Something We Always Wanted to Know about ASs: Relationships and Taxonomy

Dmitri Krioukov

dima@caida.org

X. Dimitropoulos, M. Fomenkov, B. Huffaker,
Y. Hyun, kc claffy, and G. Riley.

LIP6, June 22nd, 2006

High-level goal

Annotated topologies:

Go beyond the view of the Internet AS-level topology as an undirected unweighted graph to include information on types of links (relationships) and nodes (taxonomy).

Motivation

- Practical (providers, vendors, government)
  - Money flow
  - Traffic flow
  - Network robustness
- Theoretical (research community)
  - Routing ⇐
  - Topology ⇐
  - Modeling ⇐
  - Validation (real data)

Outline

- AS relationships
- AS taxonomy
- AS rank

Outline

- AS relationships
  - Problem formulation
  - Overview of the existing heuristics and their limitations
  - How we address these limitations
  - Validation
- AS taxonomy
- AS rank

Problem formulation

- Given: data (BGP, IRR, skitter, etc.)
- Find: business relationship between AS neighbors
- Using: a set of abstractions including these:
  - Types of relationships
    - customer-to-provider (c2p or p2c)
    - sibling-to-sibling (s2s)
    - peer-to-peer (p2p)
  - Valid paths (follows from the standard routing policies)
    - uphill: zero or more of c2p links
    - pass: zero or one p2p link
    - downhill: zero or more p2c links

Existing heuristics: Gao and SARK

- L. Gao. On inferring Autonomous System relationships in the Internet. *ToN 2001*. (Gao)
  - BGP policies $\Rightarrow$ (in)valid paths
  - AS degree-based heuristic
  - Too many invalid paths
- L. Subramanian, et al. Characterizing the Internet hierarchy from multiple vantage points. *INFOCOM 2002*. (SARK)
  - Combinatorial optimization to minimize the number of invalid paths (ToR problem)
  - Heuristic to solve it

Existing heuristics: DPP and EHS

- G. Di Battista, et al. Computing the types of the relationships between Autonomous Systems. *INFOCOM*, *2003*, (DPP); **and** T. Erlebach, et al. Classifying customer-provider relationships in the Internet. *IASTED CCN*, *2002*, (EHS).
  - No peering can be inferred in ToR
  - ToR is NP- and APX-complete
  - More rigorous approach to find an approximate solution
  - Smaller number of invalid paths (than in SARK)
  - Induced AS hierarchies are incorrect

Existing heuristics: more recent relevant papers

- J. Xia and L. Gao. On the evaluation of AS relationship inferences. *GLOBECOM 2004*.
  - Validation using IRRs
- Z. M. Mao, et al. On AS-level path inference. *SIGMETRICS 2005*.
  - Path inference based on the shorter AS-path preference assumption

Outline

- AS relationships
  - Problem formulation
  - Overview of the existing heuristics and their limitations
  - How we address these limitations in our algorithms for:
    - customer-to-provider (c2p) links
    - sibling-to-sibling (s2s) links
    - peer-to-peer (p2p) links
  - Validation
- AS taxonomy
- AS rank

Idea at the high level

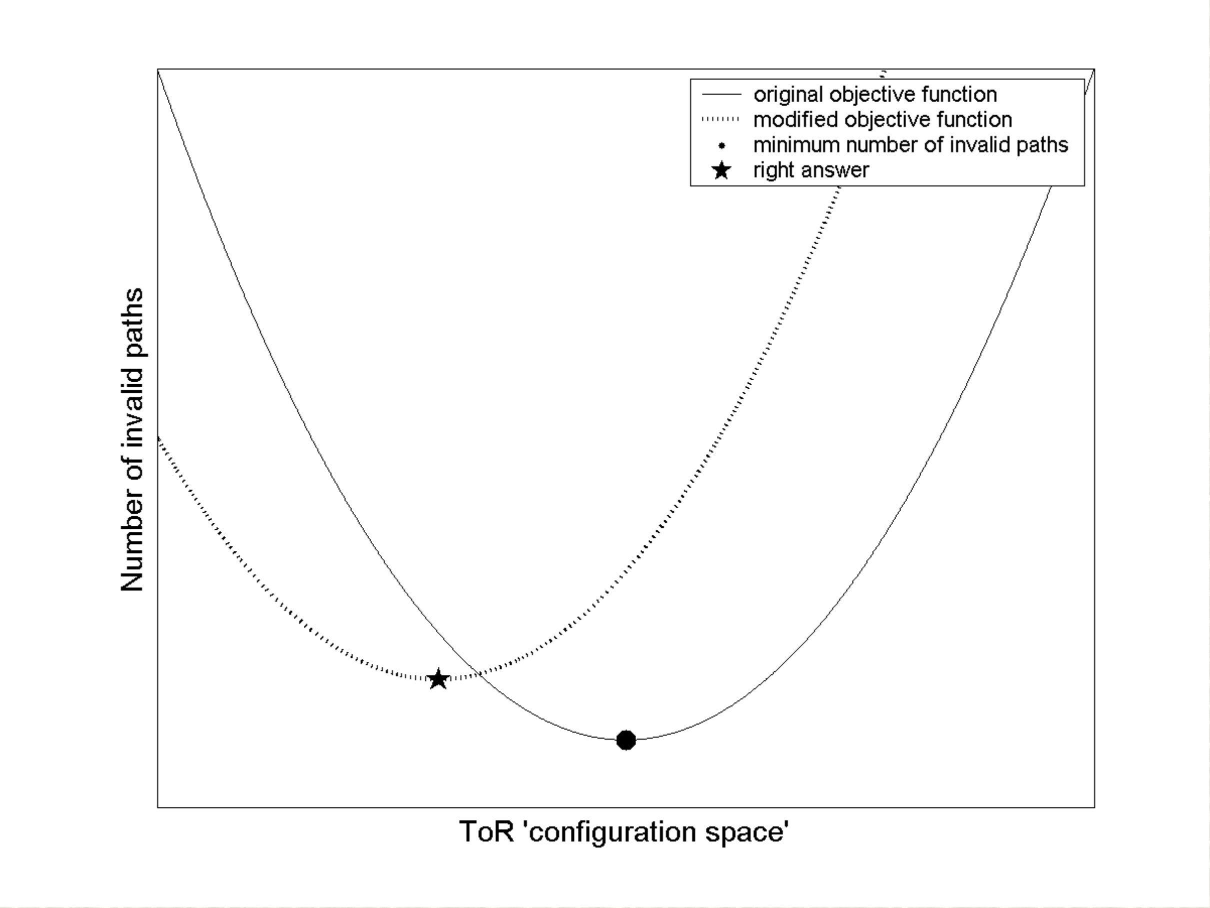

Objective function adjustment

ToR

- Given a set of BGP paths P,
- Extract the undirected AS-level graph G.
  - Every edge in G is a link between pair of ASs.
- Assuming edge direction is from customer to provider,
- Direct all edges in G (2^m combinations),
- Inducing direction of edges in P,
- Such that the number of invalid paths in P is minimized.
  - Invalid path is a path containing a provider-to-customer link followed by customer-to-provider link

ToR and MAX2SAT

- Split all paths in *P* into pairs of adjacent links (involving triplets of nodes)
- Perform mapping…

Mapping to MAX2SAT

Edges in P	2SAT clause	Edges in G_{2SAT}
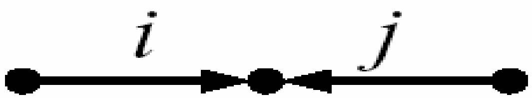	$x_i \vee x_j$	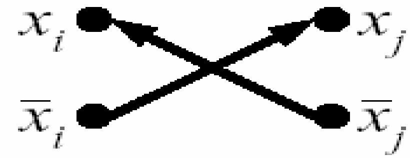
	$x_i \vee \bar{x}_j$	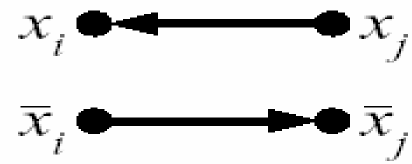
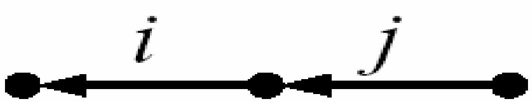	$\bar{x}_i \vee x_j$	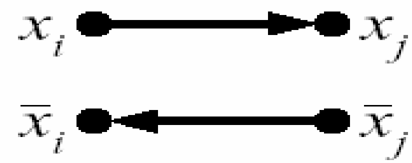
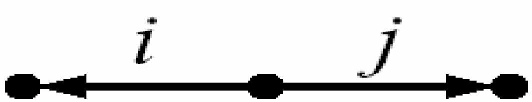	$\bar{x}_i \vee \bar{x}_j$	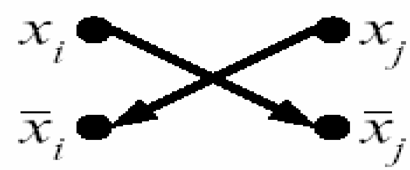

Two 2SAT observations

- All clauses can be satisfied (all paths can be made valid) if there is no variable x_i belonging with its negation to the same SCC in G_{2SAT} (conflict variable/edge)
  - SCC (strongly connected component) is a set of mutually reachable nodes in a directed graph
- Proper direction of non-conflict edges can be done via topological sorting in G_{2SAT} (if the variable negation is before the variable itself, then the variable is *true*, and vice versa)
  - Topological sorting is a natural ordering of nodes in directed acyclic graphs

MAX2SAT: DPP vs. EHS

- If P is large, not all paths (clauses) can be made valid (satisfied): 2SAT $\Rightarrow$ MAX2SAT
- DPP: find the maximum subset of paths that can all be made valid
- EHS: use known algorithms to approximate MAX2SAT
  - SDP (semidefinite programming) relaxation (with certain twists) delivers approximation ratio of 0.940
  - Inapproximability ratio is 0.954

SDP relaxation to MAX2SAT

$$\begin{aligned} \max \quad & \frac{1}{4}\sum_{k,l=1}^{2m_1} w_{kl}(3 + v_0 \cdot v_k + v_0 \cdot v_l - v_k \cdot v_l) \\ \text{s.t.} \quad & v_0 \cdot v_0 = v_k \cdot v_k = 1,\ v_i \cdot v_{m_1+i} = -1, \\ & k = 1 \ldots 2m_1,\ i = 1 \ldots m_1. \end{aligned}$$

Physical interpretation

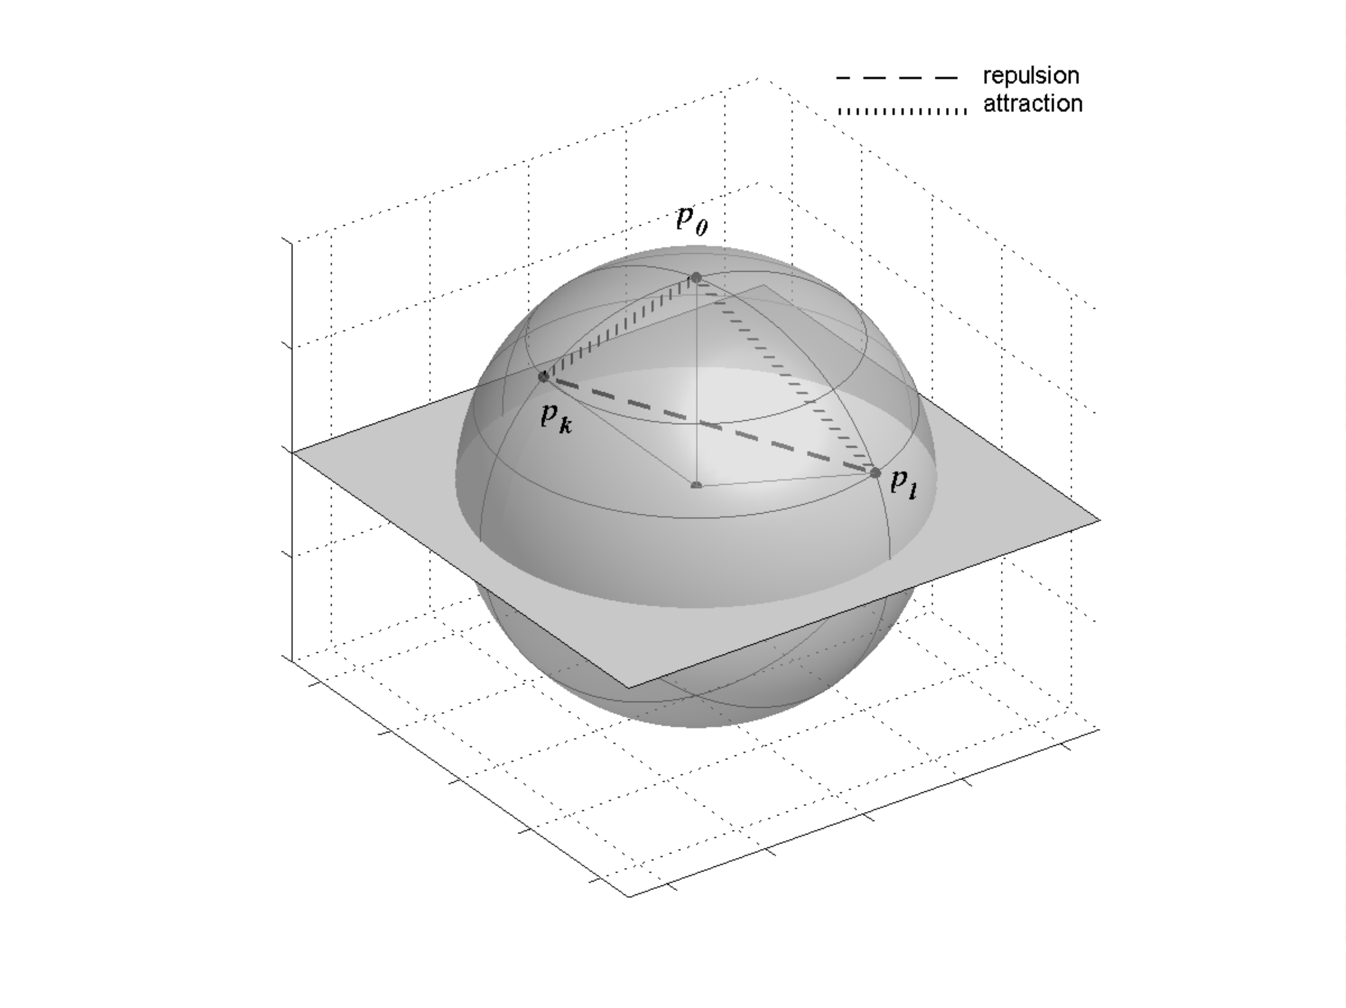

Gains and losses

- What’s good
  - Extremely small number of invalid paths
- What’s bad
  - Skewed/incorrect AS hierarchies: several small ASs are inferred as providers of large ISPs
  - But why!?

Causes of the problem and their resolutions

- **Case 1:** some edges can be directed any way without causing invalid paths
- **Fix:** introduce additional incentive to direct edge along the node degree gradient

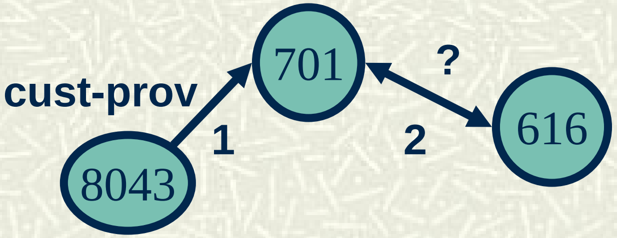

- **Case 2:** trying to infer sibling links leads to proliferation of error
- **Fix:** try to discover sibling links using the WHOIS database

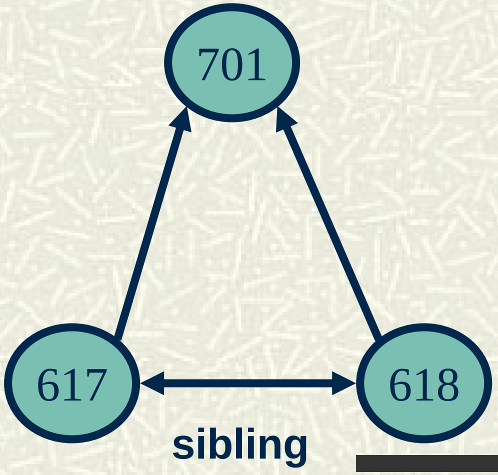

Case 1: Infer c2p links using multiobjective optimization

- Maximize number of invalid paths:
  - 2-link clauses $w_{kl}(x_k \vee x_l)$
- Direct along the node degree gradient:
  - 1-link clauses $w_{kk}(x_k \vee x_k)$

Final form of the generalized problem formulation

$$\begin{aligned} \max \quad & \frac{1}{4}\sum_{k,l=1}^{2m_1} w_{kl}(3 + v_0 \cdot v_k + v_0 \cdot v_l - v_k \cdot v_l) \\ \text{s.t.} \quad & v_0 \cdot v_0 = v_k \cdot v_k = 1,\ v_i \cdot v_{m_1+i} = -1, \\ & k = 1 \ldots 2m_1,\ i = 1 \ldots m_1. \end{aligned}$$

$$w_{kl}(\alpha) = \begin{cases} c_2\alpha & \text{if} \quad \{kl\} \in P, \\ c_1(1-\alpha) f(d_k^-, d_k^+) & \text{if} \quad k = l \leqslant m_1, \\ 0 & \text{otherwise.} \end{cases}$$

$$f(d_i^-, d_i^+) = \frac{d_i^+ - d_i^-}{d_i^+ + d_i^-} \log(d_i^+ + d_i^-).$$

Case 2: Infer s2s links using IRR data

- Hard to infer from BGP data
- Use IRRs instead
- Dictionary of organization name synonyms
- IRR data can be stale, but organization names are relatively stable

Outline

- AS relationships
  - Problem formulation
  - Overview of the existing heuristics and their limitations
  - How we address these limitations in our algorithms for:
    - customer-to-provider (c2p) links
    - sibling-to-sibling (s2s) links
    - peer-to-peer (p2p) links
  - Validation
- AS taxonomy
- AS rank

Inferring p2p links

- Find *F*: the set of links adjacent to top degree nodes in all paths
- Clean *F* with $g(d_i^-, d_i^+) < w_e$ validations: $w_e = g(3,545)$

$$g(d_i^-, d_i^+) = 1 - c_3 f(d_i^-, d_i^+)$$

- Clean "more than one p2p links per path" out of *F* with maximum weight independent set (MWIS) solver (all links are weighted by *g*)

Overview of inferring all links

- Given: graph *G(V, E)* constructed from path set *P*
- Find:
  - s2s link set *S* in *E*
  - c2p/p2c directions of links in *E* – *S*
  - p2p candidate link set *F* in *E*
- Answer:
  - s2s links are *S*
  - p2p links are *F* – *S*
  - c2p/p2c links are *E* – *S* – *F*

Results

- Input: RouteViews, 8-hour interval snapshots between 03/01/05 and 03/05/05
- Output:

	Total $\|E\|$	c2p links $\|E \setminus F \setminus S\|$	p2p links $\|F \setminus S\|$	s2s links $\|S\|$
number of links	$38,282$	$34,552$	$3,553$	177
percentage	100%	90.26%	9.28%	0.46%

AS hierarchy

				$\alpha = 0.00$		$\alpha = 0.01$		$\alpha = 0.05$		$\alpha = 0.10$		$\alpha = 0.50$		$\alpha = 1.00$	
Percentage of invalid paths															
				12.75%		1.79%		0.69%		0.46%		0.36%		0.33%	
Top of reachability based hierarchy															
	AS #	name	degree	dep.	wid.	dep.	wid.	dep.	wid.	dep.	wid.	dep.	wid.	dep.	wid.
$\alpha = 0$	701	UUNET	2334	0	1	1	1	0	105	0	120	2	201	11	319
	7018	AT&T	1911	1	1	2	1	0	105	0	120	2	201	11	319
	1239	Sprint	1703	2	1	0	1	0	105	0	120	2	201	11	319
	3356	Level 3	1228	3	1	3	1	0	105	0	120	2	201	11	319
	209	Qwest	1105	4	1	4	1	0	105	0	120	2	201	11	319
$\alpha = 1$	14551	UUNET	35	128	1	137	2	138	1	151	1	260	2	0	1
	13987	IBASIS Inc.	3	1792	955	1802	963	1830	976	1847	971	1885	966	1	2
	8631	Routing Arbiter	48	108	1	123	1	122	2	0	120	0	1	1	2
	23649	Hong Kong Teleport	4	1792	955	1802	963	899	121	916	121	967	119	3	8
	4474	Village Communications	2	2747	16136	2765	16118	2806	16077	2818	16065	2	201	3	8

Phase transition in mean field approximation

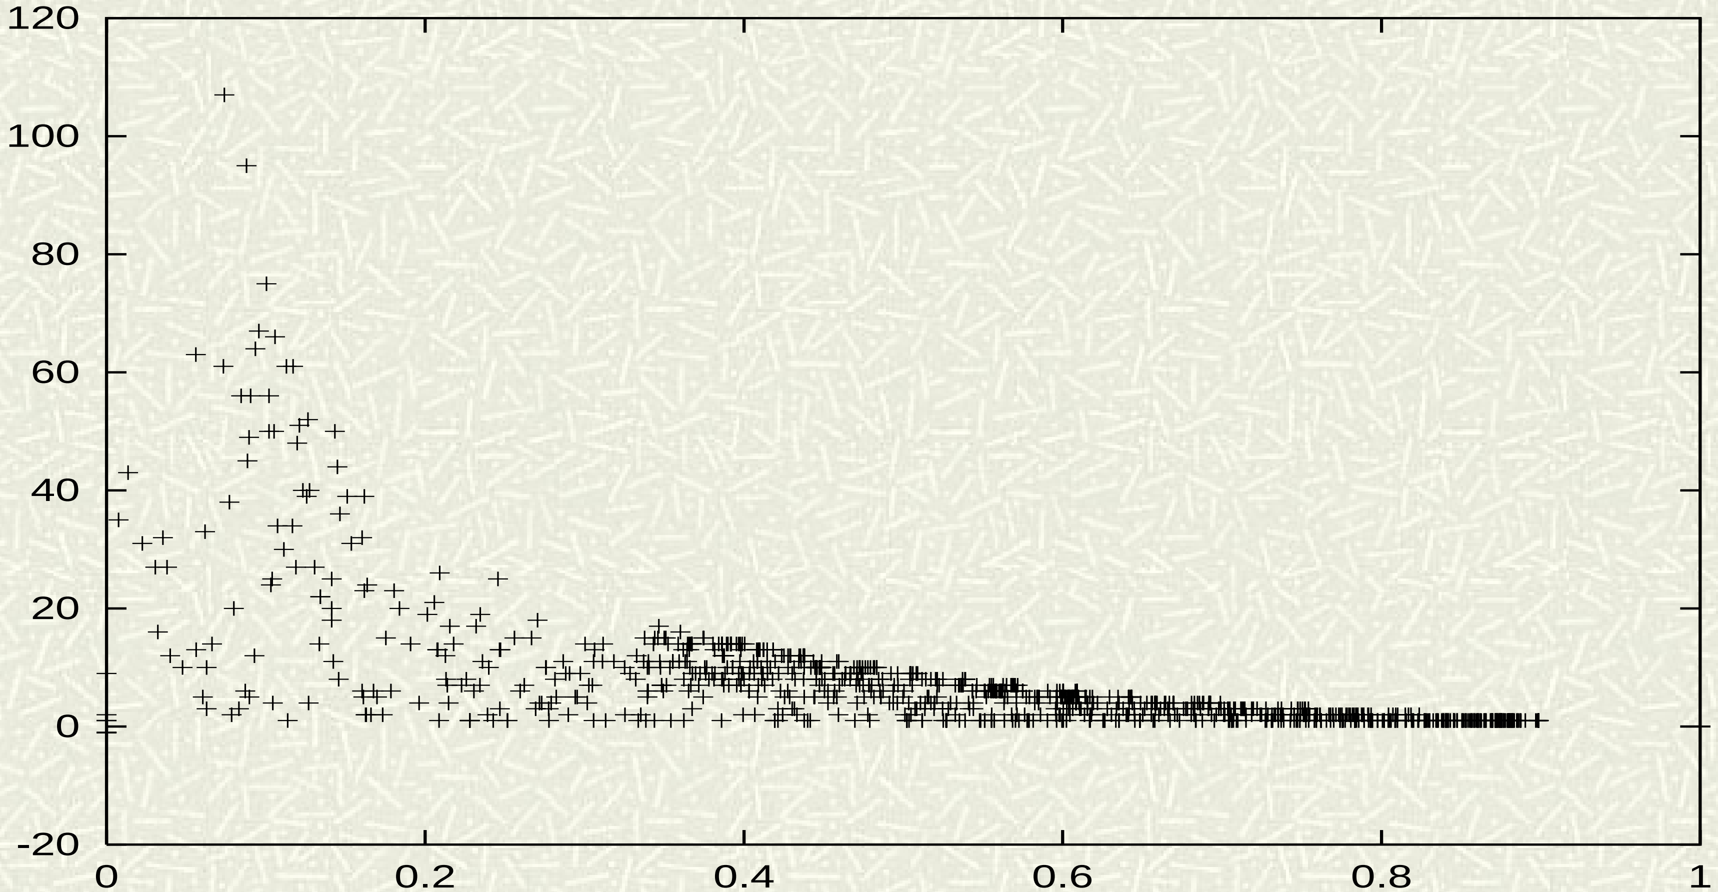

Outline

- AS relationships
  - Problem formulation
  - Overview of the existing heuristics and their limitations
  - How we address these limitations
  - Validation
- AS taxonomy
- AS rank

Validation

- Previous validation efforts
  - Gao: AT&T
  - SARK: Gao
  - Subsequent: SARK/Gao
- Our validation
  - 38 ASs (5 Tier-1 ISPs, 13 smaller ISPs, 19 universities, and 1 content provider)
  - 3,724 links (9,7% of the total)
  - 94.2% overall accuracy

	links	inferred c2p links	inferred p2p links	inferred s2s links
total number of	3,724	3,070	623	31
number of correct	3,508	2,964	516	28
percentage of correct	94.2%	96.5%	82.8%	90.3%

Questions in the questionnaire

- For the listed inferred AS relationships, specify how many are incorrect, and what are the correct types of the relationships that we mis-inferred?
- What fraction of the total number of your AS neighbors is included in our list?
- Can you describe any AS relationships, more complex than c2p, p2p, or s2s, that are used in your networks?

Missing links

- 27 (3 tier-1 ISPs) out of 38 answered the second question, too, and provided us with their full AS relationship data: 1,114 links
- Among these, we see only 552 (49.6%):
  - 38.7% out of the 865 (77.6%) p2p links
  - 86.7% out of the 218 (19.6%) c2p links
  - 93.3% out of the 30 (2.7%) s2s links
- Maximum percentage of missing links per node is 86.2% (50% of ASs miss >70% links)

Missing links visualized

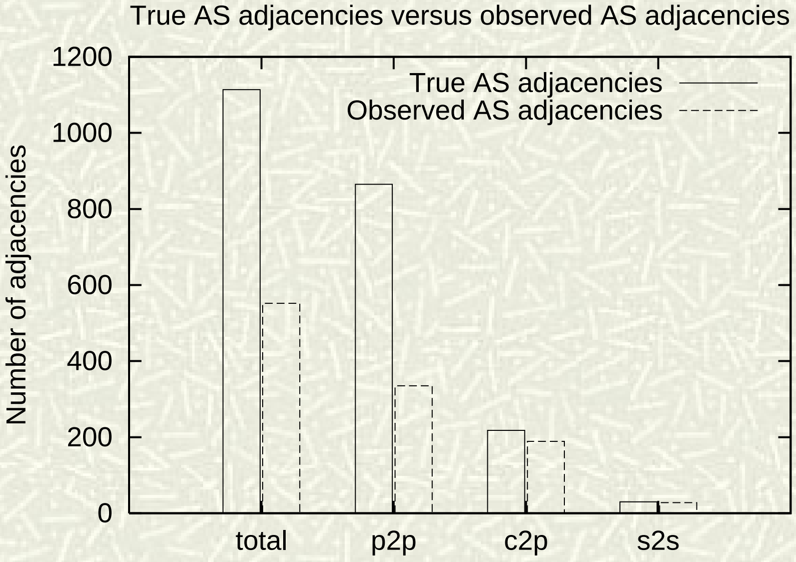

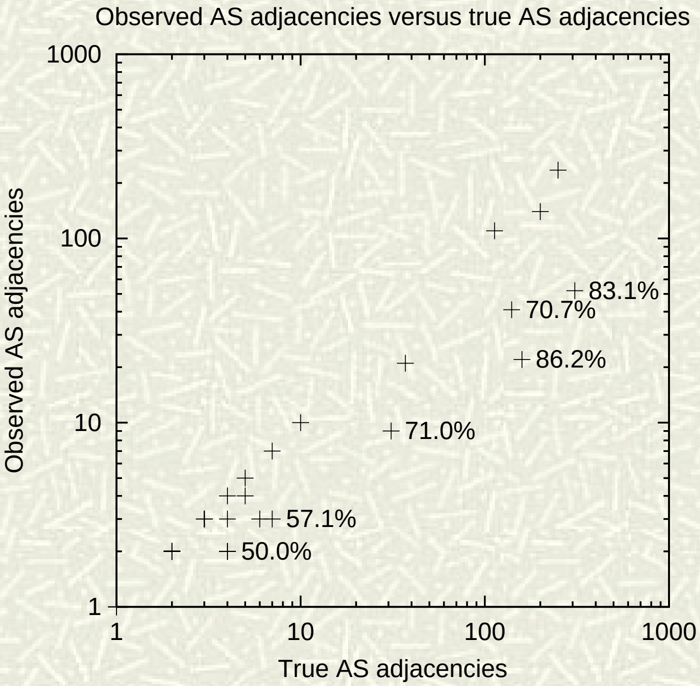

More complex policies

- Space
- Time
- Prefix

Outline

- AS relationships
- AS taxonomy
- AS rank

AS taxonomy

- Assign the following six attributes to every AS
  - organization description (IRR data, stop words are filtered out and the rest of words are stemmed)
  - number of customers
  - number of providers
  - number of peers
  - number of advertised IP prefixed
  - size of the advertised IP address space
- Feed this data into a machine learning algorithm (AdaBoost) with a training set of 1200 ASs
- Classify all ASs into the following six categories
  - Large ISPs
  - Small ISPs
  - Customer ASs
  - Universities
  - IXPs
  - NICs

AS taxonomy results

Classified 95.3% of ASs (non-abstained) with expected accuracy of 78.1%

	Large ISPs	Small ISPs	Customer ASes	Universities	IXPs	NICs
ASes	44	5,599	11,729	877	33	332
%	0.2	30.1	63.0	4.7	0.2	1.8

http://www.caida.org/data/active/as_taxonomy/

Outline

- AS relationships
- AS taxonomy
- AS rank

AS rank