The Active Traffic Control Mechanism for Layered Multimedia Multicast in Active Network

Saehoon Kang, Hee Yong Youn, Younghee Lee, Dongman Lee, and Myoungchul Kim
School of Engineering
Information and Communications University
Taejon, Korea
{kang, youn, yhlee, dlee, mckim}@icu.ac.kr

Abstract

In multicasting multimedia data, effectively adapting to heterogeneous receivers is very difficult. In this paper we propose an active traffic control mechanism for layered multimedia multicast (ATLM) in active network environment. The proposed scheme controls traffic at each router using the traffic condition of immediate children nodes. Also the traffic adjustment is much finer than existing multicast protocols of a granularity of one layer. Computer simulation reveals that the proposed scheme significantly improves the amount of delivered traffic and end-to-end delay as much as about 10% compared to layer-wise adjustment. Moreover, it does not require to maintain multiple sessions for transferring multiple layer data but only one session, and thus reduces the overhead on session management and network bandwidth.

Keywords: active network, congestion, layered multicast, multimedia, traffic control.

1. Introduction

The Internet applications have caused tremendous increase on Internet users. Among various Internet applications, multimedia systems such as multimedia conferencing and distance learning are of great importance recently. They require large network bandwidth for real-time multicast, and efficient adjustment of network operation is inevitable for supporting heterogeneous receivers. This is because the network condition significantly fluctuates in terms of available bandwidth.

The existing multicast protocols mostly consider the cost of end-to-end path and support best-effort services. Therefore they cannot accommodate different requirements of heterogeneous receivers. Recently, various mechanisms have been proposed for adjusting the transmission rate of senders in accordance with the network congestion condition [5,6]. Due to the heterogeneity of the Internet, a single transmission rate from a sender cannot satisfy the conflicting bandwidth requirements at different sites. Consequently, the transmission rate is usually decided according to the receiver with the smallest bandwidth. This results in that quality of the data received at other sites is unnecessarily degraded.

The limitation was proposed to be overcome using layered transmission mechanisms [2,3]. These proposals are based on hierarchical or layered encoding technique by which multimedia data is encoded into multiple layers - base layer and enhancement layers. The encoded data are transmitted by a sender, and each receiver decides how many layers it accepts depending on its capability or desired level of QoS. However, this approach is not very efficient since the end receivers do not have information on the entire network traffic condition. Also, if there exist more than one sender, the receivers need to react differently from sender to sender. In addition, the QoS of a receiver changes with a granularity of one layer as receivers dynamically join and leave a multicast session carrying different layers.

In this paper we propose a new scheme called active traffic control mechanism for layered multimedia multicast (ATLM) for enhancing the layered transmission schemes. This is done by dynamically controlling the transmission rate at each router according to the network congestion status that is provided by immediate children nodes. The mechanism is implemented in the routers of active network [4,8-11]. The active network concept was proposed to provide a network with programmability and flexibility by allowing applications to inject customized programs into the network nodes. With the proposed ATLM, each active router adjusts the filtering rate starting from the highest layer (the least significant layer). Due to the two distinguishing features of the ATLM, traffic adjustment at each router and fine adjustment of filtering rate instead of layer-wise dropping or adding, the proposed scheme allows significant performance improvement on earlier layered multicast protocols. Computer simulation using a GUI-based simulation tool, SES/Workbench [14,15], reveals that the proposed scheme allows about 10% improvement in the number of packets delivered and end-to-end delay. Moreover, the receivers do not need to repeatedly join and leave each session separately as in earlier layered multicast. They need to join only one multicast session carrying all the layers. This significantly reduces the overhead on the session management and network bandwidth.

The rest of the paper is organized as follows. Section 2 discusses the related works on layered multicast, which motivates the proposed mechanism. Section 3 introduces the proposed ATLM, and it is evaluated in Section 4. Section 5 concludes the paper with some remarks.

2. Related Works

In this section layered transmission is discussed first, and then active network is presented.

2.1. Layered Transmission

In the packet switched networks like the Internet, the constituent network varies in terms of bandwidth and load factors. In such heterogeneous environment, a single transmission rate decided according to the receiver with the lowest bandwidth cannot satisfy the conflicting bandwidth requirements of different sites. An approach for accommodating the network heterogeneity is to transmit data stream in different layers. The base layer provides a minimal amount of data needed for an acceptable representation of the original data stream. Each successive higher layer enhances the QoS.

Figure 1 depicts an example of a multi-layered transmission approach. The sender transmits the data split into three streams over three layers. Since a sufficient bandwidth is available on the link between the upper router and receiver *R1*, *R1* decides to subscribe to all the three layers. Since the link between the two routers has low capacity, only two lower layers of the three can be forwarded. The bandwidth on the link between the lower router and receiver *R3* is even more restricted, and thus *R3* decides to receive only the base layer.

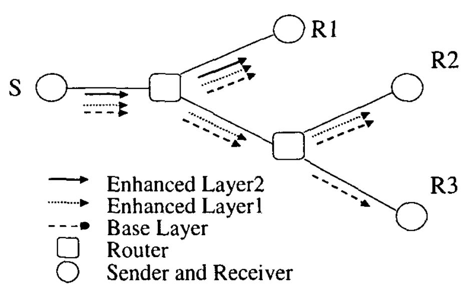

Figure 1. An example of multi-layer transmission.

The base layer is separately decoded, and it provides basic level of quality. The enhancement layers are decoded together with the base layer, and they provide improvement on the quality of the received data. Therefore packet loss in those layers does not seriously affect the quality of lower layer(s). Such layered transmission is further classified into two mechanisms, cumulative and independent layered data transmission.

In cumulative layered data transmission [2,3], each successive higher layer provides refined information on the lower layers. Therefore receivers need to listen to all the lower layers up to the highest one it wants to listen to. For example, if a receiver wants to listen to layer-3, it also has to listen to layer-1 and 2. The example presented in Figure 1 describes such an approach. This approach, however, suffers from drift and resynchronization problem. Three variations of the cumulative layered multicast have been proposed - receiver-driven layered multicast (RLM [2]), layered video multicast with retransmission (LVMR [3]), and multi-session rate control.

In RLM, a sender sends each video layer to a separate IP multicast group, and takes no active role in rate adaptation. Each receiver subscribes to several layers by joining the corresponding IP multicast groups, and the receivers drop the highest layer on congestion. On the contrary, it adds a layer when it has spare capacity. By repeating the process, each receiver finds an optimal level of subscription of video layers.

Even with no packet loss, the level of subscription is not necessarily low. To find out if the level is too low, thus, '*join-experiment*' is carried out. If congestion condition is detected by the experiment, the receiver immediately drops the newly added layer. Otherwise, the layer is kept. The failing join-experiment results in degraded video quality to both the receiver that carried out the experiment and other receivers sharing the congested link. In order to resolve this problem, a learning algorithm called '*shared learning*' is employed. With shared learning, instead of independent rate adjustment in each receiver, all the receivers are informed of the result of the join-experiment.

The idea of shared learning, although it improves scalability and interference problem, requires each receiver to maintain a variety of state information which it may or may not need. In addition, the use of multicasting required for exchanging control information may result in reduced bandwidth on low-speed links.

The layered video multicast with retransmission (LVMR) is another scheme for distributing video using layered coding over the Internet. The two key contributions of the scheme are:

- Improving the quality of reception within each layer by retransmitting the lost packets given an upper bound on recovery time and applying an adaptive playback point scheme to achieve more successful retransmission

- Adapting to network congestion and heterogeneity using hierarchical rate control mechanism

In contrast to the existing sender-based and receiver-based rate control in which the entire information on network congestion is available either at the sender or replicated at the receivers, the hierarchical rate control (HRC) mechanism distributes the information between the sender, receivers, and some agents in the network in such a way that each entity maintains only the information relevant to itself. In addition to that, the hierarchical approach enables intelligent decisions to be made in terms of conducting concurrent experiments and choosing one of several possible experiments at any instant of time based on minimal state information at the agents in the network.

In contrast with RLM, the HRC approach employs a hierarchical dynamic rate control scheme in each receiver so as to allow receivers to maintain minimal state information and decrease control traffic in the multicast session. To avoid the shortcomings of RLM, it intelligently partitions the knowledge base and distributes relevant information on the members in an efficient way. In LVMR, all the experiment results are compiled into a knowledge base that would represent the comprehensive group knowledge. LVMR also provides the schemes for partitioning the group knowledge base.

A problem with RLM and LVMR is that the protocols do not provide fair bandwidth sharing between competing video sessions or between video sessions and TCP sessions. An end-to-end control scheme was thus proposed in [7] for layer-based congestion sensitivity rate control, which improves inter-session fairness if it is used to augment video multicast protocols. The basic idea is to let higher video layers have high sensitivity to congestion. This will cause receivers to drop high layers more easily when they compete for bandwidth with the receivers receiving low layers only. It results in that the competing receivers end up with the same number of layers.

Another approach for layered transmission is to simply transmit the same data stream encoded with different quality levels for different multicast sessions [1]. This scheme is often called 'simulcast' because the source transmits multiple copies of the same signal simultaneously at different rates which results in different qualities. As the streams contain all necessary information for decompression, the receivers need to join only one multicast session. This approach avoids the resynchronization problem experienced in the layered multicast approach. However, this is achieved at the cost of sending multiple replicated streams, and thus possibility of network congestion.

We next discuss active network, which was proposed for more effective network control and management.

2.2. Active Networks

Active network is a novel approach to network architecture allowing applications to inject customized programs into the network nodes and the network to support customized services to the applications. For example, a user of an active network could send a trace program to each router, and arrange the program to be executed when their packets are processed. Figure 2 illustrates how the routers of IP network could be augmented to perform such customized processing on the datagrams flowing through them. Active networks improve the network flexibility and functionality by introducing programmability to the routers. This enables faster protocol innovation by making deployment of new network protocols easier, even over wide area. However, improvement on flexibility and functionality of the networks conflicts with the safety and security requirements. To overcome these problems, a variety of approaches such as SANE [10] and PLAN [11] are being experimented.

Active networks support dynamic control of network behavior. A variety of new schemes such as application–specific congestion control, caching [12] and active reliable multicast (ARM) [13] have been proposed for active network. In [12], intelligent discard and small network level caches using active scheme were introduced. A novel loss recovery scheme for large-scale reliable multicast was proposed in [13]. ARM utilizes soft-state storage within the network to improve the performance and scalability. In the upstream direction, the routers suppress duplicate NACKs from multiple receivers to control the implosion problem. By suppressing duplicate NACKs, ARM also lessens the traffic that propagates backward through the network. In the downstream

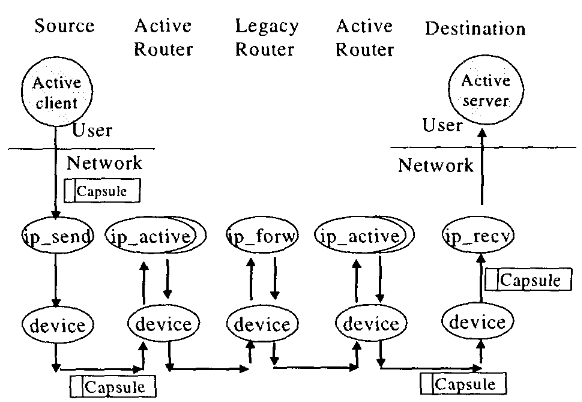

Figure 2. Application-specific processing within the nodes of an active network.

direction, the routers limit the delivery of repair packets to the receivers experiencing losses, thereby reducing network bandwidth consumption. Finally, in order to reduce wide-area recovery latency and distribute the retransmission load, the routers cache multicast data on a *'best-effort'* basis.

Therefore, active network accommodates various requirements of heterogeneous receivers at network level by running appropriate programs injected into the routers. The active routers can actively exchange the network information with each other, reroute a path of a packet, and filter packets out using the layer information in the layered multicast when the network is congested. We next present the proposed scheme based on active network.

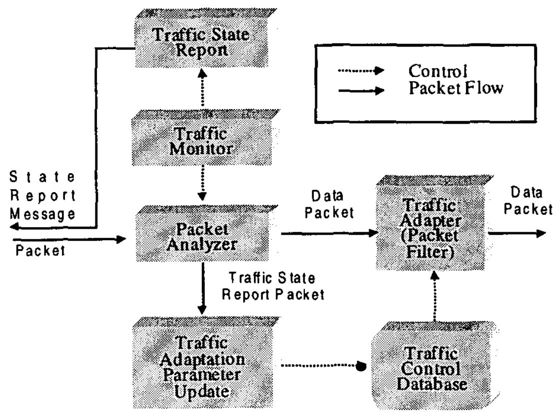

Figure 3. ATLM traffic control mechanism with active routers.

3. The Active Traffic Control for Layered Multimedia Multicast (ATLM)

In this section the proposed active traffic control scheme for layered multimedia multicast is introduced. The basic idea is presented first.

3.1. Basic Idea

As mentioned earlier, the granularity of traffic adjustment in existing layered multicast transmissions such as RLM is large. They adjust QoS with a granularity of one layer. This might cause large QoS changes at the receivers and add high load to the receivers. Moreover, the same number of sessions as the number of layers are needed for which the data are encoded, and repetitive joining and leaving the sessions are required. This may add extra load to the networks for managing the session information. Another problem with the schemes is resynchronization. As different layers might take different routes to a receiver, the transmission delays become quite different.

We thus propose an active traffic control for layered multimedia multicast (ATLM) scheme as an enhancement of layered transmission schemes such as RLM. In ATLM, only one multicast session carrying all the layered data is used. The structure of ATLM resembles the multi-core based tree. For example, the active routers in ATLM are similar to the core routers in multi-core based tree.

With ATLM, each active router or end receiver in the network monitors loss rate of incoming packets on each link. If the packet loss rate exceeds a designated value, they notify their immediate active parent nodes (they can be active routers or senders) of the congestion condition. The parent active node then reduces the transmission rate of the link from which the notification message arrives. The change in the transmission rate at each link does not affect the condition of other links. It does not lead to any significant changes in the QoS at the receivers, either.

The ATLM scheme consists of two major parts: *traffic monitoring* and *adaptation*. Figure 3 depicts the overall structure of the proposed ATLM.

3.2. Network Model

When network congestion occurs, the traffic control mechanism implemented in active router causes it to reduce the amount of traffic by increasing the filtering rate of the highest layer. If the filtering rate reaches 100% (hence no packet to that layer), then the filtering rate of the next higher layer is increased. In contrast, when the network has extra bandwidth, active router increases the amount of traffic by decreasing the filtering rate of the lowest layer among the filtered layers. If the filtering rate of that layer reaches 0% (hence no filtering), the filtering rate of the next layer is decreased. This process continues until the network reaches steady state. By applying this mechanism, the network traffic can be controlled at network level, and the receivers do not need to repeatedly join and leave each session separately as in earlier layered multicast. It is enough to join only one multicast session carrying all the layers. This will significantly reduce the overhead on the session management and network bandwidth.

In implementing an active network, it is unnecessary to convert all the routers into active routers. Some routers can be active routers while others are non-active ones (i.e., legacy routers). When we assume that active router substitutes core router in multi-core based tree, we can easily apply active network concept to the legacy network. The role of core router is very important and almost all multicast packets pass through it. Eventually the location of core router becomes a critical point in term of network traffic control, and we can exploit the advantages of active network. Unlike multi-core based tree, however, we assume tunneling between active nodes (i.e., active router or active end receiver) although all the connections between the nodes (whether they are active or inactive) are considered as the links used for multicast. This means that the tunnel is regarded as one link in terms of traffic control, and an active router knows immediately adjacent active routers and exchange messages with them.

3.3. Traffic Monitoring

When a packet arrives at an active node, the '*Packet Analyzer*' analyzes it. If it is a traffic state report packet, the active router updates '*Traffic Adaptation Parameters*'. If it is a data packet, it is transferred to '*Traffic Adapter*' and filtered according to the traffic condition of the downstream links, which is maintained in the '*Traffic Control Database*'.

Active node also monitors the incoming traffic and decides the packet loss rate. It also sends a traffic state report message to the parent active router. If the packet loss rate exceeds the upper threshold, t_u, the active node regards the link as being congested. Then the active node sends a message to its parent node in the multicast tree. The message indicates the congestion condition, and thus the parent node reduces the amount of traffic directed to the link from which the message arrives. In contrast, if the packet loss rate is lower than the lower threshold, t_l, then the active node sends a message notifying that it is in unloaded state.

3.4. Traffic Adaptation

When an active router receives a traffic state message from its child active router, it updates its traffic adaptation parameters which determine the transmission rate on each respective output link. If the traffic state message indicates congestion, the active router increases the filtering rate by the increment unit, $\Delta+$. In contrast, when the traffic state report message indicates unloaded state, the active router increases the amount of traffic by decreasing the filtering rate by the decrement unit, $\Delta-$. Since the network traffic changes continuously, the repetitive adjustment process goes on infinitely. We usually set larger $\Delta-$ value than $\Delta+$ value so that congestion can be dissolved more quickly. The proposed ATLM scheme is evaluated by computer simulation next.

4. Simulation

The sample network is introduced first. After that, the simulation tool and simulation results are presented.

4.1. Sample Network

The ATLM scheme is evaluated for a test network of Figure 4, which consists of one sender, four receivers, and three active routers.

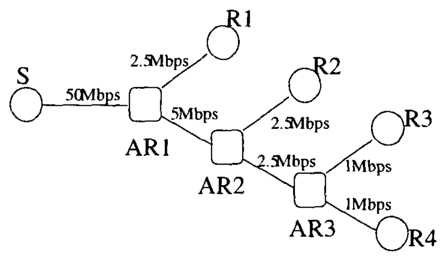

Figure 4. The sample network topology.

Note from the figure that all the routers are active routers. Node *AR1*, *AR2*, and *AR3* represent active router and node *R1*, *R2*, and *R3* represent receiver, respectively. Node *S* represents a sender. The number associated with each link is the maximum available bandwidth. Notice that the bandwidths of the links decrease towards the receivers. This arrangement is for realistically modeling a typical multicast environment. The traffic monitoring and adaptation schemes are implemented in all the active routers and receivers.

We assume that the multicast data is encoded into three layers of I, B, and P frame of MPEG. I frame is the base layer, and B and P frame are enhancement layers. The distribution of each layered data is assumed to be same even though it depends on the multimedia data type such as video conference and VOD. This is because the traffic is not controlled with a granularity of one layer but with respective filtering rate of each layer. Note here that the layer-based packet filtering allows limited link resource to be used effectively. B or P frame packet becomes useless if one of its lower layer packets is dropped.

4.2. Simulation Tool

Simulation of the sample network is done using SES/workbench [14,15]. It is a powerful general-purpose modeling and simulation tool used for designing sophisticated systems of various types. Since it supports graphical user interface and animated simulation, it is very useful for constructing evaluation models of complex systems and analyzing the performances. It is thus world-widely used for various engineering applications.

The followings are the assumed parameters and settings.

- The sender *S* generates a stream of 1Mbps traffic for each layer, that is, totally 3Mbps stream.
- The buffer in each node can hold 100 packets.
- The upper and lower threshold, t_u and t_l, are 0.1 and 0 respectively.
- Three filtering rate increment and decrement

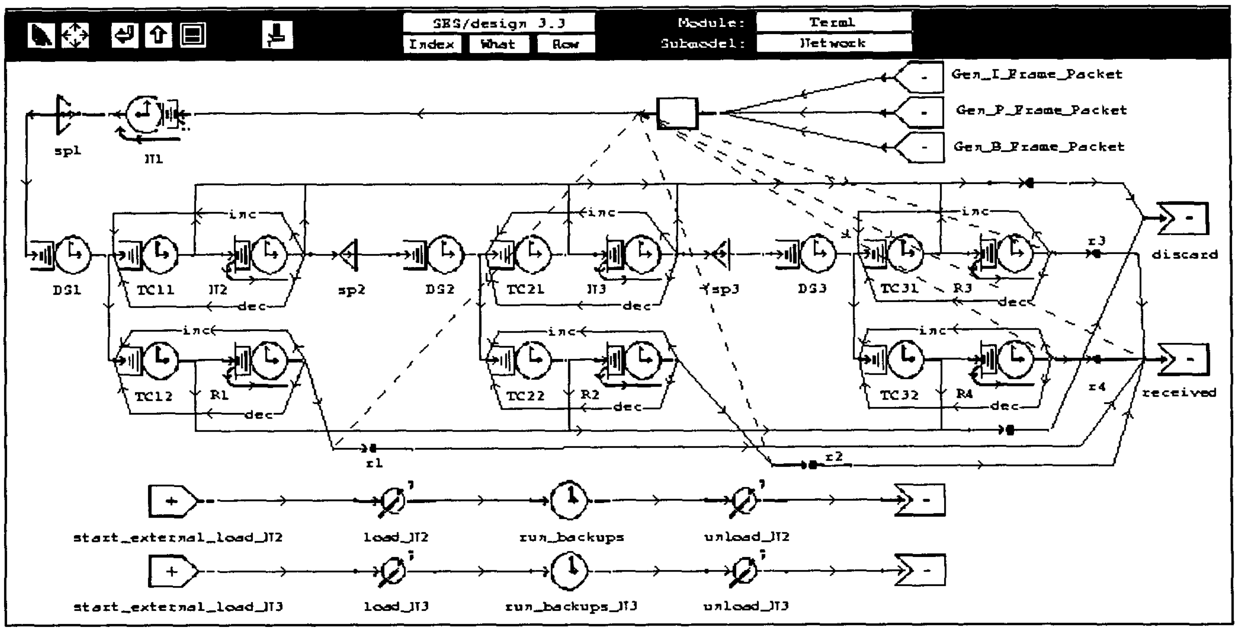

Figure 5. SES/Workbench simulation model for the network of Figure 4.

units are tested as, $(\Delta^+, \Delta^-) = (0.1, 0.1)$, $(0.3, 0.1)$ and $(1.0, 1.0)$, respectively. The three cases are named as *Ex #1, #2,* and *#3*, respectively.

Note that *Ex #1* is the case of fine adjustment with the same rate change for both increment and decrement. With *Ex #2*, increase of the filtering rate is larger than decrease for fastly reacting to the congestion condition. *Ex #3* represents layer-wise packet filtering employed in earlier layered multicast.

Figure 5 shows the simulation model of the sample network of Figure 4 developed using Workbench. In the figure, three kinds of packets are generated from *Gen_I_Frame_Packet, Gen_P_Frame_Packet* and *Gen_B_Frame_Packet* node with poisson distribution. Each active router consists of four nodes, i.e., *Nx, DSx, SPx*, and *TCxx*. Node *Nx* and *Rx* are active nodes that send traffic control messages with some delay for mimicking the packet processing time in active node. They also monitor the packet loss rate. Node *SPx* duplicates multicast packets, and node *DSx* distributes the duplicated packets to each output link. Finally, node *TCxx* controls the traffic on each output link depending on the traffic condition. If the packet loss rate exceeds the upper bound at active node *Nx* or *Rx*, it sends a message notifying the congestion state to the router in upstream through the link labeled as *inc*. On the contrary, if the packet loss rate is lower than the lower bound, it sends a message through the link *dec*. All the discarded packets sink at node *discard*, and the packets arriving at each receiver meet at node *received*. The dotted lines represent *response arcs* with which the statistical information is gathered.

Note that we vary the processing power of active routers during the simulation time for mimicking the traffic change. The processing powers are varied randomly with uniform distribution between 0 and 1. If the power of the active router is 0.5, for example, it uses 50% of its resource for serving the multicast traffic. The less the power is, the longer the packet transmission time will be.

4.3. Simulation Results

Figure 6 shows the relation among the power, population, and interarrival time at active router *AR2*. Here *Ex #2* is assumed and thus (Δ^+, Δ^-) is $(0.3, 0.1)$. Notice that the population of the active router and

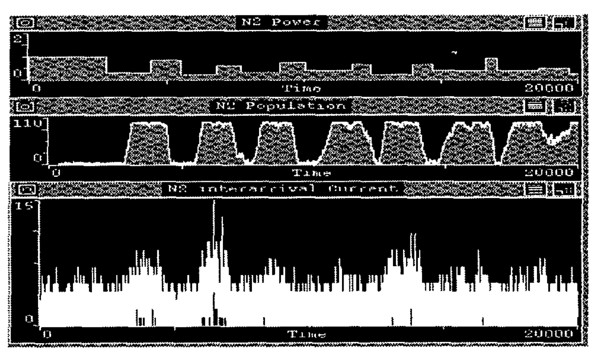

Figure 6. Relation among the power, population, and interarrival time at AR2.

interarrival time get larger when the power of the active router is low. When the packet loss rate exceeds the upper threshold, the message notifying the overflow condition is sent to *AR1*. It then starts packet filtering based on the traffic adaptation parameters for each layer to reduce the outgoing traffic. This results in increased packet interarrival time at *AR2*.

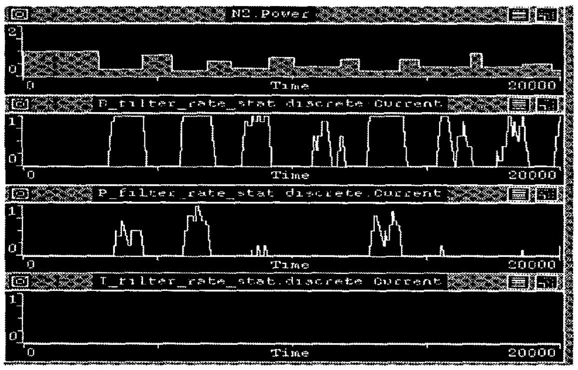

Figure 7. The variance of each filtering rate according to the variance of the power at AR2.

Figure 7 shows the variances of packet filtering rates of each layer of *AR2* as the power of it varies. We observe similar trend in other active routers while they differ in terms of the degree of variance. Notice that the filtering rate of B-frame is the highest and filtering occurs most frequently among the three layers. I-frame shows virtually no filtering. This is due to the priority assigned to the layers, with the highest priority to I-frame and lowest priority to B-frame.

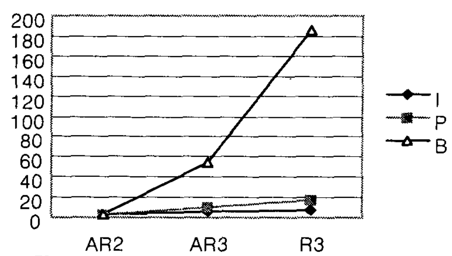

Figure 8. Interarrival time at the node AR2, AR3 and R3.

Figure 8 compares interarrival times at node *AR2*, *AR3*, and *R3*. As expected, the interarrival time of B-frame packet is the longest. This implies that B-frame packets arrive rarely. Also note that the time for I-frame linearly increases as approaching the receivers. This is because I-frame packets experience little delay in each intermediate router.

Figure 9 shows the queue population in the nodes. In each node, the number of packets in the queue increases

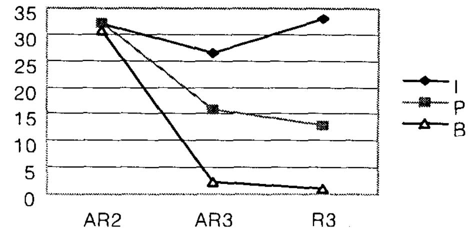

Figure 9. Queue population in the node AR2, AR3 and R3.

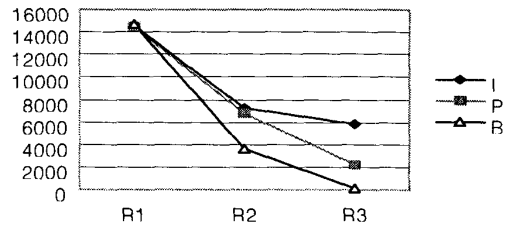

Figure 10. Number of the received packets at each end receiver.

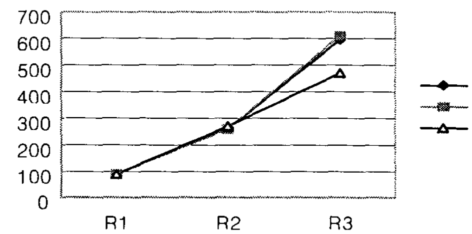

Figure 11. End-to-end delay of received packets at end receivers.

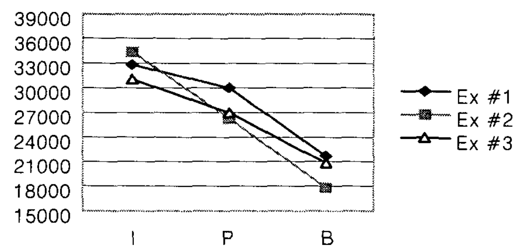

Figure 12. Total number of received packets at end receiver.

by the order of B, P and I-frame packets. In Figure 10, we can see that I-frame packets are successfully transmitted to end receivers for the most. Figure 11 reveals that B-frame layer has the shortest end-to-end delay. This is because B-frame packets are delivered only when congestion does not occur. When a link has enough bandwidth like R1, the average delays of different type packets are almost same.

Figure 12, 13, and 14 compare the results of varying filtering rate increment and decrement unit (Δ^+ and Δ^-),

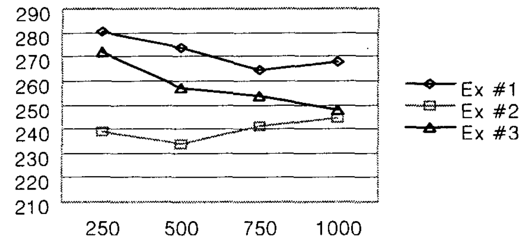

Figure 13. End-to-end delay of all received packets.

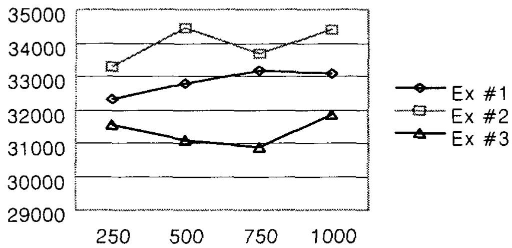

Figure 14. Total number of received I frame packets.

Ex #1, 2, and 3. Also Figure 13 and 14 show the results when the transmission period of the *state report message* to the parent is varied. In Figure 12, we can see that I-frame packets are delivered for the most with Ex #2. As shown in Figure 13, it allows the smallest end-to-end delay. This is because congestion can be quickly overcome by assigning larger Δ^- value than Δ^+, while the adjustment rate is much smaller (finer) than with Ex #3. Figure 14 reveals that the total number of received I-frame packets with Ex #2 is the largest.

5. Conclusion

Multicast is very useful for delivering data to a number of receivers. In multicasting multimedia data, adaptation to heterogeneous receivers of different bandwidths is a challenging problem. In this paper we have presented active traffic control mechanism for layered multimedia multicast (ATLM). ATLM controls traffic at each active router so that the increase of traffic on a certain link does not adversely affect the states of other links. ATLM can control the amount of layered traffic with a granularity of much smaller unit than existing protocols of a granularity of one layer. It does not require to maintain multiple sessions for transferring multiple layer data but only one session. Each node monitors its input traffic and sends a message with which the parent node decides to increase or decrease the amount of outgoing traffic.

We evaluated ATLM by simulation using SES/workbench in terms of total number of received packets, end-to-end-delay, and interarrival time for each layer. We used different parameters for each layer to filter the packets according to the congestion condition. The simulation revealed that ATLM is very effective for maximizing the bandwidth utilization and providing high QoS at the end receivers. The proposed scheme will be evaluated and compared with other schemes for more comprehensive networks.

References

[1] X. Li and M.H. Ammar, "Bandwidth control for replicated stream multicast video distribution," HPDC Focus Workshop on Multimedia and Collaborative Environments (Fifth IEEE Internal Symposium on high Performance Distributed Computing), pp. 356-363, Aug. 1996.

[2] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-driven Layered Multicast," ACM SIGCOMM, pp. 117-130, Aug. 1996.

[3] X. Li, S. Paul, and M. Ammar, "Layered Video Multicast with Retransmissions (LVMR): Evaluation of Hierarchicalrate Control," The Conference on Computer Communications (IEEE Infocom), pp. 1062-1072, March/April 1998.

[4] D.L. Tennenhouse et al., "A Survey of Active Network Research," IEEE communications magazine, pp. 80-86, Jan. 1997.

[5] J.C. Bolot, T. Turletti, and I. Wakeman. "Saclable feedback control for multicast video distribution in the internet," ACM SIGCOMM Symp. on communications architectures and protocols, pp. 58-67, England, Aug. 1994

[6] I. Busse, B. Deffner, and H. Schulzrinne, "Dynamic QoS control of multimedia applications base on RTP," Computer Communications, pp. 49-58, Jan. 1996

[7] X. Li, S. Paul, and M.H. Ammar, "Multi-Session Rate Control for Layered Video Multicast," Symp. on Multimedia Computing and Networking (MMCN'99), Jan. 1999.

[8] Y. Yemini and S. da Silva, "Towards Programmable Networks," FIP/IEEE International Workshop on Distributed Systems: Operations and Management, Oct. 1996.

[9] D.L. Tennenhouse and D.J. Wetherall, "Towards an Active Network Architecture," Computer Communication Review, Vol. 26, No. 2, pp. 5-18, April 1996.

[10] D.S. Alexander et al., "A Secure Active Network Architecture: Realization in SwitchWare," IEEE Network Special Issue on Active and Controllable Networks, Vol. 12, No. 3, pp. 37 – 45, 1998.

[11] M. Hicks et al., "PLAN: A Packet Language for Active Networks," The International Conference on Functional Programming (ICFP), pp. 86-93, 1998.

[12] S. Bhattacharjee, K.L. Calvert, and E.W. Zegura, "Congestion Control and Caching in CANES," The International Conference on Communications (ICC '98), 1998.

[13] L. Lehman, S.J. Garland, and D.L. Tennenhouse, "Active Reliable Multicast," IEEE INFOCOM'98, pp. 581-589, 1998.

[14] SES/workbench Release 3.2 Creating Models, SES.

[15] SES/workbench technical reference, SES.