ELSEVIER

Computer Communications 25 (2002) 1085–1093

computer communications

www.elsevier.com/locate/comcom

Performance analysis of multicast algorithms for mobile satellite communication networks[1]

Ryan W. Thomas, Richard A. Raines*, Rusty O. Baldwin, Michael A. Temple

Department of Electrical and Computer Engineering, Air Force Institute of Technology, 2950 P St, Bldg 640, Wright-Patterson AFB, OH 45433, USA

Received 1 August 2001; revised 30 October 2001; accepted 28 November 2001

Abstract

With the rise of mobile computing and an increasing need for ubiquitous high-speed data connections, Internet-in-the-sky solutions are becoming increasingly viable. To reduce the network overhead of one-to-many transmissions, the multicast protocol has been devised. The implementation of multicast in these low earth orbit (LEO) constellations is a critical component to achieving an omnipresent network environment. This paper examines the system performance associated with two terrestrial-based multicast mobility solutions, distance vector multicast routing protocol (DVMRP) with mobile IP and on demand multicast routing protocol (ODMRP). These protocols are implemented and simulated in a satellite LEO constellation. Results from the simulation trials show the ODMRP protocol provided greater than 99% reliability in packet deliverability, at the cost of more than 8 bits of overhead for every 1 bit of data for multicast groups with multiple sources. In contrast, DVMRP proved robust and scalable, with data-to-overhead ratios increasing logarithmically with membership levels. DVMRP also had less than 70 ms of average end-to-end delay, providing stable transmissions at high loading and membership levels. Published by Elsevier Science B.V.

Keywords: Multicasting; Low earth orbit satellite; Mobile IP; Distance vector multicast routing protocol; On demand multicast routing protocol; Mobility

1. Introduction

The Internet was initially conceived to provide a means of transferring data from one machine to another. Today, it is frequently being used as a mechanism for sending the same data to many users. In an effort to provide efficient communications, multicasting was developed to ease the pressure that duplicate message transfers place on corporate, educational, and military network bandwidth constraints.

As Internet usage becomes more ubiquitous and in demand, typical connection methods will not be able to meet the demand for either bandwidth or mobility. With the advent of low earth orbit (LEO) satellite networks such as Globalstar, Teledesic, and Iridium, the potential for having a constant broadband connection to the Internet becomes more realistic.

Transmitting multicast information through a dynamic global network is a capability these future networks will need to have. However, wireless multicasting is more complex than its wired equivalent. Whereas in the wired case, the user's location and network topology is relatively fixed, in the wireless case, the location of the multicast members is constantly changing. Adding the dynamic environment of a LEO constellation removes the assumption of a static topology.

This paper examines two very antipodal protocols for LEO multicast communication, On Demand Multicast Routing Protocol (ODMRP) and Distance Vector Multicast Protocol (DVMRP), in light of these special challenges and evaluates their appropriateness for use.

2. Internet-in-the-sky

Internet-in-the-sky is a phrase used to describe the concept of creating a TCP/IP compatible network in a satellite constellation. This can either be used as part of the larger Internet when routing between nodes or exclusively, depending on the destination and source. Most Internet-in-the-sky solutions focus on LEO satellite technology for propagation delay and availability reasons.

LEO satellites orbit close to the earth in a band between 500 and 2000 km [1]. This allows for reduced power levels

* Corresponding author. Tel.: +1-937-255-3636x4715; fax: +1-937-656-4055.

E-mail address: richard.raines@afit.edu (R.A. Raines).

[1] The views expressed in this article are those of the authors and do not reflect the official policy of the United States Air Force, Department of Defense or the U.S. Government.

0140-3664/02/$ - see front matter Published by Elsevier Science B.V.
PII: S0140-3664(02)00022-1

and propagation delays as compared with traditional geosynchronous (GEO) satellite systems. LEO satellites are not stationary however, and so it takes many LEO satellites to have constant coverage of an earth location.

The large number of satellites is not as much of a drawback as it may appear. Because of the closer proximity to earth, LEO satellites can be smaller and thus launched in multiples, a strength for commercialization of the technology. Additionally, the large satellite count permits higher cell densities than provided by sparser GEO constellations.

Among operational LEO constellations, a multiplanar approach is used most frequently. This method utilizes several orbital rings of satellites, with each ring being offset by a set spacing from the next and previous ring. Communication flows around the ring and between planes via inter-satellite links (ISLs). Because of this, the topology of a LEO constellation is constantly changing.

In many polar orbit implementations, cross-planar ISLs are turned off above 60° and below −60° latitude. This is because the rate of approach of the neighboring satellites is too great to maintain connections. This makes routing for multicast protocols in a LEO constellation particularly difficult.

3. IP multicast

Traditionally, the Internet has been set up to communicate in a one-to-one (unicast) fashion. A unicast system requires a node to send individual messages for every recipient it wishes to communicate with. This works well until a single message needs to be sent to n nodes rather than just one. In this case, n copies of the message must be sent. This approach wastes bandwidth and resources associated with the communications system.

Multicast routing [2], on the other hand, sends a single message out per link, instead of sending a copy for each node accessing the information on the link. This single copy is then reproduced and branched to individual nodes wishing to receive the message by a multicast router, instead of by the sender. Broadcasting in this manner provides an efficient, lower network load solution to the one-to-many communication problem.

In 1988, Steve Deering's PhD dissertation introduced multicasting with the 'standard model' for IP multicasting being later demonstrated at the 1992 Internet Engineering Task Force meeting in San Diego [3]. It includes a set of requirements that dictate how potential multicast protocols must behave and is summarized as follows:

- IP-style semantics. Multicast messages are to be based on user datagram protocol (UDP, a best-effort policy) and require no scheduling or registration before sending. This allows a source to transmit at any time.
- Open groups. Sources are not to be limited in quantity or scope. In other words, sources can come from outside the group, and there can be any number of sources.
- Dynamic groups. Group members can join and leave at will. They do not need to 'register, synchronize, or negotiate' with a central manager.

There are two schools of thought for handling mobility under IP multicast. The first views the network as mostly static, and assumes users will change their location in that network. The task at hand is to manage mobile user locations, not worry about routing through the network. The other school of thought assumes neither the network nor the user is fixed. This view holds that the network is 'ad hoc' and past network topology has no effect on future topology.

3.1. Mobile IP

Mobile IP is the primary method of mobility management for the first viewpoint. It is designed to deal with mobile networks operating on the fringe of a fixed network. It assumes that this fixed network has the capacity to route or tunnel IP packets across it. Furthermore, being stacked somewhere between the transport layer and the IP network layer in the OSI model, it is not concerned with how these routing mechanisms work. Mobile IP is only designed to maintain a nomadic node's connection with its fixed home Internet location.

Mobile IP is described in RFC 2002–2006 [4–8], and collectively these documents present a system that allows nodes to maintain a permanent IP address wherever they go. To accomplish this, there are three primary players in Mobile IP [4]:

- The mobile node. This is a node that has the ability to change the link used to connect into the network and still maintain communication over a permanent IP address.
- The home agent. This is a router that can access the mobile node's home link as determined by the mobile node's permanent IP address. The home agent is responsible for:
  - 2.1. Knowing the mobile node's current location,
  - 2.2. Advertising its reachability in the mobile node's absence, and
  - 2.3. Intercepting and forwarding packets destined to the mobile node to the mobile node's care-of-address.

  Foreign agent. This is a router that can access the mobile nodes current link. It is responsible for:
  - 2.1. Helping the mobile node alert its home agent of it's care-of-address,
  - 2.2. If necessary, providing a care-of-address, receive packets from the home agent, and send them to the mobile node, and
  - 2.3. Routing packets generated by the mobile node to their appropriate destinations.

The care-of-address mentioned repeatedly in the agent roles is just what it sounds like. Like a vacationing tourist might have his or her mail forwarded by a servant to a temporary address where he or she is staying, mobile