A New Delivery Scheme for 1-to-N Multicast Applications

Juyoung Park[1], Seok Joo Koh[1], Eunsook Kim[1], Shin-Gak Kang[1], and Dae Young Kim[2]

[1] Electronics and Telecommunications Research Institute,
161 Kajeong-Dong, Yuseong-Gu, Deajeon, KOREA
{eunah, jypark, sjkoh, sgkang}@etri.re.kr

[2] Chungnam National University,
200 Kung-Dong, Yusong-Gu, Daejeon, KOREA
dykim@cnu.ac.kr

Abstract. We propose a new delivery scheme for 1-to-N multicast applications such as webcasting service used for the web-based broadcasting of contents streams in Internet. The proposed scheme is based on the unicast transport from a remote sender to a subnet host and the subsequent multicast transmissions to the other receiving hosts within the subnet. The main design goal of this scheme is to improve transmission efficiency over the existing unicast transports, without help of multicast-capable networks. The proposed scheme has been tested and compared with the unicast-only transports in terms of the amount of traffic generated in the subnetworks. From experimental results, we see that the proposed scheme provides a nearly same performance as IP multicasting does in the subnetwork environments. Recognizing that IP multicast seems to be a little more delayed, it is expected that the proposed scheme can be used as an alternative short-term solution for one-to-many multicast services over the unicast networks.

1. Introduction

Webcasting, as known as Internet broadcasting, has recently been focused as a multicast killer application service. Webcasting literally means the web-based distribution of multimedia contents to Internet users. Webcasting is expected to realize a variety of the commercial multimedia services such as Internet TV and movies, remote education, and stock tickers [1].

Recognizing that the webcasting service can be viewed as one-to-many multicast application service for numerous users, it is reasonable to use IP multicasting, rather than the replicated unicast transport to each of the webcasting users. Nevertheless, IP multicasting has not yet been widely deployed in the public Internet [2]. Actually, there still exist a lot of issues to be addressed for rapid deployment of IP multicasting, as pointed out by Christope Diot, et al. in [3]. In reality, most of the current webcasting services are being provided over the unicast networks. The recently

I. Chong (Ed.): ICOIN 2002, LNCS 2343, pp. 109-118, 2002.

focused Contents Delivery Networks (CDN) service is also dependent on unicast transports in the networks [4].

In this paper, we proposed a new delivery scheme for one-to-many multicast applications. The proposed scheme is based on unicast transport from a remote sender (contents provider) to a subnet host (client) and the subsequent multicast transmissions to the other receiving hosts located in the subnet. The proposed scheme is designed to exploit the trivial multicast capability in a subnet environment which is called subnet multicast [5].

The main design goal of the scheme is to improve efficiency of the replicated unicast transports for webcasting services without modification of the current network infrastructure. The study of the scheme is also motivated from recognition that a large number of Internet users are usually located in the local area network (LAN) environments, as shown in the examples of the private enterprise networks and Digital Subscriber Lines access networks, in which the subnet multicast can easily be achieved.

This paper is organized as follows. In Section 2, the webcasting service is briefly presented as an example of one-to-many multicast application services. Section 3 describes the proposed scheme for the unicast transports with subnet multicast. In Section 4, more detailed operations for subnet multicast transport are presented along with the associated control messages. Section 5 discusses the analytical and experimental results that have been performed for comparison of the proposed scheme with the unicast-only transports. In Section 6, we conclude this paper.

2. Webcasting Services

Webcasting is a typical example of one-to-many multicast services and can be considered as a multicast killer application service. The webcasting represents the web-based distribution of streaming multimedia contents to Internet users. The webcasting system established in contents providers roughly consists of the various servers and storage equipments. Each client requests transmission of streaming contents to the sender via the web server. The requested contents are delivered over the network.

Figure 1 illustrates an example webcasting system that has widely been deployed. The contents provider generates live or video on demand (VoD) data streams by encoding raw audio/video materials. Some of those contents may be recorded into the storage devices. Each client or user contacts a web server located at the contents provider so as to request contents that it wishes to receive. The web server guides the user onto a suitable media server. Then a connection is established between user and media server to deliver the contents stream. The streaming data will be transmitted to the user by unicast or multicast transport in Internet.

Most of the current webcasting systems use the replicated unicast transports for contents delivery to numerous users. This incurs severe traffic overload at the webcasting (processing) system as well as at the network access link. This inefficiency gets severe, as the number of simultaneous access users increases. In this paper, we discuss a simple and realistic delivery scheme so as to improve efficiency of the current unicast-only transports.

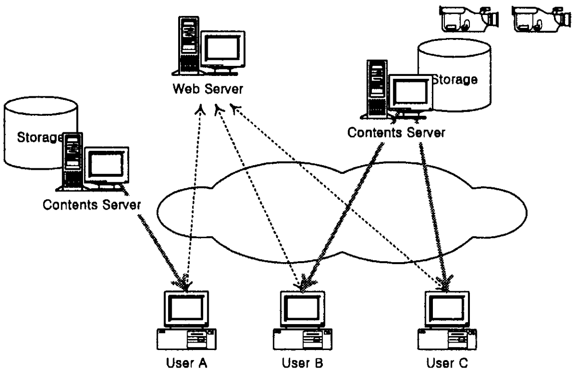

Fig. 1. Webcasting system

3. Unicast Transport with Subnet Multicast

In this section, we describe the scheme for the unicast with subnet multicast transport, and also present the associated operational procedures for the webcasting service.

Differently from the unicast-only transport, the proposed scheme exploits the multicast transmissions in a subnet. To do this, a new entity called 'feeder' is introduced. One of the receivers located in the subnet is dynamically configured as the feeder for the associated webcasting session. The feeder is a receiver and also plays a role of relaying the webcasting contents received from the remote sender to the other receivers in the subnet.

Figure 2 illustrates the unicast-only transport and the proposed delivery scheme. In the unicast-only transport, as shown in Figure 2(a), the sender has to send a data stream to each of the receivers by establishing multiple (replicated) unicast connections. In the proposed scheme, as depicted in Figure 2(b), only a single feeder receives the data stream from the remote sender, while the other receivers in the subnet receive the same data stream from the feeder, not the remote sender.

A new receiver that wishes to join a webcasting session first checks whether or not there is a feeder in the subnet. The first-joining receiver becomes the feeder for the webcasting session. The feeder establishes a unicast connection to the remote sender. The other late-joining receiver, after it realizes that there is a feeder in the subnet, will receive the streaming data from the feeder in the same subnet over the subnet multicast channel. The detailed operations related to the feeder configuration will be described in the next section.

Figure 3 compares the connection setup procedures for webcasting. In the unicast-only transport, as shown in Figure 3(a), each client downloads the requested contents from the media server, just after it gets information on the stream data via web access. In the unicast transport with subnet multicast, as illustrated in Figure 3(b), the first-arriving client becomes a feeder in the subnet by way of the feeder configuration process, and then establishes a unicast connection to the sender. The feeder relays the data stream to the other late-joining clients in the subnet, if any.

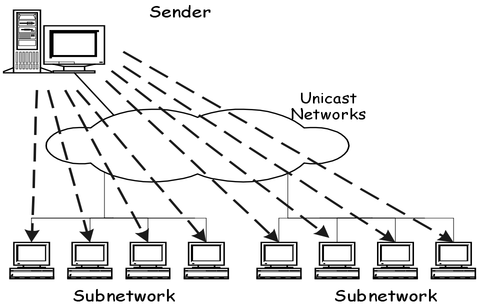

(a) Unicast-only

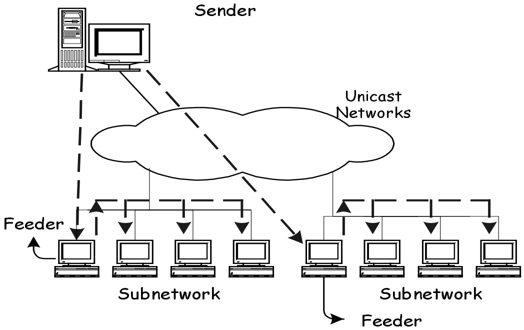

(b) Unicast with subnet multicast

Fig. 2. Transport schemes based on unicast-only and unicast with subnet multicast

4. Subnet Multicast Transport

In the proposed scheme, each new receiver checks first checks whether or not there is a feeder in the subnet, and the first-arriving receiver becomes a feeder in the subnet for the webcasting session. After that, the feeder forwards the data streams received from the remote sender toward the other receivers in the same subnet, if any. To do this, a feeder needs to be configured and announced to the other receivers in the subnet. The feeder reconfiguration is also required against the feeder release in the event that the existing feeder leaves the webcasting session.

For configuration and maintenance of a feeder, three kinds of control messages are employed: Feeder Solicitation (FS), Feeder Announcement (FA) and Feeder Release (FR). Those messages are summarized in Table 1.

All the control messages are delivered within the subnetwork via a well-known multicast address such as 224.0.0.x/24 or 224.0.1.x/24. Note that those addresses are reserved by IANA [5] for the multicast delivery of control messages in the subnetwork.

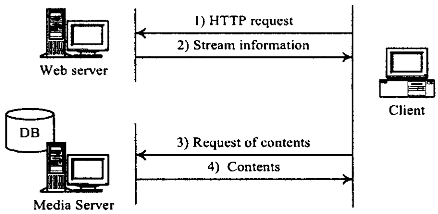

(a) Unicast-only transport

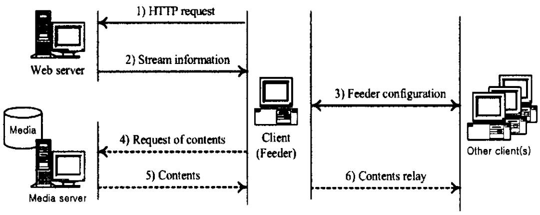

(b) Unicast transport with subnet multicast

Fig. 3. Connection setup procedures for webcasting

The data stream will be forwarded from the feeder to the other receivers by subnet multicast. For this purpose, a multicast (IP class D) address is required. For the scoping of the multicast traffic to the subnet, it is recommended to set the TTL (time to live) of the corresponding IP packets to ‘1’.

Table 1. Control Messages used for Subnet Multicast Transport

Type	Direction	Contains
FS	Receiver → 224.0.0.x	Sender’s address
FA	Feeder → 224.0.0.x	Sender’s address
FR	Feeder → 224.0.0.x	Feeder’s address Data forwarding address

4.1. Feeder Configuration

Before establishing a unicast connection to the remote sender, each new receiver sends an FS message to the subnet, and checks whether or not there exists a feeder for the webcasting session in the subnet, as shown in Figure 4. The FS message includes information on the IP address of the remote sender so as to indicate a webcasting session (see Table 1). If there is no response for the pre-configured FS time, (i.e., there is no feeder in the subnet), then the receiver becomes a feeder for the webcasting session. Otherwise, the receiver must receive an FA message from the existing feeder already configured in the subnet. In the example of Figure 4, the host F becomes a feeder after the FS timer expires. The feeder then connects to the remote sender and receives the data stream over the unicast connection.

If another receiver (host A in the figure) joins the webcast session by sending an FS message, then the feeder (host F) responds with an FA message to the host A. The FA message includes IP addresses for the sender, feeder and multicast data forwarding, as shown in Table 1. The feeder now begins to forward the data stream to the subnet. The other late-joining receivers such as host C will also receive the data stream from the feeder, after exchanging an FS message with the corresponding FA message. In case that there is no other receiver, the feeder does not need to forward the data stream.

4.2. Multicast Data Forwarding

If one or more other receivers are detected in the subnet, then the feeder begins the multicast forwarding of the data streams to the subnet. The destination address of the forwarded IP data packets will be set to the multicast address, and TTL (time to live) is set to '1' for restricting propagation to the subnet only.

During the data forwarding, the feeder continues to check whether there still exist the receivers in the subnet. If not, the feeder will not forward the data streams any more. For this purpose, each receiver sends periodic FS messages every FS time. In response to FS messages, the feeder sends the corresponding FA messages to the subnet receivers. If no FS message has arrived for some time, the feeder stops multicast forwarding of the data streams. Note that if a receiver has already sent an FS message, the other receivers in the subnet may cancel their FS messages so as to avoid implosion of the FS messages.

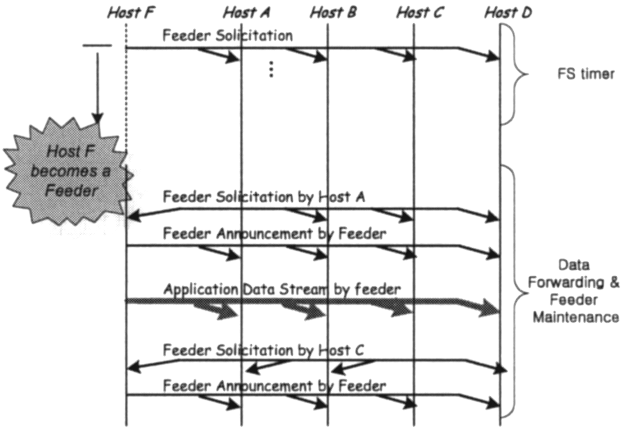

Fig. 4. Feeder Configuration and Data Forwarding

4.3. Feeder Reconfiguration

If the current feeder releases the webcasting connection to the remote sender, then a new feeder must be selected among the receivers in the subnet. To do this, the current feeder first sends an FR message to the receivers in the subnet. When an FR message arrives from the feeder, each receiver activates the 'FR timer'. One of the receivers will send the FA message, after its FR timer expires. The receiver that has first responded with an FA message will be configured as a new feeder.

The new feeder then established a unicast connection to the remote sender. Once a new feeder is configured, the other receivers cancel the FR timers. They now receive the data stream from the new feeder. Figure 5 depicts the feeder reconfiguration. In the figure, the host A is configured as a new feeder.

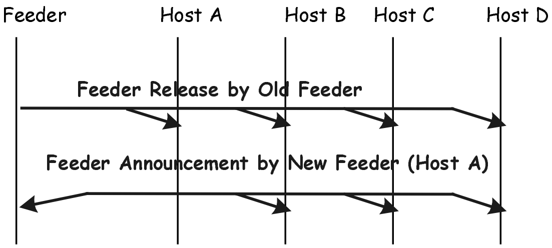

Fig. 5. Reconfiguration of a New Feeder

5. Numerical Results

Table 2 compares the proposed scheme with the unicast and multicast transports in terms of the amount of traffic generated at the sender and receiver sides. In the table, R and S represent the total number of the receivers and the number of the subnets participating in the webcasting session.

Table 2. Comparison of the schemes in terms of the traffic amount generated

	At Sender	In a Subnet
Unicast	$\propto$ R	$\propto$ R
Multicast	= 1	= 1
Proposed	$\propto$ S	$\leq$ 2

In the unicast-only transport, the traffic amount is proportional to the number of receivers at the sender and receiver sides. In IP multicasting, the traffic amount is fixed to one unit, regardless of the number of receivers. In the proposed scheme, the traffic amount at the sender depends on the number of subnets including the receivers.

The traffic amount generated in the subnet is reduced to one (for the subnet with only the feeder) or two units (for the subnet with the other receivers as well as the feeder).

From the table, we see that the proposed scheme provides a nearly same performance as the conventional IP multicasting in terms of traffic amount generated. In particular, the proposed scheme does not require any modification of the network infrastructure, and thus it can easily be deployed in the network. On the other hand, the IP multicasting requires the multicast-capable routers and multicasting routing protocols together with a reasonable control of multicast traffic in the network.

We implemented the proposed scheme on top of FreeBSD 4.1 Unix machines [6] in the form of a user-level library. The corresponding Application Programming Interface (API) was based on the Berkley socket Interfaces in the form of the wrapping function, which was employed by the each host receivers in the subnet. The API functional modules include the processing of the FS, FA, and FR messages, and the forwarding and reception of application stream over multicast data address.

The proposed scheme is compared with the unicast-only scheme. To do this, two subnetworks are configured, as shown in Figure 6. The number of receivers in a subnet is increased from one to nine. In the figure, the sender generates a random traffic stream. In Subnet A, after exchanging an FS with the corresponding FR message over a multicast address (224.0.0.18), a feeder is configured and it then begins to forward the data streams by subnet multicast. In Subnet B, the data stream is transferred over the unicast connections between the sender and all the receivers.

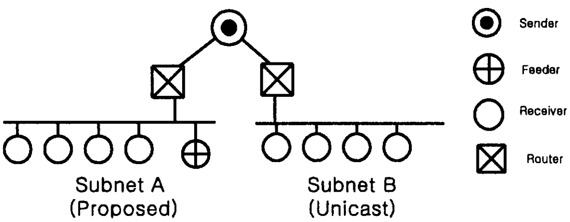

Fig. 6. Test Networks

Figure 7 shows the amount of traffic generated in the subnetworks with four receivers over a given time interval. On the average, the proposed scheme provides three or four times bandwidth utilization gains over the unicast-only scheme. Note that the sender temporarily stops transmitting data at the time of 40 and 50 seconds.

In Figure 8, the required network bandwidths for those two schemes are plotted for the different number of the receivers. For each test subnetwork, the number of receivers is increased by one, every 25 second, to nine. In the figure, we note that the proposed scheme requires a relatively fixed amount of bandwidth independently of the number of the receivers located in the subnetwork. On the other hand, the required network bandwidth in the unicast-only scheme increases linearly to the number of the receivers.

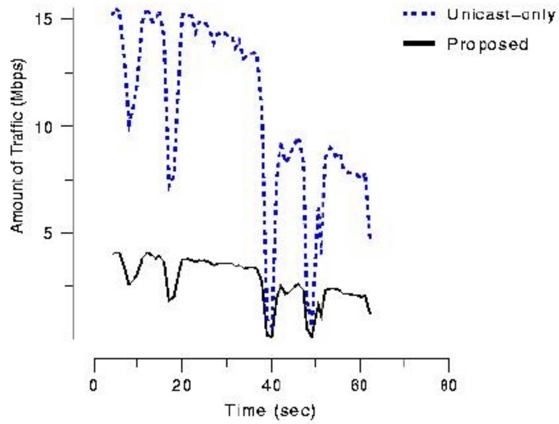

Fig. 7. Comparison in terms of traffic generated in the subnetwork

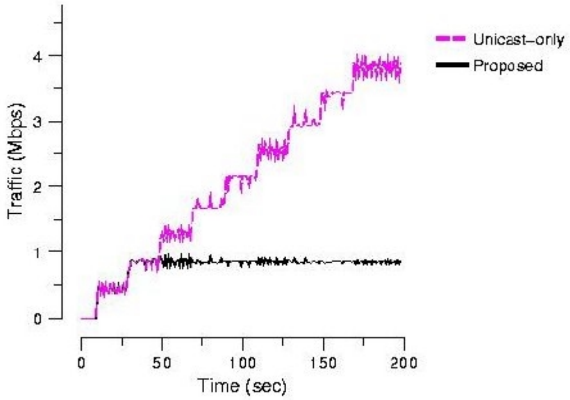

Fig. 8. Comparison of the required bandwidth as the number of receivers increases

Figure 9 plots the traffic generated by feeder(s), when the feeder reconfiguration occurs. In the figure, the feeder reconfigurations occurs at the time 50, 100, and 150 ms. In case of the first reconfiguration at the time 50 ms, the connection setup time between a new feeder and the sender elapses larger than the feeder reconfiguration time interval. At the time 100 ms, the connection setup time interval is smaller than the feeder reconfiguration interval. On the other hand, for the third reconfiguration, the connection setup time occurs exactly after the feeder reconfiguration, and thus neither data losses nor duplications are made.

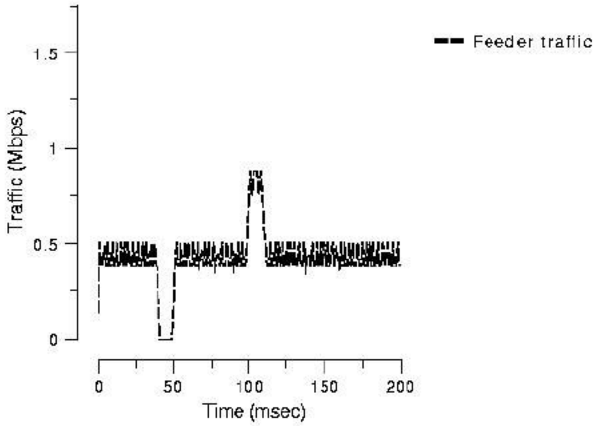

Fig. 9. Traffic discontinuity during feeder reconfigurations

6. Conclusions

Until now, we have discussed a new delivery scheme for webcasting, which is based on the unicast transport with subnet multicast. The proposed scheme has been tested and compared with the existing unicast-only transport in terms of the amount of traffic generated in the subnetworks. From the numerical and experimental results, we have seen that the proposed scheme provides a nearly same performance as the IP multicasting in the subnetwork environments.

It is noted that the proposed scheme does not require any change of the underlying network infrastructure for IP multicasting and thus can easily be deployed in the network. Recognizing that the IP multicasting still have much concern for the wide deployment, it is expected that the proposed scheme is used as an alternative short-term solution for one-to-many multicast application services.

References

1. Content Alliance, Available from http://www.content-peering.org/, 2001.
2. S. Koh, et al., "Minimizing Cost and Delay in Shared Multicast Trees", ETRI Journal, Vol. 22, No. 1, pp.30 - 37, March 2000.
3. C. Diot, et al., “Deployment Issues for the IP Multicast Service and Architecture”, IEEE Networks Magazine’s Special Issue on Multicast, January 2000.
4. J. Park, et al., "Multicast Delivery Based on Unicast and Subnet Multicast", IEEE Communications Letters, Vol. 5, No. 4, pp. 181 - 183, April 2001
5. IETF RFC 1700, Assigned Numbers, 1994.
6. W. R. Stevens, Unix Programming, 2nd Ed., Addison-Wesley, 1998.