Characterizing the Internet Hierarchy from Multiple Vantage Points

Sharad Agarwal

Lakshminarayanan Subramanian

Jennifer Rexford (AT&T Research)

Randy H. Katz

January 2002

http://www.cs.berkeley.edu/~sagarwal/research/BGP-hierarchy

Introduction

Internet

- Composed of domains (AS's)
- Connected via inter-domain routing (BGP)
- Pairwise connections not simple
  - SLA's affect routing & traffic

Goal

- Infer inter-AS relationships
- Construct AS topology of Internet

Challenges

- Each AS hop filters routes
- Inter-AS relationships often private
- Computationally complex problem

Past Work

Example Applications

- Service placement
  - Given topology & client locations
- Application level routing
- Service redirection

Past Work

- Connectivity based Internet maps
  - e.g. CAIDA maps
  - e.g. Mercator maps
- No relationship based maps

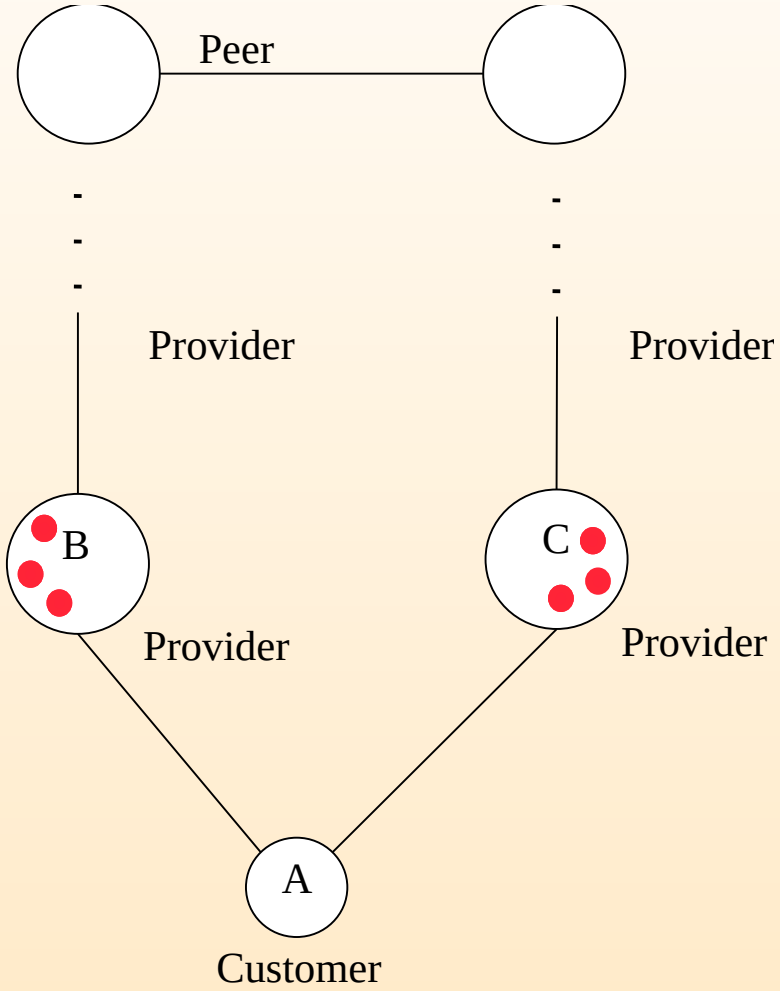

Our Work

Our Contributions

- BGP tables from multiple points
  - Completeness
  - Exploit uniqueness of each point
- Build relationship map
  - Build on Lixin Gao (UMass)'s relationship models

Experiment Data

Telnet Looking Glass Servers

AS #	Name	# Edges
1	Genuity	13419
1740	CERFnet	14287
3549	Globalcrossing	13542
3582	University of Oregon	23136
3967	Exodus Comm.	19005
4197	Global Online Japan	13474
5388	Energis Squared	13534
7018	AT&T	14160
8220	COLT Internet	11282
8709	Exodus, Europe	15519

```
ner-routes>show ip bgp
BGP table version is 6128791, local router ID is 4.2.34.165
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop            Metric LocPrf Weight Path
* i3.0.0.0         4.0.6.142             1000     50      0 701 80 i
* i4.0.0.0         4.24.1.35                0    100      0 i
* i12.3.21.0/24    192.205.32.153        1000     50      0 7018 4264 6468 i
```

Experiment Data

Telnet Looking Glass Servers

AS #	Name	# Edges
1	Genuity	13419
1740	CERFnet	14287
3549	Globalcrossing	13542
3582	University of Oregon	23136
3967	Exodus Comm.	19005
4197	Global Online Japan	13474
5388	Energis Squared	13534
7018	AT&T	14160
8220	COLT Internet	11282
8709	Exodus, Europe	15519

- Publicly available
  - Diversity
  - For how long?
- 18 Apr 2001
  - Compare to Sept and Dec

Relationship Inference

Steps

- Partial view from X
  - Directed graph, rooted at X
  - Prune, leaves get rank 1
  - Prune, leaves get rank 2 ...
- Every AS has vector of ranks
  - Length = 10 (num of partial views)
- Apply heuristics

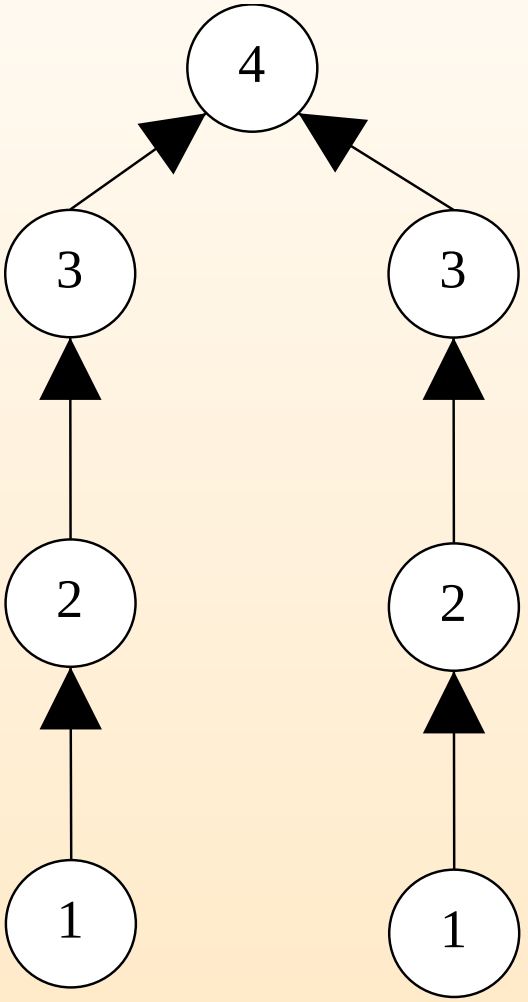

AS5444 = (4 5 4 4 3 4 4 4 5 4)

AS5665 = (3 2 3 3 3 3 3 3 4 3)

Relationship Inference

Steps contd...

- Export Rules
  - ▷ To Provider: Own, customer routes, not provider, peer routes
  - ▷ To Peer: Own, customer routes, not provider, peer routes
  - ▷ To Customer: Own, customer, peer, provider routes

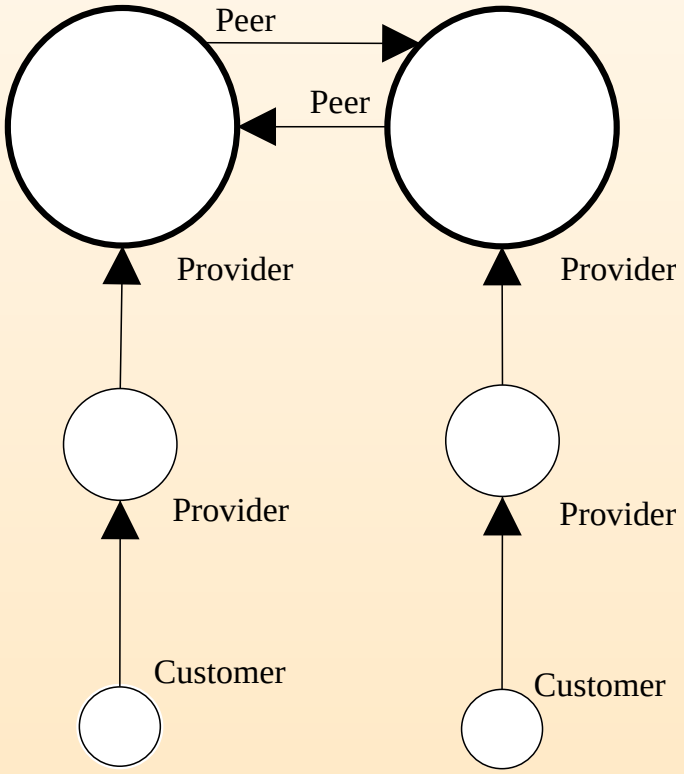

- ▷ Every path : uphill, 0-1 peering link, downhill

Relationship Inference

Steps contd...

- Heuristics
  - Complete dominance
    - provider-customer if one's entries greater than other
  - Equivalence
    - peers if ≥ 5 entries match
  - Clustering
    - peers if euclidean distance $< \sqrt{10}$
  - Probabilistic versions
    - for links not inferred above

$$\text{AS5444} = (4\ 5\ 4\ 4\ 3\ 4\ 4\ 4\ 5\ 4)$$

$$\text{AS5665} = (3\ 2\ 3\ 3\ 3\ 3\ 3\ 3\ 4\ 3)$$

Relationship Inference

Determine relationship

- Peers
- Customer-Provider
- Unknown (siblings or acquisitions)

Inferred Relationships for 24,059 AS Pairs

Relationship	# AS pairs	Percentage
Provider-customer	22,712	94.40%
Peer-peer	1,241	5.16%
Unknown	106	0.44%

Distribution of the 23,953 Inferences

Rule	Number	Percentage
Complete dominance	22,241	97.93%
Probabilistic dominance	471	2.07%
Equivalence	836	67.37%
Probabilistic equivalence	278	22.40%
Clustering	127	10.23%

Hierarchy Formation

- New directed graph
  - Based on inferred relationships
- Assign ranks
  - Customers, regional ISPs
    - leaves, 0 out-degree
  - Dense core
    - weak clique
    - N nodes, every node degree > N/2
  - Transit core
    - weak cut
  - Intra-layer links non-hierarchical

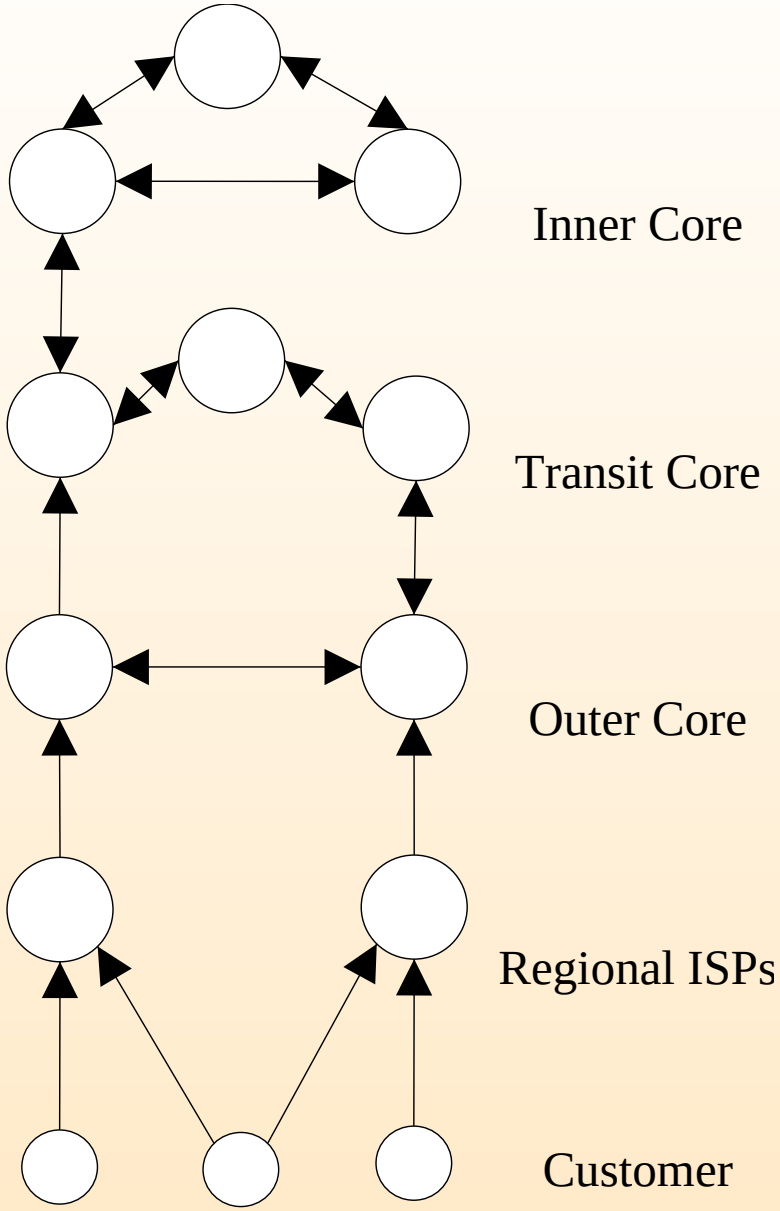

Hierarchy Formation

Distribution of AS's in Hierarchy

Level	# of AS's
Dense core (0)	20
Transit core (1)	162
Outer core (2)	675
Small regional ISPs (3)	950
Customers (4)	8852

Cumulative Distribution of AS Degree

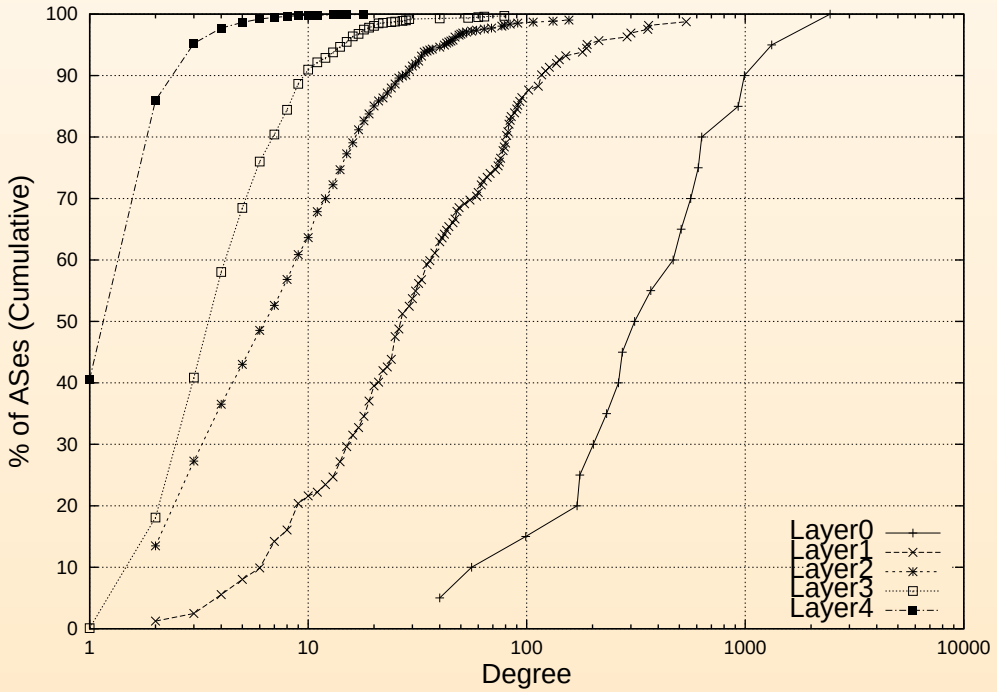

Dense Core

Dense Core

- BBN, Sprint, PSI-Net, Ebone, Telia, Qwest, Digex, Concentric, Verio, Level3, GBLX, Cable-n-Wireless, Exodus, Netreal, France Telecom, Alternet, AT&T, Carrier1
- December
  - - Ebone, Netreal, France Telecom, Alternet
  - + KDDI, Reach, Teleglobe, AOL, Tiscali
  - Alternet : 43 peer relationships down to 26
  - AOL : 0 peer relationships to 22 (RoadRunner acquisition)

Compare Partial Views

Dense Core View from BBN

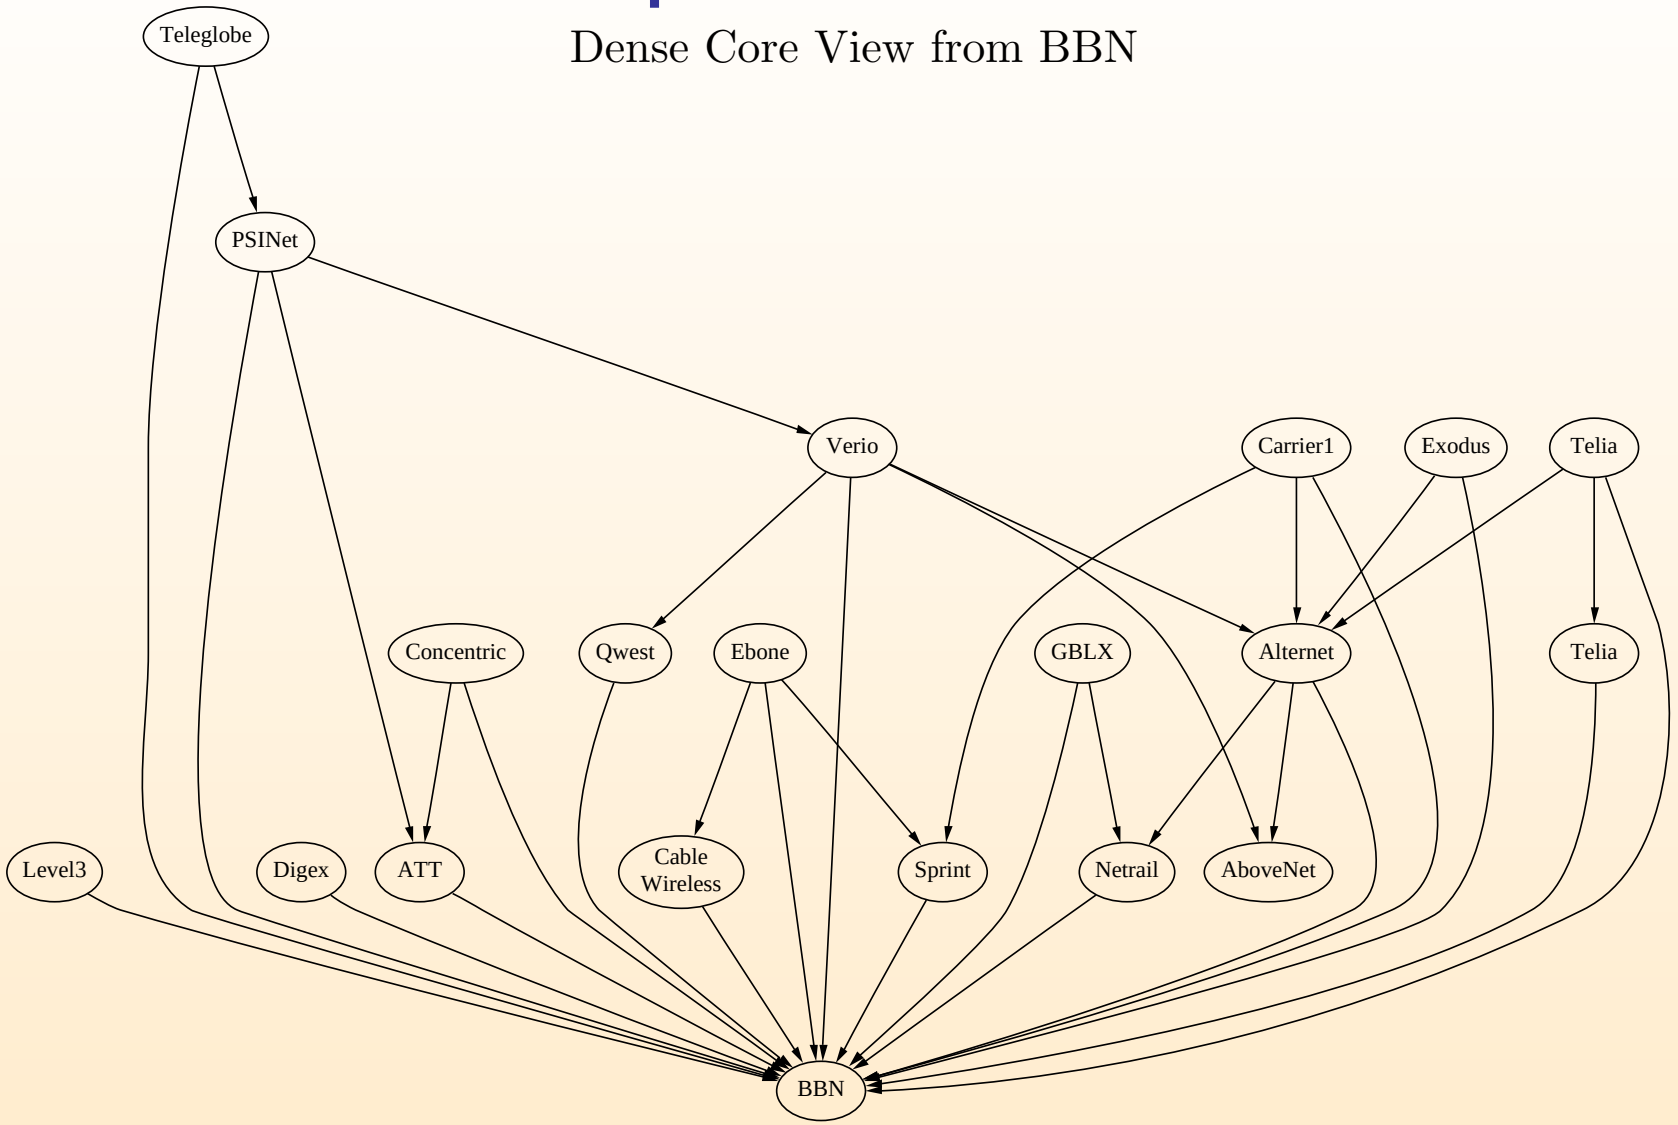

Compare Partial Views

Dense Core View from all 10 AS's

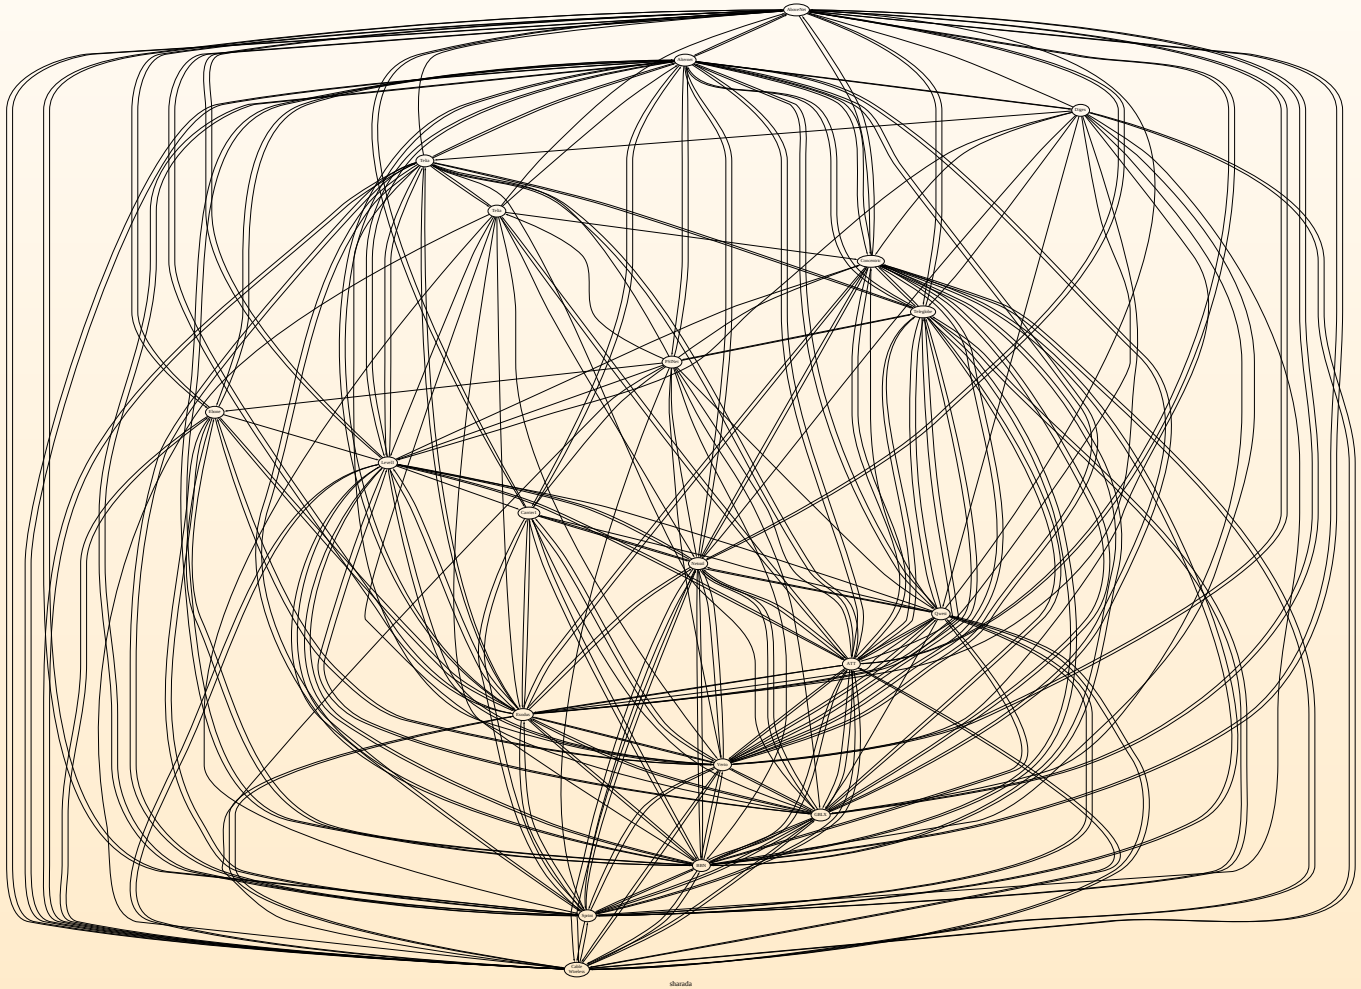

Multiple Views : Completeness

Compare Partial Views

Edge Distribution across 10 AS's

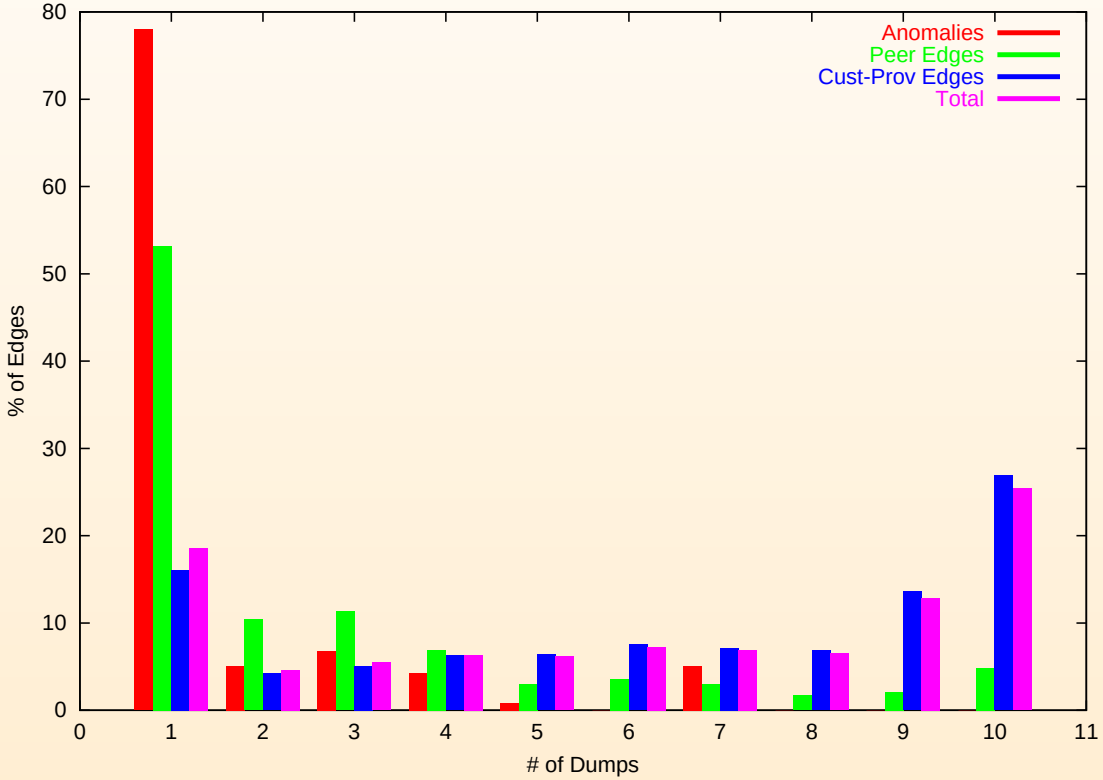

Multiple Views : Uniqueness

Summary

Infer AS relationships

- Hard problem
- Apply heuristics
- Use multiple views for completeness
- Exploit uniqueness of each view

Hierarchy

- Prior work only on connectivity
- Relationship directed graph to form hierarchy

`http://www.cs.berkeley.edu/~sagarwal/research/BGP-hierarchy`

- Raw data
- Processed data
- Upcoming Infocom 02 publication