# A Radar for the Internet

Assia Hamzaoui, Clemence Magnien, Frederic Ouedraogo and Matthieu Latapy
LIP6 – CNRS and University Pierre & Marie Curie
104 avenue du president Kennedy, 75016 Paris, France
Firstname.Lastname@lip6.fr

## ABSTRACT

In contrast with most internet topology measurement research, our concern here is not to obtain a map as complete and precise as possible of the whole internet. Instead, we claim that each machine's view of this topology, which we call ego-centered view, is an object worth of study in itself. We design and implement an ego-centered measurement tool, and perform radar-like measurements consisting of series of such views of the internet topology. We conduct long-term (several weeks) and high-speed (one round every few minutes) measurements of this kind from more than one hundred monitors, and we provide the obtained data. We also show that these data may be used to detect events in the dynamics of internet topology, at a time scale much smaller than before.

## 1. CONTEXT.

Since the seminal paper of Faloutsos et al [5], constructing maps of the internet using `traceroute`-like measurements received much attention. Such measurements are however partial and they may contain significant bias [6, 3]. As a consequence, much effort is nowadays devoted to the collection of more accurate data [8, 2], but this task is challenging.

In order to avoid these issues and obtain some insight on internet topology *dynamics*, we take here a radically different approach: we focus on what a given machine sees of the topology around itself, which we call an *ego-centered view*. These ego-centered measurements may be performed very efficiently (typically in minutes, and inducing low network load); it is therefore possible to repeat them in periodic rounds, and obtain in this way information on the *dynamics* of the topology, at a time-scale significantly higher than previous approaches.

Taking advantage of these strengths, we conduct massive radar-like measurements of the internet. We provide both the measurement tool and the collected data, and show that they reveal interesting features of the observed topology.

## 2. MEASUREMENT FRAMEWORK.

One may use `traceroute` directly to collect ego-centered views by probing a set of destinations. This approach however has serious drawbacks. First, as detailed in [4], the measurement load is highly unbalanced between nodes and there is much redundancy in the obtained data (intuitively, one probes links close to the monitor much more than others). Even worse, this implies that the obtained information is not homogeneous, and thus much more difficult to analyse rigorously (for instance, the dynamics may seem higher close to the monitor). Finally, though the measurement would intuitively produce a routing tree, the obtained view actually differs significantly from a tree. Again, this makes the analysis (visualisation of the data, for instance) more intricate.

One may avoid the issues described above by performing tree-like measurements in a backward way: given a set of destinations to probe, one 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 the probing towards all corresponding destinations, except one, stops [1].

However, naive such measurements encounter serious problems because of routing changes and other events. A solution is provided by the `tracetree` algorithm below: the tree nodes are no longer IP addresses, but pairs composed of an IP address (or a star if a timeout occurred) and the TTL at which it was observed. This is sufficient to ensure that the obtained view is a tree.

From such trees with (IP,TTL) nodes, one obtains a tree on IP addresses by applying the following filter: first merge all nodes of the tree which correspond to a same IP; remove loops (links from an IP to itself); iteratively remove the stars with no successor; merge all the stars which are successor of a same node into a unique star; construct a BFS tree of the obtained graph which leads to a tree on IP addresses; iteratively remove the leaves which are not the last nodes encountered when probing any destination.

The key point is that the obtained tree is a possible IP routing tree from the monitor to the destinations

[1] Such measurements require the distance towards each destination, which is not trivial;

(similar to a broadcast tree). It contains almost as much information as the original `tracetree` output and has the advantage of being much more simple to analyse.

With the `tracetree` tool and its filtered version, we have the ground material to conduct radar measurements: given a monitor and a set of destinations, it suffices to run periodic ego-centered measurements. The measurement frequency must be high enough to capture interesting dynamics, but low enough to keep the network load reasonable.

## 3. MEASUREMENT AND DATA.

We use a wide set of more than one hundred monitors scattered around the world, provided by PlanetLab [7] and other structures (small companies and individual DSL links) [1]. In order to be as general as possible, and to simplify the destination setup, we use destinations chosen by sampling random valid IP addresses and keeping those answering to `ping` *at the time of the list construction*. Other selection procedures would of course make sense. In our context, we will see that the choice we have just described is sufficient.

The parameters consist of a set of 3 000 destinations for each monitor, a maximal TTL of 30, a 2 seconds timeout and a 10 minutes delay between rounds. All our measurements were conducted with variations of these parameters; wherever it is not explicitly specified, the parameters were the base ones. We ran measurements continuously during several weeks, with some interruptions due to monitors and/or local network shutdowns.

The obtained data is available at [1].

## 4. TOWARDS EVENT DETECTION.

One key interest of our measurements is that they make it possible to observe the dynamics of internet topology from an ego-centered perspective, at a time scale significantly smaller than before. In particular, detecting *events* in this dynamics, *i.e.* major changes in the topology, is very appealing from a security and modeling point of view.

## 5. CONCLUSION AND PERSPECTIVES.

We propose, implement, and illustrate a new measurement approach which makes it possible to study the dynamics of IP-level internet topology at a time scale of a few minutes. We provide a rich dataset consisting in radar-like measurements from more than one hundred monitors towards thousands of destinations, conducted for several weeks in continuous.

The most important direction for further research is of course the analysis of the collected data. Among them, event detection is particularily appealing but it raises difficult fundamental questions: what dynamics should be considered as normal, as opposed to *events*? which observation time scales are the most relevant? which (graph) properties should we observe? how to interpret detected events from a network perspective? how to visualise the dynamics and capture it into relevant models ? etc. One may also conduct new kinds of measurements based on the same approach, like measurements targetting a specific network area (either by limiting the maximal TTL to a very small value, or by choosing destinations in this area), and/or combine radar-like measurements from several sources to gain deeper insight.

## 6. REFERENCES


[1] Supplementary material (programs and data). `http://www-rp.lip6.fr/~latapy/Radar-sub/`.

[2] traceroute@home project. `http://trhome.sourceforge.net/`.

[3] P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the marginal utility of network topology measurements. In *Proc. ACM SIGCOMM Internet Measurement Workshop (IMW)*, Nov. 2001.

[4] B. Donnet, P. Raoult, T. Friedman, and M. Crovella. Efficient algorithms for large-scale topology discovery. In *Proc. ACM SIGMETRICS*, Jun. 2005.

[5] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *ACM SIGCOMM*, 1999.

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

[7] PlanetLab Consortium. PlanetLab project, 2002. See `http://www.planet-lab.org`.

[8] Y. Shavitt and E. Shir. DIMES: Let the internet measure itself. *ACM SIGCOMM Computer Communication Review*, 35(5):71 – 74, October 2005.