASs



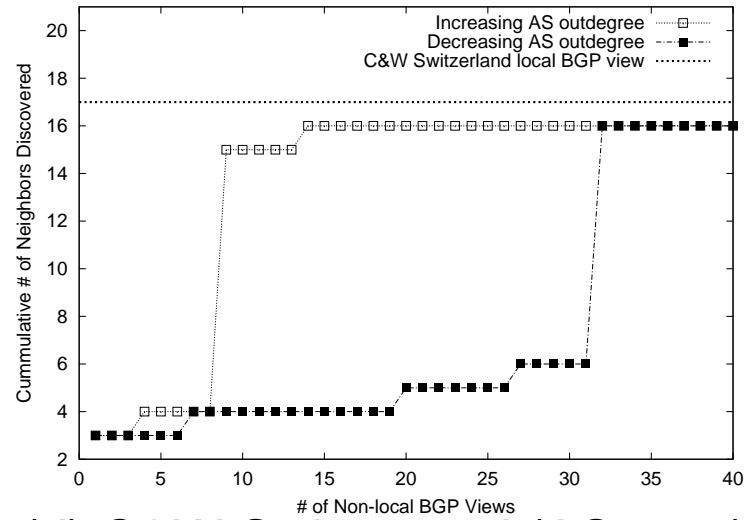
(a) STARTAP (AS10764)



(b) NEXTRANET (AS9177)



(c) IBS (AS8271)



(d) C&W Switzerland (AS6893)

Looking Glass tool

Used by individual ISPs to troubleshoot Internet-wide routing problems.

- E.g. Prefix reachability, policy routing debugging, BGP activity monitoring.

Public, but limited access to selected border routers or route servers.

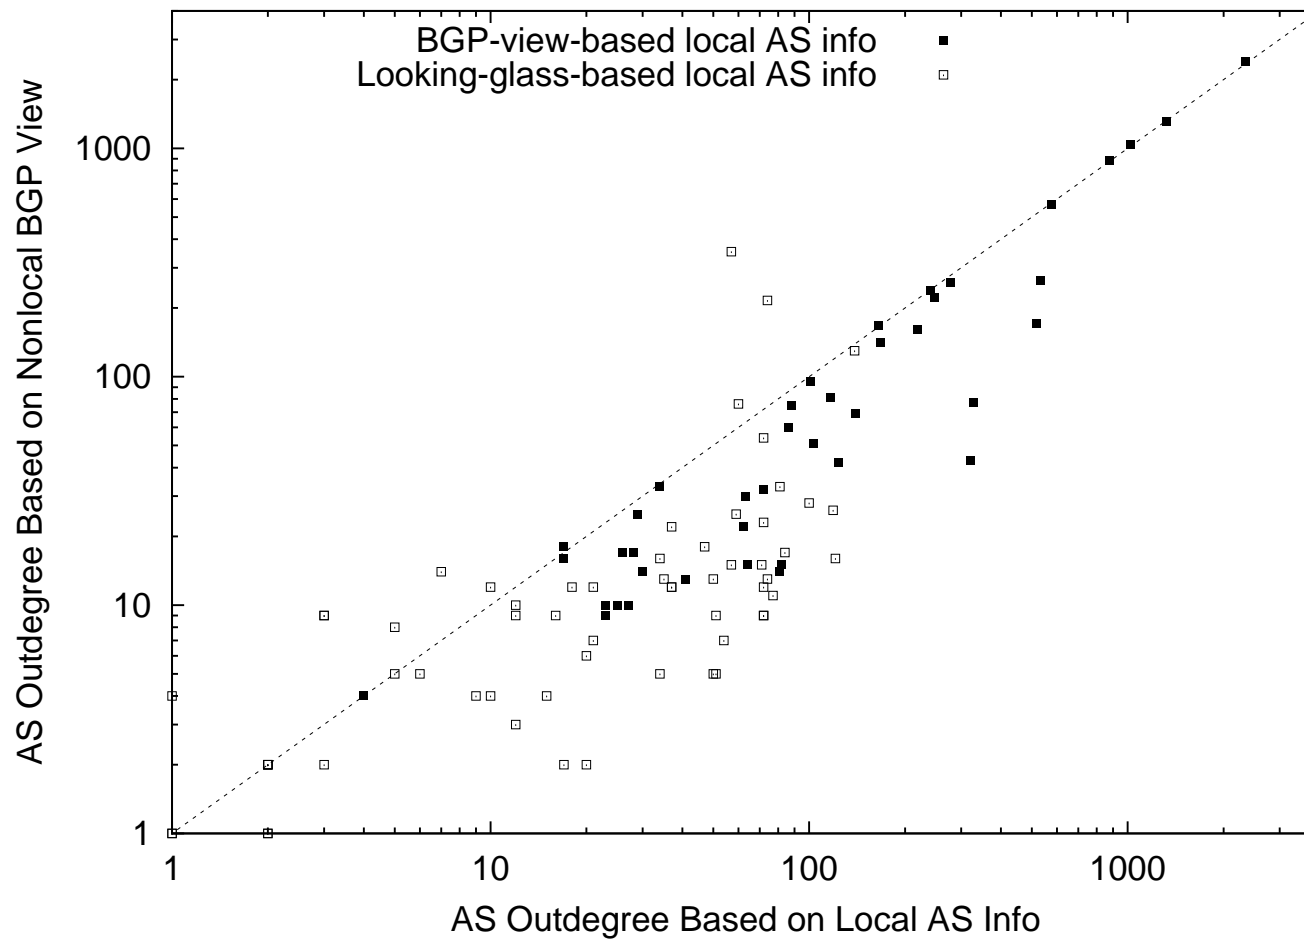
Provides BGP summary information:

- A set of active peers and current aggregated BGP update statistics.

We collected BGP summary information from 300+ routers belonging to 70+ ASs by using an automated script.

Local vs. Non-local BGP view from Looking Glass data

Visualize the incompleteness of non-local BGP views.



Local connectivity info from Routing Registry database

Background:

- Motivation: coordinate and facilitate inter-AS policy routing.
- Individual ISPs voluntarily publish their policy routing.
- Active participation of ISPs in Europe (RIPE database).

Challenges in perusing RR for collecting connectivity info:

- How can we identify an *invalid* database record?
- What is the effect of *biased* and *incomplete* AS sample space?

How can we identify invalid records in routing registry?

Consider invalid any of:

- *Void* record: a record describing a dead (or non-existent) AS,
- *Obsolete* record: a record that contains any non-existent peering relationship,
- *Incomplete* record: a record that does not contain any existing peering relationship.

Among 4,203 ASs registered in RIPE database, 1,029 ASs have *valid* records.

Source	# of nodes (%inc)	# of edges (%inc)
Oregon BGP dump	11,174	23,409
+ RSs	11,268 (0.84%)	26,324 (12.5%)
+ RSs + LG	11,320 (1.3%)	27,899 (19.2%)
+ RSs + LG + RIPE	11,456 (2.5%)	32,759 (40.0%)

What is the effect of biased and incomplete AS sample space?

Problem:

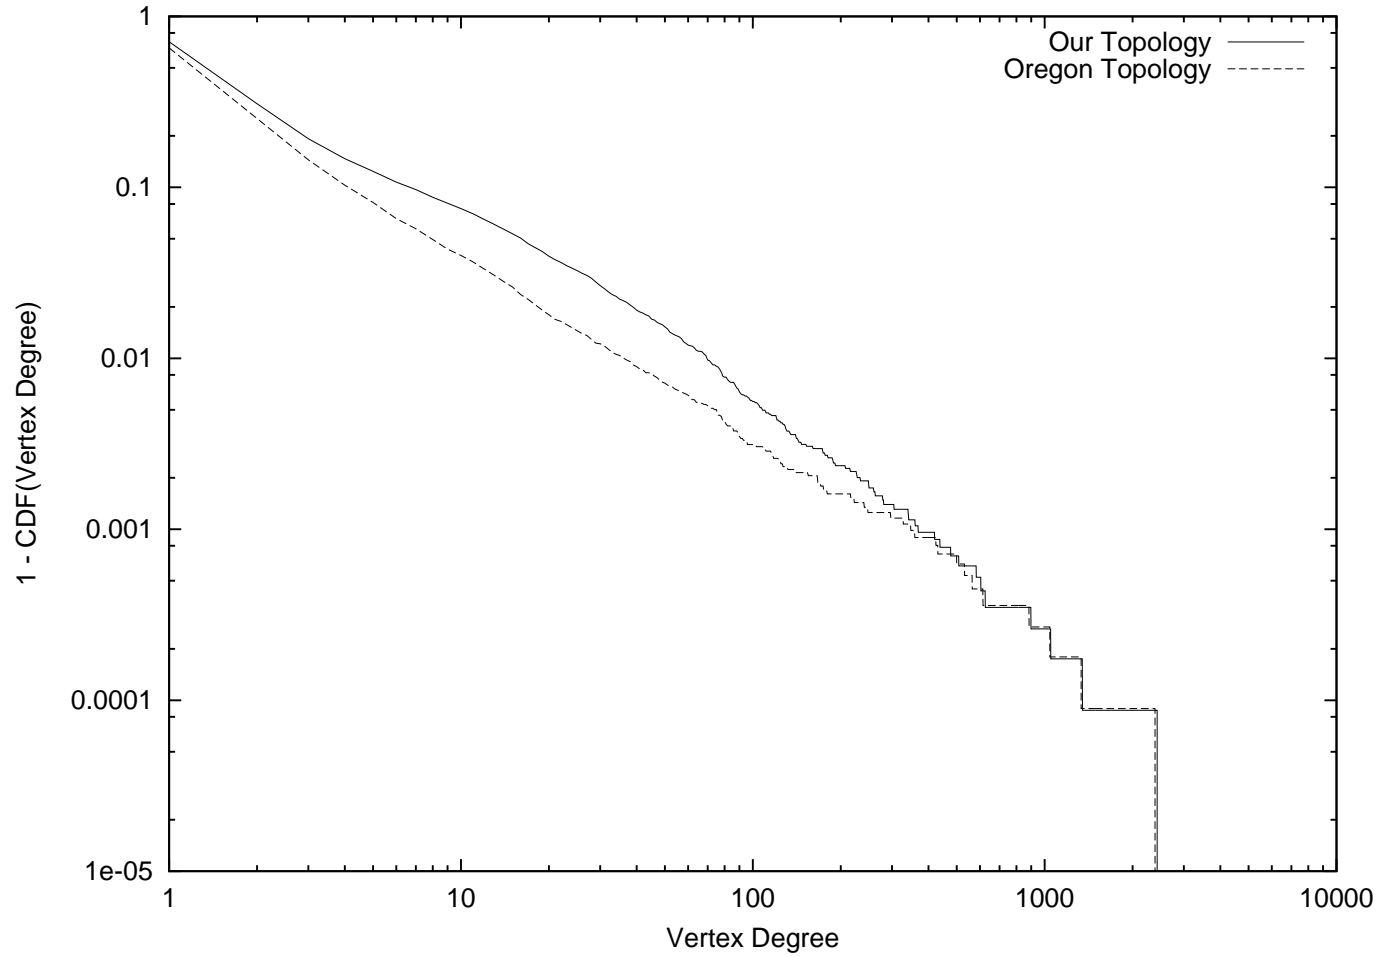
- Most of 9,350 newly found edges originate from European RIPE database.
- Even the AS connectivity of European ASs are not complete.

Then, what can we say about the *complete AS* graph with our dataset?

Observations so far:

- We know the connectivity of Tier-1 ASs.
- We know the connectivity of end-customer ASs.
- We found there exists much more connectivity among non-Tier-1 ASs.

AS Vertex Degree Distribution Revisited



Conclusion

Policy-influenced AS paths may not be able to construct *representative* AS-level topology.

There could be much *denser* inter-AS connectivity clusters bridging Tier-1 ASs and end-customers.

That may affect the well-known strict power-law degree distribution.

However, it is open to question how to *quantify* the degree of connectivity of such clusters.