



Building an Accurate Model of the Internet

Semester Thesis Proposal

The Internet is a global network consisting of many networks called autonomous systems (ASes). Its structure is highly complex: A few ASes are well-connected and make up the core of the Internet, while most of the ASes only connect to a couple of neighbors that in turn provide them access to the rest of the Internet. These connections are governed by rules, the routing policies. Both the structure and these policies determine the flow of data traffic from one AS to the other through the Internet.

Therefore, being able to precisely model and also reason about inter-domain routing is key to analyze important aspects, such as security, performance and resiliency, of the global Internet.

Unfortunately, these interconnections and their underlying agreements are usually business sensitive and therefore, the routing policies of the ASes are kept private.

Over the years, multiple research initiatives looked at the problem of building an accurate interdomain routing model. The first projects simply tried to recreate the structure of the Internet: they tried to infer all the existing interconnections among the ASes. Later projects [2–7, 10], also looked at the type of those interconnections and tried to derive their relationships.

The problem with these projects is that the employed models are too simplistic and do not allow to adequately capture the Internet's variety. Most approaches consider each AS to be a single point/router. This assumption, for example, prevents route diversity and cannot represent the internal policies (e.g., iBGP). In addition, these projects assumed over-simplistic routing policies. Generally they only allow for a few types of different policies: customer-to-provider, peer-to-peer, and sometimes also sibling-to-sibling.

One project [8] stands out. The authors recognized the above mentioned limitations and tried to drop the simplistic assumptions. Their model manages to capture well the observed routing behavior, but falls short to generalize to unobserved behavior and take dynamics of the Internet into account as it is based on a single point in time.

The goal of this thesis is to build upon the insights of Mühlbauer et al. [8] and build an accurate model of the Internet that is able to predict the outcome of inter-domain routing for unobserved events, such as a link failure or a change of a network in its routing policies.

Based on the publicly available routing data going back many years, you develop your own model.

Milestones

- Start by getting an overview of the different, existing techniques to model inter-domain routing;
- Replicate the model of Mühlbauer et al. [8] and validate their findings;
- Based on the gained insights, our suggestions and your own ideas, propose and build a model;
- Validate and evaluate your model using the publicly available routing data [1, 9].

Contact

- Rüdiger Birkner, rbirkner@ethz.ch
- Prof. Dr. Laurent Vanbever, lvanbever@ethz.ch

References

- [1] CAIDA. BGPStream. <https://bgpstream.caida.org/>.
- [2] G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 1, pages 156–165. IEEE, 2003.
- [3] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, G. Riley, et al. As relationships: Inference and validation. *ACM SIGCOMM Computer Communication Review*, 37(1):29–40, 2007.
- [4] X. Dimitropoulos, D. Krioukov, B. Huffaker, G. Riley, et al. Inferring as relationships: Dead end or lively beginning? In *International Workshop on Experimental and Efficient Algorithms*, pages 113–125. Springer, 2005.
- [5] X. Dimitropoulos, D. Krioukov, G. Riley, and k. claffy. Revealing the Autonomous System Taxonomy: The Machine Learning Approach. *arXiv preprint cs/0604015*, 2006.
- [6] T. Erlebach, A. Hall, and T. Schank. Classifying customer-provider relationships in the internet. *TIK-Report*, 145, 2002.
- [7] L. Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM ToN*, December 2001.
- [8] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. Building an as-topology model that captures route diversity. *ACM SIGCOMM Computer Communication Review*, 36(4):195–206, 2006.
- [9] U. of Oregon. Route Views Project. www.routeviews.org/.
- [10] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the internet hierarchy from multiple vantage points. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 618–627. IEEE, 2002.