Towards Capturing Representative AS-Level Internet Topologies

Presented by:
Hyunseok Chang
EECS Department
University of Michigan
hschang@eecs.umich.edu

Joint work with:
Ramesh Govindan(ICSI), Sugih Jamin(UMich),
Scott J. Shenker(ICSI/ACIRI), and Walter Willinger(AT&T-Labs)

http://topology.eecs.umich.edu

How does BGP-inferred AS graph look like?

Exhibits *highly-variable* node degree distribution expressed as:

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

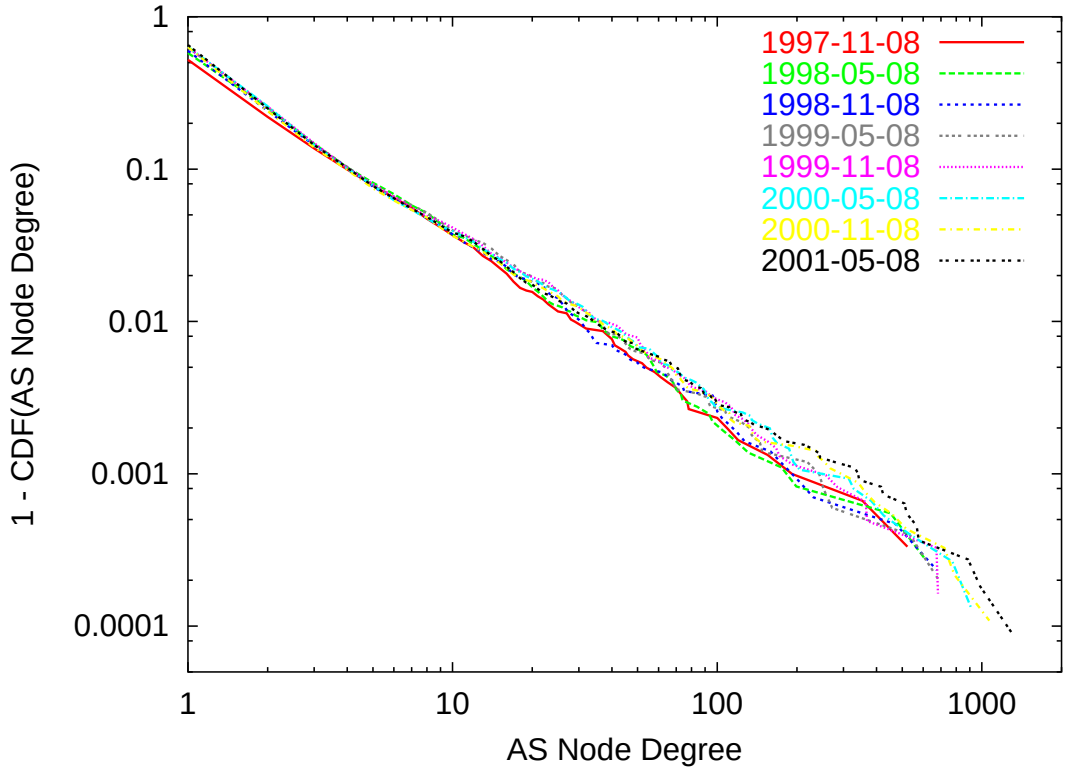

Source: BGP route-views collected by Oregon route server since Nov. 1997.

How complete is the BGP-inferred AS graph?

Can we really answer that?

Then, how about this one?

“How accurate is the number of peering relationships for AS 2500 in the inferred graph?”

We can answer that if its connectivity information is available (i.e., from its own BGP routing table or Looking Glass query).

$\Rightarrow$ A simple but important question: How well are my AS’ peering relationships observed by other ASs?

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

BGP view: an instance of the AS graph constructed from a BGP routing table.

Local & non-local BGP view: just for the definition.

Compare the peering relationships of AS X inferred from its non-local views against those inferred from its local view.

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

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

Observation #1 from local/non-local view analysis

The peering relationships of Tier-1 ASs are more easily observed than those of non Tier-1 ASs.

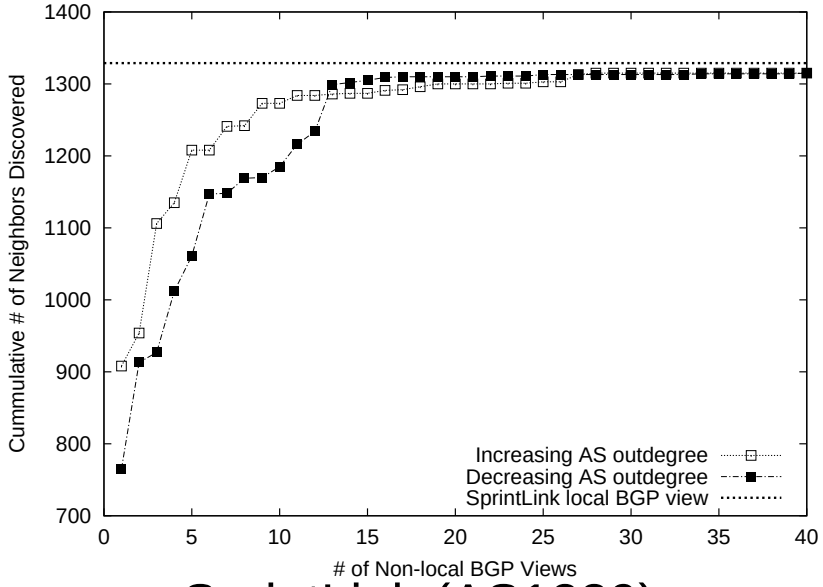

SprintLink (AS1239)

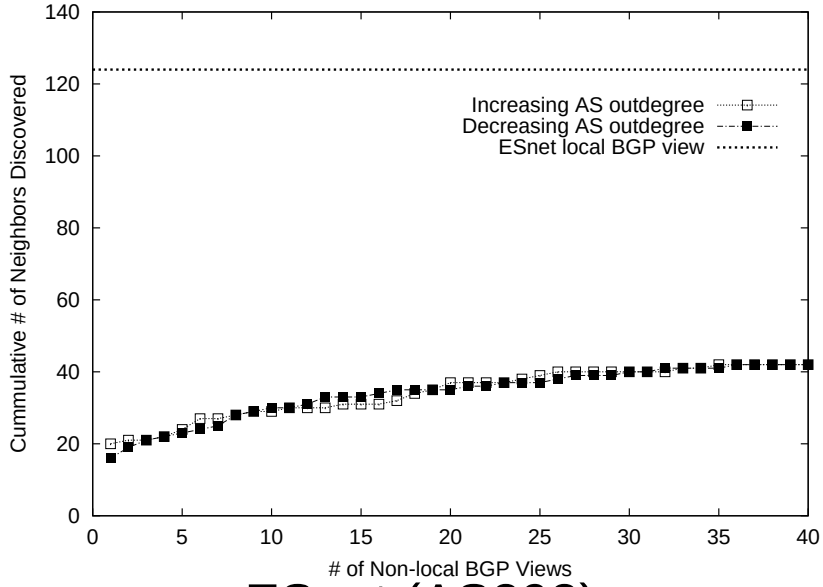

ESnet (AS293)

Observation #2 from local/non-local view analysis

The peering relationships of a given AS are more easily observed by its *customer* ASs than by its *peer* ASs.

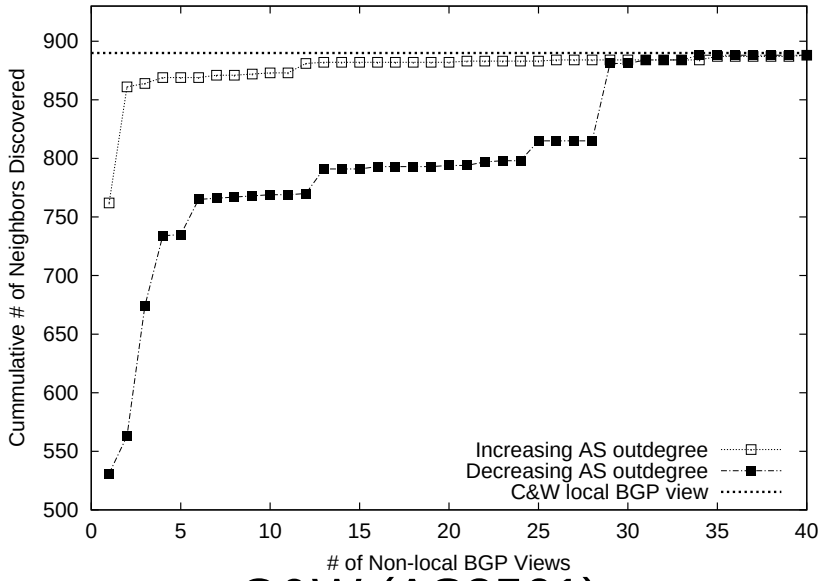

C&W (AS3561)

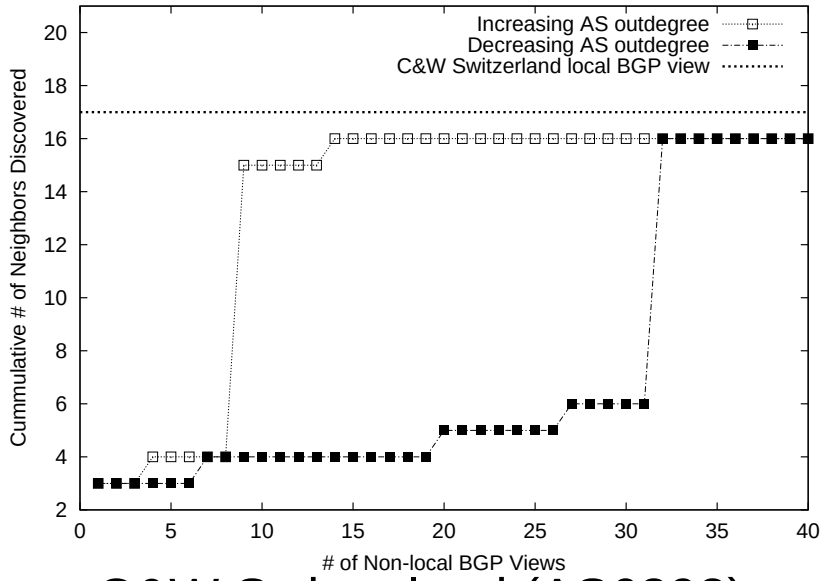

C&W Switzerland (AS6893)

Observation #3 from local/non-local view analysis

The peering relationships with upstream provider ASs are easily observed by non-local views.

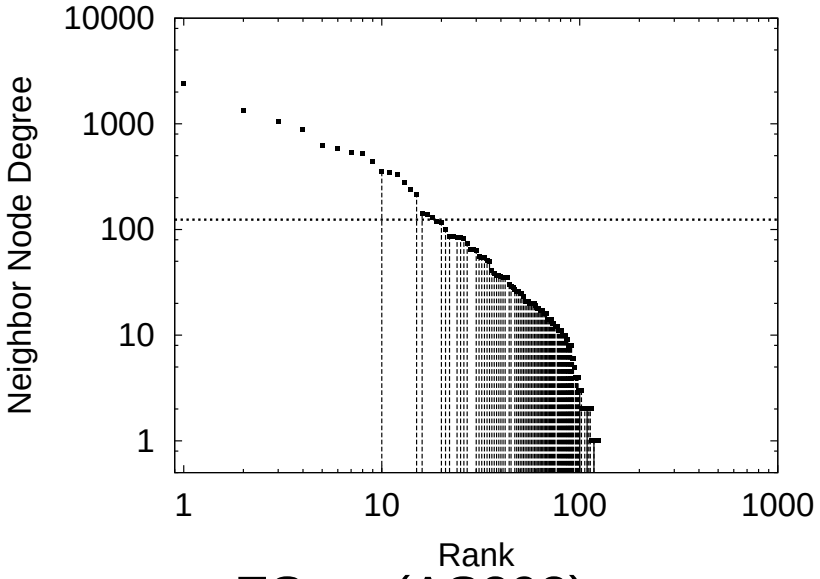

ESnet (AS293)

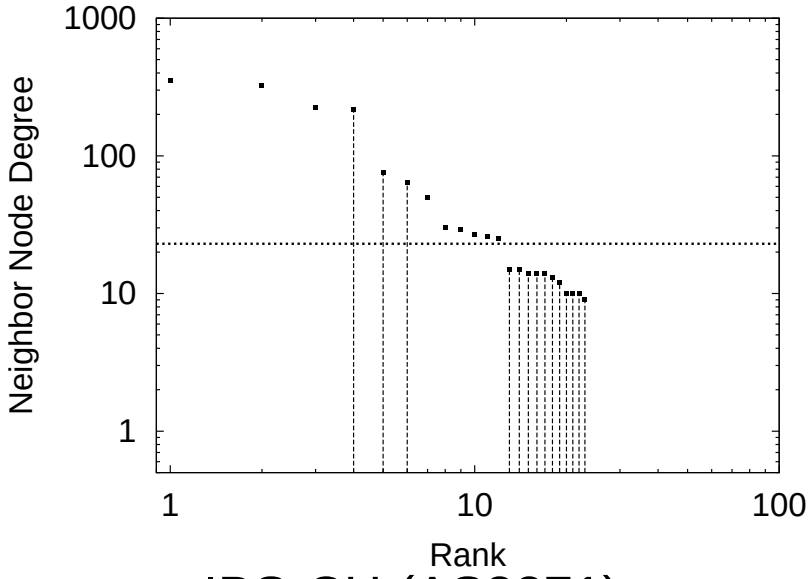

IBS-CH (AS8271)

Local vs. Non-local BGP view from BGP tables

Visualize the incompleteness of non-local BGP views.

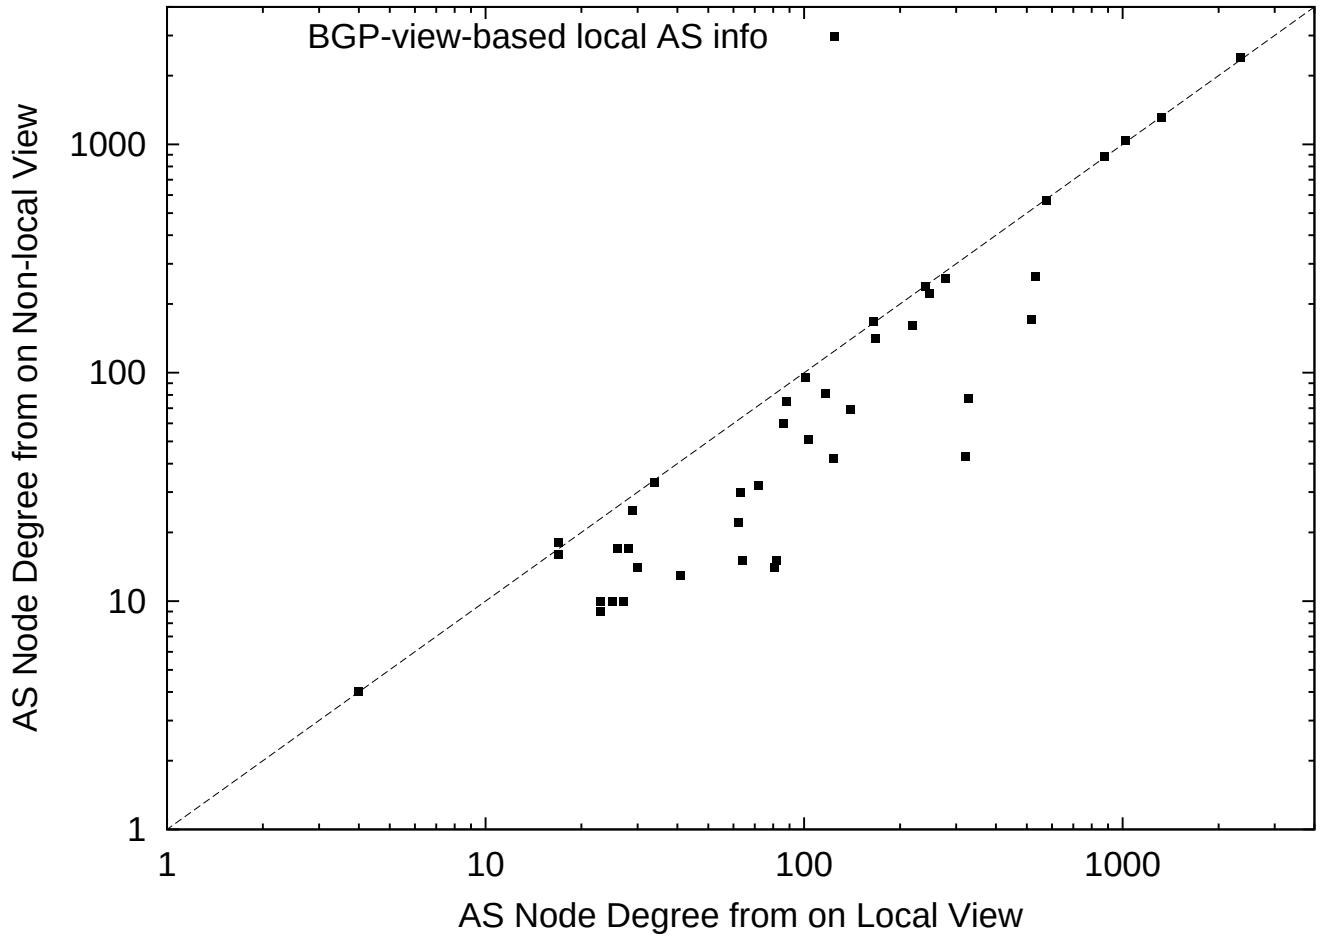

Hyunseok Chang (hschang@eecs.umich.edu)

Augmenting AS peering relationships: Looking Glass data

We collected LG BGP summary information from 300+ routers belonging to 70+ ASs.

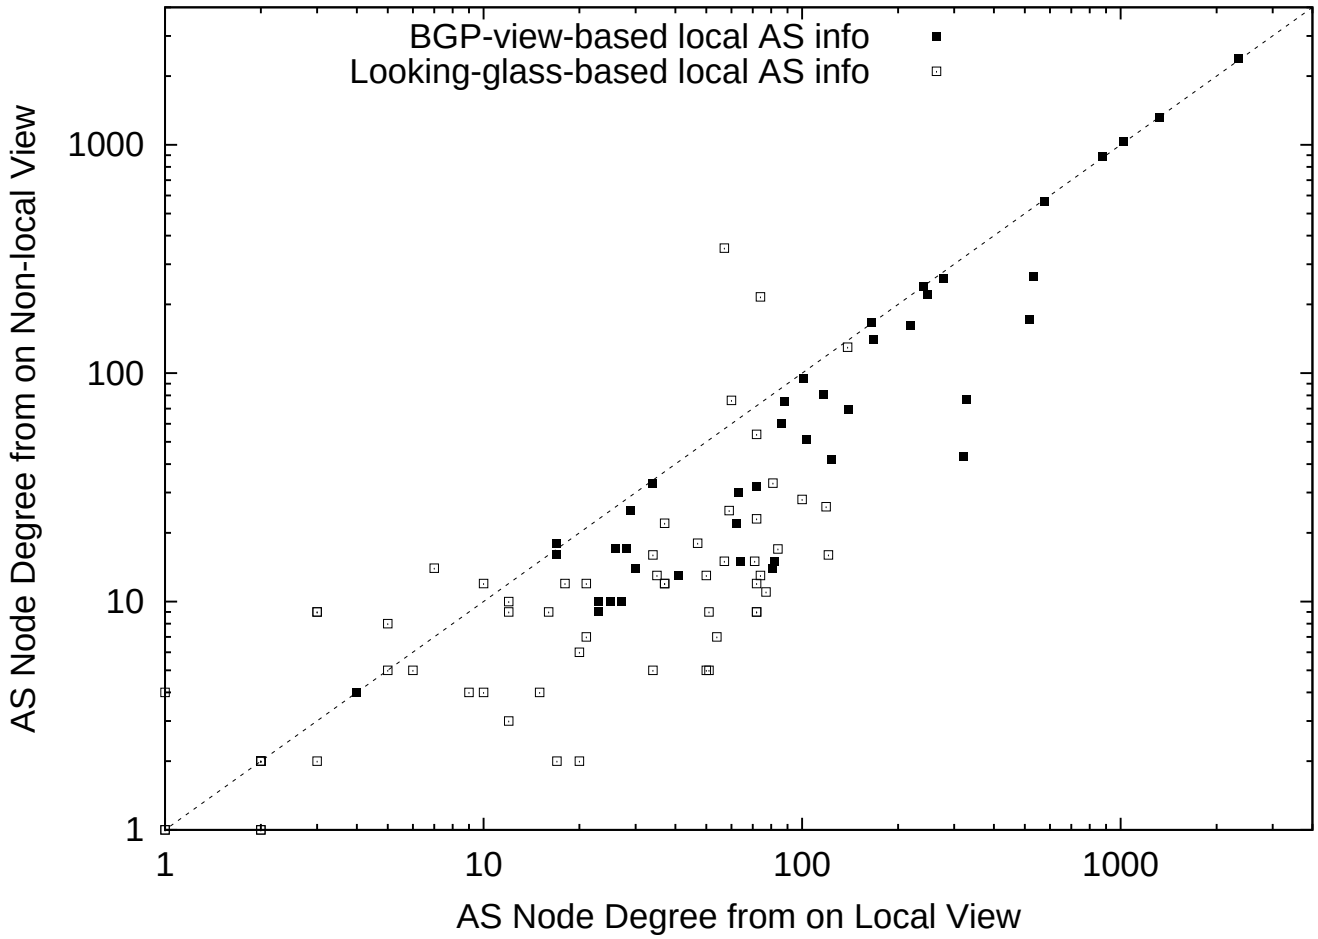

Hyunseok Chang (hschang@eecs.umich.edu)

Augmenting AS peering relationships: Internet Routing Registry (IRR)

Background:

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

Obtaining AS peering relationship info from IRR:

- Can we identify any *invalid* information?
- Furthermore, can we *validate* what's remaining?

Identifying invalid IRR records

Consider invalid any of:

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

How to find them?

- Void/incomplete record: check with Oregon BGP data.
- Obsolete record: check for peering inconsistency between two ASs.

How different is our *more* complete AS graph?

Source	# of nodes (%inc)	# of edges (%inc)
Oregon route-views	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%)

Among 4,203 ASs registered in RIPE database, we identified only 1,029 ASs as having *valid* records.

Validating the IRR data

What are all the links missing from BGP-inferred AS graphs?

- Availability and visibility of "*provider-customer*" vs. "*peer-to-peer*" links.
- Public exchange points (EPs) as connectivity hubs.

$\Rightarrow$ The AS co-location information of EPs may be useful.

By Looking Glass query, we identified the ASs *physically* co-located at 16 different European EPs.

$\Rightarrow$ More than 50% of newly-found edges occur among co-located ASs!

Validating the IRR data (cont'd)

Large-sized maximal cliques may also indicate the presence of connectivity hubs (EPs).

We consider a given maximal clique *EP-bound* if all ASs in the clique are physically co-located at *one* EP.

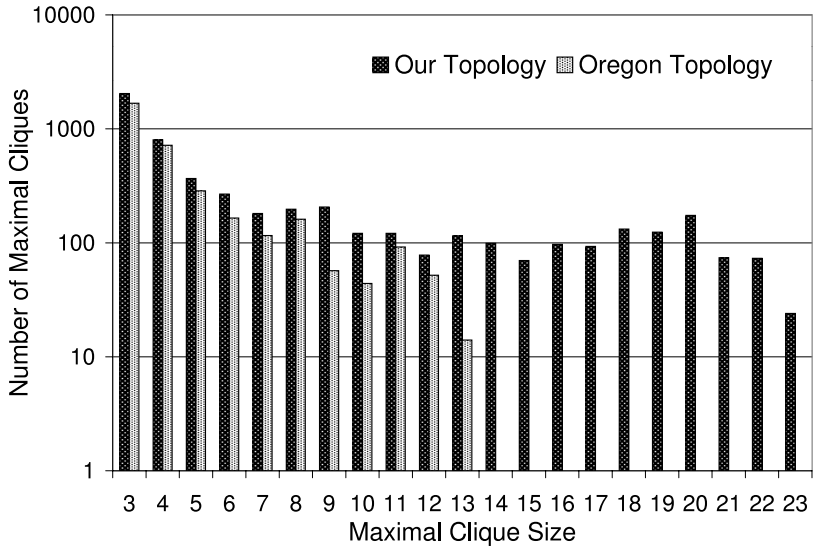

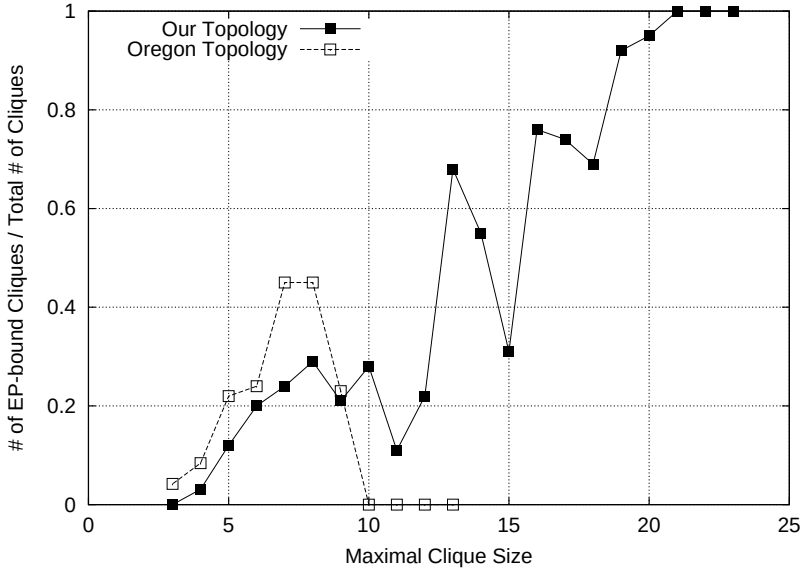

What does our *more* complete AS graph look like?

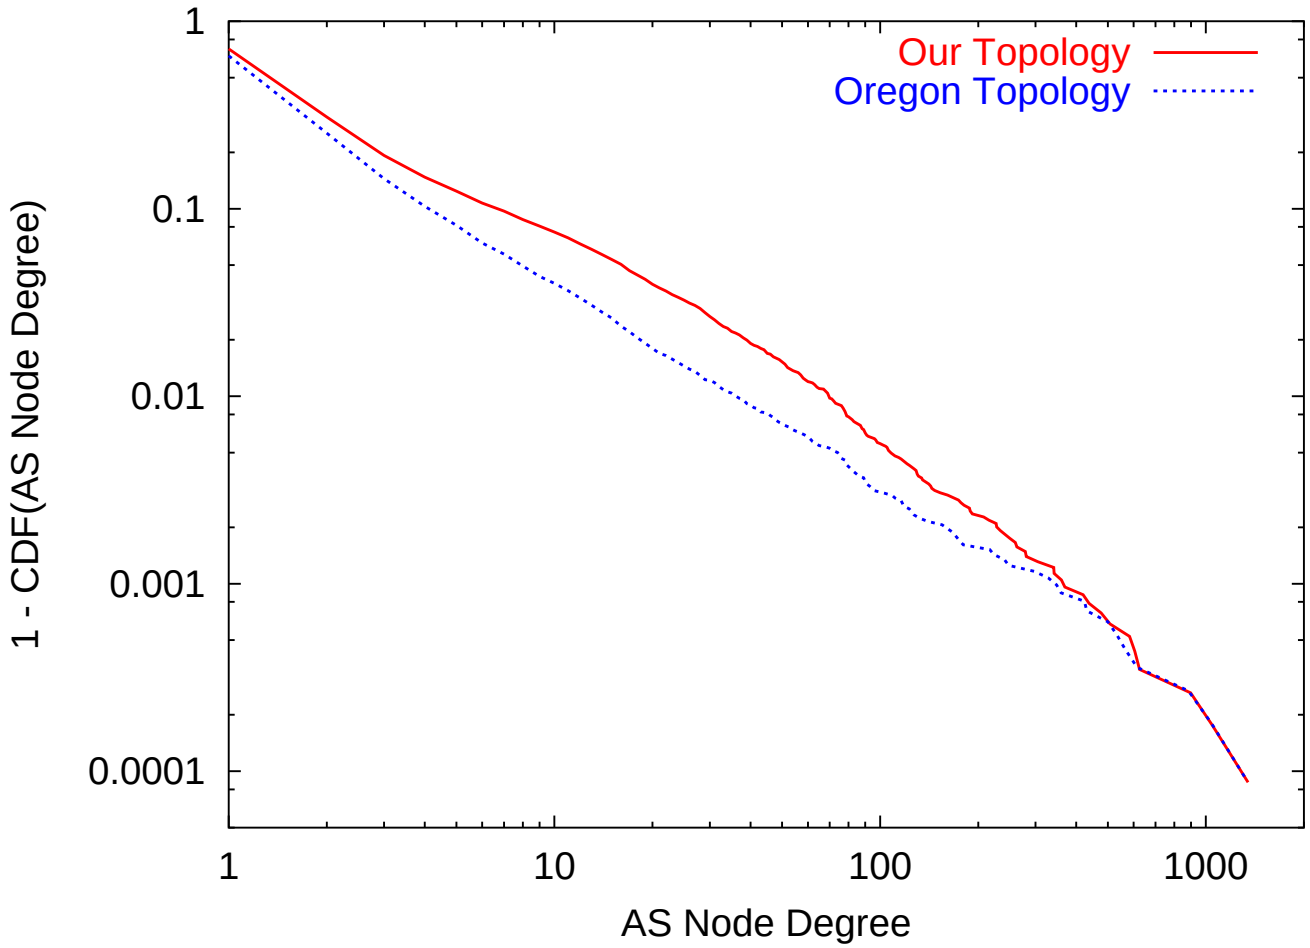

What will a *representative* AS graph look like?

Is the *renewed* node degree distribution valid on the *global* version of the complete AS graph?

Observations so far:

1. We know how many ASs exist globally.
2. We know the peering relationships of Tier-1 ASs.
3. We know the peering relationships of end-customer ASs.
4. There are many regional EPs worldwide, where we have verified a dozen to more than a hundred ASs are co-located.

Recall our AS graph looks like this.

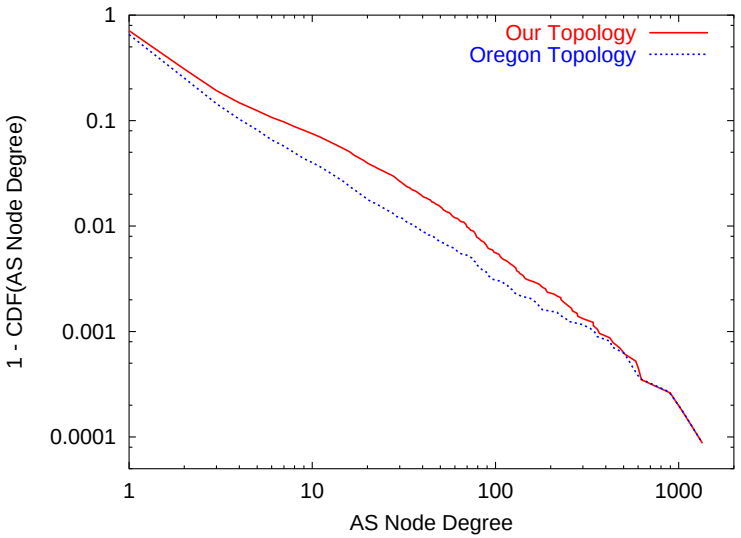

Again, our observations:

1. Total number of existing ASs.
2. Node degree of Tier-1 ASs.
3. Node degree of end-customer ASs.
4. Many connectivity hubs world-wide.

Conclusions

Collecting available BGP routing data may not be sufficient to construct a *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.