SCALABLE SURVIVABLE ATM GROUP COMMUNICATIONS

William Yurcik *
Department of Applied Computer Science
Illinois State University

David Tipper †
Department of Information Science and Telecommunications
University of Pittsburgh

Deepankar Medhi ‡
Department of Computer Networking
University of Missouri - Kansas City

Abstract

Group communications is important for command and control as well as tactical battlefield operations. Providing group communications over a network subject to failure and attack is a problem of growing interest. We examine techniques to provide survivability for ATM group communications and study the scalability of such approaches. A hop-limit constraint is a technique used to limit the number of routes considered such that routing problems of higher order complexity can be solved. While varying the hop-limit, we compare the feasibility and cost of providing survivability using working multipoint routes with disjoint dedicated backup multipoint routes, where the multipoint routes are setup using either Self-Healing Survivable Rings, Shared Multicast Trees, or VC Meshes. Based on our results, we conclude with a hop-limit heuristic that can be used in formulations to provide scalable survivable group communications.

1 Introduction

Network survivability is important due to the increasing number of information systems and society's increasing dependence on these systems for dependable service. Solutions to provide survivable ATM group communications via underlying multicasting mechanisms can take different approaches. One approach is a working set of routes protected by a standby backup set of routes. When a single fault occurs within a working set of routes, traffic flow is rerouted to the corresponding backup set of routes. The working and backup set of routes are link and node disjoint such that a single fault will not disrupt both working and backup routes simultaneously. For a typical network, finding the working and disjoint backup set of routes and provisioning dedicated bandwidth is a NP-hard integer programming optimization problem.

This research focuses on how the use of a hop-limit constraint to limit the number of routes considered effects the quality of integer optimization solutions. A hop-limit constraint is a technique used to: (1) limit the number of paths considered in network design and routing problems and (2) limit the length of routes considered to meet Quality of Service (QoS) guarantees (especially delay) [4, 5, 1, 9, 6, 7]. Experimental results from [5] report that the hop-limit approach is indeed sound as it restrains model dimensions and at the same time ensures QoS guarantees can be met. In fact [5] reports that a relatively small hop-limit provides

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

† *tipper@tele.pitt.edu*, supported in part by Defense Advanced Research Projects Agency grant # F30602-97-1-0257

‡ *dmedhi@cstp.umkc.edu*, supported in part by Defense Advanced Research Projects Agency grant # F30602-97-1-0257

0-7803-6521-6/$10.00 (C) 2000 IEEE

effective solutions while minimizing computations such that there is no real gain for setting larger hop-limits.

Several hop-limit heuristics have been proposed for providing survivability in the point-to-point context. Itai et. al., reported the first use of hop-limits combined with an algorithm to effectively calculate the number of disjoint routes if the hop-limit is less than or equal to three [7]. It is shown that except for small hop-limit values, the problem of finding disjoint paths in a general graph is NP-complete. Oki and Yamanaka propose a recursive matrix calculation approach to calculate the number of disjoint routes below a hop-limit equal to four [9]. Herzberg et. al, propose the "Extra-Hop Limit" heuristic which is the network diameter[1] in hops plus an additional hop count limit [4].

We consider the effects of hop-limits on scaling survivable ATM group communication methods to real networks. The remainder of this paper is organized as follows: Section 2 introduces proposed schemes for ATM group communications. Section 3 discusses survivability techniques specific to these proposed schemes. Section 4 presents results from numerical experiments focused on how varying the hop-limit constraint effects the quality of multipoint survivability optimization problem solutions. We close with a summary of conclusions and future research directions in Section 5.

2 Proposed ATM Group Communication Schemes

The ATM Forum MPOA (Multi-Protocol Over ATM) subworking group proposes two current approaches for ATM intracluster[2] group communications: (1) the VC Mesh Model and (2) the Multicast Server Model (MCS). As a result of the perceived complexity and inefficiency of these two approaches, other techniques have been proposed such as shared trees and rings.

[1]network diameter is defined as the maximum shortest-path between any two nodes in the network

[2]A cluster is defined as a set of ATM interfaces choosing to participate in native ATM routing. Traffic between ATM interfaces belonging to different clusters pass through an intercluster device.

2.1 VC Mesh Model

The VC Mesh Model derives its name from the criss-crossing of VCs across a network. Each sender originates a unidirectional VC to all members of the group and each group member must terminate one VC for each active source in the group. If group membership changes (add/delete members), each point-to-multipoint VC for each sender must be released and a new point-to-multipoint VC set-up. The VC Mesh is attractive because it can currently be implemented with commercial switches since the advent of the point-to-multipoint VC capabilities. However, the management of a VC Mesh is difficult due to its non-scalable signaling and its lack of support for group dynamics.

2.2 Multicast Server Model

The MCS Model is directly analogous to a conference bridge telephone session. A server is chosen to serve as a proxy for all senders in order to relieve senders from direct VC connection set and release operations due to group dynamics. All active senders establish a point-to-point VC to the MCS and the MCS establishes a point-to-multipoint VC to all members of the group. All requests to create a group, delete a group, to add a group member, or delete a group member are sent to the MCS which maintains state information for all existing groups. MCS implementations exist but are not in common use. The MCS is included in this section for completeness but will not be included in further discussion or experimentation due to its dominant single-point-of-failure vulnerability which requires a different set of restoration techniques. For a more detailed discussion of the MCS see [11].

2.3 Shared Tree

We consider two specific ATM shared multicast tree schemes: SMART and SEAM. [2, 3]. These schemes are similar in that they try to provide a general-purpose control architecture by modifying in-band control mechanisms of ATM switches. Resources are reserved in both directions on the many-to-many VC links of a shared multicast tree until the connections are released. Both schemes address two inherent problems of a shared tree architecture: (1) cell interleaving because cells from different sources may arrive interleaved at one destination and (2) resource management because resources allocated to a connection are shared between a number of different sources.

2.4 Virtual Ring

Ofek and Yener have proposed the Virtual Ring [8] for window-based packet group communications. The Virtual Ring is a unidirectional circular overlay of routes that join all desired members for a group communication. Using Virtual Rings is straightforward since a message sent from one member of the group travels around the ring and back to the sender minimizing the feedback needed for reliability and overcoming the "ACK implosion effect". In the application of Virtual Ring multicasting to ATM networks, each unicast connection between adjacent nodes can be implemented as concatenated point-to-point VCs to form an enclosed ring.

3 Providing Survivability To ATM Group Communications

The preplanned end-to-end technique to provide survivability to the VC Mesh Model is the simultaneous establishment of both a working VC and a disjoint backup VC pair for each sender. When a fault disrupts the working VC, traffic is rerouted to the disjoint backup VC. The problem with this technique is the large number of VCs which must be established and managed. There is also a potentially larger number of VCs over which a failure may occur which means an increased exposure to link and node failures. A more detailed discussion of providing survivability to the VC Mesh Model can be found in [11].

For tree-based group communication, the dominant survivability issue is handling faults in the "trunk" of the shared tree which will interrupt all group communications (a fault in a "branch" of the shared tree will only disconnect subsets of group members). A preplanned end-to-end technique to provide survivability to tree-based approaches is to restore faults on the working tree by rerouting to a preplanned disjoint backup tree. This rerouting of traffic from a working tree to a disjoint backup tree will entail a complex signaling protocol. On the other hand, sharing a common tree makes the identification and protection of links (especially the trunk) and authentication of traffic a simpler task.

Virtual Ring multicasting is extended to provide survivability for ATM group communications in [10] which describes: (1) a real-time ring reconfiguration mechanism termed the Self-Healing Survivable Ring to provide real-time restoration of single link or single node faults while maintaining the established Virtual Ring architecture and (2) a formulation of Disjoint Steiner Ring (DSR)[4] connecting all desired group members, that is *link and node disjoint*, such that a single link or node fault will only disconnect a ring in one place. The major ramification of a ring reconfiguration to provide survivability is an increase in the maximum end-to-end delay: the number of links a cell traverses in the reconfigured Virtual Ring will increase from L to $2(L-1)$ in the worst case single link failure scenario.

4 Results

Here we report the results of experimentation to compare techniques to provide survivable ATM group communications formulated under hop-limit constraints. The techniques compared are the VC Mesh Model (VCMESH), the Shared Multicast Tree (SMT), and the Self-Healing Survivable Ring (SHSR). We restrict our investigation to single fault scenarios.

We consider two actual network topologies as shown in Figure 1. Generating all possible paths between every node-pair up to the inclusion of all network nodes results in 16,178 routes for NET-1 and 63,380 paths for NET-2. The effect of varying a hop-limit to limit the search space in terms of the number of paths shows a characteristic curve common to all networks as shown in Figures 2 (despite different topologies and axis labels).

To analyze the effect of a varied hop-limit on the different survivable multicast techniques (VCMESH, SMT, and SHSR) we formed all the possible survivable groups of size three for each network. The VCMESH approach specifies that each sender establishes circuit(s) to connect with all other group members but does not specify the mechanism so we have assumed the best case scenario such that each sender optimally selects the

[4] A variant of the "Steiner Tree Problem in Graphs" (find the disjoint *least cost* ring in a general graph - A Steiner Ring - joining a subset of specified nodes).

set of point-to-multipoint/point-to-point circuits that minimizes cost.[5] The SMT approaches leave the issue of how to build the shared tree to an external routing protocol so we have assumed the best case scenario by identifying the minimum cost trees using a implementation variant of a Steiner Tree procedure due to Lawler known as the spanning tree enumeration algorithm [10]. Minimum cost working and backup rings are calculated by an implementation of the DSR formulation shown in [10].

Both the SMT and DSR depend on solving variants of the Steiner Tree Problem which is well-known to be NP-complete.[6] However, both the spanning tree enumeration algorithm and the DSR formulation are able to exactly solve SMT and SHSR ring solutions on actual networks by: (1) limiting the initial search space with preprocessed routes equal to or below a hop-limit and (2) using branch-and-bound within each algorithm implementation to further narrow the search space. While scalability problems remain, the hop-limit constraint combined with brand-and-bound can make previously intractable problems tractable by restricting the exponential rate of growth due to combinatorics. Figure 3 shows the computational complexity of solving the Steiner Tree Problem in terms of Steiner Trees and Steiner Rings on NET-1 and NET-2. The impact of a hop-limit constraint is most dramatic when considering DSR computational complexity. On NET-1, the computational complexity of the DSR solution is 10^4 Mflops[7] without a hop-limit constraint and $2\ X\ 10^1$ Mflops with a hop-limit constraint. On NET-2, the computational complexity of the DSR solution without a hop-limit was unmeasured despite lengthy attempts and $4\ X\ 10^2$ Mflops with a hop-limit constraint. It should be noted that all experiments were executed on a Sun Enterprise 4000 with ten 250 MHz ULTRASparc-II CPU's and 2.5 GB of RAM.

[5]each sender uses a least-cost routing algorithm to independently setup their VC MESH set of circuits

[6]Finding the least cost VCMESH solution is also a variant of the Steiner Tree Problem but since it is so closely related to the Steiner Tree solution it is not mentioned separately in discussion here for clarity although its computational complexity is shown separately in Figure 3.

[7]Flops are the cumulative number of executed floating point operations as measured by implementation in MATLAB version 5 (Mflops = a million flops).

As the hop-limit increases, the feasibility (i.e., ability to find both a working and disjoint backup) of all techniques increases without exception. The intuitive explanation is that some disjoint backup routes may be long and thus ineligible at lower hop-limits but become revealed as the hop-limit increases. Figure 4 shows the SHSR technique converges to 100% feasibility at a hop-limit threshold close to the network diameter on both networks while both the SMT and VCMESH techniques take more hops to converge to final feasibility levels significantly less than 100%. This difference in feasibility was first reported in [11] for the VCMESH technique and is due to 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. Traps occur because the routing algorithm for the working set of routes optimizes selection according to least cost/minimum hop or similar metric without considering the survivability provided by the selection of a non-optimal set of working routes that can be paired with a disjoint set of backup routes.

According to the "Extra-Hop Limit" heuristic proposed by Herzberg et. al., (network diameter + 2), the hop-limit would be 9 and 10 for NET-1 and NET-2 respectively. Closer examination of Figure 4 shows that these hop-limits are just at or above the knee of each curve and in some cases leave a substantial percentage of groups unable to be restored. A new hop-limit heuristic is needed specifically for the multipoint survivability context.

We define cost in terms of aggregate cost for all links involved with each link being assigned a cost weight (initially all links are assigned a cost weight of 1). As a general rule, the cost of all techniques increases as the hop-limit increases since new survivable groups that may not have previously been feasible suddenly emerge as viable with a larger hop-limit. Figure 5 shows the cost of the SHSR technique is significantly lower, at a 95% confidence level, than either the SMT and VCMESH techniques. The size of the cost differential is not small, SHSRs are about 20% less expensive than SMTs and 56% less expensive than VCMESHes for a group size of three [10]. Results from other experiments report that the ranking order of each technique by cost is consistent across different networks, group sizes, and link cost assignments: SHSR (lowest cost), SMT (middle cost), and VCMESH (highest cost) [10].

The cost of a backup ring is the same as the cost of a working ring in SHSR under no capacity constraints. It is expected that the VCMESH would be the highest cost solution due to the number of circuits required but it is non-intuitive that the SMT technique would be more costly than the SHSR technique since a Steiner Tree is the least cost solution. However, by examing numerical results one can see that although the working SMT is the least cost solution compared to its corresponding working SHSR, the corresponding disjoint backup SMT is about 160% more expensive. This same effect is amplified for the VCMESH technique.

Ultimately there is a delay tradeoff between the ring and tree-based approach. Figure 6 shows a direct comparison of the three techniques under a delay-constraint threshold (as measured in hops) on both networks. For each network, survivability for all groups of three is formulated using rings (SHSR), trees (SMT), and circuits (VCMESH) with the same worst-case delay bound and the percentage of restorable groups compared. Note that the SHSR technique does not become more feasible until the worst-case delay-constraint is more than the total number of hops in each network respectively. These results show that the superior feasibility and cost of survivability provided by the SHSR technique must be weighed against an increased worst-case delay guarantee.

5 Summary

This paper has examined the effect of a hop-limit constraint on techniques to provide survivability for ATM group communications. A hop-limit reduces the number of routes considered such that we were able to exactly solve NP-complete optimization problems to find minimum cost disjoint working and backup paths on actual networks. We found that multipoint restoration solutions are more sensitive to topology (particularly a multipoint trap topology) than to a hop-limit constraint. Results from the networks studied show that the hop-limit heuristics developed for the point-to-point restoration context are not sufficient to provide exact solutions for the multipoint restoration context.

We propose a conservative hop-limit heuristic of $\frac{N}{2}$ for the multipoint survivability context where N is the number of nodes in the network and $\frac{N}{2}$ is larger than the network diameter. This conservative hop-limit provides exact solutions for NP-complete multipoint survivability problems on real networks that would otherwise be intractable. Future research is needed to develop a lower hop-limit heuristic that would still provide exact solutions.

References

[1] S. Choudhury and L.T.M. Berry, "Two Methods for Minimum Cost Synthesis of Link-Disjoint, Hop-Limited Virtual Paths in ATM Networks," *16th Intl. Teletraffic Conf.* 1999, pp. 613-622.

[2] E. Gauthier, J-Y. Le Boudec, and P. Oechslin, "SMART: A Many-To-Many Multicast Protocol for ATM," *IEEE J. on Sel. Areas in Comm.*, Vol. 15, No. 3, April 1997, pp. 458-472.

[3] M. Grossglauser and K.K. Ramakrishnan, "SEAM: Scalable and Efficient ATM Multicast," *IEEE INFOCOM'97.*

[4] M. Herzberg, D.J. Wells, and A. Herschtal, "Optimal Resource Allocation for Path Restoration in Mesh-Type Self-Healing Networks," *15th ITC*, 1997, pp. 351-360.

[5] M. Herzberg, S.J. Bye, and A. Utano, "The Hop-Limit Approach for Spare-Capacity Assignment in Survivable Networks," *IEEE/ACM Trans. on Networking*, Vol. 3, No. 6, pp. 775-784.

[6] M. Herzberg and S.J. Bye, "An Optimal Spare-Capacity Assignment Model for Survivable Networks with Hop-Limits," *IEEE GlobeCom'94*, pp. 1601-1606.

[7] A. Itai, Y. Perl, and Y. Shiloach, "The Complexity of Finding Maximum Disjoint Paths with Length Constraints," *Networks*, Vol. 12, 1982, pp. 277-286.

[8] Y. Ofek and B. Yener, "Reliable Concurrent Multicast From Bursty Sources," *IEEE J. on Sel. Areas in Comm.*, Vol. 15, No. 3, April 1997, pp. 434-444.

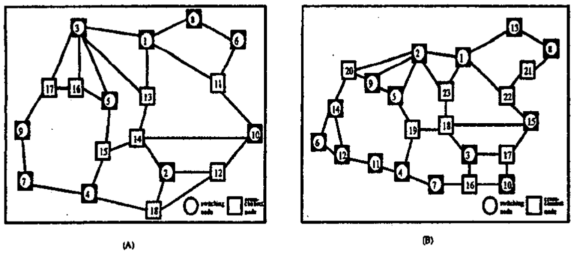

Figure 1: *(A) NET-1; (B) NET-2; (networks used in [11, 10])*

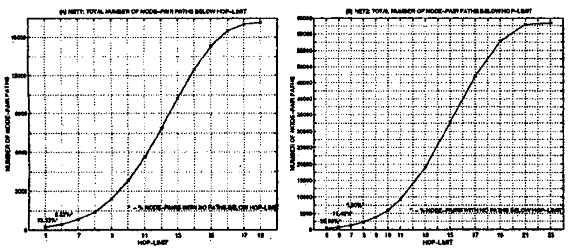

Figure 2: *Total Node-Pair Paths as a Function of Hop-Limit*

[9] Oki, E. and N. Yamanaka, "A Recursive Matrix-Calculation Method for Disjoint Path Search with Hop Link Number Constraints," *IEICE Trans. Comm.*, Vol. E78-B, No. 5, pp. 769-774.

[10] 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.

[11] W. Yurcik, "Providing ATM Multipoint Survivability Via Disjoint VC Mesh Backup Groups," *7th IEEE Intl. Conf. on Comp. Comm. and Networks (IC3N'98)*, pp. 129-136.

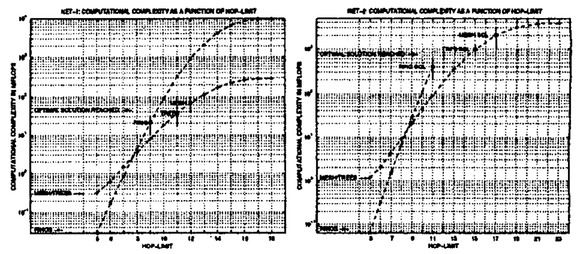

Figure 3: *Computational Complexity as a Function of Hop-Limit*

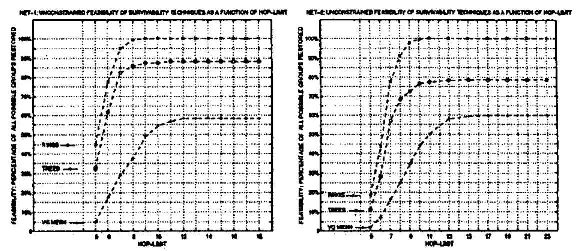

Figure 4: *Feasibility of Survivability Techniques as a Function of Hop-Limit*

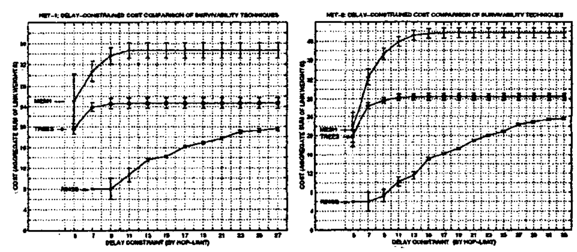

Figure 5: *Delay-Constrained Cost of Survivability Techniques as a Function of Hop-Limit*

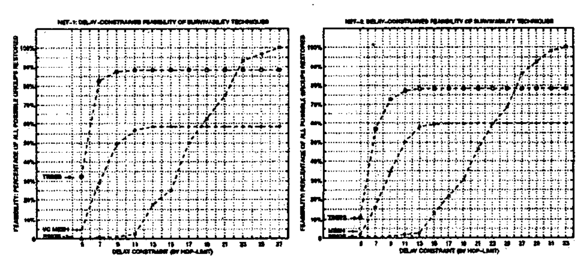

Figure 6: *Delay-Constrained Feasibility of Survivability Techniques as a Function of Hop-Limit*