# Developing a Multicast Metric

Robert C. Chalmers, Kevin C. Almeroth
Department of Computer Science
University of California
Santa Barbara, CA 93106-5110

*Abstract*—**To further advance the deployment of multicast, many consider it necessary to provide a quantitative measure of the potential benefit of employing multicast rather than unicast. We approach this problem in three ways. First, we propose a rudimentary metric that can provide a reasonable measure of the benefit achieved by an active group using multicast instead of unicast. Second, we discuss many of the issues inherent in developing multicast metrics and how these issues may be addressed to improve the accuracy of our model. Third, we consider how characteristics of network topology and group distribution might offer a predictive estimate for the proposed metric. In general, we identify two rules for evaluating the efficiency of multicast: (1)multicast efficiency increases as the average path length increases, and (2)fan-out close to the source decreases multicast's efficiency.**

## I. INTRODUCTION

Advocates for multicast have been struggling to convince the majority of the networking community to deploy multicast in their networks. One of the largest points of resistance has been that no quantitative measure exists to capture the potential benefit of using multicast rather than unicast. Many managers and other executives with the final authority to approve multicast projects may have the false impression that multicast is limited to streaming audio and video, and that deploying it will inevitably cause unwanted congestion in their networks. Both managers and network operators need to be convinced that their investment will be worthwhile in terms that they themselves can grasp and in turn relate to others in their organizations.

Monetary savings is, of course, the most widely understood measure. It is also one of the hardest to accurately quantify for aspects of network routing[1], considering the wide range of link speeds, capacities and costs, and the uncertainty present in best-effort routing. So, a simple metric is necessary to evaluate the effectiveness of deploying multicast, one that is reasonable to compute to some degree of accuracy, easily understood by those not familiar with multicast protocols, and general enough to be applied to the majority of network topologies found today. Looking a step further, a predictive metric could offer an estimate of potential savings prior to actual deployment which could serve as a basis for business case proposals, as well as a baseline for performance evaluation once multicast has been deployed.

To complicate the issue, however, different participants in the end-to-end delivery of network services, namely content providers, network providers and end-users, have differing levels of interest in how efficiently multicast performs over unicast. Implementing multicast has many implicit and explicit costs that must be considered, such as router resources and added network complexity. We will focus primarily on multicast's impact on bandwidth since it is one of the easiest factors to measure and is the common interest across participants.

In the next section, we further discuss the motivation behind defining this metric. In Section III, we formally define the metric and continue, in Section IV, to identify specific challenges present in implementing the metric. We further cover, in Section V, many of the general issues inherent in multicast metrics that make measuring multicast behavior difficult. In Section VI, we discuss predictive metrics and present an estimate for our multicast metric. We present future work in Section VII and conclude in Section VIII with a consideration of who the metric will benefit.

## II. MOTIVATION

To better understand the usefulness of our metric, it is important to provide some motivation behind its development. The primary benefit of using multicast is its scalability in terms of required resources per receiver. Since multicast groups are loosely coupled, the server does not necessarily know of individual receivers and consequently does not need to maintain a logical connection[1] for each new user. Furthermore, any link in the distribution tree carries only a single copy of a multicast packet while the number of unicast packets equals the number of downstream receivers. But who is interested in saving server and network resources, and who would benefit from multicast?

When discussing whether multicast delivery can outperform unicast, three main groups have identifiably different interests:

- *end-users.* End-users exhibit little or no interest in the level of efficiency that multicast may provide. As individuals, they care only if certain services are available. The method of delivery and its related efficiency are not considered relevant to consumer decisions unless they directly affect the cost or quality of the service.
- *service providers.* Service providers have an increased interest in a service's impact on local resources, especially server performance and LAN bandwidth, since the largest concentration of traffic crosses their networks. The bottom line is that the service providers are setting the prices and they should be interested in the relative efficiency of the delivery medium.
- *network providers.* Network providers must consider the needs of the service providers as well as ensure that their own

---

[1] A logical connection refers to the state maintained at the server for each unicast receiver such as the destination address and port.

resources, such as routers and high-traffic links, are used optimally. Unless they are charging the service providers based on actual bandwidth used, they too should consider whether implementing multicast could reduce the total amount of traffic over their networks. However, the savings in bandwidth is only one aspect of multicast's effect on the network. Network providers must also consider the actual cost of implementing and managing multicast.

The common factor across interests is bandwidth. We will focus primarily on comparing the impact of unicast and multicast on network bandwidth as a measure of efficiency. Factors such as group size and member distribution can affect how well multicast performs in relation to unicast. To accurately judge the benefit that an organization can expect to receive from implementing multicast, these factors and others must be reduced to a single value that can act as the basis for comparison. A quantifiable metric is needed.

## III. DEFINING A METRIC

In this section, we define a metric to quantify the efficiency of multicast in relation to unicast. Our rudimentary metric compares the total number of links traversed by multicast packets and unicast packets over a given topology. We will refer to each of these individual links as a **hop** in the path of a single unicast or multicast packet. Thus, the metric can be defined as a function of the ratio of multicast hops to unicast hops.

$$\delta = 1 - \frac{multicast\ hops}{unicast\ hops} \tag{1}$$

The multicast metric will be a fraction in the range $0 \leq \delta < 1$. When the value equals zero, there is no difference in the number of hops. As the value approaches one, the benefit of using multicast increases. Growth in the metric indicates higher efficiency gains when using multicast.

This metric is by its very nature specific to a particular topology and group membership. The multicast metric provides a quantitative value that represents the relative benefit of using multicast for a given configuration. As new receivers join or their relative location within the tree changes, so will the benefits provided by using multicast and consequently the value of the metric. The accuracy of the metric is dependent upon correctly determining the number of multicast and unicast hops for a given tree. As we will see, this is not always a simple task. Next, we propose a model for assessing the value of the metric on active groups, and further detail, in Section V, the challenges inherent in counting both multicast and unicast hops accurately.

## IV. A MODEL FOR ASSESSMENT

Having identified a metric, we need a model for computation. The challenges in producing this type of metric lie in accurately determining the hop counts for both unicast and multicast traffic for a particular tree. Accurately determining hop counts is difficult, in part, due to network heterogeneity and limited support for utilities. A number of factors such as multicast tunnels, asymmetric routes and dynamic group membership can complicate matters significantly. Simply determining the number of group members is nearly impossible without the use of application level protocols such as RTCP (Real-Time Control Protocol)[2]. Furthermore, the fact that most routes cross administrative boundaries makes security a primary issue that must be addressed when collecting information from individual routers.

We propose to simplify the problem: determine the multicast tree for a given group. Here, we assume that unicast and multicast paths are identical. In Section V, we will discuss several important aspects of the problem that are currently being ignored for the purpose of defining an initial model for assessment.

Within this simplified model, we do currently have tools to effectively determine the multicast tree and subsequently the multicast and unicast hops. Take a tool such as MHealth[3]. The Multicast Health Monitor gathers information about each source and receiver in a multicast group and then determines the paths between them. It uses RTCP feedback to determine each of the sources and receivers. Then mtrace (Multicast Trace)[4] is employed to perform traces from each receiver back to a given source. From this data a graphical representation of the multicast tree is constructed. Consequently, the hop counts can be computed directly by counting the links in each path of the tree.

## V. ISSUES IN MULTICAST METRICS

Having made the previous assumptions, how accurate is the metric? We feel that it is accurate enough to give an initial impression of how multicast can benefit a given scenario. We present it as a first step towards providing the information that managers need to make a quantitative decision on deploying multicast.

The next question: how can it be improved upon? The answer requires a closer look at what impact our simplification has on the problem. Next, we present a more detailed discussion of the explicit and implicit assumptions in our initial model.

### *A. Differences Between Unicast and Multicast Paths*

Using the multicast tree built with MHealth, we are assuming all unicast packets would follow the same links as found in the tree. This of course ignores the existence of indirect multicast routes due to limited infrastructure, the existence of tunnels, as well as differing unicast and multicast routing policies. Moreover, multicast routes are built using the reverse path from receiver to source while unicast routes are built from source to receiver[6]. Even in the case that all routers are capable of multicast and divergent routing policies are not a problem, the presence of asymmetric links may still cause the two paths to differ.

The most direct solution for accurately determining unicast paths is to initiate a third-party traceroute. However, this capability is not supported. Using the current implementation of traceroute one would have to initiate a trace from the source to each individual receiver. This would require access to the source which is rarely possible for actual multicast groups, but may be the case when a content provider is assessing his own service.

Another option is to use SNMP (Simple Network Management Protocol)[7] to iteratively query each router from the source to the receiver for the next hop router, but each query possibly traverses the entire routing table, loading the router. Security, which is discussed in a later subsection, is a major problem when routes cross administrative domains since rarely do routers handle SNMP queries without some form of authentication.

At the very least, an estimate of the difference between the multicast and unicast hop counts could be determined. By using unicast traces from a few hosts within a domain, an administrator could calculate a ratio to serve as a scaling factor used across all paths in the tree. This, of course, is only an estimate and if not used carefully may actually skew the results.

### *B. Multicast Tunnels*

In the MBone, the virtual multicast backbone of the Internet[5], multicast routes traverse tunnels that obscure the actual number of routers along the path. Unicast routers present in the tunnel are not reported by current tracing utilities such as mtrace. Thus, our model implicitly assumes that no tunnels exist in the multicast tree. This lowers the measured metric since duplicate unicast streams do not contribute their full cost when tunnels shorten their logical paths.

Existing tools such as MRInfo (Multicast Route Information)[2] can be used to identify the existence of tunnels by querying the router or host that is supporting the tunnel. This may require, however, that each hop be queried individually. Then, for each known tunnel a unicast trace could be performed to identify the intervening routers. Integrating this query with mtrace would provide a much clearer picture of the multicast tree and, in turn, improve the multicast metric, considering only one copy of a multicast packet will traverse the series of hops within the tunnel.

### *C. Multi-access Links*

If multicast routers are on a shared medium such as Ethernet, then many hops in the tree could possibly represent a single transmission. Multicast delivery is extremely efficient on broadcast media. Only one multicast packet must be sent for all routers on the shared link, but each unicast stream must still be duplicated since each packet holds the address of a distinct destination.

The current implementation of MHealth does not identify multi-access links and falsely assumes that each link is point-to-point. By accurately identifying these links in the multicast tree, fewer multicast hops would be counted yet the unicast hop count would remain the same, resulting in an increase in the multicast metric. Ignoring shared links and tunnels results in a worst case view of multicast efficiency.

### *D. Dynamic Membership*

Multicast group membership is dynamic. New sources and receivers can join and leave the group at any time. Consequently, our method works only on a snapshot of the multicast tree. This, however, is not really a problem. First, the metric need not be a static measure. Rather, it should be considered as a function of time. As the group membership changes, so does the metric. For real-time applications this may actually be more important than the value at any one instant in time. Second, the goal of the metric is to provide management and technical staff with a measure of how multicast is performing in relation to unicast. Using MHealth, this is accomplished without necessitating that the view of the tree be an exact representation of any one instantaneous topology. Rather, it is more important that the trend of the group's behavior and the comparison between multicast and unicast paths be represented faithfully.

### *E. Determining Receivers*

Accurately identifying each source and receiver is a necessary first step in constructing the multicast tree. Due to the dynamic quality of multicast membership, determining the exact number and location of each receiver is difficult. MHealth uses RTCP feedback to determine both the sources and receivers of a particular group. However, not all members necessarily produce RTCP packets that can be collected by MHealth. RTCP uses UDP and is inherently unreliable. Furthermore, some RTP implementations due not implement RTCP at all, and some firewalls do not allow RTCP traffic to flow across their boundaries.

### *F. Multiple Sources*

Most of our discussion has been centered around singly sourced trees. Some groups may have many sources, or in the case of teleconferencing, all receivers can also be sources. The question is, should the metric be considered on a per-source basis or on a group-wide basis? Both are equally feasible, the former being the most direct to compute. To consider the metric across the entire group simply sum the hop counts across each source and then apply the ratio accordingly.

### *G. Satellite and Dial-up Links*

The presence of satellite links and dial-up connections changes the calculation of the metric. Satellites are inherently broadcast media and in turn benefit the multicast environment much like multi-access links[8]. Over a satellite, the metric calculation simplifies to $\delta = 1 - \frac{1}{(no.\ of\ receivers)}$.

Dialup connections degrade the multicast path to individual unicast paths at the last hop. Each connection requires its own

---

[2] ftp://ftp.parc.xerox.com/pub/net-research/ipmulti/

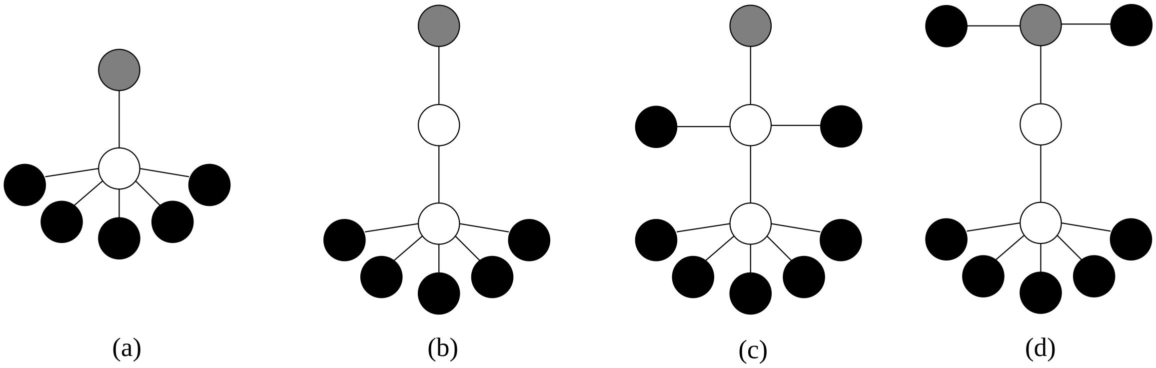

Fig. 1. Example tree topologies (grey=source router, black=receiving router).

copy of the multicast traffic resulting in a metric computation of $\delta = 1 - \frac{(multicast\ hops\ +\ PPP\ links\ -\ 1)}{(unicast\ hops)}$.

### H. Security

Performing any type of measurement across administrative domains can become difficult when security becomes an issue. Requesting information from a router or asking that a router perform some operation on your behalf uses resources that may otherwise be necessary for routing. Allowing public access to such features opens a large hole that will be exploited. Therefore, the majority of routers implement some form of security to restrict these types of access either through passwords or by disabling the service altogether.

In many cases, the metric will be used to measure multicast's impact on a single administrative domain. If the resources being queried are all within the control of the tester then any passwords can be supplied and the necessary capabilities can be enabled. For testing across domains, two possibilities remain. One, the tester can work in cooperation with the other domains thus eliminating the problem. Secondly, the networking community can come to some consensus on a set of operations that are available across domains and left unprotected.

## VI. PREDICTIVE METRICS

The use of the proposed multicast metric, so far, has been focused on comparing multicast and unicast services in active groups over existing topologies. However, a predictive measure of savings can be made. This provides managers and network engineers with an estimate prior to actual deployment. It could serve as a basis for business case proposals, as well as a baseline for performance evaluation once multicast has been deployed.

### A. Looking at Tree Shape

The behavior of the defined metric is inherently dependent upon the shape of the multicast tree. Two aspects of that shape, average path length and average fan-out (degree) with their associated standard deviations are good candidates for key parameters to a predictive metric. As the path length to a given set of receivers increases, the efficiency of multicast also increases. Take a set of five receivers all sharing a single intermediate router (see Fig. 1a). If the router were one hop away from the source, then the metric would be $\delta = 1 - \frac{6}{10} = 0.40$, and if it were two hops away (see Fig. 1b) the metric would be $\delta = 1 - \frac{7}{15} = 0.53$.

On the other hand, as the fan-out from a particular router increases, the efficiency of using multicast at that router decreases. For every outgoing link, the router must duplicate the packet it receives. However, the effect that fan-out has on the metric is dependent upon where in the tree the fan-out occurs. Take the prior example where the intermediate router was two hops from the source. If two new receivers joined through different links to the second router (see Fig. 1c), then the metric would be $\delta = 1 - \frac{9}{19} = 0.53$, but if they had joined higher in the tree (see Fig. 1d) the metric would fall to $\delta = 1 - \frac{9}{17} = 0.47$.

### B. Group Size as a Determining Factor.

One key factor that has been targeted for costing multicast traffic is group size[9]. Chuang and Sirbu have proposed a costing function which is closely related to our multicast metric and defines a direct relationship between hop counts and group size.

$$\frac{L_m}{L_u} = N^k \tag{2}$$

$L_m$ is the total length of the multicast distribution tree, $L_u$ is the average unicast routing path, $N$ is the multicast group size, and $k$ is an economies of scale factor in the range between zero and one. The interesting point made by the authors was that for the majority of topologies investigated $k$ was very near 0.8. This provides us with a formula for estimating the multicast metric assuming the total number of unicast hops is simply the average unicast path multiplied by the number of receivers.

$$\frac{multicast\ hops}{unicast\ hops} = \frac{L_m}{(L_u)(N)} \approx \frac{N^{0.8}}{N} = N^{-0.2} \tag{3}$$

$$\delta \approx 1 - N^{-0.2} \tag{4}$$

This is an appealing estimation. There are a few problems, however, that should be noted. First, Chuang and Sirbu focused on costing multicast over routed links, and considered the total number of receivers to be equivalent to the number of last hop routers. This does the multicast metric disservice. In the above estimate, many receivers on the same LAN are counted as a single receiver, but in reality, each individually contribute to the efficiency of multicast and the value of the metric. Second, it is unclear whether such an estimate accurately models real networks with real group dynamics. Although Chuang and Sirbu's work was later confirmed with a more rigorous mathematical treatment[10], both studies considered mainly generated topologies and random receiver distributions.

## VII. FUTURE WORK

The estimate presented in the previous section offers a useful characterization of multicast efficiency as a function of the group size, but can it be shown to hold for real networks. Taking a collection of traces from MHealth sessions, one could calculate the metric and compare it to the estimated value. Dependencies could be explored by varying parameters such as receiver duration, inter-arrival times and receiver distribution. We leave this for future work.

Another area of interest is to develop 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[11] may give 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 over unicast.

Characterizing multicast efficiency is only one piece of a larger problem: what constitutes a typical multicast tree? High-level metrics such as cost and efficiency are useful, but our eventual goal is to identify topological properties, such as tree height, average degree and degree frequency, which define and delineate classes of multicast trees, and ultimately to find a range of values for each property which ensures that generated trees are realistic.

## VIII. CONCLUSIONS

The motivation behind producing the multicast metric was to define a quantitative measure of the benefit of deploying multicast. For managers, the true measure should be in dollars saved. Using the multicast metric and average link cost, such a comparison is possible. The wide range of link costs, such as satellite and trans-Atlantic links, may make this comparison difficult to attain in a general sense. For a particular administrative domain, however, it is reasonable to conclude that such a measure could provide a valuable estimate for evaluating planned implementations, as well as confirm that the investment made in an existing infrastructure was sound.

Real-time applications could also benefit from the metric. To dynamically determine whether to source a particular service as unicast or multicast, the metric could be used to define an efficiency threshold above which the service provider uses multicast and below which unicast is used. As the number of active users grow and the multicast metric changes, the service could be switched to multicast where all users share the same common stream but must forfeit their individual control over that stream. This allows for a compromise between the interactive capabilities of unicast and the efficiency of multicast.

More than anything else, the multicast metric described offers a starting point for further development of more advance metrics, including a predictive metric, and related measurement techniques. Finally, it brings to light that the ultimate challenge in developing an accurate metric lies chiefly in gathering correct and meaningful statistics, and that development in this area is still needed before many of the management and measurement issues can be resolved adequately.

## ACKNOWLEDGMENT

The authors wish to thank David B. Makofske for his early contribution to the formulation of the metric and to some of the issues faced by multicast metrics in general.

## REFERENCES

[1] S. Herzog, S. Shenker, and D. Estrin, "Sharing the 'cost' of multicast trees: An axiomatic analysis," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, pp. 847–860, Dec. 1997.

[2] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," Tech. Rep. draft-ietf-avt-rtp-new-*.ps, Internet Engineering Task Force (IETF), Feb. 1999.

[3] D. Makofske and K. Almeroth, "MHealth: A real-time graphical multicast monitoring tool," in *Proceedings of Workshop on Network and Operating System Support for Digital Audio and Video (NOSSDAV'99)*, Basking Ridge, New Jersey, USA, June 1999.

[4] W. Fenner and S. Casner, "A 'traceroute' facility for IP multicast," Tech. Rep. draft-ietf-idmr-traceroute-ipm-*.txt, Internet Engineering Task Force (IETF), Aug. 1998.

[5] K. Almeroth, "The evolution of multicast: From the MBone to inter-domain multicast to Internet2 deployment," *IEEE Network*, January/February 2000.

[6] Beau Williamson, *Developing IP Multicast Networks*, Cisco Press, Indianapolis, IN, 1st edition, 2000.

[7] J. Case, K. McCloghrie, M Rose, and S. Waldbusser, "Protocol operations for version 2 of the simple network management protocol (SNMPv2)," Tech. Rep. RFC 1905, Internet Engineering Task Force (IETF), Jan. 1996.

[8] K. Almeroth and Y. Zhang, "Using satellite links as delivery paths in the multicast backbone (MBone)," in *Proceedings of ACM/IEEE International Workshop on Satellite-Based Information Services (WOSBIS'98)*, Dallas, Texas, USA, Oct. 1998.

[9] J. Chuang and M. Sirbu, "Pricing multicast communication: A cost based approach," in *Proceedings of INET'98*, Geneva, SWITZERLAND, July 1998.

[10] G. Phillips, S. Shenker, and H. Tangmunarunkit, "Scaling of multicast trees: Comments on the Chuang-Sirbu scaling law," in *Proceedings of ACM SIGCOMM'99*, Cambridge, Massachusetts, USA, Aug. 1999.

[11] A. Legout, J. Nonnenmacher, and E. W. Biersack, "Bandwidth allocation policy for unicast and multicast flows," in *Proceedings of IEEE Conference on Computer Communications (INFOCOM'99)*, New York, USA, Mar. 1999.