



HAL
open science

Performance Metrics for Multicast Flows on Random Trees

Bartłomiej Blaszczyzyn, Konstantin Tchoumatchenko

► **To cite this version:**

Bartłomiej Blaszczyzyn, Konstantin Tchoumatchenko. Performance Metrics for Multicast Flows on Random Trees. [Research Report] RR-4388, INRIA. 2002. inria-00072200

HAL Id: inria-00072200

<https://inria.hal.science/inria-00072200>

Submitted on 23 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.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

*Performance Metrics for Multicast Flows on
Random Trees*

Bartłomiej Błaszczyszyn — Konstantin Tchoumatchenko

N° 4388

Février 2002

THÈME 1

A large blue rectangular area containing the text 'Rapport de recherche' in a white serif font. To the left of the text is a large, light grey stylized 'R' logo. A horizontal grey brushstroke is positioned below the text.

*Rapport
de recherche*

Performance Metrics for Multicast Flows on Random Trees

Bartłomiej Błaszczyszyn* , Konstantin Tchoumatchenko †

Thème 1 — Réseaux et systèmes
Projet TREC

Rapport de recherche n° 4388 — Février 2002 — 22 pages

Abstract: We consider a flow of data packets from one source to many destinations in a communication network represented by a random oriented tree. The multicast mode is characterized by the ability of some tree vertices to replicate the received packets in a way which depends on the number of destinations downstream. We are interested in the cost of packet delivery and in several performance metrics associated with multicast flows on Galton–Watson trees and trees generated by point aggregates of a Poisson process. Such stochastic settings are intended to represent tree-shapes arising in the Internet and in some ad hoc networks.

The main result, in the branching process case, is a functional equation for the p.g.f. of the flow volume; we provide conditions for the existence and uniqueness of a solution and a method to compute it using Picard iterations. In the point process case, we use the stochastic comparison technique developed in percolation theory for Boolean models to provide bounds on the introduced cost functions. We use these results to derive a number of characteristics of these random trees and discuss applications to analytical evaluation of costs and loads induced on a network by a multicast session.

Key-words: random tree, multicast, Galton–Watson process, Poisson process

* ENS/INRIA and Mathematical Institute University of Wrocław, ENS, 45 rue d’Ulm, 75005 Paris, France.
E-mail: Bartek.Blaszczyszyn@ens.fr

† France Telecom R&D, FTR&D - DAC/OAT, 38-40 rue du Général Leclerc, 92131 Issy-les-Moulineaux, France. E-mail: Konstantin.Tchoumatchenko@rd.francetelecom.com

Métriques de performance pour des flots multipoint sur des arbres aléatoires

Résumé : Nous considérons un flux de paquets de données d'une source vers plusieurs destinations dans un réseau de télécommunication représenté par un arbre orienté aléatoire. Dans le mode de transmission multipoint, certains sommets de l'arbre peuvent répliquer les paquets reçus de l'amont vers plusieurs destinations en aval. Nous sommes intéressés par le coût de la transmission d'un paquet et par d'autres métriques de performance pour deux types d'arbres aléatoires: des arbres générés par un processus de Galton–Watson et ceux générés par des agrégats d'un processus ponctuel de Poisson. De telles hypothèses stochastiques apparaissent naturellement dans les arbres multipoint de l'Internet et des réseaux ad hoc.

Le résultat principal dans le cas de processus de branchement est l'équation de la fonction génératrice du volume des transmissions. On fournit une condition d'existence et d'unicité de la solution de cette équation et une méthode de calcul par des itérations de Picard. Dans le cas de processus ponctuel, nous obtenons des bornes sur les fonctions de coût en utilisant des techniques de comparaison stochastique développées dans la théorie de la percolation pour les modèles booléens. Ces résultats nous permettent de dériver certaines caractéristiques d'arbres aléatoires et d'évaluer analytiquement les coûts et les charges induites sur un réseau par une session multipoint.

Mots-clés : arbre aléatoire, session multipoint, processus de Galton–Watson, processus de Poisson.

1 Introduction

Consider a communication network which takes the form of an oriented tree $T = (V, E)$, with vertices $V = \{\mathbf{i}\}$ and edges $E = \{(\mathbf{i}_1, \mathbf{i}_2)\}$ directed from the root \mathbf{i}_0 . Suppose that a bit of data (a packet) is delivered from the source at \mathbf{i}_0 to the receivers at the vertices of T using the following scheme: every vertex gets one packet copy and every edge $(\mathbf{i}_1, \mathbf{i}_2)$ transports exactly as many copies as required to serve the receivers in the whole subtree rooted at \mathbf{i}_2 . Define two transmission modes depending on where the necessary number of packet copies is produced. Call the mode *unicast* if the packet can be replicated only at the root vertex \mathbf{i}_0 , and *multicast* if it can be replicated at any vertex. Hence, in unicast mode, the number of packet copies transported by the edge $(\mathbf{i}_1, \mathbf{i}_2)$ equals the number of receivers in the subtree rooted at \mathbf{i}_2 , whereas, in multicast mode, it never exceeds 1.

In this paper, we focus on the properties of multicast packet flows on random trees. Our aim is to assess the cost and the load induced on a network by a multicast session and to quantify the gain of multicast over unicast in terms of flow volumes and other tree-related performance metrics. This question has been addressed by many experimental studies, but only a few analytical models of multicast trees have been proposed. These latter models include the use regular k -ary trees and so-called fractal k -ary trees of fixed depth as representations for multicast infrastructure [15, 1]. We start off the assumption that multicast trees should have variable depths and branching characteristics and introduce two kinds of distributions: trees generated by Galton–Watson branching processes and by aggregates of points of a Poisson process. Both models will be defined in the next section. The first model represents the logical pattern of packet distribution, while the second reflects a physical view of network nodes and their transmission capabilities in space. Aiming for generality, we consider *mixed* unicast/multicast transmission modes assuming that only a share of vertices can replicate packets, and allow for multiple receivers per vertex.

The motivation for this study comes from telecommunications. A wide range of networking applications, such as conferencing, media streaming, and software distribution, require simultaneous delivery of data from a single source to multiple destinations. One-to-many routing protocols usually construct a distribution tree composed of paths connecting the source to all receivers. Unicast protocols treat packet delivery over different paths separately, which results in redundant transmissions of multiple packets over edges belonging to several paths. Multicast provides a more efficient alternative (at the cost of additional intelligence of network nodes): since packets can be duplicated at the vertices where routing paths diverge, sending one packet copy per tree edge is sufficient. See [2] for a survey on multicast routing techniques.

Empirical evidence shows that, in the Internet, multicast trees exhibit high variability of shapes and branching characteristics [6]. This is also true for *ad hoc* networks, which consist of mobile hosts communicating over wireless links without any centralized control, and whose topology is intrinsically subject to frequent changes [13]. This calls for the development of a stochastic modeling framework that would capture the variations of the multicast infrastructure. Regarding our models, we consider the branching process setup to be mostly fitting for the Internet case, and the point process setup for *ad hoc* networks.

We define our performance metrics as cost functions associated with multicast flows. The principal ones are the total flow volume (i.e, the total number of packets transferred over all edges) and the flow volume at a given vertex (i.e, the number of packets transferred over the edges incident from this vertex). Others include the additive and the multiplicative flow

weights. Several transmission characteristics (packet loss rate, maximal delay), and geometrical tree properties (number of vertices, summary edge length) emerge as special cases.

The main result for multicast trees generated by a branching process is a functional equation for the p.g.f.'s (probability generating functionals) of the two principal cost functions. It is based on simple conservation laws for data flows. We also provide conditions for the existence and uniqueness of a solution and a method for computing it using Picard iterations. In Section 3.4, we give an application example comparing the analytic results on the multicast/unicast gain with some experimental studies carried out for the Internet.

The aggregates of a point process are introduced in Section 2.4 to represent connected nodes of an ad hoc network. They are closely related to clumps in the Boolean model (for definition, see [17] and Example 2.2). Connectivity and clumping properties of the Boolean model have been extensively studied in percolation theory; see Chapter 4 of [11], and [14]. In the wireless context, the sets assigned to the points can be seen as communication areas of individual nodes; i.e., the areas where others can receive packets from them. Boolean models and more general coverage processes were used with similar interpretations in [8] and [9] to investigate the connectivity and the throughput capacity of ad hoc networks, and in [4] for performance evaluation of CDMA protocols.

As is the case for trees spanning the germs of Boolean clumps, trees generated by point aggregates have dependent branches, and this makes the analysis of their distributions difficult. We obtain bounds for multicast performance metrics using the technique of stochastic domination by Galton–Watson trees. This idea is due to Hall [11], who used it to give sufficient conditions for clump finiteness (a.s. and in the mean); see also [5], where higher polynomial and exponential moments of the clump size were considered. In the special case where the aggregates are generated by clumps of unit balls, we refine the bounds by constructing a tighter dominating process.

The rest of the paper is organized as follows. In Section 2, we introduce random marked trees as a model of multicast infrastructure and define the performance metrics associated with packet flows. Section 3 contains the results on metric distributions for Galton–Watson trees. Domination results and bounds for trees generated by point aggregates of a homogeneous Poisson process are derived in Section 4.

2 Multicast flows on oriented trees

2.1 Notions and notation

We model the support of a multicast session by an oriented tree $T = (V, E)$, with vertices constituting a countable set $V = \{\mathbf{i}\}$ of an arbitrary nature and edges $E = \{(\mathbf{i}_1, \mathbf{i}_2)\} \subset V \times V$ directed from the root vertex \mathbf{i}_0 . Call \mathbf{j} an immediate descendant of \mathbf{i} if $(\mathbf{i}, \mathbf{j}) \in E$, and denote by $D(\mathbf{i})$ the set of all the immediate descendants of \mathbf{i} . We call \mathbf{j} a descendant of \mathbf{i} if there exists a sequence of vertices $\pi(\mathbf{i}, \mathbf{j}) = (\mathbf{j}_0, \mathbf{j}_1, \dots, \mathbf{j}_n)$, $n > 0$, such that $\mathbf{j}_0 = \mathbf{i}$, $\mathbf{j}_n = \mathbf{j}$, and $(\mathbf{j}_{k-1}, \mathbf{j}_k) \in E$, for $k = 1, \dots, n$. This sequence $\pi(\mathbf{i}, \mathbf{j})$ is called a path from \mathbf{i} to \mathbf{j} ; if exists, it is obviously unique. Denote by $T(\mathbf{i}) = (V(\mathbf{i}), E(\mathbf{i}))$ the subtree of T generated by the vertex \mathbf{i} and by all of its descendants (hence $T(\mathbf{i}_0) \equiv T$). Let $T^{(n)}(\mathbf{i}) = (V^{(n)}(\mathbf{i}), E^{(n)}(\mathbf{i}))$ be the truncation of the tree $T(\mathbf{i})$ at the depth n ; i.e., the subtree including only those vertices \mathbf{j} of $T(\mathbf{i})$ that can be reached by a path $\pi(\mathbf{i}, \mathbf{j})$ with no more than $n + 1$ vertices, or, equivalently, in at most n hops.

Introduce marks $m(\mathbf{i}) = (r(\mathbf{i}), \sigma(\mathbf{i}), \tau(\mathbf{i}))$ representing vertex characteristics. Here $\sigma(\mathbf{i}) \in \{0, 1\}$ is the indicator of multicast ability; i.e., the ability to replicate a received packet, $\tau(\mathbf{i}) \in \{0, 1, \dots\}$ is the number of end receivers at the vertex; i.e., the number of identical packet copies requested by that vertex, and $r(\mathbf{i}) \in \mathbb{E}$ is the vertex type. We assume $\mathbb{E} = \{1, 2, \dots, \ell\}$ everywhere except Section 4.3, where we consider $\mathbb{E} = [0, 1]$.

Throughout the paper we denote by $\|\cdot\|$ the Euclidean norm in \mathbb{R}^d , by $\text{vol}(\cdot)$ the d -dimensional Lebesgue measure, and by $b_y(x)$ a closed ball in \mathbb{R}^d of radius y centered at x .

2.2 Performance metrics and flow conservation laws

Consider a multicast flow on an oriented marked tree $T = (E, V)$ in which only multicast-enabled vertices can replicate packets. As in Section 1, a packet is delivered from \mathbf{i}_0 to the end receivers at the vertices of T , so that every receiver gets one packet copy and every edge $(\mathbf{i}_1, \mathbf{i}_2)$ transports exactly as many copies as required to serve all receivers in the subtree $T(\mathbf{i}_2)$. Hence, if $\sigma(\mathbf{i}_2) = 1$ then this number equals 0 or 1 depending on the existence of a vertex $\mathbf{i} \in V(\mathbf{i}_2)$ such that $\tau(\mathbf{i}) > 0$. If $\sigma(\mathbf{i}_2) = 0$ then this number equals $\tau(\mathbf{i}_2)$ plus the number of copies sent over the edges incident from \mathbf{i}_2 .

The principal performance metrics that we consider are $K(\mathbf{i})$, the total number of packets sent by vertex \mathbf{i} to its immediate descendants, and $L(\mathbf{i})$, the total number of packet transfers within the whole subtree $T(\mathbf{i})$. Let us provide formal definitions. Denote by $R(\mathbf{i}) = \{\mathbf{j} \in V(\mathbf{i}) : \mathbf{j} \neq \mathbf{i}, \tau(\mathbf{j}) \geq 1\}$ the set of all descendants of \mathbf{i} that request at least one packet. For every $\mathbf{j} \in R(\mathbf{i})$, consider the path $\pi(\mathbf{i}, \mathbf{j})$ and denote by \mathbf{j}^* either the first occurrence of a multicast vertex $\neq \mathbf{i}$ on this path, or \mathbf{j} , if such vertex does not occur. It is easy to see that the number of packet copies sent by \mathbf{i} in order to satisfy the requests in $R(\mathbf{i})$ depends on the number of copies requested by the members of $R^*(\mathbf{i}) = \{\mathbf{j}^* : \mathbf{j} \in R(\mathbf{i})\}$. We define

$$K(\mathbf{i}) = \sum_{\mathbf{j} \in R^*(\mathbf{i})} \tau(\mathbf{j}) \mathbf{1}\{\sigma(\mathbf{j}) = 0\} + \mathbf{1}\{\sigma(\mathbf{j}) = 1\}. \quad (2.1)$$

The dual function $\overline{K}(\mathbf{i})$ is defined as the number of packet copies received by vertex \mathbf{i} in order to satisfy every packet request within $T(\mathbf{i})$. It can be expressed as

$$\overline{K}(\mathbf{i}) = \sigma(\mathbf{i})(1 - \mathbf{1}\{\tau(\mathbf{i}) = 0\}J(\mathbf{i})) + (1 - \sigma(\mathbf{i}))(\tau(\mathbf{i}) + K(\mathbf{i})), \quad (2.2)$$

where $J(\mathbf{i})$ equals $\mathbf{1}\{K(\mathbf{i}) = 0\} \equiv \mathbf{1}\{R(\mathbf{i}) = \emptyset\}$. The total multicast flow volume in $T(\mathbf{i})$ is then given by

$$L(\mathbf{i}) = \sum_{\mathbf{j} \in V(\mathbf{i}) \setminus \{\mathbf{i}\}} \overline{K}(\mathbf{j}). \quad (2.3)$$

As a refinement of $L(\mathbf{i})$, consider the total number of packets $L_{kl}(\mathbf{i})$ transmitted within the tree $T(\mathbf{i})$ over the subset of edges $\{(\mathbf{j}_1, \mathbf{j}_2) \in E(\mathbf{i}) : r(\mathbf{j}_1) = k \text{ and } r(\mathbf{j}_2) = l\}$

$$L_{kl}(\mathbf{i}) = \sum_{(\mathbf{j}_1, \mathbf{j}_2) \in E(\mathbf{i})} \overline{K}(\mathbf{j}_2) \mathbf{1}\{r(\mathbf{j}_1) = k, r(\mathbf{j}_2) = l\}. \quad (2.4)$$

More generally, let $w(\mathbf{i}, \mathbf{j}, k)$, $k = 1, 2, \dots, \overline{K}(\mathbf{j})$, be costs associated with individual packets sent over the edge (\mathbf{i}, \mathbf{j}) . The *additive* cost $W_{\Sigma}(\mathbf{i})$ of the multicast flow on $T(\mathbf{i})$ is the sum of

the costs of all packet transfers. It is given by

$$W_{\Sigma}(\mathbf{i}) = \sum_{(\mathbf{j}_1, \mathbf{j}_2) \in E(\mathbf{i})} \left[\sum_{k=1}^{\overline{K}(\mathbf{j}_2)} w(\mathbf{j}_1, \mathbf{j}_2, k) \right]. \quad (2.5)$$

The *multiplicative* cost $W_{\Pi}(\mathbf{i})$ of the multicast flow on $T(\mathbf{i})$ is defined by

$$W_{\Pi}(\mathbf{i}) = \prod_{(\mathbf{j}_1, \mathbf{j}_2) \in E(\mathbf{i})} \left[\prod_{k=1}^{\overline{K}(\mathbf{j}_2)} w(\mathbf{j}_1, \mathbf{j}_2, k) \right]. \quad (2.6)$$

Example 2.1 Obviously, if $w(\mathbf{j}_1, \mathbf{j}_2, k) \equiv 1$ then $W_{\Sigma}(\mathbf{i}) = L(\mathbf{i})$. If, moreover, $\sigma(\cdot) \equiv 1$ and $\tau(\cdot) \equiv 1$, then $L(\mathbf{i})$ equals the total number of edges in the tree $T(\mathbf{i})$. Suppose that packet transmissions may fail and let $w(\mathbf{j}_1, \mathbf{j}_2, k)$ be the number of tries necessary to successfully transmit the packet k over the edge $(\mathbf{j}_1, \mathbf{j}_2)$. Then, $W_{\Sigma}(\mathbf{i})$ is the total number of transmission attempts for the tree $T(\mathbf{i})$.

As an example of a multiplicative cost, consider

$$w(\mathbf{j}_1, \mathbf{j}_2, k) = \mathbf{1}(d(\mathbf{j}_1, \mathbf{j}_2, k) \leq t), \quad (2.7)$$

where $d(\mathbf{j}_1, \mathbf{j}_2, k)$ is interpreted as a waiting time for the delivery acknowledgment of the packet k and t is a fixed timeout. In this case, $W_{\Pi}(\mathbf{i})$ is the indicator that no timeout occurs for all packet copies sent within $T(\mathbf{i})$.

Multicast flows on oriented trees satisfy several simple conservation laws. The equality between the number of sent and received packets writes as

$$K(\mathbf{i}) = \sum_{\mathbf{j} \in D(\mathbf{i})} \overline{K}(\mathbf{j}). \quad (2.8)$$

For $L(\mathbf{i})$, the following recurrent relation holds

$$L(\mathbf{i}) = \sum_{\mathbf{j} \in D(\mathbf{i})} \overline{K}(\mathbf{j}) + L(\mathbf{j}). \quad (2.9)$$

Similarly to (2.9), we have a recursion formula for $L_{kl}(\mathbf{i})$

$$L_{kl}(\mathbf{i}) = \sum_{\mathbf{j} \in D(\mathbf{i})} \mathbf{1}\{r(\mathbf{i}) = k, r(\mathbf{j}) = l\} \overline{K}(\mathbf{j}) + L_{kl}(\mathbf{j}). \quad (2.10)$$

Denote by $K^{(n)}(\mathbf{i})$, $\overline{K}^{(n)}(\mathbf{i})$, $J^{(n)}(\mathbf{i})$ and $L^{(n)}(\mathbf{i})$ the appropriate cost functionals calculated with respect to the truncated tree $T^{(n)}(\mathbf{i})$. Relations similar to (2.8)–(2.9) hold

$$K^{(n+1)}(\mathbf{i}) = \sum_{\mathbf{j} \in D(\mathbf{i})} \overline{K}^{(n)}(\mathbf{j}), \quad (2.11)$$

where

$$\overline{K}^{(n)}(\mathbf{i}) = \sigma(\mathbf{i})(1 - \mathbf{1}\{\tau(\mathbf{i}) = 0\}J^{(n)}(\mathbf{i})) + (1 - \sigma(\mathbf{i}))(\tau(\mathbf{i}) + K^{(n)}(\mathbf{i}))$$

$$J^{(n)}(\mathbf{i}) = \mathbf{1}\{K^{(n)}(\mathbf{i}) = 0\}$$

and

$$L^{(n+1)}(\mathbf{i}) = \sum_{\mathbf{j} \in D(\mathbf{i})} \overline{K}^{(n)}(\mathbf{j}) + L^{(n)}(\mathbf{j}). \quad (2.12)$$

2.3 Galton–Watson trees

Now, we place ourself in a probabilistic settings and define random trees generated by a multi-type Galton–Watson branching process with types in $\mathbb{E} = \{1, 2, \dots, \ell\}$ (for theory of branching processes see, e.g., [12] and [3]). Individuals of the process constitute the vertices of a random tree that we denote by \mathcal{T}_B , whose edges connect every individual to its direct descendants.

The progeny of an individual \mathbf{i} is the vector $Z(\mathbf{i}) = (Z_m(\mathbf{i}))_{m \in \mathbb{E}}$, where $Z_m(\mathbf{i})$ is the number of direct descendants of \mathbf{i} having type m (the root type $r(\mathbf{i}_0)$ should be defined separately). We make a standard assumption for branching processes that, for \mathbf{i} 's of the same generation, all $Z(\mathbf{i})$ are mutually independent. Hence, the process is defined by the distribution of $Z(\mathbf{i})$, or equivalently, by the family of conditional p.g.f.'s $\psi = (\psi_k)_{k \in \mathbb{E}}$ acting on $z = (z_m)_{m \in \mathbb{E}}$

$$\psi_k(z) = \mathbf{E} \exp \left[\prod_{m \in \mathbb{E}} z_m^{Z_m(\mathbf{i})} \mid r(\mathbf{i}) = k \right]. \quad (2.13)$$

The marks of non-root vertices $r(\cdot) \in \mathbb{E}$ are thus determined by the types of the corresponding individuals. Regarding the multicast ability and the number of end receivers ($\sigma(\cdot), \tau(\cdot)$), we assume that their joint distribution depends only on $r(\cdot)$ and is given by the set of probabilities

$$p_{ij}^k = \mathbf{P}(\sigma(\cdot) = i, \tau(\cdot) = j \mid r(\cdot) = k), \quad i \in \{0, 1\}, \quad j \in \{0, 1, \dots\}. \quad (2.14)$$

Hence, the distribution of the marked tree $\mathcal{T}_B = (V_B, E_B)$ is completely defined.

2.4 Trees generated by aggregates of a Poisson point process

Trees that we introduce now have their vertex sets embedded in a homogeneous Poisson point process $\Pi = \{x_i\}_{i \in \mathbb{N}}$ on \mathbb{R}^d with intensity λ . Assume that the points of Π have i.i.d. marks $m(x_i) = (r(x_i), \sigma(x_i), \tau(x_i))$ taking values in $\mathbb{E} \times \{0, 1\} \times \{0, 1, \dots\}$, so that the distribution of $r(\cdot)$ is given by the set of probabilities

$$q_k = \mathbf{P}(r(\cdot) = k), \quad k \in \mathbb{E} \quad (2.15)$$

and the conditional distribution of $(\sigma(\cdot), \tau(\cdot))$ by (2.14). Note that, by this definition, the sets $\Pi_l = \{x_i \in \Pi \mid r(x_i) = l\}$, $l \in \mathbb{E}$, are independent Poisson processes with intensities λq_l .

Let us define bonds between points of different types using a collection of closed bounded sets $\{G_{kl} \subset \mathbb{R}^d : k, l \in \mathbb{E}\}$. Denote by $x_i \rightsquigarrow x_j$ the relation $x_j \in x_i + G_{r(x_i)r(x_j)}$. Put $A_0(x_i) = \{x_i\}$ and define by induction the set of n -accessible points $A_n(x_i)$ as all $x_j \in \Pi$, such that, for some $x_m \in A_{n-1}(x_i)$, $x_m \rightsquigarrow x_j$ and $x_j \notin A^{(n-1)}(x_i) \equiv \bigcup_{k=0}^{n-1} A_k(x_i)$. The set of all points accessible from x_i is then given by

$$A(x_i) = \bigcup_n A_n(x_i)$$

and can be seen as a mark of x_i . We will call the set $A(x_i)$ the *aggregate* associated with x_i . In what follows, we will be interested in the properties of a *typical* aggregate; i.e., having the Palm distribution with respect to the underlying Poisson point process Π . By Slivnyak's theorem (see, e.g., [17]), its distribution coincides with the distribution of the aggregate $A(0)$ constructed with respect to the process $\Pi \cup \{0\}$ with an independent $m(0)$ having the common mark distribution. Considering the Palm distribution, we will write simply A and A_n to refer,

respectively, to the typical aggregate and to the set of n -accessible points with respect to the origin.

With the typical aggregate A we associate an oriented tree $\mathcal{T}_C = (A, E_C)$ rooted at the origin and having the edge set

$$E_C = \{(v(x_i), x_i), x_i \in A \setminus \{0\}\},$$

where for every x_i such that $x_i \in A_n$, $n \geq 1$, the ancestor $v(x_i)$ is chosen by independent sampling from $\{x_j \in A_n : x_j \rightsquigarrow x_i\}$ assuming equal probability for all elements. Such construction yields a tree connecting the origin to every point of A in the least possible number of hops ($v(\cdot)$ sampling guarantees the uniqueness of ancestors in such tree). Marks $\{m(x_i), x_i \in A\}$ become the marks of the vertices of \mathcal{T}_B .

Example 2.2 Recall that a classical Boolean model in \mathbb{R}^d is a random set $\Xi = \bigcup_i (x_i + B_i)$ generated by a Poisson point process $\Pi = \{x_i\}_{i \in \mathbb{N}}$ and a sequence of i.i.d. compact subsets $\{B_i\}_{i \in \mathbb{N}}$ of \mathbb{R}^d , where $x + B = \{y + x \in \mathbb{R}^d : y \in B\}$. Maximal connected subsets of Ξ are called clumps. Consider a Boolean model with all B_i being some fixed $B \subset \mathbb{R}^d$. Then $(x_i + B) \cap (x_j + B)$ is non-empty if and only if $x_j \in x_i + G$, where $G = \{y_1 - y_2; y_1, y_2 \in B\}$. Hence, when $\mathbb{E} = \{1\}$ and $G_{11} = G$, the aggregate $A(x_i)$ consists of all points of Π covered by the Boolean clump containing x_i .

The following remark will be used in Section 4 for the construction of \mathcal{T}_C .

Remark 2.3 Note that the aggregate A is recurrently defined through its truncations $A^{(n)} \equiv \bigcup_{k=0}^n A_k$, $n \geq 0$. This sequence has the following Markov property: $A^{(n+1)}$ depends on the previous truncations $A^{(k)}$, $k \leq n$, only through $A^{(n)}$ and $A^{(n-1)}$. Indeed, $A^{(n+1)} = A^{(n)} \cup A_{n+1}$ and all points of A_{n+1} of types $l \in \mathbb{E}$ are distributed as independent homogeneous Poisson processes with intensities λq_l in $C_l^{(n+1)} \setminus C_l^{(n)}$, where

$$C_l^{(k)} = \bigcup_{y_i \in A^{(k-1)}} (y_i + G_{r(y_i)l}), \quad k = 1, 2, \dots \quad (2.16)$$

and $C_l^{(0)} = \emptyset$.

3 Flow metrics for Galton–Watson trees

The aim of this section is to characterize the distributions of our multicast performance metrics of the trees generated by branching processes, introduced in Section 2.3. To simplify the exposition, we first limit the scope to the two principal cost functions $K(\mathbf{i})$ and $L(\mathbf{i})$ defined by (2.1) and (2.3), and use the flow conservation laws from Section 2.2 to derive recurrent equations for their p.g.f.'s. Next, we extend these results to other metrics from Section 2.2. As a corollary, we derive the first moments of the flow volumes, which we use to evaluate the efficiency gain of multicast over unicast.

3.1 Recurrent relations between p.g.f.'s of multicast metrics

In the settings of Section 2.3, consider a family of functions $\phi = (\phi_k)_{k \in \mathbb{E}}$ in which the element $\phi_k = \phi_k(z_1, z_2)$ is the joint p.g.f. of the couple $(K(\mathbf{i}), L(\mathbf{i}))$ under condition $r(\mathbf{i}) = k$

$$\phi_k(z_1, z_2) = \mathbf{E} [z_1^{K(\mathbf{i})} z_2^{L(\mathbf{i})} | r(\mathbf{i}) = k].$$

Introduce also $\phi^{(n)} = (\phi_k^{(n)})_{k \in \mathbb{E}}$ as the family of the p.g.f.'s of $(K^{(n)}(\mathbf{i}), L^{(n)}(\mathbf{i}))$ given $r(\mathbf{i}) = k$.

Define an operator $\Psi[\cdot] = (\Psi_k[\cdot])_{k \in \mathbb{E}}$ acting on the family $u = (u_k)_{k \in \mathbb{E}}$ of functions $u_k = u_k(z_1, z_2)$ as follows

$$\Psi_k[u] = \psi_k(f(u)), \quad (3.1)$$

where $f(u) = (f_m(u))_{m \in \mathbb{E}}$ is given by

$$f_m(u) = (1 - z_1 z_2) u_m(0, z_2) p_{10}^m + z_1 z_2 u_m(1, z_2) \sum_{j=0}^{\infty} p_{1j}^m + u_m(z_1 z_2, z_2) \sum_{j=0}^{\infty} (z_1 z_2)^j p_{0j}^m. \quad (3.2)$$

Proposition 3.1 *The family of p.g.f.'s $\phi^{(n)}$, $n = 1, 2, \dots$, satisfies the following recursion relation*

$$\phi^{(n+1)} = \Psi[\phi^{(n)}] \quad (3.3)$$

with $\Psi[\cdot]$ given by (3.1) and (3.2).

Proof: From the relations (2.11)–(2.12) it follows that if $r(\mathbf{i}) = k$, then $K^{(n+1)}(\mathbf{i})$ and $L^{(n+1)}(\mathbf{i})$ can be represented as a sum of $Z_k(\mathbf{i})$ independent r.v.'s $\bar{K}^{(n)}(\mathbf{j})$ and $\bar{K}^{(n)}(\mathbf{j}) + L^{(n)}(\mathbf{j})$, respectively. If we show that

$$\mathbf{E} \left[z_1^{\bar{K}^{(n)}(\mathbf{j})} z_2^{\bar{K}^{(n)}(\mathbf{j}) + L^{(n)}(\mathbf{j})} \mid r(\mathbf{j}) = m \right] = f_m(\phi^{(n)}(z_1, z_2)), \quad (3.4)$$

then from (2.13) it will follow that

$$\phi_k^{(n+1)}(z_1, z_2) = \psi_k(f(\phi^{(n)}(z_1, z_2)))$$

and the proposition will be proved. Rewrite the left-hand side of (3.4) using (2.2) and conditioning on the distribution of $(\sigma(\mathbf{j}), \tau(\mathbf{j}))$ as follows

$$\begin{aligned} & \mathbf{E} \left[\left((z_1 z_2)^{1 - J^{(n)}(\mathbf{j})} p_{10}^m + z_1 z_2 \sum_{j=1}^{\infty} p_{1j}^m + \sum_{j=0}^{\infty} p_{0j}^m (z_1 z_2)^{j + K^{(n)}(\mathbf{j})} \right) z_2^{L^{(n)}(\mathbf{j})} \mid r(\mathbf{j}) = m \right] \\ &= \mathbf{E} \left[\left((1 - z_1 z_2) J^{(n)}(\mathbf{j}) p_{10}^m + z_1 z_2 \sum_{j=0}^{\infty} p_{1j}^m + (z_1 z_2)^{K^{(n)}(\mathbf{j})} \sum_{j=0}^{\infty} (z_1 z_2)^j p_{0j}^m \right) z_2^{L^{(n)}(\mathbf{j})} \mid r(\mathbf{j}) = m \right] \\ &= (1 - z_1 z_2) \phi_m^{(n)}(0, z_2) p_{10}^m + z_1 z_2 \phi_m^{(n)}(1, z_2) \sum_{j=0}^{\infty} p_{1j}^m + \phi_m^{(n)}(z_1 z_2, z_2) \sum_{j=0}^{\infty} (z_1 z_2)^j p_{0j}^m. \end{aligned}$$

Comparing with (3.2), we see that (3.4) is true. ■

Note that if there is a positive probability that $K(\mathbf{i})$ or $L(\mathbf{i})$ is infinite, then the p.g.f.'s $(\phi_k)_{k \in \mathbb{E}}$ are not well defined. A sufficient condition for the finiteness of these functionals is the a.s. extinction of the branching process generating the tree \mathcal{T}_B . Denote by $\Lambda = (\lambda_{km})_{k,m \in \mathbb{E}}$ the matrix of the first moments

$$\lambda_{km} = \mathbf{E}[Z_m(\mathbf{i}) \mid r(\mathbf{i}) = k]$$

and recall the following standard result (see, e.g., [12])

Proposition 3.2 *Suppose that the matrix of the first moments $\Lambda = (\lambda_{km})_{k,m \in \mathbb{E}}$ exists and is positively regular; i.e., for some $n > 0$, all elements of Λ^n are strictly positive, and that the branching process is non-singular; i.e., not all individuals have exactly one descendant. If the maximal eigenvalue $\rho(\Lambda) \leq 1$ then the branching process becomes extinct a.s.*

A weaker sufficient condition for the finiteness of $K(\mathbf{i})$ will be given in Remark 3.10. The following two statements summarize the influence of $\rho(\Lambda)$ on the distributions of the principle cost functionals. Consider the space of vector functions $u = (u_k)_{k \in \mathbb{E}}$ of the argument $(z_1, z_2) \in [0, 1]^2$ equipped with the norm

$$\|u\|_\infty = \sup_{k \in \mathbb{E}, (z_1, z_2) \in [0, 1]^2} |u_k(z_1, z_2)|.$$

Lemma 3.3 *Suppose that the matrix of the first moments Λ exists. Then the n -th iteration of the operator Ψ given by (3.1)-(3.2) satisfies the inequality: for $u, u' \in \{v : \|v\|_\infty \leq 1\}$,*

$$\|\Psi^n(u) - \Psi^n(u')\|_\infty \leq \|\Lambda^n |u - u'|\|_\infty, \quad n \geq 1, \quad (3.5)$$

where $|u - u'| = (|u_k - u'_k|)_{k \in \mathbb{E}}$. Thus, if any norm of Λ^n vanishes as $n \rightarrow \infty$ (in particular, if Λ is positively regular and $\rho(\Lambda) < 1$) then

$$\lim_{n \rightarrow \infty} \|\Psi^n(u) - \Psi^n(u')\|_\infty = 0.$$

Proof: Let $x = (x_k)_{k \in \mathbb{E}}$ and $x' = (x'_k)_{k \in \mathbb{E}}$ be two vectors in \mathbb{R}^ℓ such that $\max_k |x_k| \leq 1$ and $\max_k |x'_k| \leq 1$. Denoting

$$x_{|m} = (x'_1, x'_2, \dots, x'_{m-1}, x_m, x_{m+1}, \dots, x_\ell)$$

and using the absolute continuity property of p.g.f.'s (since the first moments exist), we have

$$|\psi_k(x) - \psi_k(x')| \leq \sum_{m \in \mathbb{E}} |\psi_k(x_{|m}) - \psi_k(x_{|m+1})| \leq \sum_{m \in \mathbb{E}} \lambda_{km} |x_m - x'_m|.$$

Note that by (3.2),

$$0 \leq f_m(u) \leq \sup_{(z_1, z_2) \in [0, 1]^2} u_m(z_1, z_2).$$

Therefore, for functions u and u' such that $\|u\|_\infty \leq 1$ and $\|u'\|_\infty \leq 1$,

$$\begin{aligned} \|\Psi(u) - \Psi(u')\|_\infty &= \sup_{k, z_1, z_2} |\psi_k(f_k(u)) - \psi_k(f_k(u'))| \\ &\leq \sup_{k, z_1, z_2} \sum_{m \in \mathbb{E}} \lambda_{km} |f_m(u) - f_m(u')| \\ &\leq \sup_{k, z_1, z_2} \sum_{m \in \mathbb{E}} \lambda_{km} |u_m - u'_m| \\ &\leq \|\Lambda |u - u'|\|_\infty, \end{aligned}$$

where \sup_{k, z_1, z_2} is taken over $k \in \mathbb{E}$, $(z_1, z_2) \in [0, 1]^2$. Using induction by n , we obtain (3.5). ■

Proposition 3.4 *Suppose that the matrix of the first moments Λ is positively regular, nonsingular, and $\rho(\Lambda) \leq 1$. Then for any $r(\mathbf{i}) = k \in \mathbb{E}$, the functionals $K(\mathbf{i})$ and $L(\mathbf{i})$ are a.s. finite, for any $(z_1, z_2) \in [0, 1]^2$, $\phi_k(z_1, z_2) = \lim \phi_k^{(n)}(z_1, z_2)$ as $n \rightarrow \infty$, and $\phi = (\phi_k)_{k \in \mathbb{E}}$ satisfies the functional equation*

$$\Psi[\phi] = \phi. \tag{3.6}$$

Moreover, if $\rho(\Lambda) < 1$, then the family of p.g.f.'s ϕ is the only solution of (3.6), such that $\|\phi\|_\infty \leq 1$.

Proof: The first part of the statement follows immediately from the fact that the functionals $K^{(n)}(\mathbf{i})$ and $L^{(n)}(\mathbf{i})$ grow monotonously with n and converge to $K(\mathbf{i})$ and $L(\mathbf{i})$ a.s. By the monotone convergence theorem they converge also in the mean, though the latter may be infinite. The second part follows from Lemma 3.3. ■

Corollary 3.5 *In the special case when $\sigma(\cdot) \equiv 1$ and $\tau(\cdot) \equiv 1$, $\phi_k(1, z_2)$ is the p.g.f. of the number of descendants of an individual of type k . Putting $z_1 = 1$ in (3.6), we obtain a well-known system of equations*

$$\psi_k \left(z_2 (\phi_1(1, z_2), \dots, \phi_\ell(1, z_2)) \right) = \phi_k(1, z_2) \quad k = 1, \dots, \ell.$$

If $\sigma(\cdot) \equiv 0$ and $\tau(\cdot) \equiv 1$, then $K(\mathbf{i})$ equals the total number of edges in the tree $T(\mathbf{i})$ and $L(\mathbf{i}) = \sum_{\mathbf{j} \in V(\mathbf{i})} K(\mathbf{j})$. Equation (3.6) thus becomes

$$\psi_k \left(z_1 z_2 (\phi_1(z_1 z_2, z_2), \dots, \phi_\ell(z_1 z_2, z_2)) \right) = \phi_k(z_1, z_2) \quad k = 1, \dots, \ell.$$

3.2 Multitype extension

It is straightforward to extend the basic stochastic recursion to the case of type-dependent functionals $L_{lm}(\mathbf{i})$ given by (2.4). Put $\bar{z} = (z_{lm})_{l, m \in \mathbb{E}}$. For a vertex \mathbf{i} of \mathcal{T}_B , define the family $\tilde{\phi} = (\tilde{\phi}_k)_{k \in \mathbb{E}}$ of joint p.g.f.'s of $K(\mathbf{i})$ and $(L_{lm}(\mathbf{i}))_{l, m \in \mathbb{E}}$ given $r(\mathbf{i}) = k$ as

$$\tilde{\phi}_k(z, \bar{z}) = \mathbf{E} \left[z_1^{K(\mathbf{i})} \prod_{l, m \in \mathbb{E}} z_{lm}^{L_{lm}(\mathbf{i})} \mid r(\mathbf{i}) = k \right].$$

Define an operator $\tilde{\Psi} = (\tilde{\Psi}_k)_{k \in \mathbb{E}}$ acting on the family $u = (u_k)_{k \in \mathbb{E}}$ of functions $u_k = u_k(z, \bar{z})$

$$\tilde{\Psi}_k[u] = \tilde{\psi}_k(\tilde{f}(u)),$$

where $\tilde{f}(u) = (\tilde{f}_m(u))_{m \in \mathbb{E}}$ is given by

$$\begin{aligned} \tilde{f}_m(u(z, \bar{z})) &= (1 - z z_{km}) u_m(0, \bar{z}) p_{10}^m + z z_{km} u_m(1, \bar{z}) \sum_{j=0}^{\infty} p_{1j}^m + \\ &+ u_m(z z_{km}, \bar{z}) \sum_{j=0}^{\infty} (z z_{km})^j p_{0j}^m. \end{aligned}$$

Now, relations (2.8) and (2.10) hold and the following is true.

Proposition 3.6 *Suppose that the matrix of the first moments Λ is positively regular, nonsingular, and $\rho(\Lambda) \leq 1$. Then, for any \mathbf{i} of \mathcal{T}_B , the functionals $K(\mathbf{i})$ and $L_{lm}(\mathbf{i})$ are a.s. finite, for any $(z, \bar{z}) \in [0, 1]^{1+\ell^2}$, $\tilde{\phi}_k(z, \bar{z}) = \lim_{n \rightarrow \infty} \tilde{\phi}_k^{(n)}(z, \bar{z})$ as $n \rightarrow \infty$, and $\tilde{\phi} = (\tilde{\phi}_k)_{k \in \mathbb{E}}$ satisfies the functional equation*

$$\tilde{\Psi}[\tilde{\phi}] = \tilde{\phi}. \quad (3.7)$$

Moreover, if $\rho(\Lambda) < 1$, then the family of p.g.f.'s $\tilde{\phi}$ is the only solution of (3.7) such that $\|\tilde{\phi}'\|_\infty \leq 1$, where

$$\|u\|'_\infty = \sup_{k \in \mathbb{E}, (z, \bar{z}) \in [0, 1]^{1+\ell^2}} |u_k(z, \bar{z})|.$$

The proofs are similar to those of Proposition 3.1 and 3.4.

Suppose now that the link costs $w(\mathbf{i}, \mathbf{j}, k)$ introduced in Section 2.2 are all independent r.v.'s having a common distribution for all $k \geq 1$ when $r(\mathbf{i}) = l$ and $r(\mathbf{j}) = m$ are fixed. Let $\chi(z) = (\chi_{lm}(z))_{l, m \in \mathbb{E}}$ be the family of p.g.f.'s

$$\chi_{lm}(z) = \mathbf{E}[z^{w(\mathbf{i}, \mathbf{j}, 1)} \mid r(\mathbf{i}) = l, r(\mathbf{j}) = m],$$

and let $\varepsilon = (\varepsilon_{lm})_{l, m \in \mathbb{E}}$ be the family of the first moments

$$\varepsilon_{lm} = \mathbf{E}[w(\mathbf{i}, \mathbf{j}, 1) \mid r(\mathbf{i}) = l, r(\mathbf{j}) = m].$$

Here is a simple consequence of Proposition 3.6 concerning the additive and the multiplicative tree costs $W_\Sigma(\mathbf{i})$ and $W_\Pi(\mathbf{i})$ given, respectively, by (2.5) and (2.6).

Corollary 3.7 *Assume that the conditions of Proposition 3.6 are satisfied. Then*

$$\mathbf{E}[z^{W_\Sigma(\mathbf{i})} \mid r(\mathbf{i}) = k] = \tilde{\phi}_k(1, \chi(z)) \quad (3.8)$$

$$\mathbf{E}[W_\Pi(\mathbf{i}) \mid r(\mathbf{i}) = k] = \tilde{\phi}_k(1, \varepsilon), \quad (3.9)$$

where $\tilde{\phi}_k$ is the solution of (3.7).

In particular, when $w(\mathbf{i}, \mathbf{j}, k)$ is defined as in (2.7), $\varepsilon_{lm} = \varepsilon_{lm}(t)$ is the probability distribution function (p.d.f.) of the delay $d(\mathbf{i}, \mathbf{j}, 1)$ and thus $\tilde{\phi}_k(1, \varepsilon(t))$ is the p.d.f. of the maximal delay in the tree.

3.3 First moments of the flow volumes

In this section we provide explicit expressions and upper bounds for the first moments of the multicast metrics $K(\mathbf{i})$ and $L(\mathbf{i})$. Denote by $K = (K_k)_{k \in \mathbb{E}}$ the family of conditional expectations $K_k = \mathbf{E}[K(\mathbf{i}) \mid r(\mathbf{i}) = k]$; similar notation will be used for the expectations of $K^{(n)}(\mathbf{i})$, $L(\mathbf{i})$, $L^{(n)}(\mathbf{i})$, $J(\mathbf{i})$, and $J^{(n)}(\mathbf{i})$. Introduce also the diagonal matrices $P_{ij} = \text{diag}\{p_{ij}^1, p_{ij}^2, \dots, p_{ij}^\ell\}$ and the vector $M = (M_k)_{k \in \mathbb{E}}$ with $M_k = \sum_{j=0}^{\infty} (j p_{0j}^k + p_{1j}^k)$.

Corollary 3.8 *The functionals $J^{(n)}$, $K^{(n)}$, and $L^{(n)}$, $n = 0, 1, \dots$, satisfy the following recurrent relations*

$$J^{(n+1)} = \psi_k((P_{00} + P_{10})J^{(n)}), \quad (3.10)$$

$$K^{(n+1)} = \Lambda \left[-P_{10}J^{(n)} + \sum_{j=0}^{\infty} P_{0j}K^{(n)} + M \right], \quad (3.11)$$

$$L^{(n+1)} = K^{(n+1)} + \Lambda L^{(n)}. \quad (3.12)$$

Proof: Formula (3.10) can be easily proved by setting the argument $(z_1, z_2) = (0, 1)$ in the recurrent equation (3.3). To obtain (3.11) and (3.12), take the derivatives of both sides of (3.3) with respect to z_1 and z_2 at $(z_1, z_2) = (1, 1)$. By (3.2), the derivatives of $f_m(\phi^{(n)})$ are given by

$$[f_m(\phi^{(n)})]'_{z_1}(1, 1) = -p_{10}^m J_m^{(n)} + \sum_{j=0}^{\infty} p_{0j}^m K_m^{(n)} + M_m, \quad (3.13)$$

$$[f_m(\phi^{(n)})]'_{z_2}(1, 1) = -p_{10}^m J_m^{(n)} + \sum_{j=0}^{\infty} p_{0j}^m K_m^{(n)} + L_m^{(n)} + M_m. \quad (3.14)$$

It follows from the definition that $J^{(n)}$ converges to J . The next corollary provides conditions for the limits of $K^{(n)}$, and $L^{(n)}$ to be finite. ■

Corollary 3.9 *Let $\|\cdot\|_*$ be a norm in \mathbb{R}^ℓ . If $\|\Lambda \sum_{j=0}^{\infty} P_{0j}\|_* < 1$ in the corresponding operator norm, then*

$$\|K\|_* \leq \|\Lambda M\|_* (1 - \|\Lambda \sum_{j=0}^{\infty} P_{0j}\|_*)^{-1}.$$

If $\|\Lambda\|_* < 1$, then

$$\|L\|_* \leq \|K\|_* (1 - \|\Lambda\|_*)^{-1},$$

and hence J , K , and L satisfy (3.10)–(3.12).

Proof: Since all the elements of Λ are non-negative, from (3.11) follows the component-wise inequality

$$K^{(n+1)} \leq \Lambda \left[\sum_{j=0}^{\infty} P_{0j}K^{(n)} + M \right] = \Lambda \sum_{l=0}^n \left(\Lambda \sum_{j=0}^{\infty} P_{0j} \right)^l M.$$

Hence for every n , $\|K^{(n)}\|_*$ is bounded by $\|\Lambda M\|_* (1 - \|\Lambda \sum_{j=0}^{\infty} P_{0j}\|_*)^{-1}$. In a similar manner, the inequalities

$$L^{(n+1)} \leq \sum_{l=0}^n \Lambda^l K^{(n+1-l)} \leq \sum_{l=0}^n \Lambda^l K$$

imply the second bound. ■

Note that $\rho(\Lambda) \leq \|\Lambda\|_*$ for any norm $\|\cdot\|_*$. Suppose that the distribution of the root mark $r(\mathbf{i}_0)$ is such that $\mathbf{P}(r(\mathbf{i}_0) = k) = q_k > 0$ for all $k \in \mathbb{E}$. For $z \in \mathbb{R}^\ell$, define the norm $\|z\|_*$ as $\sum_{k \in \mathbb{E}} |z_k| q_k$. Corollary 3.9 thus provides upper bounds for $\mathbf{E}K(\mathbf{i}_0) = \|K\|_*$ and $\mathbf{E}L(\mathbf{i}_0) = \|L\|_*$. In this case, the norm $\|\cdot\|_*$ of a non-negative matrix $A = (a_{km})_{k,m \in \mathbb{E}}$ equals $\sup_m \sum_{k \in \mathbb{E}} a_{km} q_k / q_m$.

Remark 3.10 Branching processes for which $\rho(\Lambda) > 1$ are usually called super-critical. The size of the tree \mathcal{T}_B generated by such process may be infinite. Suppose that every vertex of \mathcal{T}_B a.s. requests at least one packet. It is easy to see then, that in the unicast mode, when all $p_{1j}^m = 0$, the flow volume $K(\mathbf{i})$ through a vertex \mathbf{i} may also be infinite, whereas in the multicast mode, when the share of multicast vertices is such that $\|\Lambda \sum_{j=0}^{\infty} P_{0j}\|_* < 1$, the number of sent packets $K(\mathbf{i})$ is finite a.s. and in the mean.

3.4 Application to multicast efficiency estimation

We conclude this section with an example of an evaluation the efficiency gain of multicast over unicast in the special case of a Galton–Watson tree matching some real-life multicast networks.

For a fixed tree, denote by N the number of receivers on a tree and let L_M and L_U be the total number of packet transmissions required to reach all receivers on the tree assuming, respectively, pure multicast and unicast transition modes (all multicast and all unicast vertices).

The first relation between these characteristics was provided by Chuang and Sirbu [7]. They introduced a performance metric in the form of the ratio L_M/\bar{L}_U , where $\bar{L}_U = L_U/N$ is the average length of a unicast path to a receiver. Using simulations and real network samples from Mbone and Arpanet, they found out that for moderate N , the empirical power law

$$L_M/\bar{L}_U \propto N^\alpha \quad (3.15)$$

with $\alpha = 0.8$ adequately describes the efficiency gain for a wide range of network topologies. Later on, the asymptotics of the ratio L_M/L_U was explored by Chalmers and Almeroth [6], who used more Mbone data and provided a similar relation

$$L_M/L_U \propto N^{\alpha-1} \quad (3.16)$$

with α ranging between 0.66 and 0.7. They argued that the difference in the scale factors is due to the inclusion of the last hop in their measurements.

Several analytical explanations of this law have been proposed ever since. In [15] the asymptotics of the metric (3.15) were investigated for k -ary trees. It was shown that for such trees, L_M/\bar{L}_U exhibits a near-linear growth rate with respect to N , independently of the value of k . To achieve a better accordance with the real-life observations, in [1] it was suggested to replace k -ary trees with the so-called *self-similar* k -ary trees, for which, as it was shown there, the power law (3.15) holds with $\alpha \simeq 0.88$.

Let us examine the asymptotics of L_M/L_U when the trees are generated by a branching process. Consider two types of vertices representing two types of Internet hosts: the *routers* ($r(\mathbf{i}) = 1$) that can replicate packets, but do not host any receivers, and the *end receivers* ($r(\mathbf{i}) = 2$) that do request packets, but can neither replicate, nor forward them. The mark distribution for the two types is thus defined by the probabilities $p_{10}^1 = 1$, $p_{00}^1 = p_{01}^1 = p_{11}^1 = 0$; and $p_{01}^2 = 1$, $p_{00}^2 = p_{01}^2 = p_{10}^2 = 0$, respectively. The elements λ_{11} and λ_{12} of the matrix $\Lambda = (\lambda_{km})$ correspond, respectively, to the mean number of routers and receivers directly connected to a router, i.e., to the branching degree of the tree. We assume that the end receivers are located at the leaves of the multicast tree and hence $\lambda_{21} = \lambda_{22} = 0$.

It is easy to see that $\lambda_{11} < 1$ is a necessary and sufficient condition for the mean number of vertices $(\lambda_{11} + \lambda_{12})/(1 - \lambda_{11})$ in the tree \mathcal{T}_B to be finite. The mean total number of receivers $N = N(\lambda_{11}, \lambda_{12})$ in the tree \mathcal{T}_B then equals $\lambda_{12}/(1 - \lambda_{11})$.




Figure 1: Cost ratio L_M/L_U against the number of end receivers N in log-log scale. The straight lines are the results of the log-linear fitting yielding (3.16) with $\alpha = 0.77$ and 0.76 .

Since we deal with random trees, we consider mean quantities in (3.15) and (3.16). In the notation introduced in Section 3.3, the cost of multicast $L_M = L_M(\lambda_{11}, \lambda_{12})$ is expressed as $L_1 = \mathbf{E}[L(\mathbf{i}) \mid r(\mathbf{i}) = 1]$. By (3.10)–(3.12), $L_1 = (\lambda_{11}(1 - J_1) + \lambda_{12})/(1 - \lambda_{11})$, where J_1 is the solution of the equation $J_1 = \exp(\lambda_{11}(J_1 - 1) - \lambda_{12})$ in the domain $|J_1| \leq 1$.

To assess the unicast cost, consider the same distribution tree in which routers cannot replicate packets; i.e., put $p_{10}^1 = 0$, $p_{00}^1 = 1$, other conditions being unchanged. In such settings, L_1 is the mean of the summary length of all unicast paths L_U . From (3.10)–(3.12), it follows that $L_1 = \lambda_{12}/(1 - \lambda_{11})^2$. The value $L_U/N = 1/(1 - \lambda_{11})$ can be considered as an approximation to the mean length of a unicast path \bar{L}_U .

Figure 1 shows two log-log plots of the cost ratio L_M/L_U against the mean number of subscribers N with $\lambda_{11} = 0.9$ and 0.95 , respectively, and λ_{12} ranging in the interval $[1, 5.5]$. The two plots correspond thus to average unicast path lengths of 10 and 20. Characteristics of real-life multicast trees observed in the Mbone network fall within this set of parameters: in the data set used in [6], the lengths of the unicast paths do not exceed 24 hops, while the mean number of outgoing links from a router to receivers ranges between 1.28 and 5.21. The straight lines in the plots are the results of log-linear fitting, which yields (3.16) with $\alpha = 0.77$ and 0.76 . Thus, in the given range of parameters, we observe good agreement with the asymptotics found in [7] and in [6] (note that we do count the last hop to the receiver). For higher values of the branching degree of the multicast tree λ_{12} , the relation L_M/L_U tends to the constant $1 - \lambda_{11}$.

Assume that every receiver \mathbf{i} requests a random number of packets $\tau(\mathbf{i})$. In this case, λ_{12} changes for $\lambda_{12}\mathbf{E}\tau(\mathbf{i})$ in the expressions for L_M , L_U , and N . Indeed, every packet request can be considered as a separate receiver.

In [7], the last hop between a router and a receiver is excluded from the cost of the multicast tree. As shown above, the number of final hops to receivers equals $\lambda_{12}/(1 - \lambda_{11})$. Adjusted by this value, the cost functionals become $L'_M = \lambda_{11}(1 - J_1)/(1 - \lambda_{11})$ and $L'_U = \lambda_{11}\lambda_{12}/(1 - \lambda_{11})^2$, from where $L'_M/L'_U = (1 - J_1)/N$.

4 Flow metrics in the Poisson point process case

In this section we focus on trees generated by the aggregates of a Poisson point process introduced in Section 2.4. As already noticed, such trees have dependent branches, which makes the analysis of their distributions difficult. Example 2.2 shows that point aggregates have similar structure with Boolean clumps. A usual way of obtaining upper bounds for clump-related characteristics is to consider a stochastically dominating branching process, whose branches are independent. The idea goes back to [10], where it was used to obtain a bound for the clump size of a typical Boolean clump. We adopt this approach to get bounds for cost functions on the multicast trees generated by point aggregates. For this, in Section 4.1 we introduce stochastic ordering of random trees. The basic domination result is proved in Section 4.2. Finally, in Section 4.3, we show how tighter bounds can be obtained in the special case when the interaction regions defining point aggregates are all unit balls.

4.1 Stochastic comparison of random trees

For two random vectors $X = (X_n)_{n \leq N}$ and $Y = (Y_n)_{n \leq N}$ in \mathbb{R}^N , the definition of a stochastic ordering is: $X \preceq_{st} Y$ if $\mathbf{E}f(X) \leq \mathbf{E}f(Y)$ for every component-wise increasing function $f : \mathbb{R}^N \rightarrow \mathbb{R}$. For two marked oriented trees $T_1 = (V_1, E_1)$ and $T_2 = (V_2, E_2)$, with vertices in a common space, define the relation of inclusion $T_1 \subset T_2$ if $V_1 \subset V_2$, $E_1 \subset E_2$, and the marks of common vertices coincide. This relation establishes a partial ordering of marked trees. Denote by \mathcal{I} the class of all real-valued increasing functions on such trees, i.e., of all f for which $T_1 \subset T_2$ implies $f(T_1) \leq f(T_2)$. For two random trees $\mathcal{T}_1, \mathcal{T}_2$, the definition of the stochastic ordering is: $\mathcal{T}_1 \preceq_{st} \mathcal{T}_2$ if $f(\mathcal{T}_1) \leq \mathbf{E}f(\mathcal{T}_2)$ for every function $f \in \mathcal{I}$.

4.2 Upper bounds through stochastic domination

Define a marked tree \mathcal{T}_{B^*} with the vertex set $V_{B^*} \subset \mathbb{R}^d$ as follows. Place the root at the origin and let the distribution of $r(0)$ be given by (2.15). For every $y_i \in V_{B^*}$, let the immediate descendants of y_i of types $m \in \mathbb{E} = \{1, \dots, \ell\}$ form independent homogeneous Poisson point processes $\Pi(y_i, m)$ with intensities λq_m in the domain $y_i + G_{r(y_i)m}$. Let the joint distribution of $(\sigma(\cdot), \tau(\cdot))$ be given by (2.14). Apart from the fact that $V_{B^*} \subset \mathbb{R}^d$, there exists a branching process dual to the tree \mathcal{T}_{B^*} as defined in Section 2.3. Indeed, the progeny $Z(y_i) = (Z_m(y_i))_{m \in \mathbb{E}}$ of a vertex y_i such that $r(y_i) = k$ is a vector of independent r.v.'s, where the component $Z_m(y_i)$ has Poisson distribution with parameter $\lambda_{km} = \lambda q_m \text{vol}(G_{km})$. The p.g.f.'s (2.13) of $Z(y_i)$ thus have the form

$$\psi_k(z) = \exp \left[\sum_{m \in \mathbb{E}} \lambda q_m \text{vol}(G_{km}) (z_m - 1) \right], \quad k \in \mathbb{E} \quad (4.1)$$

and $\Lambda = (\lambda_{km})_{k,m \in \mathbb{E}}$ is the matrix of the first moments.

The following lemma is the main result that we will use for comparing functionals on trees generated by branching processes and by point aggregates.

Lemma 4.1 *Suppose that \mathcal{T}_C is a marked tree generated by a typical aggregate A . Then \mathcal{T}_C and \mathcal{T}_{B^*} can be constructed on the same probability space so that $\mathcal{T}_C \subset \mathcal{T}_{B^*}$ a.s., and hence $\mathcal{T}_C \preceq_{st} \mathcal{T}_{B^*}$.*

Proof: Let the tree \mathcal{T}_{B^*} be constructed. The construction of $\mathcal{T}_C = (V_C, E_C)$ will proceed by induction on tree truncations. The marks $(\sigma(\cdot), \tau(\cdot))$ can obviously be chosen coinciding for common vertices. For $n = 0$, the construction is trivial: take $V_C^{(0)} = \{0\}$, $E_C^{(0)} = \emptyset$, and $r(0)$ the same as for \mathcal{T}_{B^*} . Suppose, for some n , the tree $\mathcal{T}_C^{(n)} = (V_C^{(n)}, E_C^{(n)})$ has been constructed so that $\mathcal{T}_C^{(n)} \subset \mathcal{T}_{B^*}^{(n)}$ a.s. This means that the distributions of $V_C^{(n)}$ and of the n -th truncation of a typical point aggregate $A^{(n)}$ coincide. With a slight abuse of notation, let us write $A^{(n)}$ for $V_C^{(n)}$, and A_n for the set of leaf vertices of $\mathcal{T}_C^{(n)}$.

In order to make the induction step, recall (2.16) and define for every $y_i \in A_n$ the descendants of type l as a subset $\Pi'(y_i, l)$ of $\Pi(y_i, l)$ (which consists, by construction, of the descendants of y_i of type l in \mathcal{T}_{B^*}). Let $\Pi'(y_i, l)$ be obtained by restriction of $\Pi(y_i, l)$ to $\mathbb{R}^d \setminus C_l^{(n)}$ and subsequent random thinning with location-dependent retention probability $p_l^{(n)}(y)$ given by

$$p_l^{(n)}(y) = 1 / \text{card}\{y_j \in A_n : y \in y_j + G_{r(y_j)l}\}.$$

From Remark 2.3, it follows that the distribution of $\bigcup_l \bigcup_{y_i \in A_n} \Pi'(y_i, l)$ is the same as the conditional distribution of A_{n+1} given $(A^{(n)}, A^{(n-1)})$. Put $V_C^{(n+1)} \equiv A^{(n)} \cup A_{n+1}$, the new generation of the vertex set of \mathcal{T}_C is thus constructed. Regarding the edges, it is easy to see that a retained vertex $y_i \in A_{n+1}$ has equal chances to be an immediate descendant of every element of the set $\{y_j \in A_n : y_j \rightsquigarrow y_i\}$, which corresponds to the definition of \mathcal{T}_C in Section 2.4. ■

Corollary 4.2 *Let \mathcal{T}_C be the tree generated by the typical aggregate A , and \mathcal{T}_B , by a branching process with the distribution given by (2.15) for $r(\mathbf{i}_0)$, (4.1), and (2.14). Then the multicast cost functionals on \mathcal{T}_C are stochastically dominated by those on \mathcal{T}_B , namely*

$$(K(0), (L_{kl}(0))_{k,l \in \mathbb{E}}, W_\Sigma(0), W_\Pi(0)) \preceq_{st} (K(\mathbf{i}_0), (L_{kl}(\mathbf{i}_0))_{k,l \in \mathbb{E}}, W_\Sigma(\mathbf{i}_0), W_\Pi(\mathbf{i}_0)).$$

Proof: The multicast cost functionals introduced in Section 2.2 are increasing functions on trees. The statement thus follows from Lemma 4.1 as the tree \mathcal{T}_{B^*} has the required distribution. ■

For a tree $T = (V, E)$ such that $V \subset \mathbb{R}^d$, define the total edge length $\Sigma(T)$ as

$$\Sigma(T) = \sum_{(y_i, y_j) \in E} \|y_i - y_j\|.$$

Corollary 4.3 *Let \mathcal{T}_C and \mathcal{T}_B be the trees defined in Corollary 4.2. Then*

$$\Sigma(\mathcal{T}_C) \preceq_{st} W_\Sigma(\mathbf{i}_0),$$

where $W_\Sigma(\mathbf{i}_0)$ is the additive cost of \mathcal{T}_B with independent random weights $w(\mathbf{j}_1, \mathbf{j}_2, k)$ having the same distribution as $\|X_{r(\mathbf{j}_1)r(\mathbf{j}_2)}\|$, where $X_{kl} \in \mathbb{R}^d$ is uniformly distributed in G_{kl} .

Note that if the conditions of Proposition 3.6 are satisfied for \mathcal{T}_B , the p.g.f. of $W_\Sigma(\mathbf{i}_0)$ can be derived from Corollary 3.7 using (3.8) and the total probability formula with respect to the distribution (2.15) of $r(\mathbf{i}_0)$. *Proof:* Let \mathcal{T}_C and the corresponding dominating tree \mathcal{T}_{B^*} be defined in the same probability space as in the proof of Lemma 4.1. Then $\Sigma(\mathcal{T}_C) \leq \Sigma(\mathcal{T}_{B^*})$ a.s. Note that if $r(y_i) = k$, by the properties of Poisson processes, the distribution of $\Pi(y_i, l)$

coincides with that of a Poisson number $Z_l(y_i)$ of independently sampled points ξ_{kl} . Since $\sigma(\cdot) \equiv 1$ and $\tau(\cdot) \equiv 1$, for every $x_i \in E_{B^*}$, we have $\overline{K}(x_i) = 1$. A single mark $w(x_i, x_j, 1)$ is thus associated with every edge $(x_i, x_j) \in E_{B^*}$. Therefore, the edge lengths $\|x_i - x_j\|$ can be viewed as independent edge marks $w(x_i, x_j, 1)$. The result follows from Corollary 4.2. \blacksquare

The next example illustrates the special case when point aggregates coincide with subsets of Π covered by Boolean clumps (see Example 2.2) and some cost functionals have known distributions.

Example 4.4 In the settings of Section 2.4, assume that $r(x_i) \equiv 1$, $\sigma(x_i) \equiv 0$, and $\tau(x_i) \equiv 1$ for all $x_i \in \Pi$. If we put $G_{11} = b_2(0)$, then $x_i \rightsquigarrow x_j$ if and only if $b_1(x_i) \cap b_1(x_j)$ is non-empty. For such choice of parameters, the functional $K(0) + 1$ calculated for \mathcal{T}_C equals the size of a typical clump in the classical Boolean model (see Example 2.2) with $B_i = b_1(0)$. By Corollary 4.2, $K(0) \preceq_{st} K(\mathbf{i}_0)$, where $K(\mathbf{i}_0)$ is calculated on a tree produced by a unitype branching process having Poisson progeny distribution with parameter $\lambda' = \lambda \text{vol}(b_2(0)) = \lambda 2^d \pi^{d/2} / \Gamma(1 + d/2)$. Note that $K(\mathbf{i}_0) + 1$ equals the total population size of this process. If $\lambda' \leq 1$, then $K(\mathbf{i}_0) + 1$ is a.s. finite and has the Borel–Tanner distribution (see, e.g., Example 1, p. 188 in [16]) given by

$$\mathbf{P}(K(\mathbf{i}_0) + 1 = n) = e^{-\lambda'n} (\lambda'n)^{n-1} / n!, \quad n = 1, 2, \dots$$

4.3 Improvement of bounds for trees generated by unitype aggregates

Call a point aggregate *unitype* if in the definition of Section 2.4 all sets G_{kl} are unit balls $b_1(0)$. We derive here sharper bounds for multicast metrics on trees \mathcal{T}_C generated by unitype aggregates.

The relation $x_i \rightsquigarrow x_j$ in this case means simply $\|x_i - x_j\| \leq 1$, therefore, the definitions of the unitype aggregate A and of the corresponding tree \mathcal{T}_C do not depend on host types $r(\cdot)$. Acting as in Lemma 4.1, we will construct a dominating tree \mathcal{T}_{B^*} generated by a branching process with type-dependent progeny distribution. However, instead of discrete vertex types, we will consider here $r(\cdot) \in \mathbb{E} = [0, 1]$. From now on, for \mathcal{T}_C , let the mark $r(\cdot)$ be the distance from the vertex x_i to its ancestor $v(x_i)$, i.e., $r(x_i) = \|x_i - v(x_i)\|$. Such definition allows us to construct a tighter dominating branching process.

Define a branching process rooted at \mathbf{i}_0 in which every individual \mathbf{i} has $Z(\mathbf{i})$ direct descendants, with i.i.d. marks $r(\cdot) \in [0, 1]$ having the common distribution density

$$\mu(t) = dt^{d-1}, \quad t \in [0, 1]. \quad (4.2)$$

For every vertex \mathbf{i} such that $r(\mathbf{i}) = t$, let the p.g.f. $\psi_t(\cdot)$ of $Z(\mathbf{i})$ be given by

$$\begin{aligned} \psi_t(u) &= \mathbf{E}[u^{Z(\mathbf{i})} | r(\mathbf{i}) = t] = \exp(-\lambda S(t)(1 - z)), \quad \text{where} \\ S(t) &= \begin{cases} \pi^{d/2} / \Gamma(d/2 + 1) & \text{if } \mathbf{i} = \mathbf{i}_0, \\ \pi^{d/2-1} \arccos(-t/2) / \Gamma(d/2 + 1) & \text{else.} \end{cases} \end{aligned} \quad (4.3)$$

This means that if $r(\mathbf{i}) = t$, then $Z(\mathbf{i})$ has Poisson distribution with parameter $\lambda S(t)$. Note that if the parameters λ and d are such that

$$\lambda_0 \equiv \lambda \int_{[0,1]} S(t) \mu(t) dt \leq 1,$$




Figure 1: If $(x_k, x_i) \in E_C$, the descendants of x_i are located in $b_1(x_i) \setminus b_1(x_k)$.

the branching process becomes a.s. extinct. Indeed, the distribution of $Z(\mathbf{i})$ can be viewed as Poisson with *random* parameter $\lambda S(r(\mathbf{i}))$. Then $\lambda_0 = \mathbf{E}(Z(\mathbf{i})) \leq 1$ is a standard finiteness condition.

Lemma 4.5 *Let \mathcal{T}_C and \mathcal{T}_{B^*} be the trees generated, respectively, by a unitype aggregate, and by the branching process defined by (4.2) and (4.3), with vertices in \mathbb{R}^d . Then both trees can be constructed in the same probability space so that $\mathcal{T}_C \subset \mathcal{T}_{B^*}$ a.s., and hence, $\mathcal{T}_C \preceq_{st} \mathcal{T}_{B^*}$.*

Proof: We use a simple observation for a unitype tree \mathcal{T}_C : if $(x_k, x_i) \in E_C$, the descendants $D(x_i)$ of x_i are all located in the domain $b_1(x_i) \setminus b_1(x_k)$, see Figure 1. We thus construct a dominating tree \mathcal{T}_{B^*} embedded in \mathbb{R}^d as in the proof of Lemma 4.1 with the only difference that the descendants of every non-root vertex y_i marked with $r(y_i)$ form a Poisson process not in $b_1(y_i)$, but in the d -dimensional sector $b'_1(y_i)$ (shaded in Figure 1), with the central angle $2\alpha = 2 \arccos(-r(y_i)/2)$. Then $Z(y_i)$ has Poisson distribution with parameter $\lambda S(t)$, and for all $y_j \in D(y_i)$, the distances $\|y_i - y_j\|$ are i.i.d. r.v.'s having common density (4.2). The rest of the proof is similar to that of Lemma 4.1. \blacksquare

The following corollary follows immediately from Lemma 4.5.

Corollary 4.6 *Suppose that \mathcal{T}_C is the tree generated by a typical unitype aggregate A , and \mathcal{T}_B , by a branching process defined by (4.2) and (4.3). Then the multicast cost functionals on \mathcal{T}_C are stochastically dominated by those on \mathcal{T}_B , namely,*

$$(K(0), L(0), W_\Sigma(0), W_\Pi(0)) \preceq_{st} (K(\mathbf{i}_0), L(\mathbf{i}_0), W_\Sigma(\mathbf{i}_0), W_\Pi(\mathbf{i}_0)).$$

In the remaining part of this section, we focus on the distribution of the functionals $K(\mathbf{i})$ and $L(\mathbf{i})$ of the dominating tree \mathcal{T}_B defined by (4.2) and (4.3). Let us preserve the definitions of p.g.f. families $\phi = (\phi_t)_{t \in \mathbb{E}}$ and $\phi^{(n)} = (\phi_t^{(n)})_{t \in \mathbb{E}}$ from Section 3.1 for branching processes with types $t \in \mathbb{E} = [0, 1]$. Define operator $f(\cdot)$ acting on a function $u(z_1, z_2)$ as follows

$$f(u) = (1 - z_1 z_2) u(0, z_2) p_{10} + z_1 z_2 u(1, z_2) \sum_{j=0}^{\infty} p_{1j} + u(z_1 z_2, z_2) \sum_{j=0}^{\infty} (z_1 z_2)^j p_{0j}. \quad (4.4)$$

The next statement is an analog of Proposition 3.1.

Proposition 4.7 *The family of p.g.f. $(\phi_t^{(n)})_{t \in \mathbb{E}}$ satisfies the functional equation*

$$\phi_t^{(n+1)} = \exp \left[\lambda S(t) \left(\int_{[0,1]} f(\phi_s^{(n)}) \mu(s) ds - 1 \right) \right], \quad (4.5)$$

where $S(t)$ is given by (4.3), $f(\cdot)$ by (3.2), and $\mu(\cdot)$ by (4.2).

Proof: The argument is similar to the proof of Proposition 3.1. First, use the recurrent equations (2.11) and (2.12) to represent $K^{(n+1)}(\mathbf{i})$ and $L^{(n+1)}(\mathbf{i})$ given $r(\mathbf{i}) = t$ as sums of independent r.v.'s with the number of terms $Z(\mathbf{i})$ having Poisson distribution with parameter $\lambda S(t)$. Second, note that relation (3.4) with the right-hand side replaced by $f(\phi_m^{(n)}(z_1, z_2))$ holds. Integrating with respect to (4.2), we obtain

$$\mathbf{E} \left[z_1^{\overline{K}^{(n)}(\mathbf{j})} z_2^{\overline{K}^{(n)}(\mathbf{j}) + L^{(n)}(\mathbf{j})} \right] = \int_{[0,1]} f(\phi_s^{(n)}(z_1, z_2)) \mu(s) ds.$$

Combining the last expression with the p.g.f.'s of $Z(\mathbf{i})$, we get (4.5). ■

Recall the notation of Section 3.3. The following proposition provides closed-form expression for the first moments of multicast cost metrics on \mathcal{T}_B .

Proposition 4.8 *If $\lambda_0 \leq 1$, the joint p.g.f. ϕ of $K(\mathbf{i})$ and $L(\mathbf{i})$ satisfies equation (4.5). If, moreover, $\lambda < 1$, the first moments of $J(\mathbf{i})$, $K(\mathbf{i})$ and $L(\mathbf{i})$ have the form $J_t = \exp(-\lambda C_J S(t))$, $K_t = \lambda C_K S(t)$, and $L_t = \lambda C_L S(t)$, respectively, where C_J is the only solution of the equation*

$$C_J = 1 - (p_{00} + p_{10})I(C_J) \quad \text{with} \quad I(x) = \int_{[0,1]} \exp\{-\lambda x S(t)\} \mu(t) dt \quad (4.6)$$

and

$$C_K = \frac{1 + \sum_{j=0}^{\infty} (j p_{0j} + p_{1j}) - p_{10} I(C_J)}{1 - \lambda_0 \sum_{j=0}^{\infty} p_{0j}}, \quad C_L = \frac{1 + \sum_{j=0}^{\infty} (j p_{0j} + p_{1j}) - p_{10} I(C_J)}{(1 - \lambda_0)(1 - \lambda_0 \sum_{j=0}^{\infty} p_{0j})}. \quad (4.7)$$

Proof: As noted above, when $\lambda_0 \leq 1$, the tree \mathcal{T}_B is a.s. finite, hence $K(\mathbf{i})$ and $L(\mathbf{i})$ are proper r.v.'s. The first part of the statement can thus be proved as Proposition 4.7 using conservation laws (2.2), (2.8) and (2.9).

We now use (4.5) for ϕ to prove the second part. We obtain $J_t = \exp(-\lambda C_J S(t))$ by putting $(z_1, z_2) = (0, 1)$ in (4.5). Plugging again such representation into (4.5) yields (4.6). The uniqueness of solution follows from the fact that $1 - (p_{00} + p_{10})I(\cdot)$ is a contraction on $[0, 1]$ when $\lambda_0 < 1$. To obtain the expressions for K_t and L_t , take the derivatives of both sides of (4.5) with respect to z_1 and z_2 at $(z_1, z_2) = (1, 1)$. From (4.4) it follows that $[f(\phi)]'_{z_i}$, $i = 1, 2$, are uniformly integrable for $(z_1, z_2) \in [0, 1]^2$, thus the differentiation can be brought under the integral. Using (3.13) and (3.14), we get closed-form expressions for C_K and C_L from which (4.7) follows. ■

Example 4.9 Since $I(x)$ is a table integral, the constants in (4.6)–(4.7) can be evaluated analytically in all dimensions. For example, in \mathbb{R}^2 , $\lambda_0 = \lambda(\pi/3 + \sqrt{3}/2)$ and

$$I(x) = 1 + (p_{10} + p_{00}) \frac{4 \exp(-\frac{2}{3}\pi \lambda x) + 2\sqrt{3}\lambda x \exp(-\frac{2}{3}\pi \lambda x) - 8 \exp(-\frac{\pi}{2}\lambda x)}{(\lambda x)^2 + 4}.$$




Figure 2: Packet flow volume K_t at the root of \mathcal{T}_C and \mathcal{T}_B ; see Example 4.9.

Suppose that $\sigma(\cdot)$ and $\tau(\cdot)$ are independent, $\mathbf{P}(\tau(\cdot) > 1) = 0$, and let the parameters $\mathbf{P}(\sigma(\cdot) = 1)$ and $\mathbf{P}(\tau(\cdot) = 1)$ (respectively, the share of multicast and receiver vertices in a tree) vary in the interval $[0, 1]$. Figure 2 shows the variation of the mean packet flow volume K_t at the root vertex of \mathcal{T}_C (simulated with $\lambda = 0.5$) and of the bound (the exact value for the corresponding tree \mathcal{T}_B) obtained using Proposition 4.8.

Acknowledgments This work has been partly supported by RNRT project *GEORGES* and by KBN grant 2 P03A 000 00.

References

- [1] C. Adjih, L. Georgiadis, P. Jacquet, and W. Szpankowski. Is the Internet fractal? The multicast power law revisited. In *SIGCOMM'01*, 2001.
- [2] K. C. Almeroth. The evolution of multicast: From the MBone to inter-domain multicast to Internet2 deployment. *IEEE Network Special Issue on Multicasting*, pages 10–21, Jan-Feb 2000.
- [3] K. B. Athreya and P. E. Ney. *Branching processes*. Springer, Berlin, 1972.
- [4] F. Baccelli and B. Błaszczyszyn. On a coverage process ranging from the Boolean model to the Poisson–Voronoi tessellation, with applications to wireless communications. *Adv. Appl. Probab.*, 2000. (to appear).
- [5] B. Błaszczyszyn, C. Rau, and V. Schmidt. Bounds for clump size characteristics in the Boolean model. *Adv. Appl. Probab.*, 31, 1999.
- [6] R. C. Chalmers and K. C. Almeroth. Modeling the branching characteristics and efficiency gains in global multicast trees. In *IEEE INFOCOM*, 2001.

-
- [7] J. Chuang and M. Sirbu. Pricing multicast communications: A cost-based approach. In *INET'98*, 1998.
 - [8] P. Gupta and P. R. Kumar. Critical power for asymptotic connectivity in wireless networks. In W.M. McEneaney, G. Yin, and Q. Zhang, editors, *Stochastic Analysis, Control, Optimization and Applications*, pages 547–566. Birkhauser, Boston, 1998.
 - [9] P. Gupta and P. R. Kumar. The capacity of wireless networks. *IEEE Transactions on Information Theory*, 46(2):388–404, 2000.
 - [10] P. Hall. On continuum percolation. *Ann. Probab.*, 13:1250–1266, 1985.
 - [11] P. Hall. *Introduction to the Theory of Coverage Processes*. J. Wiley & Sons, New York, 1988.
 - [12] T. E. Harris. *The Theory of Branching Processes*. Springer-Verlag, Berlin-Göttingen-Heidelberg, 1963.
 - [13] S.-J. Lee, W. Su, J. Hsu, M. Gerla, and R. Bagrodia. A performance comparison study of ad hoc wireless multicast protocols. In *Proceedings of IEEE INFOCOM 2000*, Tel Aviv, Israel, Mar. 2000., mar 2000.
 - [14] R. Meester and R. Roy. *Continuum Percolation*. Cambridge University Press, 1996.
 - [15] G. Phillips, S. Shenker, and H. Tangmunarunkit. Scaling of multicast trees: Comments on the Chuang-Sirbu scaling law. In *SIGCOMM'99*, pages 41–51, 1999.
 - [16] B. A. Sevastjanov. *Branching processes*. Nauka, Moscow, 1971. (In Russian).
 - [17] D. Stoyan, W. Kendall, and J. Mecke. *Stochastic Geometry and its Applications*. Wiley, Chichester, 1995.



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