Providing ATM Multipoint Survivability Via Disjoint VC Mesh Backup Groups

William Yurcik *
Department of Information Science and Telecommunications
University of Pittsburgh
Pittsburgh, PA. 15260 USA
yurcik@tele.pitt.edu

Abstract

This paper presents simulation experiments investigating the feasibility and cost of providing survivability to ATM multipoint groups using dedicated bandwidth on disjoint backup groups. Disjoint dedicated backup groups are formed via the VC Mesh model by two different approaches: (1) a backup group consisting exclusively of point-to-point circuits (unicast), or (2) a backup group consisting exclusively of point-to-multipoint circuits (multicast). Results for the two networks studied show that although backup groups of point-to-multipoint circuits (multicast) are lower in cost than backup groups of point-to-point circuits (unicast) at a 95% confidence level, the feasibility of providing guaranteed survivability with backup groups via either approach begins at only about 50% for a group size of three and decreases rapidly for larger group sizes due to multipoint "trap" topologies and hop-limit constraints. The conclusion is that while use of disjoint dedicated backup groups for ATM multipoint may be viable in certain dense networks with rich connectivity, based on the two networks studied neither approach provides acceptable survivability guarantees for arbitrary networks especially noting that network congestion is not considered.

* supported in part by NASA Earth Systems Science grant # NGT-30019 and SAE International - The Engineering Society for Advancing Mobility Land Sea Air and Space; additional contact information: telephone (412) 624-9411, FAX (412) 624-2788

1 Introduction

The importance of multipoint communications[1] and the need to efficiently and reliably support it within a network spans Internet-style networking and B-ISDN technology including ATM. The majority of work on providing survivability for ATM networks focuses on extending circuit-switched facility restoration techniques to ATM using Virtual Paths (VPs). VPs can be viewed as semi-permanent routes through the network with dedicated bandwidth onto which Virtual Circuits (VCs) are grouped. VP restoration is simpler and faster than VC restoration and has been the focus of ATM restoration systems currently proposed in the literature. Several different algorithms for ATM network survivability have been proposed using either a pre-defined backup VP for each working VP and/or dynamically searching for new paths(s) after a single link failure. [10, 5] Note that the bandwidth in the predefined backup VP is reserved and not used until a fault occurs and traffic is rerouted. Several different algorithms to efficiently provision network capacity to provide the necessary reserved backup capacity at the VP layer and lower layers (e.g., transmission layer) have been proposed. [13, 8]

While network survivability has focused on point-to-point communications, there has been very little work in applying the concepts of ATM point-to-point survivability to ATM multipoint communications. To my knowledge, there has been only one academic paper on this topic. [14] The

[1] a note on terminology: Multipoint communications will be used to describe any type of communication service within a group of multiple participants (it is equivalent to group communication). Multicast will be used to describe the underlying point-to-multipoint directional transmission mechanism (i.e., delivery routing tree, multipoint circuit) which is used to create a multipoint service. Thus, multipoint has a more general connotation and will refer to the group communication network service provided to the user. In the IP datagram realm since there is no long term association between source and destination, the terminology often used to refer to the underlying transmission mechanisms are unicast (one-to-one), multicast (any-to-many), concast (many-to-one), anycast (one-to-any), and broadcast (any-to-all).

same types of vulnerabilities present in ATM unicast are also present in ATM multicast but the potential risk of these vulnerabilities is considerably greater in multipoint. For instance, ATM multicast traffic traverses potentially many more physical links than does a single unicast communication, thereby creating more of an opportunity to encounter link or node failures. To exacerbate this vulnerability, it is often more complex to determine and control which links are in use for a given multicast routing scheme as compared to the equivalent point-to-point communication. In addition, there are new vulnerabilities unique to proposed ATM which have no counterpart in ATM point-to-point such as group access control. [4]

Supporting multipoint communication in ATM networks provides its own set of unique challenges since ATM is inherently connection-oriented. Point-to-multipoint VCs were introduced in the ATM Forum's User Network Interface (UNI) Specification 3.1 in 1995. The two current models proposed for intracluster multicast by the ATM Forum MPOA (Multi Protocol Over ATM) subworking Group are: (1) the VC Mesh, and (2) the Multicast Server (MCS). The choice of either model may be made on a per-group basis and is transparent to ATM users.[2]

Preplanned, dedicated bandwidth, disjoint backup route techniques have proved to be very effective for the restoration of ATM point-to-point VPs and VCs. This paper investigates the extension of these point-to-point techniques to the restoration of ATM group communications implemented via the VC Mesh model. When a single link or single node fault occurs within a working group, traffic flow is rerouted to a corresponding backup group that is link and node disjoint with dedicated bandwidth reserved for restoration.

The remainder of this paper is organized as follows: Section 2 describes the VC Mesh ATM multipoint model. Sections 3 and 4 describe the network model and experimental design respectively. Section 5 presents the experimental results. The paper closes with a summary of conclusions and future research directions in Section 6.

2 The VC Mesh Model

In the VC Mesh model proposed by the ATM Forum, a unidirectional VC originates from each sender to all members of a group. Each group member must terminate one VC for each active source in the cluster. Each sender is essentially the root of its own source-based tree. If a member joins or leaves the group, each point-to-multipoint VC for each sender in the group must be modified (a signaling load that scales with the number of senders in a group). In the ATM Forum UNI 3.1 specification, the sender keeps track of all members in the connection in all cases, only the sender can add a new group member and either the sender can delete a member or the member can delete itself from the point-to-multipoint VC.[2] The ATM Forum UNI 4.0 specification allows a new member to join a point-to-multipoint VC without the intervention of the sender by using a Leaf-Initiated Join (LIJ) capability.[3] Receivers discover the existence of point-to-multipoint VCs through a directory and are able to signal a request to join a particular point-to-multipoint VC. LIJ does not specify where the new member should be joined to the existing point-to-multipoint VC. For example, the new member (leaf) may join at the nearest node which is already part of the point-to-multipoint VC or it may join via the shortest path to the sender (root).

For a single group containing multiple senders, a point-to-multipoint VC for each sender is established on the network with no coordination between VCs except for capacity allocation on shared links between senders. For multiple groups on the same network, a point-to-multipoint VC for each sender is handled independently except for capacity allocation on shared links between groups. This criss-crossing of VCs across a network gives rise to the name - “VC Mesh”.[1]

In the MCS model, a server is chosen within each cluster to serve as a proxy for all senders in order to relieve senders from direct VC connection setup and release operations resulting from group dynamics. Conceptually the MCS serves as a centralized connection manager mediating and disassociating senders and receivers within a group. [15] All senders establish a unidirectional point-to-point VC to the MCS. The MCS establishes a unidirectional point-to-multipoint VC to the rest of the multicast group.

The VC Mesh model has a number of survivability issues that need to be addressed. In the VC Mesh model, group dynamics drive the establishment of a new point-to-multipoint VC for each sender in a group to reflect updated membership after each join or leave operation.[3] Also in the VC Mesh model, multiple senders transmitting to the same group members use separate point-to-multipoint VCs while in the MCS model multiple senders each establish point-to-point VCs to the MCS and then the MCS has the capability to integrate point-to-multipoint transmissions to the same group members. For both of these reasons, the number of VCs utilized in the VC Mesh model increases with the number of senders at a faster rate than the MCS model. If all

[2] An ATM user is broadly defined as any entity that establishes and manages VCs and underlying ATM Adaption Layer (AAL) services to exchange data via a User-Network Interface (UNI).

[3] applies to UNI 3.1 specification without LIJ capability; the intention of the LIJ capability is to avoid all parties being aware of group dynamics but LIJ still may have a similar property depending on implementation

group members are active senders at some point over the lifetime of a particular group, then the resulting paths will be a fully-connected mesh. As a result, there is potentially a larger number of routes over which a failure may occur for the VC Mesh which means an increased exposure to link and node failures.

Other survivability issues specific to the VC Mesh model include increased signaling traffic due to more VCs and increased termination requirements (end hosts need one VC termination per source per each ATM interface). These issues both represent performance limitations due to lack of model scalability. As group sizes grow, providing survivability becomes more difficult due to the increased capacity required for signaling and the exhaustion of end host ports due to previous sender transmissions.

The VC Mesh model was chosen to investigate over the MCS model because of the obvious and dominant single-point-of-failure vulnerability of the MCS. Although the MCS constrains the number of potential paths over which a failure may occur, the MCS accomplishes this by concentrating traffic on a server which is both a performance bottleneck and single point-of-failure for a group's communication. If an MCS is disabled, all group sessions using that MCS terminate.

The failure of individual senders (or receivers) within a group will degrade but not terminate group communications if there are other senders (or receivers) still operating. We will not be addressing strategies to recover from individual sender or receiver failures but rather focus on the survivability of a group's communication.

3 Network Model

We consider the two actual network configurations shown in Figures 1A and 1B on which to perform experiments. Both configurations include switching nodes that can originate and terminate traffic, as well as cross-connect nodes that relay traffic. NET-1 has an average node degree of 3.00 and a network diameter of 6 links where the network diameter is defined as maximum shortest-path between any two nodes in the network. NET-2 has an average node degree of 2.87 and a network diameter of 7 links.

The underlying network models for NET-1 and NET-2 are composed of a set of nodes $n \subset N$, a set of bidirected links $l \subset L$, and an incidence matrix specifying how the nodes are connected by links. Each link connects two nodes for traffic flows in two directions. A link fault stops traffic flow between two nodes in both directions. A node fault stops traffic flow in both directions on all incident links. A path between nodes n_{source} and $n_{receiver}$ is a sequence of links (of length k) of the form $p = \{l_1, l_2, l_3, \ldots, l_k\}$ where $\{l_1, l_2, l_3, \ldots, l_k\}$ are distinct (no cycles). A path between nodes n_{source} and $n_{receiver}$ can equivalently be represented as a sequence of nodes of the form $p = \{n_{source}, \ldots, n_{receiver}\}$ where $\{n_{source}, \ldots, n_{receiver}\}$ are distinct. Two paths are link disjoint if their set of links are disjoint. Two paths are node disjoint if their set of nodes are disjoint except for the first (n_{source}) and last ($n_{receiver}$) nodes in the set.

Generating all possible paths between every node pair up to the inclusion of all network nodes results in approximately 16,000 paths for NET-1 and 65,000 paths for NET-2. A path length or "hop-limit" heuristic of half the network size as measured in nodes was adopted for computational feasibility and to acknowledge quality of service (QoS) limitations such as delay. This results in a hop-limit on paths to less than or equal to 9 nodes or 8 links for NET-1 and 11 nodes or 10 links for NET-2. Within this given hop-limit: for NET-1 the average number of eligible paths between node pairs is 27.69 and the average length of each of these eligible paths is 2.97 links; for NET-2 the average number of eligible paths between node pairs is 44.70 and the average length of each of these eligible paths is 3.49 links.

4 Experimental Design

The purpose of experimentation is to test the extension of preplanned dedicated backup route restoration techniques from ATM point-to-point communications to ATM group communications implemented via the VC Mesh model.[4] Disjoint dedicated backup groups are formed via the VC Mesh model by two different approaches: (1) a backup group consisting exclusively of point-to-point circuits (unicast), or (2) a backup group consisting exclusively of point-to-multipoint circuits (multicast). The term "working" refers to an individual circuit or a group of circuits initially established for communication. The term "backup" refers to an individual circuit or a group of circuits that is reserved for restoration and not used otherwise (i.e., dedicated bandwidth). When a fault occurs with a link or node within a working group, traffic flow is rerouted to a corresponding backup group. The backup group is link and node disjoint such that a single fault in the working group will not affect the backup group. The presentation order of these experiments is designed to show a logical progression from the creation of working groups to the subsequent protection of these working groups by backup groups which have been

[4]These experiments are also directly applicable to the restoration of ATM group communications implemented via the MCS model since circuits from the receiver-side of the MCS to all receivers is a logical VC Mesh model.

formed using different circuit types.

For our purposes, we make the following assumptions to enable direct comparisons: (1) the VC Mesh ATM model of forming a group is implemented; (2) only single faults (single link/node) in a single group are considered; (3) a one-to-one correspondence between each working group and a backup group such that 100% survivability can be guaranteed (i.e., no sharing of backup bandwidth reserved for restoration); (4) all links have equal cost which implies that the number of links involved in a group is linearly related to cost and that links have equal capacity; (5) the bandwidth of any traffic flow is less than or equal to link capacity; (5) all traffic flow within a group is symmetric with equal bandwidth and QoS requirements; and (6) all groups have equal traffic flow. It should be noted that these results can be extended to multiple failure scenarios, other proposed ATM models of forming groups, asymmetric group traffic, and different traffic flows for different groups as will be addressed in future research directions in the final section.

Table 1 shows the group sizes for the experiments simulated on NET-1 and NET-2. These group sizes are consistent with the range of multipoint group size relative to network size suggested in the literature (i.e., anywhere from 10-60%). A given group size for different network sizes will result in a different number of simulated groups. For example, for a group size of 3 there are 120 simulated groups for NET-1 (i.e., the combination of 10 switches taken 3 at a time) and 455 simulated groups for NET-2 (i.e., the combination of 15 switches taken 3 at a time). Within each of these groups, each member is an eligible source which can optimize circuits to the other members of the group. For a group size of 3, there are 3 possible sources and thus 3 possible source-receivers circuits. This translates to a group size of 3 in NET-1 with 120 groups and 360 source-receivers circuits (i.e., $360 = 120 \times 3$ source-receivers circuits).

For NET-1 and NET-2, ATM multipoint groups of the specified sizes are created for minimum cost (in terms of links) using combinations of point-to-point and point-to-multipoint circuits. Using the VC Mesh model to form groups does not assume groups contain only symmetrical circuits; rather each source in turn identifies the least cost circuit(s) to connect with the remainder of the group. For example, a 3-switch group may consist of the following combinations of circuits, whichever is least costly in terms of links: 6 point-to-point circuits; 3 point-to-multipoint circuits; 2 point-to-point/2 point-to-multipoint circuits; or 4 point-to-point/1 point-to-multipoint circuits (in 3 different permutations). Figure 2 shows the formation of this 3-switch group for both the asymmetric and symmetric circuit cases. In the asymmetric circuit case, traffic flow between node pairs travel along different paths. In the symmetric circuit case, traffic flow between node pairs share the same path but in different directions.

Subsequent to identification of the "working" ATM multipoint groups, "backup" groups are found that connect the same working group members, are link and node disjoint from the corresponding working group, and use either point-to-point backup circuits (unicast) exclusively or point-to-multipoint backup circuits (multicast) exclusively. A quantitative comparison of feasibility and cost is made between the disjoint backup groups using point-to-point circuits (unicast) exclusively versus the disjoint backup groups using point-to-multipoint circuits (multicast) exclusively.

The simulation experiments were implemented in MATLAB on a Sun Enterprise 4000 with ten 250 MHz ULTRASparc-II CPUs and 2.5 GB of RAM. Due to space limitations we present only the results from NET-1. The results from NET-2 were consistent with the results from NET-1 and are being prepared for a future publication.[5]

5 Experimental Results

Working groups are formed using the VC Mesh model with each group member having optimized circuit(s) to all other group members. For each group member within the group potentially functioning as a source, there are two possibilities to select working circuit(s) to join with remaining group members: (1) a single point-to-multipoint circuit (multicast) or (2) multiple point-to-point circuits (unicast). With one of the restoration criteria being cost, the source selects the optimal working circuit(s) based on minimum cost in terms of links. For each group, the minimum cost mixture of working circuits (point-to-point and point-to-multipoint) are selected. Figure 3A shows the cost in links for all working groups in NET-1 for group sizes equal to 3, 4, and 5. As would be expected, the cost increases as the group size increases. In general, the average cost of each working group is a group size multiple larger than the average cost for a working circuit.

Once the minimum cost working groups have been identified, the protection of these working groups can be attempted through restoration using rerouting. By finding disjoint backup circuits for each working circuit in a working group, backup groups can be formed for each working group using the VC Mesh model. Backup groups are formed using the VC Mesh model with each backup group consisting exclusively of either point-to-multipoint backup circuits or consisting exclusively of point-to-point backup circuits (i.e., backup groups are not formed with a mix of point-to-multipoint backup circuits with point-to-point

[5] interested readers are encouraged to contact the author for more details

backup circuits). The links and nodes that comprise each working circuit in each working group are removed from the network such that a disjoint backup circuit can be found. Figure 3B shows the cost in links for all backup groups in NET-1 for group sizes equal to 3, 4, and 5 (point-to-multipoint and point-to-point backup groups are combined). As would be expected, the cost increases as the group size increases. In general the average cost for each backup group is a group size multiple larger than the average cost for a backup circuit.

Figure 4 compares the cost of backup groups that consist exclusively of point-to-multipoint backup circuits (multicast) versus backup groups that consist exclusively of point-to-point backup circuits (unicast) within NET-1 for group sizes equal to 3, 4, and 5. While costs for the two types of backup groups are close for a group size equal to 3, the costs diverge for larger group sizes. The costs between the backup groups are statistically different for all group sizes considered as shown by non-overlapping 95% confidence intervals, with backup point-to-multipoint groups having a lower average cost. Note that the slope of the extreme costs for backup point-to-point groups are steeper than the extreme costs for backup point-to-multipoint groups.

Figure 5 shows the composite of all available backup groups for NET-1 with group sizes equal to 3, 4, and 5. The relationships graphically depicted are that for all working groups: the percentage of backup groups consisting exclusively of point-to-multipoint backup circuits that are available (the complement shows the percentage of backup point-to-multipoint groups that are not available); the percentage of backup groups consisting exclusively of point-to-point backup circuits that are available (the complement shows the percentage of backup point-to-point groups that are not available); and the percentage of backup groups that are not available (the complement shows the percentage of backup groups that are available). Backup groups are not available if they cannot be formed of backup circuits exclusively of one type, either backup point-to-multipoint (multicast) or backup point-to-point circuits (unicast). Only one backup circuit within a backup group has to be unavailable to make the entire backup group unavailable. Figure 5 shows that the feasibility of providing survivability via backup groups for a group size of 3 is approximately 50%. As the group size increases, the percentage of available backup groups that cannot be formed (working groups without protection from backup groups) increases while the percentage of available backup point-to-multipoint groups and the percentage of available backup point-to-point groups decreases commensurately. Note the percentage of available backup groups that consist exclusively of backup point-to-multipoint circuits added to the percentage of available backup groups that consist exclusively of "solitary" backup point-to-point circuits equals the percentage of available backup groups that consist exclusively of backup point-to-point circuits. Note also that the percentage of available backup groups that consist exclusively of backup point-to-point circuits is complementary to the percentage of backup groups that cannot be formed.

Closer analysis of the data shows that percentage of available point-to-point backup circuits is complementary to the percentage of working circuits which have neither type of circuit backup. Intuitively, if a backup circuit exists then a backup point-to-point circuit will exist. Figure 6A shows the converse of this with a working point-to-multipoint circuit for the group of member nodes $\{1, 2, 4\}$ in NET-1 where backup point-to-point circuits exist but a corresponding backup point-to-multipoint circuit does not exist. We will refer to this situation where backup point-to-point circuits exist but a corresponding backup point-to-multipoint circuits does not exist as "solitary" backup point-to-point circuits. Figure 5 shows the percentage of available backup groups that consist exclusively of "solitary" backup point-to-point circuits. There are no working circuits where a backup point-to-multipoint circuit exists and a backup point-to-point circuits does not exist (i.e., "solitary" backup point-to-multipoint circuits do not exist) since a point-to-multipoint circuit consists of point-to-point circuits with shared links.

A disjoint backup circuit may not be available for two reasons: (1) a multipoint "trap" may appear in the topology, or (2) there may be no disjoint backup circuits available shorter than the hop-limit bound. A multipoint "trap" arises when the removal of the links and nodes of the working circuit(s) within a working group eliminates all potential disjoint backup circuits. For example, Figure 6B shows a minimum cost working point-to-multipoint circuit for group member nodes $\{ 8, 9, 10 \}$ in NET-1. The removal of the links and nodes for this working circuit leaves no potential disjoint backup circuit(s) from source node 10 to receiver nodes 8 and 9. "Solitary" backup point-to-point circuits as discussed previously are also caused by these multipoint "trap" topologies. While the hop-limit bound is a factor in preventing available disjoint backup circuits, the trap has become the dominant factor for point-to-multipoint restoration after being a relatively rare event for point-to-point backup restoration. Point-to-point traps are discussed in more detail in [7].

Up to this point we have focused on comparing the availability and cost of backup groups consisting exclusively of point-to-multipoint backup circuits (multicast) versus backup groups consisting exclusively of point-to-point backup circuits (unicast). There is another alternative. "Hybrid" backup groups can be formed by selecting the minimum cost circuits from among both point-to-multipoint backup circuits and point-to-point backup circuits such that a minimum cost backup group may consist of a mix of point-to-multipoint backup circuits and point-to-point backup circuits. Such a hybrid backup group would theoretically have the minimum cost when compared to backup groups consisting exclusively of either point-to-multipoint backup circuits or point-to-point backup circuits.

Figure 7 shows a cost comparison between available backup groups consisting exclusively of one type of backup circuit (NO HYBRID backup groups) versus available minimum cost backup groups (HYBRID backup groups). The average cost between these backup groups is statistically different for all group sizes considered as shown by non-overlapping 95% confidence intervals, with HYBRID backup groups having a lower average cost. While the costs are similar for a group size equal to 3, the costs diverge for larger group sizes. Note in Figure 7B the steeper slope of the maximum cost line is due to NO HYBRID backup groups forcing backup circuits exclusively of one type (either point-to-multipoint or point-to-point backup circuits) while a hybrid backup group can incorporate both types of circuits in a minimum cost solution.

Despite the cost advantages, the percentage of unavailable backup groups found in the HYBRID case is exactly equal to the percentage of unavailable backup groups found in Figure 5. This is to be expected since minimum cost backup groups select from among available backup circuits and if any backup circuits are unavailable, optimization becomes irrelevant.

6 Summary

Preplanned, dedicated bandwidth, disjoint backup route techniques have proved to be very effective for the restoration of ATM point-to-point Virtual Paths and Virtual Circuits. This research has investigated the extension of these techniques for the restoration of ATM group communications implemented via the VC Mesh model. When a single link or single node fault occurs within a working group, traffic flow is rerouted to a corresponding backup group that is link and node disjoint with dedicated bandwidth reserved for restoration.

Our results show that for the two networks studied the use of disjoint dedicated backup groups is often not possible. Guaranteed survivability begins at only about 50% for a group size of the three and decreases rapidly for larger group sizes due to multipoint trap topologies and hop-limit constraints. These results are conservative because our simulations do not consider the effects of network traffic congestion which may make an otherwise available backup group become unavailable for one of the following reasons: (1) transient flows due to retransmissions of traffic backlogs at higher layers which violate reserved bandwidth restrictions and (2) traffic overload conditions which violate reserved bandwidth restrictions. Thus depending on backup groups to provide ATM multipoint survivability is infeasible due to low levels of successful restoration and the trend toward zero successful restoration as group size increases. We posit this as a general result but further research on other networks is desirable to provide independent confirmation and characterization of successful restoration rates on different network topologies.

Given the use of disjoint dedicated backup groups to provide ATM multipoint survivability is not feasible for groups of different sizes on arbitrary networks, there still exist dense networks where backup groups are available and can provide guaranteed survivability. By being identified in advance, some working groups may have the option of being protected with disjoint dedicated backup groups. For these working groups only, the question of cost of backup group restoration is relevant. Results show that for the two networks studied backup groups consisting exclusively of point-to-multipoint backup circuits (multicast) are lower in cost than backup groups consisting exclusively of point-to-point backup circuits (unicast). Introducing hybrid backup groups provides a minimum cost restoration solution that is statistically significant at a 95% confidence level. Thus for a given group in an arbitrary network, the optimal disjoint dedicated restoration in terms of cost is a minimum cost hybrid backup group (multicast and unicast), followed by an exclusively point-to-multipoint backup group (all multicast circuits), and then an exclusively point-to-point backup group (all unicast circuits). However, if only one backup group exists for a given group in an arbitrary network, it will be a point-to-point backup group.

In conclusion, disjoint dedicated backup is not feasible as a general restoration technique for ATM multipoint groups but may be used for specific groups when available. Further research is needed to propose new ATM multipoint schemes that incorporate survivability as well as investigating the feasibility and cost of providing survivability in proposed shared tree schemes.

References

[1] G. Armitage, "Support for Multicast Over UNI 3.0/3.1 Based ATM Networks," *IETF RFC*, rfc2022.txt, November 1996.

[2] ATM Forum, *ATM User-Network Interface (UNI) Specification Version 3.1*, Prentice Hall, Englewood Cliffs, New Jersey, June 1995.

[3] ATM Forum Technical Committee, *ATM User-Network Interface (UNI) Signaling Specification Version 4.0*, ATM Forum Document af-sig-0061.000, July 1996.

[4] T. Ballardie, and J. Crowcroft, "Multicast-Specific Security Threats and Counter-Measures" *Proceedings of the 1995 Symposium on Network and Distributed System Security*, IEEE Press, 1995, pp. 2-16.

[5] R. Cohen and A. Segall, "Connection Management and Rerouting in ATM Networks," in *Proceedings of IEEE INFOCOM'94*, Toronto Canada, 1994, pp. 184-191.

[6] R. Cotter, D. Medhi, and D. Tipper, "Traffic Backlog and Impact on Network Dimensioning for Survivability for Wide-Area VP-Based ATM Networks," in *Proceedings of the 15th International Teletraffic Conference* (ITC 15), Washington D.C., 1997, pp. 691-700.

[7] D. Dunn, W. Grover, and M. MacGregor, "Comparison of k-Shortest Paths and Maximum Flow Routing for Network Facility Restoration," *IEEE Journal on Selected Areas in Communications*, Vol. 12, No. 1, January 1994, pp. 88-99.

[8] O. Gerstel, I. Cidron, and S. Zaks, "The Layout of Virtual Paths in ATM Networks," *IEEE/ACM Transactions on Networking*, Vol. 4, No. 6, December 1996, pp. 873-884.

[9] A. Girard and B. Sanso, "Multicommodity Flow Models, Failure Propagation, and Reliable Loss Network Design," *IEEE/ACM Transactions on Networking*, Vol. 6, No. 1, February 1998, pp. 82-93.

[10] R. Kawamura, K. Sato, and I. Tokizawa, "Self-Healing ATM Networks Based on Virtual Path Concept," *IEEE Journal on Selected Areas in Communications*, January 1994, pp. 120-127.

[11] D. Medhi, "Models for Network Design, Servicing, and Monitoring of ATM Networks Based on the Virtual Path Concept," *Computer Networks and ISDN Systems*, Vol. 29, 1997, pp. 373-386.

[12] D. Medhi, "A Unified Approach to Network Survivability for Teletraffic Networks: Models, Algorithms and Analysis," *IEEE Transactions on Communications*, Vol. 42 No. 2/3/4, March/April 1994, pp. 534-548.

[13] A. Vakhutinsky and M. O. Ball, "Fault-Tolerant Virtual Path Layout in ATM Networks," in *Proceedings of the 15th International Teletraffic Conference* (ITC 15), Washington D.C., 1997, pp. 1031-1041.

[14] C.S. Wu, S.W. Lee and Y.T. Hou, "Backup VP Preplanning Strategies for Survivable Multicast ATM Networks," in *Proceedings of IEEE International Conference on Communications 1997* (ICC'97), Montreal, Canada, pp. 267-271.

[15] Y. Xie, M.J. Lee, and T.N. Saadawi, "Multicasting Over ATM Using Connection Server," in *Proceedings of IEEE International Conference on Communications 1997* (ICC'97), Montreal, Canada, pp. 1376-1380.

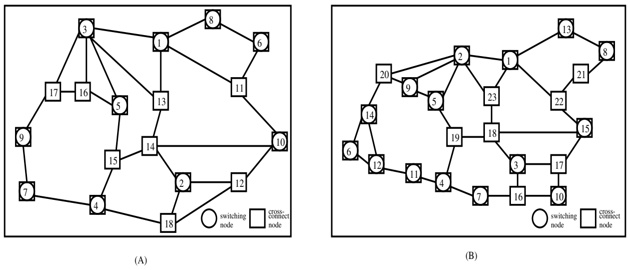

Figure 1. *Network Topologies: (A) NET-1 (10 switching nodes, 18 cross-connect nodes); (B) NET-2 (15 switching nodes, 23 cross-connect nodes); networks used in the following references [11, 6, 12, 9]*

GROUP SIZE	*NET-1* {*10 switches/18 nodes*}	*NET-2* {*15 switches/23 nodes*}
3	120 groups (360 src-rcvrs circuits)	455 groups (1,365 src-rcvrs circuits)
4	210 groups (840 src-rcvrs circuits)	1,365 groups (5,460 src-rcvrs circuits)
5	252 groups (1,260 src-rcvrs circuits)	3,003 groups (15,015 src-rcvrs circuits)

Table 1. *Simulation Experiments*

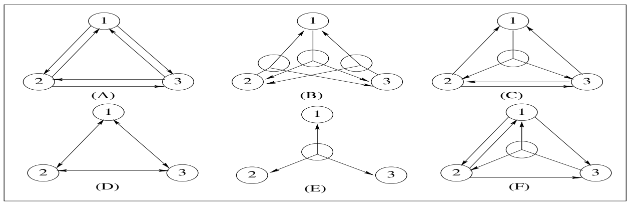

Figure 2. *The Formation of Multipoint Groups Using the VC Mesh Model: (A) a 3-switch group formed by 6 pt-to-pt asymmetrical ckts; (B) a 3-switch group formed by 3 pt-to-mpt asymmetrical ckts; (C) a 3-switch group formed by 1 pt-to-mpt and 4 pt-to-pt asymmetrical ckts; (D) a 3-switch group formed by 6 pt-to-pt symmetrical ckts; (E) a 3-switch group formed by 3 pt-to-mpt symmetrical ckts; (F) a 3-switch group formed by 1 pt-to-mpt and 4 pt-to-pt ckts (permutation of 2C)*

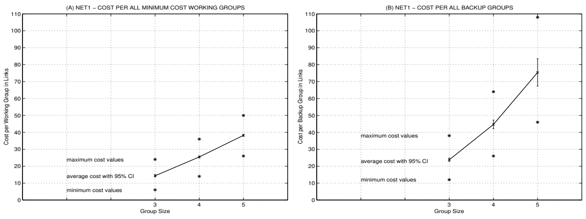

Figure 3. *Cost of VC Mesh Working/Backup Groups: (A) Cost Per All Working Groups (B) Cost Per All Available Backup Groups*

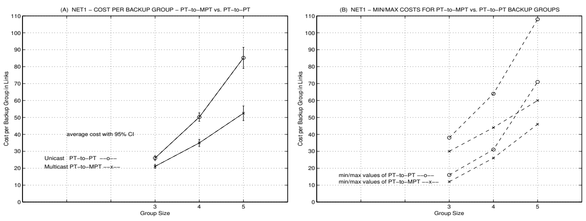

Figure 4. *Cost Comparison of Available Backup Groups - Point-to-Multipoint versus Point-to-Point: (A) comparison of average cost with a 95% confidence interval; (B) comparison of extreme (minimum/maximum) costs*

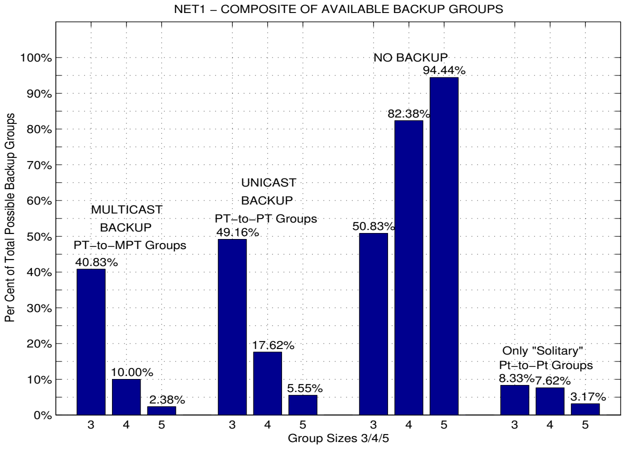

Figure 5. *Composite of Available Backup Groups*

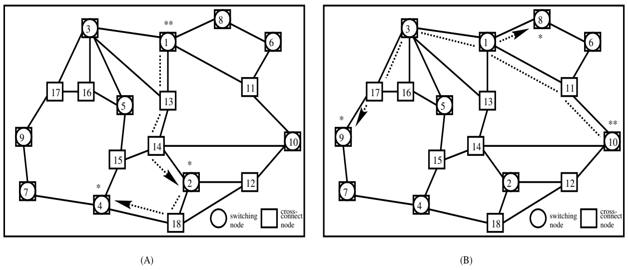

Figure 6. *Multipoint "Trap" Topologies in NET-1: (A) a multipoint "trap" topology of a working pt-to-mpt ckt for the group of member nodes {source = 1**; receivers = 2*, 4*} where "solitary" backup pt-to-pt circuits exist but a corresponding backup pt-to-mpt circuit does not exist; (B) a multipoint "trap" topology of a working pt-to-mpt ckt for the group of member nodes {receivers = 8*, 9*; source = 10**} where backup ckts do not exist (neither pt-to-pt nor pt-to-mpt backup circuits exist)*

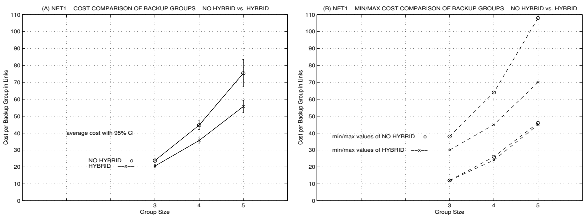

Figure 7. *Cost Comparison of NO HYBRID versus HYBRID Backup Groups: (A) comparison of average cost with a 95% confidence interval; (B) comparison of extreme (minimum/maximum) costs*