

Influence of the link weight structure on the shortest path

Piet Van Mieghem and Stijn van Langen

Delft University of Technology, P.O. Box 5031, 2600 GA Delft, The Netherlands

(Received 28 October 2004; published 20 May 2005)

The shortest path tree rooted at a source to all other nodes is investigated in a graph with polynomial link weights tunable by the power exponent α . By varying α , different types of shortest path trees, in short α trees, appear. Especially, the $\alpha \rightarrow 0$ regime that corresponds to heavily fluctuating link weights possesses a peculiar type of tree. The most important properties of this $\alpha \rightarrow 0$ tree are derived in the asymptotic limit for large N . The application of the theoretical insights to real networks (such as the Internet) are discussed: steering flow by adjusting link weights (traffic engineering), sensitivity of link weights and modeling of the network by α trees.

DOI: 10.1103/PhysRevE.71.056113

PACS number(s): 89.75.Fb

I. INTRODUCTION

The simple shortest path problem asks for the computation of the path from a source to a destination node that minimizes the sum of the weights of its constituent links. The related shortest path tree (SPT) is the union of the shortest paths from a source node to all other nodes in the graph. The SPT belongs to the fundamentals of graph theory and has many applications. Moreover, powerful shortest path algorithms like that of Dijkstra exist. Nevertheless, little seems known about the influence of the link weight structure on the properties of the SPT. The motivation to study the impact of the link weights on the path properties arose in multiconstrained routing [1]. Also from a traffic engineering perspective, a network operator may want to tune the weight of each link such that the resulting shortest paths between a particular set of ingresses and egresses follows the desirable routes in his network. Thus, apart from the topology of the graph, the link weight structure clearly plays an important role. Often as in complex molecules, global social interactions or other large infrastructures as the Internet, both the topology and the link weight structure are not accurately known. This uncertainty about the precise structure leads us to consider both the underlying graph and each of the link weights as random variables.

Since the shortest path tree problem is mainly sensitive to the smaller, non-negative link weights, the probability distribution of the link weights around zero will dominantly influence the properties of the resulting shortest path tree. A *regular* link weight distribution $F_w(x) = \Pr[w \leq x]$ has a Taylor series expansion around $x=0$,

$$F_w(x) = f_w(0)x + O(x^2)$$

since $F_w(0)=0$ and $F'_w(0)=f_w(0)$ exists. A regular link weight distribution is thus linear around zero. The factor $f_w(0)$ only scales all link weights, but it does not influence the shortest path. The simplest distribution of the link weight w with a distinct different behavior for small values is the polynomial distribution

$$F_w(x) = x^\alpha \mathbf{1}_{x \in [0,1]} + \mathbf{1}_{x \in [1,\infty)}, \quad \alpha > 0, \quad (1)$$

where the indicator function $\mathbf{1}_x$ is one if x is true else it is zero. The corresponding density is $f_w(x) = \alpha x^{\alpha-1}$, $0 < x < 1$.

The average and variance of the link weight are $\mathbb{E}[w] = \alpha/(\alpha+1)$ and $\text{var}[w] = \alpha/(\alpha+2) - [\alpha/(\alpha+1)]^2$, respectively. The exponent α is called the *extreme value index* of the probability distribution of w and $\alpha=1$ for regular distributions. By varying the exponent α over all non-negative real values, any extreme value index can be attained and a large class of corresponding shortest path trees, in short α -trees, can be generated. Finally, we assume independence of link weights which we deem a reasonable assumption in most large networks.

The main purpose of this paper is to show that, by considering a polynomial link weight structure tunable by one parameter α on, e.g., the complete graph, a broad class of shortest path trees are found as function of α . Instead of worrying about the precise topology of the graph, we show that a tunable link weight structure thins out the complete graph to the extent that a specific shortest path tree can be constructed. The approach thus eliminates the precise knowledge of the underlying graph by immediately concentrating on the tree properties induced by polynomial link weights.

Secondly, we show that relatively small variations in the link weight structures cause large differences in the properties of the SPT. In particular, the average hop count (i.e., the number of links or number of nodes minus 1 in a path) in a graph with N nodes follows a different scaling: $\mathbb{E}[H_N] = O(\ln N)$ for α around 1 while $\mathbb{E}[H_N] = O(N^{1/3})$ if $\alpha \rightarrow 0$. Moreover, for $\alpha \rightarrow 0$, all traffic in the graph routed along the SPT traverses precisely $N-1$ links and among those links there seems a large difference in the “critical backbone” over which nearly all traffic flows and the other links. The $\alpha \rightarrow 0$ regime corresponds to a strong disorder regime. In this article, many properties of the resulting tree for $\alpha \rightarrow 0$ will be derived.

The paper is outlined as follows. We first overview three types of α trees for $\alpha \rightarrow \infty$, 1, and 0 and introduce a critical $\alpha_c > 0$ for which all $\alpha < \alpha_c$ are $\alpha \rightarrow 0$ trees (with overwhelming probability). The major part (Sec. III) of the paper consists of computing properties of $\alpha \rightarrow 0$ trees. The potential applications of α trees in the setting of Internet are discussed in Sec. IV.

II. SPECIAL α TREES

Let us consider a graph $G(N, L)$ with N nodes and L links and with independent polynomial link weights specified by

Eq. (1). The shortest path tree rooted at an arbitrary node in the graph to all other nodes is called the α tree. In this section, we will meet three special α trees for $\alpha=\infty$, 1, and 0, respectively.

A. The case $\alpha \rightarrow \infty$

If $\alpha \rightarrow \infty$, it follows from Eq. (1) that $w=1$ for all links. Since all link have unit weight, the $\alpha \rightarrow \infty$ regime reduces to the computation of the shortest path tree in the underlying graph. The $\alpha \rightarrow \infty$ regime is thus entirely determined by the topology of the graph because the link weight structure does not differentiate between links. In case of the complete graph and $\alpha \rightarrow \infty$, the α tree is a star. In other underlying graphs, it is more difficult (see, e.g., Ref. [2]) to determine properties of the shortest path tree. In this paper, the $\alpha \rightarrow \infty$ regime is not further considered.

B. The case $\alpha=1$

If $\alpha=1$, the link weights are iid uniformly distributed. Earlier [3], it was shown that the shortest path tree in the complete graph with uniform (or exponential) link weights is precisely a uniform recursive tree (URT) while a URT is asymptotically the shortest path tree in the Erdős-Rényi random graph $G_p(N)$ [4] with link density p above the disconnectedness threshold $p_c \sim \ln N/N$. The interest of the URT is that analytic modeling is possible such as the computation of the degree, the hop count [3,5], the number of links in the URT or the multicast problem [6,7] and the hop count to the most nearby server or the anycast problem [8]. Comparison with Internet measurements (see Sec. IV) shows that properties of the shortest path tree in the regime around $\alpha=1$ agree reasonably well with those measured, but they remain first order estimates.

In an earlier analysis [3], it was shown that, for N large and fixed α , the hop count H_N satisfies

$$\mathbb{E}[H_N] \sim \frac{\ln N}{\alpha}, \tag{2}$$

$$\text{var}[H_N] \sim \frac{\ln N}{\alpha^2}, \tag{3}$$

and that the hop count is (at least asymptotically) independent of the link density p . That analysis has assumed independence of the links in the shortest path. Also, for large N , it follows from Eqs. (2) and (3) that

$$\alpha \sim \frac{\mathbb{E}[H_N]}{\text{var}[H_N]} \tag{4}$$

which gives a method to relate α to Internet measurements (see, e.g., Fig. 9) provided that α is around 1.

C. The case $\alpha \rightarrow 0$

If $\alpha \rightarrow 0$,

$$\frac{\sqrt{\text{var}[w]}}{\mathbb{E}[w]} \sim \frac{1}{\sqrt{\alpha}}$$

which means that, in this limit, the link weights possess strong fluctuations and that the above ratio diverges. This observation inspired by Braunstein *et al.* [9] is crucial in the analysis of the behavior of the shortest path for small α . If the set of L link weights $\{w_k\}_{1 \leq k \leq L}$ is ordered as

$$w_{(1)} < w_{(2)} < \dots < w_{(L)}$$

it is of interest to know if there exist a critical value $\alpha_c > 0$ such that, for all $\alpha < \alpha_c$, the following inequalities are obeyed (with high probability):

$$w_{(2)} > w_{(1)},$$

$$w_{(3)} > w_{(2)} + w_{(1)},$$

$$w_{(4)} > w_{(3)} + w_{(2)} + w_{(1)}, \dots,$$

$$w_{(k+1)} > \sum_{j=1}^k w_{(j)}, \dots,$$

$$w_{(L)} > \sum_{j=1}^{L-1} w_{(j)}. \tag{5}$$

In any graph with this link weight structure, the weight of a path $w(\mathcal{P}) = \sum_{i \rightarrow j \in \mathcal{P}} w(i \rightarrow j)$ consists of one link that dominates all others. Indeed, by ordering the set $\{w(i \rightarrow j)\}_{i \rightarrow j \in \mathcal{P}}$ and assuming that $\max_{i \rightarrow j \in \mathcal{P}} \{w(i \rightarrow j)\} = w_{(k)}$, then the inequalities above indicate that $w(\mathcal{P}) < w_{(k+1)}$. Hence, the set of inequalities assure that the weight of a path can be upper bounded by the next in order link weight higher than its own maximum link weight. Any path in the graph that does not contain the link with weight $w_{(L)}$ possesses a path weight that is always smaller than $w_{(L)}$.

Theorem 1. If the link weights obey the inequalities (5), the union of all shortest path trees rooted at each node in the graph is a tree.

Proof. Suppose that the shortest path graph G_U consisting of the union of all shortest paths would contain a loop, then the largest link weight of that loop is larger than the sum of all other links of that path. Since subsections of the shortest paths are again shortest paths, that largest link weight in the loop will not appear in a shortest path because it is bypassed by the other part of the loop with smaller link weight. Hence, loops do not occur in the graph G_U , and a loop-free graph is a tree.

This curious property implies that all shortest path trees rooted at different nodes are precisely the same and equal to the minimum spanning tree (MST). The coincidence of all SPT's with the MST for a general graph with strongly fluctuating weights has been reported by Dobrin and Duxbury [10].

We will now show that the inequalities (5) can be satisfied provided $\alpha_c = O(L^{-2})$. It is convenient to make the transformation $w = e^{-V/\alpha}$. Since $w \in [0, 1]$, the random variable V

ranges over all positive real numbers and V is distributed as

$$\Pr[e^{-V/\alpha} \leq x] = x^{\alpha} 1_{x \in [0,1]}$$

or, since the event $\{w = e^{-V/\alpha} \leq x\}$ is equivalent to the event $\{V \geq -\alpha \ln x\}$, and denoting $y = -\alpha \ln x$,

$$\Pr[V \geq y] = e^{-y},$$

which shows that the random variable V is exponentially distributed with mean 1. The set of inequalities generated by $w_{(k+1)} > 2w_{(k)}$ for each link $k \in \{L\}$ is more severe than the original set of inequalities (5). In other words, if there exists an α_c such that $w_{(k+1)} > 2w_{(k)}$ for each link $k \in \{L\}$ is obeyed, then also the original set of inequalities is satisfied. Since

$$\frac{w_{(k+1)}}{w_{(k)}} = \exp\left(\frac{1}{\alpha}(V_{(k)} - V_{(k+1)})\right) > 2$$

the inequalities are all satisfied if the minimum spacing $\Delta_{\min} \equiv \min|V_{(k)} - V_{(k+1)}| > \alpha \ln 2$. We further replace the exponential distribution $\Pr[V \leq y] = 1 - e^{-y}$ with the (qualitatively similar) uniform distribution on $[0, 1]$. The L random variables $V_{(k)}$ constitute a Poisson process and the spacings Δ between two consecutive V 's are independent and identically distributed as $\Pr[\Delta > t] = (1-t)^L$ with average spacing $1/(L+1) \approx 1/L$ (see, e.g., Ref. [11], Chap. 7). For large L , we have $\Pr[L\Delta > t] = (1-t/L)^L \rightarrow e^{-t}$ and the spacing Δ tends to the exponential distribution $f_{\Delta}(x) = L \exp(-Lx)$. The distribution of the minimum spacing Δ_{\min} is given by

$$\Pr[\Delta_{\min} > \alpha \ln 2] = [\Pr[\Delta > \alpha \ln 2]]^L = \exp(-\alpha L^2 \ln 2).$$

In conclusion, by choosing $\alpha_c < L^{-2}/\ln 2$, the probability of a violation of the inequalities $w_{(k+1)} > 2w_{(k)}$ for each link $k \in \{L\}$ is less than one. Although this probability rapidly decays, only if $\alpha \rightarrow 0$, the inequalities are surely satisfied. The argument has shown that there is indeed a range of $0 < \alpha < \alpha_c$ that obeys the inequalities with high probability. However, it merely demonstrates the existence of $\alpha_c > 0$. Moreover, the resulting estimate $\alpha_c = O(L^{-2})$ is much too conservative because (a) the inequalities $w_{(k+1)} > 2w_{(k)}$ for each link $k \in \{L\}$ are more stringent than the set of inequalities (5), (b) even the set of inequalities (5) is likely to impose too many restrictions, in particular on the links with larger link weights $w_{(k)}$ that are already unlikely to be part of a shortest path. Consequently, the extreme $\alpha \rightarrow 0$ regime is expected to be entered more rapidly than for $\alpha_c = O(L^{-2})$.

A better estimate for α_c should be derived from the joint density function f of order statistics $\{w_{(k)}\}$ of the set of L link weights which is

$$f(w_{(1)} = x_1, \dots, w_{(L)} = x_L) = L! \prod_{j=1}^L f_w(x_j) \times 1_{x_1 < x_2 < \dots < x_L}.$$

The factor $L!$ counts the number of permutations in which the set $\{w_j\}_{1 \leq j \leq L}$ of iid random variables w_j can be ordered to obtain the ordered set $\{w_{(j)}\}_{1 \leq j \leq L}$. Thus,

$$p_{\alpha} = \Pr \left[w_{(2)} > w_{(1)}, w_{(3)} > w_{(2)} + w_{(1)}, \dots, w_{(L)} > \sum_{j=1}^{L-1} w_{(j)} \right] \quad (6)$$

and for polynomially link weights,

$$p_{\alpha} = \alpha^L L! \int_0^1 dx_1 x_1^{\alpha-1} \int_{x_1}^1 dx_2 x_2^{\alpha-1} \int_{(x_1+x_2)1_{\{x_1+x_2 \leq 1\}}}^1 dx_3 x_3^{\alpha-1} \\ \times \dots \int_{(\sum_{j=1}^{L-1} x_j)1_{\{\sum_{j=1}^{L-1} x_j \leq 1\}}}^1 dx_L x_L^{\alpha-1}. \quad (7)$$

The largest value of α that solves the equation $p_{\alpha} = 1 - \epsilon$ for an arbitrarily small $\epsilon > 0$ can be obtained. Although we believe that the integral in (7) can be solved exactly, the result seems too cumbersome to be useful here. If all link weights are iid exponential random variables with a same mean, we find that

$$\Pr \left[w_{(2)} > w_{(1)}, w_{(3)} > w_{(2)} + w_{(1)}, \dots, w_{(L)} > \sum_{j=1}^{L-1} w_{(j)} \right] \\ = \frac{L!}{2^{L(L-1)/2}} \quad (8)$$

which shows that for independent exponential link weights in any graph the inequalities (5) are almost never satisfied for large L . Since the exponential has the same extremal index as the uniform distribution $\alpha=1$, we see that to satisfy the inequalities (5), the link weight structure must be highly fluctuating. If operators in the Internet assign link weights inversely proportional to the bandwidth [12], it is a priori not unlikely that the inequalities (5) are almost all satisfied because of the larger heterogeneity in links (roughly from 10 kbit/s up to 10 Gbit/s) in the Internet.

In summary, three distinct α regimes have been identified. Values of α around either ∞ , 1, and 0 are likely to yield α -trees with properties that are essentially the same as those of the shortest hop count tree, the URT or the MST. After a study of properties of the MST ($\alpha \rightarrow 0$ tree), we will motivate in Sec. III D a slightly better order estimate for $\alpha_c = O(N^{-2} \ln^{-2} N)$ rather than $\alpha_c = O(N^{-4})$ in worst case where $L = O(N^2)$.

D. Simulations

Figure 1 shows the average hop count $\mathbb{E}[H_N]$ as a function of the number of nodes N for various values of the exponent $\alpha \in [0, 1)$. For α near to 1, the average hop count $\mathbb{E}[H_N]$ scales linearly in $\ln N$, in agreement with Eq. (2). For smaller α , the average hop count $\mathbb{E}[H_N]$ increases faster than $\ln N$. If $\alpha \rightarrow 0$, the strong overlap (dependence) of paths (all shortest path trees are equal to the MST) radically excludes assumptions of independence made in the earlier analysis [3] that lead to a logarithmic scaling in N of the hop count $\mathbb{E}[H_N]$.

The simulations seem to indicate that α_c may be much larger than estimated in previous section which implies a higher probability that they really appear in practice. Braunstein *et al.* [9] investigated the lengths of minimum-weight

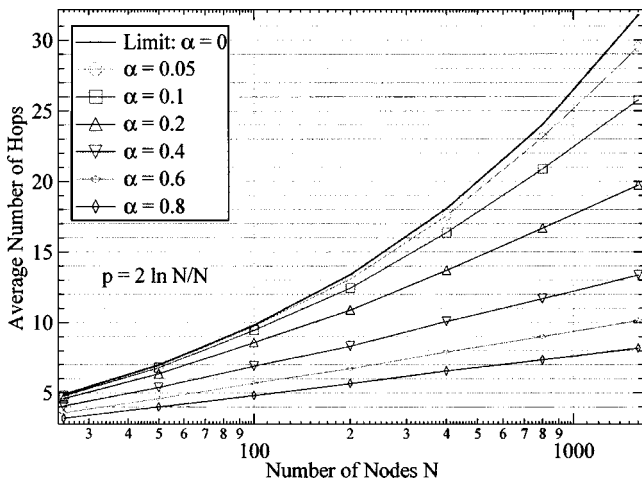



FIG. 1. The average hop count $E[H_N]$ in the random graph $G_p(N)$ with link density $p=2 \ln N/N$ as a function of the number of nodes N .

paths in the regime of large link weight fluctuations ($\alpha \rightarrow 0$) in random graphs, small-world and scale-free graphs. Using arguments from the theory of critical phenomena and numerical simulations they showed that for Erdős-Rényi random graphs the hop count scales as $\lim_{\alpha \rightarrow 0} E[H_N(\alpha)] = O(N^{1/3})$.

Figure 2 visualizes the different structure of a typical MST (a) and a typical URT (b) of the same size $N=100$. Figure 3 shows the probability that the union of all shortest



(a)





FIG. 3. The probability that the union of all shortest paths $G_{Uspst(\alpha)}$ in the complete graph with polynomial link weights is a minimum spanning tree as a function of α .

paths $G_{Uspst(\alpha)}$ between all node pairs in the complete graph with polynomial link weights is a minimum spanning tree. In real networks where almost all flows follow shortest paths through the network, the union of all shortest paths is the observable part of a network. For example, the union of all trace routes between all node pairs in the Internet, would represent the observable graph of the Internet. The real Internet is larger because it also contains dark links for backup paths needed in case of failures. We observe in Fig. 3 a phase



(b)

FIG. 2. An example in the graph with $N=100$ nodes of (a) the MST which is the SPT for $\alpha=0$ and (b) the URT which is the SPT for $\alpha=1$. Both trees are structured per level set where each level shows the number of nodes at different hop count from the root (here node with label 7).

transition around α_c , which we here define as $\Pr[G_{\text{Usp}t(\alpha_c)} = \text{MST}] = \frac{1}{2}$. For $\alpha < \alpha_c$, most graphs $G_{\text{Usp}t(\alpha)}$ are trees with high probability while for $\alpha > \alpha_c$ hardly any graph $G_{\text{Usp}t(\alpha)}$ is a tree. The width $\Delta\alpha$ of the phase transition (e.g., defined as $\Pr[G_{\text{Usp}t(\alpha_c - \Delta\alpha)} = \text{MST}] - \Pr[G_{\text{Usp}t(\alpha_c + \Delta\alpha)} = \text{MST}] = 0.9$) decreases with N . The simulations are limited to relatively small graphs because for small α the large relative variations in the link weights require a Dijkstra shortest path algorithm that runs with arbitrarily long real numbers (mantisse and exponent). Indeed, a polynomial random variable is generated as $U^{1/\alpha} = \exp(\ln U / \alpha)$ where U is a uniform random variable that is easy to generate with computers. For a small α of about 10^{-3} and for a typical value of U around $\frac{1}{2}$, we see that link weights appear of the order of 2^{-1000} and smaller.

III. PROPERTIES OF MINIMUM SPANNING TREES

As mentioned in Sec. II B, the properties of the URT ($\alpha = 1$) have been investigated earlier in detail. We devote this section to study some properties of MSTs corresponding to the $\alpha \rightarrow 0$ regime. Although some results on MSTs have been found earlier in a different setting with different methods, we present here a unified and elegant approach in the asymptotic regime for $N \rightarrow \infty$.

A. Earlier work

Frieze [13] computed the average weight of an MST on the complete graph K_N for a general weight distribution with finite $f_w(0)$. He also showed that the variance vanishes asymptotically. Aldous [14] has computed the distribution of the degree in a (single sample) MST. He has generalized Frieze’s result for the MST weight to a more general class of polynomial link weight distributions. Janson [15] shows that the distribution of the MST weight on K_N with uniform weights is asymptotically normal, and he gives an expression for the variance. Penrose [16] proves that the degree distribution of the Euclidean MST on a d -dimensional hypercube converges to Aldous’ result, Eq. (24) below, in the mean-field limit $d \rightarrow \infty$.

Barabasi [17] clarified the equivalence between invasion percolation and Prim’s algorithm for the MST (or strong-disorder SPT). Invasion percolation has been widely studied by physicists working in the area of phase transitions usually in two or three dimensions (rather than in the infinite-dimensional case that we study). It models the penetration of fluid in a porous medium saturated by another fluid. The link weight plays the role of a potential barrier for fluid to invade into a pore. Dobrin [10] shows that the MST geometry on a random network is universal (i.e., does not depend on the energy or weight distribution), which yields a simple way to compute the MST weight for general energy distributions.

B. The Kruskal growth process of the MST

Since the link weights in the underlying complete graph are chosen independently and assigned randomly to links in the complete graph, the resulting graph is probabilistically the same if we first order the set of link weights and assign

them in increasing order randomly to links in the complete graph. In the latter construction process, only the order statistics or the ranking of the link weights suffice to construct the graph because the precise link weight can be unambiguously associated (later) to the rank of a link. Hence, assume the existence of a set of $L = \binom{N}{2}$ iid polynomially link weights with $\alpha \rightarrow 0$ that are ordered, then we need only to take the rank of each link weight into account to construct the MST which coincides with the shortest path tree in the case $\alpha \rightarrow 0$. This observation immediately favors the Kruskal algorithm [18] for the MST over Prim’s algorithm. Although the Prim algorithm leads to the same MST, it gives a more complicated, long-memory growth process, where the attachment of each new node depends stochastically on the whole growth history so far. Pietronero and Schneider [19] illustrate that in our approach Prim, in contrast with Kruskal, leads to a very complicated stochastic process for the construction of the MST.

The Kruskal growth process described here is closely related to a growth process of the random graph $G(N, L)$ with N nodes and L links. The construction or growth of $G(N, L)$ starts from N individual nodes and in each step an arbitrary, not yet connected random pairs is connected. The only difference with Kruskal’s algorithm for the MST is that, in Kruskal, links generating loops are forbidden. Those forbidden links are the links that connect nodes within the same connected component or “cluster.” As a result, the internal wiring of the clusters differs, but the cluster size statistics (counted in nodes, not edges) is exactly the same as in the corresponding random graph. The metacode of the Kruskal growth process for the construction of the random $\alpha \rightarrow 0$ trees is

```

Kruskal GrowthMST
1. start with N disconnected nodes
2. repeat until all nodes are connected
3.   randomly select a node pair (i,j )
4.   if a path  $P_{i \rightarrow j}$  does not exist
5.     then connect i to j
    
```

We now relate the link density p in the random graph $G_p(N)$ to the link density in the corresponding stage of the Kruskal growth process. We first compute the size of the giant cluster in the forest as a function of the number of links added. Let $S = \Pr[n \in C]$ denote the probability that a node n belongs to the giant component C . If $n \notin C$, then none of the neighbors of node n belongs to the giant component. The number of neighbors of a node n is the degree $d(n)$ of a node such that

$$\begin{aligned} \Pr[n \notin C] &= \Pr[\text{all neighbor of } n \notin C] \\ &= \sum_{k \geq 0} \Pr[\text{all } k \text{ neighbors of } n \notin C | d(n) = k] \\ &\quad \times \Pr[d(n) = k]. \end{aligned}$$

Since in $G_p(N)$ all neighbors of n are independent, the conditional probability becomes with $1 - S = \Pr[n \notin C]$,

$$\begin{aligned} \Pr[\text{all } k \text{ neighbors of } n \notin C | d(n) = k] &= (\Pr[n \notin C])^k \\ &= (1 - S)^k. \end{aligned}$$

Moreover, this probability holds for any node in $n \in G_p(N)$ such that, writing the random variable $D_{RG} = d(n)$,

$$1 - S = \sum_{k=0}^{\infty} (1 - S)^k \Pr[d(n) = k] = \varphi_{D_{RG}}(1 - S),$$

where $\varphi_{D_{RG}}(u) = \mathbb{E}[u^{D_{RG}}]$ is the generating function of the degree D_{RG} in $G_p(N)$. For large N , the degree distribution in $G_p(N)$ is Poisson distributed with mean degree $\mu_{RG} = p(N - 1)$ and $\varphi_{D_{RG}}(u) = e^{\mu_{RG}(u-1)}$. Hence, for large N , the fraction S of nodes in the giant component in the random graph is given by

$$S = 1 - e^{-\mu_{RG}S} \quad (9)$$

and the average size of the giant component is NS . For $\mu_{RG} < 1$ the only solution is $S=0$ whereas for $\mu_{RG} > 1$ there is a nonzero solution for the size of the giant component. The solution can be expressed as a Lagrange series (Ref. [20], p. 94),

$$S(\mu_{RG}) = 1 - e^{-\mu_{RG}S} \sum_{n=0}^{\infty} \frac{(n+1)^n}{(n+1)!} (\mu_{RG} e^{-\mu_{RG}})^n. \quad (10)$$

By reversing Eq. (9), the average degree in the random graph can be expressed in terms of the fraction S of nodes in the giant component

$$\mu_{RG}(S) = -\frac{\ln(1-S)}{S}. \quad (11)$$

We will now transform the mean degree μ_{RG} in the random graph to the mean degree μ_{MST} in the corresponding stage of the MST. In early stages of the growth each selected link will be added with high probability such that $\mu_{MST} = \mu_{RG}$ almost surely. After some time the probability that a selected link is forbidden increases, and thus μ_{RG} exceeds μ_{MST} . In the end, when connectivity of all N nodes is reached, $\mu_{MST} = 2$ (since it is a tree) while $\mu_{RG} = O(\ln N)$.

Consider now an intermediate stage of the growth as illustrated in Fig. 4. Assume there is a giant component of average size NS and $n_c = N(1-S)/s_c$ small components of average size s_c each. Then we can distinguish six types of links labeled a-f in the figure. Types a and b are links that have been chosen earlier in the giant component (a) and in the small components (b), respectively. Types c and d are eligible links between the giant component and a small component (c) and between small components (d), respectively. Types e and f are forbidden links connecting nodes within the giant component (e), respectively, within a small component (f). For large N , we can enumerate how many links there are of each type L_k with $k = \{a, b, c, d, e, f\}$:

$$L_a + L_b = \frac{1}{2} \mu_{MST} N,$$

$$L_c = SN(1-S)N,$$




FIG. 4. Component structure during the Kruskal growth process.

$$L_d = \frac{1}{2} n_c^2 s_c^2,$$

$$L_e = \frac{1}{2} (SN)^2 - SN,$$

$$L_f = \frac{1}{2} n_c s_c (s_c - 1) - n_c (s_c - 1).$$

To highest order in $O(N^2)$, we have

$$L_c = N^2 S(1-S),$$

$$L_d = \frac{1}{2} N^2 (1-S)^2,$$

$$L_e = \frac{1}{2} N^2 S^2.$$

The probability that a randomly selected link is eligible is $q = (c+d)/(c+d+e+f)$ or

$$(L_c + L_d)/(L_c + L_d + L_e + L_f) \quad q = 1 - S^2. \quad (12)$$

In contrast with the growth of the random graph $G_p(N)$ where at each stage a link is added with probability p , in the Kruskal growth of the MST we are only successful to add one link (with probability 1) per $1/q$ stages on average. Thus the average number of links added in the random graph corresponding to one link in the MST is $1/q = 1/(1-S^2)$. This provides an asymptotic mapping between μ_{RG} and μ_{MST} in the form of a differential equation

$$\frac{d\mu_{RG}}{d\mu_{MST}} = \frac{1}{1-S^2}.$$

By using Eq. (11), we find




FIG. 5. Size of the giant component (divided by N) as a function of the mean degree μ_{MST} . Each simulation for different number of nodes N consists of 1 MST sample.

$$\frac{d\mu_{\text{MST}}}{dS} = \frac{d\mu_{\text{MST}}}{d\mu_{\text{RG}}} \frac{d\mu_{\text{RG}}}{dS} = \frac{(1+S)[S+(1-S)\ln(1-S)]}{S^2}.$$

Integration with the initial condition $\mu_{\text{MST}}=2$ at $S=1$, finally gives the average degree μ_{MST} in the MST as function of the fraction S of nodes in the giant component

$$\mu_{\text{MST}}(S) = 2S - \frac{(1-S)^2}{S} \ln(1-S). \quad (13)$$

As shown in Fig. 5, this result agrees well with the simulation (even for a single sample), except in a small region around the transition $\mu_{\text{MST}}=1$ and for relatively small N .

The key observation is that all transition probabilities in the Kruskal growth process asymptotically depend on only one parameter, the fraction of nodes in the giant component S . This quantity S is called an *order parameter* in statistical physics. In general, the expectation of an order parameter distinguishes the qualitatively different regimes (states) below and above the phase transition. In higher dimensions fluctuations of the order parameter around the mean can be neglected and the mean value can be computed from a self-consistent mean-field theory. In our problem, the underlying complete (or random) graph topology makes the problem effectively infinite dimensional. The argument leading to Eq. (9) is essentially a mean-field argument.

C. The average weight of the minimum spanning tree

The crucial observation [10] is that in any graph, the topology of a MST depends only on the ranks of the link weights. By definition, the weight of the MST is

$$W_{\text{MST}} = \sum_{j=1}^L w_{(j)} 1_{j \in \text{MST}}, \quad (14)$$

where $w_{(j)}$ is the j th smallest link weight. The average weight of the MST is

$$\mathbb{E}[W_{\text{MST}}] = \sum_{j=1}^L \mathbb{E}[w_{(j)} 1_{j \in \text{MST}}].$$

The random variables $w_{(j)}$ and $1_{j \in \text{MST}}$ are independent because the value $w_{(j)}$ of j th smallest link only depends on the link weight distribution and the number of edges L while the appearance $1_{j \in \text{MST}}$ of the j th edge in the MST only depends on the graph's topology. Hence,

$$\mathbb{E}[w_{(j)} 1_{j \in \text{MST}}] = \mathbb{E}[w_{(j)}] \mathbb{E}[1_{j \in \text{MST}}] = \mathbb{E}[w_{(j)}] \text{Pr}[j \in \text{MST}]$$

such that the average weight of the MST is

$$\mathbb{E}[W_{\text{MST}}] = \sum_{j=1}^L \mathbb{E}[w_{(j)}] \text{Pr}[j \in \text{MST}]. \quad (15)$$

In general for independent link weights with probability density function $f_w(x)$ and distribution function $F_w(x) = \text{Pr}[w \leq x]$, the density function of the j th order statistic (Ref. [11], Chap. 3) is

$$f_{w_{(j)}}(x) = \frac{j f_w(x)}{F_w(x)} \binom{L}{j} [F_w(x)]^j [1 - F_w(x)]^{L-j}. \quad (16)$$

The factor $\binom{L}{j} [F_w(x)]^j [1 - F_w(x)]^{L-j}$ is a binomial distribution with mean $\mu = F_w(x)L$ and variance $\sigma^2 = LF_w(x)[1 - F_w(x)]$ that, by the central limit theory, tends for large L to a Gaussian $1/(\sigma\sqrt{2\pi})e^{-(j-\mu)^2/2\sigma^2}$ which peaks at $j = \mu$. Thus for large N and fixed j/L , we have [21] $x_j = \mathbb{E}(w_{(j)}) \approx F_w^{-1}(j/L)$.

For the complete graph, we found before in Eq. (12) that the link ranked j appears in the MST with probability

$$\text{Pr}[j \in \text{MST}] = 1 - S_j^2,$$

where S_j is the fraction of nodes in the giant component during the construction process of the random graph at the stage where the number of links precisely equals j . Since links are added independently, that stage in fact establishes the random graph $G(N, L=j)$. With Eq. (11) and $\mu_{\text{RG}} = 2L/N$, it follows that

$$\frac{2j}{N} = - \frac{\ln(1 - S_j)}{S_j}. \quad (17)$$

Hence,

$$\mathbb{E}[W_{\text{MST}}] \approx \sum_{j=1}^L F_w^{-1}\left(\frac{j}{L}\right) (1 - S_j^2).$$

We now approximate the sum by an integral

$$\mathbb{E}[W_{\text{MST}}] \approx \int_1^L F_w^{-1}\left(\frac{u}{L}\right) (1 - S_u^2) du.$$

Substituting $x=2u/N$ [which is the average degree in $G(N, u)$] yields for large N , where $L \approx N^2/2$,

$$\begin{aligned} \mathbb{E}[W_{\text{MST}}] &\approx \frac{N}{2} \int_{2/N}^{N-1} F_w^{-1}\left(\frac{x}{N}\right) (1 - S_{(N/2)x}^2) dx \\ &\approx \frac{N}{2} \int_0^N F_w^{-1}\left(\frac{x}{N}\right) [1 - S^2(x)] dx. \end{aligned}$$

It is known [22] that if the number of links (edges) in the growth process of the random graph is below $N/2$, with high probability (and ignoring a small onset region just below $N/2$), there is no giant component such that $S(x)=0$ for $x \in [0, 1]$. Thus, we arrive at the general formula valid for large N ,

$$\mathbb{E}[W_{\text{MST}}] \approx \frac{N}{2} \int_0^1 F_w^{-1}\left(\frac{x}{N}\right) dx + \frac{N}{2} \int_1^N F_w^{-1}\left(\frac{x}{N}\right) [1 - S^2(x)] dx. \quad (18)$$

The first term is the contribution from the smallest $N/2$ links in the graph, which are included in the MST almost surely. The remaining part comes from the more expensive links in the graph, which are included with diminishing probability since $1 - S^2(x)$ decreases exponentially for large x as can be deduced from Eq. (10). The rapid decrease of $1 - S^2(x)$ makes only relatively small values of the argument $F_w^{-1}(x/N)$ contribute to the second integral.

At this point, the specifics of the link weight distribution needs to be introduced. The Taylor expansion of $(N/2)F_w^{-1}(x/N)$ for large N to first order is

$$\frac{N}{2} F_w^{-1}\left(\frac{x}{N}\right) = \frac{N}{2} F_w^{-1}(0) + \frac{x}{2f_w(0)} + O\left(\frac{1}{N}\right) = \frac{x}{2f_w(0)} + O\left(\frac{1}{N}\right)$$

since we require that link weights are positive such that $F_w^{-1}(0)=0$. This expansion is only useful provided $f_w(0)$ is neither zero nor infinity. These cases occur, e.g., for polynomial link weights with $f_w(x) = \alpha x^{\alpha-1}$ with $\alpha \neq 1$. Fortunately, for polynomial link weights $(N/2)F_w^{-1}(x/N) = (N^{1-1/\alpha}/2)x^{1/\alpha}$. Formally, this latter expression reduces to the first order Taylor approach for $\alpha=1$, apart from the constant factor $1/f_w(0)$. Therefore, we will first compute $\mathbb{E}[W_{\text{MST}}]$ for polynomial link weights and then return to the case in which the Taylor expansion is useful.

1. Polynomial link weights

The average weight of the MST for polynomial link weights follows [23] from Eq. (18) as

$$\mathbb{E}[W_{\text{MST}}(\alpha)] \approx \frac{N^{1-1/\alpha}}{2} \left(\frac{1}{1/\alpha + 1} + \int_1^N x^{1/\alpha} [1 - S^2(x)] dx \right).$$

Let $y=S(x)$ and use Eq. (11), then $x=S^{-1}(y)=-\ln(1-y)/y$ and $dx=-d/dy[\ln(1-y)/y]dy$ while $y=S(1)=0$ and $y=S(N)=1$, such that

$$\begin{aligned} I &= \int_1^N x^{1/\alpha} [1 - S^2(x)] dx \\ &= \int_0^1 \left(-\frac{\ln(1-y)}{y} \right)^{1/\alpha} (1-y^2) \frac{d}{dy} \left(-\frac{\ln(1-y)}{y} \right) dy. \end{aligned}$$

After partial integration, we have

$$I = -\frac{1}{\frac{1}{\alpha} + 1} + \frac{2}{\frac{1}{\alpha} + 1} \int_0^\infty x^{1/\alpha+1} \frac{e^{-x}}{(1-e^{-x})^{1/\alpha}} dx.$$

Finally, we end up with

$$\mathbb{E}[W_{\text{MST}}(\alpha)] \approx N^{1-1/\alpha} \left(\frac{1}{1/\alpha + 1} \int_0^\infty x^{1/\alpha+1} \frac{e^{-x}}{(1-e^{-x})^{1/\alpha}} dx \right). \quad (19)$$

If $\alpha < 1$, then $\mathbb{E}[W_{\text{MST}}(\alpha)] \rightarrow 0$ for $N \rightarrow \infty$, while for $\alpha > 1$, $\mathbb{E}[W_{\text{MST}}(\alpha)] \rightarrow \infty$. In particular, $\lim_{\alpha \rightarrow \infty} \mathbb{E}[W_{\text{MST}}(\alpha)] = N - 1$. Only for $\alpha=1$, $\mathbb{E}[W_{\text{MST}}(1)]$ is finite for large N . More precisely,

$$\mathbb{E}[W_{\text{MST}}(1)] = \zeta(3) = 1.202, \dots, \quad (20)$$

where we have used [Ref. [24], Eq. (23.2.7)] the integral of the Riemann Zeta function

$$\Gamma(s)\zeta(s) = \int_0^\infty \frac{u^{s-1}}{e^u - 1} du$$

convergent for $\text{Re}(s) > 1$. This particular case for $\alpha=1$ has been proved earlier by Frieze [13] based on a different method. Asymptotically, as shown in Ref. [6], the average weight of a shortest path tree is $\zeta(2) = \pi^2/6$, while here the average weight of the MST is $\zeta(3) < \zeta(2)$.

2. Generalizations

We now return to the Taylor series valid for link weights where $0 < f_w(0) < \infty$. The above result for $\alpha=1$ immediately yields

$$\mathbb{E}[W_{\text{MST}}] = \frac{\zeta(3)}{f_w(0)}. \quad (21)$$

This result is for the complete graph K_N . A random graph $G_p(N)$ with $p < 1$ and weight density $f_w(x)$ is equivalent to K_N with a fraction $1-p$ of infinite link weights. Thus the effective link weight distribution is $pf_w(x) + (1-p)\delta_{w,\infty}$, and we can simply replace $f_w(0)$ by $pf_w(0)$ in the expression (21) for the MST weight on K_N .

D. The largest link weight in a group

Consider a group (e.g., in multicasting) consisting of m members and one source node on the MST. Both the source and the m receiving group members are chosen uniformly among the N nodes of the MST. For large N , the m paths from source to each multicast group member have a same bottleneck almost surely [25] if and only if the source joins




FIG. 6. The PDF of the normalized rank R^* for various sizes m of the multicast group.

the giant component later than the m destinations which happens with probability $1/(m+1)$. We will now compute the rank of the largest link weight in the tree that spans the m members.

During the growth of the MST as explained in Sec. III B, at each step links are added and the last link that connects the m members and the source will be determined. When link weights are ranked in increasing order, the rank R of that last link corresponds to the R th order statistics $w_{(R)}$ whose distribution is specified in Eq. (16). We first determine the rank R which equals the step T in the Kruskal growth process for which $T = \max(t_1, t_2, \dots, t_m, t_{m+1})$, where t_j is the time in the growth process at which the j th multicast member is swallowed by the giant component and where t_{m+1} is the time at which the source joins the giant component. Since all m nodes and the source are uniformly selected during the growth process measured in terms of $S \in [0, 1]$, the random variables t_j are iid uniform on $[0, 1]$ and

$$\Pr[T \leq x] = x^{m+1}$$

for $x \in [0, 1]$. From the relation (11) of the average degree and the fraction of nodes in giant component S , the normalized rank R^* of the largest link weight follows from $R^* = (2/N)R = \mu_{RG}(S)|_{S=T}$ and the event $\{T \leq x\}$ is equivalent to $\{\mu_{RG}(T) \leq \mu_{RG}(x)\}$ because $\mu_{RG}(x)$ is monotonically increasing in x such that

$$\Pr[R^* \leq y] = [\mu_{RG}^{-1}(y)]^{m+1},$$

where $x = \mu_{RG}^{-1}(y)$ is the inverse function of $y = \mu_{RG}(x)$ and explicitly given in Eq. (10). The probability density function $f_{R^*}(y) = (d/dy)\Pr[R^* \leq y]$ is

$$f_{R^*}(y) = (m+1) \frac{[\mu_{RG}^{-1}(y)]^m}{\mu'_{RG}[\mu_{RG}^{-1}(y)]} \quad (22)$$

and is shown in Fig. 6 for $m=1, 2, 5$, and 10 . The mean is best computed from the tail probability formula

$$\begin{aligned} \mathbb{E}[R^*] &= \int_0^{N-1} (1 - \Pr[R^* \leq y]) dy = \int_0^1 (1 - x^{m+1}) d\mu_{RG}(x) \\ &= -(m+1) \int_0^1 x^{m-1} \ln(1-x) dx \\ &= (m+1) \sum_{k=1}^{\infty} \frac{1}{k(k+m)}. \end{aligned}$$

After partial fraction decomposition, the average normalized rank is

$$\mathbb{E}[R^*] = \frac{m+1}{m} \sum_{k=1}^m \frac{1}{k} \quad (23)$$

which shows that for large m , $\mathbb{E}[R^*] \approx \ln m$. In case $m=1$, the rank $R = (N/2)R^*$ is the rank of the largest link weight in the shortest path from the source to the only destination and has $\mathbb{E}[R^*] = 2$, thus $\mathbb{E}[R] = N$. For large N , the probability that $R < N/2$ (i.e., $R^* < 1$) is zero asymptotically because the first $N/2$ are almost surely not connected. Connectivity occurs asymptotically when the mean degree μ_{RG} exceeds 1. The probability distribution of the largest weight follows from Eqs. (16) and (22) using the law of total probability

$$\begin{aligned} f_{w_{(R^*)}}(x) &= \int_0^N f_{w_{(R^*)}}(x|R^* = y) d\Pr[R^* \leq y] \\ &= \frac{N f_w(x)}{2 F_w(x)} \int_1^N y \binom{L}{(N/2)y} [F_w(x)]^{(N/2)y} \\ &\quad \times [1 - F_w(x)]^{L - (N/2)y} f_{R^*}(y) dy. \end{aligned}$$

Observe that $w_{(R^*)}$ in contrast with R^* does depend on link probability p of the random graph. For large N , $f_{w_{(j)}}(x)$ tends to a Gaussian (as explained in Sec. III C) resulting in

$$f_{w_{(R^*)}}(x) \approx \frac{1}{\sigma \sqrt{2\pi}} \int_0^N e^{-[(N/2)y - \mu]^2 / 2\sigma^2} f_{R^*}(y) dy \approx f_{R^*}[NF_w(x)].$$

This expression shows that the probability density function (PDF) of the highest link weight in the multicast group is asymptotically distributed as the normalized rank with the rank parameter y in Eq. (22) replaced by $NF_w(x)$.

Formula (23) shows that the rank of the largest link weight in the MST ($m=N$) is about $\mathbb{E}[R_{\max}] \approx (N/2) \ln N$. Returning to the determination of α_c in Sec. II C, this observation suggests that the number of relevant equations in Eq. (5) is $L = O(N \ln N)$. In fact, if $N \rightarrow \infty$, the claim is very likely correct which leads to the critical regime $\alpha_c = O(N^{-2} \ln^{-2} N)$ in stead of $\alpha_c = O(N^{-4})$ determined in Sec. II C.

E. The degree distribution

Aldous [14] has shown that the probability distribution of the degree D_{MST} of nodes in the MST equals

TABLE I. Comparison of the degree distribution in the MST

| k | Exact Eq. (24) | Conditioned Poisson Eq. (25) |
|------------------------------|-------------------------|-------------------------------|
| 1 | 0.40658 | 0.40637 |
| 2 | 0.32429 | 0.32380 |
| 3 | 0.17112 | 0.17201 |
| 4 | 0.068353 | 0.068529 |
| 5 | 0.022006 | 0.021842 |
| 6 | 0.0059347 | 0.0058013 |
| 7 | 0.0013768 | 0.0013207 |
| 8 | 0.00028022 | 0.00026309 |
| 9 | 0.000050790 | 0.000046586 |
| 10 | 8.2970×10^{-6} | 7.4240×10^{-6} |
| $\text{var}[D_{\text{MST}}]$ | 1.1917 | $2(\lambda-1) \approx 1.1872$ |

$$\text{Pr}[D_{\text{MST}} = k] = \int_0^1 e^{-\Phi(s)} \frac{\Phi(s)^{k-1}}{(k-1)!} ds, \quad (24)$$

where $\Phi(s) = \int_0^s dt \ln t / (t-1)$. Remarkably, this expression for the probability of the degree D_{MST} is close (but not identical) to a $\text{Poisson}(\lambda)$ distribution conditioned to positive integers $k \geq 1$,

$$\text{Pr}[D_{\text{MST}} = k] \approx \frac{\lambda^k e^{-\lambda}}{k! (1 - e^{-\lambda})}, \quad (25)$$

where λ is chosen such that the mean degree equals $\mu_{\text{MST}} = \mathbb{E}[D_{\text{MST}}] = 2$. It is straight forward to show that λ is the positive solution of $e^{-\lambda} = 1 - \lambda/2$, $\lambda \approx 1.593624 \dots$ and that

$\text{var}[D_{\text{MST}}] = 2(\lambda - 1)$. Table I compares the two expressions (24) and (25).

During the growth process, we observed that the degree distribution of the forest is found to be $\text{Poisson}(x)(\mu_{\text{MST}})$ below the transition ($\mu_{\text{MST}} < 1$). Above the transition ($\mu_{\text{MST}} > 1$), the average degree of the giant component equals the final value of 2. At the transition the distribution of degree minus one in the giant component is found to be $\text{Poisson}(1)$. Thus $\text{Pr}[D_{\text{MST}} - 1 = k] = 1/(ek!)$. For larger values of μ_{MST} it evolves gradually to the asymptotic form at $\mu_{\text{MST}} = 2$. The degree distribution over all nodes *not* in the large component is Poisson for all $1 < \mu_{\text{MST}} < 2$, with average degree following from $2S + (1-S) \mu_{\text{sc}} = \mu_{\text{MST}}$:

$$\mu_{\text{sc}} = \frac{\mu_{\text{MST}} - 2S}{1 - S} = -\frac{1 - S}{S} \ln(1 - S) \quad (26)$$

F. The hop count

The hop count between two arbitrary nodes in the MST ($\alpha \rightarrow 0$ regime) has been simulated and is shown in Fig. 7. The simulation also indicated that per MST, there is a large variation in the hop count which heavily contrasts with the almost sure behavior for the hop count in the $\alpha = 1$ regime [26]. The large variation between samples is indirectly deduced from the fact that $\text{var}[H_N] > \mathbb{E}[H_N]$ in the inset of Fig. 7.

The PDF $f_h(x)$ of the scaled hop count $h = H_N / \mathbb{E}[H_N]$ versus the scaled number of hops $x = j / \mathbb{E}[H_N]$ is plotted in Fig. 8 as a full line. With the dotted line, we have added the PDF of the limit random variable W of a Poisson branching process (BP) with mean μ_{BP} . Two exact expressions for $f_W(x)$ are




FIG. 7. The probability density function of the hop count simulated for various sizes of the $\alpha=0$ tree. The inset plots the average and variance of the hop count versus the number of nodes N .

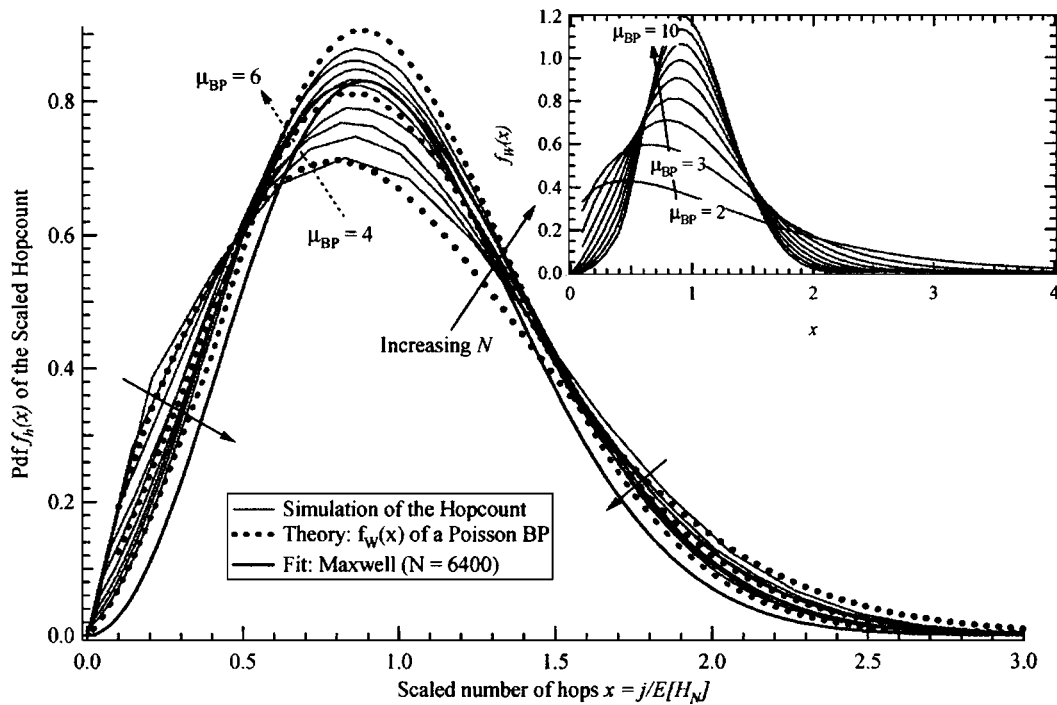


FIG. 8. The scaled hop count $h=H_N/E[H_N]$ versus the scaled number of hops $x=j/E[H_N]$ for $N=25 \times 2^k$ with $k=0, \dots, 8$. The arrows show how this scaled hop count varies with increasing N . The inset shows the PDF $f_W(x)$ of the limit random variable of a Poisson branching process.

presented in Ref. [27] and $f_W(x)$ is shown in the inset for various $2 \leq \mu_{BP} \leq 10$. We observe that the PDF of the scaled hop count for $N=50$ up to $N=6400$ lies between those of the BP corresponding to the Poisson degree distribution with mean between $\mu_{BP}=4$ and $\mu_{BP}=6$. We have also plotted the fit of a Maxwellian $f_h(x)=4/\sqrt{\pi}(2x/\sqrt{\pi})^2 \exp(-(2x/\sqrt{\pi})^2)$

which was proposed by Braunstein *et al.* in [28] purely by fitting the simulation data.

The PDF $f_W(x)$ of the Poisson BP is clearly superior to the guess of a Maxwellian. The suggestion to compare the scaled hop count (with mean 1) to the limit random variable W (also with mean 1) is explained by Aldous [14]. However, we still




FIG. 9. Trace routes from CAIDA, May 2004. Each of the 21 sources has several 10^4 destinations. The local internal hops that follow a single path towards the Internet have been subtracted. In the inset, α is computed as $\alpha=E[H]/\text{var}[H]$.

need to associate the mean μ_{BP} of the Poisson BP to the number of nodes in the MST which turns out to be a rather difficult theoretical problem.

IV. DISCUSSION AND CONCLUSIONS

Properties of both the $\alpha=1$ -tree (URT) and the $\alpha\rightarrow 0$ -tree (MST) have been characterized. We have shown that both types of trees are quite different. Based on trace-route hop count measurements as shown in Fig. 9, the regime of interest for the Internet seem [29] to lie between $0.5 < \alpha < 1.5$.

From a topological view, the Internet trees indeed seem to consist of a critical bearer tree (corresponding to the $\alpha\rightarrow 0$ tree) overgrown with URT-like small trees (influence of $\alpha = 1$). Based on Eq. (2), the latter causes that the hop count in the Internet still scales logarithmically in N , rather than polynomially as for the $\alpha\rightarrow 0$ tree (MST). This effect is similar to the small world graphs: by adding a few links in a “large average hop count graph,” the hop count may decrease dramatically.

Our theoretical study inspires some new research. If Internet trees can be modeled via α trees, insight about an *effective* link weight structure in the Internet may be gained. That effective link weight structures arises as a combination of intra-domain (shortest path) routing with interdomain (not-shortest path) routing. In spite of the relatively low number of sources, Fig. 9 seems to suggest that the effective link weight is close to a regular distribution. Although controversial to adopt link weights, from a modeling perspective, simpler tools as shortest paths, no violation of the triangular inequality, etc., can be applied to deduce first order estimates. In a next stage, the modeling of multicast trees in terms of α trees may be interesting since little about Internet multicast trees is known.

In networks where the link weights can be varied, controlled or determined independent of the topology, we have shown that if the extreme value index of the link weight distribution is larger than α_c , transport in a network is spread out over more paths, while if the extreme value index is below α_c , transports starts concentrating on very few “backbone links.” Hence, by tuning α in a same underlying topology, we may create two very different types of transport in the network. The analogy with a normal conduction (above a critical temperature T_c) and a superconducting transport (below T_c) of an electrical current in some solids (network of atoms) comes to mind.

From a control point of view in networks, operators may wish to steer flows by tuning link weights. Our study indicates that large variations in the link weights (α small) will result in overall properties close to the $\alpha\rightarrow 0$ tree (MST): many flows will traverse over a same set of links and the overall hop count will increase. From a robustness point of view, choosing α around 1 will lead to the use of more paths and, hence, a more balanced overall network load. A next step in understanding the influence of the link weight structure is to find out what the maximum amount Δw in link weight change can be in order not to modify the set of shortest paths in a network. This insight is important to estimate the topology update overhead in networks (e.g., in road traffic where the link weights may be associated with the traffic load). An accurate view of the updated topology is crucial for route planner in cars.

ACKNOWLEDGMENTS

We would like to thank Serena Magdalena for providing Figs. 1–3 and Xiaoming Zhou for Fig. 9.

-
- [1] F. A. Kuipers and P. Van Mieghem, IEEE/ACM Trans. Netw. (to be published).
- [2] R. van der Hofstad, G. Hooghiemstra, and P. Van Mieghem, Random Struct. Algorithms **26**, 598 (2005).
- [3] P. Van Mieghem, G. Hooghiemstra, and R. van der Hofstad, Delft University of Technology Report No. 2000125, 2000.
- [4] B. Bollobas, *Random Graphs*, 2nd ed. (Cambridge University Press, Cambridge, 2001).
- [5] R. van der Hofstad, G. Hooghiemstra, and P. Van Mieghem, Prob. Eng. Inf. Sci. (PEIS) **15**, 225 (2001).
- [6] R. van der Hofstad, G. Hooghiemstra, and P. Van Mieghem, Combinatorics, Probab. Comput. **14**, 795 (2005).
- [7] P. Van Mieghem, G. Hooghiemstra, and R. van der Hofstad, IEEE/ACM Trans. Netw. **9**, 719 (2001).
- [8] P. Van Mieghem, Int. J. Commun. Syst. **17**, 269 (2004).
- [9] L. A. Braunstein, S. V. Buldyrev, R. Cohen, S. Havlin, and H. E. Stanley, Phys. Rev. Lett. **91**, 168701 (2003).
- [10] R. Dobrin and P. M. Duxbury, Phys. Rev. Lett. **86**, 5076 (2001).
- [11] P. Van Mieghem, *Performance Analysis of Communications Systems and Networks* (Cambridge University Press, Cambridge, 2005).
- [12] In Cisco’s OSPF implementation, it is suggested to use $w(i \rightarrow j) = 10^8 / B(i \rightarrow j)$, where $B(i \rightarrow j)$ denotes the bandwidth (in bit/s) of the link between nodes i and j .
- [13] A. M. Frieze, Discrete Appl. Math. **10**, 47 (1985).
- [14] D. Aldous, Random Struct. Algorithms **1**, 383 (1990).
- [15] S. Janson, Random Struct. Algorithms **7**, 337 (1995).
- [16] M. D. Penrose, Ann. Prob. **24**, 1903 (1996).
- [17] A.-L. Barabasi, Phys. Rev. Lett. **76**, 3750 (1996).
- [18] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *An Introduction to Algorithms* (MIT Press, Boston, 1991).
- [19] L. Pietronero and W. Schneider, Physica A **170**, 81 (1990).
- [20] A. I. Markushevich, *Theory of Functions of a Complex Variable* (Chelsea Publishing Company, New York, 1985), Vols. I–III.
- [21] In general holds that $w_{(k)} = F_w^{-1}(U_{(k)})$ and $\mathbb{E}[w_{(k)}] = \mathbb{E}[F_w^{-1}(U_{(k)})] \neq F_w^{-1}(\mathbb{E}[U_{(k)}])$, but, for a large number of order statistics L , the central limit theorem leads to
$$\mathbb{E}[w_{(k)}] \simeq F_w^{-1}\left(\frac{j}{L}\right) \simeq F_w^{-1}(\mathbb{E}[U_{(k)}])$$
 because for a uniform random variable U on $[0,1]$ the average

weight of the j th smallest link is exactly

$$E[w_{(j)}] = \frac{j}{L+1} \approx \frac{j}{L}$$

[22] S. Janson, D. E. Knuth, T. Luczak, and B. Pittel, *Random Struct. Algorithms* **4**, 233 (1993).

[23] Since the average of the k th smallest link weight can be computed from Eq. (16) as

$$E[w_{(k)}] = \frac{L!}{\Gamma\left(L+1+\frac{1}{\alpha}\right)} \frac{\Gamma\left(k+\frac{1}{\alpha}\right)}{\Gamma(k)}$$

the exact formula (15) reduces to

$$E[W_{\text{MST}}(\alpha)] = \frac{L!}{\Gamma\left(L+1+\frac{1}{\alpha}\right)} \sum_{j=1}^L \frac{\Gamma\left(j+\frac{1}{\alpha}\right)}{\Gamma(j)} (1 - S_j^2).$$

Analogously to the above manipulations, after conversion to an

integral, substituting $x=2u/N$ and using Eq. (6.1.47) of Ref. [24], for large z , that $\Gamma(z+1/\alpha)/\Gamma(z)=(z)^{1/\alpha}[1+O(1/z)]$, we arrive at the same formula.

[24] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions* (Dover Publications, New York, 1968).

[25] The probability that m nodes become connected in a cluster different from the giant component tends to zero for large N .

[26] P. Van Mieghem, G. Hooghiemstra, and R. W. van der Hofstad, *Proceeding of Passive and Active Measurement (PAM 2001)*, April 23–24, RIPE NCC, Amsterdam (2001).

[27] P. Van Mieghem (unpublished).

[28] L. A. Braunstein, S. V. Buldyrev, S. Sreenivasan, R. Cohen, S. Havlin, and H. E. Stanley, The Optimal Path in an Erdos-Renyi Random Graph, *Proceedings of the 23rd LANL-CALS Conference on Complex Networks* (Springer Verlag, Berlin 2004).

[29] NLANR traces (Aug. 2001) with 112 sources give $E[\alpha] = 1.34$, while RIPE traces (Feb. 2004) with about 70 sources give $E[\alpha]=0.7$.