



HAL
open science

Impact of power-law topology on IP-level routing dynamics: Simulation results

Amélie Medem Kuate, Clémence Magnien, Fabien Tarissan

► To cite this version:

Amélie Medem Kuate, Clémence Magnien, Fabien Tarissan. Impact of power-law topology on IP-level routing dynamics: Simulation results. IEEE International Workshop on Network Science For Communication Networks (NetSciCom'12), Mar 2012, Orlando, United States. pp.220-225, 10.1109/INF-COMW.2012.6193495 . hal-01217947

HAL Id: hal-01217947

<https://hal.science/hal-01217947v1>

Submitted on 20 Oct 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Impact of power-law topology on IP-level routing dynamics: simulation results

Amélie Medem, Clémence Magnien and Fabien Tarissan
CNRS and UPMC Sorbonne Universités
LIP6 Laboratory
{amelie.medem, clemence.magnien, fabien.tarissan}@lip6.fr

Abstract—This paper focuses on the Internet IP-level routing topology and proposes relevant explanations to its apparent dynamics. We first represent this topology as a power-law random graph. Then, we incorporate to the graph two well known factors responsible for the observed dynamics, which are load balancing and route evolution. Finally, we simulate on the graph traceroute-like measurements. Repeating the process many times, we obtain several graph instances that we use to model the dynamics. Our results show that we are able to capture on power-law graphs the dynamic behaviors observed on the Internet. We find that the results on power-law graphs, while qualitatively similar to the one of Erdős-Rényi random graphs, highly differ quantitatively; for instance, the rate of discovery of new nodes in power-law graphs is extremely low compared to the rate in Erdős-Rényi graphs.

I. INTRODUCTION

The Internet is a living system that evolves in time. Everyday, many nodes and links are added or removed, during planned maintenance or because of unexpected network failures. It is important to map the Internet topology, in particular when designing future network protocols which can be hard to test on the real Internet. It is equally or even more important to understand its dynamics. This can be very helpful for future protocols or new types of applications to make use of its evolving nature.

Study of the dynamics of the Internet topology has been tackled both by analyzing the dynamics of individual routes [1], [2], [3], [4] and from a more global perspective, mainly at the AS- or IP-level [5], [6], [7], [8], [9]. In addition, routing changes that happen at the IP-level topology does not necessarily imply changes at the physical level and *vice-versa*. This paper focuses on the IP-level routing topology and asks the question of how it evolves over time. Instead of individual routes, we study a *tree* of IP-level routes from one monitor to a fixed set of destinations in the Internet.

In our previous work, we already analyzed the dynamic of the IP-level routing topology discovered around a single node [10]. Using a traceroute-like measurement tool, we periodically probed the route to several destinations from a single monitor in the Internet. This results in a series of routing trees which represent different *ego-centered* views of the routing topology around the monitor. Analyzing these trees, two dynamic behaviors were apparent. In particular, we observed that we never stop discovering new IP addresses over time. It is hard to understand the observed dynamics without

knowing the real Internet topology. Therefore, we relied on simulations to identify the factors behind these behaviors and to study their influence. We proposed a model whose main goal was explanatory. This model represents the Internet IP-level routing topology as an Erdős-Rényi random graph $G = (V, E)$ where vertices correspond to IP addresses and edges correspond to the IP-level connectivity or links between two IP addresses. Then, it incorporates on G well known apparent dynamic factors: load-balancing and route evolution. Finally, it simulates Internet measurements on G to create a routing tree. This process is repeated many times to create several routing trees that we use to analyze the dynamics. From this work, we learn that it is possible to reproduce on Erdős-Rényi graphs the dynamic behaviors observed on the Internet.

This paper goes further and studies the dynamic behaviors by using power-law random graphs to model the routing topology. With Erdős-Rényi random graphs, we made no supposition on the underlying topology. Here, we use a graph with a power-law degree distribution. Indeed, Faloutsos and al. [11] have shown that power-law graphs may be close to the Internet topology in term of their degree sequence, so they may well approximate its structure. We first ask the same questions as in the analysis with Erdős-Rényi graphs: (1) can we reproduce the dynamics behaviors on power-law graphs? and (2) how does the dynamic behavior depend on various simulation parameters? Then, we investigate the differences of results that appear for Erdős-Rényi and power-law graphs.

The rest of the paper is organized as follows. In Section II, we describe two characteristics of the dynamics of the IP-level routing topology around a single monitor. Section III presents the simulation model. In Section IV, we analyze the results of our experiments. Finally, Section V discusses some related works, and Section VI presents our conclusions.

II. IP-LEVEL ROUTING TOPOLOGY DYNAMICS

In our previous work [10], we already presented two main characteristics of the dynamics of the IP-level routing topology around a given monitor. To perform this study, we needed several snapshots of the IP routing topology between the monitor and a given set of destinations. We use *tracetree* [12] which is a traceroute-like measurement tool that aims at discovering a tree of routes or routing paths with the monitor as the root and the destinations as the leaves. The intermediary

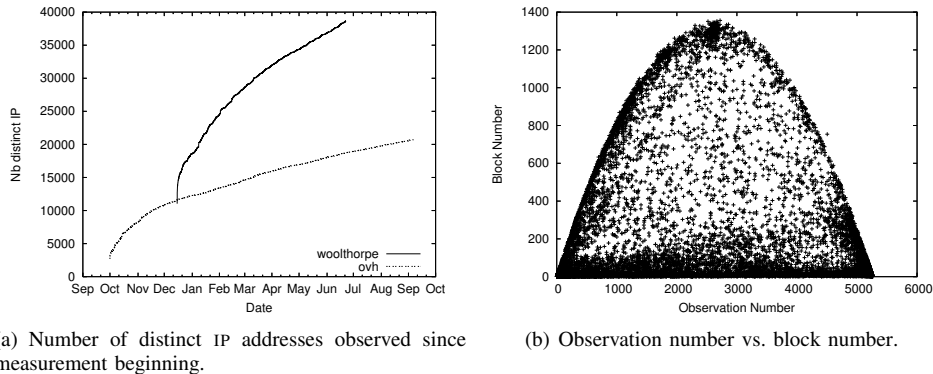


Fig. 1. Properties of the observed dynamics.

nodes of the tree are the IP addresses found on the routing path for each pair (monitor, destination). The size of a tree is then the number of all its nodes (intermediary nodes plus the monitor and destination nodes). The link between two nodes represents a hop at the IP level. One routing tree represents a subset of the IP-level routing topology between the monitor and the destinations. It is called an *ego-centered* view of this topology. Repeating many *tracetree* measurements periodically gives a set of routing trees that one can use to analyze the dynamics. Each measurement round lasts about 4 *min* and the frequency between a pair of rounds is about 15 *min*. Different datasets were collected from many monitors around the world (almost 150 monitors, mostly on PlanetLab) and are publicly available [13].

Analyzing these datasets, two main dynamic characteristics came out: (1) new IP addresses are persistently discovered around the monitor, (2) the pattern of occurrence of IP addresses (number of occurrence/observation, and numbers of blocks of consecutive observation) follows a parabolic shape. Here, we present these characteristics for two of our monitors which are *woolthorpe* and *ovh* [10]. All other monitors exhibit similar results. The collection on *woolthorpe* started in December, 2010 and ended in June, 2011 and 3,000 destinations were used. The monitor *ovh* only used 500 destinations with a higher measurement frequency. It was collected from October, 2010 to September 2011.

a) New IP addresses are persistently discovered around the monitor: Given a set of routing trees T_1, T_2, \dots, T_r , we computed the cumulative union $C_i = \cup_{k=1}^i T_k$. Fig.1(a) plots the size of all sets C_i as a function of time for *woolthorpe* and *ovh*. We observe that new IP addresses are discovered at a fast rate. In other terms, we never stop discovering new IP addresses between the monitor and the destinations over time.

b) The pattern of occurrence of IP addresses follows a parabolic shape: We defined two values that quantify the occurrence of IP addresses around a monitor. First, the *observation number* of an IP address represents the total of distinct rounds in which it occurs. Secondly, the *block number* of an IP address is the number of groups of consecutive rounds in which it is observed. As an example, an IP address which was observed on rounds 1,2,3,5,6,8,9 and 11 has an observation number of 8 and a block number of 4. Fig. 1(b) presents

the correlation between these two quantities for the monitor *woolthorpe*. The plot exhibits a clear parabolic shape, with a large number of points close to the x -axis and to the line $y = x/2$. This can be explained in the following way. The presence 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 paths¹, 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. A given round is then the first of a consecutive block of observations for one of these routers with the probability p that this router 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. This is a simplification of the real case in which a router may belong to paths used by several load balancers, themselves belonging to paths used by other load balancers. In practice, an IP address belonging to load-balanced paths can have any probability p , $0 < p < 1$, of being observed. The set of IP addresses closed to the x -axis are often observed on consecutive rounds. Finally, points on the line $y = x/2$ correspond to IP addresses that are observed only during a finite part of the measurement and have a probability of $p = 1/2$ of being observed during that time, due to load balancing.

III. MODEL

We use the same simulation model we have already proposed in [10], which incorporates the routing topology, its dynamics and the *tracetree* measurements. This model does not aim at being realistic. Its main purpose is to explain the previous dynamic characteristics observed on the *ego-centered* view around a monitor (see Section II).

We represent the IP-level routing topology of the Internet by an undirected graph $G = (V, E)$. Each vertex in V corresponds to an IP address and each edge in E corresponds to the connectivity between two IP addresses. Then, we simulate *tracetree* measurements in G . As a preliminary step, we

¹It has been shown [14] 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.




Fig. 2. Analyzing the impact of simulation parameters on *PL* graphs ($n = 500,000$).

randomly choose one node as the monitor and d nodes as the destinations. Then from G , we inferred a routing tree T with the monitor as the root, and the destinations as leaves. In practice, we simply perform a *breadth-first search (BFS)* starting from the monitor, and then discard all branches that do not lead to the destinations.

At this point, we have a routing tree of *shortest paths* from the monitor to the destinations. The next step is to repeat this procedure r times to simulate r measurement rounds. We simulate load-balancing by a *random BFS*. We model route evolution using link rewiring, or *swap*. This consists in choosing uniformly at random two links (u, v) and (x, y) ² and swap their extremities, *i.e.* replace them by (u, y) and (x, v) .

Our previous work [10] used the Erdős-Rényi random graph model [15] to produce G . Here, we use a random graph with a power-law degree distribution [16]. For power-law graph generation, we use the following procedure : (1) given an exponent α , randomly generate a list of degrees that respects the following power law [17] (d is a degree value, $f(d)$ the frequency of d):

$$f(d) = d^{-\alpha}, \alpha > 0, \quad (1)$$

(2) for each node, create as many half links as the value of its degree, (3) randomly sort the previous list and, (4) connect consecutive half links to form links.

IV. SIMULATION RESULTS

This section investigates whether it is possible to reproduce on power-law (*PL*) graphs the dynamic behaviors observed on the Internet. We perform several simulations with different values of the parameters of the model which are : the number n of nodes, the exponent α , the numbers d of destinations, s of swaps per round and m of links for Erdős-Rényi (*ER*) random graphs. We further look for relations between the simulation parameters that may lead to invariants of the dynamic behaviors, and explore the differences in the simulation results with *PL* and *ER* graphs.

A. Reproducing the evolution of IP addresses discovery

Let us first focus on the evolution of the discovery of new IP addresses over time. As a preliminary step, we vary the number s of edge swaps. Fig. 2(a) presents the simulation results on a

PL graph with $n = 500,000$ and $\alpha = 2.3$, for varying values of s . For this first step, we adopt the same number of destinations as in our measurements ($d = 3,000$).

Three main observations appear from Fig. 2(a). First, it is possible to reproduce on *PL* graphs the fast discovery of IP addresses observed on our *tracetree* data, in particular for $s = 1,000$ or $s = 10,000$ swaps. Second, the larger the number of swaps, the faster new nodes are discovered. In fact, routing paths change more quickly with a larger number of swaps (*e.g.*, for $s = 100,000$), than with a lower number of swaps (*e.g.*, for $s = 10$). Recall that we use edge swaps to simulate link changes due to route evolution. Therefore, increasing the number of swaps also naturally increases the probability for routes between the source and the destinations to change, which leads to the fast discovery of new nodes. Third, the first point of all curves are very close. This means that the number of swaps have no influence on the size of routing trees, which was expected.

For $s = 0$ swaps, the curve has a fast initial growth, and then it remains flat until the last round. In the absence of swaps, the only dynamic observed comes from *load balancing* and not from route evolution. All nodes that belong to load-balancing routing paths are quickly discovered at the beginning.

We never succeed in discovering all nodes for *PL* graphs, even when we swap almost all their links at each round. For instance, using $s = 1,000,000$ swaps leads to the quick revealing of only 60% of the nodes in the graph in less than 1,000 rounds. In the next section, we explore in depth the reasons behind this.

We also test the impact of the number of destinations on the dynamic behaviors. Fig. 2(b) shows the results on a graph with $n = 500,000$, $\alpha = 2.3$ and $s = 1,000$ swaps. We observe that the number of destinations clearly has an influence on the height of the first point of the curves, which represents the size of the first routing tree. The larger the number of destinations, the higher the first point of the curves, *i.e.* the larger routing tree. Also, for different numbers of destinations, the slope of all curves are very similar, but not exactly identical. For instance, one may assume that doubling the number of destinations (*e.g.*, from 1% to 2%), we also double the size of the resulting routing trees. However, for this to be true, two conditions need to be verified: (1) all destinations should be on strictly different routing paths from the monitor,

²We choose them such that the four nodes are distinct.

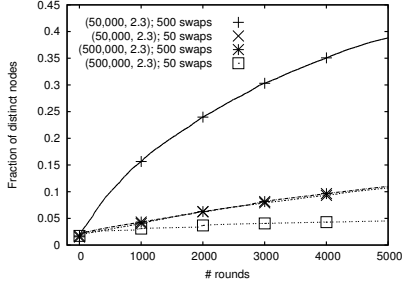


Fig. 3. Relation between size and swaps ($m/n = 2$, $d = 1\%$).

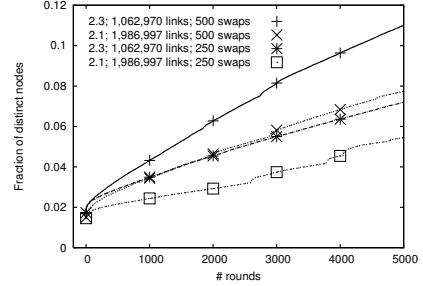


Fig. 4. Relation between links and swaps ($n = 500,000$, $d = 1\%$).




Fig. 5. Observation numbers vs. block numbers for various values of s ($n = 500,000$, $m = 2.3$, $d = 3,000$).

(2) and all destinations should be chosen at the same distance from the monitor. This is not the case in our experiments.

Finally, we vary the exponent α of PL random graphs. Fig. 2(c) presents the simulation results on a graph with $n = 500,000$, $s = 1,000$ and $d = 3,000$, for various values of the parameter α . Note that, the value of the exponent α determines the number m of links for PL graphs. From Fig. 2(c), it appears that the lower the value of α (or the higher the number m of links), the slower the rate of discovering new nodes over time. Indeed, the proportion of links affected by swaps are negligible on graphs with a high number of links. In addition, distances between pairs of nodes are highly reduced on graphs with a high number of links. Indeed, the closer the destinations are to the source, the shortest are the paths between the source and the destinations. Therefore, less new nodes will be discovered over time.

B. Finding relations between simulation parameters

To analyze the interaction between the simulation parameters, we vary several of them at the same time. The goal here is to find invariants which can be very helpful to understand our model in depth.

The first relation we explore is between the size of the graph and the number of swaps. We set $\alpha = 2.3$ and $d = \frac{n}{100}$. The value of α are chosen so that the ratio $\frac{m}{n} = 2$ is verified. Fig. 3 presents the simulation results on two graphs of different sizes $n = 50,000$ and $500,000$, for different values of swaps $s = 500$ and 50 . The y-axis on Fig. 3 represents the fraction of discovered nodes over the total of nodes, and the x-axis, the number of rounds. We observe that the two middle curves are very close and almost follow the same slope. It appears that these curves correspond to graphs with a similar ratio $\frac{s}{m}$ of the number of swaps over the number of links.

The second relation concerns the number of links and the number of swaps. Here, we fix n , and set the proportion of

destinations to 1%. We vary α and s . Fig. 4 presents the results of the simulation on graphs of $n = 500,000$ nodes, with $\alpha = 2.1$ and 2.3 , for $s = 500$ and 250 swaps. It seems that when the number of links doubles, the number of swaps also needs to be doubled as well in order to obtain curves with similar slopes. This result also tends to confirm our previous observation that simulations with the same ratio $\frac{s}{m}$ may follow a similar slope for node discovery. Sometimes, some abrupt increases may deviate a curve from its initial growth rate (e.g., the case $\alpha = 2.1$ and $s = 500$ for $d = 1\%$). We find that these events are caused by swaps that happen close to the monitor and therefore may affect a high number of paths to destinations. However, these events usually do not change the slope of curves.

These results are interesting because they imply that knowing the trend of the evolution curve for a given graph and for a given value of swaps, it can be possible to infer the slope for a range of other graphs. During our experiments, we have tested the previous two relations for other sizes of graphs and obtained the same conclusions. These analyses are at a very preliminary stage. We visually observe the similarities between different curves for various parameters of PL graphs. Later, we may need some statistical analysis to confirm our conclusions.

C. Reproducing the parabolic shape on the occurrence of IP addresses

We now turn to the correlation between the observation numbers and the block numbers. Fig. 5 illustrates our results on a PL graph with $n = 500,000$, $\alpha = 2.3$ and $d = 3,000$, for various values of s . We are clearly able to reproduce the parabola observed on our *tracetree* data (for instance, for $s = 50$ and 100). In that particular case, we also observe that a large number of points are close to the x-axis. For $s = 0$ swaps, all points strictly appear on the parabola. We already know that without swaps, the only dynamic factor in our model




Fig. 6. Comparison between *ER* and *PL* graphs of the same size ($n = 500,000$, $\alpha = 2.3$ (*PL*), $m = 1,000,000$ (*ER*)).

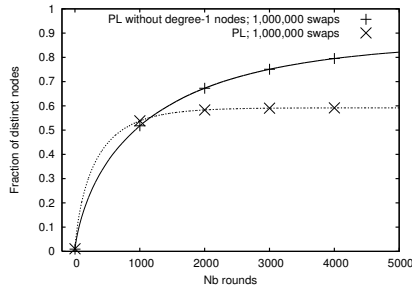


Fig. 7. Comparison between a reduced *PL* graph where all degree-1 nodes are removed and its original ($n = 500,000$, $\alpha = 2.3$).

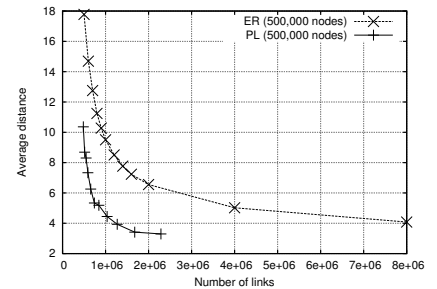


Fig. 8. Number of links vs. the average distance ($n = 500,000$).

is load balancing. This means that nodes are revealed by load-balancing related dynamics. Increasing the number of swaps, the parabola tends to vanish. For instance, for $s = 1,000$ the parabola has already started to vanish. For higher values of s , it completely loses its shape. This can be explained by the fact that with a higher number of swaps, the effects of load balancing becomes negligible. In practice, we find that if we increase the exponent of a *PL* graph, which also increases its number of links, we maintain the parabola if we increase the number of swaps as well.

D. Exploring the differences between *PL* and *ER* graphs

The evolution of node discovery on *PL* and *ER* graphs are very different. Fig. 6 shows that the majority of nodes of an *ER* graph, with $n = 500,000$, $m = 1,000,000$ and $d = 3,000$, are discovered within $r = 5,000$ rounds with only $s = 1,000$ swaps. Using the same amount of swaps on a *PL* graph with approximately the same size and the same number of destinations, we end up discovering only 12% of all its nodes. Clearly, nodes are discovered more slowly over time on *PL* graphs than on *ER* graphs. In this last set of experiments, we investigate the reasons that explain this difference.

Our first intuition concerns nodes of degree one. They represent the largest fraction of nodes on *PL* graphs and, unless they are chosen as destinations, it is impossible for them to be discovered during simulations since they are not router nodes. From Fig. 2(a), we have seen that even when we swap almost all links on a *PL* graph with $n = 500,000$ and $\alpha = 2.3$ at each round, we never succeed in discovering all its nodes. Indeed, with $s = 1,000,000$ swaps the curve tends to flatten out close to the value $y = 295,877$. Examining the remaining nodes, we find that 99.9% of them represent the degree-1 nodes. At the end, we almost discover no nodes of degree 1 on *PL* graphs.

We now ask the question whether a reduced *PL* graph in which we have removed every nodes of degree 1 will follow the same evolution of node discovery as an *ER* graph with the same dimensions; if this is true, then degree-1 nodes may be the only reason of the difference of results observed for *PL* and *ER* graphs. Fig. 6 shows that this is not the case. The reduced *PL* graph has only $n = 293,328$ nodes and $m = 841,326$ links. Therefore, we plot its curves with $d = 1,760$ to maintain the same ratio of the number of destinations over the total of nodes as in the original *PL* graph. With $s = 1,000$ swaps, the curve

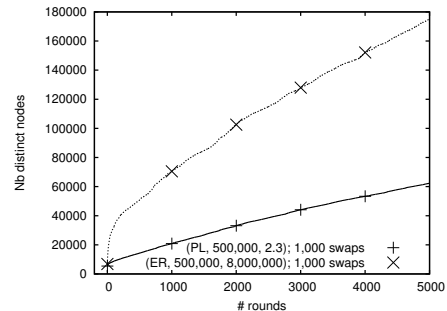


Fig. 9. Comparison between *ER* and *PL* graphs with the same average distance ($n = 500,000$, $d = 3,000$).

of the reduced *PL* graph grows more slowly than the curve of the *ER* graph, but similarly as the one of the original *PL* graph. This means that nodes of degree 1 are not the reasons behind the slow evolution. Increasing the number of swaps until we reach the deprecated case where a maximum is reached for the original *PL* graph in Fig. 7, we end up discovering the majority of nodes of the reduced *PL* graph. This confirms the fact that the flat phase on *PL* graphs are due to the non-discovery of degree-1 nodes.

Our second intuition concerns the difference in the average distance that exists between *PL* and *ER* graphs. It has been proven that the average distance is in the order of $\log \log n$ on *PL* graphs [18], while it is of $\log n$ on *ER* graphs [19]. We explore this result in Fig. 8 which plots the average distance as a function of the number of links for both *PL* and *ER* graphs with $n = 500,000$ nodes. We use the approximation proposed in [20] to compute the average distance. It appears that average distances are effectively much smaller in *PL* than in *ER* graphs. This implies that the destination nodes on *PL* graphs are closer to the monitor; therefore, the resulting routing trees on *PL* graphs will have fewer nodes. To confirm this result, we examine the size of the trees produced on *ER* and *PL* graphs with $n = 500,000$ and the same number of links. We use $d = 3,000$ destinations and $s = 0$ swaps. We find that the average size on $r = 5,000$ trees is 5,363 and 12,868 for *PL* and *ER*, respectively. We then study in Fig. 9 the evolution of node discovery on *PL* and *ER* graphs with the same average distance. We find that they still do not follow the same slope. Finally, this shows that apart from the degree-

1 nodes and the average distance, there are still other factors behind the difference between *PL* and *ER* graphs.

V. RELATED WORK

It has been said that the Internet topology dynamics should be taken into account in order to produce realistic models in the future [21], [8]. Many works empirically studied the dynamics of the Internet topology, mainly by analyzing individual routes [1], [2], [3], [4] or from a more global perspective at the AS or IP level [5], [6], [7], [8], [9], [10]. Most of them analyzed the *physical* AS topology, while we focus on the evolution of the *routing* topology at the IP level. Augustin et al. [14] found that *load-balancing* routers play an important role in the observed route dynamics as measured by traceroute-like tools. Cunha et al. [22] proposed a method for measuring load-balanced routes.

On the other hand, other works used topology evolution models to explain or reproduce the dynamic characteristics observed on the Internet [23], [24], [25], [26]. Park et al. [26] studied different evolution models in which nodes and links are progressively added over time. His goal was to rank them according to their capacity to mimic the evolution of the Internet topology. In our work, we use a model to find the factors that may explain the dynamic behaviors observed on the IP-level routing topology around a single node.

VI. CONCLUSION

This work focuses of the dynamics behaviors observed at the Internet IP-level routing topology. We use an existing model that incorporates a routing topology, its dynamics and traceroute-like measurements to explain the observed dynamics. Our former work represents the routing topology by an Erdős-Rényi random graph. Here, we use a power-law random graph and investigate the effect of its degree distribution on the dynamics. As in Erdős-Rényi graphs, we are able on power-law graphs to reproduce the dynamic behaviors observed on the Internet. However, we find that the results between the two types of graphs are quantitatively different.

Two main reasons for this difference appear: (1) it is difficult to discovered the degree-1 nodes, which represent the largest fraction of nodes on power-law graphs, (2) the average distance on power-law graphs are much smaller than for Erdős-Rényi graphs. Therefore, traceroute-like measurements on power-law graphs produce smaller routing trees, which leads to a slower discovery of new nodes over time.

Future work will proceed in three directions. We first aim to find other Internet dynamic properties. Second, it will be interesting to test our model on real topologies or on more realistic topology models. Indeed, there are many other publicly available traceroute measurements such as data from skitter, DIMES or iPlane. Third, we plan to propose a formal analysis of these dynamics. This will help to formally quantify the role played by the different factors on the dynamics of the Internet IP-level routing topology.

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

- [1] V. Paxson, "End-to-end routing behavior in the Internet," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 601–615, 1997.
- [2] Y. Schwartz, Y. Shavitt, and U. Weinsberg, "On the Diversity, Stability and Symmetry of End-to-End Internet Routes," in *Global Internet*, 2010.
- [3] C. Labovitz, G. Malan, and F. Jahanian, "Origins of Internet routing instability," in *Proc. of IEEE INFOCOM '99*, pp. 218–226 vol.1, 1999.
- [4] A. M. Kuate, R. Teixeira, and M. Meulle, "Characterizing network events and their impact on routing," in *Proc. of ACM CoNEXT 2007, student workshop*, pp. 1–2, ACM, 2007.
- [5] R. Govindan and A. Reddy, "An Analysis of Internet Inter-Domain Topology and Route Stability," in *Proc. of INFOCOM*, IEEE, 1997.
- [6] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "The Origin of Power-Laws in Internet Topologies Revisited," in *Proc. of IEEE Infocom*, IEEE, 2002.
- [7] R. Oliveira, B. Zhang, and L. Zhang, "Observing the Evolution of Internet AS Topology," in *Proc. of ACM SIGCOMM*, 2007.
- [8] S.-T. Park, A. Khrabrov, D. Pennock, S. Lawrence, C. Giles, and L. Ungar, "Static and dynamic analysis of the Internet's susceptibility to faults and attacks," *Proc. of IEEE INFOCOM*, pp. 2144–2154, 2003.
- [9] J.-J. Pansiot, "Local and Dynamic Analysis of Internet Multicast Router Topology," *Annales des télécommunications*, vol. 62, pp. 408–425, 2007.
- [10] C. Magnien, A. Medem, and F. Tarissan, "Towards realistic modeling of IP-level routing topology dynamics," in <http://arxiv.org/abs/1112.4645v1>, arXiv:1112.4645v1, Apr 2011.
- [11] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the internet topology," in *Proc. of SIGCOMM '99*, (NY, USA), pp. 251–262, ACM, 1999.
- [12] M. Latapy, C. Magnien, and F. Ouédraogo, "A radar for the internet," in *Proc. of the 2008 IEEE International Conference on Data Mining Workshops (ICDMW)*, (Washington, USA), pp. 901–908, IEEE Computer Society, 2008.
- [13] "A Radar for the Internet – Publicly available datasets." <http://data.complexnets.fr/Radar/>, 10August2011.
- [14] 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*, vol. 52, pp. 998–1018, 2008.
- [15] P. Erdős and A. Rényi, "On random graphs, I," *Publicationes Mathematicae (Debrecen)*, vol. 6, pp. 290–297, 1959.
- [16] E. A. Bender and E. R. Canfield, "The asymptotic number of labeled graphs with given degree sequences," *Journal of Combinatorial Theory (A)*, vol. 24, pp. 357–367, 1978.
- [17] W. Aiello, F. Chung, and L. Lu, "A random graph model for massive graphs," in *Proc. of STOC '00*, (NY, USA), pp. 171–180, ACM, 2000.
- [18] F. Chung and L. Lu, "The average distances in random graphs with given expected degrees," *Proceedings of the National Academy of Sciences*, vol. 99, pp. 15879–15882, Dec 2002.
- [19] B. Bollobas, *Random Graphs*. Cambridge University Press, 2001.
- [20] M. Latapy and C. Magnien, "Complex network measurements: Estimating the relevance of observed properties," in *Proc. of INFOCOM 2008*, pp. 1660–1668, 2008.
- [21] H. Haddadi, S. Uhlig, A. Andrew Moore, R. Mortier, and M. Rio, "Modeling internet topology dynamics," *ACM SIGCOMM Computer Communication Review*, vol. 38, p. 65, Mar 2008.
- [22] I. Cunha, R. Teixeira, and C. Diot, "Measuring and Characterizing End-to-End Route Dynamics in the Presence of Load Balancing," in *Passive and Active Measurement Conference*, 2011.
- [23] H. Chang, S. Jamin, and W. Willinger, "Internet connectivity at the AS-level: an optimization-driven modeling approach," in *Proc ACM SIGCOMM MoMeTools workshop*, 2003.
- [24] H. Chang, S. Jamin, and W. Willinger, "To peer or not to peer: modeling the evolution of the internet's AS-level topology," in *Proc. of IEEE INFOCOM*, 2006.
- [25] X. Wang and D. Loguinov, "Wealth-based evolution model for the internet AS-level topology," in *Proc. of IEEE INFOCOM*, 2006.
- [26] S.-T. Park, D. M. Pennock, and C. L. Giles, "Comparing static and dynamic measurements and models of the Internet's AS topology," in *Proc. of IEEE Infocom*, IEEE, 2004.