Available online at www.sciencedirect.com
ScienceDirect
Computer Communications 30 (2007) 1721–1730
computer communications
www.elsevier.com/locate/comcom

Cost analysis and optimization for IP multicast group management

Yongjun Li [a,b,1,2], James Z. Wang [c,*]

[a] School of Computer Science and Engineering, South China University of Technology, Guangzhou 510640, China
[b] Division of Science and Engineering, Guangzhou Radio and Television University, Guangzhou 510260, China
[c] Department of Computer Science, Clemson University, Box 340974, Clemson, SC 29634, USA

Received 16 August 2006; received in revised form 4 February 2007; accepted 6 February 2007
Available online 15 February 2007

Abstract

This paper studies the system parameters that affect the total cost of managing the multicast group on a router. A Petri net model is first proposed to describe the states and transitions of the multicast group management. Based on this model, a delayed vacation queue, extended from a simple M/M/1 queue, is used to analyze the total cost of the multicast group management under various system conditions. The formal analysis reveals that the total cost of the multicast group management is minimized when a router delays a certain time to send its pruning messages to upstream routers. Furthermore, a formula is derived to calculate the optimal delay time for sending the pruning messages to minimize the multicast group management cost. To validate this cost analysis model, an IP multicast group management module is added to the NS-2 simulator to simulate the IP multicast group management. The simulation results under various system parameters are consistent with the results obtained by formal analyses, validating the proposed cost analysis model.

Keywords: IP Multicast; Group management; Delayed vacation queue; Cost analysis; Optimization; NS-2

1. Introduction

Internet Protocol (IP) multicast [1,2] is a bandwidth-conserving technique that reduces network traffic by simultaneously delivering a single stream of information to a large number of network hosts. IP multicast delivers source traffic to multiple receivers without adding any additional burden on the source or the receivers while using the least network bandwidth of any competing technology. High-bandwidth applications, such as video streaming, may require a large portion of the available network bandwidth for a single stream. In these applications, an efficient way of sending data simultaneously to more than one receiver is by using IP multicast. Fig. 1 demonstrates how data are delivered from one source to a group of interested recipients by using IP multicast.

Multicast is based on the concept of grouping the receivers with a common interest to receive a single data stream. This multicast group should not have any physical or geographical boundaries, i.e., the hosts can be located anywhere on the Internet. Although senders not belonging to a multicast group can send datagrams to group members, hosts must be a member of this group to receive the multicast data stream. Since any host can join a multicast group without authentication and leave the group at any time, it is challenging to efficiently manage the dynamic multicast group. In current Internet architecture, the management and maintenance of the multicast group membership is defined in the Internet Group Management Protocol (IGMP) [1,3,4]. IGMP is a protocol used between hosts and multicast routers on a single physical network to establish membership of a particular multicast group. A host identifies its group membership by sending IGMP messages to its associated router and the router, on the other hand, listens to IGMP messages and periodically sends out

* Corresponding author. Tel.: +1 864 656 7678, fax: +1 864 656 0145.
E-mail addresses: liyj@scut.edu.cn (Y. Li), jzwang@cs.clemson.edu (J.Z. Wang).
1 Tel.: +86 20 87110840, +86 20 84272058.
2 His work is partially supported by China National Natural Science Fund (No. 10501014) and Guangdong Natural Science Fund (No. 62067).

0140-3664/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2007.02.002

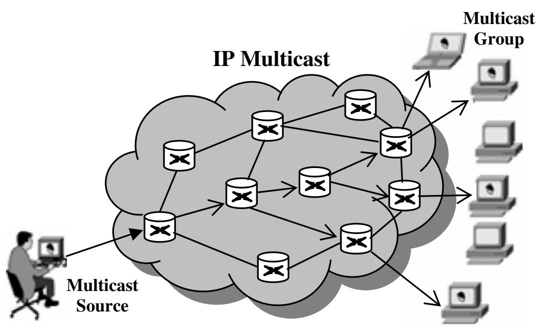

Fig. 1. IP multicast.

queries to discover which groups are active or inactive on a particular subnet. Multicast routers use the group information, in conjunction with a routing protocol [5–8], to support IP multicast over the Internet.

Most of recent multicast studies focus on application layer multicast protocols [9–11] or topology aware multicast protocols [12–17] due to the emergence of new applications, such as P2P and wireless ad hoc networks. Although there are deployment issues in IP multicast services and architectures [18], IP multicast is still the most cost-effective technique in the current Internet architecture to deliver source traffic to multiple receivers because it does not add any additional burden on the source or the receivers [19]. Thus, it is critical to study the cost efficiency of the IP multicast routing and group management, and investigate how to achieve the best IP multicast performance with the least cost in routers. Recently, only a few studies [20,21] tried to analyze the multicast performance. However, these studies focus on design different methods or tools to measure the performances of multicast protocols. No effort was made to evaluate how IGMP protocols affect the performance of the IP multicast, nor did they try to improve the performance of the IP multicast. Regarding improving the IP multicast performance, Gorinsky et al. [22] discussed the possible solutions for designing robust multicast congestion control protocols for IP multicast in the presence of untrusted hosts. Nonetheless, none of these existing studies evaluate the impact of IP multicast group management on the performance of the IP multicast.

In this paper, we evaluate how the delay of pruning the IP multicast tree impacts the cost of IP multicast group management and determine the optimal delay time for IP multicast group management under various system conditions. We first propose a Petri net to model the IP multicast group management and, in turn, develop a cost analysis model using a delayed vacation queue, extended from a simple M/M/1 queue, to evaluate the total cost of multicast group management on a router [23]. Using this delayed vacation queuing model, we derive a formula to calculate the optimal delay time for a router to send pruning message to its upstream routers so that the cost of multicast group management is minimized. Our simulation studies validate our cost analysis and optimization model.

The rest of the paper is organized as follows. After discussing our proposed Petri net model for the multicast group management in Section 2, we analyze the total cost of multicast group management on a router, using a delayed vacation queuing system, under various system conditions in Section 3. In Section 4, we use simulation studies to validate our cost analysis model. Finally, we present our concluding remarks and discuss future work in Section 5.

2. IP multicast group management model

In IP multicast, a router transmits datagrams received from the upstream routers to local multicast hosts, and manage the IP multicast group in its local network through IGMP queries. When a multicast router finds a host joining a multicast group, it broadcasts the graft message to other routers to request joining a multicast group. If a multicast router finds that the last member of a local multicast group exits, it sends a prune message to the upstream router to trim itself from the multicast tree and stop transmitting multicast datagrams to its directly attached local subnet.

However, due to the dynamic nature of the Internet, users join and leave a multicast group on their own terms. A new host may join a multicast group following the last member of this multicast group exits. If a new multicast member arrives after the router has pruned its multicast group by sending a prune message to upstream routers, this router must broadcast a graft message to other routers to rejoin the multicast group. To address this issue, most exiting IP multicast protocols choose to delay a random time before sending the prune message to its upstream routers, hoping new members will join the multicast group soon. However, if the delay time is too short, when multicast group members exit and enter a multicast group frequently, the cost of the IP multicast group management increases because the router must frequently send prune and graft messages to upstream routers. However, not pruning the inactive multicast group or delaying an excessively long time to send the prune message hinders the IP multicast performance due to the cost of maintaining the inactive multicast group. Hence, there must be an optimal delay time for sending prune messages that can minimize the IP multicast group management cost. Unfortunately, no attempt has been made to determine the optimal delay time for pruning the IP multicast tree.

To analyze the cost of the IP multicast group management and determine the optimal delay time for pruning the IP multicast tree, we propose an event/condition Petri net [24] to model the state transitions on a router during the IP multicast group management.

Definition 1 There are seven transitive states for a router during the IP multicast group management:

r_0: router initial state (no member state).
r_1: a new group member just requested for joining a multicast group.