



HAL
open science

Is the Internet Fractal? The Multicast Power Law Revisited

Cédric Adjih, Leonidas Georgiadis, Philippe Jacquet, Wojciech Szpankowski

► **To cite this version:**

Cédric Adjih, Leonidas Georgiadis, Philippe Jacquet, Wojciech Szpankowski. Is the Internet Fractal? The Multicast Power Law Revisited. [Research Report] RR-4157, INRIA. 2001. inria-00072466

HAL Id: inria-00072466

<https://inria.hal.science/inria-00072466>

Submitted on 24 May 2006

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

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

Is the Internet Fractal? The Multicast Power Law Revisited

Cédric Adjih, Leonidas Georgiadis, Philippe Jacquet, Wojciech Szpankowski

N° 4157

THÈME 1



R
apport
de recherche

Is the Internet Fractal? The Multicast Power Law Revisited

Cédric Adjih, Leonidas Georgiadis, Philippe Jacquet, Wojciech Szpankowski

Thème 1 — Réseaux et systèmes
Projet Hipercom

Rapport de recherche n° 4157 — — 21 pages

Abstract: One of the main benefits of multicast communication is the overall reduction of network load. To quantify this reduction, when compared to traditional unicast, experimental studies by Chuang and Sirbu indicated the so called *power law* which asserts that the number of links in a multicast delivery tree connecting a source to m (distinct) sites is in $m^{0.8}$. In order to explain theoretically this behavior, Phillips, Shenker, and Tangmunarunkit examined approximately the link number for a V -ary complete tree topology, and concluded that this number grows nearly linearly with m , thus not obeying the power law. We first re-examine the analysis by Phillips et.al. and provide precise asymptotic expansion for the link number that confirms the nearly linear (with some wobbling) growth. Claiming that the essence of the problem lies in the modeling assumptions, we replace the V -ary complete tree topology by a V -ary *self-similar* tree with similarity factor $0 \leq \theta < 1$. In such a tree a node at level k is replicated $CV^{(D-k)\theta}$ times, where D is the depth of the tree and C is a constant. Under this assumption, we analyze again $L(m)$ and prove that $L(m) \sim cm^{1-\theta}$ where c is an explicitly computable constant. Hence self-similar trees provide a plausible explanation of the multicast power law. Next, we examine more general conditions for general trees, under which the power law still holds. We also discuss some experimental results in real networks that reaffirm the power law and show that in these networks the general conditions hold.

Key-words: Internet, TCP, asymptotics, mean-field approximation, fairness

L' Internet est-il fractal? Révision de la loi en puissance des communications *multicast*

Résumé : Un des bénéfices principaux de la communication *multicast* réside dans la réduction de la charge sur le réseau. Dans l'objectif de quantifier cette économie par rapport aux communications classiques *unicast*, les études expérimentales de Chuang et Sirbu ont mis en évidence la loi dite *en puissance* qui conjecture que le nombre de liens dans un arbre multicast desservant m sites est en $m^{0.8}$. Dans le but d'établir une théorie pour rendre compte de cette loi, Phillips, Shenker et Tangmunarunkit ont étudié le nombre de liens dans un arbre complet de degré V , ils ont conclu que ce nombre croissait en gros linéairement avec m en contradiction avec la loi en puissance. Nous regardons d'abord l'analyse de Phillips *et al.* et nous donnons des expressions asymptotiques plus précises qui confirment la linéarité du nombre de liens. La cause du problème étant dans les hypothèses de modélisation, nous remplaçons l'arbre complet de degré V par un arbre auto-similaire de degré V et avec un coefficient d'auto-similarité $0 \leq \theta < 1$.

Mots-clés : Internet, TCP, évaluation de performances, comportement asymptotique, approximation par le champ moyen, équité.

1 Introduction

Multicast communication in the internet was proposed a decade ago in [1, 5] (cf. also [4]), and the experimental MBone network has been operational since 1992. In multicast communication senders transmit to logical address while receivers join a logical group. Multicast routing ensures that only a single copy of a packet destined to multiple destinations traverses each link, so that the overall traffic load is reduced significantly. Also, multicast alleviates the overhead on senders who can reach an entire group by the transmission of a single packet. The trade-off is that multicast requires extra control and routing overhead at the routing nodes.

In this paper we concentrate on the quantification of the main advantage of multicast routing. That is, we address the question of what is the expected traffic load reduction due to multicasting, when compared to unicast communication. Only recently some efforts were undertaken to address this question. Motivated by the problem of pricing multicast communications, Chuang and Sirbu [3] performed experiments on a number of real and generated network topologies. They measured the average number of links $L(m)$, in the multicast tree needed to reach m randomly selected routing nodes from a given source. $L(m)$ represents the average cost of multicast tree per unit of bandwidth (it is assumed that all destinations require the same bandwidth i.e., we have unirate multicast trees). If unicast communication is employed, then the communication cost per unit of bandwidth is Um where U is the average number of links in a unicast path. The efficiency gain of multicast versus unicast is reflected in how far $L(m)$ deviates from the (unicast) linear growth. Chuang and Sirbu [3], after extensive simulations, concluded experimentally that $L(m) = \Theta(m^{0.8})$. Moreover, they found that the constant in front of $m^{0.8}$ does not change significantly with network size. This naturally raises the following questions. Is this behavior specific to the chosen topologies, or is to be expected for other topologies as well? In the latter case, what is the main characteristic of the topology (or the generated multicast trees) that causes the appearance of the power law? This paper attempts to provide theoretical answers to these questions.

The most intriguing result, from our point of view, in the Chuang and Sirbu paper is presented in their Figure 7 that we reproduce here in Figure 1 in our notation. They considered a multicast tree with N routing nodes. Multiple destination hosts may be connected to the network through each routing node (e.g., each routing node may have a number of dial-in ports, or may have a LAN connected to one of its ports). A source and a number n , of multicast destination hosts, is picked randomly. A multicast tree using shortest paths from the source to each of the destinations is constructed. We denote by $L(n)$ the average number of links of the multicast trees created by the above procedure. Note that the number of destination hosts n can be either smaller or larger than the number of routing nodes N (throughout the paper we rather work with n than with distinct number of hosts m). Therefore, the ratio $a = n/N$ can vary from zero to infinity. Figure 1 shows the Chuang and Sirbu findings concerning the ratio of $L(n)$ to the average cost of unicast communication. We observe that for n small relative to N , i.e., for a small, the power law $L(n) \sim n^{0.8}$ seems to be exhibited. Deviation from the power law and a phase transition appears around $a = 1$, and saturation occurs for $a \gg 1$. Our goal is to provide a theoretical explanation of the behavior observed in Figure 1, based on topological features of multicast trees.

The first attempt to explain the power law of Chuang and Sirbu was undertaken by Phillips, Shenker and Tangmunarunkit in [9]. The authors in [9] considered several network topologies and provided approximate analysis of $L(n)$ which indicated that $L(n)$ does not grow according to the power law (cf. (17) and (18) of [9]). To make sure the approximation of [9] did not tilt final results, we provide precise analysis of the tree topology considered in [9], (cf. Theorem 1) which confirms that their result indeed gives a good approximation for the leading term for $L(n)$ when a is small. Interestingly enough, for small a we discover some small oscillations of the coefficient in front of n .

There are various ways to attempt to resolve this discrepancy between experimental and theoretical findings. One may simply give up any hope to provide an analytical insight to the experimental results by simple modeling of the multicast trees. Alternatively, one may consider the experimental results as isolated examples from which no general conclusions can be drawn. Yet another approach

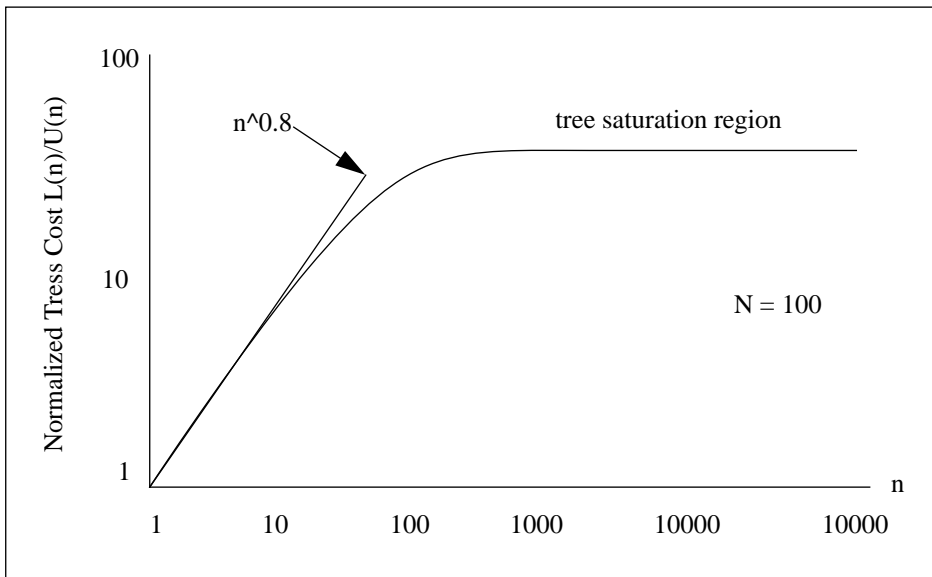


Figure 1: This is Figure 7 from Chunag and Sirbu [3] showing the phase transition of the ratio of the number of links traverse in multicast and the average path length in unicast versus the number of destinations n .

would be to investigate whether the essence of the problem lies in the modeling assumptions. The latter approach is the one we attempt to pursue in this paper. Specifically, we assume that between two successive branchings of a multicast tree there may be several concatenated unary nodes (i.e., nodes at which no branching occurs). The average number of these concatenated nodes decreases exponentially as the distance (in number of branchings) from the source increases. More precisely, a node in such a tree at level k is replicated $V^{(D-k)\theta}$ times, where D is the depth of three tree and $0 \leq \theta \leq 1$ is the *self-similarity factor*. A tree with such a property is shown in Figure 2. In fact, in our analysis we study trees of this structure which, for reasons that will become apparent in Section 2, we call *self-similar*. For self-similar trees we proceed to show that the ratio $R(n) = L(n)/U(n)$ (of the number of links in a multicast tree to the average path length in unicast) exhibits the power law. More precisely, we show that for small $a = n/N$ the ratio $R(n) \sim (c + \psi(n))n^{1-\theta}$ where c is an explicitly computable constant and $\psi(n)$ is an oscillating function of rather small amplitude (cf. Corollary 4(i)). Moreover, this constant is independent of the number of routing nodes N , which is in agreement with the observed results in [3].

Motivated by our experimental results, we concentrate on certain conditions that seem to be satisfied by multicast trees on real networks. These conditions refer to the number of routing nodes on a multicast tree that can be reached by a node on that tree. It turns out that self-similar trees satisfy these conditions, and furthermore, these conditions are sufficient for the appearance of the power law.

The paper is organized as follows. In the next section we consider a complete tree network topology with destination hosts located at the leaves and provide a precise analysis of $L(n)$ for regular and self-similar trees (cf. Section 2). In Section 3 we indicate that the same results hold true if destination hosts are located to non-leaf tree nodes as well, and we discuss more general conditions that guarantee the appearance of the power law. We provide experimental results in Section 4. In particular, we find that $\theta_{\text{exp}} \approx 0.12$ which confirms the power law $\Theta(n^{0.88})$. Finally, the Appendix contains derivations of our theoretical results. In passing, we should mention that our findings have been established by analytical techniques for the precise analysis of algorithms such

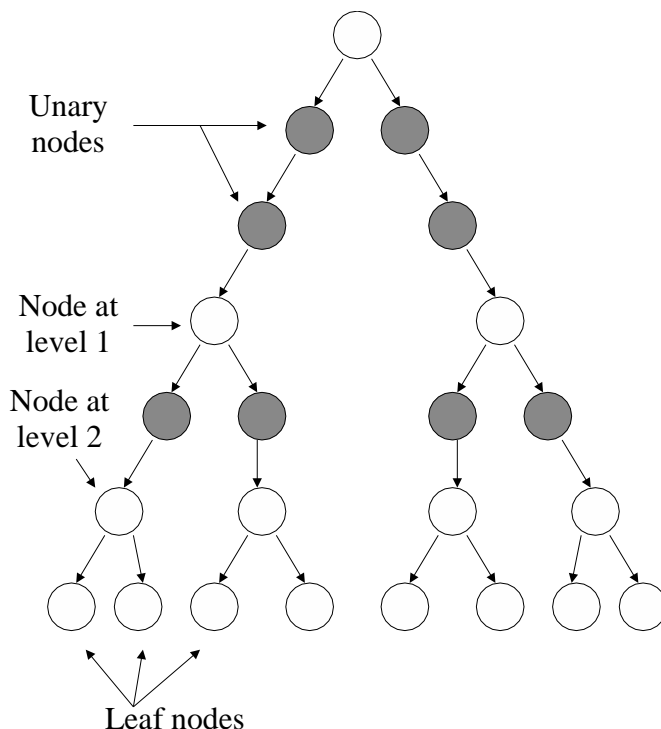


Figure 2: A Self-Similar Tree with $D = 3$ and $\theta = 1$.

as Mellin transform and complex asymptotics (cf. [8, 10]). We give a brief survey of these methods since they may be useful for the study of other problems of similar nature.

2 Main Results

In this section we present our main results. Derivations of the analytical results are provided in the Appendix. We first explore the regular tree topology and later extend our analysis to self-similar trees.

2.1 Regular Trees

As in [9], we consider a V -ary tree where the source is located at the root of the tree and all the potential destination hosts of the multicast tree are connected to the network through the leaf nodes of the tree. We assume that behind each leaf node there may be multiple destination hosts. This is the case when, for example, each of the leaf nodes connects a LAN to the network. We may think of the V -ary tree as the tree in a communication network G , composed of the shortest paths from the source to all the routing nodes of G . We refer to such a tree as the “*global multicast tree*” for the chosen source. When a number of the destination hosts in G needs to form a multicast group, they form a subtree of the global multicast tree. This subtree is used for multicast communication. Note that this formulation implicitly assumes that multicast trees are formed using shortest path trees, as in [3]. In Section 3 we consider the possibility that destination hosts may be connected to the network through other tree nodes, not just through the leaf nodes.

We consider that the V -ary tree is complete (all nodes but the leaves have degree V) and denote by D its depth, i.e., its longest (in terms of hops) path. If N is the number of leaf nodes then clearly

$$N = V^D. \quad (1)$$

Let the multicast group consist of n hosts and

$$a = \frac{n}{N} > 0. \quad (2)$$

Note that since more destination hosts may be behind each node, it is possible that $a > 1$. We need to consider this region of a in order to explain the saturation phenomena in Figure 1. We assume that the probability of a host being connected to the network through a given leaf node is uniform and independent of the manner the rest of the hosts are connected. In passing we point out, after the authors of [9], that if one insists of considering distinct hosts m , then a good approximation is obtained by setting $m = N(1 - e^{-n/N})$ (cf. (1) of [9]). We underline again that throughout this paper we will work with the destination host multiply n rather than with the number of distinct nodes m . Note that when $a \ll 1$, then from the above approximation we obtain $m \approx n$. As the power law is manifested when $a \ll 1$, the results are qualitatively the same whether we work with n or m . We choose to work with n since the analysis is simpler in this case.

Following Chuang and Sirbu [3], to quantify the reduction of traffic load in multicast over unicast, we shall analyze the average number of links $L(n)$ in the multicast tree that connects n randomly selected hosts. If U denotes the average path between the source and a host in the unicast transmission, then the reduction ratio $R(n)$ is defined as

$$R(n) = \frac{L(n)}{U}. \quad (3)$$

Observe that for the complete V -ary tree we have $U = D$.

To estimate the average number of links in the multicast tree connecting n nodes, we observed that at level k of the tree, $1 \leq k \leq D$, there are V^k links; the probability that a particular link is in the multicast tree after n destination hosts have been selected is

$$1 - (1 - 1/V^k)^n.$$

Thus the average number of links in the multicast tree is

$$L(n) = \sum_{k=1}^{D=\log_V N} V^k \left(1 - (1 - V^{-k})^n\right). \quad (4)$$

Our goal is to estimate $L(n)$ asymptotically as $n \rightarrow \infty$ for fixed a (cf. Theorem 1), as well as for $a \rightarrow 0$, $a \rightarrow 1$ and $a \rightarrow \infty$ (cf. Corollary 2).

Since $V^{-(D-k)} = aV^k/n$, in order to evaluate (4), it seems natural to replace $1 - (1 - V^{-(D-k)})^n$ by $1 - e^{-aV^k}$, take the upper sum index $D = \log_V N$ to be infinity, and then sum up. However, this direct approach leads to incorrect results since the upper index $D = \log_V N$ is at the cuff-off of a significant contribution coming from $(1 - 1/V^k)^n$. We need to be much more careful.

In the Appendix we prove the following result.

Theorem 1 *For fixed a , $0 < a = n/N < \infty$, and large n (hence large N) the average number of traversed links $L(n)$ attains the following asymptotic expansion*

$$L(n) = N \left(\frac{V}{V-1} - c_1(a) \right) - \frac{V}{V-1} - \frac{1}{2}c_2(a) + O\left(\frac{1}{\log n}\right), \quad (5)$$

Table 1: Comparison of the exact $L(n)$ with the asymptotic expansion $L^{\text{asym}}(n)$ obtained in Theorem 1 and the approximation $L^{\text{PST}}(n)$ proposed by Phillips, Shenker, and Tangmunarunkit [9] for $a = 0.5$, that is, $N = 2n$.

N	$L(n)$	$L^{\text{asym}}(n)$	$L^{\text{PST}}(n)$
2	3.250	3.202	4.885
4	7.919	7.895	9.771
8	17.295	17.283	19.541
16	36.083	36.057	39.083
32	73.609	73.605	78.166
64	148.704	148.702	156.332
128	298.896	298.896	312.665
256	599.283	599.282	625.330
512	1200.056	1200.056	1250.650
1024	2401.602	2401.602	2501.320

where

$$c_1(a) = \sum_{l=0}^{\infty} V^{-l} \exp(-aV^l), \quad (6)$$

$$c_2(a) = \sum_{l=0}^{\infty} aV^l \exp(-aV^l). \quad (7)$$

The quantities $c_1(a)$ and $c_2(a)$ converge quickly and hence (5) provides a convenient way for the approximate computation of $L(n)$. In order to verify the accuracy of the above asymptotic expansion (denoted $L^{\text{asym}}(n)$ in Table 1) we compare it to the exact formula $L(n)$ and the approximation $L^{\text{PST}}(n)$ proposed by the authors of [9]. From Table 1 one concludes that the asymptotic expansion presented in Theorem 1 is very good, even for small values of n .

>From Theorem 1 one also must conclude that $R(n) \sim \frac{N}{D}(V/(V-1) - c_1(a))$, which most definitely does not exhibit the power law for general a . But, Figure 7 of [3] (cf. Figure 1) suggests that the power law appears only for small values of a , while for large values of a one should expect a saturation. To examine these cases, in the corollary below we analyze (5) for $a \rightarrow 0$, $a \rightarrow 1$ and $a \rightarrow \infty$. This is equivalent to estimating asymptotically the constants $c_1(a)$ and $c_2(a)$ for these three regimes.

Corollary 2 *Under the same conditions as in Theorem 1, we find:*

(i) *For $a \rightarrow 0$ the quantity $L(n)$ attains the following asymptotics*

$$L(n) = n \left(D + \frac{1}{\ln V} - \frac{\ln n}{\ln V} + \left(\frac{1}{2} - \frac{\gamma}{\ln V} \right) + \psi_1(\ln a) \right) - \frac{V}{V-1} - \frac{1}{2 \ln V} + \frac{1}{2} \psi_2(\ln a) + O\left(\frac{1}{\log n}\right) \quad (8)$$

where $\gamma = 0.571\dots$ is the Euler constant, and $\psi_1(x)$, $\psi_2(x)$, are oscillation periodic functions of small amplitude for small V that can be expressed as

$$\psi_1(x) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \frac{\Gamma(-1 - (2\pi ik)/\ln V)}{\ln V} \exp\left((2\pi ik)\frac{x}{\ln V}\right), \quad (9)$$

$$\psi_2(x) = 2\pi i \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \frac{k\Gamma(-(2\pi ik)/\ln V)}{\ln^2 V} \exp\left((2\pi ik)\frac{x}{\ln V}\right). \quad (10)$$

In fact, $|\psi_1(x)| < 0.0000001725, 0.00041227, 0.0085, 0.068, 0.153$ for $V = 2, 3, 5, 100, 1000$, respectively.

(ii) For $\mathbf{a} \rightarrow \mathbf{1}$ we have

$$L(n) = N \left(\frac{V}{V-1} - C_1 - C_2(a-1) + C_3(a-1)^2 \right) - \frac{V}{V-1} - \frac{C_3}{2}a + O\left(\frac{1}{\log n}\right), \quad (11)$$

where

$$C_1 = \sum_{l=0}^{\infty} V^{-l} e^{-V^l}, \quad C_2 = \sum_{l=0}^{\infty} e^{-V^l}, \quad C_3 = \frac{\sum_{l=0}^{\infty} V^l e^{-V^l}}{2}.$$

(iii) For $\mathbf{a} \rightarrow \infty$ we arrive at

$$L(n) = N \left(\frac{V}{V-1} - e^{-a} \right) - \frac{V}{V-1} - \frac{1}{2} (ae^{-a} + aVe^{-aV}) + O\left(\frac{1}{\log n}\right). \quad (12)$$

We can now compare the theoretical results of Corollary 2 to the experimental results of Chuang and Sirbu presented in Figure 1. In particular, taking into account that $\ln N = D \ln V$ we see from (8) that the ratio $R(n)$ for $a \rightarrow 0$ can be approximated by

$$R_D(n) \approx n \left(1 + \frac{1-\gamma}{\ln N} + \frac{\ln V}{2 \ln N} - \frac{\ln n}{\ln N} + \frac{\ln V}{\ln N} \psi_1(\ln a) \right)$$

which again most decidedly is not of the power law form. However, one can argue that for large $a \rightarrow \infty$ and fixed N formulas (11) and (12) could explain respectively the transition and saturation region of Figure 1.

We note that the approximate analysis in [9] lead to the approximation

$$L(n) \approx n (D + 1/\ln V - \ln n/\ln V).$$

From (8) we see that the main term that the approximate analysis missed, is $d = (.5 - .5772/\ln V)$ plus the oscillating function $\psi_1(\ln a)$ which for small V is of small amplitude. The term d is small, i.e., $-0.33272 \leq d < .5$. Moreover, we must note that the approximation is valid for $a \ll 1$.

2.2 Self-Similar Trees

In view of the results in the previous section, we conclude that we cannot explain the multicast power law based on the adopted modeling assumptions. As discussed in the introduction, we shall argue that a possible explanation is that at the heart of the problem is our assumption regarding the structure of the global tree network. In this section we show that if the tree has a ‘‘self-similar’’ structure in the sense to be discussed below, then we indeed have the power law behavior for small a .

As in the previous subsection, consider a V -ary tree where all possible hosts are located at the leaves of the tree. However, we assume now that the link connecting a node at level k and a node at level $k-1$ consists of a concatenation of a random number of links. Let ℓ_k be the average number of these links. We postulate that ℓ_k is a fraction of ℓ_{k-1} , that is, for some A we have $\ell_1 = A$ and

$$\ell_k = \phi \ell_{k-1}, \quad 0 \leq \phi \leq 1.$$

Therefore, $\ell_k = \phi^{k-1} A$. Setting $\phi = V^{-\theta}$ we have,

$$\ell_k = \ell_D V^{(D-k)\theta}, \quad \theta > 0.$$

In the rest of paper, we assume for simplicity and without loss of generality that $\ell_D = 1$. The last equality suggests another interpretation of θ . Observe that there are $K = V^{D-k}$ leaves hanging from a node at level k ; thus we reproduce such a node K^θ times.

We call a tree with the above structure, a self-similar V -ary tree with similarity factor θ . Figure 2 shows a binary self-similar tree with similarity factor $\theta = 1$ and depth $D = 3$. Note that when $\theta = 0$, we have the regular V -ary tree. In the following we will concentrate on small values of θ , specifically, we assume that $0 \leq \theta < 1$.

We analyze now $L_\theta(n)$ and $R_\theta(n)$ for self-similar trees. In particular, as before we derive

$$L_\theta(n) = \sum_{k=1}^D \left(V^{(D-k)\theta} V^k \left(1 - (1 - V^{-k})^n \right) \right),$$

and for the average path length in a unicast connection we find

$$U_\theta = \sum_{k=1}^D V^{(D-k)\theta} = \frac{N^\theta - 1}{V^\theta - 1}. \quad (13)$$

In the Appendix we prove the following asymptotic expansions for $L_\theta(n)$.

Theorem 3 For fixed a , $0 < a = n/N < \infty$, and large n the average number of links $L_\theta(n)$ in the self-similar tree attains the following asymptotic expansion

$$L_\theta(n) = N \left(\frac{V^{\bar{\theta}}}{V^{\bar{\theta}} - 1} - c_1(a, \theta) \right) - \frac{N^\theta V^{\bar{\theta}}}{V^{\bar{\theta}} - 1} - \frac{1}{2} c_2(a, \theta) + O\left(\frac{1}{\log n}\right), \quad (14)$$

where $\bar{\theta} = 1 - \theta$ and

$$c_1(a, \theta) = \sum_{l=0}^{\infty} V^{-\bar{\theta}l} \exp(-aV^l), \quad (15)$$

$$c_2(a, \theta) = \sum_{l=0}^{\infty} aV^{(1+\theta)l} \exp(-aV^l). \quad (16)$$

Corollary 4 Under the same conditions as in Theorem 3, we find:

(i) For $a \rightarrow 0$ the quantity $L(n)$ attains the following asymptotics

$$L_\theta(n) = N^\theta \left(n^{\bar{\theta}} \left(\frac{\Gamma(\theta)}{\bar{\theta} \ln V} - \psi_3(\ln a) \right) - \frac{V^{\bar{\theta}}}{V^{\bar{\theta}} - 1} - \frac{1}{2} \frac{\theta \Gamma(\theta)}{n^\theta \ln V} - \frac{1}{n^\theta} \psi_4(\ln a) \right) + O\left(\frac{1}{\log n}\right), \quad (17)$$

where $\gamma = 0.571\dots$ is the Euler constant, $\Gamma(\theta)$ is the Gamma function, $\psi_3(a)$ and $\psi_4(a)$ are oscillating periodic functions of small amplitude for small V that can be expressed as

$$\psi_3(x) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \frac{\Gamma(-1 + \theta - (2\pi i k)/\ln V)}{\ln V} \exp\left(2\pi i k \frac{x}{\ln V}\right), \quad (18)$$

$$\psi_4(x) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \frac{(\theta - i2\pi k/\ln V)\Gamma(\theta - 2\pi i k/\ln V)}{\ln V} \exp\left(2\pi i k \frac{x}{\ln V}\right). \quad (19)$$

(ii) For $\mathbf{a} \rightarrow \mathbf{1}$ we have

$$L(n) = N \left(\frac{V^\theta}{V^\theta - 1} - C_1(\theta) - C_2(\theta)(a - 1) + C_3(\theta)(a - 1)^2 \right) - \frac{V}{V - 1} - \frac{C_3(\theta)}{2}a + O\left(\frac{1}{\log n}\right) \quad (20)$$

where

$$C_1(\theta) = \sum_{l=0}^{\infty} V^{-\theta l} e^{-V^l}, \quad C_2(\theta) = \sum_{l=0}^{\infty} V^{\theta l} e^{-V^l}, \quad C_3(\theta) = \frac{\sum_{l=0}^{\infty} V^{(1+\theta)l} e^{-V^l}}{2}.$$

(iii) For $\mathbf{a} \rightarrow \infty$ we arrive at

$$L_\theta(n) = N \left(\frac{V^\theta}{V^\theta - 1} - e^{-a} \right) - \frac{V^\theta}{V^\theta - 1} - \frac{1}{2} (ae^{-a} + aV^{1+\theta}e^{-aV}) + O\left(\frac{1}{\log n}\right) \quad (21)$$

Now we are in a position to explain Figure 1 and to justify the power law of Chuang and Sirbu [3]. Observe that for small a Corollary 4(i) and (13) suggest the following approximation.

$$R(n, \theta) = \frac{L(n, \theta)}{U_\theta} \approx n^{1-\theta} (V^\theta - 1) \left(\frac{\Gamma(\theta)}{(1-\theta)\ln V} - \psi_3(\ln a) \right) - \frac{V - V^{1-\theta}}{V^{1-\theta} - 1}. \quad (22)$$

Thus, we obtain the power law with exponent of n equal to $1 - \theta$. In addition, we observe that the constant that multiplies $n^{1-\theta}$ is independent of the tree size, which also agrees with the experimental results. Furthermore, Corollary 4(iii) predicts that for $a \rightarrow \infty$ with N fixed the ratio $R(n)$ tends to the constant equal to $N^{1-\theta}V^{1-\theta}(V^\theta - 1)/(V^{1-\theta} - 1)$. Finally, around $a = 1$ we have a transitive behavior predicted by Corollary 4(ii). Therefore, Figure 1 is explained under the assumption of self-similar multicast trees.

3 Extensions and Generalizations

In this section we extend our results to self-similar trees with hosts located at any of the non-unary node, and — more importantly — we introduce generalized self-similar trees that potentially will allow us to justify the power law for a large class of network topologies.

3.1 Nodes Located at Non-unary Nodes

In Section 2 we assumed that destination hosts are located at the leaves of the global multicast tree. If destination hosts can also be located on any of the non-unary nodes of the global multicast tree,

then in order to find the average cost of multicast we argue as in Section 2 and find

$$\begin{aligned} L_\theta(n) &= \sum_{k=1}^D \left(V^{(D-k)\theta} V^k \left(1 - \left(1 - \frac{V^{D-k+1} - 1}{V^{D+1} - V} \right)^n \right) \right) \\ &= V^{D\bar{\theta}} \frac{V^{\bar{\theta}(D+1)} - V^{\bar{\theta}}}{V^{\bar{\theta}} - 1} - V^D \sum_{k=0}^{D-1} V^{-\bar{\theta}k} \left(1 - \frac{a \frac{V^{k+1} - 1}{V-1}}{n} \right)^n \end{aligned}$$

The term of $L_\theta(n)$ that needs to be analyzed asymptotically is

$$\bar{L}_\theta(n) = \sum_{k=0}^{D-1} V^{-\bar{\theta}k} \left(1 - \frac{a \frac{V^{k+1} - 1}{V-1}}{n} \right)^n.$$

This term has the same form as the one analyzed before and hence its asymptotic expansion has also the same form. Hence, the results are qualitatively the same as in the case where destination hosts are located at nodes at the leaves of the tree.

3.2 Generalization of Self-similar Trees

Motivated by experimental results, see Figure 4 of Section 4, we provide in this section general conditions on the global multicast tree, that give rise to the multicast power law.

For a node A on the global multicast tree, define by $r(A)$ the number of destination tree nodes that can be reached by A , using the multicast tree links (the tree links are considered unidirectional). For example, in the self-similar tree in Figure 2 if destination hosts are located at leaf nodes, then we have $r(A) = 1$ if A is a leaf node or if A is a unary node between levels 2 and 3. If A is located at level 1 or if A is a unary node between levels 0 and 1 we have $r(A) = 4 = 2^{3-1}$. We call $r(A)$ the “reachability degree of A ”.

Let $Q(k)$ the number of nodes with reachability degree k . If destination nodes are at the leaves of the global multicast tree, it is easy to see that for the regular V -ary tree, we have $Q(k) = V^D k^{-1}$ if $k = V^l$, $0 \leq l \leq D$, and $Q(k) = 0$ when k is not a power of V . Similarly, taking into account that there are $V^{\theta(D-l)} - 1$ unary nodes between levels $l-1$ and l we find that for the self-similar V -ary tree we have $Q(k) = k^{-1+\theta} V^D$, if $k = V^l$, $1 \leq l \leq D$ and $Q(k) = 0$ when k is not a power of V .

Now, let us define by $F(k)$ the number of tree nodes with reachability degree at least k . That is,

$$F(k) = \sum_{l=k}^N Q(l).$$

For a self-similar tree of depth D , we have for $1 \leq k \leq N$.

$$\begin{aligned} F(k) &= \sum_{l=\lceil \log_V k \rceil}^D Q(l) \\ &= N \sum_{l=\lceil \log_V k \rceil}^D V^{l(-1+\theta)} \\ &= NV^{\lceil \log_V k \rceil(-1+\theta)} \frac{1 - V^{(D-\lceil \log_V k \rceil+1)(-1+\theta)}}{1 - V^{-1+\theta}}. \end{aligned}$$

Hence,

$$NV^{-1+\theta} k^{-1+\theta} \leq F(k) \leq \frac{N}{1 - V^{-1+\theta}} k^{-1+\theta}.$$

>From the previous discussion we conclude that for self-similar trees, $F(k)$ is decreasing according to $k^{-1+\theta}$. The experimental results in Section 4 (see Figure 4) confirm that this latter relation holds in real networks as well. Hence it seems natural to ask whether networks satisfying relations of this type give rise to the multicast power law. It turns out that indeed this is the case. There is ongoing work on this subject, however, and we report here only on a representative result which is expressed informally in the following theorem.

Theorem 5 *If for the global multicast tree and for large N it holds $F(k) \approx A_N k^{-1+\theta}$, $0 < \theta < 1$, then for $n \ll N$, we have $R(n) \approx (\theta \Gamma(\theta) / (1 - \theta)) n^{1-\theta}$.*

Proof: We sketch informally the steps of the proof, emphasizing the main ideas.

It is easy to see that under the assumptions of destination host generation in Section 2, we have

$$L(n) = \sum_{k=1}^N Q(k) \left(1 - \left(1 - \frac{k}{N}\right)^n\right) = \sum_{k=1}^N (F(k) - F(k+1)) \left(1 - \left(1 - \frac{k}{N}\right)^n\right).$$

By assumption we have that $F(k) \approx A_N k^{-1+\theta}$. Hence we can write

$$L(n) \approx N^{-1+\theta} A_N \sum_{k=1}^N \left(\left(\frac{k}{N}\right)^{-1+\theta} - \left(\frac{k+1}{N}\right)^{-1+\theta} \right) \left(1 - \left(1 - \frac{k}{N}\right)^n\right)$$

For N large compared to n , we can use the Euler-Maclaurin summation formula (cf. [8, 10]) to approximate the summation in the last equation by an integral. Hence we can write.

$$\begin{aligned} S(n) &\triangleq \sum_{k=1}^N \left(\left(\frac{k}{N}\right)^{-1+\theta} - \left(\frac{k+1}{N}\right)^{-1+\theta} \right) \left(1 - \left(1 - \frac{k}{N}\right)^n\right) \\ &\approx - \int_0^1 (1 - (1-x)^n) d(x^{-1+\theta}) \end{aligned}$$

Integrating by parts the last integral we have

$$S(n) \approx -x^{-1+\theta} (1 - (1-x)^n) \Big|_0^1 + n \int_0^1 x^{-1+\theta} (1-x)^{n-1} dx$$

Making the change of variables $y = nx$ leads to

$$S(n) \approx \int_0^n \left(\frac{y}{n}\right)^{-1+\theta} \left(1 - \frac{y}{n}\right)^{n-1} dy - 1 = n^{1-\theta} \int_0^n y^{-1+\theta} \left(1 - \frac{y}{n}\right)^{n-1} dy - 1.$$

The main contribution for the latter integral comes from small values of y relative to n . For this range of y , we can approximate $(1 - \frac{y}{n})^n \approx e^{-y}$ and hence we have

$$S(n) \approx n^{1-\theta} \int_0^n y^{-1+\theta} e^{-y} dy - 1 \approx n^{1-\theta} \Gamma(\theta) - 1$$

Hence we see that

$$L(n) \approx A_N N^{-1+\theta} \Gamma(\theta) n^{1-\theta} - A_N N^{-1+\theta}.$$

It remains to compute U . We have

$$\begin{aligned}
 U &= \sum_{k=1}^{k=N} \frac{k}{N} Q(k) \\
 &= \sum_{k=1}^{k=N} \frac{k}{N} (F(k) - F(k+1)) \\
 &= A_N N^{-1+\theta} \sum_{k=1}^{k=N} \frac{k}{N} \left(\left(\frac{k}{N} \right)^{-1+\theta} - \left(\frac{k+1}{N} \right)^{-1+\theta} \right) \\
 &\approx -A_N N^{-1+\theta} \int_0^1 x d(x^{-1+\theta}) \\
 &= A_N (1-\theta) N^{-1+\theta} \theta^{-1},
 \end{aligned}$$

and we finally get

$$R(n) = \frac{L(n)}{U} \approx \frac{\theta \Gamma(\theta)}{1-\theta} n^{1-\theta},$$

as desired. ■

Remarks:

1. It is worth noting that again the coefficient of $n^{1-\theta}$ does not depend on N , which is in agreement with the experimental observations in [3].
2. When θ is small we can approximate $V^\theta - 1$ with $\theta \ln V$. Then, comparing the coefficient of $n^{1-\theta}$ in theorem 5 with (22) we see that we recover the non-fluctuating term of the corresponding coefficient for self-similar trees. More precise analysis may reveal possible fluctuating terms in this case as well.

4 Experiments and Discussion

We conducted a set of experiments on a real network in order to provide further evidence of the power law. We used data that was graciously made available by Bill Cheswick and Hal Burch on their site Internet Mapping Project [2]: They run frequent traceroute-style path probe on registered Internet networks: these probes give a good approximation of the paths used to route packets to many Internet networks, and all together give an approximation of the tree starting from one host in the Internet to all other networks. We will use this data to construct a tree topology and a multicast tree.

The above experimentally built tree was used to generate a multicast and to compute $R(n)$. Here is a brief description of our experiment:

- The probes were taken on May 2000.
- The total number of networks probed is 103625.
- The total number of successful probes is $N = 28587$.
- The destination hosts were all drawn randomly.
- After each draw, the size of the multicast tree is computed and updated along with the size of the unicast tree.

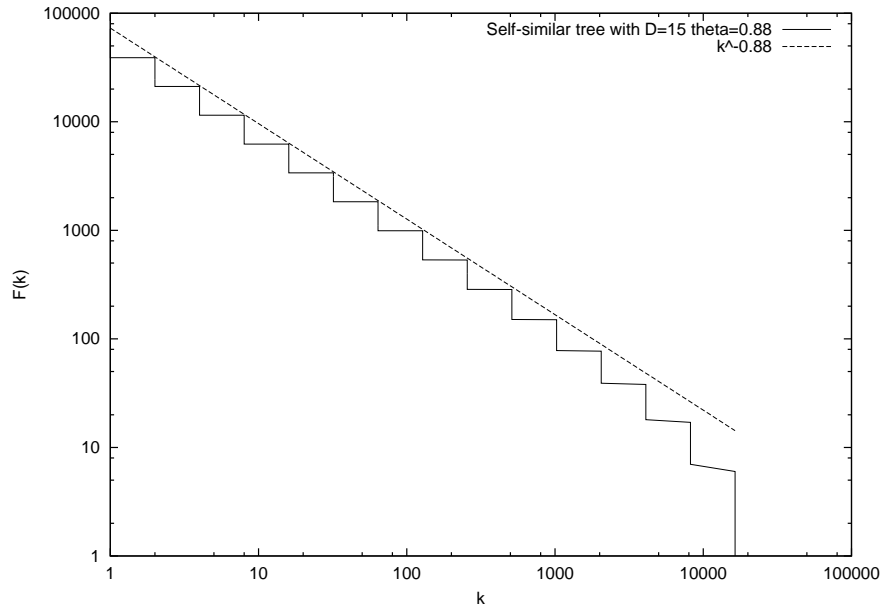


Figure 3: Log-log plots of $F(k)$ versus k for self-similar trees.

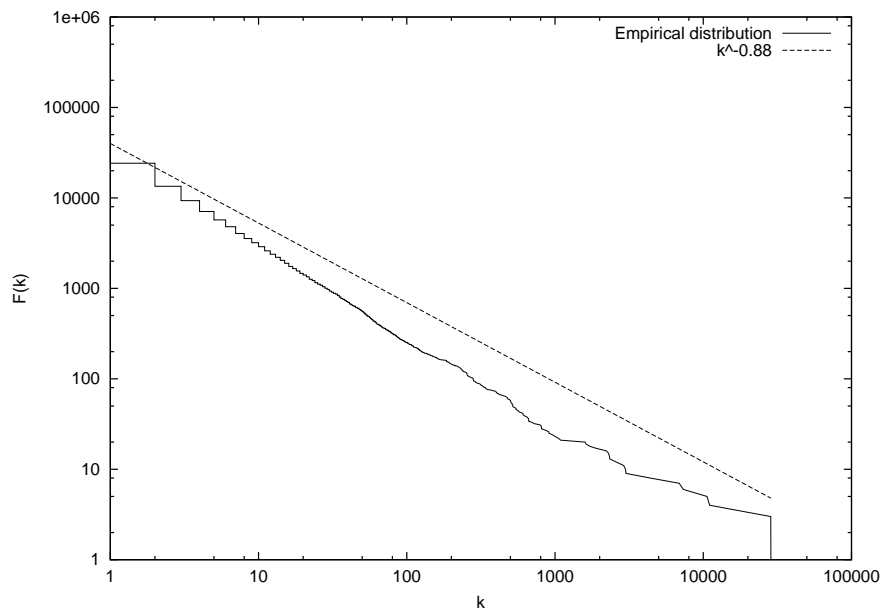
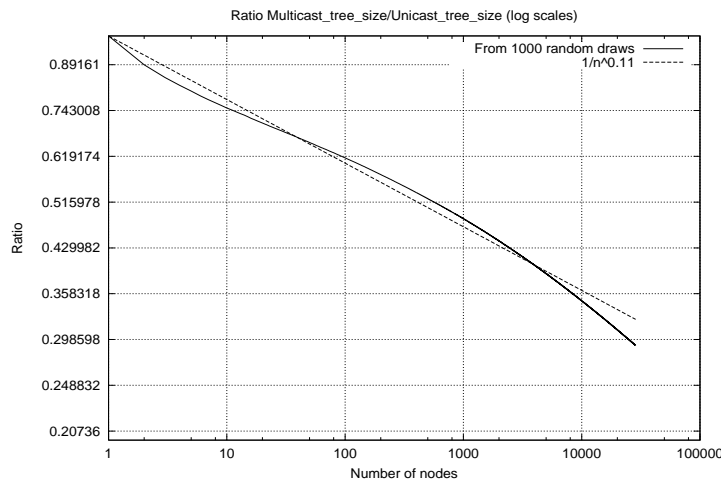


Figure 4: Log-log plots of $F(k)$ versus k for actual internet trees.

Figure 5: Experimental Ratio $R(n)/n$.

Note that not all probes are successful. This is due to the fact that probes rely on heuristics to guess what IP addresses are actually used in a given network. When an IP address is not correctly guessed, the probe isn't taken into account. For each successful probe, a complete path to some network is found that lists all the intermediary routers. Using the set of all these probes, we build a tree. For simplicity, we choose to build a tree, and not a graph, by considering the same router reached by different paths as different nodes. This tree would be a genuine multicast tree if the initial machine was multicasting to all the machines on the Internet, i.e., in our terminology, the global multicast tree. Given n , we pick a random number of n destination hosts and construct the multicast tree connecting the source to the routers at which the destination hosts are located. This multicast tree is a subset of the global multicast tree.

In Figure 4 we plot $F(k)$, i.e., the number of routers in the global multicast tree that have reachability degree k , versus k . We see that $F(k)$ is decreasing according to $k^{-.88}$. From the discussion in Section 3.2 we conclude that for this network, $\theta = 1 - .88 = .12$.

In Figure 5 we plotted $R(n)/n$ versus n together with the curve $f(n) = n^{-0.12}$. In addition, in Figure 6 we plot θ versus n . The approximation seems to be very good.

The previous experiments reaffirm the power law for multicast trees as observed in [3]. Our analysis provides a link between the structural properties of the global multicast tree and the generated multicast tree load. This leads to several new questions: Are the structural properties of the global multicast tree examined in this paper inherent to other networks as well? If so, then what causes the appearance of this structure, and what is its implication in network performance, pricing, simulation etc? Is self-similarity an inherent multicast tree structure, or networks with more general structure (e.g., appropriately behaving $F(k)$) the rule? Self-similarity implies an increased number of unary nodes (or in reality nodes with small outdegree) in the tree. Support for this property provide the experimental results in [6] where is shown that the number of network nodes with small outdegree is a significantly large proportion of the total number of nodes in the network (see Figure 6 in [6]). However, further careful experimental work is needed in order to validate self-similarity. Of course, in real networks one does not expect simple V -ary trees or the appearance of self-similarity in the exact form presented here.

The global multicast tree which was the basis for the study in this paper is a natural basic structure to consider if multicast trees are constructed based on shortest paths. If other mechanisms for the construction of multicast trees are employed, e.g. Steiner trees, then more general basic structures or even the physical network structure may need to be the starting point for the study.

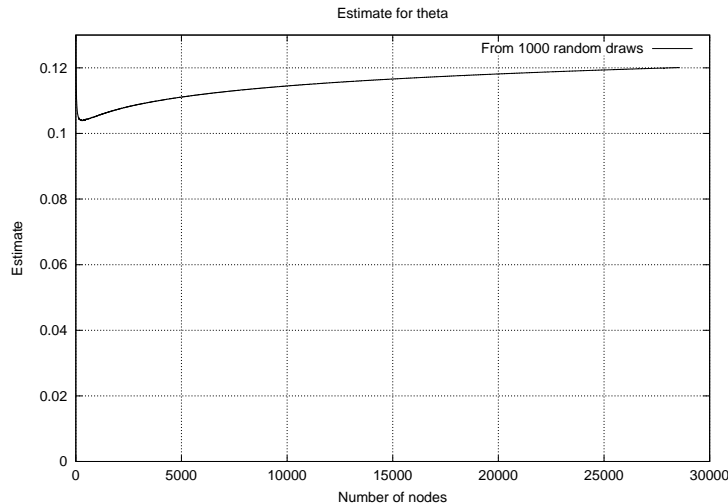


Figure 6: Experimental Curves for Estimating θ versus n .

5 Conclusions

The question whether multicast communication follows the power law arose from the attempt to price multicast communication in [3]. Power laws related to Internet topology parameters were studied in [6]. As indicated in the latter reference, power laws of this nature may be valuable for affirming how realistic simulated topologies are.

In this paper we examined structural conditions on global multicast trees that give rise to the multicast power law. Regular V -ary trees do not exhibit the power law, while self-similar trees do. In fact, the power law rises under conditions weaker than tree self-similarity. Experimental work demonstrated that these latter conditions indeed hold for the tested networks.

Further experimental work is needed to test whether the structural conditions discussed in this paper appear to other networks as well. Moreover, given the fact that several power laws related to various network parameters have been observed experimentally, the question arises as to why these laws appear and whether one law implies the other. In this paper we touched one aspect of this big problem.

A APPENDIX

In this appendix we establish our main theoretical results from Section 2. As observed before, our results for self-similar trees cover as the special case ($\theta = 0$) the finding for regular V -ary trees. Thus we concentrate here only on proving Theorem 3 and Corollary 4.

A.1 Proof of Theorem 3

We saw that the average cost of the multicast self-similar tree is

$$L_\theta(n) = \sum_{k=1}^D V^{(D-k)\theta} V^k \left(1 - (1 - V^{-k})^n \right).$$

Observe that since $D = \log_V N$ and $n = aN$, for $k = D$ the last term of the above sum is approximately equal to $V^D(1 - e^{-a})$. This term is not small in general and thus we cannot extend

the limit of the summation to infinity without introducing significant error. In order to provide an asymptotic analysis of $L_\theta(n)$, we define $\bar{\theta} = 1 - \theta$ and proceed as follows:

$$\begin{aligned} L_\theta(n) &= \sum_{k=1}^D \left(V^{(D-k)\theta} V^k \left(1 - (1 - V^{-k})^n \right) \right) = V^D \sum_{l=0}^{D-1} V^{-\bar{\theta}l} \left(1 - \left(1 - \frac{V^l}{V^D} \right)^n \right) \\ &= V^D \left(\frac{1 - V^{-D\bar{\theta}}}{1 - V^{-\bar{\theta}}} \right) - V^D \sum_{l=0}^{D-1} V^{-\bar{\theta}l} \left(1 - \frac{V^l}{V^D} \right)^n = N \left(\frac{1 - N^{-\bar{\theta}}}{1 - V^{-\bar{\theta}}} \right) - \bar{L}_\theta(n), \end{aligned}$$

where

$$\bar{L}_\theta(n) = N \sum_{l=0}^{D-1} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n} \right)^n.$$

We shall use the following expansion for $x \leq \ln n$

$$\left(1 - \frac{x}{n} \right)^n = \exp(-x) \left(1 - \frac{x^2}{2n} + \frac{x}{n} O\left(\frac{(\ln n)^3}{n} \right) \right). \quad (23)$$

After setting

$$A_n = \left\lfloor \frac{\ln\left(\frac{\ln n}{a}\right)}{\ln V} \right\rfloor,$$

(hence $\ln n < eaV^{A_n} \leq e \ln n$) we obtain

$$\bar{L}_\theta(n) = N \sum_{l=0}^{A_n} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n} \right)^n + N \sum_{l=A_n+1}^{D-1} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n} \right)^n \quad (24)$$

We analyze separately each of the terms on the right hand side of the previous equation. We first look at the term

$$\sum_{l=0}^{A_n} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n} \right)^n.$$

For $l \leq A_n$ we have $aV^l \leq \ln n$ and hence from (23),

$$\begin{aligned} \sum_{l=0}^{A_n} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n} \right)^n &= \sum_{l=0}^{A_n} V^{-\bar{\theta}l} \exp(-aV^l) \left(1 - \frac{a^2V^{2l}}{2n} + \frac{aV^l}{n} O\left(\frac{(\ln n)^3}{n} \right) \right) \\ &= \sum_{l=0}^{A_n} V^{-\bar{\theta}l} \exp(-aV^l) - \frac{a}{2N} \sum_{l=0}^{A_n} V^{(1+\theta)l} \exp(-aV^l) \\ &\quad + \frac{1}{N^\theta} O\left(\frac{(\ln n)^{3+\theta}}{n^{1+\theta}} \right) + \llbracket \theta = 0 \rrbracket \frac{\ln N}{N} O\left(\frac{(\ln n)^3}{n} \right) \end{aligned}$$

where $\llbracket \theta = 0 \rrbracket$ is equal to 1 when $\theta = 0$ and zero otherwise. We now look separately at each of the above terms. We find

$$\sum_{l=0}^{A_n} V^{-\bar{\theta}l} \exp(-aV^l) = \sum_{l=0}^{\infty} V^{-\bar{\theta}l} \exp(-aV^l) - \sum_{l=A_n+1}^{\infty} V^{-\bar{\theta}l} \exp(-aV^l). \quad (25)$$

But

$$\sum_{l=A_n+1}^{\infty} V^{-\bar{\theta}l} \exp(-aV^l) = \frac{1}{N^{\bar{\theta}}} O\left(\frac{1}{n^{\theta} (\ln n)^{\bar{\theta}}}\right), \quad (26)$$

which finally yields

$$\sum_{l=0}^{A_n} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n}\right)^n = c_1(a, \theta) - \frac{1}{2N} c_2(a, \theta) + \frac{1}{N^{\bar{\theta}}} O\left(\frac{1}{n^{\theta} (\ln n)^{\bar{\theta}}}\right), \quad (27)$$

where

$$c_1(a, \theta) = \sum_{l=0}^{\infty} V^{-\bar{\theta}l} \exp(-aV^l) \quad (28)$$

$$c_2(a, \theta) = \sum_{l=0}^{\infty} aV^{(1+\theta)l} \exp(-aV^l). \quad (29)$$

To complete the proof of Theorem 3 we must estimate the term

$$\sum_{l=A_n+1}^{D-1} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n}\right)^n$$

But for $l > A_n$ we find

$$\left(1 - \frac{aV^l}{n}\right)^n < \left(1 - \frac{\ln n}{n}\right)^n = O(1/n),$$

and hence

$$\sum_{l=A_n+1}^{D-1} V^{-\bar{\theta}l} \left(1 - \frac{aV^l}{n}\right)^n = V^{-\bar{\theta}A_n} O\left(\frac{1}{n}\right) = \left(\frac{a}{\ln n}\right)^{\bar{\theta}} O\left(\frac{1}{n}\right) = \frac{1}{N^{\bar{\theta}}} O\left(\frac{1}{n^{\theta} (\ln n)^{\bar{\theta}}}\right). \quad (30)$$

Combining our previous estimates we finally prove Theorem 3.

A.2 Derivation of Corollary 4

We now prove Corollary 4, that is, we find asymptotic expansions of $L_{\theta}(n)$ for n . Observe that we only need to analyze the quantities $c_1(a, \theta)$ and $c_2(a, \theta)$ defined in (15) and (16), respectively. The regimes: (i) $a \rightarrow 1$ and (ii) $a \rightarrow \infty$ are easy and are omitted due to lack of space.

Next we look at the regime $a \rightarrow 0$ which is the most interesting case, and the hardest. It turns out that this case can be handled by a special analytic tool, namely the Mellin transform. The Mellin transform found myriad of applications in the analysis of algorithms. The reader is referred to an excellent survey by Flajolet, Gourdon and Dumas [7] (cf. [8, 10]). For reader convenience, we collected the most important properties of the Mellin transform in Section A.3. In particular, the definition of Mellin transform is given in (33). Property (M2) defines the so called *fundamental strip* of the complex plane where the Mellin transform exists. The *harmonic sum property* (M3) and the *mapping properties* (M4) are crucial. We shall use them to derive asymptotics of $c_1(a, \theta)$ and $c_2(a, \theta)$ as $a \rightarrow 0$.

Let us first consider $c_1(a, \theta) = \sum_{l=0}^{\infty} V^{-\bar{\theta}l} \exp(-aV^l)$. Observe that by (M3) the sum in $c_1(a, \theta) := c_1(a)$ is a harmonic sum with $\lambda_k = V^{-k}$ and $g(x) = e^{-x}$ with $\mu_k = V^k$. Thus the Mellin transform $c_1^*(s)$ with respect to a of $c_1(a)$ is by (M3) (and the well known fact that the Mellin of e^{-x} is the Euler gamma function $\Gamma(s)$ for $\Re(s) > 0$):

$$c_1^*(s) = \frac{\Gamma(s)}{1 - v^{-(1+s-\theta)}}.$$

We now use (M4) to find $c_1(a)$ as $a \rightarrow 0$, that is, we shall find the inverse to the Mellin transform which according to (M1) is

$$c_1(a) = \frac{1}{2\pi i} \int_{\frac{1}{2}-i\infty}^{\frac{1}{2}+i\infty} \frac{\Gamma(s)}{1-v^{-(1+s-\theta)}} x^{-s} ds. \quad (31)$$

The goal is to apply the Cauchy residue theorem. But first we must consider a large rectangle *left* the the line, say from the line $(\frac{1}{2} - i\infty, \frac{1}{2} + i\infty)$ to $(-M - i\infty, -M + i\infty)$ for some large $M > 0$. Due to the factor x^{-s} the left line contributes $O(x^{-M})$ for any $M > 0$, which is negligible. The top and bottom lines of the big rectangle cancel out, thus the integral in (31) is equal to the residues inside the rectangle.

We now evaluate the residues. We note that the function $c_1^*(s)$ has poles at $s_k = -1 + \theta - (2\pi i k)/\ln V$, $k = \pm 1, \pm 2, \dots$. All these poles are single. The pole at 0 has residue

$$c_{0,0} = \frac{V^{\bar{\theta}}}{V^{\bar{\theta}} - 1}$$

The poles at s_k , $k \neq 1$ have residues

$$c_{0,s_k} = \frac{\Gamma(-1 + \theta - (2\pi i k)/\ln V)}{\ln V}.$$

Using now the Reverse Mapping Theorem (M4) and the property $\Gamma(-1 + \theta) = \Gamma(\theta)/(-1 + \theta)$, we see that for $a \rightarrow 0$ we have the expansion for any $M > 0$

$$c_1(a, \theta) = \frac{v^{\bar{\theta}}}{v^{\bar{\theta}} - 1} - \frac{\Gamma(\theta)}{\bar{\theta} \ln v} a^{\bar{\theta}} + a^{\bar{\theta}} \psi_3(\ln a) + O(a^M) \quad (32)$$

where

$$\psi_3(x) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \frac{\Gamma(-1 + \theta - (2\pi i k)/\ln V)}{\ln V} \exp\left((2\pi i k) \frac{x}{\ln V}\right).$$

Now we consider $c_2(a, \theta) = \sum_{l=0}^{\infty} aV^{(1+\theta)l} \exp(-aV^l)$. It is again a harmonic sum, hence by (M2) we find its Mellin transform to be

$$c_2^*(a) = \frac{s\Gamma(s)}{1-v^{-(s-\theta)}}.$$

But $c_2^*(s)$ has a single poles at θ with residue

$$c_{0,0} = \frac{\theta\Gamma(\theta)}{\ln v}$$

and single poles at $s_k = \theta - (2\pi i k)/\ln V$, $k \neq 0$ with residues

$$\frac{(\theta - i2\pi k)\Gamma(\theta - (2\pi i k)/\ln V)}{\ln V}.$$

Hence using again the Reverse Mapping Theorem (M4) we obtain

$$c_2(a, \theta) = \frac{\theta\Gamma(\theta)}{\ln V} a^{-\theta} - a^{-\theta} \psi_4(\ln a) + O(a^M)$$

where

$$\psi_4(x) = \sum_{\substack{k=-\infty \\ k \neq 0}}^{\infty} \frac{(\theta - 2i\pi k/\ln V)\Gamma(\theta - (2\pi i k)/\ln V)}{\ln V} \exp\left((2\pi i k) \frac{x}{\ln V}\right).$$

Combining everything we finally prove Corollary 4.

A.3 Main Properties of Mellin Transform

For the reader convenience, we collected here the main properties of the Mellin transform. For details and proofs see [7, 10].

(M1) DIRECT AND INVERSE MELLIN TRANSFORMS. Let c belong to the *fundamental strip* defined below.

$$f^*(s) := \mathcal{M}(f(x); s) = \int_0^\infty f(x)x^{s-1}dx \quad \Longleftrightarrow \quad f(x) = \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} f^*(s)x^{-s}ds. \quad (33)$$

(M2) FUNDAMENTAL STRIP. The Mellin transform of $f(x)$ exists in the *fundamental strip* $\Re(s) \in (-\alpha, -\beta)$, where

$$f(x) = O(x^\alpha) \quad (x \rightarrow 0), \quad f(x) = O(x^\beta) \quad (x \rightarrow \infty)$$

for $\beta < \alpha$.

(M3) HARMONIC SUM PROPERTY. By linearity and the scale rule $\mathcal{M}(f(ax); s) = a^{-s}\mathcal{M}(f(x); s)$,

$$f(x) = \sum_{k \geq 0} \lambda_k g(\mu_k x) \quad \Longleftrightarrow \quad f^*(s) = g^*(s) \sum_{k \geq 0} \lambda_k \mu_k^{-s}. \quad (34)$$

(M4) MAPPING PROPERTIES (Asymptotic expansion of $f(x)$ and singularities of $f^*(s)$).

$$f(x) = \sum_{(\xi, k) \in A} c_\xi k x^\xi (\log x)^k + O(x^M) \quad \Longleftrightarrow \quad f^*(s) \asymp \sum_{(\xi, k) \in A} c_\xi k \frac{(-1)^k k!}{(s + \xi)^k + 1}. \quad (35)$$

— (i) *Direct Mapping*. Assume that $f(x)$ admits as $x \rightarrow 0^+$ the asymptotic expansion (35) for some $-M < -\alpha$ and $k > 0$. Then for $\Re(s) \in (-M, -\beta)$, the transform $f^*(s)$ satisfies the singular expansion (35)

— (ii) *Reverse Mapping*. Assume that $f^*(s) = O(|s|^{-r})$ with $r > 1$, as $|s| \rightarrow \infty$ and that $f^*(s)$ admits the singular expansion (35) for $\Re(s) \in (-M, -\beta)$. Then $f(x)$ satisfies the asymptotic expansion (35) at $x = 0^+$.

References

- [1] S. Casner and S. Deering, First IETF internet audiocast, ACM Computer Communications, 92-97, 1992.
- [2] B. Cheswick and Hal Burch, Internet Mapping Project, <http://www.cs.bell-labs.com/who/ches/map/index.html>.
- [3] J. Chuang and M. Sirbu, Pricing Multicast Communications: A Cost-based Approach, Proceedings of the INET'98, 1998.
- [4] S. Deering, Multicast Routing in Inetrnetworks and Extended LANs, Computer Communication Review, 18, 1998.
- [5] S. Deering, and D. Cheriton, Multicast Routing in Datagram Internetworks and Extended LANs, ACM. Trans. Computer Systems, 8, 85-110, 1990.
- [6] M. Faloutsos, P. Faloutsos, and C. Faloutsos, On the Power-Law Relationship of the Internet Topology, ACM SIGCOMM, Cambridge, 1999.

-
- [7] P. Flajolet, X. Gourdon, and P. Dumas, Mellin Transforms and Asymptotics: Harmonic sums, *Theoretical Computer Science*, 144, 3–58, 1995.
 - [8] D. E. Knuth, *The Art of Computer Programming. Sorting and Searching*, Vol. 3, Second Edition, Addison-Wesley, Reading, MA, 1998.
 - [9] G. Phillips, S. Shenker, H. Tangmunarunkit, *Scaling of Multicast Trees: Comments on the Chuang-Sirbu scaling Law*, SIGCOMM, Cambridge, 1999.
 - [10] W. Szpankowski, *Average Case Analysis of Algorithms in Sequences*, John Wiley & Sons, New York, 2001.



Unité de recherche INRIA Rocquencourt

Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38330 Montbonnot-St-Martin (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur

INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)

<http://www.inria.fr>

ISSN 0249-6399