The Use of Hop-Limits to Provide Survivable ATM Group Communications

William Yurcik
Illinois State University
Dept. of Applied Computer Science
Normal IL USA
(309) 438-8016
wjyurci@ilstu.edu

David Tipper
University of Pittsburgh
Dept. Info. Sci. & Telecommunications
Pittsburgh, PA USA
(412) 624-9421
tipper@tele.pitt.edu

Deepankar Medhi
University of Missouri – Kansas City
Dept. of Computer Networking
Kansas City, MO USA
(816) 235-2006
dmedhi@umkc.edu

ABSTRACT
We examine the use of a hop-limit constraint with techniques to provide survivability for connection-oriented ATM group communications. A hop-limit constraint is an approach that has evolved from solving point-to-point routing problems but has not been fully developed for group communications. A hop-limit: (1) limits the number of routes considered such that the routing problems of higher order complexity can be solved and (2) limits the length of any individual route to meet specific Quality of Service guarantees (such as delay). This paper focuses on the former. 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 Rings, Shared Multicast Trees, or VC Meshes. We found that hop-limit constraints allowed us to exactly solve NP-complete routing problems on real networks.

Categories and Subject Descriptors
B.8.1 [**Performance and Reliability**]: Reliability, Testing, Fault-Tolerance;C.2.1 [**Computer-Communications Networks**]: Network Architecture and Design --- *asynchronous transfer mode (ATM), circuit-switching networks, network communications, network topology;* C.4 [**Performance of Systems**]: --- *fault-tolerance, reliability, availability, and serviceability.*

General Terms
Performance, Reliability, Security

Keywords
Group Communication, Survivability, Multicast, Self-Healing

1. INTRODUCTION
Multicast mechanisms provide group communications by reducing the amount of duplicate traffic in the network to conserve bandwidth and switch resources. When a sender transmits data to multiple receivers, some delivery routes may have common links that can be shared such that the data is sent only once over the common links and replicated at branch points. Thus multicast is efficient because it minimizes aggregate network load by decreasing the number of links and nodes utilized in a conversation between group members. This efficiency increases with the size of the network and the number of group members.

Network survivability refers to the ability of a network to function despite underlying faults that leave part of the network unavailable. Note these faults may be random reliability failures, damage from intentional malicious attacks, or a combination of both. Network survivability has attracted much attention recently for two reasons: (1) the need for reliable networking to support E-commerce and (2) the finite probability of an asymmetric terrorist attacks on a national critical infrastructure. In both cases (national economic and military security), it is important that networks be able to prevent, detect, and recover from a fault without a cascading impact within or between systems.

Group communications allows multiple senders and multiple receivers to engage in a simultaneous conversation. There is an increasing dependence on group communications for both E-commerce (streaming media, teleconferencing, distributed simulation, distributed computing) and military uses (command and control, battlefield sensors). In most cases the impact of a fault within a group communications scenario is much larger than point-to-point scenario since underlying multicast mechanisms efficiently aggregate traffic and by definition multiple parties are involved.

Solutions to provide survivable connection-oriented ATM group communications can take different approaches. One approach is preplanning a working set of routes protected by a standby set of routes. When a single link or node fault occurs within a working set of routes, traffic flow is rerouted (within a tight time threshold) 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 at the same time. Note that the bandwidth on the backup routes is dedicated and thus single-fault restoration can be guaranteed. For a typical network, finding the working and disjoint backup set of routes and provisioning dedicated bandwidth is a non-trivial problem.

An important factor in providing survivability is the set of routes considered. A hop-limit constraint is common technique used to:

2nd Annual Workshop on Networked Group Communications (NGC'00), November 8-10, 2000, Palo Alto, CA. USA.

(1) limit the length of routes to meet Quality of Service (QoS) guarantees (especially delay) and (2) to limit the number of routes considered in routing optimization problems [9, 10, 11, 12, 19]. Experimental results from [10] report that the hop-limit technique is indeed sound as it restrains model dimensions and at the same time ensures QoS guarantees can be met. In fact, [10] reports that a relatively small hop-limit still provides effective solutions while minimizing computations such that there is no real gain for setting larger hop-limits. Several hop-limit heuristics have been proposed. Itai et. al, report in [12] the first use of hop-limits combined with an algorithm to effectively calculate the number of disjoint routes if the hop-limit was less than or equal to three. It is shown that except for small hop-limit values, the problem of finding disjoint paths in a general graph is NP-complete.[12] Oki and Yamanaka propose a recursive matrix calculation technique to calculate the number of disjoint routes below a hop-limit equal to four [19]. Herzberg et. al., proposes the "Extra-Hop Limit" heuristic which is the network diameter in hops plus two[9] where the network diameter is defined as the maximum shortest-path between any two nodes in the network. All of this previous work on providing survivability while incorporating hop-limits has focused on the point-to-point context. To our knowledge, this paper is the first to discuss and provide quantitative results on providing survivability via hop-limits in the multipoint context.

The remainder of this paper is organized as follows: Section 2 introduces proposed schemes for connection-oriented ATM group communications. Section 3 discusses survivability approaches specifically for these proposed ATM schemes including Self-Healing Survivable Rings. Section 4 provides experimental results from varying the hop-limit constraint while solving survivability problems for real networks. In Section 5 we state our conclusions.

2. PROPOSED ATM GROUP COMMUNICATION SCHEMES

The search for an ATM group communications solution began with the introduction of ATM multicast routing capability (point-to-multipoint) in the ATM Forum's User Network Interface (UNI) specification 3.1. The ATM Forum MPOA (Multi-Protocol Over ATM) subworking group proposes two current approaches to ATM intracluster group communications: (1) the VC Mesh Model and (2) the Multicast Server Model (MCS). A cluster is defined as a set of ATM interfaces choosing to participate in native ATM routing.[5] Traffic between ATM interfaces belonging to different clusters pass through an intercluster device. As a result of the perceived complexity and inefficiency of these two models, other techniques have been proposed such as shared trees and rings. Figure 1 shows an example of each of the ATM group communication schemes for a group of three nodes.

2.1 VC Mesh Model

The VC Mesh Model, as proposed by the ATM Forum, 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 it dominant single-point-of-failure vulnerability which requires a different set of restoration techniques. For a more detailed discussion of the MCS see [25].

2.3 Shared ATM Multicast Tree Schemes

Many different schemes have been proposed for ATM group communications based on a shared ATM multicast tree approach. US patent 5,418,781 describes an architecture for maintaining the sequence of packet cells transmitted over a bidirectional ATM multicast tree. We consider three specific shared ATM multicast tree schemes: SMART, SEAM, and CRAM [6,8,14]. These schemes are similar in that they try to provide a general-purpose control architecture by modifying in-band control mechanisms of ATM switches [20]. Resources are reserved in both directions on the many-to-many VC links of a shared multicast tree until the connections are released [6]. Each of these schemes individually 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 Multicasting

Ofek and Yener have proposed the Virtual Ring [18,23] 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

3.1 Multipoint Survivability Tradeoffs

The two most relevant survivability tradeoffs that translate from point-to-point restoration to multipoint restoration are: (1) the method of identifying a backup route for restoration and (2) the location to initiate restoration. There are two methods of identifying a backup route for restoration: a preplanned backup route can be computed in advance or a dynamic search for a backup route can be computed in real-time. Survivability in the preplanned method is guaranteed because reserved bandwidth will be dedicated for restoration purposes while in the dynamic search method no bandwidth is reserved for backup routes so restoration is not guaranteed (i.e., residual bandwidth may not be available when needed). There are two locations or initiating restoration: end-to-end restoration initiated between sender and receivers and local restoration initiated between the upstream and downstream nodes next to the detected fault. While local restoration offers greater speed and ease of implementation, the further restoration is initiated away from a fault results in more efficient use of available network resources [4].

3.2 Survivability Techniques for ATM Models

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 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 [25].

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.

The formulation of the Virtual Ring as proposed by Ofek and Yener does not result in a least cost link and node disjoint ring, a single node or single link fault may disconnect a Virtual Ring in multiple places. We propose concepts similar to self-healing SONET ring systems [1,13,21,22]. Virtual Ring multicasting is extended to provide survivability for ATM group communications in [24] 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 a Disjoint Steiner Ring (DSR) connecting all desired group members, that is link and node disjoint, such that a single link or single 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 scenario (single link failure).

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, Shared Multicast Tree (SMT), and Self-Healing Survivable Ring (SHSR). To enable comparisons or feasibility and cost, we restrict our investigation to single fault scenarios (current state of the art).

We consider two actual network topologies as shown in Figure 2. Both networks 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 and NET-2 has an average node degree of 2.87 and a network diameter of 7 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 in all incident links.

Generating all possible paths between every node-pair up to the inclusion of all network nodes results in 16, 178 paths for NET-1 and 63,380 paths for NET-2. The effect of varying a hop-limit to reduce the search space in terms of the number of paths shows a characteristic curve common to all networks as shown in Figure 3 (despite different topologies and axis labels). Figure 4 shows that as the hop-limit increases below the network diameter, the number of links per node-pair path increases less than linearly until leveling off to about 2/3rds of the longest path as the hop-limit approaches the size of the network.

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 does not specify the mechanism so we have assumed the best case scenario such that each sender optimally selects the set of point-to-multipoint/point-to-point circuits that minimizes cost (each sender uses a least-cost routing algorithm to independently setup their VCMESH set of circuits). 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 an implementation variant of a Steiner Tree procedure due to Lawler known as the spanning tree enumeration algorithm [2]. Minimum cost working and backup rings are calculated by an implementation of the DSR formulation shown in [24].

The VCMESH, SMT, and DSR depend on solving variants of the Steiner Tree Problem which is well known to be NP-complete. However, both the spanning tree enumeration algorithm and the DSR formulation are able to exactly solve SMT and SHSR 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 the branch-and-bound can make previously intractable problems tractable by restricting the exponential rate of growth due to combinatorics. Figure 5 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 is most dramatic when considering DSR computational complexity. On NET-1, the computational complexity of the DSR solution is on the order of 10^4 Mflops without a hop-limit constraint and 2 X 10^1 Mflops with a hop-limit constraint. Flops are the cumulative number of executed floating point operations as measured by implementation in MATLAB version 5 (Mflops = a million flops). On NET-2, the computational complexity of the DSR solution without a hop-limit was unmeasured despite lengthy attempts and 4 X 102 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.

As the hop-limit increases, the feasibility (i.e., ability to find both working and disjoint backup route sets) of all the techniques increases without exception. The intuitive explanation is that some disjoint backup routes may be longer and thus ineligible at lower hop-limits but become revealed as the hop-limit increases. For example, 13 shows where a disjoint backup SMT and a disjoint backup VCMESH are revealed as the hop-limit increases from 15 to 17 respectively (working set of routes = solid lines, disjoint backup set of routes = dashed lines). Figure 6 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 [25] 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..

Figure 7 shows the difference in feasibility between working SMTs and disjoint backup SMTs as a function of hop-limit. As the hop-limit increases beyond the network diameter, all working SMTs are formed but disjoint backup SMTs require a larger hop-limit to converge to a lower feasibility due to multipoint traps. Multipoint traps also explain why the VCMESH technique has a dramatically lower feasibility than the SMT technique. Figure 8 shows the difference in feasibility between working and disjoint backup VCMESHes as a function of hop-limit for both networks. For VCMESH restoration to be available, all group members must form a working and corresponding disjoint backup set of circuits to all other members of the group. If a multipoint trap exists for just one of the group members then VCMESH restoration will not be available for the group as a whole.

According to the "Extra-Hop Limit" heuristic proposed by Herzberg et. al., (network diameter plus 2), the hop-limit would be 9 and 10 for NET-1 and NET-2 respectively. Closer examination of Figure 6 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 unprotected (unable to restored given a single fault). A new hop-limit heuristic is needed specifically for the multipoint survivability context.

Given the use of disjoint dedicated backup SMT/VCMESH techniques to provide survivability to ATM group communications has a lower feasibility for groups on arbitrary networks, there still exist densely connected networks (i.e., torus) where these techniques can provide guaranteed single fault restoration. For these networks the question of restoration cost is relevant. 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). Figure 9 compares the cost of the techniques as a function of hop-limit. These results show 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 [24]. Result 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) [24].

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 due to backup paths longer than the hop-limit suddenly emerge as eligible below a larger hop-limit. There are several counter examples, however, in which an increased hop-limit actually decreases the cost of restoration (see Figure 14).

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 b 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 for joining group members in a general graph given no survivability considerations. To explain this, Figure 10 shows the cost difference between working and disjoint backup SMTs as a function of hop-limit. These results show 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 than the backup SHSR. This same effect is amplified for the VCMESH technique as shown in Figure 11.

Ultimately there is a delay tradeoff between the ring and tree-based approaches. Figure 12 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 as provided by the SHSR technique must be weighed against an increased worst-case delay (recalling that the path of a reconfigured Virtual Ring doubles its length in links).

5. SUMMARY

In this paper we have examined the effect of a hop-limit constraint on restoration techniques to provide survivability for ATM group communications. Varying the hop-limit reduced the number of routes considered in routing optimization problems and limited the length of routes to enable delay-constrained comparisons between different techniques. We also 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 hop-limit heuristics developed for the point-to-point restoration context are not sufficient to provide exact solutions for the multipoint restoration context. Future research is needed to develop a lower hop-limit heuristic that still provides exact solutions. As providing QoS guarantees over a network becomes more important, it is also envisioned that the secondary effect of hop-limits as a delay-constraint will become a technique inherent to QoS-constrained survivable multipoint routing.

6. ACKNOWLEDGMENTS

The authors would like to thank the anonymous referees whose specific comments greatly improved this manuscript. We would also like to acknowledge funding support for this research from DARPA grant # F30602-97-1-0257 and NASA grant # NGT-30019.

7. REFERENCES

[1] S. Ahn, A Fast VP Restoration Scheme Using Ring-Shaped Sharable Backup VPs in Proc. of *IEEE Global Telecomm. Conf. (Globecom'97)*, 1383-1387.

[2] V.K. Balakrishnan, *Network Optimization*. Chapman & Hall, London U.K., 1995.

[3] R. Cotter, D. Medhi, and D. Tipper, Traffic Backlog and Impact on Network Dimensioning for Survivability for Wide-Area VP-Based ATM Networks in *Proc. of 15th Intl. Teletraffic Conf.*, 1997, 691-700.

[4] R. Doverspike and B. Wilson, Comparison of Capacity Efficiency of DCS Network Restoration Routing Techniques. *J. of Network and Systems Mgmt*., Vol. 2, No. 2, 1994, 95-123.

[5] Z. Dziong, *ATM Network Resource Management*. McGraw-Hill, New York, 1997.

[6] E. Gauthier, J-Y. 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, 458-472.

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

[8] M. Grossglauser and K.K. Ramakrishnan, SEAM: Scalable and Efficient ATM Multicast in *Proc. of IEEE INFOCOMM'97*, Kobe Japan, 1997.

[9] M. Herzberg, D.J. Wells, and A Herschtal, Optimal Resource Allocation for Path Restoration in Mesh-Type Self-Healing Networks in *Proc. of 15th Intl. Teletraffic Conf*., 1997, 351-360.

[10] 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, 775-784.

[11] M. Herzberg and S.J. Bye, An Optimal Spare-Capacity Assignment Model for Survivable Networks with Hop-Limits in *Proc. of IEEE Global Telecomm. Conf. (Globecom'94)*, San Francisco CA., 1994, 1601-1606.

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

[13] Y. Kajiyama, N. Tokura, and K. Kikuchi, ATM Self-Healing Rings *in Proc. of IEEE Global Telecomm. Conf. (Globecome'92)*, Orlando FL., 1992, 639-643.

[14] S. Komandur, J. Crowcroft, and D. Mosse, CRAM: Cell Re-Labeling at Merge Points for ATM Multicast in *Proc. of IEEE Intl. Conf. On ATM*, Colmar, France, 1998.

[15] 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, 373-386.

[16] D. Medhi, A Unified Approach to Network Survivability for Teletraffic Networks: Models, Algorithms and Analysis. *IEEE Trans. on Comm.*, Vol. 42, No. 2/3/4, 534-548.

[17] D. Medhi, Diverse Routing for Survivability in a Fiber-Based Sparse Network in *Proc. of IEEE Intl. Conf. on Comm. (ICC'91), Denver CO, 1991, 672-676.*

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

[19] E. Oki 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, 769-774.

[20] J.E. van der Merwe and I.M. Leslie, Service-Specific Control Architecture for ATM. *IEEE J. on Sel. Areas in Comm*., Vol 16, No. 3, April 1998, 424-436.

[21] T-H. Wu, *Fiber Network Service Survivability*, Artech House, Boston, 1992.

[22] T-H. Wu, D.T. Kong and R.C. Lau, An Economic Feasibility Study for a Broadband Virtual Path SONET/ATM Self-Healing Ring Architecture. *IEEE J. on Sel. Areas in Comm*., Vol. 10, No. 9, Dec. 1992, 1459-1473.

[23] B. Yener, Y. Ofek, and M. Yung, Combinatorial Design of Congestion-Free Networks. *IEEE/ACM Trans. on Networking*, Vol. 5, No. 6, Dec. 1997, 989-1000.

[24] W. Yurcik and D. Tipper, Providing Network Survivability to ATM Group Communications Via Self-Healing Survivable Rings in *Proc. of 7th Intl. Conf. on Telecomm. Systems*, Nashville, TN, 1999, 501-518.

[25] W. Yurcik, Providing ATM Multipoint Survivability Via Disjoint VC Mesh Backup Groups in *Proc. of IEEE 7th Intl. Conf. on Computer Comm. and Networks (IC3N'98)*, 1998, 129-136.

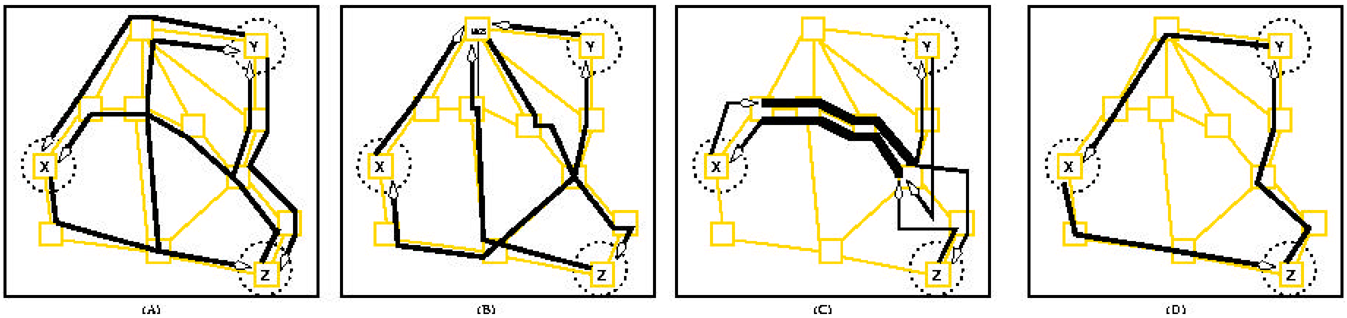

Figure 1: (A) VC Mesh; (B) Multicast Server; (C) Shared Tree; (D) Virtual Ring

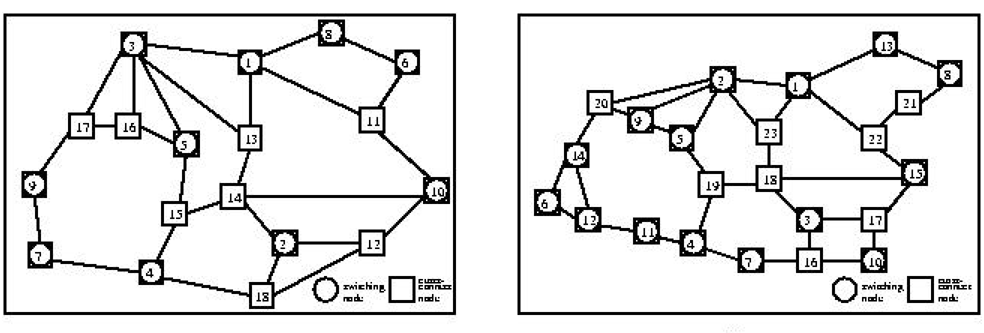

Figure 2: (A) NET-1; (B) NET-2; (networks used [3, 7, 15, 16, 24, 25])

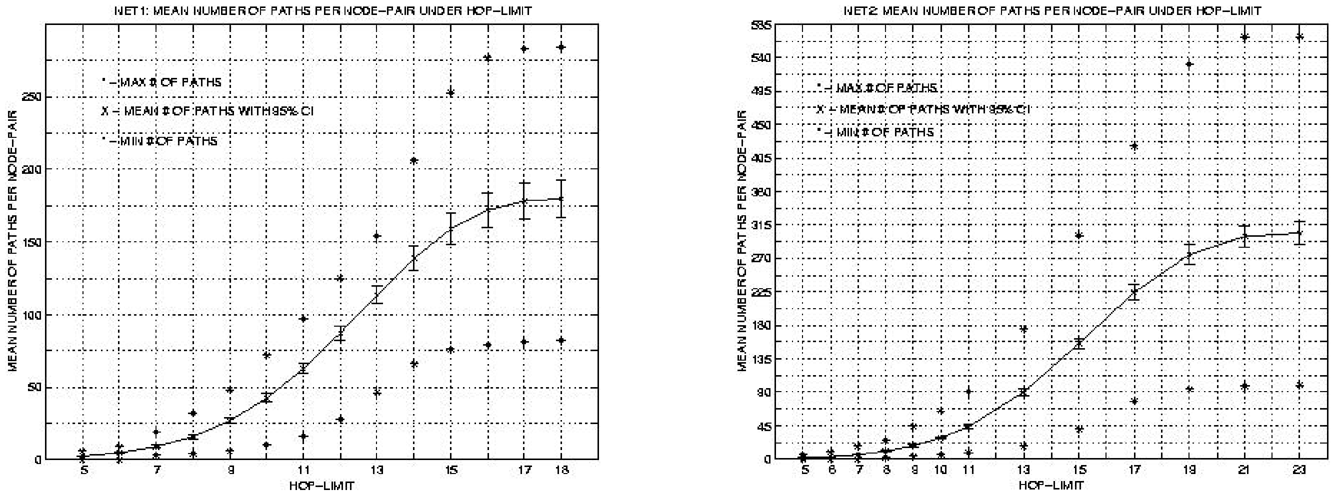

Figure 3: Mean Paths Per Node-Pair as a Function of Hop-Limit

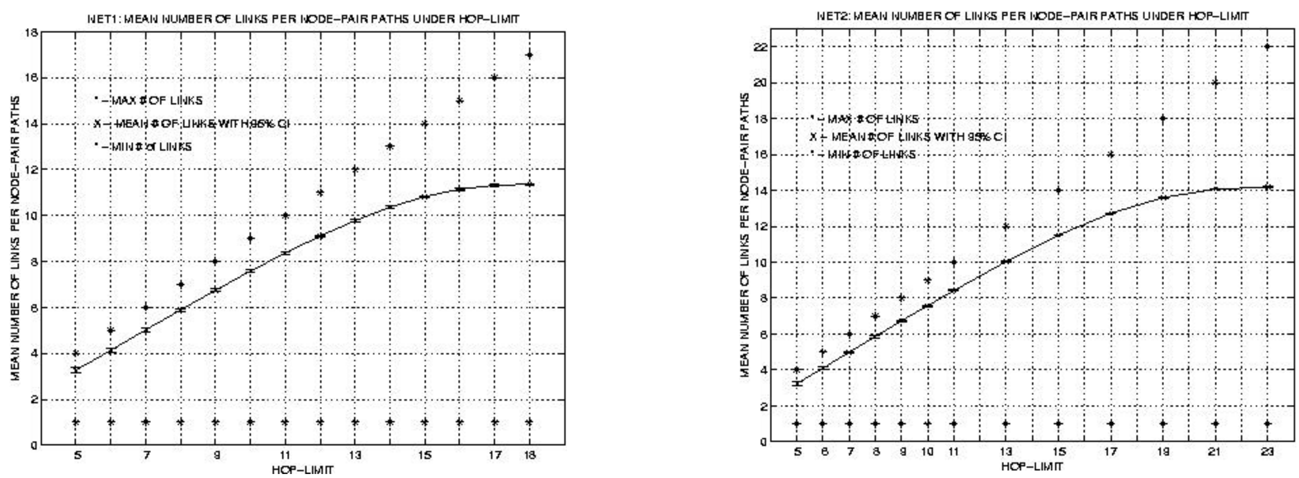

Figure 4: Mean Links Per Node-Pair as a Function of Hop-Limit

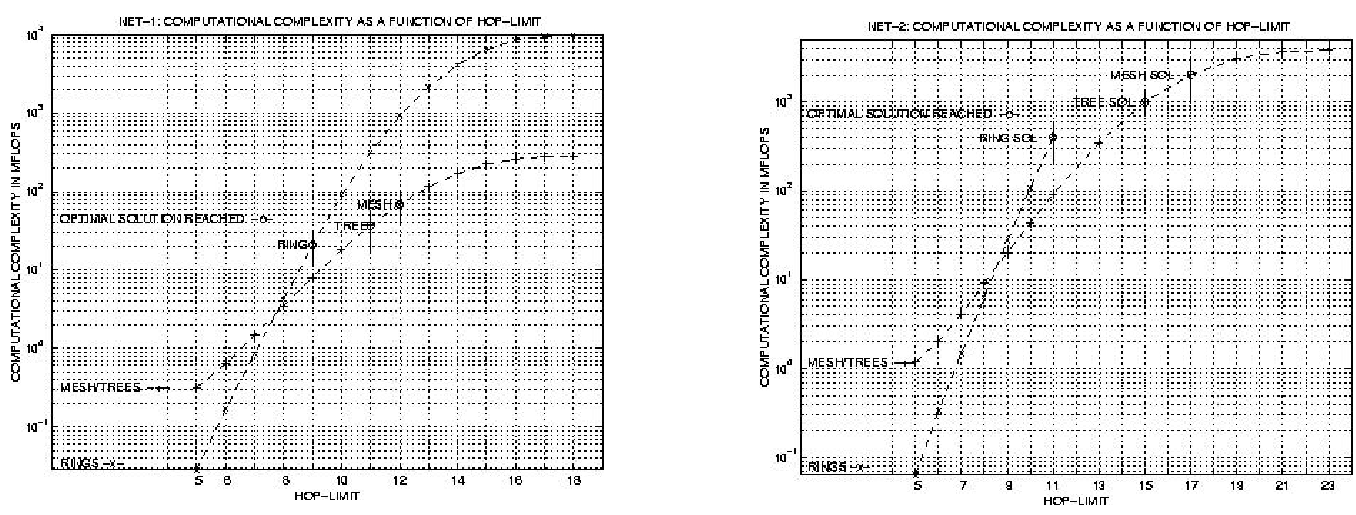

Figure 5: Computational Complexity as a Function of Hop-Limit

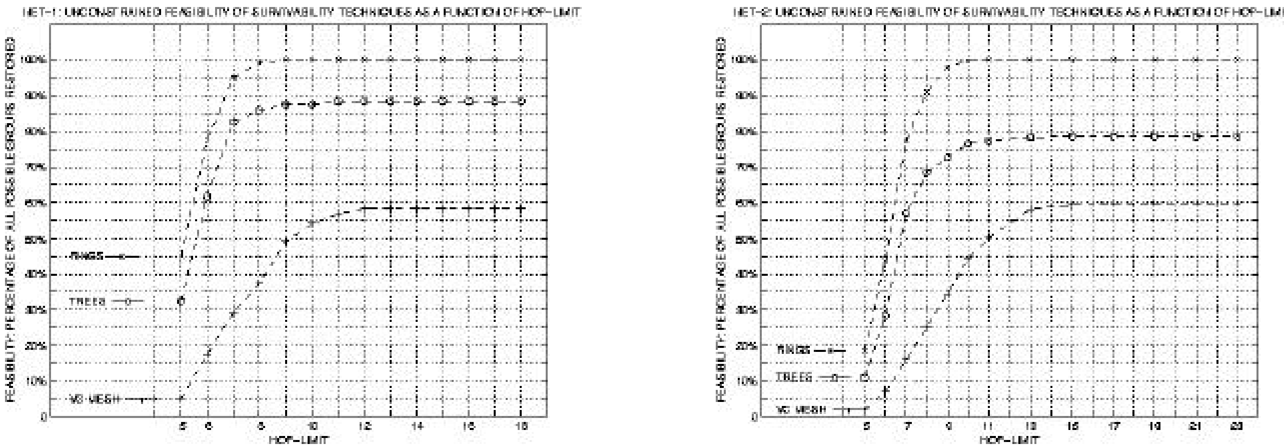

Figure 6: Feasibility of Survivability Techniques as a Function of Hop-Limit

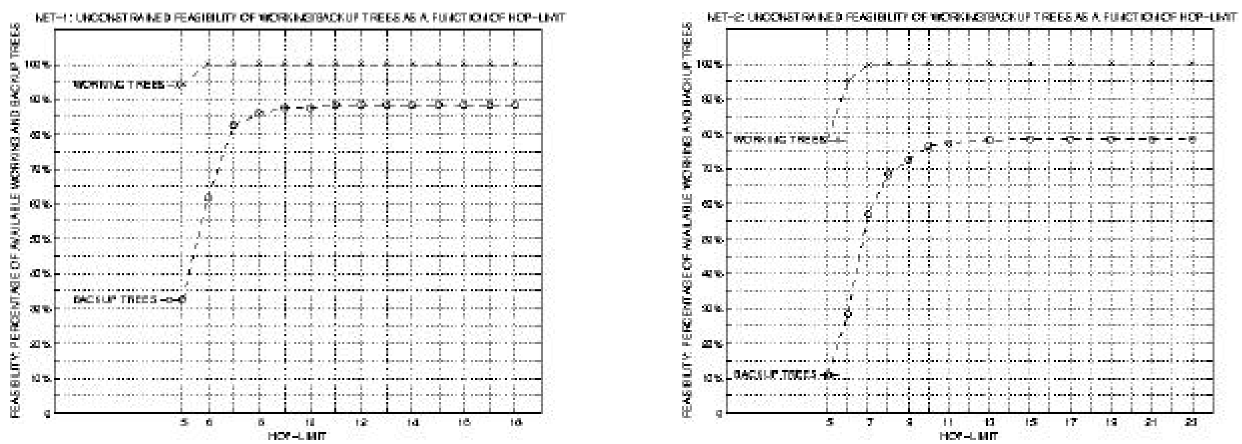

Figure 7: Working/Backup Tree Feasibility as a Function of Hop-Limit

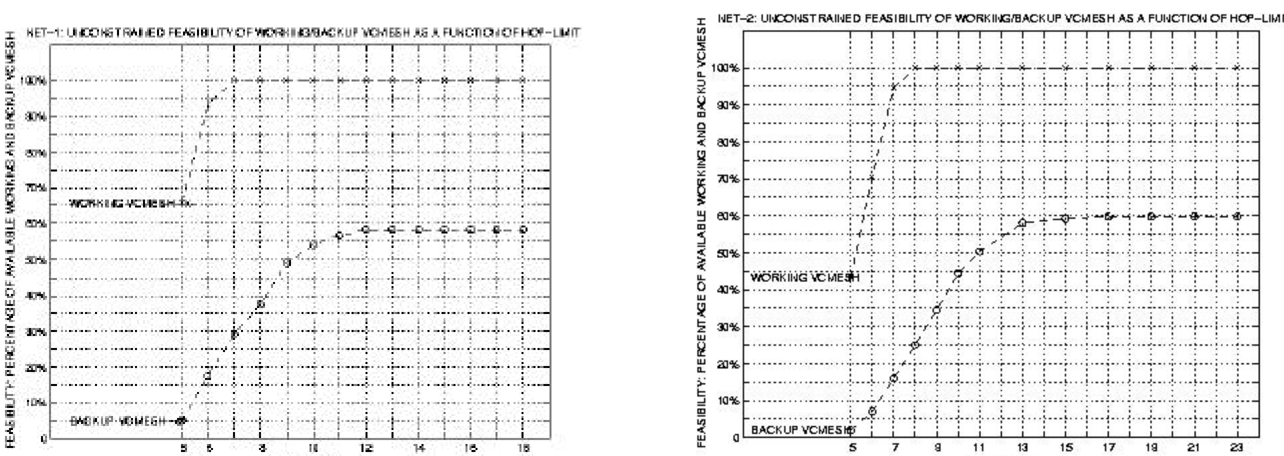

Figure 8: Working/Backup VC Mesh Feasibility as a Function of Hop-Limit

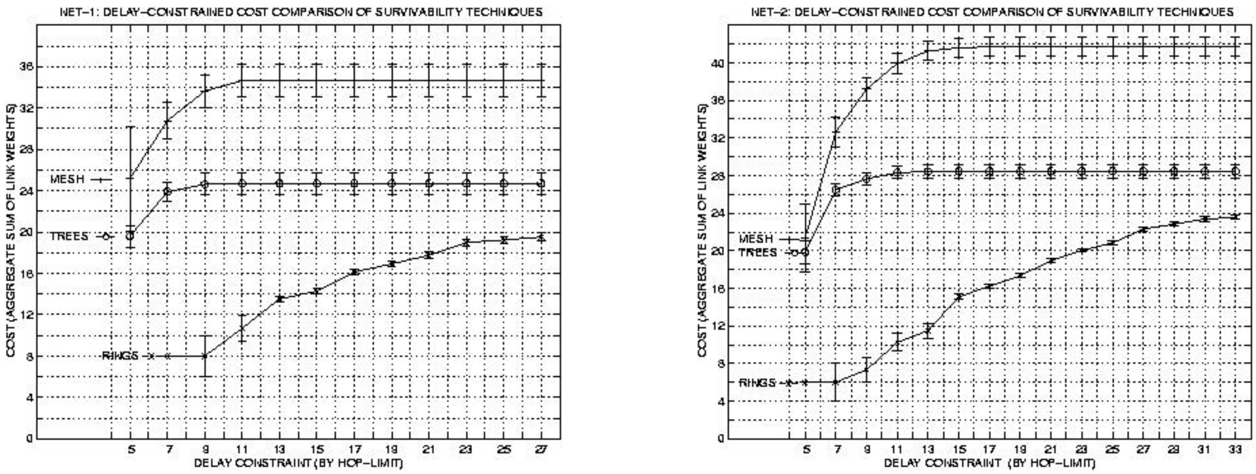

Figure 9: Delay-Constrained Cost of Survivability Techniques as a Function of Hop-Limit

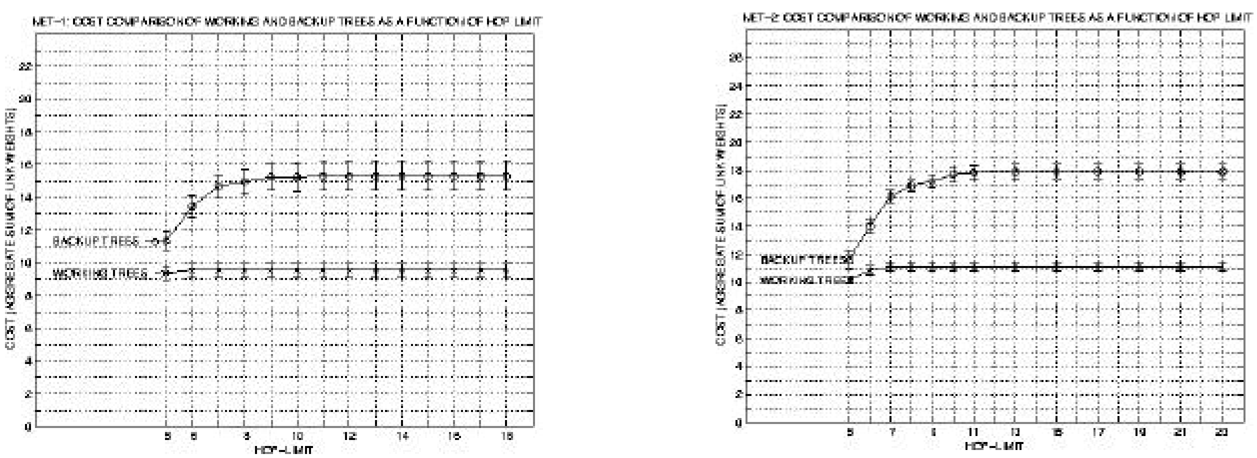

Figure 10: Working/Backup Tree Cost as a Function of Hop-Limit

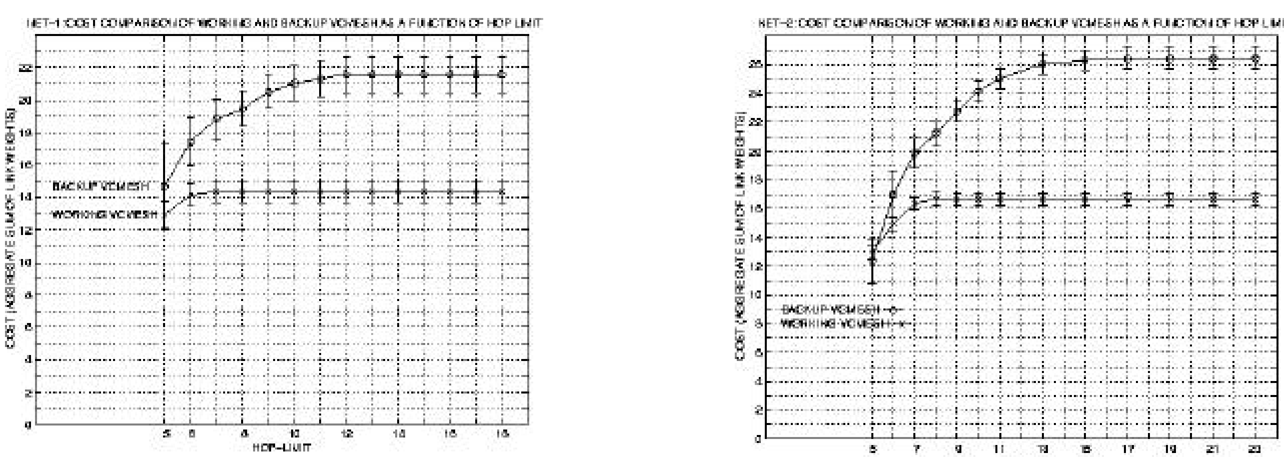

Figure 11: Working/Backup VC Mesh Cost as a Function of Hop-Limit

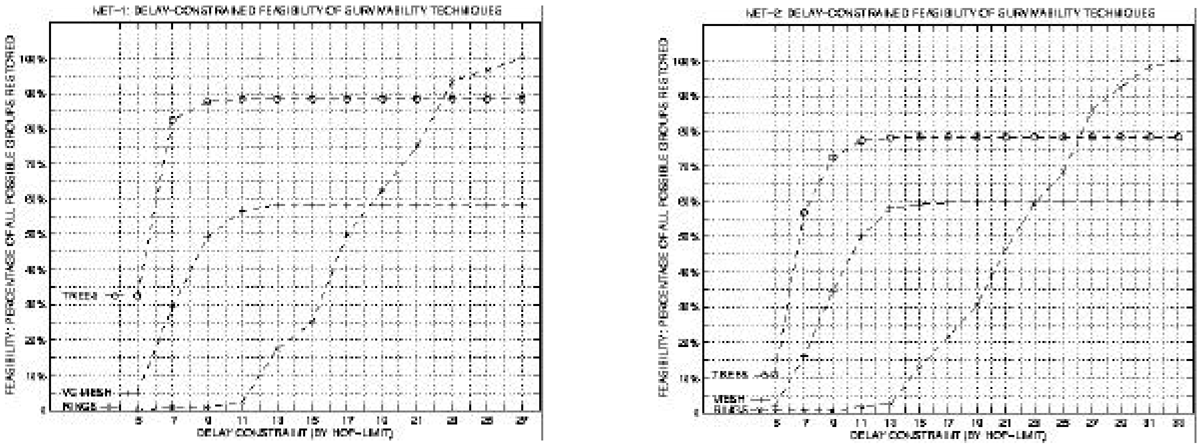

Figure 12: Delay-Constrained Feasibility of Survivability Techniques as a Function of Hop-Limit

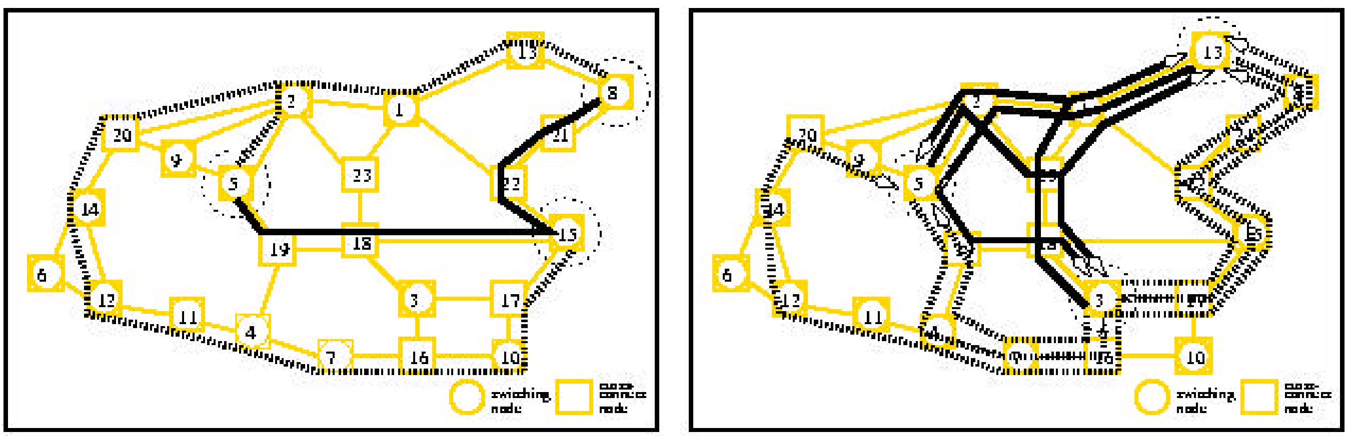

(A) (B)

Figure 13: Hop-Limits Which Reveal Disjoint Backups: (A) Hop-Limit = 15; (B) Hop-Limit = 17

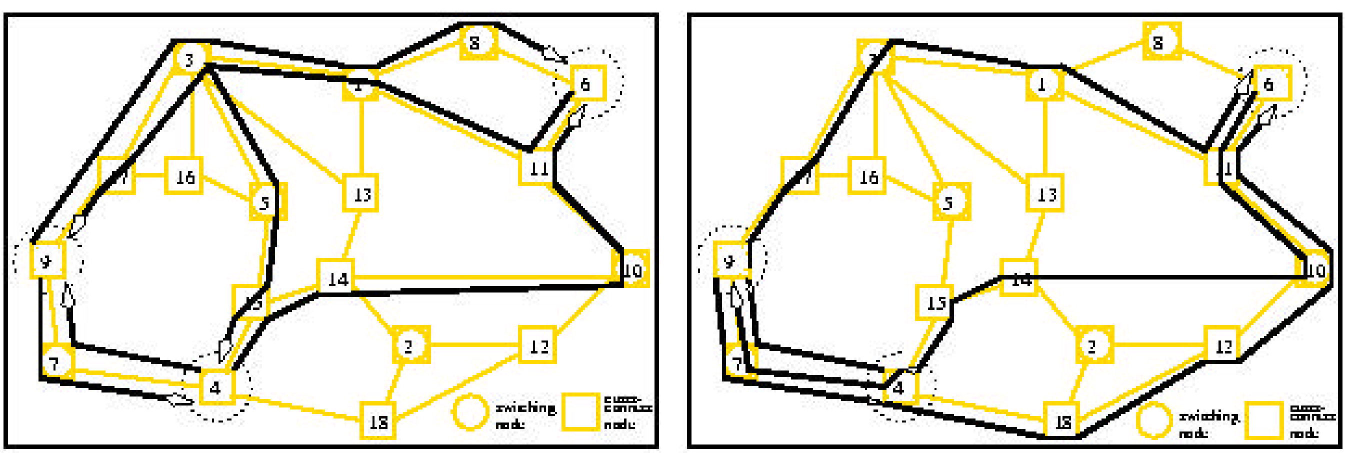

(A) (B)

Figure 14: Hop-Limits Which Reveal Disjoint Backups: (A) Hop-Limit=7, 22 links; (B) Hop-Limit=8, 21 links