

Towards realistic modeling of IP-level routing topology dynamics

Clémence Magnien¹ and Amélie Medem¹ and Fabien Tarissan¹

¹LIP6, CNRS & Université Pierre et Marie Curie - 4 place Jussieu, 75005 Paris, France - Prenom.nom@lip6.fr

Plusieurs travaux ont déjà étudié la topologie de l'Internet, mais peu ont cherché à comprendre comment celle-ci évoluait dans le temps. Ce papier se focalise sur la dynamique de la topologie de l'Internet au niveau IP. Nous analysons plusieurs mesures périodiques constituées de différents arbres de routage collectés à partir d'un unique moniteur vers un ensemble fixe de destinations et identifions des caractéristiques invariantes de leur dynamique. Nous proposons ensuite un modèle simple pour les mécanismes sous-jacents de la dynamique de la topologie de routage. Nos simulations montrent que le modèle proposé capture effectivement les caractéristiques dynamiques observées. Ceci fournit des informations pertinentes sur les mécanismes importants qui gouvernent la dynamique du routage sur Internet.

1 Introduction

Studying the structure of the Internet topology is an important and difficult question. No official map being available, researchers have to conduct costly measurement campaigns, and deal with the fact that the obtained data can be biased [LBCX03]. Studying the dynamics of this topology is therefore an equally hard, if not harder, problem.

In this paper, we use an orthogonal approach to obtain insights on the dynamics of the routing topology observed at the IP-level. We study *ego-centered views* of this topology. Given a monitor and a fixed set of destination nodes, one such view is obtained by measuring the routes from the monitor to the destinations. This can be performed quickly and with low network load with the `tracetree` tool [LMO08]. Repeating this measurement periodically therefore allows to study the dynamics of this view.

Previous work has shown that ego-centered views exhibit strong dynamics, and in particular that the set of observed nodes evolves much more quickly than what was previously expected [MOVL09]. Here, we first analyze in depth these dynamics (Sec. 2). We then rely on simulations to understand the factors behind them and to study their influence. We propose in Sec. 3 a baseline model for the routing dynamics in the Internet that incorporate routing modifications and load balancing. We use the most simple choices for modeling these two factors, and show in Sec. 4 that this model is able to accurately reproduce the behaviors observed in real data. This shows that simple mechanisms such as the ones we take into account play a key role in the Internet routing topology dynamics, giving a strong explanatory value to our model. As such, it represents a significant first step towards the modeling of the Internet IP-level topology and its dynamics.

2 IP routing topology dynamics

The `tracetree` tool [LMO08] collects the *ego-centered view* from a given monitor to a given set of destinations by measuring the routes from this monitor to each destination. This corresponds to a subset of the routing topology, in which nodes are the IP-addresses of routers, and a link exists between two nodes if the corresponding routers are connected at the IP level. Note that the routing topology is different from the physical topology, as two routers may be physically connected by a link that is not used for routing. Running the `tracetree` tool periodically allows to capture the dynamics of ego-centered views. We collected two datasets in this way. The first one, `woolthorpe`, was collected from a monitor in University Pierre and Marie Curie towards a set of 3,000 destinations. The frequency is of one measurement round every 15 *min* approximately. It started in Dec., 2010 and ended in Jun., 2011, which represents a total of 17,450 rounds.

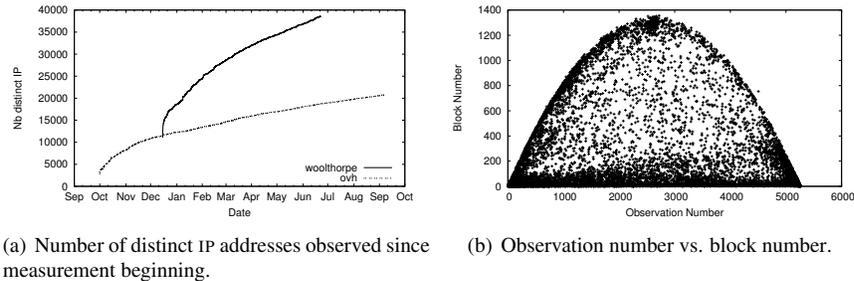


Figure 1: Properties of the observed dynamics.

The second one, `ovh`, was collected from a French server hosting company. Only 500 destinations were used, thus allowing to increase the frequency to one round every 90 seconds approximately. It started in Oct., 2010 and ended in Sep., 2011, with a total of 318,000 rounds. In both cases, the destinations were chosen by sampling random IP addresses that answered to a ping at the time of the list creation[†]. These datasets are publicly available [Rad]. We present below the main characteristics of their dynamics.

Discovery of new IP addresses. A previous study has shown that these measurements continuously discover new IP addresses that had never been observed before, at a significant rate [MOVL09]. These observations were made on two-months-long measurements. Fig. 1(a) shows that it is also true for very long measurements. It presents the number of IP addresses observed since the beginning of the measurement. A dot (x, y) in this figure means that y different addresses have been observed at least once before time x . We see that, after an initial fast growth, the plot increases significantly for extended periods of time.

Stability of IP addresses. We compute two quantities for each IP address. Its *observation number* is simply the number of distinct rounds it was observed in. Its *block number* is the number of groups of consecutive rounds in which it is observed. For example, an IP address observed on rounds 1, 3, 4, 7, 8, 9, and 10 has an observation number of 7 and a block number of 3. Fig. 1(b) presents the correlation between these quantities for the `woolthorpe` dataset. The plot presents a clear parabolic shape, with a large number of points close to the x -axis and to the line $y = x/2$. The presence of a large number of IP addresses close to the parabola can be explained by load-balancing routers. If a load-balancing router randomly spreads traffic among k paths[‡], each router belonging to any of these paths has a probability $p = 1/k$ of being observed at each round, leading to an observation number equal to rp approximately (r being the total number of measurement rounds). A given round is then the first of a consecutive block of observations for one of these routers with the probability p that it was observed in this round, multiplied by the probability $1 - p$ that it was not observed in the previous round. Multiplying this probability by r gives the expected block number, which is then equal to $rp(1 - p)$ and is the equation of the parabola. IP addresses close to the x -axis are mainly observed during blocks of consecutive rounds, with few interruptions. Points close to the line $y = x/2$ correspond to addresses that are observed only during finite part of the measurement, and have during that time a probability $p = 1/2$ of being observed, due to load balancing: an IP address, observed with $p = 1/2$ during k rounds, has an observation number of $x = k/2$, and a block number of $y = k(1/2)^2 = x/2$.

3 Model

Our purpose here is to propose relevant and simple mechanisms that reproduce the observations made in the previous section. We do not aim at proposing a realistic model, but rather at providing a first and significant step towards understanding the impact of simple mechanisms on the observed dynamics. This model incorporates four ingredients: the routing topology, the routes from the monitor to the destinations in this topology, load balancing, and routing modifications. For modeling each ingredient, we try to make the simplest choice possible, our goal being to obtain a baseline model which makes it possible to investigate the role of each component, and to which future and more realistic models should be compared.

[†] Previous work has indeed shown that tracing routes to unused IP addresses can introduce measurement artifacts.

[‡] It has been shown [ACO⁺08] that per-packet or per-flow load-balancing routers spread `traceroute` probes equally among all paths to the destination, which is roughly equivalent to randomly choosing a path.

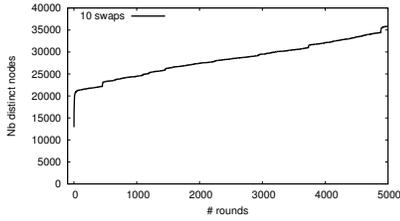


Figure 2: Node discovery for $s = 10$ ($n = 500,000$, $m = 1,000,000$, $d = 3,000$).

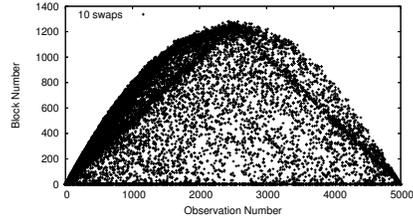


Figure 3: Observation number vs. block number for $s = 10$ ($n = 500,000$, $m = 1,000,000$, $d = 3,000$).

First, we represent the routing topology by a random graph obtained with the Erdős-Rényi model [ER59], which makes no hypothesis on the structure of the graph and is therefore the simplest model possible[§]. Given the graph, we assume that the route between the monitor and a destination is a shortest path obtained by performing a *breadth-first search (BFS)*. To simulate load balancing, the neighbors of each node are considered in a random order, which we call a *random BFS*. Second, we use a simple approach based on link rewiring, or *swap*, to model changes in the routing topology. It consists in choosing uniformly at random two links (u, v) and (x, y) and swap their extremities, *i.e.* replace them by (u, y) and (x, v) .

Finally, our simulation setup consists in the following. First, we generate a random graph G_1 . From G_1 , we randomly select one node as the monitor and d nodes as the destinations. We then simulate r measurement rounds by iterating the following steps: (1) extract a routing tree T_i from G_i ($i \in [1..r]$) by performing a random BFS from the monitor towards the destinations; (2) modify the graph G_i by performing s random swaps, which produces the graph G_{i+1} (s is a parameter of the model). This generates a series of r trees T_1, T_2, \dots, T_r that simulates periodic `tracertree` measurements.

4 Results

This section shows that our model is relevant to explain the dynamic properties presented in Sec. 2.

Node discovery. We first examine whether the simulations reproduce the evolution of the number of distinct nodes observed over time. Fig. 2 presents such an analysis on a graph with $n = 500,000$ nodes, $m = 1,000,000$ links, $d = 3,000$ and $s = 10$. It shows a similar behavior to the one we observed in real data (see Fig. 1(a)). In particular, the curve presents clearly a fast initial growth and then a linear progression. Note that we performed several simulations with varying parameter values and obtained similar results.

Observation number vs. block number. Using the same parameters, we then study the correlation between the observation number and block number (see Fig. 3). The main invariants observed in Fig. 1(b) are reproduced: the parabola, the $y = x/2$ line and a dense strip close to the x -axis[¶]. As already explained in Sec. 2, the line $y = x/2$ corresponds to nodes that are observed with probability $p = 1/2$ for a given duration, and are not observed before or after. There is also a high density of nodes on a line with equation $y = (r - x)/2$, r being the total number of rounds performed. Similarly, it corresponds to nodes which are observed with probability $p = 1/2$ for a given duration, and are observed *at all rounds* before and after that. This line is also present in Fig. 1(b), although not as clearly as here.

Impact of simulation parameters. We also analyze the impact of the simulation parameters on the dynamic properties. In particular, we find that the slope of the node discovery curves increases with the number of swaps s . When the graph does not evolve ($s = 0$), there is an initial growth in which all shortest paths are explored, then the curve becomes flat. This confirms that the constant discovery of new IP addresses in real data may stem from route dynamics. When route dynamics are very high (e.g., $s = 10,000$ for $n = 500,000$, $m = 1,000,000$ and $d = 3,000$), almost all nodes are quickly discovered: there are so many changes of the topology that eventually all nodes will appear at some point on a shortest path.

Concerning the second property, it appears that only the parabola is present when no route dynamics are simulated ($s = 0$), thus confirming that this phenomenon observed in real data is probably due to load

[§] Simulations with power-law random graphs are also available [MMT12]. They show qualitatively similar results as here.

[¶] Note that 20% of all points correspond to $y = 1$.

balancing mechanisms which are well captured by the random BFS model. At the opposite, when the number of swaps s increases, route dynamics get the better of load balancing and the parabola tends to vanish.

Relations between parameters. In order to further study the parameter ranges which allow to reproduce the invariants, we also vary several simulation parameters at the same time. Fig. 4 presents the relation between the number of swaps and the number of links for $n = 500,000$ and $d = 3,000$. We first observe that, for a given number of swaps, the larger the number of links, the smaller the slope of the corresponding curve. This comes from the fact that, when the number of links increases, a smaller fraction of them is affected by swaps. Second, different curves with a same ratio s/m will have the same slope. We can observe this in the two middle curves. Note the sharp increase in one of them close to $x = 2,000$. It is probably caused by a swap happening very close to the monitor, which therefore affects a larger part of the paths than usual. Yet, this does not affect the slope after such an event.

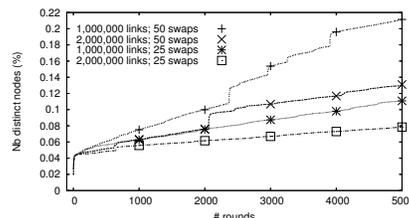


Figure 4: Relation between links and swaps ($n = 500,000$, $d = 3,000$).

5 Conclusion

In this work we conducted periodic measurements of ego-centered views of the Internet topology and isolated two main characteristics of their dynamics. We then proposed a model for the topology dynamics, which integrates simple ingredients such as load balancing and routing changes. This model is not suitable for generating realistic time-evolving topologies; in particular new IP addresses never appear. However, it captures the main characteristics of the observed ego-centered views, which shows that the factors it mimics play a strong role in the Internet routing topology dynamics. We therefore consider our model as a key step towards the realistic modeling of the Internet topology dynamics, as well as towards its understanding.

Acknowledgements: This work is partly funded by the European Commission through the FP7 FIRE project EULER (Grant No.258307). It was also supported in part by a grant from the *Agence Nationale de la Recherche*, with reference ANR-10-JCJC-0202.

References

- [ACO⁺08] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Traceroute Anomalies: Detection and Prevention in Internet Graphs. *Computer Networks*, 52:998–1018, 2008.
- [ER59] P. Erdős and A. Rényi. On random graphs, I. *Publicationes Mathematicae (Debrecen)*, 6:290–297, 1959.
- [LBCX03] A. Lakhina, J.W. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *Proc. of INFOCOM 2003 (IEEE Societies)*, volume 1, pages 332 – 341, 2003.
- [LMO08] M. Latapy, C. Magnien, and F. Ouédraogo. A radar for the internet. In *Proc. of IEEE ICDMW, 2008*, pages 901–908, Washington, DC, USA, 2008.
- [MMT12] A. Medem, C. Magnien, and F. Tarissan. Impact of power-law topology on IP-level routing dynamics: simulation results. In *Proc. of NetSciCom*, in conjunction with INFOCOM, 2012.
- [MOVL09] C. Magnien, F. Ouédraogo, G. Valadon, and M. Latapy. Fast dynamics in internet topology: Observations and first explanations. In *Proc. of IEEE ICIMP, 2009*, pages 137 –142, 2009.
- [Rad] Radar. A Radar for the Internet – Publicly available datasets. <http://data.complexnetworks.fr/Radar/>.