Relevance of Massively Distributed Explorations of the Internet Topology: Simulation Results

Jean-Loup Guillaume and Matthieu Latapy

LIAFA, CNRS – *Université Paris 7 – 2 place Jussieu, 75005 Paris, France.*
(guillaume,latapy)@liafa.jussieu.fr

Internet maps are generally constructed using the `traceroute` tool from a few sources to many destinations. It appeared recently that this exploration process gives a partial and biased view of the real topology, which leads to the idea of increasing the number of sources to improve the quality of the maps. In this short abstract, we present a set of experiments we have conduced to evaluate the relevance of this approach. It appears that the statistical properties of the underlying network have a strong influence on the quality of the obtained maps, which can be improved using massively distributed explorations. Conversely, we show that the exploration process induces some properties on the maps. We validate our analysis using real-world data and experiments and we discuss its implications.

Keywords: Internet topology, graphs, metrology, active measurements.

Introduction.

Due to its fully distributed construction and administration, mapping the Internet (in terms of machines and physical links between them) is a challenging task. It is however essential to obtain some information on its global shape which plays a central role in key problems like network robustness [AJB00, CEbAH00], simulation of future protocols and uses [MP01b], and many others.

Exploring the Internet topology is a research problem in itself [GT00, MP01a, VPSV]. Indeed, many difficulties (like the identification of the multiple interfaces of a same router) arise when one wants to map the Internet. Various techniques and methods have been introduced to achieve this goal. Some of them are very subtle, but current explorations still rely on the extensive use of the `traceroute` tool: one collects routes from a given set of sources to a given set of destinations, and then merges the obtained paths.

Two points are particularly important in the scheme sketched above. First, it must be clear that the image we obtain from the network is *partial* (except if the number of sources and destinations is huge, we certainly miss some nodes and some links) and may be *biased* by the exploration process (some properties of the obtained map may be induced by the way we explore the network). Second, whereas one only needs the IP addresses of the destinations, the number of sources cannot be increased easily, one needs direct access to the sources in order to run the `traceroute` tool. In the case of [fIDA], which is the largest exploration currently available, only a few dozens of sources are used and the number of destinations is several hundreds of thousands.

Recently, several researchers conduced experimental and formal studies to evaluate the accuracy of the obtained maps of the Internet [LBCX03, HBC, GL04b]. All these studies use simple models of networks and `traceroute` but they all give good arguments of the fact that the currently available maps of the Internet are very incomplete, and that there probably is an important bias induced by the exploration process.

In order to improve these maps, several researchers and groups now propose to deploy massively distributed measurement tools [fIDA, pro]. The basic idea is that dramatically increasing the number of sources would significantly improve the quality of the obtained maps. Our central aim in this paper is to rigorously evaluate the relevance of this approach.

To achieve this, we conduce an extensive set of experiments designed as follows. We consider a graph G representing the network to explore. We then simulate the exploration process and obtain this way a (partial and biased) view G' of the original graph. We then compare G and G' to evaluate the quality of this view. We process this simulation using all the possible numbers of sources and destinations, which makes it possible to study the impact of these numbers on the accuracy of the obtained view.

This short abstract is organized as follows. First we define the statistical properties of networks relevant to our study, we present the models we use and discuss our methodology (Section 1). Then we present and analyze the results of our simulations on various models and statistical properties (Section 2). Finally we present our conclusions and discuss them. Some results are not displayed in this short abstract but are going to be briefly discussed in the conclusion. Refer to [GL04a] for a deeper insight.

1 Preliminaries.

A network topology can naturally be represented by a graph. For our purpose, the graph does not need to be weighted nor directed. A route in the network, as given by the `traceroute` tool, is a path in the corresponding graph. Since a few years [PG98, CMM, FFF99, fIDA], a strong effort has been made to discover the topology of the Internet by extensive use of `traceroute` and other tools (BGP tables, source routing, etc).

The obtained maps give much information on the global architecture of the Internet. In particular, they gave evidence of the fact that the Internet topology has some statistical properties which make it very different from the models used until then [FFF99, CCG$^+$02]. This induced an intense activity in the acquisition of such maps [PG98, GT00, fIDA], in their analysis [FFF99] and in the accurate modelling of the Internet [Wax88, ZCD97]. See [PSV04] for a survey.

Our analysis of the exploration process will be based on these statistical properties and these models, which we present below. We also need to model the `traceroute` tool and the exploration process, which we also discuss in this section. Finally, we present our methodology, and explain how our results should be read.

Statistical properties
The Internet, at router level, is composed of several millions of nodes and dozens of millions of links. Let N denote its number of nodes and M its number of links. The density of this network defined as the number of existing edges over the number of possible ones, $\frac{2\cdot M}{N\cdot(N-1)}$, is low.

Others well known properties are the following: the average distance (length of a shortest path between two nodes) is low, and typically scales as $\log(N)$. The degree distribution of the Internet graph follows a power law [FFF99]: the proportion p_k of nodes of degree k scales as a power of k: $p_k \sim k^{-\alpha}$ with $\alpha \simeq 2.5$. Intuitively, this means that there is a high variability in degree. Finally, the clustering C of this graph, defined as the probability that two nodes are connected together given that they are both connected to a same third, is high.

Modelling networks
Hereafter are presented the models used in this short abstract. All these models (and some others) are widely used for the modelling of complex networks.

The basic model for networks is the Erdos and Rényi (ER) random graph model [Bol85]. In an ER graph with n nodes, each of the possible links exists with a given probability p. Another basic model is the Albert and Barabási (AB) one [AB99], which is based on *preferential attachment*. In this model, nodes arrive one by one and choose k neighbors among the existing ones with a probability proportional to their degree. This model yields non clusterized scale-free graphs. This model has been modified to give highly clusterized graphs: in the Dorogovtsev and Mendes (DM) model [DM02], nodes arrive one by one but at each step one chooses a random *link* $\{u,v\}$ and the new node is linked to both u and v. This implies that a node is chosen with a probability proportional to its degree and therefore the resulting graphs are scale-free. Moreover, since one forms a triangle at each step, they have a high clustering. It is also possible to sample a random graph with a prescribed degree distribution using the Molloy and Reed (MR) model [MR98]. This gives graphs with exactly the wanted degree distribution, but with low clustering.

All these models capture the low density and the logarithmic average distance of complex networks. However, ER model does not capture the power law degree distribution, nor the high clustering. The three others models (namely AB, DM and MR) capture the degree distribution but DM is the only one to capture also the clustering (notice that others models succeed in capturing all these properties).

Modelling `traceroute` and the exploration
In this short abstract, we will make the assumption that a route as obtained by `traceroute` is nothing but a shortest path between the source and the destination. It is known that this is not always true [FLLS, HBC], but the realistic modelling of routes is nowadays an open problem.

We also need a model for the exploration process. We considered two points of view: in the first one we suppose we make a *snapshot* of the network, and in the second one we make a *long-time* exploration. This leads respectively to the *unique shortest path* (USP) model, and to the *all shortest paths* (ASP) one: we either see only one route for any given source and destination, or we see all the possible ones. In the first case, in order to choose one route among all, we use the assumption that the route followed by a packet when a given router R routes it to a destination D is always the same independently of the sender.

We have studied other models of `traceroute` (random shortest path, several shortest path but not all, etc), but the results do not qualitatively vary, so we do not detail them here.

Methodology
Our global approach is as follows: we first generate a graph G using a given model with some known parameters, then we compute a view G' of G using an exploration process, and we compare the statistical properties of G' to the ones of G. In this abstract, sources and destinations are selected *at random*.

In the method sketched above, the third point (comparison of the original graph with the view we obtain) is a difficult task. To achieve it, we will make an extensive use of grayscale plots defined as follows: for a graph G of N nodes, we consider a square of size $N \times N$. Each point (x, y) of the square corresponds to a view G' of G using x sources and y destinations. The point is drawn using a grayscale representing the value of the property p under consideration: from black for $p = 0$ to white for the maximal value of p (which might be greater than the real value). To improve the grayscale plots readability, we added on each such plot the 0.25–, 0.50–, 0.75– and 0.99–level lines, where the l–level line is defined as the set of points where the value of p is about l times the maximal value.

2 Proportion discovered

We are now going to explore graphs G to get a view G' and compare both graphs in terms of *proportion of nodes, proportion of edges and average degree*. Let us insist on the fact that we seek *qualitative* results only: we want to know how qualitative properties of the network influences the exploration and how reliable are the obtained maps. It makes no sense to interpret quantitatively the results obtained with the kind of approach we use here. On the contrary, by the simplicity of the models and of the properties we use, we obtain evidences of the fact that some properties play a fundamental role in the exploration whereas others may be neglected.

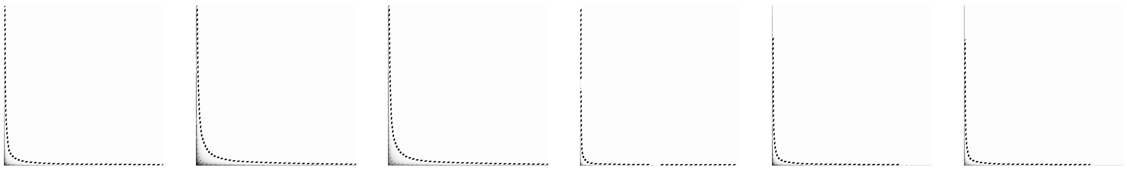

Figure 1: ER graph, $< k >= 10$, $N = 10^4$: number of vertices, edges, and average degree, for USP. Idem for ASP.

Grayscale plots in Figures 1 and 2 show the first results on ER graphs. When the average degree is quite small (Figure 1), there is no qualitative difference between ASP and USP (there exists very few shortest path between any two nodes) and the quality of the view is good even for a very small numbers of sources and destinations. On the contrary, when the average degree grows (Figure 2), so does the number of shortest paths, and the difference between ASP and USP becomes significant.

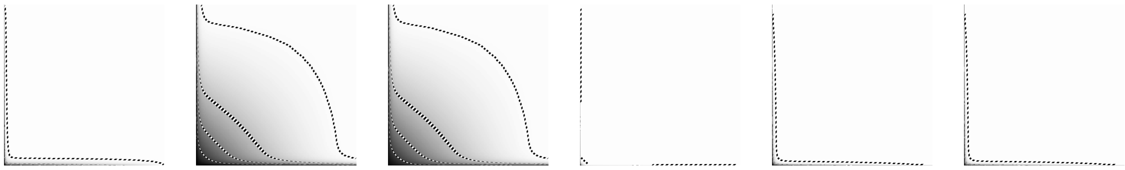

Figure 2: ER graph, $< k >= 100$, $N = 10^4$: number of vertices, edges, and average degree, for USP. Idem for ASP.

Let us now observe what happens when we consider scale-free graphs. First with the AB model which makes it possible to obtain scale-free graphs with a given average degree. In Figure 3, we can see that the efficiency of the exploration on such graphs is similar to the one on ER graphs, though it is lower. If we want a very precise map, however, we need much more sources and destinations. There is also a strong difference between USP and ASP, which tends to show that there are multiple shortest paths between nodes.

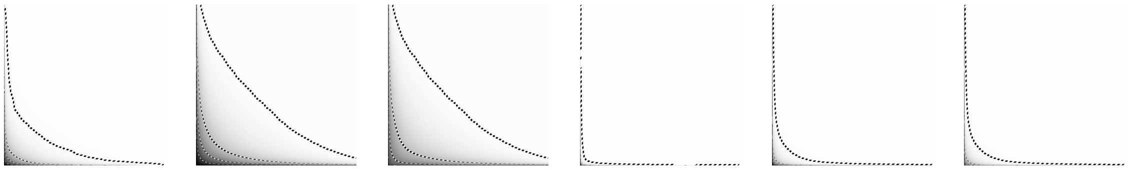

Figure 3: AB graph, $< k >= 10$, $N = 10^4$: number of vertices, edges, and average degree, for USP. Idem for ASP.

If we make the same experiments with MR graphs (Figure 4), the quality of the obtained view is much worse than for AB graphs (Figure 3). Even when considering ASP, one needs to take about half sources and destinations to view 75% of the graph (both in terms of edges and nodes).

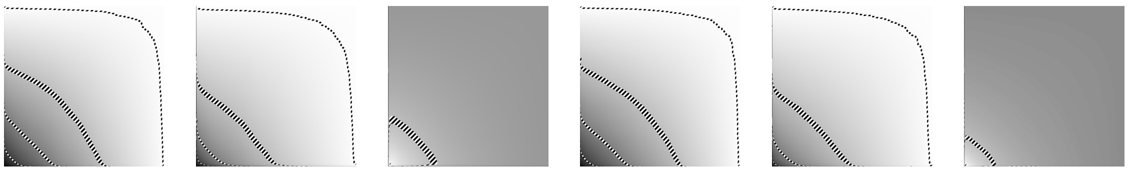

Figure 4: MR graph $\alpha = 2.5$, $N = 10^4$: number of vertices, edges, and average degree, for USP. Idem for ASP.

There is a simple explanation of this fact: in an AB graph, due to the construction process, there is no node with degree lower than $\frac{k}{2}$ (the number of links added at each step), whereas MR graphs contains a lot of low-degree nodes (including nodes of degree 1). These nodes are difficult to discover since they lie on very few shortest paths: a node of degree 1 is discovered only when it is chosen as a source or a destination. This can also be deduced from the fact that the average degree is overestimated which means that high degree nodes are discovered first. To check this hypothesis, we considered the *core* of a MR graph defined as the graph obtained by removing recursively all the nodes of degree 1. Exploration of such a graph yields results that are more in accordance with the ones for the AB graphs, which allows us to conclude that low degree nodes (and more generally small tree-like structures) are hard to discover.

Let us now consider a DM graph, which is scale-free and *clusterized*. Like in an AB graph, there is no node with only one link, therefore, the effect noticed above in MR graphs should not appear. However Figure 5 shows that we again obtain low quality maps of this kind of graphs. The fact that the plots for USP and ASP are very similar indicates that there are very few different shortest paths between nodes. This, and the fact that the quality of the obtained views is low, can be understood as follows: when one wants to explore a clique (complete graph), or a dense graph, one has to use a large number of sources and destinations. The high clustering in DM graphs is equivalent to the fact that there are many subgraphs which are cliques or almost. All these parts of the graph are difficult to explore.

Before concluding, let us discuss briefly the average degree obtained using an exploration. The fact that this property is obtained by dividing two other properties which are improved by the use of more sources

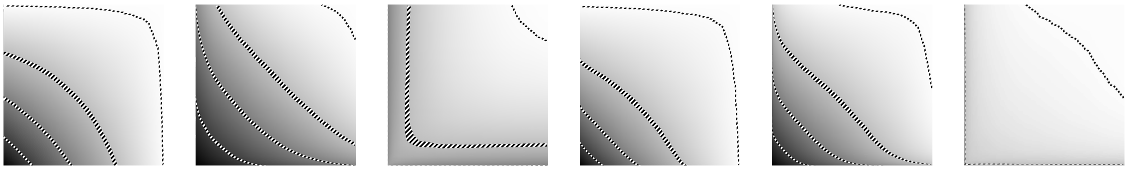

Figure 5: DM graph, $N = 10^4$: number of vertices, edges, and average degree, for USP. Idem for ASP.

and/or destinations has important consequences. If one of the two properties is highly biased and the other is not, then the average degree will have a strong bias (see Figure 2 and 3 for instance). However, if each of the two properties has the same kind of bias, this may be hidden by the division: the evaluation of the average degree is good even if the number of edges and nodes is completely inaccurate (Figure 4).

Finally, the conclusion of this section is the following: two properties of graphs make them hard to explore in different ways. The first one is the large number of tree-like structures around the core of the graph. The second one is the high clustering which induces many dense subgraphs. The two properties are complementary and act on different parts of the graph (on the border and on the core, respectively), which indicates that we should take them both into account.

Conclusion and discussion

We conduced an extensive set of simulations aimed at evaluating the quality of current maps of the Internet and understanding how to distribute explorations massively to improve it. To achieve this, we considered the most commonly used models of graphs (namely the ER, the AB, the MR and the DM ones). We constructed *views* of these graphs and compared them to the original graphs in terms of proportion of the graph discovered (both in terms of nodes, links and average degree). Other properties have been studied, such as the average distance, the degree distribution and the clustering, but not detailed here. However we are going to give some conclusions about these results:

- Two statistical properties of graphs influence strongly our ability to obtain accurate views of them: the presence of many low-degree nodes and the high clustering. These two properties act independently and their effects are combined in the case of the Internet.
- It is relevant to use massively distributed exploration schemes to obtain accurate maps of scale-free clusterized networks like the Internet, in particular if we want to discover most nodes and edges, and have an accurate estimation of the clustering.
- On the contrary, the evaluation of the degree distribution of such a network, as well as its average distance, is achieved with very good precision even for reasonably small number of sources and destinations.
- The details of the exploration scheme (for instance USP versus ASP or the behavior of `traceroute`) tends to have little importance when the number of sources and destination grows. In the case of the Internet, this means that distributing explorations can be viewed as a way to improve the independence of the results from the exploration scheme.
- Others result not shown here lead us to think that despite power-law degree distribution and high clustering play a role in the efficiency of the explorations of the Internet, other unidentified properties also influence it.

More investigations are currently in progress. First, we are considering more subtle statistical properties, like the correlations between node degrees, or the correlations between degree and clustering. and we also plan to conduct some experiments with more realistic models of `traceroute`.

Finally, let us insist on the fact that most real-world complex networks, like the World Wide Web and Peer to Peer systems, but also social or biological networks are generally not directly known. Various exploration schemes are used to infer maps of these networks, which may influence the vision we obtain. The metrology of complex networks is therefore a general scientific challenge, for which the goal is to be able to deduce properties of the real network from the ones observed.

References

[AB99] R. Albert and A.-L. Barabási. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

[AJB00] R. Albert, H. Jeong, and A.-L. Barabási. Error and attack tolerance in complex networks. *Nature*, 406:378–382, 2000.

[Bol85] B. Bollobás. *Random Graphs*. Academic Press, 1985.

[CCG$^+$02] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. Shenker, and W. Willinger. The origin of power laws in internet topologies revisited. In *INFOCOM*, 2002.

[CEbAH00] R. Cohen, K. Erez, D. ben Avraham, and S. Havlin. Resilience of the internet to random breakdown. *Phys. Rev. Lett.*, 85:4626–4628, 2000.

[CMM] K. Claffy, T. Monk, and D. McRobb. Internet tomography. Nature Magazine, Web Matters. http: //helix.nature.com/webmatters/tomog/tomog.html.

[DM02] S.N. Dorogovtsev and J.F.F. Mendes. Evolution of networks. *Adv. Phys. 51, 1079-1187*, 2002.

[FFF99] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, 1999.

[fIDA] Cooperative Association for Internet Data Analysis. http://www.caida.org/.

[FLLS] T. Friedman, M. Latapy, J. Leguay, and K. Salamatian. What is a route on the internet. preprint.

[GL04a] J.-L. Guillaume and M. Latapy. Relevance of massively distributed explorations of the internet topology: Simulation results. preprint - http://www.liafa.jussieu.fr/~latapy/Publis/, 2004.

[GL04b] J.-L. Guillaume and Matthieu Latapy. Complex network metrology, 2004. preprint - http://www.liafa.jussieu.fr/~latapy/Publis/.

[GT00] R. Govindan and H. Tangmunarunkit. Heuristics for internet map discovery. In *IEEE INFOCOM 2000*, pages 1371–1380, Tel Aviv, Israel, March 2000. IEEE.

[HBC] Y. Hyun, A. Broido, and K. Claffy. Traceroute and BGP AS path incongruities. http://www.caida.org/outreach/papers/2003/ASP/.

[LBCX03] A. Lakhina, J. Byers, M. Crovella, and P. Xie. Sampling biases in IP topology measurements. In *IEEE INFOCOM*, 2003.

[MP01a] D. Magoni and J.-J. Pansiot. Analysis of the autonomous system network topology. *ACM SIGCOMM Computer Communication Review*, 31(3):26 – 37, July 2001.

[MP01b] D. Magoni and J.-J. Pansiot. Influence of network topology on protocol simulation. In *ICN'01 - 1st IEEE International Conference on Networking*, volume Lecture Notes in Computer Science, pages 762–770, July 9-13, 2001.

[MR98] M. Molloy and B. Reed. The size of the giant component of a random graph with a given degree sequence. *Combin. Probab. Comput.*, pages 295–305, 1998.

[PG98] J. Pansiot and D. Grad. On routes and multicast trees in the internet. *ACM Computer Communication Review*, 28(1):41–50, 1998.

[pro] Traceroute@Home project. University of Paris 6, coordinator: Timur friedman.

[PSV04] R. Pastor-Satorras and A. Vespignani. *Evolution and Structure of the Internet: A Statistical Physics Approach.* Cambridge University Press, 2004.

[VPSV] A. Vazquez, R. Pastor-Satorras, and A. Vespignani. Internet topology at the router and autonomous system level. [cond-mat/0206084].

[Wax88] B.M. Waxman. Routing of multipoint connections. *IEEE Journal of Selected Areas in Communications*, pages 1617–1622, 1988.

[ZCD97] E.W. Zegura, K.L. Calvert, and M.J. Donahoo. A quantitative comparison of graph-based models for Internet topology. *IEEE/ACM Transactions on Networking*, 5(6):770–783, 1997.