# A Radar for the Internet

**Matthieu Latapy**[1]
**Clémence Magnien**[1]
**Frédéric Ouédraogo**[1,2]

[1]*LIP6 – CNRS and UPMC (University Pierre and Marie Curie – Paris 6)*
*104 avenue du Président Kennedy*
*75016 Paris, France*
*Matthieu.Latapy@lip6.fr*
*Clemence.Magnien@lip6.fr*

[2]*University of Ouagadougou*
*03 B.P. 7021 Ouagadougou 03, Burkina-Faso*
*ouedraogo.frdric@yahoo.fr*

Mapping the topology of the internet is a challenge in itself, and studying its dynamics is even more difficult. However, achieving this would provide key information on the nature of the internet, crucial for modeling and simulation. Moreover, detecting anomalies in the dynamics is a key issue for security. A new measurement approach is introduced here that makes it possible to capture the dynamics of the internet at a scale of a few minutes in a radar-like manner. By conducting and analyzing large-scale measurements of this kind, we rigorously and automatically detect events in the observed dynamics, which is totally out of reach using previous approaches.

## 1. Introduction

Since the end of the 1990s, mapping the internet as a large set of nodes and links received much attention. However, due to its distributed nature and sheer size, accurately measuring this topology is extremely difficult. The main method to do so relies on the classical traceroute tool [1], which gives a path from a machine connected to the internet (called *monitor*) to any other (called *destination*). Such paths are composed of the IP addresses of routers and links between them. A (partial) map of the internet can be obtained by running traceroute from many monitors to many destinations, and merging the obtained paths (see Figure 1). For various reasons, however, this is far from trivial and the obtained maps are not satisfactory [2–4]. Therefore, much effort is being devoted to improving available maps, in particular by increasing the number of monitors [5, 6] and by designing more accurate measurement tools [7].

In this situation, it must be clear that studying the dynamics of the internet's topology, and in particular detecting events in the dynam-

ics, is totally out of reach using current approaches. Even the study of global trends in the evolution of the internet is extremely difficult [8].

We propose here an approach that makes it possible, for the first time, to observe the dynamics of the internet's topology at the scale of a few minutes. It consists in focusing on the part of the internet's topology viewed from a single monitor, which we call an *ego-centered view* (see Figure 1). Such views are far from representative maps of the internet, but they have several key advantages. In particular, they can be obtained very rapidly with low network load, and thus may be repeated at a high frequency. This makes it possible to study their dynamics, and then to gain insight on the dynamics of the internet's topology itself.



(a) Ego-centered view from *h*.    (b) Ego-centered view from *l*.

(c) The obtained map.
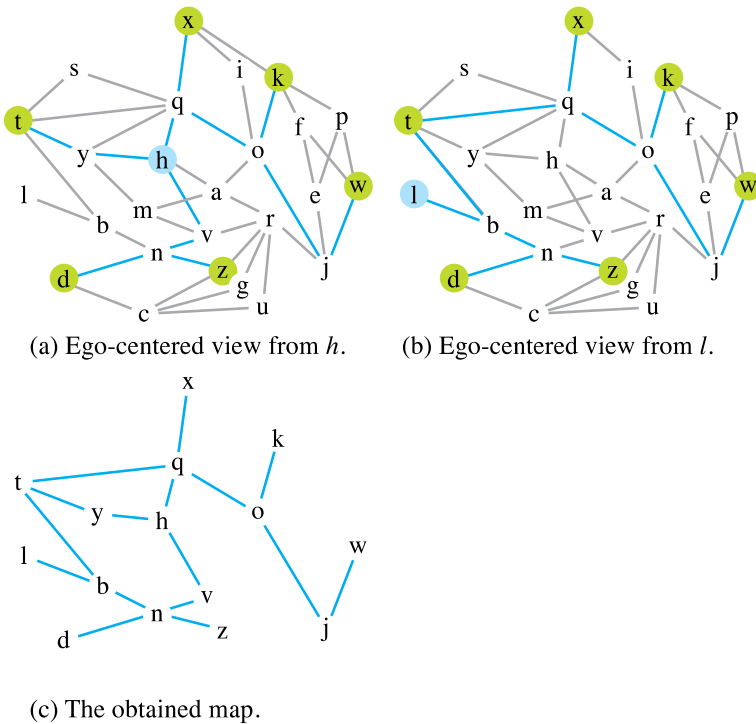
**Figure 1**. Measurement of an internet-like topology. The classical approach consists in running traceroute from a number of monitors, here *h* and *l*, toward some destinations, here *t*, *x*, *k*, *w*, *z*, and *d*. We obtain these paths from *h*: h–y–t; h–q–x; h–q–o–k; h–q–o–j–w; h–v–n–z; h–v–n–d. We obtain these paths from *l*: l–b–t; l–b–t–q–x; l–b–t–q–o–k; l–b–t–q–o–j–w; l–b–n–z; l–b–n–d. The final map of the network is obtained by merging all these paths. Instead, one may focus on the set of paths obtained from a single monitor, which we call its ego-centered view of the network.

## ▌ 2. The Measurement Tool

In principle, one may use the traceroute tool to obtain ego-centered views: it suffices to run it from a monitor toward a given set of destinations, and then merge the obtained paths, as in Figure 1. However, this approach has severe drawbacks. In particular, it is highly redundant and it induces a very heterogeneous load on links: since traceroute sends one probe for each link on the path to discover, the links close to the monitor are overloaded. For instance, the traceroute ego-centered measurement from *l* described in Figure 1 discovers link *l–b* six times, using six probes. The situation is even worse in practice, as can be seen in the supplementary material [9].

In order to perform fast ego-centered measurements with a low and balanced network load, we therefore had to design a dedicated measurement tool, called *tracetree*. The traceroute tool discovers a path by sending a series of probes toward the destination in a forward manner: the first probe discovers the first link, the second probe discovers the second link, and so on. Instead, the tracetree tool discovers a tree by sending probes toward all destinations in parallel in a backward manner and avoids redundancy by stopping probes toward some destinations when paths collapse: given a set of destinations, it first discovers the last link on the path to each of them, then the previous link on each of these paths, and so on. When two (or more) paths reach the same node, then it stops probing toward all corresponding destinations except one. See Figure 2 for an illustration, and the supplementary material for a detailed specification and implementation of tracetree [9].

The tracetree tool performs ego-centered measurements very efficiently, both regarding time and network load. It sends exactly one probe for each link to discover, and thus induces a perfectly balanced load. Radar measurements then consist of iterating ego-centered measurements with tracetree, from a monitor to a given set of destinations.

**Figure 2.** Illustration of the tracetree measurement method. (a) The first series of probes discovers the last link before each destination; tracetree stops probing toward $z$ because the paths to $d$ and $z$ collapse (at $n$). (b) The second series of probes discovers the links just before, and tracetree stops probing toward most destinations: toward $t$ and $x$ because the full path between them has been discovered; toward $k$ and $w$ because the corresponding paths collapse with previously discovered nodes. (c) The third (and last) series of probes contains only one probe, sent toward $d$, which ends the tracetree measurement. Finally, tracetree sent exactly one probe for each link, thus avoiding redundancy and reducing the network load significantly.

## 3. Measurements and Dataset

In order to conduct radar measurements relevant for our goals, several parameters have to be determined. In particular, there is a trade-off between the frequency at which these measurements are conducted (it should be as high as possible), their size (they should capture the dynamics of as many nodes as possible), and the induced network load (it must be low enough to avoid problems with network security and any bias it may induce).

Many parameters have an impact on these desirable features. As it is impossible to test all combinations to choose the best ones, we used the following approach. We first chose a set of seemingly reasonable parameters, which we call *base parameters*. Then we started measurements with these parameters from several monitors in parallel. We kept some monitors, called *control monitors*, with these parameters constant; on others, called *test monitors*, we alternated periods with base parameters and periods where we changed one parameter. The observed changes made it possible to study the influence of this parameter. Control monitors made it possible to check that the changes observed from test monitors were caused by changing parameters, not events in the network. The alternation of periods with base parameters and modified ones also made it possible to confirm this.

We provide a detailed study of parameters and their influence as supplementary material [9]. These experiments made it possible to identify several sets of parameters that reach the trade-off pointed out earlier.

We finally conducted independent measurements from more than 100 monitors scattered around the world, toward sets of 3000 destinations chosen at random among valid IP addresses. We sent probes at a maximal distance of 30 hops, and waited for answers until a timeout of 2 seconds. With these parameters, each ego-centered measurement lasted around 4 minutes. As we waited 10 minutes between two rounds (to reduce the network load), we obtained one such measurement approximately every 14 minutes, or close to 100 per day. We ran these measurements continuously during several months and obtained a very rich dataset that is provided as supplementary material [9].

## 4. Events in the Dynamics

The most natural idea to detect events in the dynamics captured by a radar measurement from a given monitor certainly is to study the number $N_i$ of nodes observed at each round $i$. A plot of a typical case is shown in Figure 3 (middle row, black plot). Clear events appear under the form of sharp decreases of $N_i$ for some values of $i$, but this brings little information: most such decreases are due to local failures of the network through which the monitor accesses the internet, which suddenly make its ego-centered view almost blank. Instead, we notice that this plot exhibits no sharp increase, which is confirmed by the distribution of the $N_i$ for all $i$, plotted in Figure 3 (top row, black plot). This is a nontrivial fact, as scenarios may very well be imagined where such increases would appear. In practice, however, the value of $N_i$ is very stable, except for sharp decreases which bring little information, if any.
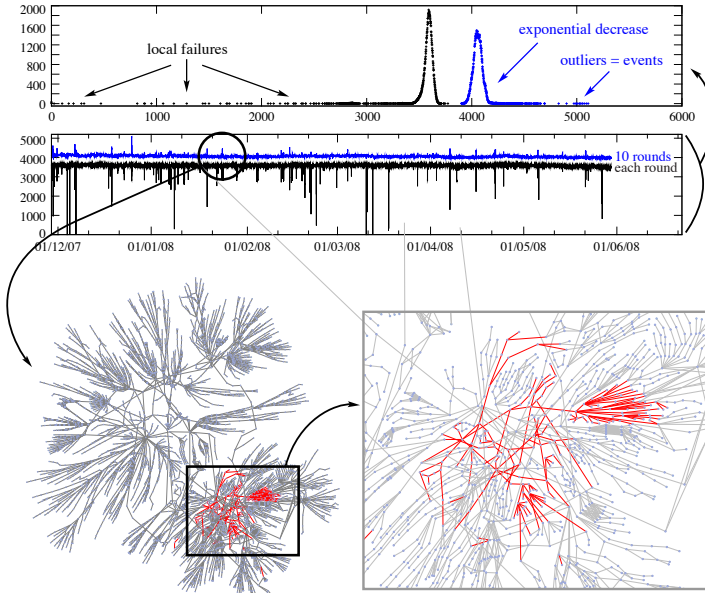
**Figure 3**. Middle row: plot of the number of distinct nodes $N_i$ (resp. $N_{(i-9)..i}$) observed during each round (resp. 10 consecutive rounds) of measurement, as a function of time. Top row: the distributions of these values, which confirms that $N_i$ exhibits abnormally small values only, never abnormally large ones, unlike $N_{(i-9)..i}$. Bottom row: topology changes observed during an event identified by an abnormally large value of $N_{(i-9)..i}$, with the part of the network where the event occurred enlarged.

The fact that the number $N_i$ of nodes observed at each round $i$ is very stable does not mean that there are no dynamics; consecutive rounds may consist of very different sets of nodes with the same size. Suppose, for instance, that we conduct a radar measurement of the network in Figure 1 from monitor $h$. Let us consider the $i^{\text{th}}$ round of measurement, for a given index $i$ at which the ego-centered measurements from $h$ is the one depicted in Figure 1(a). Then suppose that at round $i+1$ the path from $h$ to $w$ changes to $h$–$q$–$o$–$k$–$f$–$w$; and that at round $i+2$ the path from $h$ to $w$ changes to $h$–$q$–$o$–$k$–$p$–$w$ and the one from $h$ to $t$ changes to $h$–$q$–$s$–$t$. In this situation, we have $N_i = N_{i+1} = N_{i+2} = 13$, despite the fact that the ego-centered views changed significantly.

Such changes may be observed in the number of distinct nodes seen in series of consecutive rounds. Suppose, for instance, in the given scenario that the ego-centered view from $h$ did not change from round $j$ to round $i$, for $j < i$. Then the number of nodes observed during three consecutive rounds experiences an increase of more than 23%. If we

denote by $N_{x..y}$ the number of distinct nodes observed from rounds $x$ to $y$, then we have $N_{j..(j+2)} = N_{(j+1)..(j+3)} = \cdots = N_{(i-2)..i} = 13$, but $N_{(i-1)..(i+1)} = 14$ and $N_{i..(i+2)} = 16 > N_{(i-2)..i} + 23\%$.

Similarly, in Figure 3 (middle row, blue plot) we display the number of nodes observed in 10 consecutive rounds in a typical practical case, which exhibits significant increases, thus revealing events in the dynamics (significant decreases were also experienced, which we removed to improve readability, as they indicate local network failures only).

This plot has another key feature: the observed values are well centered around a typical value but also reach some extremal outlier values. See the distribution plotted in Figure 3 (top row, blue distribution). This means that the sharp increases indeed reveal events in a rigorous statistical sense.

Finally, we are able to point out precise times where events occur in the dynamics of the observed topology. This opens the way to further investigation of the shape and nature of these events, for instance by drawing the topology and the changes it experienced at these precise times. We display a typical case in Figure 3 (bottom row): it shows that, whereas the dynamics are generally scattered throughout the network, the events we detect correspond to a significant change in a specific part of the topology. This confirms that these events make sense from a networking point of view. They correspond to major changes in specific parts of the internet, which we are able to automatically detect at a time scale of a few minutes, much more precisely than all previous work in this area.

## Acknowledgments

## References

[1] J. Postel, "Internet Control Message Protocol," RFC 791, September 1981. http://www.rfc-editor.org/rfc/rfc792.txt.

[2] A. Lakhina, J. Byers, M. Crovella, and P. Xie, "Sampling Biases in IP Topology Measurements," in *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '03)*, Vol. 1, San Francisco, CA, IEEE: 2003 pp. 332–341. doi:10.1109/INFCOM.2003.1208685.

[3] L. Dall'Asta, I. Alvarez-Hamelin, A. Barrat, A. Vázquez, and A. Vespignani, "Exploring Networks with Traceroute-Like Probes: Theory and Simulations," *Theoretical Computer Science*, **355**(1), 2006 pp. 6–24.

[4] J.-L. Guillaume and M. Latapy, "Relevance of Massively Distributed Explorations of the Internet Topology: Simulation Results," in *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '05)*, Vol. 2, Miami, FL, IEEE: 2005 pp. 1084–1094. doi:10.1109/INFCOM.2005.1498336.

[5] Y. Shavitt and E. Shir, "DIMES: Let the Internet Measure Itself," *ACM SIGCOMM Computer Communication Review*, **35**(5), 2005 pp. 71–74.

[6] B. Huffaker, D. Plummer, D. Moore, and K. Claffy, "Topology Discovery by Active Probing," in *Proceedings of the 2002 Symposium on Applications and the Internet (SAINT) Workshops*, Nara, Japan, IEEE: 2002 pp. 90–96. doi:10.1109/SAINTW.2002.994558.

[7] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding Traceroute Anomalies with Paris Traceroute," in *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement (IMC '06)*, Rio de Janeiro, Brazil, ACM: New York, 2006. doi:10.1145/1177080.1177100.

[8] R. V. Oliveira, B. Zhang, and L. Zhang, "Observing the Evolution of Internet AS Topology," in *Proceedings of the 2007 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '07)*, Kyoto, Japan, ACM: New York, 2007. doi:10.1145/1282380.1282416.

[9] Supplementary material, including programs and data. http://www-rp.lip6.fr/~latapy/Radar.