University of Alberta

Availability of Service in
Mesh-Restorable Transport Networks

by

Matthieu Arnold Henri Clouqueur ©

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of *Doctor of Philosophy*

Department of *Electrical and Computer Engineering*

Edmonton, Alberta, Canada
Spring 2004

Your file Votre référence
ISBN: 0-612-96251-2
Our file Notre référence
ISBN: 0-612-96251-2

Abstract

We present a collection of studies related to the availability of service in mesh-restorable transport networks. The issue of service availability in transport networks is of increasing importance due to our always increasing reliance on telecommunication networks, which now requires that network operators give service-availability guarantees to their customers in the form of service level agreements (SLA). It is therefore important to study network architectures in terms of the level of service availability they provide, and not just on capacity efficiency, simplicity of operation, cost, etc. This work deals with this issue for transport networks using a mesh architecture and more specifically for dynamic mesh-restorable networks.

After a long era where ring-based systems dominated the world of transport networking, the mesh architecture appears to be the solution of choice for future optical transport networks. The mesh architecture is already known for its high capacity efficiency (little capacity is required to serve demands and to guarantee restorability to failures), the simplicity of service provisioning (new service paths can usually be provisioned on shortest path), and for its good scalability (capacity only needs to be added where exhausted therefore making the network scaling easy). However, the issue of service availability in mesh-restorable networks was yet to be investigated.

We first develop theoretical and practical approaches for determining the availability of service paths in mesh-restorable transport networks that are designed to be restorable to any single span-failure. This initial treatment of the problem shows the major influence that dual span-failure restorability has on service availability—the importance of considering dual span-failure scenarios is later reinforced by a study of the effects of maintenance actions on network restorability. It is also demonstrated that, despite what intuition tells us, restoration time has in fact little impact on service availability. The influence of various factors on dual-span failure restorability (and therefore on availability) is also investigated.

We then present several network capacity design methods for serving demands with various restorability requirements: first considering the case of demands with higher restorability requirements (single span-failure restorability and dual span-failure restorability), and then considering a multi quality of protection environment with demands having restorability requirements ranging from single span-failure restorability down to preemptible services.

The thesis also provides an initial theoretical treatment of the availability of service in networks using *p*-cycle protection—a new survivability scheme offering fast restoration and high capacity efficiency—, and a comparison between the availability of service with dynamic mesh-

restoration and with ring-based protection.

Experimental results are obtained using a set of detailed test networks of various sizes and connectivity. The results show the very high potential of mesh-restorable networks to simultaneously serve various types of availability requirements from low requirements to requirements of ultra high availability—higher than the availability of paths in ring-based networks—through the use of a multi QoP environment, and to benefit from that multi service aspect to achieve high capacity efficiency.

Dedication

À ma femme, Caroline, pour son soutien et son amour sans limite.

À mes parents, André et Sylvaine, à qui je dois tout.

Acknowledgements

I am deeply grateful to my wife, Caroline, whose love and support helped me every day during these years of hard work. My parents and my family are also thanked for their constant support.

I am also sincerely greateful to my supervisor, Dr. Wayne D. Grover, for providing many of the ideas that lead to this thesis, for all the trust he has placed in me and for his constant enthusiasm and inspiration.

The students of the Network Research Group at TRLabs are thanked for their constant support. In particular, I would like to thank Dr. Dave Morley and John Doucette for the helpful discussions we had on many occasions. Demetrios Stamatelakis is thanked for helping me so many times with programming, UNIX, or computer-related questions in general.

The TRLabs staff is thanked for providing a very professionnal and efficient work environment, allowing me to concentrate on my research. In particular, I wish to thank Linda Richens for being always so helpful and for her incredible patience and kindness.

I would like to thank several of the TRLabs students for making TRLabs such a pleasant and fun environment: David Mazzarese, Dion Leung, Kent Lam, Yan Xin, Fengqin Zhai, Harini Gopaluni, and P. Nakeeran, among others. Thanks to them for all the fun discussions and the good times spent in the TRLabs lunchroom.

Finally, I would like to express my sincere gratitude to the other members of my examining committee, in particular Dr. Bruce Cockburn and Dr. David Tipper, for the thesis corrections they suggested. Their help was much appreciated.

Table of Contents

List of Tables

List of Figures

List of Symbols and Abbreviations

ADM	Add/drop multiplexer
APS	Automatic protection switching
ATM	Asynchronous transfer mode
BLSR	Bidirectional line-switched ring
CLP	Cell loss probability
CO	Central office
DCS	Digital crossconnect system
DPP	Distributed pre-planning
DoS	Data over SONET
DWDM	Dense wavelength division multiplexing
FTP	File transfer protocol
GFP	Generic framing procedure
HDRP	Hypothetical digital reference path
ILP	Integer linear program
IP	Internet protocol
IP	Integer programming
ISO	International organization for standardization
IXC	Inter-exchange carrier
JCP	Joint capacity placement
ksp	k-shortest paths
LAN	Local area network
LCAS	Link capacity adjustment scheme
LEC	Local exchange carrier
LMP	Link management protocol
LP	Linear programming
MAN	Metropolitan area network
MIP	Mixed integer programming
MIPGAP	Optimality gap tolerance in CPLEX

MJCP	Modular joint capacity placement
MPLS	Multiprotocol label switching
MSCP	Modular spare capacity placement
MTBF	Mean time between failures
MTTR	Mean time to repair
OADM	Optical add/drop multiplexer
OAM&P	Operations, administration, maintenance, and provisioning
OC-N	Optical carrier signal that transmits N times 51.84 Mb/s
O-D	Origin-Destination
O/E/O	Optical-to-electrical-to-optical
OSC	Optical service channel
OSI	Open systems interconnect
OTDM	Optical time division multiplexing
OTN	Optical transport network
O/O/O	All optical
OXC	Optical crossconnect
POS	Packet over SONET
PR	Path restoration
PSTN	Public switched telephone network
QoP	Quality of protection
QoS	Quality of service
R_1	Single-failure restorability
R_2	Dual-failure restorability
SBPP	Shared backup path protection
SCP	Spare capacity placement
SDH	Synchonous digital hierarchy
SLA	Service level agreement
SONET	Synchronous optical networking
SPE	Synchronous payload envelope
SR	Span restoration

SRLG	Shared risk link group
STM	Synchronous transfer mode
STS	Synchronous transport signal
TCMF	Two-commodity maximum-flow
TDM	Time division multiplexing
UPSR	Unidirectional path-switched ring
VC	Virtual circuit
VC	Virtual concatenation
VCI	Virtual circuit identifier
VP	Virtual path
VPI	Virtual path identifier
VT	Virtual tributary
VWP	Virtual wavelength path
WAN	Wide area network
WDM	Wavelength division multiplexing
WP	Wavelength path

1. Introduction

1.1 Presentation of the Problem and Goals

1.1.1 Problem

The reliability of transport networks has been an issue for a long time. With the always increasing amount of communications being transported over telecommunication networks, and our also increasing reliance on them, network reliability has become a central issue in the way networks are being designed. Over time, network operators have started using more and more sophisticated ways to improve the reliability of their networks. This started by improving the reliability of individual networks elements by using more reliable components, or by making these network elements internally redundant. This improved the reliability of the network but did not solve the problem of failures being caused by external events (e.g. caused by humans or by nature). Among this type of failures, fibre cable cuts caused by construction crews are the most prominent [Fla90]. To withstand such failures the concept of routing diversity was developed in which diversely routed backup routes are reserved to replace the default working route in case of failure of the latter. Diverse routing brought a major improvement to the reliability of services in transport networks and it became the norm to guarantee full single failure restorability—i.e., restorability to any possible fibre cut—to all services in transport networks [Fla90][McD94][Wu95]. Various survivability schemes were developed to meet this requirement [Wu95]. Networks using such survivability schemes have been referred to as *fully restorable networks* or *fully protected networks*. Although these measures indisputably brought an improvement of the service reliability, fully restorable networks are not completely immune from failures. Multiple failures, which are less frequent than what could be thought (as will be seen later in this thesis) can still affect services. Put in simple terms, fully restorability as it is usually meant (i.e., full restorability to any single fibre cut) does not guarantee that service outages will not happen. This brings us to the central question of this thesis: "How much is the reliability of networks increased by making them restorable to single cable cuts?"

1.1.2 Goals

To answer the question posed in the previous section, we will have to first define how to characterize the network's reliability. This term, used here in its every-day English meaning has a precise technical definition, which does not in fact correspond exactly to the question raised here. In Chapter 3 we will see why the concept of *availability* is more appropriate to characterize the problem studied in this thesis. We will also explain why we are more interested in determining the availability of *service paths* rather than in the availability of the network as a whole. The question will then be re-stated more precisely as: "How much is the service path availability improved by making the transport network restorable to single cable cuts?"

Having stated the problem clearly, we will then develop the concepts and methods necessary to answer the question. Based on these methods and concepts, we will then investigate the influence of various factors on the availability of service. A first factor studied will be maintenance actions and their influence of the restorability of mesh networks. Maintenance actions are much more frequent than actual failures and it is important for network operators to be able to assess the risk of outages linked to them. Another factor will be the presence of multiple classes of service. We will investigate how multiple classes of service can be served simultaneously in transport networks and will study how their availability properties relate to each other and to the proportions of services in each class. This thesis will also investigate the question of availability of service in networks using p-cycle protection. p-Cycle protection is a new survivability scheme offering low capacity requirements and fast restoration (two characteristics that were long thought to be incompatible), and which is receiving more and more attention from industry and university researchers.[GrS98][GrS00][GrD02][Sch03a][Sch03b]

Another important part of this thesis is devoted to the development of capacity design methods for mesh-restorable networks with enhanced restorability properties. With such networks, we leave the traditional full restorable paradigm to enter a new paradigm in which some demands are now restorable to single and dual failures. The work on this topic will show how well mesh-restorable networks lend themselves to accommodating various mixes of restorability requirements and how this can be done with the means to control the increase in capacity requirements.

This thesis is specifically focused on transport networks with a mesh architecture and that use span restoration as a survivability scheme. However, other schemes will be considered and this thesis work will compare various mesh survivability schemes and also ring protection schemes, based on restorability and availability studies.

1.2 Thesis Organization

The remainder of this chapter introduces the specific notations and terminology used in this thesis and briefly presents some concepts related to optimization theory and, in particular, Mathematical Programming, which are extensively used in this thesis. The reader will be referred to additional references for more details.

Chapter 2 contains background information on transport networks. The chapter starts by explaining the concepts of client/transport relationship and of network demand which are fundamental in this research work. Then are presented the different classes of transport networks and the different technologies that have been developed including the optical transport network, which is the main type of transport network studied in this thesis work. An important part of the chapter is also devoted to survivability schemes with a special emphasis on mesh-based schemes. The chapter then covers the topics of transport network design, introducing several design methods that serve as the basis for new design methods presented in later chapters. Finally, we present what we see as the important challenges of transport networking in the future and explain how some of these issues relate to the present work.

Chapter 3 is devoted to a theoretical approach to the problem of determining the availability of service in transport networks. First, important mathematical definitions related to availability are presented to provide a solid ground for the analysis. Then general existing methods for the analysis of system availability are presented, which are then considered in the specific context of telecommunication systems. A review of the literature on the topic of network availability analysis is then presented. The chapter finishes with a new proposed approach to the availability analysis problem. This approach is then applied in following chapters.

Chapter 4 presents an approach to the problem of determining the restorability to dual-failures based on computational re-routing trials. It is first shown how difficult (if not impossible) it is to develop closed-form equations for the restorability of mesh networks using adaptive restoration mechanisms. Three pseudo-code algorithms are then presented to model various progressively more adaptive restoration mechanisms for span-restorable networks. These algorithms are then used in a series of experiments to investigate the dual-failure restorability properties of span-restorable mesh networks designed for full restorability to single failures. The influence of several factors on the restorability results are then investigated to get some insight as to which factors influence service availability the most. In particular, we investigate whether it is necessarily true that the "more capacity means higher availability."

Chapter 5 uses the methods developed in Chapter 4 to investigate the influence of maintenance actions on the availability of service. First, three models of maintenance actions are defined relative to the use of spare capacity that they require. Two of these models are then compared in terms of their effects on the restorability of the network during the maintenance action. The chapter also provides ideas on how to manage the effects of maintenance actions by identifying maintenance actions that can be conducted simultaneously and maintenance actions that should be conducted in series in order not to expose services to high risks of outage.

Chapter 6 presents extensions of the common mesh-restorable network capacity design formulation that enhance the dual-failure restorability of the designs. These new design formulations include a formulation for capacity minimization under the constraint of complete dual-failure restorability, a formulation for restorability maximization under a given total capacity cost budget, and a formulation for minimum-capacity design supporting multiple-restorability service class definitions. In the third formulation the restorability options range from no restorability guarantee to the guarantee of full restorability to any dual-failure. The work of Chapter 6 shows how to economically support an added service class in the upward quality direction and to tailor the investment in capacity to provide dual-failure restorability on a selective basis.

Chapter 7 presents a theoretical study of service availability in *p*-cycle networks. Closed-form equations are developed for the availability of paths. These equations show that there can be a significant difference in the availability of paths depending on whether they lie on the *p*-cycles or just straddle them. Based on the equations developed, we compare the influence of two factors on the availability of service in networks protected by *p*-cycles. Suggestions are provided on possible ways to control the availability of paths in these networks. The chapter also provides a restorability analysis comparison between optimal *p*-cycle designs and corresponding optimal span-restorable designs.

In Chapter 8, we introduce a capacity design formulation for span-restorable networks with multiple classes of protection. A study of the capacity requirements of these networks is presented that shows how, in many cases, it is possible to design networks with some restorable demands and without any extra protection capacity. We then investigate how a multi quality of protection (QoP) environment affects the availability of the different service classes and consider the case of four different multi-QoP restoration models. An economic interpretation of the results is also provided.

In Chapter 9, we present a study comparing the availability of service in ring and mesh-restorable networks. One of the main motivation of the study is to verify whether the higher capacity redundancy of rings compared to mesh results in higher availability. The comparison is based on detailed simulations of the network's response to random sequences of failures and repairs. Both survivable architectures are tested under the exact same conditions. The study also shows the potential of the mesh architecture to provide very high availability to a small fraction of selected high-priority service paths when prioritization in the restoration is introduced, while keeping the availability of lower-priority service paths almost unchanged.

Finally, Chapter 10 concludes by summarizing the thesis work and identifying its main contributions. The chapter also gives ideas for future research on the topic.

1.3 Definitions and Notations

1.3.1 Basic Mathematical Notations

For clarity, we will observe the following rules regarding mathematical notations as much as possible.

A set is an unordered ensemble of elements and will be denoted by a letter of the Roman alphabet, in upper case, bold and in italics like $\boldsymbol{S}$ for example.

The number of elements in a set $\boldsymbol{S}$ will be denoted $|\boldsymbol{S}|$.

Integer or real-valued variables will be represented by letters of the Roman alphabet, in lower case, regular weight and in italics like x for example.

Integer or real-valued constants will be represented by letters of the Roman alphabet, in upper case, regular weight and in italics, like C for example.

Constants or variables that are binary or only take discrete values will be represented by letters of the Greek alphabet, in lower case, regular weight and non-italicized, like λ for example.

1.3.2 Optimization and Mathematical Programming

The work presented in this thesis will make extensive use of optimization and, in particular, of mathematical programming. The use of mathematical programming is often criticized because of its limitation in terms of solving large scale problems and many researchers prefer the use of heuristics, which are considered more practical. Our aim in this work is to gain knowledge of what is fundamentally true about the ability of mesh-restorable networks to provide reliable service paths, and it is judged important to be able to present true optimal results that are not obscured by the effects of sub-optimality of results obtained using heuristic methods. We therefore choose to use mathematical programming to solve optimization problems, with sometimes the inconvenience of having to wait long hours before solutions are found. In addition, it should not always be assumed that mathematical programming could not be used to solve large problems. Today's optimization tools – CPLEX is the tool used in this work – when used with fast computers, can solve large problems historically considered as usually impossible to solve (NP-hard problems) within minutes.

A mathematical program can be generically represented as follows:

$$\text{Minimize (or Maximize): } f(x_1, x_2, \ldots, x_N) \tag{1.1}$$

Subject to:

$$g_i(x_1, x_2, \ldots, x_N) \le B_i,\ i = 1, \ldots, M \tag{1.2}$$

$$h_i(x_1, x_2, \ldots, x_N) = C_i, i = 1, \ldots, N \tag{1.3}$$

As shown in the previous equations, a mathematical program is composed of a set of *decision variables* $\{x_1, x_2, \ldots, x_N\}$ (the unknowns), an *objective function* (1.1), and sets of *constraints* (1.2) and (1.3). All functions f, g_i, h_i, and all parameters B_i and C_i are inputs to the problem. The aim of the optimization is to find the set of values of the decision variables for which all constraints are met and that minimizes or maximizes (depending on the type of problem) the objective function. If all functions f, g_i, and h_i are linear functions of the decision variables and the decision variables are allowed to take fractional values, the problem is referred to as *linear programming* (LP). If all functions are linear but the decision variables are constrained to take only integer values the problem will be referred to as *integer programming* (IP). If some variables are allowed to take fractional values and others have to be integer, the problem is referred to as *mixed integer programming* (MIP). Most optimization problems solved in this thesis are of type IP but are solved as MIP. Indeed, allowing some variables to take fractional values usually makes optimization problems solve much faster. Whenever we do this we will say that we are relaxing the integrality constraint on certain variables. We will sometimes still refer to the problem as an IP, even when some variables are relaxed.

The results of mathematical programming formulations presented in this thesis are generally presented accompanied by a certain value referred to as the MIPGAP. The MIPGAP indicates how close to the optimal solution the solution found is guaranteed to be. For example, a MIPGAP of 0.01 means that the value of the objective function for the solution found is guaranteed to be within 1% of the value of the objective function for the true optimal solution. Solving with a non-zero MIPGAP allows solutions to be found much faster than if a guaranteed optimal solution was sought. Usually, the MIPGAP will be set to 0.001 (0.1%) except when this makes the problems too difficult to solve.

1.3.3 Network Graph Terminology

A network is composed of *nodes* which are the points that we see on the graphical representation of a network (typically central offices) and of *spans* (typically cable ducts containing one or several fibres), which are the lines connecting the points on the graph. Each node can be characterized by its *nodal degree*, which corresponds to the number of spans connected to that node. For the network as a whole, we will talk about *network connectivity* to characterize the average number of spans connected to each node. Network connectivity will usually be measured by the average nodal degree.

In a capacitated network, each span is composed of a certain amount of quantized transmission capacity that we will measured in numbers of *capacity units* or *links* (very often called *channels* in the literature). Links will generally be divided into two categories: *working links* and *spare links*. Working links are used to support the transmission of signals in normal conditions. Spare links are used to support the transmission of some signals during network failures. The number of working links on a span *i* will be generally denoted w_i and the number of spare links will be denoted s_i.

A *route* is a combination of spans that are contiguous on the network graph. There is therefore no notion of capacity associated with the concept of route, it is a purely topological concept.

A *path* is what results from the cross-connection of several links in adjacent spans. A path is therefore a single-unit-capacity connection. Each path corresponds to a single route, but not vice-versa.

1.4 Research Methodology

1.4.1 Test Networks

In this thesis work, a constant effort was made to ensure that the concepts studied are investigated with a wide variety of network types. To achieve this goal, most experiments have been conducted with the large set of sixteen test networks briefly described in Table 1-1 and on Figures 1-1 to 1-3. These test networks represent various networks sizes and various levels of connectivity. Some of these networks have been obtained from journal or conference papers and others have been created manually by previous students of the Network Systems research group at TRLabs, or have been created specifically for this work. Each test network is associated with a given demand matrix that is used in every experiment. (The concept of demands is explained in detail in Section 2.2.2). Explanations of how demand matrices were generated are provided in Section 1.4.2.

The Bellcore and Bellcore Modified test networks are two networks with the same sets of nodes and the same demand matrices. The topology of the Bellcore Modified test case includes all the spans of the topology of the Bellcore test case with the addition of three spans chosen to remove the occurrence of degree-2 nodes (nodes connected to only two spans). As will shown later, degree-2 nodes are not desirable in networks with high reliability requirements. These two test networks will allow us to see some examples of the effects of removing degree-2 nodes on the reliability of networks.

Networks 16n29s1 and 16n38s1 are also two test cases with the same nodes but different numbers of spans. These two topologies will also give us some examples of how network connectivity

Table 1-1: Description of test networks

Name	Number of nodes	Number of spans	Average nodal degree	Number of degree-2 cuts	Total number of demand units	Ave. number of demand units per node pair	Source
06n14s1	6	14	4.67	0	115	7.67	
11n20s1	11	20	3.64	0	415	7.55	
11n20s2	11	20	3.64	0	448	8.15	
Bellcore	11	23	4.18	3	341	6.20	[YaH88]
Bellcore Mod.	11	26	4.72	0	341	6.20	
COST239	11	26	4.72	0	128	2.33	[Bat99]
12n20s1	12	20	3.33	3	212	3.21	
12n30s1	12	30	5	0	316	4.79	
15n28s1	15	28	3.73	2	465	4.43	[Bel93]
16n29s1	16	29	3.62	0	441	3.68	
16n38s1	16	38	4.75	0	441	3.68	
EuroNet	19	37	3.89	6	841	4.92	[LaA98]
22n41s1	22	41	3.72	2	676	2.93	
Net-A	20	40	4	0	1039	5.47	
Net-B	25	50	4	0	1615	5.38	
Net-C	30	60	4	0	2442	5.61	

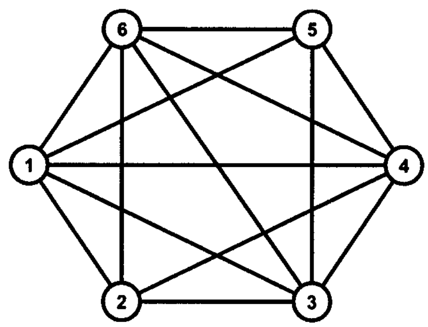

(a) 06n14s1

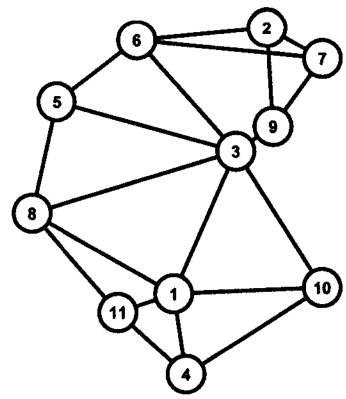

(b) 11n20s1

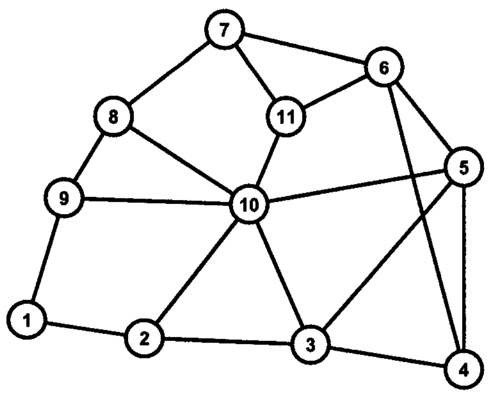

(c) 11n20s2

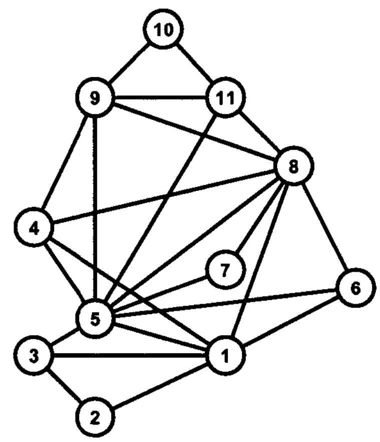

(d) Bellcore (11 nodes, 23 spans)

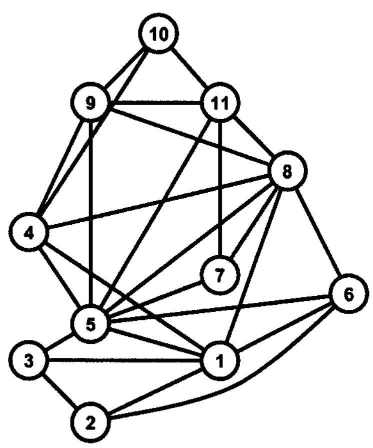

(d) Bellcore Modified (11 nodes, 26 spans)

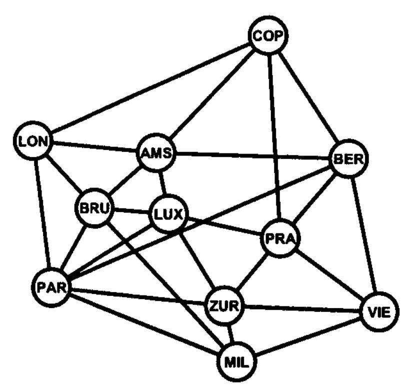

(f) COST239 (11 nodes, 26 spans)

Figure 1-1 Test networks

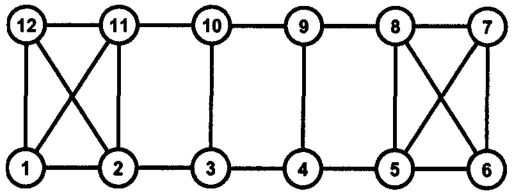

(a) 12n20s1

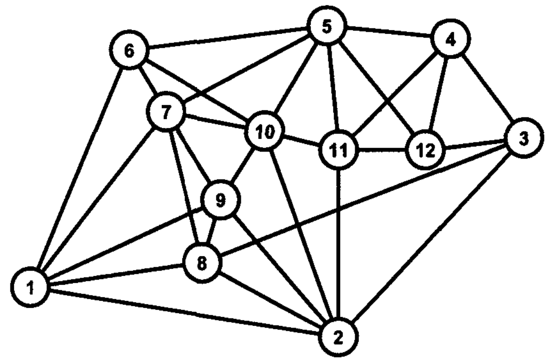

(b) 12n30s1

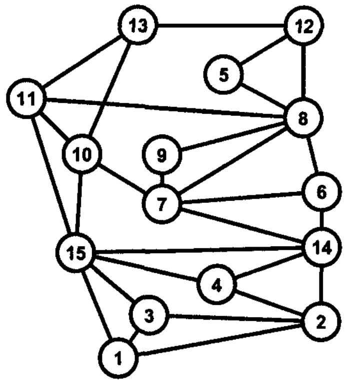

(c) 15n28s1

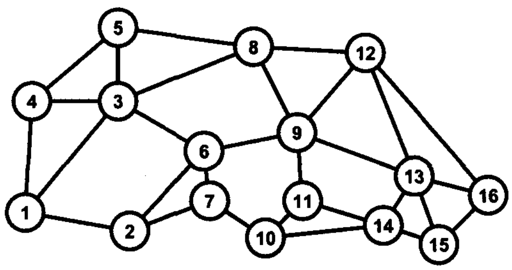

(d) 16n29s1

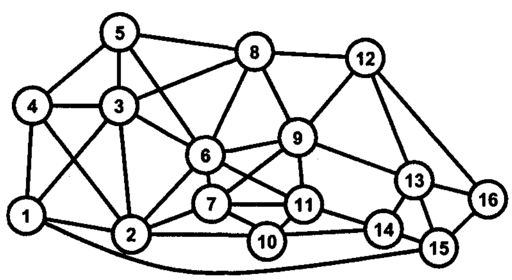

(e) 16n38s1

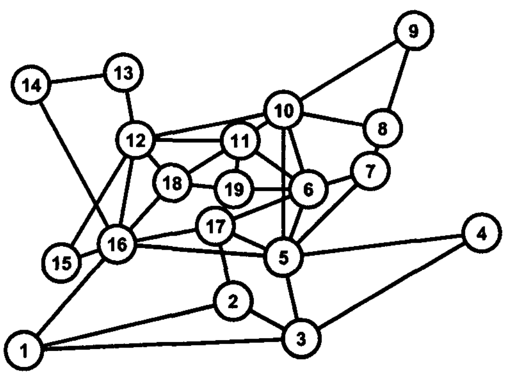

(f) EuroNet (19 nodes, 37 spans)

Figure 1-2 Test networks (continued)

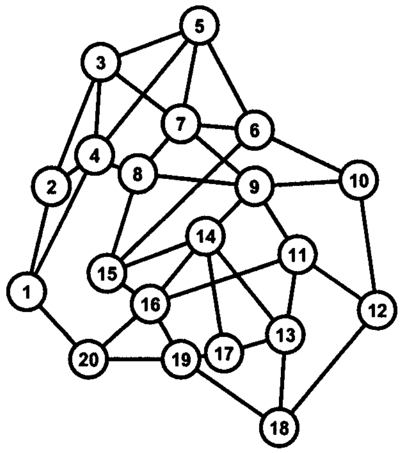

(a) Net-A (20 nodes, 40 spans)

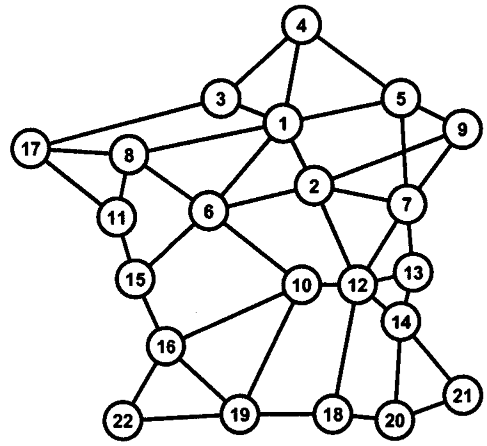

(b) 22n41s1

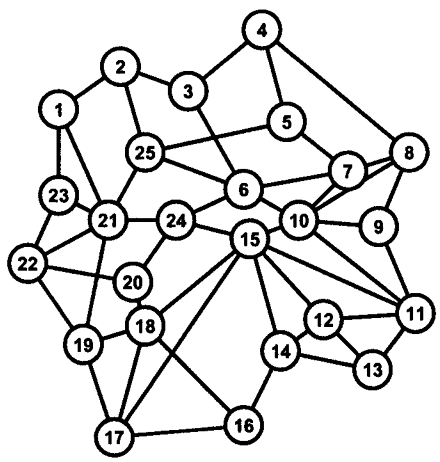

(c) Net-B (25 nodes, 50 spans)

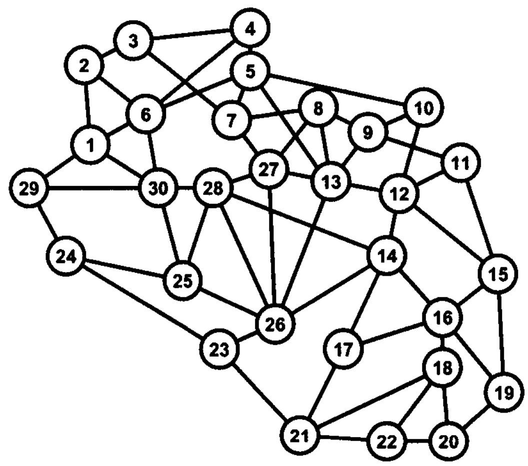

(c) Net-C (30 nodes, 60 spans)

Figure 1-3 Test networks (continued)

influences network reliability.

For a complete description of the test networks, see Appendix A.

1.4.2 Generating Test Demands

The test demands associated to the topologies presented in the previous section are generated using the gravity-based demand model of [DoG00], where the demand d(a, b) between a node pair (a,b) is calculated as:

$$d(a, b) = \left\lceil \frac{\text{nodal degree of } a \times \text{nodal degree of } b}{\text{distance from } a \text{ to } b} \times \text{constant} \right\rceil \quad (1.4)$$

This model is used since it is believed to fairly realistically represent the patterns of demand in the real world that tend to depend mostly on the importance of communication centres (represented by the product of nodal degrees) and also to a smaller extent to the distance between communication centres. Demand matrices of all test networks are detailed in Appendix A.

As communications become more and more dominated by Internet traffic, the distance between communication centres can be expected to become less and less of a factor. Future studies on this topic might therefore choose to use a model where only nodal degrees are taken into account. This would results in a higher number of demands between distant nodes and therefore in higher average path lengths, which as it will be observed later in the thesis, is expected to result in lower availability.

For test case COST 239, the demand matrix used is the same as the one used in the original publication [Bat99].

1.4.3 Experimental Equipment

Experimental results presented in this thesis have been obtained using various pieces of equipment. For optimization results, we used various generations of workstations, the most recent being a multiprocessor (4 × 900 MHz) UltraSparc III Sun Server (model v480) with 16 GB of RAM under Sun Solaris O/S 2.6. Optimization tools used are AMPL and CPLEX, the most recent version used being CPLEX 7.5. Each section presenting experimental results specifies what equipment was used for the experiments presented.

The programs of restorability and availability analysis used to produce the experimental results of Chapters 4 and 7, 8, and 9 were developed using the Borland C++ 5 environment on a PC equipped with 520 MB of RAM and an AMD processor running at 1.3 GHz. Experimental results were obtained using the same PC.

2. Background on Transport Networking

2.1 Types of Networks

There are fundamentally two types of telecommunication networks. Networks of the first type, *private networks*, are the ones that users see, like the computer network at your company for example. These are owned and used internally by private companies and can be classified into three categories: local area networks (LAN) that are contained within a building or a small area, metropolitan area networks (MAN) that cover a metropolitan area or a campus, and wide area networks (WAN) that can extend to wide areas up to thousands of kilometres [RaS02]. In the case of MANs and WANs, the company that owns the network does not always own all the land crossed by its network (especially in the case of WANs) and therefore needs to lease transmission capacity from networks of the second type, called *public networks*. Public networks are owned by companies referred to as telecommunications *carriers* or *service providers*. These networks (generally not seen from the public) offer much higher transmission capacity than private networks and nowadays use almost exclusively optical fibres as a transmission media. As shown in Figure 2-1, public networks can be geographically partitioned into three separate sub-networks: The *metropolitan access (or access) network*, the *metropolitan inter-office (or metro) network* and the *inter-exchange (or long-haul) network*. The access network connects the clients (residential or companies) to nearby *central offices* (COs), which are connected together by the metro network. The metro network usually contains one or more big hubs through which transits all the traffic that is going out of the metropolitan area into the long-haul network.

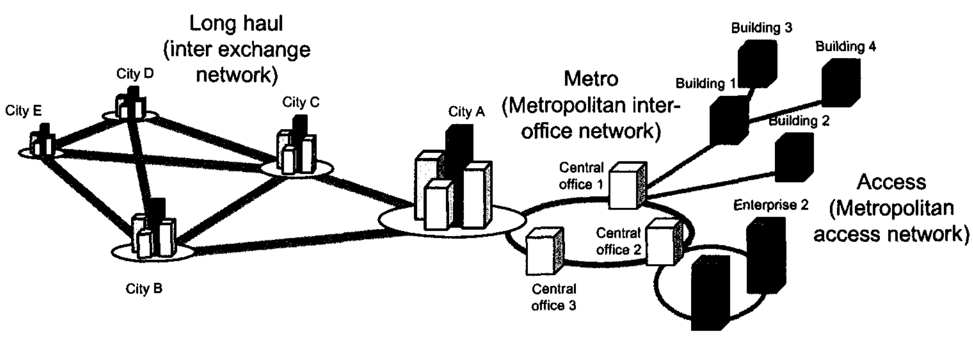

Figure 2-1 Geographically partitioned view of public networks

The access, metro and long-haul networks are usually topologically quite different. The access is usually composed of direct lines or tree-like topologies to residential customers and sometimes ring topologies to connect corporate customers. Most of the time, the metro network uses a ring topology to connect all central offices and hubs, but the trend could change and the mesh topology could become more and more common in the metro network (it is hoped that this thesis work will contribute to the understanding of why this is desirable). Long-haul networks usually have mesh topologies, in the sense that most nodes have more than two spans connected to them, but the average nodal degree varies a lot from one long-haul network to another. In Europe, where distances between major cities are not very big, long-haul networks tend to have a higher nodal degree and the topologies of these networks is very much "mesh-like." In North America, distances between important cities are much bigger and the average nodal degree is usually significantly lower, creating long chains of network spans that tend to make the networks look like collections of ring structures.

2.2 Concept of Transport Networking

The following paragraphs detail the different important aspects related to the concept of transport networking.

2.2.1 The Client Network / Transport Network Relationship

As explained in Section 2.1, public networks offer their services to private networks. In this service provider/client relation, we say that the public networks are *transport networks* and the private networks are the *client networks*. It is often said that the client network "sits above" the transport network.

Transport networks use multiplexing, switching and transmission facilities to provide end-to-end connection services to the client networks. We refer to these connection services as *service paths* or simply *connections*. Unlike in telephone networks, where many voice connections are established and taken down every minute, in transport networks connections are established for long periods of time and are therefore referred to as "semi-permanent" or "nailed-up." Also, it is usually not possible to re-arrange existing connections in order to re-distribute traffic among the network.

A given network can include several layers of client/transport relationships where the transport network at a given level is itself the client network of another lower-level transport network. Each transport network can be divided into two layers: the physical layer and the logical layer. Figure 2-2 (a) shows the physical layer view of a transport network with the description of a few end-to-end service paths. In this figure, the physical transmission facilities are represented by the thick grey lines connecting the nodes and the service paths are represented by the thin black lines connecting nodes that are not necessarily adjacent in the physical graph. Figure 2-2 (b) shows the corresponding logical layer view of the transport network, where the connections of Figure 2-2 (a) are repre-

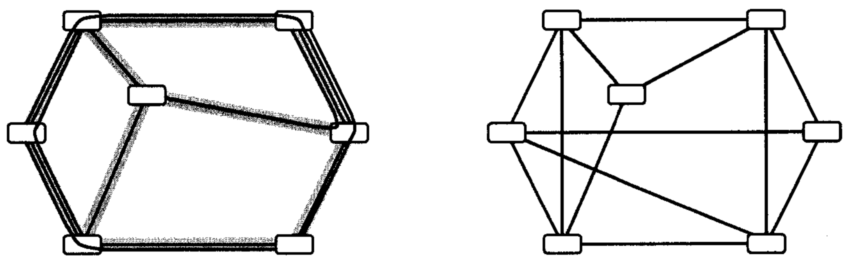

(a) (b)

Figure 2-2 Physical layer and logical layer

sented by straight lines between their end-nodes and the physical transmissions facilities are not shown. In the case of a network that includes several layers of client/transport relationships, the logical view of a transport network is what is considered as the physical network by another transport network that would be right above it. The term "physical" therefore only corresponds to the reality in the lowest client/transport relationship in the stack.

Figure 2-3 shows examples of different layer stacking models with up to four client/transport relationships. The sometimes high number of layers is due to the various types of services that are served by public networks and the fact that each type of service has a corresponding technology that best fits its needs. For example, in the five-layer IP-over-ATM-over-SONET-over-WDM-over-Fibre, the IP layer may serve datagram (packet) delay-insensitive applications connected through the Internet, the ATM layer may be used to serve the IP layer as well as delay-sensitive multimedia applications, the SONET layer may be used to serve the ATM layer as well as some companies' private WANs, and the WDM layer may serve the SONET layer as well as some very high capacity connections to some service providers. More on all these technologies follows in Section 2.3.

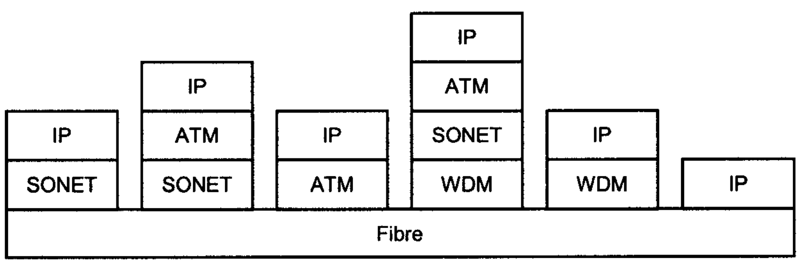

Figure 2-3 Various layer stacking models

2.2.2 *Network Demands*

In a client/transport relationship, the client network presents *demands* to the transport network.

A demand is a request for a certain amount of end-to-end transmission capacity between a pair of nodes of the network graph. Besides its end-nodes, a demand is characterized by its size or type of transmission unit, which usually belongs to set of standard transmission rates for the given type of transport network considered. Depending on the type of transport network considered, these unit designators can be DS1s, STS-1s, or even lightpaths (these will be defined in Section 2.3).

Demands presented to the physical layer of the transport network very often result from the aggregation of the traffic of different traffic sources from one or several client layers. These traffic sources can be of very different nature, like data, video conferences, and plain old telephone service, although as it will be seen later in the thesis, it is preferable to keep certain services with very different reliability requirements separate. These sources are combined together using various multiplexing techniques (more on that in the following section). The advantages of multiplexing several lower-rate traffic sources into one bigger connection is a better utilization of the capacity and a reduction of the amount of processing that is required at each node along the connection's route, since that connection is usually not decomposed until it reaches its destination. The act of combining lower-rate traffic sources and multiplexing them into bigger standard size connections is referred to as *traffic grooming*. Traffic grooming has become a topic of very much interest with the development of optical transport networks (OTN), in which demands correspond to lightpaths, which very few traffic sources can fully utilize. For studies of traffic grooming, see [BaP01], [MoL01], [ThS02], [ZZM02], and [ZhM03].

Besides point-to-point demands (demands between a pair of nodes, as defined above), point-to-multipoint (multicast) connections can also be requested for certain applications (e.g. broadcasting video over the Internet). This thesis work, however, will only consider point-to-point, symmetrical and bidirectional (the transmission capacity is required in both directions) demands.

2.2.3 Different Types of Transport Networks

Transport networks can be divided into two classes: *circuit-switched* transport networks, shown on Figure 2-4, and *packet-switched* transport networks, shown on Figure 2-5.

Circuit-switched networks use *fixed multiplexing*, which guarantees a fixed dedicated bandwidth to each connection. Circuit-switched transport networks were historically used to carry the telephone traffic of the *Public Switched Telephone Network* (PSTN), but are now used to carry both voice and data. The drawback of fixed multiplexing is that when a connection is not used, its capacity is not available to increase the capacity of the other connections in use. This can result in poor capacity utilization, especially when traffic is increasingly composed of data services

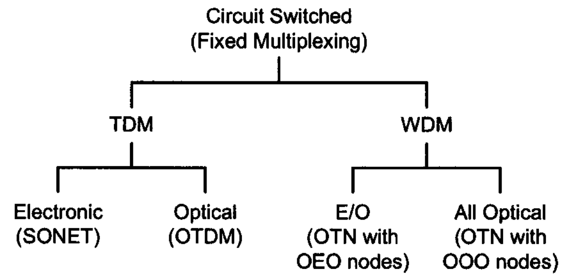

Figure 2-4 Types of circuit-switched transport networks

Packet-switched networks use *statistical multiplexing*, which takes advantage of the bursty nature of data traffic to achieve much higher utilization of the capacity. In order to achieve that, packet-switched networks decompose traffic streams into *packets* (or *cells*), which are queued at each node and sent as soon as the transmission medium is available. The trade-off for this better capacity utilization are packet delays (packets have to wait in the nodes, instead of being sent directly as in a circuit-switched network) and a non-zero *cell loss probability* (CLP) or probability of losing packets (when too many packets arrive at a node, packet buffer overrun may happen). The classic problem of packet-switched networks is that they are not well suited for applications that require guarantees in terms of delay and CLP. For this reason, new packet transport network schemes (ATM and MPLS, presented in Section 2.3) have been developed to provide circuit-like connections in a packet environment. Packet-switched networks can therefore be subdivided into *connectionless* and *connection-oriented* packet networks.

Another characteristic of each transport network is the domain in which the multiplexing is performed. Two domains are used for multiplexing in transport networking: time and frequency. In Time Division Multiplexing (TDM) signals are sent at different times and occupy the whole chan-

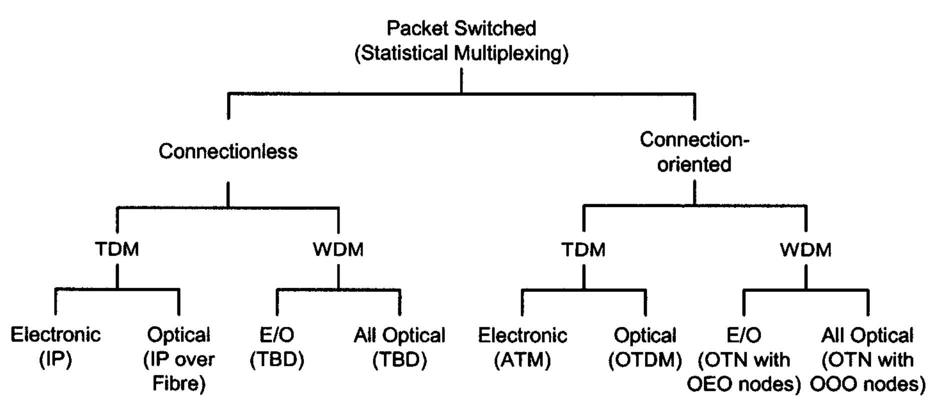

Figure 2-5 Types of packet-switched transport networks

nel. In Frequency Division Multiplexing (FDM) the channel is divided into frequency bands and signals are sent at the same time but each within the band of frequency that was allocated to it. When FDM is used with optical fibres as the transmission medium, it is usually referred to as Wavelength Division Multiplexing (WDM) instead of FDM.

With the advent of optical transport networks, an additional criterion to classify transport networks is whether the transport network functions are performed in the electrical or optical domain.

Figures 2-4 and 2-5 show how circuit-switched and packet-switched transport networks can be classified following the above criteria, and give example of transport network technologies corresponding to each category. Many of these technologies are presented in Section 2.3.

2.2.4 Relationship to the OSI Layer Model

The computer network research community often refers to the widely cited seven-layer Open Systems Interconnect (OSI) reference model developed by the International Organization for Standardization (ISO). Figure 2-6 draws the parallel between that layer model and the service layer/transport layer paradigm presented in the previous sections.

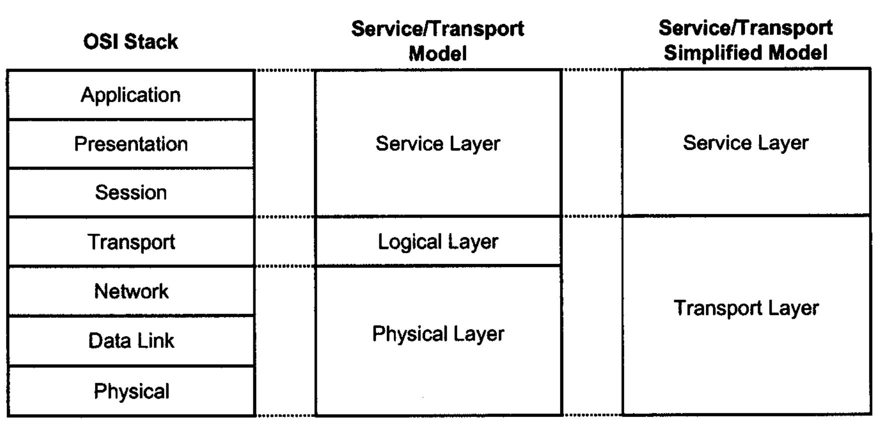

Figure 2-6 Network layer models

The transport network, as defined in the previous sections, does in fact include all four bottom layers of the OSI stack. In the case of a network that includes several levels of client/transport relationships, there are in fact several network and data link layers as shown in Figure 2-7.

2.2.5 Control Plane

Transport networks usually provide more than just end-to-end connected transmission capacity. A common functionality provided by transport networks is the monitoring of path integrity (whether service paths meet some minimum performance requirements in terms of bit error rate,

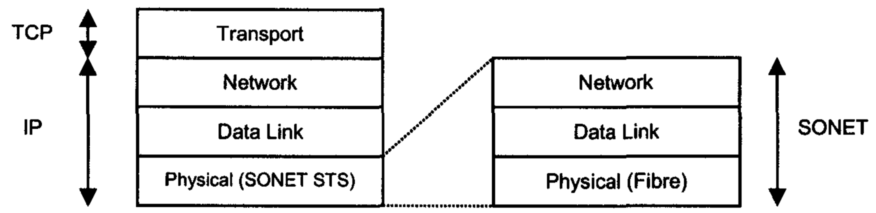

Figure 2-7 OSI stack with multiple layers of transport

signal level, etc.) When a path suffers a failure, because of a fibre cut or failed fibre amplifier for example, the network may trigger some actions so that service will be restored for that service path. This is referred to as network *survivability* and will be a central topic of this thesis (more on network survivability follows in Section 2.4). Such Operations, Administration, Management and Protection (OAM&P) mechanisms require was is called a *control plane*, which can be viewed as a separate network of supervisory channels that carry information about the network. These supervisory channels are created by adding information to the transmitted signals. That extra information is called *overhead* and the useful information (coming from the client layer) in a transmitted signal is referred to as the *payload*. In a *centralized system*, OAM&P actions are triggered as a result of a centralized decision, taken by a central network management system that collects information from all the network elements through the control plane. In a *distributed system*, OAM&P actions are triggered as a result of local decisions taken by elements of the systems based on the observation of the supervisory channels on the neighbour spans and based on simple rules. The advantage of a centralized system is that better decisions can be made since they rely on the knowledge of the state of the entire network but this is obtained at the cost of a much slower decision process and a much bigger control traffic load required to communicate the network state to the central management system. This approach is also prone to errors in case the central command point does not have an up-to-date knowledge of the complete network state. The advantage of a distributed system is the much faster decision process and the reliability of these mechanisms, based only on local knowledge of the network.

Besides network survivability, another functionality that the control plane may include is the provisioning process of new service paths. This process requires the search for available capacity, the reservation of that capacity and the establishment of the path (cross-connection of the reserved capacity at each node). This process is either automatic (as dialing in PSTN) or requires some degree of human intervention.

Other roles that the control plane is expected to play in the future are presented in Section 2.6.

2.3 Generations of Transport Network Technology

The technologies used for transport networking are evolving over time in response to changes in the types of services that need to be transported over these networks. The most important change is the increase in the demand for bandwidth. From voice circuits at an equivalent bit rate of 64 kb/s the requested connections are now sometimes as high as several Gb/s. The evolution from electrical transmission facilities to optical transmission was the answer to that trend. Another aspect that is changed about services offered by transport networks is the need to improve the availability of service. That factor resulted in the development of techniques collectively referred to as *network survivability*, especially in the context of SONET/SDH networks [Fla90][Wu95]. Section 2.4 will be specifically devoted to that topic. Also, transport networks now have to be able to sell connections for much shorter periods of time and to set-up and tear down connections quickly. These new requirements are also pushing transport networking technology to evolve.

The following sub-sections describe important transport network technologies from oldest to the most recent.

2.3.1 Public Switched Telephone Network

The first transport networks were deployed in the past century to support the *Public Switched Telephone Network* (PSTN). The core of the PSTN was composed of electromechanical switches connected together by point-to-point microwave transmission channels, coaxial cables, or even ordinary cable pairs. The PSTN was of course a circuit-switched network in which voice channels were multiplexed using FDM and all signals were transported in analog form. Although it originally carried only voice signals, the PSTN quickly started being used to support data communications as well. With the increase of the volume of voice traffic and especially of data communications, new transport technologies were developed (see Section 2.3.2) and the PSTN became a client network of bigger transport networks capable of efficiently carrying multiple types of services (voice, data, video,...). At the beginning of the twenty first century, the amount of data traffic has for the first time exceeded voice traffic and the proportion of voice traffic in the total traffic is since then quickly decreasing. The situation is now reversed and the objective is now to support voice communications on networks technologies originally developed for data. More on this topic will be presented in Section 2.3.3.

2.3.2 SONET and SDH

Synchronous optical networking (SONET) is a standard for transport networking on optical fibres that was developed in the 1980s and which is until today the most popular technology for transport networking. SONET is the standard used in North America and is compatible with the international *synchronous digital hierarchy* (SDH) standard. SONET/SDH are sometimes referred to as *synchronous transfer mode* (STM) as opposed to *asynchronous transfer mode* (ATM) presented in the following section. The goals in developing the SONET/SDH standards were to make multiplexing of independent streams easy (by using a synchronous technique), to add OAM&P capabilities to optical networks, to guarantee the compatibility of optical networking equipment from different vendors, and to improve network availability (availability will be formally defined in Chapter 3). The rest of this section presents some details of the SONET standard. The reader is referred to [ITU93] for corresponding details for SDH.

To achieve the goal of simple multiplexing, SONET defines sets of standard *synchronous transport signal* (STS) levels. The basic transport signal is the synchronous transport signal 1, STS-1. The STS-1 frame is illustrated on Figure 2-8. It is divided into overhead, which contains all the signalling necessary to perform the channel supervisory functions, path monitoring, etc., and the payload referred to as *synchronous payload envelope* (SPE).

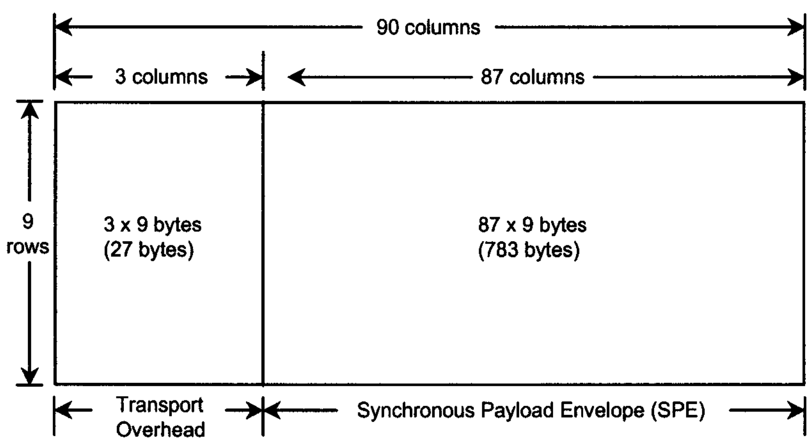

Figure 2-8 SONET STS-1 frame

Figure 2-9 shows the general organization of the SONET hierarchy. An STS-1 signal can carry one DS-3 signal (standard rate of 44.736 Mb/s inherited from a past transmission standard) or several lower rate standard signals. To be included into the STS-1 SPE, these lower rate signals have to first be mapped into standard *virtual tributary* (VT) signals. The STS-1 SPE can contain various combinations of VT signals that are combined into *VT groups* before being fitted into an STS-1 SPE. For services that require rates higher than STS-1, several STS-1 signals can be concatenated to form a higher rate signal. Concatenating M STS-1 signals will thus create an STS-Mc signal. This concatenation is made particularly easy by the fact that STS-M signal rates in the SONET

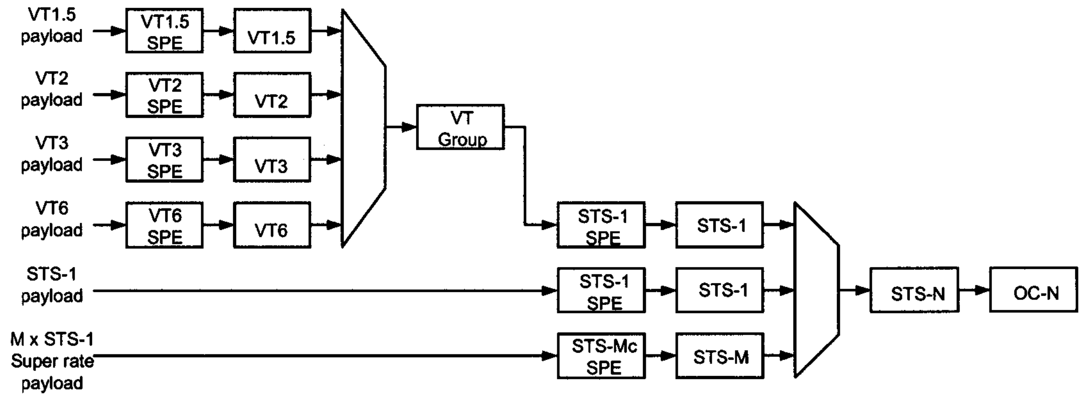

Figure 2-9 SONET protocol structure

standard are exact multiple of the STS-1 basic rate of 51.84 Mb/s. This simplicity of the multiplexing has the added advantage of reducing the cost of SONET equipment. Finally, STS-1 signals and STS-Mc signals are multiplexed to form an STS-N electrical signal that is converted to an *optical carrier signal* OC-N belonging to a set of standard optical transmission rates (OC-1, OC-3, OC-12, OC-24, OC-48, or OC-192 in the original SONET standard).

The SONET transport functions can be represented in the layered view of Figure 2-10, in which each layer corresponds to a given level of signals. Also, to each layer corresponds a set of functions related to the type of signals in that layer. The *VT path layer* is responsible for mapping VT payloads into the VT SPEs and adding overhead to the VT paths to monitor their path integrity. The *STS path layer* is responsible for mapping STS payloads into the STS-1 SPEs and STS-Mc SPEs. Overhead is also added to STS signals to monitor their integrity as well. The *line layer* is responsible for multiplexing the STS-1 and STS-Mc signals and create the line level STS-N signals.

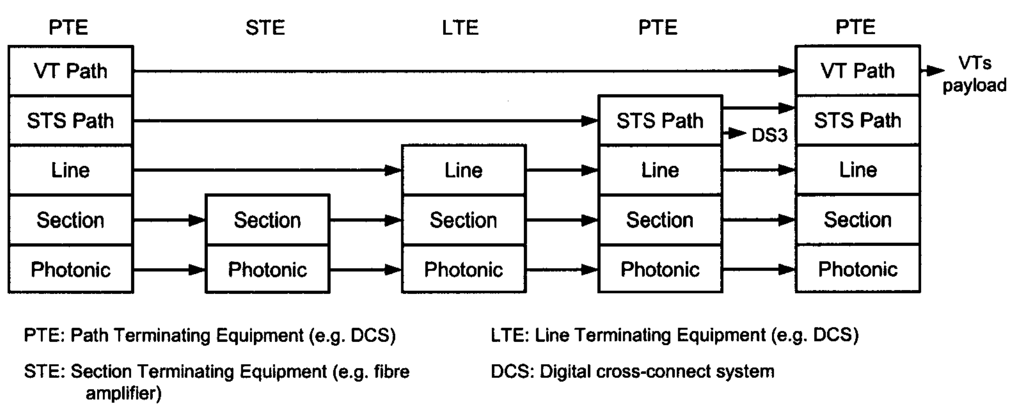

Figure 2-10 Layered view of SONET with section, line and path equipment (adapted from [ALL96])

The overhead added by the line layer is used to monitor the error performance of the STS-N signals and trigger protection actions (see Section 2.4 for more details) when signal quality falls below a certain threshold. The *section layer* operates on line level STS-N signals as well and is responsible for performing the basic functions required before SONET signals can be extracted. These functions include framing, which in simple terms is the process of locating the position of the information within a bit stream, and scrambling, which pseudo-randomizes the outgoing bit-streams in order to avoid long streams of 0's or 1's (a minimum number of 0-1 or 1-0 transitions are required for signal synchronization purposes). Finally, the *photonic layer* is responsible for the conversion of the STS-N signals to OC-N optical carrier signals and for guaranteeing the performance of the optical signal in terms of signal power level, pulse shape, timing jitter, etc.

At each signal level corresponds a type of terminating equipment. A *path terminating equipment* (PTE) terminates path level signals. These are either STS PTEs (terminating STS signals) or VT PTEs (terminating VT signals) or both. A *line terminating equipment* (LTE) terminates line level (STS-N) signals. A *section terminating equipment* (STE) terminates section level signals. Note that a piece of equipment is associated to the highest layer in which it is able to perform some function and that it also has functionalities associated to the layers below that. A "photonic terminating equipment" does not exist since all network elements have at least photonic and section layer functionalities.

An example of PTE is the *digital cross-connect system* (DCS) illustrated in Figure 2-11. DCSs

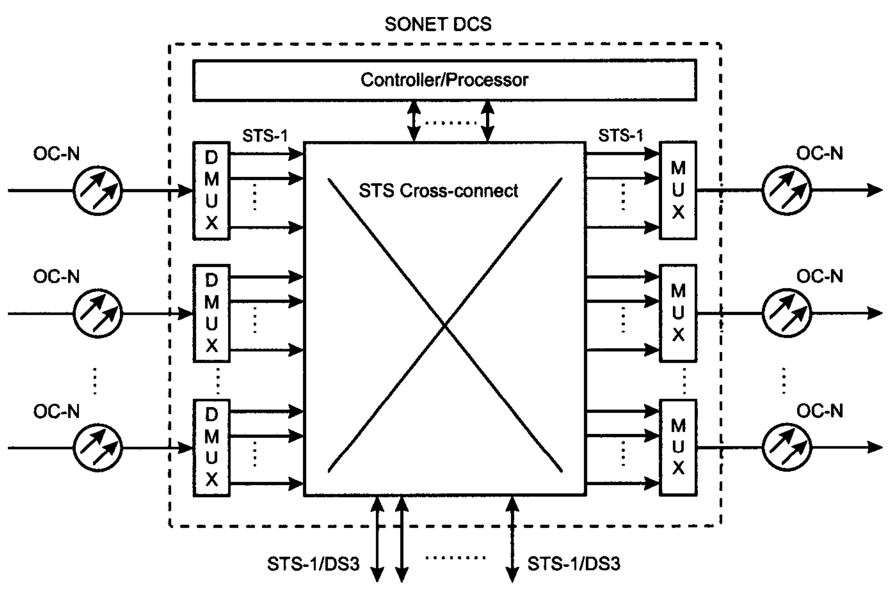

Figure 2-11 Functional block diagram of an STS DCS (from [Gro03])

are structures that have several optical line interfaces and can switch individual path layer signals between the different optical lines. In addition, DCSs can usually add/drop and multiplex/ demultiplex STS paths from the line signals therefore allowing demand grooming in these nodes. STS DCSs, also called *broadband DCS* (B-DCS), can switch signals at the STS level (STS-1 or STS-Mc). This is the type described in Figure 2-11. VT DCSs, also called *wideband DCS* (W-DCS) can switch VT path signals.

A common function of DCSs is the provisioning of new service paths. Indeed, the line layer overhead contains network management information and most DCSs are able to communicate with a centralized network management system that gives instructions about connections that need to be made and torn down. DCSs are also a key network element in the development of mesh-restorable networks that will be presented in Section 2.4.

Another important network element is the *add/drop multiplexer* (ADM), which, to simplify things, can be viewed as a special type of DCS with only two optical line interfaces. ADMs are a key element of ring-based networks (presented in Section 2.4).

2.3.3 Asynchronous Transfer Mode (ATM)

ATM was developed in an effort to integrate voice and data services into a common packet-switched network. As mentioned earlier, the major problem with serving voice applications in packet-switched networks is that these cannot easily provide QoS guarantees since they are based on statistical multiplexing of packets. Unlike traditional packet transport protocols, which are connectionless, ATM establishes connections between pairs of points, referred to as *virtual circuits* (VCs). The term "virtual" refers to the fact that unlike in the case of real circuit-switched networks, ATM uses statistical multiplexing and therefore the bandwidth of these circuits is not dedicated to them but shared with the other connections. However, ATM is able to provide some QoS guarantees to these connections. These guarantees are expressed in terms of cell loss probability (CLP), delay and jitter.

This ability of ATM to provide QoS guarantees is due to several reasons. First, rather than routing packets in the network, ATM is said to "switch packets." Indeed, packets that belong to a particular VC are identified by a *virtual circuit identifier* (VCI) in the packet header. The VCIs for a given VC can be different on each link traversed but they are unique to that VC on each of these links. Each ATM switch along the way has a forwarding table that tells it on which port packets from each incoming VC should be forwarded and what the new VCI for that packet has to be. These forwarding tables are updated with the relevant information for each VC at the time of connection establishment. Compared to pure packet routers, ATM switches have less processing to do and are therefore faster. In order to reduce the size of forwarding tables and further reduce the amount of processing performed by ATM switches, VCs that follow the same route (from end-to-end or at least partially) can be grouped together inside *virtual paths* (VPs) as shown on Figure 2-12. The result is that it is much faster for ATM switches to search their forwarding tables and it also reduces the cost of making fast switches.

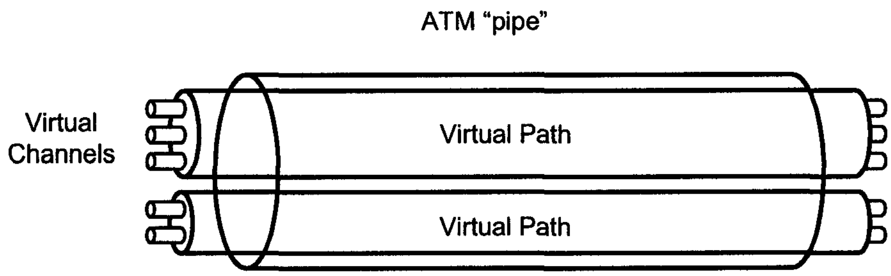

Figure 2-12 ATM VCs and VPs

A second reason why ATM can provide QoS guarantees is that the network monitors the traffic generated by each connection and makes sure that it does not exceed limits set in a contract agreed upon at the time of the connection establishment. These limits are expressed in terms of peak cell rate, average cell rate and burst size. If a connection does not respect the contract, it can either be dropped systematically or dropped in case of traffic congestion.

Another measure taken by ATM to allow the control of QoS is *call admission control.* This measure ensures that new connections are accepted only if the QoS guarantees of that new connection and of all existing connections will still be feasible after the new connection is established. This decision is based on very complex mathematical models.

QoS guarantees are also ensured through complex queuing techniques in the ATM switches and through connection *flow control*, a function that allows users of delay-insensitive applications to be notified to wait before transmitting when the network is congested, thus reducing network load.

2.3.4 Optical Transport Networks

Originally used purely as point-to-point transmission systems for SONET/SDH-based transport networks, optical networks are evolving to include more and more transport functions. The initial motivations for using optical fibres were mainly the high bandwidth they offer, the low signal attenuation, and the fact that they do not create electromagnetic interference as do electrical transmission media. Optical fibres have thus been deployed in almost all parts of the public networks: first in the long-haul and for undersea transmissions and then also in the metro (inter-office) networks. Only the access part still mostly uses copper cables (with the exception of access optical rings used for some corporate customers).

Over the time, the ever growing demand for bandwidth has required increases in the amount of data that can be transmitted though a single optical fibre. Improving the transmission technology was indeed much cheaper than laying down new fibre, which is usually very expensive. The first solution is to increase the bit rate that can be transmitted through a wavelength. The current state of the art transmission rate over a single wavelength is around 40 Gb/s and seems to be a threshold level above which electronics becomes fairly expensive to produce [CMY02]. Beyond that point, it becomes difficult and very expensive to build electronics capable of processing data fast enough. The second option to increase the data throughput on an optical fibre is to increase the number of wavelengths used for transmission. This is made possible by *wavelength division multiplexing* (WDM) and more recently *dense wavelength division multiplexing* (DWDM), which could in the long term allow the transmission of more than a thousand wavelengths on a single fibre. Knowing that each additional wavelength transmitted can effectively add 40 Gb/s of transmission capacity, one realizes the phenomenal transmission capacity that can be obtained with a single fibre.

With optical networks of the first generation, transport signals are converted back to the electrical domain at each node in order to perform switching, path monitoring, multiplexing, etc. At very high data rates, the amount of data that needs to be processed by the node becomes very high. This was the motivation for the development of new optical equipment capable of performing transport network functions in order to keep transport signals as much as possible in the optical layer and therefore reduce the amount of electronic processing required in the nodes. These advances are mostly linked to the development of the *optical cross-connect* (OXC) and the *optical add/drop multiplexer* (OADM), which are functionally very similar to their electrical layer counterparts. The result is a reduction of the amount of super fast electronics in the nodes since only the end-nodes of optical connections have to deal with the data. The first *optical transport networks* (OTN), which appeared at the end of the 1990s, offered *lightpath services* to their client layers and are referred to as *wavelength routed networks*. Lightpaths are continuous optical channels realized by cross-connecting wavelengths in the optical domain. A typical client layer of OTNs is SONET/SDH, which sees no difference between a lightpath provided by an OTN and a direct wavelength transmitted between two physically adjacent nodes connected by a fibre. Lightpaths in these networks are generally established permanently at the time the network is deployed. They are either *wavelength* *continuous* if the wavelength is the same from end-to-end, in which case the lightpath is referred to as a *wavelength path* (WP), or there can be *wavelength conversion* at some point in the path, in which case the lightpath is referred to as *virtual wavelength path* (VWP). When the lightpaths in a network are never converted back to the electrical domain (when using all optical, or *O/O/O* nodes), the network is said to be *fully transparent*. In the opposite case, when some optical-to-electrical-to-optical (O/E/O) conversion the network is said to be *opaque*. Full transparency is a highly desirable property because it means that the wavelength paths offered are compatible with any type of service and any protocol. Fully transparent networks also allow easy transport of analog signals, which is more difficult in networks where transport signals are converted to the electrical domain for regeneration for example.

To go back to the transport network classification of Section 2.2.3, wavelength-routed networks would be classified as circuit-switched networks using WDM and are either all optical or mixed electrical/optical depending on the type of nodal equipment. In fact, these networks are just the equivalent of SONET one layer below.

More recently, technology has started being developed for packet-switched OTNs [ElS02]. Similarly to what was done with ATM, the research on packet-switched OTNs is developing ways to offer VC services in the optical layer. These are made possible by technologies like *multi-protocol label switching* (MPLS) and *generalized multi-protocol label switching* (GMPLS), which use many of the same principles as ATM. Offering VC services in the optical layer is very useful because very few applications require the transmission capacity represented by a whole wavelength and presumably there would be an important demand for sub-lambda services. The technique of choice to obtain sub-lambda services is to multiplex several connections within a wavelength using TDM or optical TDM (OTDM). Techniques to support datagram services in the optical layer are still to be defined.

As the optical layer is evolving to include more and more of the functions of transport networks (multiplexing, switching, etc.), it appears that it is going to be possible to reduce the number of transport layers stacked one above the other (as shown in Figure 2-3). This is indeed one of the goals of future transport networking, motivated among other reasons by the difficulty of provisioning paths through multiple layers of client/transport relationships. However, multiple layer configurations are still likely to be used for some time because different layers are efficient at performing functions at different bit rates.

Other services that future OTNs are expected to offer are presented in Section 2.6.

2.4 Network Survivability and Survivability Architectures

2.4.1 Common Measures to Improve Service Assurance

One first obvious strategy to support high service assurance in transport networks is to simply reduce the frequency of failures occurring in the network. This requires the use of highly reliable network elements. This high reliability is obtained using high-reliability electronic components and redundant internal architectures (e.g. redundant power source, redundant wiring, etc.) Ensuring the reliability of network elements also requires extensive testing of the equipment in a simulated network environment before it is installed in order to identify possible failure situations that could be caused by the new equipment. This testing is particularly important for questions of software compatibility between the new equipment and equipment already installed. Indeed, the software of a new piece of equipment may not be compatible with other equipment's software or, worse, it could prevent other equipment's software from functioning properly.

It is also important to monitor the degradation of installed equipment and replace aging equipment before it fails. This is particularly important for optical fibres.

Other important measures to reduce the occurrence of failures include protecting central offices against building fires and floods, protecting all network infrastructures from unintentional accidents and vandalism, and informing construction companies of the location of optical fibres.

Since some network failures inevitably happen, no matter how hard we try to prevent them, another important strategy is to engineer systems in which the consequences of failures are contained as much as possible within the failed part and affect as little of the whole network as possible. Failure propagation is a major concern especially regarding software, which plays a major role in helping to bring the network back to working state in situations of failures. Making network management software redundant is therefore an important factor in improving service assurance [Ogg01]. Related to this idea of avoiding the multiplication of failures as a consequence of a single failure is the considerable importance given to the physical diversity of the transmission facilities. Ensuring the diversity of the routing of transmission cables has always been considered of vital importance in the improvement of service assurance and we will see later on that this aspect has now become a major factor in the way transport networks are designed.

The other major strategy in reducing the impact of failures is *network survivability* and it is the central topic of this thesis. Network survivability is the ability of a network to provide service replacement solutions in the event of network failures so that service may fully or partially continue for some or all of the clients that would otherwise lose service. Ideally, clients would not even notice any change, survivability happening all internally in a split-second within the transport layer. Figure 2-13 illustrates this idea. In Figure 2-13 (a), a failure occurs in the physical layer of the network that was shown on Figure 2-2 (a). In this case, all the connections that were routed through the span that is now failed (previously indicated by thin black lines) have been replaced by restoration paths indicated by the dotted black lines. All the other connections not affected by the physical failure have not changed. Notice that the logical layer shown in Figure 2-13 (b) remains identical to that shown on Figure 2-2 (b). The physical layer failure is therefore invisible to the over-lying layers.

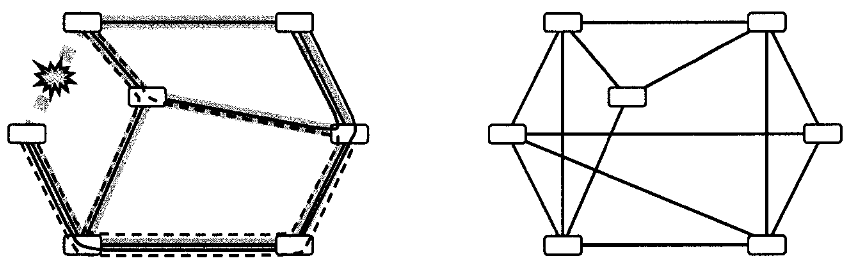

(a) (b)

Figure 2-13 Failures in the transport layer are hidden from the client layers

Depending on the type of service offered, however, survivability may mean a small interruption of service of a few seconds, referred to as *restoration time*, during the time to set up the replacement solutions, and then possibly a reduction in service level like reduced bandwidth and higher packet delay (in the case of packet-switched networks). In the case of circuit-switched networks, for a given service path, survivability will provide either complete replacement of the pre-failure bandwidth, or no replacement bandwidth will be provided at all and the service path will be considered down. The following sections present the general types of survivability schemes that can be implemented in transport networks and present a few important schemes, including "span restoration," which is the scheme that is the main focus of this thesis.

2.4.2 Classification of Survivability Schemes

As illustrated on Figure 2-14, survivability schemes can generally be classified as *pre-planned* or *protection* mechanisms, which take a proactive approach to network survivability, and *adaptive restoration* or simply *restoration* mechanisms, which take a reactive approach to network survivability. With protection mechanisms, the backup paths used to support replacement flows for the affected demands are known prior to failures. These backup paths may be dedicated to the primary paths they protect, in which case they are usually also pre-connected, or they may need to be con-

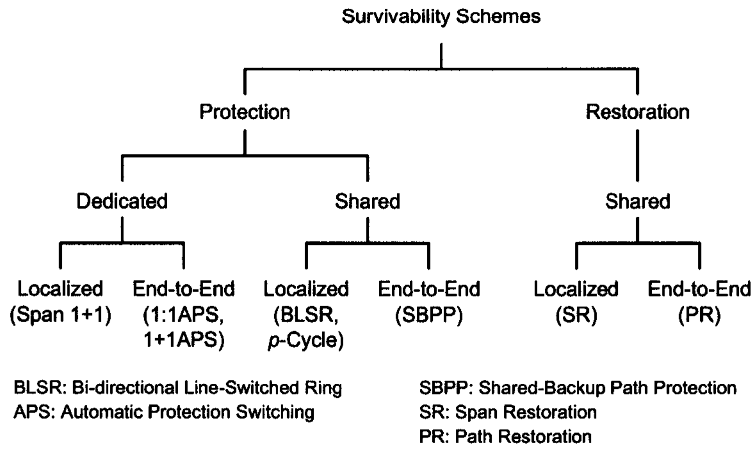

Figure 2-14 Classes of survivability schemes

nected at the time of failure in case the spare capacity used to construct them is shared by other demands or if it is used for low priority preemptible services. When a failure occurs, activation of the backup paths is fast, especially when backup paths are pre-connected and only a switching action at both end of the backup path is required, and therefore restoration time is small. With restoration mechanisms, the backup paths are searched within a layer of spare capacity shared by the whole network and constructed in a failure-specific adaptive response to the network state after a failure occurs, configuring the available spare capacity as needed at the time of failure. Because restoration paths need to be searched after a failure occurs, restoration mechanisms are usually considered as offering longer restoration times than protection mechanisms. This point, however, while it is unarguably true when comparing restoration mechanisms to pre-connected protection mechanisms, can be debated when it comes to comparing restoration mechanisms to non-pre-connected protection mechanism (all protection mechanisms allowing spare capacity sharing are in this category). Indeed, the restoration path search is not necessarily more time consuming than the path connection process (especially in the case of distributed restoration path search) and restoration mechanisms can therefore offer restoration times in the same order of magnitude as most protection mechanisms using capacity sharing [Gro97]. Restoration mechanisms also benefit from a greater flexibility that allows them to find restoration paths even when the spare capacity that would ordinarily have been used is not available due to other failures or maintenance actions or other miscellaneous events. In contrast, protection mechanisms normally only provide one possible backup path option for every failed demand in a given failure scenario.

Survivability schemes can also be further divided into *localized* and *end-to-end* response. Localized response means that the failed demands are only rerouted locally around the failure, the end-to-end after-failure path being identical to the end-to-end pre-failure path except in the vicinity of the failure. In end-to-end response, the replacement paths are found between the end-nodes of the affected demands and therefore they can be completely different from the pre-failure paths. Note that nothing prevents end-to-end survivability schemes from providing replacement paths that are identical to the ones provided by localized survivability schemes. End-to-end survivability schemes simply have more options and they therefore usually provide more optimal solutions in terms of minimizing the capacity requirements or maximizing the amount of restorable demands in failure situations. Another sub-division of end-to-end protection scheme, not shown on Figure 2-14, is between *failure-independent* and *failure-dependent* protection schemes. With failure-independent schemes, only one backup path is planned for each connection and it will be the one used no matter what failure affects the path. With failure-dependent schemes, several backup paths are planned for each connection and the backup path used in case of failure depends on where the failure has occurred on the primary path. The drawback of failure-independent schemes is that the primary and backup paths have to be completely disjoint so that a single failure may not affect both paths at the same time. Pairs of completely disjoint paths may be difficult to find in some sparsely connected graphs and this constraint may be costly in terms of capacity requirement. Failure-dependent schemes have therefore the advantage of not requiring disjoint paths to be found, however they create an added operational complexity since protection requires the identification of the failed element and transmission of that information to the end-nodes of all affected paths. This process is referred to as *fault isolation*. Failure-independent schemes are therefore much simpler from an operational viewpoint. Figure 2-14 gives example of schemes in each category. All these schemes will be presented in more or less detail in the following sections.

A third class of survivability schemes could be added to the two main classes (protection and restoration) presented above. That third class comes from the always-present relationship between any restoration scheme and a corresponding pre-planned protection scheme, which is derivable through distributed pre-planning (DPP). The idea is that restoration mechanisms can be used to exercise the network (DPP) for every failure scenario before they happen so that the network may already know what restoration paths will be used at the time of failure. This is different from protection schemes in which connection backup paths are determined at the time of connection establishment. More on that topic will be presented in Section 6.4.2.

2.4.3 Automatic Protection Switching

The simplest form of survivability scheme is *automatic protection switching* (APS), illustrated by Figure 2-15. In APS, each primary path has an associated pre-planned end-to-end backup path, which is used in case of failure on the primary path. There are multiple versions of APS depending firstly on whether the backup path is dedicated to the primary path or shared with other primary paths. When one backup path is shared between multiple primary paths we talk about 1:N APS, N referring the number of primary paths sharing the backup path. When the backup path is dedicated to a single primary path, the scheme is called either 1:1 APS if the backup channel is available to low priority services during normal operations, or 1+1 APS if the signal of the protected service is sent through both channels all the time. This last scheme is the one with the absolute fastest restoration time since a single switching action at the receiving end of the path is required when the signal is lost from the primary channel. Typical switching time is below 50 ms [Sos94].

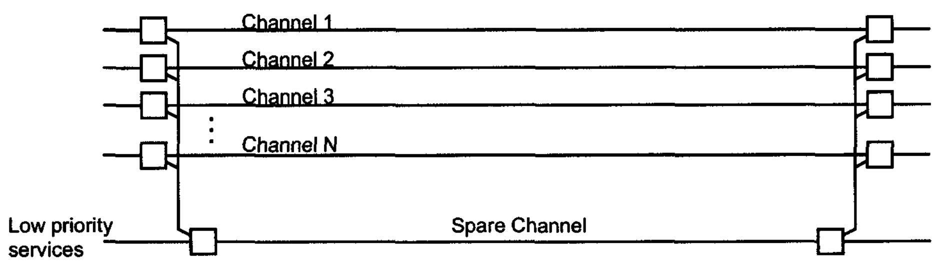

Figure 2-15 1:N automatic protection switching scheme

In order to improve the efficiency of the 1:1 APS or 1+1 APS mechanisms on could use physically fully diverse routes for the primary paths and backup paths. This prevents single physical span failures from causing both primary and backup paths to fail at the same time. In that case, the survivability scheme is referred to as 1:1 APS DP or 1+1 APS DP, where DP stands for "diverse protection." Besides being the fastest restoration mechanism, 1+1 APS is also one of the schemes that require the highest amount of extra capacity. Indeed, provisioning a dedicated backup path with full replacement of the bandwidth requires exactly 100 percent extra capacity, and in the case of 1+1 APS DP, the backup routes are in practice significantly longer than the corresponding primary routes so the required extra capacity is easily well above 100 percent. [DoG01]

2.4.4 Survivable Rings

After APS, survivable rings were the second type of protection mechanism that were developed. With survivable rings the capacity is logically associated to form cycles covering three or more spans. On each cycle an equivalent amount of working and spare capacity is placed and in the event of a span failure, the demands borne of the failed working capacity are rerouted using the spare capacity on the surviving side of the ring. There are two types of survivable rings, depending on the level at which the protection switching is performed. Figure 2-16 illustrates the most common type of survivable ring, the *bidirectional line-switched ring* (BLSR). BLSRs perform switching at the line level (e.g. STS-N line level signals for SONET rings). In the event of a span failure, the entire line level signal is switched at both ends of the failed span to the protection line, which goes all around the ring as shown on Figure 2-16 (b). Another type of survivable ring, the *unidirectional path-switched ring* (UPSR) performs protection switching at the path level (layer above the line layer). With UPSRs, protection switching is done on a per path basis and protection is done at the ingress and egress points of each path for that ring. Unlike BLSR, UPSR is therefore not a completely localized restoration scheme since the rerouting of failed demands is not done between the end-nodes of the failed span, it is in fact partway between a localized response and an end-to-end response. Not surprisingly, as a consequence of that the capacity requirement of UPSR is slightly lower than that of BLSR.

Although usually not quite as fast as APS, survivable rings still offer a very short restoration time (also considered to be in the order of 50 ms) [Wu95]. As for APS, it requires considerable extra spare capacity with the added disadvantage that its rigid structural nature (a ring has to be closed to be a ring) often results in stranded capacity that is of no use but needs to be installed to close the ring. Moreover, the routing of demands in ring-based networks is constrained by the placement of the rings and therefore demands can rarely be routed on shortest path, resulting in increased working capacity requirements [Gro92].

The total extra capacity that needs to be installed to support survivable rings operation is thus

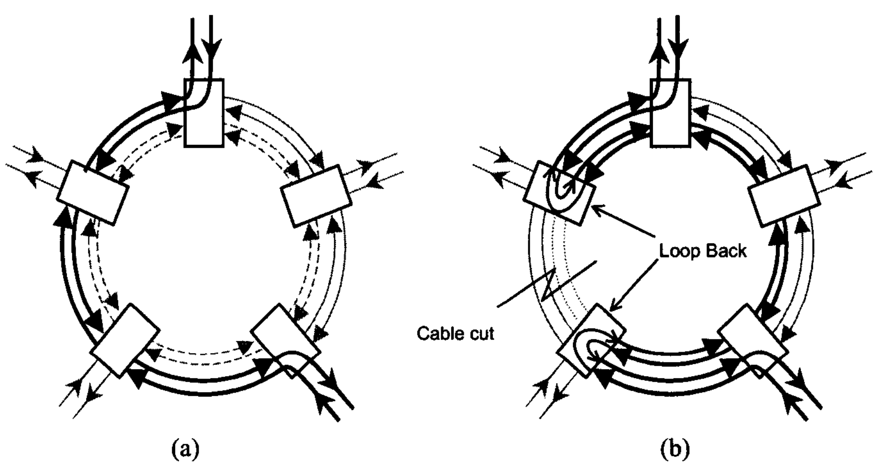

Figure 2-16 Bidirectional line-switched ring
(from [Mor01])

very often well above 100 percent compared to the minimum capacity required for a non-protected network. In fact, it is not rare to see the total capacity requirement of a ring-protected network be in the 200 percent to 300 percent range relative to the non-protected case. However, the simplicity of the ring protection mechanism and the fast restoration have made survivable rings very popular and they have been widely deployed in virtually all public transport networks, in the long-haul part as well as in the metro part.

2.4.5 Survivable Mesh Architectures

The term *mesh* is usually used in opposition to *ring*. Unlike in ring-based networks, where the routing of demands is constrained by the ring structures (demands often have to deviate from their shortest path to travel around the rings that have been installed), demands in mesh networks can usually be routed on their shortest path. The provisioning of capacity in mesh networks is also more flexible since capacity can be added wherever it has been exhausted, one capacity module at a time if needed, whereas with ring-protection entire ring structures have to be placed one at a time, making network scaling more difficult. Strictly speaking APS is a mesh survivability scheme, in the sense that it also allows demands to be routed on shortest path. However, the term mesh is usually associated with more capacity-efficient restoration schemes like span and path restoration, presented in the following section. Besides APS, another mesh protection scheme that has gained lots of interest recently is *shared backup path protection* (SBPP).

2.4.6 Span and Path Restoration

Span restoration (SR) is the localized version of mesh restoration. In SR the re-routing for survivability occurs between the immediate end-nodes of the break. This need not be via a single route, nor via only simple two-hop routes. The general idea of span restoration is illustrated in Figure 2-17. In SR restoration paths are searched dynamically within the available spare capacity at the present time. The restoration path search can be performed by a central network management system or can be performed in a distributed manner through the local broadcasting of path search information tags called *statelets* as described in [Gro94]. The great advantage of the distributed approach is a much faster and reliable restoration path search that does not rely on the integrity of a central database of network state information.

Path restoration (PR), illustrated in Figure 2-18, is the end-to-end version of mesh restoration. In PR, replacement paths are searched for each working path between the path's end-nodes. Similarly to SR, replacement paths for a given pair of end-nodes need not be all on the same route. Thanks to its greater flexibility PR can achieve better capacity efficiency than SR or higher resto-

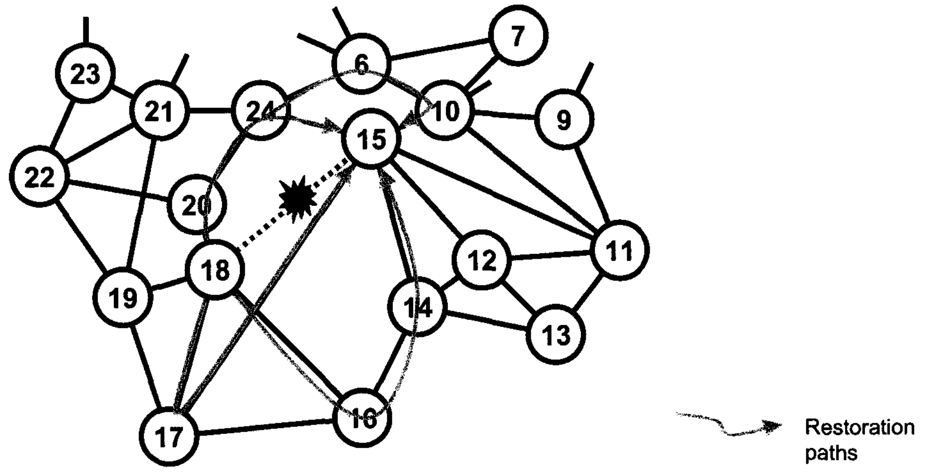

Figure 2-17 General concept of span restoration

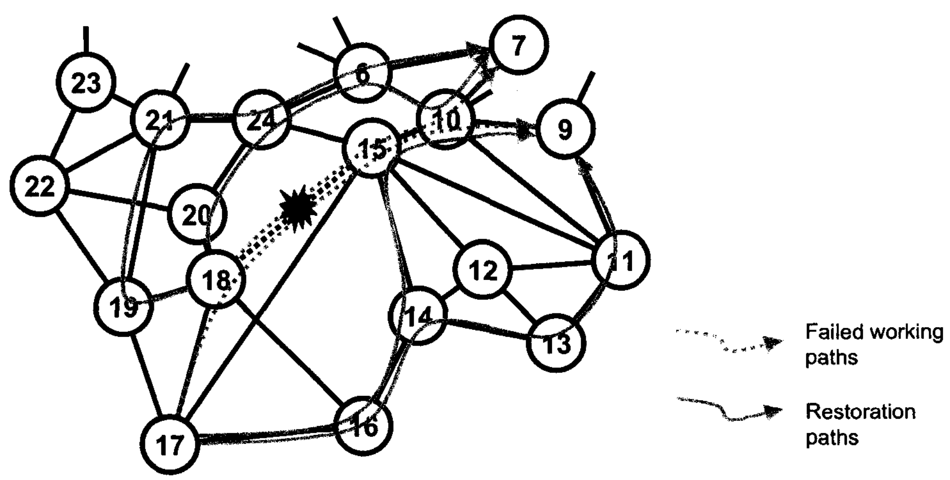

Figure 2-18 General concept of path restoration

ration levels in case of failures. The implementation of PR, however, is much more complex than SR.

2.4.7 Single vs. Multi-Layer Resilience

As explained earlier, public networks are often composed of several levels of client/server relationships. A typical stack of transport layers is IP-over-SONET/SDH-over-WDM. With the advent of optical transport networks including protection/restoration capabilities, we have now multiple layers at which survivability to failures can be performed. The question then arises: "Which layer should take care of the restoration or protection in case of failures?" A first observation is that failures occurring in a given layer cannot be restored or protected in the layers below that. For instance, if an IP interface card fails, protection switching of the SONET STS path used as the physical link for that outgoing port will not solve the problem. If a single layer had to be responsible for all restoration or protection actions it would therefore have to be the highest layer. This approach has been proposed but it is highly debatable whether this would be a viable solution.

Another approach is to let multiple layers contribute to the survivability of the network. This approach, *multi-layer resilience*, has several advantages:

- Restoration of lower-layer failures like fibre cuts is more efficient and faster in those layers.
- Each layer having different types of responses to failures (more or less fast, full bandwidth replacement or partial bandwidth replacement,...), using multi-layer resilience offers various types of service reliability.
- Multi-layer resilience may achieve higher restorability levels than if restoration or protection is handled by a single layer. Reciprocally, the capacity requirements to guarantee full single-failure restorability may be lower if using multi-layer resilience.

For example, in the event of a fibre cut in a network using multi-layer resilience, the optical transport layer could provide immediate replacement wavelengths for service paths carrying voice services. Among the services carried by the other non-protected lightpaths, STS paths carrying data services with high availability requirements would be restored after a few seconds by adaptive restoration performed at the SONET layer and restoration of lower priority data traffic would be referred to the IP layer in which the IP routing tables would be updated within a couple of minutes.

An important issue with multi-layer resilience is that of coordinating the survivability mechanisms of the different layers. For example, a problem is that of identifying the origin of a failure so that a given may know whether it should trigger its survivability mechanisms or wait for the lower layers to deal with the problem. For more on this topic, see [Dem99], [FuV00], [VCD01], [LTC01], [DeM02], [SRM02].

2.5 Transport Network Planning

This section introduces the topic of transport network planning. There are two main problems that fall under this general topic. The first one is that of capacity planning, which is the problem facing a network operator who wants to build a new network or who wants to add capacity to their existing network to support the demand growth. The second that of service provisioning, which is the problem that a network operator is facing on a daily basis as requests for new services are being presented to them.

2.5.1 *Capacity Planning*

The capacity planning problem can generally be summarized as follows: Given a fixed physical topology (set of routes on which fibres can be laid), a set of point-to-point demands with specified restorability requirements, and a particular survivability scheme, find the minimum cost capacity placement that allows all demands to be served and their restorability requirements to be guaranteed using the considered survivability scheme. Another slightly different and very practical question is the following: Given a fixed physical topology, and a set of point-to-point demands with specified restorability requirements, what survivability scheme requires the minimum cost capacity placement to serve all demands and guarantee that their restorability requirements are met?

The capacity planning problem is a long-term problem where the demands assumed are based on forecast of the future demands. All network design studies presented in this thesis work relate to this problem. The objective of these network design studies is not to present methods for the capacity design of real transport network, but to gain insight on the fundamental properties of different survivability architectures and mechanisms in terms of their ability to provide high availability services in a cost effective way. Various methods for network capacity planning are presented in [GBV91], [VGM93], [HeB94], [IMG98], [MiS98], [VVD98], [AVD00b], [DoG00].

The objective of the capacity planning problem can be refined to minimize cost instead of simply capacity. In that case, the cost function can be expressed as a function of span establishment cost (cost to have capacity on a span in the first place), capacity incremental cost, nodal equipment, etc. Revenue maximization consideration may also be included in the objective function of the capacity design problem [SrS00][SSS01].

2.5.2 *Service Provisioning*

The service provisioning problem relates to the probability of blocking of requests for connections in a transport network. With service provisioning the capacity is already placed and the problem is to find connection routing algorithms that minimize the probability of blocking of future connection requests given the statistical distributions of request arrivals and connection holding time. In optical transport networks, the routing of a new lightpath presents two sub-problems: Finding the route on which to establish the new lightpath and finding the wavelength at which the lightpath will be transmitted. This problem is referred to as the *routing and wavelength assignment* (RWA) problem. This extra dimension to the problem makes it quite complex and therefore a huge number of papers have been published on this topic (see [RaS95], [BaM96], [ZJM00], [RaM02]). Several studies have considered the possibility of introducing wavelength conversion capability in the optical network elements and shown the great reduction of the probability of blocking that this brings [SDL01][ASA02]. It was also shown that networks with sparse wavelength conversion capability can already significantly reduce the probability of blocking and that a network with full wavelength conversion capability in every node is not needed [SAS98][VSK99][LiS02][SGC02]. As in the case of the capacity planning problem, revenue maximization can be an added consideration besides minimizing the probability of blocking [AKQ00].

2.5.3 Basic IP Formulation for Optimal Capacity Design Method

The following spare capacity placement (SCP) model was introduced by Herzberg and Bye in [HeB94] and is the basis for all capacity design models presented in this thesis. The problem is that of finding a spare capacity assignment that guarantees the full restorability of all single span failures in a span-restorable mesh network. In this method, the working capacity allocations are assumed to be known and the problem is formulated as that of finding the assignment of restoration flows to the eligible routes on which restoration for each failure could be considered so as to minimize the total amount of required spare capacity. These so-called "distinct eligible routes" are obtained before running the optimization. The graph topology is first processed to find all the distinct logical routes that are "eligible" for use in the restoration routing for each failure scenario. The problem as described, where the working demands are first routed – usually on shortest path – and then the spare capacity optimized, is called the *non-joint* mesh capacity design problem. In a *joint* formulation, which will be presented later in this thesis, the routing of demands is simultaneously optimized with the placement of spare capacity so as to minimize *total* capacity.

For this and subsequent formulations, we use the following notations for parameters (inputs):

$\boldsymbol{S}$ Set of spans in the physical graph (this set is indexed by i when referring to failed spans, and by k when referring to surviving spans),

C_k Cost of each unit of capacity on span $k \in \boldsymbol{S}$,

$\boldsymbol{P}_i$ Set of eligible routes for the restoration of span $i \in \boldsymbol{S}$,

w_i Number of links carrying working demands on span $i \in \boldsymbol{S}$,

$\delta_{i,k}^p$ Equal to 1 if the p^{th} eligible route for span $i \in \boldsymbol{S}$ uses span $k \in \boldsymbol{S}$, equal to 0 otherwise $\forall (i,k) \in \boldsymbol{S}^2$, $\forall p \in \boldsymbol{P}_i$,

The following notations are used for variables:

f_i^p Restoration flow assigned to the p^{th} eligible route for span $i \in \boldsymbol{S}$

s_k Number of links allocated for spare capacity on span $k \in \boldsymbol{S}$

$$\textbf{SCP (Herzberg):} \text{ Minimize } \sum_{k \in \boldsymbol{S}} C_k \cdot s_k \tag{2.1}$$

Subject to:

$$\sum_{p \in \boldsymbol{P}_i} f_i^p = w_i, \forall i \in \boldsymbol{S} \tag{2.2}$$

$$s_k \geq \sum_{p \in \boldsymbol{P}_i} \delta_{i,k}^p \cdot f_i^p, \forall (i, k) \in \boldsymbol{S}^2, i \neq k \tag{2.3}$$

The constraint set in (2.2) ensures that restoration for span failure i meets the target level (full restoration is assumed). Constraint set (2.3) forces sufficient spare capacity on each span k such that the sum of the restoration paths routed over span k may be supported, for every failure span i. The largest simultaneously imposed set of restoration paths on a span effectively sets the minimum feasible s_k value on each span under a given assignment of restoration flows. Hence the formulation works by finding the assignment of flows that pushes up these span-wise minimums to a global minimum on total spare capacity.

2.5.4 Modularity and Economy of Scale

The concept of modularity recognizes the fact that in real life, capacity is usually installed in modules and not one "unit of capacity" at a time. In SONET for example, if 45 STS-1s have to be transmitted through a span of the physical network, it is likely that an OC-48 optical channel will be setup and not 45 OC-1s. The first reason is that OC-1 transmission modules might simply not be available and the second and main reason is that installing 45 OC-1 modules of capacity would certainly cost more than one OC-48. A typical *economy of scale* is such that two times the cost corresponds to a tripling of the capacity.

Modularity and economy of scale can easily be incorporated in the design formulations by defining a set of module types $\boldsymbol{M}$ and specifying the size Z^m of modules of type $m \in \boldsymbol{M}$ as well as the cost C_k^m of placing a module of type $m \in \boldsymbol{M}$ on span $k \in \boldsymbol{S}$ and modifying the objective function to take these parameters into account. The model was first introduced in [DoG00]. New variables n_k^m have to be added to indicate the number of modules of type $m \in \boldsymbol{M}$ placed on each span $k \in \boldsymbol{S}$. To add the modularity aspect to the general optimal spare capacity design model presented above, the objective function of (2.1) would become

$$\textbf{MSCP}\text{: Minimize } \sum_{k \in S} \sum_{m \in M} C_k^m \cdot n_k^m \tag{2.4}$$

and the following constraint would be added:

$$w_k + s_k \leq \sum_{m \in M} Z^m \cdot n_k^m,\ k \in \boldsymbol{S}. \tag{2.5}$$

The resulting formulation is referred to as *modular spare capacity placement* (MSCP). In (2.4), the spare capacity variables have been replaced by total number of modules so as to minimize *total* capacity instead of total allocated *spare* capacity. The spare capacity is not exactly minimized any more, only a "virtual" cap on allocated spare capacity is placed through (2.5). In fact, the spare capacity variables resulting of the optimization may be higher than the minimum values required to guarantee full single failure restorability. Usually the spare capacity assumed in the resulting modular designs is calculated as the difference between the total capacity placed and the working capacity allocated on that span. This effectively merges the strictly minimum spare capacity required with the extra capacity resulting from the modular aspect of the design, referred to as *slack* capacity. As it will be seen in Section 4.5, considering slack capacity as extra spare capacity has very beneficial consequences on multiple failure restorability in the context of adaptive span-restorable mesh networks.

All results of experiments using SCP or MSCP models presented later in the thesis were obtained using the AMPL model presented in Section D.1 of Appendix D. The model corresponds to the MSCP formulation but can also be used to solve the SCP problem by simply specifying a single module type of size 1. The capacity unit costs C_k used for every experiment are the ones specified in the "UNITCOST" column in the description of test network topologies in Appendix A. Although, to avoid introducing a new dimensionality to the problem, these values were simply set to be proportional to the spans' length, the design models presented in this thesis would not require any changes to consider other models where capacity cost would not be a simple linear function of span length. Indeed, the capacity cost can be specified independently for each span and can be any function of span length, geographical location and other parameters specific to each span.

2.5.5 Shared-Risk Link Groups and Fault Escalation

As we have seen in previous sections, a common goal when designing survivable transport networks is to obtain full restorability to all single span failures in the physical layer. The objective is to eliminate the possibility that a single *physical* point of failure will cause outage. Single *physical* breakage, however, can cause multiple failures in what the transport network sees as its physical –

apparently diversely-routed – layer. Although great effort is put into ensuring physical diversity of transmission facilities, it is very hard for network operator to completely avoid situations were two supposedly diversely-routed fibres have to share a common duct to cross a bridge or follow a train track. A situation like this is referred to as *shared-risk link group* (SRLG) and is illustrated in Figure 2-19.

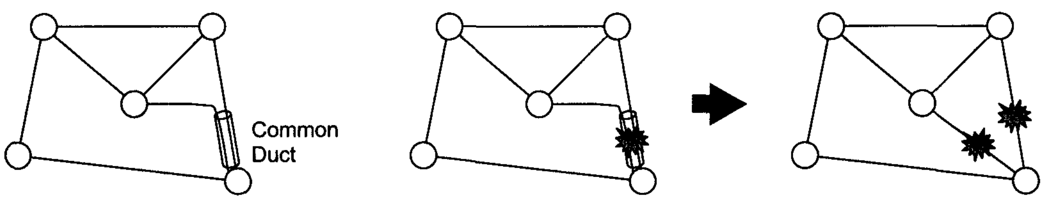

Figure 2-19 Shared-risk link group

Fault escalation is another type of situation where multiple common cause failures can occur in the physical layer of a transport network. This is becoming more and more of an issue with the advent of optical transport networks which provides lightpaths to the higher transport layers like SONET/SDH. The lightpaths that the SONET/SDH layer take as its physical transmission layer are obtained by connecting wavelengths on different physical spans of the "real" physical layer. As shown by Figure 2-20, a single failure in the physical layer of the optical transport network multiple failures in the logical layer, which is seen by the higher SONET/SDH layer as its physical layer.

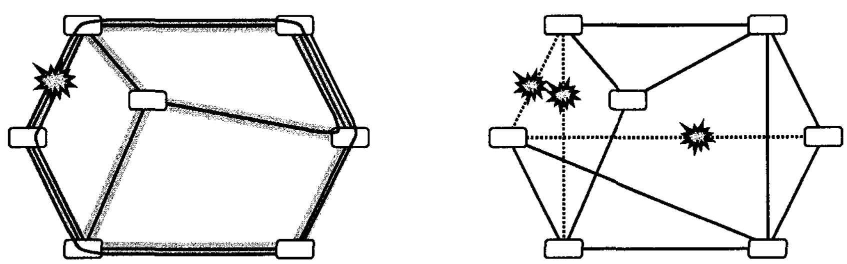

Figure 2-20 Fault escalation

SRLGs, when they cannot be avoided represent a strong motivation to investigate the possibility of providing at least a certain degree of multiple failure restorability. This aspect is largely covered by this thesis.

The problem of fault escalation gave rise to the topic of *logical topology design*, in which is concerned with the choice of topological routes used by a transport network to establish the connections offered to the higher layers. An important condition for example is to guarantee that no single failure in the physical layer of the transport network will cause disconnection of the above client networks. This topic is treated in [RaS96], [CrL98], [CLG00], [MoN01], and [LCS01].

2.6 Challenges for Future Transport Networks

2.6.1 Service Transparency

One of the most important goals of future transport networks is the ability to provide network connections for any type of client services, independently of the signal's bandwidth and nature (digital or analog). This is referred to as *service transparency* or *signal transparency* [RBB00]. The advantage of service transparency is to reduce the complexity of the transport network since it does not need to access client signals except at their network's ingress and egress points. Signals of all types are mapped into transport network's signals (usually optical signals) and the transport network handles these signals only.

One possible solution for service transparency is to use lightpath connections in fully transparent optical networks (networks with no O/E/O conversion). The advantage of this solution is that it is compatible with both digital and analog signals and that it allows the transmissions up to extremely high bit rates (or high bandwidth for analog signals). The drawback is that very few services require the amount of bandwidth that a lightpath provides so using lightpaths for all connections would result in a very high waste of bandwidth. Moreover, fully transparent networks are not likely to be developed, at least not in the near future.

Other solutions proposed for service transparency include ATM over SONET, Packet over SONET (POS) and, more recently, Generic Framing Procedure (GFP). [CMY02]

2.6.2 Reliable and Differentiated Service

An increasingly important requirement of transport networks is the ability to provide reliability guarantees with their services [RBB00]. Protection or restoration mechanisms will be needed to provide the reliability levels required by the always increasing volume of traffic being transported. But transport networks will also be required to offer various levels of reliability corresponding to different types of services. For example, client signals of multi-media applications such as live video transmissions, voice connections, will mainly require very fast restoration, whereas the private lines used for a financial company's WAN will mainly require high availability.

An important part of this thesis work deals with these issues of differentiated services (in particular Chapters 6 and 8). A strong emphasis will be placed on explaining what really characterizes the reliability of services, how much restoration speed matters in terms of service availability, how transport networks can be designed to offer service differentiation, etc.

2.6.3 Connection Service Time

Another big challenge of future transport networks will be the development of simpler and faster procedures for the establishment of new network connections (also called *service provisioning*). Ideally, the goal would be to allow near instantaneous network connection establishment (as it is the case in today's telephone network) where customers would be able to "dial in" and would not need to talk to the network operator at all, however this is not very realistic, at least not in the near future. More realistically, the goal of network operators is to be able to have ways to manage their network remotely using network management software and to simply provision new service paths by specifying the new connections end-nodes and letting the network management system set up the new connection automatically ("point and click").

The mesh network using distributed and self-organizing techniques, which is at the centre of this thesis work, looks like the most promising candidate in terms of its inherent ability to provide fast connection service time. This is due to the simplicity of routing new demands on shortest path in mesh networks, possibly within a guaranteed-restorable working capacity envelope (WCE) [Gro03].

2.6.4 Dynamically Adaptive Networks

Providing connections that can be reconfigured dynamically is another goal of future transport networks. The idea is to allow connection characteristics to be changed automatically following a change in the volume of traffic served by that connection or changed following a command initiated by the network operator. Solutions for dynamic traffic adaptation have been proposed using the techniques of *Virtual Concatenation* (VC) in conjunction with *Link Capacity Adjustment Scheme* (LCAS) [CMY02]. With VC multiple SONET standard STS-1 signals can be combined to offer various connection types with different bandwidths that will accommodate a wide range of different services. The aim is to make better use of SONET payloads since a greater range of bandwidth possibilities are offered. LCAS adds the possibility of dynamically modifying the number of concatenated SONET signals to adapt the VC connection's bandwidth to the volume of traffic served at any time. Another application of LCAS is to dynamically change bandwidth allocations to adapt to temporary restoration flows during failure or maintenance states.

2.6.5 Capacity Efficiency

Among all challenges facing future transport networks there is one common challenge which is to provide all these services and new features in a capacity efficient way. Unlike what is heard sometimes, bandwidth is not free and reducing the amount of required bandwidth by a few percent in some long haul networks can translate into huge cost savings [DoG00]. All analyses and experiments presented in this thesis will consider capacity cost as a major factor and the performance of mesh networks in terms of restorability will often be evaluated with respect to capacity requirements. An important body of literature is dedicated to cost optimal network capacity design. For example, see [VVD98], [MiS98], [IMG98], [DDH99], [DoG00].

2.7 Summary

This chapter has presented an introduction to the general concept of transport networking, explaining where transport networks fit in the general picture of telecommunication networks. We introduced the different types of transport networks (the two main categories being circuit-switched networks and packet-switched networks) and the different generations of transport networks (from the PSTN to all-optical transport networks). The emphasis was placed on SONET, which is still the main transport network technology in use today, and on WDM optical networks, which will be the basis for the development of future transport networks. We then introduced the issue of network survivability and the different classes of survivability mechanisms. The emphasis was placed on mesh-based survivability schemes and, in particular, on span-restoration, which is at the centre of this thesis work. Finally, we discussed the important issues related to transport networking, including transport network capacity design and we highlighted the challenges facing future transport networks. Another main issue related to transport networking is the analysis of the reliability of services provided by these networks. This issue is what this thesis is mainly about and it is studied in detail in the following chapters.

3. Availability Analysis

3.1 Introduction

This chapter introduces the topic of availability analysis. First, the important mathematical definitions are given with a particular emphasis on how the availability function differs from the reliability function. Then, we highlight the general existing methods to perform the analysis of a system's availability. We then consider the specific case of telecommunication networks and present various common measures of availability in networks. A review of the literature published on this topic to date is then presented. Finally, we present a new approach to the problem of determining the availability of service in transport networks using the span-restoration mechanism that was briefly presented in Section 2.4.6

3.2 Mathematical Definitions

3.2.1 Reliability

The reliability of a system is defined as the probability that the system will perform the function it is designed for during a defined period. In other words it is the probability that no failure will occur during that period. Although the term reliability is often used in the literature related to service assurance in telecommunication networks – especially for packet-switched networks – the concept of reliability relates to mission-oriented systems, like a space launcher or a Formula One car. The concept of reliability therefore relates to the question "what is the probability that the engine of this Formula One car will work for the entire race?" and not to "what is the probability each day that the engine of my car will start?"

The reliability function, as defined in (3.1), is a function of the duration T of the mission (and not a function of time.)

$$R(T) \equiv P\{\text{no failure in } [0, T]\} \tag{3.1}$$

The reliability function can be expressed in terms of the *failure density function* $f(t)$ as follows:

$$R(T) = 1 - \int_0^T f(t)dt . \tag{3.2}$$

The function $f(t)$ is in fact the probability density function of the *time to failure*, therefore integrating $f(t)$ over a certain time period gives the probability that the first failure will occur in that time period. Conversely, by differentiating (3.2) one can express $f(t)$ in terms of $R(T)$:

$$f(t) = \frac{d}{dt}R(t). \tag{3.3}$$

The expectation of $f(t)$ gives the *mean time to failure* (MTTF), which is a useful measure in availability analysis as seen in the following section:

$$MTTF = E(f(t)) = \int_0^{\infty} (t \cdot f(t))dt. \tag{3.4}$$

3.2.2 Availability

Unlike with reliability, the concept of availability is related to repairable systems and considers two equally valid ways for a system to be in working state: either it has been working without any failure since it began operating, or it has failed once or several times but has been repaired each time. Availability of a system is defined as "the probability of the system being found in the operating state at some time t in the future given that the system started in the operating state at time t = 0 and given that failures and down-states occur but maintenance or repair actions always return the system to an operating state." [BiA92]

The availability function is therefore defined as follows:

$$A(t) \equiv P\{\text{system in operating state at time } t\}. \tag{3.5}$$

The availability function is a function of time that starts from 1 and usually stays at a high level shortly after the systems starts operating and then decreases to eventually reach a steady state in which repairs compensate for failures and maintain the availability at a certain constant level A:

$$A = \lim_{t \to \infty} A(t) \tag{3.6}$$

The steady state availability A can be obtained by observing the system over a long period of time T_{obs} and calculating the fraction of the time the system is up:

$$A \equiv \lim_{T_{obs} \to \infty} \left\{ \frac{\text{Total up-time in } T_{obs}}{T_{obs}} \right\} \tag{3.7}$$

Based on (3.7), another very useful expression of A is derived in (3.8) to (3.10). In these equations, N is the number of failures in the observation period, $TTF(i)$ is the time to the i^{th} failure (time between a repair or the beginning of the observation window and the occurrence of next failure), $TTR(i)$ is the time to repair the i^{th} failure, and $MTTR$ is the mean time to repair.

$$A = \frac{\sum_{i=1...N} TTF(i)}{T_{obs}} = \frac{\sum_{i=1...N} TTF(i)}{\sum_{i=1...N} TTF(i) + \sum_{i=1...N} TTR(i)} \tag{3.8}$$

$$A = \frac{\left(\sum_{i=1...N} TTF(i)\right)/T_{obs}}{\left(\sum_{i=1...N} TTF(i)\right)/T_{obs} + \left(\sum_{i=1...N} TTR(i)\right)/T_{obs}} \tag{3.9}$$

$$A = \frac{MTTF}{MTTF + MTTR} \tag{3.10}$$

Because *MTTF* is not always known, the expression of A in (3.10) is sometimes replaced by the following expression:

$$A \approx \frac{MTBF}{MTBF + MTTR}, \tag{3.11}$$

where *MTBF* is the mean time between failures. *Time between failures* in *MTBF* refers to the time between the occurrence of failures, whereas *time to failure* in *MTTF*, as explained earlier, is the time between the repair of a failure and the occurrence of the next failure. Therefore we have:

$$MTBF = MTTF + MTTR \tag{3.12}$$

Strictly, equation (3.10) gives the correct expression although with typical values of *MTTF*, *MTBF* and *MTTR*, both equations give almost identical result as shown in (3.13) based on (3.10) and (3.14) based on (3.11) with an *MTBF* of 8766 hours (1 year) and an *MTTR* of 12 hours.

$$A = \frac{8766 - 12}{8766} = 0.998631 \tag{3.13}$$

$$A = \frac{8766}{8766 + 12} = 0.998633 \tag{3.14}$$

Very often, for reasons that are explained in the following sections, instead of working on availability values, it is easier to work with *unavailability* values. The unavailability is the complement of the availability:

$$U \equiv 1 - A, \tag{3.15}$$

and is therefore defined as the probability of finding the system in the non-operating state.

It can also be expressed in terms of *MTTF*, *MTBF* and *MTTR*:

$$U = \frac{MTTR}{MTBF} \approx \frac{MTTR}{MTTF} \tag{3.16}$$

3.2.3 *Mathematical Simplifications*

A general method to determine the availability of a system is to develop a logical block diagram with the different elements that compose the system. The block diagram can often be reduced by applying some simplifications for elements in series and elements in parallel, as detailed in the following lines.

In the case of a system composed of elements in series, all elements need to be working for the whole system to be functional. Therefore the availability of a system of N series elements E_i is:

$$A_{\text{syst}}^{\text{s}} = P\{(E_1 \text{ is up}) \cap (E_2 \text{ is up}) \cap \ldots \cap (E_N \text{ is up})\}, \tag{3.17}$$

which can be expressed in terms of the availability values A_i of the N elements in the system as follows:

$$A_{\text{syst}}^{\text{s}} = \prod_{i=1}^{N} A_i . \tag{3.18}$$

From (3.15) and (3.18), a simple expression can be derived for the unavailability of a system of N series elements:

$$U_{\text{syst}}^{\text{s}} = 1 - A_{\text{syst}}^{\text{s}} = 1 - \prod_{i=1}^{N} A_i = 1 - \prod_{i=1}^{N} (1 - U_i) \tag{3.19}$$

$$U_{\text{syst}}^{\text{s}} = \sum_{i=1}^{N} U_i - \sum_{i,j=1,\, i \neq j}^{N} U_i \cdot U_j + \sum_{\substack{i,j,k=1 \\ i \neq j,\, i \neq k,\, j \neq k}}^{N} U_i \cdot U_j \cdot U_k - \ldots \tag{3.20}$$

In general the unavailability values of elements considered are much smaller than 1 and therefore the previous equation can be simplified to give:

$$U_{\text{syst}}^{\text{s}} \approx \sum_{i=1}^{N} U_i . \tag{3.21}$$

Equation (3.21) is a widely accepted expression of the unavailability of elements in series. Accordingly, it is often said that "unavailabilities add for elements in series."

For systems composed of elements in parallel, all elements have to be failed for the system not to function, therefore:

$$U_{syst} = P\{(E_1 \text{ is down}) \cap (E_2 \text{ is down}) \cap \ldots \cap (E_N \text{ is down})\}, \tag{3.22}$$

and:

$$U_{syst} = \prod_{i=1}^{N} U_i. \tag{3.23}$$

Thus, it is often said that "unavailabilities multiply for elements in parallel."

3.3 General Methods for Availability Analysis

When the block diagram has been reduced as much as possible using the series and parallel elements simplifications presented in the previous section, two main methods can be used to complete the availability study of the system. The first one, the *tie paths method* consists in identifying all the possible configurations in which the system is in an operational state, and the second one, the *cut sets method*, involves determining all the possible configurations in which the system is failed.

3.3.1 Tie Paths Method

In the tie path approach, we are looking for the set $\boldsymbol{P}$ of all the paths between input and output in the block diagram of the system. The availability of each path p in $\boldsymbol{P}$, as explained in the previous section, can be expressed as the product of the availability values A_i of all series elements in the path:

$$A_{path}(p) = \prod_{i \in \text{path } p} A_i \tag{3.24}$$

The availability of the whole system is the union of the probabilities of all paths:

$$A_{syst} = \bigcup_{p \in \boldsymbol{P}} A_{path}(p) \tag{3.25}$$

Because there are states of the system in which multiple paths in $\boldsymbol{P}$ are simultaneously working, the union of the path probabilities cannot simply be obtained by summing the availabilities of each path. In order to get the exact system availability value, the inclusion-exclusion principle needs to be applied [Bru92]. Because of the relative complexity of having to use this principle and because there is usually a very large number of ways in which a system can be available, making the listing of all tie paths is a very tedious task, and so the cut set method presented in the following section is usually a preferred method.

3.3.2 *Cut Sets Method*

In the cut sets method we take the opposite approach: we determine the set C of all minimal combinations of elements, which if failed simultaneously will guarantee failure of the system. By minimum it is meant that if the simultaneous failure of element x and element y is sufficient to cause the system to be unavailable, then there is no need to also consider the combination of element x, element y and element z: it is already implicit in the previous combination. Each such combination that guarantees failure of the system is in fact a *cut* of the graph in the block diagram (no path is feasible between the input and the output). The probability of a cut set c in C corresponds to the unavailability of a system of parallel elements:

$$P\{\text{cut } c\} = \prod_{i \in \text{cut } c} U_i \tag{3.26}$$

The unavailability of the whole system is the union of the probabilities of all cuts:

$$U_{\text{syst}} = \bigcup_{c \in C} P\{\text{cut } c\} \tag{3.27}$$

If all cuts are considered, then the unavailability of the system is strictly equal to the sum of the probabilities of all cuts:

$$U_{\text{syst}} = \sum_{c \in C} P\{\text{cut } c\}\,. \tag{3.28}$$

However, instead of determining all cuts, a typical strategy to limit the complexity of finding cut sets is to limit the search to minimal-weight cuts, that is consider only cuts with less than a certain number of elements. In that case, summing the probability of these cuts only provides a lower bound on the unavailability of the system, but it is generally a good approximation of the real system unavailability since higher weight cuts usually have a much smaller probability and therefore contribute much less to the system's unavailability. This strategy can be referred to as "most likely path to failure." [WiS97]

The cut set method is the method used most of the time to determine a system's availability, which explains why it was said earlier that it is usually easier to work with unavailability than availability. In this thesis, the general method used to determine availability of service is close to the approach used by the cut sets methods, although slightly different to be applicable to the complex nature of networks based on distributed dynamic restoration. More on this will be presented in Section 3.6.5.

3.3.3 *Markov Modelling*

Besides the *cut sets* and *tie paths* methods, a technique sometimes used for availability analysis is *Markov modelling* [Joh89]. A Markov model is developed by identifying a set of different states of the system and determining transition rates between these states. In each of these states the system has a constant value of the metric under study and the goal of the technique is to identify the probability of each state and therefore the statistically expected value of that metric. There are two main requirements for the states defined in a Markov model. The first requirement is that the state transition rates only depend on the state the system is in and not on the previous states it was in before entering that state (condition referred to as *memorylessness*). The second requirement of Markov modelling is that state transition probabilities not change with time (the statistical transition processes are stationary).

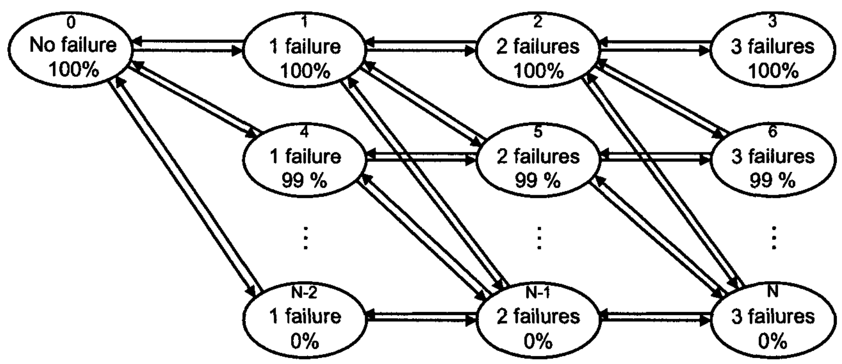

Figure 3-1 Example of Markov model for network availability analysis

An example of how Markov modelling can be used for availability analysis is shown in Figure 3-1. In this example, the metric under study is the percentage of all network connections being in operational state serving as a surrogate for network availability (we will see in Section 3.4 that other definitions of availability can be used.) That value is shown in each state. The other parameter that changes from one state to another is the number of failures outstanding in the network. In this model, up to three simultaneous failures are considered possible (for higher numbers of simultaneous failures, the probabilities are considered too small and are ignored.) The states transition rates are not shown but a specific rate would be associated with every arrow shown between two states. Determining the state probabilities is done by solving a set of equations linking the probabilities of the different states and the transition rates (see [Leo94] for details of these equations). Once the state probabilities are known, the availability of the system can be calculated using the average of the availability metric $A(i)$ in the N different states, weighted by the probability p_i of each state:

$$A_{syst} = \sum_{i=1}^{N} p_i \cdot A(i) \tag{3.29}$$

The applicability of Markov models to transport network availability analysis will be further discussed in Section 3.5.

3.3.4 *Simulation Based on the Markov Model*

Once a Markov model has been developed, an alternative to solving the equations is to run a time simulation of the evolution of the system through the set of Markov states. Each time the system enters a state, the different transition rates from that state are used to randomly determine which state the system will enter next and at what time. On top of having the advantage of not requiring to solve the Markov equations (albeit at the expense of developing a simulation tool), this method also has the advantage of producing more information than the purely mathematical approach. Indeed, as the system navigates through the state model, statistics of time spent in each state can be collected. whereas the purely mathematical method only provides the overall probability of being in each state.

Simulation using the Markov state model, however, suffers from the same limitations as the basic Markov technique, as will be explained in Section 3.5.4. Therefore, another simulation based approach will be used in this thesis.

3.4 Availability in Telecommunication Networks

3.4.1 *Network Availability vs. Service Path Availability*

Most of the time in the previous sections we have used the term *system* to designate the object under study. In network availability analysis, the first question is to decide what the system is. Then, we must decide what criterion is used to determine whether the system chosen is in the working or failed states.

There are several approaches to this question, the first being to consider the network as a whole as the system. In this case, there are several possible rules for deciding the state of the system. For example, it could be decided that the network is considered as working only when all connections of the network are themselves in working state. This approach, in fact, is not a very good one since it is rare that all connections in a network are simultaneously in a working state and such a model could therefore give very low availability results, not necessarily reflecting the reality of what is experienced by the users. Another rule could be that the network will be considered to be in working state if more than a certain percentage of all the connections are in a working state. This model brings an improvement over the previous model but there is still the difficulty of determining what the level should be. Moreover, when a certain level is chosen, the model might give the same results for networks in which the proportion of working connections is, on average, much lower or much higher than the threshold level, and networks in which the proportion is on average respectively, just below or just above the threshold.

The most important drawback of choosing the network as a whole as the system is that the results obtained are only useful to get a general feel of how available the network is but are not very useful or even meaningful from a network customer's point of view. Indeed, it does not matter to a customer buying a connection from a network whether that network as a whole has an availability of 0.95 or 0.99999 as long as the availability of the service provided to them meets their expected level. This is why, in this thesis work we choose the service path as the system under study. Another motivation for choosing service paths as the systems under study is that, as it will be seen later in the thesis, service path availability is highly dependent on the path length or on the specific spans traversed by a path. It makes therefore much more sense to study the availability of specific paths instead of determining the availability of the network and then to deduce availability values for paths in general regardless of their specific properties. The availability results presented in this thesis are usually based on the computation of the availability of some representative *hypothetical digital reference path* (HDRP) or are averages over all paths of several classes of paths with common properties (like *path length* and *protection class*, to be defined later.)

3.4.2 *Common Measures and Units Related to Availability*

Several common measures are used to express availability in the world of telecommunication networks. The most common one is the percentage measure we have presented in the previous section, where A is defined as the expected fraction of the time the system is working [Spr93]. A way used very often to express this percentage is to state the number of 9s corresponding to that availability value. For example, an availability of 0.99999 is referred to as *five 9 availability*. It should be noted that every 9 added to the availability figure corresponds to a reduction of the expected unavailability by a factor 10.

Another way to express availability is to state the expected number of minutes of outage per month or per year. For example, with an availability of 0.999 a system is expected to experience about 8.5 hours of outage per year, as shown by the following calculation:

$$\text{Expected yearly outage } (A = 0.999) = (1 - 0.999) \times 8766 \approx 8.5 \text{ hours}. \qquad (3.30)$$

This way to express the availability is motivated by the growing importance of *service level* *agreements* contracts (presented in the following section) signed between clients and network service providers. This approach, however, has some limitations as explained in Section 3.4.4 and should be used cautiously, especially when used to decide the terms of a service level agreement.

A measure found in some papers is called the *expected loss of traffic* (ELT) [AVD00a][DeM03]. It is defined as the amount of traffic that a network is expected to lose in a certain time period – usually one year. The expected loss of traffic in a time T, can be related to the availability function as follows:

$$\mathrm{ELT}(T) = \sum_{c \in \boldsymbol{C}} \mathrm{Cap}(c) \cdot (1 - A(c)) \cdot T \tag{3.31}$$

where $\boldsymbol{C}$ is the set of network connections, $\mathrm{Cap}(c)$ is the capacity of a connection c in $\boldsymbol{C}$ and $A(c)$ is the availability of connection c. This measure is interesting from a network operator's point of view since it does not only reflect the frequency of outages but also the magnitude of these outages.

A common unit related to availability and widely used in networking is the FIT, which stands for "Failure in 10^9 hours." This measure is used to characterize the reliability of network components, which are highly reliable. A value of 1 FIT means that a component is expected to fail on average once every 10^9 hours. This measure has some relevance to availability analysis that lies in the fact that the FIT value can be used to calculate the mean time to failure (in hours) of a component as follows:

$$MTTF = \frac{10^9}{FIT} \tag{3.32}$$

Finally, another commonly used measure related to network availability is the number of *defects per million* (DPM) [Ogg01]. This measure characterizes the network as a whole and indicates the number of device failures that happen in a million hours of device operation time. For example, two device failures in a month for a network with a thousand devices corresponds to 2.74 DPM.

3.4.3 Service Level Agreements

One of the major motivations in developing methods for service availability analysis in telecommunications networks is the growing importance of *service level agreements* (SLA). A service level agreement is a contract signed between a client and a network service provider, that specifies the service level guarantees that the service provider is committed to provide to the client. These guarantees can be expressed in terms of minimum average service availability over a year (or equivalently maximum total outage time per year) and/or minimum average service availability over a month (or equivalently maximum total outage time per month). The agreement can also include other details like maximum consecutive number of minutes of outage, maximum number of outage events per month or per year, etc.

Along with these guarantees, SLAs specify the type of penalties that the service provider will have to suffer if the guarantees are not respected. Such a penalty would be, for example, a free month of service for a customer whose SLA was not respected in the previous month. Given such risks linked to not respecting SLAs, network operators need to have not only an accurate knowledge of the expected total service outage in a month, but also a precise knowledge of the probability distributions of total outage, number of outages, duration of outages, etc., for all service paths in their networks. Such a knowledge will allow them to limit the risks they are taking and therefore be able offer competitive SLAs. Conversely, knowing the expected service availability of paths in the network may allow a network operator to decide when and how the network needs to be upgraded in order to be able to achieve the service levels promised by SLAs.

Developing methods for detailed service paths availability analysis and availability-based network design is a central topic of this thesis.

3.4.4 Statistically Expected Availability vs. Probability of Outage

One possible misinterpretation of the concept of availability that it is useful to clarify is the link between availability and total outage per year. It is common to read in papers on network survivability that an availability of 0.99999 (five 9s) corresponds to about 5 minutes of outage per year. This figure can indeed be obtained by the following calculation:

$$\text{outage in a year} = (1 - 0.99999) \times 525960 \approx 5.3 \text{ min} . \tag{3.33}$$

This calculation is perfectly correct but the result should be more carefully stated as "an availability of 0.99999 corresponds to an *expected total outage* of about 5 minutes per year." The apparently very subtle difference in the wording is important, especially for a network operator that wants to determine what kind of SLA can be offered to customers. To illustrate this point, consider the case of a transport network in which physical span failures happen on a regular basis and for which it takes on average 12 hours for these physical failures to be repaired by technicians. Let's say, for the sake of the argument, that this network is fully restorable to single span failures and that the main cause of outages are dual span failures (Section 3.6.4 will explain that this is indeed the main cause of outage in restorable transport networks). When a service path is failed because of a dual failure, it takes on average 6 hours before one of the two failed spans is repaired (proof of this is given in Appendix C). The average outage time for service paths put in a non-restorable state is therefore 6 hours. In these conditions, it appears that a service path is very likely to experience a total outage of either a few seconds per year (added restoration times of all the failure events in which the service was restored, if the service path was never not-restorable to a dual-failure), or a total outage in the order of 6 hours (if the service path was not-restorable once), or a total outage in the order of 12 hours (if the service path was not-restorable twice), etc. The statement "5 minutes of outage per year" therefore only has a statistical meaning in this case, but does not correspond to any event that is likely to happen. These "5 minutes of expected outage per year," which are much lower than a real outage would be, would in fact be observed if we took the average over all paths and over a few years of operation. The way this should be interpreted is thus, rather, that the probability for service paths of experiencing an outage each year is simply very low.

We therefore advocate that the calculation of total outage time per year should only be used when the probability of experiencing outage (in the sense "hard outage," not a collection of small restoration times) in a year is a significant fraction of 1, and even preferably when the probability of experiencing multiple such outages is also high. In cases when the probability of experiencing outage becomes negligible, then doing this calculation becomes meaningless and what should be considered instead is in fact the probability of experiencing outage.

The equations following in this section determine the expression of the probability of experiencing outage and link it to the availability of service. Since developing exact equations that take into account the complex nature of outage probability distributions of service paths in a restorable network would be quite complicated (the complexity of developing closed form models for statistical processes linked to restorable services will be shown in Chapter 4), we assume for simplicity that times-to-outage for restorable service paths are Poisson-distributed with a constant hazard rate and that the well-known Poisson equations can be applied. It could in fact be argued that the statistical process linked to the occurrence of outages for service paths *does* in fact intuitively correspond to the definition of a Poisson process. The goal anyway is not to claim exact results but to show how increased availability qualitatively translates into a reduced probability of outage.

For Poisson-distributed arrival times with a constant hazard rate λ_0, the probability of n arrivals in an observation window of size T is:

$$P_n(T) = \frac{(\lambda_0 \cdot T)^n}{n!} \cdot e^{-\lambda_0 \cdot T}. \tag{3.34}$$

By combining the relation between λ_0 and *MTTF*, shown in (3.35) and the relation between *MTTF*, *MTTR* and A shown in (3.36), one obtains the expression of λ_0 in terms of A and *MTTR* shown in (3.37).

$$MTTF = 1/\lambda_0 \tag{3.35}$$

$$A = \frac{MTTF}{MTTF + MTTR} \tag{3.36}$$

$$\lambda_0 = \frac{1-A}{A \cdot MTTR} \tag{3.37}$$

Equation (3.34) can be used to obtain the expression of the probability of not experiencing any outage in *T:*

$$P_0(T) = e^{-\lambda_0 \cdot T} \tag{3.38}$$

as well as the expression of the probability of experiencing at least one outage in *T*:

$$P_{n \geq 1}(T) = 1 - P_0(T) = 1 - e^{-\lambda_0 \cdot T} \tag{3.39}$$

and the probability of experiencing at least two outages in *T*:

$$P_{n \geq 2}(T) = 1 - (P_0(T) + P_1(T)) = 1 - (e^{-\lambda_0 \cdot T} + \lambda_0 \cdot T \cdot e^{-\lambda_0 \cdot T}) \tag{3.40}$$

$$P_{n \geq 2}(T) = 1 - (1 + \lambda_0 \cdot T) \cdot e^{-\lambda_0 \cdot T} \tag{3.41}$$

These equation are used to produce the results shown in Tables 3-1 and 3-2. In these tables are related different measures of availability introduced in the previous sections. In Table 3-1 the results presented are for a network in which the average outage is 6 hours. In this case the value of 6 hours is used for *MTTR* in (3.37). This *MTTR* value is not to be confused with the physical *MTTR* of spans, it is in fact the mean time to bring failed service paths back to working state, i.e. average outage time. In Table 3-2, the results presented are for a network in which the average outage is 5 minutes. An example of such network would be an IP layer transport network in which outage is

Table 3-1: Relating the different availability measures in optical transport network

Number of 9s	A	U	Expected downtime per year (stat.)	Probability of no outages in a year	Probability of 1 outage or more in a year (with 6-hour outages)	Probability of 2 outages or more in a year (with 6-hour outages)
1	0.9	10^{-1}	36.5 days	~ 0	~ 1	~ 1
2	0.99	10^{-2}	87.7 hrs.	3.90×10^{-7}	0.9999996	0.999994
3	0.999	10^{-3}	8.8 hrs.	0.232	0.768	0.430
4	0.9999	10^{-4}	52.6 min.	0.864	0.136	0.010
5	0.99999	10^{-5}	5.3 min.	0.985	0.015	1.06×10^{-4}
6	0.999999	10^{-6}	31.6 s	0.999	0.001	1.06×10^{-6}

due to router reboot or routing table re-calculation.

As shown in Table 3-1, for this type of network, an availability of 0.999 seems to represent some kind of threshold. Below that point (for A smaller than 0.999), the probability of experiencing one or even two outages is virtually 1. At an availability of 0.999 the availability of experiencing one or two outages are fairly high but it is not guaranteed anymore. For availabilities above that point, the probability of experiencing outage becomes small. In the "Expected downtime per year" column, the cells with a grey background indicate values that only have a statistical meaning as explained above and should not be considered as meaningful information when considering the case of a single service path in a single year.

Table 3-2: Relating the different availability measures in the IP layer

Number of nines	A	U	Expected downtime per year (stat.)	Probability of no outages in a year	Probability of 1 outage or more in a year (with 5-min. outages)	Probability of 2 outages or more in a year (with 5-min. outages)
1	0.9	10^{-1}	36.5 days	~ 0	~ 1	~ 1
2	0.99	10^{-2}	87.7 hrs.	~ 0	~ 1	~ 1
3	0.999	10^{-3}	8.8 hrs.	~ 0	~ 1	~ 1
4	0.9999	10^{-4}	52.6 min.	2.69×10^{-5}	0.99997	0.9997
5	0.99999	10^{-5}	5.3 min.	0.349	0.651	0.283
6	0.999999	10^{-6}	31.6 s	0.900	0.100	5.16×10^{-3}

In Table 3-2, we see that for networks in which the average outage time is lower, the probability of outage is still very high for high availability values. Here, in the case of an average outage time of 5 minutes, the probability of at least 1 outage in a year is still of 10 percent for an availability of 0.999999 (six 9s). In such networks it is therefore justifiable to talk of about 5 minutes of outage per year for an availability of 0.99999.

Results of this type obtained from network time simulations will be presented in Chapter 9.

3.5 Literature Related to Network Availability

3.5.1 Mesh Network Reliability

Judging by name only, the closest body of literature to our present concern would appear to be that of "network reliability" as surveyed in [RaA90]. However, this field is concerned with various measures of the graph disconnection probability under the assumptions of edge failure probabilities that are very high (compared to our case) and that there is no limitation on the number of simultaneous edge failures. While this is a challenging theoretical problem, the existence of a topological route is not a sufficient condition for path availability in a realistic transport network. The path availability also depends intimately on the capacitation of the network and the details of the restoration mechanism that applies. Network reliability (in the sense of [RaA90],[HKC95],[Col87]) ignores standby redundancy and the repair of physical failures. A specific sub-problem is the "terminal pair" availability problem.

Iselt [Ise00] gives an excellent survey of computational approaches to this problem, which can represent the reachability between terminal pairs in a packet-switched network or be related to the probability of blocking in a circuit-switched network. Again, however, there are no specific considerations of capacity effects or restoration mechanisms.

Spragins [SSK86][Spr93] and Colbourn [Col91] seem to have been the first ones to question the relevance of studies purely based on questions of connectivity of simplified models. What Spragins advocates instead is an approach where the actual transmission, switching and routing structures and specific fault recovery mechanisms are taken into account. Spragins recognizes the fact that incorporating more details of the network structures into already very complex problems could make them impossible to solve but argues that it is preferable to reflect all aspects of the network structures *to some extent* rather than doing an exhaustive analysis of a very unrealistic network model. What Spragins refers to by saying "to some extent" is the idea of "most likely path to failure" presented in Section 3.3.2. In [Spr93], based on these ideas, Spragins presents the general network availability analysis methodology that consists of developing a system reliability block diagram and applying the simplification rules presented in Section 3.2.3 until the block diagram is reduced to a single block. He also details all the equations needed to determine the availability and failure rate at each intermediate stage up to the final one-block system representation.

3.5.2 Analysis of Path Availability with Protection Schemes

A topic that gets us closer to the one studied in this thesis is the availability analysis of paths in transport networks using a survivability mechanism of type *protection*, as defined in Section 2.4.2. This includes automatic protection switching, self-healing rings, dual-feeding, and other specific survivability measures all based on pre-determined rerouting of service paths in the case of failures. The first study of that type is the one presented by Spragins and discussed at the end of the previous section.

Another early study on this topic is that of To and Neusy in [ToN94]. As advocated by Spragins, instead of simply considering network connectability between the end-nodes of a service path to determine whether the path is considered failed or not, these authors consider the aspect of survivability with several network protection schemes and their particular characteristics to analyze the end-to-end unavailability of service paths. The protection schemes considered in that study are called *traffic diversification*, *restoration of traffic*, and *traffic survivability*. Traffic diversification refers to dividing service paths sent from one node to an adjacent node in two groups and routing them on physically separate routes so that a single cable cut may not affect all service paths at once. Traffic diversification has two variations: one in which channels are not protected against equipment failures and one in which channels are protected by 1:N protection switching between adjacent nodes. Restoration of traffic refers to a DCS-based mechanism in which failed service paths are automatically rerouted through pre-determined backup paths after a reconfiguration time of 2 minutes. Traffic survivability refers to protection schemes in which service paths are completely protected against the effects of single failures. Restoration times for survivability schemes are considered to be zero. The survivability schemes considered are: dual-feeding, 1+1 fibre protection, and 2- and 4-fibre BLSRs.

All the schemes considered use pre-determined rerouting options and thus the integrity of service paths is completely determined by the location of failures. The paper considers the contribution of single cable cuts, single transmission equipment failures, and nodal equipment failures.

For the failure rates, the restoration speeds of the different protection schemes, and the repair rates assumed, the paper concludes that physical diversity (with or without protection) provides by far the highest unavailability. DCS-based restoration provides a reduction by a factor 70 of the unavailability compared to physical diversity with protection (and by a factor 100 compared to the non-protected case.) Survivability schemes provide the lowest unavailability with a reduction of the unavailability by a factor 3 to 4 compared to restoration for dual-feeding, 4-fibre BLSR and 2-fibre BLSR (from most available to least available), while 1+1 fibre protection only provides a slight improvement compared to restoration.

The analysis of service path availability in ring-based networks has also been investigated in great detail by Grover in [Gro99a] and [Gro99b]. In these papers, Grover develops closed form models for the availability of paths through multiple rings of a ring-based network, including all details of ring size, path routing through rings, and a variety of dual-ring interconnect arrangements. The approach taken by Grover to evaluate unavailability is similar to the study of To and Neusy [ToN94] but the comparison between different survivability options is provided based on the availability vs. cost trade-off and not just on unavailability. The cost-availability trade-off to compare survivability options is an idea that will be used later in this thesis work. An interesting result of [Gro99a] that is important for this thesis work is that it was found that, given the statistical values that were assumed for that study, the unavailability of paths in a long-haul ring-based network is dominated by the effects of non-restorable intra-ring dual span failures. This finding strongly motivates the search of new options for high availability, especially linked to the problem of responding to dual span-failures. This thesis work focuses primarily on studying the resilience of adaptive mesh restoration to dual span-failures aims to show that they provide a great level of resilience to these failures. This and their great cost effectiveness makes adaptive mesh restoration a perfect candidate to replace rings at least in the long-haul but also maybe in metro networks.

3.5.3 Availability Analysis with Adaptive Restoration Schemes

In contrast, the availability analysis of mesh-restorable networks is not as amenable to a completely analytical approach and, most often, as been approached with significant simplifying approximations. The flexible nature of a mesh-restorable network—its routing adaptability and its extensive sharing of spare capacity—make it far less clear how to directly enumerate the outage-causing failure combinations.

Rowe [Row89] and Wilson [Wil98] both provide studies that in many ways are precursors to the present work. Rowe [Row89] shows the significant improvement brought to path availability from measures such as protection switching on transmission sections, random diversification of trunk routes, and cross-connect route-diversity switching. His study used simulation of a long random sequence of failure and re-routing actions to predict availability. That early work did not, however, address fully restorable mesh-restorable networks in the sense that we now consider with fully developed capacity design theory and detailed models for the span restoration re-routing process. Wilson [Wil98] studied the impact of ring and mesh restoration mechanisms on the service availability. His 'point-to-point' mesh network model is the closest to our present work but is limited to a special configuration model that does not reflect the complete diversity, topology dependence and adaptability of the mesh networks that we consider in this study.

Arijs et al. [AVD00a] compare ring and mesh architectures from a cost versus availability point-of-view. Their availability calculations for the mesh network are, however, limited to a dedicated mesh protection model because of the complexity that mesh restoration would bring to the analysis. This aspect is precisely one of the open issues of mesh availability analysis that we seek to address in a systematic way to determine availability in the presence of truly dynamic, adaptive, mesh restoration, as opposed to pre-determined protection arrangements.

3.5.4 Service Availability Analysis using Markov Modelling

Cankaya is another investigator that found it necessary to simplify the mesh restoration model to address the availability analysis with Markov modelling methods [CaN97]. Each Markov state has an associated number of outstanding failures and a functioning state. Solving the equations gives a network-wide availability based on the assumption that below a certain functioning level the network as a whole is consider failed. Specific detailed effects of the network topology, capacity and restoration routing behaviour are, however, lost in the assignment of overall global state transition state probabilities to apply the Markov model.

3.5.5 Service Availability in Networks with Priority Classes

Barezzani et al. [BPZ92] study the availability of a network with several traffic priority classes. Network availability is defined as the probability that the proportion of end-to-end connections in the up state at any moment in time is above a certain level for each of several priority classes. The study also investigates the improvements brought to the availability of high priority traffic by allowing lower priority traffic to be dropped for its restoration.

3.5.6 Service Availability in Networks with Multi-Layer Restoration

The same definition of network availability is used in [VDL95] in the context of multi-layer restoration, in which the restoration of the traffic starts at a given layer of the transport network when the lower levels have exhausted all their restoration capability. [VDL95] recognizes that a network-wide availability definition based on a "network fully up" or "network fully down" alternative does not convey the notion of how much traffic is actually lost. That is addressed with an Expected Loss of Traffic per year metric for the whole network. Both these papers focus on a path-restoration strategy that relies on the search for alternate end-to-end paths to re-route affected circuits in the event of single and dual span-failures

3.6 Analytical Analysis of Service Availability in Mesh-Restorable Networks

3.6.1 The Issue of Restoration Time

The aim in this section is to establish that the reconfiguration time to restore single failures is essentially irrelevant in the debate about availability. Historically, great importance has been given to the issue of restoration speed with ring advocates saying "50 ms is essential" and mesh advocates saying "anything under a second is absolutely suitable." As a consequence it has often been assumed that fast restoration is required to achieve high availability. In fact, the fallacy of this assumption can be demonstrated right away. Consider, as an example, a service path that undergoes

12 successful single-failure restoration events in five years (a relatively large number of such events) but then suffers an outage of 6 hours (half of a typical physical repair time) from a dual failure. We will calculate the unavailability of the service assuming (a) 50 ms, (b) 1 second single-failure reconfiguration times. In (a) we have:

$$U_{\text{path}} = \frac{12 \cdot 50\times10^{-3} + 6 \cdot 3600}{5 \cdot 8766 \cdot 3600} = 1.36896\times10^{-4}, \tag{3.42}$$

or

$$A_{\text{path}} = 0.99986310. \tag{3.43}$$

In case (b) we have:

$$U_{\text{path}} = \frac{12 \cdot 1 + 6 \cdot 3600}{5 \cdot 8766 \cdot 3600} = 1.36969\times10^{-4}, \tag{3.44}$$

or

$$A_{\text{path}} = 0.99986303. \tag{3.45}$$

The point is that despite the intuition that fast restoration times goes hand in hand with high availability, there is virtually no practical connection between these measures. What absolutely dominates the availability of a restorable network is whether or not the service is exposed to an unrestorable dual (or higher-order) failure. Almost any imaginable number of successful reconfigurations to single failures can occur without the availability moving numerically from virtually unity. This is true whether those reconfigurations take 50 ms or 1 s. But if a single outage is experienced that relates to the time required to complete one of the physical repairs required under a dual failure scenario, then the availability is dramatically impacted. This argument, simple but counterintuitive to many, was first clearly elucidated in the literature in [ClG02a] and is gaining increased understanding in the industry today.

We are not arguing here that restoration time is not an important consideration in its own right. Obviously, if all else is equal, faster restoration is better. But the point is that if one's customer wants to talk about service availability, it has essentially nothing to do with restoration speed if any reasonable automated scheme is employed within a corresponding spare capacity environment that is designed to ensure full single failure restorability. What such clients absolutely need to avoid are multi-hour outages that could cause them losses of millions of dollars. As a performance measure, availability does reflect the latter concern but is utterly unresponsive to any practical differences in restoration switching times.

We also do not mean to imply that mesh schemes are necessarily inferior to rings in restoration speed. Modern cross-connects have been designed with mesh restoration in mind as a key application and can reduce restoration times to little more than the network propagation times involved. In fact large BLSR rings are similarly limited by signalling propagation times. In practice, 250 ms is reportedly more typical of BLSRs [Ray01], not the "50 ms" of folklore automatically associated with any kind of ring (only 1+1 APS and UPSR rings routinely achieve 50 ms.) Moreover, any remaining doubts about mesh restoration speeds are addressed by the Distributed Pre-Planning (DPP) concept [Gro94]. Under DPP mesh restoration trials are exercised constantly in the background so that efficient and up to date pre-planned paths are already known in advance of a failure. With fast cross-connects, the real-time activation phase is then limited only by physical alarm propagation speeds.

3.6.2 Path Unavailability and the Concept of Equivalent Unavailability of Links

If we first imagine a mesh network over which a path p is provisioned over N spans, but with no restoration mechanism, based on the result of equation (3.21) in Section 3.2.3 we would fairly accurately estimate:

$$A_{\text{path}}(p) \approx 1 - \sum_{i \in S(p)} U^{\text{p}}_{\text{link}}(i), \tag{3.46}$$

where $U^{\text{p}}_{\text{link}}(i)$ is the *physical unavailability* of the i^{th} link in the path.

Therefore one way of thinking about the action of span restoration is that it is a transformer of *physical* span unavailability to a lower *equivalent* unavailability of links on the span. This viewpoint, proposed in [ClG02a], argues that from the standpoint of an end-to-end path, there are two equally acceptable ways in which a link along the path can be in "up" state: either it is physically working, or it is physically "down" but transparently replaced by a restoration path between its end nodes. Thus if we define the *equivalent unavailability* of a link $U^{*}_{\text{link}}(i)$ as follows:

$$U^{*}_{\text{link}}(i) = P\{\text{link } i \text{ down} \cap \text{link } i \text{ not restored}\}, \tag{3.47}$$

then the path availability has the same form as (3.46) but is based on $U^{*}_{\text{link}}(i)$, not $U^{\text{p}}_{\text{link}}(i)$. This line of reasoning reduces the problem of calculating path availability to determining the *equivalent unavailability* of links in a span-restorable network based on the capacity in the networks and the particulars of the restoration mechanism.

3.6.3 Determining the Equivalent Unavailability of Links

Let us consider the first three orders f of failure multiplicity corresponding to single, double and triple span failures. Our viewpoint is to determine the fraction of the physical unavailability of links on span i that is converted to equivalent unavailability. Clearly, with no restoration mechanism and no spare capacity, there is a one-to-one conversion: $U^{*}_{\text{link}}(i)$ is equal to $U^{\text{p}}_{\text{link}}(i)$. But with a restoration mechanism and a given distribution of spare capacity, the fraction of $U^{\text{p}}_{\text{link}}(i)$ that comes through to $U^{*}_{\text{link}}(i)$ depends on the failure states of other spans in the order-f failure scenario. Not all demands crossing span i may be restorable if there are coincident failures outstanding on other spans. It also depends in principle on the reconfiguration time for demands that are restorable but in practice, as seen in Section 3.6.1, this is a very small effect compared to outage due to multiple failure states that are not fully restorable. Conceptually, however, the restoration time for restorable demands may be longer for higher order failures. The action of the restoration mechanism within the spare capacity environment that survives the failure scenario can then be through of as providing a mapping from physical to equivalent link unavailability. The mapping takes two effects into account:

1. First, a multiple failure state may or may not support the feasibility of restoration for all links on span i. In the absence of a priority scheme, each link on span i will have to share this exposure to a capacity-related risk of incomplete restoration. To characterize this we define the multiple-failure restorability R_f. R_f is the fraction of the total failed working capacity that can be restored averaged over all f-order scenarios. In this thesis there will be very frequent mention of R_1 to refer to the restorability to single span-failures and R_2 to refer to the restorability to dual span-failures.
2. Second, we allow for a general reconfiguration outage-time for links that *are* restored. Although in practice we assume techniques that reconfigure in a few seconds at the very most (making this factor insignificant), the general model allows that restoration time could become significant depending on the failure order. Thus $U^{*}_{\text{link}}(i)$ is shielded from $U^{\text{p}}_{\text{link}}(i)$ on a span i through the following general conceptual mapping:

$$U^{*}_{\text{link}}(i) = U^{\text{p}}_{\text{link}}(i) \times \tag{3.48}$$

$$\sum_{f=1,2,3} p(\text{state } f-1) \cdot \left[\underbrace{\underbrace{(\text{restoration time exposure}) \cdot R_f}_{\substack{\text{time exposure} \\ \text{(restorable fraction)}}} + \underbrace{1 \cdot (1-R_f)}_{\substack{\text{capacity exposure} \\ \text{(unrestorable fraction)}}}}_{\text{exposure function}} \right]$$

In equation (3.48), $p(\text{state } f-1)$ represents the probability of being in failure state f given that span i is already failed (the probability of that is already accounted for with $U^{p}_{\text{link}}(i)$), the capacity exposure relates to the first point presented above and the time exposure relates to the second point. Details of the time exposure and capacity exposure functions are given in the following section.

3.6.4 Comparison of the Contribution of Various Failure Types to the Unavailability

Table 3-3 provides expressions for the terms of equation (3.48) for single, dual and triple span-failures. This discussion is adapted from work published in [ClG02a]. The domain of the exposure function is [0,1] and it expresses the extent to which links on a failed span are exposed to outage by virtue of incomplete restorability due to coincident failure states on other spans or, if restorable, the extent to which they are exposed to the restoration time. For example, if the network is designed for full single span failure restorability, then for any single span failure scenario the capacity exposure is zero and the time exposure is the restoration reconfiguration time. Table 3-3 gives the corresponding easily derived expressions for dual and triple span-failures. For lack of better data we assume a constant reconfiguration time in Table 3-3. The time exposure values are the ratio of expected restoration time to the expected average time in the corresponding failure state. For instance the average time in a dual span-failure state will be half the *MTTR* if the failures are independent, and a third of the *MTTR* in a triple span-failure state. Mathematical proof of this is given in Appendix C.

Table 3-3: Relative contribution of failure events on the unavailability of links

Network State f	Description	Probability given failure of span i	Time exposure	Capacity exposure	Contribution in example
1	Single span-failure, i	1	$\frac{\text{Av. RestTime}}{MTTR}$	0	1.38×10^{-8}
2	Dual span-failure, i and other span j	$(\lvert S\rvert - 1)\cdot U^{p}_{\text{link}}(i)$	$\frac{\text{Av. RestTime}}{0.5\times MTTR}\cdot R_2$	$1-R_2$	8.55×10^{-7}
3	Triple span-failure, i and other spans j, k	$\frac{(\lvert S\rvert-1)(\lvert S\rvert-2)}{2}\cdot(U^{p}_{\text{link}}(i))^2$	$\frac{\text{Av. RestTime}}{0.33\times MTTR}\cdot R_3$	$1-R_3$	4.62×10^{-9}

Let us now use equation (3.48) and Table 3-3 to support an argument that, in practice, dual span-failure scenarios will dominate the unavailability. This follows because by definition there is no capacity exposure in networks designed for full single span-failure restorability. Thus for single span-failures we have only a time exposure to the restoration process. The next most likely failure scenarios are dual span-failures. For dual span-failures we may expect a significant capacity exposure in a network designed only for single span-failure restorability. Similarly for triple span-failure scenarios, we expect that R_3 is likely to be much smaller than 1. As an example, consider a 20-span network for which the restoration time is 2 seconds, the physical span *MTTR* is 12 hours, $U^{p}_{\text{link}}(i)$ equals 3×10^{-4}, R_2 is 0.5, and R_3 is 0. Based on results that follow in Section 4.5, an R_2 of 0.5 is conservative and an R_3 of 0 is the worst case assumption for this argument. Under these assumptions, Table 3-3 shows the resulting contributions to $U^{*}_{\text{link}}(i)$, indicating that dual span-failures are by far the main factor to consider.

As explained previously, the link equivalent unavailability values $U^{*}_{\text{link}}(i)$ represent the probability that any individual link (a single capacity unit) on a span is in the "failed and non-restored" state at any point in time. As a consequence, there are a small number of situations where the unavailability of a *path* may be over-estimated by (3.46). In other words, (3.46) is strictly a somewhat pessimistic estimate of the actual path availability. This arises in the rather specific circumstances of an *(i,j)* dual span failure scenario for a path that: (i) crosses both spans *i* and *j* when, (ii) both of the spans *i* and *j* have less than complete individual restorability levels under the *(i,j)* failure scenario. Under these rather specific circumstances, the path availability will be slightly under-estimated because (3.46) will sum the individual link unavailabilities on each span above as independent contributors to the path outage whereas in reality they only contribute once to the unavailability arising as a single scenario. Figure 3-2 shows an example of how a span failure could be counted twice in the unavailability of a given service path.

It is important to realize that this will not be a large numerical effect, however, and that any bias due to it is a pessimistic one. Of all dual failure combinations that only those where both failures fall on the same path can possibly have this effect, *and only if* both of the spans in those scenarios individually has less that a full restoration level. If either individually has full restoration, then $U^{*}_{\text{link}}(i)$ is zero for that span in the specific scenario, and the overestimation effect does not

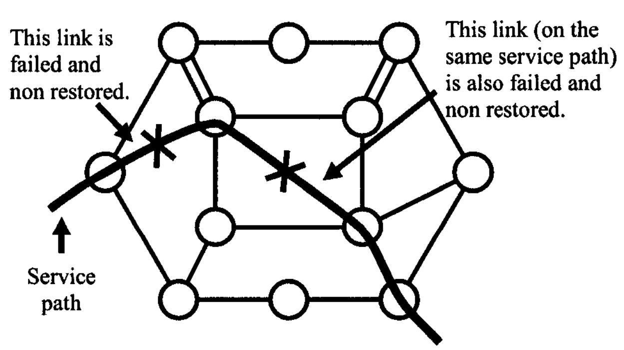

Figure 3-2 Example of how a dual span-failure can be counted twice in the unavailability of a service path

occur. For dual span failures several spans apart on the same path this is likely to be the case due to the spatial separation of their individual restoration patterns. Where the over-estimation will be more likely is for adjacent-span failures as shown in Figure 3-2.

A final point is to also observe that the overestimation is based on assuming a random allocation of restored links on each span to paths transmitting the span. In practice it could be advantageous to make a coordinated allocation of restored capacity on each span to priority-path basis. In that case a certain number of priority paths could effectively get a $U^{*}_{\text{link}}(i)$ of zero almost all the time. In fact any time the restorability $R_2(i,j)$ of a particular dual span-failure (i,j) is greater than zero we could think of the top-most priority path being preferentially allocated the restored links, in which case its path availability would also be estimated by (3.46), but with a $U^{*}_{\text{link}}(i)$ of zero any time $R_2(i,j)$ is not zero. (Methods to determine $R_2(i,j)$ and formal expression of that function are given at the beginning of Section 3.) This is an observation to keep in mind when we see in the results just how extremely rare it is ever to see an $R_2(i,j)$ of zero, and it is an important point to which we will return later in the thesis.

3.6.5 Exact and Approximate Models of Path Availability

The exact availability model for path availability is based on the cut sets approach [BiA92]:

$$A_{\text{path}}(p) = 1 - U_{\text{path}}(p) = 1 - \sum_{f \in F} P\{f\} \cdot (1 - R(f, p)) . \tag{3.49}$$

In (3.49), f denotes an element of the set F of all possible failure combinations (and not a failure order as in (3.48).) $P\{f\}$ is the probability of failure combination f and $R(f, p)$ is the restorability of path p to failure f. Following the cut set method, we are subtracting from the availability of path p the probability of each failure f in the set of all possible failures F if path p is not restorable to that failure, i.e. if $R(f, p)$ is equal to zero. Note that in the case of a non-deterministic restoration process (if the restorability of a path p to a failure f is not necessarily 0 or 1, but can be random with a real valued probability $R(f, p)$ of happening as it would be the case when multiple equal-class service paths are competing for an insufficient number of restoration paths), equation (3.49) applies too. In that case, we are subtracting the probability of each failure weighted by the probability of path p not to be restorable to that failure from the availability of the path.

As explained in Section 3.3.2, a strategy to reduce the analysis work of the cut set method is to limit ourselves to considering only failure combinations up to a certain order. Based on the results of Section 3.6.3, and assuming a network design that is fully restorable to any single span-failure, a very good approximation of path availability can be obtained by limiting the set F to all dual-span failures. In that case, the expression becomes:

$$A_{\text{path}}(p) = 1 - \sum_{i \in S} \sum_{\substack{j \in S \\ j \neq i}} P\{\text{dual span failure } (i,j)\} \cdot (1 - R_2(i,j,p)) \tag{3.50}$$

where $R_2(i,j,p)$ is 1 or 0 depending on whether path p is restorable to dual span-failure (i,j) in a deterministic restoration case, or it is the probability of path p to be restorable to dual span-failure (i,j) in a non-deterministic case.

In (3.50), the probability of dual span-failure (i,j) can be expressed in terms of the unavailability of spans i and j, respectively U_i and U_j, yielding:

$$A_{\text{path}}(p) = 1 - \sum_{i \in S} \sum_{\substack{j \in S \\ j \neq i}} U_i \cdot U_j \cdot (1 - R_2(i,j,p)) \tag{3.51}$$

A slight variation of equation (3.51), in the context of multi-class service paths, consists of considering the service class c of path p (more on service classes will be presented in Section 6.2.3 and mainly in Chapter 8) and expressing the availability of path p as:

$$A_{\text{path}}(p) = 1 - \sum_{i \in S} \sum_{\substack{j \in S \\ j \neq i}} U_i \cdot U_j \cdot (1 - R_2^c(i,j)) \tag{3.52}$$

where $R_2^c(i,j)$ denotes the restorability to dual span-failure (i,j) of affected paths in class c.

In the context of a network with only one class of service paths, equation (3.50) does not have to be used and can be replaced by:

$$A_{\text{path}}(p) = 1 - \sum_{i \in S} \sum_{\substack{j \in S \\ j \neq i}} U_i \cdot U_j \cdot (1 - R_2(i,j)) \tag{3.53}$$

where $R_2(i,j)$ is the restorability of service paths to the dual span failure (i,j).

Two even simpler availability models can be used by using the span-average restorability value $R_2(i)$, which represent the average restorability of paths routed through span i over all dual span-failures involving span i, or using the network-wide average R_2 representing the average restorability of affected-paths over all possible dual span-failures. The resulting expressions for the availability of a path p are:

$$A_{\text{path}}(p) = 1 - \sum_{i \in \text{path } p} U_i \cdot U_{\text{span}} \cdot (1 - R_2(i)) \tag{3.54}$$

and:

$$A_{\text{path}}(p) = 1 - N \cdot U_{\text{span}}^2 \cdot (1 - R_2) \tag{3.55}$$

where U_{span} denotes the average span unavailability over all spans of the network, and N is the number of spans crossed by path p.

While equation (3.54) is expected to remain fairly accurate (although the link between specific $R_2(i,j)$ and U_j values is lost in the use of an average U_{span}), equation (3.55) is likely to provide much less accurate results since the information about the specific route taken by path p is lost by using simply N to characterize the path.

3.7 Summary

This chapter has presented the important definitions and concepts of availability analysis. As it was explained, this thesis will be focusing on developing methods for the availability analysis of end-to-end service paths in span-restorable mesh networks. Special emphasis will be put on methods that take all the details of mesh networking into consideration: specific routing of each demand unit, knowledge of capacity allocation and use, and detailed simulation of restoration mechanisms. The approach taken will therefore not limit itself to considering problems of network connectivity, which, as explained, have been already intensively studied under the general topic of "network reliability."

One of the important results of this chapter has been to show the direct relation between availability and the restorability to dual-failures, and the relatively low importance of the speed of restoration to single-failures. Based on this result, the following chapter will concentrate on the problem of determining the restorability of a span-restorable network to dual-failures.

4. Computational Analysis of the Dual Span-Failure Restorability

The material in Sections 4.1 and 4.2 is an adaptation and extension of that first presented in [ClG02a].

4.1 Introduction

A dual span-failure is denoted *(i,j)*, naming the two spans involved and the order of failure. In case of a dual span-failure in a network that is "only" designed to be restorable to single failures, there may or may not be enough feasible restoration paths to restore all the affected working paths on the failed span. The restorability of a dual span-failure *(i,j)*, already referred to many times in the previous section, can be expressed formally as follows:

$$R_2(i,j) = 1 - \frac{N(i,j)}{w_i + w_j}, \; w_i + w_j \neq 0 \tag{4.1}$$

where $N(i,j)$ is the number of non-restorable working capacity units over both spans in the event of dual span-failure *(i,j)*, and w_i and w_j are the number of working capacity units on spans *i* and *j*.

The dual span-failure restorability can also be defined for a given span *i* as the average of $R_2(i,j)$ over all ordered dual-span failures involving span *i:*

$$R_2(i) = \frac{1}{2 \cdot (|S| - 1)} \cdot \sum_{\substack{j \in S, j \neq i \\ w_i + w_j \neq 0}} (R_2(i,j) + R_2(j,i)) \tag{4.2}$$

However, a preferred definition for $R_2(i)$ will be the average of $R_2(i,j)$ over all ordered dual-span failures involving span *i*, weighted by the total working capacity to be restored in each combination *(i,j)*:

$$R_2(i) = \frac{1}{2 \cdot \sum\limits_{\substack{j \in S, j \neq i \\ w_i + w_j \neq 0}} (w_i + w_j)} \cdot \sum_{\substack{j \in S, j \neq i \\ w_i + w_j \neq 0}} (w_i + w_j) \cdot (R_2(i,j) + R_2(j,i)) \tag{4.3}$$

which can also be expressed as:

$$R_2(i) = 1 - \frac{\sum\limits_{j \in S, i \neq j} (N(i,j) + N(j,i))}{2 \cdot \sum\limits_{j \in S, i \neq j} (w_i + w_j)} \tag{4.4}$$

It can easily be shown that the previous expression can be simplified to give:

$$R_2(i) = 1 - \frac{\sum\limits_{j \in S, i \neq j} (N(i,j) + N(j,i))}{2 \cdot (|S| - 2) \cdot w_i + 2 \cdot \sum\limits_{j \in S} w_j} \tag{4.5}$$

Finally, the dual span-failure restorability can be defined as a network-wide metric by averaging $R_2(i,j)$ over all ordered dual span-failures:

$$R_2 = \frac{1}{|S| \cdot (|S| - 1)} \cdot \sum_{i \in S} \sum_{\substack{j \in S \\ j \neq i \\ w_i + w_j \neq 0}} R_2(i,j) \tag{4.6}$$

or as the weighted average:

$$R_2 = \frac{1}{\sum\limits_{i \in S} \sum\limits_{\substack{j \in S \\ j \neq i \\ w_i + w_j \neq 0}} (w_i + w_j)} \cdot \sum_{i \in S} \sum_{\substack{j \in S \\ j \neq i \\ w_i + w_j \neq 0}} (w_i + w_j) \cdot R_2(i,j) \tag{4.7}$$

which can also be expressed as:

$$R_2 = 1 - \frac{\sum\sum N(i,j)}{\sum (w_i + w_j)} = 1 - \frac{\sum\sum N(i,j)}{2 \cdot (|S| - 1) \cdot \sum\limits_{i \in S} w_i} \tag{4.8}$$

where the double summation is over the same set as in (4.7).

The goal of this chapter is to present a general method for the determination of dual span-failure restorability in span-restorable mesh networks. First, we identify a set of canonical dual span-failure types. This exercise underlines the difficulty of developing closed-form models for availability in networks using adaptive restoration. Then, we present a reference optimization model that can be used to determine the maximum achievable dual-failure restorability. This is followed by the presentation of three progressively more adaptive models of the span-restoration mechanism. Finally, the last section of this chapter presents experimental results of the dual-failure restorability. Through these experiments we investigate the relative performance of the three restoration models and we identify factors that mostly affect the dual-failure restorability and therefore the availability.

4.2 Dual Span-Failure Types

When considering a span-restorable mesh network there are four logical categories that describe dual-failure combinations:

i. Failures that are spatially independent (the restoration route-sets are disjoint).

ii. Failures with individual restoration route-sets that "content" spatially.

iii. Cases where the second failure damages one or more restoration paths of the first failure.

iv. Cases where the two failures isolate a degree-2 node or effect a cut of the graph.

Figure 4-1 (i) to (iv) illustrate the four cases listed above. In these illustrations, only the network of spare links is shown and working capacities are indicated only for the failed spans. The bold lines show the restoration paths formed to restore the indicated number of working links. Where relevant, the first failure to occur is associated with w_1.

Figure 4-1 (i) shows a case of no spatial interaction between the individual restoration route-sets. Since both failures are fully restorable, $R_2(i,j)$ for the scenario is equal to 1. Whether the restoration route sets interact or not depends on the re-routing mechanism and on the working and spare capacitation of the graph.

Figure 4-1 (ii) portrays a case in which there is spatial interaction between the routes of the respective restoration path-sets. Depending on the spare capacities in Figure 4-1 (ii), outage may or

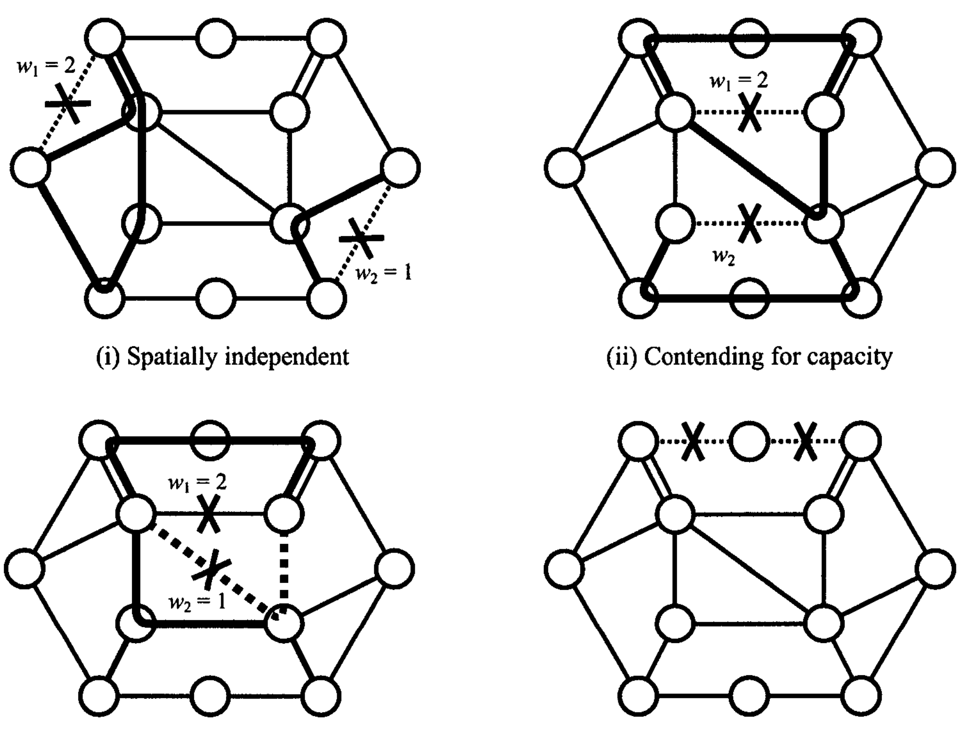

(i) Spatially independent (ii) Contending for capacity

(iii) Restoration path hit (iv) Degree-2 cut

Figure 4-1 Types of dual span-failures

may not result from contention for available spare capacity between the two failures. As drawn, after restoration of w_1, only one restoration path is feasible for the second failed span. Consequently, if w_2 is greater than 1 the second failed span will not be fully restored. If w_2 equals 2 in Figure 4-1 (ii) we would say that for the failure of the particular *ordered pair* of spans *(i,j)* the restorability on the second failed span is 1/2 and over the two failures together ($R_2(i,j)$ for the scenario) the restorability is 3/4. Note in general that if the order of the failed spans had been different, the restorability for each span could differ as well.

Figure 4-1 (iii) is a case where the second failure damages restoration paths of the first failed span. In this case the number of restored links for the second failure depends on the remaining spare capacity after the first failure and depends on the "secondary" adaptability of the restoration process, i.e. whether the restoration path is severed by a second failure. The result also depends on whether sufficient spare capacity remains for the mechanism to repair the damage to its first path-set by an updated restoration response.

Finally, Figure 4-1 (iv) portrays the case where a degree-2 node is isolated by the failure of both its adjacent spans, creating an unrestorable situation: here $R_2(i,j)$ equals 0. More generally, if there is any cut of the network graph that contains only two edges, there are two ordered pairs that will disconnect the graph. In this case, the amount of spare capacity and the adaptability of the restoration mechanism have no influence on the restorability and $R_2(i,j)$ always equals zero.

4.3 Optimal Two-Commodity Max-Flow Reference Model

As we started to explain in the previous section, in the case of a dual span-failure the number of non-restorable capacity units $N(i,j)$ depends on the location of both failures (i.e. on the type of dual-failure as described above). In the case of failures where multiple affected working capacity units are contending for spare capacity, $N(i,j)$ also depends on the adaptability of the restoration mechanism. However, even with the most adaptive and efficient mechanism, there is a theoretical minimum number of capacity units which cannot be restored due to the inherent limitation on the maximum simultaneously achievable flow between the end node of both failures. Finding this maximum flow is referred to here as the *two-commodity maximum-flow* (TCMF) problem. We do not know of other works or contexts when TCMF is specifically of interest. Let us first express this problem as an ILP formulation using the same approach as for the formulations presented earlier in this thesis, that is, as a problem of assigning flows to a set of eligible restoration routes:

$$\textbf{TCMF-Path: } \text{Minimize } N(i,j) \tag{4.9}$$

Subject to:

$$N(i,j) = w_i + w_j - \left(\sum_{p \in P_i} f_i^p + \sum_{p \in P_i} f_j^p \right) \quad (4.10)$$

$$f_i^p \le C_\infty \cdot (1 - \delta_{i,j}^p) \quad (4.11)$$

$$f_j^p \le C_\infty \cdot (1 - \delta_{j,i}^p) \quad (4.12)$$

$$w_i \ge \sum_{p \in P_i} f_i^p \quad (4.13)$$

$$w_j \ge \sum_{p \in P_i} f_j^p \quad (4.14)$$

$$\sum_{p \in P_i} \delta_{i,k}^p \cdot f_i^p + \sum_{p \in P_i} \delta_{j,k}^p \cdot f_j^p \le s_k,\ \forall k \in S,\ i \ne k,\ j \ne k \quad (4.15)$$

Constraint (4.10) relates the number of non-restored capacity units $N(i,j)$ to the total working capacity to be restored and to the total restoration flows for both failed spans. Constraints (4.11) and (4.12) use an arbitrarily high capacity constant, C_∞, to ensure that no restoration flow for any of the two failed spans is assigned to a route crossing the other failed span. Constraints (4.13) and (4.14) ensure that the total restoration flow assigned for each failed span does not exceed the total working capacity that needs to be restored on that span. This prevents a feasible excess of restoration flow for one of the two failed spans to be used to reduce $N(i,j)$ in the case where the other failed span is not fully restorable. It also prevents $N(i,j)$ from taking negative values. Finally, constraint set (4.15) forces total restoration flows on each span other than the two failed spans not to exceed the amount of spare capacity allocated on it.

To obtain the actual maximum flow using the previous formulation, we need to consider all existing restoration routes for each span, with no limitation on the number of hops or other limitation on the total number of routes provided to the ILP solver. For large or even medium networks, this becomes difficult given the astronomical number of possible restoration routes, therefore the set of eligible routes has to be limited for the solver to be able to find a solution. The problem is then that the solution found is not guaranteed to be the optimal one since a better solution may be found using longer restoration routes. In fact, the experiments that were conducted to produce the results presented in Section 4.5 showed that for some of the largest networks, the restoration flow found by experiments based on a computational analysis of the restoration exceeded the one found by the TCMF-Path formulation using a restriction on the number of eligible routes. Indeed, as it will be seen later, the restoration models do not assume any maximum restoration route length limitation, which explains this apparent contradiction. Although the TCMF-Path formulation remains useful in the case where there is, in practice, a limitation on the maximum length of restoration routes or any other constraint on what routes can be used for restoration, an alternative approach has to be used to be able to solve the true TCMF problem for medium or large networks.

This alternative approach is based on the *flow conservation* principle. In contrast with the previous approach, we are not assigning flows to end-to-end routes but to spans. For this formulation, the following parameters are defined:

- N Set of nodes in the topology
- $A(n)$ Set of nodes adjacent to node n, $\forall n \in N$
- $o(k)$ Origin node of span k, $\forall k \in \mathbf{S}$
- $d(k)$ Destination node of span k, $\forall k \in \mathbf{S}$

as well as the following new variables:

- $x^1_{m,n}$ restoration flow for the first failed span (i) from node m to node n
- $x^2_{m,n}$ restoration flow for the second failed span (j) from node m to node n

TCMF-Flow: Minimize: $N(i,j)$ (4.16)

Subject to:

$$N(i,j) = w_i + w_j - \left(\sum_{m \in A(o(i))} x^1_{o(i),m} + \sum_{m \in A(o(j))} x^2_{o(j),m} \right) \tag{4.17}$$

$$\sum_{m \in A(n)} x^1_{m,n} = \sum_{m \in A(n)} x^1_{n,m}, \forall n \in N, n \neq o(i), n \neq d(i) \tag{4.18}$$

$$\sum_{m \in A(n)} x^2_{m,n} = \sum_{m \in A(n)} x^2_{n,m}, \forall n \in N, n \neq o(j), n \neq d(j) \tag{4.19}$$

$$x^2_{m,n} = 0, \forall (m,n) \in \{(o(j), d(j)), (d(j), o(j))\} \tag{4.20}$$

$$x^2_{m,n} = 0, \forall (m,n) \in \{(o(i), d(i)), (d(i), o(i))\} \tag{4.21}$$

$$x^1_{m,o(i)} = 0, \forall m \in A(o(i)), m \neq d(i) \tag{4.22}$$

$$x^2_{m,o(j)} = 0, \forall m \in A(o(j)), m \neq d(j) \tag{4.23}$$

$$\sum_{m \in A(o(i)), m \neq d(i)} x^1_{o(i),m} \leq w_i \tag{4.24}$$

$$\sum_{m \in A(o(j)),\, m \neq d(j)} x^2_{o(j),\, m} \le w_j \tag{4.25}$$

$$x^1_{o(k),\, d(k)} + x^1_{d(k),\, o(k)} + x^2_{o(k),\, d(k)} + x^2_{d(k),\, o(k)} \le s_k,\ \forall k \in \boldsymbol{S},\ k \neq i,\ k \neq j \tag{4.26}$$

Constraint (4.17), as with (4.10), relates the number of non-restored capacity units $N(i,j)$ to the decision variables and input parameters. Constraints (4.18) and (4.19) enforce the flow conservation constraint for the restoration flows of each failed span at each node except at the end nodes of the failed spans. Constraints (4.20) and (4.21) ensure that no restoration flow for one of the failed spans can be assigned to the other failed span. Constraints (4.22) and (4.23) prevent restoration flow from coming back to the origin node of a failed span. Not having these two constraints would allow the solver to assign extra restoration flow leaving the origin node of a failed span and assign some restoration flow in the other direction to respect the flow conservation constraint at the first node en route, and therefore wrongfully decrease $N(i,j)$ through equation (4.17). Constraints (4.24) and (4.25) ensure that no more restoration flow can be credited for the restoration of each span than what is actually needed. Finally, constraint set (4.26) guarantees that the total restoration flow assigned to each span does not exceed the total spare capacity allocated on that span.

The AMPL model that was developed to solve the TCMF-Flow formulation is detailed in Section D.2 of Appendix D.

The objective of adaptive restoration mechanisms is to get as close as possible to TCMF. Adaptive restoration, however, sometimes has some additional constraints which might prevent them from achieving the maximum flow. For example, it can be a requirement that restoration paths of the first failure that are intact after the occurrence of the second failure be not touched, whereas TCMF does not consider any constraint on the routing of any of the restoration paths. The following section will present three different models of the restoration mechanism corresponding to different assumptions regarding such requirements. In Section 4.5, TCMF will be used as a reference to evaluate the performances of the various models of the restoration mechanism presented in the following section.

4.4 Various Models of the Restoration Mechanism

As conveyed by Figure 4-1 of Section 4.2, it is not straightforward to predict mesh availability analytically. The $R_2(i,j)$ of each failure scenario depends in detail on the specifics of the (i,j) pair, the failure sequence, the exact working and spare capacities, the graph topology, and the assumed restoration dynamics. In fact, the single most important factor that makes predicting mesh availability analytically more difficult than for protection mechanisms like rings is that the type of mesh restoration we are considering here is adaptive to changes. In other words, the way in which paths are rerouted in response to a failure is not pre-determined. As a result, in order to predict availability without making approximations that would over-simplify the problem, we need to consider all of the details of the restoration paths routing mechanisms. This way we can explicitly determine the outcome of all dual-failure restoration scenarios by computer-based experiments.

For the case of span-restoration, three technical models for the restoration process were developed to see how varying levels of adaptability would affect R_2. Each of the models corresponds to different technical options for engineering the restoration mechanism. The restoration routing experiments are based on *k-shortest paths* (ksp) routing behaviour for the basic single-failure response model [DGM94][MaG94]. Using *ksp* means that each restoration path set is formed by first taking all paths feasible on the shortest route, followed by all paths feasible on the next shortest route not re-using any spare capacity already seized on the shortest route, and so on, until either all required paths are found, or no more *can* be found. This is known to be extremely close to maximum flow in typical transport networks [DGM94] and can be computed in $O(n \log n)$ time [MaG94]. The ksp algorithm is also an accurate functional model for the self-organized restoration path-sets formed by the SHN™ protocol [Gro97]. The following three sub-sections look at three levels of adaptability mentioned above that can be modelled with ksp to determine R_2 by exhaustion of all dual span-failure experimental trials. These three progressively more adaptive restoration behaviours will generally be referred to as: Static (presented in Section 4.4.1), partly-adaptive behaviour (presented in Section 4.4.2) and fully-adaptive behaviour (presented in Section 4.4.3.)

4.4.1 Static Pre-plans or Protection (Model 1)

The first model for restoration behaviour is meant to represent restoration that is wholly based on centrally computed single-failure pre-plans. In this model, restoration of each span failure follows a pre-determined plan, trying to restore both spans as if each were an isolated failure. If not enough spare capacity exists to support the superposition of both static pre-plans, restoration paths of the second failure are suppressed to conform to what is feasible within the spare capacity remaining after the first failure.

Figure 4-2 shows an example of the response of the static restoration behaviour to a dual span-failure. On the figure, the first failure is represented by the explosion symbol with the number one. That failure is assumed to result in full restoration of the affected working units (three in this case) on that span using three restoration paths (each on a different route). The second failure is repre-

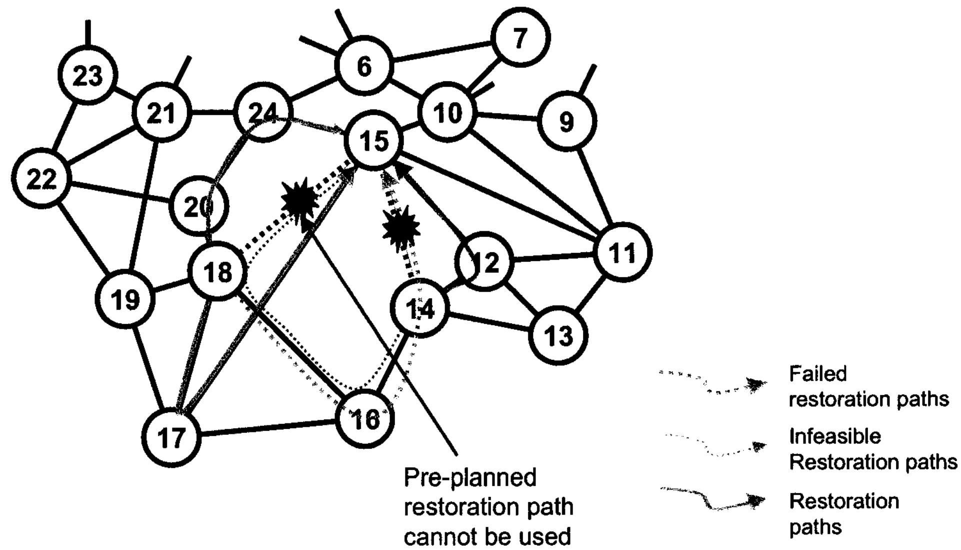

Figure 4-2 Response of the static restoration behaviour to a dual span-failure

sented by the explosion symbol with the number two. In this example, two incidents happen that are specific to the static behaviour. The first incident is the loss of a restoration path due to the fact that the second failure hits that path (here, path 18-16-14-15 is lost). With the static restoration behaviour nothing is done to repair or replace that failed restoration path. The second incident is the infeasibility of one of the pre-planned restoration paths for the second failed span. In this case, the infeasibility of that restoration path is due to the fact that the path crosses the first failed span, but another possible reason for the infeasibility of a pre-planned restoration path could be that the required spare capacity may not be present everywhere along the path (some of the required spare capacity could indeed already be used for the restoration of the first failed span).

To determine the dual span-failure restorability $R_2(i,j)$ of a mesh network using static restoration pre-plans, the algorithm described in the following lines was developed.

The algorithm uses the following input parameters:

$\boldsymbol{S}$ Set of spans in the network

$\boldsymbol{P}_i$ Set of restoration routes for the restoration of span i, $\forall i \in \boldsymbol{S}$

Δ_i^p Vector of size $|\boldsymbol{S}|$ describing p^{th} route for the restoration of span i, $\forall i \in \boldsymbol{S}, \forall p \in \boldsymbol{P}_i$ (the k^{th} component of that vector is $\delta_{i,k}^p$)

f_i^p Flow on p^{th} route for the restoration of span i, $\forall i \in \boldsymbol{S}, \forall p \in \boldsymbol{P}_i$

S Vector of size $|\boldsymbol{S}|$ describing the network's spare capacity allocation (the k^{th} component of that vector is s_k)

It is important to note that, unlike in the case of capacity design as presented in Section 2.5.3, the f_i^p values are not variables here but parameters, which clearly indicates that we are dealing with a pre-planned scheme, where the assignment of flows to the different restoration routes upon any span-failure is something known in advance.

We now describe the algorithm for determining $R_2(i,j)$ for a given failure combination (i,j). From the input information we first calculate the available spare capacity $S^*_{i,j}$ during failure of span i and j:

$$S^*_{i,j} = S - (s_i \cdot e_i + s_j \cdot e_j) \tag{4.27}$$

where e_i and e_j are the i^{th} and j^{th} unit vectors (the k^{th} unit vector contains only zeros and a one in the k^{th} position). Basically $S^*(i,j)$ is identical to S with 0 in place of s_i and s_j.

From the input information we calculate the vector of required spare capacity for the simultaneous restoration of spans i and j:

$$S^r_{i,j} = \sum_{p \in P_i} \Delta_i^p \cdot f_i^p + \sum_{p \in P_i} \Delta_j^p \cdot f_j^p \tag{4.28}$$

We then follow the algorithm described on Figure 4-3. The algorithm can be viewed as having two stages. In the first stage (where $x = j$), the number of non-restorable working capacity units on the second failed span (j) is determined. In this stage, the spare capacity of each span k is checked to see if the available spare capacity $S^*_{i,j}(k)$ is less or equal to the required spare capacity $S^r_{i,j}(k)$ on that span. Every time a span with missing spare capacity is found, some backup paths for the second failed span are suppressed if they cross that span until either the available spare capacity becomes sufficient to support all remaining paths or no paths crossing that span can be suppressed anymore. When all spans have been checked, we proceed the second stage ($x = i$) in which will be determined the number of backup paths of the first failed span that have been hit by the second failure. In this stage, these backup paths will be suppressed and counted in $N(i,j)$. When all the spans have been checked, the dual span-failure restorability can be calculated using equation (4.1).

4.4.2 *First-Event Adaptive Behaviour (Model 2)*

The second model of restoration dynamics assumes that after a first failed span has been restored (not repaired), the restoration mechanism of any second failure is aware of, and adaptive to, the changes in available spare capacity resulting from the first failure. Moreover, if any restoration path for the first failed span is routed over the second failed span, the restoration mechanism will combine the requirements for the second span failure with the failed restoration paths of the first failure. However, new restoration paths are not sought between the end-nodes of the first failed

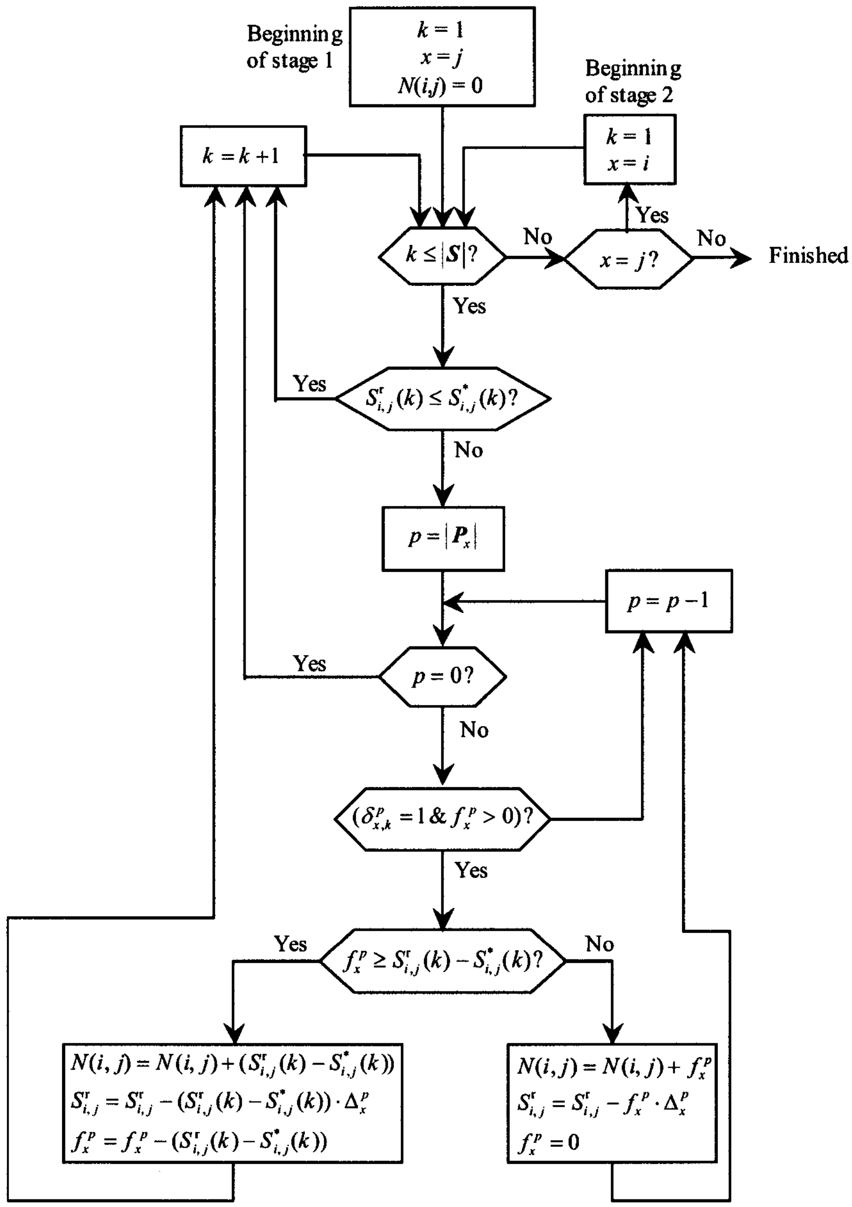

Figure 4-3 Algorithm for analyzing R_2 with the static behaviour

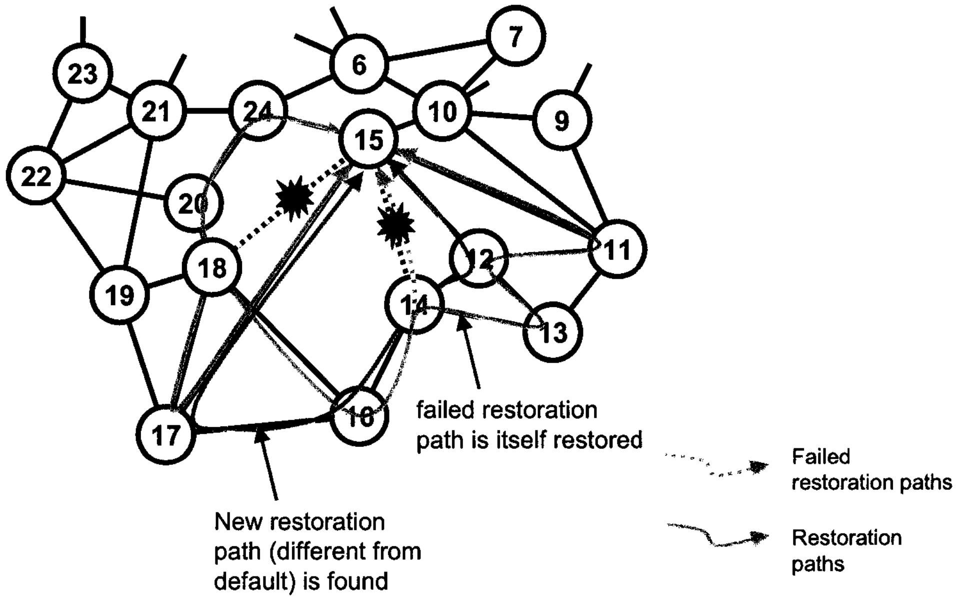

Figure 4-4 Response of the partly-adaptive restoration behaviour to a dual span-failure

span. In other words, restoration paths affected by a second failure are referred into the second failure's w_i quantity. This second model is based on the techniques of a distributed self-updating and self-organizing restoration protocol [Gro97]. Such a protocol immediately gives "working" status to any spare link it uses in a restoration path. This inherently updates itself should it be triggered to act again to protect either new capacity or prior restoration paths in the event of a second failure. The awareness of the updated spare capacity environment following the first failure is implicit. For present work, the idea is that the complete range of realistically expected restoration responses is encompassed between the static restoration pre-plan model of Section 4.4.1 and the fully adaptive behaviour presented in the following section.

Figure 4-4 shows examples of the response of the static restoration behaviour to a dual span-failure. Two major differences with the previous restoration behaviour can be noticed in this example. First, the restoration path hit by the second failure is itself restored locally between the end-nodes of the second failed span. Second, the infeasible restoration path of the previous example is replaced by a new restoration path dynamically found within the available spare capacity (not already used for restoration of the first failure) at the time of the failure.

To determine the dual span-failure restorability $R_2(i,j)$ of a mesh network using a partly-adaptive restoration behaviour, we developed the algorithm described on Figure 4-5. In this algorithm, w_j^{eq} refers to the total number of spare links on span j used for the restoration of the first

1. <u>Initialization and failure of span i</u>:

   $S^* = S$

   $w_j^{\mathrm{eq}} = 0$

   $S^*(i) = 0$

2. <u>Restoration path search for span i</u>:

   $\boldsymbol{P}_i = \mathrm{ksp}(S^*, i)$

3. <u>Restoration of span i</u>:

   $w_i^{\mathrm{nr}} = w_i$

   $p = 1$

   while $(w_i^{\mathrm{nr}} > 0 \ \& \ f_i^{|\boldsymbol{P}_i|} > 0)$

       if $(f_i^p = 0)\ p = p + 1$

       $S^* = S^* - \Delta_i^p$

       $f_i^p = f_i^p - 1$

       $w_i^{\mathrm{nr}} = w_i^{\mathrm{nr}} - 1$

       if $(\delta_{i,j}^p = 1)\ w_j^{\mathrm{eq}} = w_j^{\mathrm{eq}} + 1$

4. <u>Failure of span j</u>:

   $S^*(j) = 0$

5. <u>Restoration path search for span j</u>:

   $\boldsymbol{P}_j = \mathrm{ksp}(S^*, j)$

6. <u>Restoration of span j</u>:

   $w_j^{\mathrm{nr}} = w_j + w_j^{\mathrm{eq}}$

   $p = 1$

   while $(w_j^{\mathrm{nr}} > 0 \ \& \ f_j^{|\boldsymbol{P}_j|} > 0)$

       if $(f_j^p = 0)\ p = p + 1$

       $S^* = S^* - \Delta_j^p$

       $f_j^p = f_j^p - 1$

       $w_j^{\mathrm{nr}} = w_j^{\mathrm{nr}} - 1$

7. <u>Calculation of total non-restored working units</u>:

   $N(i, j) = w_i^{\mathrm{nr}} + w_j^{\mathrm{nr}}$

Figure 4-5 Algorithm for analyzing R_2 with the partly-adaptive behaviour failed span *i*, which as explained earlier take working status and therefore become equivalent to additional working links to restore on span *j* in case of failure of that span. In stage 1, the available spare capacity vector is initialized to the spare capacity allocated in the design and the available spare capacity on span *i* is set to zero to signify the failure of that span. Also, the equivalent working capacity w_j^{eq} on span *j* is considered to be zero initially. In stage 2, the restoration path search is performed using the ksp algorithm [Gro03] and the found paths information is returned in the form of a set $\boldsymbol{P}_i$ of restoration routes. The set $\boldsymbol{P}_i$ includes the Δ_i^p parameters defined in the previous section that detail the list of spans in each route *p* in $\boldsymbol{P}_i$, and the depth parameters f_i^p, which indicate the number of paths that are feasible on each route *p* in $\boldsymbol{P}_i$ under ksp-type restoration (f_i^1 is routed on first route in $\boldsymbol{P}_i$, then f_i^2 is routed on second route in $\boldsymbol{P}_i$, etc.)

In stage 3, restoration of span *i* is performed, using the restoration route information found in the previous stage. In this stage, w_i^{nr} represent the number of working links that are not yet restored on span *i* and it is first initialized with the value w_i. Each time a new service path is added, if it crosses span *j* ($\delta_{i,j}^p$ is in that case equal to 1) the value of w_j^{eq} is increased by 1. Failure, restoration path search and restoration of span *j* follows the same model as for span *i*. Note that the w_j^{nr} value at the beginning of stage 6 is initialized with the sum of w_j and w_j^{eq} showing that restoration paths for span *i* that cross span *j* are included in the restoration effort on span *j*. In the final stage, the total number of non-restored working units is obtained by summing the total non-restored units on span *i* and on span *j*.

4.4.3 Fully Adaptive Behaviour (Model 3)

The third model of restoration dynamics is of a completely adaptive restoration mechanism including both spare-capacity awareness following a first failure and re-restoration efforts for the first span from its original end-nodes, for any damage to previously deployed restoration paths. The restoration mechanism will try to find new restoration paths between the end-nodes of the first failed span when a second failure severs any of its initial restoration paths. This includes a release of surviving spare capacity on restoration paths of the first failure that were severed by the second and repetition of the first span's restoration effort for the newly outstanding un-restored capacity, but in recognition of the spare capacity now withdrawn by the second failure.

Figure 4-6 shows an example of the response of the static restoration behaviour to a dual span failure. The main difference with the partly adaptive behaviour is that here, instead of restoring failed restoration paths between the end-nodes of the second failed span, the restoration mechanism frees the surviving spare capacity along the failed restoration path and searches for new res-

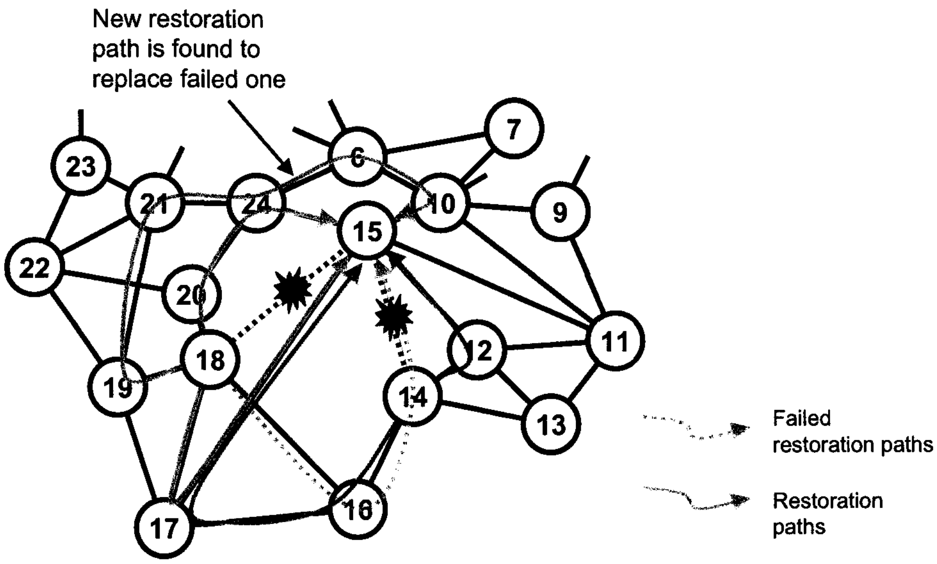

Figure 4-6 Response of the fully-adaptive restoration behaviour to a dual span-failure

toration paths for the first failed span to replace the ones affected by the second span-failure.

Detailed description of the algorithm used to calculated the restorability to a dual span-failure combination *(i,j)* in a network using a fully-adaptive restoration behaviour is given in Figure 4-7. The first stages are almost identical to those of the partly adaptive behaviour with the difference that w_j^{eq} is not used in this approach. Stage 6, corresponding to the restoration of span *j*, is where the difference with the partly adaptive restoration behaviour starts to be seen. Here the first restoration effort on span *j* only includes working link that are "native" to span *j*, restoration paths of span *i* do not benefit from this first restoration effort. In stage 7, the restoration paths of span *i* that crossed span *j* are de-allocated, which means a release of the spare capacity on the surviving spans along these paths (that is on all spans of these paths except on span *j*, justifying the term $-e_j$) and a temporary increase of the number w_i^{nr} of non restored working links of span *i* (which might decrease again in stage 9.). In stage 7, variable *c* is used to make sure we are not de-allocating paths that were not used in the first place (this is important because the number of paths found in stage 2 can be higher than the number w_i of working units to restore). Following stage 7, a new restoration path search and restoration are performed for span *i* to try to re-restore the capacity units on span *i* affected by the second failure, followed by a new restoration path search and restoration of span *j* because new restoration paths may be feasible after the release of spare capacity of stage 7. Again, in the final stage, the total number of non-restored units is calculated.

1. Initialization: Identical to 1. in Figure 4-5 except that w_j^{eq} is not used here
2. Restoration path search for span i: Identical to 2. in Figure 4-5.
3. Restoration of span i: Identical to 3. in Figure 4-5 but excluding: $w_j^{eq} = w_j^{eq} + 1$
4. Failure of span j: Identical to 4. in Figure 4-5.
5. Restoration path search for span j: Identical to 5. in Figure 4-5.
6. Restoration of span j:

$w_j^{nr} = w_j$

$p = 1$

while $(w^{nr} > 0 \ \& \ f_i^{|P_j|} > 0)$

    if $(f_j^p = 0)\ p = p + 1$

    $S^* = S^* - \Delta_j^p$

    $f_j^p = f_j^p - 1$

    $w_j^{nr} = w_j^{nr} - 1$

7. Release of affected restoration paths of span i: (re-initialize set $\boldsymbol{P}_i$ to the values found in 2.)

$c = 0$

$p = 1$

while $(c < w_i \ \& \ f_j^{|P_i|} > 0)$

    if $(f_i^p = 0)\ p = p + 1$

    if $(\delta_{i,j}^p = 1)$

        $S^* = S^* + \Delta_i^p - e_j$

        $w_i^{nr} = w_i^{nr} + 1$

    $f_i^p = f_i^p - 1$

    $c = c + 1$

8. Restoration path search for span i: Identical to 2.
9. Restoration of span i: Identical to 3. without initialization of w_i^{nr}.
10. Restoration path search for span j: Identical to 5.
11. Restoration of span j: Identical to 6. without initialization of w_j^{nr}.
12. Calculation of total non-restored working units:

$N(i,j) = w_i^{nr} + w_j^{nr}$

Figure 4-7 Algorithm for analyzing R_2 with the fully-adaptive behaviour

4.5 Experimental Results of R_2 Computational Analysis

This section presents results obtained using programs that were developed to implement the three restoration algorithms presented in the previous section, and that were applied to a series of test designs for the test networks presented in Section 1.4.1.

4.5.1 Experimental Designs

For each test network, two designs were produced: one non-modular and one modular. To generate the non-modular designs used in to obtain the experimental results presented in this section, we used Herzberg's IP formulation presented in Section 2.5.3 and, except where specified, we always used the ten shortest restoration routes based on distance as the set of eligible restoration routes for each span. When other routes had the same length as the tenth shortest, they were also included in the set of eligible routes. In practice, the non-modular designs were obtained using the MSCP AMPL formulation presented in Section D.1, with a single module type of size 1.

The modular designs used to obtain the experimental results presented in this section were created using the MSCP formulation presented in Section 2.5.4 also with ten eligible restoration routes for each span. For each modular design, the size of module chosen was in the set {3, 12, 24, 48}, the highest one that would produce a spare capacity increase no greater than 100 percent compared to the non-modular design. The idea was not to produce modular designs with too much extra spare capacity compared to the minimum required, while still reflecting the fact that in real life the amount of capacity that is installed is not exactly the minimum that is required.

4.5.2 Goals of Experiments

There are several goals in performing these experiments. The first one is to get a general idea of how much restorability to dual span-failures we get as a side effect of designing a network for single span-failure restorability. One simple reasoning about this question could be: If the network is restorable to single failures, in the case of a dual failure the first failure should be restorable and the second failure might or might not be restorable, which would imply that restorability should be at least 50 percent. Another point of view is that in case of a dual failure, the second failure might not be restorable and, in addition, it could affect the restorability of the first failed span and therefore make the restorability fall below 50 percent. Both points of view are valid and we therefore need experiments to answer this fundamental question. Another factor we will be interested in is how much variation there is between the restorability of the different dual span-failure scenarios. We will also be interested in seeing how much the adaptability of restoration mechanism affects the restorability to dual failures and how close the three restoration models get to the reference two-commodity max-flow solution.

Another factor that will be analyzed is how much the characteristics of the network topology, especially the average nodal degree, influence R_2. A motivation for investigating that aspect is that it is difficult to predict whether a higher nodal degree will actually help achieve higher R_2 or, on the contrary, will tend to cause lower R_2 to be observed. Indeed, on the one hand, a network with higher nodal degree provides more routing options, which decreases the number of service paths crossing each span, so the average number of affected service paths for dual failures decreases. Also, providing more routing options for restoration can only benefit the restorability. On the other hand, however, a network with higher nodal degree requires less spare capacity to be fully restorable to single failures and, therefore, might provide fewer restoration paths to restore dual failures. We are therefore interested to know whether the increase in routing options will have more influence on R_2 than the decrease of spare capacity.

The relation between restoration path length and restorability will also be investigated. It is expected that longer restoration paths will tend to make them more vulnerable to a subsequent failure and therefore will decrease R_2. This effect will be reinforced by the fact that allowing longer restoration paths usually allows more sharing of spare capacity and therefore lowers capacity requirements, again causing R_2 to go down. Conversely, restricting the length of restoration paths should improve R_2 at the cost of higher capacity requirements, so we will be interested in looking at the trade-off between capacity and R_2 with this approach, which will later be compared with other more sophisticated approaches to obtaining higher R_2.

Finally we will present a general qualitative analysis of the link that exists between capacity redundancy and R_2. It is often believed that adding more capacity always results in higher availability. We will argue that this is not necessarily the case and we will start to give answers the question: "In what conditions does more capacity translate into higher availability?"

4.5.3 Discussion of Experimental Results

Results of dual span-failure restorability analysis are presented in Tables 4-1 and 4-2. Table 4-1 on page 91 presents the results for the non-modular designs. For each of the three restoration models presented in Section 4.4, the table provides minimum, average and maximum results for R_2. The average results presented correspond to the definition of R_2 given in (4.8). The minimum and maximum $R_2(i,j)$ results are based on (4.1) and are results over all ordered dual span-failure pairs. The TCMF column gives results obtained using the TCMF-Flow formulation presented in Section 4.3. In practice, the formulation used minimizes the total number of non-restorable working links over all dual failure scenarios and allows direct calculation of the average R_2 through (4.8). The AMPL model used for this formulation is given in Section D.2 of Appendix D. The TCMF column of the table contains the result of the AMPL formulation and the corresponding average R_2. Table 4-2 on page 92 presents the corresponding results for the modular designs.

From the observation of Table 4-1 it appears that the average dual span-failure restorability is higher than what might have been expected. Even though the test designs are only guaranteeing full restorability to single span-failures, the restorability to dual span-failures ranges by network from 0.5543 to 0.8891 for Model 1, 0.6047 to 0.92 for Model 2, and 0.6073 to 0.9250 for Model 3. It is also interesting to note than despite these high average values, there are still cases in almost all test networks of dual span-failures for which no working links are restorable. Most of the time, these events are linked to dual failures happening on spans that are connected by a degree-2 node, but inspection of results revealed that sometimes it is due to failures happening on spans connected by a degree-3 node but whose third incident span bears no spare capacity. This finding prompts the idea that the design formulation could include some constraints that prevents such situations and, more generally, could guarantee a minimum restorability level for each dual failure. This topic is largely covered in Chapter 6. Less surprisingly than the occurrence of fully *non-restorable* dual failures, the maximum restorability over all dual span-failure scenarios is 1 for all test networks except for the non-modular design of the smallest test network. It is worth mentioning, however, that dual span-failures for which the total working capacity to restore is zero are not considered in the calculation of average R_2 or in the determination of maximum $R_2(i,j)$. There are therefore for each test network, dual span-failure scenarios for which the restoration mechanism is able to find restoration paths for all affected working links, either because failures are spatially independent (as it must be the case for Model 1) or thanks to the adaptability of the mechanism.

From the observation of Table 4-2, we see that, as expected, dual-failure restorability for all modular designs is higher than in the non-modular cases. R_2 values now range from 0.7356 to 0.9020 for Model 1, 0.8334 to 0.9727 for Model 2, and 0.8478 to 0.9739 for Model 3. With Model 2 and Model 3, the minimum restorability results are also significantly improved for networks that have no degree-2 cuts. Except for the case of three test networks with no degree-2 cuts that still experience some $R_2(i,j)$ of 0, all these events are due to the case explained earlier where three spans connect a node and one of them has no spare capacity: all the three-connected test networks always offer higher than zero restorability to dual-failures with Models 2 and 3. Some of them even guarantee minimum $R_2(i,j)$ in the order of 30 percent or higher for absolutely all possible dual span-failures. These results let us foresee the possibility of giving very high restorability and availability guarantees to mesh restorable networks with an adaptive restoration mechanism. Again, this will be covered in Chapter 6.

On Figures 4-8 to 4-11 are shown the distributions of individual $R_2(i, j)$ values over all dual failures for each test network. On these graphs it appears clearly that dual span-failures for which the restorability is low (say below 40 percent) are quite rare, especially for the bigger networks. Note that these graphs do reflect the occurrence of fully non-restorable events due to the isolation of a degree-2 node or related to the "3-spans, one with zero spare" scenario explained earlier. It appears that the frequency of such dual failures is extremely low as can be seen by looking at the "0-10%" bins of the different graphs. The following section compares the three restoration models based on the results that we started to comment in this section.

4.5.4 Comparison of the Performance of Restoration Models

Another interesting question related to the experimental results presented in the previous section is to know how the dual failure restorability of the different restoration models compare and how they compare to optimal results given by solving TCMF. In addition to Figures 4-8 to 4-11, to allow easier comparison between the results of the different restoration models and with TCMF, the results of Tables 4-1 and 4-2 are plotted on the graphs of Figures 4-12 and 4-13. Each graph shows the results for a given test network. On these graphs it appears clearly that Model 2 gives results that are very close to those given by Model 3. Model 1, as expected, gives lower restorability results than Model 2 and Model 3 (2 to 5 percent lower on average for the non-modular designs). Also, Model 2 and Model 3 appear to benefit much more than Model 1 from the spare capacity increase associated with modular designs. Indeed, the dual span-failure restorability of Model 1 becomes 5 to 10 percent lower than that of Model 2 and Model 3 in the case of modular designs. This is not surprising as the adaptability of Models 2 and 3 allows them to make better use of the extra capacity that comes with modular designs whereas Model 1 can only make use of the extra capacity on the spans where it would have looked for spare capacity in the first place.

Another striking result is how well Models 2 and 3 perform compared to the optimal TCMF solution. In all cases, the restorability of these two Models is within 1 percent of optimality.

Comparing now the three restoration models on Figures 4-8 to 4-11, it appears the adaptive effects are very active and significant in terms of raising $R_2(i, j)$ for the worst cases of low individual $R_2(i, j)$ under Model 1. This is particularly true for the modular designs of test networks 12n30s1, Net-A, Net-B, and Net-C. These graphs also confirm the fact that the performance of Model 2 is very comparable to that of Model 3.

Table 4-1: Dual span-failure restorability results for non-modular designs

Test Network	Model 1			Model 2			Model 3			TCMF	
	Min. $R_2(i,j)$	Ave. R_2	Max. $R_2(i,j)$	Min. $R_2(i,j)$	Ave. R_2	Max. $R_2(i,j)$	Min. $R_2(i,j)$	Ave. R_2	Max. $R_2(i,j)$	Ave. R_2	$\sum N(i,j)$
06n14s1	0.2632	0.5543	0.9167	0.2632	0.6047	1	0.2632	0.6073	1	0.618617	1180
11n20s1	0	0.7629	1	0	0.7753	1	0	0.7782	1	0.7812	5354
11n20s2	0.0588	0.6859	1	0.127	0.7312	1	0.127	0.7382	1	0.7480	6828
Bellcore	0	0.7179	1	0	0.7911	1	0	0.8028	1	0.806964	3958
Bellcore Mod.	0	0.771	1	0	0.8262	1	0	0.8347	1	0.84	3696
COST239	0	0.7744	1	0	0.8415	1	0	0.8494	1	0.861581	1488
12n20s1	0	0.7203	1	0	0.7508	1	0	0.757	1	0.762146	3760
12n30s1	0	0.8498	1	0	0.8807	1	0	0.883	1	0.8848	3340
15n28s1	0	0.7861	1	0	0.8206	1	0	0.8272	1	0.832796	7774
16n29s1	0	0.7938	1	0	0.8071	1	0	0.8132	1	0.818163	8686
16n38s1	0	0.8262	1	0	0.8606	1	0	0.8649	1	0.872884	6848
EuroNet	0	0.8159	1	0	0.8622	1	0	0.8701	1	0.878738	13105
Net-A	0	0.8362	1	0.1239	0.8778	1	0.1239	0.8847	1	0.893362	21909
22n41s1	0	0.8402	1	0	0.866	1	0	0.8729	1	0.879977	13577
Net-B	0.0364	0.8667	1	0.0364	0.897	1	0.0364	0.9035	1	0.911083	40162
Net-C	0	0.8891	1	0	0.92	1	0	0.925	1	0.930666	63962

Table 4-2: Dual span-failure restorability results for modular designs

Test Network	Model 1			Model 2			Model 3			TCMF	
	Min. $R_2(i,j)$	Ave. R_2	Max. $R_2(i,j)$	Min. $R_2(i,j)$	Ave. R_2	Max. $R_2(i,j)$	Min. $R_2(i,j)$	Ave. R_2	Max. $R_2(i,j)$	Ave. R_2	$\sum N(i,j)$
06n14s1	0.1538	0.7356	1	0.2857	0.8736	1	0.2857	0.8759	1	0.881707	366
11n20s1	0	0.8096	1	0.0732	0.8674	1	0.0732	0.8721	1	0.8778	2990
11n20s2	0	0.7467	1	0.2438	0.8334	1	0.2438	0.8478	1	0.8598	3798
Bellcore	0	0.7925	1	0	0.8979	1	0	0.9103	1	0.912505	1794
Bellcore Mod.	0	0.8207	1	0	0.9276	1	0	0.934	1	0.93619	1474
COST239	0	0.85	1	0	0.9398	1	0	0.943	1	0.94586	582
12n20s1	0	0.7806	1	0	0.8941	1	0	0.8957	1	0.896002	1644
12n30s1	0	0.8852	1	0.4483	0.9727	1	0.4483	0.9739	1	0.9749	728
15n28s1	0	0.8486	1	0	0.9263	1	0	0.9345	1	0.936637	2946
16n29s1	0	0.8535	1	0	0.9079	1	0	0.9126	1	0.914985	4061
16n38s1	0	0.8756	1	0.1818	0.9491	1	0.1818	0.9573	1	0.963543	1964
EuroNet	0	0.8749	1	0	0.9544	1	0	0.9571	1	0.959018	4429
Net-A	0	0.8677	1	0.1852	0.9328	1	0.1897	0.939	1	0.944191	11466
22n41s1	0	0.8971	1	0	0.9498	1	0	0.9547	1	0.956931	4872
Net-B	0	0.892	1	0.0753	0.9349	1	0.1515	0.944	1	0.94602	22764
Net-C	0	0.902	1	0.0404	0.9384	1	0.0404	0.9427	1	0.946476	49377

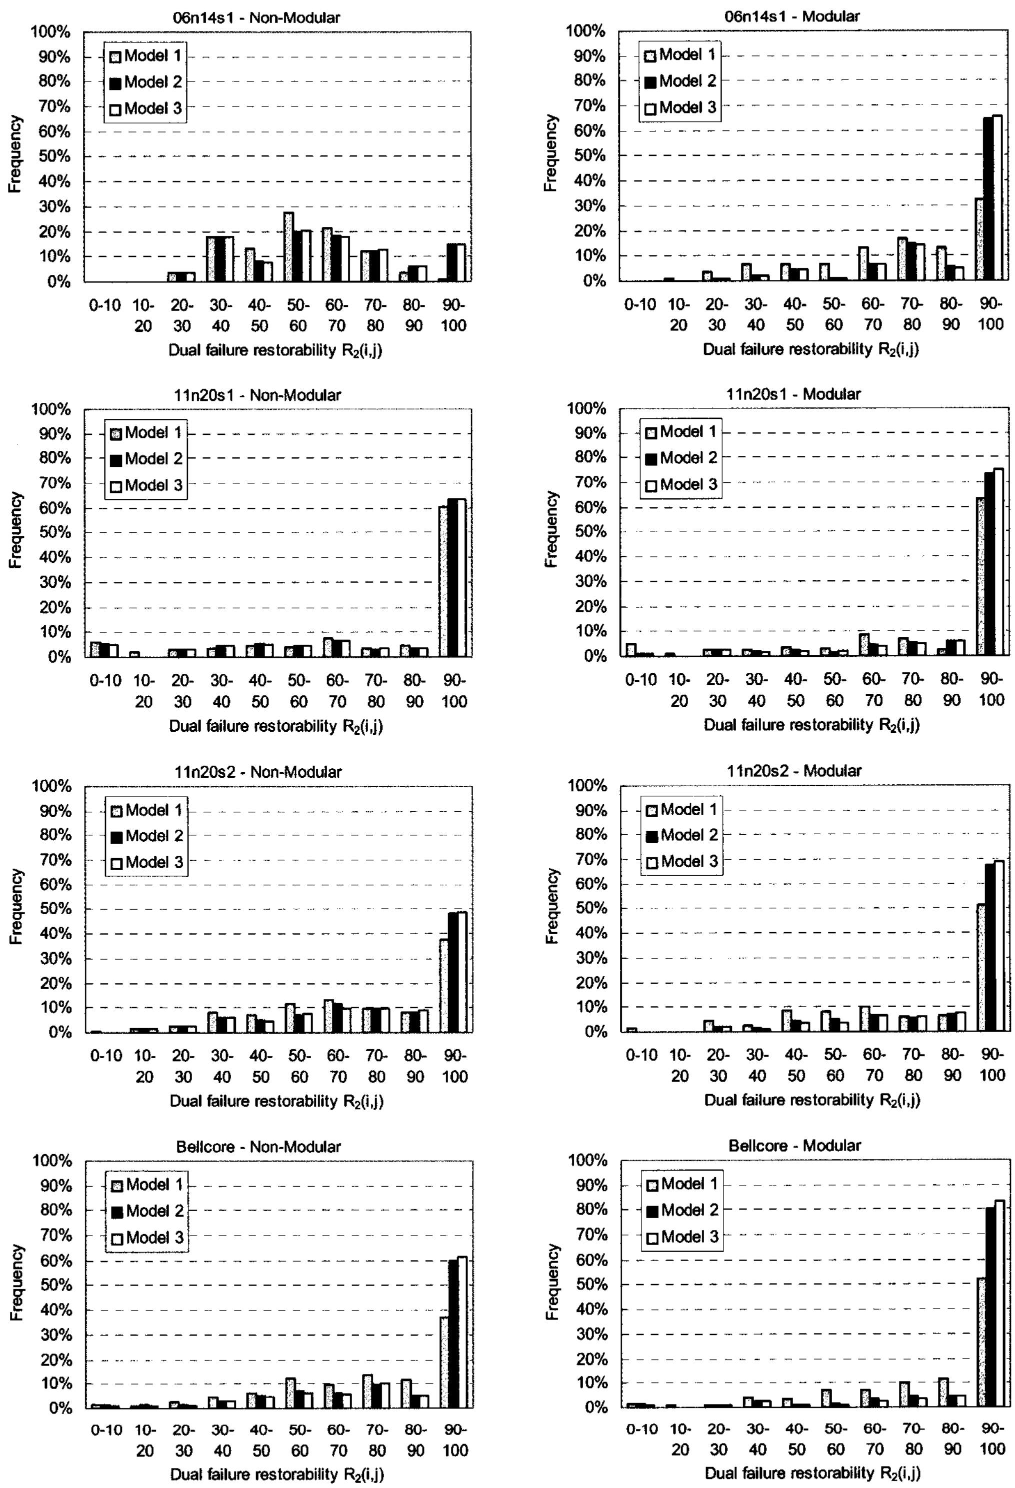

Figure 4-8 Histogram of individual $R_2(i,j)$ levels per dual failure pair

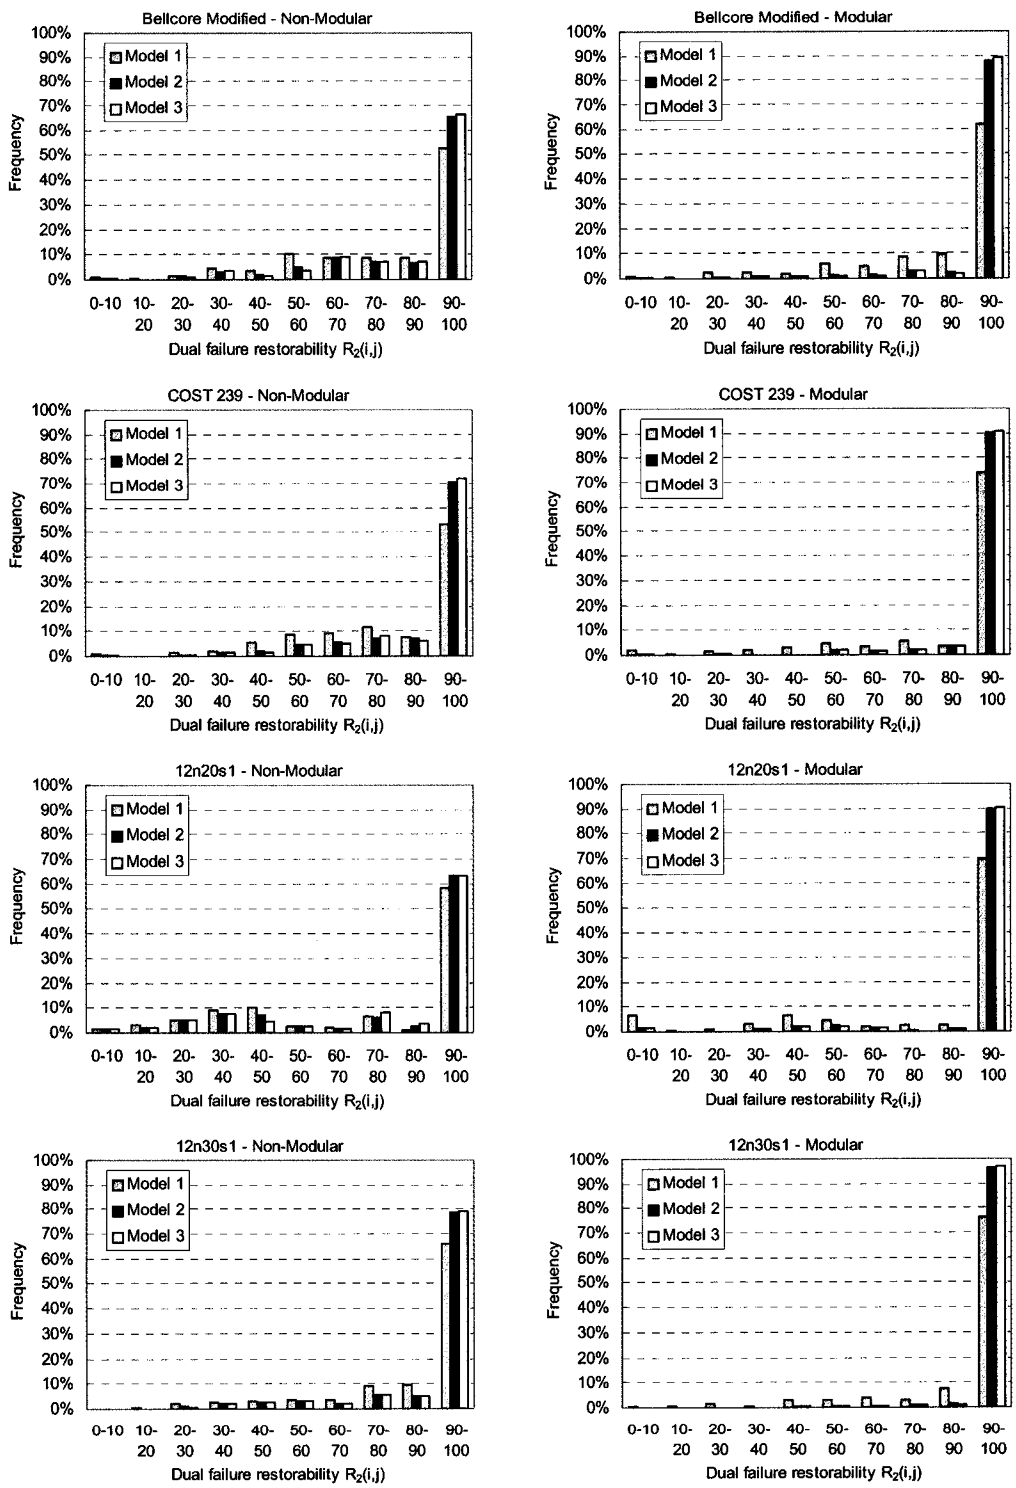

Figure 4-9 Histogram of individual $R_2(i,j)$ levels per dual failure pair (continued)

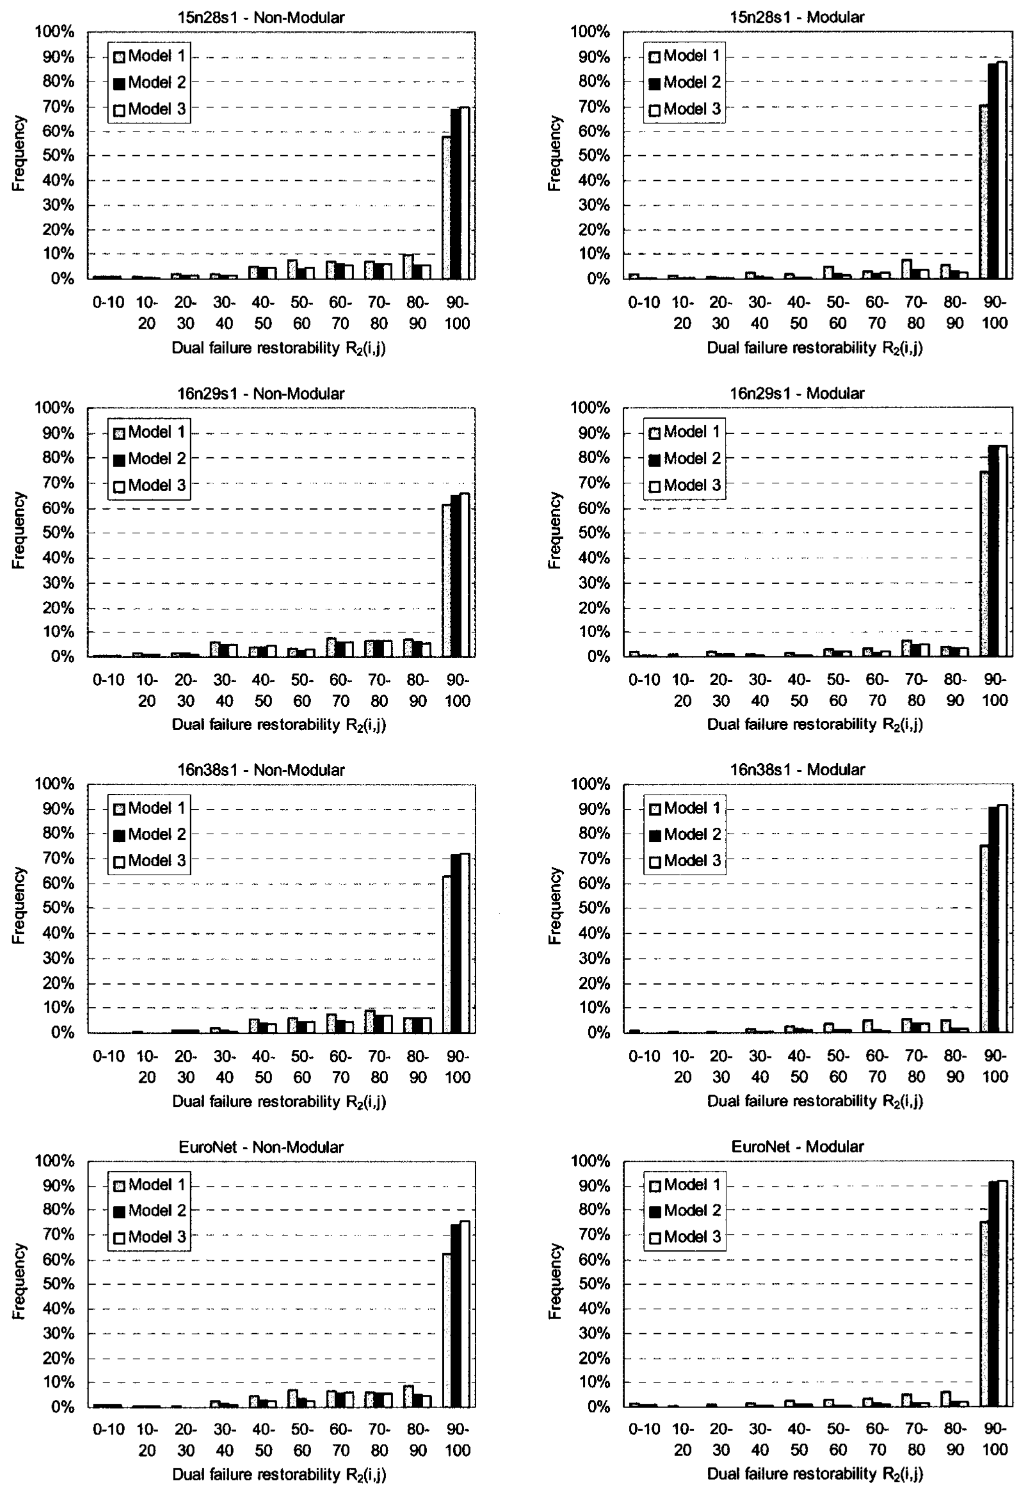

Figure 4-10 Histogram of individual $R_2(i,j)$ levels per dual failure pair (continued)

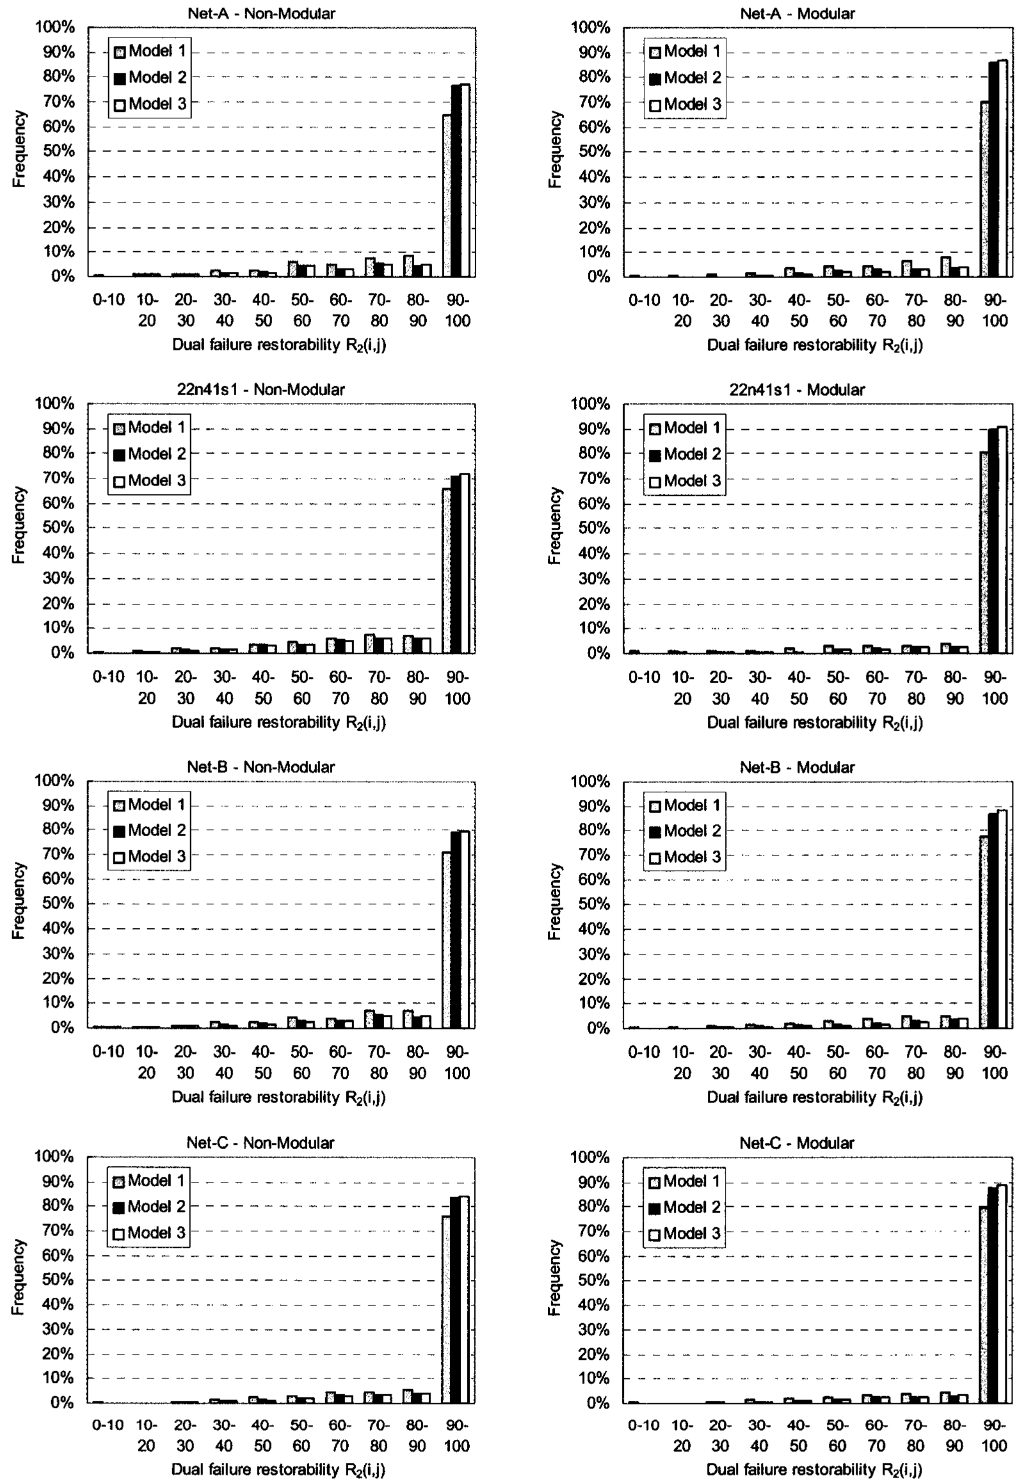

Figure 4-11 Histogram of individual $R_2(i,j)$ levels per dual failure pair (continued)

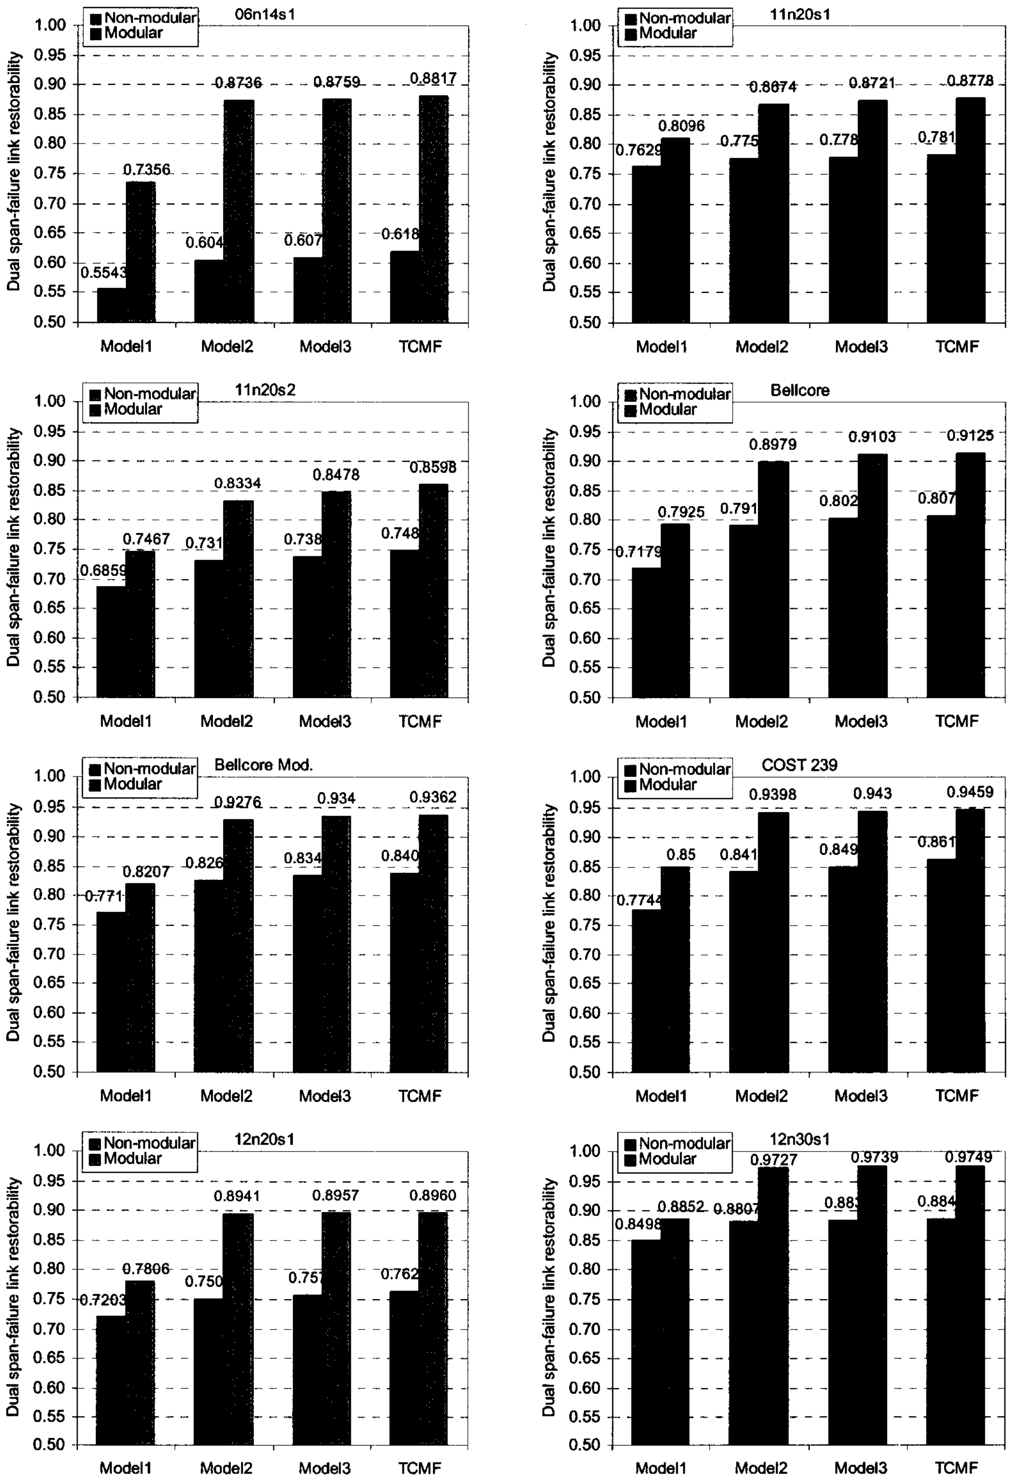

Figure 4-12 Experimental results of dual span-failure link restorability

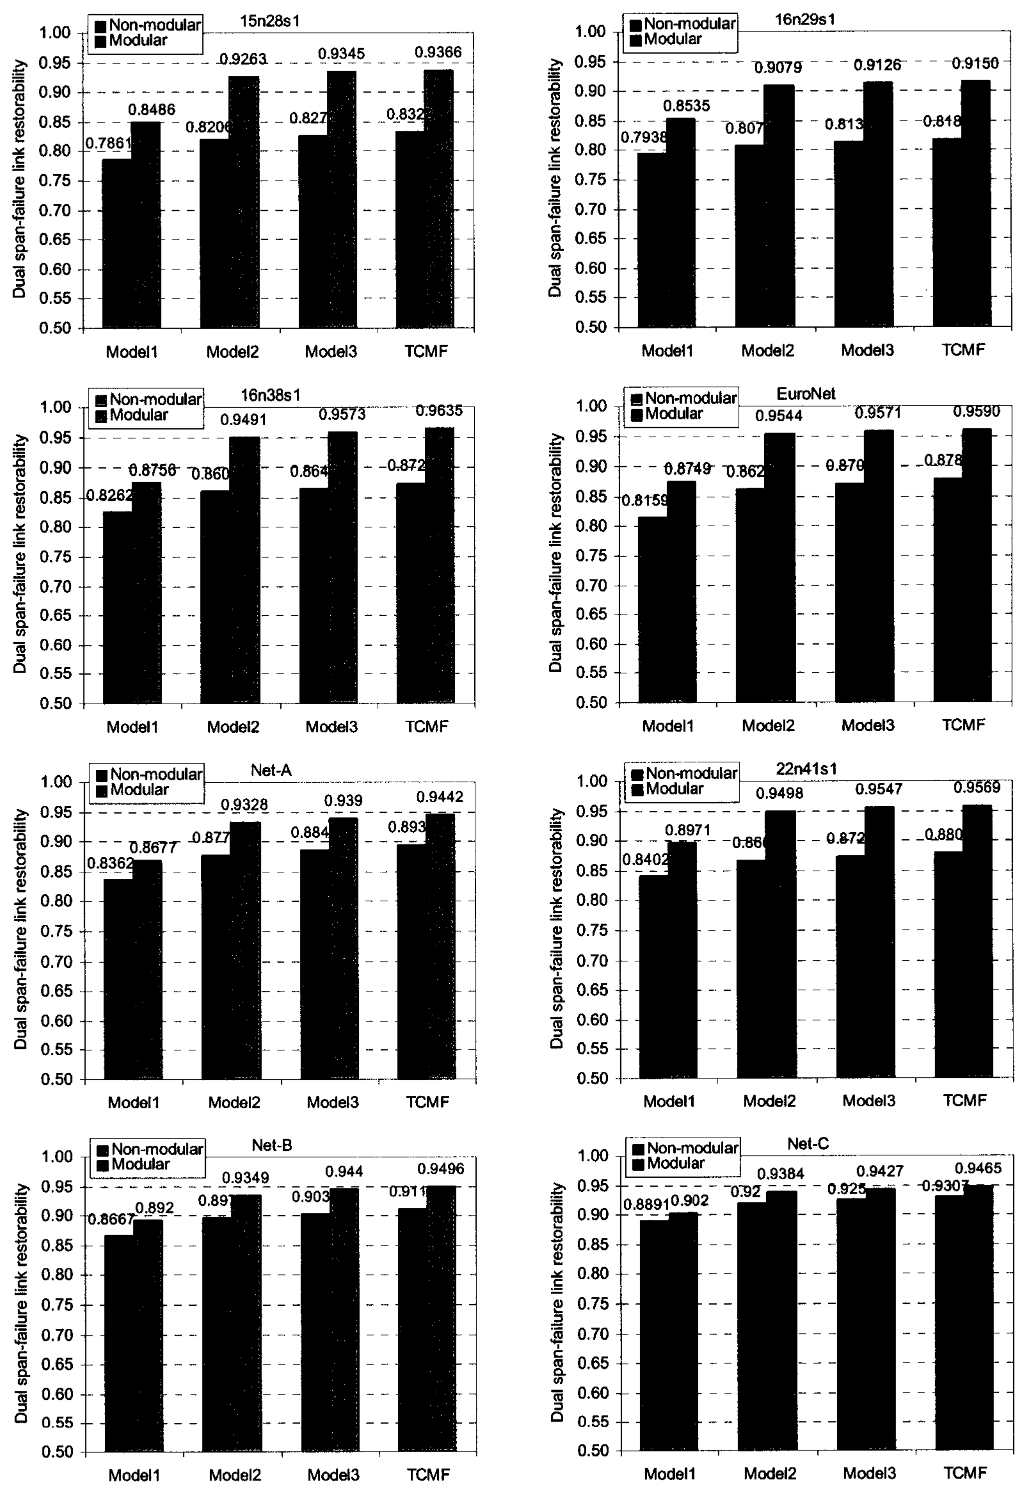

Figure 4-13 Experimental results of dual span-failure link restorability (continued)

4.5.5 Influence of Capacity Redundancy on R_2

In the results presented in the previous section we saw that the dual span-failure restorability of single-failure restorable networks is relatively high in general, but also that it varies significantly from one test case to another. In this section and in the following two, we investigate several possible factors that night influence the restorability of the network to dual span-failures.

The first factor that we investigate is the redundancy of the design. As mentioned earlier, higher redundancy is often believed to be a sufficient condition to obtain higher availability: we can test if this belief corresponds to the reality.

On Figure 4-14, each test network is represented by its capacity redundancy on the x-axis and dual span-failure restorability on the y-axis. For each test network the four points represented correspond to Model 1, Model 2, Model 3, and TCMF. The results plotted are for the non-modular designs. While the plot shows a slight tendency towards greater R_2 with increasing redundancy, it is by no means a monotonic progression. We interpret the almost flat general nature of the scatter of the test case designs as meaning that the availability depends more on the individual network and demand pattern that on the simple bulk redundancy of the network. Other possible factors are investigated in the following two sections.

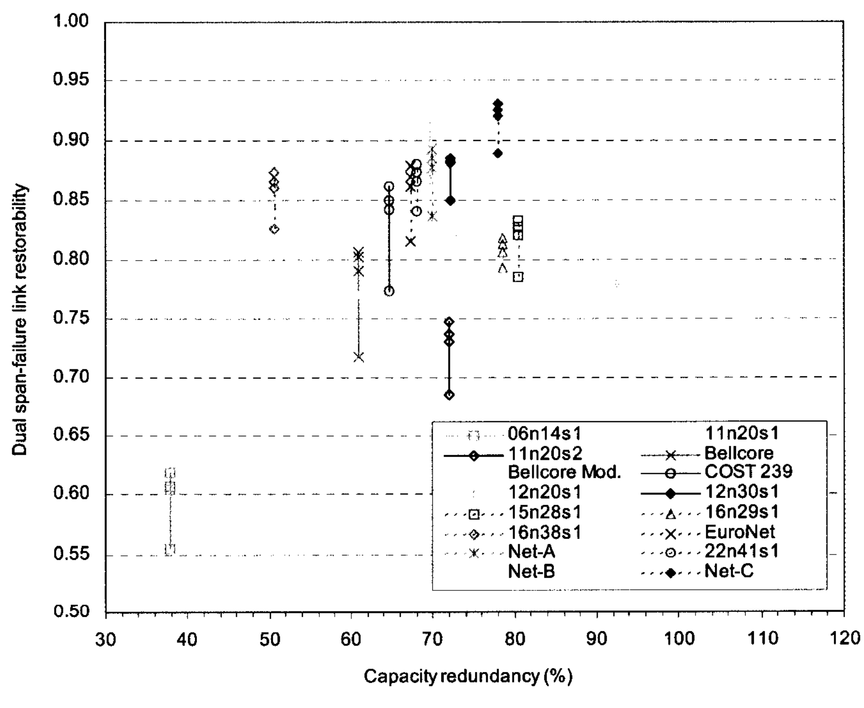

Figure 4-14 Dual span-failure restorability versus capacity redundancy for all SCP designs

4.5.6 *Influence of Network Connectivity on R_2*

The second factor tested here is the connectivity of the network that can be measured by the average nodal degree as defined in Section 1.3.3. An increase in average nodal degree corresponds also to an increase in network connectivity. On Figure 4-15 each test network is represented by its average nodal degree on the x-axis and dual span-failure restorability on the y-axis. Again, results of the three restoration models and of TCMF are shown and they correspond to the non-modular designs. This time the plot shows a slightly more pronounced tendency towards greater R_2 with increasing network connectivity but again, it is not a systematic progression. It is interesting to note that the networks that had a high capacity redundancy but a low restorability in the previous graph are among the ones with the smallest average nodal degree in this graph. This seems to indicate that the very low capacity redundancy of these networks was the factor that prevented their high redundancy to translate into a high dual-failure restorability. Conversely, here the test case 06n14s1 that has high average nodal degree but with low dual span-failure restorability is the one that had the smallest redundancy in the previous graph. This indicates that the very low redundancy of that design becomes a limiting factor that prevents its high connectivity from translating into high restorability. High connectivity and high redundancy both seem to matter in order to obtain high

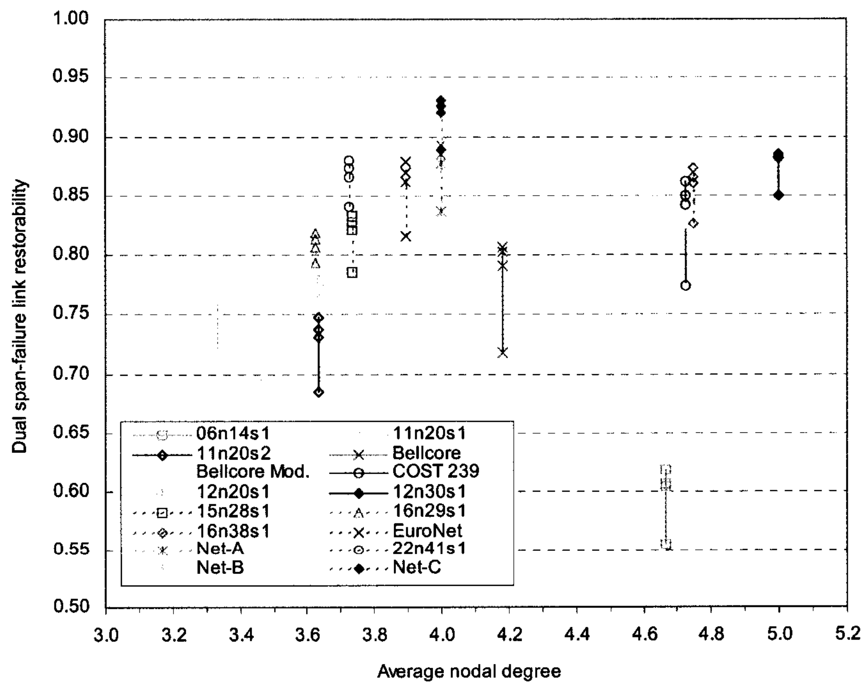

Figure 4-15 Dual span-failure restorability versus average nodal degree for all SCP designs

restorability with maybe slightly more importance related to the network connectivity. Based on this idea, we plot the results versus the capacity redundancy multiplied by the square of the average nodal degree to take into account the apparently higher importance of network connectivity. The plot, on Figure 4-16 shows a fairly clear progression towards higher R_2 for increasing values of this metric.

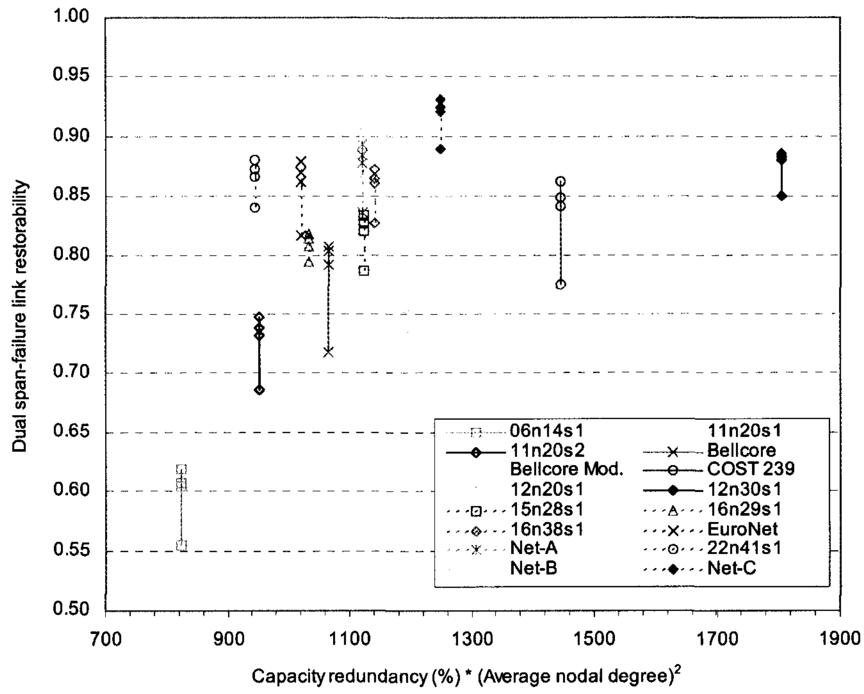

Figure 4-16 Dual span-failure restorability versus proposed metric for all SCP designs

4.5.7 *Influence of Network Size on R_2*

The last factor tested is simply the size of the network. This was motivated by inspection of the results on Figures 4-12 and 4-13, which clearly showed increasing R_2 results as network size increases. We test two measures of the network size: number of nodes and number of spans. Figures 4-17 and 4-18 show the restorability results plotted against number of nodes and number of spans, respectively. This time, the progression in very clear, particularly on Figure 4-18, where results are plotted versus the number of spans. Network size, measured by the number of spans seems, therefore, to be the factor that influences the most the restorability to dual failures. One reasoning about the reasons of this simple phenomenon could be that bigger networks offer large numbers of possible restoration routes, which can be well exploited by adaptive restoration mechanisms to achieve

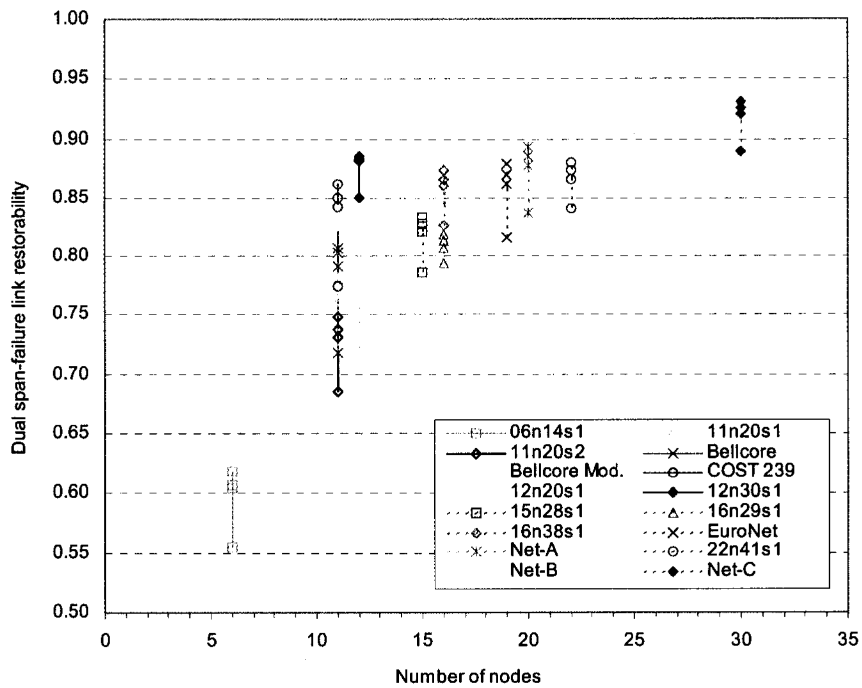

Figure 4-17 Dual span-failure restorability versus number of nodes for all SCP designs

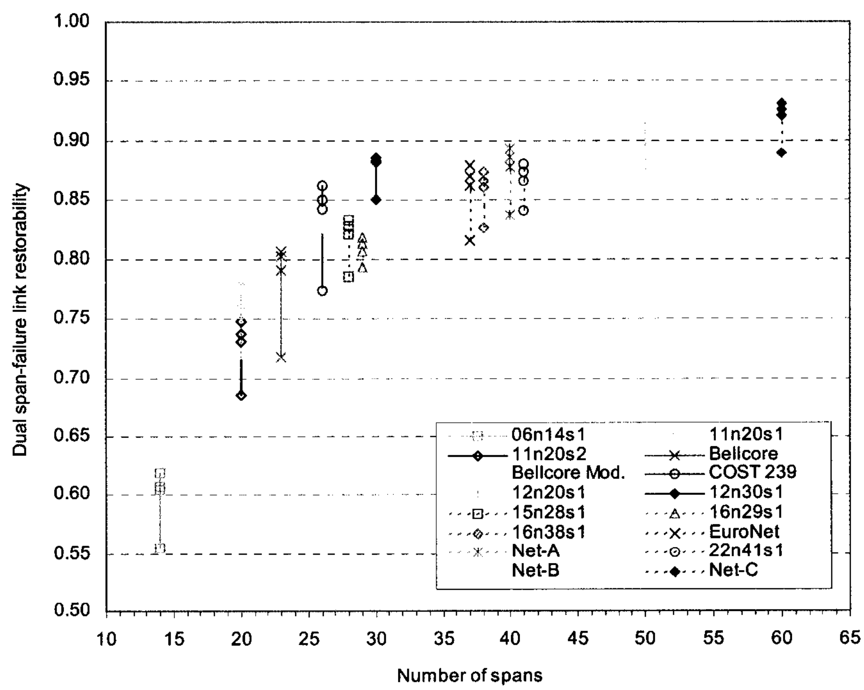

Figure 4-18 Dual span-failure restorability versus number of spans for all SCP designs

high restorability levels. This, however does not explain the fact that the static restoration behaviour (Model 1) also performs better for larger networks. In fact, the explanation is also that the large number of restoration routes makes it less likely that a dual-failure would affect both the span bearing a given working unit and the pre-planned restoration path of that capacity unit. In that case it matters less whether the restoration mechanism is very adaptive or not. Indeed, as can be seen on Figures 4-17 and 4-18, the advantage of Model 2 and Model 3 over Model 1 becomes smaller for larger networks – although it is still only about 2 to 3 percent.

4.5.8 Other Possible Factor: Demand Distribution

Another possible factor that could influence the dual-failure restorability of mesh networks is the distribution of demands in the network. For example, the demand distribution in a networks having a few hub nodes—when all demands originate or terminate at one of these nodes—could result in different restorability levels than with a demand distributions obtained using the model described in Section 1.4.2. This aspect is not covered by this thesis work and is left for future research.

4.6 Summary

This chapter described a computational approach to the problem of analyzing the dual span-failure restorability of span-restorable mesh networks. The overall method is practical to use but while not leaving out the important details of irregular topology, capacity distribution, and restoration mechanism. The results showed that networks designed for full restorability to single span-failures inherently enjoy high levels of dual failure restorability. The dual span-failure restorability levels can be over 90 percent for the combination of a fully-adaptive restoration algorithm in a modular capacity design. These findings tend to counter the qualitative expectation in some quarters that mesh-restorable networks may not give as high a service availability as ring-based networks because of their lower redundancy. What we see, however, is that despite the lower redundancy of the mesh, the generality of a highly adaptive routing process leads to an extremely high level of dual-failure restorability and hence to the high availability of paths realized over spans of the network. The fact that the dual span-failure restorability R_2 is almost never zero, and usually at least 20 percent or more, also has striking implications for the high quality of premium-path service that a mesh-restorable network can provide, at no loss of assured R_1 restorability for all other services, and with no additional spare capacity other than that required for R_1 restorability itself.

5. Influence of Maintenance Actions on Service Availability

5.1 Introduction

An operational aspect of mesh-restorable networking that has apparently not received attention in the literature is the way in which maintenance activity on mesh spans may create a theoretical exposure to reduced restorability if a failure occurs elsewhere during the maintenance state. Several cases of important network outages caused by or related to some form of network maintenance or service/software upgrades have been reported in the literature [Neu95][SnW00][Sno01]. This chapter is based on the author's involvement, along with W. Grover and J. Doucette in a novel study project on the issues related to maintenance in mesh-restorable networks following an expression of interest or concern on the topic by T. Bach and others at Nortel Networks [GCB01]. This issue is of high interest for quickly growing DWDM transport networks in which the frequency of in-service upgrades or other required maintenance activity can be fairly high. In order for service to remain uninterrupted during such actions, network operators usually use the network's protection capacity by manually rerouting signals or simply letting the network protection or restoration mechanism act upon what it interprets as being a failure. The consequence of that is that during such actions, the network is placed in a state that is to some extent equivalent to a failure state and any "real" failure that could happen during that time would effectively create a dual-failure-like situation. As we have seen in the previous section, restorability to a dual failure is not guaranteed, thus the effect of an upgrade or maintenance action is a reduction of the restorability of some other spans of the network in the event of their own failure if it happened during the time maintenance is being performed. The extent of the influence of maintenance actions on the restorability of other spans will be measured by what we will call a *risk field* as portrayed in Figure 5-1. Spans within the risk field (in grey on Figure 5-1) will be the ones for which full restorability (in the case of their own physical failure) cannot be guaranteed during the maintenance action. Each of these spans will be characterized by a *theoretical risk* value, which represents the probability for a link on that span would not be restored if that span were to fail during the maintenance action. We call these *theoretical* to emphasize that there is no actual outage unless a failure really occurs during the maintenance state. Spans outside the risk field (solid black lines on Figure 5-1) have a risk of zero, which means they would still be fully restorable even if they failed during the maintenance action. In a ring-based network, determining the structure of the risk field is straightforward: The set of spans at risk from maintenance on one span is comprised of all the other spans in the same ring and the magnitude of the risk is 100 percent loss of restorability. For span-restorable mesh networks,

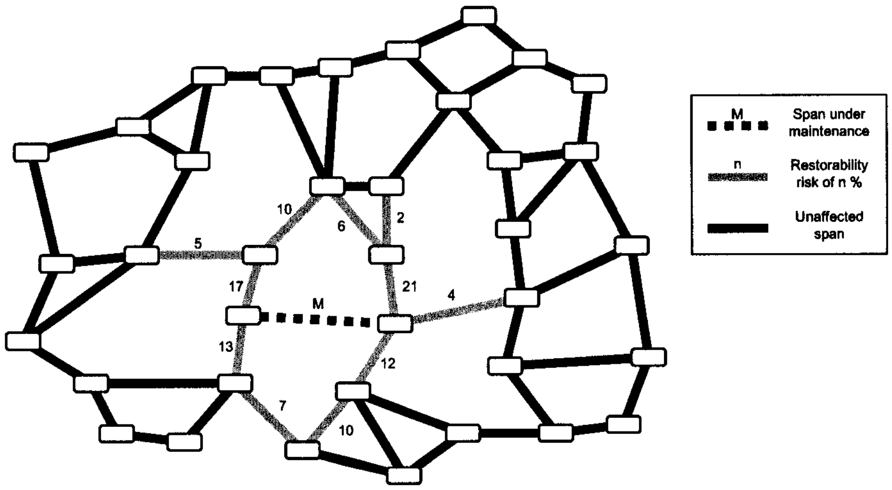

Figure 5-1 Conceptual example of risk field

the risk field will be more diffuse and generally less than 100 percent in magnitude on any span.

In this chapter, we provide methods for characterizing the corresponding theoretical risk field for span maintenance in a mesh network using span-restoration and suggest approaches to managing the theoretical risk that a "maintenance outage" may have for the rest of a mesh-restorable network. Of particular interest is the prospect that in a mesh-restorable network it may be feasible to guarantee zero risk due to maintenance actions for high priority service paths. This is an important question that the work presented in this chapter confirms.

In increasing steps of complexity, the following measures or goals were conceived in regards to the general topic defined above:

i. A first requirement is a basic analysis capability to assure that the network-wide extent of "exposure to a subsequent failure" of a proposed maintenance action is at least known in advance by some analysis. Such an initial tool or capability would simply tell a network operator what the extent of theoretical exposure is to a single-span failure anywhere else in the network, while doing a certain maintenance action. The "extent" of influence of a maintenance action can be quantified by a vector that indicates, for each other span, the loss of restorability of the other span under the condition of a single real failure arising while in the maintenance state. Given a proposed maintenance action, the tool could give a diagnosis of any particular other spans that would accrue some non-zero restorability risk or simply flag other spans sustaining risk above some threshold of significance due to the proposed maintenance.

ii. A second operational capability could be the ability to assure that multiple simultaneously scheduled maintenance outages are in some sense (to be further defined) "mesh-orthogonal" in that the geographical / topological ranges of failure exposure do not overlap or compound each other. This type of future tool or capability would tell a network operator what combinations of simultaneous maintenance activities avoid compounding any theoretical risk or possibly avoid maintenance-related risk entirely for priority service paths. If x and y are single span maintenance projects, the aim would be to know if the combined (parallel) maintenance plan (maintenance x and y at the same time) runs a significantly greater overall risk than is inherent in serial maintenance plan (maintenance x followed by maintenance y, or vice versa). This is covered further in Section 5.5. For the remainder of this chapter, the following notation will be used in discussing serial or parallel maintenance actions:

- $x \| y$: maintenance actions x and y occur simultaneously on the network (parallel maintenance),
- $x + y$: maintenance actions x and y are scheduled in succession,
- $x \perp y$: x and y are "mesh orthogonal" maintenance actions,
- $|(x \| y)| = k$: there are k spans that are in the risk field of both x and y.

Given this, $|(x \| y)|$ being equal to zero could be thought of as the definition of this type of mesh maintenance orthogonality.

iii. A third capability would be to provide support for planning operational schedules for simultaneous dispatch of lists of prioritized maintenance actions. In this concept, one would go beyond the analysis capabilities i. and ii. to the synthesis of recommended schedules for simultaneous activities. Here we envision an "input hopper" of required maintenance activities. Each input to the hopper may have an associated urgency of time priority for its completion. There may also be a designation of which regional work-crew resources will be busied out by the associated action. The aim of the scheduler would be to dispatch the items in its input hopper as quickly as possible, subject to the priorities of each job, limits on the number of simultaneous maintenance actions due to regional resource-use considerations, and subject to an ongoing assurance of mesh orthogonality of every multi-action maintenance state. Such a system might be imagined scheduling a large number of work crews on simultaneous maintenance actions continent-wide while always assuring that the collective single-failure risk exposure is no greater than if all the same activities were done one-at-a time in succession. We do not address this idea further in the present work.

This chapter is devoted primarily to the first goal. We touch briefly on considerations related to the second goal in closing paragraph. The third goal is left to future research.

5.2 Models for Restorability-Related Effects of Span Maintenance

We can see at least three functional models for the effect of maintenance-related activities from a network restorability viewpoint. This section presents them in order of increasing network influence impact.

5.2.1 Type 1: Ring-like Roll-to-Protection

In this model all of the spare capacity of the maintenance-span is withdrawn from the remainder of the network with no other effects. This assumes that a working channel block is rolled intact as a modular unit onto an identical sized protection channel block, as in a 4-fibre BLSR, to facilitate maintenance on the working channels. This is referred to as a *ring-like roll-to-protection.* The main feature is that the "roll" is completely contained within the maintenance span.

5.2.2 Type 2: Mesh Equivalent of Roll-to-Protection

More generally in a mesh-restorable network, the number of working and spare channels on each span can be separately assigned. In an optimized mesh capacity design at the per-channel capacity level, most spans will have more working capacity links than spare capacity links but the case where the number of spare capacity links is higher than the number of working capacity links and the case where no spare capacity links at all have been allocated on the span are also possible. In this case the "roll to protection" model for a span *i* under maintenance becomes generalized:

i. If $(w_m \le s_m)$:
   - The "move to protection" is completely contained within span *i* itself.
   - The network-wide effect is a withdrawal of w_m units of spare capacity from span *i* but some spare capacity may remain. In other words, the effective spare capacity s_m^{eff} is equal to $s_m - w_m$.

ii. If $(w_m > s_m)$:
   - The "move to protection" first uses up all the s_m spare capacity units on span *m*, the maintenance span, leaving w_m^{eff} (equal to $w_m - s_m$) capacity units left to restore.
   - The remaining working amount w_m^{eff} is re-routed over replacement paths through spare capacity on other spans of the network around the maintenance span.

- The network-wide effect is withdrawal of all spare capacity from the maintenance span *plus* a network withdrawal of spare capacity on other spans as required to support the creation of maintenance replacement paths for w_m^{eff}.

5.2.3 Type 3: Equivalent to Failure & Restoration

The extreme worst-case model for the maintenance action is to assume the complete withdrawal of all spare capacity on the maintenance span and the complete re-routing of all working demands crossing the span via a restoration-like set of replacement paths using spare capacity on other spans. It is difficult to say whether this model is actually required in practice as we understand most maintenance or upgrading is done under roll-and-cut type of procedures. Presumably any Type 3 maintenance action would be very carefully considered, scheduled late at night, or avoided if at all possible by other strategies because (in either a ring- or mesh-based network) such actions put the network in an initial state that is completely equivalent to already having sustained and restored a single failure. The network is then completely reliant on any inherent dual-failure restorability should a true failure arise while in this state. For any given demand in a ring, the restorability risk in such a circumstance is already known: it is all or nothing: either 100 percent or 0 percent for each other span cut. For all demands that cross a failure span while there is a Type 3 maintenance action elsewhere on the ring the risk is 100 percent. We will characterize the corresponding risk for mesh networks under this worst case model as well as the more benign Type 2 model.

5.3 The Theoretical Risk-Field of a Mesh Span-Maintenance Action

5.3.1 Transforming Type 2 Maintenance Actions to Equivalent Type 3 for Simulation

Henceforth, we will consider Type 2 and Type 3 logical models for the capacity-related effects of mesh span maintenance. Type 3 maintenance is functionally equivalent to complete failure and restoration of the maintenance span in terms of the network environment seen for restoration of any actual failure that arises while in this state. The Type2 maintenance model can be thought of as a special case of Type 3 maintenance where the actual w_m of the span is first transformed to equivalent working capacity w_m^{eff} equal to $\max(0, w_m - s_m)$, and s_m is transformed to equivalent spare capacity s_m^{eff} equal to $\max(0, s_m - w_m)$. Type 2 and Type 3 maintenance effects can therefore be conveniently approached though re-use of the dual failure restorability analysis capability presented in the previous chapter. In this view maintenance on a span is equivalent to voluntarily putting the span in a failed mode and "restoring" all its working capacity units. For Type 3 the true w_m and

s_m are used. For Type 2 modelling w_m^{eff} and s_m^{eff} are used instead.

5.3.2 Formal Definition of the Theoretical Risk Field

For both maintenance re-routing of working channels and for actual failures we assume a span-restoration mechanism that produces path-sets equivalent to what would be found by the k-shortest paths algorithm [DGM94][MaG94] through available spare capacity between the end-nodes of the failure or maintenance span involved. The theoretical risk field from maintenance on a span *m* can therefore be evaluated by means similar to computing the dual-failure restorability of the network over all dual-failure scenarios involving span *m* as one of the "failures."

Based on the analysis of the previous chapter, there are two effects through which a maintenance state coupled with an actual span failure can lead to an outcome that is not fully restorable:

1. Contention for spare capacity: When the first span is under maintenance the required number of restoration paths for restoration of a second (failure) span may not be feasible. This could be due to the fact that not enough spare capacity is left after the deployment of maintenance replacement paths or because a given restoration mechanism is not adaptive enough to changes in the spare capacity layer to find feasible paths.
2. Failure on another span used for the maintenance replacement paths: In this situation a failure occurs on a span that is currently supporting one or more of the paths used for maintenance replacement. The outcome of such a situation depends on the amount and distribution of remaining spare capacity in the network and may result in un-restorable fractions on both the maintenance span and failure span.

Under the equivalence of a Type 3 maintenance model to a restored span failure, and the simple transformation of a Type 2 maintenance action to an equivalent Type 3 action of lesser impact, we can go ahead and speak about maintenance state and failure state combinations as if two actual failures were being considered. Based on the dual span-failure restorability $R_2(i,j)$ defined by equation (4.1) of the previous chapter, we can define the *theoretical loss of restorability (or restorability risk) on a span i from maintenance on a given span m,* $L_m(i)$, in terms of the percentage of all working channels on span *i* plus span *m* that would be non-restorable if a failure occurred on *i* in the presence of the maintenance state on span *m*. $L_m(i)$ is expressed as follows:

$$L_m(i) \equiv \frac{N_m + N_i}{w_m + w_i} = 1 - R_2(m, i) \tag{5.1}$$

where N_m and N_i are the number of channels that are not restorable on span *m* and *i*, respectively. For each span taken as a maintenance span *m*, one can therefore define a vector of risk L_m whose components are the $L_m(i)$ for each other span i. As defined earlier in this chapter, the theoretical risk field caused by a maintenance action on span m is the set of all other spans i for which the theoretical loss of restorability $L_m(i)$ is not zero. Once the maintenance risk fields of all spans are quantified, an obvious use of them is to co-ordinate actions so that maintenance on a span i which is in the risk field of a current maintenance m, is deferred until m is complete.

5.3.3 Method for Evaluating the Theoretical Risk Fields due to a Span Maintenance

In this section we outline a computational procedure to define the theoretical field of risk of any given span maintenance. The procedure, fully described on Figure 5-2, computes risk fields with the following properties:

i. Type 2 or Type 3 maintenance actions are modelled. Span maintenance effectively deletes

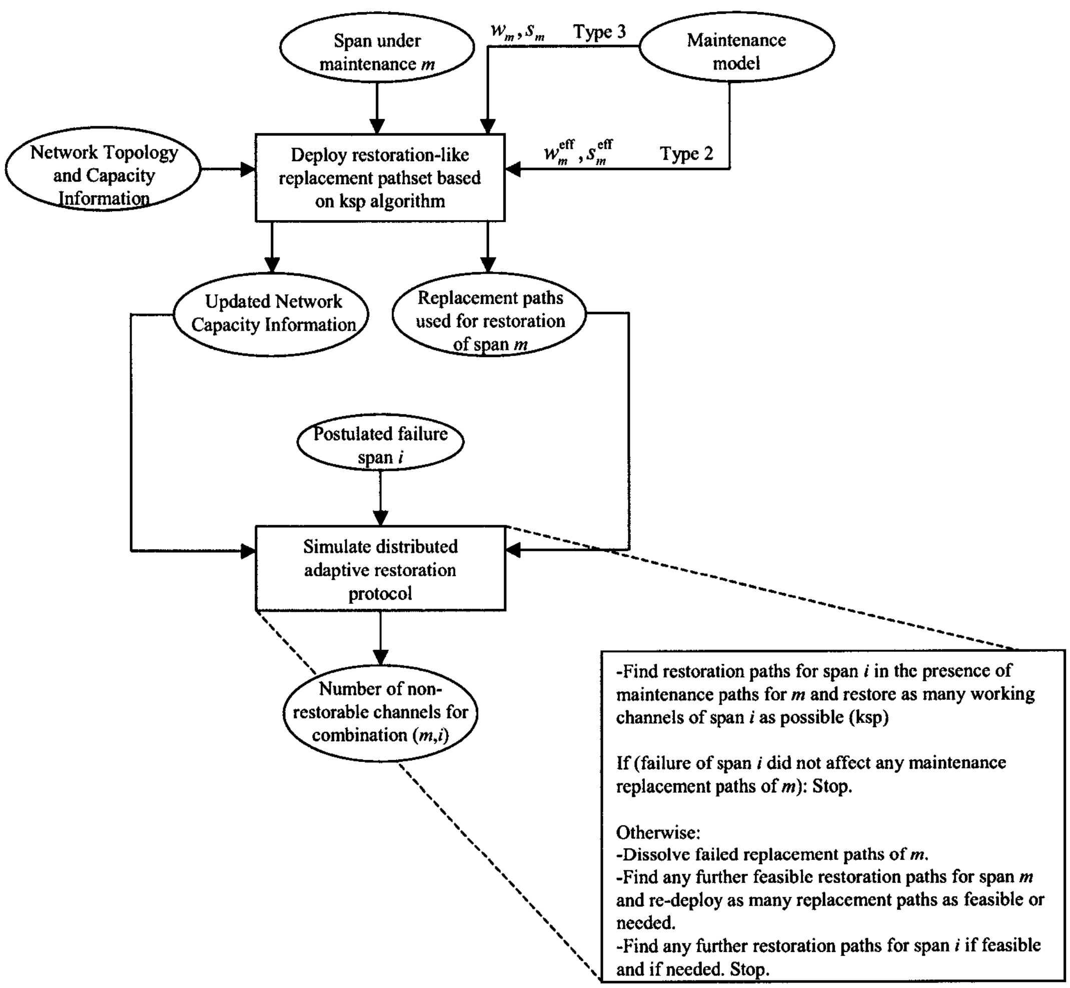

Figure 5-2 Algorithm for calculating theoretical risk fields

some or all spare channels on the span and may force some or all working channels onto restoration paths depending on the maintenance model and the working and spare capacities.

ii. We assume that the restoration mechanism is intelligent and adaptive to the extent that restoration of a span *i*, while span *m* is under maintenance, recognizes and adapts to any preceding consumption of spare capacity on the maintenance span and on the other network spans by the maintenance state.

iii. Effects of failure on both the failure and maintenance spans are considered assuming restoration Model 3 presented in Section 4.4.2: If a failure of span *i* severs one or more maintenance replacement paths for span *m*, a restoration action is triggered for span *i* followed by a restoration effort for any lost working channels from span *m*. The restoration mechanism will try to find new restoration paths between the end-nodes of the maintenance span over spare channels not used by the restoration action for the failure itself. This includes a release of surviving spare capacity on maintenance-replacement paths for the span *m* and use of the restoration process in the context of span *m* to find new maintenance replacement paths. If restoration of the failure span *i* was incomplete, any further restoration paths for span *i* that may be feasible after dissolution of the failed replacement paths from span *m*, are then sought, completing the restoration reaction to the failure.

Figure 5-3 shows an example of a calculated risk field for maintenance of the indicated span "M". The test network used is 22n41s1, presented in Section 1.4.1 and detailed in Section A.14 of Appendix A. The shape of the network has been slightly modified here to allow three graphs to be

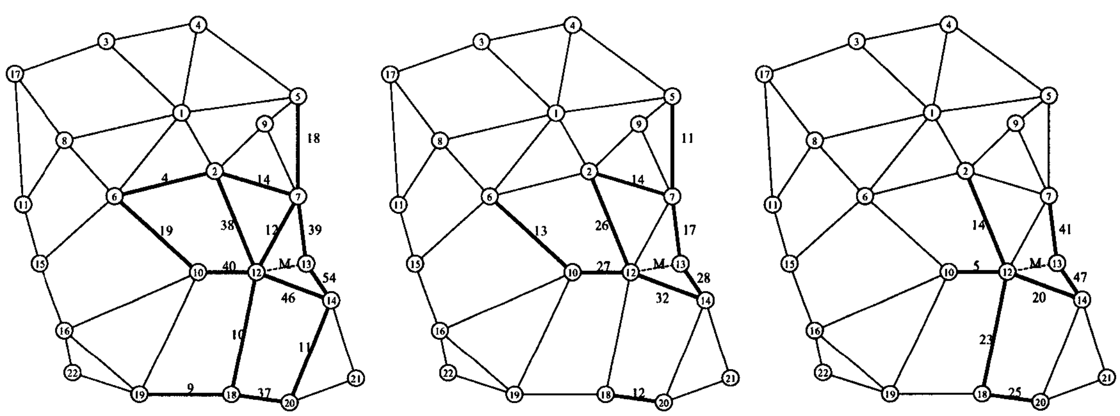

(a) Non-modular, Type 3, 14 spans, 292 channels

(b) Non-modular, Type 2, 9 spans, 157 channels

(c) Modular, Type 2, 7 spans, 129 channels

Figure 5-3 Risk field maps (22n41s1)

displayed side by side but the experimental data used was the one described in Appendix A. Demands are routed via shortest-path routing followed by an optimal spare capacity design using Herzberg's SCP formulation presented in Section 2.5.3 with a hop limit of five for eligible restoration routes. The capacity design is at the theoretical minimum of spare capacity to still be restorable. Real networks are rarely this minimally capacitated and will have smaller risk fields as a consequence. In Figure 5-3 (a) and (b) the identical unit-channel integer capacity design applies but the maintenance model differs. In Figure 5-3 (c) the corresponding minimum-capacity-cost modular capacity design is used, obtained with the MSCP formulation presented in Section 2.5.4 with a hop limit of 5 and module sizes 12, 24, and 48. The non-modular capacity design (Figure 5-3 (a) and (b)) has an average of 34.3 working and 18.4 spare channels per span and 53.5 percent redundancy. Figure 5-3 (c) has a redundancy of 73 percent due to modularity effects. The numbers on each non-maintenance span are the entries of the L_m risk vector portrayed as the percentage of the working channels on those spans plus the maintenance span that would not be restorable if the indicated span fails during the maintenance state. The maintenance span chosen for the illustration is a typical span for this network in that it has 32 working channels and 15 spare channels. The computation of risk fields for all spans based on simulation of all combinations of maintenance span and failure span combinations for Figure 5-3 takes only a fraction of a second.

Figure 5-3 (a) shows the risk field for Type 3 (worst-case) maintenance on the maintenance span. Figure 5-3 (b) shows the corresponding Type 2 maintenance result in the same network. Under a Type 2 maintenance model for this span, 15 of the 32 working channels are rolled to protection on the same span and the remaining 17 working channels are re-routed via replacement paths through the network in the vicinity of the maintenance span. Under Type 3, all 32 working channels are re-routed. The extent and magnitude of the risk field behaves as one might expect from first principles. In Figure 5-3 (a), the risk field contains 14 spans and a total of 292 working channels are non-restorable if they should fail (in a complete span cut) during the considered span maintenance. Under Type 2 maintenance the extent of the field drops to 9 spans and the total magnitude to 157. If the inevitable margin of extra capacity present in a modular design is available for re-routing, Figure 5-3 (c) shows a further contraction of the risk field to 7 spans and a magnitude total of 129 channels.

5.4 Experimental Effects of Modularity, Hop-Limit and Bi-Criterion Design

5.4.1 Test Networks

A series of experimental trials was designed to learn about the typical extent of risk fields due to maintenance in span-restorable mesh networks having a range of relevant design properties. Three basic topologies were used for these experiments, 11n21s1, 15n28s1, and EuroNet, presented in Section 1.4.1 and detailed in Appendix A. The test demand patterns used are the ones corresponding to each of these test networks and detailed in Appendix A. For each basic topology and demand matrix, five types of capacity design were produced for testing risk field properties. Three of them were produced using the MJCP design formulation presented in Appendix B, using a single module type of size 1. With the MJCP formulation working path routing and spare capacity placement are jointly optimized. Using a single module type 1 makes MJCP equivalent to the non-modular JCP formulation. The networks designs with JCP are denoted JCP "*H*," where *H* is the hop-limit of the corresponding design. The three (or in some cases four) JCP designs are complemented with a non-joint but modular capacity design at a hop limit of five, imported from the work in [DoG00]. The modular capacity designs use the same demand routing from the corresponding JCP-5 non-modular designs but min-cost modular capacity decisions are made such that the working and spare total placement on each span is modular. Capacity module sizes were 12, 24, and 48 channels. The relevance of the modular designs is that modularity should tend to reduce the extent of the risk fields. Finally, a corresponding design for each network was imported from a recent study of a bi-criterion optimization approach that allows a trade-off between total capacity and average restoration path lengths [DGB01]. The bi-criterion networks tested here were those at which the greatest tightening of restoration paths has occurred before any increase in capacity investment is made in the method of [DGB01] and are non-modular. The bi-criterion designs are also relevant to this study of risk fields as the "tightness" of the restoration paths should also affect risk fields.

5.4.2 *Methods and Results*

Each span in each test network was tested for its effect as a maintenance span under Type 2 and Type 3 models for the maintenance process. For each span, in the context of a maintenance span, m, we compute the theoretical loss of restorability $L_m(i)$ for each other span i as a prospective failure span during the maintenance state. This is done for both maintenance models. We summarize the results for each test case and each maintenance model by computing the average total theoretical risk $L(m)$ for each prospective maintenance scenario, defined as:

$$L(m) \equiv \sum L_m(i) \qquad (5.2)$$

These data are given as the absolute average number of channels in total that undergo a risk of restorability loss due to each maintenance action. For illustration, a $L(m)$ of 100 would be equivalent to a risk field of 10 spans in extent each exposed to a 10-channel restorability path shortage if a failure occurred on the span while in a maintenance state. The absolute values are not of primary significance since this depends on the total demand of the test network design, but it is not clear which case could be justified as a normalizer, so we present the absolute data for comparative inspection. A separate analysis of the span-by-span fractional restorability loss percentages (Figure 5-4) addresses the question of how significant a risk exposure these absolute totals represent. The second summary characteristic of each trial in Table 5-1 is the average logical extent of the risk field, $S(m)$. This is the average number of spans undergoing any non-zero risk over all maintenance span scenarios. The results are summarized in this way in Table 5-1 on p. 115.

While Table 5-1 gives the extent and absolute magnitude averages for each risk field, Figure 5-4 shows the results from a converse standpoint which is the distribution of risk exposure on non-maintenance spans in the network due to all possible single-span maintenance actions in its network. The data in Figure 5-4 pools individual trial $L_m(i)$ values for the modular designs with a restoration hop limit of five over all networks for each maintenance type. The three test networks provide a total of 86 instances (sum of the number of spans on each network) of maintenance-spans to test for the risk imposed on each other span in the test case.

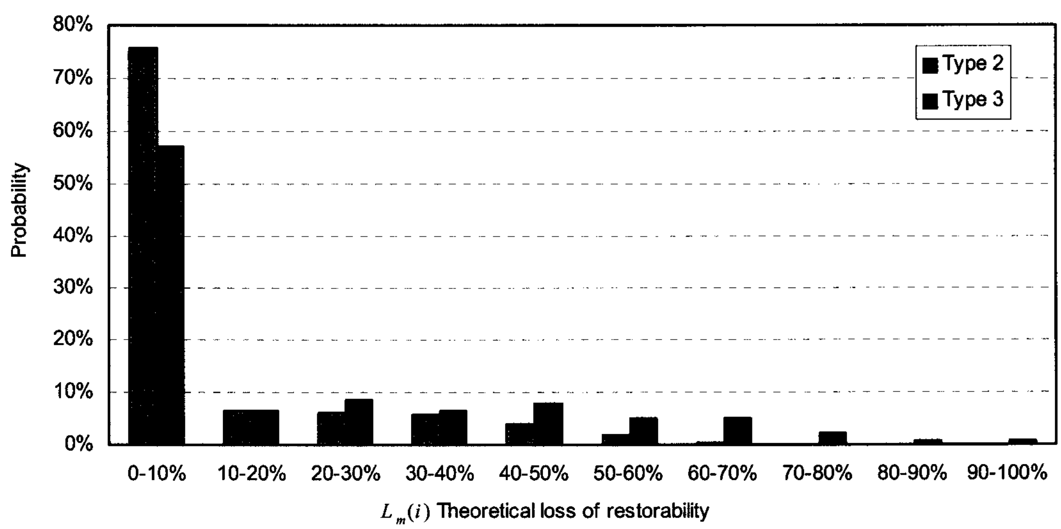

Figure 5-4 Histogram of individual span restorability risks

The total population of $L_m(i)$ risk exposure values for Figure 5-4 is therefore:

$$21 \cdot 20 + 28 \cdot 27 + 37 \cdot 36 = 2508 \text{ instances.}$$

The modular designs are chosen for this exhibit of results because they are the most realistic family of capacity designs. Figure 5-4 is a histogram which shows the probability of experiencing

Table 5-1: Average risk field extent (spans) and total risk exposure (channels) in trials

Topology	Maintenance Model	JCP 4		JCP 5		JCP 6		JCP 7		Modular		Bi-criteria	
		$L(m)$	$S(m)$	$L(m)$	$S(m)$	$L(m)$	$S(m)$	$L(m)$	$S(m)$	$L(m)$	$S(m)$	$L(m)$	$S(m)$
11n21s1	Type 3	247.7	14.2	294.7	16	330.1	16.7	n/a	n/a	152.1	9.2	294.1	16
	Type 2	128.4	10.2	169	12.4	202.8	13.5	n/a	n/a	57	5.2	169.1	12.5
15n28s1	Type 3	352.4	14.5	401.8	15.7	406.2	16	490.4	18.3	263.2	11.7	401.1	16
	Type 2	172.5	9.3	209.1	10.6	215.8	11	274.2	13.5	99.7	6	208.9	10.6
EuroNet	Type 3	530.6	18	577.2	19.5	713	23.4	535.8	20	398.5	15.1	586.9	19.8
	Type 2	297.3	12.8	323.3	13.5	420.7	16.9	276.2	13.5	181.2	8.9	330.5	13.9

a restorability risk of the percentage magnitude shown in each x-axis category. As produced the results inherently assume that all spans are equally likely maintenance candidates. Reading one point for illustration Figure 5-4 says, for instance, that *under Type 2 maintenance there is 75 percent chance that any given other span experiences zero risk.* Or, reading from, say, the sixth x-axis category as a further example, we would say that *under Type 3 maintenance there is 5 percent chance of any given other span being exposed to a risk of 50 to 60 percent loss of restorability.*

5.4.3 Discussion of Results

First, and not surprisingly, Table 5-1 shows that in all trial networks the Type 3 maintenance model engenders much greater risk of restorability loss than with Type 2. Second, in comparing both risk field extent and total risk exposure as the design eligible restoration-route hop-limit decreases from 6 down to 4, we see the expected contraction of extent and risk magnitude, due to changes in the design that increasingly keep restoration flows (and maintenance replacement paths) closer to home. Only the results for EuroNet JCP 7 break the trend and are more difficult to explain in that we see a lower risk field extent and magnitudes than JCP 6 for that network. Being jointly optimized capacity designs, we surmise that this must be a network-specific effect, i.e., that in that particular topology at a hop limit of seven, some significant shift in the basic solution structure is enabled which is enough to overcome the otherwise general trend in $S(m)$ and $L(m)$ as the design H increases.

Some important practical observations can be made from Figure 5-4. First, under 2508 instances of spans at potential risk due to 86 instances of Type 2 span maintenance, in three different networks, there *was not one case of 100 percent risk exposure* as in rings. In fact, there was only a 25 percent chance of any risk and the risk exposure was less than 50 percent of the number of working channels on the span in 99.6 percent of the trial cases. Moreover, under Type 2 no spans ever experience risk exposures over 70 percent. Even under Type 3 maintenance, 91.4 percent of the time the risk imposed on other spans is below 50 percent. Clearly this means that *priority signal paths could be given an assurance, not possible in rings, of zero risk of restorability loss due to maintenance.* As expected there are more instances of higher-risk situations under Type 3 maintenance, including a small percentage of 100 percent restorability risks. The results contain exactly 20 cases where $L_m(i)$ was 0 under Type 3, but every one of these is easily explained by its correspondence to a degree-2 node in the respective test network. Under Type 3, maintenance to one span at a degree-2 node, the opposite span at the node undergoes a 100 percent restorability risk, as in a ring, purely from topological considerations. However, this is only true for the mesh under Type 3 maintenance. Under Type 2 maintenance there are no $L_m(i)$ of 1. In contrast, in a ring, $L_m(i)$ is 1 for either Type 2 or Type 3 maintenance models.

5.5 Considerations for Simultaneous Maintenance Actions

In this section we give some preliminary considerations about risk-related issues of conducting more than one maintenance action at a time on different parts of the network. A central concern is on ways to manage multiple maintenance actions so that the total risk exposure is not compounded to be more than the sum of the individual maintenance actions risk magnitudes. A simple solution of the span maintenance scheduling problem would be to schedule each maintenance one after the other so as to minimize the consequences of possible true span failures happening while these maintenance actions are being conducted, but in a large continental or even metro network, workforce efficiency and total maintenance intensity might both require that simultaneous maintenance actions be routinely performed.

5.5.1 The Notion of Maintenance Orthogonality

Obviously if two maintenance actions generate completely disjoint risk fields, then it should be possible to conduct them at the same time without running a higher risk of outage than if they are conducted one after the other. For example, if a span maintenance conducted in the north-west part of the city affects only spans in that part of the network and a span maintenance conducted in the south-east part of the city affects only the spans in that part of the network, then doing the maintenance on the span in the south-east while the span in the north-west is also under maintenance will not aggravate the risk in the north-west region and vice-versa. But more generally the risk fields may not be disjoint. However, in any case where the total risk of the combined action is not greater than sum of the risks of the same actions in series then we say the two maintenance projects are mesh-orthogonal. A way to define this approximate notion of orthogonality is:

A set of spans $\boldsymbol{U}$ is said to respect the mesh *maintenance orthogonality* condition if:

$$\sum_{i \in \boldsymbol{S}} L_U(i) \cdot \Delta T_U = \sum_{m \in \boldsymbol{U}} \sum_{i \in \boldsymbol{S}} L_m(i) \cdot \Delta T_m \tag{5.3}$$

where $L_U(i)$ is the number of channels at risk on span i under the simultaneous maintenance of all spans in $\boldsymbol{U}$, ΔT_U is the elapsed time in the composite maintenance scenario, $L_m(i)$ is as defined earlier, and ΔT_m is the elapsed time in individual maintenance scenario m.

Conversely the difference between terms in (5.3) is a measure of the compounded risk of doing the actions in parallel. To compute the test of orthogonality as defined in (5.3), we need a method to compute the left-hand side, i.e. the composite risk of the simultaneous maintenance states. Note the role that elapsed time can play as well as the more subtle topological effects that are involved. If the parallel actions can be done in less time, compound risk may be effectively avoided. Obviously it is preferable to conduct parallel operations that are truly topologically orthogonal, however. Obvious extensions to the procedure in Figure 5-2 can be used to compute $L_U(i)$ values by conducting the test of span i restorability in the presence of the set of $|U|$ specific sets of maintenance replacement paths and span spare capacity withdrawals.

5.5.2 *Managing Avoidable Risk for Priority Paths*

Unless we were willing to pay the quite significant price for a complete dual-failure restorable capacity design (as it will be seen in the following chapter), we have seen above that a certain risk due to maintenance is unavoidable. But what about trying to do so just for a subset of premium service customers? Let us here consider that in a network where much maintenance activity is required, it may be unavoidable that multiple simultaneous actions have to be undertaken without any assurance of orthogonality in the above sense. In such circumstances a reasonable viewpoint might be to attempt to at least ensure the avoidance of compound risk for priority paths. Thus, the idea here is to consider a selective measure of the total risk for a class of priority paths crossing each span. This would lead to a corresponding definition of maintenance orthogonality with respect to priority paths. Ensuring orthogonality of maintenance risks for a subset of premium services may be much more feasible than trying to achieve it for all working capacity.

A proposed measure of priority-services total risk caused by a maintenance action is the following:

$$\text{Priority Risk }(m) = \sum_{i \in S} (\max(0, P_m(i) - w_p(i)) + \max(0, P^*_m(i) - w_p(m))) \tag{5.4}$$

where $P_m(i)$ is the number of paths found for restoration of channels on span i when span m is in the maintenance state, $P_m(i)$ is the number of paths found for replacement of channels on span m when it is under maintenance and span i fails, and $w_p(k)$ is the number non-priority service paths crossing any span k.

Based on the above definition of the priority risk, and similarly to what has been introduced for the general definition of the risk, a total "serial risk" $\text{Risk}^+(U)$ can be defined for multiple maintenance actions when done individually in series as well as a total "parallel risk" $\text{Risk}^{||}(U)$ for multiple maintenance actions done in parallel. The difference in these two computed risk measures is a more specialized assessment of whether a proposed set of simultaneous maintenance actions will deleteriously expose priority paths to compound risk.

As seen in the previous sections, the risk values calculated for Type 2 are almost always below 50 percent and always below 70 percent. Consequently, if priority channels never represent more than 30 percent of the channels on any span, then the priority risk for *single* maintenance actions is already null. The priority risk for parallel maintenance actions, however, might be non-zero for some maintenance combinations. Calculation of the total parallel risk $\text{Risk}^{\|}(\boldsymbol{U})$ will identify such situation for avoidance.

5.6 Summary and Conclusion

The work presented in this chapter has developed and applied methods for the analysis of the risk to restorability engendered by maintenance actions of spans of a mesh-restorable network. It is now clear that while rings have a contained risk field of $(S-1)$ spans—the spans on the same ring as the maintenance span—that are at 100 percent restorability risk, mesh networks exhibit a more extended risk field but with much lower risk of restorability loss on any one span. From the fact that over 90 percent of spans have restorability risks under 50 percent for either Type 2 or 3 maintenance models, it seems valid to conclude that mesh networks would be able to support a large fraction of priority services which can be *guaranteed zero risk to their normal restorability* due to the operator's maintenance activities. The assumptions are only that maintenance actions are done one span at a time and that priority status is taken into account in the restoration process.

6. Mesh-Restorable Networks with Enhanced Availability Properties

6.1 Introduction

As already explained in previous chapters, dual-failure situations, or the equivalent, are not as rare as might be thought in an extensive national or regional optical network. Certainly, the bulk of transport requirements will be well served by designing for them to withstand any one failure at a time, but the idea has been raised by industry people that there is a significant number of other applications and customers for whom there really is a desire to support dual-failure restorability. Considering rates of fibre cuts in some networks, span maintenance operations which create situations similar to dual-failure (as seen in the previous chapter), and shared-risk link groups, we are motivated to study and understand issues and phenomena surrounding dual-failure restorability in mesh networks and to see what steps can be taken to enhance or even design for certain specific abilities to withstand dual failure scenarios.

In this chapter, we present some extensions of the mesh network capacity design formulations presented earlier in the thesis that enhance the dual-failure restorability of the designs. The chapter is based primarily on adaptation and extension of collaboration in [ClG02b].

6.1.1 Goals of this Section

In chapters 3 and 4, we have presented a method for analyzing the mesh-restorable networks from a service path availability standpoint considering all possible dual-failures as the primary contributor to unavailability. The aim in these chapters was to develop a means of determining the availability of service paths through mesh-restorable networks that are designed for full restorability to single failures and may involve an active, state-adaptive, recovery mechanism, but are then presented with dual-failure situations. These chapters were thus essentially based on a work of *analysis* whereas the work presented in this chapter is to be one of design *synthesis* to specific new goals. One of the most striking findings of chapters 3 and 4, was that a span-restorable mesh network that is designed for full restorability to any single span failure is rather naturally able to protect a high average proportion of working capacity against *dual* span failures. This was found to be especially true when the restoration mechanism is adaptive to changes in the spare capacity layer of any prior failure when the second failure arises. It was also noted that the dual-failure restorability increases with the size of the network. These basic findings give considerable motivation to the idea of trying to specifically harness or otherwise exploit this basic property by design.

However, an important consideration about the high average-case dual-failure restorability found in Chapter 4 is that while this is of obvious benefit to the *overall* network availability, there is in principle no way to target or otherwise structure the dual-failure restorability so that specific paths, such as those of premium customers, would be the ones assured to benefit. Therefore, we are motivated to set the next two goals:

1. How can we take advantage of the naturally high restorability of mesh networks to dual failures to specifically serve a new category of extreme high availability service paths that would enjoy full dual span-failure restorability?
2. How can we specifically design the network capacity to serve any given mixture of single failure and dual failure restorability requirements, designated on a path-by-path basis, at minimum total cost?

The work presented in this chapter combines the concepts and observation of the previous chapters with ILP techniques to address these questions and to use the resulting capability to observe the overall trade-off between availability and network capacity. We also provide new design methods for support of restorability-related service level agreements.

6.1.2 Prior Work on Availability-Oriented Design

There has been prior work both on methods to design for complete dual failure restorability. As early as 1992, Sakauchi et al. [SOH92] gave a formulation for spare capacity design of a span-restorable mesh network to withstand all possible dual span failure scenarios, although they did not produce results with the model.

Work on the issue of making a network restorable to dual-failures was also conducted by H. Choi, S. Subramaniam and H.-A. Choi in [CSC01] and [CSC02]. The approach taken by these authors considers link protection methods – which we would call here span protection or localized protection. In their first method, each span is assigned two span-disjoint backup paths. After a first failure, the name of the failed span is broadcast to all nodes so that they know what backup path to use in case of a second failure. If the failure of a second path hits the backup path used for restoration of the first failure, the restoration of the first failure is switched to the secondary backup path – a process which can be quite slow. In a second method, each span is still assigned two backup paths but when a second failure hits the backup path used for the first failure, that backup path is itself restored locally around the second failed span using either the second failed span's primary backup path or secondary backup path (at least one of the two is guaranteed not to cross the first failed span). This concept is similar to the partly adaptive restoration behaviour (Model 2) introduced in Section 4.4.2 in that both methods give "working" status to any spare capacity used for restoring a failed working unit. These first two methods require that the first failed span be known by all network nodes to allow them to decide what backup path to use in case of a second failure. In contrast, the third method they present does not require the identification of the first failed span. This third method is based on finding a single backup path for each span such that no span *i* has a backup path on which one of the spans has itself a backup path containing span *i*. This way any second span failure could be fully restored (including any backup flow) using the backup path for that second failed span. [CSC01] and [CSC02] then investigate the feasibility of finding the set of backup paths corresponding for their third method. Although they claim that full dual failure restorability could be obtained with a "modest increase in backup capacity," no study of capacity requirements is provided.

In contrast to the work of [SOH92], [CSC01] and [CSC02], the study presented in this chapter is based on a very concrete and detailed analysis of the required capacity to make a network restorable to dual span-failures. The results of restorability requirements presented here are determined using powerful optimization tools therefore they are guaranteed to be optimal solutions. As far as we know, this work was the first to provide quantitative results for how much extra spare capacity would be required for complete dual failure restorability. Also, the restoration mechanisms considered are not limited to one or two pre-determined backup paths per span but are based on dynamic mesh-based restoration as portrayed in Figure 2-17 of Section 2.4.6.

6.1.3 Link to Multiple Quality of Protection Concept

The work presented here can also be linked to the topic of multiple quality of protection (QoP). What distinguishes the present work from prior considerations of multi QoP in both industry standards forums and in the research literature is that we enable a new QoP class which stands above the previously considered hierarchies of gold, silver, bronze, etc. For instance gold usually means assured single-failure restorability, silver is best effort, bronze is un-protected, and there may be another lower class of preemptible economy class services as well. But here we find an economical way to design-in support for what could be called a platinum service class that stacks above the existing QoP paradigms, in the sense illustrated in Figure 6-1. The particular contribution made by the end of the paper is thus of showing how a new super-high availability service class can be added to the range of possible QoP classes and that it can be supported with remarkably little additional cost within span restorable networks that are already efficiently designed only to support single failure restorability, i.e., the typical gold QoP class. [IMG98][DoG00]. Some work on the more traditional multi QoP paradigm is presented in Chapter 8.

Figure 6-1 New and existing protection options

6.1.4 Outline of Chapter

In Section 6.2 we present three extensions of the SCP formulation presented in Section 2.5.3 that enhance the dual-failure restorability of span-restorable mesh networks. As a reference case, the first extension is for complete dual-failure restorability. This will show extremely high capacity penalties to support 100 percent dual-failure restorability. The second design model will allow us to explore the trade-off between capacity cost increase and dual-failure restorability improvement starting from a single-failure restorable design. Finally, a third design strategy that supports multiple-restorability service class definitions at minimum cost will be presented.

In Section 6.3, we will discuss experimental results obtained with the three formulations and the test networks presented in Section 1.4.1.

Section 6.4 will present a strategy for a particular type of network survivability measure we call "first-failure Protection, second-failure Restoration." With such a strategy, paths with high availability requirements can not only be offered the very high availability that results from the guarantee to be protected against all dual span-failures, but they can also benefit from the fast restoration that is traditionally only associated to protection schemes.

6.2 Optimizing the Capacity Design for High Availability

We now develop three specific optimization models through which we can explore and understand the trade-offs and opportunities to design for dual failures and to specifically structure the dual failure restorability so that it is targeted onto desired paths. As mentioned, Chapter 4 showed that surprisingly high levels of dual-failure restorability tend to arise simply as a side effect of placing spare capacity to assure the restorability of single failures, a condition we can denote as "$R_1 = 1$." In other words "$R_1 = 1$" tends to beget high average R_2 without further special effort of any type. This prompts a natural first exercise which is to simply see how much it would cost to directly design for restorability to all dual-failures. This we will refer to as the problem of *dual-failure minimum capacity* design (DFMC). As it will be seen in Section 6.3, although the average

$R_2(i,j)$ may be relatively high, it turns out that achieving strictly an $R_2(i,j)$ of 1 for all (i,j) failure combinations is very expensive. The next question is therefore to see just how high an R_2 can be achieved for a given set limit on additional capacity expenditure over the "$R_1 = 1$" condition. This is called the dual-failure maximum restorability (DFMR) problem. These preliminary investigations lead to the final goal of deliberately placing capacity at minimum cost to serve certain paths at "$R_2 = 1$" and others at only "$R_1 = 1$." This is called multi-restorability capacity placement design (MRCP). To summarize we will now consider each of the following problems:

i. Dual-failure minimum capacity design (DFMC): What is the minimum spare capacity assignment that guarantees full dual-failure restorability?

ii. Dual-failure maximum restorability (DFMR): Given a finite budget of total spare capacity, find the spare capacity assignment that maximizes R_2 in the network, given the available spare budget.

iii. Multi-restorability capacity placement design (MRCP): Find the minimum total (spare plus working) capacity assignment that satisfies a mixture of end-to-end path restorability objectives (R_0, R_1 or R_2) for each demand.

In the first two problems DFMC, DFMR, the working demands are first shortest-path routed, generating the working capacity values w_i, which are inputs to the problem. MRCP, however, is a type of joint optimization problem and is based on the JCP formulation presented in Appendix B.

6.2.1 Dual-Failure Minimum Capacity (DFMC)

This DFMC formulation finds a minimum total spare capacity assignment that guarantees full restorability of all dual span-failure scenarios. This formulation will tell us what the minimum "price" is for reaching full dual-failure restorability. For obvious reasons, a feasible solution to this problem cannot be found for a network with degree-2 nodes or 2-edge cuts of the network graph. Therefore our tests of this formulation are limited to the test networks of Section 1.4.1 having a graph topology that qualifies. In practice few transport networks have cuts of only 2 edges not involving degree-2 nodes but many have degree-2 nodes. For example, among all the test networks of Section 1.4.1, only 12n20s1 has 2-edge cuts that are not linked to degree-2 nodes. For all the other networks with some 2-edge cuts, these cuts are always linked to degree-2 nodes. In practice DFMC can therefore be considered a formulation primarily for the assurance of full dual-failure restorability on the "mesh backbone" component of the overall transport network. In this context the formulation can easily be applied by logical removal of a degree-2 node between spans i and j, and assertion that the working capacity $w_{i,j}'$ of the single logical edge arising from the degree-2 node removal is given by the maximum of w_i and w_j.

There are no new parameters or variables to introduce at this stage, other than:

$f_{i,j}^{p}$ Restoration flow assigned to p^{th} restoration route of span i when span j has failed simultaneously (integer), $\forall i, j \in \boldsymbol{S}$, $i \neq j$

$$\textbf{DFMC: } \text{Minimize } \sum_{k \in \boldsymbol{S}} C_k \cdot s_k \tag{6.1}$$

$$\sum_{p \in \boldsymbol{P}_i} f_{i,j}^{p} = w_i \quad \forall i, j \in \boldsymbol{S},\ i \neq j \tag{6.2}$$

$$f_{i,j}^{p} \leq C_{\infty} \cdot (1 - \delta_{i,j}^{p}) \quad \forall i, j \in \boldsymbol{S},\ i \neq j,\ \forall p \in \boldsymbol{P}_i \tag{6.3}$$

$$\sum_{p \in \boldsymbol{P}_i} f_{i,j}^{p} \cdot \delta_{i,k}^{p} + \sum_{p \in \boldsymbol{P}_j} f_{j,i}^{p} \cdot \delta_{j,k}^{p} \leq s_k, \quad \forall i, j, k \in \boldsymbol{S},\ i \neq j,\ i \neq k,\ j \neq k \tag{6.4}$$

The constraint set (6.2) ensures full restoration in all dual span-failures. Constraints (6.3) use an arbitrarily high capacity constant, C_{∞}, to ensure that span *j* can support restoration flow when it is not part of the failure scenario, but that it will not be used for restoration of span *i* in the *(i,j)* scenario. Constraints (6.4) ensure that there is enough spare capacity on each span of the network to support all the restoration flows assigned for the restoration of any dual span failure. No explicit statement of the single-failure restorability requirement is needed. This is implicitly provided for in specifying the restorability of all *dual* failures.

The AMPL model that was developed to solve the DFMC formulation is detailed in Section D.3 of Appendix D.

6.2.2 *Dual-Failure Maximum Restorability (DFMR)*

The second formulation is the converse problem of maximizing the achievable R_2 level with a *given* total spare capacity investment. The same parameters and variables apply as defined so far with the addition of the following parameter:

B Total budget available for spare capacity

and the variable $N(i, j)$ with the same definition as given in Section 4.1.

$$\textbf{DFMR: } \text{Minimize } \sum_{i \in \boldsymbol{S}} \sum_{\substack{j \in \boldsymbol{S} \\ j \neq i}} N(i, j) \tag{6.5}$$

$$\sum_{p \in \boldsymbol{P}_i} f_i^{p} = w_i, \quad \forall i \in \boldsymbol{S} \tag{6.6}$$

$$N(i,j) = w_i + w_j - \left(\sum_{p \in \boldsymbol{P}_i} f_{i,j}^p + \sum_{p \in \boldsymbol{P}_j} f_{j,i}^p \right) \quad \forall i,j \in \boldsymbol{S},\ i \neq j \tag{6.7}$$

$$\sum_{p \in \boldsymbol{P}_i} f_{i,j}^p \le w_i \quad \forall i,j \in \boldsymbol{S},\ i \neq j \tag{6.8}$$

$$f_{i,j}^p \le C_\infty \cdot (1 - \delta_{i,j}^p) \quad \forall i,j \in \boldsymbol{S},\ i \neq j,\ \forall p \in \boldsymbol{P}_i \tag{6.9}$$

$$\sum_{p \in \boldsymbol{P}_i} f_i^p \cdot \delta_{i,k}^p \le s_k \quad \forall i,k \in \boldsymbol{S},\ i \neq k \tag{6.10}$$

$$\sum_{p \in \boldsymbol{P}_i} f_{i,j}^p \cdot \delta_{i,k}^p + \sum_{p \in \boldsymbol{P}_j} f_{j,i}^p \cdot \delta_{j,k}^p \le s_k, \quad \forall i,j,k \in \boldsymbol{S},\ i \neq j,\ i \neq k,\ j \neq k \tag{6.11}$$

$$\sum_{k \in \boldsymbol{S}} C_k \cdot s_k \le B \tag{6.12}$$

Constraints (6.6) ensure that every single span-failure is fully restorable. Equation (6.7) expresses $N(i,j)$ in terms of the other variables of the problem. By minimizing the sum of $N(i,j)$ over all dual failure scenarios, one is maximizing R_2 (through Equation (4.8) and thereby also maximizing the availability through Equation (3.49)). Constraints (6.8) ensure that the number of paths assigned to the restoration of a span in a dual span-failure scenario is at most equal to the number of working links to be restored. Constraint set (6.9) forces the exclusion of restoration routes for span i from using span j and vice-versa during the (i,j) scenario, while allowing use of span j for all other scenarios. Constraints (6.10) ensure adequate spare capacity for every single failure case. Constraints (6.11) is identical to (6.4) but here it serves to assure adequate spare capacity only for the dual failure scenarios that the budget-limited formulation *chooses to cover*. The aspect of selectivity, not present in DFMC, is effected through (6.8), which – in contrast to (6.6) – is not required to provide fully adequate restoration flows for all dual-failure scenarios. This is also why an explicit assurance of single failure restorability is present through (6.6) whereas it is not in DFMC. Constraint (6.12) imposes the budget limit on the total cost of spare capacity.

The AMPL model that was developed to solve the DFMR formulation is detailed in Section D.4 of Appendix D.

6.2.3 Multi-Restorability Capacity Placement (MRCP)

The last model is the most computationally challenging but perhaps also the most practically interesting and directly useful in the business of network operators. The idea is to explicitly design the allocation of spare capacity and the routing of demands to support *a multi-service restorability* *classification of the demands served.* The previous formulation maximizes the network R_2 as an average over all spans of the network. All demands served in the network will share in the increased availability in a general way. The MRCP formulation will, however, allow us to target and structure the high availability investment *specifically to the intended services or customers.* Every demand will receive a specific class of restorability guarantee on every span end-to-end over its route. In the first two formulations a given demand might cross a mixture of spans with various R_2 levels, whereas here we will be able to stipulate that a priority service will exclusively travel over spans where its restorability to dual span failures is guaranteed, thus deriving an overall R_2-guarantee for entire paths in the service class. More generally, to each demand unit we assign one of the following restorability service class designations:

i. R_0 restorability: no assured restorability – best effort in both single and dual failures

ii. R_1 restorability: assured restorability to any single span-failure, best effort for dual failures.

iii. R_2 restorability: assured restorability to any single and dual span-failure.

The formulation finds the minimum total cost of capacity (working plus spare) and the routing of each working demand so as to satisfy the restorability class-of-service of each demand. For MRCP a demand group is now defined as one or more demand units on the same O-D pair and in the same service class. Demand groups must be defined now to distinguish between demands of different service classes on each O-D pair. The sum of the demand quantities in all demand groups on an O-D pair here equals the prior single-service demand quantities for the corresponding O-D pair. r now indexes not just over all O-D pairs but over all demand groups.

The MRCP formulation uses the following new parameters:

$\boldsymbol{D}$ Set of demand groups, indexed by r,

d^r Size of the r^{th} demand group (integer number of individual demand units), $\forall r \in \boldsymbol{D}$

$\boldsymbol{Q}^r$ Set of eligible working routes for demand group r, $\forall r \in \boldsymbol{D}$

ψ_1^r Equal to 1 if r^{th} demand group is in the R_1 service class, 0 otherwise

ψ_2^r Equal to 1 if r^{th} demand group is in the R_2 service class, 0 otherwise

$$\textbf{MRCP:} \text{ Minimize } \sum_{k \in \boldsymbol{S}} C_k \cdot (w_k + s_k) \tag{6.13}$$

Subject to:

$$\sum_{q \in \boldsymbol{Q}^r} g^{r,q} = d^r \quad \forall r \in \boldsymbol{D} \tag{6.14}$$

$$\sum_{r \in D} \sum_{q \in Q^r} \zeta_i^{r,q} \cdot g^{r,q} = w_i \quad \forall i \in \boldsymbol{S} \tag{6.15}$$

$$\sum_{p \in P_i} f_i^p = \sum_{r \in D} \sum_{q \in Q^r} \zeta_i^{r,q} \cdot (\psi_1^r + \psi_2^r) \cdot g^{r,q} \quad \forall i \in \boldsymbol{S} \tag{6.16}$$

$$\sum_{p \in P_i} f_{i,j}^p = \sum_{r \in D} \sum_{q \in Q^r} \zeta_i^{r,q} \cdot \psi_2^r \cdot g^{r,q} \quad \forall i,j \in \boldsymbol{S},\ i \neq j \tag{6.17}$$

$$f_{i,j}^p \le C_\infty \cdot (1 - \delta_{i,j}^p) \quad \forall (i,j) \in \boldsymbol{S}^2,\ i \neq j,\ \forall p \in \boldsymbol{P}_i \tag{6.18}$$

$$\sum_{p \in P_i} f_i^p \cdot \delta_{i,k}^p \le s_k \quad \forall i, k \in \boldsymbol{S},\ i \neq k \tag{6.19}$$

$$\sum_{p \in P_i} f_{i,j}^p \cdot \delta_{i,k}^p + \sum_{p \in P_j} f_{j,i}^p \cdot \delta_{j,k}^p \le s_k,\ \forall i,j,k \in \boldsymbol{S},\ i \neq j,\ i \neq k,\ j \neq k \tag{6.20}$$

Constraint set (6.14) ensures that each demand group is fully served. Constraint set (6.15) ensures that there is enough working capacity on each span to support the routing of demands. Constraint set (6.16) ensures that adequate restoration flow exists for each single-span failure affecting demands in the R_1 or R_2 service class. Constraint set (6.17) ensures that adequate restoration flow exists for each dual-failure scenario affecting demand groups in the R_2 service class. The rest of the constraints are as previously explained.

6.3 Experimental Results

The three formulations were solved on the test networks presented in Section 1.4.1. Each formulation was implemented as an AMPL model and solved using the CPLEX 7.5 MIP solver on a 4 × 900 MHz Sun UltraSparc III processor running the Sun Solaris Operating System 2.6 with 16 GB of RAM. All DFMC and DFMR problems solved in at most a few hours; most of them were solved in seconds. Good feasible solutions to the MRCP problems were typically found in a few minutes, but complete solution could take several hours to a maximum of 2 days for the biggest test network. All results are based on a full CPLEX termination or a MIPGAP of under 0.001 (i.e. solutions are provably within 0.1 percent of optimal.) for all networks except for 30n60s1, for which the MIPGAP had to be increased to 0.005.

In all cases, problems were formulated with the twenty shortest restoration routes based on distance and included any restoration route that had the same length as the twentieth shortest one. In MRCP problems, the sets of eligible working routes were composed of the five shortest routes by distance between each O-D pair and included any additional route that had the same length as the fifth shortest route. The sets of eligible restoration routes for each span were composed of the ten shortest routes and included any route with the same length as the tenth shortest one. The reason for restricting the number of eligible restoration routes with the MRCP formulation was to make it possible to solve the problem reasonably fast for all test networks including the largest ones (the maximum was two days) and use the same parameters for all test networks. There is no problem with results not being comparable with previous results since the MRCP formulation is a joint formulation, whereas DFMC and DFMR are non-joint, so it would not have been able to compare the results of MRCP with results of DFMC and DFMR anyway. In any case, the goal of the MRCP formulation is to demonstrate the possibility of serving multiple classes of protection in mesh networks and observe the link between the proportion of the different protection classes and the capacity requirements. Results of MRCP will therefore not need to be compared to any previous results. In case R_0 demand groups are present in the MRCP problem, these would always be explicitly routed over shortest paths because this is optimal when no corresponding investment in spare capacity is made for them.

Table 6-1 on p. 130 shows the results obtained with the DFMC formulation. In this table, three factors are compared between the capacity requirements of the SCP single-failure restorable designs and of the DFMC dual-failure restorable designs. These are: total spare capacity, total capacity and total capacity cost. For these three factors, the Table 6-1 gives the total for the SCP design, the total for the DFMC design and the relative increase of the DFMC design compared to the SCP design. Looking at the total spare capacity results, it appears that the amount of spare capacity that is needed for dual span-failure restorability is between 1.91 and 3.09 times higher than the amount needed for single failure restorability. These somewhat surprising results indicate that although high R_2 levels arise just by designing a network for full single span-failure restorability, the price strictly to assure full dual span-failure restorability is very high. Achieving the "last few percent" in R_2 is extremely expensive. It is also interesting to observe how much this increase in spare capacity affects the total capacity requirement of the networks. Going from the single failure restorability constraint to the dual failure restorability constraint requires between 35% and 92% of total capacity increase. Since the SCP and DFMC formulation optimize total capacity cost, it is also interesting to see how much the capacity cost is increased. For this measure, the increase from SCP to DFMC ranges between 43% and 103%.

Table 6-1: DFMC Results

Test Network	Total spare capacity			Total capacity			Total capacity cost		
	SCP (R_1)	DFMC (R_2)	Relative incr. (R_2/R_1)	SCP (R_1)	DFMC (R_2)	Relative incr. (R_2/R_1)	SCP (R_1)	DFMC (R_2)	Relative incr. (R_2/R_1)
06n14s1	45	102	2.267	164	221	1.348	4753	6802	**1.431**
11n20s1	560	1548	2.764	1108	2096	1.892	146737.6	274571.1	1.871
11n20s2	188	540	2.872	483	835	1.729	60285.54	110238.1	1.829
Bellcore Mod.	282	540	**1.915**	744	1002	**1.347**	26546	39143	1.475
COST239	139	269	1.935	354	484	1.367	175850	254980	1.450
12n30s1	262	511	1.950	660	909	1.377	53954.29	80681.63	1.495
16n29s1	670	2068	**3.087**	1523	2921	**1.918**	169924.8	345096.7	**2.031**
16n38s1	368	809	2.198	1096	1537	1.402	91053	129090.3	1.418
Net-A	1843	3927	2.131	4477	6561	1.465	513957.6	807744.5	1.572
Net-B	3218	7427	2.308	7827	12036	1.538	912985.9	1437320	1.574
Net-C	6095	12713	2.086	13913	20531	1.476	1504465	2345433	1.559

It appears clearly that the capacity requirements associated with DFMC designs could hardly be seriously considered by any network operator and that another approach would have to be considered if one wanted to improve the network's dual span-failure restorability.

We therefore turn our attention to the results of the DFMR formulation in which the capacity cost budget can be decided and provided as input and the formulation gives us the achievable R_2 for that budget. What the DFMR formulation minimizes in practice is the total of non-restorable working units over all possible ordered dual span-failure scenarios. This total can be then used in Equation (4.8) to obtain the network's average dual span-failure restorability R_2. Table 6-2 on p. 132 shows the results obtained using the DFMR formulation with all test networks. For each test network, the DFMR formulation was run several times with increasing capacity cost budgets. The sequence of capacity cost budget increase used was: 0%, 3%, 6%, 12%, 24% and 48% except when the total capacity cost increase needed to achieve full dual span-failure restorability was below 48%. The case of 0% capacity cost increase can be interpreted as a way to let the DFMR formulation "re-distribute" the spare capacity corresponding to the minimum spare capacity requirement for single span-failure restorability. This "pure re-distribution" tries to enhance the achievable R_2 without adding any capacity cost to the design. Table 6-2 gives the results of total number of non-restored working units over all dual span-failure scenarios and the corresponding R_2 values calculated with Equation (4.8). Note that the DFMR formulation does not require that the test networks contain no 2-edge cuts and therefore results were produced for all test networks, including the ones with some 2-edge cuts. The results are also shown on Figures 6-2 and 6-3. Each curve on Figures 6-2 and 6-3 shows the improvement in R_2 as the capacity cost budget is increased and its allocation optimized under DFMR. In each case there is an initially better-than-linear growth of R_2 as the capacity cost allowance is increased but this slows greatly as R_2 nears unity and merges with the very high capacity results of the limiting DFMC case. For the test networks with some 2-edge cuts, the same thing is observed except that R_2 reaches a limit below unity. Indeed, with these networks there is a limit to the minimum achievable sum of non-restorable working units that is independent of the capacity cost budget specified. This limit is given by twice the sum of working capacity over all degree-2 cuts. For each of these cases it is possible to calculate the maximum achievable R_2 using the minimum sum of non-restorable working units in Equation (4.8). In all cases, this maximum R_2 is extremely close to unity, which indicates that high dual-failure restorability is achievable even in networks with 2-edge cuts.

Finally, we can observe the effects of the "zero-cost" redistribution of capacity which is visible on some test networks but stays below 5% in all cases.

Table 6-2: DFMR Results

Test Network	R_2 @ X% total capacity cost increase											
	$X = 0$		$X = 3$		$X = 6$		$X = 12$		$X = 24$		$X = 48$	
	$\sum N(i,j)$	Ave. R_2	$\sum N(i,j)$	Ave. R_2	$\sum N(i,j)$	Ave. R_2	$\sum N(i,j)$	Ave. R_2	$\sum N(i,j)$	Ave. R_2	$\sum N(i,j)$	Ave. R_2
06n14s1	1180	0.618617	908	0.706529	750	0.757595	518	0.832579	218	0.929541	n/a	n/a
11n20s1	3798	0.817614	3076	0.852286	2700	0.870342	2114	0.898483	1366	0.934403	672	0.96773
11n20s2	3232	0.711686	2594	0.768599	2124	0.810526	1434	0.872079	740	0.933988	258	0.976985
Bellcore	3262	0.840909	2494	0.878365	1924	0.906165	1122	0.945279	470	0.977078	200	0.990246
Bellcore Mod.	3208	0.861126	2310	0.9	1710	0.925974	958	0.958528	336	0.985455	n/a	n/a
COST239	1500	0.860465	968	0.909953	744	0.930791	476	0.955721	210	0.980465	n/a	n/a
12n20s1	3760	0.762146	3034	0.808072	2682	0.830339	2192	0.861336	1556	0.901569	820	0.948128
12n30s1	2536	0.89014	2054	0.911021	1690	0.926789	1136	0.950788	492	0.978687	n/a	n/a
15n28s1	6912	0.851336	5772	0.875855	4890	0.894825	3598	0.922614	2034	0.956252	630	0.98645
16n29s1	8400	0.82415	6968	0.854128	5996	0.874477	4534	0.905083	2814	0.94109	1098	0.977014
16n38s1	6172	0.885432	4392	0.918473	3308	0.938595	1970	0.963432	584	0.989159	n/a	n/a
EuroNet	12142	0.887649	8750	0.919035	6442	0.940392	3670	0.966041	1664	0.984603	658	0.993911
Net-A	21020	0.897689	16804	0.91821	13694	.93347	9474	0.953887	4642	0.977406	660	0.996788
22n41s1	11742	0.896199	9308	0.917716	7408	0.934512	4886	0.956807	2290	0.979756	364	0.996782
Net-B	33256	0.926373	26898	0.940449	21730	0.951891	14448	0.968013	6466	0.985685	922	0.997959
Net-C	52310	0.943297	42568	0.953857	34824	0.962251	23280	0.974765	10344	0.988787	658	0.999289

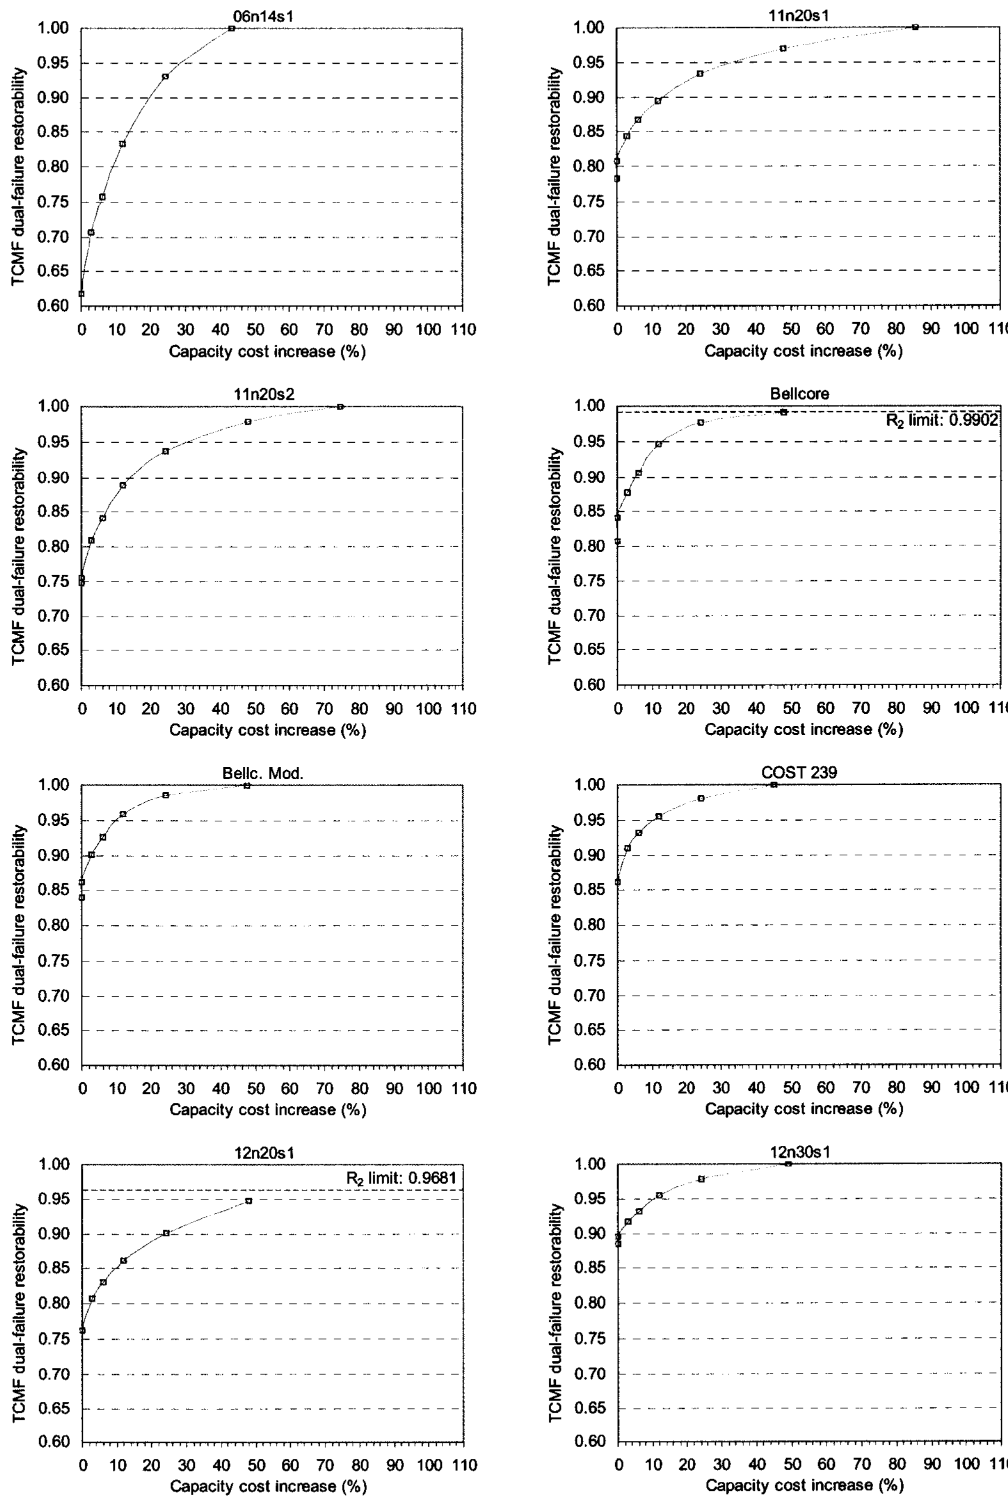

Figure 6-2 R_2 versus capacity trade-off obtained with DFMR

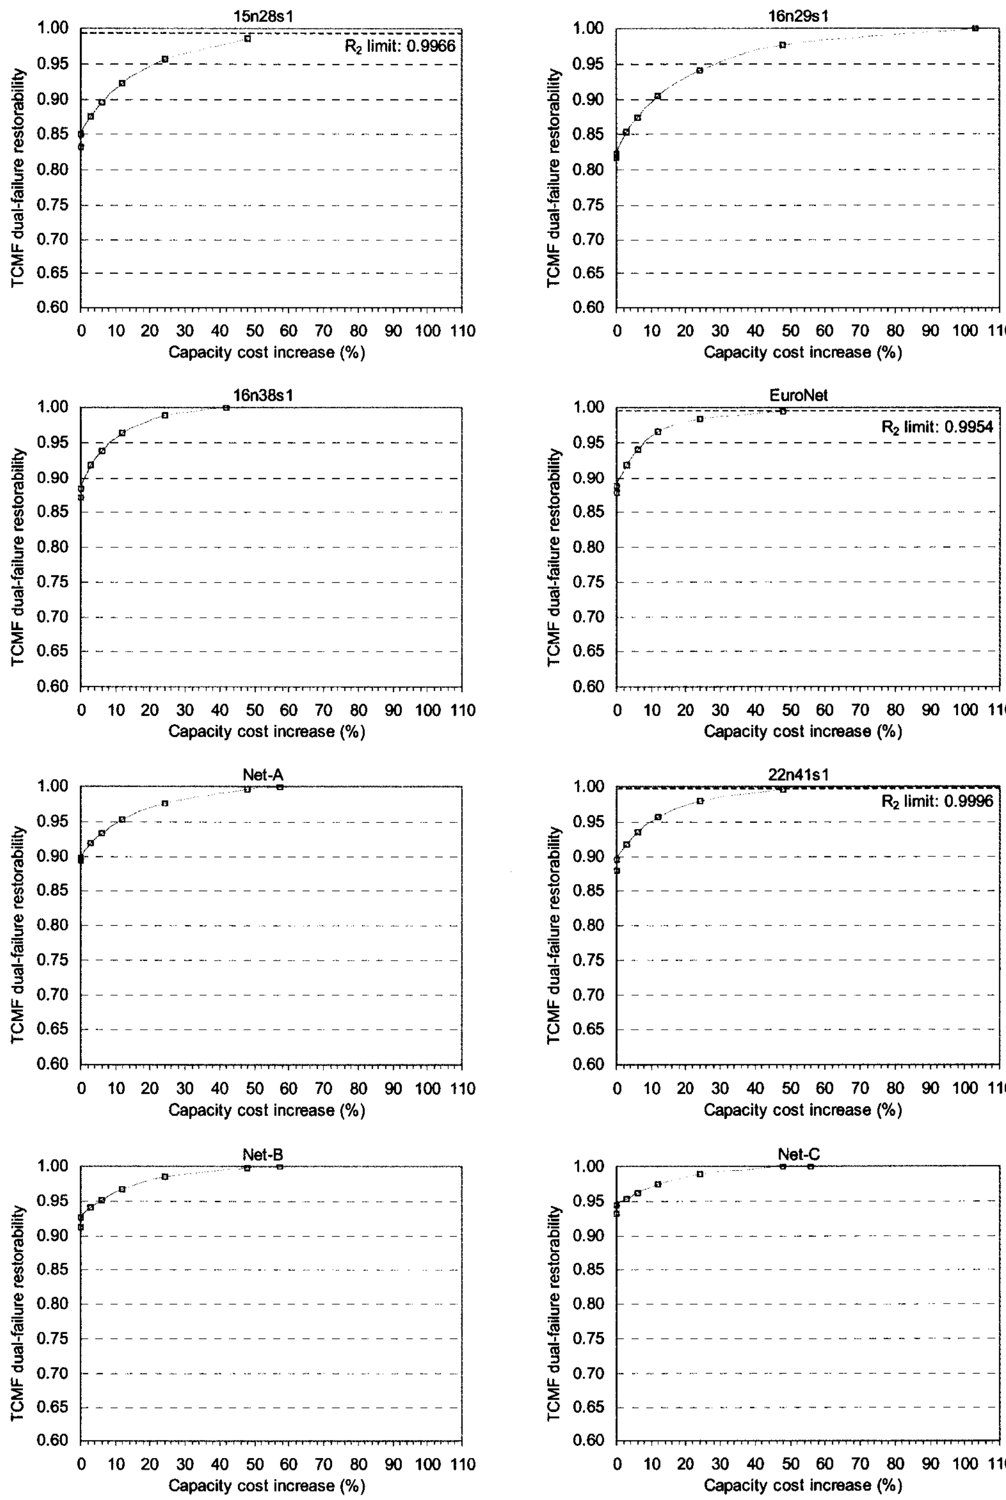

Figure 6-3 R_2 versus capacity trade-off obtained with DFMR (continued)

Results of the MRCP formulation are presented in Tables 6-3 and 6-4 and on Figure 6-4. As in the case of DFMC, since the MRCP formulation guarantees full dual span-failure restorability to some demands, experiments can only be conducted on test networks with no 2-edge cuts. For the test networks that qualify we are still using the demand matrices detailed in Appendix D for all test networks, but we overlay a status classification of individual demands on each O-D pair with respect to the multi-service restorability categories. Although the general formulation allows R_0, R_1, and R_2 categories, it was thought more meaningful and stringent for test purposes to retain R_1 as a minimum requirement for all demands in the tests here. This is conservative in terms of assessing the potential R_2-serving ability and it reduces the dimensionality of the results to be presented. For tests of MRCP the fraction of demands on each O-D pair that were given R_2 status was varied from 0 to 100 percent, as detailed in Tables 6-3 and 6-4. Note that the "0 percent" case is equivalent to the traditional JCP model detailed in Appendix B, and the "100 percent" case corresponds to a joint working and spare capacity placement version of the DFMC formulation. Tables 6-3 and 6-4 show the total capacity (working plus spare) cost requirements to support the given mix of "R_1-assured" and "R_2-assured" services (R_2 class implies an R_1 assurance too). Figure 6-4 shows the increase in total capacity cost relative to that of the single-class R_1 restorability requirements.

The results suggest that through this formulation *a significant class of R_2 customers could be served with relatively small increases in total capacity cost.* For instance, with test networks 06n14s1 and COST 239, there is no required capacity cost increase to serve up to 30% of demands with R_2 status, and for test networks Bellcore Modified, 12n30s1, 16n38s1, Net-A, Net-B and Net-C, the required capacity cost increase at this level is only around 1 percent. On Figure 6-4, it appears clearly that all networks can serve up to 40 percent of demands with modest capacity cost increase. Then, above 40 percent the capacity cost is subject to a strong increase for all networks and seems to reach a plateau between 50 and 60 percent and then starts increasing again without interruption until 100 percent of demands with R_2 status is reached. Our interpretation of these curves is the following: at first networks are capable of naturally serving up to 20 or 30 percent of demands with R_2 status, simply as a side-effect of being designed to support full single-failure restorability. At this level, in the event of dual failures the spare capacity is probably not completely utilized and it is likely that small islands of spare capacity still remain but are disconnected and cannot be used to form additional restoration paths that could bring the restorability up to 50 or 60 percent. The effect of the first capacity cost increase is to add spare capacity to connect these pockets of spare capacity. We then benefit from this remaining spare capacity that was not usable at first and we see the fraction of demands with R_2 status increase from 50 to 60 percent with a smaller ca-

Table 6-3: MRCP Results

Test Network	0% R_2 (JCP)	10% R_2		20% R_2		30% R_2		40% R_2		50% R_2	
	Total cap. cost	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.
06n14s1	4710	4710	0.00	4710	0.00	4710	0.00	4794	1.78	5146	9.26
11n20s1	143394	143811.2	0.29	144708	0.92	147963.2	3.19	158096.2	10.25	181747	26.75
11n20s2	131216.5	131322.4	0.08	131866.2	0.50	135980	3.63	145002.4	10.51	162256.4	23.66
Bellcore Mod.	25290	25346	0.22	25393	0.41	25651	1.43	25959	2.65	29114	15.12
COST239	153840	153840	0.00	153840	0.00	153840	0.00	155700	1.21	178315	15.91
12n30s1	58392.38	58407	0.03	58484.26	0.16	58899.24	0.87	59791.78	2.40	64358.72	10.22
16n29s1	156964.4	156973.9	0.01	157527.7	0.36	165670.2	5.55	177794.9	13.27	203962	29.94
16n38s1	85401.62	85460.61	0.07	85523.83	0.14	86682.26	1.50	88583.11	3.73	97362.9	14.01
Net-A	460159.5	460159.5	0.00	460844.8	0.15	465520.2	1.16	482669.8	4.89	539912.3	17.33
Net-B	783105.3	783835.4	0.09	785730.7	0.34	794351.5	1.44	830426.7	6.04	926300.1	18.29
Net-C	1273762	1276119	0.19	1275796	0.16	1289778	1.26	1329732	4.39	1473905	15.71

Table 6-4: MRCP Results (continued)

Test Network	60% R_2		70% R_2		80% R_2		90% R_2		100% R_2	
	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.	Total cap. cost	Perc. cap. cost incr.
06n14s1	5342	13.42	5735	21.76	6100	29.51	6454	37.03	6694	42.12
11n20s1	189561.6	32.20	211789	47.70	225200	57.05	254700	77.62	266536	85.88
11n20s2	169999.4	29.56	188010.1	43.28	199416	51.97	216885	65.29	228811	74.38
Bellcore Mod.	30160	19.26	32012	26.58	34302	35.63	36635	44.86	38221	51.13
COST239	183755	19.45	191525	24.50	205880	33.83	218665	42.14	228730	48.68
12n30s1	65712.93	12.54	69838.71	19.60	73527.4	25.92	81085.6	38.86	82976.4	42.10
16n29s1	209814.2	33.67	228600.1	45.64	253955	61.79	280498	78.70	289882	84.68
16n38s1	99397.38	16.39	104267.3	22.09	114693	34.30	124105	45.32	126343	47.94
Net-A	565734.9	22.94	611897.8	32.98	621595	41.60	701343	52.41	741280	61.09
Net-B	965023	23.23	1050436	34.14	1101518	40.66	1186155	51.47	1249066	59.50
Net-C	1529409	20.07	1664447	30.67	1755131	37.79	1893212	48.63	1992820	56.45

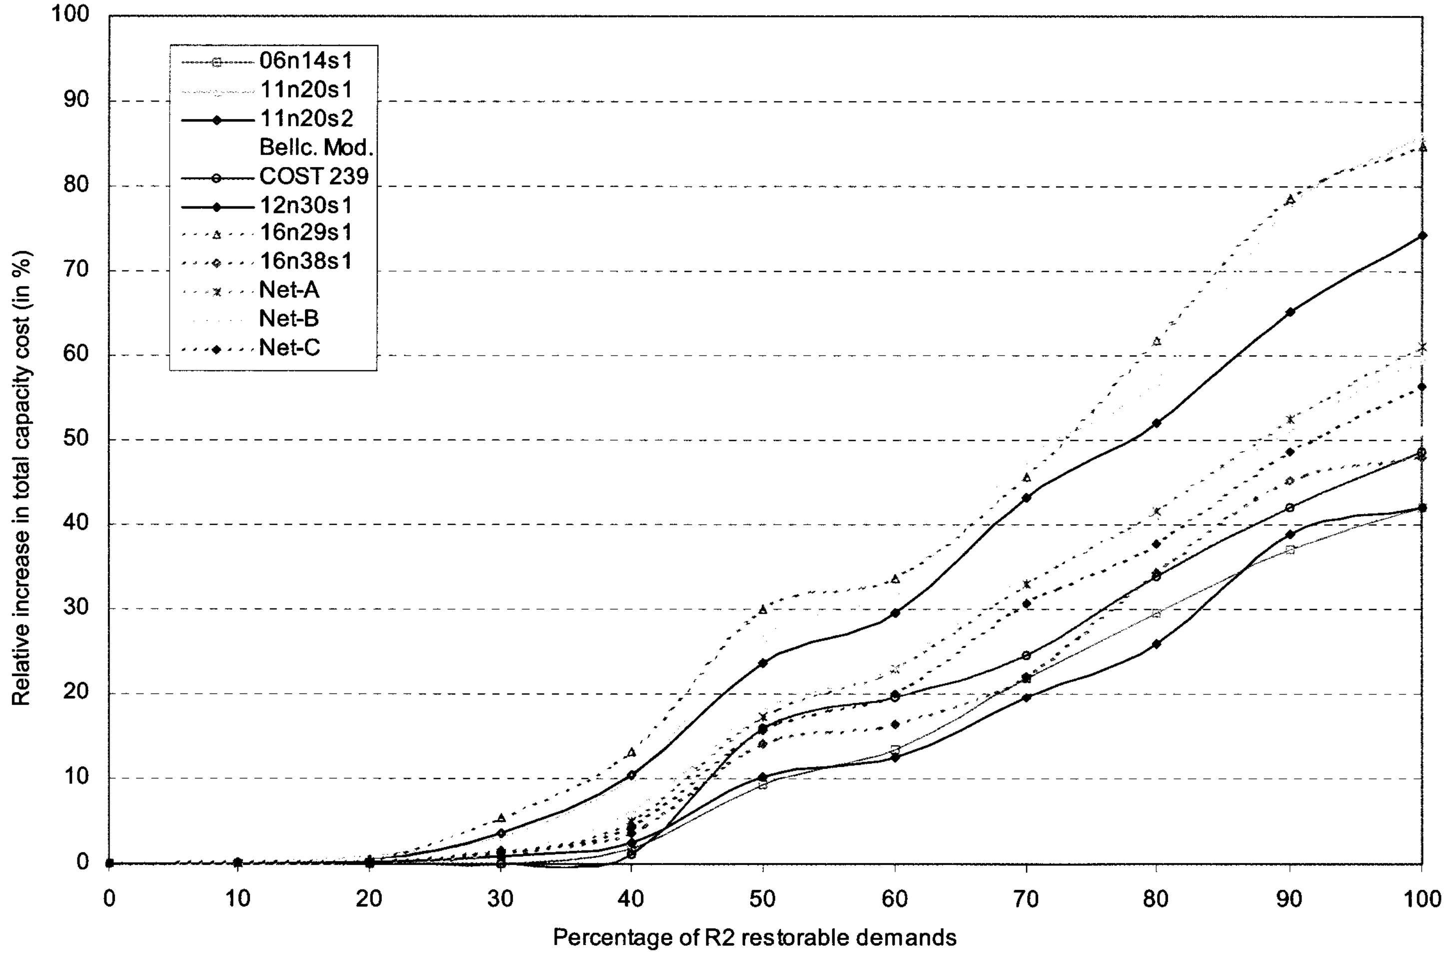

Figure 6-4 Percentage of R_2 demands versus capacity (results of MRCP)

pacity cost increase compared to the rest of the curve. In passing, the levels of capacity cost increase at 100 percent of R_2-demands observed on these curves are consistent with the results obtained with the DFMC formulation, although – as expected – they are not exactly equal since the MRCP formulation is joint while DFMC was a non-joint formulation strictly based on shortest path routing.

Overall, what we see from these results is that there is a tendency of mesh networks to be capable of supporting a select class of "ultra high availability" R_2-guaranteed customers without the huge capacity requirement for the network to have an R_2 of 1 as a whole. Inspection of the designs show that the satisfaction of the R_2 service class arises not only through re-distribution of spare capacity to target high R_2 on the key spans, but also from changes in the routing of working paths that bring them through regions of the network that are more easily made R_2 restorable.

These results are probably also fairly conservative estimates of the operational potential for multi-service class design because the present MRCP formulation is non-modular and because we present results with no R_0 service class members. As seen in Chapter 4, modularity only increases the inherent R_2 level. In addition R_0 class services require no spare capacity so can only increase the fraction of R_2-servable customers at the same total capacity cost investment.

6.4 High Availability and Fast Restoration

6.4.1 The Issue of Restoration Time Revisited

As explained in Section 3.6.1, despite the common belief that high availability requires fast restoration, the restorability to dual span-failures is really what matters when determining the availability of service paths. We have seen in the previous section that it is possible to offer the dual span-failure restorability guarantee to selected demands in a mesh network, and that this could even be achieved with very modest capacity cost increase when the proportion of demands in the R_2 class does not exceed 30 or 40 percent. As we have said, the availability of service paths in that class of highly restorable demands is extremely high since it now takes a triple failure to affect them. However, now that dual-failures do not need to be considered any more for these services, the contribution of restoration times to unavailability may become relevant in the calculation of the unavailability. As a simple exercise, let us revisit Table 3-3 of Section 3.6.4 and recalculate the contributions of the different failure types with an the values with an R_2 of 1. All the other values are kept identical to the ones used in Table 3-3. The new results are shown in Table 6-5. It appears that the contribution of time exposure to single span failures – i.e. the contribution of restoration time to the unavailability – now becomes the dominant factor in the unavailability of the links that are guaranteed to be restorable to dual failures. The following section presents a simple concept that shows how restoration in mesh networks can in fact be performed with speeds almost comparable to those usually associated to protection schemes.

Table 6-5: Relative contribution of failure events on the unavailability of links (with $R_2 = 1$)

Network State f	Description	Probability given failure of span i	Time exposure	Capacity exposure	Contribution in example
1	Single span-failure, i	1	$\frac{\text{Av. RestTime}}{MTTR}$	0	1.38×10^{-8}
2	Dual span-failure, i and other span j	$(\lvert S\rvert - 1) \cdot U_{\text{link}}^{\text{p}}(i)$	$\frac{\text{Av. RestTime}}{0.5 \times MTTR} \cdot R_2$	0	1.58×10^{-10}
3	Triple span-failure, i and other spans j, k	$\frac{(\lvert S\rvert - 1)(\lvert S\rvert - 2)}{2} \cdot (U_{\text{link}}^{\text{p}}(i))^2$	$\frac{\text{Av. RestTime}}{0.33 \times MTTR} \cdot R_3$	$1 - R_3$	4.62×10^{-9}

6.4.2 First-Failure Protection Second-Failure Restoration (1FP-2FR) Concept

During the development of this thesis, a concept emerged within our research group, closely related to the above findings. It is the notion of "1FP-2FR" for ultra-high availability services. The main ideas are as follows. Much is made of the distinction between "protection" and "restoration," it being often asserted in a very general way that protection is fast and restoration is slow [ZhS00]. The view is too simplified however because there are really three categories of scheme to consider and the perception of how fast each can be depends whether it is assumed that path-finding is time consuming or if it is path cross-connection that is assumed to be slow. Moreover, it misses the always-present relationship between any restoration scheme and a corresponding pre-planned protection scheme, which is derivable through distributed pre-planning (DPP).

The three basic possibilities are:

1. schemes where the protection routes are known in advance and cross-connection is not required to use the backup path(s), e.g.; UPSR, BLSR, 1+1 APS;
2. schemes where the protection routes are known in advance but cross-connection is required in real-time to use the backup path(s), e.g. MPLS shared backup path protection, or ATM Backup VP protection schemes;
3. schemes where the routes are found adaptively based on the state of the network at the time of failure, and the cross-connections to put the restoration routes into effect are also made in real time, e.g., Self-healing networks, distributed self-organizing restoration schemes.

Schemes of type 3) are only slower than type 2) schemes if it is assumed that the restoration path-finding process is time consuming. But in such cases, distributed pre-planning can create – and frequently update – a corresponding type 2) scheme where the protection routes are known in advance of failure. This is done by distributed pre-planning with mock-failure trials responded to by the embedded restoration protocol. The concept is described more fully in [Gro94] or [Gro97]. It is quite a simple technique that retails all of the generality and database freedom of a distributed restoration algorithm, but provided a "protection" scheme of the type 2) above.

Consequently any type 3) scheme, which tend to be called restoration schemes, need to be seen as actually providing the option of both a self-planning protection scheme and an on-demand dynamic adaptive real-time restoration scheme. The relevance to dual-failure recovery is that the response to a first failure can be based on a pre-planned protection reaction, and only in the event of a second failure, is the truly adaptive but possibly slower restoration protocol itself executed directly in real time. Seen in this light, an adaptive restoration protocol is ideally suited for both requirements in a mesh network. First, it can serve as the engine for constant background re-planning of a fast "protection" reaction against single failures. Second, in the event of a second failure (more generally, any time the protection level is not 100%), it then executes directly in real time where its completely adaptive nature is exactly what is required to produce the highest possible overall recovery level.

6.5 Summary

The work presented in this chapter has considered three ways in which the basic formulations for span-restorable mesh network design can be extended to design for higher availability through strategies for control of R_2. To start with we have shown that designing to support complete dual-failure restorability requires very large additions of spare capacity relative to the single-failure design case. Assuming full dual span-failure restorable design may often be too expensive, the DFMR formulation can be used to at least maximize the achievable network-wide level of R_2 subject to any desired capacity cost budgetary allocation for this purpose. Results showed that the overall average R_2 could be significantly increased for a small increase of capacity cost. In addition as a check on results, the budget requirement to reach an R_2 of 1 is consistent with that predicted by DFMC. DFMR can also be used in many planning contexts to study the shape of the R_2 (and hence availability) versus cost curve of an overall network design.

The MRCP method may, however, be of the most practical importance because it directly supports the desired business model of multiple restoration service classes. Results with MRCP also show that a relatively high number of customers can be supported in an ultra-high availability service class that withstands all dual span-failures, but with little or no increase of the total network capacity cost. The key to how MRCP does this is that the inherently high average R_2 levels of ordinary single-failure designs are in effect given internal structuring so that this initially incoherent or random attribute property is coherently targeted on the priority paths, rather than being a network-average property only. This makes it possible to guarantee extremely high availability levels to selected demands, with minimal added capacity. As such MRCP directly addresses a long-held issue for network operators: "How to earn new revenue from investing in survivability."

7. Availability of Paths in *p*-Cycle Networks

7.1 Introduction

The study presented in this chapter was first proposed by Dr. Wayne Grover as a final project in the course E E 681 "Survivable Networks." Several students in the past few years have worked under my supervision on developing equations for the availability of service paths in *p*-cycles. The following equations are partly inspired by the course project of H. Gopaluni and P. Nakeeran who showed the feasibility of taking a *p*-cycle by *p*-cycle approach. The idea is to follow a "polynomial type" of analytical approach, analogous to that in [Gro99a] for rings, to see how far corresponding models for *p*-cycles can be developed. In such an approach the *p*-cycle structure and path routing will be known for the analysis

7.2 The *p*-Cycle Concept

p-Cycle protection is a recently developed survivability scheme that offers fast restoration and low capacity requirements [GrS98][GrS00]. Figure 7-1 portrays the main principles of *p*-cycle pro-

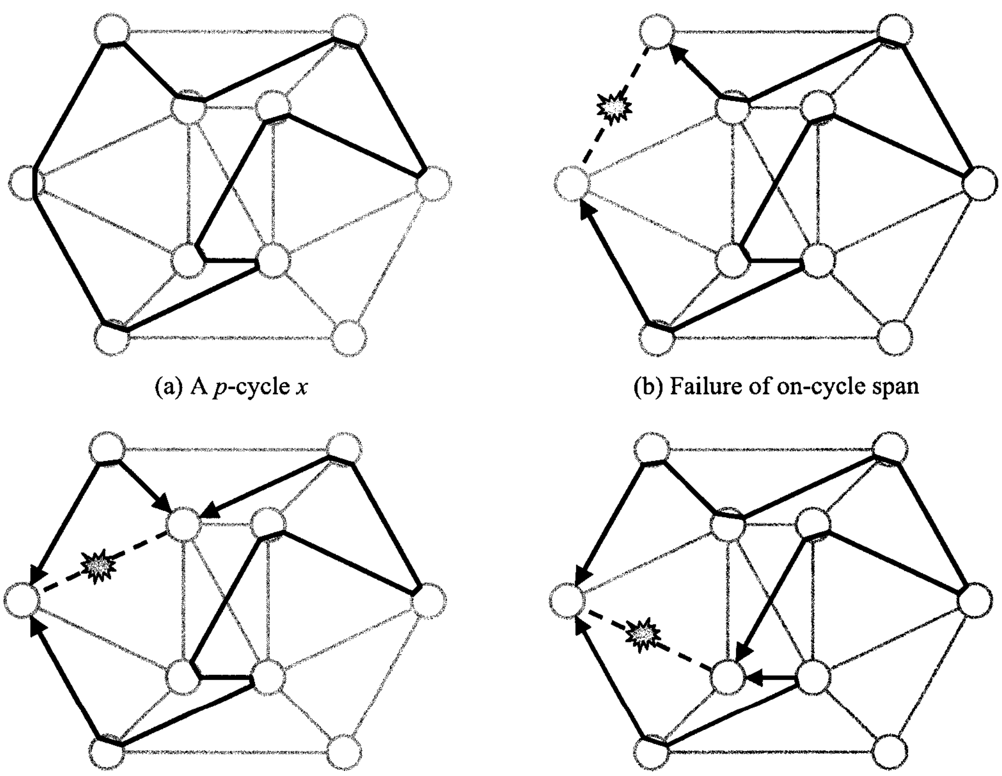

(a) A *p*-cycle *x* (b) Failure of on-cycle span

(c) Failure of straddling span (d) Failure of straddling span

Figure 7-1 *p*-Cycle restoration
(from [GrS98])

tection. A p-cycle, as shown on Figure 7-1 (a) is a closed path composed of spare capacity links. When a failure occurs on a span covered by the cycle, the p-cycle provides one backup path for restoring one link on the failed span, as shown on Figure 7-1 (b). So far, the p-cycle concept is identical to that of a protection ring, except maybe that a p-cycle represents only one unit of protection capacity whereas a ring provides protection capacity either at the line level (BLSR) or at the path level (UPSR) but even in that case offers entire modules of protection capacity. The main difference between rings and p-Cycles is that p-Cycles can also protect straddling spans – the spans that have both ends on the cycle but do not themselves lie on the cycle – as shown on Figure 7-1 (c) and (d). This apparently minor difference has in fact a great impact on the capacity requirements of p-cycle protection, which achieves spare capacity requirements almost as low as those usually seen only with span-restoration.

What this chapter focuses on is the question of the availability of service in p-cycle based networks. Of particular interest is how p-cycle availability compares to the availability of span-restorable mesh networks. In the following section, we develop a set of closed-form equations for the availability of paths in p-cycle protected networks. These equations will provide some insights about the factors that influence the most the availability of service in p-cycle based networks. Due to the lack of closed-form equations for the availability of a path in span-restorable mesh networks, a comparison between the two schemes based on dual-failure restorability calculations will be provided in Section 7.4. The remainder of the chapter provides suggestions of ways to improve service availability in p-cycle based networks.

7.3 Theoretical Treatment of Path Availability in p-Cycle Networks

7.3.1 Protection Domains

To develop the equations of path availability in a p-cycle protected network, we will first explain the concept of protection domain. Between its origin node and destination node, a path crosses several protection domains, each of which is associated with a given cycle. For example, to say that a path p crosses a protection domain associated to cycle x means that at some point path p is protected by a p-cycle lying on cycle x. For each protection domain crossed and its associated cycle, the path is either a straddler or on-cycle but not both. What we mean is that if path p is protected by a p-cycle on cycle x as an on-cycle path and also as a straddling path, these will count as two separate protection domains. If, however, on multiple on-cycle spans (or respectively straddling spans) for cycle x, path p is protected by p-cycles on cycle x, this will count as only one protection domain.

For a path p crossing a protection domain associated to cycle x, the following notations will be used:

S_x^c Set of spans on cycle x.

S_x^s Set of spans straddling cycle x.

$S_{p,x}^c$ Set of spans on cycle x and on which path p is protected by a p-cycle lying on cycle x.

$S_{p,x}^s$ Set of spans straddling cycle x and on which path p is protected by a p-cycle lying on cycle x.

7.3.2 *Outage-Causing Dual Span-Failure Sequences*

We now present the set of all dual span-failure sequences that lead to outage for a given path in a p-cycle protected network. Each of these sequences is independent from the others, that is, a dual span-failure can only belong to one of these sequences. This will later facilitate the calculation of path availability since we will be able to add the probability of each of these sequences to determine the path unavailability.

The first two dual span-failure sequences are for an on-cycle path p in a given protection domain. On Figure 7-2, we show the case of a first failure on the path itself followed by a second failure elsewhere on the cycle. This is guaranteed to cause outage for the path. The case where these two failures happens in reverse order is also guaranteed to cause outage for the path. The combined probability of this failure sequence (and the reversed one) is given by:

$$P(\text{seq. }1) = |S_{p,x}^c| \cdot (|S_x^c| - 1) \cdot U_s^2 \tag{7.1}$$

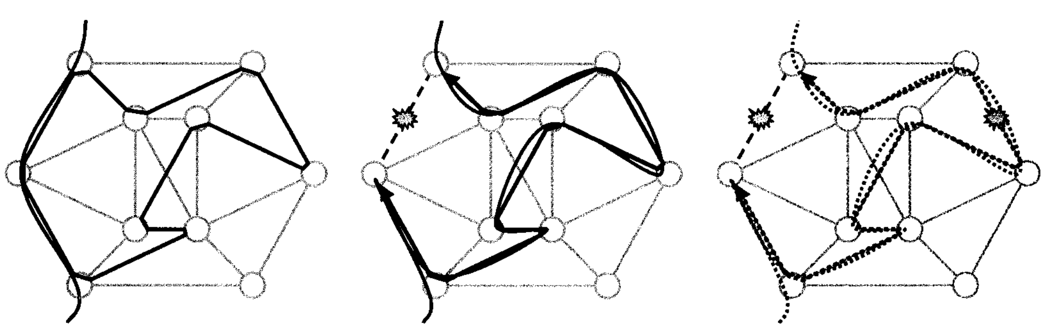

(a) On-cycle path, no failure (b) On-cycle path is restored (c) Backup path is hit by a failure

Figure 7-2 Dual-failure sequence leading to outage for on-cycle path, sequence 1

On Figure 7-3 we show the second type of dual span-failure sequence that is guaranteed to cause outage for an on-cycle path p. In this case, a first span-failure happens on a straddling span and that span requires one or two protection paths from the p-cycle protecting path p. Note that un-

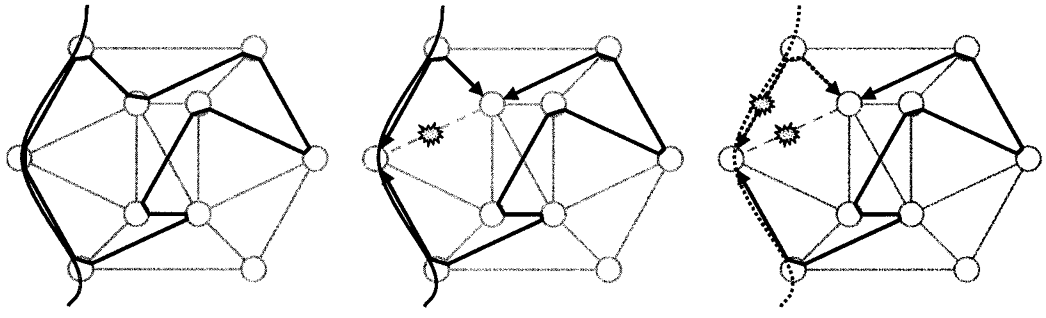

(a) On-cycle path, no failure (b) Straddling link(s) are restored (1 or 2 backup paths used) (c) On-cycle path fails and is not restorable

Figure 7-3 Dual-failure sequence leading to outage for on-cycle path, sequence 2

like in the previous case, the same failure combination would not cause outage for path p if the failures happened in the reverse order.

The probability of this second span-failure sequence is given by (7.2). In that equation, $P_x^{s,1}$ represents the probability that the failure of a span straddling cycle x will require the use of exactly one backup path on a given p-cycle on cycle x. Similarly, $P_x^{s,2}$ represents the probability that the failure of a span straddling cycle x will require the use of two backup paths on a given p-cycle on cycle x. The 0.5 factor represents the probability that the straddling span failure happens first.

$$P(\text{seq. }2) = \left|S_x^{\text{s}}\right| \cdot \left|S_{p,x}^{\text{c}}\right| \cdot \frac{1}{2} \cdot U_{\text{s}}^2 \cdot (P_x^{\text{s},1} + P_x^{\text{s},2}) \tag{7.2}$$

The following four dual span-failure sequences are for a straddling path in a given protection domain. In the first sequence, portrayed on Figure 7-4, the first span failure affects the straddling path and the second failure hits the backup path of the straddling path. If this failure sequence happens there is on average a 50 percent chance that the second failure will hit the backup path used to restore the straddling path. If the same failure combination happens but the order of the failures is

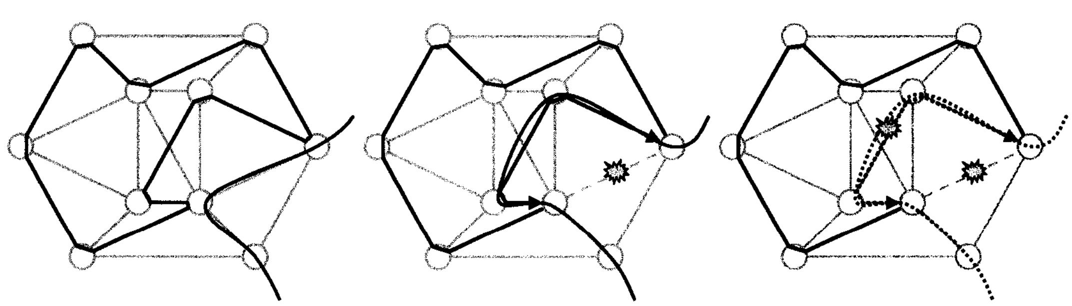

(a) Straddling path, no failure (b) Straddling path is restored (c) Backup path is hit by a failure

Figure 7-4 Dual-failure sequence leading to outage for p-cycle straddling path, sequence 3 reversed, then outage for the straddling path will be guaranteed. Given this (un-ordered) combination of failures, the probability of outage for the considered path is therefore 75 percent. The combined probability of this failure sequence and the reverse one is given by (7.3).

$$P(\text{seq. }3) = \left|S^{\text{s}}_{p,x}\right| \cdot \left|S^{\text{c}}_{x}\right| \cdot \frac{3}{4} \cdot U_{\text{s}}^{2} \qquad (7.3)$$

Figure 7-5 portrays the case where the first span-failure happens on a straddling span not crossed by the straddling path and that span requires two backup paths on the p-cycle protecting the straddling path. The second span failure hits the straddling path and restoration of the straddling path is not possible. For this failure combination, the reverse order would not cause outage for the straddling path. The probability of this failure sequence is given by (7.4). In this expression, the negative term in the left brackets accounts for the fact that if the path crosses multiple straddling spans of the p-cycle, some of these dual-straddling span failure combinations (the ones involving two straddling spans crossed by the considered path) will in fact be counted twice.

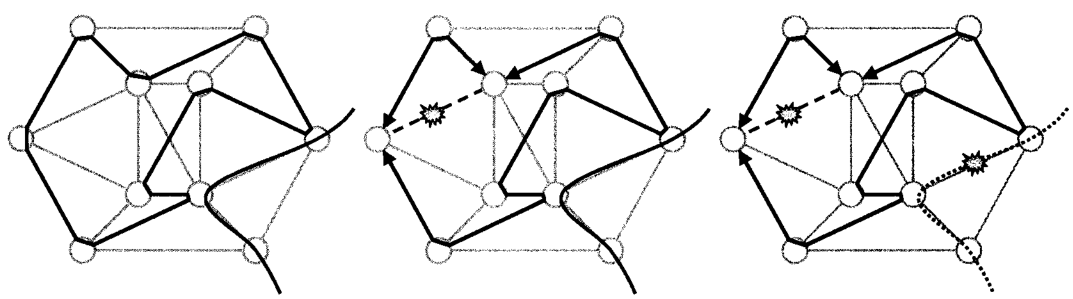

(a) Straddling path, no failure (b) Straddling links are restored (2 backup paths used) (c) Straddling path fails and is not restorable

Figure 7-5 Dual-failure sequence leading to outage for p-cycle straddling path, sequence 4

$$P(\text{seq. }4) = \left[\left|S^{\text{s}}_{p,x}\right| \cdot \left(\left|S^{\text{s}}_{x}\right| - 1\right) - \binom{\left|S^{\text{s}}_{p,x}\right|}{2}\right] \cdot \frac{1}{2} \cdot U_{\text{s}}^{2} \cdot P_{x}^{\text{s},2} \qquad (7.4)$$

The sequence illustrated on Figure 7-6 shows a first span-failure affecting a straddling span not crossed by the straddling path and that span requires only one backup path on the p-cycle protecting the straddling path. The second span failure affects the straddling path but the configuration of the backup path used by the first failure prevents a backup path for the straddling path to be formed. In this particular case, the reason why a backup path cannot be found for the straddling path is that the first failed span is in a crossing situation with the second failed straddling span. The idea of crossing straddler is illustrated by Figure 7-7. On the left part of Figure 7-7, the p-cycle is represented as usually and none of the straddling spans seem to cross each other. On the right part

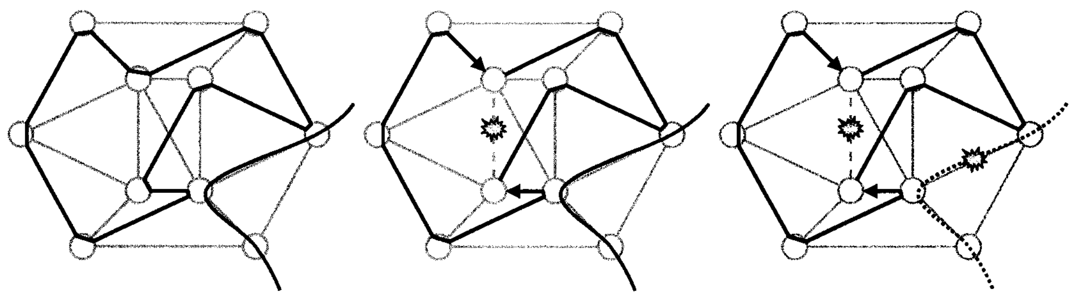

(a) Straddling path, no failure (b) Crossing straddling link is restored (c) Straddling path fails and is not restorable

Figure 7-6 Dual-failure sequence leading to outage for p-cycle straddling path, sequence 6

of Figure 7-7, the p-cycle is re-drawn in expended form and it appears that many of the straddling span are in fact crossing each other. For obvious reasons, when a straddling span uses a backup path on the p-cycle, another straddling span that crosses it cannot be restored at the same time, even if it requires only one backup path. A link on a straddling span that does not cross it will have 25 percent chance of being restorable (assuming each straddling link has a pre-determined backup path and cannot use the backup path on the other side of the p-cycle if the pre-determined backup path is not available).

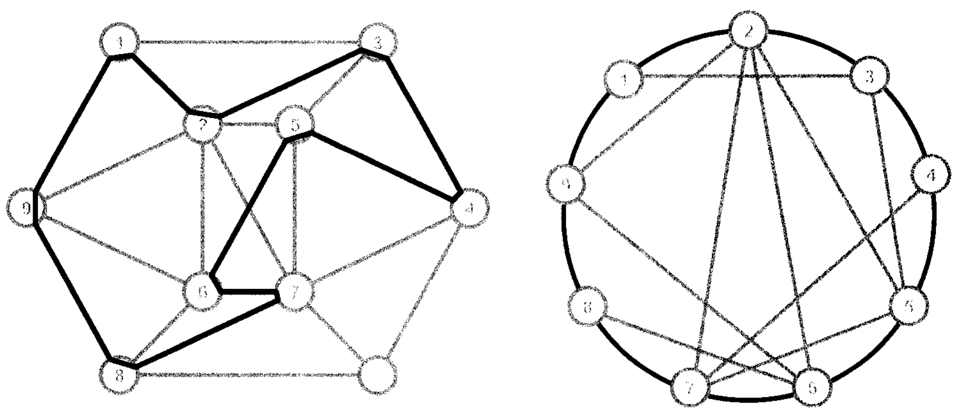

Figure 7-7 Illustration of the concept of crossing straddlers

The probability of the failure sequence shown on Figure 7-6 is given by (7.5). In that equation, K_x represents the percentage of crossing straddler combinations out of all straddler combinations of cycle x.

$$P(\text{seq. } 5) = \left[\left|S^{s}_{p,x}\right| \cdot \left(\left|S^{s}_{x}\right| - 1 \right) - \binom{\left|S^{s}_{p,x}\right|}{2} \right] \cdot \frac{1}{2} \cdot U^{2}_{s} \cdot K_{x} \cdot P^{s,1}_{x} \qquad (7.5)$$

Finally, Figure 7-8 represents the sequence of a first failure occurring on a straddler not

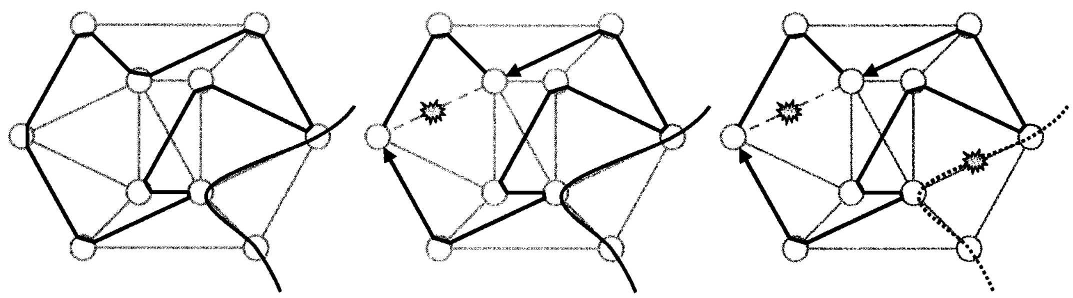

(a) Straddling path, no failure (b) Non-crossing straddling link restored on "bad" side (c) Straddling path fails and is not restorable

Figure 7-8 Dual-failure sequence leading to outage for p-cycle straddling path, sequence 5

crossed by the straddling path followed by the failure of the straddling path on a span that is in a non-crossing situation with the first failed span. As mentioned above, in this situation there is only 25 percent chance for the second failed link to be restorable.

The probability of the failure sequence portrayed by Figure 7-8 is given by:

$$P(\text{seq. }6) = \left[\left|S^{\text{s}}_{p,x}\right| \cdot \left(\left|S^{\text{s}}_{x}\right| - 1\right) - \binom{\left|S^{\text{s}}_{p,x}\right|}{2}\right] \cdot \frac{1}{2} \cdot U_{\text{s}}^{2} \cdot (1 - K_x) \cdot \frac{3}{4} \cdot P_x^{\text{s},1} \tag{7.6}$$

7.3.3 *Path Unavailability*

The unavailability of a path in a p-cycle protected network can be expressed as the sum of the unavailability of the path in the different protection domains crossed:

$$U_{\text{path}}(p) = \sum_{o \in O(p)} U_{p,o} \tag{7.7}$$

where $O(p)$ is the set of protection domains crossed by path p and $U_{p,o}$ is the unavailability of path p in protection domain o.

For a protection domain in which path p is an on-cycle path, we have:

$$U^{\text{c}}_{p,o} = P(\text{seq. }1) + P(\text{seq. }2) \tag{7.8}$$

For a protection domain in which path p is a straddling path, we have:

$$U^{\text{s}}_{p,o} = P(\text{seq. }3) + P(\text{seq. }4) + P(\text{seq. }5) + P(\text{seq. }6) \tag{7.9}$$

Some simplifications of the probability expressions can be made if we assume that $P_x^{\text{s},2}$ is 1. In that case, (7.2) becomes:

$$P(\text{seq. }2) = \left|S^{\text{s}}_{x}\right| \cdot \left|S^{\text{c}}_{p,x}\right| \cdot \frac{1}{2} \cdot U_{\text{s}}^{2}. \tag{7.10}$$

Equation (7.4) becomes:

$$P(\text{seq. }4) = \left[\left|S_{p,x}^{s}\right| \cdot (\left|S_{x}^{s}\right| - 1) - \binom{\left|S_{p,x}^{s}\right|}{2}\right] \cdot \frac{1}{2} \cdot U_{s}^{2}, \tag{7.11}$$

which can be simplified to:

$$P(\text{seq. }4) = \left|S_{p,x}^{s}\right| \cdot \left[(\left|S_{x}^{s}\right| - 1) - \frac{(\left|S_{p,x}^{s}\right| - 1)}{2}\right] \cdot \frac{1}{2} \cdot U_{s}^{2}. \tag{7.12}$$

Also, in that case, we have:

$$P(\text{seq. }5) = P(\text{seq. }6) = 0 \tag{7.13}$$

Finally, for a protection domain in which path p is an on-cycle path, we have:

$$U_{p,o}^{c} = \left|S_{p,x}^{c}\right| \cdot (\left|S_{x}^{c}\right| - 1) \cdot U_{s}^{2} + \left|S_{x}^{s}\right| \cdot \left|S_{p,x}^{c}\right| \cdot \frac{1}{2} \cdot U_{s}^{2} \tag{7.14}$$

$$U_{p,o}^{c} = \left|S_{p,x}^{c}\right| \cdot \left[\left|S_{x}^{c}\right| + \frac{1}{2} \cdot \left|S_{x}^{s}\right| - 1\right] \cdot U_{s}^{2} \tag{7.15}$$

where x is the cycle associated to protection domain o.

For a protection domain in which path p is a straddling path, we have:

$$U_{p,o}^{s} = \left|S_{p,x}^{s}\right| \cdot \left|S_{x}^{c}\right| \cdot \frac{3}{4} \cdot U_{s}^{2} + \left|S_{p,x}^{s}\right| \cdot \left[(\left|S_{x}^{s}\right| - 1) - \frac{(\left|S_{p,x}^{s}\right| - 1)}{2}\right] \cdot \frac{1}{2} \cdot U_{s}^{2} \tag{7.16}$$

$$U_{p,o}^{s} = \left|S_{p,x}^{s}\right| \cdot \left[\left|S_{x}^{c}\right| \cdot \frac{3}{4} + \left((\left|S_{x}^{s}\right| - 1) - \frac{(\left|S_{p,x}^{s}\right| - 1)}{2}\right) \cdot \frac{1}{2}\right] \cdot U_{s}^{2} \tag{7.17}$$

$$U_{p,o}^{s} = \left|S_{p,x}^{s}\right| \cdot \left[\frac{3}{4} \cdot \left|S_{x}^{c}\right| + \frac{1}{2} \cdot \left|S_{x}^{s}\right| - \frac{1}{4} \cdot \left|S_{p,x}^{s}\right| - \frac{3}{4}\right] \cdot U_{s}^{2} \tag{7.18}$$

7.3.4 On-Cycle vs. Straddling Path Unavailability Comparison

In this section, we provide a comparison of the unavailability of a path in a given domain depending on whether the path is an on-cycle path for the cycle associated to that domain or a straddling path for that cycle. This comparison is based on Equations (7.15) and (7.18). To allow comparison of the unavailabilities, we assume that $\left|S_{p,x}^{c}\right|$ is 1 for the on-cycle path and $\left|S_{p,x}^{s}\right|$ is 1 for the straddling path. What this means is that we assume that in both cases the path stays in the domain on a single span. As shown by (7.15) and (7.18), in both cases the unavailability of the path in the domain is proportional to the number of spans the path stays in the domain. The results of the comparison will therefore remain valid for other values of $\left|S_{p,x}^{c}\right|$ and $\left|S_{p,x}^{s}\right|$.

The first factor we investigate is the number of on-cycle spans for cycle x associated to the do-

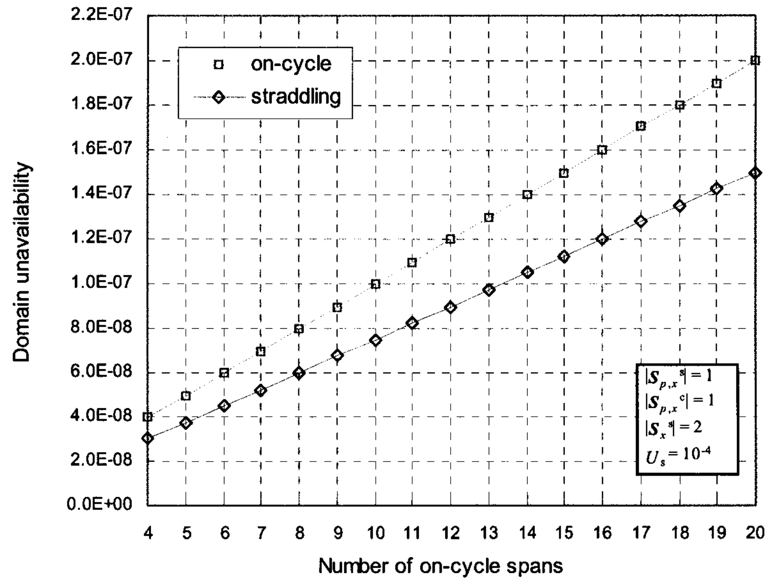

Figure 7-9 Comparison of on-cycle and straddling unavailability for varying number of on-cycle spans

main. The unavailability of the on-cycle path and of the straddling path were calculated for a number of on-cycle spans varying from 4 to 20. For these calculations, the number of straddling spans was fixed to 2 and the average span physical unavailability was assumed to be 10^{-4}. Results are shown on Figure 7-9. Results show that for both the on-cycle path and the straddling path, the unavailability grows proportionally to the number of on-cycle spans. For both cases, the unavailability with 20 on-cycle spans is 6.67 times higher than with 4 on-cycle spans. Results also show a significantly lower unavailability for the straddling path. The absolute difference between the availability of the on-cycle path and that of the straddling path increases with increasing number of on-cycle spans. In terms of percentage, the availability of the straddling path is 25 percent lower than that of the on-cycle path for all values of the number of on-cycle spans. The difference can be explained by the fact that on-cycle paths have longer backup paths and are therefore two times more exposed to the risk of having a secondary failure hit their backup path. As the number of on-cycle spans increases the difference in average backup path length of on-cycle spans increases twice as fast as that of straddling spans, explaining the increasing absolute difference in unavailability.

Figure 7-10 shows the effects of varying the number of straddling spans on the same two unavailability measures. The number of straddling spans was varied between 2 and 18 and we assumed 10 on-cycle spans and the same span physical unavailability as in the previous calculation. Here

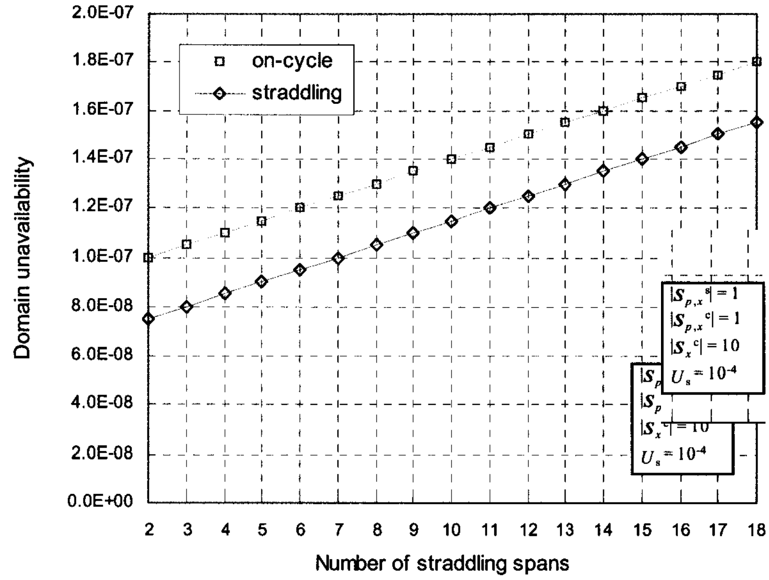

Figure 7-10 Comparison of on-cycle and straddling unavailability for varying number of straddling spans

again, the unavailability of paths increases linearly with increasing number of straddling spans but the increase is not as fast as in the case of increasing number of on-cycle spans. Between 2 and 18 straddling spans the unavailability of the on-cycle path increases by a factor 1.89 and the unavailability of the straddling path increases by a factor 2.21. The absolute difference between the unavailability of the on-cycle path and that of the straddling path remains constant in this case, which means that the relative difference decreases as the number straddling spans increases, but at 18 straddling spans, the difference is still of 13.9 percent.

What these results tell us is that the size of *p*-cycles plays a major role in determining the availability of service paths. Small *p*-cycles will allow much higher availability to be offered to paths. Not surprisingly, limiting the size of *p*-cycles will cause them to be less capacity efficient as shown in [Sch03b]. The number of straddling spans also matters but not as much.

The other important finding is that the unavailability of paths that transit on the cycles that protect them is higher than paths straddling the *p*-cycles. This difference could be used to provide some kind of service differentiation by routing service paths with higher availability requirements on straddling spans instead of on-cycle spans.

7.3.5 Proposed New Metric for Controlling the Unavailability in p-Cycle Networks

Based on the results of the previous section, it appears that a possible approach for controlling the unavailability of paths in *p*-cycle protected networks could be to put a limit on the maximum value of the following metric:

$$\lambda^x = (\text{number of on-cycle spans}) + \frac{1}{2} \cdot (\text{number of straddling spans}) \qquad (7.19)$$

for eligible cycles selected as input to the design formulation. This was not tested in this thesis work and is left for future research.

7.4 Experiments of Dual-Failure Restorability of *p*-Cycles

This section presents experimental results of the dual span-failure restorability of *p*-cycle based networks. A study of *p*-cycle restorability was presented in [Sch03a] but did not provide a comparison with the restorability of span-restorable designs for the same test cases, which is what this section does. Results of dual span-failure restorability for *p*-cycles networks were obtained using a custom-made restorability analysis program similar to the one used in Chapter 4. The restorability analysis results for span-restoration were obtained using the same tool as in Chapter 4.

In order to provide a meaningful comparison of the two schemes, the designs tested are based on the exact same topologies and demand patterns. These are the Net-A, Net-B, and Net-C test cases presented in Section 1.4.1 and described in details in Appendix A. The span-restorable designs used are the same as the non-modular ones used in Chapter 4. The *p*-cycle designs used are based on the sets of optimal *p*-cycles produced using the *p*-cycle SCP design formulation presented in [GrS98]. In each case, the set of 1000 best eligible cycles based on the metric introduced in [GrD02] was used in the capacity design optimization. Each resulting design was tested for full restorability to single failures first.

In Table 7-1 we present the results of dual span-failure restorability for the *p*-cycle designs and for the span-restorable designs. The results presented are path restorability values: they indicate the average fraction of affected service paths that are restored in case of a dual-failure. For the case of span-restoration, Table 7-1 gives the results obtained with the three restoration models presented in Section 4.4. Results show much higher restorability levels for the span restoration mechanism. The difference between p-cycle dual failure restorability and the span-restoration dual failure restorability is on average around 35 percentage points in favour of span-restoration. The adaptability of span-restoration could be thought as being the major cause of this difference, but even the purely static span restoration behaviour achieves much higher dual-failure restorability. What explains the great difference in restorability is the average size of the *p*-cycles used in the three *p*-cycle designs. The second column of Table 7-1 shows the average size of *p*-cycles in each design.These very high size averages are due to the fact that long *p*-cycles allow very efficient capacity sharing to be obtained and therefore are more likely to be used in optimal capacity designs. As a result, backup paths used to restore a first span failure are very highly exposed to subsequent failures. That, combined with the inability of the *p*-cycle mechanism to provide replacement solutions when pre-planned backup paths are not available explains the significantly lower dual-failure restorability.

Table 7-1: Experimental path restorability results for *p*-cycle

Test Network	Average *p*-cycle length	*p*-cycle protection	Span Restoration		
			Model 1	Model 2	Model 3
Net-A	17.1	0.5242	0.8375	0.8886	0.8773
Net-B	19.8	0.5416	0.8681	0.8967	0.9075
Net-C	24.7	0.5747	0.8901	0.9197	0.9278

The fact that the fully static model of span-restoration (Model 1) does not suffer as much from its inability to provide a revised reaction to a failure when the pre-planned restoration paths are not available is due to the lower average size of restoration paths and a potentially more diversified set of restoration routes as was shown on Figure 2-17. Also, Model 1 of the span-restoration mechanism is a bit more flexible than a purely fully static protection mechanism in the sense that it does not presume what specific spare capacity link must be used on a given span for restoration of a given failed working link. This flexibility was not assumed for *p*-cycle restoration, which is based on protection switching into pre-connected backup paths.

On Figure 7-11 is shown the distribution of the frequency of the various levels of dual failure path restorability for the three tests networks for the *p*-cycle designs and for the span-restorable designs with the partly adaptive behaviour. The distributions observed for the *p*-cycle designs are much wider than for span-restoration and show that the path restorability of different dual-failures are varied, ranging from 0 to 100 percent restorability. It should be noted than none of the test networks contain 2-edge cuts.

What results from these experiments is that very capacity efficient *p*-cycle designs will tend to provide much lower and somewhat unpredictable dual span-failure restorability than designs of span-restorable networks. To limit this phenomenon, the best approach seems to impose a limit on the maximum length of *p*-cycles in the designs. A study of this is provided in [Sch03b].

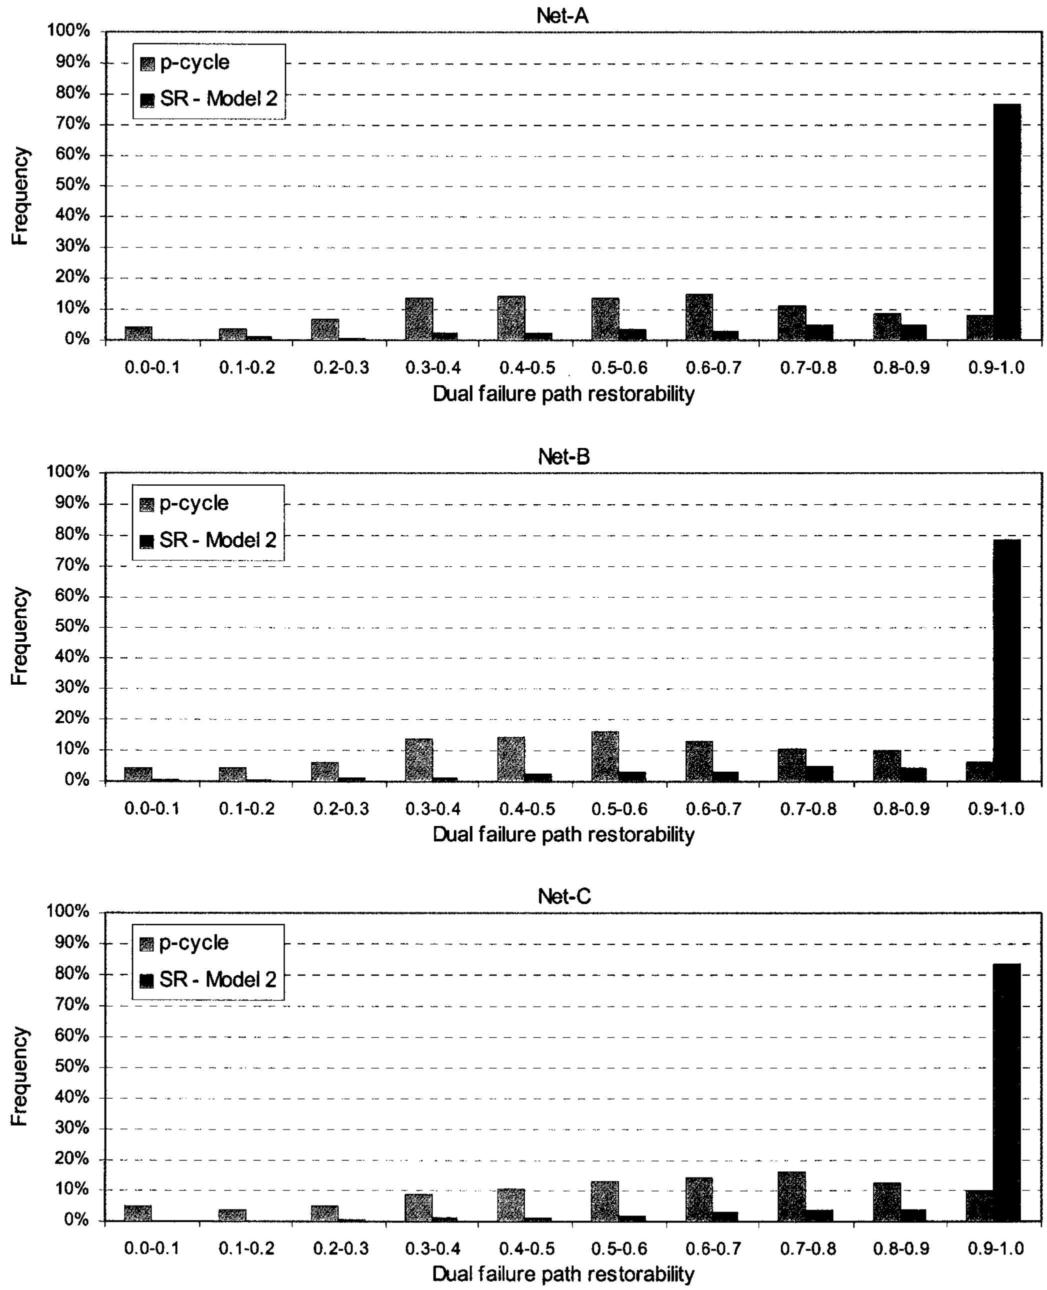

Figure 7-11 Distributions of dual-failure path restorability for *p*-cycle and span-restoration for Net-A

7.5 Summary

This chapter has presented a short theoretical study of the availability of service paths in networks protected by *p*-cycles. Results of the analysis show that the most important factor in determining the availability of service with *p*-cycles is the size of these *p*-cycles as measured by the number of on-cycle spans. The number of protected straddling spans also plays a role but has less influence than the number of on-cycle spans. This confirms the findings of recent studies that showed that dual-failure restorability of *p*-cycles could be improved by limiting the size of the *p*-cycles selected in the design.

8. Service Availability in Multi Quality of Protection (QoP) Designs

8.1 Introduction

The work presented in this chapter was originally started as a project to produce numerical results to accompany an advanced chapter draft on the topic in the recent book on mesh-based survivable networks [Gro03]. The goal of this project was to study the capacity requirements of span-restorable networks with various mixes of service types. The project led to the development of a complete study of design and restorability analysis offering a lot of interesting results and also leading to the publication in [GrC02] and [GrC04]. This chapter includes some more extensive results and more detailed documentation of the algorithms and networks than was contained in either [GrC02] or [GrC04].

An important shift in thinking, coming with a data-centric orientation in recent years, is the notion that not all demand flows in a transport network necessarily need the same level of protection or restoration. The prior view tended to be that all transport had to be equally protected, as evidenced by solutions such as SONET rings. Indeed (other than for strictly preemptible "extra traffic" on protection) that is the only paradigm that SONET rings support. The more recent view is that there is a business opportunity to support a whole range of survivability requirements, in a flexible and efficient manner. For example, lightpaths used as inter-router IP pipes may not need to be automatically protected within the optical networking layer because the IP service-layer relies on routing re-convergence or other means (MPLS for instance) within the router-peer layer for fault recovery. Other customers or operators may, however, choose to refer survivability into the transport layer for faster response and/or simpler service-layer administration and software management. Although the idea of multiple quality of protection (QoP) classes itself is not new [GeS02], we believe the work presented here is the first study of how the optimal design of span-restorable mesh network capacity would be altered to accommodate a multi-QoP demand environment and to look at quantitative aspects of how the different service classes interact in the resultant designs. Our primary interest is in how the capacity design problem changes to both recognize and exploit the presence of up to four simultaneously present service classes. The generic service classes are similar to those in [GeS02]: "gold" (assured single-failure restorability), "silver" (best effort restorability), "bronze" (non-protected, but also not interrupted), and "economy" (preemptible). More details of the rules for each class follow.

There are many performance and network strategy studies that can be conducted with the Multi-QoP capacity design models that follow. One question of interest is to characterize the levels of best effort restorability that are inherently achievable as a side-effect of designing for 100 percent span-restorability for a given volume of gold-class demand requirements. Another interesting question is the extent to which gold-class restoration requirements wind up being met by preemption of economy-class services, as opposed to the usual notion of "spare capacity" per-se. At what percentage of economy-class services, for instance, could one expect that no spare capacity, in the usual sense, is needed at all to protect gold-class services? Such networks would be of obvious commercial interest because they would contain literally no idle capacity (other than modularity overheads). All capacity would be revenue earning to some extent although economy capacity would earn less by virtue of its preemptability for the protection of gold-class services. On the other hand, one would want to assess how often the economy class services were actually interrupted as a result. These are some of the questions we address, hopefully illustrating a timely and useful additional advantage to the business case for mesh-based optical networking.

8.1.1 Prior Work on Multiple Quality of Protection Paradigms

In capacity design models [MiS98][VVD98][HeB94][IMG98][DoG00] for almost any variety of transport architecture to date, the most commonly studied type of capacity design problem assumes one service class that must be fully restorable (or protected) against any single failure scenario. Some works where multiple service classes in general have been considered include [GeS02][AKQ00][YoQ99][SSS01]. In [GeS02] Gerstel and Sasaki reiterate the same four basic QoP classes we consider and mention they go back to "reliability of service" classes first proposed for ATM. The interest in [GeS02] is, however, limited to considering an extension of rings in which the probability with which a given "best effort" connection will be protected can be assigned, allowing for reduced protection capacity in the ring. In effect the protection capacity of the ring is oversubscribed and a "rigged lottery" is run for access to protection when needed based on the customer-purchased level or probability of protection. In [AKQ00] demands for point-to-point lightpaths in a WDM network are treated in a way that maximizes profitability, i.e. the difference between revenue and routing costs. This implicitly involves a kind of multiple service quality notion but one in which protection is not considered and services are either provided or not based on operator profitability. In [YoQ99] multiple service classes are considered in an IP over WDM context but the aim is to create a higher priority packet delivery service class within an otherwise purely best effort IP packet delivery environment. Again, there is no consideration of protection aspects, only the main payload delivery QoS issues. Closest to the work presented in this chapter is [SSS01] which involves multi-QoP aspects involving what we would call the gold, silver and bronze service classes using shared backup path protection (also called disjoint path "backup multiplexing" in [SSS01]) for gold services. Best effort services are assigned a backup path only if resources are available and there is no preemptible service class. It is also assumed that wavelength conversion is not possible. The work presented here complements that in [SSS01] by now looking at the combination of four QoP classes in conjunction with dynamic adaptive span restoration based on OXC nodes that either have an OEO construction or have a pool of enough wavelength converters so that any wavelength blocking effects are negligible.

8.2 Concept of Multi-QoP

Table 8-1 lists three different schemes of QoP service-class offerings that are of interest based on the simultaneous offering of different combinations of the four QoP classes defined above. In Scheme 1, only two of the four QoP classes are actually offered: traffic is either protected or unprotected. In Scheme 2, we add a best effort class of traffic. Scheme 3 represents a capacity design environment where all four QoP service offerings are simultaneously available. The rest of this section discusses how each scheme can be supported operationally in the network. These are the operational principles that would go hand-in-hand with the capacity designs presented in Section 8.3.

Table 8-1: Three schemes of multiple QoP service classes

Scheme 1 (two levels)	Scheme 2 (three levels)	Scheme 3 (four levels)
Gold: Guaranteed protection Bronze: Un-protected	Gold: Guaranteed protection Silver: Best effort protection Bronze: Un-protected	Gold: Guaranteed protected Silver: Best effort protection Bronze: Un-protected Economy: Preemptible

First, each lightpath would be labelled as to its service class when provisioned. This is easily accommodated in concatenated OC-n payloads on lightpaths or more generally through digital wrapper or generic framing protocol overheads, or, can be conveyed out of band on an associated optical service channel (OSC). The cross-connects adjacent to each transport span would use these designations to classify the working capacity on the span. The IETF Link Management Protocol (LMP) is perfectly suited to extract this information from each lightpath at each OXC en route so the information does not need to be held in any network database or disseminated by any protocols. By being present in the lightpath overheads itself, it is self-updating at every cross-connect. If the service class changes, or the routing of a path is changed at any time, the custodial cross-connects for a span see the updated working capacity priority-composition of the span essentially the instant the new path or status change is made.

In the context of a span-restorable network (or the corresponding link-protected network derivable through distributed pre-planning [Gro94]) this means that what has previously been a single "w_i" quantity for the working capacity to be protected on span i will now be decomposed in up to four constituent types of working capacity, w_i^g, w_i^s, w_i^b, w_i^e, which the custodial cross-connects will treat functionally as follows:

- **Gold** (assured restoration): these are working capacity units that must be restored.
- **Silver** (best effort restoration): these are working capacity units that should be restored if possible within existing spare capacity, following restoration of any gold-class service capacity.
- **Bronze** (non-protected service): these are working capacity units that are ignored from a restoration viewpoint, but they are also not interrupted if they do not themselves experience a failure.
- **Economy** (preemptible services): these are working capacity units that are not protected and moreover may be seized (interrupting their service paths) and logically converted to spare capacity on spans that are not affected by a current failure if needed to satisfy gold class restoration requirements.

Note that several variations are conceivable within the basic framework of priorities. In particular economy (preemptible) capacity might alternately be defined as being available to cover both gold or silver class requirements. The only difference between gold and silver then is that gold gets first access before silver to spare and preemptible capacity. If silver were allowed to also preempt economy capacity, we suspect its QoP would be much higher than in the opposite case and there might be little to actually distinguish these two service offerings. In addition, the frequency of preemption of economy services would be even higher. Based on these considerations, we initially define the silver class as not preempting economy. Subsequent tests reverse this to assess the implications.

In yet another structure economy class capacity could itself also be defined as receiving secondary best effort restoration upon its own failure, while still being preemptible for other failures. These and other variations make for more complex operational and capacity design models but can be worked through with the same principles as follow for the four service classes as they are defined above, to serve our present discussion. Note, however, that when one is familiar with the basic multi-QoP model that follows it is not difficult to adapt the general model to study these or many other variations of interest as well.

8.3 Multi-QoP Design

8.3.1 Changes to MJCP for Scheme 1 (gold and bronze only)

The capacity design to support Scheme 1 in Table 8-1 is the easiest to handle: One uses the MJCP formulation presented in Appendix B with sets $\boldsymbol{S}$, $\boldsymbol{M}$, $\boldsymbol{D}$, $\boldsymbol{Q}^r$, and $\boldsymbol{P}_i$, parameters C_k^m, Z^m, d^r, $\zeta_i^{r,q}$, $g^{r,q}$, and $\delta_{i,j}^p$, and variables f_i^p, s_k, and n_k^m still as defined in that section.

The difference with MJCP is that the ordinary span working capacity quantities (w_i) are replaced by w_i^{g} in all but the modularity constraint and a constraint set similar to constraint set (B.3) is added for the generation of the w_i^{b} values. In effect this makes only "gold" capacity visible to the restoration process and the related spare capacity allocation considerations in the design model. Bronze capacity is essentially ignored from a restorability standpoint – it is neither protected, but nor is it cannibalized under any circumstances. It does, however, contribute ultimately to the sizing of the modular capacity requirements. Thus the MJCP model changes to become:

$$\textbf{Multi-QoP MJCP (Scheme 1):}\text{ Minimize } \sum_{k \in \boldsymbol{S}} \sum_{m \in \boldsymbol{M}} C_k^m \cdot n_k^m \tag{8.1}$$

$$\sum_{q \in \boldsymbol{Q}^r} g^{r,q} = d^r \quad \forall r \in \boldsymbol{D} \tag{8.2}$$

$$w_k^{\mathrm{g}} = \sum_{r \in \boldsymbol{D}^{\mathrm{g}}} \sum_{q \in \boldsymbol{Q}^r} \zeta_k^{r,q} \cdot g^{r,q} \quad \forall k \in \boldsymbol{S} \tag{8.3}$$

$$w_k^{\mathrm{b}} = \sum_{r \in \boldsymbol{D}^{\mathrm{b}}} \sum_{q \in \boldsymbol{Q}^r} \zeta_k^{r,q} \cdot g^{r,q} \quad \forall k \in \boldsymbol{S} \tag{8.4}$$

$$\sum_{p \in \boldsymbol{P}_i} f_i^p = w_i \quad \forall i \in \boldsymbol{S} \tag{8.5}$$

$$s_k \geq \sum_{p \in \boldsymbol{P}_i} \delta_{i,k}^p \cdot f_i^p \quad \forall i, k \in \boldsymbol{S},\ i \neq k \tag{8.6}$$

$$w_k^{\mathrm{g}} + w_k^{\mathrm{b}} + s_k \leq \sum_{m \in \boldsymbol{M}} n_k^m \cdot Z^m \quad \forall k \in \boldsymbol{S} \tag{8.7}$$

Compared to MJCP, only constraint sets (8.3), (8.4), and (8.7) have changed. Constraints (8.3) and (8.4) generate the replace the w_i^{g} and w_i^{b} values and replace the single constraint set (B.3) that generated the single class w_i values. Constraint set (8.7) replaces (B.6) to take into account the different types of working capacity.

The same changes apply to other standard design models such as for non-joint spare capacity placement (SCP, presented in Section 2.5.3), or modular non-joint spare capacity placement (MSCP, also presented in Section 2.5.3).

8.3.2 Changes to Accommodate Schemes 2 and 3

In approaching the other mixtures of QoP classes, we can first make some helpful observations about how the basic model (MJCP) would be altered to reflect the different treatments of working capacity. By this we mean, for instance, that in one constraint system preemptible capacity is equivalent to spare capacity, and so on. Table 8-2 is provided to guide and summarize our discussion of these observations on the basic model. In Table 8-2 we identify each constraint system in the above MJCP model by the role it plays and with 1, 0, or -1 notations to indicate how each class of working capacity is to be accounted for in the corresponding multi-QoP design model. For instance, the column for economy capacity w_i^e in Table 8-2 shows that in any design problem with an economy class the corresponding capacity is:

i. omitted from the restorability constraint (e.g., given a weight of 0 in row 2)
ii. treated as a credit against needed spare capacity in the spare capacity constraint (e.g., given a weight of -1 in row 3), and
iii. included in the modular capacity constraint (e.g., given a weight of 1).

Table 8-2: Three schemes of multiple QoP service classes

	Design Consideration	Example constraint(s)	w_i^g	w_i^s	w_i^b	w_i^e	s_i	Comment
1	Routability and working capacity	Equation (B.2) and (B.3)	Generated				n/a	Corresponding QoP class demand matrices
2	Restorability	Equation (B.4)	1	0	0	0	n/a	only gold is assured of restorability
3	Spare capacity	Equation (B.5)	1	0	0	-1	Generated	preemptible working capacity reduces spare capacity
4	Total capacity	Equation (B.6)	1	1	1	1	1	all capacity contributes to modular sizing

In this way, Table 8-2 actually specifies many different multi-QoP design models all of which can be derived from MJCP for any subset of service classes that is present in a given design context.

In row 1 of Table 8-2, we indicate that as with the non-priority model, the basic routing of working paths that is implied under constraints such as Equation (B.2) and (B.3) in a joint formulation, will now generate all of the respective individual w_i^g, w_i^s, w_i^b, w_i^e working capacity requirements on each span. Thus, in a revised (multi-QoP) model we can expect to see copies of Equations (B.2) and (B.3) for each specific class of working capacity in the priority scheme.

In row 2, we record the reasoning that in any of the multi-QoP models restorability will only ever be explicitly asserted for w_i^g quantities. Note in this regard that the difference between best effort (silver) and strictly non-protected (bronze) services is really only an operational distinction; there is no distinction between them from a capacity-design standpoint. Neither silver nor bronze classes strictly require any assured or built-in restoration considerations in the basic design, but silver will receive operational best effort to exploit any available and otherwise unneeded spare capacity for its restoration under the given failure scenarios, whereas non-protected bronze does not receive this effort at all. Thus, the difference between non-protected and best effort is much more one of pricing and business strategy, but it is not an issue of capacity design because neither class has any capacity built into the design for its restorability. Viewed another way, the point is that the basic spare capacity design to assure gold-class restorability will inherently support a certain amount of additional restoration capability. The distinction between best effort and non-protected is only a matter of who gets access to this limited extra restoration capability, not any structural difference in the capacity design itself. This reasoning also suggests that we can simply merge the silver and bronze working capacity requirements.

In row 3 of Table 8-2, we address the changes to the spare capacity generating constraint, Equation (B.5), indicating that restoration flows corresponding to gold-class working flows need to be fully supported by the spare (+ preemptible) capacity dimensioning and no explicit considerations of spare capacity are made for silver, bronze or economy-class demands crossing the span. On the other hand any preemptible capacity is completely inter-changeable for spare capacity on the corresponding spans. (Hence the -1 weighting on w_i^e). Finally, the last row of Table 8-2 simply recognizes that all forms of working capacity, plus spare capacity required for gold-class assured restoration, must be supported by the final modular capacity placement decisions.

As a result of these considerations we can state a generalized model for joint modular capacity design with an arbitrary mixture of demands in the four service classes. All variables and parameters are the same as in MJCP above, with the exception that the prior set of demand requirements is now represented by four constituent demand sets, one for each priority class. In practice, for the capacity design problem, this can be reduced to three by merging the best effort and non-protected requirements, based on the considerations above. Thus we have the new demand subsets:

- $\boldsymbol{D}^g$, $\boldsymbol{D}^{s,b}$, $\boldsymbol{D}^e$ are the sets of demand requirements for gold, silver and bronze (merged), and economy-class services, respectively. Individual values of these sets are still referred to as d^r values (integer), and index r continues to be used in enumerating members of any of these sets.

The sets $\boldsymbol{P}_i$ and $\boldsymbol{Q}^r$ remain unchanged because the eligible restoration routes and eligible choices for working routes remain properties of the topology alone, unaffected by the structuring of demands into different service classes. The multi-QoP capacity-design model thus becomes:

Multi-QoP MJCP: Minimize $\sum_{k \in \boldsymbol{S}} \sum_{m \in \boldsymbol{M}} C_k^m \cdot n_k^m$ (8.8)

$$\sum_{q \in \boldsymbol{Q}^r} g^{r,q} = d^r \quad \forall r \in \boldsymbol{D}^{\mathrm{g}} \cup \boldsymbol{D}^{\mathrm{s,b}} \cup \boldsymbol{D}^{\mathrm{e}} \tag{8.9}$$

$$w_k^{\mathrm{g}} = \sum_{r \in \boldsymbol{D}^{\mathrm{g}}} \sum_{q \in \boldsymbol{Q}^r} \zeta_k^{r,q} \cdot g^{r,q} \quad \forall k \in \boldsymbol{S} \tag{8.10}$$

$$w_k^{\mathrm{s,b}} = \sum_{r \in \boldsymbol{D}^{\mathrm{s,b}}} \sum_{q \in \boldsymbol{Q}^r} \zeta_k^{r,q} \cdot g^{r,q} \quad \forall k \in \boldsymbol{S} \tag{8.11}$$

$$w_k^{\mathrm{e}} = \sum_{r \in \boldsymbol{D}^{\mathrm{e}}} \sum_{q \in \boldsymbol{Q}^r} \zeta_k^{r,q} \cdot g^{r,q} \quad \forall k \in \boldsymbol{S} \tag{8.12}$$

$$\sum_{p \in \boldsymbol{P}_i} f_i^p = w_i^{\mathrm{g}} \quad \forall i \in \boldsymbol{S} \tag{8.13}$$

$$s_k \ge \sum_{p \in \boldsymbol{P}_i} \delta_{i,k}^p \cdot f_i^p - w_k^{\mathrm{e}} \quad \forall i, k \in \boldsymbol{S},\ i \ne k \tag{8.14}$$

$$w_k^{\mathrm{g}} + w_k^{\mathrm{s,b}} + w_k^{\mathrm{e}} + s_k \le \sum_{m \in \boldsymbol{M}} n_k^m \cdot Z^m \quad \forall k \in \boldsymbol{S} \tag{8.15}$$

$$f_i^p \ge 0 \quad \forall i \in \boldsymbol{S},\ \forall p \in \boldsymbol{P}_i \tag{8.16}$$

$$g^{r,q} \ge 0 \quad \forall r \in \boldsymbol{D}^{\mathrm{g}} \cup \boldsymbol{D}^{\mathrm{s,b}} \cup \boldsymbol{D}^{\mathrm{e}} \quad \forall q \in \boldsymbol{Q}^r \tag{8.17}$$

A non-modular but still jointly optimized multi-QoP model is obtained from by dropping Equation (8.15) and using the cost-weighted sum of all types of working and the spare capacity as the objective function. A multi-QoP model corresponding to SCA is further obtained by also dropping Equations (8.9) through (8.12) and generating working capacities through shortest-path routing (or any other procedural routing method), prior to optimizing only the spare capacity. Note that once the problem is reduced back to its SCP variant it is identical to an instance of conventional SCP for the gold working capacity only with effectively "existing" spare capacity quantities on spans representing the preemptible working capacities generated by pre-routing the economy class services.

8.3.3 "Extra traffic" Concept: an Approximate Analogy for the Mesh Economy Class

In row 3 of Table 8-2 we see an expression of interchangeability between spare capacity and economy class capacity in the sense that economy capacity subtracts (the -1 weighting) from the otherwise needed spare capacity on each span. In prior practice with APS systems or rings, a feature called "extra-traffic" lets an operator put temporary or low-priority traffic on the protection channel of a ring or APS system. It is interesting to digress a bit therefore to consider whether the economy class capacity is the exact mesh analogy to the "extra traffic" concept or not. After all, the argument would go: preemptible services are realized over what is for all other intents, simply the spare capacity of a span-restorable mesh designed only for the gold-class services. But it is not strictly accurate because the multi-QoP design formulation will also try to accommodate the routing of economy flows as well as trying to use their capacity for gold protection, and this is not part of the logic of "extra traffic." The point is that being a form of joint optimization problem, in multi-QoP the placement of economy class and spare capacity may differ from the distribution and amount of spare capacity alone in a pure all-gold class design. Even though the economy services will never get to use the spare capacity themselves for restoration, the placement of what we would otherwise identify as regular spare capacity in the mixed QoP case will differ so as to better accommodate the economy demand flows.

Nonetheless, thinking of the preemptible services as "extra traffic on the spare capacity of an ordinary mesh design" does suggest a useful simplification of the multi-QoP design problem, applicable when one or the other volume of gold or economy demands tends to dominate the other.

i. If the fraction of all demand that is preemptible is relatively small, one could initially ignore the preemptible service demands, generate the routing and spare capacity to serve (and protect) gold demands only, then afterwards seek routings of the economy-class service paths through the spare capacity of the gold-class restorable design. If there are economy-class demands that are un-served within the initial spare capacity then the more general multi-QoP design model should be run.

Or, conversely:

ii. If the fraction of all demand that is preemptible is relatively high, one could initially ignore the gold-class service requirements and generate a w_i^e capacitation of spans based on shortest-path routing only of the economy-class services only. The capacity design and routing for gold-class service paths then follows, viewing all already-placed w_i^e capacity as equivalent to existing spare capacity in the graph. This can be done with the conventional JCP model, solved for gold-class demands only, with the simple addition of lower bounds on logical spare capacity on each span to represent the preemptible working capacity already present on those spans. If the total spare capacity added by the instance of JCP with the added lower bounds is significant, then that is an indication that the full multi-QoP design model may be preferable. But if little or no spare capacity is added by the JCP solution, it is an indication that a fully restorable routing of gold-class working demand requirements is feasible using only preemptible working capacity for the protection of gold-class services.

8.4 Multi-QoP Experimental Design Studies

As mentioned, there are many performance and network strategy studies that can be conducted with an implementation of Multi-QoP MJCP. Our aim in this section is to summarize a set of test case studies that allow us to look at some of the interesting effects and trade-offs that exist in a multi-QoP environment depending on the service mix served by the network. We are especially interested in quantifying the best effort restorability that the silver class can expect and in trying to assess the service mix at which conventional spare capacity might be almost completely replaced by preemptible economy-class services over a network as a whole. We are also interested in measuring the exposure of the economy class to preemption. To provide a sample of results related to these networking questions involving multiple QoP strategies, we compared a conventional (single-priority) fully restorable joint but non-modular design against the corresponding non-modular multi-QoP designs with four different scenario profiles of multi-priority demand mix. The non-modular context is chosen to observe intrinsic effects of shifting priority mixes on total required channel capacity and spare capacity, without effects from assumptions of specific module sizes or economy of scale. In practice, as mentioned earlier, a non-modular design can also be obtained using a modular design model (e.g. MJCP or multi-QoP MJCP) using a single module type of unit capacity.

The four mixed demand scenarios are represented in the star-plots of Figure 8-1. The test case demand patterns are constructed by assuming a total of 20 units of demand of all types between each node-pair. For the benchmark capacity design all 20 are gold class on each node pair. The oth-

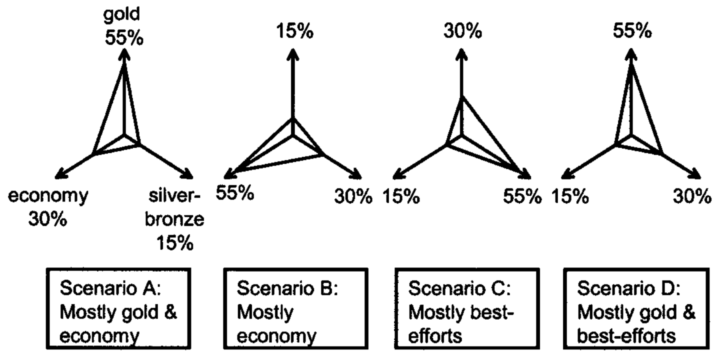

Figure 8-1 Four multi-QoP mixed demand scenarios for testing

er (mixed-class) demand scenarios are generated by allocating different numbers of the 20 demand units on each pair to each class: 55 percent corresponds to 11 units, 30 percent is 6 units and 15 percent is 3 units. Tests were conducted using each demand scenario on each of the test networks Net-A, Net-B, and Net-C, presented in Section 1.4.1 and fully described in Appendix A. Test network A has 20 nodes and 40 spans, network B has 25 nodes and 50 spans, and network C has 30 nodes and 60 spans. In each of these random test-case networks, edge distances are in proportion to the graphical presentation of each network exactly as shown and are used as the basis for (capacity times distance) measures of cost.

For each of the four multi-QoP demand scenarios, on each test network, the corresponding design model based on multi-QoP MJCP for each priority scheme was implemented in AMPL and solved with Parallel CPLEX 7.1 MIP Solver on a multiprocessor (4 × 450 MHz) UltraSparc Sun Server with 4 GB of RAM (under Sun Solaris O/S 2.6). In all formulations the set of 10 shortest distinct eligible restoration routes between the end-nodes of a failure span were represented for each failure scenario. For routing of working paths the set of 10 shortest distinct eligible routes between each end-node pair was represented. As is common in such problems we kept demand parameters and capacity variables integer but relaxed restoration flows. (It is known that although relaxed, restoration flows very seldom become fractional under the integer capacity and moreover that even if they do, an equivalent-cost solution exists with fully integral restoration flows). Results are based on a full CPLEX termination with a MIPGAP under 0.001 (i.e. solutions are provably within 0.1 percent of optimal). Run times were under a minute for all solutions.

The capacity design results are summarized in Tables 8-3 to 8-5 for test networks A, B, C, respectively, under each of the mixed-demand scenarios. For each scenario plus a conventional ("all gold") reference design, the corresponding row gives the total distance-weighted capacity costs of working capacity for each protection class (and its percentage breakdown), followed by the total amount of explicit spare capacity provisioned, and the total of all capacity (working plus spare). These percentage values are generally close to the proportions of the corresponding protection class in the total demand mix, but in some cases they are slightly different and – as it will be explained later – this gives us some insight about the way these demands are routed in the different network designs.

One of the first things we notice is that there is a major reduction of the spare capacity requirements for all four scenarios compared to the conventional case. In demand scenarios A, B and C in particular, there is literally none or very little conventional spare capacity required at all. Instead of spare capacity per-se we see the bulk of the gold restoration requirements being provided by preemption of economy capacity. This means that more of the total capacity investment is in use for working capacity of some type or other. This is obviously a preferable situation from an operator's standpoint compared to having large amounts of non-revenue producing spare capacity. An economic interpretation of the benefit follows in Section 8.6.3.

Table 8-3: Test case results comparing multi-priority span-restorable designs for Net-A

Demand mix scenario	Total working capacity	Gold class working capacity	Silver class working capacity	Econ. class working capacity	Total spare	Total capacity
Conventional	11890.80	11890.80 (100.0%)	0.00 (0.0%)	0.00 (0.0%)	6095.33	17986.12
A – Mostly gold and economy	11896.37	6450.85 (54.2%)	1698.26 (14.3%)	3747.26 (31.5%)	0.00	11896.37
B – Mostly economy	11321.75	1698.26 (15.0%)	3396.52 (30.0%)	6226.96 (55.0%)	0.00	11321.75
C – Mostly best effort	11708.95	3538.85 (30.2%)	6226.96 (53.2%)	1943.14 (16.6%)	0.00	11708.95
D – Mostly gold and best effort	12701.22	6540.95 (51.5%)	3396.98 (26.7%)	2763.28 (21.8%)	588.97	13290.19

Table 8-4: Test case results comparing multi-priority span-restorable designs for Net-B

Demand mix scenario	Total working capacity	Gold class working capacity	Silver class working capacity	Econ. class working capacity	Total spare	Total capacity
Conventional	21267.19	21267.19 (100.0%)	0.00 (0.0%)	0.00 (%)	9258.42	30525.60
A – Mostly gold and economy	20749.49	11414.09 (55.0%)	3051.30 (14.7%)	6284.11 (30.3%)	0.00	20749.49
B – Mostly econ-omy	20341.99	3051.30 (15.0%)	6102.60 (30.0%)	11188.09 (55.0%)	0.00	20341.99
C – Mostly best effort	20648.13	6255.55 (30.3%)	11188.21 (54.2%)	3204.37 (15.5%)	0.00	20648.13
D – Mostly gold and best effort	22485.46	11699.47 (52.0%)	6102.99 (27.1%)	4683.01 (20.8%)	407.05	22892.50

Table 8-5: Test case results comparing multi-priority span-restorable designs for Net-C

Demand mix scenario	Total working capacity	Gold class working capacity	Silver class working capacity	Econ. class working capacity	Total spare	Total Capacity
Conventional	32518.27	32518.27 (100.0%)	0.00 (0.0%)	0.00 (0.0%)	15800.16	48318.43
A – Mostly gold and economy	32450.38	17551.27 (54.1%)	4617.84 (14.2%)	10281.28 (31.7%)	104.61	32555.00
B – Mostly econ-omy	30783.22	4617.49 (15.0%)	9234.96 (30.0%)	16930.76 (55.0%)	0.00	30783.22
C – Mostly best effort	31909.28	9634.37 (30.2%)	16930.80 (53.1%)	5344.12 (16.7%)	17.28	31926.56
D – Mostly gold and best effort	33699.69	17886.72 (53.1%)	9235.15 (27.4%)	6577.81 (19.5%)	2110.75	35810.44

While spare capacity decreases, the working capacity totals can in fact rise under the multi-QoP designs. This is because in a jointly optimized capacity design, demands are not necessarily routed on their shortest path. Even in the conventional single QoP case, a lightpath may deviate from its shortest path if by doing so it can increase the sharing of spare capacity with other working demands. In the mixed QoP scenarios, the reasons to divert from shortest paths can be even greater. In particular, economy and gold class demands may both deviate so as to enhance the synergy between themselves even more under the multi-QoP optimization model. (The routing of an economy-class service path will be fairly strongly influenced to go over spans where its channels can be best used for gold restoration, as opposed to a shorter route without this added cost-leverage). Silver (and bronze) demands, however, have no reason not to continue to be routed on their shortest routes as they do not require any (built-in) restorability, nor do they contribute to the restorability of any other demand class – although if specific modularity assumptions are introduced they will tend to deviate to coordinate routes so as to use fewer capacity modules.

Depending on the mix of demands the tendency to deviate demands from their shortest path is more or less pronounced. In the case of scenario B, it appears that demands are essentially all routed on their shortest paths. This is inferred because the working totals are equal to the minimum possible (as shown in the different rows of the silver class working column) and the proportions are exactly equal to the proportions of each demand class. Scenario B therefore has the minimum working capacity requirement possible. At the other extreme scenario D shows the highest levels of deviation (as can be seen right away from the percentages of each working capacity total). Interestingly, it appears that the economy-class demands are the ones that are most deviated (between 42 and 62 percent higher total working capacity than in the shortest path case). In the case of scenario A and C, the working capacity total is generally closer to the conventional case but we again see that the economy-class demands are significantly more deviated than the gold-class demands. Overall this confirms that, not too surprisingly, the routes for economy-class services have to be specifically coordinated with the current set of gold class demands for best economic benefit in multi QoP operation where gold class can preempt economy class. Importantly, however, even though working capacity may rise for the reasons given, it only ever does so when the net effect is a reduction in total capacity cost, including spare. The total capacity values are in all cases much lower than in the conventional case (especially for scenario A, B and C).

8.5 Multi-QoP Restoration Schemes

We identify two main options regarding restoration with multiple QoP classes:

i. *Direct use of preemption* vs. *use of spare links only first*: With *direct use of preemption,* restoration paths are searched directly using spare links and/or economy-class working links, whereas with *use of spare links only first* restoration paths are first searched using spare capacity only and a second restoration path search allowing the preemption of economy-class working links is performed only if additional restoration paths are needed. The advantage of preempting economy class working links directly will be a more efficient use of the total (spare + economy) capacity available for forming restoration paths. This could result in higher restorability levels for the gold-class. The advantage of using spare links only first will be to avoid preempting economy-class working links when it is not strictly needed, therefore improving the availability of economy class service paths. This, however could result in smaller numbers of restoration paths found. Use of spare links only first could also potentially make restoration slower. Assessing the speed of restoration, however, is outside of the scope of this study. Before moving to the next option, it should be mentioned that the approach that uses direct preemption of economy class services will still favour the use of true spare capacity links over preemption of economy links whenever it is possible (for example if two equally short restoration paths are available with different numbers of preempted economy-links on them).

ii. *No preemption for best effort restoration* vs. *preemption allowed for best effort restoration*: As mentioned earlier, we have the choice of whether to limit preemption to only benefit restoration of services in the Gold class or to allow it also for restoration of services in the Silver class. In the case of preemption being allowed for best effort restoration, the only difference between gold and silver-class working units would be that gold-class channels are restored first in real time, silver-class subsequently. In both cases the capacity design model only guarantees the restorability of gold.

By combining these options we obtain four different multi-QoP restoration models as described in Table 8-6:

Table 8-6: Three schemes of multiple QoP service classes

	Direct use of Preemption	Spare links only first
No preemption to restore Silver services	Multi-QoP Model 1	Multi-QoP Model 3
Preemption to restore Silver services	Multi-QoP Model 2	Multi-QoP Model 4

As already mentioned, the two main questions we have in this study are how restorable to failures best effort services will be and how much Economy-class services will suffer from preemption. Before running any experiments we can already devise about how the differences between the four models should affect these two outcomes. For example, we can expect that Model 3 will be the worse for the restorability of silver class services. Indeed, restoration with that model will make use of as much spare capacity as it can to restore Gold-class services without considering keeping some spare capacity to later restore best effort services. If any spare capacity is left after restoration of the Gold-class services, it can then potentially be used in a second round of restoration path search using preemption but still also using spare capacity where possible. Model 1 should be the second worse model in terms of the restorability of Silver-class services. Although presumably using less spare capacity than Model 3, the restorability of Silver-class is likely to be much higher with Models 2 and 4 that allow preemption of economy services to restore them. In terms of Silver-class restorability, the ranking of the models in order of increasing restorability should be: Model 3, Model 1, Model 4, Model 2. The reason for placing Model 4 before Model 2 is that by allowing direct use of preemption we expect that a more efficient use of combined spare capacity and economy links will be made to form restoration paths and that it should therefore benefit to the restorability of Silver-class restorability. In terms of the frequency of preemption of the Economy-class services it seems logical to expect that the ranking of the models in order of increasing frequency will be: Model 3, Model 1, Model 4, Model 2. Not surprisingly the two rankings are the same, indicating that there is a clear relation between the frequency of Economy-class preemption and the restorability of Silver-class services.

Experimental results presented in Section 8.6 will confirm whether the above logical reasoning is correct and will inform clarify how much the type of restoration model used will influence the two factors under study. But first, we describe the four models in details.

8.5.1 Description of Multi-QoP Restoration Model 1

Model 1 for multi QoP restoration is specified in Figure 8-2. As in previous chapters, the restoration path search is performed using the ksp algorithm. What changes between the different stages at which restoration path search is performed is the type of capacity links within which the search is performed. With Model 1, the search is performed twice since the paths found during the first search may contain some preempted economy-class links that would not be suitable to restore Silver-class service. The second path search only considers true spare capacity links as eligible to form restoration paths.

1. Search for restoration paths using spare and/or economy working links.
2. Restore Gold-class working links until either all these links are restored or all found restoration paths have been used.
3. If all Gold-class working links have been restored, search for restoration paths using spare links only, otherwise stop.
4. Restore as many Silver-class working links as possible using restoration paths found in 3.

Figure 8-2 Details of Multi-QoP Restoration Model 1

8.5.2 Description of Multi-QoP Restoration Model 2

This model is identical to Model 1 except that the third and fourth points are replaced by:

3. If all Gold-class working links have been restored, restore as many silver-class working links using the remaining restoration paths found in 1.

8.5.3 Description of Multi-QoP Restoration Model 3

The third model is described in Figure 8-2. Here, restoration paths are first searched using spare links only. The resulting restoration paths are used to restore Gold-class (or higher) services until either all Gold-class services have been restored or all restoration paths have been used. If all not all Gold-class services have been restored, a new restoration path search is performed, this time allowing preemption of Economy-class service where needed. These restoration paths are used to restore as many Gold-class working links as possible. Following that, there is no need to make any attempt at restoring Silver-class working links since no restoration paths using only spare capacity can possibly be found. If all Gold-class working links were restored using restoration paths using

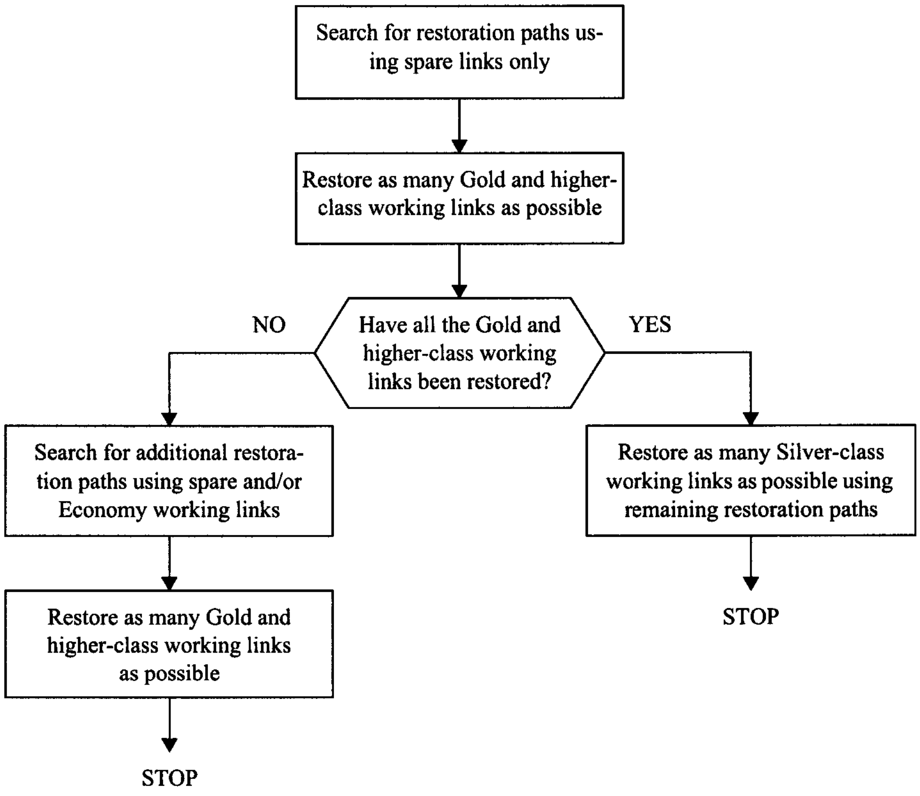

Figure 8-2 Description of Multi-QoP Restoration Model 3

spare links only and some restoration paths are left, these are used to restore as many Silver-class working links as possible.

8.5.4 Description of Multi-QoP Restoration Model 4

Despite the apparent greater complexity of restoration Model 4, shown on Figure 8-3, compared to Model 3, this algorithm is in fact very simple. In this model, a first restoration path search is performed using spare capacity only and the resulting paths are used to restore as many of the Gold and Silver-class working links. If there are not enough restoration paths to restore links in both classes a new restoration path search is performed using spare and Economy-class working links and the restoration paths found are used to restored as many remaining working links.

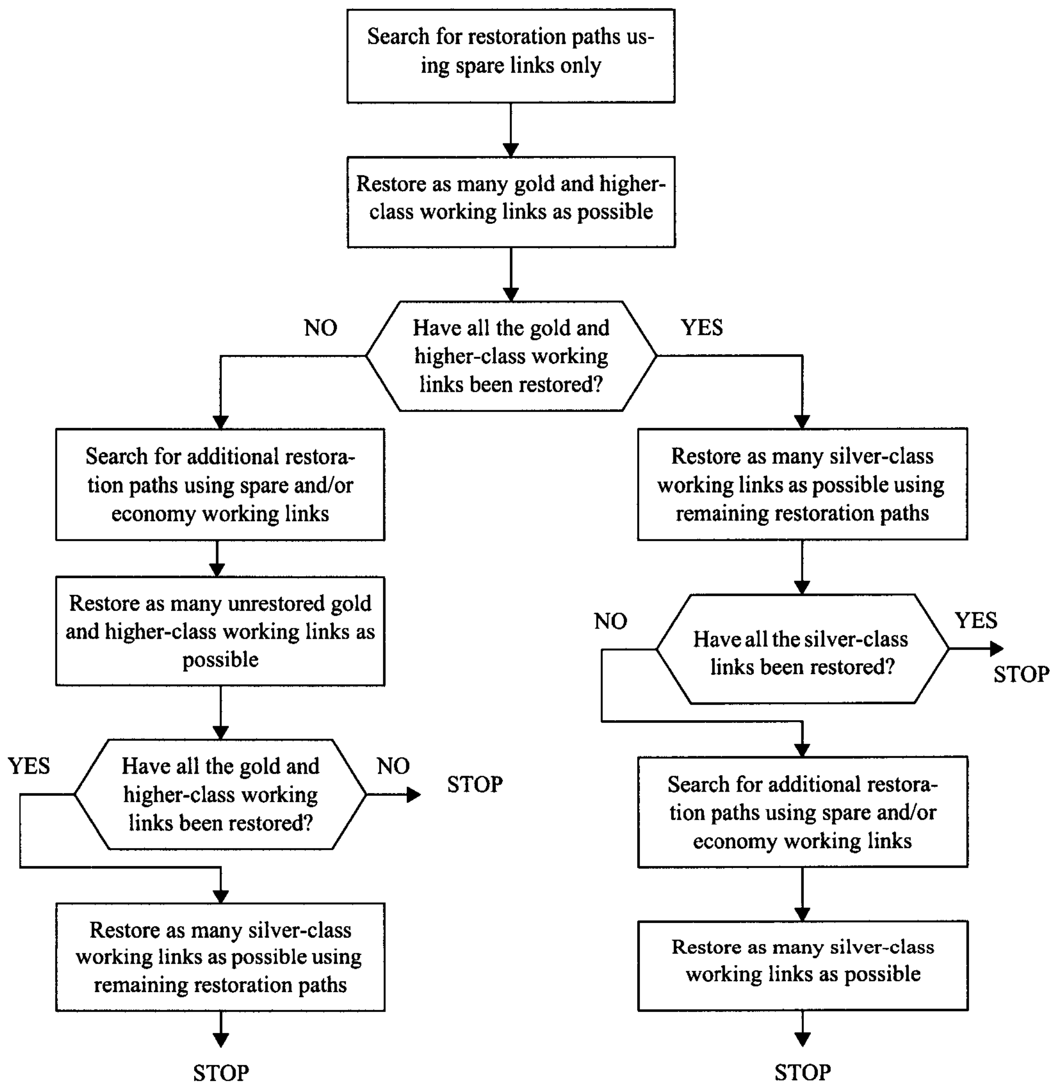

Figure 8-3 Description of Multi-QoP Restoration Model 4

8.6 Experimental Results of Multi-QoP Restoration Analysis

8.6.1 Initial Experiments

For experiments of multi-QoP restorability analysis, we first tested the designs presented in Section 8.4 for their various service-class restorability levels.

As discussed previously, the capacity design model itself has no direct concern with the restorability of the silver class of service. Silver-class restorability and other restoration related results were obtained by experiments on the resulting network designs using a custom-made program. That program uses the k-shortest paths (ksp) algorithm to find the maximum feasible restoration flow through the spare capacity of the surviving graph between the end-points of each span failure. The restorability levels of each service class are summarized in Tables 8-7 to 8-9. Each table presents the restorability results for one of the test networks with the four different restoration models. The R_1^s column gives measures of the average restorability levels of silver-class demands following restoration of all gold-class capacity on the same span. Restorability results of the gold-class are omitted as they were always 100 percent by design (and separately validated as such).

The P_{pre}^e column gives the probability of preemption of economy-class working channels. This is defined as the average over all span failures of the fraction of economy-class working channels on all other spans that are taken over for gold (or silver) restorability requirements. In other words this allows us to answer the question: "if an economy class path y crosses span x, what is the probability that, given a span failure elsewhere in the network, path y will be preempted on span x?" Thus, an economy path experiences an end-to-end probability of a preemption that is the union of the preemption probabilities on the spans it crosses. The absolute frequency of such preemption outages thus depends also on the absolute failure rate of spans, and may thus remain acceptably infrequent for useful "economy" applications.

More precisely, if we assume there are H spans in the economy service path, and S spans in the network, each span in the economy path contributes 1/MTBF to the frequency of outage (since the failure of a span on an economy path is certain to cause outage). But in addition an economy path runs a risk of preemption outage from the $(S-H)$ other spans which each contribute another 1/MTBF weighted by the probability that path y is preempted on at least one of its spans. This probability can be expressed as follows:

$$P^{pre}(y) = 1 - \prod_{i \in \text{path } y} P(y \text{ not preempted on span } i) \quad (8.18)$$

$$P^{pre}(y) = 1 - (1 - P^{pre})^H \quad (8.19)$$

Table 8-7: Test case results comparing multi-priority span-restorable designs for Net-A

Demand mix scenario	Restoration Model 1 (no preemption for Silver, direct preemption)				Restoration Model 2 (preemption for Silver, direct preemption)				Restoration Model 3 (no preemption for Silver, spare only first)				Restoration Model 4 (preemption for Silver, spare only first)			
	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.
Conventional	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
A – Mostly gold and economy	0.0	14.2	100.0	0.0	37.0	16.3	100.0	100.0	0.0	14.2	100.0	0.0	37.0	16.3	100.0	100.0
B – Mostly economy	0.0	1.7	100.0	0.0	97.1	5.9	100.0	100.0	0.0	1.7	100.0	0.0	97.1	5.9	100.0	100.0
C – Mostly best effort	0.0	15.1	100.0	0.0	13.0	20.0	100.0	98.2	0.0	15.1	100.0	0.0	13.0	20.0	100.0	98.2
D – Mostly gold and best effort	0.9	15.3	77.7	0.0	24.0	18.0	77.7	75.0	0.9	15.2	76.6	0.0	24.0	18.0	76.6	75.0

(all values are percentages)

Table 8-8: Test case results comparing multi-priority span-restorable designs for Net-B

Demand mix scenario	Restoration Model 1 (no preemption for Silver, direct preemption)				Restoration Model 2 (preemption for Silver, direct preemption)				Restoration Model 3 (no preemption for Silver, spare only first)				Restoration Model 4 (preemption for Silver, spare only first)			
	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.
Conventional	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
A – Mostly gold and economy	0.0	11.6	100.0	0.0	48.4	14.1	100.0	100.0	0.0	11.5	100.0	0.0	48.4	14.1	100.0	100.0
B – Mostly economy	0.0	1.4	100.0	0.0	98.1	4.7	100.0	100.0	0.0	1.4	100.0	0.0	98.1	4.7	100.0	100.0
C – Mostly best effort	0.0	12.6	100.0	0.0	12.2	16.6	100.0	100.0	0.0	12.6	100.0	0.0	12.2	16.6	100.0	100.0
D – Mostly gold and best effort	0.0	15.3	92.7	0.0	13.6	17.1	92.7	95.4	0.0	15.3	92.7	0.0	13.6	17.1	92.7	95.4

(all values are percentages)

Table 8-9: Test case results comparing multi-priority span-restorable designs for Net-C

Demand mix scenario	Restoration Model 1 (no preemption for Silver, direct preemption)				Restoration Model 2 (preemption for Silver, direct preemption)				Restoration Model 3 (no preemption for Silver, spare only first)				Restoration Model 4 (preemption for Silver, spare only first)			
	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.
Conventional	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
A – Mostly gold and economy	0.0	9.8	98.4	0.0	41.7	11.2	98.4	98.0	0.0	98.8	98.4	0.0	41.7	11.2	98.4	98.0
B – Mostly economy	0.0	1.2	100.0	0.0	93.8	4.3	100.0	100.0	0.0	1.2	100.0	0.0	93.8	4.3	100.0	100.0
C – Mostly best effort	0.0	10.4	99.5	0.0	17.0	14.6	99.5	100.0	0.0	10.4	99.5	0.0	17.0	14.6	99.5	100.0
D – Mostly gold and best effort	0.8	11.7	75.0	0.0	28.8	14.6	75.0	74.6	0.2	11.7	74.3	0.0	28.8	14.6	74.3	74.1

(all values are percentages)

and thus, the average frequency of outage of an economy service path (due to both failure and preemption risk) is:

$$Freq(y) = \frac{H}{MTBF} + \frac{S-H}{MTBF} \cdot (1 - P^{\text{pre}})^{H} \qquad (8.20)$$

The "Propor. preemp. for gold" (and "for silv.") columns give the average percentage of the capacity required for gold (and silver) restoration paths that is met by preemption of economy capacity. Of course, under restoration Model 1 and Model 3, this is always 0 by definition.

Looking at the restorability results of Tables 8-7 to 8-9 it appears right away that the results obtained with restoration Model 1 and restoration Model 2 are almost identical to those of restoration Model 3 and restoration Model 4 respectively. This indicates that in the designs used for these experiments it makes very little difference whether "direct preemption" or use of "spare links only first" is chosen. This is not surprising given the details of the capacity design we showed in Tables 8-3 to 8-5. There is indeed so little spare capacity in the designs that the initial search of restoration paths using spare only is bound to make absolutely no difference. In real life though, due to modularity, networks usually have some slack capacity that can be considered as true spare capacity. In Section 8.6.2 we will present restorability results obtained with the same designs with modularity taken into account. For now we will comment on the differences between the results for Model 1 and Model 2, knowing that the same comments in this case apply to Model 3 and Model 4 respectively.

Another fact that appears immediately is that Silver-class restorability with restoration Model 1 is virtually zero. Again, this is not surprising because the designs include little or no ordinary spare capacity (which Silver can access on a best effort basis) and Silver cannot preempt economy in thus model. In designs that do have some spare capacity, for instance Net- C, under scenario D, (Table 8-5), the Silver restorability (Table 8-9) is still very low (~1 percent). This is not only because the bulk amount of spare capacity left over from gold restoration is relatively small, but also because what remains available is sparsely distributed, so that piecing it together to form coherent restoration paths for Silver is extremely difficult.

Under restoration Model 2, however, Silver class is allowed to preempt economy-class working capacity (following Gold requirements) and this leads to very much higher best effort restorability, especially demand scenarios A and B which have the most economy class demand. The higher the proportion of economy services the higher the best effort restorability following gold. Scenario A, with an average best effort restorability of 42 percent may be indicative of truly viable commercial situations. At the same time it is noteworthy that this major improvement of Silver class restorability (over restoration model 1) does not require a huge increase of the probability of preemption of economy-class capacity. P_{pre}^{e} rises an average from 11.9 percent to only 13.9 percent. This is explained by the relatively small proportion of silver class demands in that scenario and the synergy of being able to preempt economy channels selectively to "bridge together" remaining amounts of the true spare capacity still available following gold restoration.

Most notably, under scenario B, the best effort restoration enjoyed by Silver-class services is 94 to 98 percent, almost as good as the 100 percent restorability guarantee that gold services have presumably paid more to obtain. This might not be desirable from a marketing point of view because gold and silver service are not sufficiently distinguished. This could be where pricing strategy comes into play to encourage a demand mix more like scenario D than B. Or conversely, although its treatment is beyond the scope of the present paper, Gold-class customers may remain satisfied with paying a premium when they consider how they will fare under dual-failure scenarios relative to best effort services. From Chapter 3, we know that the unavailability of a path that is fully restorable to single failures will in effect "jump" to dependence on the probability of a dual-span failure. A service path that is not quite completely restorable to single failures will continue to have an unavailability that remains proportional to the probability of a single span failure. This means that as soon as a demand becomes fully restorable to single failures, its unavailability is reduced in a quantum step-like way, whereas partial restorability yields only a proportional benefit. The fact only 15 percent of customers are Gold-class in scenario B also makes it more plausible that these are the customers who want the highest assured availability possible. Finally, note that in scenario B, P_{pre}^{e} for economy-class is quite low (an average of 4.9 percent) which will tend to make that service class relatively attractive. It is also interesting to note that the higher the proportion of economy-class demands, the lower the probability of preemption for that class will be and therefore the more attractive that class will appear (from the standpoint of preemption outage frequency).

On the other hand, under scenarios C and D, the best effort Silver restorability even under restoration Model 2, could easily be judged as being too low to be of much significance (12 to 29 percent). Ironically the Economy-class probability of preemption in both these scenarios is also noticeably higher than in scenario A and much higher than in scenario B. Of course, one option that always exists is to tailor the best effort restorability level with a designed-in increment of true spare capacity. The multi-QoP MJCP design formulation given can be fairly easily extended in this regard to add a constraint asserting fractional target restorability levels for Silver-class demands.

8.6.2 Experimental Results with Post-Modularized Designs

In order to observe the effects of searching restoration paths using spare links only first instead of doing direct preemption, we created modular test designs based on the designs presented in the previous sections. Instead of generating new designs using the modular capacity placement formulation, we simply "post modularized" all the designs presented so far using capacity modules of size 48 and 192. Concretely, what this means is that for each design we calculated the total capacity on each span and we added on each of them as many modules of size 192 as possible without exceeding that total capacity and completed with the minimum required number of modules of size 48. The amounts of spare capacity and of extra capacity that this post-modularization produced on the different designs are given in Table 8-10. The working capacity values are still the same as in Tables 8-3 to 8-5.

Table 8-10: Percentage of extra capacity in post-modularized designs

	Scenario A		Scenario B		Scenario C		Scenario D	
	Total spare	Total capacity increase	Total spare	Total capacity increase	Total spare	Total capacity increase	Total spare	Total capacity increase
Net-A	1315.35	11.1%	1379.04	12.2%	1231.45	10.5%	1853.83	9.5%
Net-B	1797.59	8.7%	1516.50	7.5%	1315.14	6.4%	2044.86	7.2%
Net-C	1777.06	5.1%	1530.16	5.0%	1653.77	5.1%	3705.85	4.5%

The restorability analysis program was used again to produce experimental results for these new designs. Results of these experiments are presented in Tables 8-11 to 8-13.

First off, the effects of having added some spare capacity are obviously that the restorability of Silver-class services has increased and the probability of preemption of Economy-class services has decreased. Also, looking at the proportions of preempted capacity in restoration paths for Gold and Silver, we can observe that these values have now fallen from 100 percent (or close to 100 percent) down to values in the range 30 to 70 percent in most cases. Ironically but understandably, the proportions of preempted capacity in restoration paths for the Silver-class (when allowed) is usually higher than for the Gold-class. This time, some effects of using "spare only first" as opposed to "direct preemption" can be observed. The proportions of preempted capacity in Gold restoration paths for Model 3 and Model 4 are on average 19.7 percentage points below those observed with Model 1 and Model 2. For the proportion of preempted capacity in Silver restoration paths, however, the effects of using the "spare only first" strategy are not as what could have been expected. In-

Table 8-11: Test case results comparing multi-priority span-restorable designs for post-modularized Net-A

Demand mix scenario	Restoration Model 1 (no preemption for Silver, direct preemption)				Restoration Model 2 (preemption for Silver, direct preemption)				Restoration Model 3 (no preemption for Silver, spare only first)				Restoration Model 4 (preemption for Silver, spare only first)			
	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.
Conventional	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
A – Mostly gold and economy	28.3	8.1	62.7	0.0	96.0	11.8	62.7	73.9	3.4	6.8	47.2	0.0	96.0	10.7	47.2	83.3
B – Mostly economy	40.5	0.8	46.2	0.0	100	3.7	46.2	75.0	21.0	0.1	4.5	0.0	100	1.6	4.5	35.5
C – Mostly best effort	11.9	6.7	49.7	0.0	44.7	16.8	49.7	59.7	3.5	4.4	26.4	0.0	44.7	14.4	26.4	63.2
D – Mostly gold and best effort	20.4	8.3	49.7	0.0	70.5	14.0	49.7	56.7	4.0	6.2	32.6	0.0	70.5	12.6	32.6	67.8

(all values are percentages)

Table 8-12: Test case results comparing multi-priority span-restorable designs for post-modularized Net-B

Demand mix scenario	Restoration Model 1 (no preemption for Silver, direct preemption)				Restoration Model 2 (preemption for Silver, direct preemption)				Restoration Model 3 (no preemption for Silver, spare only first)				Restoration Model 4 (preemption for Silver, spare only first)			
	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.
Conventional	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
A – Mostly gold and economy	21.2	7.1	67.3	0.0	88.6	9.9	67.3	70.7	2.3	6.5	56.7	0.0	88.6	9.4	56.7	77.7
B – Mostly economy	24.5	0.8	59.3	0.0	100	3.4	59.3	82.9	12.2	0.5	27.1	0.0	100	2.7	27.1	68.3
C – Mostly best effort	8.5	6.2	57.7	0.0	38.9	13.7	57.7	63.6	2.9	5.2	43.5	0.0	38.9	12.5	43.5	62.4
D – Mostly gold and best effort	7.7	8.4	53.3	0.0	58.8	13.9	53.3	61.9	0.6	7.4	43.6	0.0	58.8	13.3	43.6	70.8

(all values are percentages)

Table 8-13: Test case results comparing multi-priority span-restorable designs for post-modularized Net-C

Demand mix scenario	Restoration Model 1 (no preemption for Silver, direct preemption)				Restoration Model 2 (preemption for Silver, direct preemption)				Restoration Model 3 (no preemption for Silver, spare only first)				Restoration Model 4 (preemption for Silver, spare only first)			
	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.	R_1^s	P_{pre}^e	Propor. preemp. for gold	Propor. preemp. for silv.
Conventional	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a	n/a
A – Mostly gold and economy	11.5	6.9	74.0	0.0	82.8	9.5	74.0	80.8	2.6	6.4	63.1	0.0	82.8	9.0	63.1	82.0
B – Mostly economy	18.1	0.8	67.9	0.0	98.3	3.4	67.9	88.7	4.2	0.5	29.9	0.0	98.3	2.8	29.9	81.0
C – Mostly best effort	4.6	5.8	62.6	0.0	37.3	12.6	62.6	72.4	1.1	5.0	48.3	0.0	37.3	11.7	48.3	71.3
D – Mostly gold and best effort	10.9	7.4	49.7	0.0	56.2	11.8	49.7	58.6	3.3	6.6	40.8	0.0	56.2	11.2	40.8	62.1

(all values are percentages)

deed, unlike for Gold restoration paths, using the "spare only first" strategy does not necessarily reduce the proportion of preempted capacity in Silver restoration paths. This is, however, fairly easily explained by the fact that most of the true spare capacity in that case is used for the formation of Gold restoration paths and the strong decrease of the proportion of preempted capacity in Gold restoration paths creates an increase of the proportion of the preempted capacity in Silver restoration paths as a side effect. The net effect of using "spare only first" on the probability of preemption of the Economy-class is however positive as can be seen by comparing values in the P^{e}_{pre} columns for Model 1 and Model 3 to those in the P^{e}_{pre} columns for Model 2 and Model 4, respectively. What results from choosing the "spare only first" strategy is the reduction of P^{e}_{pre} by an average of 1.01 percentage. Choosing the "spare only first" strategy also has an impact on the restorability of the Silver class in the case of Model 3. What is observed in a significant reduction of the Silver-class restorability when "spare only first" is used. This phenomenon makes sense since the effect of using "spare only first" is that less spare capacity is left after restoration of the Gold-class and the Silver-class is not allowed to preempt Economy-class service in Model 3. Using "spare only first" in Model 4, however, does not affect the restorability of the Silver-class compared to Model 2.

The results presented in Tables 8-11 to 8-13 thus confirm our forecast about the ranking of the restoration Models in terms of their effects on Silver-class restorability and Economy-class probability of preemption. The trade-off between Silver-class restorability and Economy-class probability of preemption is shown on Figures 8-4 to 8-5. On these figures, each point corresponds to one of the four restoration models. From left to right, the models always appear in the same order as predicted: Model 3, Model 1, Model 4, and Model 2. From the observation of these curves, it appears clearly that restoration Model 4 is preferable to restoration Model 2 since it brings some improvement to the probability of preemption of the Economy-class (up to 2.4% in the best case) without degrading the restorability of the silver class. Except Model 2, all restoration models could be considered valid choices depending on whether a high Silver-class restorability or a low Economy-class probability of preemption is judged more important. The choice will obviously depend on the traffic mix scenario.

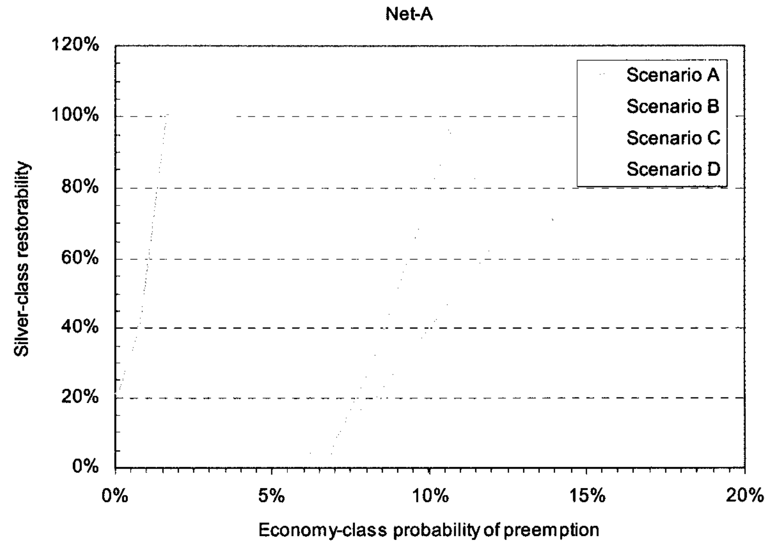

Figure 8-4 Trade-off between Silver-class restorability and Economy-class probability of preemption

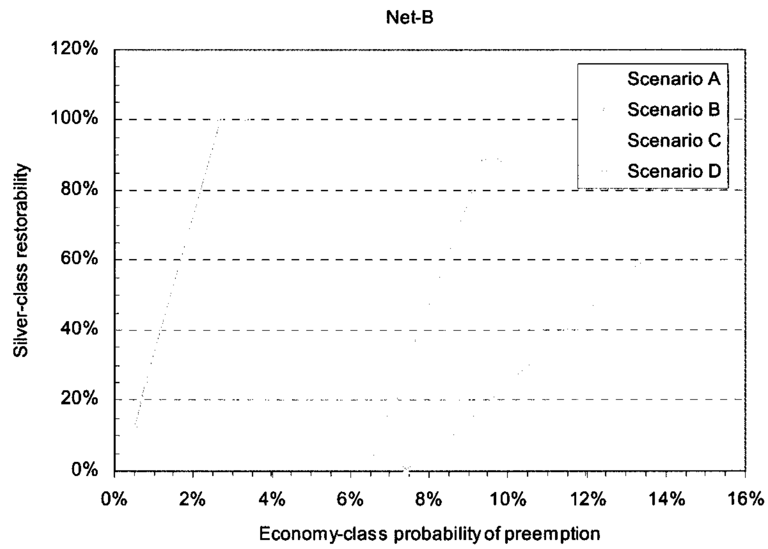

Figure 8-5 Trade-off between Silver-class restorability and Economy-class probability of preemption

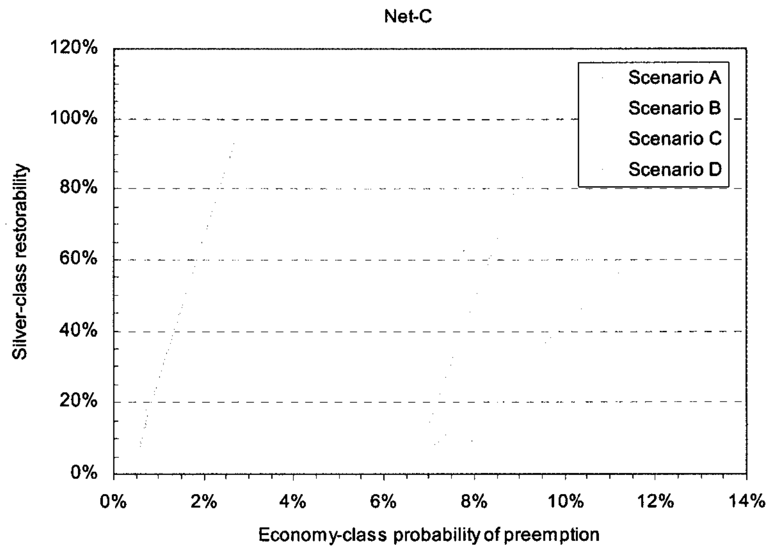

Figure 8-6 Trade-off between Silver-class restorability and Economy-class probability of preemption

8.6.3 Economic Interpretation

The previous sections have shown the attractiveness of operating under a multi QoP environment from a design efficiency point of view. However, to strictly optimize profitability we would have to also optimize the relative pricing of the four QoP service classes based on data for the price-elasticity and mutual displacement for service type. This is beyond the scope of the present paper but we can at least consider the intermediate question of at least assessing (as opposed to optimizing) net revenue and revenue gain (or loss) by going to a multi QoP environment from a fixed "gold only" environment. In general we could be losing relative to the all-gold case because, although we are saving total capacity, the non-gold services are also presumably earning less revenue. Clearly if one could retain all customers in a higher-paying gold class it may not make sense to split this voluntarily into some lower-priced categories. In reality, however, it can make sense for competitive reasons and because price stimulation of volumes may be expected so that the total demands rise.

In general, assume that the revenue generated by each demand served is determined by the product of a revenue coefficient (a^g, a^s, a^b and a^e, for gold, silver-bronze, and economy class respectively) times the number of hops for the working path. (This is easily generalized to a distance-dependent earnings model). With these assumptions the total earnings of the network in an multi-QoP scenario is:

$$\text{M-QoP earnings} = \sum_{k \in \{g, s, b, e\}} \left(a^k \cdot \sum_{i \in S} w_i^k \right) \tag{8.21}$$

whereas the "all-gold" earnings (if that amount of gold demand is achievable) would be:

$$\text{All-gold earnings} = a^g \cdot \sum_{i \in S} w_i^g \tag{8.22}$$

To assess the net benefit of the multi-QoP scenario relative to the all-gold service benchmark, we have to compare the change in earnings to the change in capacity cost of the complete network. Thus,

$$\text{Gain} = \underbrace{\sum_{i \in S} C_i \cdot \left((w_i + s_i)^{\text{gold}} - \left(\sum_{k \in \{g, s, b, e\}} w_i^k + s_i \right)^{\text{M-QoP}} \right)}_{\text{change in network cost}} - \underbrace{(\text{All-gold earnings} - \text{M-QoP earnings})}_{\text{change in earnings}} \tag{8.23}$$

where w_i and s_i values are obtained from the multi QoP solution for the given demand mix on the given topology. While not optimizing the revenue per-se, these considerations, combined with the multi QoP MJCP model, can be used in a planning tool to evaluate the potential earnings gain or loss of any particular planning scenario involving an assumed multi QoP demand mixture and volume, and assumed pricing coefficients.

8.7 Summary

This work developed an optimal capacity design model for span-restorable mesh networks that may have mixtures of up to four different service classes in terms of how each class is treated for restoration purposes. Gold class demands correspond to the basic assumption in most prior work of demands that must be fully restorable to any single failure scenario. Gold demands may use either built-in spare capacity or preempt economy class service capacity. To this we have added a generic silver class which receives best effort at restoration after restoration of gold services in the same failure scenario. Silver is considered in cases where it may, and may not, also preempt economy capacity. A bronze service class is neither restored nor preempted. Economy class services are not restored and may be interrupted to use their working capacity to protect gold demands and optionally silver demands, if needed.

We found that the most substantive interactions affecting capacity design occur between the gold and economy class services. Purely from a capacity design standpoint, both silver and bronze services only need their own working capacity to be provided for. Neither generates spare capacity or uses economy capacity for restoration. Silver, however, enjoys whatever best effort restorability arises as a side-effect of the allocations of spare capacity for gold service restoration. Two of the most interesting findings from studies of the multi-QoP design model on test cases is that over a wide range of demand mixes (here, scenarios A, B, and C) there is almost no requirement for dedicated spare capacity: the routing of gold and economy services can be worked out so that gold restorability is almost, or completely, provided through preemption of economy-class services. We also found that to give real meaning to a silver "best effort" class under these circumstances, one really has to also allow preemption of economy-class by silver class; otherwise, the restorability of silver demands is nearly zero. Nonetheless in no cases does the preemption frequency of economy services rise to exceedingly high levels.

We've seen also that scenarios A and B both present very interesting characteristics and can be considered as architectural "sweet spots" in the space of different demand mixes. Under scenario A, there is a lot of top-paying gold service, we require essentially no conventional spare capacity, we still provide a usefully high level of best effort silver restorability but we do so without extremely high probability of preemption for economy services. Under scenario B, the high restorability of the silver class makes it very attractive for clients needing good transport layer availability like ISPs but are not critical services like 911 call centres, the probability of preemption for the economy class is very low and will tend to make it look more like an unprotected class rather than extra traffic, which could be sufficient for services relying on higher layer survivability, and only a small proportion of gold services are provided for critical services that need absolute guarantee of all single failure restorability. As for scenario A, scenario B requires no conventional spare capacity.

The above points show that whether the service mix is more like scenario A or scenario B (the latter will in fact become more and more likely with the increasing predominance of data traffic relying on IP layer survivability) span-restorable mesh networks have a natural ability to take advantage of the mix of different services and to accommodate the needs of different clients.

We think this work contributes to providing a greater range of tools and knowledge that a transport network operator can use to enhance their business strategies and provide more customer options.

9. A Comparison of Ring and Mesh Service Path Availability

9.1 Introduction

Industry interest in mesh-based optical transport has increased greatly in recent years, from an era where ring-based transport dominated. Many network operators are now considering mesh as the way to go in the future [Muk00]. The reason for this change is a growing appreciation of the advantages of the mesh architecture. The most widely recognized advantage is that mesh-based transport networks are considerably more capacity efficient than their ring counterparts [Gro92]. This is attributable to the network-wide sharing of protection capacity over non-simultaneous failure scenarios and to the fact that service paths are not ring constrained – they can follow true shortest path routes over the graph. Rings not only require a minimum of 100 percent redundancy but also suffer from "stranded capacity" effects in which working capacity cannot be as highly utilized as when capacity is placed along pure mesh growth principles. Mesh-based survivable networking also supports simpler provisioning, especially of multiple service classes, and easier accommodation of growth where it was not foreseen [DoW94][GeR00b][Man01][GrL99]. Mesh networks also scale more easily in the sense that growing a mesh network only requires the addition of capacity wherever capacity has been exhausted whereas rings require the addition of entire rings.

The work presented in the previous sections has also shown how gracefully and efficiently a span-restorable mesh can support multiple classes of protection and provide extremely high restorability to priority service paths. In Chapter 4 it was shown that mesh-restorable networks exhibit very high average restorability against dual failures even when strictly designed only for single-failure restorability. This is attributable to the ability of mesh networks to use spare capacity in a very general way and, if also adaptive, to find restoration paths to the greatest extent possible under any circumstances. In contrast, rings (and 1+1 APS) provide only one pre-defined protection option and can never guarantee full dual-failure restorability to any paths because there are always dual failures that are guaranteed to bring down both the service path and its backup path.

This brings us to an interesting remaining question that industry colleagues have recently posed: "Which provides higher service-path availability – a ring or a mesh-based network?" So far the question seems to have been debated only qualitatively. One view is that rings are more redundant, spatially localized, and self-contained, so they should automatically be higher in availability. The opposite view is that although mesh is less redundant, it is also more general in its rerouting characteristics and can employ a failure-adaptive backup (even if more time is taken to do so), if initial pre-plans are overwhelmed. So in this viewpoint, mesh may do as well or even better. Another a priori view is that because rings switch in "50 ms," and mesh may take longer, rings should exhibit less outage time. This, however, is an easily corrected misunderstanding about what dominates the unavailability of either scheme. As explained in Section 3.6.1, while the duration of a restoration "hit" may be of concern in its own right, it has virtually no influence on service availability, which depends almost wholly on the likelihood of an unrestorable dual-failure, not the speed of response to restorable single failures. With this background, Dr. Grover and I realized in late 2002 that the timing was right for a comprehensive quantitative comparison between ring and mesh availability and that perhaps for the first time all the tools needed to do this job were also available together to do the first such careful and exact comparison. This project was to become the last major mandate of the thesis work and resulted in publication in [ClG03]. Here that work is documented in greater detail and extent in terms of results and methods.

9.1.1 Prior Work on Ring vs. Mesh Availability

A few past studies can be found addressing availability analysis of ring or mesh [AVD00a][BBB00][Gro99a][Sch00][CaN97] but these studies are either devoted to one or the other architecture (not a comparison between them) or, where they do offer comparison, there are approximating simplifications that one or the other side of the comparison would find objectionable. Typical simplifying approximations are for instance that all paths in a ring are down if two failures hit the same ring (which is not the case), or, to model shared mesh restoration as approximately equivalent to 1+1 APS, or even to "smooth out" all details of capacity, topology, and mechanism to employ average-case Markov state transition modelling. The aim here is to deliberately avoid any such simplifying approximations to provide an extremely precise and fair comparison using exact implementations of both survivability mechanisms. The study is thus essentially experimental in nature and is scrupulously detailed as to the exact re-routing mechanisms involved and how they interact with the topology, the spare capacity present, and the network state due to any prior failures.

9.2 Qualitative Comparison of the Two Mechanisms

The simulation of each survivability mechanism relies on the precise knowledge of the routing of all working paths and the reaction of all structures to single, dual, and higher order failures. In the ring case we know precisely which segment of each end-to-end path traverses each ring. When failures occur on a particular span, each affected service path is restored following a BLSR-protection mechanism provided that another failure is not already employing the ring's protection. Note that this does permit certain paths through a ring to survive dual failures, depending on how the two failures hit the ring. More formally the precise logic governing the outage of a path transiting a ring is as follows:

Let the set of W spans and (W+1) nodes in the normal route of a signal path segment X within the ring be called the forward path of X, For{X}. The set of other nodes and spans in the ring is then defined as Rev{X}, i.e. Rev{X} = {{R} - For{X}} where set {R} is the complete ring. Rev{X} contains S-W spans, where S is the number of spans on the ring. For outage of path X in the ring, it is then necessary and sufficient that one failed element belong to For{X} and the other to Rev{X}. This follows the detailed considerations given in [Gro99a]. Note, however, that unlike [Gro99a], this work is not concerned with special measures such as matched node interconnection between rings to avoid node failure. In the language of [Gro99a] this study considers only single-fed paths through ring-based networks, subject only to span failures.

The mesh restoration mechanism is implemented as follows. Upon a single failure, restoration is effected by rerouting affected paths around the failed span in a k-shortest paths like manner through the exact spare capacity present. This means that first all feasible paths on the shortest route are taken. Next, all paths feasible through the second shortest route around the failure span are taken, not using any spare channels of the first set of paths, and so on. This functional routing model goes on until either all failed paths are restored or no more restoration paths are feasible. This is known to be extremely close a single commodity maximum-flow solution for span restoration [DGM94]. When a second failure happens, the reaction to that failure is functionally identical but takes into account any spare capacity usage from the first failure. Moreover, if a subsequent failure directly hits restoration paths of a first failure, the affected working paths and restoration paths are viewed as a single revised number of failed working channels for which restoration is sought in the reaction to the second event. Overall this behaviour corresponds to the partly adaptive span-restoration behaviour described in Chapter 4. All channels bearing working service are initially considered equally important for assignment of the available restoration paths. Later, tests with priority assignment give available restoration paths to priority service paths first. In all cases the environment of spare capacity in use and spare capacity available for restoration of any failure is always updated to reflect current usage arising from any previous still outstanding span failure(s). When any failure state terminates, all related working paths are returned to normal routing, the associated restoration paths are collapsed, and their spare capacity is returned to an available spare state. Figure 9-1 illustrates some of the key concepts of the two survivability mechanisms under comparison.

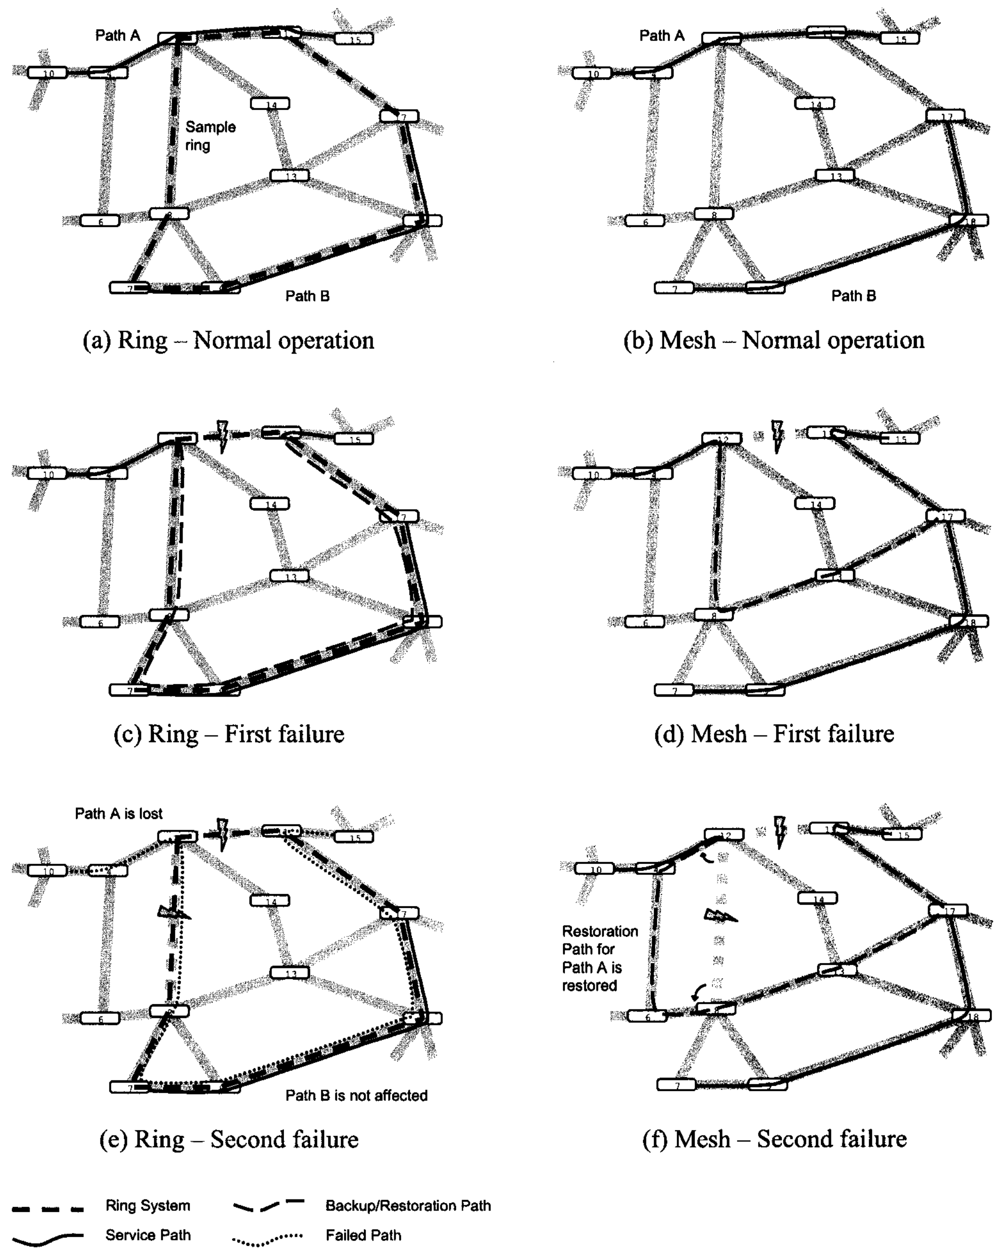

Figure 9-1 Comparison of the response of ring and mesh architectures to a dual failure

Parts (a) and (b) of Figure 9-1 show both architectures under normal operations with two identical service paths for illustration of effects. In the ring architecture, path A traverses one span of the ring before travelling on and path B is completely contained between nodes of the ring. Part (c) and (d) of Figure 9-1 show the reaction of both mechanisms to the same first failure. In the ring architecture, For$\{A\}$ of path A is the single-hop segment between the upper two ring nodes. When failure hits that span, it is restored using path Rev{A} through the backup capacity around the rest of the ring. In the mesh, path A is restored by assigning it to one of the set of dynamically found "ksp" restoration paths described above. For clarity in Figure 9-1 we only show one path on the affected span but there is no implication intended that the mechanism restore all affected service paths over a single restoration route (this is sometimes a misconception about span restoration.) At the same time as the one rerouted path shown takes its route, many other paths (not shown but affected by the same failure) would also be restored over other routes of the ksp path-set.

Parts (e) and (f) of Figure 9-1 show the reaction of both mechanisms to a second failure. This second failure hits a span in Rev{A} and thus brings down service path A. That path is now experiencing hard outage, although path B remains unaffected in the same ring and is still routed over its normal working channels. The second failure also hits the restoration path of path A for the mesh architecture. In this case the failed restoration path is unified with the rest of the failed working paths on the second failed span and another maximum feasible ksp-type restoration path-set is developed in response to the second failure in the presence of the first span failure and the already used spare capacity removals from the network. Now in the mesh architecture, path A may or may not be restored following the second failure – it all depends on what is feasible within the remaining spare capacity following the first failure. If the number of restoration paths for the second failure is less than fully required, it also depends on whether path A has either luck or priority in the assignment by the failure end nodes of such restoration paths as are feasible for the second failure. As drawn we illustrate the possibility that the restoration path of path A is itself re-restored on the second failed span, showing how service paths can survive dual failures in the mesh. The example shows that the two architectures have very different responses to failures and that in neither architecture does a dual or triple etc., failure necessarily mean outage for service paths. Not shown are the cases of dual failures that hit different rings or are spatially separated enough in the mesh that they do not even interact. Neither of these cases is outage-causing to any paths.

9.3 Simulation Method and Test Networks

9.3.1 General Methodology

For the simulation, importance is placed on making a true "apples-to-apples" comparison between the two architectures. To achieve this the following conditions were set:

1. Ring and mesh are compared on identical facilities graphs serving identical end-to-end demands: Two test topologies were used. The first one, displayed on Figure 9-2, is the net32 topology with 32 nodes and 45 spans. It was taken from [Mor01] and was tested

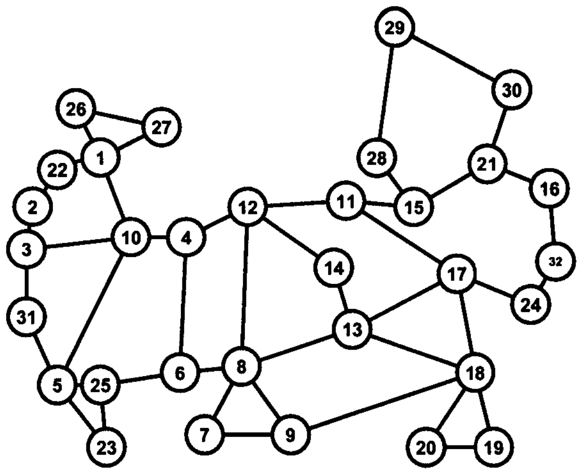

Figure 9-2 Net-32 test network
(from [MoG01])

with two different demand matrices. The first demand matrix and the corresponding ring design were also taken from [Mor01]. This first demand matrix corresponds to a hub-type of demand where all demands originate or terminate at one of three different hub nodes. The corresponding test case is denoted net32-A. The second demand matrix was created for the present study by generating random demands between nodes in proportion to the product of their nodal degrees (gravity-based approach) and serving them within the same ring set as net32-A following ring-constrained shortest path routing until the first demand request was blocked. This resulted in a demand pattern that achieves about 78 percent loading of working capacity in the corresponding ring set. The corresponding test case is denoted net32-B. The same end to end demand pattern is used for the corresponding mesh network design on the topology of net32. The second topology is the arbitrary manually designed topology Net-B (25 nodes, 50 spans) presented in Section 1.4.1 and described in details in Appendix A. As explained in Section 1.4.2, the corresponding demand matrix was generated using a purely deterministic gravity-based model.

2. Efficient fully restorable designs are used for both architectures: The ring networks are based on efficient sets of BLSR ring placements and routing of demands through the rings. The ring design used for net32 (identical for net32-A and net32-B) is taken from the work of [Mor01] and is based on the fcrip design method. The ring design for Net-B was produced for this study using the Tabu Search method [Mor01][MoG01] starting from a solution obtained with the fcrip design method. The mesh designs were produced using the Modular Joint Capacity Placement (MJCP) linear programming formulation from

[DoG00] (and presented in Appendix B) using ten eligible working routes for each demand pair and ten eligible restoration routes for each span-failure scenario and the same capacity modularity and end-to-end demands used by the ring designs (module size of 48 capacity units in all designs).

3. Exact survivability mechanisms are emulated: As explained in the previous section, exact emulation of both survivability mechanism is performed, taking into account the capacity usage of all capacity units and producing a detailed report of the status of all service paths after each failure or repair event.
4. Both architectures experience identical span-failure sequences: A sequence of random physical layer failures and repairs is generated and both architectures are tested with it. As illustrated by Figure 9-3, the sequence is obtained by first generating a unique timeline of failures and subsequent repairs for each span as an independent entity and then combining them all into one composite history of failure and repair for the whole network. Note that under these circumstances, nothing prevents the network from experiencing dual, triple or higher order failures. In addition, this type of simulation is effectively at steady state the instant it starts because the up/down state probability of each span already reflects the long-term average when its individual failure/repair history was generated. Thus, there is no need for a period of transient simulation to reach steady state in this approach to simulating the failure environment. For simplicity, it was assumed that each span in the network had the same MTBF of 1 year with times to failure being negative-exponentially distributed. Times-to-repair were also negative exponentially distributed with a mean of 12 hours. The simulation tool behind the work can however be easily employed to model

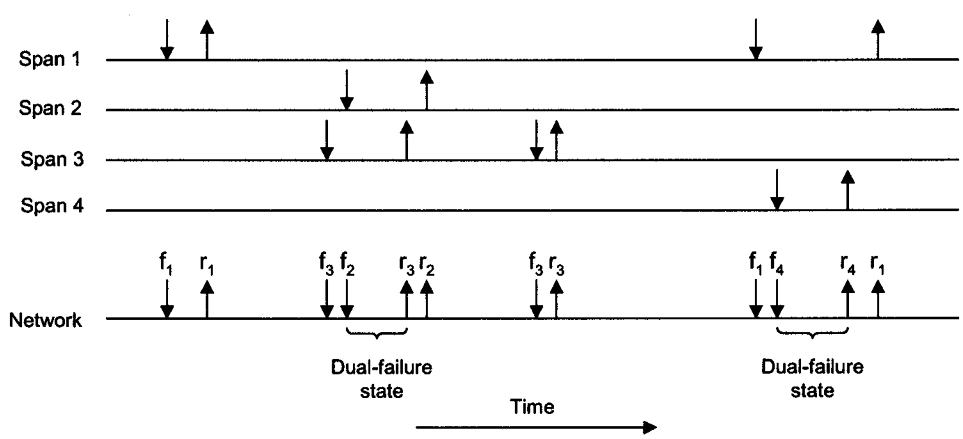

Figure 9-3 Generation of network-wide failures and repairs sequence

length-dependent failure rates on each span and/or any particular distribution of repair times.

Table 9-1 summarizes some details of the test networks used. The working capacity totals correspond to the sum of working channels along all service paths in the network. The total capacity values correspond to the total number of capacity modules in each design multiplied by the module capacity. Confirmation appears right away of the great difference in capacity efficiency between ring and mesh designs. The ring designs require much more capacity than the corresponding mesh designs. It is particularly noticeable with the Net-B test case, in which the topology is favourable to the mesh architecture. With the less connected net32 topology, mesh is also less efficient (especially with the hubbed demand) but still requires significantly less capacity than ring. Also, path lengths are slightly higher for mesh in the net32 topology. This sparse topology tends to force paths to go more out of their way to enhance spare capacity sharing under MJCP design, which is a joint optimization model. In contrast in the highly connected Net-B topology paths are longer for the ring design. It will be seen later that path length has a considerable importance in terms of availability. With such total capacity differences, one could wonder whether mesh can provide comparable availability to that of the ring architecture. Section 9.4 answers the question.

Table 9-1: Details of Test Network Designs

Network Design	Demand	Working capacity	Total capacity	Average path length	Network redundancy
net32-A Ring	Hubbed	1824	8448	5.15	363%
ne32-A Mesh	Hubbed	1943	4656	5.48	140%
net32-B Ring	Random / gravity based	3304	8448	2.01	156%
net32-B Mesh	Random / gravity based	3500	7152	2.14	104%
Net-B Ring	Gravity based	4966	17520	3.07	253%
Net-B Mesh	Gravity based	4574	6960	2.83	52%

9.3.2 Availability Simulator

The same main program is used for availability simulation of mesh and rings. The general architecture of the availability simulator is shown on Figure 9-4. The architecture shown can be used for the availability simulation of any network type. The particulars of a given network reside in the Failure/Repair Events Generation Module, the Restoration/Protection Analysis module and the Reversion Module. For this study, the same Failure/Repair Events Generation Module is used for both

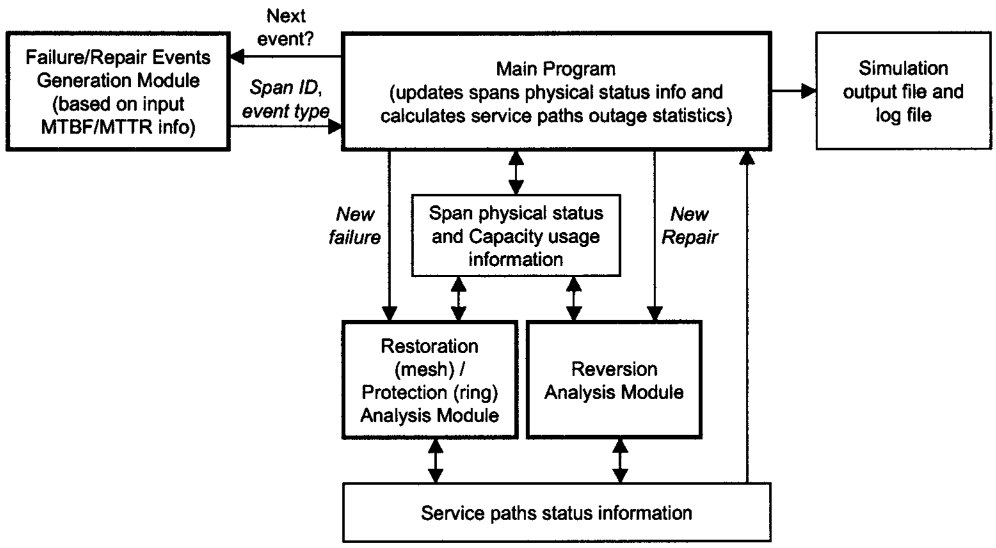

Figure 9-4 Architecture of availability simulator

mesh and rings, therefore the only differences between the mesh availability simulator and the ring availability simulator are in the details of the restoration or protection and reversion routines. The restoration analysis module is in fact identical to the one used in the restorability analysis program used to produce results of Chapter 4. The availability simulator therefore offers the choice between the three restoration models described in that chapter, but only the partly adaptive model (Model 2) is used in this section. Also, for the mesh case, the simulator offers the possibility to perform a capacity-based availability analysis (when only capacity allocation is known but not the details of demand routing) or a end-to-end demand-based availability analysis. In the case of the capacity-based availability analysis, results are presented in the form of equivalent link unavailabilities as defined in Section 3.6.2. In the case of the end-to-end demand-oriented analysis, results are presented in the form of end-to-end service path availability. The demand-oriented analysis is the approach taken in this chapter.

9.4 Experimental Results

As the results are essentially experimental in nature, let us first address issues of statistical confidence. All results are based on a thousand simulations of one year of the network's life. As a check on confidence intervals, we inspected the variance of the point estimation of one of the lowest overall unavailability predictions resulting from the study. This is for the Net-B test network in the mesh case for the average path unavailability independent of path length. The result was that for the ensemble of 1000 one-year simulations, an average unavailability of 3.297×10^{-5} was found with a standard deviation of 8.91×10^{-7} which is only 2.7 percent of the average unavailability. For higher unavailability values the accuracy of estimation is only increased because such cases are based on more outage-contributing events. In addition, the total number of dual or higher-order failures (over the 1000 one-year trials) was 2619 for the net32 topology and 3180 for the 25n50s1 topology. Each of the individual span failure/repair sequences was produced using the negative exponential random number generator (RNG) in [Pre93]. With suitable scaling, the generator is called alternately to generate the time-to-next-failure and the time-to-repair for each span until the total simulation period has been reached. The number of calls that this produces to the RNG is much lower than its period (in the order of 108) so we perceive no risk of correlation between sequences for the different spans.

For each design, the following measures of the resulting availability were computed:

i. Average path unavailability: This is fraction of time that a given service path is experiencing outage. Because it is reasonable to expect that individual path availability depends on path length, results are presented as a function of path length in terms of number of hops in the path. Each data point presented is the average over all paths of a given length and over all 1000 test periods.

ii. Average number of outages per year (network total) and per path: These values indicate the probability of each service path of experiencing an outage during a given one-year period and characterize the total number of outage-causing events the operator could expect over the network as a whole per year. This is also presented versus path length.

iii. Statistical frequency of total path outage times per year: This data summarizes the proportions of service paths expected to undergo no outage over a one-year period and the fraction of all paths that experience various higher levels of annual outage. This also allows visibility of the worst case scenarios experienced by any path at all.

In Figure 9-5 we can see that in all cases the unavailability of paths in ring networks is significantly higher than the paths of same lengths in the mesh networks. The difference in unavailability also increases as path length increases. The difference is about a factor of two at the longest paths. For the net32 topology, this difference is reduced slightly if we take into account the fact that service paths are on average slightly longer in the mesh design. For the Net-B topology however, which already shows the biggest advantage for mesh, the difference between the two architectures is even higher if we consider the fact that service paths are on average longer in the ring architecture. The curves on Figure 9-6 are almost identical in shape but they show the expected number of outages per year and indicate an expected number of outages per path per year reaching a maximum of about 0.17 for mesh and 0.26 for ring. If, for the sake of the argument, we neglect the probability of two outages happening to the same service path in a year, this is close to saying that there is about a 17 percent chance of the longest mesh paths, and about a 26 percent chance in the ring designs, to experience an outage during any given year.

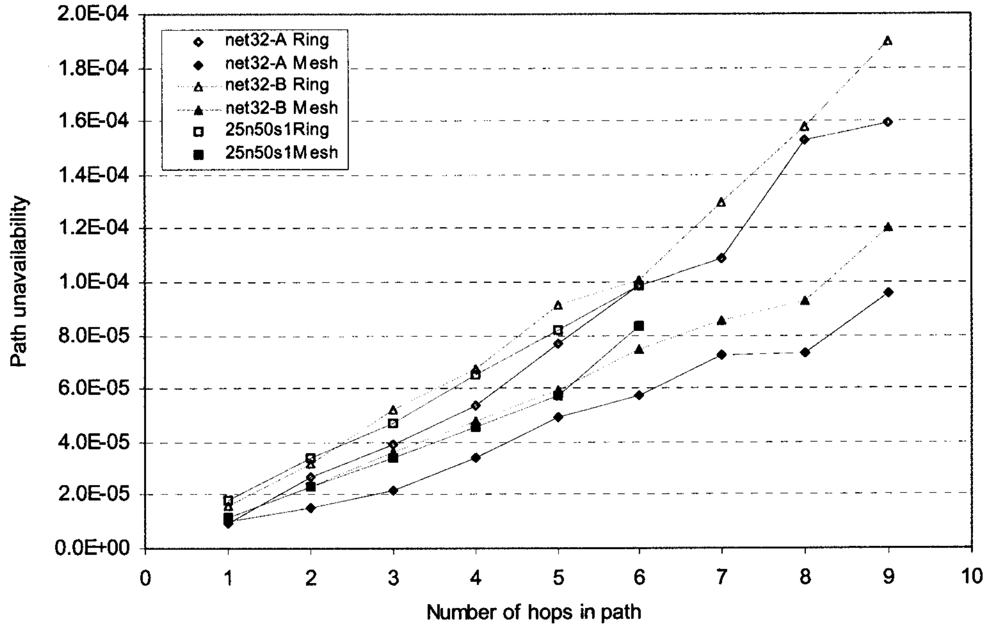

Figure 9-5 Results of path unavailability for the three test networks

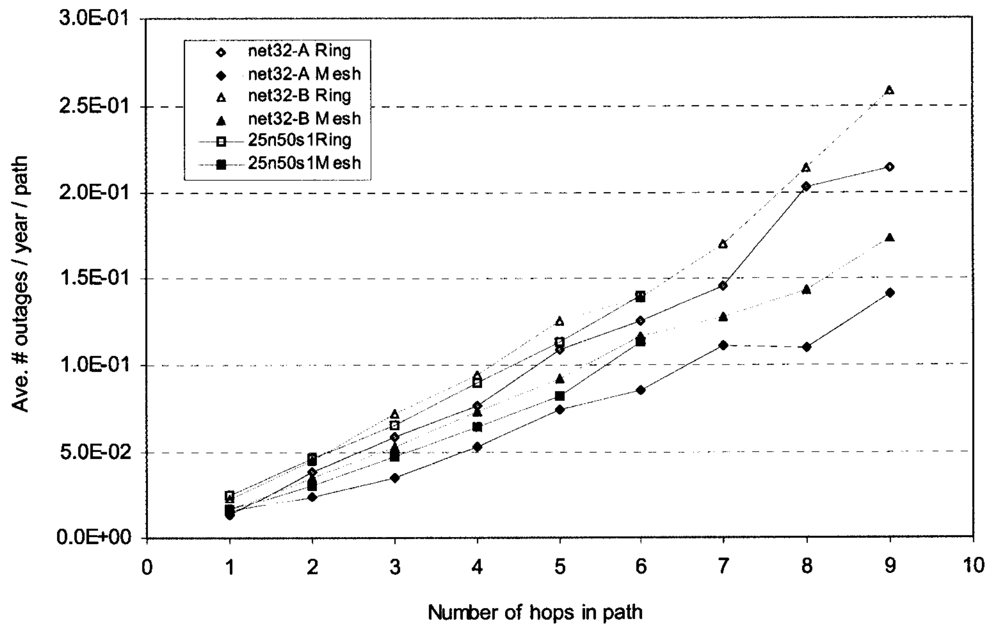

Figure 9-6 Results of expected number of outages per year for the three test networks

Figure 9-7 shows the distribution of actual outage times experienced by paths in both architectures for the 25n50s1 test case with annotation of the expected fraction of paths that experience no outage in any one year. Results show a slightly higher proportion of paths expected to experience no outage in a year with mesh (95.4 percent for mesh vs. 93.7 percent for ring.) The worst-case experience of any one path is also 72 hours of total outage in one year compared to 48 in total for the worst-case mesh path. Similar results are observed for the other test cases.

Overall, these results show an advantage for mesh on all availability measures. Note, however, that even with mesh restoration, long paths experience relatively high probabilities of an outage during any given year. For critical services it would clearly be desirable to have a probability of outage lower than 17 percent per year. This is the motivation for introducing prioritization in mesh restoration in the next section. This is an option that only has meaning in the context of the mesh restorable architecture. With ring protection, when failure occurs on one span of a given ring, all affected paths are either restorable or not because protection switching inherently occurs at the optical line rate, not at the individual channel level where priorities can be implemented in cross-connect based mesh restoration.

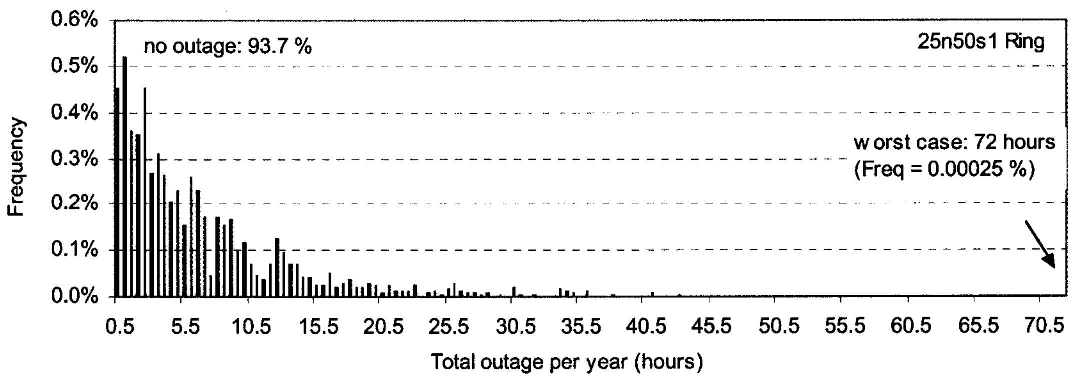

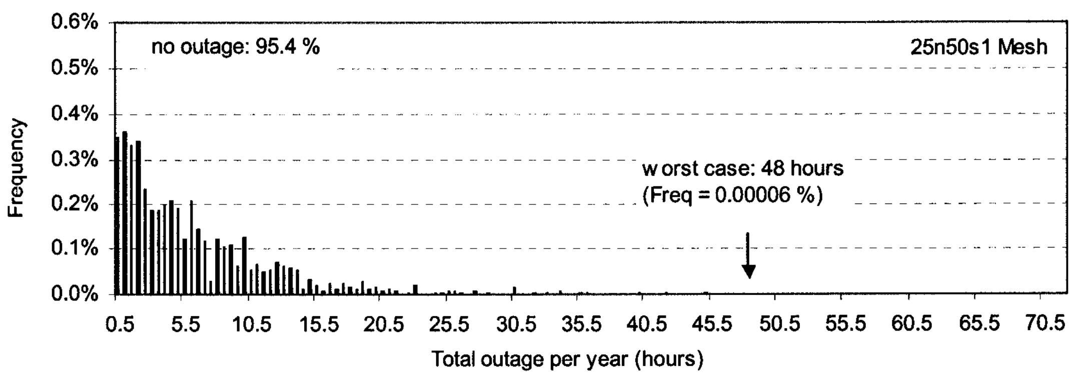

Figure 9-7 Distribution of outage times for test case 25n50s1

9.5 Effect of Mesh Restoration with Priorities

Unlike in the studies presented in Chapter 6 and Chapter 8, where different classes of service were treated differently from the point of view of the capacity design itself, the present distinction between the high-priority and the low-priority paths is purely from an operational point of view. Both classes are still guaranteed full restorability to any single failure by design. The difference is only that the high priority service paths are always considered first for restoration in any circumstances where full restoration may not be possible for all affected paths. High priority service paths are still not guaranteed dual-failure restorability by design (as in Chapter 6) but they have a higher chance of being restored. To test the effect of applying priority distinctions, we tag a certain fraction of the service paths on each demand pair as "high priority." We tested three service mixes: "10/90," "30/70," "50/50" where the first number is the proportion of high-priority service paths.

Figures 9-8 to 9-13 show the effects of priority for the three test cases. In all cases we see a clear improvement in availability for the high-priority service paths. As expected, the improvement is greater for the priority class when fewer paths have high-priority status. A more surprising outcome is that giving high priority status to some service paths does not affect the availability of the remaining paths very much. In fact, when the proportion of high-priority services is only 10 percent, the priority group benefits very noticeably, but the availability of the low priority paths is almost unchanged. In other words, the Mesh LP 90 curves in Figures 9-8 to 9-13 are almost identical to the basic un-prioritized results for mesh in Figures 9-5 and 9-6. This is an interesting result from a revenue point of view because it indicates the prospect of selling a fraction of services at a higher price without having to reduce quality to other services. One explanation for this effect is that the small high priority group benefits mainly from the coherence of applying available restoration path resources in a consistent, as opposed to random way. In other words to provide the best service for a few, opportunities to avoid outage must be consistently applied to them, but this makes little difference to the remainder who receive their opportunities (in the event of any shortage) on an essentially random basis of allocating the hardship. In this thinking the randomness has more to do with the experience of the non-priority group than does the preference given to the smaller priority group when it is possible and necessary to do so for them.

Another interesting result is that the effects of priorities appear to depend on the type of topology. The two test cases using the net32 topology show an improvement for the priority services that is less pronounced than in the Net-B test-case. The difference in availability improvement between the demand mix scenarios is also not as high as in the case of Net-B. The most dramatic effect is seen in the highly connected Net-B network. In particular, the 10/90 demand mix shows an

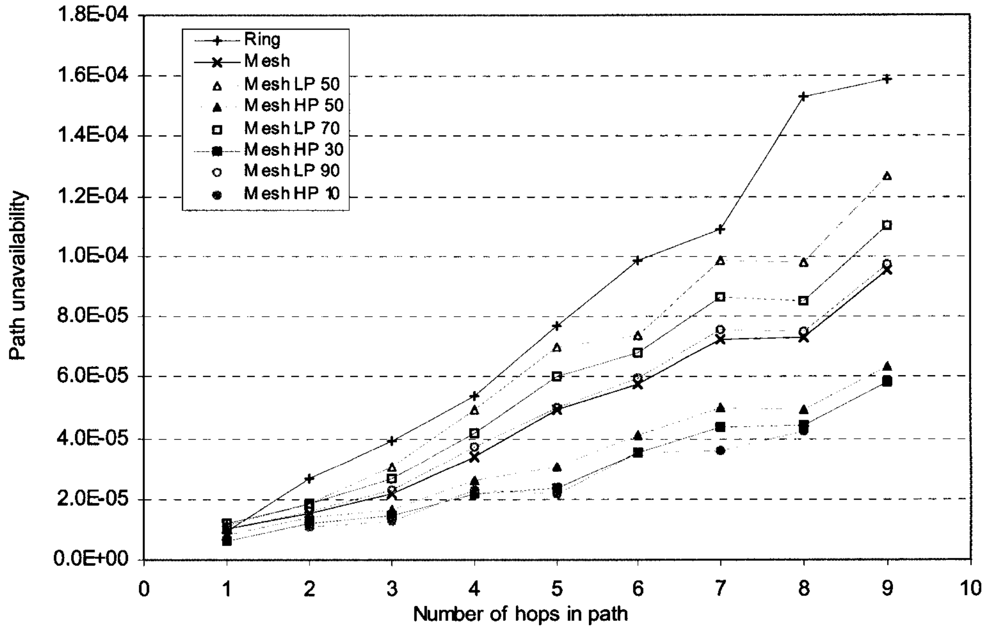

Figure 9-8 Path unavailability for test network net32-A with priorities

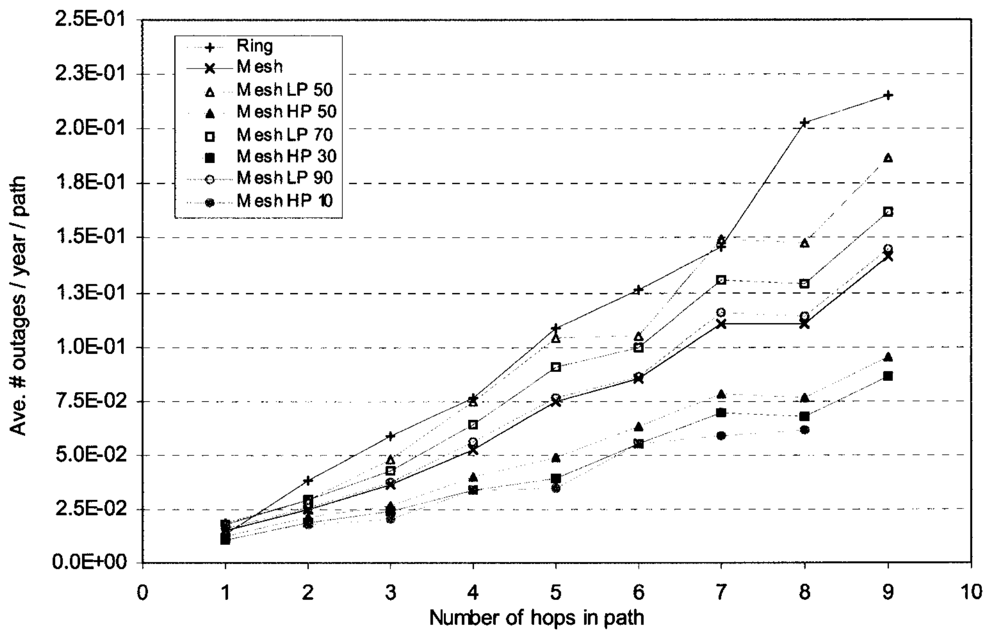

Figure 9-9 Expected number of outages per year for test network net32-A with priorities

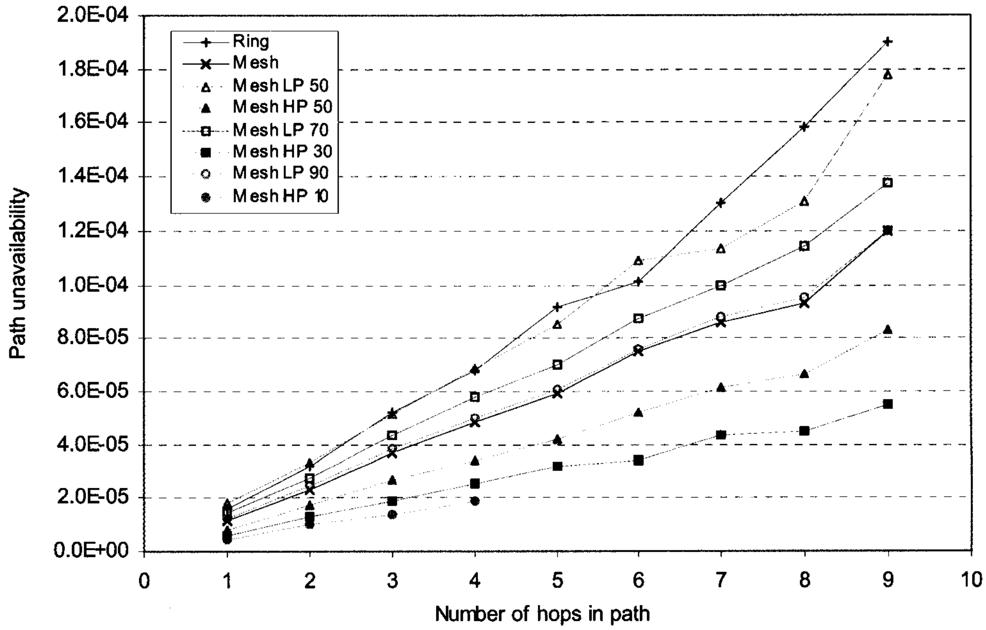

Figure 9-10 Path unavailability for test network net32-B with priorities

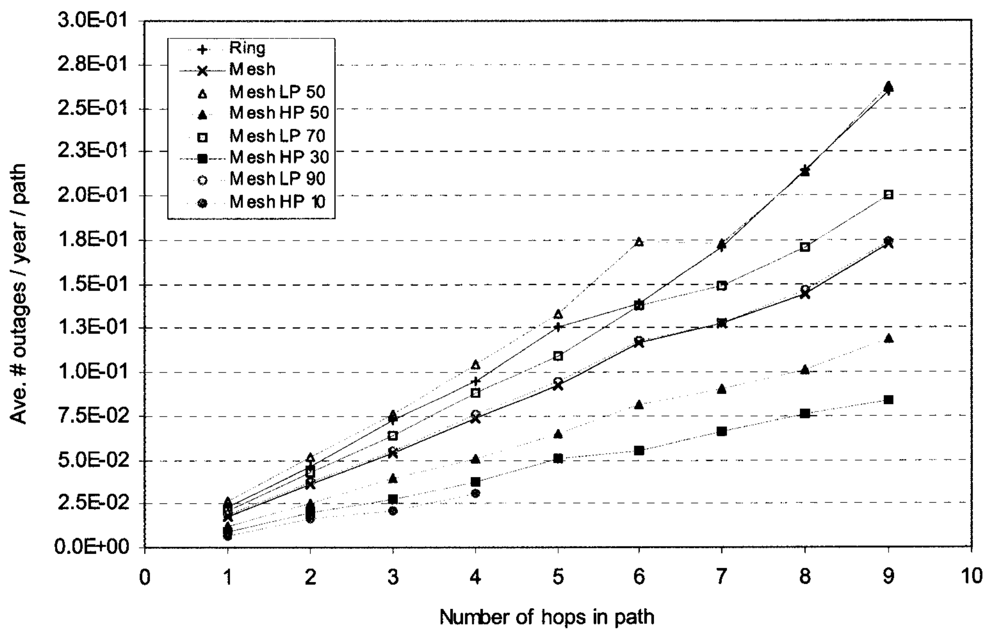

Figure 9-11 Expected number of outages per year for test network net32-B with priorities

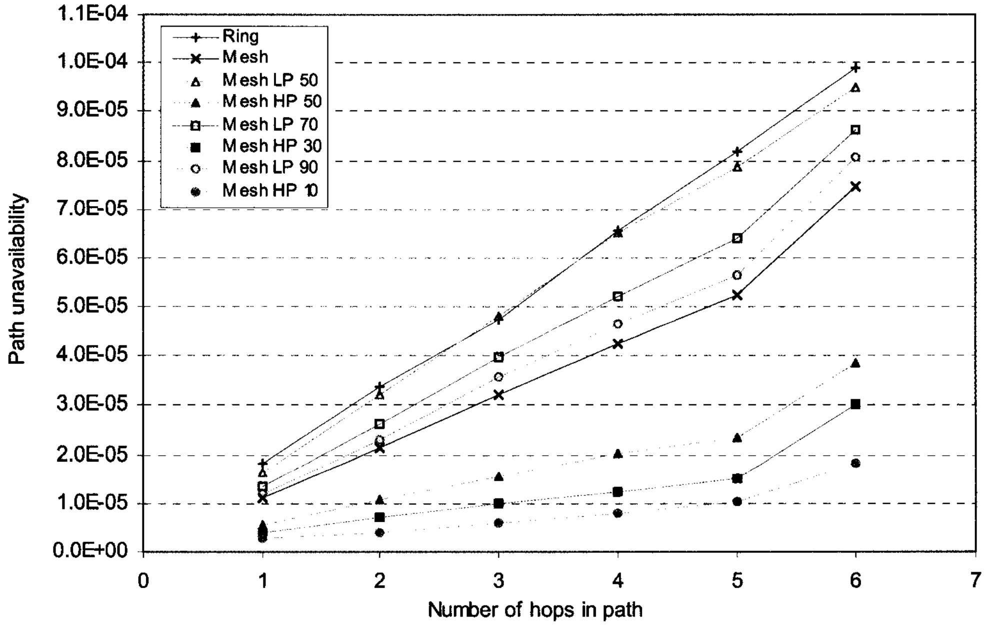

Figure 9-12 Path unavailability for test network Net-B with priorities

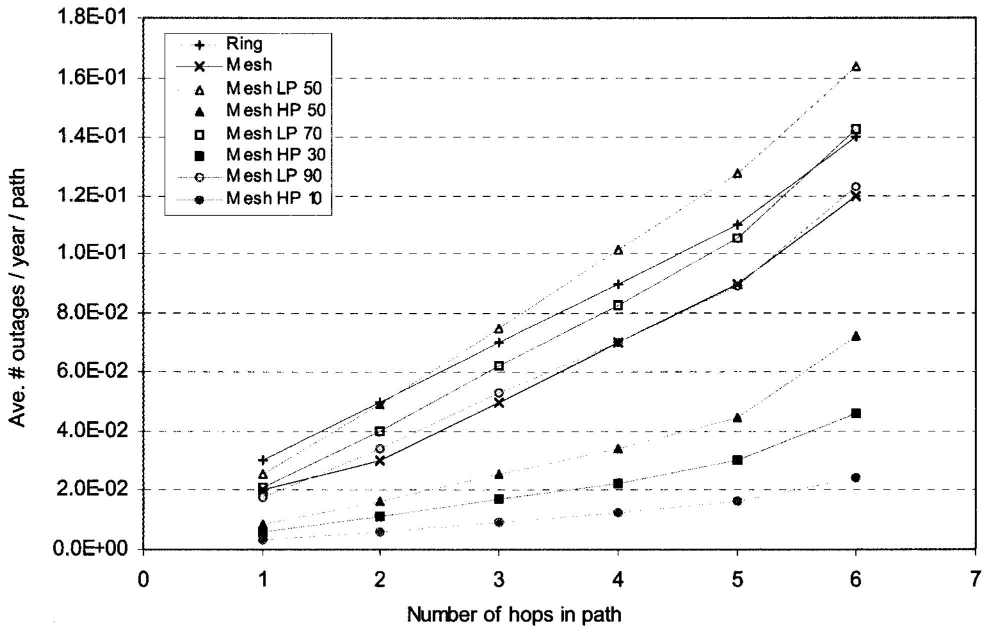

Figure 9-13 Expected number of outages per year for test network Net-B with priorities

almost six-fold reduction in unavailability in the best case and the availability of the 90 percent of low-priority paths is virtually unchanged. The probability of experiencing an outage in a given year is in that case reduced to about 2 percent in the worst case (compared to about 12 percent without priorities).

Finally Figure 9-14 shows the distribution of outage times for the high-priority class in the Net-B test case with the 10/90 demand mix. The improvement is very clear: the expected proportion of paths in that class experiencing no outage in a given year is now 99.2 percent (compared to 95.4 percent without priorities) and the worse case scenario of total annual outage is reduced from 48 to 40 hours (happening with a probability of 0.0017 percent).

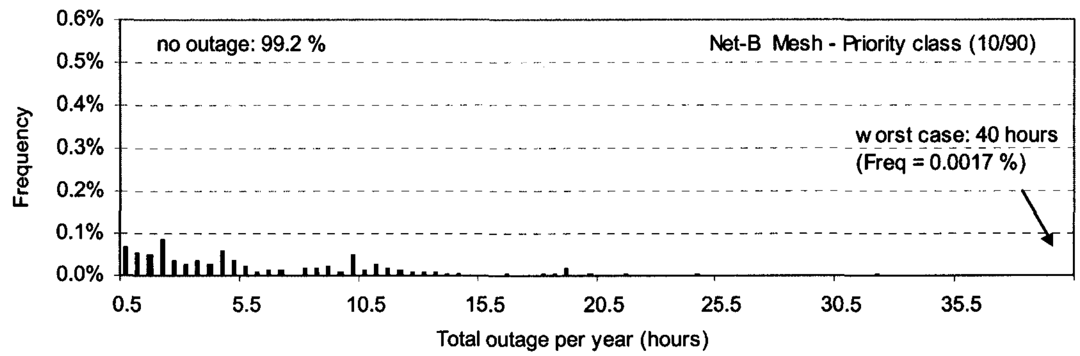

Figure 9-14 Result of total outage distribution for 10% priority class in mesh

9.6 Conclusion and Summary

The work presented in this section has sought to address an open question about ring and mesh network availability through a carefully controlled study in which the network and failure environments are absolutely identical for each alternative and where we mechanize the exact reaction of each architecture in the face of failures. Results are sufficient to put aside the argument that because of its lower capacity requirement, mesh-restorable networks cannot be as high-availability as rings. In fact, what we actually see is that mesh flexibility wins out over ring redundancy largely because of the "locked up" nature of the ring investment in protection capacity. Results show that service paths enjoy significantly higher availability in the mesh architecture despite the much lower capacity requirements of the mesh designs. The advantage is even more pronounced in the more connected test topology. The study also investigated the availability of mesh service paths with two levels of priority. This can only be implemented with the mesh architecture and has the advantage of enhancing the availability of high-priority paths even further. If the proportion of high-priority paths is about 10 percent, the availability of low-priority paths is almost unaffected while the availability of high-priority paths is greatly increased.

It must be acknowledged that node failures were omitted from this study. This is partly justified by the high internal redundancy already provided in optical cross-connects and ADMs and the security measures taken to enforce the very high reliability of network nodes. Taking node failures into account would therefore not have significantly affected the results and the main finding that mesh does at least as well as rings in terms of availability.

We believe that after the previous chapters, this study provides yet another confirmation of the great potential of mesh-restorable networks and clearly shows that there is no simple link between network redundancy and availability. High redundancy is not sufficient to provide high availability. Based on this work an increasingly clear view is emerging that what is ideal for very high availability is to be a priority service path in a highly connected mesh-based restorable network that employs an adaptive response to any circumstances where a pre-planned single failure response is overwhelmed. When these conditions are met then more redundancy could reasonably be indeed expected to translate into higher availability. This is a subject of interest in further studies in this area.

10. Thesis Conclusion

10.1 Introduction

This chapter reviews the work developed in this thesis, summarizes the main contributions, and suggests future work on the topic of availability analysis in mesh transport networks.

10.2 Review of Thesis

The main goal of the thesis was to develop theoretical and practical methods for determining the availability of service paths in mesh-restorable transport networks. These methods were used to increase our understanding of what factors influence the availability of service and to assess the ability of mesh networks to provide reliable services. Another goal of the thesis was to investigate ways of providing different restorability options other than just the single-failure restorability option traditionally offered in survivable transport networks.

Chapter 1 gave some reference material describing mathematical notations and terminology used in the thesis, and presented some concepts related to optimization theory and, in particular, Mathematical Programming, which were then extensively used in the following chapters.

Chapter 2 presented background information on transport networks, including the concepts of client/transport relationship and of network demand, which are fundamental in this research work. We also presented different classes of transport networks and the different technologies. The chapter devoted an important part to survivability schemes and, in particular, to mesh-based schemes. The chapter also introduced several design methods that serve as the basis for new design methods presented in following chapters. The end of the chapter was dedicated to what we see as the future challenges of transport networking and we explained how some of these challenges relate to the present work.

In Chapter 3, we developed a theoretical approach to the problem of determining the availability of service in transport networks. The chapter started with a review of important mathematical definitions related to availability followed by a presentation of general existing methods for the analysis of system availability. These methods were then considered in the specific context of telecommunication systems. The chapter then surveyed the literature on the topic of network availability analysis and presented a new approach to the availability analysis problem. That approach showed the major role played by dual failures in determining the availability of service paths. Analyzing the effects of dual-failures was then identified as the primary objective of service availability in mesh-restorable networks.

Chapter 4 presented an approach to the problem of determining the restorability to dual-failures based on computational re-routing trials. It was shown how difficult it is to develop closed-form equations for the restorability of mesh networks using adaptive restoration mechanisms. Three algorithms were presented to model various progressively more adaptive restoration mechanisms for a span-restorable network. Experimental results investigating the dual-failure restorability properties of span-restorable mesh networks design for full restorability to single failures were then presented. The influence of several factors on the restorability results was then investigated and it was found that there is no simple direct relation between the proportion of spare capacity and the availability of service paths. In fact, the factor that seems to influence the dual-failure restorability the most seemed to be the size of the network (represented by the number of spans or nodes), with larger networks having higher dual-failure restorability.

Chapter 5 presented a study showing the influence of maintenance actions on the restorability of the network to single span-failures. This chapter gave additional support for the relevance of studying the effects of dual-failures. It was explained that under some maintenance models, dual-failure situations can result from the occurrence of a physical failure during a maintenance action. The chapter defined three models of maintenance actions with respect to the use of spare capacity that they require. Two of these models were compared in terms of their effects on the restorability of the network during the maintenance action. The chapter also provided ideas on how to manage the effects of maintenance actions by identifying maintenance actions that can be conducted simultaneously and maintenance actions that should be conducted in series in order not to expose services to higher risks of outage.

Chapter 6 presented extensions of the common mesh-restorable network capacity design formulation that enhance the dual-failure restorability of the designs. These new design formulations included a formulation for capacity minimization under the constraint of complete dual-failure restorability, a formulation for restorability maximization under a given total capacity cost budget, and a formulation for minimum-capacity design supporting multiple-restorability service class definitions. In the third formulation the restorability options range from no restorability guarantee to the guarantee of full restorability to any dual-failure. Results obtained with the first formulation showed how expensive it is to design for full dual-failure restorability. Indeed, to go from a full single-failure restorable design to a full dual-failure restorable design it is common to see a tripling of the required spare capacity, making this option unrealistic. These results were confirmed by the results obtained with the second formulation allowing us to observe the trade-off between dual-failure restorability and capacity requirements. These results showed that the dual-failure restorability could be improved significantly by adding a little bit of extra spare capacity but that quickly small improvements in dual-failure restorability require large additions of spare capacity. The third formulation presented in Chapter 6 showed that it is possible to economically support an added service class in the upward quality direction and tailor the investment in capacity to provide ultra-high availability on a selective basis. In most cases up to 20 or 30 percent of demands could be offered dual-failure restorability guarantee for almost no capacity requirement increase, or even no increase at all. These results showed the potential of mesh-restorable networks to fulfill the requirements of service differentiation, presented as one of the important challenges of future transport networks in Section 2.6.

Chapter 7 presented a theoretical study of service availability in *p*-cycle networks. Closed-form equations were developed for the availability of paths and these equations showed that there can be a significant difference in the availability of paths depending on whether they lie on the *p*-cycles or just straddle them. Based on the equations developed, two factors were investigated for their influence on the availability of service in networks protected by *p*-cycles. The chapter provided suggestions for possible ways to control the availability of paths in these networks. The chapter also provided a restorability analysis comparison between optimal *p*-cycle designs and corresponding optimal span-restorable designs. Results showed that the dual-failure restorability in *p*-cycle networks suffers greatly from the high exposure of backup paths to secondary failures because of the large size of *p*-cycles in optimal capacity design. Networks protected by *p*-cycles therefore offered significantly lower dual-failure restorability than networks protected by a span-restoration mechanism. These results give some motivations for developing new *p*-cycle capacity design methods that allow to control the trade-off between capacity efficiency and the exposure of backup paths to secondary failures.

In Chapter 8, we introduced a capacity design formulation for span-restorable networks with multiple classes of protection. A study of the capacity requirements of these networks was presented and showed how, in many cases, it is possible to design networks with some restorable demands and without any extra protection capacity. We then investigated how a multi quality of protection (QoP) environment affects the availability of the different service classes and considered the case of four different multi-QoP restoration models. Several multi-QoP restoration models were considered and their influence of the reliability of the different QoP classes was studied. An economic interpretation of the results was also provided, which showed how a multi QoP environment allows the prospect for network operators to earn more revenue from their existing capacity investments.

In Chapter 9, we compared the availability of service in ring and mesh-restorable networks. One of the main motivations of the study was to verify whether the higher capacity redundancy of rings compared to mesh results in higher availability. The comparison was based on detailed simulations of the network's response to random sequences of failures and repairs. Both survivable architectures were tested under the exact same conditions. The results showed that mesh-restorable networks, despite their lower capacity requirements, provide higher availability than their ring counterparts. The key aspect that explained why availability is higher in mesh was the ability of these networks to provide more diversified restoration options, which proves very efficient when it comes to restoring dual failures. The study also showed the potential of the mesh architecture to provide very high availability to a small fraction of selected high-priority service paths when prioritization in the restoration is introduced, while keeping the availability of lower-priority service paths almost unchanged.

10.3 Contribution of the Thesis to the Research Field

10.3.1 Key Contributions

The most important contributions of this thesis work were as follows:

- Development of generic tools for the analysis of service path restorability and availability in survivable transport networks.
- Theoretical treatment of the availability of service paths in mesh-restorable networks and understanding of the major importance of dual-failure effects.
- Development of three progressively more adaptive computational models for the analysis of dual span-failure restorability in span-restorable mesh networks.
- Development of a method for determining the restorability risk fields associated with network maintenance actions.
- Optimal capacity design formulation for capacity cost minimization under the full dual span-failure restorability constraint in span-restorable mesh networks (DFMC).
- Optimal capacity design formulation for maximization of the dual-failure restorability under given capacity cost budget in span-restorable mesh networks (DFMR).
- Optimal capacity design formulation minimum capacity under multiple restorability requirements for span-restorable mesh networks (MRCP).
- Closed-form models for the availability of service paths in networks protected by the *p*-cycle protection mechanism.
- Four models for the analysis of restoration in a span-restorable mesh networks with multiple QoP classes.

10.3.2 Publications

The following journal, magazine, or conference papers, based on this thesis work have been published or have been accepted for publication:

1. [ClG00] M. Clouqueur and W. D. Grover, "Computational and design studies on the unavailability of mesh-restorable networks," in *Proceedings of the Second International Workshop on the Design of Reliable Communication Networks, DRCN 2000*, Munich, Germany, April 2000, pp. 181-186.

2. [GCB01] W. D. Grover, M. Clouqueur, T. Bach, "Quantifying and managing the influence of maintenance actions on the survivability of mesh-restorable networks," in *Proceedings of the 17th National Fiber Optic Engineers Conference, NFOEC 2001*, Baltimore, MD, USA, July 2001, vol. 3, pp. 1514-1525.

3. [ClG02a] M. Clouqueur, W. D. Grover, "Availability analysis of span-restorable mesh networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, May 2002, pp. 810-821.

4. [ClG02b] M. Clouqueur, W. D. Grover, "Mesh-restorable networks with complete dual failure restorability and with selectively enhanced dual-failure restorability properties," in *Proceedings of SPIE OptiComm 2002*, Boston, MA, USA, July 2002, pp. 1-12. (recipient of Best Student Paper Award)

5. [GrC02] W. D. Grover, M. Clouqueur, "Span-restorable mesh networks with multiple quality of protection (QoP) service-classes," in *Proceedings of the International Conference on Optical Communications and Networks, ICOCN 2002*, Singapore, November 2002, pp. 321-323.

6. [ClG03] M. Clouqueur, W. D. Grover, "Quantitative comparison of end-to-end availability of service paths in ring and mesh-restorable networks," in *Proceedings of the 19th National Fiber Optic Engineers Conference, NFOEC 2003*, Orlando, FL, USA, September 2003.

7. [DCG03] J. Doucette, M. Clouqueur, W. D. Grover, "On the service availability and respective capacity requirements of shared backup path-protected mesh networks," *SPIE Optical Networks Magazine, Special Issue on Engineering the Next Generation Optical Internet*, vol. 4, no. 6, November 2003.

8. [ClG04] M. Clouqueur, W. D. Grover, "Mesh-restorable networks with enhanced dual-failure restorability properties," *SPIE Optical Networks Magazine*, accepted for publication. To appear in 2004.

9. [GrC04] W. D. Grover, M. Clouqueur, "Span-restorable mesh networks with multiple quality of protection (QoP) service classes," *SPIE Optical Networks Magazine*, accepted for publication. To appear in 2004.

10.3.3 TRLabs MeshAnalyzer

In addition to the work presented here, a related network availability analysis and simulation software was developed. The tool, TRLabs MeshAnalyzer, combines the different restorability and availability analysis programs developed for this thesis [TRL03]. More information about MeshAnalyzer can be obtained by contacting:

TRLabs
7^{th} Floor – 9107 116 Street NW
Edmonton, Alberta T6G 2V4
Canada

Tel: +1 780 441 3800
Fax: +1 780 441 3600

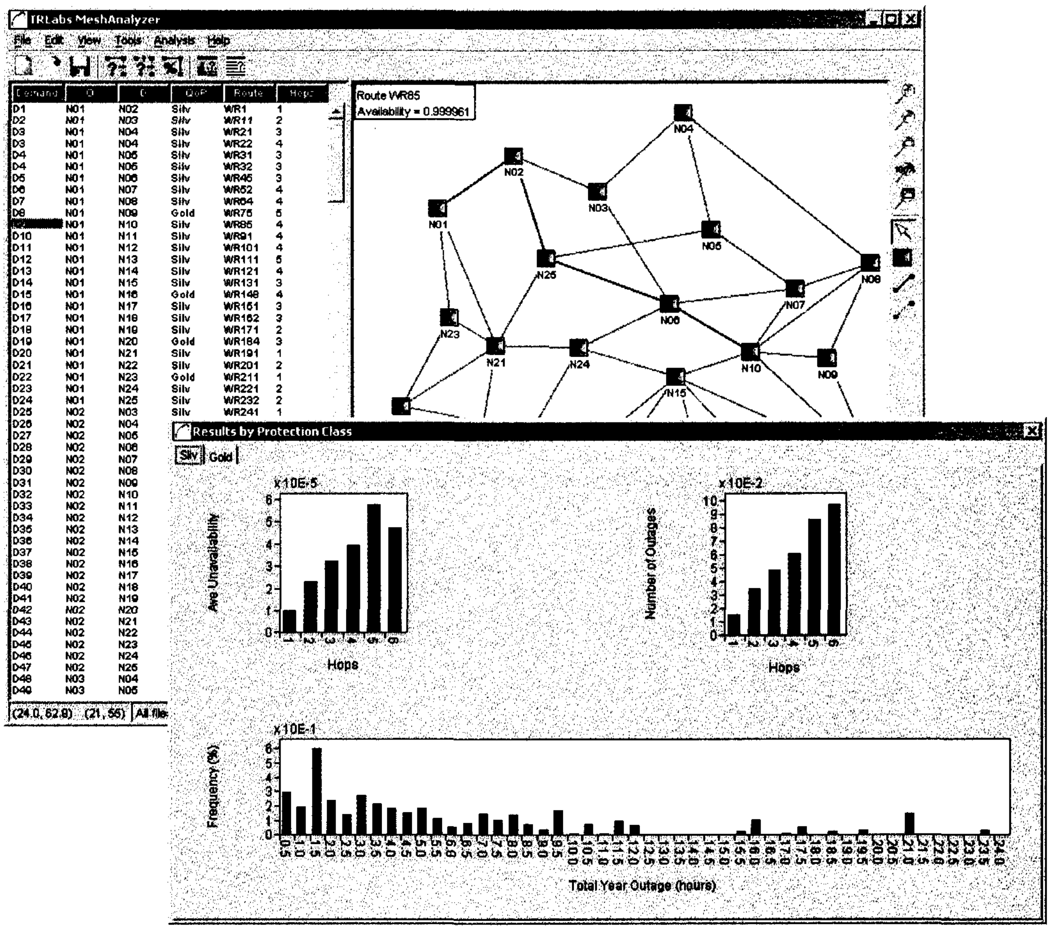

Figure 10-1 Screenshot of TRLabs MeshAnalyzer

10.4 Ideas for Future Research on the Topic

10.4.1 Extension to Path Restoration

This thesis has studied the restorability and availability of paths in span-restorable mesh networks. A possible extension of this work would be to consider path restoration as the survivability mechanism. Path restoration is known for having lower capacity requirements than span restoration [DoG01], which should not benefit the restorability, however it provides more restoration options than span restoration so this might compensate for the lower available capacity. Also, the fact that paths can be restored end-to-end could be an important advantage for rerouting failed paths around regions of the network affected by two failures interacting spatially—those that do cause lower restorability levels with span restoration. A possible aspect to investigate would be the dual-failure restorability of span restoration and of path restoration under the same capacity design. Also, path restoration could be used as a secondary measure to revert to in case a path with high availability requirements is not restorable using span-restoration.

10.4.2 Availability Analysis in Metro Networks

Discussion with our industry colleagues has recently revealed the importance of considering some types of node failures in the context of metropolitan access networks. Future work on the availability of service paths in metro networks could consider new failure types such as failures of line cards, network elements' chassis or shelf, failure of power supplies, shelf controller cards, forwarding engine, switching matrix, cooling fans, etc.

10.4.3 Restorability Analysis with Allowed Over-Subscription

With the increasing importance of IP traffic being transported over optical networks, the requirement that restored connections need full bandwidth replacement could be relaxed in the future and some level of over-subscription of the restoration bandwidth could be tolerable. Future work on the topic of restorability and availability analysis could take that aspect into account in the context of a network using MPLS as a routing mechanism and investigate the effects of allowing a certain level of bandwidth over-subscription during restoration. Of interest would be the trade-off between the maximum tolerable over-subscription factor (and corresponding characteristics of packet loss probability, average packet delay, etc.) and the average single or dual-failure restorability of connections.

10.4.4 p-Cycles in a Multi-QoP Environment

p-Cycle protection is a very promising survivability scheme that offers many interesting features including fast restoration speed, low capacity requirements, operational simplicity, and easy deployment starting from a ring-based architecture. As seen in Chapter 7, p-cycle protection potentially provides lower availability than do span-restorable mesh networks. Moreover, to date p-cycle protection has only been studied in the context of a single protection class. Future work on p-cycle protection could investigate the possibility of using p-cycle protection in a multi-QoP environment, including Gold, Silver, Bronze, and Economy services as defined in Chapter 8, as well as higher availability service as studied in Chapter 6.

Bibliography

[AKQ00] V. Anand, T. Katarki and C. Qiao, "Profitable connection assignment in all optical WDM networks," in *CD-ROM Proceedings of Optical Networks Workshop*, Richardson, TX, USA, January/February 2000.

[ALL96] B. Allen, SONET, in Wired Communications Systems, CRC Press, 1996, pp. 1524-1537.

[AVD00a] P. Arijs, W. Van Parys, P. Demeester, "Cost and availability comparison of WDM mesh and ring network architectures," in *Proceedings of DRCN 2000*, Munich, Germany, April 2000, pp. 33-38.

[AVD00b] P. Arijs, B. Van Caenegem, P. Demeester, P. Lagasse, W. Van Parys, "Design of ring and mesh based WDM transport networks," *SPIE Optical Networks Magazine*, vol. 1, no. 3, July 2000, pp. 25-40.

[ASA02] G. M. Assi, A. A. Shami, M. A. Ali, Z. Zhang, X. Liu, "Impact of wavelength converters on the performance of optical networks," *SPIE Optical Networks Magazine*, vol. 3, no. 2, March/April 2002, pp. 22-30.

[BaM96] D. Banerjee, B. Mukherjee, "A practical approach for routing and wavelength assignment in large wavelength-routed optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, June 1996, pp. 903-908.

[BaP01] R. Barr, R. A. Patterson, "Grooming telecommunications networks," *SPIE Optical Networks Magazine*, vol. 2, no. 3, May/June 2001, pp. 20-23.

[Bat99] P. Batchelor et al., Ultra high capacity optical transmission networks : Final report of action COST 239 (Book), ISBN 953-184-013-X, Faculty of Electrical Engineering and Computing, Zagreb, Croatia, 1999.

[BBB00] M. Bolduc, S. Baddipudi, J. Babbitt, A. Wright, "Optical backbone architecture comparison," in *Proceedings of NFOEC 2000*, Denver, CO, USA, August 2000.

[BBI00] S. Binetti, A. Bragheri, E. Iannone, F. Bentivoglio, "Mesh and multi-ring optical networks for long-haul applications," *IEEE/OSA Journal of Lightwave Technology*, vol. 18, no. 2, December 2000, pp. 1677-1684.

[BDL01] A. Banerjee, J. Drake, J. Lang, B. Turner, D. Awduche, L. Berger, K. Kompella, Y. Rehkter, "Generalized multiprotocol label switching: an overview of signalling enhancements and recovery techniques," *IEEE Communications Magazine*, vol. 39, no. 7, July 2001, pp. 144-151.

[Bel93] Bellcore, Digital Cross-Connect Systems in Transport Network Survivability, SR-NWT-002514, Issue 1, 1993.

[BFP98] M. Bettin, G. Ferraris, and G. Pignari, "Comparison of protection and restoration schemes for SDH networks," in *Proceedings of DRCN 1998*, Brugge, Belgium, May 1998.

[Bha98] R. Bhandari, Survivable Networks: Algorithms for Diverse Routing. Norwell, MA: Kluwer Academic Press, Nov. 1998.

[BiA92] R. Billinton, R. N. Allan, Reliability Evaluation of Engineering Systems, 2nd Edition, Plenum Press, New York and London, 1992.

[BPZ92] M. Barezzani, S. Pupolin, and M. Zorzi, "A new definition of transmission network availability with applications," *European Transactions on Telecommunications*, vol. 3, no. 4, July/August 1992, pp. 349-357.

[Bru92] R. A. Brualdi, Introductory Combinatorics, 2nd Edition, Englewood Cliffs, New Jersey: Prentice-Hall, 1992.

[CaN97] H. C. Cankaya, V. S. S. Nair, "Reliability and availability evaluation of self-healing SONET mesh networks," in *Proceedings of IEEE GLOBECOM 1997*, Phoenix, AZ, USA, November 1997, pp. 252-256.

[CFZ96] I. Chlamtac, A. Faragó, T. Zhang, "Lightpath (wavelength) routing in large WDM networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 5, June 1996, pp. 909-913.

[CGK92] I. Chlamtac, A. Ganz, G. Karmi, "Lightpath communications: An approach to high bandwidth optical WANs," *IEEE Transactions on Communications*, vol. 40, no. 7, July 1992, pp. 1171-1182.

[CGV01] D. Colle, A. Groebbens, P. Van Heuven, S. D. Maesschalk, M. Pickavet, P. Demeester, "Porting MPLS-recovery techniques to the MPλS paradigm," *SPIE Optical Networks Magazine*, vol. 2, no. 4, July/August 2001, pp. 29-47.

[CLG00] O. Crochat, J.-Y. Le Boudec, O. Gerstel, "Protection interoperability for WDM optical networks," *IEEE/ACM Transactions on Networking*, vol. 8, no. 3, June 2000, pp. 384-395.

[ClG00] M. Clouqueur, W. D. Grover, "Computational and design studies on the unavailability of mesh-restorable networks," in *Proceedings of DRCN 2000*, Munich, Germany, April 2000, pp. 181-186.

[ClG02a] M. Clouqueur, W. D. Grover, "Availability analysis or span-restorable mesh networks," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, May 2002, pp. 810-821.

[ClG02b] M. Clouqueur, W. D. Grover, "Mesh-restorable networks with complete dual failure restorability and with selectively enhanced dual-failure restorability properties," in *Proceedings of SPIE OptiComm 2002*, Boston, MA, USA, July 2002, pp. 1-12.

[ClG03] M. Clouqueur, W. D. Grover, "Quantitative comparison of end-to-end availability of service paths in ring and mesh-restorable networks," in *Proceedings of the 19th National Fiber Optic Engineers Conference, NFOEC 2003*, Orlando, FL, USA, September 2003.

[ClG04] M. Clouqueur, W. D. Grover, "Mesh-restorable networks with enhanced dual-failure restorability properties," *SPIE Optical Networks Magazine*, accepted for publication, to appear in 2004.

[CMY02] D. Cavendish, K. Murakami, S.-H. Yun, O. Matsuda, M. Nishihara, "New transport services for next-generation SONET/SDH systems," *IEEE Communications Magazine*, vol. 40, no. 5, May 2002, pp. 80-87.

[Col87] C. J. Colbourn, The Combinatorics of Network Reliability. Oxford, U.K.: Oxford University Press, 1987.

[Col91] C. J. Colbourn, "Network reliability: numbers or insight? A disussion paper," *Annals of Operations Research (Kluwer Academic)*, vol. 33, 1991, pp. 87-93.

[CrL98] O. Crochat, J.-Y. Le Boudec, "Design protection for WDM optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 15, no. 7, September 1998, pp. 1158-1165.

[CSC01] H. Choi , S. Subramaniam, H.-A. Choi, "Loopback methods for double-Link failure recovery in optical networks," in *Proceedings of 39th Annual Allerton Conference on Communications, Control, and Computing*, Monticello, IL, USA, October 2001.

[CSC02] H. Choi, S. Subramaniam, H.-A. Choi, "On double-link failure recovery in WDM optical networks," in *Proceedings of IEEE INFOCOM 2002*, New York, NY, USA, June 2002.

[DaS93] M. Daneshmand, C. Savolaine, "Measuring outages in telecommunications switched networks," *IEEE Communications Magazine*, vol. 31, no. 6, June 1993, pp. 34-38.

[DCG03] M. Dahlin, B. B. V. Chandra, L. Gao, A. Nayate, "End-to-end WAN service availability," *IEEE/ACM Transactions on Networking*, vol. 11, no. 2, April 2003, pp. 300-313.

[DDH99] B. T. Doshi, S. Dravida, P. Harshavardhana, O. Hauser, Y. Wang, "Optical network design and restoration," *Bell Labs Technical Journal*, vol. 4, no. 1, January-March 1999, pp. 58-83.

[Dem99] P. Demeester, et al. "Resilience in multilayer networks," *IEEE Communications Magazine*, vol. 37, no. 8, August 1999, pp. 70-76.

[DeM02] S. De Maesschalck et al., "Intelligent optical networking for multilayer survivability," *IEEE Communications Magazine*, vol. 40, no. 1, January 2002, pp. 42-49.

[DeM03] S. De Maesschalck et al., "Pan-european optical transport networks: an availability-based comparison," *Photonic Network Communications (Kluwer Academic)*, vol. 5, no. 3, May 2003, pp.203-225.

[DGB01] J. Doucette, W. D. Grover, T. Bach, "Bi-criteria studies of mesh network restoration path-length versus capacity tradeoffs," in *Proceedings of OSA/IEEE OFC 2001*, Anaheim, CA, USA, March 2001, pp. TuG2-1-TuG2-3.

[DGM94] D. A. Dunn, W. D. Grover, M. H. MacGregor, "A comparison of k-shortest paths and maximum flow methods for network facility restoration," *IEEE Journal on Selected Areas in Communications*, vol. 12, no. 1, January 1994, pp. 88-99.

[DoG00] J. Doucette, W. D. Grover, "Influence of modularity and economy-of-scale effects on design of mesh-restorable DWDM networks," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, October 2000, pp. 1912-1923.

[DoG01] J. Doucette, W. D. Grover, "Comparison of mesh protection and restoration schemes and the dependency on graph connectivity," in *Proceedings of DRCN 2001*, Budapest, Hungary, October 2001, pp. 121-128.

[DoW94] R. D. Doverspike and B. Wilson, "Comparison of capacity efficiency of DCS network restoration routing techniques," *Journal of Network and Systems Management (Plenum Publishing)*, vol. 2, no. 2, June 1994, pp. 95-123.

[EAP03] T. S. El-Bawab, A. Agrawal, F. Poppe, L. B. Sofman, D. Papadimitriou, B. Rousseau, "The evolution to optical-switching-based core networks," *SPIE Optical Networks Magazine*, vol. 4, no. 2, March/April 2003, pp. 7-19.

[ElS02] T. El-Bawab, J.-D. Shin, "Optical packet switching in core networks: Between vision and reality," *IEEE Communications Magazine*, vol. 40, no. 9, September 2002, pp. 60-65.

[FaL99] N. S. Fard, T.-H. Lee, "Cutset enumeration of network systems with link and node failures," *Reliability Engineering and System Safety (Elsevier)*, vol. 65, no. 2, August 1999, pp. 141-146.

[Fla90] T. Flanagan, "Fiber network survivability," *IEEE Communications Magazine*, vol. 28, no. 6, June 1990, pp. 46-53.

[Fre96] R. L. Freeman, *Telecommunication System Engineering*, 3rd Ed. New York: Wiley, 1996, pp. 444-447.

[FTU02] A. Fumagalli, M. Tacca, F. Unghváry, A. Faragó, "Shared path protection with differentiated reliability," in *Proceedings of IEEE ICC 2002*, New York, NY, USA, April/May 2002, pp. 2157-2161.

[FuV00] A. Fumagalli, L. Valcarenghi, "IP restoration vs. WDM protection: Is there an optimal choice?," *IEEE Network*, vol. 14, no. 6, November/December 2000, pp.34-41.

[GBV91] W. D. Grover, T. D. Bilodeau, and B. D. Venables, "Near optimal spare capacity planning in a mesh restorable network," in *Proceedings of IEEE GLOBECOM 1991*, Phoenix, AZ, USA, December 1991, pp. 2007-2012.

[GCB01] W. D. Grover, M. Clouqueur, T. Bach, "Quantifying and managing the influence of maintenance actions on the survivability of mesh-restorable networks," in *Proceedings of NFOEC 2001*, Baltimore, MD, USA, July 2001, vol. 3, pp. 1514-1525.

[GeS02] O. Gerstel, G. Sasaki, "Quality of protection (QoP): a quantitative unifying paradigm to protection service grades," *SPIE Optical Networks Magazine*, vol. 3, no. 3, May/June 2002, pp. 40-49.

[GeR00a] O. Gerstel and R. Ramaswami, "Optical layer survivability: a services perspective," *IEEE Communications Magazine*, vol. 38, no. 3, March 2000, pp. 104-113.

[GeR00b] O. Gerstel and R. Ramaswami, "Optical layer survivability—An implementation perspective," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, October 2000, pp.1885-1899.

[Glo93] K. C. Glossbrenner, "Availability and reliability of switched services," *IEEE Communications Magazine*, vol. 31, no. 6, June 1993, pp. 28-32.

[GNS00] N. Golmie, T. D. Ndousse, D. H. Su, "A differentiated optical service for WDM networks," *IEEE Communications Magazine*, vol. 38, no.2, February 2000, pp. 68-73.

[GrC02] W. D. Grover, M. Clouqueur, "Span-restorable mesh networks with multiple quality of protection (QoP) service-classes," in *Proceedings of the International Conference on Optical Communications and Networks, ICOCN 2002*, Singapore, November 2002, pp. 321-323

[GrC04] W. D. Grover, M. Clouqueur, "Span-restorable mesh networks with multiple quality of protection (QoP) service classes," *SPIE Optical Networks Magazine*, accepted for publication, to appear in 2004.

[GrD02] W. D. Grover, J. Doucette, "Advances in Optical Network Design with *p*-Cycles: Joint Optimization and Pre-Selection of Candidate *p*-Cycles," in *Proceedings of IEEE/LEOS Summer Topicals 2002*, paper WA2, Mont Tremblant, Quebec, Canada, July 2002, pp. 49-50.

[Gre01] P. Green, "Progress in optical networking," *IEEE Communications Magazine*, vol. 39, no. 1, January 2001, pp. 54-61.

[GrL99] W. D. Grover, D. Li, "The forcer concept and applications to express route planning in mesh survivable networks," *Journal of Network and Systems Management (Plenum Publishing)*, vol. 7, no. 2, June 1999, pp. 199-223.

[Gro87] W. D. Grover, "The selfhealing network: A fast distributed restoration technique for networks using digital crossconnect machines," in *Proceedings of IEEE GLOBECOM 1987*, Tokyo, Japan, November 1987, pp. 28.2.1-28.2.6.

[Gro89] W. D. Grover, Selfhealing Networks – A distributed algorithm for k-shortest link-disjoint paths in a multi-graph with applications in real time network restoration, Ph.D. thesis, University of Alberta, Edmonton, Alberta, Canada, Department of Electrical Engineering, 1989.

[Gro92] W. D. Grover, "Case studies of survivable ring, mesh, and mesh-arc hybrid networks," in *Proceedings of IEEE GLOBECOM 1992*, Orlando, FL, USA, December 1992, pp. 633-638.

[Gro94] W. D. Grover, "Distributed restoration of the transport network," in *Telecommunications Network Management into the 21st Century*, S. Aidarous and T. Plevyak, Eds. New York: IEEE Press, 1994, ch. 11, pp. 337-419

[Gro97] W. D. Grover, "Self-organizing broad-band transport networks," *Proceedings of the IEEE: Special Issue on Communications in the 21st Century*, vol. 85, no. 10, October 1997, pp. 1582-1611.

[Gro99a] W. D. Grover, "High availability path design in ring-based optical networks," *IEEE/ ACM Transactions on Networking*, vol. 7, no. 4, August 1999, pp. 558-574.

[Gro99b] W. D. Grover, "Resource management for fault tolerant paths in SONET ring networks," *Journal of Networks and Systems Management (Plenum Publishing)*, vol. 7, no. 4, December 1999, pp. 373-394.

[Gro02] W. D. Grover, et al., "New options and insights for survivable transport networks," *IEEE Communications Magazine*, vol. 40, no. 1, January 2002, pp. 34-41.

[Gro03] W. D. Grover, *Mesh-based Survivable Networks: Options and Strategies for Optical, MPLS, SONET and ATM Networking*, Prentice Hall, 2003.

[GrS98] W. D. Grover, D. Stamatelakis, "Cycle-oriented distributed preconfiguration: ring-like speed with mesh-like capacity for self-planning network restoration," in *Proceedings of IEEE ICC 1998*, Atlanta, GA, USA, June 1998, pp. 537-543.

[GrS00] W. D. Grover, D. Stamatelakis, "Bridging the ring-mesh dichotomy with *p*-Cycles," in *Proceedings of DRCN 2000*, Munich, Germany, April 2000, pp. 92-104.

[GVM91] W. D. Grover, B. D. Venables, M. H. MacGregor, and J. H. Sandham, "Development and performance verification of a distributed asynchronous protocol for real-time network restoration," *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 1, January 1991, pp. 112-125.

[HKC95] D. D. Harms, M. Kraetzl, C. J. Colbourn, and J. S. Devitt, Network Reliability. Boca Raton, FL: CRC Press, 1995.

[HeB94] M. Herzberg, S. Bye, "An optimal spare-capacity assignment model for survivable networks with hop limits," in *Proceedings of IEEE GLOBECOM 1994*, San Francisco, CA, USA, November 1994, pp. 1601-1607.

[IMG98] R. R. Iraschko, M. H. MacGregor, W. D. Grover, "Optimal capacity placement for path restoration in STM or ATM mesh-survivable networks," *IEEE/ACM Transactions on Networking*, vol. 6, no. 3, June 1998, pp. 325-335.

[IrG00] R. R. Iraschko, W. D. Grover, "A highly efficient path-restoration protocol for management of optical network transport integrity," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 5, May 2000, pp. 779-794.

[Ise00] A. Iselt, "Terminal pair availability calculation algorithms for communication networks," in *Proceedings of DRCN 2000*, Munich, Germany, April 2000, pp. 233-247.

[ITU93] ITU-T, Network Node Interfaces for the Synchronous Digital Hierarchy (SDH), ITU-T Recommendation, G.707, 1993.

[JAM98] T. Jakab, D. Arató, D. D. Marcenac, "Availability analysis of some optical ring network architectures," in *Proceedings of DRCN 1998*, Brugge, Belgium, May 1998, pp. 043/1-6.

[Joh89] B. W. Johnson, *Design and Analysis of Fault-Tolerant Digital Systems*, Addison-Wesley, 1989.

[JJU02] L. Jereb, T. Jakab, F. Unghváry, "Availability analysis of multi-layer optical networks," *SPIE Optical Networks Magazine*, vol. 3, no. 2, March/April 2002, pp. 84-95.

[JUJ01] E. Jereb, F. Unghváry, T. Jakab, "A methodology for reliability analysis of multi-layer communication networks," *SPIE Optical Networks Magazine*, vol. 2, no. 5, September/October 2001, pp. 42-51.

[KCO90] H. Komine, T. Chujo, T. Ogura, K. Miyazaki, and T. Soejima, "A distributed restoration algorithm for multiple-link and node failures of transport networks," in *Proceedings of IEEE GLOBECOM 1990*, San Diego, CA, USA, December 1990, pp. 459-463.

[LaA98] A. Lardiés and A. Aguilar, "Planning methodology for SDH + optical networks," in *Proceedings of DRCN 1998*, Brugge, Belgium, May 1998.

[LCS01] H. Lee, H. Choi, S. Subramaniam, H.-A. Choi, "Survivable logical topology design in WDM optical ring networks," in *Proceedings of 39th Annual Allerton Conference on Communications, Control, and Computing*, Monticello, IL, USA, October 2001.

[Leo94] A. Leon-Garcia, Probability and random processes for electrical engineering, 2nd Edition, Addison-Wesley Publishing, 1994.

[LiS02] L. Li, A. K. Somani, "Efficient algorithms for wavelength converter placement," *SPIE Optical Networks Magazine*, vol. 3, no. 2, March/April 2002, pp. 54-62.

[LMT00] S. Lumetta, M. Médard, Y.-C. Tseng, "Capacity versus robustness: a tradeoff for link restoration in mesh networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 18, no. 12, December 2000, pp. 1765-1775.

[LTC01] K. Long, R. Tucker, S. Cheng, J. Ma, R. Zhang, "An new approach to multi-layer network survivability: strategies, model and algorithm," *Journal of High Speed Networks (IOS Press)*, vol. 10, no. 2, October 2001, pp. 127-134.

[MaG94] M. H. MacGregor, W. D. Grover, "Optimized k-shortest paths algorithm for facility restoration," Software-Practice & Experience, vol. 24(9), Wiley & Sons, 1994, pp. 823-834.

[Man01] Nitish Mandal, "Ringbusters: with optical mesh carriers can stop going in circles for 'new-world bandwidth'," Fiber Optics News, June 2001, http://www.tellium.com/documents/RINGBUSTERS_FINAL.pdf.

[MBF02] M. Médard, R. A. Barry, S. G. Finn, W. He, S. S. Lumetta, "Generalized loop-back recovery in optical mesh networks," *IEEE/ACM Transactions on Networking*, vol. 10, no. 1, February 2002, pp. 153-164.

[McD94] J. C. McDonald, "Public network integrity—Avoiding a crisis in trust," *IEEE Journal on Selected Areas in Communications*, vol. 12, no. 1, January 1994, pp. 5-12.

[MFB99] M. Medard, S. G. Finn, R. A. Barry, "WDM loop-back recovery in mesh networks," in *Proceedings of IEEE INFOCOM 1999*, New York, NY, USA, March 1999, vol. 2, pp. 752-759.

[MGM93] M. H. MacGregor, W. D. Grover, U. M. Maydell, "Connectability: A performance metric for reconfigurable transport networks," *IEEE Journal on Selected Areas in Communications*, vol. 11, no. 9, December 1993, pp. 1461-1469.

[MiI00] B. Mikac, R. Inkret, "Availability Models of WDM optical networks," in *Proceedings of DRCN 2000*, Munich, Germany, April 2000, pp. 80-85.

[MiS98] Y. Miyao, H. Saito, "Optimal design and evaluation of survivable WDM transport networks," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 7, September 1998, pp. 1190-1198.

[MoG01] G. D. Morley, W. D. Grover, "Tabu Search Optimization of Optical Ring Transport Networks," in *Proceedings of IEEE GLOBECOM 2001*, San Antonio, TX, USA, November 2001, vol. 4, pp. 2160-2164

[MoL01] E. Modiano, P. J. Lin, "Traffic grooming in WDM networks," *IEEE Communications Magazine*, vol. 39, no. 7, July 2001, pp. 124-129.

[MoN01] E. Modiano, A. Narula-Tam, "Survivable routing of logical topologies in WDM networks," in *Proceedings of IEEE INFOCOM 2001*, Anchorage, Alaska, USA, April 2001, pp. 348-357.

[Mor01] G. D. Morley, Analysis and Design of Ring-Based Networks, Ph.D. thesis, Department of Electrical and Computer Engineering, University of Alberta, Edmonton, Alberta, Canada, Spring 2001.

[MoS00] G. Mohan, A. Somani, "Routing dependable connections with specified failure restoration guarantees in WDM networks," in *Proceedings of IEEE INFOCOM 2000*, Tel-Aviv, Israel, March 2000, pp. 1761-1770.

[Muk00] B. Mukherjee, "WDM optical communication networks: progress and challenges," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 10, October 2000, pp. 1810-1824.

[NeH99] P. Neusy, R. Habel, "Availability analysis of optical shared protection rings for long haul networks," in *Proceedings of OSA/IEEE OFC 1999*, San Diego, CA, USA, February 1999, pp. TuL5-1-TuL5-4.

[Neu95] Peter Neumann, *Computer-Related Risks*, Addison-Wesley/ACM Press, ISBN 0-201-55805-X, 1995.

[NoW93] S. Nojo, H. Watanabe, "Incorporating reliability specifications in the design of telecommunication networks," *IEEE Communications Magazine*, vol. 31, no. 6, June 1993, pp. 40-43.

[Oco91] P. O'Connor, "Practical reliability engineering," 3rd Edition, John Wiley & Sons, 1991.

[Ogg01] C. Oggerino, High Availability Network Fundamentals, Cisco Press, April 2001.

[Pre93] W. H. Press et al., Numerical Recipes in C – The Art of Scientific Computing, 2nd Edition, ISBN 0521431085, Cambridge University Press, January 1993.

[RaA90] S. Rai and D. Agrawal (Eds.), Distributed Computing Network Reliability. Piscataway, NJ, IEEE Press, 1990.

[RaM99a] S. Ramamurthy, B. Mukherjee, "Survivable WDM mesh networks, Part I –Protection," in *Proceedings of IEEE INFOCOM 1999*, New York, March 1999, pp. 744-751.

[RaM99b] S. Ramamurthy, B. Mukherjee, "Survivable WDM mesh networks, Part II –Restoration," in *Proceedings of IEEE ICC 1999*, Vancouver, BC, Canada, June 1999, pp. 2023-2030.

[RaM02] R. Ramamurthy, B. Mukherjee, "Fixed-alternate routing and wavelength conversion in wavelength-routed optical networks," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, June 2002, pp. 351-367.

[RaS95] R. Ramaswami, K. N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Transactions on Networking*, vol. 3, no. 5, October 1995, pp. 489-500.

[RaS96] R. Ramaswami, K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 6, June 1996, pp. 840-851.

[RaS02] R. Ramaswami, K. N. Sivarajan, Optical Networks: A Practical Perspective, Second Edition, Morgan Kaufmann Publishers, ISBN 1-55860-655-6, 2002.

[Ray01] K. Raychaudhuri (Lucent Technologies), Panel Discussion at IEEE Bandwidth Management Workshop IX, Montebello, Quebec, Canada, 25-28 June 2001.

[RBB00] A. Rodriguez-Moral, P. Bonenfant, S. Baroni, R. Wu, "Optical data networking: Protocols, technologies, and architectures for next generation optical transport networks and optical internetworks," *IEEE/OSA Journal of Lightwave Technology*, vol. 18, no. 2, December 2000, pp. 1855-1870.

[Row89] J. W. Rowe, "A study of the path availability improvements in a large network for various route diversity protection strategies, using failure event simulation," in *Proceedings of 2nd IEE National Conference On Telecommunications*, York, UK, April 1989, pp. 167-172.

[SaA01] G. Sahin, M. Azizoglu, "Optical layer survivability for single and multiple service classes," *Journal of High Speed Networks (IOS Press)*, vol. 10, no. 2, October 2001, pp. 91-108.

[SaM02a] C. V. Saradhi, C. S. R. Murthy, "Routing differentiated reliable connections in WDM optical networks," *SPIE Optical Networks Magazine*, vol. 3, no. 3, May/June 2002, pp. 50-67.

[SaM02b] C. V. Saradhi, C. S. R. Murthy, "Dynamic establishment of segmented protection paths in single and multi-fiber WDM mesh networks," in *Proceedings of SPIE OptiComm 2002*, Boston, MA, USA, July 2002, pp.211-222.

[SAS98] R. Subramaniam, M. Azizoglu, A. K. Somani, "On the optimal placement of wavelength converters in wavelength-routed networks," in *Proceedings of IEEE INFOCOM 1998*, San Francisco, CA, USA, March-April 1998, vol. 2, pp. 902-909.

[Sch00] D. A. Schupke, "Reliability models of WDM self-healing rings," in *Proceedings of DRCN 2000*, Munich, Germany, April 2000, pp. 174-180.

[Sch03a] D. A. Schupke, "Multiple failure survivability in WDM networks with *p*-cycles," in *Proceedings of IEEE International Symposium on Circuits and Systems, ISCAS 2003*, Bangkok, Thailand, May 2003.

[Sch03b] D. A. Schupke, "The tradeoff between the number of deployed *p*-cycles and the survivability to dual fiber duct failures," in *Proceedings of IEEE ICC 2003*, Anchorage, AK, USA, May 2003, pp. 1428-1432.

[SDL01] J. Strand, R. Doverspike, G. Li, "Importance of wavelength conversion in an optical network," *SPIE Optical Networks Magazine*, vol. 2, no. 3, May/June 2001, pp. 33-44.

[Sen00] A. Sengupta, "On fault management in all-optical networks," *SPIE Optical Networks Magazine*, vol. 1, no. 4, October 2000, pp. 33-40.

[SGC02] G. Shen, W. D. Grover, T. H. Cheng, S. K. Bose, "Sparse placement of electronic switching nodes for low blocking in translucent optical networks," *OSA Journal of Optical Networking*, vol. 1, no. 12, December 2002, pp. 424-441.

[SMD81] J. D. Spragins, J. D. Markov, M. W. Doss, S. A. Mitchell, D. C. Squire, "Communication network availability predictions based on measurement data," *IEEE Transactions on Communications*, vol. COM-29, no. 10, October 1981, pp. 1482-1491.

[SNH90] H. Sakauchi, Y. Nishimura, S. Hasegawa, "A self-healing network with an economical spare channel assignment," in *Proceedings of IEEE GLOBECOM 1990*, San Diego, CA, USA, December 1990, pp. 438-443.

[Sno01] A. P. Snow, "Network reliability: the concurrent challenges of innovation, competition, and complexity," *IEEE Transactions on Reliability*, vol. 50, no. 1 , March 2001, pp. 38 -40.

[SnW00] A. P. Snow, T. M. Whiting, "Defeating Telecommunication System Fault-Tolerant Designs," in *Proceedings of third IEEE Computer Society's Information Survivability Workshop, ISW 2000*, October 2000, pp. 143-146.

[SOH92] H. Sakauchi, Y.Okanoue, S. Hasegawa, "Spare-channel design schemes for self-healing networks," *IEICE Transactions on Communications*, vol. E75-B, no. 7, July 1992, pp. 624-633.

[Sos94] J. Sosnosky, "Service applications for SONET DCS distributed restoration," *IEEE Journal on Selected Areas in Communications*, vol. 12, no. 1, January 1994, pp. 59-68.

[Spr84] J. D. Spragins, "A fast algorithm for computing availability in networks with dependent failures," in *Proceedings of IEEE INFOCOM 1984*, San Francisco, CA, USA, April 1984, pp. 36-42.

[Spr93] J. D. Spragins, "Telecommunication network reliability models based on network structures," in *Proceedings of 1st International Conference on Telecomunication Systems, Modeling and Analysis*, Nashville, TN, USA, February 1993, pp. 79-90.

[SRM02] L. Sahasrabuddhe, S. Ramamurthy, B. Mukherjee, "Fault management in IP-over-WDM networks: WDM protection versus IP restoration," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 1, January 2002, pp. 21-33.

[SrS00] M. Sridharan, A. K. Somani, "Revenue maximization in survivable WDM networks," in *Proceedings of SPIE OptiComm 2000*, Dallas, TX, USA, August 2000, pp. 291-301.

[SSG91] B. Sansó, F. Soumis, M. Gendreau, "On the evaluation of telecommunications network reliability using routing models," *IEEE Transactions on Communications*, vol. COM-39, no. 10, October 1991, pp. 1494-1501.

[SSK86] J. D. Spragins, J. C. Sinclair, Y. J. Kang, H. Jafari, "Current telecommunication network reliability models: a critical assessment," *IEEE Journal on Selected Areas in Communications*, vol. SAC-4, no. 7, October 1986, pp. 1168-1173.

[SSS01] M. Sridharan, A. K. Somani, M. V. Salapaka, "Approaches for capacity and revenue optimization in survivable WDM networks," *Journal of High-Speed Networks (IOS Press)*, vol. 10, no. 2, October 2001, pp. 109-125.

[StG00] D. Stamatelakis, W. D. Grover, "Theoretical underpinnings for the efficiency of restorable networks using preconfigured cycles (*p*-cycles)," *IEEE Transactions on Communications*, vol. 48, no. 8, August 2000, pp. 1262-1265.

[SYH01] P. Sebos, J. Yates, G. Hjalmtysson, A. Greenberg, "Auto-discovery of shared risk link groups," in *Proceedings of OSA/IEEE OFC 2001*, Anaheim, CA, USA, March 2001, paper WDD3.

[ThS02] S. Thiagarajan, A. K. Somani, "Traffic grooming for survivable WDM mesh networks," *SPIE Optical Networks Magazine*, vol. 3, no. 3, May/June 2002, pp. 88-98.

[ToN94] M. To and P. Neusy, "Unavailability analysis of long-haul networks," *IEEE Journal on Selected Areas in Communications*, vol. 12, no. 1, January 1994, pp. 100-109.

[TRL03] TRLabs MeshAnalyzer 1.0, Mesh Network Restorability and Availability Analysis Tools, User's Guide (version1), TRLabs Technical Report TR-03-01, October 2003.

[VCD01] K. Vinodkrishnan, N. Chandhok, A. Durresi, R. Jain, R. Jagannathan, S. Seetharaman, "Survivability in IP over WDM networks," *Journal of High Speed Networks (IOS Press)*, vol. 10, no. 2, October 2001, pp. 79-90.

[VDL95] D. Vercauteren, P. Demeester, J. Luystermans, E. Houtrelle, "Availability analysis of multilayer networks," in *Proceedings of 3rd International Conference on Telecomunication Systems, Modeling and Analysis*, Nashville, TN, USA, March 1995, pp. 483-493.

[VGM93] B. D. Venables, W. D. Grover, and M. H. MacGregor, "Two strategies for spare capacity placement (SCP) in mesh restorable networks," in *Proceedings of IEEE ICC 1993*, Geneva, Switzerland, May 1993, pp. 267-271.

[VSK99] K. Venugopal, M. Shivakumar, P. Kumar, "A heuristic for placement of limited range wavelength converters in all-optical networks," in *Proceedings of IEEE INFOCOM 1999*, New York, NY, USA, March 1999, vol. 2, pp. 908-915.

[VVD98] B. Van Caenegem, W. Van Parys, F. De Turck, P. M. Demeester "Dimensioning of survivable WDM networks," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 7, September 1998, pp. 1146-1157.

[WAV01] G. Willems, P. Arijs, W. Van Paris, P. Demeester, "Capacity vs. availability tradeoffs in mesh-restorable networks," in *Proceedings of DRCN 2001*, Budapest, Hungary, October 2001, pp. 107-112.

[Wil98] M. R. Wilson, "The quantitative impact of survivable network architectures on service availability," *IEEE Communications Magazine*, vol. 36, no. 5, May 1998, pp. 122-126.

[WiS97] M. Willebeek-LeMair, P. Shahabuddin, "Approximating dependability measures of computer networks: An FDDI case study," *IEEE/ACM Transactions on Networking*, vol. 5, April 1997, pp. 311-327.

[WoS00] L. Wosinska, T. Svensson, "Analysis of connection availability in all-optical networks," in *Proceedings of NFOEC 2000*, Denver, CO, USA, August 2000, pp. 409-419.

[Wu92] T.-H. Wu, *Fiber Network Service Survivability*, Artech House, Inc., Norwood, MA, 1992.

[Wu95] T.-H. Wu, "Emerging technologies for fiber network survivability," *IEEE Communications Magazine*, vol. 33, no. 2, February 1995, pp. 58-74.

[YaH88] C. H. Yang and S. Hasegawa, "FITNESS: A failure immunization technology for network service survivability," in *Proceedings of IEEE GLOBECOM 1988*, Hollywood, FL, USA, November 1988, pp. 1549-1544.

[YoQ99] M. Yoo, C. Qiao, "Supporting multiple classes of services in IP over WDM networks," in *Proceedings of IEEE GLOBECOM 1999*, Rio de Janerio, Brazil, December 1999, pp. 1023-1027.

[ZaM01] H. Zang, B. Mukherjee, "Connection management for survivable wavelength-routed WDM mesh networks," *SPIE Optical Networks Magazine*, vol. 2, no. 4, July/August 2001, pp. 17-28.

[ZhM03] K. Zhu, B. Mukherjee, "A review of traffic grooming in WDM optical networks: architectures and challenges," *SPIE Optical Networks Magazine*, vol. 3, no. 2, March/April 2003, pp. 55-64.

[ZhS00] D. Zhou, S. Subramaniam, "Survivability in optical networks," *IEEE Network*, vol. 14, no. 6, November/December 2000, pp. 16-23.

[ZJM00] H. Zang, J. P. Jue, B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *SPIE Optical Networks Magazine*, vol. 1, no. 1, January 2000, pp. 47-60.

[ZZM02] K. Zhu, H. Zang, B. Mukherjee, "Design of WDM mesh networks with sparse grooming capability," in *Proceedings of IEEE GLOBECOM 2002*, Taipei, Taiwan, November 2002, vol. 3, pp. 2696-2700.

Appendix A: Detailed Description of Test Networks

A.1 06n14s1

A.1.1 Topology

NODE	X	Y	SIZE
N01	10	50	5
N02	20	70	4
N03	40	70	5
N04	50	50	5
N05	40	30	4
N06	20	30	5

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	22	22
S02	N01	N03	36	36
S03	N01	N04	40	40
S04	N01	N05	36	36
S05	N01	N06	22	22
S06	N02	N03	20	20
S07	N02	N04	36	36
S08	N02	N06	40	40
S09	N03	N04	22	22
S10	N03	N05	40	40
S11	N03	N06	45	45
S12	N04	N05	22	22
S13	N04	N06	36	36
S14	N05	N06	20	20

A.1.2 Demand Matrix

Table A-1: Demand matrix for 06n14s1

	N02	N03	N04	N05	N06
N01	9	7	7	6	12
N02		10	6	4	5
N03			12	5	6
N04				9	7
N05					10

A.2 11n20s1

A.2.1 Topology

NODE	X	Y	SIZE
N01	162	291	5
N02	268	7	3
N03	235	137	6
N04	178	383	3
N05	40	80	3
N06	122	15	4
N07	325	33	3
N08	8	202	4
N09	276	105	3
N10	333	284	3
N11	105	316	3

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N03	170.426	170.426
S02	N01	N04	93.381	93.381
S03	N01	N08	177.868	177.868
S04	N01	N10	171.143	171.143
S05	N01	N11	62.241	62.241
S06	N02	N06	146.219	146.219
S07	N02	N07	62.65	62.65
S08	N02	N09	98.326	98.326
S09	N03	N05	203.16	203.16
S10	N03	N06	166.292	166.292
S11	N03	N08	236.123	236.123
S12	N03	N09	52.01	52.01
S13	N03	N10	176.672	176.672
S14	N04	N10	183.918	183.918
S15	N04	N11	99.086	99.086
S16	N05	N06	104.637	104.637
S17	N05	N08	126.127	126.127
S18	N06	N07	203.796	203.796
S19	N07	N09	87.092	87.092
S20	N08	N11	149.683	149.683

A.2.2 Demand Matrix

Table A-2: Demand matrix for 11n20s1

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11
N01	5	18	16	6	7	5	11	7	9	24
N02		13	2	4	8	14	4	9	3	3
N03			7	9	14	13	10	35	10	8
N04				3	3	2	5	3	5	9
N05					11	3	10	4	3	4
N06						6	7	7	4	4
N07							3	10	4	3
N08								4	4	8
N09									5	3
N10										4

A.3 11n20s2

A.3.1 Topology

NODE	X	Y	SIZE
N01	9	253	3
N02	107	269	3
N03	245	269	4
N04	375	292	3
N05	375	123	4
N06	310	42	4
N07	179	9	4
N08	82	82	3
N09	41	147	3
N10	196	155	6
N11	229	82	3

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	99.298	99.298
S02	N01	N07	297.382	297.382
S03	N01	N09	110.725	110.725
S04	N02	N03	138	138
S05	N02	N10	144.627	144.627
S06	N03	N04	132.019	132.019
S07	N03	N05	195.489	195.489
S08	N03	N10	124.085	124.085
S09	N04	N05	169	169
S10	N04	N06	258.312	258.312
S11	N05	N06	103.856	103.856
S12	N05	N10	181.838	181.838
S13	N06	N07	135.093	135.093
S14	N06	N11	90.338	90.338
S15	N07	N08	121.4	121.4
S16	N07	N11	88.482	88.482
S17	N08	N09	76.851	76.851
S18	N08	N10	135.37	135.37
S19	N09	N10	155.206	155.206
S20	N10	N11	80.112	80.112

A.3.2 Demand Matrix

Table A-3: Demand matrix for 11n20s2

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11
N01	10	6	3	4	4	5	5	9	9	4
N02		9	4	4	4	5	5	7	13	5
N03			10	9	7	6	5	6	20	7
N04				8	5	4	3	3	8	4
N05					16	8	5	4	14	8
N06						12	6	5	15	14
N07							10	7	17	14
N08								12	14	7
N09									12	5
N10										23

A.4 Bellcore Network

A.4.1 Topology

NODE	X	Y	SIZE
N01	240	360	6
N02	120	420	2
N03	60	360	3
N04	60	240	4
N05	120	330	8
N06	360	300	3
N07	240	282	2
N08	300	180	7
N09	120	120	5
N10	180	60	2
N11	240	120	4

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	20	20
S02	N01	N03	30	30
S03	N01	N04	75	75
S04	N01	N05	35	35
S05	N01	N06	50	50
S06	N01	N08	60	60
S07	N02	N03	20	20
S08	N03	N05	40	40
S09	N04	N05	35	35
S10	N04	N08	70	70
S11	N04	N09	25	25
S12	N05	N06	55	55
S13	N05	N07	45	45
S14	N05	N08	40	40
S15	N05	N09	45	45
S16	N05	N11	50	50
S17	N06	N08	40	40
S18	N07	N08	15	15
S19	N08	N09	60	60
S20	N08	N11	30	30
S21	N09	N10	18	18
S22	N09	N11	30	30
S23	N10	N11	20	20

A.4.2 Demand Matrix

Table A-4: Demand matrix for Bellcore Network

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11
N01	5	6	7	23	8	9	13	7	2	6
N02		4	3	11	1	1	3	2	1	1
N03			6	21	2	2	4	4	1	2
N04				18	2	3	7	9	2	4
N05					6	7	14	11	3	8
N06						3	9	3	1	3
N07							7	3	1	3
N08								11	5	20
N09									7	10
N10										6

A.5 Modified Bellcore Network

A.5.1 Topology

NODE	X	Y	SIZE
N01	240	360	6
N02	120	420	2
N03	60	360	3
N04	60	240	4
N05	120	330	8
N06	360	300	3
N07	240	282	2
N08	300	180	7
N09	120	120	5
N10	180	60	2
N11	240	120	4

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	20	20
S02	N01	N03	30	30
S03	N01	N04	75	75
S04	N01	N05	35	35
S05	N01	N06	50	50
S06	N01	N08	60	60
S07	N02	N03	20	20
S08	N02	N06	75	75
S09	N03	N05	40	40
S10	N04	N05	35	35
S11	N04	N08	70	70
S12	N04	N09	25	25
S13	N04	N10	80	80
S14	N05	N06	55	55
S15	N05	N07	45	45
S16	N05	N08	40	40
S17	N05	N09	45	45
S18	N05	N11	50	50
S19	N06	N08	40	40
S20	N07	N08	15	15
S21	N07	N11	27	27
S22	N08	N09	60	60
S23	N08	N11	30	30
S24	N09	N10	18	18
S25	N09	N11	30	30
S26	N10	N11	20	20

A.5.2 Demand Matrix

The demand matrix of the Modified Bellcore Network is identical to that of the Bellcore Network presented in Section A.4.2.

A.6 COST 239

A.6.1 Topology

NODE	X	Y	SIZE
PAR	140	281	6
MIL	315	350	4
ZUR	304	298	5
PRA	356	235	5
VIE	447	308	4
BER	417	161	5
AMS	242	159	5
LUX	250	214	5
BRU	185	208	5
LON	128	143	4
COP	344	50	4

SPAN	O	D	LENGTH	UNITCOST
S01	PAR	MIL	820	820
S02	PAR	ZUR	600	600
S03	PAR	BER	1090	1090
S04	PAR	LUX	400	400
S05	PAR	BRU	300	300
S06	PAR	LON	450	450
S07	MIL	ZUR	320	320
S08	MIL	VIE	820	820
S09	MIL	BRU	930	930
S10	ZUR	PRA	565	565
S11	ZUR	VIE	730	730
S12	ZUR	LUX	350	350
S13	PRA	COP	740	740
S14	PRA	VIE	320	320
S15	PRA	BER	340	340
S16	PRA	LUX	730	730
S17	VIE	BER	660	660
S18	BER	COP	390	390
S19	BER	AMS	660	660
S20	AMS	COP	760	760
S21	AMS	LUX	390	390
S22	AMS	BRU	210	210
S23	AMS	LON	550	550
S24	LUX	BRU	220	220
S25	BRU	LON	390	390
S26	LON	COP	1310	1310

A.6.2 Demand Matrix

Table A-5: Demand matrix for COST 239

	MIL	ZUR	PRA	VIE	BER	AMS	LUX	BRU	LON	COP
PAR	5	6	1	2	11	5	1	7	10	1
MIL		6	1	3	9	2	1	2	3	1
ZUR			1	3	11	3	1	6	3	1
PRA				1	2	1	1	1	1	1
VIE					9	1	1	1	2	0
BER						0	0	0	0	0
AMS							0	0	0	0
LUX								0	0	0
BRU									0	0
LON										0

A.7 12n20s1

A.7.1 Topology

NODE	X	Y	SIZE
N01	9	130	3
N02	90	130	4
N03	171	130	3
N04	252	130	3
N05	333	130	4
N06	414	130	3
N07	414	9	3
N08	333	9	4
N09	252	9	3
N10	171	9	3
N11	90	9	4
N12	9	9	3

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	81	81
S02	N01	N11	145.609	145.609
S03	N01	N12	121	121
S04	N02	N03	81	81
S05	N02	N11	121	121
S06	N02	N12	145.609	145.609
S07	N03	N04	81	81
S08	N03	N10	121	121
S09	N04	N05	81	81
S10	N04	N09	121	121
S11	N05	N06	81	81
S12	N05	N07	145.609	145.609
S13	N05	N08	121	121
S14	N06	N07	121	121
S15	N06	N08	145.609	145.609
S16	N07	N08	81	81
S17	N08	N09	81	81
S18	N09	N10	81	81
S19	N10	N11	81	81
S20	N11	N12	81	81

A.7.2 Demand Matrix

Table A-6: Demand matrix for 12n20s1

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12
N01	6	3	2	2	1	1	2	2	2	4	3
N02		6	3	3	2	2	3	3	4	6	4
N03			5	3	2	2	3	3	3	4	2
N04				6	3	2	4	3	3	3	2
N05					6	4	6	4	3	3	2
N06						3	4	2	2	2	1
N07							6	3	2	2	1
N08								6	3	3	2
N09									5	3	2
N10										6	3
N11											6

A.8 12n30s1

A.8.1 Topology

NODE	X	Y	SIZE
N01	9	180	5
N02	212	212	6
N03	334	74	4
N04	285	17	4
N05	204	9	6
N06	74	25	4
N07	98	66	6
N08	123	163	5
N09	131	123	5
N10	163	74	6
N11	212	90	5
N12	269	90	4

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	205.507	205.507
S02	N01	N06	168.077	168.077
S03	N01	N07	144.627	144.627
S04	N01	N08	115.261	115.261
S05	N01	N09	134.659	134.659
S06	N02	N03	184.196	184.196
S07	N02	N08	101.597	101.597
S08	N02	N09	120.341	120.341
S09	N02	N10	146.441	146.441
S10	N02	N11	122	122
S11	N03	N04	75.166	75.166
S12	N03	N08	229.002	229.002
S13	N03	N12	66.94	66.94
S14	N04	N05	81.394	81.394
S15	N04	N11	103.238	103.238
S16	N04	N12	74.733	74.733
S17	N05	N06	130.981	130.981
S18	N05	N07	120.354	120.354
S19	N05	N10	76.851	76.851
S20	N05	N11	81.394	81.394
S21	N05	N12	103.856	103.856
S22	N06	N07	47.508	47.508
S23	N06	N10	101.597	101.597
S24	N07	N08	100.17	100.17
S25	N07	N09	65.863	65.863
S26	N07	N10	65.49	65.49
S27	N08	N09	40.792	40.792
S28	N09	N10	58.523	58.523
S29	N10	N11	51.546	51.546
S30	N11	N12	57	57

A.8.2 Demand Matrix

Table A-7: Demand matrix for 12n30s1

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12
N01	3	2	2	3	3	5	5	4	4	3	2
N02		3	3	4	3	4	6	5	5	5	4
N03			5	4	2	3	2	2	3	4	5
N04				6	2	3	2	3	4	4	5
N05					4	6	4	5	10	8	5
N06						11	3	4	5	3	2
N07							6	10	11	6	3
N08								13	7	5	3
N09									11	6	3
N10										12	5
N11											8

A.9 15n28s1

A.9.1 Topology

NODE	X	Y	SIZE
N01	175	420	3
N02	380	385	4
N03	210	380	3
N04	280	345	3
N05	285	135	2
N06	380	255	3
N07	220	270	5
N08	365	180	6
N09	220	215	2
N10	140	220	4
N11	85	165	4
N12	365	90	3
N13	195	90	3
N14	380	310	5
N15	135	315	6

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	207.966	207.966
S02	N01	N03	53.151	53.151
S03	N01	N15	112.361	112.361
S04	N02	N03	170.074	170.074
S05	N02	N04	107.703	107.703
S06	N02	N14	75	75
S07	N03	N15	99.247	99.247
S08	N04	N14	105.948	105.948
S09	N04	N15	148.071	148.071
S10	N05	N08	91.788	91.788
S11	N05	N12	91.788	91.788
S12	N06	N07	160.702	160.702
S13	N06	N08	76.485	76.485
S14	N06	N14	55	55
S15	N07	N08	170.66	170.66
S16	N07	N09	55	55
S17	N07	N10	94.34	94.34
S18	N07	N14	164.924	164.924
S19	N08	N09	149.164	149.164
S20	N08	N11	280.401	280.401
S21	N08	N12	90	90
S22	N10	N11	77.782	77.782
S23	N10	N13	141.156	141.156
S24	N10	N15	95.131	95.131
S25	N11	N13	133.135	133.135
S26	N11	N15	158.114	158.114
S27	N12	N13	170	170
S28	N14	N15	245.051	245.051

A.9.2 Demand Matrix

Table A-8: Demand matrix for 15n28s1

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12	N13	N14	N15
N01	3	8	3	1	2	5	3	1	3	2	1	1	3	8
N02		4	6	1	5	5	6	2	3	2	2	2	13	5
N03			6	1	2	7	4	2	3	2	1	2	4	9
N04				1	3	8	5	2	3	2	2	2	7	6
N05					2	3	7	2	2	2	3	3	3	3
N06						5	12	2	2	2	3	2	14	4
N07							9	9	11	6	3	4	8	16
16								4	5	4	10	5	11	7
N09									5	3	2	2	3	5
N10										10	2	4	4	13
N11											2	5	3	8
N12												3	3	3
N13													3	4
N14														6

A.10 16n29s1

A.10.1 Topology

NODE	X	Y	SIZE
N01	24	256	3
N02	157	281	3
N03	145	109	5
N04	36	109	3
N05	145	12	3
N06	255	171	4
N07	267	243	3
N08	316	37	4
N09	376	147	5
N10	340	281	3
N11	389	243	3
N12	462	49	4
N13	523	207	5
N14	486	268	4
N15	559	293	3
N16	620	231	3

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	135.329	135.329
S02	N01	N03	190.394	190.394
S03	N01	N04	147.489	147.489
S04	N02	N06	147.323	147.323
S05	N02	N07	116.379	116.379
S06	N03	N04	109	109
S07	N03	N05	97	97
S08	N03	N06	126.27	126.27
S09	N03	N08	185.54	185.54
S10	N04	N05	145.911	145.911
S11	N05	N08	172.818	172.818
S12	N06	N07	72.993	72.993
S13	N06	N09	123.357	123.357
S14	N07	N10	82.298	82.298
S15	N08	N09	125.3	125.3
S16	N08	N12	146.492	146.492
S17	N09	N11	96.876	96.876
S18	N09	N12	130.384	130.384
S19	N09	N13	158.773	158.773
S20	N10	N11	62.008	62.008
S21	N10	N14	146.578	146.578
S22	N11	N14	100.17	100.17
S23	N12	N13	169.366	169.366
S24	N12	N16	241.015	241.015
S25	N13	N14	71.344	71.344
S26	N13	N15	93.231	93.231
S27	N13	N16	99.925	99.925
S28	N14	N15	77.162	77.162
S29	N15	N16	86.977	86.977

A.10.2 Demand Matrix

Table A-9: Demand matrix for 16n29s1

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12	N13	N14	N15	N16
N01	6	4	4	2	4	3	2	2	2	2	1	2	2	1	0
N02		7	4	3	9	8	4	5	4	4	2	3	3	2	0
N03			7	8	11	6	5	5	3	4	3	3	2	2	0
N04				4	4	3	3	3	2	2	2	2	2	1	0
N05					5	3	4	3	2	2	2	2	2	1	0
N06						15	8	11	7	8	4	4	4	2	0
N07							4	7	8	7	3	3	3	2	0
N08								8	3	4	5	3	3	2	0
N09									6	10	6	6	5	3	0
N10										11	2	4	4	2	0
N11											3	6	6	3	0
N12												4	3	2	0
N13													9	5	0
N14														5	0
N15															0

A.11 16n38s1

A.11.1 Topology

The nodes of 16n38s1 are identical to that of test network 16n29s1 detailed in Section A.10.1. The spans are the following:

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	90.427	90.427
S02	N01	N03	127.781	127.781
S03	N01	N04	98.326	98.326
S04	N01	N15	357.74	357.74
S05	N02	N03	114.215	114.215
S06	N02	N04	139.846	139.846
S07	N02	N06	98.412	98.412
S08	N02	N07	77.162	77.162
S09	N02	N10	123.004	123.004
S10	N03	N04	74	74
S11	N03	N05	65	65
S12	N03	N06	83.726	83.726
S13	N03	N08	122.772	122.772
S14	N04	N05	98.494	98.494
S15	N05	N08	114.272	114.272
S16	N05	N06	128.705	128.705
S17	N06	N07	48.508	48.508
S18	N06	N08	97.576	97.576
S19	N06	N09	81.786	81.786
S20	N06	N11	99.765	99.765
S21	N07	N09	97.745	97.745
S22	N07	N10	55.462	55.462
S23	N07	N11	81.006	81.006
S24	N08	N09	82.365	82.365
S25	N08	N12	97.329	97.329
S26	N09	N11	64.498	64.498
S27	N09	N12	85.703	85.703
S28	N09	N13	106.231	106.231
S29	N10	N11	39.825	39.825
S30	N10	N14	96.255	96.255
S31	N11	N14	67.446	67.446
S32	N12	N13	112.721	112.721
S33	N12	N16	160.863	160.863
S34	N13	N14	48.021	48.021
S35	N13	N15	60.926	60.926
S36	N13	N16	66.94	66.94
S37	N14	N15	51.245	51.245
S38	N15	N16	57.28	57.28

A.11.2 Demand Matrix

The demand matrix used for 16n38s1 is identical to that of test network 16n29s1 detailed in Section A.10.2

A.12 EuroNet

A.12.1 Topology

NODE	X	Y	SIZE
N01	20	385	3
N02	270	330	3
N03	345	370	4
N04	550	255	2
N05	325	280	7
N06	350	205	6
N07	420	185	3
N08	435	135	3
N09	470	20	2
N10	325	115	6
N11	280	145	5
N12	160	145	6
N13	145	70	2
N14	40	85	2
N15	80	305	2
N16	135	265	7
N17	250	245	4
N18	200	195	4
N19	270	200	3

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	255.979	255.979
S02	N01	N03	325.346	325.346
S03	N01	N16	166.208	166.208
S04	N02	N03	85	85
S05	N02	N17	87.321	87.321
S06	N03	N04	235.053	235.053
S07	N03	N05	92.195	92.195
S08	N04	N05	226.385	226.385
S09	N05	N06	79.057	79.057
S10	N05	N07	134.35	134.35
S11	N05	N10	165	165
S12	N05	N16	190.591	190.591
S13	N05	N17	82.765	82.765
S14	N06	N07	72.801	72.801
S15	N06	N10	93.408	93.408
S16	N06	N11	92.195	92.195
S17	N06	N17	107.703	107.703
S18	N06	N19	80.156	80.156
S19	N07	N08	52.202	52.202
S20	N08	N09	120.208	120.208
S21	N08	N10	111.803	111.803
S22	N09	N10	173.349	173.349
S23	N10	N11	54.083	54.083
S24	N10	N12	167.705	167.705
S25	N11	N12	120	120
S26	N11	N18	94.34	94.34
S27	N11	N19	55.902	55.902
S28	N12	N13	76.485	76.485
S29	N12	N15	178.885	178.885
S30	N12	N16	122.577	122.577
S31	N12	N18	64.031	64.031

S32	N13	N14	106.066	106.066
S33	N14	N16	203.531	203.531
S34	N15	N16	68.007	68.007
S35	N16	N17	116.726	116.726
S36	N16	N18	95.525	95.525
S37	N18	N19	70.178	70.178

A.12.2 Demand Matrix

Table A-10: Demand matrix for EuroNet

	N02	N03	N04	N05	N06	N07	N08	N09	N10	N11	N12	N13	N14	N15	N16	N17	N18	N19
N01	2	2	1	3	2	1	1	1	2	2	3	1	1	3	6	2	2	1
N02		7	1	14	6	2	2	1	4	4	4	1	1	2	7	7	4	3
N03			2	15	7	3	2	1	5	4	4	1	1	1	6	5	4	3
N04				3	3	2	2	1	2	2	1	0	0	0	2	1	1	1
N05					27	8	6	2	13	12	10	3	2	3	13	17	9	11
N06						12	8	3	19	16	9	2	2	2	9	11	8	11
N07							9	2	8	5	3	1	1	1	4	3	3	3
N08								2	8	5	3	1	1	1	3	3	2	3
N09									3	2	2	1	0	0	2	1	1	1
N10										28	11	3	2	2	9	8	8	9
N11											13	3	2	2	9	10	11	13
N12												8	4	3	17	9	19	7
N13													2	1	4	2	3	2
N14														1	3	2	2	1
N15															10	2	2	1
N16																12	15	7
N17																	11	12
N18																		9

A.13 Net-A

A.13.1 Topology

NODE	X	Y	SIZE
N01	183	456	3
N02	222	322	3
N03	275	163	4
N04	266	297	5
N05	403	116	4
N06	470	253	4
N07	378	241	5
N08	331	314	4
N09	476	318	5
N10	607	311	3
N11	533	410	5
N12	634	482	3
N13	513	516	4
N14	406	389	6
N15	285	437	4
N16	338	473	5
N17	433	538	3
N18	510	637	3
N19	380	549	4
N20	260	543	3

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	139.56	139.56
S02	N01	N04	179.36	179.36
S03	N01	N20	116.181	116.181
S04	N02	N03	167.601	167.601
S05	N02	N04	50.606	50.606
S06	N03	N07	129.201	129.201
S07	N04	N03	134.302	134.302
S08	N04	N05	227.002	227.002
S09	N04	N08	67.186	67.186
S10	N05	N03	136.356	136.356
S11	N05	N06	152.506	152.506
S12	N05	N07	127.475	127.475
S13	N06	N10	148.772	148.772
S14	N06	N15	260.923	260.923
S15	N07	N06	92.779	92.779
S16	N07	N08	86.822	86.822
S17	N07	N09	124.631	124.631
S18	N08	N09	145.055	145.055
S19	N08	N15	131.32	131.32
S20	N09	N11	108.227	108.227
S21	N09	N14	99.705	99.705
S22	N10	N09	131.187	131.187
S23	N10	N12	173.118	173.118
S24	N11	N12	124.036	124.036
S25	N11	N13	107.87	107.87
S26	N12	N18	198.497	198.497
S27	N13	N17	82.97	82.97
S28	N14	N11	128.725	128.725
S29	N14	N13	166.066	166.066
S30	N14	N17	151.427	151.427

S31	N15	N14	130.173	130.173
S32	N15	N16	64.07	64.07
S33	N16	N11	204.924	204.924
S34	N16	N14	108.074	108.074
S35	N16	N20	104.805	104.805
S36	N17	N19	54.129	54.129
S37	N18	N13	121.037	121.037
S38	N19	N16	86.833	86.833
S39	N19	N18	156.984	156.984
S40	N20	N19	120.15	120.15

A.13.2 Demand Matrix

Table A-11: Demand matrix for Net-A

	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	9	1	7	3	3	5	9	4	7	2	6	2	3	1	2	6	1	3	10
2		6	7	3	6	3	2	9	1	10	2	8	6	3	10	10	4	2	3
3			2	4	5	8	4	9	5	8	7	2	3	8	2	1	3	5	8
4				5	2	5	8	6	5	3	7	10	6	2	9	5	9	9	6
5					9	5	7	6	8	5	5	5	8	1	7	10	9	8	8
6						10	5	2	3	3	4	7	10	6	1	1	10	7	2
7							4	4	9	3	4	8	4	1	6	9	3	10	1
8								4	3	6	10	5	8	8	5	9	6	4	3
9									8	8	8	3	2	10	10	7	5	4	10
10										8	5	4	10	1	8	9	5	1	1
11											2	2	7	8	10	5	1	10	2
12												8	6	7	9	2	4	7	9
13													9	2	3	8	7	5	8
14														3	9	3	8	6	1
15															4	6	10	1	6
16																1	9	6	5
17																	4	7	1
18																		2	1
19																			9

A.14 22n41s1

A.14.1 Topology

NODE	X	Y	SIZE
N01	243	292	6
N02	276	235	5
N03	186	316	3
N04	260	382	3
N05	358	309	4
N06	178	210	5
N07	358	210	5
N08	97	260	4
N09	414	283	3
N10	260	137	4
N11	89	202	3
N12	317	137	6
N13	366	146	3
N14	357	106	4
N15	105	145	3
N16	130	80	4
N17	8	267	3
N18	293	15	3
N19	203	16	4
N20	349	7	3
N21	414	33	2
N22	89	7	2

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	65.863	65.863
S02	N01	N03	61.847	61.847
S03	N01	N04	91.591	91.591
S04	N01	N05	116.25	116.25
S05	N01	N06	104.637	104.637
S06	N01	N08	149.466	149.466
S07	N02	N06	101.139	101.139
S08	N02	N07	85.726	85.726
S09	N02	N09	146.11	146.11
S10	N02	N12	106.231	106.231
S11	N03	N04	99.156	99.156
S12	N03	N17	184.621	184.621
S13	N04	N05	122.201	122.201
S14	N05	N07	99	99
S15	N05	N09	61.741	61.741
S16	N06	N08	95.189	95.189
S17	N06	N10	109.786	109.786
S18	N06	N15	97.745	97.745
S19	N07	N09	92.005	92.005
S20	N07	N12	83.726	83.726
S21	N07	N13	64.498	64.498
S22	N08	N11	58.549	58.549
S23	N08	N17	89.275	89.275
S24	N10	N12	57	57
S25	N10	N16	141.947	141.947
S26	N10	N19	133.754	133.754
S27	N11	N15	59.203	59.203
S28	N11	N17	103.856	103.856

S29	N12	N13	49.82	49.82
S30	N12	N14	50.606	50.606
S31	N12	N18	124.338	124.338
S32	N13	N14	41	41
S33	N14	N20	99.323	99.323
S34	N14	N21	92.617	92.617
S35	N15	N16	69.642	69.642
S36	N16	N19	97.082	97.082
S37	N16	N22	83.726	83.726
S38	N18	N19	90.006	90.006
S39	N18	N20	56.569	56.569
S40	N19	N22	114.355	114.355
S41	N20	N21	70.007	70.007

A.14.2 Demand Matrix

Table A-12: Demand matrix for 22n41s1

	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
1	14	9	6	7	9	7	5	4	5	4	7	3	4	3	3	3	2	3	2	1	0
2		4	4	6	8	9	4	4	7	3	9	4	4	3	3	2	3	3	2	1	0
3			3	3	5	3	4	2	2	2	3	2	2	2	2	2	1	2	1	1	0
4				3	3	3	2	2	2	2	3	2	2	1	2	1	1	1	1	1	0
5					3	7	2	6	3	2	5	3	3	2	2	2	2	2	1	1	0
6						5	7	2	6	6	6	3	3	5	5	3	2	4	2	1	0
7							3	5	5	2	11	7	6	2	3	2	3	3	2	1	0
8								2	3	7	3	2	2	4	3	5	2	2	1	1	0
9									2	1	4	2	2	1	2	1	1	2	1	1	0
10										2	13	4	5	3	4	2	3	4	2	1	0
11											3	1	2	5	3	3	1	2	1	1	0
12												11	15	3	4	2	5	5	3	2	0
13													9	2	2	1	2	2	2	1	0
14														2	3	1	4	3	3	2	0
15															6	2	2	3	1	1	0
16																2	3	5	2	1	0
17																	1	2	1	1	0
18																		4	4	1	0
19																			2	1	0
20																				1	0
21																					0

A.15 Net-B

A.15.1 Topology

NODE	X	Y	SIZE
N01	92	136	3
N02	175	78	3
N03	266	117	3
N04	359	32	3
N05	390	159	3
N06	344	239	5
N07	480	223	4
N08	561	195	4
N09	515	297	3
N10	432	290	6
N11	564	411	5
N12	446	414	4
N13	504	482	3
N14	390	454	4
N15	351	316	7
N16	337	556	3
N17	168	571	4
N18	212	427	5
N19	127	451	4
N20	193	375	3
N21	155	283	6
N22	52	349	4
N23	105	254	3
N24	245	286	4
N25	210	190	4

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	101.257	101.257
S02	N01	N21	159.931	159.931
S03	N02	N03	99.005	99.005
S04	N02	N25	117.341	117.341
S05	N03	N04	125.992	125.992
S06	N03	N06	144.803	144.803
S07	N05	N04	130.729	130.729
S08	N05	N07	110.436	110.436
S09	N06	N24	109.59	109.59
S10	N07	N06	136.938	136.938
S11	N07	N08	85.703	85.703
S12	N07	N10	82.42	82.42
S13	N08	N04	259.563	259.563
S14	N08	N09	111.893	111.893
S15	N08	N10	160.206	160.206
S16	N09	N10	83.295	83.295
S17	N09	N11	124.085	124.085
S18	N10	N06	101.71	101.71
S19	N11	N10	179.067	179.067
S20	N11	N13	92.957	92.957
S21	N11	N15	233.225	233.225
S22	N12	N11	118.038	118.038
S23	N12	N13	89.376	89.376
S24	N13	N14	117.388	117.388
S25	N14	N12	68.819	68.819

S26	N15	N10	85.071	85.071
S27	N15	N12	136.488	136.488
S28	N15	N14	143.405	143.405
S29	N15	N17	313.869	313.869
S30	N16	N14	114.948	114.948
S31	N16	N18	179.627	179.627
S32	N17	N16	169.664	169.664
S33	N17	N19	126.811	126.811
S34	N18	N15	177.882	177.882
S35	N18	N17	150.572	150.572
S36	N19	N18	88.323	88.323
S37	N20	N18	55.362	55.362
S38	N21	N19	170.317	170.317
S39	N21	N22	122.332	122.332
S40	N21	N24	90.05	90.05
S41	N22	N19	126.606	126.606
S42	N22	N20	143.377	143.377
S43	N23	N01	118.714	118.714
S44	N23	N21	57.801	57.801
S45	N23	N22	108.784	108.784
S46	N24	N15	110.164	110.164
S47	N24	N20	103.078	103.078
S48	N25	N05	182.65	182.65
S49	N25	N06	142.678	142.678
S50	N25	N21	108.046	108.046

A.15.2 Demand Matrix

Table A-13: Demand matrix for Net-B

	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25
1	2	1	9	7	1	2	4	10	1	3	4	3	4	3	10	1	6	1	10	3	9	10	9	4
2		8	1	5	8	4	2	6	3	9	6	3	10	4	6	3	4	10	6	3	7	7	1	1
3			6	4	3	6	3	5	9	1	3	1	2	6	6	9	3	5	1	5	10	6	8	3
4				1	5	7	6	5	4	5	3	1	3	3	8	2	7	9	9	8	5	3	8	1
5					5	4	4	10	8	8	4	2	8	9	8	8	6	7	2	4	6	1	4	6
6						5	4	1	4	2	6	6	9	1	9	3	3	5	5	8	3	10	8	2
7							9	4	1	7	2	8	3	7	6	3	9	5	1	2	5	7	4	1
8								8	5	3	1	5	7	10	7	5	6	9	5	10	7	4	6	10
9									2	10	5	5	10	1	4	8	2	4	5	5	3	5	9	3
10										7	5	4	8	2	10	2	5	3	10	6	7	9	9	9
11											9	2	4	7	10	4	2	8	8	7	10	2	2	5
12												4	4	10	2	7	3	4	3	9	8	8	6	10
13													1	4	4	8	9	5	9	3	6	4	8	6
14														5	9	3	9	8	6	5	9	8	5	9
15															4	1	5	4	5	9	6	1	8	5
16																4	10	8	1	6	5	7	1	7
17																	1	3	3	9	10	2	10	3
18																		3	9	1	4	5	9	10
19																			8	4	5	9	10	4
20																				1	1	1	6	10
21																					3	6	10	8
22																						2	5	8
23																							7	1
24																								7

A.16 Net-C

A.16.1 Topology

NODE	X	Y	SIZE
N01	155	169	4
N02	145	72	3
N03	206	40	3
N04	350	30	3
N05	350	76	5
N06	220	129	5
N07	322	134	4
N08	437	120	4
N09	493	153	4
N10	567	133	3
N11	613	191	3
N12	534	228	5
N13	451	223	6
N14	520	309	5
N15	658	331	4
N16	587	390	5
N17	466	422	3
N18	589	445	4
N19	664	475	3
N20	594	537	3
N21	412	533	4
N22	519	539	3
N23	320	428	3
N24	117	311	3
N25	270	339	4
N26	381	391	6
N27	370	201	5
N28	301	229	5
N29	72	227	3
N30	234	221	5

SPAN	O	D	LENGTH	UNITCOST
S01	N01	N02	97.514	97.514
S02	N01	N30	94.578	94.578
S03	N03	N02	68.884	68.884
S04	N03	N07	149.305	149.305
S05	N04	N03	144.347	144.347
S06	N04	N06	163.404	163.404
S07	N05	N04	46	46
S08	N05	N07	64.405	64.405
S09	N05	N13	178.354	178.354
S10	N06	N01	76.322	76.322
S11	N06	N02	94.202	94.202
S12	N06	N05	140.389	140.389
S13	N07	N08	115.849	115.849
S14	N08	N09	65	65
S15	N08	N13	103.947	103.947
S16	N08	N27	105.119	105.119
S17	N09	N10	76.655	76.655
S18	N09	N11	125.873	125.873
S19	N09	N13	81.633	81.633
S20	N10	N05	224.361	224.361

S21	N11	N15	147.054	147.054
S22	N12	N10	100.568	100.568
S23	N12	N11	87.235	87.235
S24	N13	N12	83.15	83.15
S25	N14	N12	82.201	82.201
S26	N14	N26	161.385	161.385
S27	N15	N12	161.199	161.199
S28	N15	N16	92.315	92.315
S29	N16	N14	105.119	105.119
S30	N16	N19	114.691	114.691
S31	N17	N14	125.24	125.24
S32	N17	N16	125.16	125.16
S33	N17	N21	123.438	123.438
S34	N18	N16	55.036	55.036
S35	N18	N20	92.136	92.136
S36	N19	N15	144.125	144.125
S37	N20	N19	93.509	93.509
S38	N21	N18	197.669	197.669
S39	N22	N18	117.201	117.201
S40	N22	N20	75.027	75.027
S41	N22	N21	107.168	107.168
S42	N23	N21	139.603	139.603
S43	N24	N23	234.303	234.303
S44	N24	N25	155.541	155.541
S45	N25	N26	122.577	122.577
S46	N25	N30	123.369	123.369
S47	N26	N13	182	182
S48	N26	N23	71.344	71.344
S49	N26	N27	190.318	190.318
S50	N27	N07	82.42	82.42
S51	N27	N13	83.934	83.934
S52	N28	N14	233.154	233.154
S53	N28	N25	114.285	114.285
S54	N28	N26	180.677	180.677
S55	N28	N27	74.465	74.465
S56	N29	N01	101.257	101.257
S57	N29	N24	95.294	95.294
S58	N29	N30	162.111	162.111
S59	N30	N06	93.059	93.059
S60	N30	N28	67.476	67.476

A.16.2 Demand Matrix

The demand matrix of test network Net-C is shown in Table A-14 on the following two pages.

Table A-14: Demand matrix for Net-C

	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
1	5	4	4	4	8	10	10	10	6	2	3	5	6	8	8	6	7	10	2	1	6	3	10	4	2	3	10	3	8
2		3	6	9	1	4	10	2	1	5	1	8	2	1	5	10	5	10	4	7	7	10	4	2	4	10	8	2	9
3			9	3	9	3	7	2	10	1	3	7	2	6	4	2	9	3	5	3	2	5	6	8	1	5	5	3	5
4				1	2	4	10	1	8	1	4	1	7	8	3	1	9	10	5	4	8	3	10	10	9	9	7	8	10
5					4	7	8	4	7	2	8	5	8	7	6	9	5	9	8	4	8	1	7	10	9	10	2	8	3
6						7	4	10	1	4	1	1	8	9	10	5	5	5	1	3	1	8	4	3	6	6	9	5	4
7							10	1	5	5	1	10	7	2	5	3	10	10	8	3	9	8	3	5	6	5	5	4	8
8								5	7	6	8	7	5	5	8	7	5	10	9	2	3	10	1	6	4	9	3	4	10
9									7	8	4	10	2	10	7	9	9	5	5	5	8	9	8	4	7	9	1	4	9
10										1	9	4	5	4	8	6	9	2	1	2	5	1	3	5	7	7	5	9	6
11											8	6	5	5	2	7	7	2	5	6	1	3	2	9	2	7	5	4	6
12												6	5	2	8	5	3	2	7	3	1	8	10	3	9	4	8	5	7
13													8	10	9	1	5	3	2	5	9	5	4	7	3	2	8	4	4
14														8	4	2	3	4	6	7	1	9	1	6	3	2	10	2	2
15															9	2	10	7	8	7	3	10	5	4	10	5	3	2	4
16																7	10	2	1	9	4	8	1	9	1	3	10	10	2
17																	5	10	9	6	2	9	7	1	1	7	10	7	4
18																		7	5	3	7	5	9	3	2	6	10	8	7
19																			3	7	2	3	8	4	6	9	5	4	10
20																				1	5	4	1	4	8	10	3	9	9
21																					2	3	9	9	5	6	7	10	8

Table A-14: Demand matrix for Net-C

	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
22																						7	3	5	8	10	6	5	10
23																							5	7	7	8	6	9	3
24																								3	6	1	8	3	7
25																									1	4	10	4	7
26																										5	10	1	6
27																											10	10	7
28																												5	7
29																													7

Appendix B: Joint Capacity Placement Formulations

B.1 MJCP

The MJCP formulation was originally presented in [DoG00]. Given the high number of times it is used as a base model in this thesis, it is provided in this section as a reference.

$$\text{Minimize: } \sum_{k \in \boldsymbol{S}} \sum_{m \in \boldsymbol{M}} C_k^m \cdot n_k^m \tag{B.1}$$

$$\sum_{q \in \boldsymbol{Q}^r} g^{r,q} = d^r, \forall r \in \boldsymbol{D} \tag{B.2}$$

$$w_k = \sum_{r \in \boldsymbol{D}} \sum_{q \in \boldsymbol{Q}^r} \zeta_k^{r,q} \cdot g^{r,q}, \forall k \in \boldsymbol{S} \tag{B.3}$$

$$\sum_{p \in \boldsymbol{P}_i} f_i^p = w_i, \forall i \in \boldsymbol{S} \tag{B.4}$$

$$s_k \geq \sum_{p \in \boldsymbol{P}_i} \delta_{i,k}^p \cdot f_i^p, \forall i, k \in \boldsymbol{S}, i \neq k \tag{B.5}$$

$$w_k + s_k \leq \sum_{m \in \boldsymbol{M}} n_k^m \cdot Z^m \quad \forall k \in \boldsymbol{S} \tag{B.6}$$

where $\boldsymbol{S}$, $\boldsymbol{P}_i$, w_i, $\delta_{i,k}^p$, f_i^p, and s_k are as defined in Section 2.5.3, $\boldsymbol{M}$, Z^m, C_k^m, and n_k^m are as defined in Section 2.5.4, and the following new parameters are used:

- $\boldsymbol{D}$ Set of demand pairs
- $\boldsymbol{Q}^r$ Set of eligible working routes for demand pair r, $\forall r \in \boldsymbol{D}$
- $\zeta_k^{r,q}$ Equal to 1 if eligible route q for demand pair r crosses span k, equal to 0 otherwise, $\forall r \in \boldsymbol{D}$, $\forall q \in \boldsymbol{Q}^r$, $\forall k \in \boldsymbol{S}$.

as well as the following new variable:

- $g^{r,q}$ working flow on route q serving demand pair r, $\forall r \in \boldsymbol{D}$, $\forall q \in \boldsymbol{Q}^r$.

The objective function (B.1) minimizes the total cost of modular capacity. Constraint set (B.2) ensures that for each traffic demand, there is enough flow over all eligible working routes to fully route all traffic, and constraints (B.3) place enough associated working capacity. Constraint set (B.4) guarantees that there is enough restoration flow over all eligible restoration routes to fully restore any single physical span cut (or some other failure of a span), while constraint set (B.5) provisions enough associated spare capacity on each span to accommodate all restoration flows simultaneously routed over it. Finally, constraint set (B.6) places enough total modular capacity on each span to carry all working and spare capacities on it.

B.2 JCP

As explained several times throughout this thesis, all modular design formulations can also be used to produce non-modular designs. All that is needed is to define the set $\boldsymbol{M}$ as being composed of a single module type of size 1. The corresponding problem that is solved in that case can be expressed by using equations as given in the previous section, dropping (B.6) and replacing (B.1) by:

$$\text{Minimize: } \sum_{k \in \boldsymbol{S}} C_k \cdot s_k \tag{B.7}$$

where C_k is the cost of a single unit of capacity on span k.

Appendix C: Calculation of Average Time in a Given Failure State

In this section, we proove formal results concerning the average duration of a dual-failure state and a triple failure state.

C.1 Average Duration of a Dual-Failure State

Since we are looking for an average value of the state duration, we will assume that the duration of each failure is *MTTR*. Based on this assumption, the duration of the dual failure state is equal to the expected value of random variable X represented on Figure C-1.

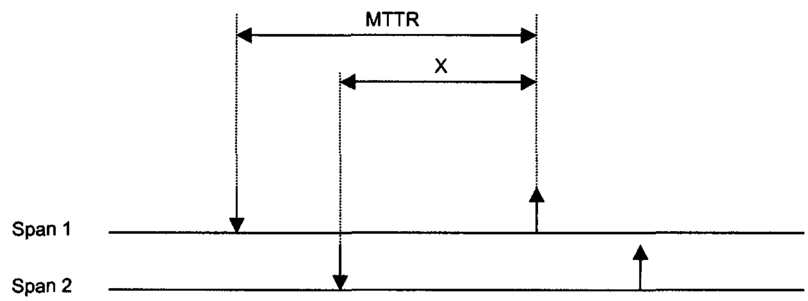

Figure C-1 Chronological view of a dual-failure

This random variable is uniformly distributed between 0 and *MTTR*, its probability density function is therefore:

$$f_X(x) = \frac{1}{MTTR},\ x \in [0, MTTR] \tag{C.1}$$

$$f_X(x) = 0,\ x \notin [0, MTTR] \tag{C.2}$$

The expectation of X is:

$$E(X) = \int_0^{MTTR}\left(\frac{1}{MTTR}\cdot x\right)dx = \frac{MTTR}{2} \tag{C.3}$$

The average time spent in a dual failure state is therefore half of *MTTR*.

C.2 Average Duration of a Triple-Failure State

For a triple-failure state, the average duration is given by the minimum of random variables X and Y, shown on Figure C-2. These two random variables are both uniformly distributed between 0 and *MTTR*, and they both have the same probability density function as given in (C.1) and (C.2). Let us define variable Z as:

$$Z = \min(X, Y) \tag{C.4}$$

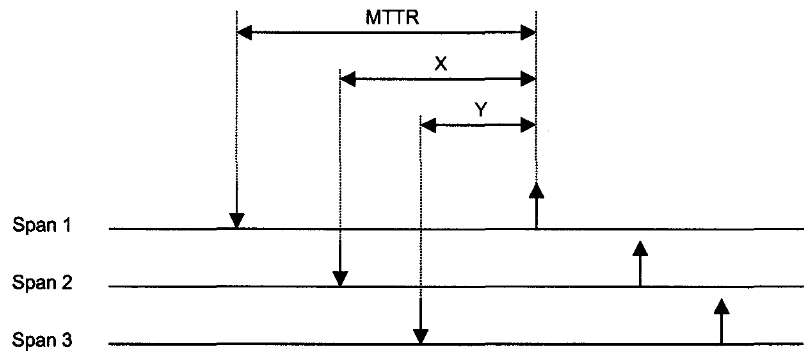

Figure C-2 Chronological view of a triple-failure

To simplify the calculations, we can define new variables X^*, Y^* and Z^* as follows:

$$X^* = \frac{X}{MTTR} \tag{C.5}$$

$$Y^* = \frac{Y}{MTTR} \tag{C.6}$$

$$Z^* = \frac{Z}{MTTR} \tag{C.7}$$

in which case:

$$Z^* = \min(X^*, Y^*) \tag{C.8}$$

and variables X^* and Y^* are two random variables uniformly distributed between 0 and 1.

$$P(Z^* \le z) = P(X^* \le z \cap Y^* \le z) = P(X^* \le z) \cdot P(Y^* \le z) \tag{C.9}$$

The cumulative density function of Z^* is defined as:

$$F_{Z^*}(z) = P(Z^* \ge z) = 1 - P(Z^* \le z) = 1 - (P(X^* \le z) \cdot P(Y^* \le z)) \tag{C.10}$$

$$F_{Z^*}(z) = 1 - (1 - F_{X^*}(z)) \cdot (1 - F_{Y^*}(z)) \tag{C.11}$$

Since X^* and Y^* are both uniformly distributed between 0 and 1, we have:

$$F_{X^*}(z) = F_{Y^*}(z) = z \text{, } z \in [0, 1] \tag{C.12}$$

$$F_{X^*}(z) = F_{Y^*}(z) = 1 \text{, } z > 1 \tag{C.13}$$

Thus:

$$F_{Z^*}(z) = 1 - (1 - z)^2 \text{, } z \in [0, 1] \tag{C.14}$$

$$F_{Z^*}(z) = 1 \text{, } z > 1 \tag{C.15}$$

$$f_{Z*}(z) = \frac{d}{dz}F_{Z*}(z) = 2 \cdot (1-z), z \in [0, 1] \quad (C.16)$$

$$f_{Z*}(z) = 0, z > 1 \quad (C.17)$$

$$E(Z^*) = \int_0^{\infty} (z \cdot f_{Z*}(z))dz = 2 \cdot \int_0^1 (z \cdot (1-z))dz \quad (C.18)$$

$$E(Z^*) = 2 \cdot \int_0^1 (z - z^2)dz = \left[z^2 - \left(\frac{2}{3} \cdot z^3\right)\right]_0^1 = 1 - \frac{2}{3} = \frac{1}{3} \quad (C.19)$$

$$E(Z) = E(MTTR \cdot Z^*) = MTTR \cdot E(Z^*) = \frac{MTTR}{3} \quad (C.20)$$

The average duration of a triple failure is therefore a third of *MTTR*.

Appendix D: AMPL Formulations

D.1 MSCP

```
####################################################################
# sr-mscp.mod                                                      #
# Span-restorable mesh networks modular spare capacity placement   #
# includes multi-QoP capability                                    #
#                                                                  #
# Written by Matthieu Clouqueur (based on Formulation by           #
# J. Doucette in JSAC vol. 18, no. 10, Oct. 2000, pp. 1912-1923.)  #
# TRLabs, Network Systems                                          #
# February 2002                                                    #
####################################################################
############################## SETS ################################
# set of all spans:
set SPANS;

# set of all types of capacity modules:
set MODULETYPES;

# set of eligible restoration routes for each span i:
set RROUTES{i in SPANS};

########################### PARAMETERS #############################
# cost of each module of type m on span j:
param Cost{k in SPANS, m in MODULETYPES};

# number of gold working capacity units on span j:
param work_g{k in SPANS};

# number of silver or bronze working capacity units on span j:
param work_s_b{k in SPANS};

# number of economy working capacity units on span j:
param work_e{k in SPANS};

# capacity of module type m:
param Z{m in MODULETYPES};

# equal to 1 if pth restoration route for failure of span i uses
# span k and 0 otherwise:
param Delta{i in SPANS, k in SPANS, p in RROUTES[i]} default 0;

############################ VARIABLES #############################
# flow on route p for restoration of span i:
var f{i in SPANS, p in RROUTES[i]} >=0, <=2500;

# number of spare capacity units on span j
var spare{j in SPANS} >=0, <=2500;

# number of modules of type m placed on span k:
var n{k in SPANS, m in MODULETYPES} >=0, <=5000 integer;

######################## OBJECTIVE FUNCTION ########################
minimize TOTCOST: sum{k in SPANS, m in MODULETYPES} Cost[k,m] * n[k,m];
```

```
########################### CONSTRAINTS ############################
subject to RESTORING{i in SPANS}: sum{p in RROUTES[i]} f[i,p] = work_g[i];

subject to SPARE{i in SPANS, k in SPANS: i<>k}: spare[k] + work_e[k] >= sum{p in
RROUTES[i]} Delta[i,k,p] * f[i,p];

subject to NUMMODULES{k in SPANS}: work_g[k] + work_s_b[k] + work_e[k] + spare[k]
= sum{m in MODULETYPES} n[k,m] * Z[m];
```

D.2 TCMF-Flow

```
####################################################################
# sr_tcmf_flow.mod                                                 #
# Span-restorable mesh networks - Two commodity max flow analysis  #
# (Flow conservation approach)                                     #
#                                                                  #
# Written by Matthieu Clouqueur                                    #
# TRLabs, Network Systems                                          #
# August 2003                                                      #
####################################################################
############################### SETS ###############################
# set of all spans:
set SPANS;

# set of all nodes;
set NODES;

############################ PARAMETERS ############################
# number of working links allocated on span j:
param Work{k in SPANS};

# number of spare links allocated on span k:
param Spare{k in SPANS};

# Adjancency: 1 if node n is adjacent to node m, 0 otherwise:
param Adjacency{n in NODES, m in NODES} default 0;

# Relation: 1 if node m and node n are connected by span k
param Relation{n in NODES, m in NODES, k in SPANS} default 0;

# IncidenceS: 1 if span k starts at node n, 0 otherwise:
param IncidenceS{k in SPANS, n in NODES} default 0;

# IncidenceE: 1 if span k ends at node n, 0 otherwise:
param IncidenceE{k in SPANS, n in NODES} default 0;

############################ VARIABLES #############################
# backup flow from node n to node m for restoration of span i,
# during dual span-failure (i,j)
var x_i{n in NODES, m in NODES, i in SPANS, j in SPANS: n<>m and Adjacency[n,m]=1
and i<>j} >=0, <=1000 integer;

# backup flow from node n to node m for restoration of span j,
# during dual span-failure (i,j)
var x_j{n in NODES, m in NODES, i in SPANS, j in SPANS: n<>m and Adjacency[n,m]=1
```

```
and i<>j} >=0, <=1000 integer;

# number of non restored units in the case of failure of span i and
# span j simultaneously:
var N2{i in SPANS, j in SPANS: i<>j} >=0, <= 3500 integer;

######################## OBJECTIVE FUNCTION ########################
minimize totunrestorable: sum{i in SPANS, j in SPANS: i<>j} N2[i,j] ;

########################### CONSTRAINTS ###########################
subject to FLOW_CONSERVATION1{i in SPANS, j in SPANS, n in NODES: i<>j and Inci-
denceS[i,n]=0 and IncidenceE[i,n]=0}:
sum{m in NODES: Adjacency[n,m]=1} x_i[m,n,i,j] = sum{o in NODES: Adjacen-
cy[n,o]=1} x_i[n,o,i,j];

subject to FLOW_CONSERVATION2{i in SPANS, j in SPANS, n in NODES: i<>j and Inci-
denceS[j,n]=0 and IncidenceE[j,n]=0}:
sum{m in NODES: Adjacency[n,m]=1} x_j[m,n,i,j] = sum{o in NODES: Adjacen-
cy[n,o]=1} x_j[n,o,i,j];

subject to SPECIAL1{i in SPANS, j in SPANS, m in NODES, n in NODES: i<>j and Re-
lation[m,n,j]=1}:
x_i[m,n,i,j] = 0;

subject to SPECIAL2{i in SPANS, j in SPANS, m in NODES, n in NODES: i<>j and Re-
lation[m,n,i]=1}:
x_j[m,n,i,j] = 0;

subject to NO_FLOW_BACK1{i in SPANS, j in SPANS,  n in NODES, m in NODES: i<>j and
IncidenceS[i,n]=1 and Adjacency[n,m]=1}:
x_i[m,n,i,j] = 0;

subject to NO_FLOW_BACK2{i in SPANS, j in SPANS,  n in NODES, m in NODES: i<>j and
IncidenceS[j,n]=1 and Adjacency[n,m]=1}:
x_j[m,n,i,j] = 0;

subject to SPARE_ALLOC_LIMIT{i in SPANS, j in SPANS, m in NODES, n in NODES, k in
SPANS: i<>j and i<>k and j<>k and Relation[m,n,k]=1}:
x_i[m,n,i,j] + x_i[n,m,i,j] + x_j[n,m,i,j] + x_j[m,n,i,j] <= Spare[k];

subject to RESTORATION_LIMIT1{i in SPANS, j in SPANS, n in NODES: i<>j and Inci-
denceS[i,n]=1}:
sum{m in NODES: Adjacency[n,m]=1 and IncidenceE[i,m]=0} x_i[n,m,i,j] <= Work[i];

subject to RESTORATION_LIMIT2{i in SPANS, j in SPANS, n in NODES: i<>j and Inci-
denceS[j,n]=1}:
sum{m in NODES: Adjacency[n,m]=1 and IncidenceE[j,m]=0} x_j[n,m,i,j] <= Work[j];

subject to DEFINITION_N2{i in SPANS, j in SPANS, n in NODES, m in NODES: i<>j and
IncidenceS[i,n]=1 and IncidenceS[j,m]=1}:
N2[i,j] = Work[i] + Work[j] - sum{o in NODES: Adjacency[o,n]=1 and Inciden-
ceE[i,o]=0} x_i[n,o,i,j] - sum{o in NODES: Adjacency[o,m]=1 and Inciden-
ceE[j,o]=0} x_j[m,o,i,j];
```

D.3 DFMC

```
#####################################################################
# sr_dfmc.mod                                                       #
# Span-restorable mesh networks dual failure minimize capacity      #
#                                                                   #
# Written by Matthieu Clouqueur                                     #
# TRLabs, Network Systems                                           #
# December 1999                                                     #
#####################################################################
############################### SETS ################################
# set of all spans:
set SPANS;

# set of all module types:
set MODULETYPES;

# set of all restoration routes for each span failure i:
set RROUTES{i in SPANS};

############################ PARAMETERS #############################
# cost of each capacity module of type m on span j:
param Cost{j in SPANS, m in MODULETYPES};

# number of working links placed on span j:
param Work{j in SPANS};

# size of modules of type m:
param Z{m in MODULETYPES};

# equal to 1 if pth restoration route for failure of span i uses
# span j and 0 otherwise:
param Delta{i in SPANS, j in SPANS, p in RROUTES[i]} default 0;

############################ VARIABLES ##############################
# rest flow on route p span i when span j has failed simultaneously:
var f{i in SPANS, j in SPANS, p in RROUTES[i]} >=0, <=5000 integer;

# number of spare links place on span j:
var spare{j in SPANS} >=0, <=1000 integer;

# number of modules of type m on span j:
var n{j in SPANS, m in MODULETYPES} >=0, <= 2000 integer;

######################## OBJECTIVE FUNCTION #########################
minimize totcost: sum{j in SPANS, m in MODULETYPES} Cost[j,m] * n[j,m];

########################### CONSTRAINTS #############################
subject to restn{i in SPANS, j in SPANS: i <> j}:
sum{p in RROUTES[i]} f[i,j,p] = Work[i];

subject to security{i in SPANS, j in SPANS, p in RROUTES[i]: i <> j}:
f[i,j,p] <= (1-Delta[i,j,p])*5000;

subject to sparasst{i in SPANS, j in SPANS, k in SPANS: i<>j && i <> k && j <> k}:
spare[k] >= sum{p in RROUTES[i]} Delta[i,k,p] * f[i,j,p] + sum{p in RROUTES[j]}
Delta[j,k,p] * f[j,i,p];
```

```
subject to enoughcap{j in SPANS}:
Work[j] + spare[j] <= sum{m in MODULETYPES} n[j,m] * Z[m];
```

D.4 DFMR

```
####################################################################
# sr_dfmr.mod                                                      #
# Span-restorable mesh networks dual failure maximize              #
# restorability                                                    #
#                                                                  #
# Written by Matthieu Clouqueur                                    #
# TRLabs, Network Systems                                          #
# December 1999                                                    #
####################################################################
############################### SETS ###############################
# set of all spans:
set SPANS;

# set of all module types:
set MODULETYPES;

# set of all restoration routes for each span failure i:
set RROUTES{i in SPANS};

############################ PARAMETERS ############################
# Maximum total capacity cost allowed:
param MaxCapCost;

# cost of each capacity module of type m on span j:
param Cost{j in SPANS, m in MODULETYPES};

# number of working links placed on span j:
param Work{j in SPANS};

# size of modules of type m:
param Z{m in MODULETYPES};

# equal to 1 if pth restoration route for failure of span i uses
# span j and 0 otherwise:
param Delta{i in SPANS, j in SPANS, p in RROUTES[i]} default 0;

############################ VARIABLES #############################

# restoration flow through pth restoration route for failure of span
# i only:
var f1{i in SPANS, p in RROUTES[i]} >=0, <=3500 integer;

# restoration flow through pth restoration route for failure of span
# i when span j has failed simultaneously:
var f2{i in SPANS, j in SPANS, p in RROUTES[i]} >=0, <=3500 integer;

# number of non restored units in the case of failure of span i and
# span j simultaneously:
var N2{i in SPANS, j in SPANS} >=0, <= 3500 integer;
```

```
# number of spare links place on span j:
var spare{j in SPANS} >=0, <=3500 integer;

# number of modules of type m on span j:
var n{j in SPANS, m in MODULETYPES} >=0, <= 2000 integer;

######################## OBJECTIVE FUNCTION ########################
minimize totunrestorable: sum{i in SPANS, j in SPANS: i<>j} N2[i,j] ;

########################### CONSTRAINTS ###########################
subject to MAX_CAP_COST:
sum{j in SPANS, m in MODULETYPES} Cost[j,m] * n[j,m] <= MaxCapCost;

subject to SPARE_ALLOC1{i in SPANS, k in SPANS: i<>k}:
spare[k] >= sum{p in RROUTES[i]} Delta[i,k,p] * f1[i,p];

subject to SPARE_ALLOC2{i in SPANS, j in SPANS, k in SPANS: i<>j && i <> k && j
<> k}:
spare[k] >= sum{p in RROUTES[i]} Delta[i,k,p] * f2[i,j,p] + sum{p in RROUTES[j]}
Delta[j,k,p] * f2[j,i,p];

subject to RESTORATION{i in SPANS}:
sum{p in RROUTES[i]} f1[i,p] = Work[i];

subject to SPECIAL_CONST1{i in SPANS, j in SPANS, p in RROUTES[i]: i <> j}:
f2[i,j,p] <= (1-Delta[i,j,p])*1000;

subject to SPECIAL_CONST2{i in SPANS, j in SPANS: i <> j}:
sum{p in RROUTES[i]} f2[i,j,p] <= Work[i];

subject to DEFINITION_N2{i in SPANS, j in SPANS: i <> j}:
N2[i,j] = Work[i] + Work [j] - sum{p in RROUTES[i]} f2[i,j,p] - sum{p in
RROUTES[j]} f2[j,i,p];

subject to TOTAL_CAP{j in SPANS}:
Work[j] + spare[j] <= sum{m in MODULETYPES} n[j,m] * Z[m];
```

D.5 MRCP

```
######################################################################
# sr_mrcp.mod                                                        #
# Span-restorable mesh networks multiple restorability capacity     #
# placement                                                          #
#                                                                    #
# Written by Matthieu Clouqueur                                      #
# TRLabs, Network Systems                                            #
# February 2001                                                      #
######################################################################
################################ SETS ################################
# set of all spans
set SPANS;

# set of all capacity module types:
set MODTYPES;

# set of all eligible restoration routes for each span failure i:
set RROUTES{i in SPANS};

# set of all demand relations:
set RELATIONS;

# set of all eligible working routes for each demand pair r:
set WROUTES{r in RELATIONS};

############################# PARAMETERS #############################
# cost of each capacity module of type m on span k:
param Cost{k in SPANS, m in MODTYPES};

# number of demand units for demand r:
param DemUnits{r in RELATIONS};

# equal 1 if rth demand requires R1 restorability:
param R1{r in RELATIONS};

# equal 1 if rth demand requires R2 restorability:
param R2{r in RELATIONS};

# size of modules of type m:
param Z{m in MODTYPES};

# equal to 1 if pth restoration route for failure of span i uses
# span k and 0 otherwise:
param Delta{i in SPANS, k in SPANS, p in RROUTES[i]} default 0;

# equal to 1 if qth working route for demand relation r
# uses span k and 0 otherwise:
param Zeta{k in SPANS, r in RELATIONS, q in WROUTES[r]} default 0;

############################# VARIABLES ##############################
# restoration flow through pth restoration route for failure of
# span i:
var f1{i in SPANS, p in RROUTES[i]} >=0, <=500 integer;

# restoration flow through pth restoration route for failure of span
# i when span j has failed simultaneously:
```

```
var f2{i in SPANS, j in SPANS, p in RROUTES[i]} >=0, <=500 integer;

# working capacity required by qth working route for demand between
# node pair r:
var g{r in RELATIONS, q in WROUTES[r]} >=0, <=500 integer;

# number of MODTYPES of type m placed on span j:
var n{j in SPANS, m in MODTYPES} >=0, <=1000 integer;

# number of working links placed on span j:
var work{k in SPANS} >=0, <=500 integer;

# number of spare links place on span j:
var spare{k in SPANS} >=0, <=500 integer;

######################## OBJECTIVE FUNCTION ########################
minimize totcost: sum{j in SPANS, m in MODTYPES} Cost[j,m] * n[j,m];

########################### CONSTRAINTS ###########################
subject to SERVE_DEM{r in RELATIONS}:
sum{q in WROUTES[r]} g[r,q] = DemUnits[r];

subject to WORK_ALLOC{k in SPANS}:
work[k] = sum{r in RELATIONS, q in WROUTES[r]} Zeta[k,r,q] * g[r,q];

subject to RESTORATION1{i in SPANS}:
sum{p in RROUTES[i]} f1[i,p] = sum{r in RELATIONS, q in WROUTES[r]} Zeta[i,r,q] *
g[r,q] * R1[r];

subject to RESTORATION2{i in SPANS, j in SPANS: i <> j}:
sum{p in RROUTES[i]} f2[i,j,p] = sum{r in RELATIONS, q in WROUTES[r]} Zeta[i,r,q]
* g[r,q] * R2[r];

subject to SPECIAL_CONST{i in SPANS, j in SPANS, p in RROUTES[i]: i <> j}:
f2[i,j,p] <= (1-Delta[i,j,p])*1000;

subject to SPARE_ALLOC1{i in SPANS, k in SPANS: i <> k}:
spare[k] >= sum{p in RROUTES[i]} Delta[i,k,p] * f1[i,p];

subject to SPARE_ALLOC2{i in SPANS, j in SPANS, k in SPANS: i<>j && i <> k && j
<> k}:
spare[k] >= sum{p in RROUTES[i]} Delta[i,k,p] * f2[i,j,p] + sum{p in RROUTES[j]}
Delta[j,k,p] * f2[j,i,p];

subject to TOTAL_CAP{k in SPANS}:
spare[k] + work[k] <= sum{m in MODTYPES} n[k,m] * Z[m];
```