Rigorous Measurement of the Internet Degree Distribution

Matthieu Latapy, Elie Rotenberg, Christophe Crespelle, Fabien Tarissan

► To cite this version:

Matthieu Latapy, Elie Rotenberg, Christophe Crespelle, Fabien Tarissan. Rigorous Measurement of the Internet Degree Distribution. *Complex Systems*, 2017, 26 (1). hal-01500371

HAL Id: hal-01500371

<https://hal.science/hal-01500371>

Submitted on 3 Apr 2017

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.

Rigorous Measurement of the Internet Degree Distribution

Matthieu Latapy¹, Élie Rotenberg^{1,2}, Christophe Crespelle², Fabien Tarissan¹

Abstract

The degree distribution of the internet, i.e. the fraction of routers with k links for any k , is its most studied property. It has a crucial influence on network robustness, spreading phenomena, and protocol design. In practice, however, this distribution is observed on partial, biased and erroneous maps. This raises serious concerns about the true knowledge we actually have of this key property. Here, we design and run a drastically new measurement approach for the reliable estimation of the degree distribution of the internet, without resorting to any map. It consists in sampling random core routers and precisely estimating their degree with probes sent from many monitors scattered over the internet. Our measurement shows that the true degree distribution significantly differs from classical assumptions: it is heterogeneous but it decreases sharply, in a way incompatible with a heavy-tailed power law.

1 Introduction

The internet has become a crucial infrastructure sustaining our social, economic, cultural and scientific lives at both local and worldwide scales. Despite this, our understanding of its structure remains very limited. To gain more insight, the internet is often modeled as a network where nodes represent routers and links represent direct connections between them (wires, satellites, etc). The degree distribution (*i.e.* for each integer k , the fraction p_k of nodes having k links) of this network is particularly important: it plays a key role for resilience to failures and attacks [3, 20], cascade and spreading phenomena [6, 24], as well as protocol and network design [2, 13]. As a consequence, it is an essential building block of most modern models of the internet [15, 16, 29, 31].

However, current knowledge of this degree distribution is far from satisfactory, and it is at the core of a lively scientific controversy [1, 5, 17, 18, 21, 22, 27, 32]. Indeed, the degree distribution is known only from internet maps obtained through intricate measurement procedures giving partial, biased and erroneous views. These measurements generally rely on the use of the *traceroute* tool which provides in principle a route in the network from the monitor running the measurement to a given target. By collecting and merging many such routes, one obtains a map of the internet. See Figure 1 (left) for an illustration. However, the *traceroute* tool is prone to numerous errors [12, 22, 30, 33] and, most importantly, the procedure itself is intrinsically biased [1, 10, 14, 18].

For all these reasons, much effort is devoted to the design of more accurate internet measurement tools and to the collection of larger and larger maps [8, 9, 11, 19, 23, 26, 28, 30]. However, as measurement capabilities remain limited and as the internet evolves faster than our ability to measure it, this may very well be a dead-end.

¹Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR 7606, 4 place Jussieu 75005 Paris. First-name.Lastname@lip6.fr

²Université Claude Bernard Lyon 1, DANTE/INRIA, LIP UMR CNRS 5668, ENS de Lyon, Université de Lyon

We present here a drastically new approach able to reliably estimate the degree distribution of the internet *without resorting to any map*. We probe randomly chosen routers from monitors scattered all over the internet and obtain an accurate estimate of their degree. We infer from these degrees a rigorous estimate of the internet degree distribution, far more reliable than previous knowledge. See Figure 1 (middle and right) for an illustration. This methodological shift raises challenging questions, that we address here. We conclude that, contrary to what most current studies assume, the degree distribution is heterogeneous but is not a heavy-tailed power law.

Figure 1: Comparison of our method to the classical traceroute method. Left: a traceroute measurement from 1 monitor (square node) towards 25 targets (bullet nodes). This measurement needs 97 probes. Middle: measurement with our method from 9 monitors (square nodes) towards 10 targets (bullet nodes). Figure 2 details how the links of each target are discovered. This measurement needs 90 probes. Right: the true degree distribution of the network together with the estimates obtained by both methods.

Figure 2: Measurement of the degree of 10 targets using our method. We display 10 copies of the network, one for each target measurement. On each copy we show the routes followed by the probes sent from our 9 monitors (square nodes) toward the corresponding target (bullet node).

2 Our measurement method

A machine in the internet (a router or an end-host) may have several interfaces, each corresponding to a connection to a neighbor machine. Each interface has its own address, and the degree of a router is nothing but its number of interfaces/addresses.

Let us consider an address t , which we call *target*, and let us denote by $r(t)$ the node (router or end-host) to which t belongs. Internet protocol specifications [4, 7] state that when a monitor m sends a packet to destination t on an unallocated port, then $r(t)$ should answer to m with an error packet (ICMP Destination Unreachable, Code 3/Port unreachable). An important detail is

that the source of this error packet is in principle the address of the interface i by which $r(t)$ sent it, see Figure 3.

Figure 3: Monitor m sends a packet to destination address t on an unallocated port; the node $r(t)$ answers with an error packet with source address i , and thus m discovers interface i of $r(t)$.

Let us temporarily assume that $r(t)$ implements this feature correctly (we handle other cases in Section 3). Now consider a set M of monitors which all send such a probe towards address t . If for each interface i of $r(t)$ there is a monitor m in M to which $r(t)$ answers using i , then one obtains the set of *all* interfaces of $r(t)$, and so its degree. This constitutes our basic measurement primitive: 1) from each monitor of a set M , we send a packet to an unallocated port of target address t and 2) we collect the set $M(t)$ of all addresses used by $r(t)$ to answer to monitors in M .

Depending on the target t and on the set of monitors M this measurement primitive may succeed or fail to discover all the interfaces of $r(t)$. In particular, one has to distinguish between two very different kinds of targets: 1) the target node $r(t)$ is in the *core* internet (Figure 4, middle) or 2) the target node $r(t)$ is in the *border* (Figure 4, right).

Figure 4: Left: the core and the border of the network; the border is the set of all trees connected to the network, the core is the part remaining when one removes these trees. Middle: a set of monitors (the squared nodes) send probes towards a target address t belonging to a core router $r(t)$ and obtain its four core interfaces of $r(t)$. Right: the same monitors send probes towards another target t' belonging to a border router $r(t')$ and miss most interfaces of $r(t')$.

As illustrated in Figure 4 (right), when the target address belongs to a border node our measurement primitive may miss many of its interfaces, and most likely discovers only the interface directed towards the core. The situation regarding core interfaces of core routers is quite different (see Figure 4 middle). Indeed, such interfaces route traffic toward a non-negligible part of the internet, and one may therefore expect that a reasonably large and well distributed set M of monitors discovers them. Of course, this highly depends on the considered set of monitors, and we explain in supplementary material how to assess the quality of a monitor set in practice.

We focus here on the core, which is the key part of the network: it performs the non-trivial routing of packets from one point to another point of the network, while the border is a set of trees connected to this core, where packets are just forwarded upward or downward the tree (see

Figure 4, left). We therefore discard target addresses which belong to border nodes, see Section 3.

In summary, we expect a good enough set of monitors M to be able to discover all or almost all core interfaces of any core router, leading to an estimate of its degree in the core internet. Once this measurement primitive is implemented, one may use it to observe the degree of all targets in a set T . If T is a set of core routers sampled uniformly at random (which means that all core routers have the same probability to appear, independently from their degree), then, the distribution of degrees observed for T is an estimate of the degree distribution of core routers.

3 Measurement

We present in this section practical measurements we conducted following our approach. We describe the whole procedure step by step, as well as the obtained dataset.

We first built an initial **target set** by sending (from a machine in our lab) a probe to addresses corresponding to 32 bit integers sampled uniformly at random. We stopped this process when we obtained correct answers (*i.e.* ICMP Destination Unreachable (Code 3/Port unreachable) error packets) from 3 million such targets (we considered that no answer would arrive after 1 minute). This took approximately 10 hours.

Our initial **monitor set** was composed of the approximately 700 machines of the PlanetLab platform [9], which is a distributed infrastructure provided to researchers to conduct network experiments. Some of these potential monitors are of little interest (they have very poor connections, for instance, or they belong to networks that filter our probes) and some are colocated, therefore providing redundant information in our measurement. However, we present in the supplementary material several assessments of this monitor set, which all show that it fits our needs.

Given these initial target and monitor sets, we uploaded the target set to each monitor and remotely asked them to **send probes** to all targets (in a random order to avoid situations where targets would receive many probes in a short period of time). This lasted approximately 4 hours (and so each target received at most 700 probes during this period, which is a reasonable load). In order to explore the stability of our measurements, we repeated this operation three times in a row. The whole measurement (building the target set and probing each of them from each monitor three times) took less than 24 hours, with a very reasonable load for targets and monitors. At this stage, we obtained for each target its answers to the probes from all monitors (repeated three times), which we gathered onto a local machine for analysis.

We then applied a drastic **filtering** process (detailed in supplementary material) in order to ensure we kept only data relevant to our needs: we removed monitors and targets that behaved incorrectly, as well as border nodes. We also conducted an auxiliary measurement able to obtain the set of all border interfaces visibles from our monitors. Thanks to this, we were able to keep only target addresses that were core interfaces of core routers answering correctly to our probes. Unsurprisingly, most target addresses belonged to border nodes. We finally obtained for each of our three measurements approximately 5 600 targets belonging to reliable core routers. The output of our measurements is the observed degree of these routers, from which we will estimate the degree distribution of internet core routers.

We provide our measurement tools (source code and documentation) as well as the raw dataset at <http://rmidd.complexnetworks.fr>

4 Unbiased estimation

Based on the procedure above, we can achieve the crucial point of our method, namely estimating the degree of core routers sampled uniformly at random. Note that until now, we only sampled uniformly at random the *addresses* of their interfaces, not core routers themselves. Indeed, one has k possibilities to sample a router with k interfaces, so high-degree routers appear in our target list with probability higher than low-degree ones. In order to correct this bias *a posteriori*, we discard from the result of the measurement the core routers whose address t present in the target set turns out to be the address of one of their border interfaces. After this discarding step, the probability that a core router has been sampled is proportional to its number k of core interfaces (which is precisely what we measure). Then, the observed fraction p'_k of routers of core degree k sampled with this bias is proportional to k times the fraction p_k of routers of core degree k sampled uniformly at random: $p'_k \sim k \cdot p_k$. As a consequence, we obtain:

$$p_k = \frac{p'_k}{k} \cdot \frac{1}{\sum_{i>1} \frac{p'_i}{i}}$$

where the second term is nothing but a normalization constant to ensure that $\sum_k p_k = 1$.

We then use this formula to estimate the true degree distribution p_k from the observed one p'_k .

5 Obtained degree distribution

The degree distributions observed from our three measurements after bias correction following the formula above are given in Figure 5 (left). We plot the inverse cumulative distributions in Figure 5 (right).

Figure 5: Left: the degree distributions obtained from our three measurements (after bias correction): for each degree k , we give the estimated fraction p_k of core routers with degree k . Right: plot of the inverse cumulative degree distributions obtained from our three measurements, after bias correction: for each value x on the horizontal axis, we plot the fraction of core routers having degree higher than or equal to x (log-log scale). We also plot the power law of exponent $\alpha = 3.8$ to show that obtained distributions are incompatible with a power law of exponent lower than this.

First notice that the results of all measurements are very similar, which confirms that our results are stable in this setup. We present in supplementary material several other assessments of the quality of our final observation, all confirming that the obtained distributions are good approximations of the true one.

Obtained distributions show clearly that low-degree core routers are prevalent: approximately 75% of them have degree 2 only, and almost 20% have degree 3. This is not surprising, as we observe core interfaces only: these routers certainly have other interfaces connected to border routers and/or end-hosts. The number of interfaces they use to actually *route* traffic in the core internet, however, is very low.

On the other hand, some core routers have much larger degrees, and the highest one we observe is 29. We may possibly miss a few interfaces of this router but there is little chance that the true largest degree is much higher: we perform measurements from a much larger number of monitors and so the fact that observed degrees are bounded by this number plays no role. Of course, core routers with degree significantly higher than 29 may exist, and they probably do. There is however none in our random target set and we therefore expect them to be extremely rare (which is reinforced by the sampling bias towards high-degree routers explained in Section 4).

Going further, we observe that the first values of the obtained distribution (p_k for $k < 10$) are reasonably well fitted by a power law (a straight line in a log-log plot of the distribution). After that, the distribution experiences a sharp decrease. The first values are the ones that our method estimates best, and so one may ask if the obtained distribution is compatible with a power law. As highest degree may be under-estimated, this may even be in accordance with the shape of the whole obtained distribution.

In order to explore this question, we compute a lower bound α for power law exponents compatible with the first values (the most reliable ones). It is the slope of a straight line fitting the distribution in log-log scale. The exponent would clearly be larger than $\alpha = 3.8$, see Figure 5 (right), which discards the usual assumption of an exponent close to 2. This also shows that *if the true degree distribution is a power law*, it is hardly distinguishable from an exponential decrease in practice [25] even for a system the size of the internet.

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):1–28, 2009.
- [2] Aditya Akella, Shuchi Chawla, Arvind Kannan, and Srinivasan Seshan. On the scaling of congestion in the internet graph. *SIGCOMM Comput. Commun. Rev.*, 34(3):43–56, July 2004.
- [3] Réka Albert, Hawoong Jeong, and Albert-László Barabási. The internet’s Achilles’ heel: Error and attack tolerance of complex networks. *Nature*, 406:378–382, 2000.
- [4] Fred Baker. Requirements for IP Version 4 Routers, 1995.
- [5] Paul Barford, Azer Bestavros, John W. Byers, and Mark Crovella. On the marginal utility of network topology measurements. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, IMW ’01, pages 5–17, 2001.

- [6] Alain Barrat, Marc Barthélemy, and Alessandro Vespignani. *Dynamical Processes on Complex Networks*. Cambridge University Press, New York, NY, USA, 1st edition, 2008.
- [7] Robert Braden. Requirements for Internet Hosts - Communication Layers, 1989.
- [8] CAIDA. Caida, macroscopic topology measurement projects. <http://www.caida.org/projects/macroscopic/>, 2015.
- [9] PlanetLab Consortium. Planetlab: An open platform for developing, deploying and accessing planetary-scale services. <http://www.planet-lab.org/>, 2009.
- [10] 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.
- [11] Benoit Donnet and Timur Friedman. Internet topology discovery: A survey. *IEEE Communications Surveys and Tutorials*, 9(1-4):56–69, 2007.
- [12] Benoit Donnet, Matthew J. Luckie, Pascal Mérindol, and Jean-Jacques Pansiot. Revealing MPLS tunnels obscured from traceroute. *Computer Communication Review*, 42(2):87–93, 2012.
- [13] Christos Gkantsidis, Milena Mihail, and Amin Saberi. Conductance and congestion in power law graphs. In *Proceedings of the 2003 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, volume 31 of *SIGMETRICS '03*, pages 148–159, New York, NY, USA, 2003. ACM.
- [14] 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.
- [15] Hamed Haddadi, Miguel Rio, Gianluca Iannaccone, Andrew W. Moore, and Richard Mortier. Network topologies: Inference, modeling, and generation. *IEEE Communications Surveys and Tutorials*, 10(1-4):48–69, 2008.
- [16] Hamed Haddadi, Steve Uhlig, Andrew W. Moore, Richard Mortier, and Miguel Rio. Modeling internet topology dynamics. *Computer Communication Review*, 38(2):65–68, 2008.
- [17] Balachander Krishnamurthy, Walter Willinger, Phillipa Gill, and Martin F. Arlitt. A socratic method for validation of measurement-based networking research. *Computer Communications*, 34(1):43–53, 2011.
- [18] Anukool Lakhina, John W. Byers, Mark Crovella, and Peng Xie. Sampling biases in IP topology measurements. In *Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 1 of *INFOCOM'03*, pages 332–341, 2003.
- [19] Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas E. Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iplane: An information plane for distributed services. In *7th Symposium on Operating Systems Design and Implementation (OSDI'06)*, November 6-8, Seattle, WA, USA, pages 367–380, 2006.

- [20] Clémence Magnien, Matthieu Latapy, and Jean-Loup Guillaume. Impact of random failures and attacks on poisson and power-law random networks. *ACM Comput. Surv.*, 43(3):13:1–13:31, April 2011.
- [21] Aniket Mahanti, Niklas Carlsson, Anirban Mahanti, Martin F. Arlitt, and Carey Williamson. A tale of the tails: Power-laws in internet measurements. *IEEE Network*, 20(1):68–88, 2013.
- [22] Pascal Mérindol, Benoit Donnet, Olivier Bonaventure, and Jean-Jacques Pansiot. On the impact of layer-2 on node degree distribution. In *Proceedings of the 10th ACM SIGCOMM Internet Measurement Conference, IMC 2010, Melbourne, Australia - November 1-3, 2010*, pages 179–191, 2010.
- [23] Reza Motamedi, Reza Rejaie, and Walter Willinger. A survey of techniques for internet topology discovery. *IEEE Communications Surveys and Tutorials*, 17(2):1044–1065, 2015.
- [24] Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic spreading in scale-free networks. *Phys. Rev. Lett.*, 86(14):3200–3203, 2001.
- [25] Richard Perline. Strong, weak and false inverse power laws. *Statistical Science*, 20(1), 2005.
- [26] RIPE-NCC. Ripe atlas. <https://atlas.ripe.net>.
- [27] Matthew Roughan, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 10 lessons from 10 years of measuring and modeling the internet’s autonomous systems. *IEEE Journal on Selected Areas in Communications*, 29(9), 2011.
- [28] Yuval Shavitt and Eran Shir. DIMES: let the internet measure itself. *Computer Communication Review*, 35(5):1810–1821, 2005.
- [29] Fabien Tarissan, Bruno Quoitin, Pascal Mérindol, Benoit Donnet, Jean-Jacques Pansiot, and Matthieu Latapy. Towards a bipartite graph modeling of the internet topology. *Computer Networks*, 57(11):71–74, 2013.
- [30] Fabien Viger, Brice Augustin, Xavier Cuvellier, Clémence Magnien, Matthieu Latapy, Timur Friedman, and Renata Teixeira. Detection, understanding, and prevention of traceroute measurement artifacts. *Computer Networks*, 52(5), 2008.
- [31] Xiaoming Wang and Dmitri Loguinov. Understanding and modeling the internet topology: economics and evolution perspective. *IEEE/ACM Trans. Netw.*, 18(1):998–1018, 2010.
- [32] 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.
- [33] Yu Zhang, Ricardo V. Oliveira, Yangyang Wang, Shen Su, Baobao Zhang, Jun Bi, Hongli Zhang, and Lixia Zhang. A framework to quantify the pitfalls of using traceroute in as-level topology measurement. *IEEE Journal on Selected Areas in Communications*, 29(9):1822–1836, 2011.

– Supplementary material –
for
**Rigorous Measurement
of the Internet Degree Distribution**

Matthieu Latapy², Élie Rotenberg^{1,2}, Christophe Crespelle², Fabien Tarissan¹

6 Proof of concept

In order to assess the relevance of our approach, we conducted a comprehensive set of simulations which we present in this section. Assuming that we are able to build appropriate sets of monitors and targets, the key questions we want to answer are: what is the risk that our estimate of a node's degree is different from its true degree, and how many monitors do we need to have an accurate estimate of the degree distribution?

To investigate this, we have conducted simulations as follows (see [6] for more details): we considered different kinds of artificial graphs to represent the network; we used as monitors random nodes with degree one (representing end-hosts); and we used *all* core targets (*i.e.* nodes in the graph obtained by iteratively removing degree-one nodes). We then assumed that each target answers to probes from each monitor using one (randomly chosen) of its interfaces that starts a shortest path from the target to the monitor. We used two different kinds of graphs: one with Poisson degree distribution, which is a typical homogeneous distribution, and one with a power law degree distribution, which is a typical heterogeneous distribution. These two kinds of distributions are considered as extreme cases for what the true degree distribution may be.

Figure 6 shows the results of the simulations for Poisson and power law graphs of 2.5 million nodes. Figure 6(a) presents the degree distribution observed with respectively 12, 25, 50, 100, 200, 400 and 800 monitors. As one could expect, with 12 monitors the degree distribution is poorly estimated in the two cases. Nevertheless, it is remarkable that, even with this poor level of quality, the nature of the distribution (*i.e.* homogeneous or heterogeneous) appears clearly. When the number of monitors grows, so does the quality of the observed degree distribution.

With 200 monitors in particular, the observed and the true distributions become visually indistinguishable in the homogeneous case (left). For the heterogeneous case (right), one can observe a cut-off for very large degrees. This comes from a limitation of our method: the observed degree cannot exceed the number of monitors, and more generally, the estimate becomes inaccurate for targets whose degree is close to the number of monitors. On the other hand, for reasonably low-degree targets, up to approximately 20, the observed distribution and the true one are visually indistinguishable with 200 monitors.

These last statements are strengthened by the plots on Figure 6(b) which shows the scatter plot of true degree (on the x-axis) and observed degree (on the y-axis) for all targets and for the two kinds of graphs. We can see that with 200 monitors, the estimate degree of each node is quite close to its true degree for the Poisson graphs, thus proving that our method performs very well on

²Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR 7606, 4 place Jussieu 75005 Paris. First-name.Lastname@lip6.fr

²Université Claude Bernard Lyon 1, DANTE/INRIA, LIP UMR CNRS 5668, ENS de Lyon, Université de Lyon

(a) Observed degree distribution

(b) Scatter plot of the true degree vs. the observed degree

Figure 6: Simulations with different number of monitors (12, 25, 50, 100, 200, 400 and 800) over graphs of 2.5×10^6 nodes whose degree distribution follows either a Poisson law with average degree 25 or a power law with exponent 2.1.

this kind of graphs. As regards power law graphs, we can see that using 200 monitors, the estimate degree of low-degree nodes is quite close to the true one. More than 95% of degree-2 nodes are correctly observed and this proportion drops to 85% when considering all nodes whose degree is lower than 10. This shows that, for this type of nodes at least, our method performs also very well on power law graphs.

Therefore, the only limitation of our method in this theoretical setup seems to be the estimation of the degree of high-degree nodes in power law graphs. Indeed, an intrinsic limitation of our method is that we cannot obtain a degree estimate larger than the number of monitor $|M|$. However, this limitation has to be put in perspective as Figure 6(b) shows that, even if poorly estimated, they still cannot be confused with low-degree nodes. Whatever the number of monitors, the worst estimation (lower point on the y-axis) grows as the true degree grows.

In conclusion, both for Poisson graphs and power law graphs, the nature and the shape of

the degree distribution are correctly observed even with a low number of monitors. In addition, the observed distribution quickly converges to the true one when the number of monitors grows. The true degree of low-degree nodes is correctly observed (also true for high-degree nodes in the homogeneous case), and a high-degree node is never observed as a low-degree node.

One may wonder if these results still hold for graphs of different sizes and with different parameters, average degree for Poisson graphs and exponent for power law graphs. These questions were investigated in [6], as well as the influence of some other parameters of the simulations. It turns out that the conclusions we derive here are still valid for different sizes and parameters. In particular, [6] shows that the size of the graph has very little importance, if any, for the quality of the observation with a given number of monitors. Then, the conclusion obtained by simulations on graphs of a few millions of nodes still holds for graphs of the size of the internet.

7 Comparison with traceroute measurements

In this section we deepen the comparison between our method and the classical traceroute method, with regard to two criteria: the correctness of observed degree distribution and the load induced on the network by the measurement (number of probes sent). We simulate measurements with our method like in Section 6 but using only a restricted set of targets. To simulate traceroute measurements, we follow the method of [19]: we give to each link a weight $1 + \epsilon$, where ϵ is uniformly randomly chosen in $[-1/n, 1/n]$, which ensures with very high probability that there is a unique shortest path between any two nodes. Then, for each monitor we compute the shortest path tree from this monitor to all the other nodes of the network using Dijkstra's algorithm. From these trees, we extract the set of shortest paths from all monitors to all targets in the target list and we aggregate all these paths together into one graph which is the map resulting from the traceroute measurement, and on which the degree distribution is then observed.

We present here the results for two graphs on 5 million nodes of the same kind as those of Section 6: a Poisson graph of average degree 25 and a power law graph of exponent 2.1. For clear comparison, we use the same set of monitors for both methods, composed of 200 monitors for Poisson graphs and 800 monitors for power law graphs. Our method always uses 5 000 targets, which is close to the number of correct core routers in the real-world measurement presented in Section 3. We simulate traceroute measurements with various numbers of targets, resulting in different number of probes sent on the network (from the same number of probes as our method, to a number 2 000 times larger).

Figure 7 shows results for the Poisson graph. First notice that our method accurately estimates the true degree distribution, while traceroute with the same number of probes completely fails. The distribution obtained with traceroute even looks closer to a power law distribution than to a Poisson distribution, which is a known bias of the traceroute method [1, 19]. Using 10 times more probes does not significantly improve this situation. Only when using 500 times more probes than the number used by our method, the distribution observed by traceroute starts to look like a Poisson distribution, even though it remains far from the true one. Still, even with a load 2 000 times larger, i.e. probing about 40% of all nodes of the network, which is hardly possible to achieve in practice, the degree distribution observed by traceroute is clearly distinct from the true one.

Figure 7 right explains this situation: it gives the average observed degree (y-axis) for nodes of given true degree (x-axis). The average degree observed by our method is almost indistinguishable from the true degree. Instead, degrees observed by traceroute measurement with up to ten times

Figure 7: Comparison between our method and traceroute on a 5 million nodes Poisson graph of average degree 25. We use 200 monitors and the comparison is done with regard to the number of probes sent. Our method requires 1 000 000 probes for 5 000 targets (similarly to our real-world measurement). We compare it to the traceroute methods when it is allowed to send the same number of probes (which results in 986 targets) and up to 2 000 times more probes (which results in 1 972 000 targets). Left: true degree distribution and estimates obtained by both methods. Right: average observed degree (y-axis) as a function of true degree (x-axis).

more probes is barely correlated to the true degree: the average observed degree remains very low almost independently of the true degree. When the number of probes grows, the situation gradually improves, but even with a load 2 000 times larger than the one of our method, traceroute still is less accurate than our method, therefore explaining that the distribution itself is not correctly observed.

Going further, let us mention that we pushed the number of targets used by traceroute up to 90% of all nodes. For this huge value only, which is infeasible in practice, the traceroute method performs as well as our method: average error made on observed degree of a node is 0.03 (0.04 with our method) and 97% of all nodes have their degree perfectly measured (96% with our method). With this number of targets, traceroute uses a number of probes 4 500 times larger than our method (and the same number of monitors).

Figure 8 shows results for the power law graph. Our method observes the degree distribution accurately for degrees up to 60, whereas traceroute obtains much poorer estimate even with up to 100 times more probes. However, when traceroute is allowed 500 times more probes than our method, it obtains a better estimate, which is visually almost perfect. Surprisingly, using even more probes then reduces the quality of the estimate, which finally becomes less accurate than our method. This is explained by Figure 8 (top-right): even when traceroute obtains a good estimate of the distribution, this is not the consequence of an accurate estimate of individual node degrees. Therefore the good performances of the traceroute method is for some specific values of the number of targets only, and it is a side-effect of its own bias. One cannot rely on such artifacts to properly estimate the distribution.

To deepen this, we show in Figure 8 (bottom-left) the converse statistics: for each observed degree (x-axis) we plot the average true degree (y-axis) of nodes that were observed with this degree. Figure 8 (bottom-right) gives the ratio between the average true degree and the observed

Figure 8: Comparison between our method and traceroute on a 5 million nodes power law graph of exponent 2.1. We use 800 monitors and the comparison is done with regard to the number of probes sent. Our method requires 4 000 000 probes for 5 000 targets (similarly to our real-world measurement). We compare it to the traceroute methods when it is allowed to send the same number of probes (which results in 1 602 targets) and up to 2 000 times more probes (which results in 3 204 000 targets). Top left: true degree distribution and estimates obtained by both methods. Top right: average observed degree (y-axis) as a function of true degree (x-axis). Bottom left: average true degree (y-axis) as a function of observed degree (x-axis). Bottom right: ratio between average true degree and observed degree (y-axis) as a function of observed degree (x-axis).

degree. For traceroute with a load 500 times higher than our method, this shows that nodes of a given degree in the observed distribution have a true degree which is in average 2.5 to 3 times higher: nodes observed with degree 5 have in average a true degree above 12, nodes observed with degree 10 have in average a true degree above 25, etc. Our methods performs much better: the ratio between the average true degree and the observed degree remains close to 1 for most nodes, in particular for those having degree up to 20.

When traceroute uses more than 500 times as many probes as our method, the obtained degrees become more accurate, see Figure 8 (top-right). However, the accuracy of the observed distribution decreases at the same time. This is due to the fact that, even with 64% of the nodes as targets (i.e. a number of probes 2 000 times larger than our method), the quality of degree estimates

remains poor, and the ratio between average true degree and observed degree remains close to 2 for most observed degrees. Then, nodes of observed degree d with traceroute have a very different true degree. This is yet another demonstration of the issues of traceroute for degree distribution estimation.

Back to the comparison between the two methods, let us mention that even when traceroute targets 90% of all nodes (*i.e.* 4 500 000 nodes here) it does not reach the accuracy of our method (although both use the same monitors, and our method uses 5 000 targets only). For nodes of degree at most 60 (which represents 98% of all nodes), the average error made by traceroute on the degree of individual nodes is 1.52 and it perfectly measures the degree of 55% of these nodes. The average error for our method is 0.74 only and it perfectly measures the degree of 71% of these nodes.

8 Core vs border

Intuitively, the border of the Internet is the part of the network made of all trees connected to the rest of the network, which is called the core. More formally, the core, also called 2-core in graph theory, may be defined as follows. Consider the pruning process that iteratively removes all nodes of the network having degree exactly one, until there remains no such node. Border routers are the nodes removed during this process when it is applied to the physical internet graph, while core routers are the nodes that remain when the process terminates.

Then, by definition, core routers necessarily have more than one interface linking them to other core routers, and we call such interfaces *core interfaces*, their other interfaces being called *border interfaces*. On the other hand, note that any border node has exactly one interface directed toward the core of the network, namely the one which is linked to its unique neighbour when it is removed from the network during the pruning process. We also call *core interface* this unique interface of a border node and we call *border interfaces* all its other interfaces. The *core degree* (resp. *border degree*) of a node is its number of core (resp. border) interfaces.

8.1 Distinguishing between core and border interfaces

Beside our main measurement, we also conduct an auxiliary measurement in order to obtain for each monitor m the set of border interfaces it may see. To that purpose, monitor m iteratively sends k packets to k random addresses (for a given integer k) with increasing TTLs: the first k packets are sent with TTL 1, the k next packets with TTL 2, and so on. Thanks to the ICMP Time-Exceeded packets issued by the nodes at distance d from m (we discuss below the case of machines that do not send such packets), for each value d of the TTL m discovers a set of interfaces at distance d from m . We denote this set of interfaces by $I_d(m)$. Let us denote by $\delta(m)$ the smallest d such that $|I_d(m)| > 1$, that is the first TTL at which m discovers more than one interface. We have by definition $|I_{\delta(m)}(m)| > 1$ and $|I_j(m)| = 1$ for all $j < \delta(m)$. The set of border interfaces visible from m is then precisely $\bigcup_{j < \delta(m)} I_j(m)$. All other interfaces visible from m are core interfaces belonging to some core router or belonging to some border node which is not on the path between m and the core of the network. Proceeding in this way for all monitors, we build the set $B(M) = \bigcup_{m \in M} \bigcup_{j < \delta(m)} I_j(m)$ of all border addresses that can be seen from them. Consequently, for each interface seen in the measurement we are able to determine whether it is a core interface

or a border interface: it is in the core if and only if it is not in the set $B(M)$ of border interfaces visible from M .

8.2 Recognizing core routers

Figure 9: If we target an interface i that belongs to a border router $r(i)$ then our measurements may see more than one interface for $r(i)$, here two. However, only one of them does not belong to $B(M)$, as displayed in this picture where all interfaces of $B(M)$ are marked with a small dash.

Once we can distinguish between core and border interfaces, we can also distinguish between core and border routers. If a target address t belongs to a border node $r(t)$, our measurements are likely to see only one interface of $r(t)$. Though, in some cases, we may see more than just one interface, see Figure 9. Indeed, $r(t)$ may be a router on the route between some monitor $m \in M$ and the core of the network. In this case, our measurement will also discover the border interface i of $r(t)$ which is directed toward m . By definition, this interface i belongs to the set $B(M)$ of border interfaces that are visible from the set M of monitors (see Section 8.1 above). The key point here is that, if $r(t)$ is a border router then, from what precedes, it follows that our measurements see only one interface not in $B(M)$ for $r(t)$. On the other hand, if $r(t)$ is a core router, our measurement will discover at least two core interfaces of $r(t)$ (provided that M is of sufficient quality), which does not belong to $B(M)$ by definition.

Thus, in the result of our measurement, we are able to distinguish which answering addresses belong to a core router and which of them belong to a border router. And from Section 8.1, we can also determine for each answering address whether it is the address of a core interface or a border interface. This plays a key role for our unbiased estimation, detailed in next section.

9 Uniformly sampling core routers

Being able to sample a core router uniformly at random in the internet³ is at the core of our approach. Unfortunately, there is no direct way to do so. Instead, it is straightforward to get addresses uniformly at random, as they are nothing but 32 bit integers. Of course, sampling such a random integer does not necessarily give a relevant address with regards to our measurement needs: this address may for instance be unallocated or belong to an end-host or a router that does not answer our probes.

³Recall that *uniformly at random* means that all possible elements are sampled with the same probability.

In this section, we first show how to sample uniformly at random an interface of an internet core router that correctly answers our probes, which we call a *correct core router*. From this sampling, which is not a *uniform* sampling of core routers themselves but only of their interfaces, we rigourously deduce an estimate of the degree distribution of all internet core routers. In other words, from the observed distribution resulting from a uniform sample of the interfaces of core routers, we deduce what is the observed distribution resulting from a uniform sample of the core routers themselves.

The general scheme of the way we proceed is as follows. We first perform a measurement on a target list uniformly sampled at random among 32-bits integer (see Section 3). Afterwards, we select into this target list the addresses of core routers correctly answering our probes, which we determine thanks to the result of the measurement. Then, we restrict these results to the set of target interfaces that belong to correct core routers. Finally, we use the results for this set of targets only to infer the degree distribution of Internet core routers.

For the rest of this section, we assume that given an address t , we are able to decide whether t belongs to an host that correctly answer our probes. We show how to do so in Section 10 below. From Section 8.2, we are also able to decide on the result of the measurement whether a given address belong to a core router or not. Consequently, extracting from our uniformly randomly generated target list the addresses that belong to a host that correctly answers our probes and that is a core router, we obtain a uniform sample of the interfaces of correct core routers.

This is not enough for our goal as we need a uniform sample of (correct) core routers themselves, not just of their interfaces, as it turns out that when interfaces are uniformly sampled, routers are not. Indeed, one has k possibilities to sample a router with k interfaces, so high-degree routers appear in our target list with probability higher than low-degree ones. This introduces a bias in the sampling of routers that one can correct if one knows for each router its number of interfaces. Unfortunately, our measurement does not provide this information but instead gives the number of core interfaces of each router (provided that the set of monitors is of sufficient quality).

In order to correct the bias on the number of interfaces of routers introduced by our target selection method, we introduce a supplementary bias in this method: we discard all the target addresses that are not core interfaces of core routers. We are able to do so as we can distinguish between core and border interfaces (see Section 8.1). Then the number of possible addresses to select a router so that it will still be in the target list after this last discarding step is no longer its number of interfaces but instead its number of core interfaces. The great benefit here is that, since our measurement determines the number of core interfaces of each core router, we are now able to correct the bias introduced by this target selection procedure.

The observed fraction p'_k of routers of core degree k sampled with this bias is proportional to k times the fraction p_k of routers of core degree k sampled uniformly at random: $p'_k \sim k \cdot p_k$. As a consequence, we obtain:

$$p_k = \frac{p'_k}{k} \cdot \frac{1}{\sum_{i>1} \frac{p'_i}{i}}$$

where the second term is nothing but a normalization constant to ensure that $\sum_k p_k = 1$. We may therefore use this formula to infer the true degree distribution p_k from the observed one p'_k .

In summary, our method to build target sets is as follows. We sample random 32 bit integers and we select the addresses that are core interfaces of core routers that correctly answer our probes. This procedure and the result of its application on our sample measurement is further detailed and

discussed in Section 10 below.

10 Data filtering and processing

In this section, we describe step by step the way we process the raw data obtained from our measurement, containing some irrelevant or unappropriate data, in order to extract from it the part we use to faithfully estimate the degree distribution of Internet core routers. The key numbers encountered during the different steps of this processing are summarized in Table 1.

		1-st	2-nd	3-rd
Raw data	Nb running monitors	619	625	622
	Nb answering targets	2 849 740	2 734 548	2 699 642
Step 1	Nb targets giving multiple answers	10 150	9 842	11 048
Step 2	Nb monitors receiving answers from $\leq 80\%$ of targets	198	183	180
	Nb targets answering to $\leq 80\%$ of monitors	590 605	527 346	544 252
Step 3	Nb interfaces in $B(M)$	1 040	1 107	1 097
	Nb targets having ≤ 1 interface not in $B(M)$	2 842 481	2 727 422	2 692 135
Step 4	Nb targets t such that $t \notin M(t)$	2 634 226	2 519 320	2 484 483
Processed data	Final number of monitors	421	442	442
	Final number of targets	5 593	5 623	5 619

Table 1: Key post-processing steps for our three measurements.

Step 0. Reserved addresses. As explained in Section 3, before the measurement starts, we build the list of targets by sampling uniformly at random addresses corresponding to 32 bit integers and by keeping the first 3 millions of these addresses that answered to the probe we send to each of them from a single monitor. For sake of completeness, let us mention that we actually apply one additional filter at this step: if the address sampled at random belongs to a known class of reserved addresses [10], then we simply discard it and pick another one at random. Thus, in the measurement itself, all the targets we use do not belong to such a reserved class of addresses (and they correctly answered to the monitor we use in this step).

All the subsequent filters are based on the result of the measurement and are therefore applied afterwards. The numbers of targets and monitors they apply to are given in Table 1.

Step 1. Targets giving multiple answers. Some targets in our list behaved incorrectly: they sent several answers to a unique probe sent by one monitor. As these targets do not behave correctly with regard to our measurement primitive, we simply discarded them and kept only those that sent a single answer to the probe of each monitor. The number of discarded targets is given in Table 1.

Step 2. Targets and monitors with only few answers. Some targets answered to a few monitors only, probably because of shutdowns during measurements, very low ICMP rate limiting, or other specific reasons. Conversely, some monitors received surprisingly few answers, probably due to a very poor local connections, shutdowns, or to the fact that PlanetLab machines may be overloaded (they are shared by numerous users). We plot these numbers in Figure 10, which shows that most monitors received answers from most targets, as we expected. In practice, we discarded

monitors that received answers to less than 80% of their probes, and conversely all targets that sent answers to less than 80% of monitors. See numbers in Table 1.

Figure 10: Left: for each number x on the horizontal axis, we plot the number y of targets that sent at least x answers to our probes. Right: for each number x on the horizontal axis, we plot the number y of monitors that received at least y answers to our probes.

Step 3. Recognizing core routers. The aim of this filtering step is to select only the addresses of the target list that belong to correct core routers, with the method presented in Section 8: i) we build the set $B(M)$ of the border interfaces visible from our set M of monitors (see Section 8.1) and ii) we keep only the target addresses t such that the set of interfaces $M(t)$ observed for t contains at least two interfaces that do not belong to $B(M)$ (see Section 8.2). The number of interfaces in $B(M)$ as well as the number of targets filtered, i.e. that do not satisfy Condition ii), are given in Table 1.

Step 4. Uniform sampling of core routers. In order to correct the bias due to the fact that we uniformly sample interfaces instead of uniformly sampling routers, we perform a supplementary filter. This filter consists in discarding all addresses of the target list that are not addresses of core interfaces. The effect of this filter is to replace the bias mentioned above by another one that we can rigourously correct (see Section 9). Note that this filtering step is independant of Step 3 above: they can be perform in any order, and even simultaneously, on the dataset. Table 1 gives the number of addresses in the target list that are filtered at Step 4, independently of Step 3.

It must be clear that a core router r may give incorrect answers to our probes. In particular, r may give no answer at all, or it may always answer using the same interface independently of the monitor⁴. In the former case, there is only very little chance that an address of r is in our target list as we target only addresses that answered to one probe some hours before. Nevertheless, it may still happen that a router behaves this way during our measurement and in this case, it will be removed from the target list at Step 2. In the latter case, where the router r always answers using

⁴Of course, more intricate behaviors are also possible, but they are very unlikely [17] and we ignore them here.

the same interface independently of the monitor, it will be filtered at Step 4. On the opposite, if an address t of a correct core router r is in our target list, then our measurement sees at least two of the interfaces of r (as long as monitors are reasonably well distributed) and therefore t will successfully pass all filters. Then, our filtering process is satisfying in the sense that it is able to distinguish between correct core routers and other core routers.

Finally, note that there is no reason to assume that the degree of core routers is correlated to whether they answer correctly to our probes or not. Indeed, low-degree core routers may *a priori* misbehave as well as high-degree ones, and conversely. As a consequence, the degree distribution of correct core routers, which we estimate here, is the same as the degree distribution of all core routers.

11 Quality of the monitor set

Our method relies on the use of a large set M of monitors scattered over the internet. It is crucial that this set is large enough since the accuracy of the estimation of the degrees of targets highly depends on this number (see Section 6). On the other hand, having several monitors in the same location has limited interest: it is probable that most targets use the same interface to answer probes coming from these monitors (see Figure 11). Assessing the quality of a given set M of monitors (regarding our measurement goals) is therefore crucial, and we propose here three different and complementary approaches to do so.

11.1 Colocated monitors

Figure 11: Three monitors, m_1 , m_2 and m_3 are actually colocated, and therefore they may observe a unique interface for any given target router $r(i)$. They are redundant regarding the quality of the measurement.

When a packet sent from one monitor m , which is an end-host, goes through the core of Internet, by definition of the core and the border (see Figure 4 left), it always enter the core through the same router, which we call the *branching point* of m . Thanks to the auxiliary measurement method described in Section 8.1, any monitor may identify its branching point: the unique interface in $I_{\delta(m)-1}(m)$ is the (unique) interface of this branching point which is directed toward m .

Now, let us consider two monitors m and m' such that $I_{\delta(m)}(m) = I_{\delta(m')}(m')$. In other words, the first time m and m' see several interfaces they see the exact same ones. Then certainly having both m and m' in the monitor set has little interest for our measurements: m and m' enter in the core internet through very close routers (probably through the same branching point, see Figure 11).

We say that such monitors are *colocated*. The number of non-colocated monitors in M is a key value for estimating the quality of M : it basically represents the number of significantly different locations hosting monitors in M .

In the analysis above, we ignored machines that do not send ICMP Time-Exceeded packets. Because of them, we may erroneously decide that some monitors are colocated; this means that we under-estimate the quality of our monitor set, which has no important consequence in our context: the quality is only under-estimated. Similarly, it is possible that two monitors m and m' have different branching points but satisfy $I_{\delta(m)}(m) = I_{\delta(m')}(m')$. Again, this would make us under-estimate the quality of the monitor set and therefore we may safely ignore this. Conversely, some monitors m and m' may have different but similar sets $I_{\delta(m)}(m)$ and $I_{\delta(m')}(m')$, indicating that they are not colocated but located close from each other. It may be interesting to use this for a more subtle assessment of the level of distribution of monitors, but we leave this for further work.

We used the method we just described and the auxiliary measurement described in Section 8.1 to identify classes of colocated monitors, which provide basically redundant information. We obtained 203 different classes, each containing in average 2.11 monitors. This is consistent with the fact that each institution involved in PlanetLab often contributes with several monitors located at the same place. Examination of the DNS names of monitors belonging to a same class confirmed this: they typically match the same *.domain.tld pattern.

11.2 Diversity of views

In the approach above, we estimate an intrinsic quality of a monitor set M as the number of different locations hosting a monitor. A complementary view is obtained by evaluating the quality of a measurement from M towards targets in a set T . For instance, one may evaluate the quality of M as the number of distinct interfaces observed from M : $Q_0(M) = \sum_{t \in T} |M(t)|$. Clearly, if $Q_0(M') > Q_0(M)$ then M' may be considered as better than M . More subtle quality functions may be defined. In particular, it is interesting to take into account the fact that interfaces of low-degree routers are easier to observe than the ones of high-degree routers. This leads to the quality function $Q_1(M) = \sum_{t \in T} |M(t)|d(t)$ where $d(t)$ stands for the degree of target router $r(t)$. Of course we do not have the value of $d(t)$ and approximate it using the results of our measurements.

Given a quality function Q like the ones above, one may assess the impact of the addition of a new monitor m to the current monitor set, by calculating $Q(M)$ and $Q(M \cup \{m\})$. Ideally, one wants to maximize Q to collect the most accurate set of observed interfaces while keeping M as small as possible to prevent redundant measurements (which may be costly).

In practice, we want to assess a given monitor set M , and to do so we start from an empty monitor set and compute the expected quality improvement when monitors are added one by one, in a random order. The quality is expected to grow with the number of monitors, and then to reach a steady or almost steady regime meaning that adding more monitors would not improve the measurement significantly. Of course, if many monitors are colocated (for instance, if they are all at the same location), the quality will have precisely this behavior (as adding more monitors at the same location does not significantly improve the measurement). This is why this quality function approach is *complementary* to the colocation-based one: we perform first the colocation and then plot the behavior of the quality function when non-colocated monitors are added.

More precisely, once colocated monitors are identified, we proceed as follows: we first estimate the quality of the monitor set when only one colocation class is used, then two colocation classes, etc,

until all colocation classes (and thus all monitors) are used. We add colocation classes in a random order and average the obtained quality. The result is displayed in Figure 12 (left). As expected, for both quality functions, the quality grows sharply at the beginning and rapidly converges. This indicates that adding more monitors at more locations would not improve the results much, and so that our monitor set and the number of locations hosting them are reasonable.

Figure 12: Left: evolution of the quality of the monitor set when we add colocation classes. Right: ratio of the observed fraction $p_k(x)$ of routers of degree k with x colocation classes over the final obtained value p_k (with all classes).

11.3 Convergence of observations

Last but not least, a clear way to assess the quality of a given monitor set regarding our measurements objectives is to directly observe how the estimated fraction p_k of routers of degree k converges when the number of monitors grows, for all k . Here again, we expect these fractions to converge rapidly to a steady value, which is our final estimate. This would indicate that the last monitors we added were not necessary, and thus that we obtain an accurate view. For the same reasons as above, this is complementary to colocation analysis.

In order to examine the impact of adding more monitors at more locations on the estimated fraction p_k of core routers with degree k (which is what we are interested in), we proceed as follows: we add colocation classes one by one like above and observe how p_k evolves. Results are depicted in Figure 12 (right). The estimates for small degrees rapidly converges, which was expected as only few monitors (and locations) are needed to correctly estimate them. Interestingly, only very few locations (approximately 10) are needed to obtain an estimate of p_k for $k < 5$ with a 80% precision. Increasing the number of monitors rapidly improves the quality of the estimate. Even for large degrees, the estimate rapidly reaches a value comparable to the final one, despite the fact that it only slowly converges after that.

11.4 Conclusion

Finally, this work on the monitor set shows that we have around 200 significantly different locations hosting monitors, and that this is sufficient to ensure a reasonable quality for our results. It is clear

however that increasing the number of monitors and the number of locations hosting them would improve both accuracy and reliability of our estimates.

12 Related work

The physical and IP-level internet structures are extensively studied since the seminal papers of Pansiot *et al.* [23] and Faloutsos *et al.* [9]. The most classical approach consists in building maps from traceroute-like measurements. However, several studies have shown that obtained maps are intrinsically biased [1, 7, 13, 19, 20, 22, 25, 29], and even that traceroute outputs are unreliable [8, 25, 28]. The hope that increasing the size and quality of maps would overcome these issues has led to much effort, but the situation remains far from satisfactory [2, 20, 29].

Conducting precise measurements of the degree of random nodes to obtain a reliable estimate of the degree distribution was first suggested in [19]. We explored the possibility to do so at IP level in [5] but we only partly succeeded and we conducted thorough simulations in [6]. Property-driven network measurement are also developed in other contexts, in particular Online Social Networks (OSNs) [11, 18] and P2P overlay measurements [27].

Our work is also closely related to alias resolution (which plays a key role in the building of maps): while we seek all (unknown) interfaces of a given router identified by one of its interfaces, alias resolution aims at identifying in a given set of interfaces the ones that belong to a same router [12, 14, 15, 17]. Probes similar to ours are used in this context, in particular by the *iffinder* tool [16], as well as other techniques. Our use of such probes was clearly inspired by these works.

Finally, important efforts are devoted to the deployment of large and distributed measurements infrastructures, which are crucial for this field of research [3, 4, 21, 24, 26]. Some of them distribute monitoring capabilities at a huge scale (typically onto thousands of end-hosts) and so are particularly promising for extending the work we present here [24, 26].

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):1–28, 2009.
- [2] Paul Barford, Azer Bestavros, John W. Byers, and Mark Crovella. On the marginal utility of network topology measurements. In *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, IMW ’01, pages 5–17, 2001.
- [3] CAIDA. Caida, macroscopic topology measurement projects. <http://www.caida.org/projects/macroscopic/>, 2015.
- [4] PlanetLab Consortium. Planetlab: An open platform for developing, deploying and accessing planetary-scale services. <http://www.planet-lab.org/>, 2009.
- [5] Christophe Crespelle, Matthieu Latapy, and Élie Rotenberg. Rigorous measurement of IP-level neighborhood of internet core routers. In *Second International Workshop on Network Science for Communication Networks*, NetSciCom’10, pages 499–504, 2010.

- [6] Christophe Crespelle and Fabien Tarissan. Evaluation of a new method for measuring the internet degree distribution: Simulation results. *Computer Communications*, 34(5):635–648, 2011.
- [7] 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.
- [8] Benoit Donnet, Matthew J. Luckie, Pascal Mérindol, and Jean-Jacques Pansiot. Revealing MPLS tunnels obscured from traceroute. *Computer Communication Review*, 42(2):87–93, 2012.
- [9] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, volume 29 of *SIGCOMM ’99*, pages 251–262. ACM, 1999.
- [10] E. Gerich. Guidelines for Management of IP Address Space. RFC 1466 (Informational), 1993.
- [11] Minas Gjoka, Maciej Kurant, Carter T Butts, and Athina Markopoulou. Walking in facebook: A case study of unbiased sampling of osns. In *Proceedings of the 29th Conference on Information Communications*, INFOCOM’10, pages 2498–2506. IEEE, 2010.
- [12] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for internet map discovery. In *Proceedings of the 19th Conference on Information Communications*, volume 3 of *INFOCOM’00*, pages 1371–1380, 2000.
- [13] 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.
- [14] Mehmet Hadi Gunes and Kamil Saraç. Resolving IP aliases in building traceroute-based internet maps. *IEEE/ACM Trans. Netw.*, 17(6):1738–1751, 2009.
- [15] M.H. Gunes and K. Sarac. Importance of IP alias resolution in sampling internet topologies. In *IEEE Global Internet Symposium*, pages 19–24, 2007.
- [16] B. Huffaker, D. Plummer, D. Moore, and K. Claffy. Topology discovery by active probing. In *Symposium on Applications and the Internet*, SAINT, pages 90–96, 2002.
- [17] Ken Keys. Internet-scale IP alias resolution techniques. *ACM SIGCOMM Computer Communication Review*, 40(1):50–55, 2010.
- [18] Maciej Kurant, Athina Markopoulou, and Patrick Thiran. Towards unbiased BFS sampling. *IEEE Journal on Selected Areas in Communications*, 29(9):1799–1809, 2011.
- [19] Anukool Lakhina, John W. Byers, Mark Crovella, and Peng Xie. Sampling biases in IP topology measurements. In *Twenty-Second Annual Joint Conference of the IEEE Computer and Communications Societies*, volume 1 of *INFOCOM’03*, pages 332–341, 2003.

- [20] Matthieu Latapy and Clémence Magnien. Complex network measurements: Estimating the relevance of observed properties. In *Proceedings of the 27th Conference on Computer Communications*, INFOCOM'08, pages 1660–1668, 2008.
- [21] Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas E. Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iplane: An information plane for distributed services. In *7th Symposium on Operating Systems Design and Implementation (OSDI'06)*, November 6-8, Seattle, WA, USA, pages 367–380, 2006.
- [22] Pascal Mérindol, Benoit Donnet, Olivier Bonaventure, and Jean-Jacques Pansiot. On the impact of layer-2 on node degree distribution. In *Proceedings of the 10th ACM SIGCOMM Internet Measurement Conference, IMC 2010, Melbourne, Australia - November 1-3, 2010*, pages 179–191, 2010.
- [23] Jean-Jacques Pansiot and Dominique Grad. On routes and multicast trees in the internet. *SIGCOMM Comput. Commun. Rev.*, 28(1):41–50, 1998.
- [24] RIPE-NCC. Ripe atlas. <https://atlas.ripe.net>.
- [25] Matthew Roughan, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 10 lessons from 10 years of measuring and modeling the internet's autonomous systems. *IEEE Journal on Selected Areas in Communications*, 29(9), 2011.
- [26] Yuval Shavitt and Eran Shir. DIMES: let the internet measure itself. *Computer Communication Review*, 35(5):1810–1821, 2005.
- [27] Daniel Stutzbach, Reza Rejaie, Nick Duffield, Subhabrata Sen, and Walter Willinger. On unbiased sampling for unstructured peer-to-peer networks. *IEEE/ACM Trans. Netw.*, 17(2):377–390, 2009.
- [28] Fabien Viger, Brice Augustin, Xavier Cuvellier, Clémence Magnien, Matthieu Latapy, Timur Friedman, and Renata Teixeira. Detection, understanding, and prevention of traceroute measurement artifacts. *Computer Networks*, 52(5), 2008.
- [29] 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.