Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 51 (2007) 1421–1443
Computer Networks
www.elsevier.com/locate/comnet

Generalized multicast congestion control

Jiang Li [a], Murat Yuksel [b,*], Xingzhe Fan [c], Shivkumar Kalyanaraman [d]

[a] Howard University, Department of Systems and Computer Science, 2300 6th Street NW, Washington, DC 20059, United States
[b] University of Nevada – Reno, Computer Science and Engineering Department, 171, Reno, NV 89557, United States
[c] University of Miami, Electrical and Computer Engineering Department, Room 411, McArthur Engineering Building, Coral Gables, FL 33124, United States
[d] Rensselaer Polytechnic Institute, Electrical Computer and Systems Engineering Department, 110 8th Street, Troy, NY 12180, United States

Received 13 December 2005; received in revised form 24 July 2006; accepted 31 July 2006
Available online 30 August 2006

Responsible Editor: Nelson Fonseca

Abstract

Efficient multicast congestion control (MCC) is one of the critical components required to enable the IP multicast deployment over the Internet. Previously proposed MCC schemes can be categorized in two: single-rate or multi-rate. Single-rate schemes make all recipients get data at a common rate allowed by the slowest receiver, but are relatively simple. Multi-rate schemes allow of heterogeneous receive rates and thus provide better scalability, but rely heavily on frequent updates to group membership state in the routers. A recent work by Kwon and Byers, combined these two methods and provided a multi-rate scheme by means of single-rate schemes with relatively low complexity.

In this paper, we propose a new scheme called generalized multicast congestion control (GMCC). GMCC provides multi-rate features at low complexity by using a set of *independent* single-rate sub-sessions (a.k.a layers) as building blocks. The scheme is named GMCC because single-rate MCC is just one of its special cases. Unlike the earlier work by Kwon and Byers, GMCC does not have the drawback of *static* configuration of the source which may not match with the *dynamic* network situations. GMCC is *fully* adaptive in that (i) it does not statically set a particular range for the sending rates of layers, and (ii) it eliminates redundant layers when they are not needed. Receivers can subscribe to different subsets of the available layers and hence can always obtain different throughput. While no redundant layers are used, GMCC allows receivers to activate a new layer in case existing layers do not accommodate the needs of the actual receivers.

Keywords: Multicast; Congestion control; Multi-rate

1. Introduction

In multicast congestion control (MCC), satisfying demands of several heterogeneous receivers while maintaining scalable and efficient operation

* Corresponding author. Tel.: +1 775 784 6974; fax: +1 775 784 1877.

E-mail addresses: lij@scs.howard.edu (J. Li), yuksem@cse.unr.edu (M. Yuksel), fanx@miami.edu (X. Fan), shivkuma@ecse.rpi.edu (S. Kalyanaraman).

1389-1286/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2006.07.014

has been one of the major research problems. Researchers have developed various schemes that work effectively with different situations. *Single-rate* MCC schemes are simple and easy to deploy, but they only work well with small number of receivers or high number of receivers with less heterogeneity. In single-rate protocols such as ERMCC [2], PGMCC [3] and TFMCC [4], the source sends data to all receivers at a dynamically adjusted rate. The rate has to be adapted to the slowest receiver to avoid consistent congestion at any part of the multicast tree. Therefore, faster receivers suffer. Still, single-rate protocols have advantages because they are simple.

For cases where receivers are high in number or significantly different in their bandwidth and congestion circumstances, single-rate schemes do not scale. Hence, by adding more implementation complexity, *multi-rate* MCC schemes that are able to operate under a wider range of network conditions have been developed.

In multi-rate schemes, the source maintains several layers each with different transmission rate, and receivers subscribe to different subsets of these layers depending on their and network's bandwidth and congestion circumstances. In a multi-rate multicast session, each layer uses a separate multicast group address. In most multi-rate protocols, the sending rates in these layers are not fully adaptive. They are either static, such as in RLM [5] and PLM [6], or dynamic but are defined by a carefully designed pattern, such as in RLC [7], FLID-DL [8], FLGM [9], STAIR [10] and WEBRC [11]. Recipients have to increase or decrease their receiving rates by joining or leaving some groups.[1] To perform join and leave operations, they send IGMP messages to routers. Upon the receipt of these IGMP messages, routers update their multicast group states to allow traffic through (for join) or stop traffic forwarding (for leave), which allows adjusting throughput for receivers. To quickly react to congestion, these operations have to occur frequently. As a result, a large volume of control traffic is introduced into the network, and the routers are heavily loaded because all the rate control burden has been shifted to them. Moreover, according to IGMP [12], the join and leave operations (especially leave) need time to take effect. Therefore, the number of these operations is limited during a period and restricts the effectiveness of these multicast congestion control schemes. These schemes are also called receiver-driven schemes.

A concurrently proposed scheme SMCC [1] is a hybrid of single-rate and multi- rate multicast congestion control. It combines a single-rate scheme TFMCC [4] with the receiver-driven idea. In each layer, the source adjusts sending rate *within a certain limit* based on TFMCC, and receivers join or leave layers cumulatively according to their estimated maximum receiving rates using TCP throughput formula [13]. Since the flows in each layer are adaptive to network status, the number of join and leave operations are greatly reduced. The congestion control is more effective.

SMCC requires static configuration of the maximum sending rates for each layer. This requirement makes SMCC not capable of accommodating receivers with variant bandwidth circumstances. In the case when many or all of the receivers fall into the lowest layer, SMCC cannot provide new layers with smaller sending rates. Again, when many or all of the receivers subscribe to the very highest layer(s), then lower layers become redundant, thereby causing the scheme to spend extra effort to maintain those unnecessary layers.

In this paper we propose a new scheme GMCC that solves these problems while having the merits of SMCC. In the remainder of this section, we will briefly describe GMCC and summarize key contributions and properties of it. Then, in the rest of the paper, we will describe the details of GMCC, and show simulation results to demonstrate the performance of GMCC. In Section 3, we will provide a detailed explanation of GMCC components at the source and receivers. Finally, we will show our simulation-based performance evaluation of GMCC in Section 4, and conclude in Section 5.

1.1. Brief description of GMCC

The functions of the source and the receivers in GMCC can be decoupled into two main categories: intra-layer, and inter-layer. GMCC uses single-rate MCC to manage intra-layer activities at the source and the receivers. In particular, rate adaptation and congestion representative (CR) selection are totally left to the single-rate MCC scheme that is being used. Similarly, creation and management of feedback packets at the receivers are done by the

[1] Joining a layer is also called subscription, leaving a layer is also called unsubscription. In this paper we will use both sets of terms ex-changeably.