

Towards realistic modeling of IP-level routing topology dynamics

Clémence Magnien · Amélie Medem · Sergey Kirgizov · Fabien Tarissan

Received: date / Accepted: date

Abstract Many works have studied the Internet topology, but few have investigated the question of how it evolves over time. This paper focuses on the Internet routing IP-level topology and proposes a first step towards realistic modeling of its dynamics. We study periodic measurements of routing trees from a single monitor to a fixed destination set and identify invariant properties of its dynamics. Based on those observations, we then propose a model for the underlying mechanisms of the topology dynamics. Our model remains simple as it only incorporates load-balancing phenomena and routing changes. By extensive simulations, we show that, despite its simplicity, this model effectively captures the observed behaviors, thus providing key insights on relevant mechanisms governing the Internet routing dynamics. Besides, by confronting simulations over different kinds of topology, we also provide insights on which structural properties play a key role to explain the properties of the observed dynamics, which therefore strengthens the relevance of our model.

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 [19, 1]. Studying the dynamics of this topology is therefore an equally hard, if not harder, problem.

Instead of trying to obtain a complete view of the Internet topology dynamics, it is possible to use an orthogonal approach to obtain insight on the dynamics

of the routing topology observed at the IP-level [20]. In this paper, we follow this approach and study *ego-centered views* of this topology. Given a monitor and a fixed set of destinations, one such view is obtained by measuring the routes from a monitor to a set of destinations. This can be performed quickly and with low network load with the `tracetree` tool [20]. 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 [23]. Here, we analyze in depth this dynamics (Section 2), and find that two factors play a key role in the observed dynamics: load-balancing routers, and the evolution of the routing topology (Section 3).

Based on these observations, we propose (Section 4) a baseline model for the routing dynamics in the Internet that incorporates routing modifications and load balancing, using simple choices for modeling these factors. We perform an in-depth study of the model behavior (Section 5) and study the impact of the underlying topology structure by comparing the results obtained for random graphs and graphs with a power-law degree distribution. Finally, we show that this model is able to accurately reproduce the behaviors observed in real data. Our results show 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 Dynamics analysis

The `tracetree` tool [20] 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. Note also that we only observe a subset of the whole routing topology, as measuring the routes from a single monitor to a limited set of destinations certainly does not allow to discover all nodes and links in this topology. Moreover, this subset is not representative of the whole topology, see for instance [19, 1]. Keeping this in mind, we will see that we are still able to make interesting observations about the *dynamics* of this topology.

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 in Paris towards a set of 3,000 destinations. The collection frequency is of one measurement round every 15 minutes approximately. It started in September, 2011 and lasted approximately a year with some small interruptions due to power shortage. This represents a total of 32,018 rounds. The second one, `ovh`, was collected from a French server hosting company. Only 500 destinations were used in order to increase the measurement frequency, which is of one round every one and a half minute approximately. It started in October, 2010 and ended in September, 2011, which represents 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 creation¹. These datasets are publicly available [30].

It is possible that, at a given time, several routes to a same destination co-exist, in particular because of load balancing. Therefore, two consecutive measurement rounds may capture different routes to a same destination even if no routing change has occurred. We study this in the next section, and present below the main characteristics of the *observed* dynamics. Notice that previously measured datasets are available, for different durations, at different times since 2008 [30]. We performed our analysis on a representative set of these datasets, and made similar observations to the ones we

present here. This shows the generality of our observations.


Discovery of new IP addresses. A previous study of the same type of data has shown that these measurements continuously discover new IP addresses that had never been observed before, at a significant rate [23]. 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, for both datasets. A dot (x, y) in this figure means that y different addresses have been observed at least once before time x ². We see that, after an initial fast growth, the plot increases significantly for extended periods of time.

This plot presents the number of distinct IP addresses observed, and not the number of distinct routers, as in general several IP addresses, or interfaces, correspond to a same router. Detecting which interfaces correspond to which routers is a difficult task. Though several methods exist, none is 100% accurate. We used the MIDAR tool developed by CAIDA [5], and studied the number of discovered *routers* observed since measurement beginning. The corresponding plot, not presented here due to lack of space, clearly displays the same shape as those of Figure 1(a). Moreover, previous work has studied the number of distinct ASes discovered by such measurements, and showed that it also increases significantly [23]. All in all, there is a good evidence that new routers are actually discovered at a significant rate, even if part of the observed growth may be caused by discovering new interfaces for already observed routers. As there is no method that allows to know with certainty which interfaces correspond to a same router, we limit ourselves to the study of interfaces in the rest of the paper.


Stability of IP addresses. To analyze more in depth the dynamics of the ego-centered views, we compute two quantities for each IP address. Its *observation number* is simply the number of distinct rounds it was observed in. An IP address is in general observed in blocks of several consecutive rounds, preceded and followed by one or more rounds during which it is not observed. More precisely, the *block number* of an IP address is the number of groups of consecutive rounds in which it is observed. For example, an IP address which was observed

¹ Previous work has indeed shown that tracing routes to unused IP addresses can introduce measurement artifacts [37].

² Since the `woolthorpe` dataset was collected after the end of the `ovh` dataset measurement, we shifted x -axis one year for the plot for the `woolthorpe` dataset, so that both plots appear in the same time span.



(a) Number of distinct IP addresses observed since measurement beginning.



(b) Observation number vs. block number.

Fig. 1 Properties of the observed dynamics.

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` dataset³. Each dot corresponds to an IP address, and its coordinates are its observation number on the x -axis and its block number on the y -axis. The plot presents a clear parabola shape. This can be explained by load-balancing routers. If a load-balancing router randomly spreads traffic among k paths⁴, 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, where r is the total number of rounds performed. 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.

We can also observe a large number of dots close to the $y = x/2$ line. They correspond to addresses that are observed only during a finite part of the measurement, and have during that time a probability $p = 1/2$ of being observed, due to load balancing. If such an IP address is observed with a probability $1/2$ during k

rounds, its observation number will indeed be $x = k/2$, and its expected block number will be $y = k(1/2)^2 = x/2$.

Finally, a large number of IP addresses are close to the x -axis. This means that, whether they are observed in a large or small number of rounds, they are mainly observed during blocks of consecutive rounds, with few interruptions.

3 Causes of the observed dynamics

It is acknowledged that load-balancing routers play a significant role in the observed dynamics of routes with `traceroute`-like measurements [9, 2]. Previous work also suggests that routing dynamics play a key role in the continuous discovery of new IP addresses in our measurements [23]. This section identifies the strong role played by these factors in our observations.


These two factors play different roles. Suppose first that there is no load balancing. In this case, a measurement will discover routing changes as they occur, and the longer a measurement lasts, the more IP addresses it will observe (because more changes will occur). If on the contrary there are no routing changes but load balancing is used, then performing more measurement rounds will lead to observe more IP addresses, independently of the time elapsed between consecutive rounds⁵. The observed dynamics is a combination of these factors.

In order to study this rigorously, we use the `woolthorpe` data set and simulate slower measurements by considering only one out of every two rounds. Fig. 2 presents the number of distinct IP addresses observed with both these measurements, as a function of time elapsed since

³ We computed this plot for the longest uninterrupted part of the measurement, which represents 25322 rounds.

⁴ It has been shown [2] 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.

⁵ This is of course only true under certain conditions on the number of measurement rounds and the time elapsed between consecutive rounds.



(a) Number of discovered IP addresses as a function of time. (b) Number of discovered IP addresses as a function of the number of measurement rounds.

Fig. 2 How the frequency impacts the number of discovered IP addresses.

the beginning of the measurement, and the number of measurement rounds performed.

As expected, less IP addresses are observed over time with the slow measurements than with the faster ones. Fig. 2(a) shows that in a given time interval, performing more measurement rounds therefore allows to discover more IP addresses. This confirms that several measurement rounds are needed to discover all existing routes. This is caused by factors such as load balancing. Conversely, Fig. 2(b) shows that the slow measurements discover more IP addresses *at each round* than the faster ones. Therefore if more time elapses between two consecutive rounds, then each round discovers more IP addresses. This indicates that routes evolve with time.

In both cases, the gap between the plots for the slow and faster measurements are significant, which shows that both factors play an important role in the observed dynamics. This is why we propose a model that incorporates load balancing and route dynamics.

4 Model

Our purpose here is to propose relevant and simple mechanisms that reproduce the observations made in Section 2. 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 changes. 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.

We represent the topology by a random graph. In order to strengthen the conclusions drawn from our

study, we used two different models generating different topologies: the Erdős-Rényi model [12] which makes no hypothesis on the structure of the graph and is therefore the simplest model possible and the configuration model [3] in order to generate graphs with power-law degree distributions. The random graph model has two parameters: the numbers n of nodes and m of links. The configuration model has two parameters: the number n of nodes and the exponent γ of the power-law. As we will see in the next section, the comparison between results obtained with both generation processes gives insights on the impact of the topology on the observed dynamics.

Given a graph representing the topology, we assume that the route between the monitor and a destination is a shortest path, which can be obtained by performing a *breadth-first search (BFS)*. In order to simulate load balancing, each node chooses at random the next node on a shortest path to the destination, and we therefore implement a *random BFS*. It generates a shortest-path tree from the monitor to the destinations by considering the neighbors of explored nodes in a random order. These routing trees will therefore be different from one random BFS to the next, even if the underlying graph does not change.

Second, we need to model changes in the routing topology. We use a simple approach based on link rewiring, or *swap*. It consists in choosing uniformly at random two links (u, v) and (x, y) ⁶ 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 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:

⁶ We choose them such that the four nodes are distinct.




Fig. 3 Number of nodes observed since measurement beginning for various values of s (random graphs, $n = 500,000$, $m = 1,000,000$, $d = 3,000$).

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 process generates a series of r trees T_1, T_2, \dots, T_r , which simulate periodic `tracertree` measurements, on which we can conduct similar analysis as those we performed on real data.

5 Results

In this section we show that this model is relevant to explain the dynamic properties presented in Section 2. To that purpose, we perform several simulations varying the parameters of the model: the numbers n of nodes, m of links, d of destinations, and s of swaps per round. Our goals are to find (1) whether the simulations reproduce the observations and (2) how the different parameters impact the results and what are the relations between them.

5.1 First observations

Evolution of the number of distinct nodes. In order to answer the first question, we present Fig. 3 the evolution of the number of distinct nodes observed over time for Erdős-Rényi random graphs with $n = 500,000$, $m = 1,000,000$, $d = 3,000$ and various values of the number s of swaps. It shows a similar behavior to the one we observed in real data (see Fig. 1(a)). In particular all the curves present clearly a fast initial growth⁷ and then a more or less linear progression. More-

⁷ this phase lasts more than 1 round, although it is difficult to visualize it on the plot.

over, the slopes of the curves increase with the number of swaps. This is due to the fact that with a higher number of swaps, the paths to the destinations change more quickly and thus more nodes are discovered at each step.

This figure also show that when the underlying graph does not evolve ($s = 0$), there is only an initial growth in which all shortest paths are explored. Once all nodes on these paths have been discovered, the curve becomes flat. This confirms that the regular discovery of new IP addresses in real data may stem from route dynamics.

Observation number vs. block number. We also present in Fig. 4 the correlations between the observation number and block number for the same simulations.

For $s = 10$ (Fig. 4(b)), the main invariants we 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 Section 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. We also observe a high density of nodes on a line with equation $y = (r - x)/2$, r being the total number of rounds performed. This line has a similar explanation: it corresponds to nodes which are observed with probability $p = 1/2$ for a given duration, and are observed *in all rounds* before and after that. Although this line is not present in Fig. 1(b), it sometimes can be observed in other datasets, although not as clearly as here.

When no route dynamics is simulated ($s = 0$, Fig. 4(a)), only the parabola is present, thus confirming that this phenomenon observed in real data is due to load balancing mechanisms which are well captured by the random BFS model. At the opposite, when the number of swaps is too high (Fig. 4(c) and 4(d)), route dynamics get the better of load balancing phenomena and the parabola tends to vanish.

Assessing the results. The simulations presented above indicate that the model succeeds in reproducing the two main characteristics identified in the data. In order to assess more formally this statement, we performed extensive analyses that we present below.

The first question that arises is whether the linear progression is truly linear or just seems so, both for real data and simulations. In order to answer that question, we performed several tests to check that this is indeed the case. We only describe the following one, due to space limitations: given a plot that goes from round 1 to round r , we (1) manually identify a value r_0 such




Fig. 4 Observation number vs. block number for various values of s (random graphs, $n = 500,000$, $m = 1,000,000$, $d = 3,000$).




Fig. 5 Impact of the number of swaps on the slope (random graphs, $n = 100,000$, $d = 300$).




Fig. 6 Relation between links and swaps (random graphs, $n = 100,000$, $d = 300$).

that the initial fast increase is over after r_0 rounds⁸; (2) consider several intervals of different lengths over the remaining $r - r_0$ rounds; (3) perform linear fits with the *least square method* on all these intervals; and (4) finally check that all the obtained values for the slopes are consistent.

The second issue comes from the fact that results naturally vary from one simulation to another: the slope may vary and the plot may present some sharp increases at some points, as can be seen in Fig. 3. In order to capture the notion of *typical* slope of the plot, we therefore run a large number of simulations with a given set of parameters and consider the plot of the average result over all simulations. We then perform a linear fit over this plot which gives the typical slope α .

We applied this methodology to explore the simulations over Erdős-Rényi graphs, which we call *random graphs* (Section 5.2) and we then compared the obtained results with simulations performed on random graphs with a power-law degree distributions, which we call *power-law graphs* (Section 5.3).

5.2 Random graphs

Node discovery. To confirm the results of Section 5.1 we applied the methodology presented above on simulations for random graphs with $n = 100,000$ and $d = 300$ for various numbers of links and swaps. It makes it

possible to study more precisely the impact of the dynamics (number of swaps s) over the node discovery (slope α). Results are presented Fig. 5. We observe that the slope increases almost linearly with the number of swaps. This indicates a strong correlation between the observation of new nodes and the underlying dynamics, which confirms the intuition given by Fig. 3. With a higher number of swaps, the topology changes more frequently and, consequently, more paths are affected at each round.

Besides, the plot also shows that the relation between the two quantities is affected by the total number of links in the graph. Intuitively, the swaps are less likely to impact the paths from the monitor to the destinations if the graph is more dense. This is confirmed on the plot: for a given number of swaps, the slope is higher for graphs with 200,000 links than for graphs with 800,000 links.

Swaps vs. links. We studied more deeply the relation between the number of swaps s and the number of links m by making the two parameters vary at the same time. We set $n = 100,000$ and $d = 300$ and made simulations with several values of s and m . The results are presented on Fig. 6. We first observe that, for a given number of swaps, the larger the number of links, the smaller the slope of the corresponding curve, which confirms that when the number of links increases, a smaller fraction of them is affected by the swaps. Notice however that, for different simulations with a same ratio s/m , the corresponding slopes are not equal. For instance, we can observe that for $s/m = 10^{-5}$ (the two curves marked

⁸ We are not interested here in the smallest such value, making finding a relevant value very easy.




Fig. 7 Impact of the number of destinations (random graphs, $n = 100,000$, $m = 800,000$, $s = 2$).

with triangles), the slopes are equal to 0.46 and 0.38, respectively. We will study more in depth the relationship between m and α in Section 5.3.

Number of destinations. Finally, we also studied the impact of the number of destinations d for graphs with $n = 100,000$, $m = 800,000$ and $s = 2$ (Fig. 7). Intuitively, increasing the number of destinations causes the number of nodes on the shortest paths to the destinations to increase. Indeed, we observe that the initial growth phase, which corresponds to the discovery of all nodes on all shortest paths to the destinations, reaches a higher value when the number of destinations increases. As before, this phase is followed by a linear progression. Notice that increasing the number of destinations also increases the slope. This is clearly due to the fact that since the size of routing trees from the monitor to the destinations increases with the number of destinations, the probability for a swap to affect such a routing tree increases likewise.

5.3 Impact of the underlying structure

In order to study the impact of the underlying topology's structure, we compared the behaviors observed above for random graphs to those obtained for power-law graphs.

The observations made for power-law graphs are the same, qualitatively, as the ones made for random graphs. The number of nodes observed since measurement beginning displays a linear progression after a fast initial growth, and the correlations between block number and observation number have the same characteristics as the ones for random graphs, depending on the number of swaps.

We do however observe a *quantitative* difference: new nodes are observed at a slower rate for power-law graphs than for random ones. This can be observed in Figure 8, which presents the number of nodes observed since measurement beginning for a power-law graph




Fig. 8 Comparison between random and power-law graphs ($n = 100,000$, $d = 300$, $swaps = 50$).

with exponent 2.3, which corresponds to approximately 200,000 links, and two random graphs. The slope of the curve for the power-law graph is indeed much smaller than the one for a random graph with the same number of nodes and links.

The average distance may play a role in this. It has indeed been proven that the average distance is smaller for power-law graphs (for which it is in the order of $\log \log n$ [8]), than for random graphs (for which it is in the order of $\log n$ [4]). This implies that shortest path trees from the monitor to the destinations will have fewer nodes in power-law graphs than in random graphs, naturally inducing the observation of fewer new nodes. However, though this certainly plays a role, this is not enough to explain the observed difference. Figure 8 shows that the slope of the curve for the power-law graph is also smaller than the one for a random graph with the same average distance (and hence with $m = 1,200,000$ links).

As we can observe from the three graphs used above as an example, the structural differences between random and power-law graphs lead to important differences in the average distance and/or in the number of links. As observations made in the previous section suggest, the slope of the curve for the number of nodes should intuitively be proportional to the probability that a given swap will change the shortest path tree from the monitor to the destinations. Let us call spt_l the *typical*⁹ number of *links* in a shortest path tree from the monitor to the destinations.

Table 1 compares the value of spt_l for a power-law graph and two random graphs: one with the same average distance and one with the same number of links. We observe that it highlights additional structural differences between these two types of graphs: though the

⁹ In the same way that we performed several simulations and averaged the results in order to obtain the typical behavior for given model parameters, we generate several graphs with the same size and number of destinations to compute the typical size of a shortest path tree.

	Avg. dist.	spt_l
power-law, $m = 2 \cdot 10^5$	4.0	545
random, $m = 12 \cdot 10^5$	3.9	792
random, $m = 2 \cdot 10^5$	8.4	1420

Table 1 Average distance and size of shortest path trees for different graphs. $n = 100,000$, $d = 300$.

the average distance is approximately equal in a power-law graph with exponent 2.3 and a random graph with $12 \cdot 10^5$ links, the numbers of links in shortest path trees from the monitor to the destinations are significantly different.




Fig. 9 Slope of the curve of the observed number of links vs spt_l/m for different graphs and different destination numbers. $s = 2$. For all graphs we used several values of d : $d = 100, 200, 300, 400, 500$

Finally, Fig. 9 plots the slope of the curve of the observed number of *links*¹⁰ since measurement beginning vs spt_l/m for different types of graphs. For each type of graph, we perform simulations using different numbers of destinations, which induces different values of spt_l .

We observe a strong correlation between these two quantities, meaning that spt_l/m plays a key role in the observed behavior, for all types of graph. Though this does not fully allow us to understand the model’s behavior, this shows that the number of links and the size of a shortest path routing tree are key parameters for understanding the quantitative difference between random and power-law graphs.

As a conclusion, by exploring the impact of the parameters, we showed that a wide range of their values are relevant. In particular, they all produce a behavior qualitatively similar to what we observed in real data. The fact that the observed behavior for random and power-law graphs is qualitatively the same also strengthens the relevance of the model.

¹⁰ Since swaps affect links, the probability that a shortest path tree is affected by a swap depends on its size in terms of links, and it is consistent to study the number of *links* observed since measurement beginning, which has the same behavior as the number of observed nodes.

6 Related work

Study of the dynamics of the Internet topology has been tackled both by analyzing the dynamics of individual routes [29,31,18,17,36] and from a more global perspective, mainly at the AS- or IP-level [14,7,21,25,27,26,15,32,11,10]. Load balancing has also been acknowledged for playing an important role in the dynamics of routes as measured with traceroute-like tools [2]. Cunha *et al.* [9] used a method for measuring load-balanced routes, i.e. routes containing one or more load-balancing routers, and study their dynamics.

Some works [16,27] argue that the topology dynamics should be taken into account in order to produce realistic models for the Internet topology. Work on modeling this topology and its dynamics can be roughly divided between approaches aiming at realistically mimicking the evolution mechanisms of the topology, e.g. reproducing the criteria taken into account by ASes for creating peering or customer-provider links, see for instance [13,6,35,28,22], and approaches aiming at reproducing global network characteristics through simple mechanisms, thus exhibiting simple causes for more complex observations. This paper belongs to this second approach. Tangmunarunkit *et al.* [33] showed that this approach is relevant by establishing that network generators based on local properties, such as the degree distributions of nodes, can capture global properties of the topology, such as its hierarchical structure. Most related to our characterization and modeling of the evolution of the Internet topology is the work by Oliveira *et al.* [25], which analyzes the AS topology and shows that real topology dynamics can be modeled as constant-rate births and deaths of links and nodes. In a similar spirit, Valler *et al.* model BGP routing churn by a process similar to an epidemic spreading on a network [34]. Park *et al.* [28] studied several growing models for the Internet topology, i.e. models in which nodes and links are progressively added over time. They compared the evolution of these induced networks with the evolution of the real topology, and use this to distinguish between the quality of the different models.

Whereas most existing works focus on the long-term evolution (e.g. from the Internet birth to current times) of the *physical* AS topology, we are concerned here with the short- to medium-term evolution of the *routing* topology at the IP-level. The routing and physical topology are closely linked but not identical objects. In particular, routing changes can occur in the absence of physical changes. They are closely linked to BGP dynamics which have been studied for instance in [18,21,34,36]. Finally, our model does not take into account node ap-

pearance and disappearance, which would be necessary for modeling the long-term topology evolution.

Finally, note that this paper is an extended version of an earlier work [24] in which we performed a preliminary analysis of the different behaviors observed for simulations with random and power-law graphs. This version gives evidence that load-balancing and routing dynamics play a key role in the observed properties of the IP-level routing topology dynamics. It presents a detailed and rigorous analysis of the impact of the model parameters, as well as an explanation of the differences observed between random and power-law graphs.

7 Conclusion

In this work we conducted periodic measurements of ego-centered views of the Internet topology and studied their dynamics. We isolated invariant characteristics of these dynamics, and identified load balancing and evolution of the routing topology as key factors in the observed properties.

Based on this observation, we proposed a model for the dynamics of the topology, which integrates both load balancing and routing changes. Simulations show that this model captures the main characteristics of the dynamics of the ego-centered views. We performed an analysis of the underlying topology structure by comparing random and power-law graphs. We showed that there is a quantitative difference in the corresponding behaviors, which depends mostly on the difference in the relative size of a shortest path tree with respect to the total number of links.

Our model is based on simple mechanisms, both for the topology generation and the characteristics of route dynamics. It is therefore not suitable for generating realistic time-evolving topologies. However, the fact that it captures the main characteristics of the observed ego-centered views shows that the factors it mimics play a strong role in the Internet routing topology dynamics, which offers key insight on the understanding of these dynamics. We therefore consider this model as a key step towards the realistic modeling of the Internet topology dynamics, as well as towards its understanding.

Future work lies in two main directions. First, we strongly believe that this model can be used to estimate some properties of the actual IP-level routing topology that are not directly available through measurements. For instance, performing more extensive studies of the relations between the model's parameters and the observed behavior, such as the one presented in Fig-

ure 5, would allow to infer the parameters from the observed behavior. Applying this knowledge to real-world data would allow to estimate the real-world values corresponding to these parameters, such as for instance the frequency of link changes in the whole topology. Moreover, since our model is based on random graphs and simple mechanics for load balancing and routing dynamics, it lends itself well to formal analysis. This would allow to obtain formal proofs for such results.

Second, the field of Internet topology modeling is very active, and models far more realistic than random graphs are available. One should explore the combination of our routing mechanisms principles with these topology models, to investigate the role played by the topology structure on the observed dynamics. In particular, our model does not take into account the long term topology evolution, since it does not model node birth or death. Coupling the ingredients of our routing dynamics with, e.g., a growing model for the Internet topology which would reflect its long term dynamics would surely lead to insightful results.

Acknowledgements:

This work was 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. Finally, we thank Vincent Cohen-Addad, Louis Fournier, Antoine Javelot and Sura-bhi Sankhla for their fruitful collaboration, and Matthieu Latapy for enlightening comments.

References

1. Achlioptas, D., Clauset, A., Kempe, D., Moore, C.: On the bias of traceroute sampling. *Journal of the ACM* **56**(4) (2009)
2. Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., Teixeira, R.: Traceroute Anomalies: Detection and Prevention in Internet Graphs. *Computer Networks* **52**, 998–1018 (2008)
3. Bender, E.A., Canfield, E.R.: The asymptotic number of labeled graphs with given degree sequences. *Journal of Combinatorial Theory (A)* **24**, 357–367 (1978)
4. Bollobás, B.: *Random Graphs*. Academic Press (1985)
5. CAIDA: MIDAR antialiasing tool. <http://www.caida.org/tools/measurement/midar/>
6. Chang, H., Jamin, S., Willinger, W.: To peer or not to peer: modeling the evolution of the internet's AS-level topology. In: *Proceedings of IEEE INFOCOM* (2006)
7. Chen, Q., Chang, H., Govindan, R., Jamin, S., Shenker, S., Willinger, W.: The Origin of Power-Laws in Internet Topologies Revisited. In: *IEEE Infocom* (2002)

8. Chung, F., Lu, L.: The average distances in random graphs with given expected degrees. *Proceedings of the National Academy of Sciences* **99**(25), 15,879–15,882 (2002). DOI 10.1073/pnas.252631999. URL <http://dx.doi.org/10.1073/pnas.252631999>
9. Cunha, I., Teixeira, R., Diot, C.: Measuring and Characterizing End-to-End Route Dynamics in the Presence of Load Balancing. In: *Proceedings of Passive and Active Measurement Conference* (2011)
10. Dhamdhere, A., Cherukuru, H., Dovrolis, C., Claffy, K.: Measuring the evolution of internet peering agreements. In: *Proceedings of IFIP Networking* (2012)
11. Dhamdhere, A., Dovrolis, C.: Twelve years in the evolution of the internet ecosystem. *IEEE/ACM Transactions on Networking* **19**(5), 1420–1433 (2011)
12. Erdős, P., Rényi, A.: On random graphs. *Publicationes Mathematicae Debrecen* **6**, 290 (1959)
13. Fabrikant, A., Koutsoupias, E., Papadimitriou, C.H.: Heuristically optimized trade-offs: A new paradigm for power laws in the internet. In: *Proceedings of ICALP* (2002)
14. Govindan, R., Reddy, A.: An Analysis of Internet Inter-Domain Topology and Route Stability. In: *Proceedings of IEEE INFOCOM* (1997)
15. Haddadi, H., Fay, D., Uhlig, S., Moore, A., Mortier, R., Jamakovic, A.: Mixing biases: Structural changes in the AS topology evolution. In: *Proceedings of the second international workshop on Traffic and Measurements Analysis (COST-TMA 2010)* (2010)
16. Haddadi, H., Uhlig, S., Andrew Moore, A., Mortier, R., Rio, M.: Modeling internet topology dynamics. *ACM SIGCOMM Computer Communication Review* **38**(2), 65 (2008). DOI 10.1145/1355734.1355745
17. Kuate, A.M., Teixeira, R., Meulle, M.: Characterizing network events and their impact on routing. In: *CoNEXT '07: Proceedings of the 2007 ACM CoNEXT conference*, pp. 1–2 (2007)
18. Labovitz, C., Malan, G., Jahanian, F.: Origins of Internet routing instability. In: *Proceedings of IEEE INFOCOM*, pp. 218–226 vol.1 (1999)
19. Lakhina, A., Byers, J., Crovella, M., Xie, P.: Sampling biases in IP topology measurements. In: *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications*. IEEE Societies, vol. 1, pp. 332 – 341 vol.1 (2003). DOI 10.1109/INFCOM.2003.1208685
20. Latapy, M., Magnien, C., Ouédraogo, F.: A radar for the internet. *Complex Systems* **20**(1), 23–30 (2011)
21. Li, J., Guidero, M., Wu, Z., Purpus, E., Ehrenkrantz, T.: BGP routing dynamics revisited. *SIGCOMM Comput. Commun. Rev.* **37**(2), 5–16 (2007)
22. Lodhi, A., Dhamdhere, A., Dovrolis, C.: Genesis: An agent-based model of interdomain network formation, traffic flow and economics. In: *Proceedings of the IEEE INFOCOM conference* (2012)
23. Magnien, C., Ouédraogo, F., Valadon, G., Latapy, M.: Fast dynamics in internet topology: Observations and first explanations. In: *Proceedings of the Fourth IEEE International Conference on Internet Monitoring and Protection*, pp. 137–142 (2009)
24. Medem, A., Magnien, C., Tarissan, F.: Impact of power-law topology on IP-level routing dynamics: simulation results. In: *Proceedings of the Fourth International Workshop on Network Science for Communication Networks (NetSciCom)* (2012)
25. Oliveira, R., Zhang, B., Zhang, L.: Observing the Evolution of Internet AS Topology. In: *Proceedings of ACM SIGCOMM* (2007)
26. Pansiot, J.J.: Local and Dynamic Analysis of Internet Multicast Router Topology. *Annales des télécommunications* **62**, 408–425 (2007)
27. Park, S.T., Khrabrov, A., Pennock, D., Lawrence, S., Giles, C., Ungar, L.: Static and dynamic analysis of the Internet’s susceptibility to faults and attacks. In: *Proceedings of IEEE INFOCOM*, pp. 2144–2154 (2003)
28. Park, S.T., Pennock, D.M., Giles, C.L.: Comparing static and dynamic measurements and models of the Internet’s AS topology. In: *Proceedings of IEEE Infocom* (2004)
29. Paxson, V.: End-to-end routing behavior in the Internet. *IEEE/ACM Transactions on Networking* **5**(5), 601–615 (1997). DOI 10.1109/90.649563
30. A Radar for the Internet – Publicly available datasets. <http://data.complexnetworks.fr/Radar/>
31. Schwartz, Y., Shavitt, Y., Weinsberg, U.: On the Diversity, Stability and Symmetry of End-to-End Internet Routes. In: *Proceedings of Global Internet* (2010)
32. Shavitt, Y., Weinsberg, U.: Topological trends of internet content providers. In: *Proceedings of the Fourth Annual Workshop on Simplifying Complex Networks for Practitioners (SIMPLEX)* (2012)
33. Tangmunarunkit, H., Govindan, R., Jamin, S., Shenker, S., Willinger, W.: Network topology generators: degree-based vs. structural. In: *Proceedings of ACM SIGCOMM* (2002)
34. Valler, N., Butkiewicz, M., Prakash, B.A., Faloutsos, M., Faloutsos, C.: Non-binary information propagation: Modeling BGP routing churn. In: *IEEE INFOCOM Workshop on Network Science for Communication Networks (NetSciCom)* (2011)
35. Wang, X., Loguinov, D.: Wealth-based evolution model for the internet AS-level topology. In: *Proceedings of IEEE INFOCOM* (2006)
36. Watari, M., Tachibana, A., Ano, S.: Inferring the origin of routing changes based on preferred path changes. In: *Passive and Active Measurement Conference* (2011)
37. Xia, J., Gao, L., Fei, T.: Flooding attacks by exploiting persistent forwarding loops. In: *Proceedings of the ACM SIGCOMM Internet Measurement Conference* (2005)