Distributed Models and Algorithms for Survivability in Network Routing

(Extended Abstract)

Fred S. Annexstein and Kennth A. Berman
Department of ECE & Computer Science
University of Cincinnati
Cincinnati, OH 45221
{fred.annexstein, ken.berman}@uc.edu

Abstract

We introduce a natural distributed model for analyzing the survivability of distributed networks and network routing schemes based on considering the effects of local constraints on network connectivity. We investigate the computational consequences of two fundamental interpretations of the meaning of local constraints. We consider a full-duplex model in which constraints are applied to edges of a graph representing two-way communication links and a half-duplex model in which constraints are applied to edges representing one-way links. We show that the problem of determining the survivability of a network under the full-duplex model is NP-hard, even in the restricted case of simply defined constraints. We also show that the problem of determining the survivability of network routing under the half-duplex model is highly tractable. We are able to effectively determine fault-tolerant routing schemes that are able to dynamically adapt to withstand any connectivity threats consistent with the constraints of the half-duplex model. The routing scheme is based on a multitree data structure, and we are able to generate optimal multitrees of minimum weighted depth. We also investigate an optimization problem related to achieving survivable network routing using a small sets of retransmission or landmark sites. Although the associated optimization problem is NP-hard, we show that sufficiently dense graphs can achieve survivable routing schemes using a small sets of landmarks. We prove an associated extremal result that is optimal over all graphs with minimum degree δ.

1 Introduction

In this paper we address the problem of fault-tolerant network communications in terms of effective strategies for providing routing schemes that are able to tolerate network faults and/or the reduction of communication bandwidth. The survivability of network routing schemes must be approached somewhat differently from traditional topological notions of survivability of networks based on graph connectivity. Cutsets of a network have long been used as a tool for analyzing network reliability and survivability[7, 8]. Characterizing properties of classes of fault configurations is an interesting yet difficult problem, since in most reasonable fault models, the complete enumeration of (cutset) fault configurations results in combinatorial explosion.

In this paper we investigate the consequences of fault configuration models characterized by local distributed properties. We analyze the computational complexity of two fundamental variants on the interpretation of fault configurations defined by local constraints. In the first interpretation, called the *full-duplex model*, edges model two-way communication and edge faults produce communication loss in both directions. In the second interpretation, called the *half-duplex model*, edges model one-way communication and edge faults produce communication loss in only one direction.

We show that problem of determining the survivability of network routing under the full-duplex model is NP-hard, even in the case of highly restricted local constraints. In fact, the problem remains NP-hard for local constraints that permit at most one edge fault at each node. However, for the half-duplex model we can completely and effectively determine network routing survivability given *any* set of local constraints. In addition, we can determine the effect local constraints have on shortest fault-tolerant path distance between any pair of vertices.

Our approach is to provide an algorithm that generates a survivable routing scheme based on a multitree data structure. A multitree routing scheme stores a set of parent pointers for forwarding packets at each node of the network. Multitree routing schemes based on independent trees were studied in [9, 2], and multitree routing schemes based on acyclic orientations were studied in [1]. In this paper we show how to generate optimal multitrees of minimum depth with respect to any given collection of locally defined constraints in the half-duplex model. Our results generalize the results reported in [1] on algorithms for maximum capacity multitree routing scheme. The algorithm we give here for obtaining optimal depth multitree routing tables is based on a generalized version of Bellman-Ford shortest-path algorithm, and due to its localized label-correcting nature is applicable in dynamic network environments.

Finally, we investigate an optimization problem related to finding small sets of landmark sites that can facilitate fault-tolerant multitree-based communication schemes. These landmark sites are analogous to retransmission nodes used in certain network applications such as multicasting. The problem we study is that of finding a small number of landmarks so that fault-tolerant communication is guaranteed to the landmarks from each of the remaining fault-free nodes in the network. The landmarks must then communicate among themselves, through other channels or by retransmission, to achieve global communication. Although the general optimization problem is NP-hard, we are able to show that sufficiently dense graphs always have small sets of landmarks. We prove an extremal result that is optimal over all graphs with minimum degree δ.

The results reported in this paper on network routing survivability using local constraint models can be interpreted and applied to a number of areas of network management; particularly for networks which can be usefully modeled with half-duplex links, e.g., certain internets and mobile ad-hoc networks. First, by analyzing a range of possible local constraint or bandwidth reduction scenarios, we can determine a worst-case level of integrity in a given network. This is possible by testing whether a certain set of constraints can produce disconnection. Conversely, the techniques can determine an operational level which must be maintained to guarantee connectivity (or distance bounds), for example when determining schedules for downtimes of network components. This may be particularly useful in large distributed networks where it may be impossible to precisely characterize the global data traffic, but where local traffic statistics may be readily available.

2 Formal Framework

We formalize our distributed, local-constraint model for survivability as follows. We consider networks as multigraphs G, or equivalently simple graphs with edges labeled with positive integer weights. We begin with a collection of local constraints $\Lambda = \{\Lambda_v | v \in V\}$, one constraint function Λ_v for each node v in the network. In general, constraint functions may be arbitrary locally defined (boolean) functions; however, for simplicity we will often focus on an important special case in which each Λ_v is defined by a threshold function on the capacity of local edge sets. When Λ_v is such a threshold function, we write $| \Lambda_v |$ as the value of the threshold function.

A set of edges E' is said to *satisfy* Λ if the set E' is consistent with Λ (that is, the set does not violate any local constraint Λ_v, for any v); otherwise we say the set E' *fails to satisfy or violate* Λ. With threshold functions we use the interpretation that local edges sets of cardinality at or above the value $| \Lambda_v |$ violate the constraint, and satisfy otherwise.

A network G is Λ-*survivable* if it has the property that any edge set E' that satisfies Λ can be removed from the network without disconnecting G, i.e., $G - E'$ is a connected subgraph of G. We define the Λ-*distance* from vertex v to vertex t to be the worst case over all edge sets E' satisfying Λ of the shortest vt-path distance in the graph $G - E'$. If there is a edge set E' satisfying Λ whose removal separates v and t, then we say the Λ-distance from vertex v to vertex t is infinite.

In the full-duplex (resp., half-duplex) model networks are composed of edges representing two-way (resp., one-way) communication links. In the next two sections we show that the computational complexity of determining whether a network is Λ-survivable is critically different for these two models.

3 Survivability in Full-duplex Model

In this section we consider the computational complexity of determining whether a network is Λ-survivable in the full-duplex model. We show that the Λ-SURVIVABILITY problem, that is the problem of testing whether a network is Λ-survivable, is co-NP-hard, and this result holds when Λ is a highly restricted collection of functions. In particular the problem remains co-NP-hard when each Λ_v is a threshold of at most one on the cardinality of (faulty) edges at each vertex.

Theorem 1 *Given a collection Λ of local constraints, the problem of determining whether a network is Λ-survivable is co-NP-complete. Further, the problem remains co-NP-complete when local constraints are restricted to constant threshold functions.*

This result follows immediately by showing that the problem of determining whether there exists a set E' of independent (matching) edges in graph, such that E' is also a disconnecting set is an NP-complete problem. Let MATCHING-CUT define the following decision problem: given a graph G and pair of vertices s, t, does there exist an independent set of edges that separate s and t.

Proposition 1 *The MATCHING-CUT problem is NP-complete.*

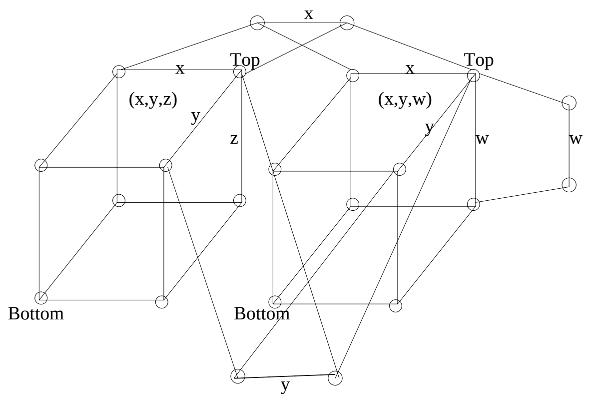

Figure 1. The partial structure of G_ϕ, for $\phi = (x, y, z) \vee (x, y, w)$.

Proof. Clearly, MATCHING-CUT is in NP. We now show that a restricted EXACT3-SAT problem reduces to the MATCHING-CUT problem. EXACT3-SAT asks if a 3-sat formula has a truth assignment in which exactly 1 of the 3 literals per clause is true. EXACT3-SAT remains NP-complete when instances are restricted to contain only positive literals [6].

Take such an instance ϕ, from which we will construct a graph G_ϕ. For each clause (x, y, z) create a graph of the boolean 3-cube where each dimension of the cube is associated with a unique variable of the clause. Consider antipodal corners of each cube, and denote these corners as top $t_{(x,y,z)}$ and bottom $b_{(x,y,z)}$. Link all top corners (one for each clause) together in a large clique with a vertex s, and likewise do this for all bottom corners with vertex t. For each variable x, create an edge labeled e_x. For each 3-cube associated with a clause containing x, link the edge e_x in a four cycle containing the top corner vertex and the vertex along the dimension associated with x; see for example Figure 1. This completes the construction of G_ϕ

We now claim that ϕ is in EXACT3-SAT iff G_ϕ is in MATCHING-CUT. Suppose we have a 1 in 3 satisfying assignment for ϕ. Then we can obtain a matching-cut of G_ϕ as follows: in each 3-cube remove all edges associated with the unique true variable. This edge set is a matching that separates s and all top corners from t and all bottom corners.

Conversely, suppose we are given a matching cut of G_ϕ. This cut cannot separate t and the set of bottom corners, nor can it separate s and the set of top corners, since they are contained in large cliques. Hence, the cut must separate all bottom corners from all top corners. But this can only be done by separating every cube with a matching. But such a matching naturally defines an exact 1 in 3 truth assignment for ϕ. $\Box$

4 Survivability in Half-duplex Model

Although the problem of determining whether a network is Λ-survivable is co-NP-hard even for highly restricted constraint functions, we are able to show in this section that for local constraints Λ defined in the half-duplex model (that is, where local constraints are restricted to one-way communication links) we can efficiently solve the problem of determining if a network is Λ-survivable, even when the local constraints are arbitrarily complex. Furthermore, we determine the precise effect such local constraints have on the length of shortest paths in the network (in the worst case).

Given a set of local constraints, we provide an algorithm that generates a fault-tolerant routing scheme that is able to withstand any collection of simultaneous faults, so long as they are consistent with the local constraints. The routing scheme is based on a *multitree* data structure. Such structures can effectively be used for fault-tolerant routing in communication networks that rely on local routing tables that associate with each destination a parent link on which to forward messages to the destination. Typically in network routing, a single set of parent links is used to form a directed *sink tree* for each destination sink [3, 4, 5]. However, tables based on a single tree suffer from the fact that they do not provide a full representation of the entire network, and therefore are vulnerable to faults and other dynamic network changes. In previous work [1] we studied multitree generalizations of sink tree schemes, and examined how acyclic orientations of the underlying network could be used to derive such multitree schemes. The multitree schemes achieve fault-tolerance by applying routing tables that store sets of parent pointers and, unlike the the sink-tree model, is able to dynamically and obliviously route around faults and congestion. In this paper we extend the work in [1] by showing how to generate optimal, survivable multitrees of minimum depth given any collection of locally defined constraints in the half-duplex model. We say a multitree scheme is Λ-*survivable*, if in the presence of any fault configuration consistent with Λ, there exists a fault-free tree path to the sink remaining in the collection. We prove the following result.

Theorem 2 *Given a graph G with a destination vertex t, and a collection Λ of local constraints in the half-duplex model, there is an $O(n^2)$ algorithm that computes the Λ-distance from each vertex v to t, and furthermore yields an Λ-survivable multitree routing scheme which achieves the minimum depth Λ-distance for routing.*

Proof. To determine the distance optimal Λ-survivable multitree routing scheme, we introduce the notion of an Λ-complement dag, or simply a Λ^c-dag. An Λ^c-dag D to t is defined as a st-connected subdigraph of G with the property that the out-neighborhood $N_D^+(v)$ of each vertex v in D fails to satisfy the constraints defined by Λ_v. The depth $d(D)$ of an Λ^c-dag D is defined as the longest vt-path length in D over all vertices v.

The association of Λ^c-dags with multitree routing schemes is an obvious one. An Λ^c-dag D to t of depth d, can be used to implement a Λ-survivable routing scheme that always successfully routes along paths of distance at most d. This follows since any fault configuration consistent with Λ could never isolate a node of D, and furthermore, any path from a vertex v to t generated using D has length at most d.

Let d^* be the minimum depth over all spanning Λ^c-dag to t. We now give an algorithm which returns a spanning Λ^c-dag D to t with depth d^*. The algorithm works by computing d_v, the Λ-distance from vertex v to vertex t, for each vertex v in the graph. At the same time we compute d_v, we also generate the edges in an Λ^c-dag which are used to achieve fault-tolerant routing paths of length at most d_v from v to t. To do this we use a generalization of traditional shortest path algorithms. We now present a generalization based on the Bellman-Ford algorithm, however, we remark that a generalization based on Dijkstra's algorithm is also possible. Analogous to the traditional algorithms, our algorithm works successfully for positive edge weighted graphs; however for simplicity we sketch the steps for the unweighted version.

First, initialize $d_t = 0$, and set $d_v = 1$ for all v for which $\{vt\}$ is a singleton edge that violates the constraints of Λ_v, and set $d_v = \infty$ otherwise. Now loop for $|V| - 1$ iterations: in each iteration select those vertices v that are not finitely labeled. For each such vertex, consider the set of local edges $N'(v)$ joining v to vertices that are finitely labeled, $N'(v) = \{vv' | v' \text{ is finitely labeled }\}$. If $N'(v)$ violates the constraints of Λ_v then set $d_v = i$.

The proof of correctness of the procedure follows by induction on the number of vertices. The procedure maintains the invariant that after the i^{th} iteration all vertices v that are labeled $d_v \leq i$ have the property that the Λ-distance d_v from v to t is no more than i.

This invariant is clearly true after Step 1. Suppose this invariant is true after the i^{th} iteration. Now suppose, for the sake of contradiction, that there exists a vertex w which was not labeled in the $i+1^{\text{st}}$ iteration, but for which the Λ-distance d_w (from w to t) is i or less. The out-neighbors of w in D must have Λ-distance d_w equal to i or less. By the inductive hypothesis these neighbors are labeled by the algorithm with values i or less. Hence, the algorithm would have labeled w with $i+1$ or less, a contradiction. The proof of the theorem now follows. □

The following minimax result is an immediate corollary of Theorem 2.

Corollary 1 *The minimum over all $\Lambda^c - dags D$ of the maximum length of an vt-path in D is equal to the maximum over all edge sets E' that are consistent with the constraints Λ of the minimum length of vt-paths in the graph G with E' removed,*

$$\min_{\Lambda^c - dags\ D} \{\max \textit{length of vt-path in } D\}$$
$$= \max_{E' \models \Lambda} \{\min \textit{length of vt-path in } G - E'\}.$$

□

5 Landmark Location for Survivable Routing

In this section we investigate an optimization problem related to achieving survivable network routing using a small collections of retransmission or landmark sites. Here we restrict ourselves to the problem of achieving the survivable routing to the landmarks from each of the remaining nodes in the network. Once the set of landmark sites are identified, completing global communication then requires only that the landmarks communicate among themselves, which can be accomplished either through auxiliary channels or by retransmissions through the existing network. The problem of locating such landmark sites is particularly important for vulnerable networks, i.e., networks which may become disconnected by faults consistent with the local constraints of a given model.

We show in this section that the general landmark location optimization problem is NP-hard. However, we are able to show that sufficiently dense graphs always have small sets of landmarks that achieve survivable routing. To show this we prove an extremal result that is optimal over all graphs of with minimum degree δ.

Let $G = (V, E)$ be a graph equipped with a collection of local constraints Λ. Let $C \subset V$, and define G_C to be the graph obtained from G by adding a new vertex v_C and new edges (v_C, u), for each $u \in C$. We say that C is a Λ-*survivable cover* if the graph G_C has a Λ-survivable multitree scheme. Our interest is finding the smallest sets C such that C is a Λ-survivable cover.

Let *SURVIVABLE COVER* denote the decision problem for graphs that asks the question: for a given integer k and a given collection of constraints Λ, does the graph G have a Λ-survivable cover C such that $|C| \leq k$?

Proposition 2 *The problem* SURVIVABLE COVER *is* NP-*complete. The problem remains* NP-*complete when Λ is restricted to a collection of threshold functions.*

Proof: Clearly *SURVIVABLE COVER* belongs to the class *NP*, since for a given C we can verify if C is a Λ-survivable cover by computing a Λ-survivable multitree scheme for G_C in polytime using the algorithm of the previous section.

We reduce the vertex cover problem for regular graphs which is NP-complete [6] to *SURVIVABLE COVER*. An instance of the vertex cover problem is a graph G and an integer k; the problem is to determine if there exists a subset of vertices V' of size at most k such that each edge is incident with a vertex in V'. Consider an instance of the SURVIVABLE COVER problem given by the graph G and collection of constraints Λ consisting of uniform threshold functions Λ_v, where $|\Lambda_v| = \delta - 1$. The Proposition follows with the observation that G has an Λ-survivable cover C iff C is a vertex cover for G. □

Although the previous result shows finding a minimum set of landmarks for survivable routing is NP-hard, we can in polytime find a *minimal* Λ-survivable cover. Such minimal sets are readily computable by repeatedly removing arbitrary elements until we determine the remaining set is not a Λ-survivable cover. We now prove an upper bound on the size of any minimal Λ-survivable cover in terms of the minimum degree δ of the graph and the sum of the threshold values $|\Lambda| = \sum_{v\in V} |\Lambda_v|$.

Theorem 3 *If $G = (V, E)$ is a strict graph with minimum degree δ, and Λ is a collection of locally defined threshold functions, then there exists a set $C \subset V$, such that C is a Λ-survivable cover and $|C| \leq |\Lambda|/(1+\delta)$. Furthermore, this bound is best possible over all graphs of minimum degree δ.*

Proof. The proof of the upper bound is an immediate consequence of Lemma 1 given below. To see that the result is best possible, consider the graph G that is the disjoint union of $n/(\delta+1)$ complete graphs $K_{\delta+1}$, and let Λ be a uniform collection of thresholds with each $|\Lambda_v| = k$. Clearly, any Λ-survivable cover must have k vertices in each complete graph of the disjoint union. Hence, the minimum cardinality set C such that C is a Λ-survivable cover is $kn/(\delta+1) = |\Lambda|/(1+\delta)$. □

We remark that in the following Lemma we prove a somewhat stronger result then the theorem requires by considering the degrees of vertices in the minimal set rather than just the minimum degree δ in the graph.

Lemma 1 *If $C \subset V$ is any minimal Λ-survivable cover, then $|C| \leq |\Lambda| - \sum_{u\in C} deg(u)$.*

Proof. Suppose $C = \{u_1, u_2, \ldots, u_{|C|}\}$ is a minimal survivable cover. For each $1 \leq i \leq |C|$, define $G_i \subseteq G$ to be the largest subgraph of G for which the initial subset $C_i = \{u_1, u_2, \ldots, u_i\}$ is a Λ-survivable cover for G'. Note that $G_0 = \phi$, and $G_{|C|} = G$.

We now define a sequence of mappings $l_0, l_1, \ldots, l_c$ that are used to count how well vertices are covered by each initial subset C_i of C. Each $l_i : V \to \{0, 1, 2, \ldots\}$, where $l_0 \equiv 0$, and each l_i is defined inductively as follows: for each $i \geq 1$:

$$l_i(v) = \begin{cases} |\Lambda_v| & \text{if } v \in V(G_i) \\ |N(v) \cap \{u_1, u_2, \ldots, u_i\}| & \text{otherwise} \end{cases}$$

Since $G_{|C|} = G$, we have $l_{|C|}(v) = |\Lambda_v|$, for each $v \in V - C$. Also for each $v \in V - C$, the difference $l_i(v) - l_{i-1}(v) = 1$ if $v \notin G_{i-1}$ and $vu_i \in E(G)$.

Since C is minimal, it follows that $u_i \notin G_{i-1}$, and furthermore that u_i has fewer than $|\Lambda_{u_i}|$ neighbors in the set $G_{i-1} \cup C$. Therefore u_i has at least $deg(u_i) - |\Lambda_{u_i}| + 1$ neighbors outside the set $G_{i-1} \cup C$. Each of these neighbors contributes to a difference in l values. Thus we have the sum of differences

$$\sum_{v\in V-C} (l_i(v) - l_{i-1}(v)) \geq deg(u_i) - |\Lambda_{u_i}| + 1.$$

Now we sum this over all $1 \leq i \leq |C|$, and we obtain the the following inequality

$$\sum_{1\leq i\leq |C|} \Big(\sum_{v\in V-C} (l_i(v) - l_{i-1}(v)\Big) \geq \sum_{1\leq i\leq |C|} (deg(u_i) - |\Lambda_{u_i}| + 1)$$

The double sum on the left is telescoping and equal to $\sum_{v\in V-C}(l_{|C|}(v)) = \sum_{v\in V-C}(|\Lambda_v|)$. Hence we have that

$$\sum_{v\in V-C} (|\Lambda_v|) \geq \sum_{1\leq i\leq |C|} (deg(u_i) - |\Lambda_{u_i}| + 1).$$

Simplifying and collecting terms we have

$$|C| \leq |\Lambda| - \sum_{u\in C} deg(u).$$

The lemma and the theorem now follow. □

6 Conclusion

We introduced in this paper a new distributed model for network survivability. We showed that the general problem of determining the survivability of a network under simple local constraints is co-NP-hard. We showed that a closely related model of local constraints, called the half-duplex model is highly tractable. In fact, we showed how to effectively construct optimal strategies for fault-tolerant routing with respect to such constraints. The routing strategy consisted of the construction and application of multi-parent routing tables that are survivable under all fault configurations satisfied in the model. The work in this paper can be easily extended to deal with more complex fault models by analyzing the effects of a combination of several different local constraint models $\Lambda_1, \Lambda_2, \ldots, \Lambda_k$ using a probability distribution over these models.

6.1 Game Theoretic Interpretation

The problem of determining survivable routing schemes in the half-duplex model can be cast in terms of a two-player game played on a network involving a router and a blocker. Consider an *st-path game*, where the router has the objective of moving a packet from a source node s to a destination node t using the fewest (or least cost) edges; and, the blocker has the objective of inhibiting (or maximimizing the cost of) the movement of that packet. The payoff of a play of the game is the length of the st-path (or walk) achieved by the router. We define the *adversarial distance* to be the value of the game, i.e., the expected payoff under optimal mixed strategies. Particular games are defined by the set of prescribed rules that the players must follow in the st-path game; here one can associate the (directed) local constraints as the set of rules on the blocker. Once a game is defined, we consider the algorithmic complexity of determining a solution for the st-path game, or in other words, determining optimal strategies for the router and the blocker.

Given a graph $G = (V; E)$, a source node s and a destination node t, and a collection of local constraints Λ, consider the st-path game defined by $\Gamma = (G; s; t; \lambda)$. A fixed blocker strategy bs and a fixed router strategy rs determine a play of Γ that achieves a payoff denoted by $p(bs; rs)$. The routing distance (or upper value) of the game Γ is given by

$$rd_\Gamma = \min_{rs} \max_{bs} p(bs; rs);$$

that is, the minimum over all routing strategies rs of the maximum over all blocking strategies bs of the payoff $p(bs; rs)$. Likewise, the blocking distance (or lower value) of the game Γ is given by

$$bd_\Gamma = \max_{bs} min_{rs}\, p(bs; rs);$$

that is, the maximum over all blocking strategies bs of the minimum over all routing strategies rs of the payoff $p(bs; rs)$. Since the game is played with perfect information, it follows from basic game theory that these distance are all equal, i.e., $ad_\Gamma = bd_\Gamma = rd_\Gamma$. The results of Section 4 allow us to algorithmically determine this adversarial distance.

Corollary 2 *There is an efficient $O(n^2)$ algorithm that computes the adversarial distance ad_Γ from s to t in any st-path game.* $\square$

Acknowledgment

This research has been supported in part by NSF Grant No. CCR-9877139 and an Ohio Board of Regents' Research Investment Grant.

References

[1] F. Annexstein, K. A. Berman, T. Hsu, and R. Swaminathan. A multi-tree generating routing scheme using acyclic orientations. In *Computing and Combinatorics, Proceedings of the third annual international conference, Lecture Notes in Computer Science, 0302-9743.*, 1997. To appear in Theoretical Computer Science.

[2] F. Annexstein, K. A. Berman, and R. Swaminathan. Computing nearly optimal multi-tree paths for reliable routing. In *Proceedings of 4th International Colloquium on Structural Information and Communication Complexity (SIROCCO '97)*, 1997.

[3] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, New York, 1992.

[4] R. Cohen, B. Patel, F. Schaffa, and M. Willebeek-LeMair. The sink tree paradigm: connectionless traffic support on atm lans. *IEEE/ACM Trans. on Networking*, 4:363–374, 1996.

[5] E. Gafni and D. Bertsekas. Distributed algorithms for generating loop-free routes in networks with frequently changing topology. *IEEE Transactions on Communications*, COM-29(1):11–18, 1981.

[6] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to NP-Completeness*. W.H. Freeman and Company, San Francisco, California, 1979.

[7] F. Granot, M. Penn, and M. Queyranne. Disconnecting sets in single and two-terminal-pair net works. *Networks*, 27(2):117–123, 1996.

[8] D. Gusfield and D. Naor. Extracting maximal information about sets of minimum cuts. *Algorithmica*, 10(1):64–89, 1993.

[9] A. Itai and M. Rodeh. The multi-tree approach to reliability in distributed networks. *Information and Computation*, 79:43–59, 1985.