Progress in inferring business relationships between ASs

Dmitri Krioukov

dima@caida.org

4th CAIDA-WIDE Workshop

Outline

- Overview of ongoing projects
- AS relationship inference
  - Motivation
  - Problems with recent heuristics
  - How we fix them
  - Future work

Outline

- Overview of ongoing projects
- AS relationship inference
  - Motivation
  - Problems with recent heuristics
  - How we fix them
  - Future work

Projects

- Routing is dormant
- Comparative analysis of Internet topologies (with P. Mahadevan, UCSD)
- Internet topology evolution modeling (with R. Liu, UCLA)
- Improving AS business relationship inference heuristics (with X. Dimitropoulos, Georgia Tech, and B. Huffaker, CAIDA)

Outline

- Overview of ongoing projects
- AS relationship inference
  - Motivation
  - Problems with recent heuristics
  - How we fix them
  - Future work

Motivation

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

Outline

- Overview of ongoing projects
- AS relationship inference
  - Motivation
  - Problems with recent heuristics
  - How we fix them
  - Future work

Previous results: 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

Previous results: 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

Outline

- Overview of ongoing projects
- AS relationship inference
  - Motivation
  - Problems with recent heuristics
  - How we fix them
  - Future work

Idea at the high level

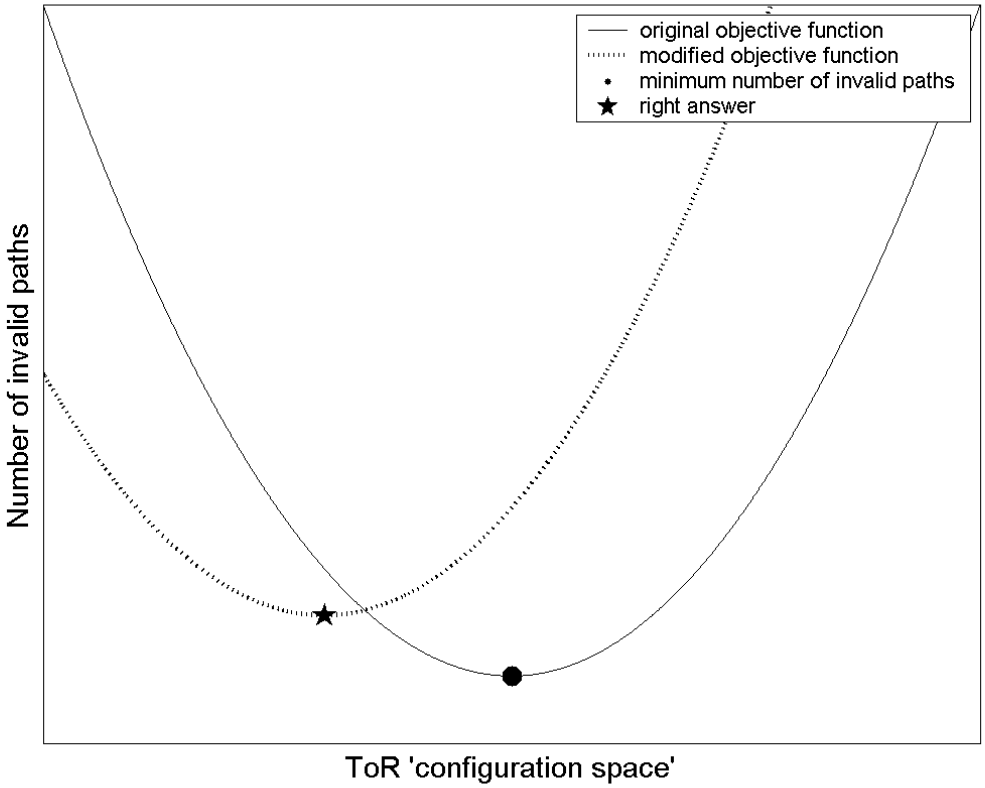

Objective function adjustment

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}
	$x_i \vee x_j$	
	$x_i \vee \overline{x}_j$	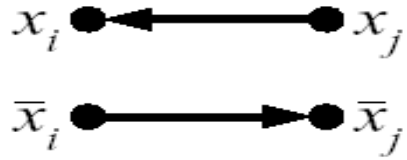
	$\overline{x}_i \vee x_j$	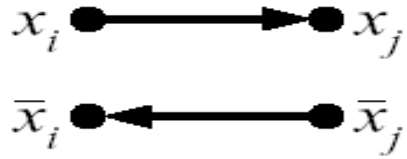
	$\overline{x}_i \vee \overline{x}_j$	

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

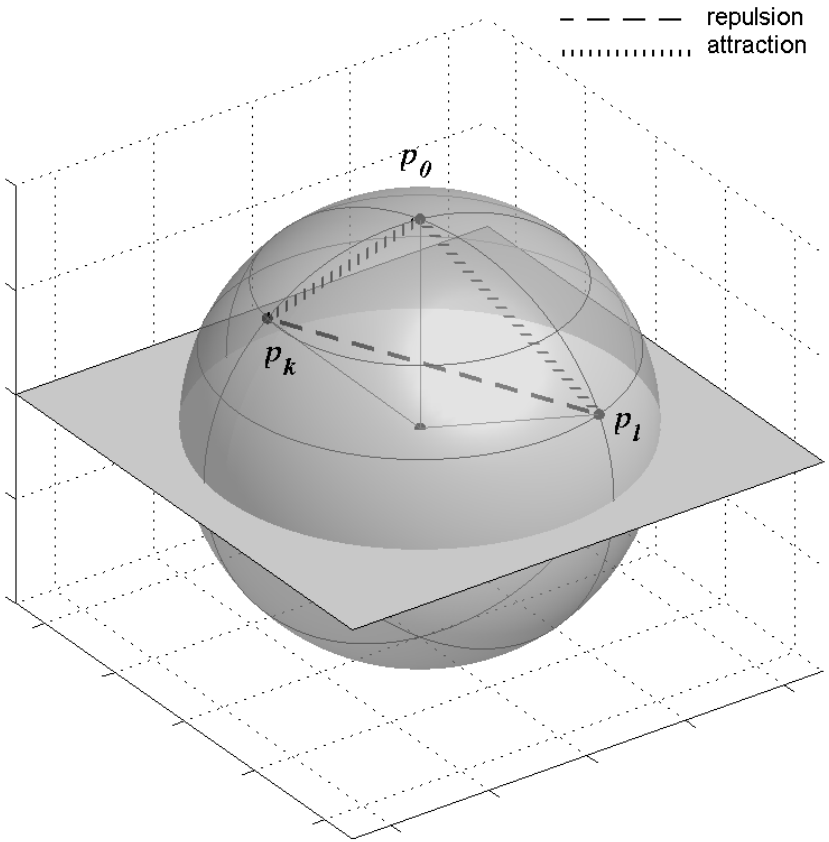

Gains and losses

- What's good
  - Extremely small number of invalid paths (*99.7%*)
- 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

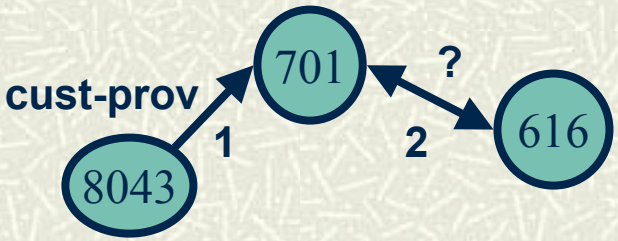

- **Case 2:** trying to infer sibling links leads to proliferation of error
- **Fix:** try to discover sibling links using the WHOIS database

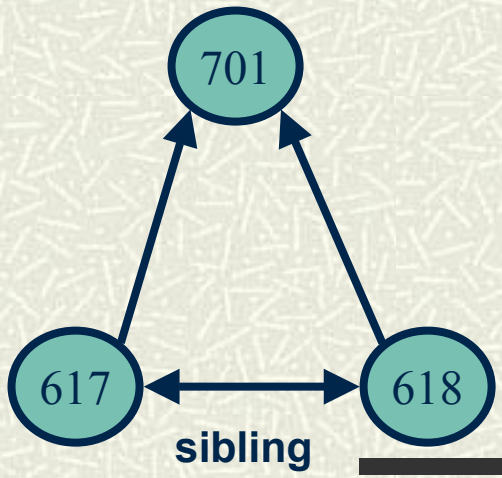

Case 1: Multiobjective optimization

- 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

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

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

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

Execution process

- BGP data: RouteViews and 18 route servers from traceroute.org (1025775 paths, 17557 ASs, 37021 1-link clauses/edges, 382917 2-link clauses/pairs of adjacent edges)
- Resolve siblings that can be found via WHOIS
- Direct all edges along the node degree gradient
- Strip non-conflict edges whose non-conflict direction, inferred by topological sorting in G_{2SAT} , satisfy all clauses containing corresponding variables (remains: 1590 ASs, 4249 1-links, 23460 2-links)
- Execute the algorithm for different values of α: intensive usage of:
  - LEDA library
  - DSDP solver

First results: invalid paths

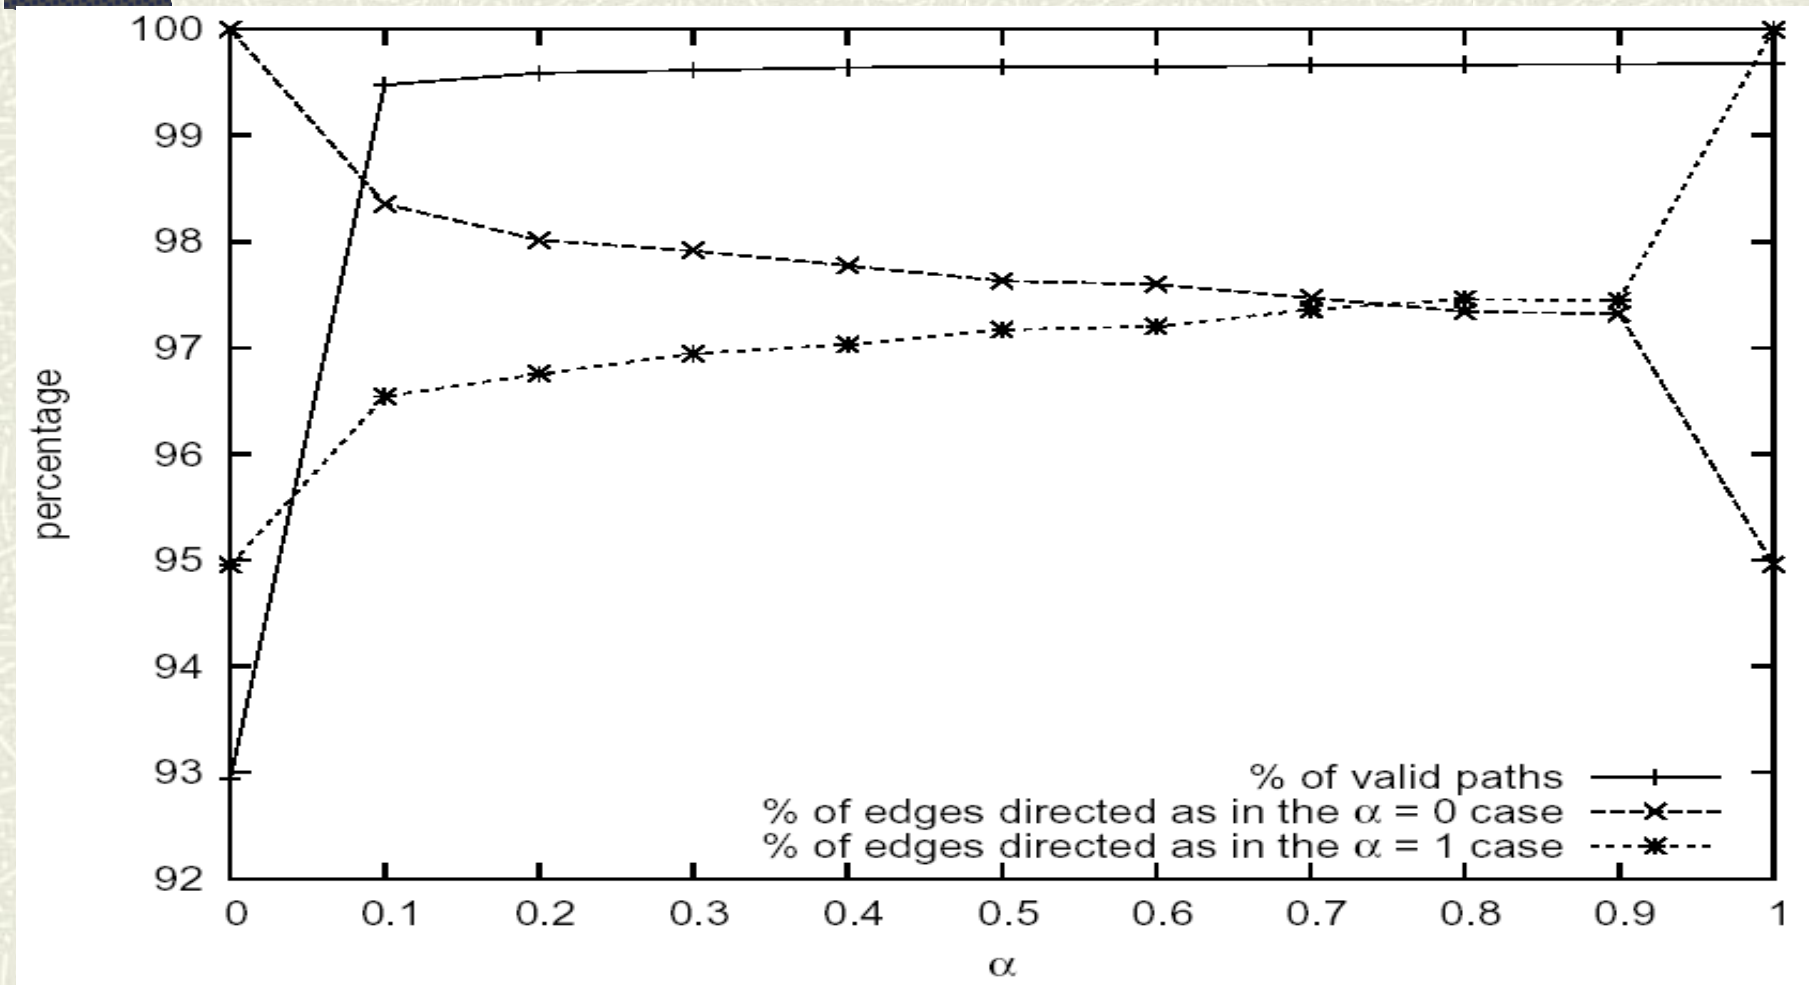

First results: AS hierarchy

Table 2: Ranking of ASs induced by our inference algorithm. The ranks of the top five ASs for $\alpha = 0$ and $\alpha = 1$ are shown for different values of α. The AS numbers are matched to an AS name using the WHOIS databases.

AS #	AS name	AS outdegree	$\alpha = 0.0$	$\alpha = 0.2$	$\alpha = 0.4$	$\alpha = 0.6$	$\alpha = 0.8$	$\alpha = 1.0$
701	UUNET	2373	0-1	0-173	0-217	1-242	1-252	17-476
1239	Sprint	1787	1-1	0-173	0-217	1-242	1-252	17-476
7018	AT&T	1723	2-1	0-173	0-217	1-242	1-252	17-476
3356	Level 3	1085	3-1	0-173	0-217	1-242	1-252	17-476
209	Qwest	1072	4-1	0-173	0-217	1-242	1-252	17-476
11551	Pressroom Services	2	1742-941	1419-398	1435-391	1449-390	1457-386	0-4
6721	Nextra Czech Net	3	1742-941	833-88	853-90	874-90	884-89	0-4
3643	Sprint Australia	17	194-1	222-1	233-1	261-1	268-1	0-4
1243	Army Info. Systems	2	2683-62853	2753-14655	1435-391	1449-390	1457-386	0-4
6712	France Transpac	2	2683-62853	2753-14655	2774-14634	298-2	1-252	4-13

Conclusion and future work

- What's done
  - Using a standard multiobjective optimization method, resolved issues associated with two polar approaches
- What's to be done
  - Peering
  - Further analysis of the 'trade-off surface'
  - Further analysis of the directed AS graph
  - Validation
  - Other AS-ranking mechanisms
  - AS-ranking webpage update http://www.caida.org/analysis/topology/rank_as/