Sampling of Networks with Traceroute-Like Probes

Alain Barrat^a Ignacio Alvarez-Hamelin^a
Luca Dall'Asta^a Alexei Vázquez^b
Alessandro Vespignani^{a, c}

^aLaboratoire de Physique Théorique (UMR 8627 du CNRS), Université de Paris-Sud, Orsay, France;

^bNieuwland Science Hall, University of Notre Dame, Notre Dame, Ind., and

^cSchool of Informatics and Department of Physics, Indiana University, Bloomington, Ind., USA

Key Words

Network sampling • Traceroute • Internet exploration • Topology inference

Abstract

A large part of the recent development of the interest in complex networks has been triggered by the observation of particular characteristics of real world networks, such as the small-world properties or the heavy-tailed distributions of degrees. Many datasets are, however, the result of an incomplete sampling of the underlying real networks, and it has been argued that sampling procedures might introduce uncontrolled biases in the statistical properties of the sampled graph. In this paper, we explore this issue in the case of the Internet, which is generally mapped from a limited set of sources by using traceroute-like probes. The origin of the biases introduced by such a sampling process is investigated

Simplexus

Despite the pervasive influence of today's Internet, for business and science, government, and just about every other area of human activity, it is not at all easy to get a good snapshot of what the network looks like; that is, of the pattern by which its computers are 'wired' together by telephone lines, satellite links and so on. The Internet has grown up organically, and without any central authority making it conform to some rational design. Hence, its structure can only be determined through careful measurement. The usual technique is a kind of Internet X-ray, carried out by sending information packets from a handful of Internet sources, targeted towards many destinations. Along the way, such packets record the pathways defined by the links on which they travel, and so return narrow 'slices' of the overall network. By putting many such slices together, researchers have assembled crude maps of the Internet, and discovered what seem to be a number of interesting topological features.

For example, the Internet appears to be a 'small world' network in that any pair of node computers (or 'routers') can be linked by short paths having only a handful of links. The Internet also appears to have a highly 'heterogeneous' character, in that a few 'hub' nodes have hundreds or thousands of links, while most have only one or a few. The latter feature appears in the particular mathematical form of the distribution of nodes according to the number of links they have (known as the 'degree distribution'). For the Internet, as for many other complex networks in physics or biology, this seems to follow a power law form, $P(k) \sim k^{-\gamma}$, where γ is between 2.0 and 2.5. The small world and heterogeneous character of the Internet appear to be key to its capacity for stable and resilient operation even in the presence of frequent computer failures, and understanding its topology in detail is crucial for the intelligent design of

and related with the global topological properties of the underlying network. We complement the analytical discussion with a thorough numerical investigation of simulated mapping strategies in network models with different topologies.

Copyright © 2006 S. Karger AG, Basel

1 Introduction

A significant research and technical challenge in the study of large information networks is related to the lack of highly accurate maps providing information on their basic topology. This is mainly due to the dynamical nature of their structure and to the lack of any centralized control resulting in a self-organized growth and evolution of these systems. A prototypical example of this situation is faced in the case of the physical Internet. The topology of the Internet can be investigated at different granularity levels such as the router and autonomous system (AS) level, with the final aim of obtaining an abstract representation where the set of routers (ASs) and their physical connections (peering relations) are the vertices and edges of a graph, respectively. In the absence of accurate maps, local views are obtained by evaluating a certain number of paths to different destinations by using specific tools such as traceroute or by the analysis of BGP tables. At first approximation these processes amount to the collection of shortest paths from a source vertex to a set of target vertices, obtaining a partial spanning tree of the network. The merging of several of these views provides the map of the Internet from which the statistical properties of the network are evaluated.

This strategy has led to obtaining various maps of the Internet [1–5] which have been used for the statistical characterization of the network. Defining $G = (V, E)$ as the sampled graph of the Internet with $N = |V|$ vertices and $|E|$ edges, it is quite

intuitive that the Internet is a sparse graph with a much lower number of edges than in a complete graph: $|E| \ll N(N - 1)/2$. Moreover, the average distance, measured as the shortest path between vertices is very small. This is the so called *small-world* property that is essential for the efficient functioning of the network. Most surprising is the evidence of a skewed and heavy-tailed behaviour for the probability that any vertex in the graph has degree k defined as the number of edges linking each vertex to its neighbours. In particular, the degree distribution appears to be approximated by $P(k) \sim k^{-\gamma}$ with $2 \leq \gamma \leq 2.5$ [6]. Evidence for the heavy-tailed behaviour of the degree distribution has been collected in several other studies at the router and AS level [7–12] and has generated a large activity in the field of network modelling and characterization [13–17].

The obtained maps are, however, undoubtedly incomplete. Along with technical problems such as the instability of paths between routers and interface resolutions [18], typical mapping projects are run from relatively small sets of sources whose combined views are missing a considerable number of edges and vertices [12, 19]. In particular, the various spanning trees are specially missing the lateral connectivity of targets and sample more frequently vertices and links which are closer to each source, introducing spurious effects that might seriously compromise the statistical accuracy of the sampled graph.

These *sampling biases* have been explored in numerical experiments of synthetic graphs generated by different algorithms [20–23]. Very interestingly, it has been shown (numerically and analytically) that apparent degree distributions with heavy tails may be observed even from homogeneous topologies such as in the classic Erdős-Rényi (ER) graph model [20, 21, 24]. These studies thus point out that the evidence obtained from the analysis of the Internet sampled graphs might be insuffi-

better techniques for managing the flow of information.

Unfortunately, studies to date that reveal these properties are only as trustworthy as the maps on which they are based. All studies so far have probed the network from only a few source locations, and researchers know that such probes introduce biases into the data. They discover links along the rough direction between the source and target more easily than links in other directions, and also find nodes close to the source more easily than those further away. Most worryingly, recent studies have found that these biases mean that probing measurements of a network can return power law degree distributions – suggesting heterogeneous topology and the presence of dominant hubs – even if the real underlying network lacks such hubs.

In this paper, Barrat and colleagues attempt to clarify the nature and consequences of such sampling biases in the probe-based mapping of complex networks. As they argue, some nodes (equivalently, ‘vertices’) and links (or ‘edges’), depending on their placement in the network, will naturally have a higher chance of being detected by probing techniques than others. Specifically, the likelihood of detection should depend on the ‘betweenness’ of a node or link, a technical measure of how frequently shortest paths in the network run through these elements. Studies using probes to sample a network can be expected to find most elements of high betweenness, while missing those of lower betweenness. This natural bias introduces errors into any mapping project.

But the authors also use their analysis to estimate how extensive a sampling study has to be (in terms of the number of source and target vertices used) to give a reliable picture of a network. Encouragingly, they conclude that studies showing power-law degree distributions for the Internet can probably be trusted. More generally, the analysis provided here will help to estab-

cient to draw conclusions on the topology of the actual Internet network.

This issue may be tackled through a mean-field statistical analysis and extensive numerical study of shortest path routed sampling, considered as the first approximation to traceroute sampling (see section 2), in different network models. We recall in section 3 the theoretical arguments leading to an approximate expression for the probability of edges and vertices to be detected. The analytical study provides a general understanding of which kind of topologies yields the most accurate sampling. In particular, the map accuracy depends on the underlying network *betweenness centrality* distribution: the heavier the tail the higher the statistical accuracy of the sampled graph.

Numerical investigation of maps obtained varying the number of source-target pairs on networks models with different topological properties provides support to the analytical analysis. In particular, we consider networks with degree distribution with Poissonian, Weibull and power-law behaviour. We study the fractions of discovered vertices and edges as a function of the degree (section 4), stressing the agreement with the theoretical predictions, as well as the degree distributions obtained in the sampled graph (section 5). Single source mapping processes are shown to face serious limitations: even the targeting of the whole network results in a very partial discovery of its connectivity. In contrast, the use of multiple sources promptly leads to obtained maps fairly consistent with the original sample.

In section 6, we also inspect quantitatively the portion of discovered network in different mapping strategies for the deployment of sources that, however, impose the same density of probes to the network. A region of low efficiency (less vertices and edges discovered) is found, depending on the relative proportion of sources and targets. This low efficiency region, however, corresponds to the optimal estimation of

the network average degree. This finding calls for a ‘trade-off’ between the accuracy in the observation of different quantities and hints at possible optimization procedures in the traceroute-driven mapping of large networks.

2 Network Models and Traceroute-Like Processes

In a typical traceroute study, active sources deployed in the network send traceroute probes to a set of destination vertices. Each probe collects information on all the vertices and edges traversed along the path connecting the source to the destination [18]. By merging the information collected on each path it is possible to reconstruct a partial map of the network. The edges and the vertices discovered by each probe will depend on the ‘path selection criterion’ used to decide on the path between a pair of vertices. In the real Internet, many factors, including commercial agreement, traffic congestion and administrative routing policies, contribute to determine the actual path, which may differ even considerably from the shortest path. Despite these local, often unpredictable path distortions, a reasonable first approximation of the route traversed by traceroute-like probes is the shortest path between the two vertices. This assumption, however, is not sufficient for a proper definition of a traceroute model in that equivalent shortest paths between two vertices may exist. For the sake of simplicity, we can thus define three selection mechanisms defining different ideal paths that may account for some of the features encountered in real Internet discovery:

- Unique shortest path (USP) probe. In this case the shortest path route selected between a vertex i and the destination target T is always the same independently of the source S (the path being initially chosen at random among all the equivalent ones).
- Random shortest path (RSP) probe. The shortest path between any source des-

lish guidelines for the future design and interpretation of probing studies on large networks, not only in technological systems, but also in biology and other settings where network topology has to be inferred from incomplete measurements.

As the authors point out, many studies of Internet topology have relied on the ‘traceroute’ facility – a software utility that records the paths of packets through the network. The actual path taken depends on many factors, including the commercial interests of businesses and service providers, but to a rough approximation resembles the shortest path (in a network sense, meaning the fewest number of links) between a source and a target. Depending on the technical details in the design of traceroute studies, they may return just one unique shortest path, select randomly among several equivalent paths (if there is more than one shortest path) or return information on all such paths. Over time, the former procedures should generally give less information on the network, so Barrat *et al.* take this interpretation – the most pessimistic, they note – as a standard for their analyses. To set notation, they call the path returned by a specific probe as an *M-path*.

To analyze the effectiveness of a set of probes, the authors suppose that the true graph (or ‘network’) G is a set of N vertices (or ‘nodes’) V and some edges (or ‘links’) E . The probes run from a set of sources S to a set of targets T , and the data collected come from the entire set of source-target pairs, $\Omega = \{S, T\}$. Putting together all the edges and vertices discovered by the probes for all pairs gives the detected graph G^* consisting of vertices V^* and edges E^* ; this is a subset of the true graph G . The aim of their analysis is to see how the topology of the discovered graph G^* depends on that of the true graph G and on the details of the sampling procedure. In section 3, the authors derive analytical results showing how the probability for any node or edge to be detected depends on the densities of

termination pair is chosen randomly among the set of equivalent shortest paths. This might mimic different peering agreements that make the paths among couples of vertices independent.

- All shortest path (ASP) probe. The selection criterion discovers all the equivalent shortest paths between source-destination pairs. This might happen in the case of repeated probing in time (long-time exploration), so that back-up paths and equivalent paths are discovered in different runs.

We will generically call $\mathcal{M}$ -path the path found using one of these measurement or path selection mechanism. Actual traceroute probes contain a mixture of the three mechanisms defined above. We do not attempt, however, to account for all the subtleties that real studies encounter, i.e. IP routing, BGP policies, interface resolutions and many others. In fact, in the real mapping process, many effective heuristic strategies are commonly applied to improve the reliability and the performances of the sampling. However, it turns out that the different path selection criteria have only little influence on the general picture emerging from our results. Moreover, the USP procedure clearly represents the worst case scenario since, among the three different methods, it yields the minimum number of discoveries. For this reason, if not otherwise specified, we will report the USP data to illustrate the general features of our synthetic exploration. The interest of this analysis resides largely in the choice of working in the most pessimistic case, being aware that path inflations should actually provide a more pervasive sampling of the real network.

More formally, the experimental setup for our simulated traceroute mapping is the following. Let $G = (V, E)$ be a sparse undirected graph with vertices (vertices) $V = \{1, 2, \dots, N\}$ and edges (links) E . Then let us define the sets of vertices $\mathcal{S} = \{i_1, i_2, \dots, i_{N_S}\}$ and $\mathcal{T} = \{j_1, j_2, \dots, j_{N_T}\}$ spec-

ifying the random placement of N_S sources and N_T destination targets. For each ensemble of source-target pairs $\Omega = \{\mathcal{S}, \mathcal{T}\}$, we compute with our path selection criteria the paths connecting each source-target pair. The sampled graph $\mathcal{G} = (V^*, E^*)$ is defined as the set of vertices V^* (with $N^* = |V^*|$) and edges E^* induced by considering the union of all the $\mathcal{M}$ -paths connecting the source-target pairs. The sampled graph is thus analogous to the maps obtained from real traceroute sampling of the Internet.

In our study the parameters of interest are the densities $\rho_T = N_T/N$ and $\rho_S = N_S/N$ of targets and sources. In general, traceroute-driven studies run from a relatively small number of sources to a much larger set of destinations. For this reason, it is appropriate to work with the density of targets ρ_T while still considering N_S instead of the corresponding density. In many cases, an appropriate quantity representing the level of sampling of the networks is $\varepsilon = N_S N_T / N$: it represents the density of traceroute probes in the network and therefore a measure of the load provided to the network by the measuring infrastructure.

In the following, our aim is to evaluate to which extent the statistical properties of the sampled graph $\mathcal{G}$ depend on the parameters of our experimental setup and are representative of the properties of the underlying graph G . The analytical insights of section 3 will be complemented by a numerical investigation of the traceroute-like exploration process on various graph models endowed with very well-defined topological properties, so as to give a clear result on which kind of topologies are related to good sampling performances and vice versa. Starting from this first investigation, further studies could deal with more realistic models such as those created using Internet topology generators [13, 14]. In particular, we will consider two main classes of graphs.

sources ρ_S and targets ρ_T , and then verify these through extensive simulations using several representative examples of graphs likely to be found in the real world.

To derive this result for edges, the authors first define an ‘indicator’ function, π_{ij} , equal to one if an edge in the real graph between nodes i and j shows up in at least one of the shortest paths sampled by probes from sources in $\mathcal{S}$ to targets in $\mathcal{T}$. Their equation 1 gives an expression for this quantity in terms of the function $\pi_{ij}^{l,m}$, defined as equal to one if the edge linking nodes i and j shows up in an M -path corresponding to a particular source l and target m . In this equation, the terms in parenthesis will vanish for some l and m pair only if the edge between i and j appears somewhere, in which case the product will also be zero, giving $\pi_{ij} = 1$.

The authors’ idea is to estimate the average probability of discovering a link between i and j by averaging over all possible choices of source and target nodes. Calculating this average $\langle \pi_{ij} \rangle$ naturally involves taking the average of a product of terms, which is, of course, not generally equal to a product of the averages of those terms. But the authors rely on this approximation, arguing that it should be okay as long as changes in the source and target nodes make the different terms in the product fluctuate independently about their averages. Making this seemingly plausible assumption, and also taking note of the fact that the product $\rho_S \rho_T$ should be much less than one (because probe studies in practice typically use only a handful of sources), this leads the authors to their equation 4 for the discovery probability in terms of the betweenness of the edge, b_{ij} .

This equation makes intuitive sense. The betweenness b_{ij} reflects the probability that if one selects any two nodes at random, anywhere in the network, the shortest path running between them will pass between nodes i and j . Literally, the betweenness measures how frequently the link between i and j stands ‘between’ pairs of

(1) *Homogeneous graphs* in which the degree distribution $P(k)$ has small fluctuations and a well-defined average degree. In this context, the *homogeneity* refers to the existence of a meaningful characteristic average degree that represents the typical value in the graph. The most widely known model for homogeneous graphs is given by the classical ER model [25]: in such random graphs $G_{N,p}$ of N vertices, each edge is present in E independently with probability p . The expected number of edges is therefore $|E| = pN(N - 1)/2$. In order to have sparse graphs one thus needs to have p of order $1/N$, since the average degree is $p(N - 1)$. ER graphs are typical examples of homogeneous graphs, with degree distribution following a Poisson law. Since $G_{N,p}$ can consist of more than one connected component, we consider only the largest of these components.

Another important characteristic discriminating the topology of graphs is the clustering coefficient c_i that, giving the fraction of connected neighbours of a given node i , measures the local cohesiveness of nodes. The average clustering coefficient

$$C = \frac{1}{N} \sum_i c_i$$

provides an indication of the global level of cohesiveness of the graph. This number is generally very small in random graphs that lack correlations. In many real graphs, however, the clustering coefficient appears to be very high and opportune models have been formulated to represent this property, both for homogeneous and heterogeneous graphs. In particular, we consider the construction algorithm proposed by Watts and Strogatz [29] for smallworld networks: starting from a regular network (e.g. a one-dimensional lattice with connections to the $\bar{k}$ nearest neighbours along the chain), each link is rewired with a certain probability p . The resulting degree distribution has a shape similar to the ER graphs, peaking around its average value. The clustering

coefficient, however, is large if $p \ll 1$, making this network a typical example of a clustered homogeneous network.

(2) *Heterogeneous graphs* for which $P(k)$ is a broad distribution with heavy-tail and large fluctuations, spanning various orders of magnitude. The prototype of a scale-free graph is the growing network model by Albert and Barabási (BA) [30]. The preferential attachment mechanism (each new node is connected to m already existing nodes chosen with a probability proportional to their degree) yields a connected graph of $|V| = N$ nodes with $|E| = mN$ edges, having a power-law degree distribution $P(k) \sim k^{-\gamma}$ with $\gamma = 3$, and a small clustering coefficient. Another growing model was introduced by Dorogovtsev, Mendes and Samukhin (DMS) [31]: at each time step, a new node is introduced and connected to the two extremities of a randomly chosen edge, thus forming a triangle. A given node is thus in fact chosen with a probability proportional to its degree, which corresponds to the preferential attachment rule. The resulting graphs have a large clustering coefficient (≈ 0.74) along with a power-law degree distribution $P(k) \sim k^{-\gamma}$ with $\gamma = 3$.

Such graphs can be considered as particular since they are constructed with the preferential attachment mechanism, and we also consider random graphs with given broad degree distributions. In the literature, different definitions of heavy-tailed-like distributions exist. While we do not want to enter the detailed definition of a heavy-tailed distribution we considered two classes of such distributions: (1) *scale-free* or Pareto distributions of the form $P(k) \sim k^{-\gamma}$ (RSF), and (2) Weibull distributions (WEI) $P(k) = (a/c) (k/c)^{a-1} \exp(-(k/c)^a)$. The scale-free distribution has a diverging second moment and therefore virtually unbounded fluctuations, limited only by eventual size cut-off. The Weibull distribution is akin to power-law distributions truncated by an exponential cut-off which are often encountered in the

nodes along the shortest path in the network. It is a measure of the importance of that link as a connector. Not surprisingly, then, probes are more likely to find high betweenness links. Also not surprisingly, this probability increases with $\rho_S \rho_T = N_S N_T / N^2$, the ratio of sampled network pairs to all possible network pairs. The authors note that a similar analysis for the probability of discovering not an edge but a particular node i leads to a similar formula, equation 6, where b_i is the betweenness of that node.

Barrat and colleagues also go on to calculate the expected measured degree $\langle k_i^* \rangle$ of a node, finding their equation 9. (Here $\tilde{b}_i$ is a ‘rescaled’ betweenness equal to b_i/N , and also $\varepsilon = \rho_T N_S$.) Finally, they also report simple formulae in equations 11 and 12 for the ‘redundancy’ of an edge or node, this being defined as the expected number of probes that will discover that edge or node.

To test the validity of these analytic results, the authors ran simulations of traceroute-like studies of several qualitatively different kinds of networks. A network might be assembled in such a way that it makes sense to talk about the average number of links for its nodes. This is the case if the variation in the number of links is small, so that all nodes have roughly the same number of links. Simple and influential models for networks within this category go under the names Erdős-Rényi (ER) or Watts-Strogatz (WS). In contrast, a network may be highly heterogeneous and have some nodes with far more links than others. In this case, the average number of links says little about the network, and the full degree distribution is a more instructive quantity. Barabási and Albert (AB) introduced influential models of this class of networks, as have other authors. In their figures 1 and 2, Barrat and colleagues compare the effectiveness of traceroute-like sampling for estimating the properties of networks within these two different and very broad categories.

Table 1. Main characteristics of the graphs used in the numerical exploration

	ER	ER	WS	BA	DMS	RSF	Weibull
N	10^4	10^4	10^4	10^4	10^4	10^4	10^4
$ E $	10^5	$5 \cdot 10^5$	10^5	$4 \cdot 10^4$	$2 \cdot 10^4$	22,000	55,000
$\bar{k}$	20	100	20	8	4	4.4	11
$\mathbf{C}$	0.002	0.01	0.52	0.006	0.74	0.067	0.12
$k_{\max}$	40	140	26	334	346	3,500	2,000

analysis of scale-free systems in the real world. Indeed, a truncation of the power-law behaviour is generally due to finite-size effects and other physical constraints. Both forms have been proposed as representing the topological properties of the Internet [8]. We generated the corresponding random graphs by using the algorithm proposed by Molloy and Reed [32,33]: the vertices of the graph are assigned a fixed sequence of degrees $\{k_i\}$, $i = 1, \dots, N$, chosen at random from the desired degree distribution $P(k)$, and with the additional constraint that the sum $\sum_i k_i$ must be even; then, the vertices are connected by $\sum_i k_i/2$ edges, respecting the assigned degrees and avoiding self- and multiple connections. The parameters used are $a = 0.25$ and $c = 0.6$ for the Weibull distribution, and $\gamma = 2.3$ for the RSF case.

The main properties of the various graphs are summarized in table 1. In all numerical studies we used networks of $N = 10^4$ vertices. It is noteworthy that the maximum value of the degree ($k_{\max}$) is of the same order as the average for homogeneous graphs, but much larger for heterogeneous ones.

3 Mean-Field Theory of Simulated Mapping Process

We begin our study by recalling briefly the mean-field statistical analysis of the simulated traceroute mapping done by Dall'Asta et al. [34]. The aim is to provide a statistical estimate for the probability of edge and vertex detection as a function of

N_S, N_T and the topology of the underlying graph.

Let us define the quantity

$$\sigma_{i,j}^{(l,m)}$$

that takes the value 1 if the edge (i, j) belongs to the selected $\mathcal{M}$ -path between vertices l and m , and 0 otherwise. For a given set of sources and targets $\Omega = \{S, T\}$, the indicator function that a given edge (i, j) will be discovered and belongs to the sampled graph is simply $\pi_{i,j} = 1$ if the edge (i, j) belongs to at least one of the $\mathcal{M}$ -paths connecting the source-target pairs, and 0 otherwise. We can obtain an exact expression for $\pi_{i,j}$ by noting that $1 - \pi_{i,j}$ is 1 if and only if (i, j) does not belong to any of the paths between sources and targets, i.e. if and only if

$$\sigma_{i,j}^{(l,m)} = 0$$

for all $(l, m) \in \Omega$. This leads to

$$\pi_{i,j} = 1 - \prod_{l \neq m} \left(1 - \sum_{s=1}^{N_S} \delta_{l,i_s} \sum_{t=1}^{N_T} \delta_{m,j_t} \sigma_{i,j}^{(l,m)} \right), \quad (1)$$

where $\delta_{i,j}$ is the Kronecker symbol and selects only vertices belonging to the set of sources or targets.

Starting from the above exact formula, it is interesting to look at the process on a statistical ground by studying the average over all possible realizations of the set $\Omega = \{S, T\}$, identified by $\langle \dots \rangle$

Figure 1 shows the quantity N_k^*/N_k , giving the probability of actually detecting a vertex of degree k if it is present in the network. For both homogeneous and heterogeneous networks, this probability (displayed in the upper row, first and third graphs) increases with the density of probes $\varepsilon = \rho_T N_S$. For the homogeneous networks, N_k^*/N_k remains relatively flat across nodes of a different degree. In contrast, for the heterogeneous networks, the same quantity increases significantly for nodes of higher degree. In the former case, nodes of different degree have roughly the same chance of being detected. In the latter, nodes of high degree will almost certainly be detected, while nodes of low degree will not be. As the authors note, these results can be tied directly to the 'heavy tails' of the power-law degree distribution for the heterogeneous networks. Nodes with high degree tend also to be nodes of high betweenness, and therefore easier to discover. This natural bias also shows up in their figure 2, showing the redundancy as a function of node degree. Again, nodes of higher degree tend to be discovered more frequently.

Figures 3 and 4 finally look at the consequences of these sampling biases for estimates of the degree distribution of a network. For homogeneous networks, figure 3 shows that the mapping strategy returns estimates of the true degree distribution. [Note that here the authors plot the 'cumulative distribution', rather than the density $P(k)$; the cumulative distribution, obtained by integrating $P(k)$, gives the likelihood of finding a vertex of degree greater than or equal to k]. Naturally, these estimates grow better as the density of source and target nodes increases. A significant result, however, is that the estimated distribution only has a spurious resemblance to a power law when only a single source is used. Even the use of two sources is sufficient to avoid this qualitative error in the measured data. Moreover, the measured data has a cutoff at a value determined by the average de-

An uncorrelation assumption allows to obtain the average discovery probability of an edge as

$$\langle \pi_{i,j} \rangle \simeq 1 - \prod_{l \neq m} (1 - \rho_T \rho_S \langle \sigma_{i,j}^{(l,m)} \rangle), \quad (2)$$

where we take advantage of neglecting correlations by replacing the average of the product of variables with the product of the averages. In the case of the ASP probing,

$$\langle \sigma_{i,j}^{(l,m)} \rangle$$

is just 1 if (i,j) belongs to one of the shortest paths between l and m , and 0 otherwise. In the case of the USP and the RSP, in contrast, only one path among all the equivalent ones is chosen. If we denote by $\sigma^{(l,m)}$ the number of shortest paths between vertices l and m , and by

$$x_{i,j}^{(l,m)}$$

the number of these paths passing through the edge (i,j) , the probability that the traceroute model chooses a path going through the edge (i,j) between l and m is

$$\langle \sigma_{i,j}^{(l,m)} \rangle = x_{i,j}^{(l,m)} / \sigma^{(l,m)}.$$

The standard situation we consider is the one in which $\rho_T \rho_S \ll 1$ and since

$$\langle \sigma_{i,j}^{(l,m)} \rangle \leq 1,$$

we have

$$\begin{aligned} & \prod_{l \neq m} (1 - \rho_T \rho_S \langle \sigma_{i,j}^{(l,m)} \rangle) \\ & \simeq \prod_{l \neq m} \exp(-\rho_T \rho_S \langle \sigma_{i,j}^{(l,m)} \rangle), \end{aligned} \quad (3)$$

that inserted in equation 2 yields

$$\langle \pi_{i,j} \rangle \simeq 1 - \exp(-\rho_T \rho_S b_{ij}), \quad (4)$$

where

$$b_{ij} = \sum_{l \neq m} \langle \sigma_{i,j}^{(l,m)} \rangle.$$

In the case of the USP and RSP probing, the quantity b_{ij} is by definition the edge *betweenness centrality*

$$\sum_{l \neq m} x_{i,j}^{(l,m)} / \sigma^{(l,m)}$$

[26, 27], sometimes also refereed to as 'load' [28]. (In the case of ASP probing, it is a closely related quantity.) Indeed the vertex or edge betweenness is defined as the total number of shortest paths among pairs of vertices in the network that pass through a vertex or an edge, respectively. If there are multiple shortest paths between a pair of vertices, the path contributes to the betweenness with the corresponding relative weight. The *betweenness* gives a measure of the amount of all-to-all traffic that goes through an edge or vertex, if the shortest path is used as the metric defining the optimal path between pairs of vertices, and it can be considered as a non-local measure of the *centrality* of an edge or vertex in the graph.

The edge betweenness assumes values between 2 and $N(N-1)$ and the discovery probability of the edge will therefore depend strongly on its betweenness. In particular, for edges with minimum betweenness $b_{ij} = 2$ we have

$$\langle \pi_{i,j} \rangle \simeq 2\rho_T \rho_S,$$

that recovers the probability that the two end vertices of the edge are chosen as source and target. This implies that if the densities of sources and targets are small but finite in the limit of very large N , all the edges in the underlying graph have an appreciable probability to be discovered. Moreover, for edges with high betweenness the discovery probability approaches one. A fair sampling of the network is thus expected. In most realistic samplings, however, we face a very different situation. While it is reasonable to consider ρ_T a small but finite value, the number of source-

gree of nodes in the network. As the authors show, the only way to get a false power-law extending over several orders of magnitude is to have a single source probing an unusual network where the average node degree is itself extremely high, at least 100 or more. This does not seem relevant to the Internet.

Figure 4 finally gives similar results for heterogeneous networks, where nodes of high degree are detected more readily than those of low degree. The results here suggest two conclusions. First, that the heterogeneous character of the network, in a qualitative sense, can be reliably established even with relatively low densities of sources and targets. Second, that distinguishing between various different heterogeneous networks – where the degree distributions show subtle differences that may however indicate differences in the dynamics of network growth, for example – requires a higher level of probing.

The authors go on to explore how mapping studies may be designed specifically so as to return optimal results for different aspects of network topology. It might be more important in some applications, for example, to detect almost all vertices, even if some edges remain undiscovered. Interestingly, they conclude that no one style of study will be best for all situations; rather, there is an inherent trade-off in the design of probes, especially in the selection of the numbers of sources and targets, depending on which elements of the network's topology need to be identified most accurately.

Overall, what is most important in this work is the evident conclusion that probe-based studies of the Internet – or other complex networks – can be carried out and assessed on an informed and rational basis. In particular, the discovery of power-law degree distributions, which seem to imply scale-free organization in the underlying graph, can be made on a legitimate basis. This study appears to be of great practical importance for anyone designing

es is not extensive ($N_S \sim \mathcal{O}(1)$) and their density tends to zero as N^{-1} . In this case it is more convenient to express the edge discovery probability as

$$\langle \pi_{i,j} \rangle \simeq 1 - \exp(-\varepsilon \widetilde{b}_{ij}), \quad (5)$$

where $\varepsilon = \rho_T N_S$ is the density of probes imposed to the system and the rescaled betweenness

$$\widetilde{b}_{ij} = N^{-1} b_{ij}$$

is now limited in the interval $[2N^{-1}, N-1]$. In the limit of large networks $N \rightarrow \infty$ it is clear that edges with low betweenness have

$$\langle \pi_{i,j} \rangle \sim \mathcal{O}(N^{-1}),$$

for any finite value of ε . This readily implies that in real situations the discovery process is generally not complete, a large part of low betweenness edges being not discovered, and that the network sampling is made progressively more accurate by increasing the density of probes ε .

A similar analysis can be performed for the discovery probability of vertices, leading to the average

$$\langle \pi_i \rangle \simeq 1 - (1 - \rho_S - \rho_T) \exp(-\rho_T \rho_S b_i), \quad (6)$$

where b_i is the vertex betweenness centrality, that is limited in the interval $[0, N(N-1)]$ [26–28]. The betweenness value $b_i = 0$ holds for the leafs of the graph, i.e. vertices with a single edge, for which we recover

$$\langle \pi_i \rangle \simeq \rho_S + \rho_T.$$

Indeed, this kind of vertices are dangling ends discovered only if they are either a source or target themselves. As discussed before, the most usual setup corresponds to a density $\rho_S \sim \mathcal{O}(N^{-1})$ and in the large N limit we can conveniently write

$$\langle \pi_i \rangle \simeq 1 - (1 - \rho_T) \exp(-\varepsilon \widetilde{b}_i), \quad (7)$$

where we have neglected terms of order $\mathcal{O}(N^{-1})$ and the rescaled betweenness

$$\widetilde{b}_i = N^{-1} b_i$$

is now defined in the interval $[0, N-1]$. This expression points out that the probability of vertex discovery is favoured by the deployment of a finite density of targets that defines its lower bound.

We can also provide a simple approximation for the effective average degree $\langle k_i^* \rangle$ of vertex i discovered by our sampling process. Each edge departing from the vertex contributes proportionally to its discovery probability, yielding

$$\langle k_i^* \rangle = \sum_j (1 - \exp(-\varepsilon \widetilde{b}_{ij})) \simeq \varepsilon \sum_j \widetilde{b}_{ij}. \quad (8)$$

The final expression is obtained for edges with

$$\varepsilon \widetilde{b}_{ij} \ll 1.$$

Since the sum over all neighbours of the edge betweenness is simply related to the vertex betweenness as $\sum_j b_{ij} = 2(b_i + N-1)$, where the factor 2 considers that each vertex path traverses two edges and the term $N-1$ accounts for all the edge paths for which the vertex is an endpoint, this finally yields

$$\langle k_i^* \rangle \simeq 2\varepsilon + 2\varepsilon \widetilde{b}_i. \quad (9)$$

Finally, the analysis allows to compute the edge redundancy $r_e(i, j)$ of an edge (i, j) , defined as the number of probes passing through the edge (i, j) . This quantity is indeed written for a given set of probes and targets as

$$r_e(i, j) = \sum_{l \neq m} \left(\sum_{s=1}^{N_S} \delta_{l, i_s} \sum_{t=1}^{N_T} \delta_{m, i_t} \sigma_{i,j}^{(l,m)} \right). \quad (10)$$

probe-based mapping efforts, as it shows how the efficacy of such efforts can be scientifically assessed and planned to achieve desired results.

Mark Buchanan

Averaging over all possible realizations and assuming the uncorrelation hypothesis, we obtain

$$\langle r_e(i, j) \rangle \simeq \sum_{l \neq m} \rho_T \rho_S \langle \sigma_{i,j}^{(l,m)} \rangle = \rho_T \rho_S b_{ij}. \quad (11)$$

This result implies that the average redundancy of an edge is related to the density of sources and targets, but also to the edge betweenness. For example, an edge of minimum betweenness $b_{ij} = 2$ can be discovered at most twice in the extreme limit of an all-to-all probing. In contrast, a very central edge of betweenness b_{ij} close to the maximum $N(N-1)$ would be discovered with a redundancy close to $(N-1)$ by a traceroute probing from a single source to all the possible destinations.

Similarly, the redundancy $r_n(i)$ of a vertex i , intended as the number of times the probes cross the vertex i , can be obtained:

$$\langle r_n(i) \rangle \simeq 2\varepsilon + \rho_S \rho_T b_i. \quad (12)$$

In this case, a term related to the number of traceroute probes ε appears, showing that a part of the mapping effort unavoidably ends up in generating vertex detection redundancy.

The present analysis shows that the measured quantities and statistical properties of the sampled graph strongly depend on the parameters of the experimental setup and the topology of the underlying graph. The latter dependence appears

Fig. 1. Frequency N_k^*/N_k of detecting a vertex of degree k , frequency N_b^*/N_b of detecting a vertex of betweenness b and proportion of discovered edges k^*/k as a function of the degree and as a function of the betweenness for the WS (a) and the BA (b) models. The exploration setup considers $N_S = 2$ and increasing probing level ε obtained by progressively higher density of targets ρ_T .

through the key role played by edge and vertex betweenness in the expressions characterizing the graph discovery. The betweenness is a non-local topological quantity whose properties change considerably depending on the kind of graph considered. This allows an intuitive understanding of the fact that graphs with diverse topological properties deliver a different answer to sampling experiments.

4 Numerics

The analytical findings of the previous section may be tested and used as guidance in the numerical analysis of simulated mapping experiments of network models. In particular we will consider the graph topologies defined in section 2. Let us first consider the case of homogeneous graphs (ER and WS model): the vertex and edge betweennesses are homogeneous quantities and their distributions are peaked around their average values $\bar{b}$ and $\bar{b}_e$, respectively, spanning only a small range of variations. These values can thus be considered as typical values. We can thus use equations 5 and 7 to estimate the order of magnitude of probes that allows a fair sampling of the graph. Indeed, both $\langle \pi_{i,j} \rangle$ and $\langle \pi_i \rangle$ tend to 1 if

$$\varepsilon \gg \max[\bar{b}^{-1}, \bar{b}_e^{-1}].$$

In this limit all edges and vertices will have a probability to be discovered very close to 1. At a lower value of ε , obtained by varying ρ_T and N_S , the underlying graph is only partially discovered. Figure 1 shows for the WS model the behaviour of the fraction N_k^*/N_k of discovered vertices of degree k , where N_k is the total number of vertices of degree k in the underlying graph, and the fraction of discovered edges $\langle k^* \rangle / k$ in vertices of degree k . N_k^*/N_k naturally increases with the density of targets and sources, and it is slightly increasing with k . The latter behaviour can be easily understood by noticing that vertices with larger degree have on average a larger betweenness. On the other hand, the range of variation of k in homogeneous graphs is very narrow and only a large level of probing may guarantee very large discovery probabilities. Similarly the behaviour of the effective discovered degree can be understood by looking at equation 9. Indeed the initial decrease of $\langle k^* \rangle / k$ is finally compensated by the increase of $b(k)$.

The situation is different in graphs with heavy-tailed connectivity distributions (BA, DMS, RSF and WEI models), with an

appreciable fraction of vertices and edges with very high betweenness [35]. In particular, in scale-free graphs the site betweenness is related to the vertices degree as $\bar{b}(k) \sim k^\beta$, where β is an exponent depending on the model [35]. Since in heavy-tailed degree distributions the allowed degree is varying over several orders of magnitude, the same occurs for the betweenness values, and the tail of the distribution is broader the broader the connectivity distribution. In such a situation, even in the case of small ε , vertices whose betweenness is large enough ($b_i \varepsilon \gg 1$) have $\langle \pi_i \rangle \simeq 1$. Therefore all vertices with degree $k \gg \varepsilon^{-1/\beta}$ will be detected with probability 1. This is clearly exemplified for the BA model in figure 1 where the discovery probability N_k^*/N_k of vertices with degree k saturates to one for large degree values. Consistently, the degree value at which the curve saturates decreases with increasing ε . A similar effect is appearing in the measurements concerning $\langle k^* \rangle / k$. After an initial decay (fig. 1) the effective discovered degree is increasing with the degree of the vertices. This qualitative feature is captured by equation 9 that gives

$$\langle k^* \rangle / k \simeq \varepsilon k^{-1} (1 + \bar{b}(k)).$$

At large k the term

$$k^{-1} \bar{b}(k) \sim k^{\beta-1}$$

takes over and the effective discovered degree approaches the real degree k . Figure 1 also displays the frequency N_b^*/N_b and the discovered degree of vertices with betweenness b , showing in a more direct way the qualitative agreement with the analytical predictions.

In figure 2 we also report the behaviour of the average vertex redundancy as a function of the degree k for both homogeneous (ER) and heterogeneous (RSF) graphs. For both models, the behaviours are in good agreement with the mean-field prediction, showing the tight relation between redun-

Fig. 2. Average vertex redundancy as a function of the degree k for RSF (a) and ER (b, c) model ($N = 10^4$). For the ER model, two blocks of data are plotted, for $\bar{k} = 20$ (b) and for $\bar{k} = 100$ (c). The target density is fixed ($\rho_T = 0.1$), and $N_S = 2, 10$, and 20 . The dashed lines represent the analytical prediction $2\epsilon + \rho_S \rho_T \bar{b}(k)$ in perfect agreement with the simulations.

Fig. 3. Cumulative degree distribution of the sampled ER graph for USP probes. a, b Correspond to $\bar{k} = 20$. c, d Correspond to $\bar{k} = 100$. a, c Sampled distributions obtained with $N_S = 2$ and varying density target ρ_T . In the insets we report the peculiar case $N_S = 1$ that provides an apparent power-law behaviour with exponent -1 at all values of ρ_T , with a cut-off depending on $\bar{k}$. The insets are in lin-log scale to show the logarithmic behaviour of the corresponding cumulative distribution. b, d Correspond to $\rho_T = 0.1$ and varying number of sources N_S . The solid lines are the degree distributions of the underlying graph. For $\bar{k} = 100$, the sampled cumulative distributions display plateaus corresponding to peaks in the degree distributions, induced by the sampling process.

dancy and betweenness centrality. In the case of heavy-tailed underlying networks, the vertex redundancy typically grows as a power law of the degree, while the values for random graphs vary on a smaller scale. This behaviour points out that the intrinsic hierarchical structure of scale-free networks plays a fundamental role even in the process of path routing, resulting in a huge number of probes iteratively passing through the same set of few hubs. On the other hand, for homogeneous graphs the total number of vertex discoveries is quite uniformly distributed over the whole range of connectivity, independently of the relative importance of the vertices.

5 Degree Distribution Measurements

A very important quantity in the study of the statistical accuracy of the sampled graph is the degree distribution. Figure 3 shows the cumulative degree distribution $P_c(k^* > k)$ of the sampled graph defined by the ER model for increasing density of targets and sources. Sampled distributions are only approximating the genuine distribution; however, for $N_S \geq 2$ they are far from true heavy-tail distributions at any appreciable level of probing. Indeed, the distribution runs generally over a small range of degrees, with a cut-off that sets in at the average degree $\bar{k}$ of the underlying graph. In order to stretch the distribution range, homogeneous graphs with a very large average degree $\bar{k}$ must be considered; however, other distinctive spurious effects appear in this case. In particular, since the best sampling occurs around the high degree values, the distributions develop peaks that show in the cumulative distribution as plateaus. Note that, in the case of RSP and ASP model, the obtained distributions are closer to the real one since they allow a larger number of discoveries.

Only in the peculiar case of $N_S = 1$ an apparent scale-free behaviour with slope -1 is observed for all target densities ρ_T , as analytically shown by Clauset and Moore [21]

Fig. 4. Cumulative degree distributions of the sampled RSF, WEI and DMS graphs for USP probes. The top figures show sampled distributions obtained with $N_S = 5$ and varying density target ρ_T . The figures at the bottom correspond to $\rho_T = 0.25$ and a varying number of sources N_S . The solid lines are the degree distributions of the underlying graph.

and Achlioptas et al. [24]. Also in this case, the distribution cut-off is consistently determined by the average degree $\bar{k}$. The present analysis shows that in order to obtain a sampled graph with apparent scale-free behaviour on a degree range varying over n orders of magnitude we would need the very peculiar sampling of a homogeneous underlying graph with an average degree

$$\bar{k} \simeq 10^n,$$

a rather unrealistic situation in the Internet and many other information systems where $n \geq 2$.

Since, in heterogeneous graphs, vertices with high degree are efficiently sampled with an effective measured degree that is rather close to the real one, the degree distribution tail is fairly well sampled, while deviations should be expected at lower degree values. This is indeed what we observe in numerical experiments on graphs with heavy-tailed distributions (see fig. 4). Despite both RSF and WEI underlying graphs have a small average degree; the observed degree distribution spans more than two orders of magnitude. The distribution tail

is fairly reproduced even at rather small values of ε . The data shows clearly that the low degree regime is instead undersampled. This undersampling can either yield an apparent change in the exponent of the degree distribution (as also noticed in Petermann and De Los Rios [22] for single source experiments), or, if N_S is small, yield a power-law-like distribution for an underlying Weibull distribution. Furthermore, as figure 4 shows, an increase in the number of sources starts to discriminate between scale-free and Weibull distributions by detecting a curvature in the second case even at small values $\rho_T = 0.25$. It is, however, fair to say that while the experiments clearly point out a broad and heavy-tailed distribution, the distinction between different types of heavy-tailed distribution needs an adequate level of probing.

In conclusion, graphs with heavy-tailed degree distribution allow a better qualitative representation of their statistical features in sampling experiments. Indeed, the most important properties of these graphs are related to the heavy tail part of the statistical distributions that are indeed well

discriminated by the traceroute-like exploration. On the other hand, the accurate identification of the distribution forms requires a fair level of sampling that it is not clear how to determine quantitatively in the case of an unknown underlying network. We will discuss the implications of these results in real Internet measurements in section 7.

6 Optimization of Mapping Strategies

In the previous sections we have shown that it is possible to have a general qualitative understanding of the efficiency of network exploration and the induced biases on the statistical properties. The quantitative analysis of the sampling strategies, however, is a much harder task that calls for a detailed study of the discovered proportion of the underlying graph and the precise deployment of sources and targets. In this perspective, very important quantities are the fraction N^*/N and E^*/E of vertices and edges discovered in the sampled graph, respectively. Unfortunately, the mean-field approximation breaks down when we aim at a quantitative representation of the results. The neglected correlations are in fact very important for the precise estimate of the various quantities of interest. For this reason we performed an extensive set of numerical explorations aimed at a fine determination of the level of sampling achieved for different experimental setups.

In figure 5 we report the proportion of discovered nodes and edges in the numerical exploration of the graph models defined previously for increasing the level of probing ε . The level of probing is increased either by raising the number of sources at fixed target density or by raising the target density at a fixed number of sources. As expected, both strategies are progressively more efficient with increasing levels of probing. In heterogeneous graphs, it is also possible to see that when the number of sources is $N_S \sim \mathcal{O}(1)$ the increase of the

Fig. 5. Behaviour of the fraction of discovered nodes and edges in explorations with increasing ε for RSF, WEI and ER graphs. For each underlying graph studied we report two curves corresponding to larger ε achieved by increasing the target density ρ_T at constant $N_S = 5$ or the number of sources N_S at constant $\rho_T = 0.1$. Curves similar to ER are obtained for WS, and to RSF for BA and DMS.

Fig. 6. Explorations with fixed ε (here $\varepsilon = 2$). Behaviour as a function of ρ_T of the fraction of discovered vertices N^*/N and edges E^*/E , of the normalized average degree $\bar{k}^*/\bar{k}$ and of the fraction of the normalized average clustering coefficient C^*/C . Since $\varepsilon = \rho_T N_S$, the increase of ρ_T corresponds to a lowering of the number of sources N_S .

number of targets achieves better sampling than increasing the deployed sources. On the other hand, it is easy to perceive that the shortest path route mapping is a symmetric process if we exchange sources with targets. This is confirmed by numerical experiments in which we use a very large number of sources and a density of targets $\rho_T \sim \mathcal{O}(1/N)$, where the trends are opposite: the increase of the number of sources achieves better sampling than increasing the deployed targets.

In figure 6, we report the behaviour of E^*/E and N^*/N at fixed ε and varying N_S and ρ_T . Very interestingly, the curves show a structure allowing for local minima and maxima in the discovered portion of the underlying graph: at fixed levels of probing ε , different proportions of sources and targets may achieve different levels of sampling. This hints at the search for optimal strategies in the relative deployment of sources and targets. The picture, however, is more complicated if we look at other quantities in the sampled graph. In figure 6 we show the behaviour at fixed ε of the average degree $\bar{k}^*$ measured in sampled graphs normalized by the actual average degree $\bar{k}$ of the underlying graph as a function of ρ_T . The plot in this case also shows a symmetric structure. By comparing the data of figure 6 we notice that the symmetry point is of a different nature for different quantities: the minimum in the fraction of discovered edges corresponds to the best estimate of the average degree. This implies that at the symmetry point the exploration discovers less edges than in other setups, however, achieving a more efficient sampling of the effective degree for the discovered vertices. A similar problem is obtained by studying the behaviour of the ratio C^*/C between the clustering coefficient of the sampled and the underlying graphs: the best level of sampling is achieved at particular values of ε and N_S that are conflicting with the best sampling of other quantities.

The evidence purported in this section hints at a possible optimization of the sampling strategy. The optimal solution, however, appears as a trade-off strategy between the different levels of efficiency achieved in competing ranges of the experimental setup. In this respect, a detailed and quantitative investigation of the various quantities of interest in different experimental setups is needed in order to pinpoint the most efficient deployment of source-target pairs depending on the underlying graph topology. While such a detailed analysis lies beyond the scope of the present study, an interesting hint comes from the analytical results of section 3: since vertices with large betweenness have typically a very large probability of being discovered, placing the sources and targets preferentially on low-betweenness vertices (the most difficult to discover) may have an impact on the whole process. The usual correlation between connectivity and betweenness thus indicates that the exploration of a real network could be improved by a massive deployment of sources using low-connectivity vertices.

7 Conclusions and Outlook

The rationalization of the sampling biases at the statistical level provides a general interpretative framework for the results obtained from the numerical experiments on graph models. The sampled graph clearly distinguishes between homogeneous and heavy-tailed topologies. This is due to the exploration process that statistically focuses on high betweenness vertices, thus providing a very accurate sampling of the distribution tail. In graphs with heavy tails, such as scale-free networks, the main topological features are therefore easily discriminated since the relevant statistical information is encapsulated in the degree distribution tail which is fairly well captured. Quite surprisingly, the sampling of homogeneous graphs appears more cumbersome than those of heavy-tailed graphs. Dramatic ef-

fects such as the existence of apparent power laws, however, are found only in very peculiar cases. In general, exploration strategies provide sampled distributions with enough signatures to distinguish at the statistical level between graphs with different topologies.

This evidence might be relevant in the discussion of real data from Internet mapping projects. Indeed, data indicate the presence of heavy-tailed degree distribution both at the router and AS level. The present discussion indicates that it is very unlikely that this feature is just an artefact of the mapping strategies. The upper degree cut-off at the router and AS level runs up to 10^2 and 10^3 , respectively. A homogeneous graph should have an average degree comparable to the measured cut-off, which is hardly conceivable in a realistic perspective (for instance, it would be necessary that nine routers over ten would have more than 100 links to other routers). In addition, the major part of mapping projects are multi-source, a feature that readily washes out the presence of spurious power-law behaviour. In contrast, heterogeneous networks with heavy-tailed degree distributions are sampled with particular accuracy for the large degree part, generally at all probing levels. This makes it very plausible, and a natural consequence, that the heavy-tail behaviour observed in real mapping experiments is a genuine feature of the Internet.

On the other hand, it is important to stress that while at the qualitative level the sampled graphs allow a discrimination of the statistical properties, at the quantitative level they might exhibit considerable deviations from the true values such as size, average degree, and the precise analytic form of the heavy-tailed degree distribution. For instance, the exponent of the power-law behaviour appears to suffer from noticeable biases. In this respect, it is of major importance to define strategies that optimize the estimate of the various parameters and quantities of the underly-

ing graph. In this paper we have shown that the proportion of sources and targets may have an impact on the accuracy of the measurements even if the number of total probes imposed to the system is the same. For instance, the deployment of a highly distributed infrastructure of sources probing a limited number of targets may turn out to be as efficient as few very powerful sources probing a large fraction of the addressable space [36, 37]. The optimization of large network sampling is therefore an open problem that calls for further work aimed at a more quantitative assessment of the mapping strategies both on the analytic and numerical side.

Acknowledgements

We are grateful to M. Crovella, P. De Los Rios, T. Erlebach, T. Friedman, M. Latapy and T. Petermann for very useful discussions and comments. This work has been partially supported by the European Commission – Fet Open project COSIN IST-2001-33555 and contract 001907 (DELIS).

References

- 1 National Laboratory for Applied Network Research (NLNR), sponsored by the National Science Foundation (see <http://moat.nlanr.net/>).
- 2 Cooperative Association for Internet Data Analysis (CAIDA), located at the San Diego Supercomputer Center (see <http://www.caida.org/home/>).
- 3 Topology Project, Electric Engineering and Computer Science Department, University of Michigan (<http://topology.eecs.umich.edu/>).
- 4 SCAN Project at the Information Sciences Institute (<http://www.isi.edu/div7/scan/>).
- 5 Internet Mapping Project at Lucent Bell Labs (<http://www.cs.belllabs.com/who/ches/map/>).
- 6 Faloutsos M, Faloutsos P, Faloutsos C: On power-law relationships of the Internet topology. ACM SIGCOMM '99. Comput Commun Rev 1999; 29: 251–262.
- 7 Govindan R, Tangmunarunkit H: Heuristics for Internet Map Discovery. Proceedings of IEEE Infocom 2000. Los Alamitos, IEEE Computer Society Press, 2000, vol 3, pp 1371–1380.
- 8 Broido A, Claffy KC: Internet Topology: Connectivity of IP Graphs. San Diego Proceedings of SPIE International Symposium on Convergence of IT and Communication, Denver, 2001.
- 9 Caldarelli G, Marchetti R, Pietronero L: The fractal properties of Internet. Europhys Lett 2000; 52: 386.
- 10 Pastor-Satorras R, Vázquez A, Vespignani A: Dynamical and correlation properties of the Internet. Phys Rev Lett 2001; 87: 258701.

- 11 Vázquez A, Pastor-Satorras R, Vespignani A: Large-scale topological and dynamical properties of the Internet. *Phys Rev E* 2002; 65: 066130.
- 12 Chen Q, Chang H, Govindan R, Jamin S, Shenker SJ, Willinger W: The Origin of Power Laws in Internet Topologies Revisited. *Proceedings of IEEE Infocom* 2002, New York, 2002.
- 13 Medina A, Matta I: BRITE: a flexible generator of Internet topologies. Tech Rep BU-CS-TR-2000-005, Boston University, 2000.
- 14 Jin C, Chen Q, Jamin S: INET: Internet topology generators. Tech Rep CSETR-433-00, EECS Department, University of Michigan, 2000.
- 15 Dorogovtsev SN, Mendes JFF: *Evolution of Networks: from Biological Nets to the Internet and WWW*. Oxford, Oxford University Press, 2003.
- 16 Baldi P, Frasconi P, Smyth P: *Modeling the Internet and the Web: Probabilistic Methods and Algorithms*. Chichester, Wiley, 2003.
- 17 Pastor-Satorras R, Vespignani A: *Evolution and Structure of the Internet: a Statistical Physics Approach*. Cambridge, Cambridge University Press, 2004.
- 18 Burch H, Cheswick B: Mapping the Internet. *IEEE Comput* 1999; 32: 97–98.
- 19 Willinger W, Govindan R, Jamin S, Paxson V, Shenker S: Scaling phenomena in the Internet: critically examining criticality. *Proc Natl Acad Sci USA* 2002; 99: 2573–2580.
- 20 Lakhina A, Byers JW, Crovella M, Xie P: Sampling biases in IP topology measurements. Tech Rep BUCS-TR-2002-021, Department of Computer Sciences, Boston University, 2002.
- 21 Clauset A, Moore C: Accuracy and scaling phenomena in Internet mapping. *Phys Rev Lett* 2005; 94: 018701.
- 22 Petermann T, De Los Rios P: Exploration of scale-free networks – do we measure the real exponents? *Eur Phys J B* 2004; 38: 201–204.
- 23 Guillaume J-L, Latapy M: Relevance of Massively Distributed Explorations of the Internet Topology: Simulation Results. *Proceedings of Infocom 2005*, in press.
- 24 Achlioptas D, Clauset A, Kempe D, Moore C: On the Bias of Traceroute Sampling or Power-Law Degree Distributions in Regular Graphs. *condmat/0503087*. STOC 2005, in press.
- 25 Erdős P, Rényi P: On random graphs I. *Publ Math Debrecen* 1959; 6: 290.
- 26 Freeman LC: A Set of Measures of Centrality Based on Betweenness. *Sociometry* 1977; 40: 35–41.
- 27 Brandes U: A faster algorithm for betweenness centrality. *J Math Soc* 2001; 25: 163–177.
- 28 Goh K-I, Kahng B, Kim D: Universal behavior of load distribution in scale-free networks. *Phys Rev Lett* 2001; 87: 278701.
- 29 Watts DJ, Strogatz SH: Collective dynamics of small-world networks. *Nature* 1998; 393: 440–442.
- 30 Barabási A-L, Albert R: Emergence of scaling in random networks. *Science* 1999; 286: 509–512.
- 31 Dorogovtsev SN, Mendes JFF, Samukhin AN: Size-dependent degree distribution of a scale-free growing network. *Phys Rev E* 2001; 63: 062101.
- 32 Molloy M, Reed B: A critical point for random graphs with a given degree sequence. *Random Struct Algorithms* 1995; 6: 161.
- 33 Molloy M, Reed B: The size of the giant component of a random graph with a given degree distribution. *Combinatorics Probab Comput* 1998; 7: 295.
- 34 Dall'Asta L, Alvarez-Hamelin I, Barrat A, Vázquez A, Vespignani A: Statistical theory of Internet exploration. *Phys Rev E* 2005; 71: 036135.
- 35 Barthélemy M: Betweenness centrality in large complex networks. *Eur Phys J B* 2004; 38: 163–168.
- 36 <http://www.netdimes.org/>; Shavitt Y, Shir E: DIMES: Let the Internet Measure Itself. Preprint cs.NI/0506099.
- 37 <http://www.tracerouteathome.net/> 21.