On the Degree of Multicast Bandwidth Sharing in Existing Internet Routing

K. Schuh, D. Panchal, and H. Fujinoki
Department of Computer Science
Southern Illinois University at Edwardsville
Edwardsville, IL 62026-1584

Abstract-**In this study, the degree of bandwidth sharing (i.e., how much bandwidth could be shared by using multicasting) in the Internet is examined for multicast network applications, where receivers dynamically join and leave a multicast session. The degree of bandwidth sharing is studied by measuring the extent to which the links could be shared by multiple receivers to common multicast senders. Paths to imaginary multicast senders are traced from multiple receivers using the traceroute servers. It is found that the degree of bandwidth sharing ranges from 0.63 to 0.88, suggesting that from 63 to 88% of bandwidth consumption by unicasting could be saved by multicasting. Results of tracing 49,720 paths suggest that multicasting will be efficient in bandwidth saving for multicast sessions where receivers are connected to a multicast sender by long paths.**

I. INTRODUCTION

One of the largest advantages in multicasting is bandwidth sharing by multiple receivers. Different from legacy unicast (i.e., point-to-point) data transmission, data packets for multiple receivers are duplicated only at a router where a path is split for multiple receivers in the downstream. This implies that a multicast sender does not have to transmit multiple copies of identical packets to receivers, as long as their paths are overlapped. This results in bandwidth sharing of links up to a bifurcation point, which is usually a router [1]. The current routing algorithm in the Internet finds the best path between two end hosts to minimize the hop count in terms of the number of intermediate Autonomous Systems (ASes), but not to minimize the total bandwidth consumption for multiple receivers [2]. Biersack studied the effect of bandwidth sharing for n receivers where an average degree of connectivity at each node is k. Biersack [1] suggested that the average degree of bandwidth sharing by n receivers will be:

$$\log_k(n)\,\frac{n}{n-1}\,\frac{k-1}{k} \quad (1)$$

In this paper, the term, "the current routing algorithm in the Internet", means the AS-level routing (also called exterior routing) used in the current Internet. A link is meant to be a logical link that inter-connects two routers. Therefore, a logical link may consist of multiple physical links. The term "link" is assumed to mean a logical link hereafter.

II. BACKGROUND

Since the expected multicast applications are real-time multimedia applications, reducing bandwidth consumption by bandwidth sharing has importance. In this respect, various multicast routing algorithms and protocols have been proposed to maximize bandwidth sharing in multicast routing [3, 4, 5]. These algorithms and protocols were proposed based on the assumption that the current unicast routing algorithm for the Internet is inefficient in bandwidth sharing. However, to the best of the authors' knowledge, there does not exist any work to study how inefficient the current unicast routing algorithm in the Internet will be as a multicast routing algorithm in terms of bandwidth sharing (Doar estimated the efficiency of multicasting by simulation experiments, but not by examining actual networks [6]). It could happen that the current unicast routing algorithm in the Internet may perform well enough in terms of bandwidth sharing, which might negate the raison d'être of the multicast routing algorithms and protocols proposed so far. However, how bad (or how good) the current routing algorithm in the Internet will be is one of the unanswered questions in the area of multicast routing.

III. A METHOD FOR MULTICAST PATH-TRACES

To accurately measure the sharing of the links by multiple receivers, paths to imaginary multicast senders are simultaneously traced. "Imaginary multicast senders" mean that the hosts used as multicast senders in this study are not multicast senders in reality but simply world-wide-web file servers distributed throughout the world in the Internet. To perform this task, ***traceroute servers*** distributed in the Internet are used as multicast receivers. Paths to imaginary multicast senders are traced using the current routing unicast minimum hop-count based algorithm in the Internet. In order to trace from multiple receivers to a common imaginary multicast sender (an imaginary multicast sender is called a multicast sender hereafter for simplicity), we used traceroute servers. A traceroute server is a web site that accepts traceroute requests to a particular destination in the Internet and returns a traceroute result as an html file.

To examine how much bandwidth saving would be possible by using multicasting, the number of links shared by multiple receivers is counted. To accurately count the number of the links dynamically shared by receivers, paths from receivers to a multicast sender should be traced simultaneously. To achieve this goal, a program called ***requestor*** was developed. Fig. 1 shows the procedures of this operation. The program, ***requestor***, is executed at a computer in the Department of Computer Science at Southern Illinois University Edwardsville (SIUE). The program simultaneously sends traceroute requests to multiple multicast receivers, which are traceroute servers, to request them to perform `traceroute` to a common multicast sender in the Internet (shown as ① in Fig. 1). Each request to traceroute servers is a command for CGI (Common Gateway Interface) script that is to be executed at a traceroute server. Then those traceroute servers simultaneously perform `traceroute` to a common multicast sender (shown as ② in Fig. 1). After the traceroute servers completed path traces to a common multicast sender, they return their traceroute results to the ***requestor*** program running at SIUE (shown by ③ in Fig. 1). The ***requestor*** program receives the results from the traceroute servers and stores the results in its local hard disk for later analysis.

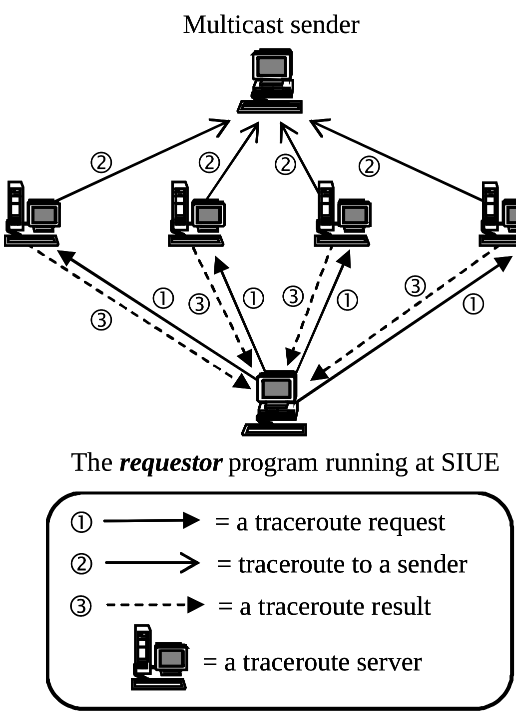

Fig. 1. Procedures for multicast path traces

In analyzing the collected traceroute results, the number of links from all receivers to a common multicast sender is counted. It is called "the total link count (represented by C_{Total})". Then, the number of links actually used is counted. This is called "the actual link count (represented by C_{Actual})". For the actual link count, if a link is shared by more than one receiver, it is counted as one. The degree of bandwidth sharing is defined as:

$$1-\frac{C_{Actual}}{C_{Total}} \qquad (2)$$

Fig. 2 shows an example of how to calculate the total link and the actual link count. In Fig. 2, node ***S*** is the router where a multicast sender is located. The black nodes indicate the routers where at least one multicast receiver exists and the white nodes are the intermediate multicast routers where no multicast receiver exists. Intermediate multicast routers simply relay multicast traffic. In the figure, the three routers, ***R2***, ***R4*** and ***R5*** are the routers where at least one active multicast receiver exists. Routers, ***R1*** and ***R3*** are intermediate multicast routers. Since the total link count is simply the sum of the number of links used by the all multicast receivers, it is 8 (2 links for ***R2***, 3 for ***R4***, and 3 for ***R5***). Since the bandwidth of a link can be shared if a link on a path overlaps in more than one multicast receiver, only 5 links are actually used (links ***S-R1***, ***R1-R2***, ***R1-R3***, ***R3-R4*** and ***R3-R5***). Therefore, the actual link count is 5. Calculating the actual link count requires $n?(n+1)/2$ comparisons of paths, where n is the total link count. In this particular example, the degree of bandwidth sharing is $1.0-5/8$, which is 0.375. The degree of bandwidth sharing represents the efficiency of multicasting for bandwidth saving. If the degree of the bandwidth sharing is high, that means multicasting potentially saves much link bandwidth by sharing link bandwidth for many receivers.

IV. RESULTS OF MULTICAST PATH TRACES

We found 194 traceroute servers in twenty two different countries in the world (currently we are looking for more traceroute servers). Those 194 traceroute servers are classified in terms of their geographical location. The traceroute servers

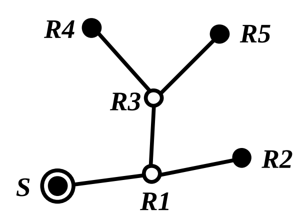

Fig. 2. Example multicast session

are classified as the United States, Europe, Asian traceroute servers or traceroute servers from other regions. For the United States (will be abbreviated as US hereafter), 73 traceroute servers have been found. For Europe, 70 traceroute servers have been found and they are from thirteen different countries; Czech Republic, Denmark, Estonia, Finland, France, Germany, Italy, Latvia, Netherlands, Sweden, Switzerland, United Kingdom and Ukraine. For Asia, 51 traceroute servers have been found and they are from five different countries; China, Japan, Korea, Taiwan and Thailand. To examine geographical location of multicast senders and traceroute servers, we used GTrace [7]. From the found traceroute servers, 45, 25 and 30 traceroute servers were randomly chosen for the US, Europe and Asian traceroute servers, respectively. Requests to traceroute servers are prepared as a command for traceroute CGI script running at a traceroute server. Paths to multicast senders in the US, Europe and Asian countries were then traced from the traceroute servers within the US, Europe and Asian countries. The lists of traceroute servers and the multicast senders used in this study are available at the authors' homepage [8].

The path traces were conducted as follows. The requestor program was executed for each group of traceroute servers and multicast senders pair. A ten-second wait was inserted for every five traceroute servers. This means the requestor program first sends a traceroute request to each of the five traceroute servers. Then the requestor program waits for the traceroute results from the five traceroute servers. After the requestor program receives the traceroute results from all the five traceroute servers, it waits for ten seconds. After the ten-second wait, the requestor program issues a traceroute request to another five traceroute servers. This is to avoid overloading common multicast senders.

The traceroute results contain a list of IP addresses of the routers on a path from traceroute servers to multicast senders. Each link is identified by two IP addresses, the IP address of the beginning router and the IP address of the ending router of a particular link. If a link that has the same IP addresses for both the beginning and ending router appears in another path, it is considered as shared.

The requestor program was executed from February 11, 2002 through February 18, 2002. The number of multicast senders that have been traced is shown in Table 1. Total of 1,361 multicast senders were traced to. The number of traceroute servers used for tracing paths to the multicast senders is shown in Table 2. Table 3 shows the numbers of traced paths for different traceroute servers and multicast senders pair, which are the products of the corresponding entries in Table 1 and Table 2 (e.g., the entry for US traceroute servers and US multicast senders in Table 1 multiplied by the entry for the US traceroute servers and US multicast senders in Table 2). Table 4 shows the average path length (in the number of links in a path) for different traceroute servers and multicast senders pair. Table 5 shows the average degree of

TABLE 1
NUMBER OF MULTICAST SENDERS

Traceroute Servers	Locations of multicast senders		
	US	Europe	Asia
US	365	188	103
Europe	53	150	35
Asia	125	150	200

TABLE 2
NUMBER OF TRACEROUTE SERVERS

Traceroute Servers	Locations of multicast senders		
	US	Europe	Asia
US	45	45	45
Europe	25	25	25
Asia	30	30	30

TABLE 3
MUMBER OF TRACED PATHS

Traceroute Servers	Locations of multicast senders		
	US	Europe	Asia
US	16,425	8,460	4,635
Europe	1,325	3,750	875
Asia	3,750	4,500	6,000

TABLE 4
MEAN PATH LENGTH

Traceroute Servers	Locations of multicast senders		
	US	Europe	Asia
US	12.2	11.3	14.9
Europe	12.6	10.4	12.4
Asia	14.3	16.5	14.1

TABLE 5
AVERAGE DEGREE OF BANDWIDTH SHARING

Traceroute Servers	Locations of multicast senders		
	US	Europe	Asia
US	0.78	0.88	0.85
Europe	0.76	0.76	0.68
Asia	0.71	0.87	0.63

bandwidth sharing for different traceroute servers and multicast senders pairs.

Throughout the path traces, 45, 25 and 30 traceroute servers in US, Europe and Asian were used, respectively. From the 45 US traceroute servers, the paths to 365, 188 and 103 common multicast senders within US, Europe and Asian, respectively, have been traced. Similarly, from the 25 Europe traceroute servers, the paths to 53, 150 and 35 common multicast senders have been traced. Finally, from the 30 Asian traceroute servers, the paths to 125, 150 and 200 common multicast senders have been traced.

As a result of these traceroute activities, 49,720 individual paths have been traced. In calculating the actual link count, the total of 224,890,710 link comparisons were performed. The average path lengths in terms of the number of links on a path range from 10.4 links (from European traceroute servers to European multicast senders) to 16.5 links (from Asian traceroute servers to European multicast senders). Table 4 shows these results.

In Table 5, the average degrees of bandwidth sharing for different traceroute servers and multicast senders pair are shown. The average degrees of bandwidth sharing for different traceroute servers and multicast senders pair were calculated using all the path trace results for each traceroute and multicast sender pair. For example, for the US traceroute servers and the US multicast senders pair, the average degree of bandwidth sharing was calculated from the 16,425 paths. It is found that the mean degrees of bandwidth sharing range from 0.63 to 0.88 for the 1,361 multicast senders. Fig. 3. shows the results of the average degrees of bandwidth sharing and the average path lengths for the nine different combinations of the locations of traceroute servers and multicast senders. The left Y-axis in the figure indicates path length in terms of the number of links in a path, while the right Y-axis indicates the degree of bandwidth sharing. The nine numbers (from 1 to 9) in the X-axis indicate the nine different combinations of locations of traceroute servers and multicast senders. The nine number are defined as follow: 1, 2 and 3 for US traceroute servers and US, European and Asian multicast senders, respectively, 4, 5 and 6 for European traceroute servers and US, European and Asian multicast senders, respectively, and 7, 8 and 9 for Asian traceroute servers and US, European and Asian multicast senders, respectively. In Fig. 3. it can be seen that as the path length increases, the degree of bandwidth sharing also increases. It was observed that for the two combinations, US traceroute servers and Asian multicast senders (shown by "3" in the X-axis in Fig. 3), and Asian traceroute servers and European multicast senders

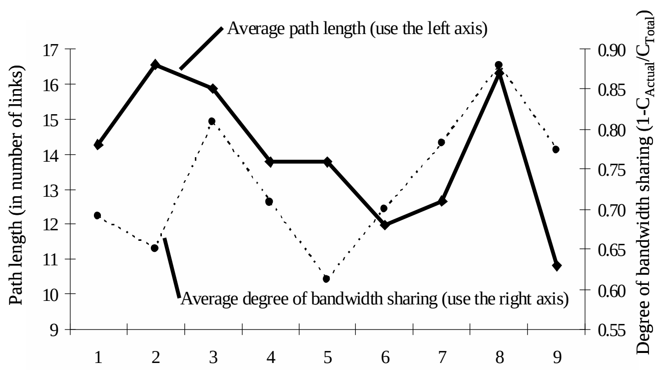

Fig. 3. Average degree of bandwidth sharing for various trace-servers and multicast senders pairs

(shown by "8" in Fig. 3), both path length and the degree of bandwidth sharing are high. For European traceroute servers and multicast senders pair (shown by "5" in Fig. 3), both path length and the degree of bandwidth sharing are low. One of the possible interpretations of these observations could be that if path length is long, multicasting can efficiently save bandwidth consumption by effectively sharing the link bandwidth since multicast tree is deep and narrow, while if path length is short, multicast tree is shallow and wide, for which existing shortest-path based unicast routing will result in relatively efficient routing for bandwidth sharing. From these observations, it can be concluded that existing unicast routing is inefficient for bandwidth sharing if receivers are distributed in geographically wide area and path lengths are long. In another word, multicasting will be efficient in bandwidth saving for multicast sessions where receivers are distributed in a wide area and are connected to a multicast sender by long paths.

V. CONCLUSIONS AND FUTURE WORK

In this paper, the paths in the existing Internet have been etraced to study the efficiency of the multicast routing for bandwidth sharing. The total of 49,720 paths have been examined to calculate the degree of bandwidth sharing, which is an index of how much link bandwidth could be saved over existing unicast routing by using multicasting. The following observations were made:

- If a multicast sender and receivers are located within the geographical locations nearby, the paths from multicast receivers to a multicast sender will result in short paths.
- If path length increases, the degree of bandwidth sharing also increases while if path length decreases, the degree of bandwidth sharing decreases.

The above two observations suggest that multicasting will be efficient for multicast sessions where multicast receivers are distributed all over the Internet with long paths. This conclusion also suggests that multicasting will be a more efficient technique for the Internet routing than for local area routing for bandwidth sharing.

Currently, there are four ongoing activities in this project. First, multicast senders are classified by ASes with similar size, instead of geographical locations. The size of an AS is measured in terms of the number of inter-AS connections at an AS, since the way an AS is inter-connected to other ASes was found to have a significant impact on the performance of Internet routing in terms of bandwidth consumption and on path lengths [9]. Second, multicast senders are classified by the Internet domain groups. Multicast senders will be chosen from commercial sites, educational sites and government sites. As the third activity, we are planning to compare the bandwidth consumption of multicast paths calculated by the existing Internet routing algorithm to that of the paths with the minimal bandwidth routing in the Internet. Finally, the number of traceroute servers will be changed to see the effect of different multicast receiver densities. Currently, we are working to collect more traceroute results for a target of ten thousand paths and to analyze the traceroute results to identify the most significant factor that determines the degree of bandwidth sharing.

REFERENCES

[1] A. Legout, J. Nonnenmacher, and E. Biersack, "Bandwidth allocation policies for unicast and multicast flows," ***Proceedings of IEEE INFOCOM***, pp. 254-261, 1999.

[2] Y. Rekhter and T. Li, "A border gateway protocol 4 (BGP-4)", ***RFC 1771***, March 1995.

[3] B. Waxman, "Routing of multipoint connections," ***IEEE Journal of Selected Areas in Communications***, vol. 6, no. 9, pp. 1617-1622, December 1988.

[4] X. Jia, "A distributed algorithm of delay-bounded multicast routing for multimedia applications in wide area networks," ***IEEE/ACM Transactions on Networking***, vol. 6, no. 6, pp. 828-837, 1998.

[5] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C. Liu, and L. Wei, "The PIM architecture for wide-area multicast routing," ***Proceedings of IEEE Journal on Selected Areas in Communications***, vol. 15, no. 3, pp.277-290, April, 1997.

[6] M. Doar and I. Leslie, "How bad is naive multicast routing?" ***Proceedings of IEEE INFOCOM***, pp. 82-89, January 1993.

[7] GTrace - A Graphical Traceroute
URL: ***http://www.caida.org/tools/visualization/gtrace/***

[8] H. Fujinoki, Homepage
URL: ***http://www.siue.edu/~hfujino***

[9] H. Fujinoki and K. Christensen, "The directed reverse path join (DRPJ) protocol: an efficient multicast routing protocol," ***Computer Communications***, vol. 24, no. 12, pp. 1121-1133, July 2001.