Available online at www.sciencedirect.com

SCIENCE DIRECT®

Computer Communications 29 (2006) 1201–1212

computer communications

www.elsevier.com/locate/comcom

Overlay subgroup communication in large-scale multicast applications☆

Jangwon Lee[a,*], Gustavo de Veciana[b]

[a]Qualcomm Inc., 92121, San Diego, USA
[b]Electrical and Computer Eng. Dept. University of Texas at Austin

Received 10 May 2004; revised 22 July 2005; accepted 25 July 2005
Available online 19 August 2005

Abstract

In the paper, we consider a *preference heterogeneity* problem in large-scale multicast applications. Abundant content, data type and diverse members' interests naturally lead to preference heterogeneity within a multicast session requiring frequent communication within subgroups of members sharing common interests/requirements. In this paper, we take an overlay approach which builds topology-sensitive subgroup communication (TSC) structures to support efficient subgroup communications in large-scale multicast applications. Our TSC mechanism completely eliminates additional creation of multicast groups while minimizing the exposure of unnecessary packets to members and links. Our mechanism exploits the spatial locality of members within a given subgroup, and enables members to autonomously build a TSC structure consisting of multiple unicast and scoped multicast connections.

Simulations using real topology data show that TSC mechanism performs well for diverse configurations with different densities and distributions of nodes in a subgroup.
© 2005 Elsevier B.V. All rights reserved.

Keywords: Multicast; Subgroup; Topology-sensitive; Overlay

1. Introduction

IP multicast is an efficient one-to-many or many-to-many delivery method which can provide a number of operational advantages for content and network providers by reducing the overall resources consumed to achieve such distribution. A single packet transmitted by the source traverses each link in the multicast distribution tree to all receivers in the multicast group. Due to intensive needs for high bandwidth requirement, large-scale interactive applications such as distributed interactive simulations (DIS), video conferencing tools and multi-player games can benefit from IP multicast. Although the members in such applications join a multicast session for some common goal, abundant content, data type and heterogeneity in members' interests naturally lead to *preference heterogeneity* within sessions [1], requiring frequent communication within *subgroups* of members sharing common interests/requirements.

A multicast session shared by all members (referred to as the global multicast session) can be used to support subgroup communication. However, this may lead to inefficiency, i.e., packets are delivered to the entire tree, which results in wasted bandwidth and CPU processing power to transmit and handle unnecessary packets. This is referred to as the *exposure* problem. The exposure problem can be completely eliminated if data is only forwarded along a tree induced by the members of each subgroup, as required. This can be achieved by creating a new multicast session for each subgroup. However, this requires routers to store multicast forwarding state information for each subgroup, which can cause a significant scalability problem as the number of subgroups increases [2,3]. Thus, mechanisms to handle preference heterogeneity should consider the both exposure problem and the scalability of multicast forwarding state problem.

Most existing approaches to the preference heterogeneity problem focus on developing *clustering* frameworks, i.e., given a limited number of multicast sessions, determine how to best cluster multiple subgroups into multicast

☆ Partial work of this paper was presented at proceedings of the 4th international workshop on networked group communication (NGC), Oct. 2002

* Corresponding author. Tel.: +858-350-9755; fax: 858-845-2651.
E-mail addresses: jangwonl@qualcomm.com (J. Lee), gustavo@ece.utexas.edu (G. Veciana).

0140-3664/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2005.07.005

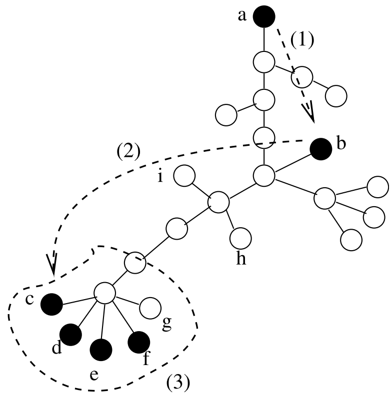

Fig. 1. An illustration of a subgroup communication in a TSC mechanism.

sessions based on a preference matrix [1] or players' positions in a virtual cell [4].

In the paper, we propose a topology-sensitive subgroup communication (TSC) mechanism to support efficient subgroup communication in large-scale multicast applications. Our TSC mechanism allows members in a subgroup to autonomously build a TSC structure consisting of multiple unicast and scoped multicast connections.

For example, consider a distribution tree for a multicast session, G, in Fig. 1. All end nodes are members of G and the black end nodes are members of a subgroup, S. In our scheme, when a wishes to send packets to other members in S, packets will be delivered as follows: (1) $a \rightarrow b$ via unicast; (2) $b \rightarrow c$ via unicast; and, (3) $c \rightarrow \{d, e, f\}$ via multicasting with a TTL scope of 2 as shown in Fig. 1. We assume that the multicast tree for G is a bidirectional shared one[1].

Note that in the example, the use of unicast can suppress the exposure and the use of scoped multicast can reduce duplicate packets traversing the same link. Our approach does not require the creation of new multicast sessions, which can completely eliminate any additional multicast forwarding state except those of the global session. It tries to minimize the exposure by exploiting spatial locality among members within a given subgroup.

Throughout simulations, we study which environments are advantageous to apply the proposed mechanism and other existing approaches, e.g., global multicast tree or unicast. Simulation results show that under various configurations of density and distribution modes of a subgroup, the sensitivity of our TSC mechanism is small compared to others. This is especially beneficial in the case where information about subgroups is not available - as is likely to be the case in practice.

The paper is organized as follows. We discuss related work and contrast them with our work in Section 2. In Section 3 we discuss how to construct and maintain a TSC structure. In Section 4, we evaluate and compare the proposed TSC mechanism with other schemes in various environments. Section 5 concludes the paper.

2. Related work

The scalability of state associated with multicast forwarding by routers has been one of the challenges in a wide deployment of IP multicast. Reduction of multicast forwarding state at routers can be achieved through aggregation or elimination of non-branching approaches. In [2], multiple multicast forwarding entries are aggregated if entries have adjacent group address prefixes and matching incoming and outgoing interfaces. The goal of dynamic tunnel multicast [8] and REUNITE [9] is to reduce multicast states by eliminating non-branching point. That is, only fan-out (branching) points keep state information, which is mostly beneficial in a sparse distribution of members.

The clustering schemes aim to efficiently cluster members into a limited number of multicast sessions based on a preference matrix [1] or players' position in a virtual cell [4]. Note that the first two approaches (aggregation and non-branching elimination) are at the routing level, that is, trying to eliminate multicast forwarding state at each router. However, the clustering schemes are at the application level, i.e., aim at reducing the number of multicast groups using application specific information. Thus, the first two approaches can be applied to any single multicast group and the clustering schemes are for large-scale multicast applications consisting of lots of subgroups, which is our target in the paper.

The proposed TSC scheme completely eliminates creation of additional multicast groups and takes a full end-to-end approach for subgroup communication. A single multicast channel is efficiently used in the scheme for multiple roles: (1) data forwarding for the entire multicast session G, (2) providing a control channel for discovery of subgroup members and constructing a TSC forwarding structure, and (3) forwarding data to subgroup members via scoping. Unlike existing clustering frameworks, it does not require correlated information among subgroups, which eliminates the need for a central point where the information is collected and grouping decisions are made.

Overlay solution in our approach has a similarity with a number of recent application-level multicast studies, e.g., Ref. [10–14]. However, our approach should be contrasted with them. First, the goal of application-level multicast is the replacement of IP multicast due to a number of challenges such as infrastructure modification, reliability,

[1] Our mechanism targets many-to-many large-scale multicast applications where each member can be a sender and/or receiver. For such applications, it is generally agreed that shared multicast routing protocols are more efficient than source based ones. Even though PIM-SM [5], widely deployed for shared multicast routing, takes a unidirectional forwarding mechanism, we argue that bidirectional forwarding mechanisms are more efficient. The larger the multicast session and the more the demand for local communication, the larger the overhead incurred by using a unidirectional tree. Reflecting these observations, the long term inter-domain routing solution, Border Gateway Multicast Protocol(BGMP) [6] currently under development, constructs bidirectional shared trees like CBT [7].

Download English Version:

https://daneshyari.com/en/article/449863

Download Persian Version:

https://daneshyari.com/article/449863

Daneshyari.com