SURVIVABLE NETWORKS

Algorithms for Diverse Routing

Ramesh Bhandari

KLUWER ACADEMIC PUBLISHERS

SURVIVABLE NETWORKS
Algorithms for Diverse Routing

by

Ramesh Bhandari, Ph.D.
AT&T Laboratories
New Jersey

Kluwer Academic Publishers
Boston/Dordrecht/London

Distributors for North, Central and South America:
Kluwer Academic Publishers
101 Philip Drive
Assinippi Park
Norwell, Massachusetts 02061 USA
Telephone (781) 871-6600
Fax (781) 871-6528
E-Mail <kluwer@wkap.com>

Distributors for all other countries:
Kluwer Academic Publishers Group
Distribution Centre
Post Office Box 322
3300 AH Dordrecht, THE NETHERLANDS
Telephone 31 78 6392 392
Fax 31 78 6546 474
E-Mail <orderdept@wkap.nl>

Electronic Services <http://www.wkap.nl>

Library of Congress Cataloging-in-Publication Data

Bhandari, Ramesh.
Survivable networks : algorithms for diverse routing / by Ramesh Bhandari.
p. cm. -- (The Kluwer international series in engineering and computer science : SECS 477)
Includes bibliographical references and index.
ISBN 0-7923-8381-8
1. Computer networks--Reliability. 2. Telecommunication--Traffic. 3. Algorithms. I. Title. II. Series.
TK5105.5.B48 1999
004.6--dc21 98-45585
CIP

1) R. Bhandari, "Optimal Diverse Routing in Telecommunication Fiber Networks," Proceedings of IEEE Infocom, pages 1498-1508 (1994)
2) R. Bhandari, "Optimal Physical Diversity Algorithms and Survivable Networks," Proceedings of the Second IEEE Symposium on Computers and Communications, pages 433-441 (1997)
3) R. Bhandari, "On Congestion Problems in Digital Cross-Connect (DCS) Transport Networks," Proceedings of the Third IEEE Symposium on Computers and Communications, pages 503-509 (1998).
Published by Kluwer Acacemic Publishers

Printed on acid-free paper.

Printed in the United States of America

SURVIVABLE NETWORKS
Algorithms for Diverse Routing

THE KLUWER INTERNATIONAL SERIES
IN ENGINEERING AND COMPUTER SCIENCE

NOTE TO TEACHERS

As an aid to teachers interested in teaching this book as a course, exercises suitable for classroom instruction and sample lecture notes are available by contacting the author at *rbhandari@att.com* or *rbhandari@worldnet.att.net*.

To

My Parents and My Family

PREFACE

With the advent of fiber and its increasing deployment in networks, the risk of losing large volumes of traffic data due to a span cut or node failure has escalated. Because of fierce competition among service providers and customers' intolerance of disruption of service, survivability of a network has assumed great importance. Survivability refers to the ability of a network to provide continuity of service with no disruption, no matter how much the network may be damaged due to events such as fiber cable cuts or node failures (due to equipment breakdown at a central office or other disasters such as fires, flooding, etc).

The book considers the common mesh-type telecommunication network and describes in detail the construction of physically disjoint paths algorithms for diverse routing. The key features of the algorithms are optimality as well as simplicity. The algorithms are a consequence of research work emanating from a real project related to enhancing the survivability of a large network. As a result, the algorithms are not only applicable to the standard graph-theoretic networks described by nodes and links, but also to the actual real-life networks described by nodes, links, span nodes, and spans; algorithms for such types of networks are *not* found in current books.

The book was motivated by the constant interest shown in the diversity algorithms and the need expressed for them by telecommunication network planners and designers. They have often been discouraged by the lack of simplicity in previous work. The repeated requests for the author's recently published work, its translation into Hungarian, and requests for the developed computer codes gave the final impetus for the writing of this book.

Although the algorithms have been developed from the view point of survivability of communication networks, they are in a generic form, and thus applicable in other scientific and technical disciplines to problems that can be modeled as a network. Some examples of other physical networks to which the current work may apply are networks of computer chips, electric power systems, highways, etc.

The book utilizes fundamental arguments in the development of the algorithms. The use of mathematical symbols is minimized to keep the book readable. Yet proofs are given, when necessary, to put the algorithms on a firmer footing. Overall, the approach is one of pedagogy, where concepts and algorithms are frequently illustrated with figures and examples. As a result, this book is intended for 1) practicing network designers and planners, who can skip the development procedures and simply refer to the stated algorithms and examples 2) researchers who want to explore the subject matter further 3) teachers and students in computer science, telecommunications, and related disciplines 4) any layman familiar with the rudiments of graph theory.

The book is an outgrowth of author's personal research and investigation, originating in a real project. He thanks colleagues for useful conversations at the inception of the project; summer students, Casimir Wierzynski, Trac-Duy Tran, and Katherine Nguyen, who worked with the author to develop computer codes; Dr. Art O'Leary and Dr. Gerry O'Reilly for their readings of the manuscript and their helpful comments and suggestions; friends and relatives, especially his mother, for their encouragement. Last but not least, the author is grateful to his wife, Savita, for not only her understanding and moral support during the writing of the book, but also for generating most of the line drawings; to his sons, Simit and Rohan, for understanding their father's inability to be with them on numerous occasions.

Ramesh Bhandari
AT&T Laboratories,
New Jersey.

CONTENTS

PREFACE **xi**

1 INTRODUCTION **1**
1.1 Overview 1
1.2 Definitions and Notations 14
1.3 Types of Graphs Considered 18

2 SHORTEST PATH ALGORITHMS **21**
2.1 Introduction 21
2.2 The Dijkstra Algorithm 22
2.3 The Modified Dijkstra Algorithm 29
2.4 The Breath-First-Search (BFS) Shortest Path Algorithm 33
2.5 Concluding Remarks 35

3 SHORTEST PAIR OF DISJOINT PATHS ALGORITHMS **39**
3.1 Introduction 39
3.2 The Simple Two-Step-Approach Algorithms and their Shortcomings 40
3.3 Edge-Disjoint Shortest Pair Algorithm Development 46
3.4 Vertex-Disjoint Shortest Pair Algorithm Development 68
3.5 Suurballe's Disjoint Pair Algorithms 86
3.6 Comparison of the Developed Algorithms and the Suurballe's Algorithms 92

4 MAXIMALLY DISJOINT PATHS AND PHYSICAL DIVERSITY VERSUS COST ALGORITHMS **93**
4.1 Implementation of Edge-Disjoint Paths Algorithms 94
4.2 Implementation of Vertex-Disjoint Paths Algorithms 99
4.3 Physical Disjointness versus Cost 108

5 PHYSICALLY DISJOINT PATHS IN REAL-LIFE TELECOMMUNICATION FIBER NETWORKS **117**
5.1 Fiber Network Description 118
5.2 Shortest Pair of Physically Disjoint Paths Algorithm 123
5.3 Application to Other Possible Configurations 148
5.4 Concluding remarks 155

6 MAXIMALLY DISJOINT PATHS ALGORITHMS FOR ARBITRARY NETWORK CONFIGURATIONS **157**
6.1 General Approach 158
6.2 Heuristic Based on Vertex/Link-Commonness Penalty 158
6.3 Optimal Algorithm for Maximal Disjointness 162

7 K(> 2) DISJOINT PATHS ALGORITHMS **175**
7.1 K(> 2) Edge-Disjoint Paths Algorithms 175
7.2 K(> 2) Vertex-Disjoint Paths Algorithms 179
7.3 Implementation of K(> 2) Disjoint Paths Algorithms 182

8 DISJOINT PATHS: MULTIPLE SOURCES AND DESTINATIONS **183**
8.1 Single Source, Two Destinations 183
8.2 Two Sources, Single Destination 184
8.3 Two Sources, Two Destinations - Unfixed Source/Destination Pairs 185
8.4 Two Sources, Two Destinations - Fixed Source/Destination Pairs 186

FURTHER RESEARCH **189**

REFERENCES **191**

INDEX **195**

1

INTRODUCTION

This chapter provides an overview of the contents of the book as well as the definition of the graph-theoretic terms used within the book.

1.1 OVERVIEW

As networks modernize and expand with the increasing deployment of optical technology, the large bandwidth offered by the optical fiber offers tremendous potential for exploitation. The number of services offered to customers over a fiber network is proliferating. But at the same time in today's highly competitive environment, where customers have a large choice of service providers, the customers have come to expect the highest quality of service, including sustained continuity of service during the time they pay for the service. Service disruption due to a network failure is no longer being tolerated, since it can cause the customers significant loss of revenue during the network down time. Such loss of revenue can lead to bad publicity for the service provider and erosion of customer base due to customer dissatisfaction.

Major network failures are essentially of three types:

- Node failure due to equipment breakdown or equipment damage resulting from an event such as an accidental fire, flood, or earthquake; as a result, all or some of the communication links terminating on the affected node may fail.

- Link failure due to inadvertent fiber cable cut; the fiber cable carrying traffic from one telecommunication office to another is buried approximately 3 feet underground in a conduit, but due to the ubiquitous construction activities as world economies grow rapidly, accidental fiber cuts occur frequently, despite increased network care and maintenance efforts.

- Software failure that can impact a large portion of the given network, and is, in general, hard to identify and recover from; this type of failure is relatively rare and not considered in this book.

Thus, in what follows, the failure of a communication path is assumed to be due to a node or link breakdown along the path.

Network survivability against physical network failures is a broad and widely discussed area, covering a gamut of network designs ranging from a restorable network based on the digital cross-connect systems (DCS's) to self-healing fiber rings (see, e.g., [1-16]). Despite the recent interest and progress in rings [8-16], mesh networks based on the DCS architecture are still the most prevalent globally, and will continue to be so due to the heavy expense and effort involved in conversion to rings. This book therefore considers the common mesh-type network based on the DCS architecture.

Restorability of affected traffic within a DCS-based network has been the subject of immense investigation in the past [1-8]. Several techniques including a large number of diverse mathematical models have been employed (see, e.g., [6]) in detailed network designs. Some of the methods [6, 7] have focused on creating networks that are K-connected, i.e., for every pair of nodes, the network sought has K ($\geq$2) disjoint paths. The intent of this book is not to describe these methods or techniques. Rather, the book focuses on enhancing the survivability of a given DCS-based network via provisioning of services over disjoint paths so that if one path fails, communication can still proceed by cross-connecting the affected traffic over the predetermined alternate (disjoint) paths. It explores the different network configurations inherent in real-life networks, and provides algorithms for diverse routing, i.e., for finding physically disjoint paths between a pair of nodes, the two nodes being the two endpoints between which continuity of service is desired. Emphasis is on optimality, i.e., on finding the shortest (or least cost) set of disjoint paths, since use of such optimal paths reduces network costs for providing such services.

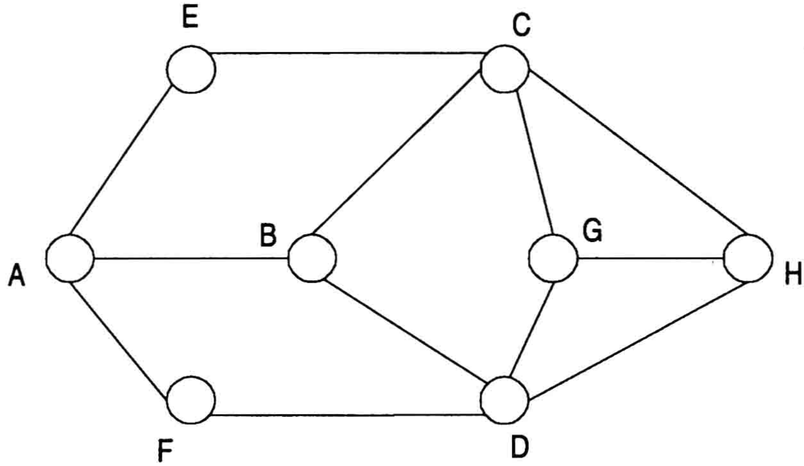

Figure 1.1 A network of nodes and links

In its simplest form, a DCS-based transport network is composed of nodes connected by physical links such as fiber cables in today's modern network. Figure 1.1 shows such a network; circles represent office sites or nodes where communication equipment such as the DCS for rerouting of traffic resides whereas straight lines represent links or physical connections such as the fiber cables between the different pairs of nodes. Traffic flows within the network from one node to another along the connecting link, being cross-connected at a node to another link, depending upon its final destination.

Figure 1.2 shows a possible scenario at a network node. Assuming this node to be node B in Figure 1.1, traffic from node A reaches node B on the fiber cable optically, which is then converted to an electric form for rerouting by a DCS (a light-terminating equipment, indicated by a triangle in Figure 1.2, converts an optical signal into an electric signal, and vice versa). Some of this traffic reaching node B on the fiber from node A is rerouted on the fiber cable connected to node C and some on the fiber cable connected to node D, with the remaining traffic terminating at node B itself (note that some traffic could also originate at node B, but this is not shown). Thus, within the network of Figure 1.1, if there is traffic originating at node A and destined for node C, it may take the path ABC, being cross-connected to link BC by an appropriate equipment such as the DCS at node B. The same traffic could also be routed along path AEC, in which case it would be cross-connected at the intermediate node E, or path AFDGC, in which case it would be cross-connected at the intermediate nodes F, D, and G, and so on. What path the traffic from A to C in Figure 1.1 takes depends upon such factors as the availability of circuits along a given link and the cost of traversing the link. When

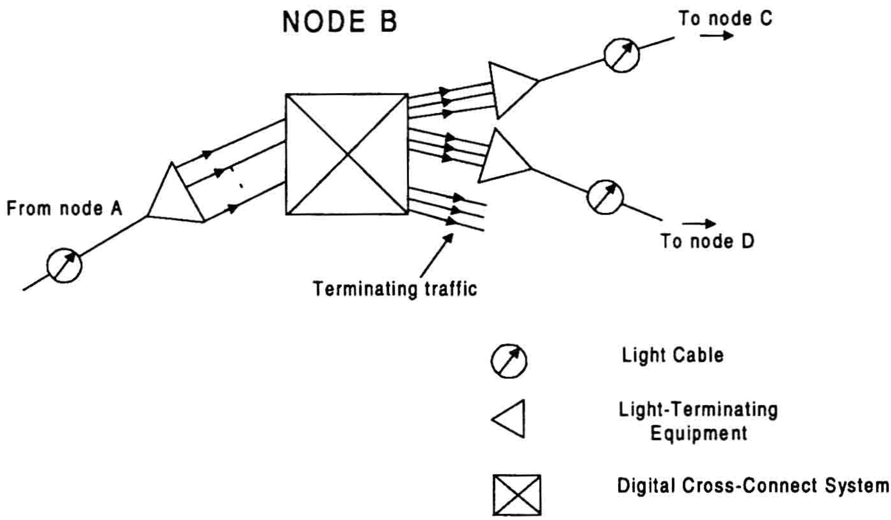

Figure 1.2 A possible scenario at node B of Figure 1.1; © 1998 IEEE.

no circuit on a given link is available for routing of traffic, we consider the link unavailable within the network.

In this book, we shall assume that a network is specified by nodes and only those links that are available for routing. Furthermore, consideration for routing traffic along a given path of the network is the cost of routing along that particular path. As a result, we focus on *least cost* paths. The cost of using a path for routing is composed of the individual costs of the links making up that path. The cost of a link is determined essentially from the physical length of the link, the capacity of the link, and the cost of terminating or cross-connecting the traffic at the end of a link, the latter cost being dependent upon the cost of light-terminating equipment and cross-connect equipment such as the DCS.

Similarly, when diversity or disjointness is desired, we attempt to select that pair of paths whose total cost is a minimum. For example, between nodes A and C in Figure 1.1, one observes that there are several pairs of disjoint paths possible: (AEC, ABC), (AEC, AFDBC), (ABC, AFDGC), (ABC, AFDHC), etc. The most desirable pair of paths from the standpoint of minimizing costs would be the least cost pair of paths, i.e., the pair of disjoint paths whose sum is a minimum. Refer to Figure 1.3, which is the network of Figure 1.1 with the cost of each link specified in some units. The shortest or least cost path between nodes A and C is ABDGC of

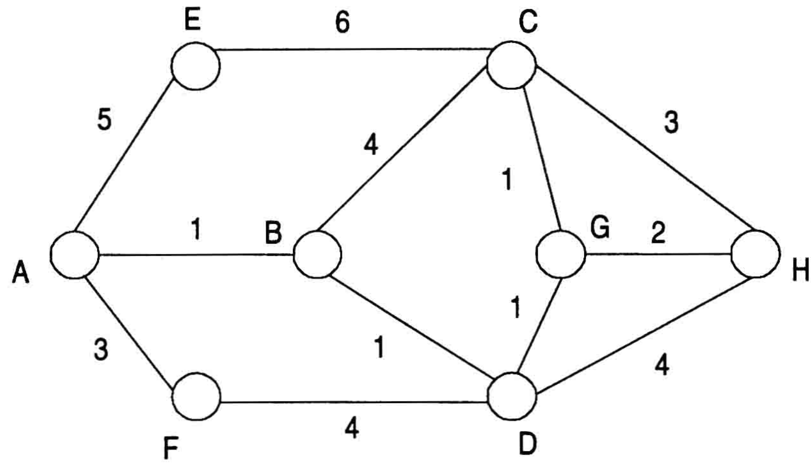

Figure 1.3 The network of Figure 1.1; the numbers specify the cost of traversing the links.

length = 1 + 1 + 1 + 1 = 4, while the least cost pair of disjoint paths between A and C is (ABC, AFDGC) of length = 5 + 9 = 14. It is interesting to note that the least cost pair does not include the *least cost* path.

The book is divided into 8 chapters. Chapter 2 focuses on finding the shortest (or least cost) path between the given pair of nodes in a given network; the network is described by nodes, links, and their traversal costs; the cost of traversing a link generally includes the cost of the node on which the link terminates (refer to remarks in Section 4.3.3, when the node costs are specified separately). A plethora of shortest path algorithms and their implementations exists in the literature [17-29]. Emphasis here is on description of those that are simple in form, yet suitable in the construction of disjoint paths. One such algorithm, called the modified Dijkstra algorithm [29], obtains from a slight tweak of the standard Dijkstra algorithm [18, 21].

Chapter 3 deals with the development of the algorithm for the shortest pair of disjoint paths between the given pair of nodes. As noted in our earlier discussion with respect to Figure 1.1, several pairs of disjoint paths can be possible for a given pair of nodes in the network. The major aim here is to construct algorithms for finding the optimal pair, i.e., the shortest (or least cost) pair of disjoint paths. In sharp contrast to an earlier method due to Suurballe [30, 31] of employing a special network transformation, the development procedure described here avoids such transformations; it further circumvents the need of the general (relatively slow) shortest path algorithm like the Ford's algorithm [17, 21, 30], invoking instead the simple and straightforward alternatives introduced in Chapter 2. The adopted approach yields simple forms of the disjoint paths algorithms, facilitating their easy application by practicing engineers. Further, they lend themselves to easy extensions to related problems, including the case of K disjoint paths (K > 2), dealt with in subsequent chapters.

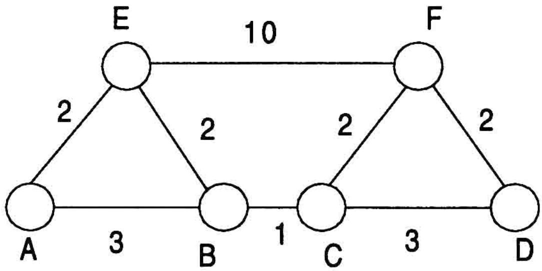

Figure 1.4 A 6-node, 8-link network

Chapter 4 focuses on the practical implementation of the disjoint pair algorithms developed in Chapter 3, and discusses the situation when full disjointness of the pair of paths does not exist, or is not desired by customers to reduce their own costs [32]. Refer to Figure 1.4, which shows a network of 6 nodes and 8 links; as before, the numbers specify link costs, i.e., the cost of traversing the links; e.g., the cost of routing traffic on a circuit along link EF is 10 units, and along link AB is 3 units, and so on; normally, shorter the physical length of the link, smaller the cost of a circuit over that link. For traffic flowing between the pair of nodes A and D, ABCD is the shortest (least cost) path with a cost of 3 + 1 + 3 = 7 units. The two physically disjoint paths between nodes A and D are ABCD and AEFD with a total cost of 7 + 14 = 21. If we permit some diversity violation, then provisioning of two circuits, one on the path ABCD and the other on the path AEBCFD, would result in commonness of the short link BC, with a total cost for the two paths = 7 + 9 = 16 units, which is 5 units less than the cost for full diversity. Similarly, for traffic from B to F, the cost of the shortest pair of disjoint paths BCF and BEF is 3 + 12 = 15 units, while the two paths BCF and BCDF, having the short link BC in common, cost only 3 + 6 = 9 units. Thus, provisioning of two circuits on the latter pair of paths is cheaper by (15 - 9)/15 x 100% = 40%, and the customer by accepting the diversity violation of link BC commonness saves money. Clearly, the cost of two paths is maximal when complete (100%) diversity is desired, and is minimal in the extreme case of the two circuits traversing the same (shortest) path, which is a case of minimal (zero) diversity. When the level of diversity is less than 100%, the cost lies between the above two extremes. Chapter 4 provides an algorithm that permits evaluation of cost savings when the two paths sought are allowed to have partial overlaps.