논문 2004-41TC-8-4

SIP과 Xcast를 이용한 이동 멀티캐스트 지원 방안

(Multicast support for mobile nodes by using SIP and Xcast)

남 세 현*

(Seahyeon Nam)

요 약

본 논문에서는 이동 노드들 사이의 실시간 통신을 보다 효율적으로 제공하기 위한 새로운 멀티캐스트 방안을 제안한다. 제안된 방안에서는 SIP과 Xcast를 통합하여 이동 노드들에게 멀티캐스트를 제공한다. 시뮬레이션 결과를 통하여 제안된 방식은 불필요한 망의 트래픽을 줄일 수 있을 뿐만 아니라 패킷의 지연도 낮출 수 있음을 알 수 있다.

Abstract

A new multicast scheme for mobile nodes is proposed to support real-time communication in a more efficient way. In the proposed multicast scheme, the Xcast (a very flexible data plane mechanism) is integrated with the SIP (a very flexible control plane protocol) to support multicast for mobile nodes. The simulation results verify that the proposed scheme reduces unnecessary network traffic and achieves low latency of packets in the network.

Keywords : Mobile multicast, Explicit multicast, Session initiation protocol

Ⅰ. Introduction

Recently, mobility support for Internet access has created significant interest among researchers as wireless/mobile communications and networking proli-ferate, especially boosted by the widespread use of laptops and handheld devices. IP multicast, the ability to efficiently send data to a group of destinations, is becoming increasingly important for applications such as IP telephony and multimedia conferencing. Provi-ding multicast support for mobile nodes in an IP internetwork is a challenging problem and the solu-tions rely on the underlying mechanism of mobility support.

Currently, there are two basic approaches to support mobility in IP networks. The first one seeks to solve the mobility problem in the network layer by using Mobile IP[1] and related proposals. The multi-cast schemes over Mobile IP and related proposals include remote subscription, bi-directional tunneling, MoM[2], RBMoM[3], and so on. All of these schemes rely on two components: the host group model and the multicast routing protocol. In the host group model, a group of hosts is identified by a multicast group address, which is used both for subscriptions and forwarding. The multicast address allocation leads to quite complex procedures and introduces additional state information in the network. The multicast routing protocol is required to maintain the member state and the multicast delivery tree. Those traditional multicast schemes were designed to handle very large multicast groups. These work well if one is trying to distribute broadcast-like channels all around the world. However, they have scalability problems when there is very large number of small

* 정회원, 대구대학교 정보통신공학부
(School of Computer and Communication Engineering, Daegu University)
※ This paper was supported in part by the Daegu University Research Grant.
접수일자: 2004년4월13일, 수정완료일: 2004년8월2일

groups. In addition, for delay-sensitive multimedia applications, Mobile IP has some limitations, including triangle routing, triangle registration, encapsulation overhead, and the need for home addresses.

The other approach is to solve the mobility problem in the application layer by using Session Initiation Protocol (SIP)[4]. The SIP is an application layer protocol used for establishing and tearing down multimedia sessions, both unicast and multicast. It has been standardized within the Internet Engineering Task Force (IETF) for the invitation to multimedia conferences and Internet telephone calls. Three types of multiparty sessions can be supported by SIP: full mesh, mixer, and network-layer multicast. Both full mesh and mixer are not true multicast schemes. They deliver datagrams by using multi-unicasting. SIP should rely on the host group model and the traditional multicast routing protocols to support network-layer multicast. Thus, this solution shares the same disadvantages as the Mobile IP in providing multicast service for mobile nodes.

Explicit multicast (Xcast)[5] is a new scheme for Internet multicast that complements the traditional multicast schemes. In Xcast, the source node keeps track of the destinations in the multicast session and encodes the list of destinations in the Xcast header. Unlike traditional multicast schemes, Xcast does not specify a "control plane". It relies on neither IGMP nor multicast routing protocols to support multicast. With Xcast, the means by which multicast sessions are defined is an application level issue.

In this paper, a new multicast scheme for mobile nodes is proposed to support real-time communication in a more efficient way. In the proposed multicast scheme, the Xcast (a very flexible data plane mechanism) is integrated with the SIP (a very flexible control plane protocol) to support multicast for mobile nodes. Since both Xcast and SIP address super-sparse multicast sessions, it turns out that the Xcast can be easily integrated with SIP. Also, by using Xcast for multicast traffic, the proposed scheme can reduce unnecessary network traffic and achieve low latency of packets in the network.

The remainder of this paper is divided as follows. A brief description of the previous works is presented in Section Ⅱ. The proposed multicast scheme for mobile nodes is presented in Section Ⅲ. The discrete-event simulation model to evaluate the performance of the proposed scheme is presented in Section Ⅳ. The comparative simulation results and discussions are presented in Section Ⅴ. Finally, Section Ⅵ presents the conclusions.

Ⅱ. Previous works

1. SIP

The SIP is an application layer protocol used for establishing and tearing down multimedia sessions, both unicast and multicast. It has been standardized within the IETF for the invitation to multimedia conferences and Internet telephone calls. Entities in SIP are user agents, proxy servers, and redirect servers.

The SIP user agent has two basic functions: listening to the incoming SIP messages, and sending SIP messages upon user actions or incoming messages. The SIP proxy server relays SIP messages, so that it is possible to use a domain name to find a user, rather than knowing the IP address or name of the host. A SIP proxy can thereby also be used to hide the location of the user. On the other hand, the SIP redirect server returns the location of the host rather than relaying the SIP messages. This makes it possible to build highly scalable servers, since it only has to send back a response with the correct location, instead of participating in the whole transaction. The SIP redirect server has properties resembling those of the home agent (HA) in Mobile IP with route optimization, in that it tells the caller where to send the invitation. Both the redirect and proxy servers accept registrations from users, in which the current location of the user is given. The location can be stored either locally at the SIP server, or in a dedicated registrar.

One of the central tasks of SIP is to locate one or more IP addresses where a user can receive media

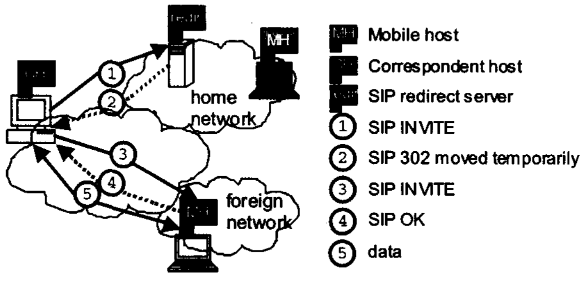

그림 1. SIP 호-전 이동성
Fig. 1. SIP pre-call mobility.

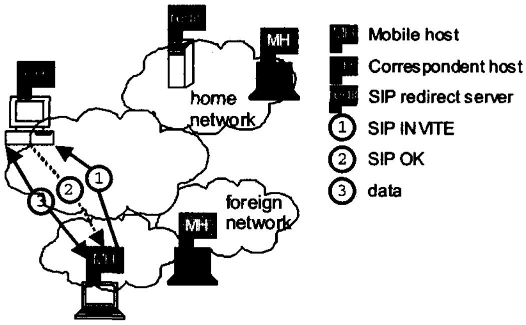

그림 2. SIP 호-중 이동성
Fig. 2. SIP mid-call mobility.

streams, given only a generic, location-independent address identifying a domain. This mechanism makes it easy to offer pre-call mobility. The mobile host (MH) simply re-registers with its home registrar each time it obtains a new IP address. When the correspondent host (CH) sends an INVITE message to the MH, the SIP server can redirect (or relay) the INVITE message (see Fig. 1).

If the MH moves during a session (mid-call mobility), the moving MH sends another INVITE request to the CH using the same call identifier as in the original call setup (see Fig. 2). It should put the new IP address in the "Contact" field of the SIP message, which tells the CH where it wants to receive future SIP messages. Finally, the MH should update its registration at the home SIP server, so that new calls can be correctly redirected.

Three types of multiparty sessions can be support-ed by SIP: full mesh, mixer, and network-layer multicast. In a full mesh, every participant establishes sessions with every other participant and sends an individual copy of the media stream to the others. In a mixer, participants either call into the mixer or are called by the mixer. A mixer or bridge takes several media streams and replicates them to all participants. Both full mesh and mixer are not true multicast schemes. They deliver datagrams by using multi-unicasting. Thus, they are inefficient in terms of network resources utilization and only scale to very small groups. In network-layer multicast, the initiator of a session simply invites the others to join the multicast session. Inviting a person to a multicast session is not different from any other invitation. However, before the multiparty session can be established, the initiator has to obtain a multicast address that will be used by the session. Multicast address allocation leads to quite complex schemes and introduces additional state information in the network. Also, SIP should rely on the traditional multicast routing protocols to support network-layer multicast. Thus, this solution is not suitable for highly mobile hosts due to the overhead associated with the resubscription and reconfiguration of the multicast delivery tree.

2. Xcast

The Xcast is a new scheme for Internet multicast that complements the traditional multicast schemes. Whereas the traditional multicast schemes can support a limited number of very large multicast sessions, Xcast can support a very large number of small multicast sessions. In traditional multicast schemes, the packet carries a multicast address as a logical identifier of all group members. However, in Xcast, the source node keeps track of the destinations in the multicast session that it wants to send packets to. The source node encodes the list of destinations in the Xcast header, and then sends the packet to a router. Each router along the way parses the header, partitions the destinations based on each destination's next hop, and forwards a packet with an appropriate Xcast header to each of the next hops. When there is only one destination left, the Xcast packet can be converted into a normal unicast packet, which can be unicasted along the remainder of the route.

Unlike traditional multicast schemes, Xcast does not specify a "control plane". There is no IGMP, and there are no intradomain or interdomain multicast routing protocols. With Xcast, the means by which multicast sessions are defined is an application level issue and applications are not confined to the model in which hosts use IGMP to join a multicast session. Thus, the application developer is not limited to the receiver-initiated joins of the IGMP model.

III. Proposed multicast scheme

For real-time communications such as IP telephony, multimedia conferencing, collaborative applications, and networked games, there are typically very large numbers of small multicast groups. Since the traditional multicast routing protocols impose limitations on the number of groups and the size of the network in which they are deployed, they have scalability problems. However, Xcast eliminates the membership management and routing information exchange in the intermediate routers, so that it can support very large numbers of small multicast groups. For real-time traffic, it is more common to use the Real-Time Transport Protocol (RTP) over UDP, and important issues are fast handoff, low latency, and high bandwidth utilization especially for wireless networks. Therefore, it is desirable to introduce mobility awareness on a higher layer, where we can utilize knowledge about the traffic to make decisions on how to handle mobility in different situations. Also, application-layer mobility does not require any changes to the operating system of any of the participants and thus can be deployed widely much easier than the Mobile IP.

In the proposed multicast scheme, the Xcast (a very flexible data plane mechanism) is integrated with the SIP (a very flexible control plane protocol) to support multicast for mobile nodes. Specifically, by using SIP, a full or partial mesh of RTP sessions is established to provide connectivity among multicast members. If a full mesh of sessions is established, then every hosts can send their messages to every other participants. Thus, every hosts can be a multicast sender of the group. This is useful for applications such as multimedia conferences and multi-player games. When a multicast application requires only one sender, a partial mesh that connects one sender to all receivers is established. By using SIP, a host takes the initiative to set up sessions for multicast. When a MH (multicast sender or receiver) moves during a session, we do not need to reconfigure the multicast delivery tree or rely on the tunneling service between HA and foreign agent (FA). With the assistance of SIP servers, sessions for multicast are created and maintained to support pre-call and mid-call mobility. The session states are kept in the hosts.

After a full or partial mesh of RTP sessions is established, Xcast is used to deliver identical RTP datagrams sent from a sender to multiple receivers. As long as the current SIP/SDP syntax and semantics are used, one has to rely on the UDP-enhanced version of Xcast with 6 bytes overhead (4 bytes for IP address and 2 bytes for UDP port number) for each destination. When an application decides to use Xcast forwarding, it does not affect its interface to the SIP agent and it can use the same SIP messages as it would for multi-unicasting. Since the application in the sender host keeps track of the participants' unicast addresses, it is a simple matter to replace multi-unicast code of SIP with Xcast code. All that the developer has to do is to replace a loop that sends a unicast to each of the participants by a single "Xcast_send" that sends the data to the participants. Thus, it is easy to incorporate the Xcast into the SIP.

IV. Simulation model

To assess the effects of integrating Xcast into SIP, a series of simulations have been performed. The simulations involve a network with a simple tree topology (Fig. 3). With a tree topology, only one shortest path exists between any two nodes, so routing issues do not affect the simulation results. The tree consists of 85 network nodes where each

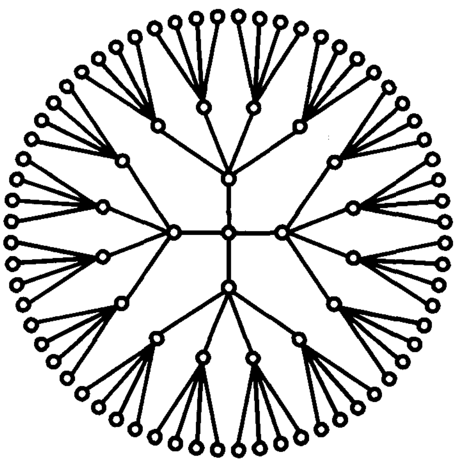

그림 3. 시뮬레이션에 사용된 망의 구조
Fig. 3. Network topology used in simulation.

internal node has four children. The 64 leaf nodes are considered as networks (or subnets) where stationary and mobile hosts are connected. In the simulation, it is assumed that a multicast sender is a stationary host located at randomly selected network and multicast receivers are all mobile hosts. At the beginning of a simulation, each MH is connected to its home network. After receiving 10 multicast packets in a network, the MH moves to other network with probability 0.1. The foreign networks to visit are chosen equiprobably at random. Thus, the residency time for each visit to a home or foreign network is geometrically distributed. When a MH moves to a foreign network, it can receive the multicast packets by mid-call mobility mechanism of SIP without the services of HA and FA. In the simulation, the multicast group size was varied from 1 to 20 and the multicast source node appends 6 bytes overhead per each destination to the Xcast header.

During a simulation run, one host on the network attempts to send a stream of data simultaneously to different destinations. Depending on the scenario being tested, the stream of data is sent using one of two distribution methods: multi-unicasting or xcasting. The data stream that is sent simultaneously to multiple destinations is a CBR stream of 200 byte UDP packets sent every 5 ms. The group of destina-tions to which the stream is sent does not change during the course of a simulation run. All links in the network have transfer rates of 2.048 Mbps and queues that can hold 50 packets. The simulation time for a run was selected as the time for a multicast sender to generate 20000 CBR packets for the multicast group.

V. Results

The simulation measures the number of packets successfully received in the network, the latency of those packets, and the number of packet transmi-ssions in the network links. Initially a single mobile destination for the data stream is randomly chosen, and the simulation is run. Then, an additional mobile destination for the stream is randomly selected, and the simulator is run again. This cycle continues until the stream is being sent to 20 mobile destinations. The measurements from ten such simulation sequen-ces are averaged together to produce the final results.

Figure 4 shows the total number of successfully received packets during the simulation. In the simulation scenarios, xcasting is able to deliver all of the packets sent into the network successfully. Since the amount of data sent into the network is approximately the same for the two different distribu-tion methods, if the multi-unicasting distribution method results in fewer successfully received packets than xcasting, the difference is due to lost packets. Xcasting clearly dominates in its ability to deliver the stream data to a large number of destinations. By contrast, unicasting can not handle the number of packets generated by a data stream sent to more than seven destinations. This sudden degrade of performance occurs because the majority of data stream packets must travel from the stream source to the center of the network in order to reach their destinations. If any link along this path becomes saturated, that link will act as a bottleneck and prevent any additional data stream packets from reaching their destinations.

Figure 5 shows the average latency of packets in the network. Again, xcasting exhibits better latency

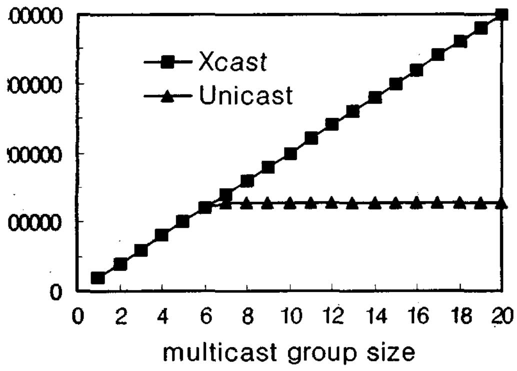

그림 4. 성공적으로 수신된 패킷의 수
Fig. 4. Number of packets successfully received.

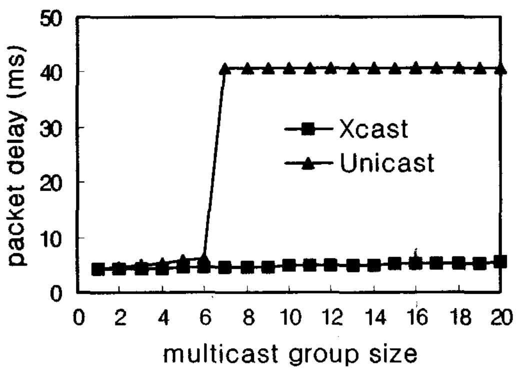

그림 5. 패킷 지연
Fig. 5. Packet delay.

performance than unicasting in situations where unicasting experiences saturated links. For xcasting, the delay increases slightly as the multicast group size increases. This is because a Xcast packet can have multiple destinations, so no redundant data travels over links in the network. However, for unicasting, the delay increases slightly at first, but it shoots up dramatically when a stream is sent to more than seven destinations. This is because the network queues overflow and the packets are discarded after the saturation point.

Figure 6 shows the number of packet transmi-ssions in the network links. For example, when the multicast group size is six, the multicast source generates 20000 packets during the simulation time and the six mobile hosts totally receive 120000 packets. If unicasting is used to deliver the multicast packets, about 651600 packet transmissions are required in the network links. However, if xcasting is used, about 322400 packet transmissions occur in the

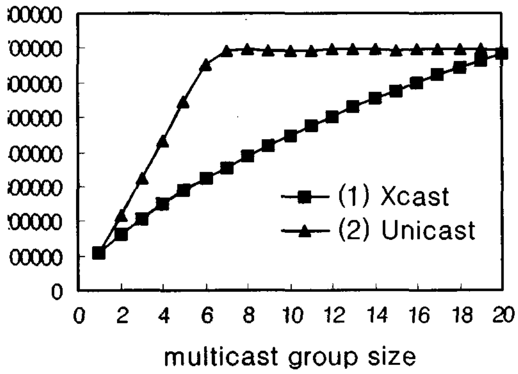

그림 6. 패킷 전송 수
Fig. 6. Number of packet transmissions.

network links to deliver the same packets. In Xcast, the number of packet transmissions in the network links increases sublinearly as the multicast group size increases. Also, the graph of unicasting shows similar behavior, but the increasing rate is much higher than the xcasting. In unicasting, the number of packet transmissions in the network links is saturated when a stream is sent to more than seven destinations.

VI. Conclusions

In this paper, a new multicast scheme for mobile nodes was proposed to support real-time communica-tion in a more efficient way. In the proposed multi-cast scheme, the Xcast (a very flexible data plane mechanism) is integrated with the SIP (a very flexible control plane protocol) to support multicast for mobile nodes. By using SIP, a full or partial mesh of RTP sessions is established to provide connecti-vity among multicast members. After establishing connections for multicast, Xcast is used to deliver identical RTP datagrams sent from a sender to multiple receivers. The proposed multicast scheme complements the existing multicast schemes for mobile nodes in that it can efficiently support very large numbers of distinct (small) multicast groups and thus can play an important role in making multicast practical for real-time applications. Through the simulation study, it is verified that the proposed scheme reduces unnecessary network traffic and achieves low latency of packets in the network.

Rererences

[1] C. Perkins, "IP mobility support," RFC 2002 (Proposed Standard), IETF, Oct. 1996.

[2] T. G. Harrison, C. L. Williamson, W. L. Mackrell, and R. B. Bunt, "Mobile Multicast (MoM) Protocol: Multicast Support for Mobile Hosts," *Proc. of ACM/IEEE Mobicom'97*, pp. 151-160, Sept. 1997.

[3] R. Lin and K. M. Wang, "Mobile Multicast Support in IP Networks," *Proc. of IEEE Infocom 2000*, pp. 1664-1672, March 2000.

[4] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg, "SIP: session initiation protocol," RFC 2543 (Proposed Standard), IETF, March 1999.

[5] R. Boivie, Y. Imai, W. Livens, D. Ooms, and O. Paridaens, "Explicit Multicast Basic Specifica tion," draft-ooms-xcast-basic-spec-04. txt, Jan. 2003.

저 자 소 개

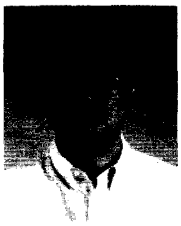

Seahyeon Nam(정회원)

He received the B. S. degree in electronics engineering from Yonsei University, in 1985, and the M. S. and Ph. D. degrees in electrical engineering from the KAIST in 1987 and 1991, respectively. Since 1994, he has been with the School of Computer and Communication Engineering at Daegu University, where he is currently an Associate Professor. His current research interests are multicast protocols, lightwave networks, and queueing theory.