Fast dynamics in Internet topology: observations and first explanations

Clémence Magnien[1,2], Frédéric Ouédraogo[1,2,3], Guillaume Valadon[1,2], Matthieu Latapy[1,2]
1: UPMC Univ Paris 06, UMR 7606, LIP6, F-75016, Paris, France
2: CNRS, UMR 7606, LIP6, F-75016, Paris, France
3: University of Ouagadougou, LTIC, Ouagadougou, Burkina Faso

E-mail: First-name.Last(-)name@lip6.fr

Abstract

By focusing on what can be observed by running `traceroute`*-like measurements at a high frequency from a single monitor to a fixed destination set, we show that the observed view of the topology is constantly evolving at a pace much higher than expected. Repeated measurements discover new IP addresses at a constant rate, for long periods of times (up to several months).*

In order to provide explanations, we study this phenomenon both at the IP, and at the Autonomous System levels. We show that this renewal of IP addresses is partially caused by a BGP routing dynamics, altering paths between existing ASes. Furthermore, we conjecture that an intra AS routing dynamics is another cause of this phenomenon.

1 Introduction

Most works aimed at mapping the Internet IP-level topology rely on `traceroute`-like probes, for instance [2, 19]. These probes are repeated periodically for large amounts of time, each round of measurement leading to a partial and biased view of the topology. It is indeed known that, because of phenomena such as load balancing [24], it is not possible to see everything that can be seen from a monitor in a single round. One round discovers only one path among several between the monitor and a destination. Snapshots of the Internet topology are therefore constructed by merging series of measurement rounds. This relies on the assumption that it is possible to explore a given part of the topology with a finite number of probes.

We focus here on what can be observed by running `traceroute`-like probes at a high frequency from a single monitor to a constant destination set [11]. We show that the observed view of the topology is constantly evolving at a rate much higher than expected. For instance, during the last week of two-months measurements, we discovered 1 118 new IP addresses (on a total of 29 100) that had never been observed before.

These observations imply in particular that it is never possible to discover everything that can be seen from a monitor; also, aggregating data from such measurements leads to topology maps with much obsolete information.

In this paper we describe and study this phenomenon. Though we do not obtain a conclusive explanation, we show that a fast routing dynamics is the cause.

2 Data set

We use the data described in [11]. Measurements were conducted from more than 150 monitors (mainly from PlanetLab) scattered around the world. Each monitor had a *destination set* that stayed the same for the whole duration of the measurements. The measurements then consisted in periodically running the `tracetree` tool, which collects a routing tree from a given monitor to a set of destinations in a `traceroute`-like manner. The measurements were conducted with a high frequency (typically about 100 measurement rounds per day), for a long period of time (from weeks to several months, depending on the monitor). For more details, see [11].

Our goal is not to study the data in detail or compare all the data sets obtained from all monitors. On the contrary, we insist on the fact that we observed *similar* phenomena for each of them: while the exact details do of course depend on the particular monitor under study, our observations were *qualitatively* the same in all cases.

In this paper, we have therefore chosen to illustrate our results by using a single monitor and the corresponding data set. It consists of a single two-months measurement (June and July 2007) from a monitor located in Japan, at a rate of approximately 100 rounds per day, leading to 5 891 rounds in total. The destination set consists of 3 000 IP addresses

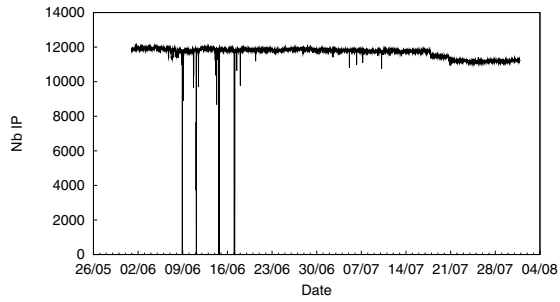

Figure 1. Number of IP addresses observed in each measurement round.

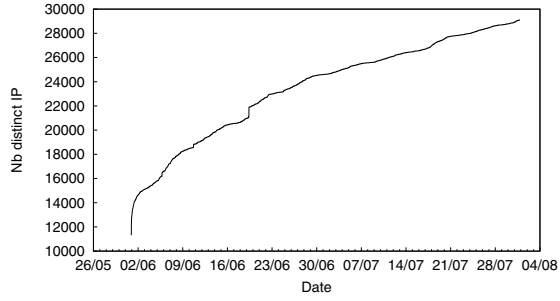

Figure 2. Number of IP addresses observed since the beginning of the measurements as a function of time.

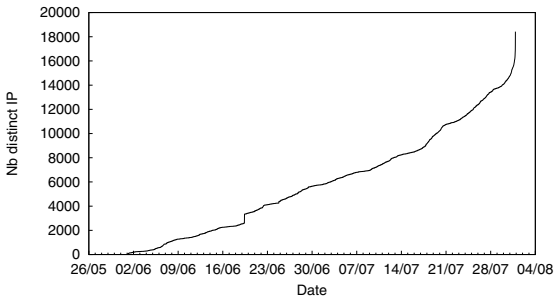

Figure 3. Number of IP addresses that were observed before time t and are never observed after t as a function of t.

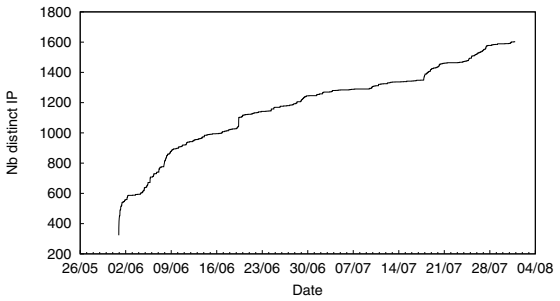

Figure 4. Number of distinct IP addresses seen with stable destinations only.

chosen randomly that replied to `ICMP echo request` messages at selection time. We kept all the observed IP addresses (and did not remove the first hops needed to get out of the monitor's network).

3 IP addresses renewal

In this section, we describe the evolution of the set of observed IP addresses. Figure 1 presents the number of IP addresses observed in each measurement round. All values in this plot are centered around a same value (close to 12 000) except some downward peaks which indicate rounds with *less* IP addresses than usual. These peaks could indicate a loss of connectivity at or near the monitor, or an event such as a major routing change or failure. Studying this is however out of the scope of this paper.

The next question is whether we observe the same IP addresses in all rounds. This leads to the plot in Figure 2, which presents the number of distinct IP addresses observed since the beginning of the measurements as a function of time. This plot gives evidence for a striking fact: measurements continuously discover new IP addresses never seen before [1].

[1] Other, six-months long, measurements exhibit the same behavior.

Though it seems natural to observe *some* new IP addresses after some measurement time, this happens here with a surprisingly high rate: during the second month of the measurements, around 150 new IP addresses are discovered *each day*.

We have seen (Figure 1) that the number of IP addresses seen at each round is not increasing. The continuous discovery of new IP addresses must therefore come together with a continuous disappearance of addresses that we cease to observe after some time. Figure 3 presents this. The disappearances are indeed symmetric with the observation of new IP addresses [2].

One possible cause for these observations would be that some routers reply with random IP addresses; we will show in the next section that it is not the case. Another possible cause would be that some of our destinations are dynamic addresses, *i.e.* dynamically allocated to different hosts over time. Since such hosts could be in different locations, depending on network operation, these dynamic addresses could lead us to discover new paths and as a result new addresses in the measurements.

Figure 4 shows that this is not the case. The idea is to select the destinations that were *stable during the mea-*

[2] The plots of Figures 2 and 3 have a similar shape if rotated at a 180° angle.

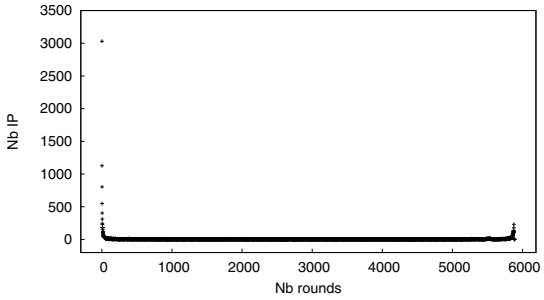

Figure 5. Distribution of the number of rounds in which each IP address was observed.

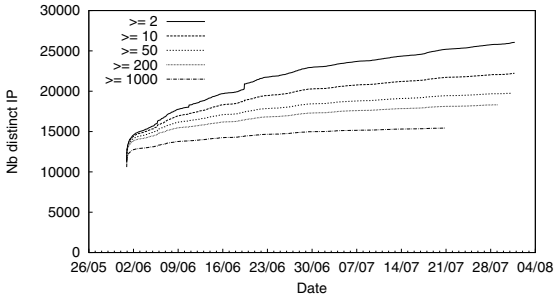

Figure 6. Number of IP addresses observed in at least 2, 10, 50, 200 or 1 000 different rounds (top to bottom) since the beginning of the measurements.

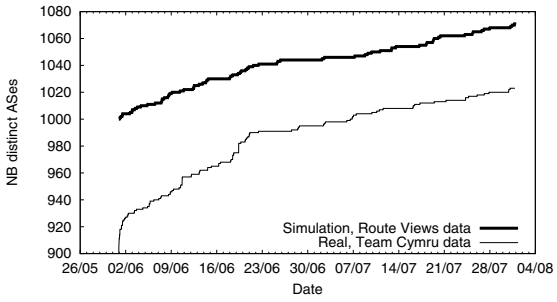

Figure 7. Number of distinct ASes observed since the beginning of the measurements: real number (thin line), and estimated with a simulation from Route Views data (thick line).

surements. Using a similar approach to geolocation studies (see for instance [14]), we considered that a destination address is not dynamic if the address immediately before it in the measurements is always the same; 35 out of 3 000 destinations satisfied this condition. We then simulated the measurements by keeping only these stable addresses, and we still clearly see a constant appearance of new IP addresses: dynamic addresses are therefore not the cause of this renewal. Note that our criterion for characterizing stable addresses is very restrictive; we do not imply that the addresses that do not satisfy it are dynamic. We tested other criteria, which provided the same results.

In summary, we observe a continuous, high-rate renewal of the set of IP addresses observed from a monitor, and showed that it is not a measurement artifact, but an actual property of the IP-level topology. This implies that repeating measurements, even for long periods of time, cannot converge to a full view of what can be observed. Moreover, aggregating data obtained during consecutive rounds to construct a topology map is not satisfying, because this means grouping up-to-date data together with obsolete one.

4 Recurring IP addresses

In this section we ask whether we observe IP addresses with consistency, or if we only see them in a very small number of rounds. For any number of rounds x, Figure 5 presents the number of IP addresses that were observed in exactly x different rounds during the measurements.

This distribution shows that a large number of IP addresses are very volatile: 3 030 IP addresses are indeed observed only once during these two-month measurements. On the other hand, a significant number of IP addresses appear recurrently: they are seen in almost each round during the measurements.

The presence of a large number of highly volatile IP addresses naturally induces the question of whether these addresses are the cause of the renewal of the observed IP addresses. Figure 6 answers this question. It presents the number of distinct IP addresses observed since the beginning of the measurements, restricted to recurring addresses that were observed in at least 2, 10, 50, 200 or 1 000 different rounds. Though the slope of these plots are smaller than the one in Figure 2, we continuously observe new *recurring* addresses. As a matter of fact, if we only consider IP addresses observed in least 1 000 rounds (out of 5 891 rounds), we still observe a non-negligible renewal. This shows that the constant observation of new IP addresses is not caused by volatile addresses only, and that recurring addresses are also renewed. Moreover, this means that routers replying with random addresses are not the cause of this renewal: the corresponding addresses would only be observed a very small number of times, and we showed that such addresses are not the main cause of our observations.

5 Autonomous Systems

We now study the same question on a different scale: do we observe the same type of behavior when we consider ASes rather than IP addresses? We associate each IP ad-

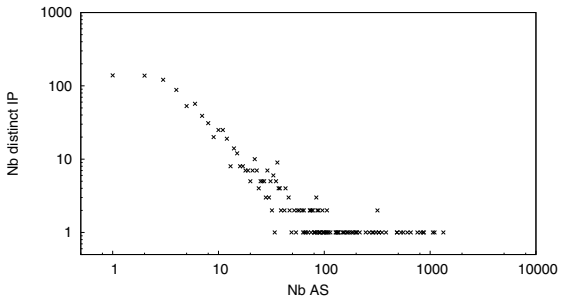

Figure 8. Distribution of the observed size of ASes.

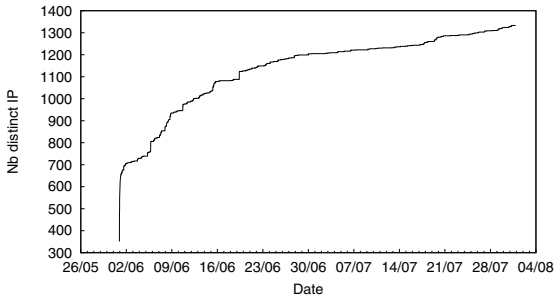

Figure 9. Number of IP addresses observed in the largest AS as a function of time.

dress seen in the measurements to its AS using the Team Cymru database [21]. Figure 7 (thin line) then presents the number of distinct ASes observed during the measurements. As expected, we see fewer ASes than IP addresses. However, we observe the same *behavior* as before: each round sees a more or less constant number of ASes (close to 950, not shown here), and we continuously discover new ASes during the measurements.

This can be considered as a partial explanation of what we observe at the IP level: if we discover new ASes, it is only natural that we should discover new IP addresses within them.

To study this question further, we asked if all the ASes are equivalent in the measurements. Figure 8 shows the distribution of the observed size of ASes: for each AS seen, we computed how many different IP addresses we observed within it, and plotted the corresponding distribution. As we can see, this distribution is highly heterogeneous: for more than 100 ASes (on a total of 1 023), we observe only a single IP address, while 3 ASes contain more than one thousand of observed IP addresses.

The presence of ASes with very large observed sizes in the measurements naturally led to the question of whether we also observe a renewal of observed IP addresses *within* a single AS. Figure 9 shows the number of IP addresses observed since the beginning of the measurements in the largest observed AS [3]. Again in this case we observe the same type of behavior: we continuously observe new IP addresses in this AS.

From these observations we can derive the following conclusion: we observe a constant renewal, both at the AS level, with the constant discovery of new ASes, and within single ASes, discovering new IP addresses in already seen ASes. This therefore allows us to break down the constant appearance of new IP addresses between these two factors.

However, this does not explain *why* we should observe new ASes, or new IP addresses within previously observed ASes. In particular, the question that naturally arises about the newly discovered ASes is whether they are *new*, *i.e.* created after the beginning of the measurements [4], or if they are pre-existing ASes that become visible to the measurements due to BGP routing dynamics.

To study this question further, we used data from the Route Views project [23]. This project makes publicly available a recording of BGP routing tables from several hosts. This data allowed us to simulate the measurements from an AS/BGP point of view: we chose a Route Views monitor located close to our monitor [5], then selected the routing tables corresponding to the period of the measurements. For each routing table, we extracted the ASes belonging to BGP paths corresponding to IP prefixes of the destinations. We thus obtained the set of all possible *observable* ASes for each routing table.

Figure 7 (thick line) then presents the number of distinct observable ASes since the beginning of the measurements, obtained through our simulations. We obtain a similar slope with both methods, which confirms their validity. Note that the numbers obtained with the Route Views data are larger than the ones obtained from Team Cymru. This is due to the fact that the Route Views data represents, at each moment, the set of *all* possible AS paths allowing to reach the destinations; instead, the Team Cymru data is directly extracted from the measurements, and therefore provides only a single path to each destination.

The use of the Route Views data moreover allows us to go further. We observed 1 072 ASes in total, 72 of which were discovered after the beginning of the measurements. Out of these ASes, we found out that 70, *i.e.* all but two of them, were present in the *first* routing table (but did not belong to AS paths leading to the destinations). This means that these 70 ASes were already existing at the beginning of the measurements, but became visible because of BGP routing changes.

Finally, we are able to conclude that, at the AS level, our

[3] This is the Level3 Communications AS (number 3 356), containing 1 333 IP addresses (on a total of 29 100).

[4] During the year 2007, around 250 ASes were created every month, see `http://www.cidr-report.org/as2.0/`

[5] This is host `route-views.wide.routeviews.org`, located in one peering point of the AS where our monitor is located.

observations are caused by a dynamics of the BGP routing, causing pre-existing ASes to become visible on the paths between the monitor and the destinations.

6 Related work

Much work has focused on the measurement bias created by mapping the Internet topology with `traceroute`-like probes. The majority of these works concern the fact that running probes from a limited number of monitors misses some links and/or creates a bias on the observed degrees of the nodes, see for instance [9, 1, 3, 20, 5]. Others have studied the fact that tools such as `traceroute` may report incomplete and/or false information, see for instance [24, 12, 18, 25, 6, 22, 20].

It is an acknowledged problem in the field that the Internet topology evolves with time and that this may create a bias in the measurements. However, though some works have studied the dynamics of the topology, at the IP or AS level (see for instance [3, 18, 25, 15, 13, 17, 16, 4, 7, 8]), up to our knowledge only one paper has attempted to study the bias caused by this dynamics on the measurements [13]. The authors of this paper study the AS-level topology, and design methods for evaluating with a certain degree of confidence if an *observed* topology change is a *real* change or not. Though their approach and some of their observations are similar to ours, they study the AS-level topology whereas we study the IP-level topology, and consider time-scales much longer, and hence a much coarser time resolution, than we do. The phenomena playing a role in their observations are therefore different than in our case: they decompose their observations into a birth/death process, coupled with transient routing dynamics.

Finally, another work [10] studied the measurement process of different complex networks, including Internet maps. They observed that, for the skitter data [2], measurements continuously discover new IP addresses. This is similar to our observations, though other causes probably play a role in this: the skitter data is collected from several monitors, at a lower frequency and for larger time scales than the data we study here.

7 Conclusion and perspectives

In this paper, we bring to light a surprising phenomenon: when performing periodic `traceroute`-like measurements from a single monitor to a fixed set of destinations, the obtained view of the topology never stabilizes. On the contrary, we continuously observe new IP addresses at a rate much higher than expected. This phenomenon is observed with various monitors and destination sets, and seems to be universal.

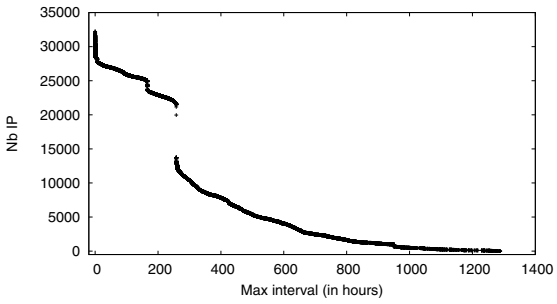

Figure 10. Complementary cumulative distribution of the time elapsed between the first and last discovery of IP addresses over all monitors.

We described this phenomenon in details and attempted to determine its cause. We first ruled out some possible explanations: dynamic IP addresses among the destinations, and routers answering with many random addresses. This showed that this phenomenon is not a measurement artifact, but an actual property of the IP-level topology.

We were able to break down the observation of new IP addresses into two factors: a constant observation of new ASes, coupled with an observation of new IP addresses within already discovered ASes.

From the record of routing tables by the Route Views project, we concluded that the discovery of new ASes is caused by the dynamics of the BGP routing. These ASes were in fact created before the beginning of the measurements, and became visible as a consequence of a change in routing paths.

Following these conclusions on AS renewal, we conjecture that the same cause holds for the IP address renewal, and that newly discovered addresses were already allocated at the beginning of the measurements, and became visible because of routing changes.

Some preliminary results on this question confirm this hypothesis. We combined measurements performed from several monitors in order to test if all monitors discover the new IP addresses at the same time. We chose 11 monitors that used the same destination set, and studied addresses observed with two monitors or more. For each such address, we wrote down the time it was discovered by each monitor individually and then computed the interval between the first and the last of these discoveries. For instance, an address seen with two monitors, first observed with the first monitor at 7 AM and then discovered by the second monitor at 11 AM the same day, will give an interval of four hours.

Figure 10 presents the complementary cumulative distri-

bution of these interval sizes. We observe that a large number of IP addresses discovered by a given monitor were in fact observed a significant duration before with other monitors. Among the 32 228 (out of 40 076) IP addresses seen with at least two monitors, 22 897 were observed with one monitor more than 200 hours before they were discovered by another one, which means that these addresses existed for a long time before they were discovered. Note that this does not tell us whether other addresses existed before their discovery or were created at this time. This indicates that a large number of the IP addresses discovered by a given monitor existed in fact for a significant time before their discovery, and that a routing dynamics between existing addresses plays a strong role in our observations.

This work should be pursued in several directions. First, we want to fully characterize and understand the renewal of observed IP addresses. We think that it is possible to perform new measurements, specifically designed for answering the question of whether newly discovered IP addresses existed prior to their discovery or not. Another direction would be to perform simulations, which would open the way to an accurate modeling of the phenomena causing the renewal of IP addresses.

Second, our work indicates that there is no perfect solution for constructing maps of the Internet topology while a single measurement round does not discover everything that can be seen from a monitor (because of load balancing for instance), aggregating data from several consecutive rounds puts together obsolete and up-to-date data. It would be of prime interest to determine if a best compromise exists, allowing to construct the most accurate maps, for instance by tuning the number of measurement rounds and their frequency.

References

[1] D. Achlioptas, A. Clauset, D. Kempe, and C. Moore. On the bias of traceroute sampling. In *ACM STOC*, 2005.

[2] CAIDA – Skitter project. `http://www.caida.org/tools/measurement/skitter/`.

[3] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power laws in internet topologies revisited. In *Proc. IEEE INFOCOM*, Jun. 2002.

[4] R. Govindan and A. Reddy. An analysis of internet inter-domain topology and route stability. In *Proc. IEEE INFOCOM*, 1997.

[5] J.-L. Guillaume and M. Latapy. Relevance of massively distributed explorations of the internet topology: Simulation results. In *Proc. IEEE infocom*, 2005.

[6] B. Huffaker, D. Plummer, D. Moore, and k. claffy. Topology discovery by active probing. In *Proc. Symposium on Applications and the Internet*, Jan. 2002.

[7] C. Labovitz, G. R. Malan, and F. Jahanian. Origins of internet routing instability. In *Proc. IEEE INFOCOM*, pages 218–226, 1999.

[8] M. Lad, D. Massey, and L. Zhang. Visualizing internet routing changes. *IEEE Transactions on Visualization and Computer Graphics, special issue on Visual Analytics*, 2006.

[9] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in ip topology measurements. In *Proc. IEEE INFOCOM*, 2003.

[10] M. Latapy and C. Magnien. Measuring fundamental properties of real-world complex networks. In *Proc. IEEE INFOCOM*, 2008.

[11] M. Latapy, C. Magnien, and F. Ouédraogo. A radar for the internet. In *Proc. first International Workshop on Analysis of Dynamic Networks (ADN), in conjunction with IEEE ICDM 2008*, 2008. To appear. Available at `http://arxiv.org/abs/0807.1603`.

[12] T. Moors. Streamlining traceroute by estimating path lengths. In *Proc. IEEE Workshop on IP Operations and Management*, Oct. 2004.

[13] R. Oliveira, B. Zhang, and L. Zhang. Observing the evolution of internet AS topology. In *ACM SIGCOMM*, 2007.

[14] V. N. Padmanabhan and L. Subramanian. An investigation of geographic mapping techniques for internet hosts. In *SIGCOMM*, 2001.

[15] J.-J. Pansiot. Local and dynamic analysis of internet multicast router topology. *Annales des télécommunications*, 62:408–425, 2007.

[16] S.-T. Park, A. Khrabrov, D. Pennock, S. Lawrence, C. L. Giles, and L. Ungar. Static and dynamic analysis of the internet's susceptibility to faults and attacks. In *Proc. IEEE Infocom*, 2003.

[17] 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. IEEE Infocom*, 2004.

[18] V. Paxson. End-to-end internet packet dynamics. *IEEE/ACM Trans. Networking*, 7(3):277–292, June 1999.

[19] Y. Shavitt and E. Shir. DIMES: Let the internet measure itself. *ACM SIGCOMM Computer Communication Review*, 35(5), 2005.

[20] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proc. ACM SIGCOMM*, Aug. 2002.

[21] Team Cymru. `http://www.team-cymru.org/`.

[22] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker. In search of path diversity in ISP networks. In *Proc. ACM SIGCOMM Internet Measurement Conference*, Aug. 2003.

[23] University of Oregon. Route Views, University of Oregon Route Views project. `http://www.antc.uoregon.edu/route-views/`.

[24] F. Viger, B. Augustin, X. Cuvellier, B. Orgogozo, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Detection and prevention in internet graphs. *Computer Networks*, 52, 2008.

[25] F. Wang, N. Feamster, and L. Gao. Quantifying the effects of routing dynamics on end-to-end internet path failures. *Technical report (TR-05-CSE-03) in Department of Electrical and Computer Engineering in the University of Massachusetts at Amherst*, 2006.