



HAL
open science

Evaluation of a new method for measuring the internet degree distribution: Simulation results

Christophe Crespelle, Fabien Tarissan

► **To cite this version:**

Christophe Crespelle, Fabien Tarissan. Evaluation of a new method for measuring the internet degree distribution: Simulation results. *Computer Communications*, 2011, 34 (5), pp.635-648. 10.1016/j.comcom.2010.06.006 . hal-01146072

HAL Id: hal-01146072

<https://hal.science/hal-01146072v1>

Submitted on 23 Oct 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Evaluation of a New Method for Measuring the Internet Degree Distribution: Simulation Results

Christophe Crespelle and Fabien Tarissan

LIP6 – CNRS and Université Pierre et Marie Curie – Paris 6

104 avenue du président Kennedy, 75016 Paris, France

`Firstname.Lastname@lip6.fr`

Abstract

Many contributions rely on the degree distribution of the Internet topology. However, current knowledge of this property is based on biased and erroneous measurements and is subject to much debate. Recently, in [7], a new approach, referred to as the *Neighborhood Flooding* method, was proposed to avoid issues raised by classical measurements. It aims at measuring the neighborhood of Internet core routers by sending `traceroute` probes from many monitors distributed in the Internet towards a given target router. In this paper, we investigate the accuracy of this method with simulations. Our results show that Neighborhood Flooding is free from the bias highlighted in the classical approach and is able to observe properly the exact degree of a vast majority of nodes in the core of the network. We show how the quality of the estimation depends on the number of monitors used and we carefully examine the influence of parameters of the simulations on our results. We also point out some limitations of the Neighborhood Flooding method and discuss their impact on the observed distribution.

1 Introduction

”A very general but largely ignored fact about Internet-related measurements is that what we can measure in an Internet-like environment is typically not the same as what we really want to measure (or what we think we actually measure).” [17] This is particularly true for `traceroute` measurements of the Internet topology. Maps are obtained by merging paths acquired from a limited number of monitors to many targets [15, 10, 5]. It has been shown

that doing so leads to intrinsically biased views [13, 1, 8, 11, 12, 14], on which is based our current knowledge of Internet topology properties.

It is often assumed that the degree distribution of the Internet is heterogeneous and close to a power-law distribution. Since seminal papers of Pansiot *et al.* [15] and Faloutsos *et al.* [10] emphasized this property, it has been extensively studied. It appeared to have a great impact on fundamental properties of the network, such as robustness and quality of service.

In [13], Lakhina *et al.* experimentally showed that estimating the degree distribution by the classical method with 10 monitors and around 1 000 targets (which are parameters similar to those of the data set used in [15, 10]) leads to the observation of an heterogeneous distribution even if the underlying topology has a Poisson degree distribution. This result was proved formally in [1] in the case of a single source, and extended to the case of other graph topologies. This raises the question of which part of the observed distribution results from properties of the network itself and which part stems from the measurement process.

It is striking to realize that whereas the degree distribution of the Internet is known as one of its fundamental property, and even though it is also widely admitted that our knowledge of this property relies on biased measurement processes, no alternative method has been proposed in the last decade. Instead, most effort has been done to lead wider measurements, following – as named in [3] – a “more is better” approach. Very recently, a radically different method, which we refer to as *Neighborhood Flooding* in the rest of the paper, was proposed [7]. It aims at rigorously estimating the degree distribution of the central part of the Internet, classically referred to as the *core*. Instead of gathering `traceroute` outputs together in order to obtain a sample of the network topology whose representativity is uncertain, this method consists in dedicating all measurement resources to independently determine the exact degree of a random node in the core of the network, as suggested in [13]. In [7], the authors claimed that this approach is free from the usual bias induced by the collection of a sample of the graph. In this paper, we investigate in depth, using simulations, the quality of the estimation of the distribution given by Neighborhood Flooding.

Contribution The main contribution of the paper is to show experimentally that the Neighborhood Flooding method is free from the bias outlined in [13]: it does not modify the nature of the considered distribution, both for the Poisson and the power-law case, and gives a correct estimation of its main characteristics. Furthermore, we give evidence that the method is

even able to determine the exact degree of the majority of nodes of the network. We also point out its limitations and study their impact on observed distribution.

Outline of the paper Section 2 describes the Neighborhood Flooding method and its expectations. Section 3 presents our simulation framework, the graph models we use and the modeling choices we made. Section 4 is devoted to the presentation of our main results concerning quality of the estimations of the degree distribution given by Neighborhood Flooding. Section 5 studies the sensitivity of those estimations to graph models parameters. In Section 6 we complete our analysis by determining the impact of modeling choices we made on these results. Finally, Section 7 draws some concluding remarks and outlines perspectives opened by the present work.

2 Neighborhood Flooding

Unlike the classical approach, Neighborhood Flooding does not intend to collect any sample of the graph. It is only dedicated to the estimation of the degree distribution. Given a node in the core of the network, the method attempts to discover all of its neighbors in order to deduce its exact degree. Doing so for many random core nodes allows to infer the actual degree distribution. Clearly, the higher the number of nodes is, the more precise the inference will be.

In [7], this approach was implemented and used to perform a sample measurement. In order to get the degree of a given target router, ICMP `traceroute` probes are launched from all available monitors towards this target. The last hop of each `traceroute` gives an interface (IP address) of a neighbor of the target¹. This sample measurement was performed from hundreds of monitors towards 10 000 random targets within a very short time (about 30 hours). A key point is that this atomic operation can be repeated successively and independently on different sets of targets in order to make the size of the sample used to infer the distribution grow. This means that the Neighborhood Flooding method is practically feasible and efficient. Therefore, the question of whether it does properly estimate the degree of a target in the core of the network is of prime importance.

Classically, the Internet topology is divided into two parts: the *core* and the *border* (see Figure 1). The core is defined as the network obtained by

¹This implies the use of anti-aliasing techniques in order to detect different interfaces advertised by a single router.




Figure 1: Left: an example of target t in the core. The routes coming from the set of monitors go through all of its neighbors. Right: an example of target t' in the border. Only the neighbor a of t' is accessible from the core. In the two figures, the light colored squares stand for the monitors, the big nodes denote a node in the core while the small ones represents the nodes in the border and the black links belong to a route from a monitor to the target.

removing all trees attached to it (*i.e.* iteratively removing all nodes with only one link). The border is then defined as the union of the removed trees. As noted in [7], it is inappropriate to observe the degree of nodes located in the border using Neighborhood Flooding. Indeed, if the target is in a tree structure – which is the case of t' in Figure 1 (right) –, then all the routes coming from the monitors that are not in the subtree of the target will go through the neighbor of the target with higher level in the tree – namely node a in the figure. Since the number of monitors one can expect to use is very low compared to the number of nodes in the border of the network, with high probability, the chosen target does not lie on the route between any monitor and the core of the network. Thus, the degree of a target in the border will be evaluated to one with high probability though it can be much higher: t' has many links to the end-hosts it provides traffic to. Those links are impossible to detect with monitors out of the tree. On the opposite, if the target is in the core, it is reasonable to expect reaching it by going through all of its neighbors – as for t in Figure 1 (left). The crucial point here lies in the fact that the set of monitors is numerous and distributed enough in the network so that the probes reach the target by coming from all the directions. This is the main question addressed by the present paper.

3 Simulation setup

Our simulations consist in constructing an artificial graph, then randomly selecting a set of nodes which will serve as monitors, and simulating `traceroute` measurements from these monitors to all core nodes (called targets). In this way, we are able to observe the quality of the estimation of the degree for each possible target. In this section, we detail the modeling choices we adopted, the parameters of the simulations and the settings we used for deriving our results.

3.1 Choice of graph topologies

In order to study the behavior of Neighborhood Flooding on different topologies, we performed the simulations both on graphs with a Poisson degree distribution (referred to, in the following, as *Poisson graphs*) and on graphs with a power-law degree distribution (referred to as *power-law graphs*). Those are typical representatives for respectively homogeneous and heterogeneous distributions, and studying the behavior of the method on these two topologies is meaningful. As we could hardly perform simulations on graphs having the size of the actual Internet, we considered different graph sizes, from 1 million to 20 million nodes, and study the effect of size variations on the results. We also varied graph model parameters, respectively the mean degree for Poisson graphs and the exponent for power-law graphs. This allowed us to determine how sensitive is the method according to these parameters. Finally, in order to check that our results are not biased by a particular randomly generated graph, we picked 3 graphs for each choice of parameters. It turned out that the variations induced by these different choices are marginals. This is the reason why we will study the impact of the variations induced by different choices once and for all in Section 4.3. All the other results of the paper are presented on a single arbitrarily chosen graph, while, in fact, simulations are always run on 3 different choices of graph (unless explicitly mentioned otherwise).

Poisson graphs were generated using the $G_{n,m}$ Erdős-Rényi graph model [9]. Basically, it starts with creating an empty graph of n isolated vertices, then chooses uniformly at random m pairs of vertices, without repetition of the pairs, and adds a link between each of them. These graphs are known for having a Poisson degree distribution and they are completely described by their two parameters n and m , or equivalently by n and their mean degree δ , where $\delta = \frac{2m}{n}$. The later system is the one we use in the following. We study sensitivity of the method to this parameter by generating graphs with

different δ 's ranging over $\{6, 12, 25, 50\}$. This set spans a range of values from the supposed mean degree of the core of the Internet – roughly around 6 or 12 – to higher values, which should not favor the method. Poisson graphs we use have sizes of 1 million, 2.5 millions, 5 millions and 10 millions of nodes. Thus, totally, our simulations involve 48 different Poisson graphs.

Power-law graphs were generated according to an extension [16] of the well known Power-Law Random Graph model introduced in [2]. Given two integers a and b standing for respectively the minimum and maximum degree in the graph, the process mainly consists in choosing at random n integers between a and b , with repetitions, following a power-law distribution. In other words, integer k is chosen with probability $C k^{-\alpha}$, where $\alpha > 0$ is the exponent of the power-law and C a constant fixed so that $\sum_{k=a}^b C k^{-\alpha} = 1$. When generating a graph, n is set to its desired number of vertices, the minimum degree a is set to 1, the maximum degree b is set to $n - 1$, and the generated integers define the degree sequence of the graph. Each vertex is assigned to a number of semi-links according to its degree and pairs of semi-links are randomly chosen to be connected together. In [16], this process was implemented² so that it generates uniformly at random a connected graph containing no loops nor multiple edges. As for Poisson graphs, in order to study sensitivity to graph model parameters (α in this case) we generate graphs with different α 's ranging over $\{2.1, 2.2, 2.3, 2.4\}$. This range of values was chosen to contain the exponent of the degree distribution of the Internet observed in [15, 10], namely 2.1, and which is commonly assumed to lie between 2 and 3. Power-law graphs we use have sizes of 5 millions, 10 millions and 20 millions of nodes. Totally, we use 36 different power-law graphs.

Note that Poisson graphs and power-law graphs have a very different behavior with respect to core extraction. In the set of Poisson graphs we used, the core is often the whole graph and contains at least 98 % of the vertices in the other cases (this actually depends on the mean degree). On the contrary, cores of power-law graphs we used contain roughly one third of the vertices of the graph. In the rest of the paper, unless explicitly mentioned otherwise, size of graphs, both for Poisson and power-law, refers to the number of nodes in the whole graph, not in the core.

²see <http://fabien.viger.free.fr/liafa/generation/>

3.2 Modeling traceroute

Following classical modeling approaches [13, 12], supported by empirical observations [12], we assume that `traceroute` delivers a random shortest path. This makes the simulation simpler to run and interpret, while having a limited impact on the results presented here. One could think of several ways to choose a random path, maybe not all equally pertinent with regard to modeling purposes. Taking this issue into account is out of the scope of our work and we avoid it by choosing a path between a monitor and a target uniformly at random among all possible shortest paths. Yet, the reader may refer to Section 6.2 to have an insight on how the method performs according to some other possible choices.

The uniform random choice can be performed between a single monitor and all possible targets in the graph at once very efficiently, namely within a linear complexity. Note that the specificity of the Neighborhood Flooding method does not require the whole path but only the last link. Thus, we only need to determine for each neighbor of a given target the probability that it lies on the shortest path chosen uniformly at random among all shortest paths between the considered monitor and the target. We determine these probabilities for all targets and all of their neighbors at the same time. To that purpose, we use a routine which is a slight modification of the Breadth-First Search algorithm (BFS for short), in which we add information on the number of path between the root and each node of the BFS tree (see [6] for details on the classical BFS algorithm). The routine takes as input the graph and a vertex s serving as monitor (the root of the tree) and outputs for each vertex $t \neq s$ the number $\#path(t)$ of shortest paths between s and t as well as the list $L(t)$ of neighbors of t belonging to at least one shortest path between s and t . Thanks to this information, it is straightforward to derive the probability that vertex $x \in L(t)$ lies on a uniformly randomly chosen shortest path from s to t : this probability is $\#path(x)/\#path(t)$.

3.3 Choice of monitors and targets

Since the Neighborhood Flooding method is dedicated only to nodes in the core of the network, all targets will be chosen in the corresponding part of the graph. Here, we do not intend to study the impact of the sample of nodes used to infer the distribution. We are only interested in the ability of the method to discover all neighbors of a given target. This is the reason why we take all nodes as targets and determine for each of them the difference between their observed and actual degree. Nevertheless, for

sake of completeness, we present in Section 6.3 some simulations enabling to determine the order of magnitude of the number of targets that should be used in order to be close enough from the overall distribution.

In real networks, monitors are end-hosts and should therefore be chosen among degree-one nodes of the graph, also called *leaves*. However, since we aim only at targets in the core, the neighbors collected by a leaf monitor are the same as those who would be collected by the core node which is connected to the border tree to which the leaf monitor belongs. Moreover, there is no interest in choosing several monitors in a same tree-like part of the network (a same academic institution for example) as they would see the whole core in the same way. Then, the number of connection points of the set of monitors with the core of the network seems more relevant than the number of monitors itself. For these reasons, and for sake of computational efficiency, we restrict the considered graph to its core and choose monitors directly in the core rather than choosing them among leaves. Indeed, this allows us to save considerable amount of computation time, since the size of the core, for power-law graphs, is significantly lower than the size of the graph itself. Thus we are able to deal with power-law graphs whose core sizes are comparable to those of Poisson graphs we use. Despite the fact that we still call them monitors in the rest of the paper, one should be aware that vertices we choose in the core actually correspond to connection points of monitors with the core rather than monitors themselves. It remains that choosing monitors in the core could favor Neighborhood Flooding compared to choosing them among leaves : we precisely study this impact in Section 6.1.

One question still has to be answered : how should we choose monitors among core nodes? We are interested in determining the ability of the Neighborhood Flooding method to discover all neighbors of targets. One can guess that it must mainly depend on two characteristics of the set of monitors used : its cardinality and its dispersion in the network. Here, we focus on the former question; more precisely, provided that the set of monitors we use is dispersed enough, how does the accuracy of the observation vary with respect to the number of monitors employed? Therefore, we choose monitors uniformly randomly among core nodes. Since the maximal number of monitors we use (800) is small compared to the size of the graph (at least 1 million), we can expect that a uniform random choice must result in a well dispersed set of monitors. However, we do not study this question in details, nor the more general question of the impact of dispersion on the quality of our estimation. Instead, we concentrate on how this quality varies with regard to the number of random monitors used. More precisely, we choose 800 monitors and successively extract subsets of 400, 200, 100,

50, 25 and 12 monitors. As previously, in order to ensure that obtained results do not depend on a particular choice of monitors, for each considered graph, we always perform 3 different choices of the 800 base monitors and one series of extraction for each of these 3 choices. It appears that these random choices are equivalent (see Section 4.3).

The rest of the paper is dedicated to the presentation of the simulation results according to the settings presented above. Each section will focus on a specific parameter of the problem. As said before, we performed each simulation for different sets of graphs and different choices of sets of monitors. However, in order to ease the reading, we will only present one case of study each time. This means that the figures we are presenting in this paper have to be taken as representative of all the simulations dedicated to the corresponding parameter but not as the mean of all the simulations. Each section will then be divided into two parts corresponding to whether we are discussing Poisson graphs or power-law graphs.

4 Quality of the observed distribution

The main question addressed by our paper is whether the Neighborhood Flooding method is able to observe properly degrees of nodes in the core of the network, using a reasonable number of monitors. As an example, the set of monitors used in [7] is spread over about four hundred different institutions in the Internet and authors claimed that it would be possible to extend this number to roughly one thousand using other sets of monitors. Then, it must be clear that if the target has a low number of neighbors, say for instance 2 or 3, we will very likely discover all of them with these monitors. On the contrary, we cannot observe more neighbors than the number of monitors, and more generally if the target has many neighbors we will probably miss some of them. The key question then is to determine the limit between the two situations: how many monitors do we need compared to the degree of the target, or equivalently, given a number of monitors, up to what degrees are we able to capture the whole neighborhood? Answering these questions is the purpose of the present section.

4.1 Poisson graphs

Our first case of study is Poisson graphs. One of the major results on the bias of the classical method for measuring degree distribution was highlighted in [13] by Lakhina *et al.*. They generated random graphs with 100 000 nodes

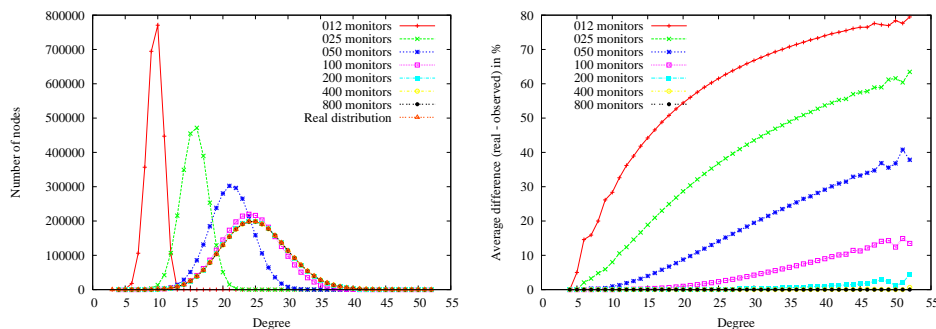


Figure 2: On the left: Degree distributions observed in our simulations for a Poisson graph with 2.5 million nodes and mean degree 25. Each curve corresponds to a different number of monitors used: 12, 25, 50, 100, 200, 400 and 800. On the right: Mean difference (in %) between real and estimated degrees of nodes according to the real degree.

and mean degree 15, having Poisson degree distribution, and merged all together random shortest paths obtained from 10 sources to 1000 destinations. Their results showed a striking fact: the observed degree distribution they obtained is not a Poisson distribution but rather looks like a power-law distribution. This shows that the classical method is intrinsically biased towards heterogeneous distributions.

In our case, we used a Poisson graph of 2.5 million nodes with mean degree 25 and we plot degree distributions observed with different numbers of monitors, from 12 to 800 (Figure 2 left). As one could expect, using 12 monitors, the distribution is poorly estimated : the mean observed degree is lower than 10 (which is close to the number of monitors), while the actual one is 25. This is not surprising, as the number of monitors itself is significantly lower than the mean degree. Nevertheless, it is remarkable that though its parameters appear modified, the observed distribution is still Poisson-like. This shows that the Neighborhood Flooding method does not suffer from the bias towards heterogeneous distributions highlighted in [13] for classical method. Instead, it qualitatively preserves the nature of the distribution, which constitutes a major improvement compared to the classical approach.

Pushing further their analysis, Lakhina *et al.* show that, in the classical method, observed degrees are not correlated to actual degrees in the graph. They plot, for each node in their sample, a point whose coordinate on the x-axis is its actual degree in the graph and its coordinate on the y-axis is its observed degree. Their plot (Figure 5(c) of [13] which is here reproduced

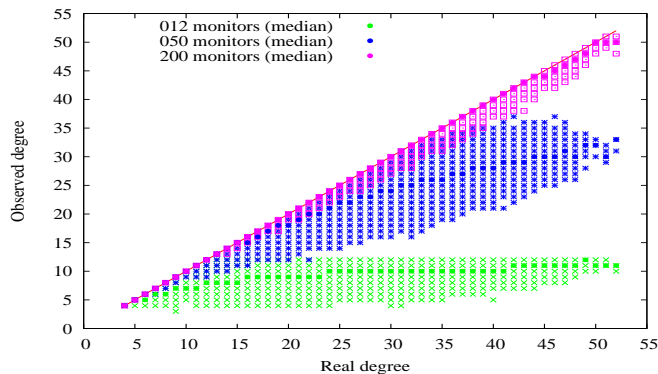


Figure 3: Comparison between observed and real degrees for all nodes of a Poisson graph with 2.5 million nodes and mean degree 25. Each group of points corresponds to a different number of monitors used: 12, 50 and 200. We also plot the median observed degree for each actual degree and each number of monitors (dark line).

in Figure 4) shows that nodes whose degree is at most 20 appear with all possible degrees between 2 and their actual degree. Opposite is the case of nodes whose degree is greater than 20: exceptions apart, they appear only with degree significantly lower than their actual degree. Even more surprising, the median observed degree is 2 for all true degrees between 8 and 21 (recall that mean degree is 15), indicating that the observed degree is very low in average independently of the actual degree of nodes.

This is not the case with Neighborhood Flooding. On Figure 3, we plot the same scatter plot than the one in [13], for 12, 50 and 200 monitors. Let us focus on the results with the 12-monitors set. A striking difference with the plot of [13] is that points are not spread anarchically. They are agglomerated close to the median value, and this median value increases according to the actual degree, at least for degrees lower than 12, which is the number of monitors. Below this value, the median curve is growing almost constantly, then the curve bends and stays stable around a value of 10, close to the number of monitors, which is the maximum degree observable by Neighborhood Flooding. Thus, unlike with the classical approach, observed degree is approximately proportional, at least for low degree nodes, to the actual degree. Moreover, high degree nodes tend to be observed with a degree close to the maximal observable degree (*i.e.* the number of monitors). This constitutes a great improvement over the former situation.




Figure 4: Correlation between observed degree and actual degree for the classical method when using 10 monitors and 1000 targets over 100 000 nodes. This figure is extracted from [13] (Figure 5(c)).

Besides these qualitative results, when the number of monitors grows, the estimation rapidly becomes sharp. With 100 monitors, the observed distribution seems quite close to the actual one, and with 200 monitors and more, the two distributions are visually indistinguishable. This shows that not only the Neighborhood Flooding method is able to qualitatively determine the nature of the distribution, but with a few hundred monitors, it is also able to determine its exact values very precisely. In order to appreciate the accuracy of the observed degree for each node, we plot on Figure 2 (right), the mean difference between real and observed degree (as a percentage of the real degree) for each real degree in the graph. Confirming the first intuition, it suggests that nodes with a low degree, at most four, are very well observed even with a low number of monitors, e.g. 12. The plot shows clearly that the maximal degree properly observed increases when the number of monitors increases. For example, this degree is about 10 with 50 monitors and about 25 with 200 monitors. Even though we can not detect it on the plot of the observed distribution, there is still a gain using more than 200 monitors for nodes of degrees greater than 30. In particular, with 400 monitors and more, the error is close to 0: more than 99.9 % of node degrees are correctly estimated. And it is remarkable that, even for the nodes with the largest degrees, 50 and more, the observation is excellent.

These results indicate therefore that, unlike the classical approach, the Neighborhood Flooding method does not denature the shape of a Poisson distribution, even with a low number of monitors, and the observed degree

is approximately proportional to the actual one. Moreover, at least for small degrees (less than 50) and for graphs of this size, a few hundred monitors (like used nowadays in real-world measurements) is sufficient for estimating very precisely the degree of each node and the underlying Poisson distribution.

4.2 Power-law graphs

We turn now to the case of power-law graphs. Figure 5 presents the results of our simulations on a power-law graph of 10 million nodes with exponent 2.1, which was chosen to fit the exponent observed in [15, 10] (results for different exponents are presented in Section 5.2). Notice that, unlike in the Poisson case where the core of the graph is very close to the whole graph, here, the core involves only about 3.3 million nodes ($\sim 33\%$ of nodes in the graph).

On Figure 5 (top/left), we present the degree distributions observed with different numbers of monitors (inverse cumulative distributions are presented in the bottom/left plot of the figure). In our analysis, we do not take into account nodes observed with degree one, since we know they are mis-observed as there are no such nodes in the core of the network. In particular, we removed them to compute inverse cumulative distributions. Though, we let them appear on the non-cumulative distribution. Be aware that, because of the logarithmic scale, some points appearing to be exactly the same are actually only very close to each other. We can see that, for small degrees, the estimated values of the distribution are close to the real ones but the curves suddenly bend down when they reach a specific threshold for each set of monitors. This is clearly implied by the fact that we cannot observe degrees higher than the number of monitors we use : curves have to stay below these values, which explains the shape of their tail. It remains that, again, similarly to the Poisson case, the power-law shape of the distribution is qualitatively preserved, at least for low degrees, even with a small number of monitors. Furthermore, not only the shape is preserved but the exponent is also well estimated.

As for Poisson graphs we give on Figure 6 a scatter plot where point coordinates are actual and observed degrees of nodes in the network. It shows that for power-law distributions too, the observed degree is related to the actual one. For each number of monitors employed (12, 100 and 800) points are gathered around the median value which is, for actual degrees up to the number of monitors, approximately proportional to the actual degree. As we pointed out earlier, this constitutes a significant improvement over the classical approach, in which the observed degree is not related to the

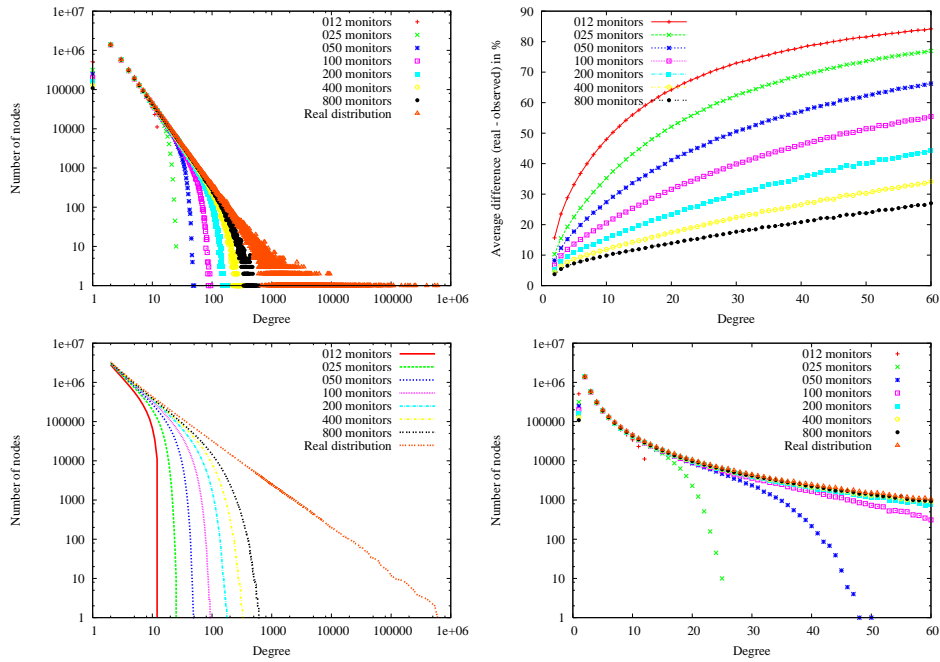


Figure 5: Top/left: degree distributions observed in our simulations for a power-law graph with 10 million nodes and an exponent of 2.1. Each curve corresponds to a different number of monitors used: 25, 50, 100, 200, 400 and 800 and the plot is displayed on a logarithmic scale. Top/right: Mean difference (in %) between real and estimated degrees of nodes according to the real degree. Bottom/left: inverse cumulative distribution for the same setting. Bottom/right: zoom of the degree distribution plot for degrees values less than 60 and displayed on a lin-log scale.

actual one (see [13]).

Concerning the quality of estimation of the distribution, it appears that even with a low number of monitors, the values of the distribution for small degrees are correctly estimated. In order to appreciate this quality more precisely, we plot on Figure 5 (bottom/right) the distributions in lin-log scale³, restricted to degrees less than 60. In order to give an idea of the degree until which the precision of the estimated distribution is satisfying, we use the following criterion. Thanks to the inverse cumulative distribution (Figure 5 bottom/left), we determine the maximum degree d such that the

³Note that the wide range of values of the distribution for these degrees does not allow to visualize it with a linear scale on the vertical axis.




Figure 6: Comparison between observed and real degrees for all nodes of a power-law graph with 10 million nodes and exponent 2.1. Each group of points corresponds to a different number of monitors used: 12, 100 and 800.

estimated number of nodes having degree greater than d is at least 50% of the actual number. This degree d is 5 with 12 monitors, 9 with 25, and goes up to 59 and 90 with respectively 400 and 800 monitors. Until this degree d , the estimation of the distribution (non-cumulative) is quite good. For example, with 12 monitors, the estimated value of the distribution for degree 5 is 90% of the actual value, and with 800 monitors, the estimated value for degree 90 is 85% of the actual one.

However, unlike for Poisson graphs, and even though it is quite good, the observation does not become perfect for degrees less than 50 even with 800 monitors. Moreover, the fact that the distribution is well observed does not come from a very precise observation of node degrees, as shown on Figure 5 (top/right). Indeed, even though 81.4% of nodes having degree at most 10 are properly observed and the mean error on the degree of degree-ten nodes is 1, with 800 monitors, it remains that only 11.5 % of nodes having degree between 10 and 50 have their degree correctly estimated and the mean error for degree 50 is 12. Even if the obtained precision for degrees at most ten is acceptable, it cannot explain the much better precision we get for the estimation of the degree distribution, until degree 50 and more. This is actually due to a self-compensation phenomenon: while some degree- d nodes are estimated with a degree lower than d , in the same time, some nodes having degree greater than d are estimated with degree d , compensating the diminution due to mis-observed degree- d nodes. In particular, this phenomenon is responsible for the fact that observed degree distribution of the core contains degree-one nodes (see Figure 5 top/left and bottom/right)

while we know that there is no such nodes in the core of the graph. Thus, despite the good results given by the Neighborhood Flooding method for observing the distribution of low-degree nodes in power-law graphs, the self-compensation phenomenon should be quantified and controlled in order to give stronger warranty on the precision of the observation.

Besides, Figure 6 shows also that even for a high number of monitors, the precision of the estimation is decreasing when the degree of the nodes increases. For instance with 800 monitors nodes of degree 80, the mean error drops to 35%. Although this quality is clearly related to the number of monitors used (and thus can be overcome by increasing this number), this behavior remains a limitation of the method which should be addressed in the futur.

As a conclusion, the Neighborhood Flooding method gives a correct estimation of the degree distribution of power-law graphs for relatively low degrees, at least for graphs of the considered size and with exponent 2.1. The power-law shape of the distribution is properly observed, as well as its exponent. Nevertheless, these good results partly relies on a self-compensation phenomenon that should be controlled.

4.3 Note on choices of graphs and monitors

In what precedes, for sake of clarity, we presented results on a single graph and with a single choice of monitors (for each number of them). In fact, in order to ensure that our results do not depend on a particular choice of graph or monitors, we conducted all our simulations on three different graphs and using three different choices of monitors for each of them. These variations turned out to have negligible impact on the results we obtained. We show below, as an example, how results vary according to different choices of graphs and monitors. We show independently two situations: three different choices of monitors on a fixed graph, and a single choice of monitors on three different graphs.

Figure 7 shows results for Poisson graphs of 2.5 million nodes and mean degree 25. The left plot gives the mean error for three different choices of monitors, on a fixed graph. Obviously, the three different choices of monitors depicted lead to the same observation of degrees. This is true for 25, 100 and 800 monitors. The slight variations observed on high degree nodes (greater than 45) come from the fact that these nodes are difficult to observe, because of their high degree, and there are only a few of them in the graph. This makes this population of nodes very sensitive to a change in the estimation of the degree of one of them. Nevertheless, for degrees




Figure 7: Mean difference (in %) between real and estimated degree of nodes according to real degree, on Poisson graphs, for different sets of monitors and different graphs. The graphs have 2.5 million nodes and mean degree 25. The left plot shows results for three different choices of monitors on a fixed graph, and the right plot shows results on three different randomly generated graphs.

up to 45, the three curves are almost indistinguishable. As far as choice of graphs is concerned, the right plot of the Figure 7 shows that the situation is quite similar. The three different choices of graphs (with same parameters) give the same curves, indicating that the conclusions obtained on a specific graph remain valid for all random graphs having same parameters.




Figure 8: Mean difference (in %) between real and estimated degree of nodes according to real degree, on power-law graphs, for different sets of monitors and different graphs. The graphs have 10 million nodes and exponent 2.1. The left plot shows results for three different choices of monitors on a fixed graph, and the right plot shows results on three different randomly generated graphs.

Figure 8 shows the same curves for power-law graphs of 10 million nodes (~ 3.3 millions in the core) and exponent 2.1. Again, we come to the conclusion that the observation of degrees is not sensitive to choices of monitors : for each number of monitors (25, 100 and 800), the curves corresponding to the three different choices are almost indistinguishable. On the other hand, the result of the simulation seems more sensitive to choice of the graph. Indeed, on the right plot, the three curves are not exactly the same for the three graphs, and for each number of monitors. Nevertheless, they are very close from each other and the conclusions derived on one of them still hold for the others. Note that this is particularly true for low degrees.

5 Sensitivity to graph parameters

In this section, we explore the sensitivity of the Neighborhood Flooding method to variations of parameters of graph models, namely size and characteristic parameter.

5.1 Impact of the size of the graph

The results presented above were obtained on graphs of large size: 2.5 million nodes for Poisson graphs, and 10 millions (about 3.3 millions after extraction of the core) for power-law graphs. One may however argue that the Internet is much larger than the graphs used here. Clearly, for computational reasons, it is difficult to run simulations on graphs whose size is comparable to the actual Internet. Then, we run simulations on different sizes of graphs in order to study how our previous results are affected by size effect.

Poisson graphs For Poisson graphs, we considered four different sizes: 1 million, 2.5 millions, 5 millions and 10 millions of nodes. We ran simulations with the four different mean degrees in our settings, but gives only results with mean degree 25 (other results are identical). Figure 9 shows results for 1 million nodes (left) and 10 million nodes (right). Both estimations are made with 50, 100 and 200 monitors; 200 being the least number of monitors in our settings such that the distribution is well estimated. It appears that for both graphs (as well as for 2.5-million-node and 5-million-node graphs, omitted here), the estimated distributions are almost identical for each number of monitors. Therefore, it seems that, for Poisson graphs, size has little impact, if any, on the quality of the observation.

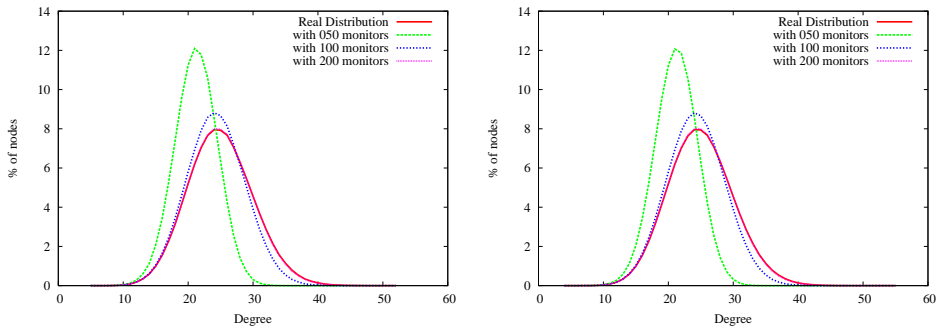


Figure 9: Degree distributions observed in our simulations for different sizes of Poisson graphs with mean degree 25 : 1 million of nodes (left) and 10 million of nodes (right). In each plot, we display the true degree distribution and the one observed with 50, 100 and 200 monitors. The curves for the real distribution and the one observed with 200 monitors are indistinguishable.

Power-law graphs We conducted simulations on three sizes of graphs – 5 millions, 10 millions and 20 millions of nodes – and the four values of exponent used in our settings. Figure 10 presents results only for exponent 2.1 (others are similar). After core extraction, sizes of the restricted graphs are respectively 1.6 millions, 3.3 millions and 6.6 millions. The situation is slightly different than for Poisson graphs. Size effect is perceptible on the observed distributions: when the size of the graph grows, the tail of the observed distribution bends down slightly later but steeper. This means, surprisingly, that the bigger the graph is, the better the distribution is estimated. This difference is slight and can be better appreciated on the second row of Figure 10 (displayed in lin-log scale for degrees between 30 and 60) and on Figure 11 (mean error according to actual degree). This fact may seem counter-intuitive. But taking into account, as suggested by results on Poisson graphs, that the size may not have any impact on the quality of the observation, results on power-law graphs may simply denote a slight modification of the structure of power-law graphs according to their size. Anyway, we can observe, at least for graphs presented here, that size has a very limited impact, and even a positive one if any, on our ability to properly observe the degree distribution.

This is our general conclusion for both types of distributions studied here. This fits the intuition that the key point is the target degree, not the graph size: as long as monitors are well distributed in the graph, the size has a limited impact on our ability to reach the target by all its neighbors, provided

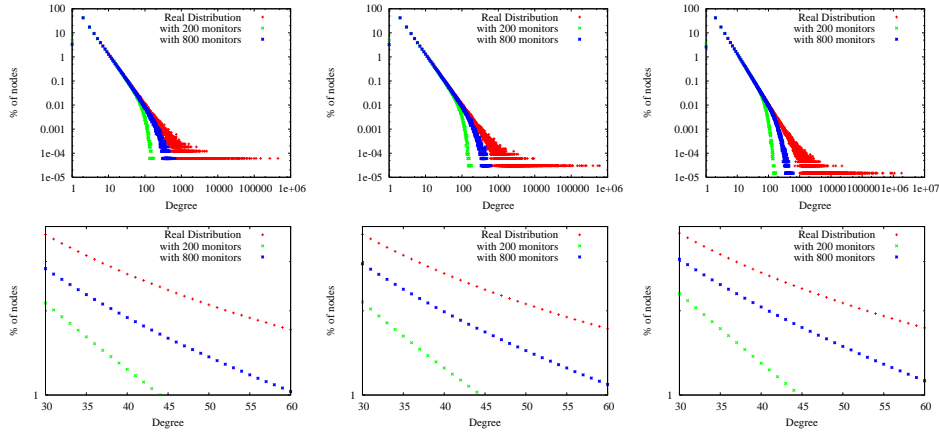


Figure 10: Degree distributions (first row) and inverse cumulative degree distributions (second row) observed in our simulations for different sizes of power-law graphs with exponent 2.1. From left to right 5 million of nodes, 10 million of nodes and 20 million of nodes. In each plot, we display the true degree distribution and the ones observed with 200 and 800 monitors.

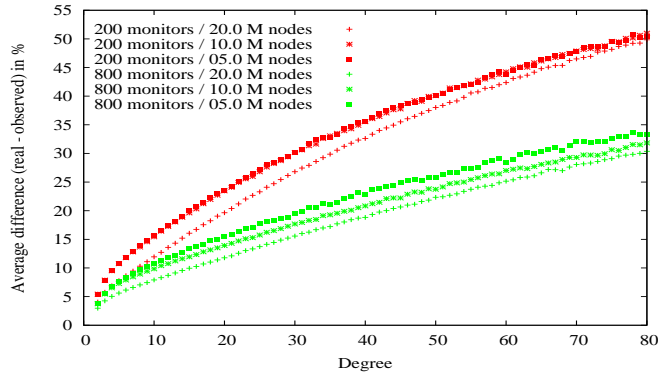


Figure 11: Mean difference (in %) between real and estimated degrees of nodes according to the real degree for power-law graphs with 5, 10 and 20 million nodes and exponent 2.1. Each group of curves corresponds to a different number of monitors used: 200 and 800.

that the number of these neighbors is reasonably small. We consider the formal analysis of this empirical observation as a particularly interesting perspective.

5.2 Sensitivity to characteristic parameter

The results presented so far were obtained on graphs with fixed characteristic parameter: mean degree 25 for Poisson graphs and exponent 2.1 for power-law graphs. In this section, we study how variations of this parameter impact the quality of the observation.




Figure 12: Degree distributions for Poisson graphs with 2.5 million nodes and different mean degree observed with different number of monitors. From left to right: mean degree 6 with 12, 25 and 50 monitors; mean degree 12 with 25, 50 and 100 monitors; mean degree 50 with 50, 100 and 200 monitors;

Poisson graphs We generated Poisson graphs of 2.5 million nodes with various mean degrees: 6, 12, 25 and 50. Results of simulations are depicted on Figure 12 for graphs with mean degree 6, 12 and 50 (results for mean degree 25 can be found in Section 4). Note that the numbers of monitors used are not the same in the three plots. As expected, it shows that the number of monitors required to observe properly the distribution grows when the mean degree of the graph grows. It is remarkable that this fact is due only to the presence of higher degree nodes in the graph: for low degree nodes, the quality of the observation is the same independently from the mean degree. This can be seen on Figure 13: for each number of monitors, the four curves corresponding to the four values of the parameter gives the same estimation quality. The only difference between these four curves stands in that the higher the mean degree is, the further the curve goes in the high degrees. But on their common parts, curves are indistinguishable.

Even if the number of monitors required slightly grows in order to get a precise estimation, as long as the graph has a Poisson degree distribution, the Neighborhood Flooding method remains appropriate to properly observe this distribution. For example, for graphs used here, the distribution is always properly observed with 400 monitors; even for mean degree 50, which is much greater than the assumed mean degree of the actual Internet.




Figure 13: Mean difference (in %) between real and estimated degrees of nodes according to the real degree for Poisson graphs with 2.5 million nodes and different mean degree (6, 12 and 50) with different number of monitors (12, 50 and 200).

Power-law graphs We generated power-law graphs with different exponents: 2.1, 2.2, 2.3 and 2.4. Figure 14 shows results of the observation with 800 monitors (skipping the 2.1 case as it can be found in Section 4). Unlike for Poisson graphs, variations of the characteristic parameter lead to variations in the quality of the observed distribution. When the exponent grows, the observed distribution move slightly away from the actual one. The estimated distribution still remain a power-law, but its exponent is slightly underestimate. Even the low-degree part of the distribution is affected. This is confirmed by Figure 15 (mean error according to actual degree) which shows that the quality of the observation diminishes for all degrees, including lower ones (see degree-ten nodes for example). This suggests that the structure of the graph varies with the exponent leading to the situation that some links attached to low-degree nodes become hardly observable. This may be due to the fact that, for greater exponents, some links in the core are not central in the graph. This notion of centrality can be formally captured by the *betweenness centrality* of links [4], which evaluates the proportion of shortest paths of the graph going through a given link. Studying connections between betweenness centrality of links and performance of the Neighborhood Flooding method, as well as studying the impact of the exponent of the power-law on the distribution of betweenness centrality of links in the network are major perspectives raised by our work. We discuss them further in the general conclusion of the paper.

Nevertheless, note that for all values of the parameter we tested, the

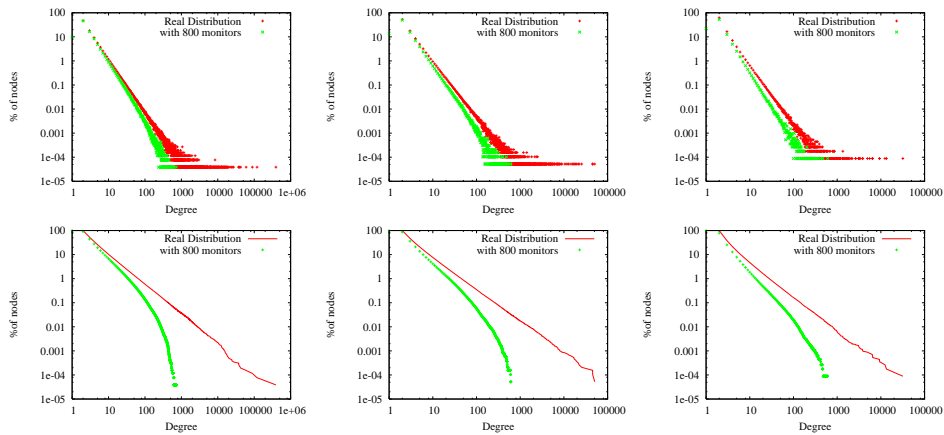


Figure 14: Degree distributions (first row) and inverse cumulative degree distributions (second row) observed in our simulations for different exponent of power-law graphs of 10 million nodes. From left to right exponent 2.2, 2.3 and 2.4. In each plot, we display the true degree distribution and the ones observed 800 monitors.




Figure 15: Mean difference (in %) between real and estimated degrees of nodes according to the real degree for power-law graphs with different exponents (2.2, 2.3 and 2.4) using 800 monitors.

observed distribution stays close to the actual one (see Figure 14 first row) and that it is also true for the slope of the first part of the distribution. Moreover, an interesting feature is that when the exponent grows, the curve of the observed distribution does not bend down as it previously used to do for lower exponents: this can be used, at least for such exponents, to know

whether we estimate properly the distribution.

6 Impact of simulation choices

Since the purpose of the present paper is to investigate the relevance of the Neighborhood Flooding method to measure the Internet degree distribution, it is therefore natural to wonder how choices made for simulations influence the obtained results. We do not intend to argue the relevance of our choices but, instead, to explore whether alternative choices that one could think of could alter the results. More precisely, we study the influence of three modeling issues which have potentially an impact on the behavior of the method: the location of monitors in the network (Section 6.1), the modeling of `traceroute` (Section 6.2) and the use of all core nodes as targets instead of a sample of random core nodes to infer the degree distribution (Section 6.3).

6.1 Choice of monitors on leaves

All simulation results presented before were obtained as follows : the referred graph was first restricted to its core, then monitors were chosen among remaining (core) nodes and the simulation was performed in the restricted graph. Running simulation directly in the core was achieved for computational reasons: it allows us to treat power-law graphs whose size of the core is comparable to the one of the largest Poisson graphs we could handle. However, as explained in Section 3, choosing monitors among leaves of the graph would be more realistic since, in a real environment, monitors are end-hosts. One could even suspect that choosing uniformly among core nodes could favor the Neighborhood Flooding method compared to a uniform choice among leaves.

In order to determine whether this is the case or not, we performed, on several graphs, simulations choosing monitors as leaves and we compared results obtained on the same graph with a choice of monitors in the core. More precisely, we compute the core and the set of leaves of the graph, without restricting it. Then, we chose monitors uniformly randomly among leaves and perform simulations taking all nodes of the network as targets. At the end we keep only observation results of core nodes and we discard their neighbors that are not in the core. We then compare results to those obtained by first restricting the graph and choosing monitors in the core.

This comparison is meaningless for Poisson graphs as the core is the whole graph itself. On Figure 16, we present results obtained for a power-

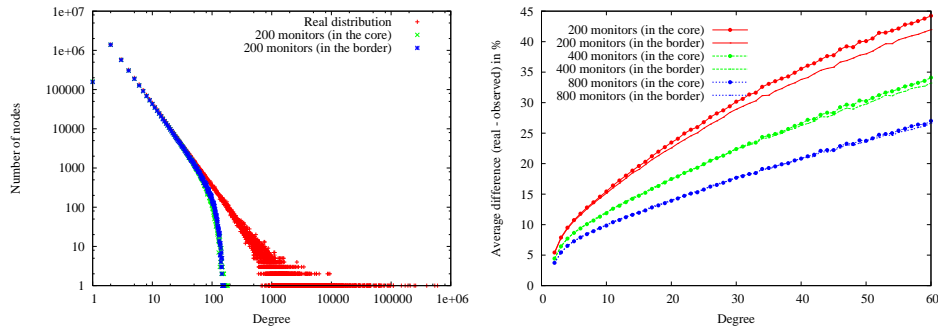


Figure 16: On the left: Degree distributions observed for a power-law graph with 10 million nodes and exponent 2.1 using 200 monitors. On the right: Mean difference (in %) between real and estimated degrees of nodes according to the real degree using 200, 400 and 800 monitors. For each set of monitors, we plot the result obtained when choosing monitors in the border and in the core.

law graph⁴ of 10 million nodes and exponent 2.1 using 200, 400 and 800 monitors. As usual, we plot the mean error according to actual degree in the graph. For each number of monitors, the two curves corresponding to the two different methods are very close to each other. They are almost indistinguishable for 400 and 800 monitors, while for 200 monitors we even obtain a slightly better estimation choosing monitors among leaves. However, the two curves remain very close to each other and the difference between them may not be meaningful (we note, nevertheless, that we observe it in all the simulations we ran). The conclusion of these tests is that choosing monitors in the core is equivalent to choosing them among leaves and do not favor the Neighborhood Flooding method. This reinforces the validity of the results we presented here.

6.2 Different traceroute modelings

As announced in Section 3, in what precedes, a `traceroute` from node A to node B was modeled by a shortest path between A and B chosen uniformly randomly among all possible shortest paths. This choice will be referred to in the following as the *path-uniform* choice. Here, we will not discuss the possibility of modeling `traceroute` by a path different from a

⁴similar results, not presented here, were obtained for the other graphs of our set of 36 power-law graphs, using three different choices of monitors among leaves for each graph.

shortest one, though this question is of course of high interest. Instead, we investigate the question of the method used for randomly choosing (not necessarily uniformly) among all shortest paths. Indeed, the uniform choice is the default one but it has no particular relevance with regard to modeling issues, and one may want to use others. Here, we investigate the impact of some other possible choices on obtained results. Note that we do not argue for the validity of the choices we tested with regard to modeling purposes: we are only interested in quantifying the impact of these choices on simulation results.

For example, because of the specificity of the Neighborhood Flooding method, which uses only the last hop of `traceroute`, we only need to determine the last link of the shortest path. That is, the neighbor of the target lying on this path. Therefore, a simple choice would be to uniformly choose among all neighbors of the target lying on a shortest path; we call it the *neighbor-uniform* choice. Regardless of the relevance of such a choice, we note that, since we aim at discovering all neighbors of the target with various monitors, intuitively, this choice seems to be the most favorable to the Neighborhood Flooding method. Therefore, we include it in our tests as reference.

Another possible way of choosing a shortest path, which we will refer to as the *local-uniform* choice, consists in going from the monitor to the target by determining at each step which shortest path will be followed for next hop. This seems to us a possible way of choosing as it echoes the locality of routing process: each router has to choose a possible neighbor to which it will forward the packets to a given destination. Again, we do not argue for the validity of this choice but only study its impact on simulation results.

The motivation for studying the influence of shortest path choice method is that for a given pair of monitor M and target T , and for some topology of the network, the probabilities to discover a given neighbor of T by a probe sent from M varies considerably according to the chosen method. On the example of Figure 17, it varies from $1/2$ for the neighbor-uniform choice to $(1/2)^n$ for the local-uniform choice, where n is the length of shortest paths from M to T in the given topology. To investigate the resulting impact on quality of observation, we plot on Figure 18 and 19 the results obtained with the three methods mentioned above. Note that while path-uniform choice and neighbor-uniform choice can be computed very efficiently (within the same complexity as a Breadth-First Search), the complexity of the algorithm we devised for the local-uniform choice is higher. This is the reason why results are presented with a set of only 100 monitors and on smaller graphs : respectively 1 million nodes for Poisson graphs and 5 million nodes for




Figure 17: Repartition of different probabilities for choosing randomly a neighbor on the same topology depending on the method used. The (directed) graph corresponds to all possible shortest paths between a monitor M and a target T . Each case gives the probability of discovering the left and the right neighbor of T for a different method. Fig (a) corresponds to the *neighbor-uniform* method. Fig (b) corresponds to the *path-uniform* method. Fig (c) corresponds to the *local-uniform* method.

power-law graphs, restricted to 1.6 millions after core extraction. This size limitation has negligible impact on the simulations as shown in Section 5.1.

Figure 18 (right) shows mean error according to actual degree for a Poisson graph of one million nodes and mean degree 12 observed with 100 monitors. For degrees up to 23, the three curves are almost indistinguishable. For higher degrees, up to 28, they remain very close to each other and no significant differences appear between them. For degrees beyond 28, the behavior of the curves become chaotic because only very few nodes are estimated with such a degree, which makes mean values irrelevant. These observations are confirmed by the left part of Figure 18 which shows that the observed distribution is very similar to the real one. Thus, on Poisson graphs, it seems that the three methods have the same behavior, which is not really surprising considering that the structure of those graphs, very homogeneous, is well balanced.

The situation is slightly different for power-law graphs, as shown on Figure 19 (right). The three curves are still close to each other, especially for low degrees, but it appears that they are slightly different for higher degrees. As one could expect, the most balanced choice for the discovered neighbor (uniform-neighbor choice) gives the best estimation of degrees. It is followed by the uniform-path method, which we used in all results presented in the paper, and finally by the uniform-local choice, which is, as shown on the example of Figure 17, the method most likely to produce unbalanced probabilities on the discovered neighbor of the target. However, we can see

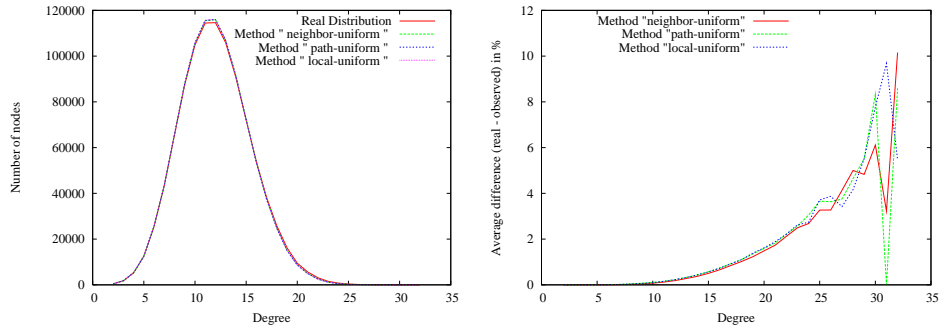


Figure 18: Left: Degree distributions observed for a Poisson graph with 1 million nodes and mean degree 12 using 100 monitors. Each curve corresponds to a different method used for choosing randomly the discovered neighbors. Right: Mean difference (in %) between real and estimated degrees of nodes according to the real degree.

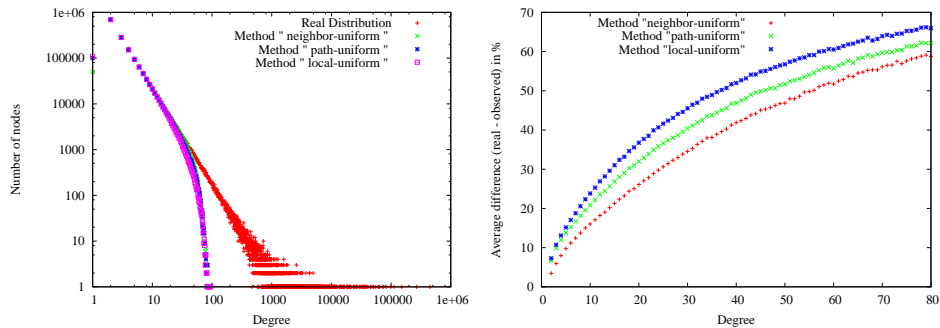


Figure 19: Left: Degree distribution for a power-law graph of 5 million nodes (1.6 millions in the core) with an exponent of 2.10 observed with 100 monitors. Each curve corresponds to a different methods for randomly choosing a neighbor. Right: Mean difference (in %) between real and estimated degrees of nodes according to the real degree.

clearly that the method employed has only limited impact on the quality of the estimation (see Figure 19 left), even for heterogeneous graph structures such as those of power-law graphs.

6.3 Inference of the distribution from a sample of nodes

In simulations presented so far, the set of targets was the whole core of the network. The reason is that we aim at rigorously evaluating the performance of Neighborhood Flooding for observing the degree of any core node. It remains that, in practice [7], the distribution is inferred from a sample of core nodes and one may wonder what size of sample is needed to retrieve the same approximation abilities than those obtained when considering all core nodes as targets. This question is not directly related to Neighborhood Flooding, it is a classical question of statistics. Nevertheless, for sake of completeness, we integrate this parameter in our experimental analysis.

We extracted some samples of core nodes of different sizes and plotted the normalized obtained distributions, both for a Poisson graph and a power-law graph. The left part of Figure 21 presents the results for Poisson graphs and for samples of 1 000, 10 000, 100 000 and 200 000 core nodes. With 1 000 targets, the general shape of the distribution (left part) is retrieved but it appears much disturbed by the fact that proportions in the sample are only an approximation of true ones. Nevertheless, with 10 000 targets, the distribution appears well estimated and the estimation seems perfect with 100 000 nodes, which represents only 2% of core nodes. One may be surprised that degrees whose proportions appears more difficult to estimate are those around the mean. Indeed, since nodes of such degrees are numerous they should be the easiest proportions to be determined by a sampling. This is indeed the case, relatively to their number, which is hidden by the linear scale of the y-axis of the plot. On the logarithmic scale of Figure 20 (right), it appears clearly that, relatively to their number, proportions of nodes whose degree is close to the mean are quite well estimated even with 1 000 targets, and that this estimation is more difficult to be precisely obtained for degrees away from the mean, whose proportion in the graph is low.

Figure 21 depicts results for power-law graphs with 20 million nodes (6.6 after core extraction) when using samples of 2 000, 20 000 and 200 000 nodes. As expected, degrees whose proportion are very low in the core cannot be correctly sampled with a low number of targets. This is the reason why the distribution for high degree nodes is very badly estimated: there are almost no such nodes in the samples used here. However, it appears that the distribution is well estimated up to degree 12 even with only 2 000

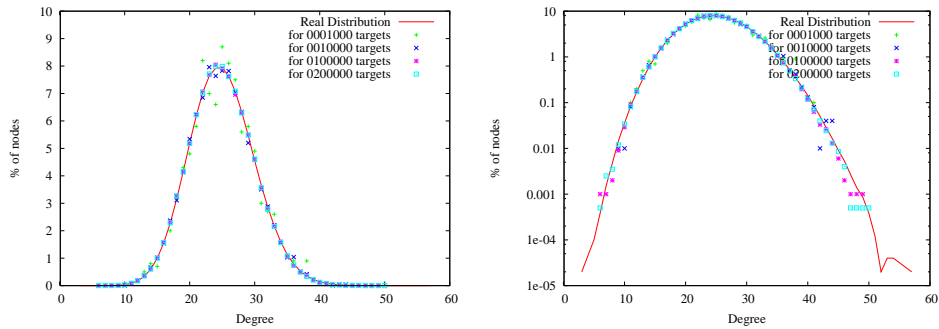


Figure 20: Degree distribution displayed in lin-lin scale (left) and lin-log scale (right) for a Poisson graph with 5 million nodes using 200 monitors. Different curves correspond to different size of samples: 1 000, 10 000, 100 000, 200 000.




Figure 21: Degree distribution for a power-law graph with 20 million nodes (6.6 after core extraction) using 800 monitors and for different sample size. From left to right: 2 000, 20 000, 200 000.

targets. With 20 000 targets (which represents only 0.3 % of core nodes) the first part of the distribution seems properly observed, and in particular, the coefficient of the power-law is correctly approximated. When the number of targets grows, the estimation becomes less sensitive to variations of the proportions in the sample and the curve in log-log scale goes thinner and closer to a straight line. As previously, with a number of targets around 200 000 (about 3% of core nodes) the estimation of the distribution on low degree nodes is quite satisfying.

This shows that, both for Poisson and power-law distributions, even a low number of targets in the core, about 2 – 3%, gives a good view of the actual distribution.

7 Conclusion and perspectives.

The simulations we conducted came to a clear conclusion: the Neighborhood Flooding method is free from the bias introduced by the classical method consisting in merging paths to obtain a sample of the graph, which was highlighted in [13]. Indeed, we showed that Neighborhood Flooding respects the nature of the distribution both for Poisson graphs and power-law graphs, even with a low number of monitors. Moreover, increasing this number to some hundreds (which is the number of monitors used in [7]) leads to a perfect observation of the degrees of all nodes for Poisson graphs having mean degree up to 50, which is far greater than the assumed mean degree of the Internet. For power-law graphs, using such number of monitors leads to a correct estimation of the distribution for exponents slightly greater than 2. When the exponents grows, the quality of the observed distribution remains correct, for reasonably low degrees, but the estimation is less precise. In particular, the exponent of the power-law is slightly underestimated.

Another way of study enabling to appreciate the performance of Neighborhood Flooding would be to investigate the question related to the dispersion of the set of monitors. In particular a criterion for measuring this dispersion would be highly desirable, as well as investigating the impact of this dispersion on the quality of the observation.

For power-law graphs, despite the fact that the distribution of low degrees is correctly estimated, it appears that some nodes, even some having degree less than 10, are not perfectly observed with 800 monitors. This fact is surprising, especially compared to the situation for Poisson graphs. It suggests that some links in the core of a power-law graph are not really central in the graph, which is opposite to what we expect for the core of a graph. Classically, this notion of centrality is formally captured by the betweenness centrality of links [4], which evaluates the proportion of shortest paths of the graph going through a given link. Therefore, studying the distribution of the betweenness centrality of links in the core of random power-law graphs would certainly lead to great insight on this extensively used class of graphs.

In a more practical perspective, the reader might have noticed that the results presented in this paper strongly rely on the non dynamicity of the graphs. Indeed we performed the simulations on networks whose topology do not change over time. This is of course a strong assumption. To this regard, the current results will hold true as long as the measure of the network is close to an atomic operation compared to its dynamicity. Another way to state this remark is that when gathering monitors and targets and sending probes, one should be carefull in doing so in as fast as possible. This is a

concern that any implementation of the method should address.

Our work also raises several key questions which may be tackled formally. In particular, one may formally establish the results we obtained on the ability of the Neighborhood Flooding method to discover the neighbors of nodes in Poisson and power-law graphs, or the limited impact of the size of the underlying graph, which we empirically observe here. We let that for further investigations.

References

- [1] Dimitris Achlioptas, Aaron Clauset, David Kempe, and Cristopher Moore. On the bias of traceroute sampling: or, power-law degree distributions in regular graphs. *J. ACM*, 56(4), 2009.
- [2] W. Aiello, F. Chung, , and L. Lu. A random graph model for massive graphs. In *32nd Annual Symposium in Theory of Computing (STOC)*, pages 171–180, 2000.
- [3] Paul Barford, Azer Bestavros, John W. Byers, and Mark Crovella. On the marginal utility of network topology measurements. In Vern Paxson, editor, *Internet Measurement Workshop*, pages 5–17. ACM, 2001.
- [4] U. Brandes. A faster algorithm for betweenness centrality. *Journal of Mathematical Sociology*, 25(2):163–177, 2001.
- [5] Caida, macroscopic topology measurement projects. <http://www.caida.org/projects/macroscopic/>, 2010.
- [6] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, second edition, 2001.
- [7] Christophe Crespelle, Matthieu Latapy, and Élie Rotenberg. Rigorous measurement of ip-level neighborhood of internet core routers. In *2nd IEEE International Workshop on Network Science and Communication Networks (NetSciCom'10)*, 2010. To appear.
- [8] Luca Dall’Asta, J. Ignacio Alvarez-Hamelin, Alain Barrat, Alexei Vázquez, and Alessandro Vespignani. Exploring networks with traceroute-like probes: Theory and simulations. *Theor. Comput. Sci.*, 355(1):6–24, 2006.
- [9] P. Erdős and A. Rényi. On the evolution of random graphs. *Math. Inst. Hung. Acad. Sci.*, 5:1761, 1960.

- [10] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, 1999.
- [11] Jean-Loup Guillaume and Matthieu Latapy. Relevance of massively distributed explorations of the internet topology: simulation results. In *INFOCOM*, pages 1084–1094. IEEE, 2005.
- [12] Jean-Loup Guillaume, Matthieu Latapy, and Damien Magoni. Relevance of massively distributed explorations of the internet topology: Qualitative results. *Computer Networks*, 50(16):3197–3224, 2006.
- [13] Anukool Lakhina, John W. Byers, Mark Crovella, and Peng Xie. Sampling biases in ip topology measurements. In *INFOCOM*, pages 332–341, 2003.
- [14] Matthieu Latapy and Clémence Magnien. Complex network measurements: Estimating the relevance of observed properties. In *INFOCOM*, pages 1660–1668. IEEE, 2008.
- [15] Jean-Jacques Pansiot and Dominique Grad. On routes and multicast trees in the internet. *SIGCOMM Comput. Commun. Rev.*, 28(1):41–50, 1998.
- [16] Fabien Viger and Matthieu Latapy. Efficient and simple generation of random simple connected graphs with prescribed degree sequence. In *Computing and Combinatorics, 11th Annual International Conference (COCOON)*, volume 3595 of *Lecture Notes in Computer Science*, pages 440–449, 2005.
- [17] Walter Willinger, David Alderson, and John C. Doyle. Mathematics and the internet: A source of enormous confusion and great potential. *Notices of the AMS*, 56(5):586–599, May 2009.