

Workshop name: Computer Networks

Paper title: Multicast Congestion Control Strategy based on Call Admission Control

Abstract—In most of multicast congestion control mechanisms, the sender of a multicast session will adapt its sending rate to satisfy the network congestion status and acceptable receiving rates of all receivers. However, most of these mechanisms suffer the “drop-to-zero” problem where the sender adapts its sending rate to the acceptable receiving rate of the slowest receiver in the multicast group. As the size of the multicast group grows, it is very likely that the slowest receiver will have a very poor network status, either highly congested or low bandwidth, which leads the throughput of the multicast group to almost zero. The essence of this problem is that, as network traffic varies from time to time, any receiver is very likely to suffer network congestion at some instance of time. When the network status is very bad, a receiver really should not deserve to receive normal service. In other words, receivers with normal network status should not be affected by receivers with poor network status. Therefore, in this paper, we advocate the integration of multicast congestion control mechanisms with a call admission control (CAC) scheme to avoid the “drop-to-zero” problem. The proposed mechanism uses a multicast congestion control algorithm to gather acceptable receiving rate of each receiver and adapt the sending rate of the sender to the slowest receiving rate. When there is a change of receiving rate or a join of new receiver, the CAC mechanism is used to determine whether to kick out receivers with slow receiving rate or a new comer with a slow receiving rate based on a reward function. Our simulation results show that the proposed CAC scheme can effectively drop slow receivers in multicast group and, thus, serve the majority of group receivers at a much higher throughput.

Authors:

Yu-Shun Lo

Email: saint@exodus.cs.ccu.edu.tw

Phone number: 05-2720411 ext. 231

Ren-Hung Hwang

Email: rhhwang @cs.ccu.edu.tw

Phone number: 05-2720411 ext. 33112

Jeng-Muh Hsu

Email: hsjm @cs.ccu.edu.tw

Phone number: 05-2720411 ext. 33113

Fax : 05-2720859

Postal address: No. 160, Sanshing Tsuen, Minshiung Shiang, Chiai-Yi, Taiwan 621, R.O.C.

Contact author: Ren-Hung Hwang

Keywords: multicast, congestion control, call admission control

Multicast Congestion Control Based on Call Admission Control

Yu-Shun Lo, Ren-Hung Hwang and Jenq-Muh Hsu

Dept. of Communication Engineering

National Chung Cheng University

Chiayi, Taiwan, R.O.C.

{saint, rhhwang, hsujm}@exodus.cs.ccu.edu.tw

Abstract — In most of multicast congestion control mechanisms, the sender of a multicast session will adapt its sending rate to satisfy the network congestion status and acceptable receiving rates of all receivers. However, most of these mechanisms suffer the “drop-to-zero” problem where the sender adapts its sending rate to the acceptable receiving rate of the slowest receiver in the multicast group. As the size of the multicast group grows, it is very likely that the slowest receiver will have a very poor network status, either highly congested or low bandwidth, which leads the throughput of the multicast group to almost zero. The essence of this problem is that, as network traffic varies from time to time, any receiver is very likely to suffer network congestion at some instance of time. When the network status is very bad, a receiver really should not deserve to receive normal service. In other words, receivers with normal network status should not be affected by receivers with poor network status. Therefore, in this paper, we advocate the integration of multicast congestion control mechanisms with a call admission control (CAC) scheme to avoid the “drop-to-zero” problem. The proposed mechanism uses a multicast congestion control algorithm to gather acceptable receiving rate of each receiver and adapt the sending rate of the sender to the slowest receiving rate. When there is a change of receiving rate or a join of new receiver, the CAC mechanism is used to determine whether to kick out receivers with slow receiving rate or a new comer with a slow receiving rate based on a reward function. Our simulation results show that the proposed CAC scheme can effectively drop slow receivers in multicast group and, thus, serve the majority of group receivers at a much higher throughput.

Keywords — multicast, congestion control, call admission control

I. Introduction

Recently, more and more Internet applications, such as video conference and distance learning, rely on group communication for transmitting data to a group of members. Compared to multiple point-to-point connections, multicast is a more efficient way to transmit data to a large number of receivers as it reduces the transmission load both on the data sender and the network. However, for following two reasons, multicast applications are vulnerable to network congestion. First, a multicast session usually consists of members spread over the Internet. Second, receivers of a multicast session may be quite heterogeneous, e.g., the path to each receiver may have different capabilities and error characteristics. Therefore, the research on multicast congestion control has become more and more important.

In general, there are four issues in multicast congestion: congestion identification, feedback implosion, responsiveness and fairness. Congestion identification refers to defining a metric for identifying the degree of congestion condition at a receiver over a measurement period. When congestion occurs, how to deal with the potentially huge volume of the feedbacks from all receivers is called the feedback implosion problem. As congestion is identified, how quickly actions can be taken to react to the congestion is the issue of responsiveness. Finally, the fairness issue characterizes how the protocol shares bandwidth among multiple flows and how the protocol co-exists with existing protocols, especially with TCP.

Based on the number of multicast sessions used, existing multicast congestion control protocols can be classified into two broad categories. First, multi-group protocols [1-4] require setting up multiple multicast groups. The data stream is organized into multiple sub-streams in an incremental way and each sub-stream is transmitted on a multicast session. A receiver then incrementally joins to higher multicast groups according to its available bandwidth. Second, single-group protocols [5-11] establish a single multicast group with rate-based or window-based feedback control algorithms. These protocols allow a sender to update its transmission rate or congestion window size according to congestion information obtained on-line. In most cases, when the sender detects congestion, it reduces its rate or window size multiplicatively; otherwise, it increases its rate or window size linearly. Different protocols differ in their ways of determining the time interval to update, how to aggregate congestion information, and the criteria to determine congestion.

In this paper, we focus on the problem of single-group multicast congestion control. In general, in this type of protocols, a sender aggregates the feedback information from all receivers and adapts its sending rate according to the slowest receiver. As a consequence, most of single-group protocols suffer the “drop-to-zero” problem, i.e., the throughput of the multicast group will drop to zero due to adapting the

transmission rate to the slowest receiver. Therefore, special treatment is required to avoid the “drop-to-zero” problem.

In [5], the concept of virtual grouping was first introduced to partition the multicast group members into several subsets. The “drop-to-zero” problem is then solved by following two approaches. First, the sender can tune its window size (sending rate) according to the fastest receiver in a virtual group instead of the slowest receiver. Second, the sender can explicitly ask the slowest group to merge with other multicast groups. Although the multicast performance gets improved, tuning the window size to the fastest receiver may lead to link congestion.

Another approach to the “drop-to-zero” problem is to let the sender adjust its rate to the needs of majority receivers. For example, the LTRC (Loss Tolerant Rate Controller) [11] algorithm tries not to adapt the sending rate to congestion on a per receiver’s report basis, but to maintain the sending rate such that the loss rate is kept within a minimum and a maximum loss rate threshold. As a result, the sending will not overly reactive to spurious independent loss or a receiver with poor receiving rate. To prevent the “drop-to-zero” problem more effectively, it may be necessary for applications that require minimum transmitting rate guarantee to require receivers on congested paths to leave the multicast group [12].

The essence of the “drop-to-zero” problem is that, as network traffic varies from time to time, any receiver is very likely to suffer network congestion at some instance of time. When the network status is very bad, a receiver really should not deserve to receive normal service. In other words, receivers with normal network status should not be affected by receivers with poor network status. Therefore, we believe that forcing slow receivers to leave the multicast group is the key to solve the problem. In this paper, we propose to integrate single-group multicast congestion control with call admission control (CAC) scheme to avoid the “drop-to-zero” problem. The basic idea is two-folded. First, the sender adapts its sending rate to the slowest receiver according to feedback information from receivers. Second, a reward model is defined for a sender to determine whether to admit a new receiver or to kick out an existing slow receiver. Our simulation results show the proposed CAC scheme can serve the majority of group receivers by effectively dropping slower receivers in multicast group such that the network reward can be maximized.

The remainder of this paper is organized as follows. Section 2 presents the framework of CAC scheme. Section 3 describes the policy our CAC scheme and its implementation detail. Section 4 evaluates the performance of the proposed scheme via simulations. Conclusion and future work of this paper is discussed in section 5.

II. Call Admission Control Framework

In this section, we integrate the multicast congestion control protocol with CAC scheme to avoid the “drop-to-zero” problem. In the following, we describe the multicast congestion control protocol required and the proposed CAC scheme based on a reward model.

2.1 Multicast congestion control protocol

In this paper, we assume a rate-based single-group multicast congestion control protocol is adopted to adjust the sending of the sender to the slowest receiver of the multicast group. Examples of this kind of protocols are [7, 9]. Specifically, we assume the multicast congestion control protocol has following features:

- (1) Receivers measure their receiving rate periodically.
- (2) Receivers report their receiving rate to the sender periodically.
- (3) The sender adapts its sending rate to the receiving rate of the slowest receiver.

The measurement period is called the round trip time, RTT. That is, we assume that the sender will obtain receiving rate of each receiver and adjust its sending rate every RTT.

Other features, such as TCP friendly, can be adopted also. However, the detail mechanism for achieving these features is not the focus of this paper.

2.2 The CAC scheme

The CAC scheme is performed at the sender utilizing the receiving rates received by the multicast congestion control protocol, as shown in Figure 1. The CAC scheme uses a policy called On-line Membership Management Strategy (OMMS) to determine whether to accept a new receiver or kick out an existing receiver. Therefore, the OMMS policy is invoked when a receiver joins the multicast group or at the end of the measurement period of the multicast congestion control protocol. The receiver will be notified via a signaling protocol if the policy decides to ask it to leave.



Fig.1 The CAC Signaling Protocol

III. On-line Membership Management Strategy

Let us describe the reward model used in the OMMS policy. Let N be the number of receivers in the multicast group and AR_i be the receiving rate reported from receiver i . Based on the multicast congestion control protocol, the sending rate, R , of the multicast group will be

$$R = \min\{AR_1, AR_2, \dots, AR_N\}.$$

Let S be the reward rate (per unit of time) of a multicast group. S can be defined in a number of ways. In this paper, we define S as the summation of the throughput of all receivers. That is, for a multicast group with N receivers and a sending rate of R , S is given by:

$$S = N \cdot R.$$

Since the OMMS policy will reject a slow receiver based on the reward model, we also introduce a penalty function for rejecting the i th receiver, $P(i)$. In this paper, we will consider two penalty functions: $P(i)=0$ and $P(i)=R$.

The OMMS policy decision is invoked at following two instances: first, a new receiver joins the group; second, at the end of each RTT. Without loss of generality, let us assume that there are N receivers in the multicast group and the reward of the multicast group is $S_N = N \cdot R$. When the $(N+1)$ st receiver joins the group, it measures its receiving rate, AR_{N+1} , during the next RTT and reports the receiving rate to the sender. The sender then invokes the OMMS policy to determine whether accepting the new receiver will increase the reward of the multicast group by computing the following equations:

$$R = \min(AR_1, \dots, AR_{N+1}) \quad (1)$$

$$S_{N+1} = (N+1) \cdot R \quad (2)$$

$$\text{If } [S_{N+1} - (S_N - P(N+1))] \begin{cases} \geq 0: & \text{accept} \\ < 0: & \text{reject} \end{cases} \quad (3)$$

It is possible that an accepted receiver becomes the bottleneck of the multicast group as more and more receivers join the group. It is also possible that, due to the dynamics of the network traffic, the receiving rate of a receiver drops to a very low rate. Therefore, at the end of each RTT, the sender could invoke the OMMS policy to check if some receivers are too slow to degrade the throughput of the group. (To reduce the processing overhead, the sender could invoke the OMMS policy only at the multiplicity of the RTT.) Without loss of generality, let us assume that there are N receivers in the multicast group with receiving rate sorted non-increasingly such that $AR_1 \geq AR_2 \geq \dots \geq AR_N$. The reward of the group is then given by $S_N = N \cdot AR_N$. Now OMMS policy will be invoked to decide whether to kick out the N th receiver by following equations:

$$S_{N-1} = (N-1) \cdot AR_{N-1}$$

$$\text{If } [S_N - (S_{N-1} - P(N))] = \begin{cases} \geq 0: & \text{keep the } N\text{th receiver} \\ < 0: & \text{kick out the } N\text{th receiver} \end{cases} \quad (4)$$

The OMMS policy repeats the procedure until the decision is to keep the receiver under consideration.

We assume a receiver that receives the leave message from the sender will leave the multicast group through IGMP [15] protocol. A receiver can join the group later on if the network traffic becomes less congested. (For more efficient membership management, IGMPv3 [16] is suggested.)

IV. Numerical Results

In this section, we evaluate the performance of the proposed scheme via simulations using the ns simulator [13]. The topology simulated is a 200-node graph generated by the BRITE network topology generator [14]. The 200 nodes are grouped into 10 domains (AS's), each with 20 nodes (routers). The bandwidth between two AS nodes is 45.4Mbps while the bandwidth of the link within an AS is uniformly distributed between 128k and 100Mbps. (The link bandwidth within an AS can be viewed as the residual bandwidth, i.e., excluding the background traffic of unicast, for multicast traffic.) There are 8 multicast groups in the simulation with one sender in each group. Receivers of each multicast group arrive according to a Poisson process. After a receiver joins a multicast group, it will stay forever unless kicked out by the CAC scheme. A sender or receiver is connected to a randomly selected router from the 200 nodes. The time simulated by each simulation is 2 hours. The CBT protocol is used as the multicast routing protocol.

Three schemes are simulated in the simulations: the multicast congestion control protocol with no CAC scheme (NO_CAC), the CAC scheme with no penalty (CAC1) and the CAC scheme with penalty (CAC2). The penalty is set to the receiving rate (AR) of the receiver rejected or kicked out. Performance metrics we are interested include group sending rate (R), the number of receivers (N), and the reward of each group ($S = R \times N$).

Since the performance of different groups resembles each other, we only show performance of some groups. Figure 2-10 shows the performance of the whole network and group 2, 4, 6, respectively. From these figures, we can clearly observe the “drop-to-zero” problem of the multicast congestion control protocol with no CAC scheme. In all groups, if the CAC scheme is not adopted, the sending rate of each group all drops to a very low rate as the number of receivers increases. On the other hand, the sending rate maintains at a reasonably high rate for all groups when the

CAC scheme is applied. CAC schemes, even with penalty, also yield much higher reward than the NO_CAC scheme. That is, rejecting slow receivers not only improves the throughput of a multicast session but also increase the reward of the network. In other words, it is beneficial to the majority of group members and the network.

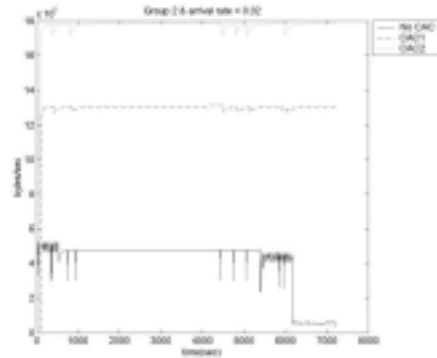


Fig.2 Sending rate of group 2 (R)

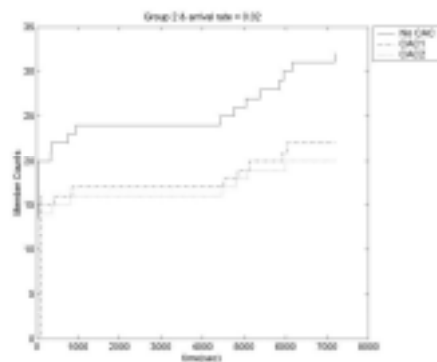


Fig.3 Numbers of group 2 receivers (N)

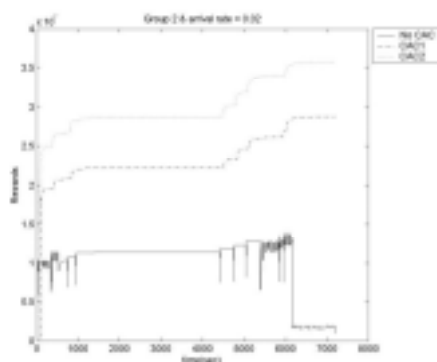


Fig.4 Group 2 reward (S)

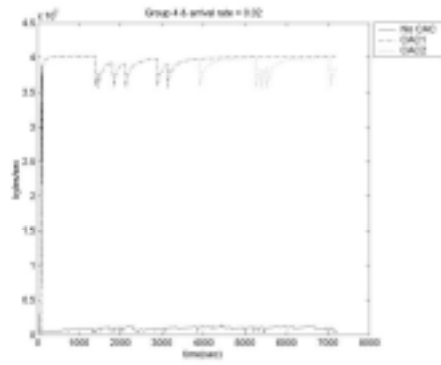


Fig.5 Sending rate of group 4 (R)




Fig.6 Numbers of group 4 receivers (N)

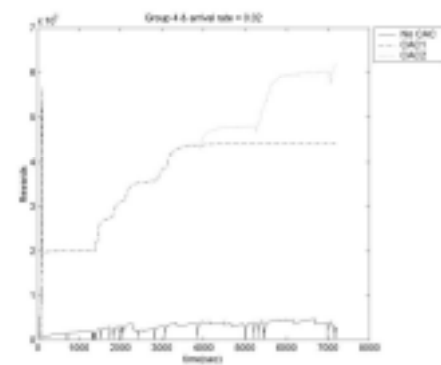


Fig.7 Group 4 reward (S)

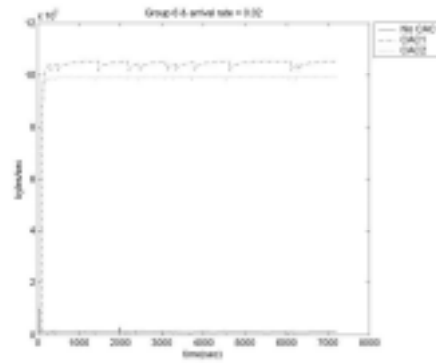


Fig.8 Sending rate of group 6 (R)

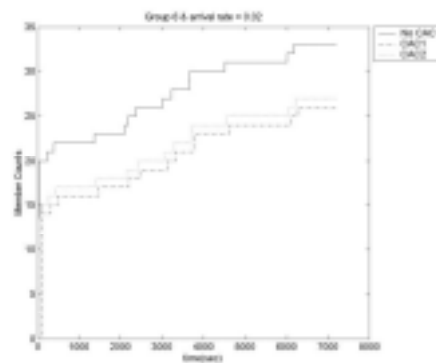


Fig.9 Numbers of group 6 receivers (N)




Fig.10 Group 6 reward (S)

The effectiveness of the proposed CAC scheme can also be verified by observing that the rewards of the two CAC schemes do not decrease as more and more receivers arrived. In other words, the CAC scheme will try to maximize the network reward according to the dynamics of receivers, either the arrival of new receiver or the changes of receiving rate of each receiver.

It is interesting to observe that the size of each multicast group maintains roughly at a fixed number under the CAC schemes. This suggests that there may exist an optimal

group size for certain network conditions. Currently, we are formulating a Markov Decision Process (MDP) to explore this observation.

V. Conclusions and Future Work

In this paper, we propose to integrate the multicast congestion control protocol with a CAC scheme to avoid the “drop-to-zero” problem. Under the multicast congestion control protocol, a sender adapts its sending rate to the slowest receiver according to feedback information from receivers. With the CAC scheme, a slow receiver will be rejected or kicked out based on a reward model. Simulation results show that the CAC scheme can effectively serve the majority of group receivers, improve the network reward, and resolve the “drop-to-zero” problem.

There are several issues of the CAC scheme require further investigation. In particular, we are developing a mathematic model to analyze the performance of the CAC scheme and solve the optimal group size problem.

Reference

- [1] L. Vicisano, J. Crowcroft and L. Rizzo, “TCP-like congestion control for layered multicast data transfer,” IEEE INFOCOM’98, 1998, pp. 996–1003.
- [2] D. Sisalem and A. Wolisz, “MLDA: A TCP-friendly Congestion Control Framework for Heterogeneous Multicast Environments,” IWQoS 2000, pp. 65-74.
- [3] J. Byers, *et al.*, “FLID-DL: Congestion Control for Layered Multicast,” in Proc. 2nd Int’l Workshop of Networked Group Communication, 2000.
- [4] J. Byers, M. Luby, and M. Mitzenmacher, “Fine-Grained Layered Multicast,” IEEE INFOCOM 2001, pp. 1143-1151.
- [5] S. Chang, H. Jonathan Ghao and X. Guo., “TCP-friendly window congestion control with dynamic grouping for reliable multicast,” IEEE GLOBECOM '00, 2000, pp. 538-547.
- [6] I. Rhee, N. Balaguru and G.N. Rouskas, “MTCP: scalable TCP-like congestion control for reliable multicast,” INFOCOM '99, 1999 , pp. 1256-1273.
- [7] S. Shi and M. Waldvogel, “A Rate-based End-to-end Multicast Congestion Control Protocol,” in Proc. 5th IEEE symposium on Computers and Communications, 2000, pp. 678-686.
- [8] I. Rhee, V. Ozdemir and Y. Yi, “TEAR: TCP Emulation at Receivers – Flow Control for Multimedia Streaming,” Tech. Rep., Dept. of Comp. Sci., NCSU, 2000.
- [9] L. Rizzo, “pgmcc: a TCP-friendly single-rate Multicast congestion control scheme,” ACM SIGCOMM, 2000, pp. 17-28.

- [10] S. Paul, K.K. Sabnani, J.C.-H. Lin and S. Bhattacharyya, "Reliable Multicast Transport Protocol (RMTP)," IEEE Journal on communications, vol. 15, Issue:3 , 1997, pp. 407-421.
- [11] T. Montgomery, "A loss tolerant rate controller for reliable multicats," Technical Report NASA-IVV-97-011, West Virginia University, August 1997.
- [12] S. Floyd, V. Jacobson, C.G. Liu, S. McCanne and L. Zhang, "A Reliable Multicast Framework for Light-Weight Sessions and Application Level Framing," IEEE/ACM Transactions on networking, vol. 5, Issue: 6, 1997, pp. 784-803.
- [13] UCB/LBNL Network Simulator – ns – version 2, 1997, <http://www.isi.edu/nsnam/ns/>.
- [14] A. Medina, A. Lakhina, I. Matta and J. Byers, "BRITE: an approach to universal topology generation," in Proc. of Ninth International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems, 2001, pp. 346-353.
- [15] W. Fenner, "Internet Group Management Protocol, Version 2", RFC 2236, Nov 1997.
- [16] B. Cain, S. Deering, B. Fenner, I. Kouvelas, A. Thyagarajan, "Internet Group Management Protocol, Version 3", INTERNET-DRAFT(working draft), May 2002, <http://www.ietf.org/internet-drafts/draft-ietf-idmr-igmp-v3-11.txt>.
- [17] Y.R. Yang and S.S. Lam, "Internet Multicast Congestion Control: A Survey," in Proc. ICT 2000, Acapulco, Mexico, May 2000.
- [18] Z. Qi, Z. Mingyu, W. Guoxin and G. Guanqum, "Analysis and Taxonomy of Congestion Control Mechanisms for Internet Reliable Multicast," in Proc. IEEE ICCS'2000, Singapore, Nov 2000.