

Accuracy and Scaling Phenomena in Internet Mapping

Aaron Clauset^{1,*} and Cristopher Moore^{1,2}

¹Computer Science Department, University of New Mexico, Albuquerque New Mexico 87131, USA

²Department of Physics and Astronomy, University of New Mexico, Albuquerque New Mexico 87131, USA

(Received 4 October 2004; published 6 January 2005)

It was recently argued that sampling a network by traversing it with paths from a small number of sources, as with *traceroutes* on the Internet, creates a fundamental bias in observed topological features like the degree distribution. We examine this bias analytically and experimentally. For Erdős-Rényi random graphs with mean degree c , we show analytically that such sampling gives an observed degree distribution $P(k) \sim k^{-1}$ for $k \lesssim c$, despite the underlying distribution being Poissonian. For graphs whose degree distributions have power-law tails $P(k) \sim k^{-\alpha}$, sampling can significantly underestimate α when the graph has a large excess (i.e., many more edges than vertices). We find that in order to accurately estimate α , one must use a number of sources which grows linearly in the mean degree of the underlying graph. Finally, we comment on the accuracy of the published values of α for the Internet.

DOI: 10.1103/PhysRevLett.94.018701

PACS numbers: 89.75.Hc, 02.50.Ey, 89.20.Hh, 89.75.Da

The Internet is a canonical complex network, and a great deal of effort has been spent measuring its topology. However, unlike the Web where a page's outgoing links are directly visible, we cannot typically ask a router who its neighbors are. As a result, studies have sought to infer the topology of the Internet by aggregating paths or *traceroutes* through the network, typically from a small number of sources to a large number of destinations [1–5], routing decisions like those imbedded in Border Gateway Protocol (BGP) routing tables [6–8], or both [9–11]. Although such methods are known to be noisy [7,12–14], they strongly suggest that the Internet has a power-law degree distribution at both the router and domain levels.

However, Lakhina *et al.* [15] recently argued that traceroute-based sampling introduces a fundamental bias in topological inferences, since the probability that an edge appears within an efficient route decreases with its distance from the source. They showed empirically that traceroutes from a single source cause Erdős-Rényi random graphs $G(n, p)$, whose underlying distribution is Poissonian [16], to appear to have a power-law degree distribution $P(k) \sim k^{-1}$. Here, we prove this evocative result analytically by modeling the growth of a spanning tree on $G(n, p)$ using differential equations.

Although it is widely accepted that the Internet, unlike $G(n, p)$, has a power-law degree distribution $P(k) \sim k^{-\alpha}$ with $2 < \alpha < 3$ [9], we may reasonably ask whether traceroute sampling accurately estimates the exponent α . Petermann and de los Rios [17] and Dall'Asta [18] considered this question, and found that because low-degree vertices are undersampled relative to high-degree ones, the observed value of α is lower than the true exponent of the underlying graph. We explore this idea further, and find that single-source traceroute sampling only gives a good estimate of α when the underlying graph has a small excess, i.e., has average degree close to 2 and is close to a tree. As the average degree grows, so does the extent to which traceroute sampling underestimates α .

Since single-source traceroutes can significantly underestimate α , we then turn to the question of how many sources are required to obtain an accurate estimate of α . We find that the number of sources needed increases linearly with the average degree. We conclude with some discussion of whether the published values of α for the Internet are accurate, and how to tell experimentally whether more sources are needed.

Traceroute spanning trees: analytical results.—The set of traceroutes from a single source can be modeled as a spanning tree [19]. If we assume that Internet routing protocols approximate shortest paths, this spanning tree is built breadth-first from the source. In fact, the results of this section apply to spanning trees built in a variety of ways, as we will see below.

We can think of the spanning tree as built step-by-step by an algorithm that explores the graph. At each step, every vertex in the graph is labeled *reached*, *pending*, or *unknown*. Pending vertices are the leaves of the current tree; reached vertices are interior vertices; and unknown vertices are those not yet connected. We initialize the process by labeling the source vertex pending, and all other vertices unknown. Then the growth of the spanning tree is given by the following pseudocode:

```

while there are pending vertices:
    choose a pending vertex  $v$ 
    label  $v$  reached
    for every unknown neighbor  $u$  of  $v$ ,
        label  $u$  pending.
  
```

The type of spanning tree is determined by how we choose the pending vertex v . Storing vertices in a queue and taking them in FIFO (first-in first-out) order gives a breadth-first tree of shortest paths; if we like we can break ties randomly between vertices of the same age in the queue, which is equivalent to adding a small noise term to the length of

each edge as in [15] (see also [20]). Storing pending vertices on a stack and taking them in LIFO (last-in first-out) order builds a depth-first tree. Finally, choosing v uniformly at random from the pending vertices gives a “random-first” tree.

Surprisingly, while these three processes build different trees, and traverse them in different orders, they all yield the same degree distribution when n is large. To illustrate this, Fig. 1 shows the degree distributions for each type of spanning tree for a random graph $G(n, p = c/n)$ where $n = 10^5$ and $c = 100$. The three degree distributions are indistinguishable, and all agree with the analytic results derived below.

We now show analytically that building spanning trees in Erdős-Rényi random graphs $G(n, p = c/n)$ using any of the processes described above gives rise to an apparent power-law degree distribution $P(k) \sim k^{-1}$ for $k \leq c$. To model the progress of the **while** loop described above, let $S(T)$ and $U(T)$ denote the number of pending and unknown vertices at step T respectively. The expected changes in these variables at each step are the following:

$$\begin{aligned} E[U(T+1) - U(T)] &= -pU(T) \\ E[S(T+1) - S(T)] &= pU(T) - 1. \end{aligned} \quad (1)$$

Here the $pU(T)$ terms come from the fact that a given unknown vertex u is connected to the chosen pending vertex v with probability p , in which case we change its label from unknown to pending; the -1 term comes from the fact that we also change v 's label from pending to reached. Moreover, these equations apply no matter how we choose v ; whether v is the “oldest” vertex (breadth-




FIG. 1 (color online). Sampled degree distributions from breadth-first, depth-first, and random-first spanning trees on a random graph of size $n = 10^5$ and average degree $c = 100$, and our analytic results (black dots). For comparison, the black line shows the Poisson degree distribution of the underlying graph. Note the power-law behavior of the apparent degree distribution $P(k) \sim k^{-1}$, which extends up to a cutoff at $k \sim c$.

first), the “youngest” one (depth-first), or a random one (random-first). Since edges in $G(n, p)$ are independent, the events that v is connected to each unknown vertex u are independent and occur with probability p .

Writing $t = T/n$, $s(t) = S(Tn)/n$ and $u(tn) = U(T)/n$, the difference equations (1) become the following system of differential equations:

$$\frac{du}{dt} = -cu, \quad \frac{ds}{dt} = cu - 1. \quad (2)$$

With the initial conditions $u(0) = 1$ and $s(0) = 0$, the solution to (2) is

$$u(t) = e^{-ct}, \quad s(t) = 1 - t - e^{-ct}. \quad (3)$$

The algorithm ends at the smallest positive root t_0 of $s(t) = 0$; using Lambert’s function W , defined as $W(x) = y$ where $ye^y = x$, we can write

$$t_0 = 1 + \frac{1}{c} W(-ce^{-c}). \quad (4)$$

Note that t_0 is the fraction of vertices which are reached at the end of the process, and this is simply the size of the giant component of $G(n, c/n)$.

Now, we wish to calculate the degree distribution $P(k)$ of this tree. The degree of each vertex v is the number of its previously unknown neighbors, plus one for the edge by which it became attached (except for the root). Now, if v is chosen at time t , in the limit $n \rightarrow \infty$ the probability it has k unknown neighbors is given by the Poisson distribution with mean $m = cu(t)$, $\text{Poisson}(m, k) = e^{-m} m^k / k!$. Averaging over all the vertices in the tree and ignoring $o(1)$ terms gives

$$P(k+1) = \frac{1}{t_0} \int_0^{t_0} dt \text{Poisson}[cu(t), k].$$

It is helpful to change the variable of integration to m . Since $m = ce^{-ct}$ we have $dm = -cm dt$, and

$$\begin{aligned} P(k+1) &= \frac{1}{t_0} \int_{c(1-t_0)}^c dm \frac{\text{Poisson}(m, k)}{cm} \\ &\approx \int_{ce^{-c}}^c dm \frac{\text{Poisson}(m, k)}{cm} \\ &= \frac{1}{ck!} \int_{ce^{-c}}^c dm e^{-m} m^{k-1}. \end{aligned} \quad (5)$$

Here in the second line we use the fact that $t_0 \approx 1 - e^{-c}$ when c is large (i.e., the giant component encompasses almost all of the graph).

The integral in (5) is given by the difference between two incomplete Gamma functions. However, since the integrand is peaked at $m = k - 1$ and falls off exponentially for larger m , for $k \leq c$ it coincides almost exactly with the full Gamma function $\Gamma(k)$. Specifically, for any $c > 0$ we have

$$\int_0^{ce^{-c}} dm e^{-m} m^{k-1} < ce^{-c}$$

and, if $k - 1 = c(1 - \epsilon)$ for $\epsilon > 0$, then

$$\begin{aligned} \int_c^\infty dm e^{-m} m^{k-1} &= e^{-c} c^{k-1} \int_0^\infty dx e^{-x} (1 + x/c)^{k-1} \\ &< e^{-c} c^{k-1} \int_0^\infty dx e^{-x} e^{x(k-1)/c} \\ &= \frac{e}{-c} c^{k-1} \epsilon < \frac{e}{-(k-1)} (k-1)^{k-1} \epsilon \\ &\approx \frac{\Gamma(k)}{\epsilon \sqrt{2\pi(k-1)}}. \end{aligned} \quad (6)$$

This is $o[\Gamma(k)]$ if $\epsilon \geq 1/\sqrt{k}$, i.e., if $k < c - c^\alpha$ for some $\alpha > 1/2$. In that case we have

$$P(k+1) = [1 - o(1)] \frac{\Gamma(k)}{ck!} \sim \frac{1}{ck} \quad (9)$$

giving a power law k^{-1} up to the mean degree c .

We note that this derivation can be made mathematically rigorous, at least for constant c . Wormald [21] showed, under fairly generic conditions, that discrete stochastic processes like this one are well modeled by the corresponding differential equations. Specifically, we can show that if the initial source vertex is in the giant component, then with high probability, for all t such that $0 < t < t_0$, $U(tn) = u(t)n + o(n)$ and $S(tn) = s(t)n + o(n)$. It follows that with high probability our calculations give the correct degree distribution of the spanning tree within $o(1)$.

Power-law degree distributions.—While the result of the previous section shows that power-law degree distributions can be observed even when none exist, the evidence seems overwhelming that the Internet does, in fact, have a power-law degree distribution $P(k) \sim k^{-\alpha}$. However, as shown in [17,18], traceroute sampling on graphs of this kind can underestimate the value of α by undersampling the low-degree vertices relative to the high-degree ones. Here we show experimentally that the extent of this underestimate increases with the average degree of the underlying graph. We performed experiments on both the preferential attachment model of Barabási and Albert [22] and the configuration model [23].

The preferential attachment model of [22] gives each new vertex m edges, and so has minimum degree m and average degree $2m$. In the extreme case $m = 1$, the graph is a tree, and traceroutes from a single source will sample every edge. However, as m increases the fraction of edges sampled by a given source decreases. Figure 2 shows the observed and underlying degree distributions for different values of m . For $m = 2$, for instance, the observed slope is $\alpha_{\text{obs}} \approx 2.7$ instead of the correct value $\alpha = 3$.

It is worth pointing out that the average degree, and therefore α_{obs} , is highly sensitive to the low-degree part of the degree distribution, not just the shape of its high-degree tail. For instance, we used the configuration model [23] to construct random graphs with minimum degree k_{min} and a power-law tail, i.e., $P(k) = 0$ for $k < k_{\text{min}}$ and $P(k) \propto k^{-\alpha}$ for $k \geq k_{\text{min}}$. [Note that the normalization of

$P(k)$ then depends on k_{min} .] Here we found that α_{obs} is a function of k_{min} , not just of α . We are currently extending our analytic calculations to this and other degree distributions.

Building unbiased maps.—Since single-source traceroutes can significantly underestimate α , especially for graphs of large average degree, we now turn to the question of how many sources are needed to obtain a good estimate of α . In Fig. 3, we show the observed exponent (estimated by performing a fit to the high-degree tail $k \gg m$) for preferential attachment networks as a function of the number of sources divided by m ; it also shows the fraction of edges included in the sample. The collapse of the data clearly shows that the number of sources s we need to converge to within a given error from the true exponent grows linearly in m , and the error decreases rapidly as s/m increases. For instance, with m sources we see 41% of the edges and $\alpha_{\text{obs}} \approx 2.82$; with $10m$ sources, 5 times the average degree, we see 94% of the edges and our estimate improves to $\alpha_{\text{obs}} \approx 2.99$.

Traceroute-based studies [1–5,9–11,13] suggest an average degree for the Internet of 2.8 ± 0.5 . (Of course, it may be higher since these studies do not see all the edges of the graph.) However, none of these studies use more than 12 sources, suggesting that the published values of α may still be somewhat low.

For the Internet, gaining access to an increasing number of sources in order to sample traceroutes from them can present practical difficulties. However, even if the measured exponent increases with each additional source—indicating that we still do not have the correct value of α , and the “marginal value” [13] of each source is nonzero—it may be possible to extrapolate the true α from the rate of convergence.




FIG. 2 (color online). Single-source traceroute sampling for preferential attachment networks with $n = 5 \times 10^5$ and varying values of the minimum degree m . The extent to which traceroute sampling underestimates α increases with m .




FIG. 3 (color online). Performance of multisource traceroute sampling in preferential attachment networks as a function of the number of sources divided by m . On the left, the convergence of α_{obs} to the correct value $\alpha = 3$; on the right, the fraction of edges observed at least once. Both curves collapse, showing that the number of sources necessary to counter the sampling bias grows linearly with the average degree.

Conclusions.—Unlike the World Wide Web where links are visible, the Internet’s topology must be sampled indirectly, e.g., by traceroutes. Lakhina *et al.* [15] noted that this process can significantly bias the observed degree distribution, and may create the appearance of a power law where none exists. We have proved this result analytically for random graphs $G(n, p = c/n)$, showing that single-source traceroutes yield an observed distribution $P(k) \sim k^{-1}$ up to the mean degree c . Other mechanisms for observing apparent power laws in $G(n, p)$ include gradient-based flows [24], probabilistic pruning [17], and minimum weight spanning trees [25]; however, these are rather different from our analysis.

For graphs with a power-law distribution $P(k) \sim k^{-\alpha}$ traceroute sampling underestimates α by undersampling low-degree vertices [17,18], and we have found that the extent of this underestimate increases with the network’s average degree. To compensate for this effect, we have found that to estimate α within a given error it is necessary to use a number of sources that grows linearly with the average degree. Given the small number of sources used in existing studies, it seems possible to us that the published values of α for the Internet are somewhat low. In future work, we will measure whether α_{obs} for the Internet increases with the number of sources, and if it does, attempt to extrapolate the correct value of α .

The authors are grateful to David Kempe, Mark Newman, Mark Crovella, Paolo De Los Rios, Michel Morvan, Todd Underwood, Dimitris Achlioptas, Nick Hengartner, and Tracy Conrad for helpful conversations, and to ENS Lyon for their hospitality. This work was funded by NSF Grants No. PHY-0200909, No. ITR-0324845, and Hewlett-Packard Gift 88425.1 to Darko

Stefanovic.

*Corresponding author.

Electronic address: aaron@cs.unm.edu

- [1] J.-J. Pansiot and D. Grad, ACM SIGCOMM Computer Commun. Rev. **28**, 41 (1998).
- [2] R. Govindan and H. Tangmunarunkit, in Proceedings of IEEE INFOCOM, Tel Aviv, 2000.
- [3] Internet Mapping Project, Lumeta Corp., <http://research.lumeta.com/ches/map/>
- [4] CAIDA skitter Project., <http://www.caida.org/>
- [5] Opte Project, <http://www.opte.org/>
- [6] Autonomous Systems (ASs) use the Border Gateway Protocol (BGP) to route packets between networks. BGP routing tables represent logical domain adjacencies rather than physical connectivity. Tables contain advertised paths through this logical network; routing within domains is typically done via Open Shortest Path First or another distributed spanning tree protocol.
- [7] L. Amini, A. Shaikh, and H. Schulzrinne, in Internet Performance and Control of Network Systems III, Proceedings of SPIE ITCOM, 2002.
- [8] University of Oregon Route-Views Project, <http://www.anc.uoregon.edu/route-views/>
- [9] M. Faloutsos, P. Faloutsos, and C. Faloutsos, in ACM SIGCOMM Computer Commun. Rev. **29**, 251 (1999).
- [10] N. Spring, R. Mahajan, and D. Wetherall, ACM SIGCOMM Computer Commun. Rev. **32**, 133 (2002).
- [11] Traceroute Organization, <http://www.traceroute.org/>
- [12] Q. Chen *et al.*, in Proceedings of IEEE INFOCOM, New York, New York, 2002.
- [13] P. Barford, A. Bestavros, J. Byers, and M. Crovella, in Proceedings of the ACM SIGCOMM Internet Measurement Workshop, San Francisco, CA, 2001.
- [14] R. Teixeira, K. Marzullo, S. Savage, and G.M. Voelker, in Proceedings of the ACM SIGCOMM Internet Measurement Workshop, Miami Beach, FL, 2003.
- [15] A. Lakhina, J. Byers, M. Crovella, and P. Xie, in Proceedings of IEEE INFOCOM, San Francisco, CA, 2003.
- [16] P. Erdős and A. Rényi, Publ. Math. Inst. Hung. Acad. Sci. **5**, 17 (1960).
- [17] T. Petermann and P. De Los Rios, Eur. Phys. J. B **38**, 201 (2004).
- [18] L. Dall’Asta *et al.*, cond-mat/0406404.
- [19] While there are typically multiple shortest paths from the source to a given destination, their union forms a tree if we break ties, e.g., by giving each vertex an index and taking the first shortest path in lexicographic order.
- [20] L. A. Braunstein *et al.*, Phys. Rev. Lett. **91**, 168701 (2003).
- [21] N. C. Wormald, Ann. Appl. Probab. **5**, 1217 (1995).
- [22] A.-L. Barabási and R. Albert, Science **286**, 509 (1999).
- [23] B. Bollobás, *Random Graphs* (Academic Press, London, New York, 1985).
- [24] Z. Toroczkai and K. E. Bassler, Nature (London) **428**, 716 (2004).
- [25] P. J. Macdonald, E. Almaas, and A.-L. Barabasi (to be published).