On the Topology of Multicast Trees

Robert C. Chalmers, *Student Member, IEEE,* and Kevin C. Almeroth, *Member, IEEE*

***Abstract*—The benefit derived from using multicast is seemingly dependent upon the shape of the distribution tree. In this paper, we attempt to accurately model interdomain multicast trees. We measure a number of key parameters, such as depth, degree frequency, and average degree, for a number of real and synthetic data sets. We find that interdomain multicast trees actually do share a common shape at both the router and autonomous system levels. Furthermore, we develop a characterization of multicast efficiency which reveals that group sizes as small as 20 to 40 receivers offer a 55%–70% reduction in the total number of links traversed when compared to separately delivered unicast streams. A final contribution of our work consists of a number of data sets, compiled from multicast group membership and path data, that can be used to generate large sample trees, representative of the current multicast infrastructure.**

***Index Terms*—Efficiency, modeling, multicast, topology.**

I. Introduction

MULTICAST provides a conceptual advantage for applications with a large, distributed set of clients. With logical addressing, multicast offers a mechanism to loosely assemble groups with dynamic membership, freeing the application from the task of transmitting to each individual receiver. More importantly, however, multicast provides an operational advantage for content and network providers by reducing the overall resource demands of the application. Only a single multicast packet is sent to the group regardless of how many receivers have joined. Multicast reduces the overall bandwidth demand of content transmission since packet duplication only occurs when paths to multiple receivers diverge. The problem, however, is that the actual benefit derived by any application is seemingly dependent upon the shape of the distribution tree, and a tree's shape changes over time with the arrival and departure of receivers.

This dependency is not a primary concern for most application developers since they are, in large part, interested in leveraging the group model of multicast rather than guaranteeing some level of efficiency. For network providers, though, implementing multicast comes with certain costs. Network providers are faced with a decision which depends on a so-called *sweetspot* where the bandwidth cost of using unicast out-weighs the increased overhead of deploying multicast [1]. In other words, multicast's efficiency must overcome the additional state requirements placed on routers, as well as, the possibly more important costs of installing and managing multicast.

From a research perspective, efficiency is only one part of a broader concern: to understand the shape, or topology, of multicast trees in real networks, particularly the Internet. When designing protocols to provide reliability, security, and congestion control for multicast applications, it is often necessary to make assumptions about the topology [2], [3]. For example, where should one place retransmission or key distribution agents within the tree to maximize their utility? Certain protocols may actually depend upon particular characteristics of the topology, such as maximum fan-out, in order to keep overhead manageable. Furthermore, when evaluating and simulating these protocols, it is critical that sample topologies are realistic if the conclusions should be considered valid [4].

Prior to this study, little quantitative work had been presented which provides an accurate picture of real multicast topologies. In this paper, we identify a number of key properties which we believe sufficiently describe the shape of interdomain multicast trees. Our analysis is conducted at both the router and autonomous system (AS) levels. We also consider multicast efficiency; we present an estimate to model the bandwidth gains of multicast with respect to separate unicast streams. Our efficiency estimate is defined as a simple function of the number of receivers in the multicast group.

First, we define a metric with which to compare the bandwidth utilization of multicast and unicast in terms of the total number of links traversed; i.e., a count of duplicate packets across all links of the distribution tree. Next, we provide an estimate for our metric that accurately characterizes multicast efficiency over a range of tree topologies and group dynamics. The estimate is validated by calculating the metric for a set of real groups collected from the Internet. We then loosen *temporal* and *spatial* constraints to investigate the estimate's dependence upon receiver duration, inter-arrival time and receiver distribution. Consistently, this characterization shows that, for even a small number of receivers, multicast out-performs unicast in terms of bandwidth utilization. Group sizes as small as 20 to 40 receivers offer a 55%–70% reduction in the number of links traversed when compared to separately delivered unicast streams.

With respect to tree topologies, we focus primarily on how and where branching occurs within the tree, and conclude that multicast trees do share several common properties such as average degree and degree frequency. When studying the fan-out of a node, we distinguish between links that connect a router to a receiver from those between neighboring routers. We find that the latter measure of *interior* degree provides a much clearer picture of the shape of a multicast tree. Although theoretically

Manuscript received February 19, 2002; revised May 5, 2002; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Jamin.

The authors are with the Department of Computer Science, University of California, Santa Barbara, CA 93106-5110 USA (e-mail: robertc@cs.ucsb.edu; almeroth@cs.ucsb.edu).

Digital Object Identifier 10.1109/TNET.2002.804835

a multicast tree could take on any shape, in real networks, a tree's shape is constrained by the underlying network connectivity. The result is a fairly consistent model of the topology of interdomain multicast trees.

The paper is organized as follows. In the next section, we present related work to prepare an appropriate foundation for our analysis. Then, in Section III, we describe our experimental data and methodology. We define the efficiency characterization in Section IV and discuss the results of measuring efficiency for each of our data sets. We look more closely at the shape of multicast trees in Section V. In Section VI, we outline future work, and we present our conclusions in Section VII.

II. Related Work

To provide sufficient background for the remainder of the paper, we first present an overview of previous work relevant along four lines:

- understanding the evolution of multicast deployment;
- measuring multicast characteristics, e.g., bandwidth utilization, loss, delay and traffic concentration;
- analyzing the topology of the Internet (unicast) and the shape of multicast trees; and
- comparing multicast to unicast in terms of cost and efficiency.

A. Multicast Deployment

In order to place this work in context, it is important to understand the evolution of multicast deployment in the Internet. From the first large-scale experiments in 1992 until 1997, deployment consisted of a flat, overlay network referred to as the multicast backbone (MBone) [5]. The MBone connected a number of small multicast-capable networks through unicast-encapsulated tunnels. Each tunnel, possibly consisting of multiple unicast routers, acted as a virtual link in the construction of multicast trees. This flat routing topology was inefficient and became difficult to manage as multicast deployment increased.

In 1997, work began to develop a hierarchical multicast infrastructure. Analogous to interdomain unicast routing, AS's are allowed to deploy separate intra-domain multicast routing protocols. Each AS exchanges with its peers information concerning the reachability and activity of multicast sources. Based on these exchanges, a global, interdomain distribution tree is constructed for a multicast group, connecting the individual intra-domain trees.

By mid-1999, the two Internet2 backbone networks, vBNS and Abilene, had deployed interdomain multicast with peering points across the US. Around this time, the existing MBone and its collection of tunnels were relegated to a special AS, AS10888. Since then, the size of the old MBone has diminished significantly [6]. Although native multicast support in commercial backbones has been slow to evolve, recent developments in Sprint's backbone and several other major ISP networks have followed the deployment model of Internet2.

B. Multicast Characteristics

There are many interesting characteristics of multicast trees and protocols. A number of studies have investigated multicast in terms of bandwidth utilization, loss, delay and traffic concentration [7]–[11]. The authors also address the impact of routing protocol overhead such as dense-mode flooding, and illustrate the differences between the construction of shortest-path and core-based trees. In this paper, we choose not to re-address these issues. Rather, we focus on the general shape of multicast trees, and how that shape impacts the efficiency of packet distribution.

C. Topology and Shape

Recent studies have attempted to model important properties of network topologies with distributions that obey power-laws [12]–[15]. Power-laws are functions of the form, $y \propto x^b$, where b is a constant, and x and y represent the variables of interest. These relationships describe heavily skewed distributions.

Faloutsos, Faloutsos, and Faloutsos propose a number of power-law relationships describing several properties of the Internet at both the router and AS levels. In particular, one formula relates out-degree and frequency, $f_d \propto d^{\mathcal{O}}$ ($\mathcal{O}$ is negative) [12]. The function maintains that the majority of nodes, both AS's and routers, have few out-going links, and that only a very small number of nodes have high degrees. Recently, Chen *et al.* have disputed the accuracy of this model when considering AS-level topologies [16]; power-laws are not appropriate representations at this level, although heavily skewed distributions are still evident.

Looking strictly at routers, Pansiot and Grad construct a router-level graph consisting of nearly 4000 nodes and 5000 edges collected from Internet traces [17]. The authors note that the majority of nodes, more than 70%, have a degree of only one or two. Applying this observation to multicast, one may infer that trees built over such an underlying network infrastructure would also consist of many nodes with small degrees since branching is inherently limited by the out-degree of each router.

D. Multicast versus Unicast

Chuang and Sirbu provide the first definitive comparison of multicast and unicast while defining a measure for pricing multicast traffic [18]. They focus on the ratio between the total number of multicast links in the distribution tree and the average unicast path length in the network. The authors determined that the relationship can be concisely expressed as a power-law in terms of the number of receivers. They produced the formula

$$\frac{L_m}{\overline{L_u}} = N^k \tag{1}$$

where L_m is the total number of multicast links; $\overline{L_u}$ is the average path length between any two nodes in the network; N is the number of receivers in the tree; and k is an economies-of-scale factor, ranging between 0 and 1, which expresses the *affinity* of a new receiver to share some portion of its path with the existing tree. Using samples of the old ARPANET, the early MBone, and a number of generated topologies, they concluded that the value of k was consistently near 0.8.

This provides a simple, yet nonintuitive, relationship. Could cost be dependent solely upon the number of receivers? Should it matter where in the tree the receivers are located? Is not the

TABLE I
REAL GROUP DATA SETS USED IN EVALUATION

Name	*Description*	*Trace Period*	*Receivers*	
			Total	*Traced*
IETF43-A	43rd IETF Audio	Dec. 7-11, 1998	257	107
IETF43-V	43rd IETF Video	Dec. 7-11, 1998	305	129
NASA-A	NASA Shuttle Launch Audio	Feb. 14-21, 1999	144	43
NASA-V	NASA Shuttle Launch Video	Feb. 14-21, 1999	209	58

TABLE II
SYNTHETIC GROUP DATA SETS USED IN EVALUATION

Name	*Source*	*Trace Period*	*Receivers*
SYNTH-1	UC Santa Barbara	Jan. 6-10, 2000	1,871
SYNTH-2	Georgia Tech	Jul. 12-25, 2001	1,497
SYNTH-3	University of Oregon	Dec. 18-19, 2001	1,019
SYNTH-4	UC Santa Barbara	Dec. 19-22, 2001	1,018

shape of the tree just as important as how many receivers it serves? Either the shape of most multicast trees is, in some way, constrained, or a wide range of shapes exhibit similar costs.

This cost relationship (1) was later confirmed with a more rigorous mathematical treatment [19]. Although the authors produced a cost function that was logarithmic rather than a power-law, it was found to behave very similarly over a wide range of generated topologies. An ensuing study specifically considered multicast efficiency [20]. The authors determined that a power-law is indeed applicable when the size of the network remains below 10^6 nodes, and the number of group members constitutes only a small fraction of that size. Moreover, the economies-of-scale factor, k, was proven to be a function of the network size, rather than a simple constant.

All three studies considered mostly generated networks and uniformly distributed receivership [18]–[20]. None investigated how well this relationship holds in the current Internet for real groups.

III. EXPERIMENTAL METHODOLOGY

In this section, we describe the data sets collected from the Internet that we use as a basis for our initial analysis. We also explain how we generated synthetic data sets in order to explore random receiver distributions. Finally, we detail how each data set is processed to derive our final results.

A. Real Data Sets

We have collected topologies from four live multicast sessions. Using the MHealth tool [21], multicast tree data were recorded during the 43rd meeting of the IETF and the NASA shuttle launch in February of 1999. Both events consisted of separate audio and video channels generated by a single source. The logs of these four sessions form the source of our initial data sets: IETF43-A, IETF43-V, NASA-A, and NASA-V (Table I).

There are some problems with these data sets worth mentioning. MHealth joins a particular multicast group and collects real-time control protocol (RTCP) [22] packets. These packets are used to identify the source(s) and receivers of the group, and to subsequently track each member's activity. The tool then uses the *mtrace* (multicast traceroute) utility [23] to trace the path from each receiver back to the source. Because MHealth traces each receiver sequentially, it may miss receivers who are part of the group for only a short duration.

Thus, only a portion of the total number of receivers in the group were successfully traced, 43% for IETF43 and 29% for NASA. Furthermore, not all receivers are necessarily known since some decoding tools do not implement RTCP. Also, some RTCP packets may be filtered by firewalls. It is possible that the failed or missing traces are obscuring irregularities in the tree structure, or that entire subnets hidden behind firewalls have extremely different properties than what is observed in the rest of the tree. However, we feel that the results of the experiments presented in the following sections are definitive enough to consider these possibilities as remote.

B. Synthetic Data Sets

In addition to studying real multicast groups, we wanted to investigate how efficiency was affected by generalizing group size, member composition and receiver distribution. We attempted to trace multicast paths using a number of sources for a series of 22 000 IP addresses that were known to have participated in multicast groups over a two year period, June 1997–June 1999 (Table II). The IP addresses were collected by the *mlisten* tool [9] which joins multicast groups advertised over the MBone's session directory, *sdr*, and collects RTCP packets from group members.

While most of the traces were completed recently and reflect the latest multicast infrastructure, the group members represent a relatively random sample taken from the older MBone. This dichotomy partially accounts for the rather low success rate in tracing receivers; less than two thousand of the twenty-two thousand IP addresses were reachable via multicast. Significant changes in the multicast routing infrastructure occurred during the period (1997–1999) in which the candidate addresses were collected. Many networks that were once reachable via the MBone had not yet converted to native, interdomain multicast routing [6]. Therefore, the population of traced receivers that composes our synthetic data sets represent rather long-lived multicast participants. The effect that this may have on the results of our analysis are unknown.

C. Data Processing

For each of the data sets, the RTCP and mtrace data are used to re-construct multicast trees as accurately as possible. However, since the multicast distribution tree changes as receivers join and leave the group, there is no single tree that represents the entire session. Our solution was to develop a new tool, *mwalk*, that builds an **activity graph** of all possible trees over time. Ultimately, these trees are traversed, or **walked**, in order to measure properties such as average degree.

The RTCP data from the MHealth logs are used to build an **activity table** for each receiver which lists the intervals over which the receiver is known to have been a group member. Explicit RTCP *BYE* packets are considered as the end of an active period. Otherwise, a configurable timeout limits active periods. From the mtraces, paths are built from each receiver back to the source, adding intermediate routers to the *mwalk* graph. The timestamp of each trace determines which of a number of possible paths are active at any given time during the life of the session. The result is a graph in which each edge is associated with its own activity table. For each node, only one parent link is active at any moment in time, providing a snapshot of the multicast tree.

The entire session time is partitioned into 10 000 sample periods.[1] The tree is walked to collect pertinent statistics such as the total number of multicast and unicast links,[2] or the out-degree of each node. To capture real group dynamics, we use the link and receiver activity tables to determine the *active* tree for each period. When randomizing receiver activity, however, we establish a single **primary tree** by selecting the most frequently used path between any two nodes. Then, ten sets of n receivers are chosen from a uniform distribution for each possible value of n, $1 <= n <= N$, where N is the total number of distinct receivers throughout the session.

To associate individual nodes with their AS's, we utilize a combination of tools from CAIDA [24] and the University of Michigan [25] to process BGP snapshots taken from the University of Oregon's *Route Views* project [26]. The AS mapping is not perfect. A small number of nodes from each data set were left unmapped, and a few resulted in multiple matching AS's. For the most part, though, these problematic nodes are within the same address range and contiguous within the tree. So, we treat each group as a single, unique AS.

IV. Measuring Multicast Efficiency

The goal of this section is to characterize the bandwidth efficiency of interdomain multicast in the current Internet. To this end, we define a simple metric to measure the reduced packet duplication achieved by multicast. Drawing from prior work [18], we then present an estimate for this metric based solely on the number of group receivers. Finally, we use our proposed metric to evaluate real and synthetic multicast trees to determine whether the estimate is indeed valid.

[1]The resulting granularity is from 30 to 60 s. Shorter periods produce equivalent results.

[2]Since the actual unicast paths are unknown, we make a simplifying assumption that the length of the unicast and multicast paths are equivalent. We address this issue in Section IV-D.

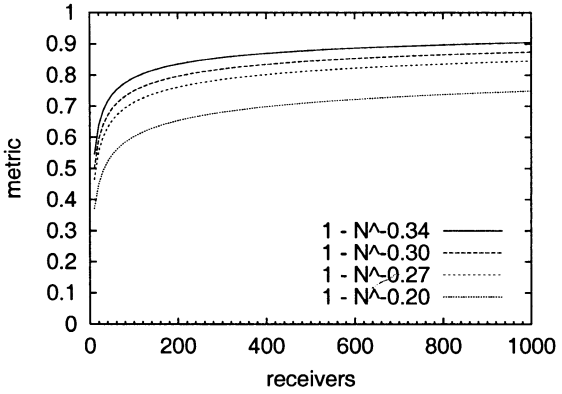

Fig. 1. Efficiency estimate shown over 1000 receivers for a range of efficiency factors.

A. Defining a Metric

In order to effectively evaluate the behavior of multicast efficiency in real networks, we must first define a metric to use as a basis for comparison. In previous work, we introduced a metric for efficiency using multicast and unicast hop counts [27]. The metric is defined as

$$\delta = 1 - \frac{L_m}{L_u} \tag{2}$$

where L_m is the total number of multicast links in the distribution tree and L_u is the sum of all unicast hops. δ represents the percentage gain in multicast efficiency over unicast. As δ approaches zero, multicast and unicast are nearly equal, and little to no savings in bandwidth is achieved. As δ approaches one, all receivers share a single multicast path, resulting in the maximum possible bandwidth efficiency.

B. Defining the Estimate

With a metric defined, we now develop an estimate for multicast efficiency in terms of the number of receivers, N, and an associated *efficiency factor*, ε. First, we express the cost function (1), defined by Chuang and Sirbu, with respect to L_u rather than $\overline{L_u}$

$$\frac{L_m}{\overline{L_u}} = N^k$$
$$\frac{L_m}{L_u} = s_\varepsilon N^{k-1}, \qquad \text{where } \overline{L_u} = s_\varepsilon \frac{L_u}{N}. \tag{3}$$

We have introduced a scaling factor, s_ε, since the relationship between L_u and $\overline{L_u}$ is not exact. Chuang and Sirbu considered $\overline{L_u}$ as the expected distance between any two nodes in the network [18] while we treat it as the average path from source to receiver. Next, we substitute the ratio of multicast and unicast hops (3) into the multicast metric (4)

$$\delta = 1 - \frac{L_m}{L_u}$$
$$= 1 - s_\varepsilon N^\varepsilon, \qquad \text{where } \varepsilon = k - 1. \tag{4}$$

For an economies-of-scale factor (k) of 0.8, the efficiency factor, $(\varepsilon) = -0.2$. As will be seen in the following analysis, however, ε tends to range between -0.38 and -0.27 for real and synthesized receiver distributions. Fig. 1 shows how the estimate behaves as the number of receivers increases. The shape of

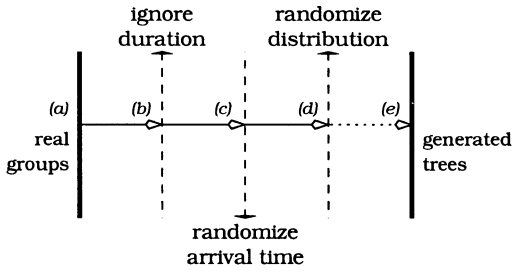

Fig. 2. An illustration of the path followed by the analysis in this section: a transition from real groups to generated trees.

the curve implies that multicast outperforms unicast with even just a few receivers; between 20 and 40 receivers, multicast provides 55%–70% increased efficiency over unicast. The estimate predicts greater than 75% savings for 150 users, reaching as much as 90% for large groups with 1000 receivers.

The estimate provides us with a characterization of multicast efficiency as a function of the number of members in the multicast group. This characterization seemingly ignores the influence of receiver dynamics, such as duration and distribution. Having derived the estimate from previous work, it is expected to behave well with generated trees and random distributions [18].

In the next section, we attempt to validate this characterization for real multicast trees by measuring the efficiency metric for sample multicast groups. We then further extend our confidence in the estimate by evaluating synthesized group distributions.

C. Analysis

Using our metric for multicast efficiency, we must now determine how the estimate reacts to real group dynamics. We consider the efficiency characterization valid if the measured values of L_m (multicast links), L_u (unicast links), and N (number of receivers) for real multicast groups fit the power-law in equation (4) for some values of the efficiency factor, ε. Ultimately, to be useful, ε should be similar across the entire sample range.

Using the values collected from the *mwalk* trees, we analyze how temporal and spatial constraints affect the estimate, $\delta = 1 - s_\varepsilon N^\varepsilon$ (4). We specifically look at four key variables:

1) receiver duration;
2) receiver arrival time;
3) receiver distribution;
4) number of receivers.

Starting from the most restrictive environment, we incrementally loosen the constraints placed on each tree instance. At each stage, we test whether the efficiency metric still conforms to the estimate, determining whether a dependence exists upon the variable of interest.

Fig. 2 illustrates the path of our analysis. We begin (step a) with our real data sets (Table I). When evaluating the efficiency metric, we initially take into account the actual receiver distribution, as well as the join time and duration of each member. Then, we stretch the duration (step b) until the receiver eventually remains joined from his/her initial arrival through the end of the session. Next, the receiver activity is randomized (step c), thus ignoring the time domain completely. Finally, we stretch the spatial domain (step d) by randomizing the receiver distribution with synthesized data sets.

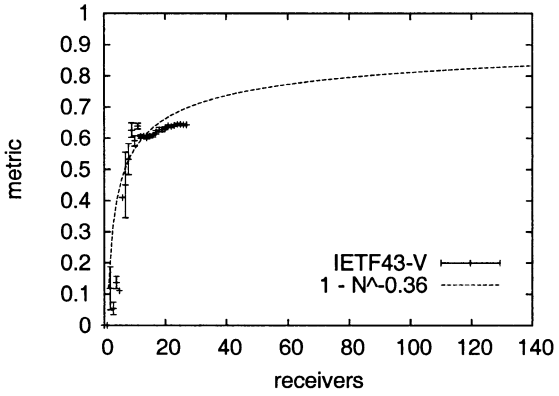

Fig. 3. (step a) Average metric calculated using a 1 minute time-out in the IETF43-V data set (99% confidence intervals).

Throughout our analysis, tree topologies are derived from the actual multicast infrastructure of the Internet. The final step (e) considers only randomly generated topologies, but this was the focus of the initial work used to derive the estimate (see Section IV-A). Hence, our study bridges the gap between real and generated trees.

1) Receiver Duration: Our initial objective, then, is to validate the efficiency estimate in the most restrictive case (step a). With this in mind, we collect L_m, L_u, and N using the actual inter-arrival pattern of the receivers, and a small 1-min timeout to minimize any effects of artificially extending receiver durations. Using the IETF43 and NASA data sets, we plot, for each receiver count, the metric averaged over 10 000 sampled periods. Fig. 3 shows results for the IETF43-V data set. Though not shown, the other data sets exhibit similar behavior. As can be seen from the graph, the data roughly fit the characterization with an efficiency factor, ε, of -0.36 (error bars indicate 99% confidence intervals).[3]

The resulting graph is not smooth, however, due to the rather low number of simultaneous receivers present in any one sample period (about 25). This leads to an interesting observation about Fig. 3. The higher-than-estimated efficiency gains noticeable around ten receivers is likely due to receiver clustering. Several receivers may join the group from the same or a relatively nearby leaf router. Consequently, they share a large portion of the tree, and artificially inflate the metric. However, as the number of receivers increases beyond fifteen, we can see that the efficiency metric begins to follow the characterization more closely.

The next step [Fig. 2(b)] in our analysis is to lengthen the time-out value to determine whether receiver duration is a contributing factor to the behavior of the metric. This also allows us to incrementally increase the number of active receivers since longer durations mean that more receivers are likely to be joined in any given period.

Lengthening receiver duration to 5 min, and then 20 min results in graphs similar to Fig. 3, which are not included due to space constraints. We ignore the duration completely in Fig. 4 by

[3]In an earlier publication [28], we reported a slightly lower value for ε due to an error in our linear regression; however, the conclusions drawn from the analysis have not changed.

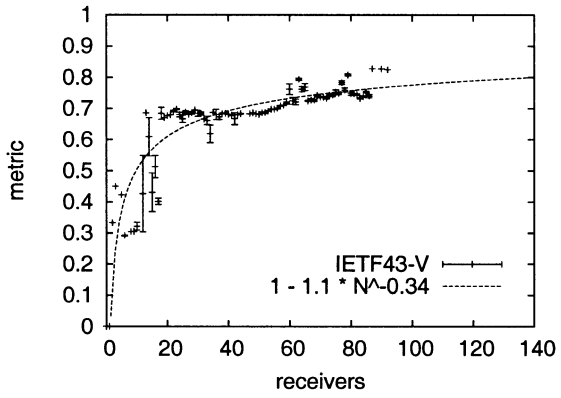

Fig. 4. (step b) Average metric calculated ignoring duration in the IETF43-V data set (99% confidence intervals).

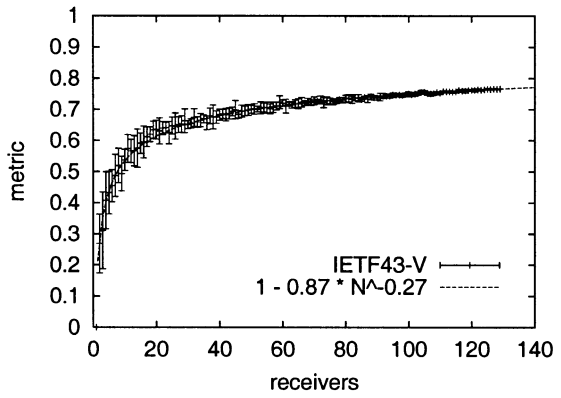

Fig. 5. (step c) Average metric calculated using random receiver activity in the IETF43-V data set (99% confidence intervals).

extending the time-out beyond the length of the actual session; i.e., once receivers join, they remain members for the duration of the session. Each receiver still joins the group according to the original data set. Thus, inter-arrival times are not affected. In all three cases, the metric stays fairly consistent with the efficiency characterization. The coefficient of determination[4] actually increases with each experiment, from 0.75 to 0.83. From this, we conclude that multicast efficiency is not adversely affected by the total amount of time receivers remain in the group.

2) Receiver Arrival Time: Next, we dismiss the time element entirely (step c from Fig. 2) by randomizing receiver activity. The results are shown in Fig. 5. By ignoring the receiver arrival times, we explore whether the pattern in which the receivers originally joined the group has an effect on multicast efficiency. Comparing Figs. 4 and 5 (transition from step b to c), the magnitude of ε drops from 0.34 to 0.27 and the coefficient of determination rises to 0.94. Randomization in the time domain reduces outliers, and seems to lessen the potential benefit of multicast efficiency. As we randomly distribute active receivers, we eliminate temporal clustering, and, in turn, normalize the multicast benefit.

3) Receiver Distribution: Working with real data offers confidence in the efficiency estimate. The estimate properly characterizes the data from real multicast groups. It does not appear to depend heavily on group dynamics such as duration and inter-arrival time. However, to verify that the characterization is truly consistent with what might be found for most groups in the Internet, it is necessary to further investigate how receiver distribution might affect multicast efficiency. In truth, one might expect receiver distribution to be one of the major factors governing how well multicast performs in respect to unicast. Where the receivers are placed in the tree dictates how much sharing can possibly occur. In other words, changing the placement of receivers in the network changes where branching (packet duplication) occurs in the multicast tree.

[4]The coefficient of determination is the square of the correlation coefficient, in the range 0 to 1. High values reflect a good linear fit on a log–log scale.

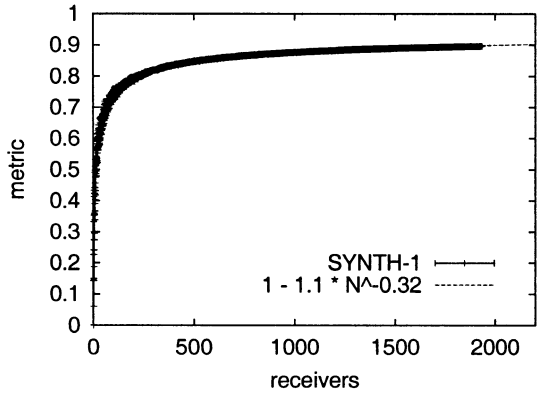

Fig. 6. (step d) Average metric calculated using the random receiver distribution of SYNTH-1 (99% confidence intervals).

The next step (d in Fig. 2) in our analysis, then, is to consider synthetic data, ignoring both temporal and spatial domains. The previous work, cited earlier, was based on random distributions for networks other than the modern Internet [18]–[20]. Should their conclusions necessarily hold for receivers randomly distributed throughout the Internet? The SYNTH-1 through SYNTH-4 data sets (Table II) consist of trees constructed within the current multicast infrastructure (see Section III-B for details on the procedure used to select receivers). This approach allows for much larger group trees to be constructed, nearly 1900 receivers in the case of SYNTH-1. More importantly, the topology of the underlying routing infrastructure is represented in the traces. As can be seen in Fig. 6, the estimate holds extremely well with a coefficient of determination of 0.99. All synthetic data sets produce almost identical results.

D. Discussion

What exactly have we accomplished? Most importantly, we have identified a characterization of multicast efficiency that is independent of group distribution and behavior. The efficiency estimate, $\delta = 1 - s_\varepsilon N^\varepsilon$ (4), has been shown to hold rather well for real group dynamics, real group distributions and random distributions in the modern multicast infrastructure. Furthermore, we have found that the efficiency factor, ε, tends to stay within a fairly limited range for real networks, between -0.27 and -0.38.

This characterization is a benefit to current ISPs which are considering whether to implement multicast. Using the estimate to calculate possible bandwidth savings, a network provider can compare the savings afforded by multicast against the cost of

its implementation. In this vein, we have focused only on the bandwidth efficiency of data transmission. Other problems that face network providers deploying multicast, such as increased router state and other overhead specific to routing protocols, domain independence, address allocation and billing, have been covered in detail in other works [1], [7], [8], [18], [29]. Furthermore, other aspects of efficiency, such as the delay or quality of content as perceived by a receiver, have not been addressed herein.

Another result of validating the efficiency metric has been to indirectly support the conclusions of previous work with real data [18], [19]. However, the economies-of-scale factor, k, of Equation (1) seems to be nearer 0.7 than 0.8 for our experiments. The trouble with comparing these two numbers directly, however, is that Chuang and Sirbu computed L_m, $\overline{L_u}$, and N somewhat differently [18]. They were interested in the cost of providing multicast to the leaf router, and, consequently counted any router with at least one member to be a single receiver, no matter how many actual receivers it served. For this paper, we chose to look at the efficiency of the entire multicast tree. Ignoring a large number of receivers at the last-hop seriously underestimates multicast's efficiency. This difference in approach explains much of the difference in the measured values of k.

A point that still needs to be addressed is the assumption mentioned in Section III-C. We simplified the calculation of L_u by assuming that the unicast and multicast paths were of equivalent length. In real topologies, this of course may not necessarily be the case. To test the assumption, we performed a series of unicast traceroutes [30] for the receivers in the SYNTH-4 data set. Of the 1018 receivers, 770 unicast paths were successfully recorded. Comparing these paths to the multicast paths used in the prior analysis, we find that our assumption holds. The average ratio of multicast to unicast path-length is 1.03 $\pm$0.03 with 99% confidence. An older experiment using the SYNTH-1 data set resulted in a lower average ratio of 0.97, an indication of the reduced role that tunnels play in the current multicast infrastructure.

A final consideration, then, is the effect that changes in the multicast infrastructure, like the migration from tunnels to native support, have on the efficiency characterization. Is the characterization simply a snapshot of the current topology or is it more fundamental? Can we expect the characterization to hold as multicast deployment continues? We note that the efficiency estimate was derived from previous work which included the early MBone architecture (1996) as one sample topology [18]. Moreover, our data sets were gathered over a three-year period in which multicast deployment has both evolved and consolidated considerably. Our initial traces were taken during a period in which the MBone was transitioning from tunnels to native multicast routing [5]. Since then, the majority of the Internet's multicast infrastructure has transitioned to hierarchical, interdomain routing, but the total number of connected networks has actually decreased [6]. It is difficult to predict the future of multicast deployment; however, due to the similarity across data sets of the efficiency factor ε we conjecture that future evolution in the infrastructure will not cause significant change in how multicast efficiency compares to unicast.

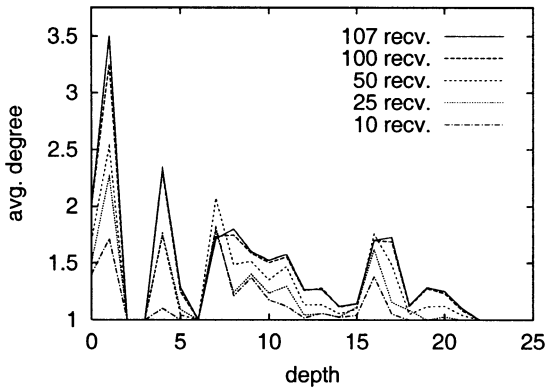

Fig. 7. Average **total** out-degree distribution for nodes at differing depths across a range of receiver-set sizes of IETF43-A.

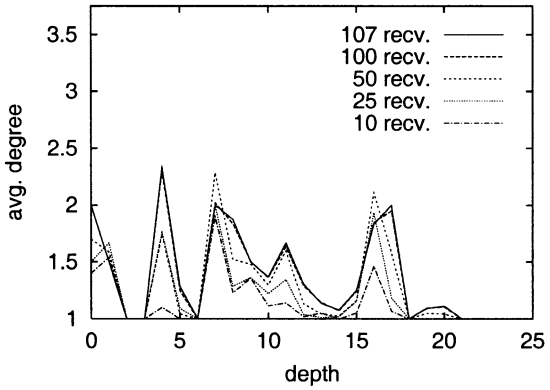

Fig. 8. Average **internal** out-degree distribution for nodes at differing depths across a range of receiver-set sizes of IETF43-A.

V. Modeling Multicast Trees

High-level relationships, such as cost and efficiency, can model the *behavior* of interdomain multicast trees, but does this imply that all real trees are similarly shaped or that different shapes produce relatively similar behaviors? In this section, we will explore the shape of the multicast trees from our data sets. Generally, we will focus on branching and where it occurs within the tree since a tree's topology is best described in terms of the degree of each of its nodes.

A. Degree Distribution

Throughout the discussion, we separate a node's *total* out-degree, the number of outgoing edges, into two distinct components: (a) *internal* degree and (b) *leaf* degree. The internal degree counts the number of out-going links between neighboring routers. The leaf degree refers to the number of incident links directly connecting receivers. It is important to note that receivers do not occur just in the deepest part of the tree.

Fig. 7 depicts the average of the total out-degree for nodes occurring at different levels of the distribution tree. Each of the real and random group distributions produces a similar graph. The majority of branching appears to occur near the top of the tree, but focusing only on the total out-degree can be misleading. Fig. 8 shows a similar graph for the internal degree rather than total out-degree. Comparing the two graphs, it is clear that the initial spike in Fig. 7 is due primarily to receiver clustering near

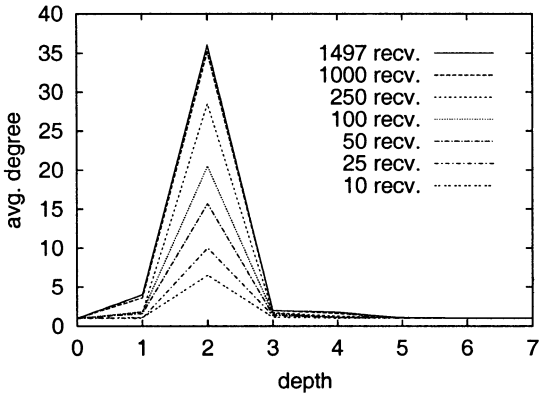

Fig. 9. Average **internal** out-degree distribution for ASs at differing depths across a range of receiver-set sizes of SYNTH-2.

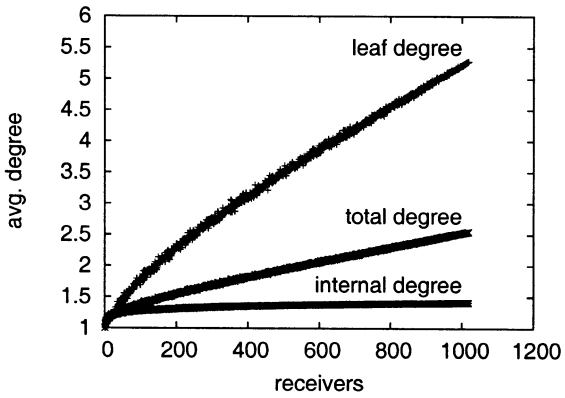

Fig. 10. Distribution of total, internal and leaf out-degrees as the receiver-set size grows from the SYNTH-3 data set.

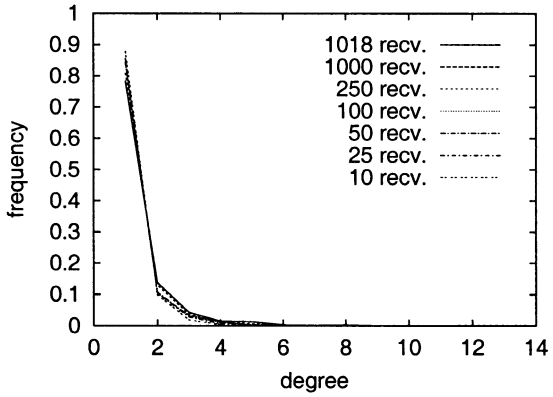

Fig. 11. Frequency of internal out-degrees across nodes for a number of receiver set sizes of SYNTH-4.

the source and is not representative of the body of the tree. Although this observation in itself is not quantitatively compelling, it does have some qualitative value. Considering internal and leaf degrees independently provides a better view of the body of multicast trees, in contrast to receiver clustering at the leaves.

At the AS level, branching is more concentrated. Fig. 9 shows that inter-AS branching actually does occur early in the tree. Each of the more recent data sets (SYNTH-2-4) reveals a large backbone AS just two hops from the source, where the majority of branching takes place. In all data sets, the backbone in question is the Abilene network, AS 11 537. It is important to note that all of our synthetic data sets were collected from universities in the United States. The specific backbone network would probably differ for sessions sourced from Europe or Asia, but we believe the result will be similar; a high degree of branching along a primary backbone network.

It is also interesting to compare this result over time. Looking back two years (SYNTH-1), we can still identify the Abilene backbone, but the difference is less prominent. We find significant branching in the vBNS network, as well as a few European providers. Shortly after this data set was traced, around the Fall of 2000, Internet2 stabilized. Many universities initiated connections to the Abilene backbone, and a number of international peers were added at STARTAP and NGIX-AMES. Since then, Abilene has become a significant backbone provider for multicast traffic in the U.S. This helps to explain the difference between the current and earlier data sets.

Fig. 10 illustrates how the average total out-degree depends on its internal and the leaf components. The leaf degree grows almost linearly with the number of receivers in the tree. The internal degree grows logarithmically, tapering off quickly.[5] The total out-degree shadows the internal degree until the two components pass around 25 receivers; then, the leaf degree dominates. The character of this relationship holds for all data sets, the primary difference being the steepness of the leaf degree distribution. The relationship is less apparent in the real data sets (IETF43-A, IETF43-V, and NASA-AV) due to the low number of receivers.

Focusing on internal branching, we observe that the average internal degree is consistent across data sets and quite low, ranging between 1.3 and 1.6 for routers and between 2.1 and 2.5 for ASs (see Table IV in the Appendix). Although the average remains fairly stable as the number of receivers increases, the actual range of internal degrees grows with the receiver count. One explanation for this variability may be that, as the tree grows, we increase the likelihood of encountering singular nodes with extremely high degrees, such as peering points along the backbone.

[5]We do not feel that modeling the average internal degree as a power-law is appropriate since the standard deviation is quite large.

B. Degree Frequency

So, if average internal degrees are low, how frequently do large degrees occur within the tree? Fig. 11 shows internal degree plotted against frequency for the SYNTH-4 data set. The graph is similar for every data set, both real and generated, as well as across degree components; the majority of nodes in multicast trees have an out-degree of one. This supports previous work which claims that a large percentage of nodes in real multicast trees are routers with a single out-going link [31]. These *relay* nodes simply pass the packet along the path to the receiver.

Skewed frequency distributions occur at the AS level as well. Compared to router degrees, AS distributions are slightly less skewed, but exhibit a much longer tail (Fig. 12). In other words, we observe more ASs with high degrees, and a lower percentage of relay nodes occurring in the AS-path from source to receiver. Both observations can be attributed to the fact that an AS is a superset of the routers within its domain. If the majority of

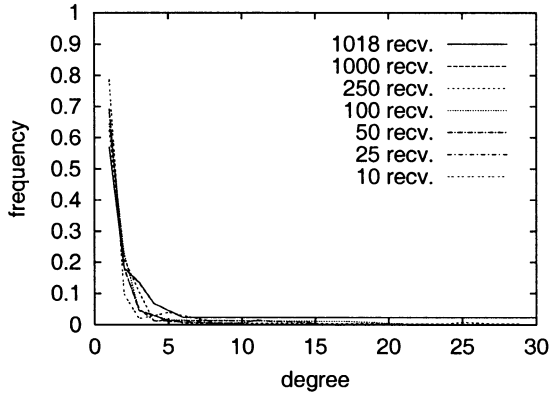

Fig. 12. Frequency of internal out-degrees across ASs for a number of receiver set sizes of SYNTH-4.

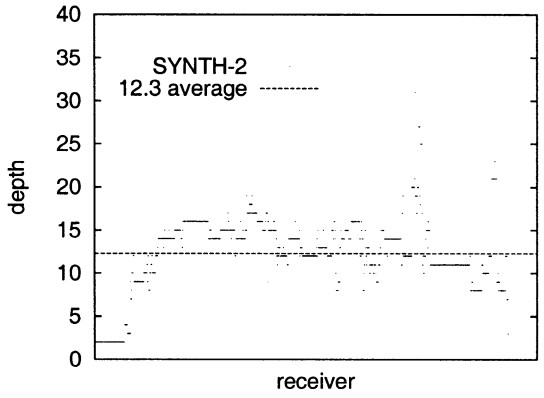

Fig. 13. Scatter plot of receiver depths for SYNTH-2; receivers are unordered. The average depth is indicated by the dotted line.

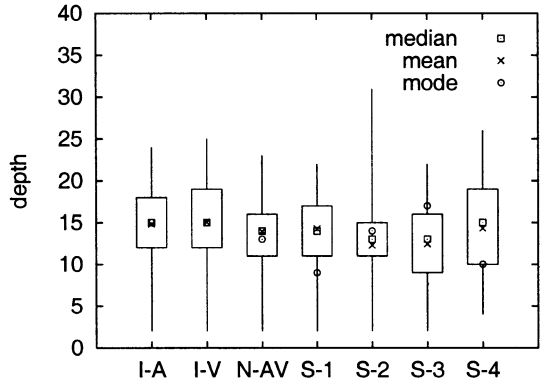

Fig. 14. Comparing receiver depth across data sets. Each box represents the inter-quartile range; the whiskers indicate the minimum and maximum values.

branching occurs at the edge of a domain, then the AS will likely have a higher degree than any of its constituent routers. Moreover, we expect a larger percentage of domains to incur some branching as a tree grows since there are far fewer ASs than individual routers.

So, why is there such a high frequency of relay nodes in multicast trees? We believe that the majority of path divergence occurs in a limited number of peering points along the key backbone networks. We found that the maximum internal degree within a domain (intra-AS) is less than that at the general router level (Table IV). This means that much of the branching is occurring at the borders of ASs.

Although the existence of a few high-degree branching points may be accounted to the currently sparse deployment of multicast, it has been shown that the majority of routers, themselves, in the Internet have a very low degree [12];[6] most routers have only one or two neighbors. Thus, it is our claim that *the possible range of tree shapes is necessarily constrained by the underlying network connectivity.*

C. Receiver Depth

As discussed in Section V-A, receivers are distributed throughout the height of the tree, although by no means uniformly. Fig. 13, a scatter plot of receiver depths, further demonstrates this with clusters of receivers present at varying depths with an average depth around twelve hops. The average receiver depth remains consistent as we vary the number of receivers in a given tree, and is similar across trees, ranging from 12.3 to 15.1 (see Table VI in the Appendix). This trend is also apparent at the AS level except that in a scatter plot, similar to Fig. 13, receivers are clustered much closer to the mean, which ranges between 4.8 and 6.6 for the synthetic data sets.

Simply using the average depth as a measure of central tendency, however, could be misleading since the router-level distributions appear to be primarily bi-modal. Fig. 14 presents a box-and-whiskers plot of receiver depth for the complete primary tree from each data set. This graph shows the total range of receiver depths (whiskers), as well as the inter-quartile range (box) where the middle fifty percent of the receivers reside. The median, mean and mode are also provided for comparison. Across data sets, the majority of receivers reside between 10 and 20 hops from the source.

D. Node Classification

When describing a multicast tree, we can classify each node as either: 1) a receiver; 2) a transit router; 3) a stub router; or 4) a *split* router that is directly connected to both receivers and other routers. An interesting question is: how does each class of node contribute to the overall structure of the tree?

When separating receivers from routers, we can see a clear trend across data sets. Fig. 15 plots the average total nodes (M) in the tree versus the number of receivers (N). Surprisingly, this relationship can be succinctly modeled as a power-law:

$$M = s_\tau N^\tau. \tag{5}$$

Across data sets, real and synthetic, the *saturation factor*, τ, falls within a tight range, 0.66–0.72 (see Table III in the Appendix). The function indicates that a large number of internal routers are required to support a small number of receivers, nearly three to one for 100 receivers. For sparse, interdomain trees, this is to be expected since receivers rarely join from the edge of their own ASs. Packets leaving the source must travel a number of hops to exit the source's domain, must cross the backbone(s) connecting source and receivers, and then must pass between the nearest backbone and each receiver. The early hops between the source and major backbones are shared between most receivers, but for truly sparse multicast groups, the later hops nearest the receivers are less likely to be shared. As the tree grows denser, however,

[6]Although the distribution is heavily skewed, our experiments do not indicate the existence of a power-law relationship for degree frequency in multicast trees.

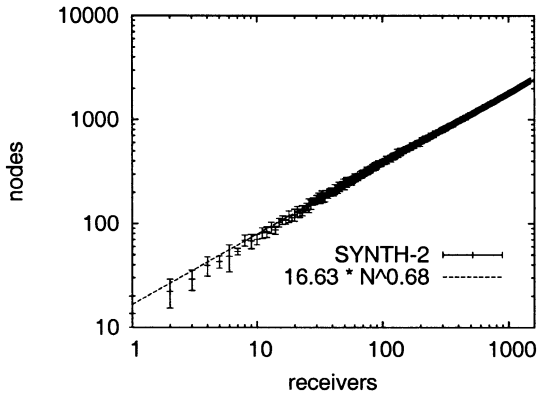

Fig. 15. Average total nodes versus receiver count for SYNTH-2, plotted in log–log scale (99% confidence intervals).

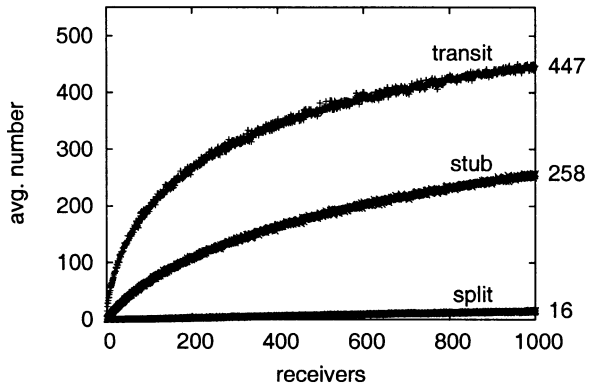

Fig. 16. Node class distribution for SYNTH-1. Right-hand labels indicate plotted values at 1000 receivers.

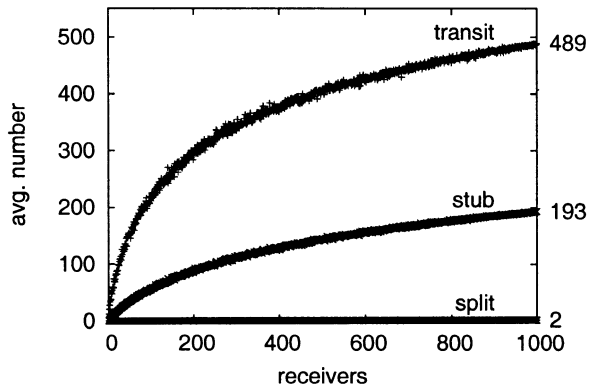

Fig. 17. Node class distribution for SYNTH-4. Right-hand labels indicate plotted values at 1000 receivers.

fewer internal nodes are necessary to support each receiver. At a thousand nodes, the ratio is closer to five receivers for every four routers.

Looking strictly at routers, Figs. 16 and 17 show the distribution of node classes as the number of receivers increases. Both graphs were limited to 1000 receivers for a better visual comparison. The two data sets were taken from the same source, two years apart. Comparing the two graphs, it is clear that the difference between the number of transit nodes and the number of stub nodes has increased with time. With 1000 receivers in SYNTH-4 (Fig. 17), purely transit nodes account

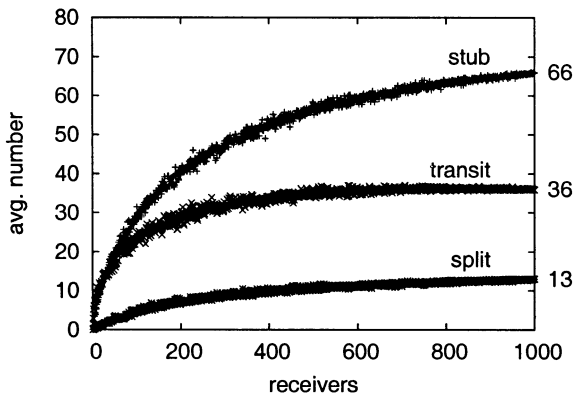

Fig. 18. AS-level class distribution for SYNTH-3. Right-hand labels indicate plotted values at 1000 receivers.

for more than 70% of all routers, compared to approximately 60% in SYNTH-1 (Fig. 16). Moreover, split routers are almost nonexistent in any of the recent traces (SYNTH-2-4). Both phenomena may be due to the overall reduction in multicast deployment witnessed over the last few years [6].

In all cases, including the real data sets, transit nodes out-pace stub routers with a low number of receivers. As the total number of receivers increases, the two sets begin to grow at similar rates. As we have already seen, a larger percentage of new receivers re-use much of the existing multicast tree. In a sense, the tree's backbone becomes saturated [18]; most of the possible branching has already occurred.

For ASs, we can make a similar classification. Fig. 18 is the class distribution for transit, stub and split ASs. A different picture emerges at this level. Here, stub ASs dominate throughout. The total number of transit domains tapers off, another sign of saturation. We also see a much higher occurrence of split domains than split routers. Considering the average AS path length is around five (Table VI), it seems reasonable that all but the true backbone domains would contain at least one receiver.

E. Discussion

Much work has been performed to determine the factors important in constructing realistic Internet topologies [4], [12]–[15], [32], [33], but little effort has been made to classify the shape of real multicast trees [17], [31]. Existing graph generation techniques already consider a number of factors when determining whether a graph is realistic; e.g., hierarchy, average node degree, network diameter, and the number of bi-connected components [34]. When generating random trees for testing multicast protocols, these factors can be extended with those discussed in this section in order to better evaluate the realism of generated tree topologies. In short, real multicast trees appear to exhibit:

- a high frequency of *relay* nodes through the body of each path, both at the router and AS levels;
- low average internal degrees that grow logarithmically with the number of receivers in the tree;
- average leaf degrees that grow almost linearly with the number of receivers in the tree;
- fewer internal nodes per receiver as the tree becomes saturated; and

TABLE III
PROPERTIES OF COMPLETE PRIMARY TREES

Data Set	*Receivers*	*Nodes*	*Graph Edges*	*Efficiency* ε	*Efficiency* s_ε	*Saturation* τ	*Saturation* s_τ	
IETF43-A	107	474	553	-0.28	0.91	0.72	13.74	
IETF43-V	129	530	612	-0.27	0.87	0.72	13.87	*Inter-AS*
NASA-AV	73	333	397	-0.29	0.98	0.71	13.74	*Routers*
SYNTH-1	1871	2764	2800	-0.32	1.10	0.69	15.26	182
SYNTH-2	1497	2426	2452	-0.32	1.40	0.68	16.63	178
SYNTH-3	1019	1677	1678	-0.32	1.24	0.68	15.49	111
SYNTH-4	1018	1710	1720	-0.34	1.20	0.66	17.34	115

TABLE IV
INTERNAL DEGREE WITH 90% CONFIDENCE INTERVALS

	Router Level								
Data Set	*Max*	*Median*	*Mean*	*Inter-AS Level*			*Intra-AS Level*		
IETF43-A	6	1	1.4 ±0.03						
IETF43-V	8	1	1.4 ±0.03	*Max*	*Median*	*Mean*	*Max*	*Median*	*Mean*
NASA-AV	5	1	1.3 ±0.03						
SYNTH-1	13	1	1.6 ±0.03	17	1	2.5 ±0.19	9	1	1.4 ±0.02
SYNTH-2	15	1	1.4 ±0.02	36	1	2.1 ±0.28	11	1	1.3 ±0.02
SYNTH-3	16	1	1.4 ±0.03	34	1	2.3 ±0.35	8	1	1.3 ±0.02
SYNTH-4	14	1	1.4 ±0.02	30	1	2.4 ±0.34	8	1	1.3 ±0.02

- average receiver depths of 12–15 hops at the router level or 5–6 ASs.

Finally, a high-level characterization such as cost (1) or efficiency (4) can be employed as a metric for validating the behavior of the tree.

It is important to note, however, that the state of multicast deployment is on-going. The observations made herein are applicable to the current multicast infrastructure. Whether new deployment efforts in commercial networks will result in similar properties for future multicast trees is an open question. It is our hypothesis, though, that many of the properties of multicast trees explored here, such as degree frequency and receiver depth, are actually artifacts of the underlying unicast topology. The branching behavior of a multicast tree is necessarily constrained by the available connectivity of each router within that tree. Moreover, interdomain multicast protocols have borrowed heavily from the success and scalability of hierarchical unicast routing. Both make use of peering points, such as MAE-EAST, where a number of backbone providers and local ISPs inter-connect, providing much larger branching points than are found within any single network. Unless the unicast topology undergoes some radical change in characterization for a sustained period of time, the multicast topology is unlikely to experience serious changes. Therefore, future multicast tree topologies should look similar to those seen today.

VI. FUTURE WORK

With any useful investigation, there are often more new questions raised than existing questions answered. The more we come to understand multicast trees, the more we feel there is to know. As a result, there are a number of open issues left for future work.

A. Degenerate Cases

So far, we have focused on rather sparse receiver distributions, attempting to characterize multicast's behavior on a wide scale. We have yet to look closely at truly degenerate cases, such as extreme affinity between receivers. It has been postulated that affinity and dis-affinity generally do not affect the form of the cost function (1) [19]. One might assume that it would affect the economies-of-scale factor while maintaining the power-law relationship of (1). However, looking at a case from the SYNTH-1 data set where a collection of 133 receivers were chosen from a single subnet, the ratio of multicast links to unicast links is no longer linear when plotted on a log-log scale, i.e., no longer a power-law. This implies that extreme affinity does affect the general form of the relationship. Degenerate behavior like this might possibly occur in real multicast groups if extreme temporal or spatial dependencies exist between members of the group.

B. Advanced Metrics

Another area of interest is in developing more advanced metrics for measuring multicast efficiency. In particular, we plan to look at techniques for weighting unicast streams to better capture the efficiency gains available through multicast. In the current model, each additional unicast stream that passes over a given link has an additive impact. In actuality, duplicating a stream over a link has implications beyond the extra bandwidth allocated to the stream since that bandwidth is no longer available to other multicast and unicast streams. Applying metrics that more aggressively penalize duplicate unicast streams, such as multiplicative or logarithmic metrics [35], may provide a more appropriate view of multicast's benefits. Another possible avenue is to weight the links themselves. Rather than simply using hop counts (weight of one), properties of the individual links, such as capacity and delay, may give a more accurate view of the impact of using multicast rather than unicast. Can these advanced metrics be characterized in the same way as the simple metric we have introduced here? Do changes in weighting simply affect the efficiency factor, or do they radically change the form of the characterization?

TABLE V
LEAF DEGREE WITH 90% CONFIDENCE INTERVALS

	Router Level					
Data Set	*Max*	*Median*	*Mean*			
IETF43-A	4	1	1.3 ±0.04			
IETF43-V	5	1	1.3 ±0.04	*Intra-AS Level*		
NASA-AV	5	1	1.3 ±0.05	*Max*	*Median*	*Mean*
SYNTH-1	133	2	5.3 ±0.29	133	2	5.3 ±0.32
SYNTH-2	108	2	5.6 ±0.35	108	2	5.9 ±0.39
SYNTH-3	89	2	5.3 ±0.37	89	2	5.5 ±0.42
SYNTH-4	91	2	5.2 ±0.37	91	2	5.4 ±0.41

TABLE VI
RECEIVER DEPTH WITH 95% CONFIDENCE INTERVALS

	Router Level									
Data Set	*Min*	*Max*	*Mode*	*Median*	*Mean*					
IETF43-A	2	24	15	15	14.8 ±1.17					
IETF43-V	2	25	15	15	15.1 ±1.11	*Inter-AS Level*				
NASA-AV	2	23	13	14	13.9 ±1.15	*Min*	*Max*	*Mode*	*Median*	*Mean*
SYNTH-1	2	22	9	14	14.3 ±0.21	1	11	5	6	6.6 ±0.11
SYNTH-2	2	31	14	13	12.3 ±0.28	1	9	5	5	5.0 ±0.10
SYNTH-3	2	22	17	13	12.5 ±0.36	1	8	5	5	4.8 ±0.12
SYNTH-4	4	26	10	15	14.3 ±0.36	1	8	5	5	5.1 ±0.08

VII. CONCLUSIONS

The focus of this paper has been to better understand the factors that determine the realism of multicast trees, and to quantify multicast's bandwidth gain over unicast. To this end, we have accomplished the following:

- developed a characterization of multicast that accurately describes its efficiency over a wide range of group dynamics and receiver distributions;
- validated, with real group data, previous work defining cost metrics for pricing multicast [18]–[20];
- discovered that the range of possible shapes of interdomain multicast trees is constrained by the underlying network connectivity;
- expanded on research investigating skewed distributions in networks [12], [14] to include properties of multicast trees;
- identified several properties of our sample groups that are representative of real trees and that could prove useful for tree-generation tools; and
- collected a number of synthetic data sets, with more than a thousand receivers, that have been shown to be representative of real multicast groups in both their efficiency characterization and general shape. This collection has been made available for other researchers to generate sample trees of varying sizes.[7]

APPENDIX I
DETAILED PROPERTIES OF DATA SETS

Tables III–VI present more detailed measurements for the primary trees from each of our data sets.

[7]Available at http://www.nmsl.cs.ucsb.edu/mwalk.

ACKNOWLEDGMENT

The authors wish to acknowledge J. Jaeggli and H. Kuhn, University of Oregon, for their assistance in expanding the data sets of the authors, and H. Chang, University of Michigan, for his code which maps IP addresses to their respective ASs.

REFERENCES

[1] C. Diot, B. N. Levine, B. Lyles, H. Kassem, and D. Balensiefen, "Deployment issues for the IP multicast service and architecture," *IEEE Network*, vol. 14, no. 1, pp. 77–88, Jan. 2000.
[2] A. Fei, J. Cui, M. Gerla, and M. Faloutsos, "Aggregated multicast: An approach to reduce multicast state," in *Proc. IEEE Global Internet (GLOBECOM'01)*, San Antonio, TX, Nov. 2001.
[3] S. Mittra, "Iolus: A framework for scalable secure multicasting," in *Proc. ACM SIGCOMM'97*, Cannes, France, Sept. 1997.
[4] C. R. Palmer and J. G. Steffan, "Generating network topologies that obey power laws," in *Proc. IEEE Global Internet (GLOBECOM'00)*, San Francisco, CA, Nov. 2000.
[5] K. Almeroth, "The evolution of multicast: From the MBone to interdomain multicast to Internet2 deployment," *IEEE Network*, pp. 10–20, Jan./Feb. 2000.
[6] P. Rajvaidya and K. Almeroth, "Analysis of routing characteristics in the multicast infrastructure," University of California, Santa Barbara, Tech. Rep., Jan. 2002.
[7] L. Wei and D. Estrin, "Multicast routing in dense and sparse modes: Simulation study of tradeoffs and dynamics," in *Proc. IEEE Conf. Computer Communications and Networks (ICCCN'95)*, Las Vegas, NV, Sept. 1995.
[8] K. L. Calvert, E. W. Zegura, and M. J. Donahoo, "Core selection methods for multicast routing," in *Proc. IEEE Conf. Computer Communications and Networks (ICCCN'95)*, Las Vegas, NV, Sept. 1995.
[9] K. Almeroth, "A long-term analysis of growth and usage patterns in the multicast backbone (MBone)," in *Proc. IEEE Conf. Computer Communications (INFOCOM'00)*, Tel Aviv, Israel, Mar. 2000.
[10] M. Handley, "An examination of MBone performance," USC/ISI, Marina del Rey, CA, Tech. Rep., Jan. 1997.
[11] M. Yajnik, J. Kurose, and D. Towsley, "Packet loss correlation in the MBone multicast network," in *Proc. IEEE Global Internet Conf. (GLOBECOM'96)*, London, U.K., Nov. 1996.
[12] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proc.ACM SIGCOMM'99*, Cambridge, MA, Aug. 1999.

[13] A. L. Barabasi, R. Albert, and H. Jeong, "Scale-free characteristics of random networks: The topology of the World-Wide Web," *Physica A*, vol. 281, no. 1-4, pp. 69–77, June 2000.
[14] A. Medina, I. Matta, and J. Byers, "On the origin of power laws in Internet topologies," *ACM Comput. Commun. Rev.*, vol. 30, no. 2, pp. 18–28, Apr. 2000.
[15] W. Aiello, F. Chung, and L. Lu, "Random evolution of massive graphs," in *Handbook on Massive Data Sets*, J. Abello, P. Pardalos, and M. Resende, Eds. Dordrecht, The Netherlands: Kluwer , May 2002, vol. 4, pp. 97–122.
[16] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "The origin of power laws in Internet topologies revisited," in *Proc. IEEE Conf. Computer Communications (INFOCOM'02)*, New York, June 2002.
[17] J. Pansiot and D. Grad, "On routes and multicast trees in the Internet," *ACM Computer Communication Review*, vol. 28, no. 1, pp. 41–50, Jan. 1998.
[18] J. Chuang and M. Sirbu, "Pricing multicast communication: A cost based approach," in *Proc. INET'98*, Geneva, Switzerland, July 1998.
[19] G. Phillips, S. Shenker, and H. Tangmunarunkit, "Scaling of multicast trees: Comments on the Chuang-Sirbu scaling law," in *Proc. ACM SIGCOMM'99*, Cambridge, MA, Aug. 1999.
[20] P. Van Mieghem, G. Hooghiemstra, and R. van der Hofstad, "On the efficiency of multicast," *IEEE/ACM Trans. Networking*, vol. 9, pp. 719–732, Dec. 2001.
[21] D. Makofske and K. Almeroth, "Real-time multicast tree visualization and monitoring," *Soft. Practice Exp.*, vol. 30, no. 9, pp. 1047–1065, 2000.
[22] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," Internet Engineering Task Force (IETF), Reston, VA, Tech. Rep. draft-ietf-avt-rtp-new-*.ps, Feb. 1999.
[23] W. Fenner and S. Casner, "A 'traceroute' facility for IP multicast," Internet Engineering Task Force (IETF), Reston, VA, Tech. Rep. draft-ietf-idmr-traceroute-ipm-*.txt, Aug. 1998.
[24] Cooperative Association for Internet Data Analysis (CAIDA). CoralReef. [Online]. Available: http://www.caida.org/tools/measure-ment/coralreef/.
[25] H. Chang, S. Jamin, and W. Willinger, "Inferring AS-level internet topology from router-level path traces," in *Proc. SPIE ITCom'01*, Denver, CO, Aug. 2001.
[26] University of Oregon. Route Views Project. Univ. of Oregon, Eugene. [Online]. Available: http://www.routeviews.org.
[27] R. Chalmers and K. Almeroth, "Developing a multicast metric," in *Proc. IEEE Global Internet (GLOBECOM'00)*, San Francisco, CA, Nov. 2000.
[28] ——, "On modeling the branching characteristics and efficiency gains of global multicast trees," in *Proc. IEEE Conf. Computer Communications (INFOCOM'01)*, Anchorage, AK, Apr. 2001.
[29] M. B. Doar and I. Leslie, "How bad is naive multicast routing," in *Proc. IEEE Conf. Computer Communications (INFOCOM'93)*, San Francisco, CA, Mar. 1993.
[30] V. Jacobson. (1989, Feb.) Traceroute. Lawrence Berkeley Lab. (LBL), Berkeley, CA. [Online]. Available: ftp://ee.lbl.gov/traceroute.tar.Z.
[31] I. Stoica, T. S. E. Ng, and H. Zhang, "REUNITE: A recursive unicast approach to multicast," in *Proceedings of IEEE Conference on Computer Communications (INFOCOM'00)*, Tel Aviv, Israel, Mar. 2000.
[32] C. Jin, Q. Chen, and S. Jamin, "Inet: Internet Topology Generator," University of Michigan, Department of EECS, Ann Arbor, Tech. Rep. UM-CSE-TR-433-00, 2000.
[33] A. Medina, A. Lakhina, I. Matta, and J. Byers, "BRITE: An approach to universal topology generation," in *Proc. 9th Int. Symp. Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS'01)*, Cincinnati, OH, Aug. 2001.
[34] E. W. Zegura, K. L. Calvert, and M. J. Donahoo, "A quantitative comparison of graph-based models for Internet topology," *IEEE/ACM Trans. Networking*, vol. 5, pp. 770–783, Dec. 1997.
[35] A. Legout, J. Nonnenmacher, and E. W. Biersack, "Bandwidth allocation policy for unicast and multicast flows," in *Proc. IEEE Conf. Computer Commun. (INFOCOM'99)*, New York, Mar. 1999.

Robert C. Chalmers (S'00) received the B.S. degree in computer science from the California State Polytechnic University, Pomona, in 1997. He is working toward the Ph.D. degree at the University of California, Santa Barbara, where his main research interests focus around leveraging intelligence within the network.

He has studied multicast and its effect on resource utilization, and how to provide services for small devices in edge networks.

Mr. Chalmers was awarded the Ericsson Fellowship in 2001 and is currently a Eugene Cota-Robles Fellow.

Kevin C. Almeroth (M'93) received the Ph.D. degree in computer science from the Georgia Institute of Technology, Atlanta, in 1997.

He is currently an Associate Professor at the University of California, Santa Barbara (UCSB), where his main research interests include computer networks and protocols, multicast communication, large-scale multimedia systems, and performance evaluation. At UCSB, he is a Founding Member of the Media Arts and Technology Program (MATP), Associate Director of the Center for Information Technology and Society (CITS), and on the Executive Committee for the University of California Digital Media Innovation (DiMI) program.

Dr. Almeroth is on the Editorial Board of IEEE Network, has co-chaired NGC 2000, Global Internet 2001, NOSSDAV 2002, and MMNS 2002; has served as tutorial chair for several conferences, and has been on the program committee of numerous conferences. He is serving as the chair of the Internet2 Working Group on Multicast, and is a Member of the IETF Multicast Directorate (MADDOGS). He is also serving on the advisory boards of several startups including Occam Networks, NCast, Hidden Footprint, and the Santa Barbara Technology Incubator.