

Providing Guaranteed Assurance to Connection-Oriented Group Communications Using Disjoint Routing

William Yurcik*

*Illinois State University
Dept. of Applied Computer Science
wjyurci@ilstu.edu*

David Tipper

*University of Pittsburgh
Dept. of Information Science & Telecommunications
tipper@tele.pitt.edu*

Abstract

This paper compares different approaches which provide guaranteed assurance to connection-oriented group communications based on the use of working and backup disjoint route sets. Specifically, we present experimental results showing the effect of disjoint backup route sets on mesh/tree/ring: (1) feasibility and (2) cost.

1. Introduction

Providing group communications over a network subject to failure and attack is a problem of growing interest. Multicast mechanisms provide group communications by reducing the amount of duplicate traffic in a network to conserve bandwidth and switch resources. In order to manage large bandwidth multimedia feeds and to slow congestion by eliminating waste, multicasting concentrates resources to gain efficiency. However, this concentration amplifies the impact of failures such that a random fault or malicious attack will affect a large number of users.

The connectionless IP multicast model is based on the abstract concept of multicast group addresses. A sender transmits a packet to a multicast group address and underlying layers route it transparently. Establishing a multicast tree spanning the group members has emerged from the Internet community as one efficient solution for IP multicast.[2] There have been several types of delivery trees proposed for different IP multicast algorithms. These trees are either source-based trees, whereby the source is built from an active sender, or shared tree, which disseminate multicast traffic using one delivery tree spanning all group members.

Supporting connection-oriented group communications is more difficult than connectionless

* corresponding author; additional contact information: voice/fax 309-438-8016/5113, hard copy: Campus Box 5150, 202 Old Union, Normal IL 61790 USA; supported in part by NASA Earth Systems Science grant # NGT-30019, Defense Advanced Research Projects Agency grant # F30602-97-1-0257, and SAE International - The Engineering Society for Advancing Mobility Land Sea Air and Space.

group communications because: (1) Steiner Tree heuristics developed for bidirected networks do not distinguish sending nodes from receiving nodes, and (2) connection-oriented communication requires setup establishment which is not compatible with group dynamics of participants joining and leaving. New Internet schemes to provide QoS guarantees based on end-to-end allocation of resources (e.g., RSVP) are converging connectionless IP multicast solutions into connection-oriented group communications solutions such as we describe in this paper.

The focus of our investigation is the extension of point-to-point restoration techniques to the assurance of connection-oriented group communications. The basic concept is each group communication session requiring assurance will establish a disjoint backup set of routes with reserved bandwidth for restoration. When a fault occurs within a "working" set of routes, traffic flow is guaranteed assurance by being rerouted to the corresponding disjoint "backup" set of routes under a specified timing threshold such that any single link or node fault (link or node) are transparent to users at the application layer.

The remainder of the paper is organized as follows: Section 2 reports comparative experimental results on the relative feasibility and cost of the proposed approaches for connection-oriented group communications. Section 3 briefly summarizes conclusions.

2. Feasibility and Cost

To compare the feasibility and cost of using disjoint working and backup route sets to provide single-fault assurance to connection-oriented group communications, we modeled two actual network topologies.[4,5] We then considered three different approaches for connection-oriented group communications: (1) MESH, (2) Shared Tree (ST), and (3) RING and formed all groups of size three/four/five for each network.

The MESH approach specifies that each sender establishes circuit(s) to connect all other group members but does not specify the mechanism so the "best" case

scenario was selected such that the sender optimally selects the minimum cost set of point-to-multipoint/ point-to-point circuits that connect all group members. The ST approach leaves the issue of how to build the source-based shared multicast tree to an external routing protocol so the “best” case scenario identifies the minimum cost tree using an implementation variant of a Steiner Tree procedure due to Lawler known as the spanning tree enumeration algorithm.[4] Minimum cost working and backup rings are calculated by an implementation of the DSR formulation shown in [4]. In the RING approach, a real-time preplanned reconfiguration around a single fault (similar to FDDI or SONET) can be achieved such that working and backup RINGS do not need to be disjoint.

For the networks studied, the feasibility of restoration using the different approaches is not equal. While the RING displayed 100% feasibility on the networks tested, the feasibility of the ST approach peaks at 75%. The feasibility of both the ST and MESH approaches are limited by the availability of disjoint backup route sets and both rapidly decrease in feasibility as group size increases. The explanation for the poor feasibility of the ST and MESH approaches is the existence of “multipoint traps”. A “trap” is a topology where a corresponding set of backup routes are not available due to the disjointness constraint although disjoint working and backup routes may be available if selected differently.[3]

Before this investigation, it was assumed that the ST approach would be least cost since it is well-known that the Steiner Tree solution is least cost. However, we found that the cost of a disjoint backup ST is 150% greater than the cost of a working ST (when a disjoint backup is available) and thus the total cost of combined working/disjoint backup STs can be greater than the working/backup rings. Experimental results show that the cost of the RING approach is lower, with statistical significance, than either the ST or MESH approaches, with a cost differential that increases as the group size increases. The size of the cost differential is not small; RINGS are 20% less expensive than STs and 56% less expensive than MESHes.

3.0 Conclusions

With the commercialization of the Internet, the need for assured group communications is growing. This paper has explored the use of different preplanned assurance approaches for connection-oriented group communications. From our preliminary results we conclude the following:

- connection-oriented group communications can be characterized as either mesh, tree, or ring.

- the preplanned ring approach is often more feasible and less costly than either preplanned tree or mesh approaches
- topology (in terms of node degree and link density architecture) is a determining factor in survivability
- while point-to-point “traps” exist, “traps” have a greater impact on group communications

The routing scalability of these approaches is an important issue to address. Routing complexity may dictate that the exact implementation of these approaches on real networks will need to be modified. A large number of working routes additionally combined with backup routes for each may give further advantage to the ring approach which can define just one bidirectional ring of routes for each group session.

Attacks to oppose the type of group communications assurance we have presented include group eavesdropping and attacks on the dedicated bandwidth backups.[1] Encryption techniques to prevent eavesdropping and active network monitoring for redundant links (that may be “dark”) to prevent attacks on backups are being developed.

4.0 Acknowledgments

This paper is a summary of Ph.D. research in progress. We would like to specially acknowledge Deep Medhi/University of Missouri-Kansas City who has also made significant contributions to this work.

5.0 References

- [1] R. Canetti and B. Pinkas, “A Taxonomy of Multicast Security Issues,” IETF Draft draft-irtf-smug-taxonomy-01.txt, August 2000.
- [2] S. Deering and D. Cheriton, “Multicast Routing in Datagram Internetworks and Extended LANs,” ACM Trans. On Computer Systems, Vol. 8, No. 2, pp. 85-110.
- [3] D. Dunn et. al., “Comparison of k-Shortest Paths and Maximum Flow Routing for Network Facility Restoration,” IEEE J. on Sel. Areas in Comm., Vol. 12, No. 1, pp. 88-99.
- [4] W. Yurcik and D. Tipper, “Providing Network Survivability to ATM Group Communications Via Self-Healing Survivable Rings,” 7th Intl. Conf. On Telecomm. Systems, 1999, pp. 501-518.
- [5] W. Yurcik, “Providing ATM Multipoint Survivability Via Disjoint VC Mesh Backup Groups,” 7th IEEE Intl. Conf. On Computer Comm. And Networks, 1998, pp. 129-136.