Explicit Multicast Extension (Xcast+) for Efficient Multicast Packet Delivery

Myung-Ki Shin, Yong-Jin Kim, Ki-Shik Park, and Sang-Ha Kim

ABSTRACT— In this letter, we propose a new multicast scheme, named Xcast+, which is an extension of Explicit Multicast (Xcast) for an efficient delivery of multicast packets. The mechanism incorporates the host group model and a new control plane into existing Xcast, and not only does it provide the transparency of traditional multicast schemes to senders and receivers, but it also enhances the routing efficiency in networks. Since intermediate routers do not have to maintain any multicast states, it results in a more efficient and scalable mechanism to deliver multicast packets. Our simulation results show distinct performance improvements of our approach compared to Xcast, particularly as the number of receivers in a subnet increases.

I. INTRODUCTION

IP multicast [1], the ability to efficiently send data to a group of destinations, is becoming increasingly important for applications such as IP telephony and video-conferencing. However, while traditional multicast schemes are scalable in the sense that they can support very large multicast groups, there are scalability issues when a network needs to support a very large number of distinct multicast groups. Explicit Multicast (Xcast) [2]-[5] is a newly proposed multicast scheme to support a very large number of small multicast groups. Xcast uses explicit encoding of a destination list in the data packets instead of multicast address. Xcast scheme has a number of advantages. For example, it can save bandwidth between routers similar to traditional multicast schemes, even if the routers do not need to maintain states. However, it may suffer from a scalability problem as the number of receivers in a subnet increases, since it is applicable only to a small group. This paper proposes a new multicast scheme, named Xcast+ (Explicit Multicast Extension), which is an extension of Xcast for an efficient delivery of multicast packets. The mechanism provides an enhanced scheme supporting host group model in Xcast scheme. This is achieved by adding an IGMP (S, G) join at receivers' side and sending the join request through Registration Request Message toward a sender, and by explicitly encoding the addresses of the Designated Routers (DRs) at the receivers' side instead of receiver addresses. Xcast+ is not intended to replace the existing multicast schemes. Moreover, it complements the existing multicast schemes.

Manuscript received September 17, 2001; revised October 23, 2001.

Myung-Ki Shin (phone: +82 42 860 4847, e-mail: mkshin@pec.etri.re.kr), Yong-Jin Kim (e-mail:yjkim@pec.etri.re.kr), and Ki-Shik Park (e-mail: kipark@pec.etri.re.kr) are with the Protocol Engineering Center, ETRI, Daejeon, Korea.

Sang-Ha Kim (e-mail:shkim@cclab.cnu.ac.kr) is with the ChungNam National University in Daejeon, Korea.

II. XCAST+ SCHEME

In Xcast+, a receiver initiates IGMP (S, G) join. When a Designated Router (DR) at receivers' side receives the report, it sends Xcast+ Registration Request Message containing sender address, S, group address, G, and its own address, DR, toward the sender. These procedures imply the addition of control plane in Xcast. Thus, when the DR at the sender's side receives the message, it can keep track of the addresses of all DRs at the receivers' side involved in the multicast session (S, G) in its cache table. When the sender sends multicast packets, a DR at the senders' side which receives the multicast packets explicitly encodes the addresses of the DRs at the receivers' side in Xcast header, and sends the packets as Xcast+ packets on unicast path (M2X: Multicast to Xcast). Each router along the way parses the header partitions the destinations based on each destination's next hop and forwards the packet with an appropriate Xcast header to each of the next hops. These procedures comply with the data plane for existing Xcast, except for encoding addresses of the DRs at the receivers' side in the data packets instead of addresses of receivers. When the DRs at the receivers' side receive the Xcast+ packets, they send the packets as standard multicast packets to receivers (X2M: Xcast to Multicast).

For example, suppose that B, C, D, E, F and H are trying to receive multicast packets (G) distributed from A in Fig. 1. This is accomplished by B, C, D, E, F and H initiating IGMP (A, G) join. When R4, R8 and R9 receive the request, they send a Registration Request Message toward the sender. Meanwhile, when R1 receives the message, it sends a Registration Reply Message (A message acknowledging a safe delivery of the Registration Request) and does not forward this message to the sender. Therefore, R1 identifies a set of all DRs at the receivers' side dynamically <R4, R8 and R9>. Thus, when R1 receives multicast packets (G) from A, R1, by M2X processing algorithm, it sends the packets as Xcast+ packets with the list of <R4, R8 and R9> in its data packets to the next hop router, R2. Xcast+ encoding of the destination list in IPv4 and IPv6 are the same as Xcast [2]. Therefore, ignoring details, the packet that R1 sends to R2 looks like this: [src = R1 | dest = R4 R8 R9 | payload].

When R2 receives this packet, it needs to properly process the Xcast+ header. The processing that a router does upon receiving one of these Xcast+ packets is exactly the same as in Xcast except for the X2M processing that occurs when an Xcast+ packet is received at a DR processing. Therefore, in the example above, R3 will send one copy of packet to destination R5 with Xcast+ list of <R8 and R9> and one copy of packet to destination R4 with Xcast+ list of <R4>. When R4 receives the packet, it will, by the X2M processing algorithm, send the packet as a standard multicast packet to the receivers <B and C>.

III. XCAST+ COST EFFECTIVENESS

In the Xcast+ scheme, there are a few extensions to Xcast. They are the use of encoded addresses of DRs in the data packets instead of addresses of receivers, and the use of X2M and M2X at the DRs. These extensions to Xcast bring following benefits:

- From the viewpoint of the receivers, procedures in the control plane are the same as existing ASM (Any-Source Multicast) and SSM (Source-Specific Multicast). Therefore, Xcast+ receivers do not need to use an additional control to join in a session. This means that the control plane of Xcast+ is compatible with the existing ASM and SSM. A receiver that is an IGMP capable host does not need to be an Xcast+ capable host.
- Similar to ALM (Application Level Multicast) and Overlay Multicast schemes, Xcast+ supports both multicast and unicast, where multicast is used within a subnet and unicast is used between routers. Therefore, intermediate routers do not have to maintain multicast states, but there still exists benefits of multicasting.

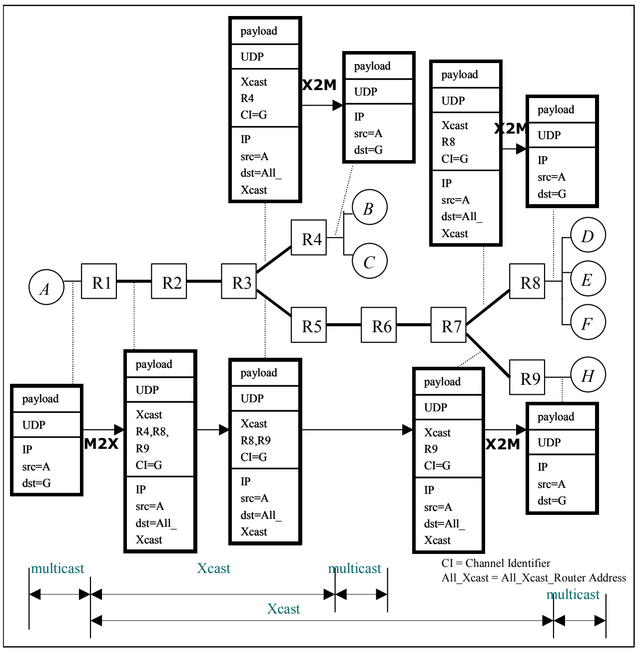

Fig. 1. Example of multicast packet delivery in Xcast+.

- There can be an increase of the number of receivers in a subnet, which means Xcast+ can support larger number of members compared to that of existing Xcast. Whereas Xcast can support a very large number of small multicast groups, Xcast+ can support a very large number of medium size multicast groups.
- When the scalability in the ASM scheme is considered, one of the main issues may be a complexity of multicast tree construction between multicast routers on the Internet backbone. Because Xcast+ uses the multicast scheme in a subnet (the use of the multicast address is so simple in all other cases), deployment and management are easy and simple even if a multicast scheme is used.

We evaluate the performance of Xcast+ by comparing it with Xcast. Assuming that the cost of Xcast header encoding per destination is "1" and the cost of Xcast header decoding per destination is "1", we employ a real network topology, which is the map of major MBone routers [6]. We ignore additional control overheads. Figure 2 shows the results of evaluation on extra header processing overhead in Xcast+ and Xcast on the MBone topology [6]. We assume that there are no receivers in intermediate nodes, which means edge nodes that have only one link can have their receivers. Among 82 nodes, the number of DRs is 35 and the number of intermediate routers is 47. In order to show distinct performance advantages of our approach compared to existing Xcast, particularly as the number of receivers in a subnet increases, we run every experiment when

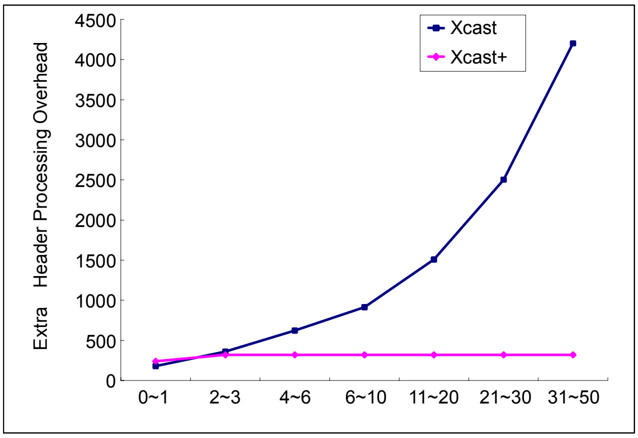

Fig. 2. Results of evaluation on extra header processing overhead in Xcast+ and Xcast based on MBone Topology [6].

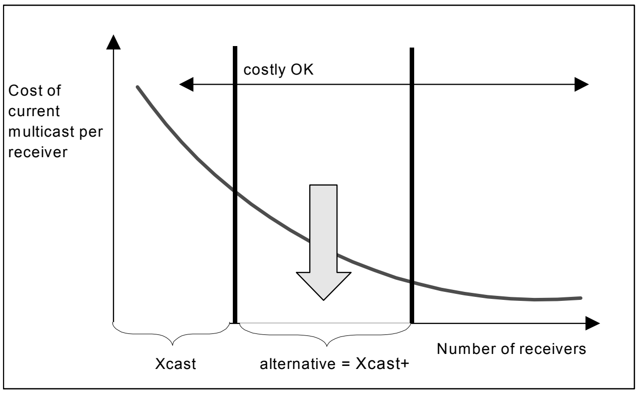

Fig. 3. Xcast+ cost effectiveness as the number of receivers in a subnet increases.

the number of receivers in subnets randomly increases (0~1, 2~3, 4~6, 6~10, 11~20, 21~30, and 31~50). We observed that Xcast+ gets more profit than Xcast as the number of receivers in a subnet increases. In fact, in an example of MBone topology [6], Xcast cannot be used because of excessive Xcast packet processing overheads when there are 50 receivers in a subnet. However, Xcast+ can endure the Xcast packet processing overheads because there are at most 35 DRs.

Due to the cost of multicast address allocation, multicast routing state management, control overhead, and scalability issues of traditional IP multicast scheme, ASM leads to a search for other multicast schemes. SSM avoids a multicast address allocation. However, SSM still creates state and signaling per multicast channel in each on-tree node. Both ASM and SSM become expensive for its members if the groups are small. These applications are well-served by Xcast-like schemes [2]. Since Xcast+ incorporates the host group model into Xcast, it can support an efficient multicast delivery for medium size multicast groups. Figure 3 depicts the costs as a function of the number of receivers in the session or channel. Table 1 shows cost analysis of ASM, SSM, Xcast and Xcast+ schemes. As a

Table 1. Cost Analysis of ASM, SSM, Xcast and Xcast+.

	ASM	SSM	Xcast	Xcast+
Multicast Address Allocation	H	M	NA	M
Multicast Routing State Management	H	H/M	L	L
Control Overhead	H	H/M	N	M
Overhead by Increase of Receivers	L	L	H	L
Extra Header Processing Overhead	L	L	H	H/M
Deployment	H	M	L	L

H: High, M : Medium, L : Low or none and NA : Not applicable

result of analysis, while Xcast+ has some control plane overheads compared to Xcast, its cost of extra header processing can be saved as the number of receivers increases in a subnet. For examples, whereas the cost of Xcast and Xcast+ is the same if there is one receiver in a subnet, theoretically the cost of Xcast+ is reduced by n times than that of Xcast when the number of receivers increases to n (assuming that overheads of control plane are ignored).

IV. CONCLUSIONS

This paper describes a new multicast scheme, called Xcast+ that combines the advantages of Xcast with the strengths of traditional multicast schemes. Like Xcast, Xcast+ avoids the use of per flow state in the core of the network and thus Xcast+ can support an unlimited number of distinct multicast groups. And like traditional multicast schemes, Xcast+ can support (relatively) large multicast groups. We believe the advantages of Xcast+ will be important in mobile networks and in heterogeneous QoS environments.

REFERENCES

[1] S. Deering, *Multicast Routing in a Datagram Internetwork*, Stanford University, Ph.D. thesis, Dec. 1991.

[2] R. Boivie *et al.*, *Explicit Multicast (Xcast) Basic Specification*, IETF Internet-Draft, http://www.ietf.org/internet-drafts/draft-ooms-xcast-basic-spec-01.txt, 2001.

[3] R. Boivie, *Small Group Multicast*, IETF Internet-Draft, http://www.ietf.org/internet-drafts/draft-boivie-sgm-01.txt, 2000.

[4] D. Ooms, *Connectionless Multicast*, IETF Internet-Draft, http://www.alcatel.com/xcast/draft-ooms-cl-multicast-02.txt, 2000.

[5] Y. Imai, *Multiple Destination Option on IPv6 (MDO6),* IETF Internet-Draft, http://www.ietf.org/internet-drafts/draft-imai-mdo6-02.txt, 2000.

[6] S. Casner, *Major MBONE Routers and Links,* ftp://ftp.isi.edu: mbone/mbone.topology.ps, 1994.