Near-Deterministic Inference of AS Relationships

Udi Weinsberg Yuval Shavitt Eran Shir
School of Electrical Engineering, Tel-Aviv University, Israel

I. Introduction

The commercial Internet consists of three major type-of-relationships (ToR) between neighboring ASes: customer-to-provider (c2p), peer-to-peer (p2p), and sibling-to-sibling (s2s). Since ToRs are regarded as proprietary information, deducing them is an important yet difficult problem. Gao [7] was the first to study the AS relationships inference problem and deduced that every BGP path must comply with a hierarchical *valley-free* (or valid) structure: an uphill segment of zero or more c2p or s2s links, followed by zero or one p2p links, followed by a downhill segment of zero or more p2c or s2s links.

Current relationships inference algorithms solve the ToR problem either using heuristic assumptions or by optimizing some aspects of ToR assignments. Optimization is usually achieved by minimizing the number of paths that violate valley free routing [10] while not allowing cycles to be created [4] in the directed relationships' graph. Inference using heuristic assumptions causes the erroneous ToRs to be spread over all inferred links. Optimization models might fail to capture the true Internet hierarchy [6] and have a relatively low p2p inference accuracy [12]. Hence, both solutions fail to provide an insight or a bound on inference errors.

This paper aims to improve on existing methods by providing a near-deterministic inference scheme (*ND-ToR*) for solving the ToR problem. The input for ND-ToR is the Internet *Core*, a sub-graph that consists of the globally top-level providers of the Internet and their interconnecting links with their already inferred relationship types. Theoretically, given an accurately classified core, the algorithm *deterministically* infers most of the remaining AS relationships using the AS-level paths relative to this core, without incurring additional inference errors. In real-world scenarios, where the core and AS-level paths can contain errors (due to misconfigurations or measurements mistakes), the algorithm introduces minimal inference mistakes. We show that ND-ToR has relaxed requirements from the core, and proves to be robust under changes in its definition, size and density.

II. Near Deterministic Inference

The deterministic algorithm receives as input the undirected AS graph and the core graph. Prior to starting the relationships inference algorithm, we infer s2s relationships using s2s data collected from [2], since ignoring these relationships might cause proliferation of erroneous inference [5]. Once classified, s2s links are removed from the graph and the two adjacent vertices are united to form a single vertex that inherits the connectivity of both. To further avoid incorrect inferences resulting from transient routing effects we use voting [7] instead of direct relationship inference, where each method "votes" for the ToR it infers. Once the phases are finished, each edge is classified to the type that received a relative votes count that passes a fixed threshold.

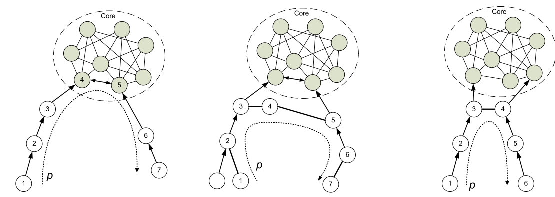

(a) Phase 1 (b) Phase 2 (c) Phase 3

Fig. 1. Deterministic ToR inference algorithm

Phase 1. All paths that traverse are classified using the valley-free rule (see Fig. 1(a)), starting with the uphill segment of the path, classifying each edge as c2p, until reaching the core. Inside the core edges edges are classified as p2p. Downhill links from the core are classified as p2c. Invalid paths are detected when an edge is directed towards the core (uphill) during the downhill segment.

Phase 2. For a given path (see Fig. 1(b)), edges that precede a c2p edge are in an uphill segment, and are of type c2p. Edges that follow a p2c edge are in a downhill segment, and are of type p2c. This phase is repeated until there are no more edges that can be classified using this method.

Phase 3. The deterministic phases (1 and 2) fail to classify edges that appear in paths that do not traverse the core, and reside between a c2p edge and a p2c edge (see Fig. 1(c)). The algorithm heuristically classifies a *single* remaining unclassified edge between a c2p and a p2c edges as a p2p edge.

III. Experimental Results

We evaluate NDToR using data merged data from Route-Views (RV) [8] and DIMES [9] from the first five weeks of 2007. The data set consists of over 24,000 AS vertices and approximately 58,000 undirected links. Two core graphs that vary in size and spatial properties are evaluated. Greedy-Max-Clique core (GMC) [11] holds the clique of descending degree ASes holding mostly tier-1 ASes, resulting in an average of 17 ASes and 272 links. *k*-Core is created from the nucleus [3] extracted from k-pruning algorithm. Both techniques provide consistent cores between consecutive weeks. However, while most GMC ASes have high degrees and are located in North America and Europe, the ASes in *k*-Core include all the ASes

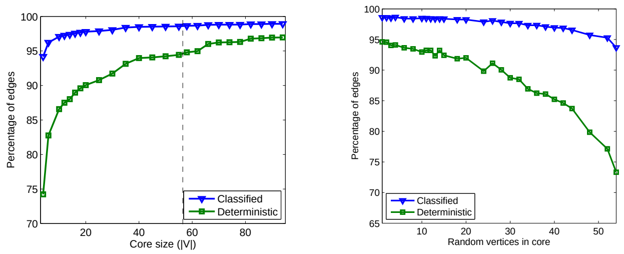

(a) Size (b) Error

Fig. 2. Robustness of ND-ToR

in GMC and additional ASes with more diverse geographic and degree distributions, showing that k-Core has a more global view of the Internet AS-level graph. An important feature of k-Core is that it holds ASes that lie in a large fraction of the paths that connect different ASes, resulting in better deterministic inference percentage. Additionally, both k-Core and GMC exhibit relatively consistent results over the five weeks period. Breaking voting ties when using GMC, is done by comparing adjacent ASes' degrees [7] and using k-Core by comparing the k-Shell index of two adjacent ASes.

Fig. 2(a) shows the robustness of ND-ToR to the size of the input core (k-Core). The overall inference converges within 10% of the maximal success after 20 ASes in the core. However, further increasing the core also increases the number of non-valley-free paths, implying that the core must be kept small enough to decrease the number of invalid paths.

Finding a time frame for which ND-ToR captures best the relationships between ASes was done using data from an increasing time frame starting from a single day until reaching 10 consecutive weeks. We found that using data from a single week results in over 90% of the edges being classified for both core types, and over 98% of the inferences remain constant between consecutive time frames, which suggests that there are only a few commercial relationships that change over time.

Estimating the accuracy and robustness of ND-ToR is done by increasing the "mistake" in the core by randomly replacing ASes in the core. Fig. 2(b) shows the percentage of overall classified edges and deterministically classified edges (phases 1 and 2) using k-Core. While ND-ToR's performance decreases as randomness increases, the overall degradation is not as high as one would expect. However, as more errors are injected, ND-ToR needs to use more heuristics reaching to roughly 20% of all edges using a completely random core.

IV. Validation and Discussion

Validating the accuracy of the results was done by comparing ND-ToR to perviously suggested heuristic algorithms (GAO [7], CAIDA [5] and BPP [1]). The p2c links are similar between compared algorithms, having ND-ToR disagrees on only 0.05% of the p2c links, showing that it is possible to deterministically infer p2c regardless of the input core. However, the algorithms agree on only 21% of the p2p links that are inferred by ND-ToR. Using k-Core, p2p inferences made by ND-ToR are similar to CAIDA while using GMC, p2p inferences are similar to GAO. This can be attributed to the fact that GAO uses very local-based heuristics, which are similar to the local properties of GMC, while CAIDA takes a more global approach, which is closer in spirit to k-Core.

Analysis of the correlation between inference types comparing ND-ToR to GAO, CAIDA and BPP reveals that although the majority of links are classified the same, there is almost no distinct agreement. Even for the p2c links, there is a small number of links (0.3% with GAO, 1.5% with CAIDA, 4.3% with BPP) that are inferred in the wrong "direction". Additionally, the overall percentage of unresolved edges is very small (less than 0.7%) for all algorithms except CAIDA (that has less edges, since it uses only BGP paths). Among CAIDA's inferences, p2p relationship has the highest percentage of unresolved edges (compared to ND-ToR, over 32% using k-Core and over 42% using GMC) which strengthen the claim that most of the links missing from BGP paths are of type p2p, therefore strengthens the need for inclusion of active probing data.

V. Conclusion

The common weakness of current AS relationships inference algorithms is their lack of guarantee on inference errors caused by extensive usage of heuristics. This work improves on existing methods by providing a near-deterministic algorithm that, given a classified error-free input core, does not introduce additional inference errors. We show that the proposed algorithm provides accurate inferences that are robust under changes in the core's size and spatial properties. Moveover, using k-Core that contains as little as 20 almost fully-connected ASes is sufficient for good inference results.

References

[1] G. D. Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. Technical Report RT-DIA-73-2002, Dipartimento di Informatica e Automazione, Universita di Roma Tre, 2002.

[2] CAIDA. Automated Autonomous System (AS) ranking. Research Project. http://as-rank.caida.org.

[3] S. Carmi, S. Havlin, S. Kirkpatrick, Y. Shavitt, and E. Shir. A model of Internet topology using k-shell decomposition. *Proceedings of the National Academy of Sciences USA (PNAS)*, 104(27), July 2007.

[4] R. Cohen and D. Raz. Acyclic type of relationships between autonomous systems. In *IEEE INFOCOM*, 2007.

[5] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, kc claffy, and G. Riley. AS relationships: Inference and validation. *ACM SIGCOMM Computer Communications Review*, 37:2007, 2006.

[6] X. Dimitropoulos, D. Krioukov, B. Huffaker, kc claffy, and G. Riley. Inferring as relationships: Dead end or lively beginning? *LNCS*, 3503:113, 2005.

[7] L. Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, 2001.

[8] University of Oregon RouteViews Project (http://www.routeviews.org/).

[9] Y. Shavitt and E. Shir. DIMES: let the internet measure itself. *ACM SIGCOMM Computer Communications Review*, 35(5):71–74, 2005.

[10] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the internet hierarchy from multiple vantage points. In *IEEE INFOCOM*, New York, NY, USA, Apr. 2002.

[11] L. Tauro, C. Palmer, G. Siganos, and M. Faloutsos. A simple conceptual model for the Internet topology. In *Global Internet*, Nov. 2001.

[12] J. Xia and L. Gao. On the evaluation of AS relationship inferences. In *IEEE Globecom*, Dallas, TX, USA, Nov. 2004.