Anders Mykkeltveit

Dependability differentiation in communication networks

Thesis for the degree of Philosophiae Doctor

Trondheim, December 2008

Norwegian University of Science and Technology
Faculty of Information Technology, Mathematics
and Electrical Engineering
Department of Telematics

NTNU
Norwegian University of
Science and Technology

NTNU
Norwegian University of Science and Technology

Thesis for the degree of Philosophiae Doctor

Faculty of Information Technology, Mathematics
and Electrical Engineering
Department of Telematics

ISBN 978-82-471-1330-1 (printed ver.)
ISBN 978-82-471-1331-8 (electronic ver.)
ISSN 1503-8181

Doctoral theses at NTNU, 2008:312

Printed by NTNU-trykk

Abstract

Unintentional failures affect links and nodes in communication networks. Recovery mechanisms are the key tool for achieving the dependability required by the services using the network. However, high dependability in communication networks comes at a high cost in terms of the capacity needed by these mechanisms. The traffic from all services and users is carried by the same backbone network. Since the users and services have different requirements, and users have different willingness to pay for a high quality of service, it is desirable to have methods that enable provision of different levels of dependability in the same network, i.e. dependability differentiation.

The thesis addresses dependability differentiation in connection-oriented backbone communication networks. Two methods to provide connections meeting differentiated guarantees on the asymptotic availability are proposed. The first of these uses a novel flexible arrangement for dedicated protection denoted a protection pattern. The protection pattern is used in a proposed distributed connection management system. The system is compared with alternative proposals based on centralized management and shows good performance. The second proposal uses shared protection, which may potentially use less resources in terms of bandwidth, but has higher complexity than dedicated protection. The proposed system is based on rules to control the sharing to enable provision of guarantees. Simulation results show that the proposed method performs significantly better than an alternative strategy based on dedicated protection.

A different approach to availability-guaranteed services is to offer guarantees on the interval availability which is a measure commonly used in Service Level Agreements (SLAs). The thesis contains a proposal of using adaptive management to increase compliance with interval availability guarantees. Different adaptive management policies are proposed and compared to alternative static provisioning policies in a case study.

The thesis also addresses the problem of measuring dependability by simulation. To reduce the simulation effort needed to obtain precise estimates of dependability attributes, a rare-event simulation technique has been applied to the well-known Network Simulator 2 (NS2). The results show that the technique is applicable to this types of simulation scenario, but the gain is modest.

The thesis also contains a broad literature survey of dependability differentiation research. This is the first survey of the topic. Hence, it is in itself a significant contribution. A classification scheme for how to approach differentiation is proposed and a critical evaluation of the state of art is given. This thesis contributes to fill in some of the “gaps” identified, but there are still significant challenges ahead before differentiation may be deployed in operational networks.

Preface

This thesis is submitted in partial fulfillment of the requirements for the degree of philosophiae doctor (PhD) at the Norwegian University of Science and Technology (NTNU). The work was performed at the Centre for Quantifiable Quality of Service in Communication Systems (Q2S), Centre of Excellence (CoE), during 2004-2008, and has been supervised by Professor Bjarne E. Helvik.

Parts of this work were conducted within the EU FP7 Network of Excellence framework Euro-NGI and its successor Euro-FGI. In particular, some of the work included in this thesis has been presented in preliminary versions at workshops organized by Euro-NGI.

Numerous people have directly or indirectly contributed to the work presented in this thesis. First of all I would like to thank my thesis advisor Professor Bjarne E. Helvik, for all the help and support you have given over the last four years. I would also like to thank all co-authors of the papers I have been involved in preparing. In particular, I thank Dr. Piotr Chołda for the fruitful discussions on differentiation. I thank my office mate Laurent Paquereau for valuable discussions, helpful feedback on the papers, and for being a great roommate. I would also like to thank all colleagues who have created a very enjoyable atmosphere at Q2S in coffee breaks and lunches. Finally, I would like to thank my family for their support throughout my education, from I first entered primary school and up until today, and most of all my girlfriend Solvor for cheering me up in troublesome times during the work on this thesis.

Contents

Abbreviations

AN	Adaptive, no preemption
AP	Adaptive with preemption
APS	Automatic Protection Switching
AS	Autonomous System
BER	Bit Error Rate
BGP	Border Gateway Protocol
CDF	cumulative distribution function
CE	Cross-Entropy
CEAS	Cross-Entropy Ant System
DiR	Differentiated Reliability
DPM	Disjoint Path-pair Matrix
DPP	Dedicated Path Protection
DWDM	Dense Wavelength Division Multiplexing
GMPLS	Generalized Multi-Protocol Label Switching
i.i.d.	independent and identically distributed
IP	Internet Protocol
ISP	Internet Service Provider
LSP	Label Switched Path
MDT	Mean Down Time
MFP	Maximum Failure Probability
MPLS	Multi-Protocol Label Switching

MTBI	Mean Time Between Interruptions
MTBF	Mean Time Between Failures
MTTF	Mean Time To Failure
MTTR	Mean Time To Repair
MUT	Mean Up Time
NA	Nonadaptive
NS2	Network Simulator 2
OSNR	Optical Signal-to-Noise Ratio
OSPF	Open Shortest Path First
pdf	probability density function
PSTN	Public Switched Telephone Network
QoP	Quality of Protection
QoS	Quality of Service
RESTART	Repetitive Simulation Trials After Reaching Thresholds
RON	Resilient Overlay Network
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
SLA	Service Level Agreement
SP	Static priorities
SPP	Shared Path Protection
SRLG	Shared Risk Link Group
TE	Traffic Engineering
WDM	Wavelength Division Multiplexing

List of Papers

Publications Included in the Thesis

These papers are included as Part II of this thesis. Note that some of the papers have been subject to minor editorial changes since their publication.

- PAPER A:
  Piotr Chołda, Anders Mykkeltveit, Bjarne E. Helvik, Otto J. Wittner, Andrzej Jajszczyk. *A Survey of Resilience Differentiation Frameworks in Communication Networks*. IEEE Communications Surveys and Tutorials. Vol. 9, No. 4, Oct, 2007.

- PAPER B:
  Anders Mykkeltveit, Bjarne E. Helvik. *Provision of Connection-Specific Availability Guarantees in Communication Networks*. Proceedings of the 6th International Workshop on Design of Reliable Communication Networks (DRCN). La Rochelle, France, Oct, 2007.

- PAPER C:
  Anders Mykkeltveit, Bjarne E. Helvik. *Comparison of Schemes for Provision of Differentiated Availability-guaranteed Services Using Dedicated Protection*. Proceedings of the Seventh International Conference on Networking (ICN). Cancun, Mexico, April 2008.

- PAPER D:
  Anders Mykkeltveit, Bjarne E. Helvik. *On Provision of Availability Guarantees Using Shared Protection*. Proceedings of the 12th Conference on Optical Network Design and Modelling (ONDM). Vilanova i la Geltru, Spain, March 2008.

- PAPER E:
  Anders Mykkeltveit, Bjarne E. Helvik. *Application of the* RESTART*/Splitting Technique to Network Resilience Studies in NS2*. Proceedings of the 19th IASTED International Conference on Modelling & Simulation. Quebec City, Canada, May 2008.

- PAPER F:
  Anders Mykkeltveit, Bjarne E. Helvik. *Adaptive Management of Connections*

to Meet Availability Guarantees in SLAs. Accepted at the "IM 2009 Mini-Conference" track at the IFIP/IEEE International Symposium on Integrated Network Management (IM 2009)

Other Papers by the Author

These papers were also prepared while working with this thesis.

- Piotr Chołda, Andrzej Jajszczyk, Bjarne E. Helvik, Anders Mykkeltveit. *Service Differentiation Based on Recovery Methods*. 2nd EuroNGI Workshop on Traffic Engineering, Protection and Restoration. Rome, Italy, April 2005.
  – This paper is an early version of Paper A.

- Anders Mykkeltveit, Bjarne E. Helvik. *Prospects for dependability differentiation in interdomain networks*. 3rd EuroNGI Workshop on Traffic Engineering, Protection and Restoration. Kraków, Poland, May 2006.
  – A revised version of this paper is included as Appendix B.

- Norvald Stol, Harald Øverby, Steinar Bjørnstad, Andreas Kimsås Anders Mykkeltveit. *Differentiated survivability in the OpMiGua Hybrid Optical network*. Proceedings of the 10th Conference on Optical Network design and Modelling (ONDM). Copenhagen, Denmark, May 2006.
  – This paper has a different focus than the work included as Part II of this thesis.

- Anders Mykkeltveit, Bjarne E. Helvik. *Application of the RESTART/Splitting technique in object-oriented simulation*. 6th International Workshop on Rare Event Simulation (RESIM). Bamberg, Germany, Oct 2006.
  – This paper is an early version of Paper E.

- Anders Mykkeltveit, Bjarne E. Helvik. *On provision of guaranteed availability using dedicated protection*. Euro-FGI Workshop on Traffic Engineering, Protection and Restoration. Oslo, Norway, Sep, 2007.
  – This paper is an early version of Paper C.

Part I

THESIS INTRODUCTION

Introduction

People use several services delivered by communication networks. These services have different importance. Examples are:

- Surfing on the web or chatting with family or friends. The usage of these services as leisure activities means that the consequences of service failures are small.
- Emergency voice services such as emergency telephone calls to the police, medical personnel or the fire department. This may be some of the same services as above, but the consequences of service failures may be high.
- Background bulk data transfers such as system backup. Service failures may have small consequences as the transfer may be completed when the service are restored.
- Business-critical financial services such as bank transfers of stock trade. Service failures may result in large economic losses.
- Grid computing applications that require the transfer of enormous amounts of data between different sites to complete an experiment.

The *dependability* of a system is defined as “its ability to deliver service that can be justifiably trusted” [ALRL04]. The services depend on delivery of data between different points in the network. If the delivery of data needed by a service can be justifiably trusted, it is here said that the service is dependable. The examples above illustrate that there may be different requirements for the dependability, depending on the type of service, e.g. voice vs. background bulk data transfer, and the usage of the same service, e.g. emergency voice services vs. leisure voice services.

Historically, different communication networks were built for one or a limited number of services. There was the Public Switched Telephone Network used mainly for voice services, cable networks used for TV, and the Internet providing a best-effort data delivery service. The users did not have much choice for the dependability level of a service. However, the current all-over-IP trend means that many services will be carried over the same (IP) network. A typical example is *triple play* which is a marketing term for the provision of voice, TV and broadband Internet access over a single broadband access connection. This integration is a trend in the backbone networks as well. The trend allows rapid introduction of new services and for reduced cost, but also introduces the possibility of providing services with different dependability over the same networks, i.e. *dependability differentiation* which is the topic of this thesis.

The main part of this thesis, Part II, is a collection of six papers. Part I gives an introduction to the material covered in the papers.

The introduction is organized as follows. The general background for the work is presented in Section 1. Then, issues relevant for dependability differentiation are presented in Section 2. The research goals for the work are presented in Section 3. The research methodology followed is presented in Section 4. The contributions of the thesis are presented in Section 5. Finally, a summary and a conclusion are provided in Section 6 while future and ongoing work is listed in Section 7.

1. Background

In an ideal world, a communication network would be working perfectly at any time. In the real world, this is not the case. Random failures affect the network and may cause service outages. Failures may be physical, like fiber cuts, power outages, fires and earthquakes. Other types of failures may be software failures or failures resulting from unintentional human errors. This section gives a short background for the work presented in this thesis. First, the concept of network dependability is presented in Section 1.1. Then, the main concept addressed in the thesis, dependability differentiation is presented in Section 1.2. Finally, a short overview of recovery mechanisms which are the building blocks used to provide dependability and differentiation is given.

1.1 Network dependability

The *resilience* of a network may be defined as the ability of a network to automatically react to failures and to redirect the traffic from paths affected by failures to alternative, failure-free paths. The enabling mechanisms are denoted *recovery mechanisms*. These mechanisms make use of the redundancy in the network topology to deal with failure situations, and are the keys to providing dependable services.

The core backbone networks are typically built with redundancy to cope with failures, for example by ensuring that there are at least two disjoint paths between any two nodes in the network, and that there is spare capacity in the network to be used in case of failures. This will greatly increase the dependability of the network, compared to the situation without any redundancy. However, there is still a (small) chance that both the paths between a node pair will have failures. There may also not be sufficient capacity to carry all traffic interests under some failure scenarios.

Dependability is an integrating concept. As already mentioned it is related to the trustworthiness of a system. Moreover, dependability is characterized by the threats to the dependability of a system, the means by which dependability is obtained, and the attributes related to the dependability [ALRL04, Hel04]. These attributes are stochastic random variables which quantitatively characterize the dependability of the system. Based on the attributes of dependability from [ALRL04] and the ITU-T

E.800 standard [ITU94], the following concepts, denoted in this thesis as dependability attributes, are important in characterizing the dependability of the network associated with the delivery a service:

Availability The availability attribute is related to the probability that a system is able to deliver a service. For communications networks, both the asymptotic availability and the interval availability may be of interest. The asymptotic availability is the probability that the network is able to deliver a service (according to some requirements) at some point in time in the future when the network is in steady state. The interval availability is the fraction of time in a specified interval the network is able to deliver a given service [ITU94].

Continuity This term is related to the reliability attribute of dependability and is related to the time during which a service is delivered without interruption.

Downtimes This is the duration of the outages, i.e. the time period the system is not able to deliver a service due to a failure.

The remaining attributes of dependability from [ALRL04], maintainability, safety, confidentiality and integrity are not treated in this thesis. Maintainability is somewhat related to downtimes, safety is often not an issue unless the consequences of failures may be catastrophic, while confidentiality and integrity are related to security.

If sufficient information about the behavior of the network elements, i.e. time between failures and repair times, is known, it may be possible to obtain estimates of these attributes and thereby quantify the dependability.

1.2 Dependability differentiation

As illustrated above, people have different expectations or requirements to the dependability associated with the delivery of different services. In communication networks, the dependability of a service is closely related to the cost of providing the service, i.e. high dependability comes at a high cost. Also, the willingness to pay for the quality a service varies with the user. Hence, the different users' and services' requirements and users' different willingness to pay together call for different degrees of dependability.

To take advantage of the potential economic benefit associated with dependability differentiation, the same network must be able to provide different levels of dependability to different users and for different services. Different users generate traffic, i.e. data to be transfered through the network. First, the traffic from the different users and services must be separated and appropriately classified when it enters the network. Then, the recovery mechanisms must differentiate the treatment of the traffic from different customers and services.

Two fundamentally different approaches to dependability differentiation are found in the literature. The first approach is denoted *structural differentiation* and bases the differentiation on the type of recovery mechanism applied. This means that the customer knows what will happen to his traffic in case of a failure. A disadvantage with structural differentiation is that it does not guarantee a certain level on the dependability attributes since a given recovery mechanism does not directly translate to a given attribute. The second approach is denoted *guaranteed differentiation* and is the focus of the work presented in this thesis. With guaranteed differentiation, the recovery methods are selected to provide the service with a statistical guarantee on one or more dependability attributes. The guarantees may be stated in a contract between the customer and the provider, denoted a Service Level Agreement (SLA), or they may be determined by the provider for internal use as a method to keep the customers satisfied.

Technically, guaranteed differentiation aims at giving the network operator the possibility to deliver arbitrary numerical values of the guaranteed attribute(s). However, the operator may choose to present a set of service classes to the users, where each service class is associated with some given numerical attribute value(s). The asymptotic availability may be used as an example. This attribute is often specified as a number of nines, e.g. an availability of 0.99999 is denoted “five nines”. One possible menu of classes may for instance be five nines of availability for the “Gold” class, three nines for the “Silver” class and two nines for the “Bronze” class.

Among the dependability attributes, availability is often perceived as the most important. There may be several reasons for this. Availability may be claimed to have the following properties:

Intuitive The availability is a measure that may be intuitive. An availability of 0.99999 translates to an average unavailability of 5 minutes per year.

Frequently used The availability over the contract period is commonly used in SLAs.

Convenient Availability is a relatively convenient measure to calculate under a set of assumptions. This means that the availability is relatively easy to predict.

A weakness of the availability measure is that it does not include information neither about how often a service is interrupted, nor about the downtimes. An availability of 0.9999 could be obtained by ten outages of 5.2 minutes every year or by a single outage of 10.8 hours every ten years. This may be two situations which are perceived as very different, depending on the service/user.

Note on terminology In Paper A, differentiation has a broader meaning than in this introduction, since other features than availability, continuity and downtimes are included in guaranteed differentiation frameworks. The dependability attributes here are identical to the reliability attributes in Paper A.

1.3 Recovery mechanisms overview

In a failure-free situation, the traffic between two nodes in a network normally follows a certain path through the network. When a failure happens, the recovery mechanisms attempt to recover the data delivery by making use of an alternative path. There exists a large number of recovery mechanisms, and a detailed presentation would be very long and is not the purpose of this introduction. This section provides only a short overview. The discussion in Section 3.2 of Paper A provides more details, while more complete discussions may be found in textbooks such as [VPD04] and [Gro04].

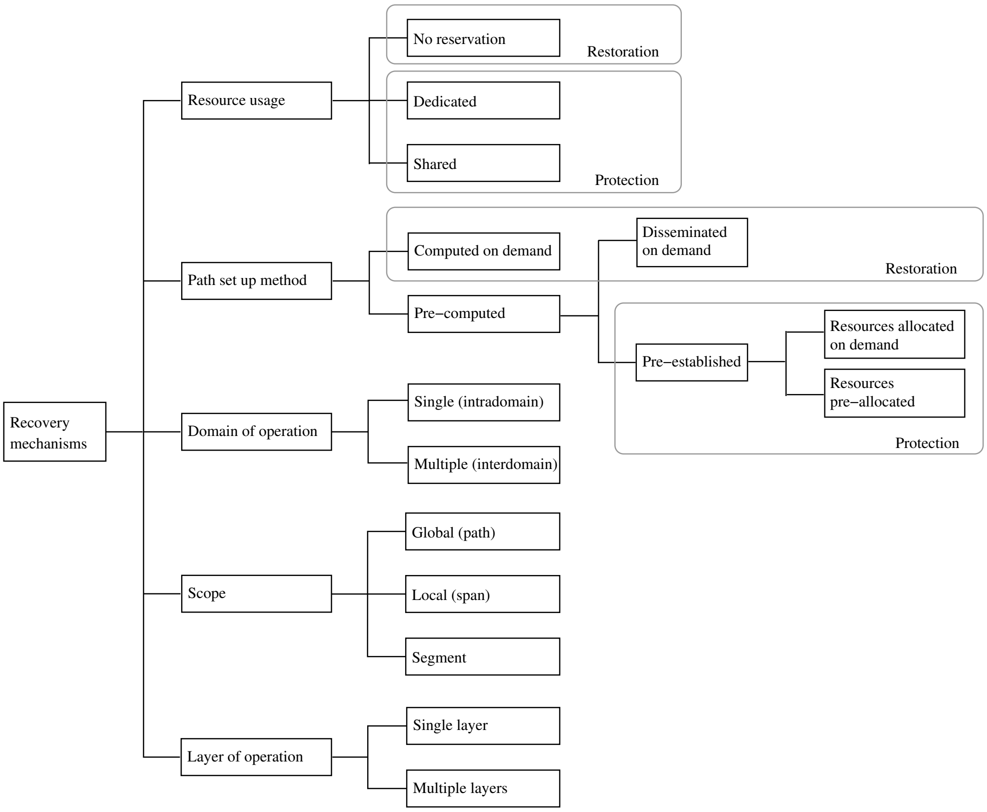

Figure 1. Classification of recovery mechanisms

A partial classification of recovery mechanisms is shown in Figure 1. The recovery mechanisms are classified according to five criteria. First, the resources used for the recovery may be reserved prior to the failure or not. Second, the recovery path may be set up before the failure, or determined after the failure. In the latter case it may not be possible to find such a path if there is not sufficient capacity. A central distinction is made between protection mechanisms and restoration mechanisms. For protection mechanisms, the paths used for recovery is ready for use before the failure, and no signaling is required before the path may be used to transit data [VPD04]. The resources in the recovery paths are with protection either reserved exclusively for one connection (dedicated) or shared with other connections. Third, a recovery mechanism may operate between two nodes in a single domain (intradomain) or it may be effective between nodes in different domains, across one or more domain borders (interdomain). Fourth, the scope of recovery identifies what part of the working path is bypassed by the recovery path. With global (path) recovery, the entire working path between the source and destination nodes is bypassed. With local recovery, only a single link or a single node is bypassed. The solutions in between are denoted segment recovery. Fifth, since most networks have multiple layers, the recovery mechanisms may be distinguished as either single layer mechanisms or multi-layer mechanisms.

The different mechanisms have different properties with regard to what combinations of failures (failure scenarios) they may provide recovery in, as well as the time it takes to recover after a failure. Furthermore, the Quality of Service (QoS) after recovery may be different from the QoS before the failure.

2. Issues in dependability differentiation

This section deals with issues related to the implementation of the functionality of the network that enable delivery of dependability differentiated services. First, the distinction between connection-oriented and connectionless networks is discussed in the context of differentiation in Section 2.1. Then issues in providing end-to-end differentiation are discussed in Section 2.2. Based on the considerations in these two sections, a reference scenario, which is used as a basis for most of the work in this thesis, is presented in Section 2.3. Finally, a discussion of the implementation of a management system for providing differentiation is given in Section 2.4.

2.1 Connection-oriented versus connectionless networks

Communication networks may be divided in two fundamentally different groups; connection-oriented and connectionless networks [Tan03]. When guarantees on services are to be offered, connection-oriented networks seem better suited.

In connection-oriented networks, the data is routed through pre-established virtual connections referred to as working paths. In some networks, capacity for the connections may be reserved, guaranteeing that a service is working properly as long as the working path is up. The dependability of the working path can in many cases be predicted. If this dependability does not meet the requirement for the service it provides, protection mechanisms may be used to increase the dependability by guaranteeing that the connection will be almost instantly restored after failures. Examples of connection-oriented networks are optical Wavelength Division Multiplexing (WDM) networks and Multi-Protocol Label Switching (MPLS) networks. It is foreseen by some that the most likely candidate for the Next Generation Internet is an IP/MPLS over WDM network, i.e. a connection-oriented network using the IP protocol.

In connectionless networks, there are no virtual circuits, and other recovery mechanisms than those used for connection-oriented networks are used. The connectionless networks divide the data transfered through the network into packets which are routed independently to the destination. An example of connectionless networks are the pure IP networks. They are based on the best-effort principle where all the traffic receives services without any guarantees. Today, connectionless networks are for instance IP-based networks using intradomain routing protocols such as Open Shortest Path

First (OSPF) and the currently deployed interdomain routing protocol Border Gateway Protocol (BGP). In the last years, there have been many advances in research on improving the dependability of connectionless networks, for instance the different proposals for IP Fast Reroute, see for instance [SB08]. It is, however, not clear how different service classes may be differentiated using these methods, and whether it will be possible to provide guarantees on dependability attributes. Therefore, in this thesis, differentiation is mainly approached in connection-oriented networks.

2.2 End-to-end differentiation

The Internet is a global network which is composed of thousands of smaller networks. Each of these networks is managed and operated by a company or organization and is denoted a domain or an Autonomous System (AS). Network operators have full control over the resources in their own networks, and are free to engineer their network as they like. Hence, when dependability differentiation is provided inside the borders of a domain, i.e. intradomain differentiation, the challenges are mostly technical. Many connections, however, span multiple domains. Hence, there is also a need for differentiated dependability at the interdomain level. At this level, there are additional challenges due to the reluctance of network operators to share information of their network internals and to allow external management of their networks.

For true end-to-end differentiation, access networks should also be taken into account. The access networks may in some cases be performance and dependability bottlenecks. However, with multiple access technologies involved, possibly from different providers, this may be an interesting research field. It is, however, most likely that differentiation will first be introduced as an intradomain mechanism adapted to the network of each operator. Only after the ASes support differentiation internally, migration to a scenario with real end-to-end differentiation may start. Therefore, in this thesis, and in most of the literature surveyed in Paper A, differentiation is approached from an intradomain perspective. A discussion about the prospects for interdomain differentiation is, however, included as Appendix B.

Figure 2 shows the different areas in end-to-end differentiation. The smaller networks provide access to end-customers and the operators of these networks are denoted Internet Service Providers (ISPs). The two networks B1 and B2 are backbone networks providing transport services. The customers of these networks may be other backbone networks, smaller ISPs networks or even large companies or organizations. The backbone networks are typically large and transport great amounts of data. They are the main focus of much of the work in this thesis. Note that the term end-to-end is sometimes used to denote the provision of connections between an ingress and egress router in a domain when only a single network is considered.

2.3 Reference scenario

Based on the above considerations, this section presents a scenario illustrating a system providing differentiated dependability. The scenario is typical for the majority

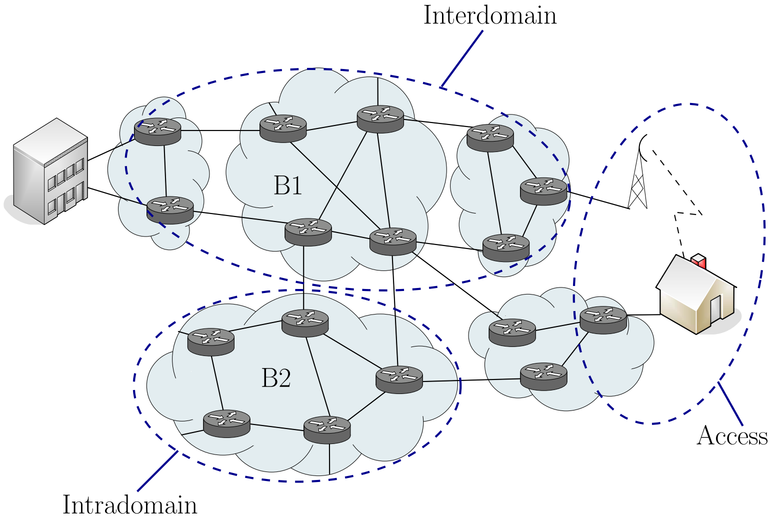

Figure 2. Research areas in end-to-end differentiation.

of the work in this thesis. More specifically, this scenario is assumed for most of the papers reviewed in Paper A, and used in Papers B, C, D and F.

This scenario considers a connection-oriented network where different choices for the recovery mechanisms are associated with the connections. A conceptual model of a the functionality and information needed to provide differentiated dependability is shown in Figure 3.

In this thesis, the term connection management is used to denote the functionality needed to handle connections with dependability requirements. Connection requests and releases arrive during the operation of the system, i.e. it is an on-line connection management system. Each connection is to be set up between two nodes, denoted the source and destination, and it has a specified bandwidth requirement. In addition, a dependability requirement is associated with each connection request. The network supports different protection options which can be used when establishing a connection. Connections may be unprotected, protected with a dedicated backup path or protected by a shared backup path. There may be many possible ways to establish the connection with a given protection arrangement using different links and nodes in the working and eventual backup paths. All these options are the candidate arrangements for a connection, and all have a certain dependability and cost. The task of the connection management system is to choose an appropriate arrangement for each arriving request if such an arrangement exists.

The choice should be "optimized" according to some criteria such as meeting the requested dependability attributes while at the same time using the minimum number of resources. Depending on the choice of connection arrangement among the candidates, the problem of finding the optimal solution may be NP-hard. To let the service be established with an acceptable delay to the users, the arrangement must be

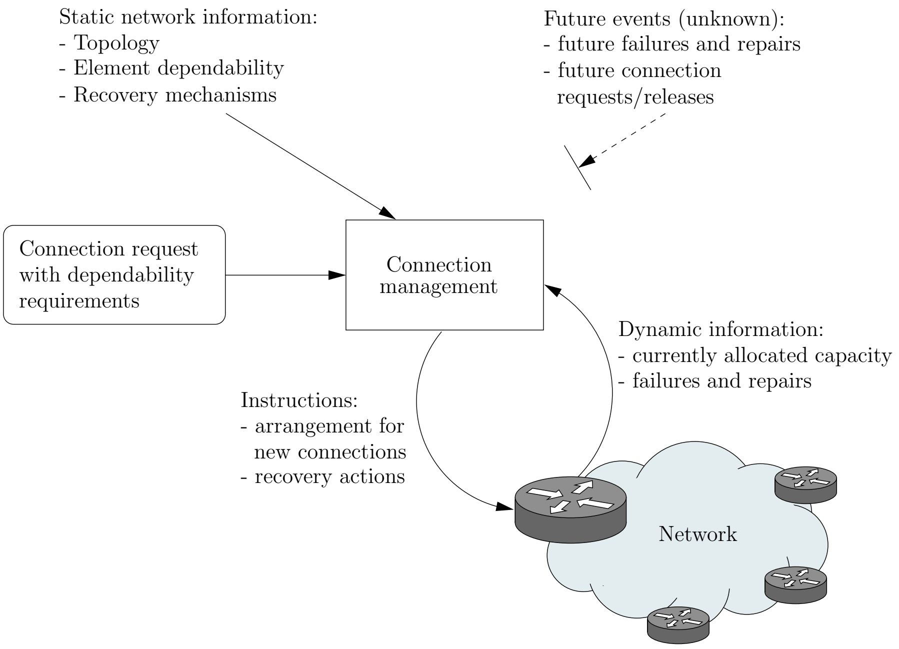
Static network information:
- Topology
- Element dependability
- Recovery mechanisms
Future events (unknown):
- future failures and repairs
- future connection requests/releases
Connection request with dependability requirements
Connection management
Dynamic information:
- currently allocated capacity
- failures and repairs
Instructions:
- arrangement for new connections
- recovery actions
Network

Figure 3. Conceptual model of management functions and the information used to provide differentiated dependability.

identified in a short time. This means that the algorithm used needs to be of limited complexity.

The system uses available information to identify a good connection arrangement. Some of this information is relatively static, such as information on the network topology, the capacity of links and the dependability of the network elements. Other static information is related to which type of recovery mechanisms may be used for the connection arrangements. Some of the information may be dynamic, i.e. subject to change on a shorter timescale, for instance the current traffic situation. As a minimum, the management system needs information on spare capacity so the candidate arrangements may be determined. When the connection arrangement has been identified, the connection management system instructs the network to set up a new connection. When failures affect the connection, the recovery mechanisms attempt to recover the connection by moving the traffic to an alternative path included in the arrangement.

Some of the information needed to find a truly optimal connection arrangement may be inaccessible to the management system. Typical examples of this is knowledge about the future. The information about the dependability of the network is typically only stochastic, and there is no way of knowing when or where the next failure will occur. Normally, exact information of how long a repair takes will not be accessible either. Furthermore, the future connection requests and releases may also be unknown. This may mean that serving one connection in a seemingly optimal way may actually make serving a later request impossible without changing the allocation of already established connections.

In summary, the connection management functions must be able to establish connections meeting differentiated requirements on the dependability attributes with a short delay and at a low cost in terms of bandwidth.

2.4 Implementation of management system

This section considers the system implementing the functionality needed for connection management. The system may be implemented as a centralized or a distributed system. With a centralized system, a single management system has all available information and full control of the entire network. Since managing networks of some size is very complex, the system needs very high performance to avoid being a performance bottleneck. Additionally, the centralized management system is a single point of failure, which means that a failure in the system will mean that the network is unmanaged. With a distributed solution, there are many autonomous connection management entities that manage part of the network. This may be done by having one management system for each ingress router, controlling connections which originate in that router. Distributed solutions may be more scalable than centralized solutions, and may also be more robust. There are, however, additional challenges associated with a distributed solution. The information is typically only local, meaning that it may be more difficult to find optimal solutions due to the lack of global knowledge [RS01, VPD04].

3. Research goals

As presented in the two previous sections, dependability differentiation in communication networks is a rather broad topic. In the work presented in this thesis, three goals have been pursued. The first goal, new methods, is the most intuitive. During the initial work, two other goals were discovered. First, a survey of the state of the art in differentiation was missing. Second, the problem of measuring the dependability of networks by simulation calls for new approaches.

3.1 New methods

New methods are needed to enable delivery of differentiated dependability. The term method is to be understood in a broad sense as any systematic plan or procedure that can be used to provide differentiated dependability. The key issues or ideas of the methods may be techniques, schemes, rules or policies used for providing the services. Although some methods have been proposed in the literature, this research area is still open. Different methods have different strengths and weaknesses with respect to cost, flexibility, complexity and performance, and different methods target different types of networks and different services. The work presented in this thesis focuses on connection-oriented backbone networks. The key design goal pursued was to develop simple but flexible and effective methods.

3.2 State of the art survey, classification and evaluation

When the work with this thesis started, a significant number of publications which could be classified as related to differentiated dependability existed, but no survey had ever been done. Such a survey is needed to be able to classify, group, compare, discuss, and finally build further on results and experience.

3.3 Simulation tools to analyze differentiated dependability

It is relatively difficult to quantitatively measure the dependability of the services delivered by a network. The best option may be to simulate the network and analyze the results. This is due to the complexity of the problem which leaves mathematical analysis infeasible. Also, no real networks were available for implementing the proposed methods.

Service outages are often the result of multiple simultaneously failed links or nodes. These situations are *rare events*, since the simulations will most of the time alternate between situations with zero or a single failure. Hence, simulation of a sufficient number of outages to obtain firm estimates of the dependability attributes may become very time-consuming.

Different techniques to speed up simulations exist, see for instance [Hee95]. Among these are rare-event simulation techniques which aim at driving the simulations toward the rare events of interest. One approach followed was to see if it was possible to enhance a commonly used network simulator, Network Simulator 2 (NS2), with the capability to use a rare-event simulation technique to speed up the simulation of outages. The goal was to be able to use NS2 for simulations of attributes of dependability in the context of differentiation. The result of this work was mainly Paper E which showed that the RESTART/Splitting method could be applied to NS2. This did, however, not give sufficient reduction in simulation time to use the tool as a basis for evaluation of differentiated dependability. A lesson from this work was that it was necessary to simplify the models in order not to require excessively long simulations. Often the only possibility is to use a purpose-built simulator tool which models the details needed and nothing more. This approach is taken for the simulator used in Paper F which models a set of disjoint path between two nodes in a network. In the work presented in papers B, C and D, the problem of long simulation times was omitted because only the estimated dependability was measured.

4. Research methodology

To approach different goals, different research methodologies had to be followed. The two first goals, new methods and simulation tools (pursued in paper B-F), follow a traditional scientific methodology starting with a hypothesis which is tested before the results from the test are validated. The last research goal, state of the art survey and classification followed the methodology outlined below.

4.1 State of the art survey, classification and evaluation

The work towards this goal consisted in a search for published literature followed by the identification of the contribution of each work. Based on the gathered work, a classification of the literature was developed. This classification was based on what was addressed in the literature and what seemed to be missing. Finally, qualitative conclusions could be drawn from the state of the art.

4.2 New methods and simulation tools

The traditional research methodology starts with the formulation of a *working hypothesis*. Here, the hypotheses are formulated as descriptions of new methods that are intended to be able to provide differentiated dependability at a low cost in terms of resource usage. The specific hypotheses/ideas of in papers B-F are indicated with bold letters in Table 1 on page 18.

The second step in the research methodology is the *hypothesis testing*. In this thesis, all proposed methods have been evaluated by simulation. It should be noted that there are limitations to which conclusions may be drawn based on simulation results. One limitation is that only a few scenarios may be explored by a simulation experiment, but also other problems exist, see [FP01] for a discussion.

The third step in the research methodology is the *result validation*. It is known that most software contains errors. Therefore, testing, tracing and feasibility analysis of the results are necessary. To give high statistical confidence in the results, the simulations must be repeated where appropriate. In papers D-F, estimates for mean values are given with confidence intervals. The simulations in this thesis are based on reference topologies from [SND] and use the best practice models of failures and repairs from the literature.

5. Contributions

This section presents the main contributions of each of the six papers included in Part II of this thesis. Figure 4 shows which topics related to dependability differentiation are addressed by the different papers. Note that, although many of the topics are addressed in the literature surveyed in Paper A, this does not mean that the topic is fully explored.

For Papers B to F, the thesis author had the original idea, performed the work and wrote the papers, under supervision of and in cooperation with Professor Helvik. Paper A is a joint work between AGH University of Science and Technology, Kraków, Poland and Q2S. The contributions to this paper by the thesis author are mentioned specifically below.

5.1 Contributions of the papers

This section reviews the papers that constitute Part II of this thesis and identifies the main contributions of each paper.

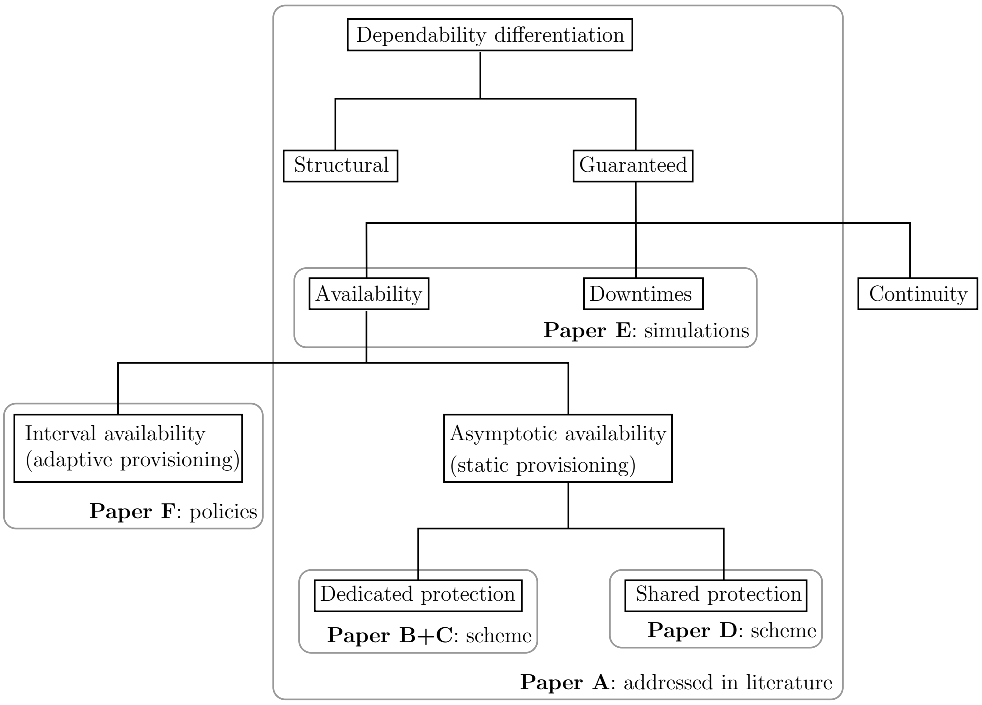

Figure 4. Topics addressed in the different papers of this thesis.

PAPER A

A Survey of Resilience Differentiation Frameworks in Communication Networks

This paper is the first extensive literature survey on the topic of resilience differentiation. More than 100 publications are included. When reviewing the literature, and by creating graphs of the citations between the papers, it was found that many authors were working on similar problems seemingly without being aware of each other's work. This indicates the need for this type of survey. To make it easier to get an overview of the state of the art, the publications were grouped according to a proposed classification scheme. Based on the reviewed literature, the paper discusses the state of the art and identifies a number of issues that must be addressed before differentiation can be adopted in real networks.

Piotr Chołda originally initiated the work on the survey. All five authors participated in discussions of the totality of the literature. The proposed literature classification scheme, assessment of state of the art and challenges were a result of these discussions. The literature study in itself, i.e. reading the papers and classifying them, was done by Chołda and the thesis author. The initial version of the paper was mostly written by Chołda and the thesis author with inputs and editorial changes from the other authors. A first, much shorter, version of the survey was presented at a workshop [CJHM05].

PAPER B

Provision of Connection-Specific Availability Guarantees
in Communication Networks

The paper proposes a new protection arrangement denoted a protection pattern which is more flexible than link and path protection. This flexibility can be used to better tailor the dependability and the cost of the different connections. In this paper, protection patterns are used to find arrangements with guarantees on the asymptotic availability. The Inclusion-Exclusion method is used to calculate this availability. To find near-optimal protection patterns meeting availability guarantees, a distributed management system is proposed. This system is a new version of an emergent behavior-based system denoted the Cross-Entropy Ant System (CEAS). The system is stochastic and completely distributed. The system is implemented as an extension to NS2, and simulation results demonstrate that it is able to find feasible solutions in a relatively short time.

PAPER C

Comparison of Schemes for Provision of Differentiated Availability-guaranteed Services Using Dedicated Protection

The main contribution of this paper is to compare the scheme for finding connection arrangements based on dedicated protection proposed in Paper B with two other schemes proposed in the literature. This paper addresses performance comparisons of different differentiation schemes which were found to be missing in the literature according to the findings in Paper A. The schemes are quantitatively compared with regards to their ability to meet availability requirements and their resource usage in a simulation study. While the system proposed in Paper B is distributed and based on stochastic optimization, the two other schemes are based on a centralized management system executing deterministic algorithms. The results indicate that the scheme proposed in Paper B performs well compared to the other two systems. Moreover, it performs similar to a system where the two other schemes are combined by using the best solution they can find.

PAPER D

On Provision of Availability Guarantees Using Shared Protection

Methods for providing availability guarantees based on dedicated protection, like the ones dealt with in papers B and C, are popular candidates for use in real networks due to their simplicity. Shared protection is potentially much more resource-efficient, but must be restricted to control the availability of the connections.

The main contribution of this paper is the proposal of a scheme for controlling the sharing of backup resources to ensure guarantees on asymptotic availability requirements. The scheme uses rules based on temporal priorities and preemption to control the availability. The performance of the scheme is compared with a scheme based on dedicated protection and a second proposed scheme for shared protection denoted the ultra-conservative scheme. The simulation results show that the preemptive sharing scheme performs best for high availability requirements, while the two schemes based on shared protection both perform better than the dedicated protection-based scheme.

PAPER E

Application of the RESTART/*Splitting Technique to Network Resilience Studies in NS2*

The original motivation of this work was to create a simulator that could be used to investigate the dependability attributes for individual connections or traffic flow in the commonly used NS2 simulator. It was understood that direct simulations would be too consuming in terms of simulation time to obtain useful results, and the hope was that some rare-event simulation method could be used. The two similar methods RESTART and Splitting seemed most suited for the task. A crossover-variant of the two techniques was used to enable measurements of the availability and the downtime durations. NS2 was selected as simulator tool since it is widely used and provides various protocol implementations. Working with the implementation revealed that NS2 had not been used for measuring these attributes, and this required extensions to the simulator. The main contribution of this paper is the application of the RESTART/Splitting technique to simulate network dependability in a large and complex network simulator. Experiments show that the RESTART/Splitting method gave reduced simulation times in scenarios where more than two failures were necessary to observe outages. However, the gain is moderate, and the simulation time needed is still very long. The added complexity of the approach when it comes to analyzing the output data counter-balance this gain. In conclusion, using the method for measuring dependability attributes is not unconditionally recommended.

PAPER F

Adaptive Management of Connections to Meet Availability Guarantees in SLAs

While papers B to D deal with schemes that can provide a guaranteed asymptotic availability, this paper instead focuses on the interval availability. This is motivated by the fact that SLAs typically involve guarantees on the availability over a limited contract period. The paper proposes to use an adaptive management approach to meet different customers' requirements on the interval availability guaranteed in their SLAs. This means that when there is reduced capacity in the network due to failures, the selection of which connections are affected is based on the goal of maximizing the compliance with the customers' SLAs. To the authors' knowledge, this is a new idea. The paper presents a set of policies for managing the connections and compares the results with respect to the risk, i.e. the probability of violating the different SLAs. The adaptive policies are compared to traditional static policies. Simulation of a relatively simple model shows that adaptive management may in many cases significantly reduce this risk.

5.2 Summary and evaluation of the contributions

An overview of the contributions of the papers and which research goals they contribute to is given in Table 1.

Research goal	New methods (key issues/ideas)	State of the art survey, classification and evaluation	Simulation tools
Paper A		**Broad and extensive survey**	
Paper B	**Protection patterns, Distributed scheme**		NS2-based tool to evaluate proposed scheme
Paper C		Identify candidate schemes	**Purpose-built tool to evaluate schemes**
Paper D	**Rules for shared protection**		Purpose-built tool to evaluate proposed scheme
Paper E			**Enhanced version of NS2 which uses the RESTART/Splitting technique**
Paper F	**Adaptive management, Propose policies**		Purpose-built tool to evaluate risk
Appendix B		Survey in interdomain context	

Table 1. Summary of research goals and paper contributions. Main contributions in **bold** letters.

With regard to the first research goal, new methods, Papers B and D propose methods for providing availability guarantees. These methods are based on the assumption that the failures and repairs of the links are stochastically independent. These assumptions are necessary in order to predict the availability of the configurations of the connections. These assumptions may or may not be realistic. These assumptions are not made for papers E and F, since the papers measure the availability delivered instead of giving guarantees. Papers B and D propose new approaches to a problem which is already addressed in the literature, i.e. provisioning of connections with guarantees on the asymptotic availability. The adaptive management strategy proposed in Paper F may be seen as the most innovative among the papers in this thesis that propose new methods.

6. Summary and conclusion

Different users and services have different dependability requirements and different willingness to pay for the services. Differentiation based on dependability attributes can be provided by assigning different recovery mechanisms to the traffic from different customers and services. Since offering a high level of dependability is associated with a high cost, schemes that can provide the requested dependability at a low cost are interesting.

The work presented in this thesis has focused on connection-oriented backbone networks, which seems to be the most likely place to start when differentiation is to be applied to the global communication infrastructure.

This thesis contributes to different aspects of dependability differentiation in six papers. First, the literature has been surveyed and grouped according to a proposed classification. Second, a speed-up technique for simulation has been modified to reduce the time needed to measure dependability attributes by simulation. Third, novel methods for providing services with guarantees on availability have been proposed. Two methods are designed to offer statistical guarantees on the asymptotic availability. The first method is to use a proposed flexible arrangement for dedicated protection, denoted a protection pattern. A distributed system has been proposed for finding near optimal protection patterns. The proposed system is compared to other methods proposed in the literature through a simulation study. The second method is based on a proposed sharing rule for Shared Path Protection. It is shown how this sharing rule can be used to establish connections using distributed signaling. Simulations show that the proposed scheme performs better than a scheme based on Dedicated Path Protection and an alternative ultra-conservative scheme. Finally, it is proposed to use adaptive management strategies to meet interval availability guarantees given in SLAs. Different policies for adaptive management are proposed and compared by simulation. The results show that adaptive management may in many cases significantly reduce the risk of violating the SLAs.

7. Ongoing and future work

There are several options for future work. There are options for direct continuations and extensions of the work in this thesis, as well as issues related to the introduction of dependability differentiation in general.

7.1 Continuation of the work in this thesis

There are options for further studies of the papers B, D and F that propose new methods to provide differentiation. Some of the most important are as follows.

- The protection pattern concept from Paper B is studied for dedicated protection only. In Paper D, shared protection is used. It would be interesting to combine the flexibilities of the protection pattern with the sharing rules proposed in Paper D. Finding near-optimal configurations in this respect would be very computationally challenging, but it may be possible by extending the emergent behavior-based system proposed in Paper B.

- The results from Paper D illustrate that fairness among connections with different requirements should be addressed. It is difficult to establish connections with high requirements since they typically require more resources than connections with low requirements. This is a general problem associated with many availability-aware schemes. Finding methods to ensure that connections with high requirements are successfully established should therefore be investigated. This increased success ratio may come at the expense of connections with lower requirements.

- Paper F proposes to use adaptive provisioning to increase compliance with SLAs. The case studied is relatively simple. The complexity of adaptively managed networks where the capacity is shared between connections with different source and destination nodes could be assessed. If this is found to be feasible, different management strategies could be investigated.

- The schemes proposed are all based on the availability attribute. As mentioned, the availability attribute alone may not be a sufficient indicator of the dependability for all services. Therefore, the differentiation should be able to provide guarantees on the other attributes of dependability in cases where this is required. To the thesis author's knowledge, there has not been any work that uses the continuity attribute as a basis for differentiation. This attribute may be relevant for a number of applications. There is some work in the literature using multiple attributes as a basis for differentiation. It should be further investigated if more than one attribute is necessary, and if so, which attributes should be used. In general, it may be useful to use all three attributes discussed in this introduction. Moreover, it should be investigated if the schemes proposed in this thesis could be applied to continuity-based differentiation.

7.2 General issues related to dependability differentiation

As identified in Paper A and also in Appendix B, there are a number of challenges before the types of differentiation discussed here will be ready to be deployed in operational networks and before differentiated services may generally be available to the users. Some of the most important are as follows.

- The proposed frameworks in this thesis and most of the frameworks proposed in the literature are relatively limited. They relate to a single domain and a single network layer. There is a need for work taking into account access, interdomain and multi-layer networks.

- The approach taken in this thesis, as well as in many other works found in the literature, is to provide communication with guarantees on dependability attributes. A common consensus of which attribute(s) should be the basis for the differentiation is needed.

- There is need for adequate measurement techniques for determining if the requirements of a customer have been met. If the attributes are long-term statistical measures, this may be challenging since the measurement period must be very long.

- There is need for extended SLAs which specify reaction patterns when the contracts are not met. This issue is complicated when multiple operators are involved in the provisioning of a service.

Part II

INCLUDED PAPERS

PAPER A

A Survey of Resilience Differentiation Frameworks in Communication Networks

Piotr Chołda, Anders Mykkeltveit, Bjarne E. Helvik, Otto J. Wittner and Andrzej Jajszczyk

IEEE Communications Surveys and Tutorials

Vol. 9, No. 4, Oct, 2007

PAPER B

Provision of Connection-Specific Availability Guarantees in Communication Networks

Anders Mykkeltveit and Bjarne E. Helvik

Proceedings of the 6th International Workshop on Design of Reliable Communication Networks (DRCN)

La Rochelle, France, Oct, 2007

PAPER C

Comparison of Schemes for Provision of Differentiated Availability-guaranteed Services Using Dedicated Protection

Anders Mykkeltveit and Bjarne E. Helvik

Proceedings of the Seventh International Conference on Networking (ICN)

Cancun, Mexico, April 2008

Is not included due to copyright

PAPER D

On Provision of Availability Guarantees Using Shared Protection

Anders Mykkeltveit and Bjarne E. Helvik

Proceedings of the 12th Conference on Optical Network Design and Modelling (ONDM)

Vilanova i la Geltru, Spain, March 2008

PAPER E

Application of the Restart/Splitting Technique to Network Resilience Studies in NS2

Anders Mykkeltveit and Bjarne E. Helvik

Proceedings of the 19th IASTED International Conference on Modelling & Simulation

Quebec City, Canada, May 2008

Is not included due to copyright

PAPER F

Adaptive Management of Connections to Meet Availability Guarantees in SLAs

Anders Mykkeltveit and Bjarne E. Helvik

Accepted at the “IM 2009 Mini-Conference” track at the IFIP/IEEE International Symposium on Integrated Network Management (IM 2009)

New York, NY, USA, June 2009

Is not included due to copyright

Part III

THESIS APPENDIX

APPENDIX A: ADDITIONAL RESULTS TO PAPER F

Anders Mykkeltveit
Centre for Quantifiable Quality of Service in Communication Systems
Norwegian University of Science and Technology,
Trondheim, Norway
{mykkeltv, bjarne}@q2s.ntnu.no

Abstract This appendix provides additional results to Paper F. It can be seen as an extended version of Section 5 in Paper F.

1. Introduction

The effect of adaptive management is studied in three scenarios. In the first two scenarios all the customers have equal requirements, and the effect of the different policies for management of connections from the same class are investigated. In the first scenario, the connections are partially protected, while in the second scenario the connections are fully protected. In the third scenario there are two classes of customers with different requirements to study how the two classes affect each other under different management policies.

In the scenarios, the n paths used to provide the service are disjoint and have independent failure and repair processes. Since the number of disjoint paths is usually limited in typical backbone networks, the cases $n = 2$ and $n = 3$ are regarded. Furthermore, the equal bandwidth case is studied, i.e. $\forall i, B_i = B$, and the bandwidth of all paths is set to $4B$, i.e. all paths can serve up to four connections.

Two different models of network failure and repair processes are considered. The first model is based on parameters from an optical network and the second is based on parameters from an IP network which are relevant for connections in IP/MPLS networks.

The parameters for failures and repairs in the optical network model are based on [VCD$^+$05] and are shown in Table 1.

Very different failure and repair times are found in IP networks. Based on [MIB$^+$04], a log-normal distribution was selected for the path repair times. The time between failures was selected to give the same asymptotic availability as for the case with optical networks. The parameters for failures and repairs in the IP/MPLS network model are shown in Table 2

Parameter	Value
No. of paths, n	3 and 5
Contract period, T	1 and 5 years
Transient period	1 year
MTTR	ned(12h)
MTTF	ned(5256h = 219 days)
Path availability	0.9977

Table 1. Parameters for optical network scenario

Parameter	Value
No. of paths, n	3 and 5
Contract period, T	1 and 5 years
Transient period	5 years
MTTR	Lognormal with mean 0.2h
MTTF	ned(88h)
Path availability	0.9977

Table 2. Parameters for IP/MPLS network scenario

For both cases, the probability for a given number of failed paths at a given point in time will be the same. The numerical values for i simultaneously failed paths, denoted p_i are shown in Table 3.

n	p_0	p_1	p_2	p_3
2	0.99545	0.004545	5.18885E-6	
3	0.99318	0.006802	1.55311E-5	1.18197E-8
5	0.98866	0.011286	5.15347E-5	1.17659E-7

Table 3. Probability for different number of simultaneously failed paths

It is assumed that path failures can be detected and the connections rearranged in a negligible time. In order to start the observation in a steady state, a transient period was used. In the plots below, the left plots show the results for 1-year contract periods, while the right plots show the results for 5-year contract periods.

The observations for each A_c is based on simulation of 10000 intervals. Confidence intervals are for all plots small and are omitted for sake of readability.

2. Partially protected connections

In this case there are two disjoint paths between the source and destination, i.e. $n = 2$. Not all the connections can be carried when one path has failed, and therefore this case is denoted partially protected. Figure 1 shows a scenario where $z = 6$ customers are using two paths. In this case, two customers will be down when there is one failed path, i.e. 33% of the connections are not protected.

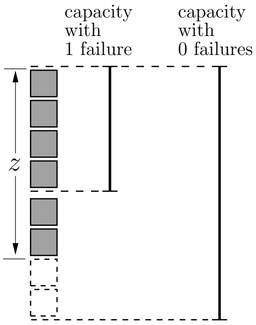

Figure 1. Assignment of connections to paths in partially protected scenario

2.1 Optical network parameters

Figure 2 shows the cumulative distribution function (CDF) of interval unavailability for 6 customers with optical network parameters. This is the distribution of the interval availability for all connections independently of the individual random priority r_i. With the parameters given in this study, the distribution of interval availability is independent of A_c, the availability specified in the Service Level Agreement (SLA), for all policies except Adaptive with preemption (AP). In this particular case, $A_c = 0.998$. The figure shows that there is a significant difference between all four policies for both values of T. There is a significant jump for the AP policy at the requirement for $T = 1$ year, while the jump is much smaller for $T = 5$ years. This indicates that reconfiguration after budget expiration is used more rarely for the longer contract period. Among the others, the Static priorities (SP) policy performs best at the specified A_c. The vertical lines in the plot indicate the average interval unavailabilities, i.e. $E(U_i)$.

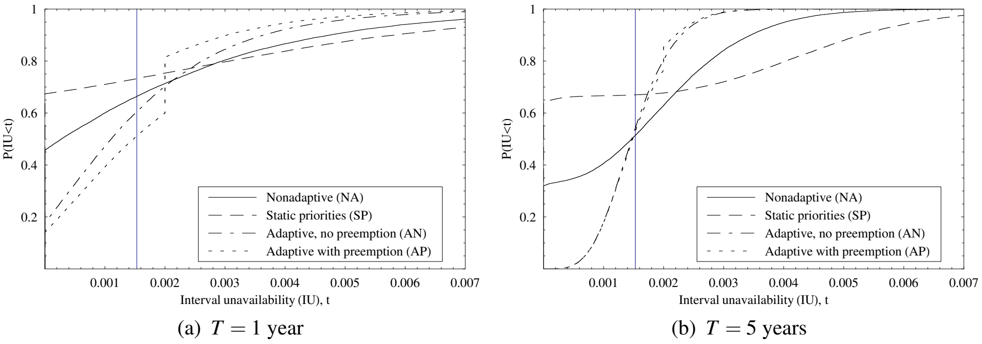

(a) $T = 1$ year

(b) $T = 5$ years

Figure 2. Cumulative distributions of interval unavailability, optical network parameters, $n = 2$

It is anticipated that network operators have as their objective to offer guarantees that are met most of the time, i.e. having a low risk. The risk for different levels of U_i is shown in Figure 3. The figure shows that the AP policy has lowest risk for all $U_c \geq 0.002$.

The risk for different levels of U_i is shown in Figure 3. The figure shows that the AP policy has lowest risk for all $U_c \geq 0.002$ for both values of T. As could be expected, the risk tends faster to zero for the longer contract period. The AN policy performs second best. For $T = 5$ years, the difference to the AP policy is small.

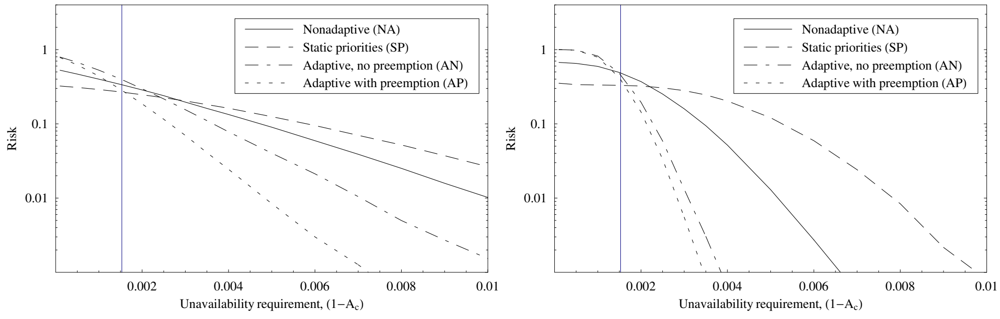

(a) $T = 1$ year

(b) $T = 5$ years

Figure 3. Risk for optical network parameters, $n = 2$, and 6 customers

2.1.1 Effect of traffic load

The risk ratios between both the adaptive policies and the NA policy for $n = 2$ are shown in Figure 4. The two upper plots compare the AN and NA policy, i.e. $\frac{\text{Risk for NA}}{\text{Risk for AN}}$, while the lower plots compare the AP and NA policy. The number of connections, i.e. the value z of Figure 1 is varied. A given z corresponds to a certain percentage of unprotected connections, this percentage is given in parenthesis in the legends of Figure 4. It is seen that a low percentage gives much higher risk ratio when the AP policy is used. Furthermore, the AP policy performs better than the AN policy for $T = 1$ year, while the two schemes perform equally well for $T = 5$ years.

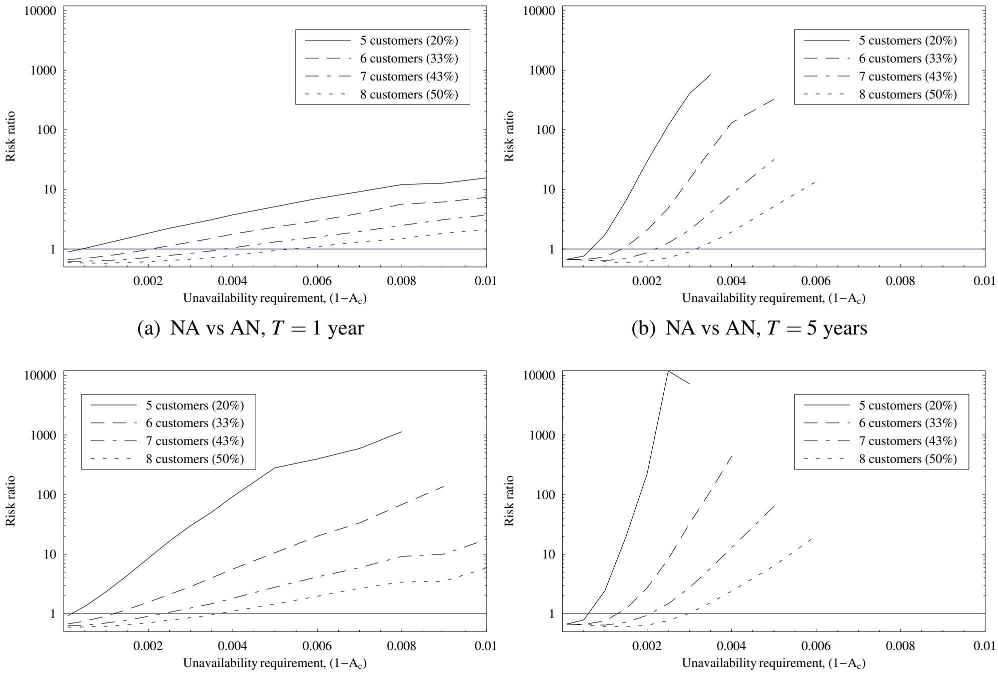

(a) NA vs AN, $T = 1$ year

(b) NA vs AN, $T = 5$ years

(c) NA vs AP, $T = 1$ year

(d) NA vs AP, $T = 5$ years

Figure 4. Risk ratio for optical network parameters, $n = 2$ for variable traffic loads.

2.1.2 Special case: Differing path availabilities

The cases above are all for the case when the availability of the paths are i.i.d. To assess the effect of different path failure probabilities, a case was constructed for $n = 2$ where path π_1 and path π_2 have different failure rates. Let λ denote the "usual" failure intensity, while λ_1 and λ_2 denote the failure intensities of path 1 and 2 respectively. It is required that $\lambda_1 + \lambda_2 = \lambda$. Now let $\lambda_1 = k\lambda_2$. This gives the new failure rates

$$\lambda_1 = \frac{2\lambda}{k+1} \text{ and } \lambda_2 = \frac{2k\lambda}{k+1}$$

Note that this means that the state probabilities (probability of 0, 1 and 2 failures) will all be altered slightly from the original model with identical path availabilities.

Figure 5 shows the effect of having a bad link with $k = 5$ times the failure intensity of the other link.

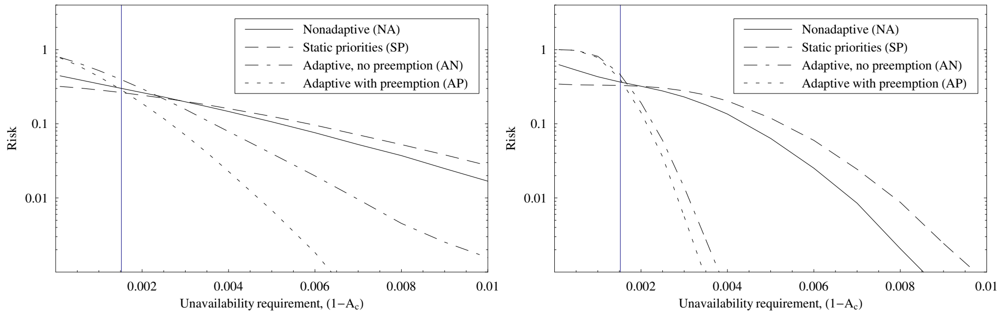

(a) $T = 1$ year (b) $T = 5$ years

Figure 5. Risk for optical network parameters, $n = 2$ with one bad path.

2.2 IP/MPLS network parameters

Figure 6 shows the CDF of interval unavailability for $z = 6$ customers with IP/MPLS network parameters. It is seen that the SP policy performs worse now than for the optical network parameters. The difference between the AP and SP policy is very small here, but this may be due to the strict U_c specified here which is smaller than $E(U_i)$.

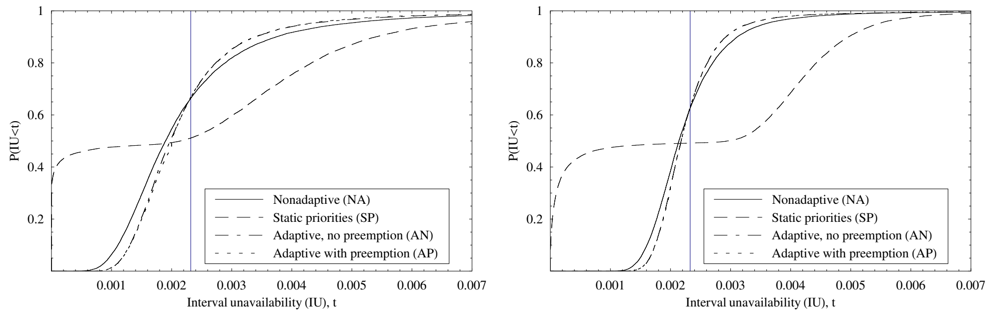

(a) $T = 1$ year (b) $T = 5$ years

Figure 6. Cumulative distributions of interval unavailability, IP/MPLS network parameters, $n = 2$

Figure 7 shows the risk for $n = 2$. The gain from the AP policy compared with the Adaptive, no preemption (AN) policy is smaller than for the optical network parameters. This is due to the short outages observed which often do not need the preemption mechanism since the path is restored before the budget is expired.

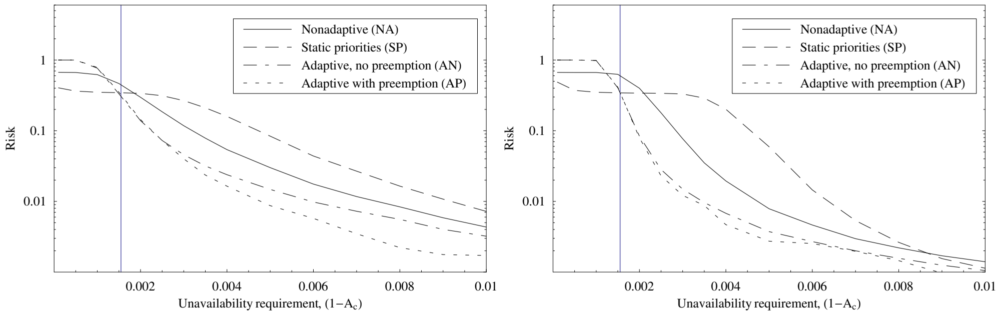

(a) $T = 1$ year

(b) $T = 5$ years

Figure 7. Risk for IP/MPLS network parameters, $n = 2$

3. Fully protected connections

In this case $n = 3$ and all the connections can be carried when one path has failed. This scenario is denoted fully protected. If two paths are down simultaneously, only part of the connections can be carried. Figure 8 shows a scenario where $z = 6$ connections may use three paths. In this case, two connections will be down when there are two failed paths.

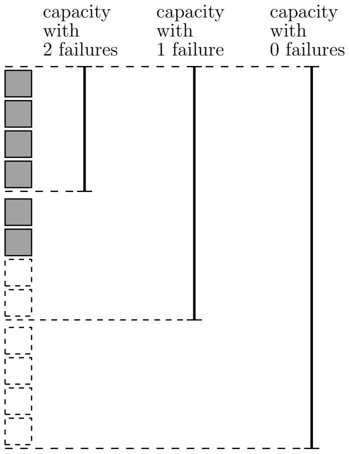

Figure 8. Assignment of connections to paths in fully protected scenario

3.1 Optical network parameters

Figure 9 shows the CDF of interval unavailability for 6 connections for $n = 3$ with optical network parameters. The figure shows that there is a high probability (about 98%) that there will be zero outages for a connection in the contract period for $T = 1$ year and a somewhat smaller probability for no outages for $T = 5$ years. This is because it takes two simultaneous failures to cause an outage. Also, $E(U_i)$ is close to zero. The interval unavailability for the connections that do experience outages is also more than one order of magnitude smaller than for $n = 2$. All the policies except the AP policy have more or less identical distributions. In the considered case, the gain with the AP policy is moderate.

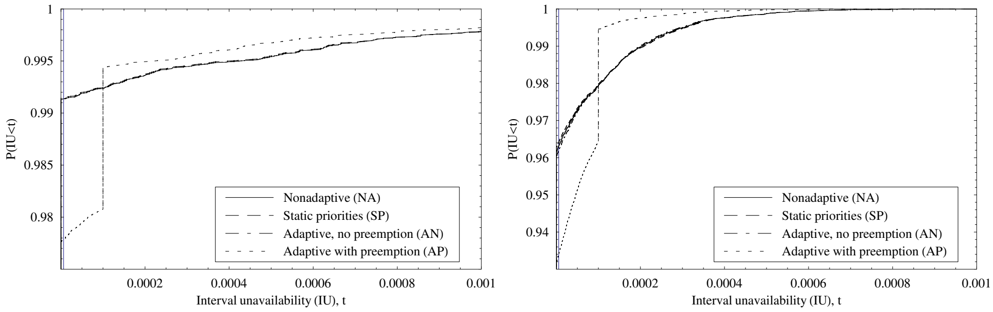

(a) $T = 1$ year (b) $T = 5$ years

Figure 9. Cumulative distributions of interval unavailability, optical network parameters, $n = 3$

Figure 10 shows the risk for optical network parameters with $n = 3$. If the operator is interested in providing services with very low levels of risk, the AP policy is better than the others. If a risk value of 0.01 is acceptable, any policy may be used for $T = 1$ year and correspondingly for $T = 5$ years with a risk value of 0.04.

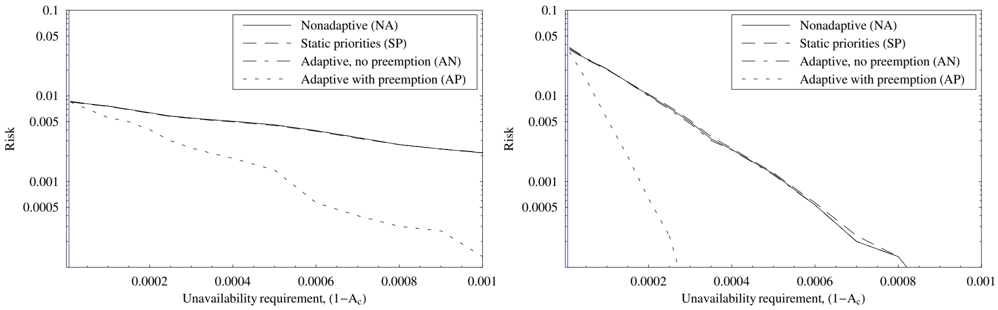

(a) $T = 1$ year (b) $T = 5$ years

Figure 10. Risk for optical network parameters, $n = 3$

3.1.0.1 Special case: Different Start and end times for customer contracts.

If the customers have different start and end times for their contract period, it will be possible to meet some contracts by moving the outage to a connection which is relatively new. The effect of this could be expected to be largest when the number of expected outages in an interval is small.

Figure 11 shows the effect of having customers with start times distributed uniformly over the contract period. The case is the optical network model for $n = 2$. When comparing with Figure 3, no significant differences are seen.

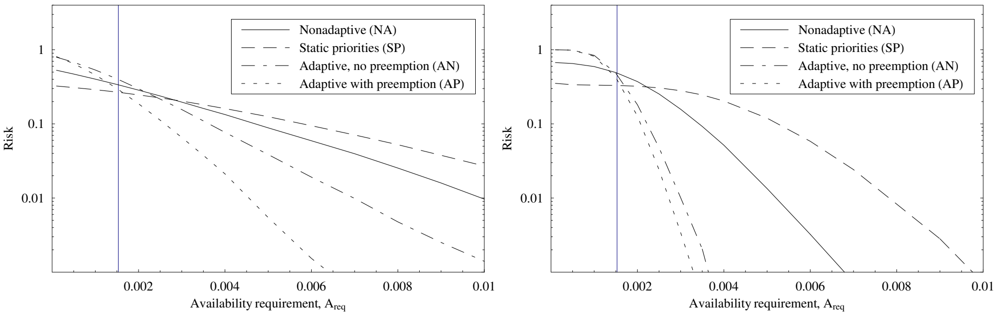

(a) $T = 1$ year

(b) $T = 5$ years

Figure 11. Risk with different start times for connections, optical network parameters, $n = 2$.

Figure 12 shows the effect of having customers with start times distributed uniformly over the contract period. The case is the optical network model for $n = 3$ and T of one year, chosen for minimum number of outages during an interval. When comparing with Figure 10, no significant differences are seen.

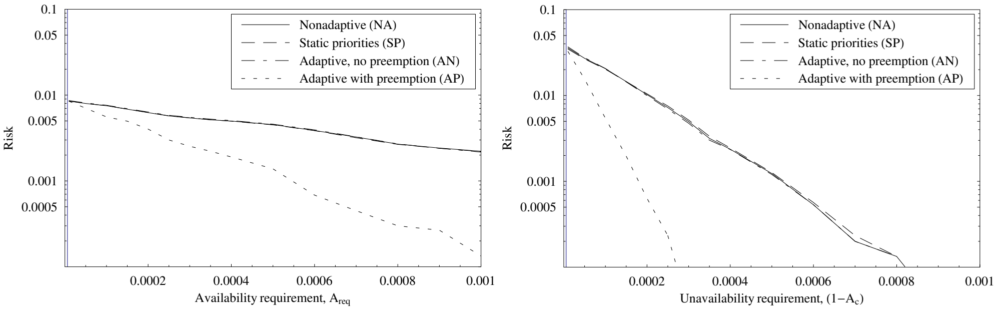

(a) $T = 1$ year

(b) $T = 5$ years

Figure 12. Risk with different start times for connections, optical network parameters, $n = 3$.

3.2 IP/MPLS network parameters

Figure 13 shows the CDF of interval unavailability for 6 connections for $n = 3$ for the IP/MPLS network parameters. When comparing with the optical network parameters in Figure 9, it is seen that there is a lower probability for observing a period without any failures, however, the failures are in general shorter and the CDF tend to 1 faster for IP/MPLS network parameters than for optical network parameters.

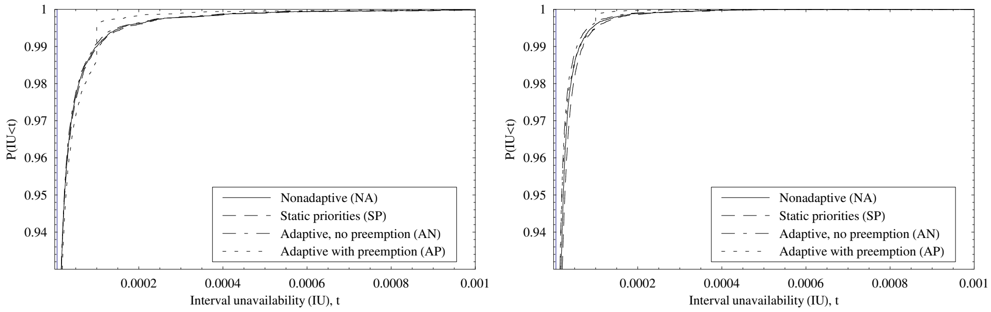

(a) $T = 1$ year

(b) $T = 5$ years

Figure 13. Cumulative distributions of interval unavailability for IP/MPLS network parameters, $n = 3$

Figure 14 shows the risk for $n = 3$ with 6 connections. When comparing with the risk for optical network parameters in Figure 10, it is seen that the difference between the risk for the AP policy and the other policies was smaller in the IP/MPLS model. This was also observed in the partially protected scenario.

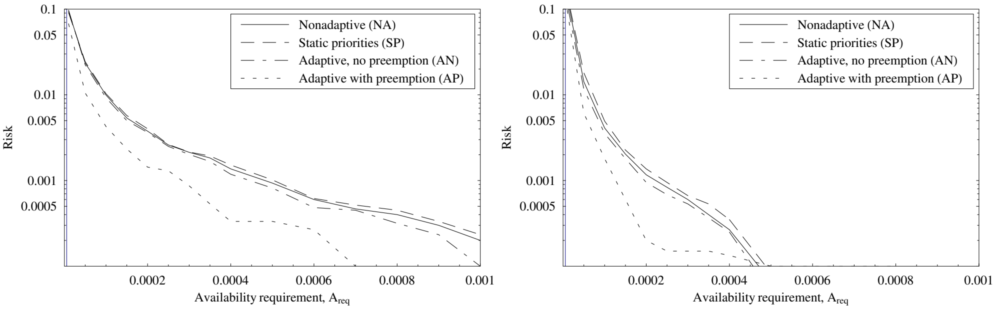

(a) $T = 1$ year

(b) $T = 5$ years

Figure 14. Risk for IP/MPLS network parameters, $n = 3$

4. Two service classes

In this case, $n = 3$ and there are two service classes with two different requirements. The scenario is illustrated in Figure 15.

There are two classes, the High and Low class, where the connections from the High class have requirement A_h and the connections from the Low class have requirement A_l, $A_h > A_l$. There are five connections from each class. A series of simulations are performed to investigate the effect of different ratios between A_h and A_l. A "base requirement" of $A_b = 0.999$ is used, and the requirements of the High and Low class are varied as $1 - A_h = (1 - A_b)/\sqrt{k}$ and $1 - A_l = (1 - A_b)\sqrt{k}$ respectively and simulated for different k. For instance, for $k = 100$, $A_h = 0.9999$ and $A_l = 0.99$. Parameters for the optical network model are used. The objective of this case is to show how the difference in requirements affect the risk of the two classes.

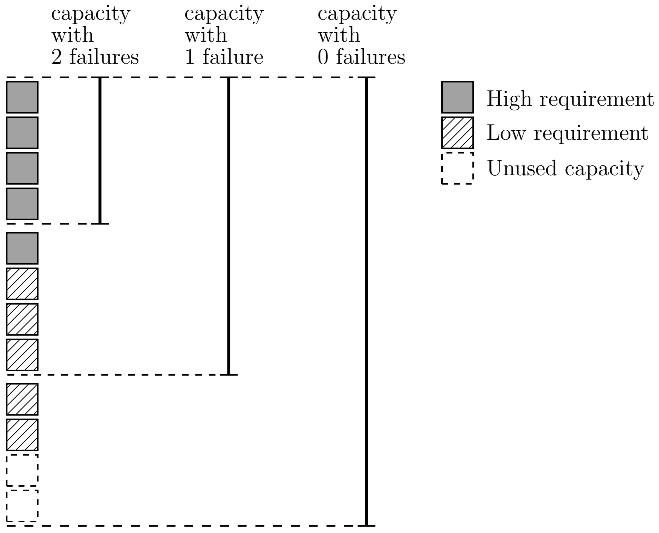

Figure 15. Assignment of connections with two service classes

The “NA, spare resources” performs bad for the High class for all requirements and is not included in the plots below. The “NA, all resources” policy serves as a references as a fair nonadaptive scheme and is included in both figures. It is seen that in all cases, the Low class has practically zero risk for $k > 100$ since the requirements for these k are low and easy to meet.

Figure 16 and Figure 17 shows the risk of the High and Low class separately for the considered policies.

Figure 16 shows the risk of the High and Low class separately for the “AN, class before budget”, the “AN, budget before class” policy and the “NA, all resources” policy. The risk for the High class with “AN, class before budget” is almost identical to the risk for the High class with the “NA, all resources”. This agrees with the results from the Fully Protected scenario since the High class connections are unaffected by the Low class connections with these policies and it takes two simultaneous path failures to cause an outage.

Figure 17 shows the risk of the High and Low class separately for the “AP, class before budget”, the “AP, budget before class” policy and the “NA, all resources” policy. It is seen that the two AP-based policies perform uniformly better than the corresponding AN-based policies. This agrees with the results from the two previous subsections. The two AP-based schemes perform better than the NA-based scheme for all k. For $k > 100$, the two AP-based schemes perform comparable to each other, but there is significant difference between them for $k < 100$. For low k, the “AP, budget before class” policy has high risk for both High and Low connections, while the “AP, class before budget” has a low risk for High connections. The “AP, budget before class” policy gives a more balanced risk for the two classes, since, when k is

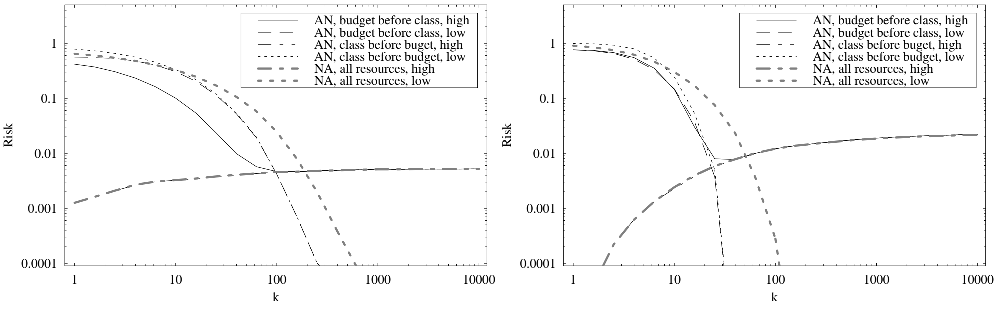

(a) $T = 1$ year

(b) $T = 5$ years

Figure 16. Risk for AN and nonadaptive policies with optical network parameters, $n = 3$ and two service classes

low, the difference in requirements is not very large and the policy gives similar risk for customers with similar requirements.

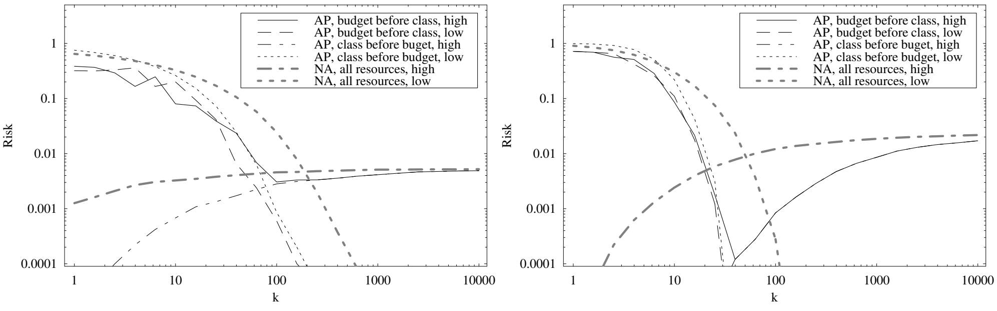

(a) $T = 1$ year

(b) $T = 5$ years

Figure 17. Risk for AP and nonadaptive policies with optical network parameters, $n = 3$ and two service classes

In Figure 17 for $T = 5$ years, the curves for the High class for the two policies “AP, budget before class” and “AP, class before budget” overlap from $k > 40$, but the risk for $k < 40$ is smaller than 0.0001 for “AP, class before budget”.

Which policy is the best depends on the cost associated with violating the SLAs for the different classes of customers. The expected total loss is the risk of violating an SLA multiplied by the cost of an eventual violation. Figure 18 shows the expected total loss where the cost of violating the SLA for High and Low class connections is the same, i.e. with cost factors $\alpha_h = \alpha_l = 1$. The results show that the two variants of the AP policy perform equally well. For $k > 10$, the AP-based policies perform best, while the the “NA, all resources” policy performs worst of the three considered policies.

An alternative cost-function is to let the cost of violating an SLA be proportional to the unavailability requirement. This means that for a given k, the cost of violating an SLA from the High class is $\alpha_h = \sqrt{k}$ while the cost for the Low class is $\alpha_l =$

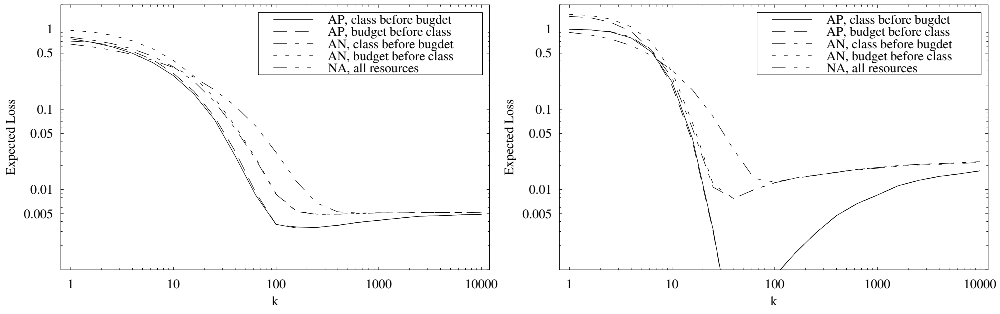

(a) $T = 1$ year

(b) $T = 5$ years

Figure 18. Expected loss with equal cost for optical network parameters, $n = 3$ with two service classes

$\frac{1}{\sqrt{k}}$. Figure 19 shows the expected loss for the proportional cost function. The most notable result is that the “AP, class before budget” policy now performs better than the “AP, budget before class” policy. This is because the High class connections have a lower risk with the class before budget policy, while it is more important to meet the requirements of the High class with the proportional function.

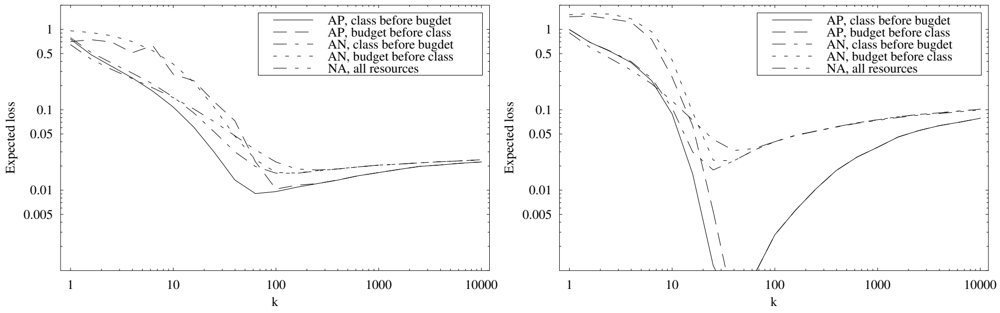

(a) $T = 1$ year

(b) $T = 5$ years

Figure 19. Expected loss with proportional cost for optical network parameters, $n = 3$ with two service classes

APPENDIX B: PROSPECTS FOR DEPENDABILITY DIFFERENTIATION IN INTERDOMAIN NETWORKS

Anders Mykkeltveit, Bjarne E. Helvik
Centre for Quantifiable Quality of Service in Communication Systems
Norwegian University of Science and Technology,
Trondheim, Norway
{mykkeltv, bjarne}@q2s.ntnu.no

Abstract Traditionally, the Internet has been used for best-effort traffic. The introduction of new services that have different Quality of Service (QoS)-requirements yields a need for the network operators to be able to offer different service levels. An important QoS-factor is dependability, e.g. the provision of services with a certain maximum accumulated downtime per year or guaranteed long periods of uninterrupted service. In this paper we investigate whether it is feasible for a provider to offer his customers the choice of differentiated levels of dependability. Most recent research on dependability differentiation focuses on intradomain networks. In order to introduce end-to-end differentiation over the Internet, differentiation must also be supported at the interdomain level. Studies on the behavior of interdomain routing show that the Internet cannot currently offer the level of dependability needed by customers or applications with high requirements. We therefore need mechanisms that can increase the availability for interdomain routes. These techniques may be used as a basis for differentiation by offering the improved level of dependability exclusively for high priority traffic. In this paper we present the challenges related to interdomain routing and review the literature on techniques to improve dependability at the interdomain level. We discuss how differentiation may be introduced at the interdomain level using the proposed improvements to interdomain dependability.

1. Introduction

In the future next generation Internet, different applications with different Quality of Service (QoS) requirements are expected to be carried in the same network. An important QoS-factor is dependability, e.g. that the service should be available when needed, be delivered without interruptions as long as needed and that down periods have a low upper bound. For communication services the dependability is related to the network's ability to tolerate faults, such as cable cuts and equipment failures, and it relates to the correctness of network protocols, operation and maintenance.

Since different customers and different services have different dependability requirements, the provider is interested in mechanisms that allow him to differentiate the level of dependability offered to different customers. Improving the dependability is expensive since redundant resources are needed, but if some customers are willing to pay a higher price for better service, this may generate revenue for the operator.

The Internet is composed of a large number of networks. Each network is managed and operated by a company or organization and is denoted an Autonomous System (AS). The network operator has full control over the resources in his own network, and he may provide different levels of QoS for the traffic inside the AS. Several frameworks for providing differentiated dependability inside an AS have been proposed in the literature, see [CJHM05] for an overview. A customer with high dependability requirements is likely to expect the same level of dependability for all his connections, whether they are with other customers connected to the same network, or if they are connected to a network operated by another provider. The frameworks for differentiated dependability only consider a single AS and are not easily extended to the challenges imposed by interdomain routing. In particular, the frameworks will suffer severe scalability problems, and they rely on a global view of the network they operate in, which is not consistent with the network operators' decisions about keeping information on their network hidden from the outside world. There is therefore a need for new frameworks that can allow the introduction of differentiated dependability in the Internet.

The dependability of interdomain paths has been studied for many years. The availability is the probability that a service is working at a random point of time. The results show that the Internet cannot currently offer the level of availability needed by customers with high dependability requirements. A measurement study [Pax97] found that the probability of encountering a routing pathology was between 1.5% and 3.3% in 1994–1995. A more recent study indicated that the situation has not improved much since then [GMG$^+$04]. In addition to availability, some services will have requirements on the duration of the outages, and the time between failures. While some services are unaffected by short failures, other services like tele-surgery will have very strong requirements on continuity of the service. Unfortunately, failures are reported to be frequent, and although many failures are short a significant percentage of them last more than 2 minutes [FABK03].

Recently, some research has been devoted to techniques for improving the dependability for interdomain paths. The adoption of some of these suggestions would introduce the possibility to give a positive differentiation if the mechanisms were only applied for traffic with high dependability requirements.

The objective of this paper is to investigate the challenges related to providing improved and differentiated dependability in the interdomain level and to investigate mechanisms proposed in the literature for achieving this. We discuss the feasibility of using these mechanisms as a basis for dependability differentiation at the interdomain level. To the authors' knowledge, this is the first paper to compare and investigate the existing literature in this area. The paper is organized as follows. In Section 2, we present the concept of interdomain differentiation. Section 3 introduces the reader to the state-of-art in interdomain routing. Section 4 presents an overview and comparison of the literature of improving interdomain dependability. Section 5 gives concluding remarks.

2. Dependability Differentiation in interdomain networks

In this paper, we are interested in investigating whether an Internet Service Provider (ISP) can offer his customers the choice between multiple levels of dependability. The dependability seen from a customer is the end-to-end dependability for all his connections. These connections span access networks, and the networks of one or more ASes. The routing inside an AS is referred to as intradomain routing, and the routing crossing multiple ASes is denoted interdomain routing. Intradomain, interdomain and access networks can all be regarded as different research areas for dependability differentiation. Figure 1 shows the scopes of these three areas.

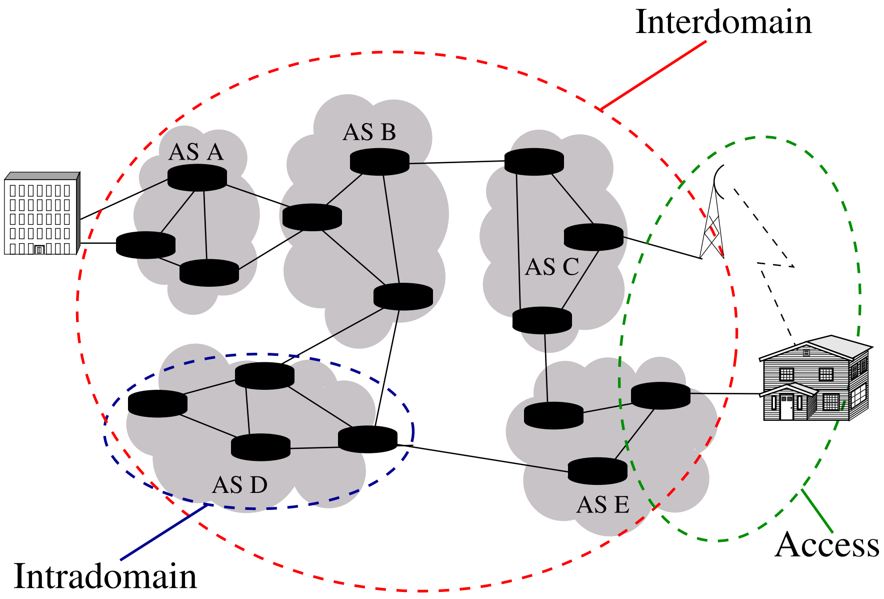

Figure 1. Research areas in end-to-end dependability differentiation.

The need for dependability differentiation has spurred a lot of research, but so far this research focuses on a single AS. It is convenient to limit the scope of differentiation to the intradomain level, since the operator has full control over his own network However, since most connections in the Internet span several ASes, we believe that dependability differentiation is also needed at the interdomain level in order to provide a real difference to the users. As a minimum, it should be possible to establish some paths with a higher level of dependability through the Internet. Interdomain routing poses a number of additional challenges compared to intradomain routing, and these challenges have to the authors' knowledge not been directly addressed in the context of dependability differentiation.

For real end-to-end dependability, the access network all the way to the customer premises should be included since it will probably be the dependability bottleneck for the Internet connection. The customer can be multihomed either to two service providers or with multiple links to the same provider. We focus on a situation where one ISP provides a dependable service to the end-users. We therefore assume that the customer is connected with at least two independent physical links to his ISP, and that the ISP has implemented support for providing a dependable service of packet transfer through his network from the user's access links and to the network border toward other ASes.

3. Interdomain routing

The Internet is composed of a large number of networks, denoted ASes. Each AS is managed independently by a company or organization. Internally, the operator runs an interior routing protocol. In pure IP-networks, protocols such as OSPF [Moy98] or IS-IS [ISO90] are used, and packets are forwarded on a hop-by-hop basis. Other networks use connection-oriented technologies such as Multi-Protocol Label Switching (MPLS) or Generalized Multi-Protocol Label Switching (GMPLS) to establish Label Switched Paths (LSPs). With respect to differentiation, the connection-oriented approach seems most suited and has been the focus for most research.

To allow traffic to traverse network boundaries, operators must connect their networks and agree on the terms for data transfer. The two most common relations between ASes are provider-customer and peer-to-peer. In a provider-customer relationship, a provider transports traffic to and from its customers to give them access to the rest of the Internet. The customer AS is charged for the transport service. ASes that have a peer-to-peer relationship have agreed to forward traffic between their customers, usually without charging each other.

Figure 2 shows a graph representing six ASes and the relations between them. Provider-customer relations are indicated by solid-lined arrows. In the figure, AS C is a customer of AS E and a provider for AS A. The dotted arrow between AS C and AS D indicates that they have a peer-to-peer relation, allowing traffic between AS A to AS B to traverse peering link. The relations between the ASes give the Internet a hierarchical structure. ASes that do not have any customer ASes are denoted stub ASes. The stub ASes are typically ISPs.

Multihoming at the AS-level is frequently used to increase the availability of a network or for load-sharing. This means that a customer is connected to multiple providers. In Figure 2, AS A is connected to two providers, a special case of multihoming denoted dual homing.

The business-relationships between the ASes are reflected in routing policies which are configured in the routers. A multihomed provider may choose to forward most of its traffic to the one of his provider that offers the lowest price for the transit service. This will affect which routes the operator advertises to his customer networks, and the outcome of the policy could be that the path toward the destination is inflated [GW02]. An important result of policy configuration is that not all paths that are physically possible may be used for transporting data.

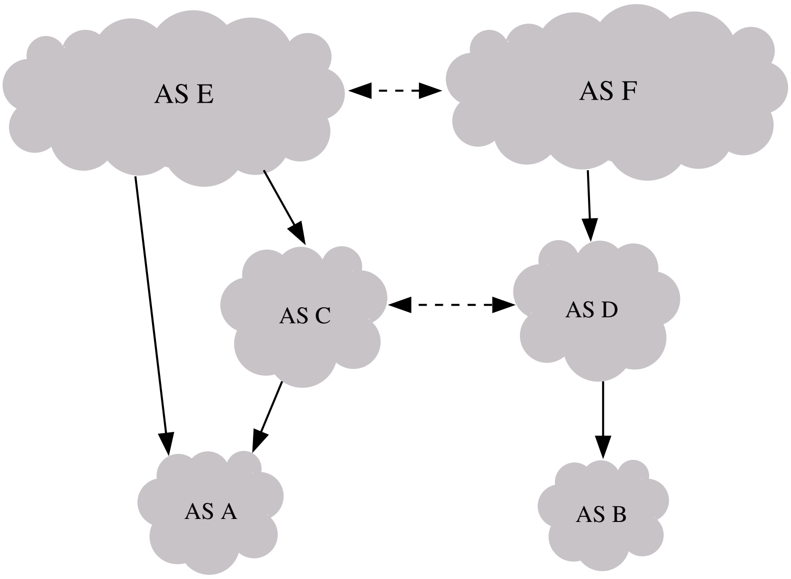

Figure 2. ASes and relations

Studies on the Internet topology have found that the AS-degree (the number of connection points between an AS and other ASes) follows a power law [DKR05]. This means that there is a large number of networks with low AS-degree and a small number of transit networks at the top of the AS-hierarchy with very high AS-degree.

Today the internals of the ASes are hidden from other ASes for many reasons, and this makes it difficult to introduce mechanisms to enhance the performance of the network at large since no information on capacities, redundancy and congestion is exchanged between the ASes.

The huge number of ASes in the Internet today, in addition to the expected continued growth, indicates that scalability is one of the most important requirements for the interdomain routing protocols [Yu00, FBR04]. It seems unlikely that one entity will be able to have or allowed to have a global control over the routing of the Internet.

3.1 Current practice - BGP

The current de facto standard interdomain routing protocol is the Border Gateway Protocol (BGP) [RLH06]. The key idea of BGP is to allow exchange of network reachability information between BGP routers. The destination for a BGP route announcement is expressed with an Internet Protocol (IP) address prefix. The announcement contains the path of ASes used by the announcing router to reach the destination prefix. BGP is divided into an internal and an external protocol. The internal protocol is run between routers in the same domain. They exchange information on reachability they have received from other networks and decide on which is the best if a network can be reached via multiple paths. Normally, all BGP routers have peering sessions with all the other BGP routers in the same domain. In large networks, the number of internal sessions is reduced by applying route reflectors that mediate this communication.

Each network has one or more edge routers that communicate with edge routers of neighbor networks. These routers run an external routing protocol which they use to exchange routes to other networks and routers in their own network, all depending on their policies and agreements. In Figure 3, internal routing sessions are indicated by dotted arrows while external sessions are indicated by solid arrows.

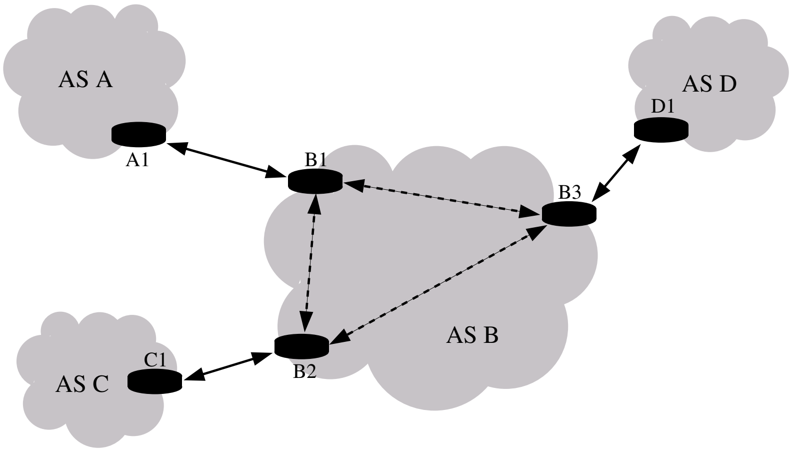

Figure 3. Internal and external BGP

When a BGP router looks up the route for a packet, the longest matching prefix is always selected. If more than one path to a destination is found in the router's database, the routing policies are checked to see if one route is preferred over the other. It there is still a tie, the shortest path, measured in the number of ASes it traverses is chosen. Finally some tie-breaking rules are applied. This is in effect a random process. Each BGP router normally only announces the routes they use themselves, so each router has an incomplete view of the overall network topology.

With BGP, failures are detected by a BGP router when it looses contact with one of its peers. It then sends out UPDATE messages to its remaining peers. The update message could contain a withdrawal of the routes that are affected by the failure. The news about the route withdrawal must be propagated throughout the Internet. If an AS is multihomed, its BGP routers may have information on other routes to the destination. They can send out UPDATE messages containing the new routes to their peers.

Measurement studies have found that the time taken for all the BGP routers to converge to a consistent view of the topology can be in the order of several minutes. During this convergence process, packets are highly likely to be dropped at some point, but they may also reach the destination via fluctuating paths [LABJ01]. The QoS for applications using paths that are converging can be so degraded that the route should be considered as down.

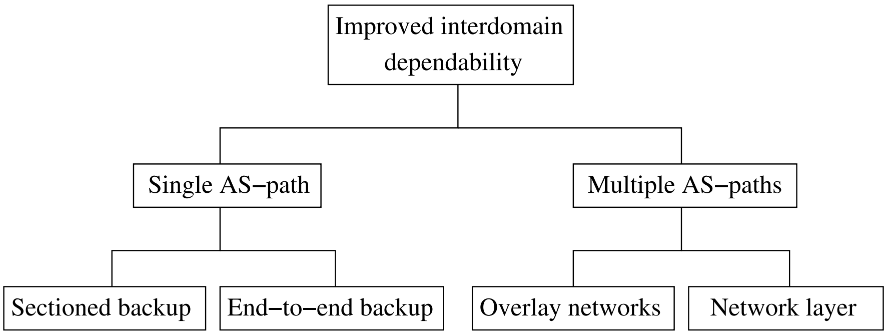

Figure 4. Proposed methods to improve interdomain dependability.

It has been observed that a single failure is typically announced many times, and that several rounds of route updates are needed before convergence is reached.

Today it is common for peers to extend the peering agreement to include forwarding of traffic as a backup solution if the provider link fails [GGR01], [CB96]. BGP can announce backup routes that span one or more peer links if both ASes have agreed on this.

3.2 Interdomain traffic engineering

Network operators have the possibility to manually configure routing in their ASes to minimize the cost of traffic forwarding. One common strategy is early-exit (or hot-potato) routing which is used if a multihomed AS has multiple routes to a destination. The idea is simply to forward traffic to the nearest egress router from the incoming router. While this type of routing can reduce the costs for one operator, it can certainly increase the cost for other operators, and might also result in a higher global cost of routing.

Announcement of redundant routes is used for interdomain traffic engineering. A multihomed stub AS may prefer to share the load of incoming traffic over its provider links. This can be done by announcing an address prefix corresponding to one half of the address space to one provider and the other half to the other provider in addition to announcing the full AS-prefix. With this arrangement, the load is shared under normal operation since the ASes will route packets based on the most specific prefix. If one of the provider links fail, the traffic going over that link can be sent over the other link since the full AS-prefix is now the only available route to the destination AS [QPS$^+$03].

A framework where neighbor ASes negotiate on the use of the links between them is suggested. The idea is that it should be possible to find better routing configurations for both networks than the default (early-exit routing) case where no negotiation is used [MWA05].

There has been some research on how to provide differentiated traffic handling performance in the Internet [XWLN04, HFP$^+$05]. The challenges faced here are similar to what we face with dependability differentiation. The most important requirements to new architectures are scalability in addition to the requirements from the business relations between the ASes in the Internet.

4. Improving interdomain dependability

Some research has been directed toward improving dependability at the interdomain level. Differentiation is not mentioned explicitly, but some techniques can be applied to only parts of the traffic, thereby implicitly introducing (at least) two dependability levels. Figure 4 shows a decomposition of the principal methods used to improve interdomain dependability. The presentation of the literature on interdomain dependability is organized according to this figure. The first classification is whether the proposed methods use a single or multiple interdomain (AS-level) routes to improve dependability. In the first case, the suggested methods try to improve the dependability of routes so that failures of links and routers do not result in a route failure. In the second case, the methods take advantage of the existence of multiple paths at the AS-level so that when one path fails, data can be switched to another path.

4.1 Single AS-path

These approaches aim to reduce the number of route changes that are announced globally by strengthening the paths. If faults can be handled locally by the ASes that are members of the path, the failure may not need to be announced outside the neighborhood. This will reduce the frequency of BGP updates which lead to the slow convergence process.

Methods for strengthening single AS-paths may be implemented using different underlying technologies. In the literature, both connection-oriented approaches using MPLS LSPs and methods based on IP-restoration have been proposed. In general, MPLS protection will give faster recovery than IP-restoration. In MPLS networks, capacities might be dedicated so congestion is avoided when switching to the backup path. If the connection-oriented (MPLS) approach is used, all the ASes must have the same underlying technology. This will typically limit the possibilities of using these techniques end-to-end in the Internet, and these techniques are therefore most easily applied inside large networks that are divided into multiple domains, or in a small neighborhood of cooperating ASes.

The single AS-path approaches can be divided into two classes depending on the scope of the backup paths that are used in case of failure. With sectioned backup, the backup paths have limited scope, restricted to two neighbor ASes in the path. With global backup, the backup paths are disjoint from the source AS to the destination AS.

4.1.1 Single AS-path with sectioned backup

A scheme using MPLS to provide local backup LSPs in case of failure to the infrastructure connecting two ASes was proposed by Huang and Messier [HM03]. They propose to improve the reliability of BGP-connections by allowing adjacent ASes to cooperate by establishing local MPLS backup paths crossing the domain boundary between them. The approach assumes that there exists at least two independent links between each neighboring AS in the path. An example configuration of the LSP arrangement resulting from this approach is shown in Figure 5. The primary path (solid line) is protected by different backup LSPs (dotted lines). The backup path arrangement is able to protect against all single failures of a link or router in the primary LSP. The minimal number of backup paths that are needed for each primary path in each transit AS, like AS B in the figure, is two.

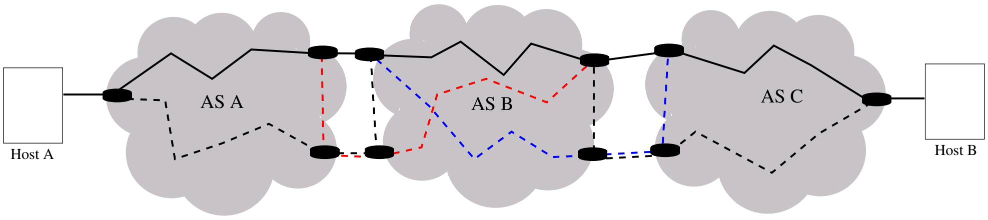

Figure 5. Single AS-path with sectioned backup.

It has been suggested to allow backup LSPs via ASes that are not in the single path announced at the BGP level [RLC+05]. It is suggested to attach a *recovery IP address* to each AS in the AS-path in the BGP route announcements. The recovery address belongs to a router in an AS that is different from the next hop in the AS-path. Figure 6 shows an example where a backup path for failures in AS B goes via AS D.

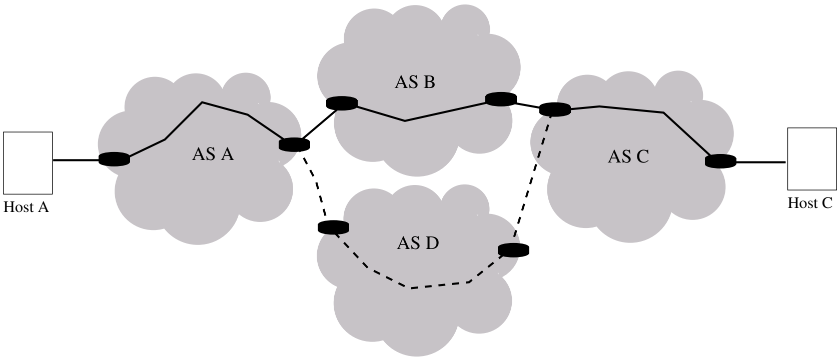

Figure 6. Backup paths via AS not in currently used route.

In [ELS+05], several ideas for strengthening a single AS-path using IP restoration are presented. First of all, failures inside an AS should be recovered by IP rerouting without altering the AS-path announced by BGP. Second, the interconnection infrastructure between all neighbor ASes should be redundant, so that failures at the inter-AS links and edge routers can be recovered by IP-level mechanisms. When the two former approaches fail, a temporary reroute through a third AS could be allowed, resulting in the same route pattern as shown in Figure 6. The information on the route change is not propagated via BGP unless it persists for some time. A potential problem with this approach is that it may cause policy violations for a short time period since some ASes will not accept to have their traffic sent through the third AS that provides the local backup. In cases where the failure cannot be recovered locally, ordinary BGP UPDATE messages with withdrawals of the affected routes are initiated.

4.1.2 Single AS-path with end-to-end backup

These approaches use BGP to find AS-paths, but MPLS is used to set up disjoint primary and backup paths through the AS-path. Currently, there exist drafts that suggest how to establish link and node disjoint paths through an AS-chain [RMA05]. The resulting primary and backup paths are shown in Figure 7.

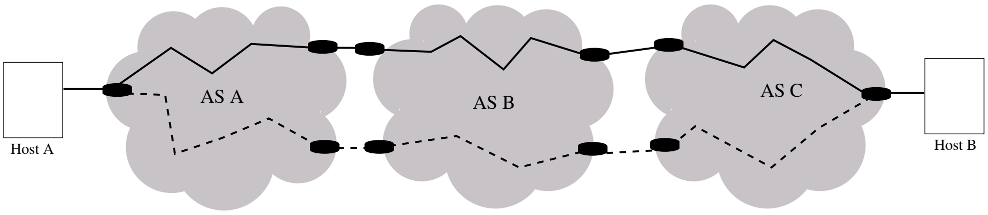

Figure 7. Single paths with end-to-end backup paths.

It should be noted that segments of the backup paths might be allowed to traverse different ASes than the primary path. In [RMD04], the existence of two partially AS-disjoint BGP routes in the originating AS is used to set up LSPs that are AS-disjoint until they merge.

4.2 Multiple AS-paths

These approaches aim to take advantage of the existence of disjoint paths at the AS-level. By assuming that the failure processes of the ASes are independent, the probability of finding two (or more) paths unavailable is reduced. We divide the approaches using multiple AS-paths into two categories; overlay networks that operate on top of the current Internet and approaches that operate at the network layer.

4.2.1 Overlay Networks

Overlay networks are "networks on top of networks". An overlay network consists of nodes that communicate with a protocol using the underlying network. Overlay networks is a general technique that does not intervene with the underlying infrastructure.

An overlay network denoted Resilient Overlay Network (RON) is presented in [ABKM01]. The network consists of a set of 2–50 host computers that are connected to different providers. A host may act as a relay for other hosts if they can reach a server that is unavailable to other hosts or if this relaying will improve performance. The improved dependability obtained by using a RON is due to the existence of multiple paths at the interdomain level are available in the RON, and even routes that are not advertised to a given AS may in some cases be used when relaying through a node in the overlay network.

In Figure 8, the primary path for the connection between Host A and Host C is indicated with the solid line. The backup path, indicated with dotted line, goes via Host B. Note that the underlying AS-structure is not part of the routing information in the overlay network.

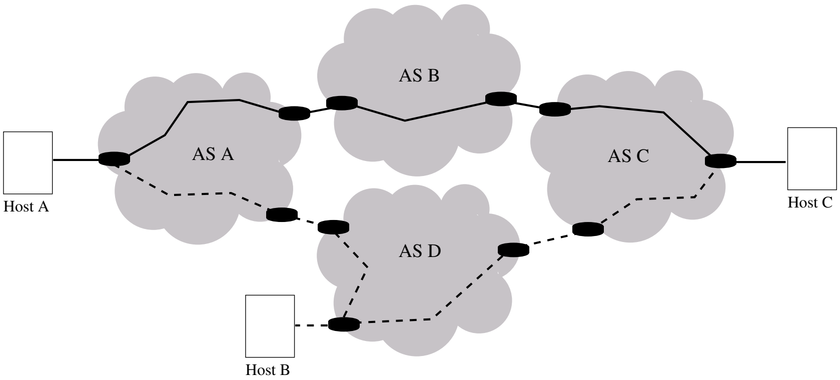

Figure 8. Overlay network.

The diversity of the paths found by nodes in RON networks has been studied in [HJ04], and it is found that the probability of finding two disjoint paths is small if the overlay nodes are chosen at random.

To improve the dependability of overlay networks such as RON, [HWJ05] suggests a method for placing the nodes so that the paths become less overlapping. This is combined with dynamic selection of backup routes via nodes that give "most disjoint" routes. Using the new strategy, many failures can be covered by going via only one node in the overlay network.

In [Jan02], it was suggested that the access routers could forward packets after instructions from the host that is member of an overlay networks in order to reduce the load on the access links. The overlay nodes can be imagined moved further into the networks. An arrangement with one or more multi-homing routers placed in each domain is presented in [LMC05]. The overlay routes and the data packets can be source routed in the overlay network.

4.2.2 Network layer approaches

These approaches try to find disjoint AS-paths by finding and using disjoint paths at the network layer.

4.2.2.1 Multihoming. Multihoming of the provider AS is currently used to have access to multiple routes. Commercial solutions to take advantage of multihoming exist. For instance, a service that monitors the connections of the user and dynamically selects the best is offered by Internap [Int06]. Unfortunately, multihoming does not always give disjoint routes, so the dependability improvement is limited. There seems to be potential for methods to improve the diversity of the paths that are used, but this will require modifications to the BGP protocol.

4.2.2.2 Source routing. The IP protocol offers source routing so that a user may specify a loose or strict sequence of addresses a packet should visit before it reaches its destination. For security reasons, IP source routing is generally disabled in routers for forwarding traffic other than ICMP [HH99, SR04]. A variant of source routing techniques could have potential to be used for improved dependability, but the security problems must be avoided.

4.2.2.3 AS-level route selection. In the literature, several authors have proposed architectures that allow stub ASes or even end users to choose the routes they want their packets to be forwarded over. The dependability might increase if the senders of the packets are able to choose from more paths than those available today. We consider AS-level route selection approaches to be different from IP source routing since the users are only allowed to choose from routes that are suggested by the ASes that will provide the transit service. To introduce AS-levels route selection, new interdomain routing protocols and architectures are needed.

Feedback based routing has recently been proposed as a network layer mechanism to improve dependability [ZGC03]. The border routers in transit networks forward information on available interdomain links and which ASes they connect. Each stub AS has an access router that calculates two disjoint routes to each destination AS based on topology information received from the routers in the transit networks. The access routers specify the path for each packet sent, and the intermediate routers forward packets according to this annotation. The quality of the connections is monitored and the best path is used as primary. For applications that demand zero downtime, the data can be sent along both paths. The link state information is updated periodically with the period being at least an hour. The access routers have timers associated with each link, and if the timer expires, the routes associated with that link are attempted replaced by other routes.

Other frameworks for allowing the users of stub ASes to choose the routes are Nimrod [CCS96], NIRA [Yan03] and BANANAS [KKW$^+$03].

4.3 Comparison and evaluation

Table 1 compares the proposed strategies for improved interdomain dependability. The currently deployed BGP protocol does not offer potential for differentiated dependability, and suffers from long service interruption times. All the other strategies considered in the table have some potential for differentiation. A multihomed ISP is currently able to offer a limited form of differentiated dependability for outgoing traffic.

Strategy		Guaranteed protection from single failure	Service interruption time	Continuous service (1+1)	Differentiation	Comment	Selected references
Situation today	BGP today	no	long	not possible	no	slow convergence after failures	[RLH06]
	Multihoming	no	shorter than above	possible	limited possibilities	paths not necessarily disjoint	[HJ04]
Single AS-path	End-to-end backup	yes	short	yes	yes	need redundant inter-AS links	[RMA05]
	Sectioned backup	yes	short	no	yes	need redundant inter-AS links	[ELS$^+$05, HM03]
Multiple AS-pats	Overlay networks	no	short	possible	yes	paths not necessarily disjoint	[ABKM01]
	Network layer	yes	short	possible	yes	need new interdomain routing architecture	[ZGC03, Yan03]

Table 1. Proposed strategies for improving interdomain dependability and their properties

Upon failure of the link to the ISP's primary provider AS, the link to his remaining providers could be used exclusively to send packets for customers or services with high dependability requirements. Overlay networks are also possible to establish today as they do not require any changes to the BGP protocol. However, if the overlay approach is taken, the underlying infrastructure is hidden, and it will not be possible to provide guaranteed protection from a single failure since infrastructure may be shared between any pair of paths in the overlay network.

The single-AS strategies have potential to provide routes that can handle all single failures for an end-to-end connection. The service disruption time depends on the recovery mechanisms employed, and can be short if MPLS is used. The strategies, however demand that all ASes that are part of a dependable path are upgraded with redundant inter-AS links and implement the needed functionality.

The network layer approaches using multiple AS-paths require new interdomain routing architecture to be able to be used for providing differentiated dependability.

5. Concluding remarks

The provision of differentiated dependability over multiple ASes in the Internet is a challenging research topic. In addition to technical challenges such as scalability, interdomain routing must also take the business relationships like information hiding and service level agreements between the ASes into account. The proposed frameworks for intradomain differentiation are not directly extensible to the interdomain level, and new frameworks are needed to cope with these challenges. Based on an evaluation on the literature on improved interdomain dependability, we have outlined two possible approaches toward interdomain differentiation. The first approach is to strengthen the interconnection between the ASes so that failures of the interdomain paths are avoided. The strengthened paths could be used to provide a service with better dependability than what is offered today to users and services with high dependability requirements. The second approach to dependability differentiation is to take advantage of the existence of multiple paths at the AS-level to maintain disjoint primary and backup paths which can be used by the users and services with high dependability requirements.

References

[ABKM01] David Andersen, Hari Balakrishnan, Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of the eighteenth ACM symposium on Operating systems principles (SOSP)*, pages 131–145, New York, NY, USA, 2001. ACM Press.

[CB96] E. Chen and T. Bates. An application of the BGP community attribute in multi-home routing. RFC 1998, August 1996.

[CCS96] I. Castineyra, N. Chiappa, and M. Steenstrup. The Nimrod routing architecture. RFC 1992, August 1996.

[CJHM05] Piotr Chołda, Andrzej Jajszczyk, Bjarne E. Helvik, and Anders Mykkeltveit. Service differentiation based on recovery methods. In *EURO NGI Workshop on Traffic Engineering, Protection and Restoration*, Rome, Italy, April 21-22 2005.

[DKR05] Xenofontas A. Dimitropoulos, Dmitri V. Krioukov, and George F. Riley. Revisiting internet AS-level topology discovery. In *Proceedings of Passive & Active Measurement (PAM) workshop*, pages 177–188, Boston, USA, March 31–April 1 2005.

[ELS$^+$05] T. Engel, G. Lichtwald, K. Schrodi, T. Schwabe, and B. Toedtmann. Increasing end-to-end availability over multiple autonomous systems. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications*, Las Vegas, USA, June 27–30 2005.

[FABK03] Nick Feamster, David G. Andersen, Hari Balakrishnan, and M. Frans Kaashoek. Measuring the effects of internet path faults on reactive routing. In *Proceedings of the ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 126–137, New York, NY, USA, 2003. ACM Press.

[FBR04] Nick Feamster, Hari Balakrishnan, and Jennifer Rexford. Some foundational problems in interdomain routing. In *3rd ACM SIGCOMM Workshop on Hot Topics in Networks (HotNets)*, San Diego, California, USA, November 2004.

[GGR01] Lixin Gao, Timothy G. Griffin, and Jennifer Rexford. Inherently safe backup routing with BGP. In *Proceedings of IEEE INFOCOM*, pages 547–556, Anchorage, Alaska, April 22–26 2001.

[GMG$^+$04] Krishna P. Gummadi, Harsha V. Madhyastha, Steven D. Gribble, Henry M. Levy, and David Wetherall. Improving the reliability of Internet paths with one-hop source routing. In *Proc. 6th USENIX OSDI*, San Francisco, CA, December 2004.

[GW02] Lixin Gao and F. Wang. The extent of AS path inflation by routing policies. In *Proceedings of IEEE Global Internet Symposium*, pages 2180–2184, November 17–21 2002.

[HFP$^+$05] M.P. Howarth, P. Flegkas, G. Pavlou, Ning Wang, P. Trimintzios, D. Griffin, J. Griem, M. Boucadair, P. Morand, A. Asgari, and P. Georgatsos. Provisioning for interdomain quality of service: the MESCAL approach. *IEEE Communications Magazine*, 43(6):129–137, 2005.

[HH99] B. Harris and R. Hunt. TCP/IP security threats and attack methods. *Computer Communications*, 22(10):885–897, June 1999.

[HJ04] Junghee Han and Farnam Jahanian. Impact of path diversity on multi-homed and overlay networks. In *Proceedings of International Conference on Dependable Systems and Networks (DSN)*, pages 29–38, Florence, Italy, June 28–July 01 2004.

[HM03] Changcheng Huang and Donald Messier. Inter-domain MPLS restoration. In *Proceedings of Fourth International Workshop on Design of Reliable Communication Networks (DRCN)*, pages 341–348, Ottawa, Canada, October 19–22 2003.

[HWJ05] Junghee Han, David Watson, and Farnam Jahanian. Topology aware overlay networks. In *Proceedings of IEEE Infocom*, volume 4, pages 2554–2565, Miami, Florida, March 2005.

[Int06] Internap. http://www.internap.com, March 2006.

[ISO90] Intermediate system to intermediate system intra-domain routing exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473), February 1990.

[Jan02] John Jannotti. Network layer support for overlay networks. In *Proceedings of IEEE OPENARCH*, pages 3–13, New York, New York, June 2002.

[KKW$^+$03] H. Tahilramani Kaur, S. Kalyanaraman, A. Weiss, S. Kanwar, and A. Gandhi. BANANAS: an evolutionary framework for explicit and multipath routing in the Internet. In *Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture*, pages 277–288, New York, NY, USA, 2003. ACM Press.

[LABJ01] Craig Labovitz, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. Delayed Internet routing convergence. *IEEE/ACM Tansactions on Networking*, 9(3):293–306, 2001.

[LMC05] Zhi Li, Prasant Mohapatra, and Chen-Nee Chuah. Virtual multi-homing: On the feasibility of combining overlay routing with BGP routing. In *IFIP Networking Conference*, volume 3462 of *Lecture Notes in Computer Science (LNCS)*, pages 1348–1352. Springer, May 2005.

[Moy98] J. Moy. OSPF version 2. RFC 2328, April 1998.

[MWA05] Ratul Mahajan, David Wetherall, and Thomas Anderson. Negotiation-based routing between neighboring ISPs. In *Networked Systems Design and Implementation (NSDI)*, May 2005.

[Pax97] V. Paxson. End-to-end routing behavior in the Internet. *IEEE/ACM Transactions on Networking*, 5(5):601–615, 1997.

[QPS$^+$03] Bruno Quoitin, Cristel Pelsser, Louis Swinnen, Olivier Bonaventure, and Steve Uhlig. Interdomain traffic engineering with BGP. *IEEE Communications Magazine*, 41(5):122–128, May 2003.

[RLC$^+$05] Ricardo Romeral, David Larrabeiti, Miguel Couto, Macelo Bagnulo, and Alberto Garca. MPLS-supported interdomain recovery in the public internet. In *Proceedings of IV Workshop in MPLS/GMPLS networks*, pages 103–114, Girona, Spain, April 21–22 2005.

[RLH06] Y. Rekhter, T. Li, and S. Hares. A border gateway protocol 4 (BGP-4). RFC 4271, January 2006.

[RMA05] Fabio Ricciato, Ugo Monaco, and Daniele Ali. Distributed Schemes for Diverse Path Computation in Multidomain MPLS Networks. *IEEE Communications Magazine*, 43(6):138–146, June 2005.

[RMD04] Fabio Ricciato, Ugo Monaco, and Alessio D'Achille. A novel scheme for end-to-end protection in a multi-area network. In *Proceedings of 2nd International Workshop on Inter-domain Performance and Simulation (IPS)*, pages 89–96, Budapest, Hungary, March 2004.

[SR04] Alex C. Snoeren and Barath Raghavan. Decoupling policy from mechanism in internet routing. *SIGCOMM Comput. Commun. Rev.*, 34(1):81–86, 2004.

[XWLN04] Li Xiao, Jun Wang, King-Shan Lui, and Klara Nahrstedt. Advertising interdomain QoS routing information. *IEEE Journal on Selected Areas in Communications*, 22(10):1949–1964, December 2004.

[Yan03] Xiaowei Yang. NIRA: A new Internet routing architecture. In *Proceedings of the ACM SIGCOMM workshop on Future directions in network architecture (FDNA)*, pages 301–312, New York, NY, USA, 2003. ACM Press.

[Yu00] J. Yu. Scalable routing design principles. RFC 2791, July 2000.

[ZGC03] Dapeng Zhu, Mark Gritter, and David R. Cheriton. Feedback based routing. *SIGCOMM Computer Communications Review*, 33(1):71–76, January 2003.

Bibliography

[AHPH05] M. Amin, K.-H. Ho, G. Pavlou, and M. Howarth. Improving Survivability through Traffic Engineering in MPLS Networks. In *Proc. 10th IEEE Symposium on Computers and Communications ISCC'2005*, La Manga del Mar Menor, Cartagena, Spain, June 27-30, 2005.

[AK02a] Achim Autenrieth and Andreas Kirstaedter. Engineering end-to-end IP resilience using resilience-differentiated QoS. *IEEE Communications Magazine*, 40(1):50–57, January 2002.

[AK02b] Achim Autenrieth and Andreas Kirstaedter. RD-QoS - the integrated provisioning of resilience and QoS in MPLS-based networks. In *Proceedings of IEEE International Conference on Communications (ICC '02)*, volume 2, pages 1174–1178, April 28–May 2 2002.

[AKM03] Shinichi Arakawa, Junichi Katou, and Masayuki Murata. Design method of logical topologies with quality of reliability in WDM networks. *Photonic Network Communications*, 5(2):107–121, March 2003.

[ALRL04] Algirdas Avizienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1):11–33, January–March 2004.

[AS05] C. Awad and B. Sanso. Network reliability under mixed IP and optical protection. In *Design of Reliable Communication Networks, 2005. (DRCN 2005). Proceedings.5th International Workshop on*, page 8pp., 16-19 Oct. 2005.

[Aut03] Achim Autenrieth. Recovery Time Analysis of Differentiated Resilience in MPLS. In *Proc. 4th International Workshop on the Design of Reliable Communication Networks DRCN 2003*, pages 333–340, Banff, Alberta, Canada, October 19-22, 2003.

[BBC$^+$98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An Architecture for Differentiated Services. IETF RFC 2475, December 1998.

[BH03] Marcus Brunner and Charlotte Hullo. GMPLS fault management and impact on service resilience differentiation. In *Proceedings of IFIP/IEEE Eighth International Symposium on Integrated Network Management*, pages 665–678, Colorado Springs, USA, March 24–28, 2003. Kluwer Academic Publishers.

[Bha98] Ramesh Bhandari. *Survivable Networks: Algorithms for Diverse Routing*. Kluwer Academic Publishers, Norwell, MA, USA, 1998.

[Bir03] G. M. Birtwistle. Demos - a system for discrete event modelling on Simula, 2003.

[BKLS01] E. Bouillet, K. Kumaran, G. Liu, and I. Saniee. Wavelength usage efficiency versus recovery time in path-protected DWDM mesh networks. In *Optical Fiber Communication Conference and Exhibit, 2001. OFC 2001*, volume 2, pages TuG1–1–TuG1–3 vol.2, 2001.

[BP81] Richad E. Barlow and Frank Proschan. *Statistical theory of reliability and life testing*. To begin with, 1981. Reprint of original edition from 1975.

[BRS04] Greg Bernstein, Bala Rajagopalan, and Debanjan Saha. *Optical Network Control. Architecture, Protocols, and Standards*. Addison-Wesley, Boston, MA, 2004.

[Cal04] Eusebi Calle. *Enhanced fault recovery methods for protected traffic services in GMPLS networks*. PhD thesis, Universitat de Girona, Department of Electronics, Computer Science and Automatic Control, February 2004.

[CdCRd06] Noélia Susana Costa Correia and Maria do Carmo Raposo de Medeiros. Protection Schemes for IP-over-WDM Networks: Throughput and Recovery Time Comparison. *Photonic Network Communications*, 11(2):127–149, March 2006.

[CDD$^+$02] Didier Colle, Sophie De Maesschalck, Chris Develder, Piet Van Heuven, Adelbert Groebbens, Jan Cheyns, Ilse Lievens, Mario Pickavet, Paul Lagasse, and Piet Demeester. Data-Centric Optical Networks and Their Survivability. *IEEE Journal on Selected Areas in Communications*, 20(1):6–19, January 2002.

[Cho05] Piotr Chołda. *The Reliability Analysis of Recovery Procedures in GMPLS-Based Optical IP Networks*. PhD thesis, AGH University of Science and Technology, Krakow, Poland, 2005.

[CJHM05] Piotr Chołda, Andrzej Jajszczyk, Bjarne E. Helvik, and Anders Mykkeltveit. Service differentiation based on recovery methods. In *EURO NGI Workshop on Traffic Engineering, Protection and Restoration*, Rome, Italy, April 21-22 2005.

[CJW05] P. Chołda, A. Jajszczyk, and K. Wajda. A unified framework for the assessment of recovery procedures. In *Workshop on High Performance Switching and Routing (HPSR)*, pages 269–273, 2005.

[ČKHG05] Tarik Čičić, Amund Kvalbein, Audun Fosselie Hansen, and Stein Gjessing. Resilient Routing Layers and *p*-Cycles: Tradeoffs in Network Fault Tolerance. In *Proc. IEEE 2005 Workshop on High Performance Switching and Routing HPSR 2005*, Hong Kong, China, May 12-14, 2005.

[CMH$^+$07] Piotr Chołda, Anders Mykkeltveit, Bjarne E. Helvik, Otto J. Wittner, and Andrzej Jajszczyk. A survey of resilience differentiation frameworks in communication networks. *IEEE Communications Surveys & Tutorials*, 9(4):32–55, October 2007.

[CMU04] Eusebi Calle, José L Marzo, and Anna Urra. Protection performance components in MPLS networks. *Elsevier Computer Communications*, 27(12):1220–1228, July 2004.

[CMUV03] Eusebi Calle, José L Marzo, Anna Urra, and Pere Vila. Enhancing MPLS QoS routing algorithms by using the network protection degree paradigm. In *Proceedings of IEEE Global Telecommunications Conference, (GLOBECOM '03)*, volume 6, pages 3053–3057, San Francisco, USA, December 1–5 2003.

[CSDC03] Roberto Clemente, Laura Serra, Giancarlo D'Orazio, and Giuseppe Cosmo. A framework for class of service definition in GMPLS-based meshed ASTN. In *Proceedings* *of Fourth International Workshop on Design of Reliable Communication Networks*, pages 93–100, Banff, Alberta, Canada, October 19–22, 2003.

[CWJ+05] P. Chołda, K. Wajda, A. Jajszczyk, J. Tapolcai, T. Cinkler, S. Bodamer, D. Colle, and G. Ferraris. Considerations about service differentiation using a combined QoS/QoR approach. In *Proceedings of the International Workshop on Design of Reliable Communication Networks (DRCN)*, pages 345–352, October 2005.

[DAR] DARPA: VINT Project. The network simulator - ns (version 2). http://nsnam.isi.edu/nsnam/.

[DCD98] Gianni Di Caro and Marco Dorigo. AntNet: Distributed stigmergetic control for communications networks. *Journal of Artificial Intelligence Research*, 9:317–365, 1998.

[DDTT03] Mathilde Durvy, Christophe Diot, Nina Taft, and Patrick Thiran. Network availability based service differentiation. *Lecture Notes in Computer Science*, 2707:305–324, January 2003.

[DGA+99] Piet Demeester, Michael Gryseels, Achim Autenrieth, Carlo Brianza, Laura Castagna, Giulio Signorelli, Roberto Clemente, M. Ravera, Andrzej Jajszczyk, Dariusz Janukowicz, Kristof Van Doorselaere, and Yohnosuke Harada. Resilience in multilayer networks. *IEEE Communications Magazine*, 37(8):70–76, 1999.

[DGH02] D. Dutta, A. Goel, and J. Heidemann. Faster network design with scenario pre-filtering. In *Proc. IEEE MASCOTS'02*, pages 237–246, October 11-16 2002.

[Dij59] E. W. Dijkstra. A note on two problems in connexion with graphs. *Numerische Mathematik*, 1(1):269–271, December 1959.

[DPF06] Song Dong, C. Phillips, and R. Friskney. Differentiated-resilience provisioning for the wavelength-routed optical network. *Lightwave Technology, Journal of*, 24(2):667–673, February 2006.

[FP01] S. Floyd and V. Paxson. Difficulties in simulating the internet. *IEEE/ACM Transactions on Networking*, 9(4):392–403, 2001.

[FT01a] A. Fumagalli and M. Tacca. Differentiated reliability (DiR) in WDM rings without wavelength converters. In *IEEE International Conference on Communications (ICC)*, volume 9, pages 2887–2891, 11-14 June 2001.

[FT01b] Andrea Fumagalli and Marco Tacca. Optimal design of optical ring networks with differentiated reliability (DiR). In *Proceedings of International Conference on Quality of Service in Multiservice IP Networks (QoS-IP)*, volume 1989 of *LNCS*, pages 299–313, Rome, Italy, January 2001. Springer.

[FT06] A. Fumagalli and M. Tacca. Differentiated reliability (DiR) in wavelength division multiplexing rings. *Networking, IEEE/ACM Transactions on*, 14(1):159–168, 2006.

[FTUF02] A. Fumagalli, M. Tacca, F. Unghvary, and A. Farago. Shared path protection with differentiated reliability. In *IEEE International Conference on Communications (ICC)*, volume 4, pages 2157–2161, 28 April-2 May 2002.

[FY94] Hiroyuki Fujii and Noriaki Yoshikai. Restoration Message Transfer Mechanism and Restoration Characteristics of Double-Search Self-Healing ATM Network. *IEEE Journal on Selected Areas in Communications*, 12(1):149–158, January 1994.

[GC02] W. D. Grover and M. Clouqueur. Span-restorable mesh network design to support multiple quality of protection (QoP) service classes. In *Proc. 1^{st} International Conference on Optical Communications and Networks ICOCN'02*, pages 305–308, Singapore, November 11-14, 2002.

[GC05] Wayne D. Grover and Matthieu Clouqueur. Span-restorable mesh networks with multiple quality of protection (QoP) service classes. *Photonic Network Communications*, V9(1):19–34, January 2005.

[GHSZ99] Paul Glasserman, Philip Heidelberger, Perwez Shahabuddin, and Tim Zajic. Multilevel splitting for estimating rare event probabilities. *Operations Research*, 47(4):585–600, 1999.

[GJS03] Janusz Gozdecki, Andrzej Jajszczyk, and Rafał Stankiewicz. Quality of Service Terminology in IP Networks. *IEEE Communications Magazine*, 41(3):153–159, March 2003.

[GNS00] N. Golmie, T.D. Ndousse, and D.H. Su. A differentiated optical services model for WDM networks. *Communications Magazine, IEEE*, 38(2):68–73, Feb. 2000.

[GR00] O. Gerstel and R. Ramaswami. Optical Layer Survivability: a Services Perspective. *IEEE Communications Magazine*, 38(3):104–113, March 2000.

[Gro04] Wayne D. Grover. *Mesh-Based Survivable Networks. Options and Strategies for Optical, MPLS, SONET, and ATM Networks.* Prentice Hall PTR, Upper Saddle River, NJ, 2004.

[GS01] Ornan Gerstel and Galen Sasaki. Quality of protection (qop): a quantitative unifying paradigm to protection service grades. In *Proc. Optical Networking and Communications Conference OptiComm*, Denver, CO, USA, August 21–23 2001.

[GS02] Ornan Gerstel and Galen Sasaki. Quality of Protection (QoP): A quantative unifying paradigm to protection service grades. *Optical Networks Magazine*, 3(3):40–49, May/June 2002.

[GSKG] D. Griffith, K. Sriram, S. Klink, and N. Golmie. Optimal Mixtures of Different Types of Recovery Schemes in Optical Networks.

[GT88] A. Goyal and A.N. Tantawi. A measure of guaranteed availability and its numerical evaluation. *IEEE Transactions on Computers*, 37(1):25–32, Jan. 1988.

[Hee95] Poul E. Heegaard. Speed-up techniques for high-performance evaluation. *Telektronikk*, 91(2/3):195–207, 1995.

[Hel04] Bjarne E. Helvik. Perspectives on the dependability of networks and services. *Telektronikk*, 100(3):27–44, 2004.

[HF06] Poul E. Heegaard and Ingebrigt Fuglem. Demonstrator 1: Ant-based monitoring on software IP routers. Deliverable D14 (IST-2001-38923), BISON, 2006.

[HLS07] Changcheng Huang, Minzhe Li, and Anand Srinivasan. A scalable path protection mechanism for guaranteed network reliability under multiple failures. *IEEE Transactions on Reliability*, 56(2):254–257, June 2007.

[HM03] Changcheng Huang and Donald Messier. Inter-domain MPLS restoration. In *Proceedings of Fourth International Workshop on Design of Reliable Communication Networks (DRCN)*, pages 341–348, Ottawa, Canada, October 19–22 2003.

[HMM97] H. Harai, M. Murata, and H. Miyahara. Performance of alternate routing methods in all-optical switching networks. In *Proceedings of IEEE INFOCOM*, volume 2, pages 516–524, April 7-11 1997.

[HS96] Bjarne E. Helvik and Norvald Stol. QoS differentiation in ATM networks; a case study. In *Proceedings of the 13th Nordic Teletraffic Seminar (NTS-13)*, pages 237–250, Trondheim, Norway, August 20–22 1996.

[HTC04] Pin-Han Ho, János Tapolcai, and Tibor Cinkler. Segment Shared Protection in Mesh Communications Networks with Bandwidth Guaranteed Tunnels. *IEEE/ACM Transactions on Networking*, 12(6):1105–1118, December 2004.

[HW05] Bjarne E. Helvik and Otto Wittner. Network resilience by emergent behavior from simple autonomous agents. In *Dependable Computing Systems: Paradigms, Performance Issues, and Applications*, chapter 17. Wiley, October 2005.

[ITU94] ITU-T. Terms and Definitions related to Quality of Service and Network Performance Including Dependability. ITU-T Rec. E.800, August 1994.

[ITU00] ITU-T. B-ISDN Semi-permanent Connection Availability. ITU-T Rec. I.357, November 2000.

[ITU02a] ITU-T. Framework of a Service Level Agreement. ITU-T Rec. E.860, June 2002.

[ITU02b] ITU-T. Internet Protocol Data Communication Service — IP Packet Transfer and Availability Performance Parameters. ITU-T Rec. Y.1540, December 2002.

[JMvA98] A. Jukan, A. Monitzer, and H.R. van As. QoS-restorability in optical networks. In *Optical Communication, 1998. 24th European Conference on*, volume 1, pages 711–712, 20-24 Sept. 1998.

[JMvA99] A. Jukan, A. Monitzer, and H.R. van As. Service-specific recovery of wavelength connections in WDM networks. In *Optical Fiber Communication Conference, 1999, and the International Conference on Integrated Optics and Optical Fiber Communication. OFC/IOOC '99. Technical Digest*, volume 1, pages 164–166, 21-26 Feb. 1999.

[JT03] Bjørn Jæger and David Tipper. Prioritized Traffic Restoration in Connection Oriented QoS based Networks. *Computer Communications*, 26(18):2025–2036, December 2003.

[Kel96] Christian Kelling. A framework for rare event simulation of Stochastic Petri Nets using "RESTART". In *Proceedings of the Winter Simulation Conference*, pages 317–324, Coronado, California, USA, December 1996.

[KG05] A. Kodian and W. D. Grover. Multiple-Quality of Protection Classes Including Dual-Failure Survivable Services in *p*-Cycle Networks. In *Proc. 2nd International Conference on Broadband Networks BROADNETS 2005*, Boston, MA, October 3-7, 2005.

[KO99] Ryutaro Kawamura and Hiroshi Ohta. Architectures for ATM Network Survivability and their Field Deployment. *IEEE Communications Magazine*, 37(8):88–94, August 1999.

[LDT07] P. L'Ecuyer, V. Demers, and B. Tuffin. Rare-events, splitting, and quasi-monte carlo. *ACM Transactions on Modeling and Computer Simulation*, 17(2), April 2007.

[LGS04] SuKyoung Lee, David Griffith, and Nah-Oak Song. A new analytical model of shared backup path provisioning in GMPLS networks. *Photonic Network Communications*, 4(3):271–283, 2004.

[LKL+03] Jae-Dong Lee, Sung-Un Kim, Sun-Seok Lee, Jae-Il Jung, and David H. Su. Differentiated wavelength assignment with QoS recovery for DWDM next generation internet backbone networks. *Photonic Network Communications*, 5(2):163–175, March 2003.

[LM04] Youngseok Lee and Biswanath Mukherjee. Traffic Engineering in Next-Generation Optical Networks. *IEEE Communications Surveys & Tutorials*, 6(3):16–33, July/September 2004.

[LS75] Stephen S. Lavenberg and Donald R. Slutz. Introduction to regenerative simulation. *IBM Journal of Research and Development*, 19(5):458–462, 1975.

[LSO05] H. Lønsethagen, A. Solem, and B. Olsen. easibility of bandwidth on demand. case study approach, models and issues. EU FP6 IP IST-NOBEL Project internal presentation, September 19–25 2005.

[Man04] Eric Mannie. Generalized Multi-Protocol Label Switching (GMPLS) Architecture. IETF RFC 3945, October 2004.

[MCSA03] José L. Marzo, Eusebi Calle, Caterina Scoglio, and Tricha Anjali. Adding QoS protection in order to enhance MPLS QoS routing. In *Proceedings of IEEE International Conference on Communications (ICC '03)*, volume 3, pages 1973–1977, Anchorage, Alaska, May 11–15 2003.

[MFB99] M. Médard, S. G. Finn, and R. A. Barry. Redundant Trees for Preplanned Recovery in Arbitrary Vertex-Redundant or Edge-Redundant Graphs. *IEEE/ACM Transactions on Networking*, 7(4):641–652, October 1999.

[MH07] Anders Mykkeltveit and Bjarne E. Helvik. Provision of connection-specific availability guarantees in communication networks. In *Proceedings of 6th International Workshop on Design of Reliable Communication Networks (DRCN)*, La Rochelle, France, October 2007.

[MH08] Anders Mykkeltveit and Bjarne E. Helvik. On provision of availability guarantees using shared protection. In *Proceedings of International Conference on Optical Network Design and Modeling (ONDM)*, pages 76–81, Vilanova i la Geltru, Spain, March 2008. IFIP.

[MIB+04] Athina Markopoulou, Gianluca Iannaccone, Supratik Bhattacharyya, Chen-Nee Chuah, and Christophe Diot. Characterization of failures in an IP backbone. In *Proceedings of IEEE INFOCOM*, Hong Kong, China, March 2004.

[MKMU05] T. Miyamura, T. Kurimoto, A. Misawa, and S. Urishidani. A Disjoint Path Selection Scheme Based on Enhanced Shared Risk Link Group Management for Multi-reliability Service. In *Proc. 2005 IEEE Global Telecommunications Conference GLOBECOM'05*, St. Louis, MO, 27 November – 2 December, 2005.

[ML05] F. Michaut and F. Lepage. Application-Oriented Network Metrology: Metrics and Active Measurement Tools. *IEEE Communications Surveys & Tutorials*, 7(2):2–24, April/June 2005.

[MLG05] C. Ming, Z. Luying, and M. Gurusamy. Dynamic Routing of Dependable Connections with Different QoP Grades in WDM Optical Networks. In *Proc. 10th IEEE Symposium on Computers and Communications ISCC'2005*, La Manga del Mar Menor, Cartagena, Spain, June 27-30, 2005.

[MM99] G. Mohan and C. Siva Ram Murthy. Routing and Wavelength Assignment for Establishing Dependable Connections in WDM Networks. In *Proc. 29th International Symposium on Fault-Tolerant Computing FTCS-29*, Madison, WI, June 15-18, 1999.

[MPR$^+$05] D.A.A. Mello, J.U. Pelegrini, R.P. Ribeiro, D.A. Schupke, and H. Waldman. Dynamic provisioning of shared-backup path protected connections with guaranteed availability requirements. In *Proceedings of International Conference on Broadband Networks*, volume 2, pages 1320–1327, 2005.

[MS00] G. Mohan and A. K. Somani. Routing Dependable Connections with Specified Failure Restoration Guarantees in WDM Networks. In *Proc. IEEE Conference on Computer Communications INFOCOM 2000*, Tel Aviv, Israel, March 26-30, 2000.

[MSRMS01] G. Mohan, C. Siva Ram Murthy, and A.K. Somani. Efficient algorithms for routing dependable connections in WDM optical networks. *Networking, IEEE/ACM Transactions on*, 9(5):553–566, October 2001.

[MZM05] P. Ma, L. Zhou, and G. Mohan. Reliability and Recovery Time Differentiated Routing in WDM Optical Networks. In *Proc. 2005 IEEE Global Telecommunications Conference GLOBECOM'05*, St. Louis, MO, 27 November – 2 December, 2005.

[NG99] Thomas D. Ndousse and Nada Golmie. Differentiated Optical Services: a Quality of Optical Service Model for WDM Networks. In *Proc. SPIE Conference in All-Optical Networking 1999: Architecture, Control, and Management Issues*, Boston, MA, September 19-21, 1999.

[NM04] Hassan Naser and Hussein T. Mouftah. A multilayer differentiated protection services architecture. *IEEE Journal on Selected Areas in Communications*, 22(8):1539–1547, October 2004.

[NSN01] V.F. Nicola, P. Shahabuddin, and M.K. Nakayama. Techniques for fast simulation of models of highly dependable systems. *IEEE Transactions on Reliability*, 50(3):246–264, 2001.

[NSO$^+$95] Leo Nederlof, Kris Struyve, Chris O'Shea, Howard Misser, Yonggang Du, and Braulio Tamayo. End-to-end Survivable Broadband Networks. *IEEE Communications Magazine*, 33(9):63–70, September 1995.

[NTBT04] Antonio Nucci, Nina Taft, Chadi Barakat, and Patrick Thiran. Controlled use of excess backbone bandwidth for providing new services in IP-over-WDM networks. *IEEE Journal on Selected Areas in Communications*, 22(9):1692–1707, November 2004.

[Ogg01] C. Oggerino. *High Availability Network Fundamentals*. Cisco Press, Indianapolis, IN, 2001.

[OM04] C. Ou and B. Mukherjee. Differentiated Quality-of-Protection Provisioning in Optical/MPLS Networks. In *Proc. 3rd International IFIP-TC6 Networking Conference NETWORKING 2004*, Athens, Greece, May 9-14, 2004.

[ORM05] C. Ou, S. Rai, and B. Mukherjee. Extension of Segment Protection for Bandwidth Efficiency and Differentiated Quality of Protection in Optical/MPLS Networks. *Optical Switching and Networking*, 1(1):19–33, January 2005.

[OSB05] H. Overby, N. Stol, and S. Bjornstad. Dependability differentiation in optical packet switched networks. In *Proceedings of 7th International Conference on Transparent Optical Networks (ICTON)*, volume 1, pages 385–388, 2005.

[OSH91] Y. Okanoue, H. Sakauchi, and S. Hasegawa. Design and Control Issues of Integrated Self-Healing Networks in SONET. In *Proc. 1991 IEEE Global Telecommunications Conference GLOBECOM'91*, Phoenix, AZ, December 2-5, 1991.

[OW03] Sebastian Orlowski and Roland Wessäly. Comparing Restoration Concepts using Optimal Network Configurations with Integrated Hardware and Routing Decisions. In *Proc. 4th International Workshop on the Design of Reliable Communication Networks DRCN 2003*, pages 39–46, Banff, Alberta, Canada, October 19-22, 2003.

[OYP95] Eiji Oki, Naoaki Yamanaka, and Francis Pitcho. Multiple-availability-level ATM network architecture. *IEEE Communications Magazine*, 33(9):80–88, September 1995.

[PK03] P. Pongpaibool and H.S. Kim. Novel algorithms for dynamic connection provisioning with guaranteed service level agreements in IP-over-optical networks. In *Proc. IEEE GLOBECOM*, volume 5, pages 2643–2648, 2003.

[PKS02] Jigesh K. Patel, Sung U. Kim, and David H. Su. Qos recovery schemes based on differentiated MPLS services in all-optical transport next generation internet. *Photonic Network Communications*, 4(1):5–18, January 2002.

[PM04] Michał Pióro and Deepankar Medhi. *Routing, Flow and Capacity Design in Communication and Computer Networks*. Morgan Kaufmann Publishers — Elsevier, San Francisco, CA, 2004.

[PM06] Dimitri Papadimitriou and Eric Mannie. Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanism (including Protection and Restoration). IETF RFC 4428, March 2006.

[PR02] A. Paradisi and S.M. Rossi. Differentiated reliability (DiR) in mesh networks with shared path protection: theoretical and experimental results. In *Optical Fiber Communication Conference and Exhibit, 2002. OFC 2002*, pages 490–492, 17-22 Mar 2002.

[PTF05] Z. Pandi, M. Tacca, and A. Fumagalli. A threshold based on-line RWA algorithm with end-to-end reliability guarantees. In *Proceedings of International Conference on Optical Network Design and Modeling (ONDM)*, pages 447–453, Feb. 7-9, 2005.

[RJBS04] F. Rosenbaum, S. Jha, P. Boustead, and F. Safei. Resilience-Differentiation in Programmable Virtual Networks. In *Proc. IEEE International Conference on Communications ICC 2004*, Paris, France, June 20-24, 2004.

[RLaBP03] Fabio Ricciato, Marco Listanti, angelo Belmonte, and Daniele Perla. Performance evaluation of a distributed scheme for protection against single and double faults for MPLS. In *nternational Workshop on Quality of Service in Multiservice IP Networks (QoS-IP)*, LNCS, pages 218–232. Springer, January 2003.

[RLS04] Fabio Ricciato, Marco Listanti, and Stefano Salsano. An architecture for differentiated protection against single and double faults in GMPLS. *Photonic Network Communications*, 8(1):119–132, 2004.

[RM99] S. Ramamurthy and Biswanath Mukherjee. Survivable WDM Mesh Networks. Part II — Restoration. In *Proc. IEEE International Conference on Communications ICC'99*, Vancouver, Canada, June 6-10, 1999.

[RMA05] Fabio Ricciato, Ugo Monaco, and Daniele Ali. Distributed Schemes for Diverse Path Computation in Multidomain MPLS Networks. *IEEE Communications Magazine*, 43(6):138–146, June 2005.

[Rob04] James W. Roberts. Internet Traffic, QoS, and Pricing. *Proceedings of the IEEE*, 92(9):1389–1399, 2004.

[RR01] Daniel Rossier-Ramuz. Dynamic protection set-up in optical VPN using mobile agent ecosystem. In *Proceedings of International Workshop on the Desin of Reliable Communication Networks (DRCN 2001)*, pages 189–196, Budapest, Hungary, October 2001.

[RR02] Daniel Rossier-Ramuz. *Towards Active Network Management with Ecomobile, an Ecosystem-inspired Mobile Agent Middleware*. PhD thesis, University of Freiburg, Freiburg, Switzerland, October 2002.

[RS94] K. Ramamritham and J. A. Stankovic. Scheduling algorithms and operating systems support for real-time systems. *Proceedings of the IEEE*, 82(1):55–67, 1994.

[RS01] D. Raz and Y. Shavitt. Toward efficient distributed network management. *Journal of Network and Systems Management*, 9(3):347–361, September 2001.

[RSM03] S. Ramamurthy, L. Sahasrabuddhe, and B. Mukherjee. Survivable wdm mesh networks. *Lightwave Technology, Journal of*, 21(4):870–883, April 2003.

[Rub99] Reuven Y. Rubinstein. The cross-entropy method for combinatorial and continous optimization. *Methodology and Computing in Applied Probability*, 1(2):127–190, September 1999.

[SAG06] B. Sansò, C. Awad, and A. Girard. Can DiffServ Guarantee IP QoS Under Failures? *IEEE Network*, 20(4):32–40, July/August 2006.

[SB08] M. Shand and S. Bryant. IP fast reroute framework. Internet-Draft draft-ietf-rtgwg-ipfrr-framework-08.txt, February 2008.

[Sch06] Matthias Scheffel. Adaptation of failure scenario based resilience schemes toward availabilityguarantees. *Journal of Optical Networking*, 5(7):521–531, July 2006.

[SGZ04] Chava Vijaya Saradhi, Mohan Gurusamy, and Luying Zhou. Differentiated QoS for Survivable WDM Optical Networks. *IEEE Optical communications*, 42(5):S8–S14, May 2004.

[SH03] Vishal Sharma and Fiffi Hellstrand. Framework for Multi-Protocol Label Switching (MPLS)-based Recovery. IETF RFC 3469, February 2003.

[SHBR97] Ruud Schoonderwoerd, Owen Holland, Janet Bruten, and Leon Rothkrantz. Ant-based load balancing in telecommunications networks. *Adaptive Behaviour*, 5(2):169–207, 1997.

[SM02a] Chava Vijaya Saradhi and C. Siva Ram Murthy. A Framework for Differentiated Survivable Optical Virtual Private Networks. *Photonic Network Communications*, 4(3/4):457–487, July 2002.

[SM02b] Chava Vijaya Saradhi and C. Siva Ram Murthy. Routing Differentiated Reliable Connections in WDM Optical Networks. *Optical Networks Magazine*, 3(3):50–67, May/June 2002.

[SM04] Chava Vijaya Saradhi and C. Siva Ram Murthy. Dynamic establishment of differentiated survivable lightpaths in WDM mesh networks. *Computer Communications*, 27(3):273–294, February 2004.

[SND] SNDlib. Library of test instances for Survivable fixed telecommunication Network Design.

[SOB+06] Norvald Stol, Harald Overby, Steinar Bjrnstad, Andreas Kimss, and Anders Mykkeltveit. Differentiated survivability in the OpMiGua hybrid optical network. In *Proceedings of Conference on Optical Network design and Modeling (ONDM)*, May 2006.

[Som05] Arun K. Somani. *Survivability and Traffic Grooming in WDM Optical Networks*. Cambridge University Press, Cambridge, UK, 2005.

[SS00] Murari Sridharan and Arun K. Somani. Revenue maximization in survivable WDM networks. In *Optical Networking and Communications Conference (OptiComm)*, volume 4233, pages 291–302, Richardson, TX, USA, 2000. SPIE.

[SS03a] G. Sahin and S. Subramaniam. Online Control-Message Scheduling fir Quality of Protection (QoP) in DWDM Mesh Networks. In *Proc. Optical Fiber Communication Conference and Exhibition OFC 2003*, Atlanta, GA, March 23-28, 2003.

[SS03b] Gokan Sahin and Suresh Subramaniam. Quality of protection through control-message scheduling in optical mesh networks. In *Proceedings of International Workshop on the Desin of Reliable Communication Networks (DRCN 2003)*, pages 39–46, Banff, Alberta, Canada, October 19–22 2003.

[SS04] G. Sahin and S. Subramaniam. Providing Quality-of-Protection Classes Through Control-Message Scheduling in DWDM Mesh Networks with Capacity Sharing. *IEEE Journal on Selected Areas in Communications*, 22(9):1846–1858, November 2004.

[Suu74] J. W. Suurballe. Disjoint paths in a network. *Networks*, 4:125–145, 1974.

[SZM05] L. Song, J. Zhang, and B. Mukherjee. Dynamic Provisioning with Reliability Guarantee and Resource Optimization for Differentiated Services in WDM Mesh Networks. In *Proc. Optical Fiber Communication Conference and Exhibit OFC 2005*, Anaheim, CA, March 6-11, 2005.

[Tan03] Andrew S. Tanenbaum. *Computer Networks*. Prentice-Hall, fourth edition, 2003.

[TCC+05] J. Tapolcai, P. Chołda, T. Cinkler, K. Wajda, A. Jajszczyk, A. Autenrieth, S. Bodamer, D. Colle, G. Ferraris, H. Lonsethagen, I.-E. Svinnset, and D. Verchere. Quality of resilience (QoR): NOBEL approach to the multi-service resilience characterization. In *Broadband Networks, 2005 2nd International Conference on*, pages 405–414, 2005.

[TCC+06] J. Tapolcai, Piotr Chołda, T. Cinkler, K. Wajda, A. Jajszczyk, and D. Verchere. Joint quantification of resilience and quality of service. In *Proc. IEEE International Conference on Communications ICC 2006*, Istanbul, Turkey, June 11-15 2006.

[TFP+03] Marco Tacca, Andrea Fumagalli, Alberto Paradisi, Ferenc Unghvary, Kishore Gadhiraju, Sanjeev Lakshmanan, Sandro Marcelo Rossi, Antonio de Campos Sachs, and Dhruvish S. Shah. Differentiated reliability in optical networks: Theoretical and practical results. *Journal of Lightwave Technology*, 21(11):2576–2586, November 2003.

[TFU03] M. Tacca, A. Fumagalli, and F. Unghvary. Double-fault shared path protection scheme with constrained connection downtime. In *Proceedings. Fourth International Workshop on Design of Reliable Communication Networks (DRCN)*, pages 181–188, 19-22 Oct. 2003.

[TMF04] M. Tacca, P. Monti, and A. Fumagalli. The disjoint path-pair matrix approach for online routing in reliable WDM networks. In *IEEE International Conference on Communications (ICC)*, volume 2, pages 1187–1191, 20-24 June 2004.

[TN94] M. To and P. Neusy. Unavailability analysis of long-haul networks. *IEEE Journal on Selected Areas in Communications*, 12(1):100–109, Jan. 1994.

[TTD$^+$01] P. Thiran, N. Taft, C. Diot, H. Zang, and R. Mac Donald. A Protection-Based Approach to QoS in Packet over Fiber Networks. In *Thyrrhenian International Workshop on Digital Communications*, volume 2170, pages 266–278, 2001.

[VA98] José Villén-Altamirano. RESTART method for the case where rare events can occur in retrials from any threshold. *International Journal of Electronics and Telecommunication (AEU)*, 52(3):183–189, 1998.

[VAVA91] Manuel Villén-Altamirano and José Villén-Altamirano. RESTART: A method for accelerating rare event simulations. In *Proceedings of 13th International Teletraffic Congress*, pages 71–76, Copenhagen, Denmark, 1991. North-Holland.

[VAVA06] Manuel Villén-Altamirano and José Villén-Altamirano. On the efficiency of RESTART for multidimensional state systems. *ACM Transactions on Modeling and Computer Simulation*, 16(3):251–279, 2006.

[VCD$^+$05] S. Verbrugge, D. Colle, P. Demeester, R. Huelsermann, and M. Jaeger. General availability model for multilayer transport networks. In *Proceedings of International Workshop on Design of Reliable Communication Networks (DRCN)*, pages 85–92, October 16–19 2005.

[Ver04] D.C Verma. Service level agreements on IP networks. *Proceedings of the IEEE*, 92(9):1382–1388, 2004.

[VHS96] Paul Veitch, Ian Hawker, and Geoffrey Smith. Administration of restorable virtual path mesh networks. *IEEE Communications Magazine*, 34(12):96–101, December 1996.

[VPD04] Jean-Philippe Vasseur, Mario Pickavet, and Piet Demeester. *Network Recovery. Protection and Restoration of Optical, SONET-SDH, IP, and MPLS*. Morgan Kaufmann Publishers, San Francisco, CA, 2004.

[VVA02] Manuel Villn-Altamirano and José Villén-Altamirano. Analysis of RESTART simulation: Theoretical basis and sensitivity study. *European Transactions on Telecommunications*, 13(4):373–385, 2002.

[VVV94] Thierry Van Landegem, Patrick Vankwikelberge, and Hans Vanderstraeten. A Self-Healing ATM Network Based on Multilink Principles. *IEEE Journal on Selected Areas in Communications*, 12(1):149–158, January 1994.

[WBYM04] Wushao Wen, S. J. Ben Yoo, and Biswanath Mukherjee. Quality-of-Service Based Protection in MPLS Control WDM Mesh Networks. *Photonic Network Communications*, 4(3):297–320, 2004.

[WH02] Otto Wittner and Bjarne E. Helvik. Cross-Entropy Guided Ant-like Agents Finding Dependable Primary/Backup Path Patterns in Networks. In *Proceedings of Congress on Evolutionary Computation (CEC2002)*, Honolulu, Hawaii, May 12-17 2002.

[WVF03] Kai Wu, L. Valcarenghi, and A. Fumagalli. Restoration schemes with differentiated reliability. In *Communications, 2003. ICC '03. IEEE International Conference on*, volume 3, pages 1968–1972, 11-15 May 2003.

[WZW04] Wei Wei, Qingji Zeng, and Yun Wang. Multi-layer differentiated integrated survivability for optical internet. *Photonic Network Communications*, 8(3):267–284, 2004.

[WZWY04] Wei Wei, Qingji Zeng, Hongquan Wei, and Hongtao Yu. Integrated Survivable QoS Routing in Metro IP/WDM Networks. In *Proc. 13th IEEE Workshop on Local and Metropolitan Area Networks LANMAN 2004*, San Francisco Bay Area, CA, April 25-28, 2004.

[XCT03] Guoliang Xue, Li Chen, and Krishnaiyan Thulasiraman. Quality-of-service and quality-of-protection issues in preplanned recovery schemes using redundant trees. *IEEE Journal on Selected Areas in Communications*, 21(8):1332–1345, October 2003.

[XYWL04a] Bing Xiang, Hongfang Yu, Sheng Wang, and Lemin Li. A Differentiated Shared Protection Algorithm Supporting Traffic Grooming in WDM Networks. In *Proc. 2004 International Conference on Communications, Circuits and Systems ICCCAS 2004*, Chengdu, China, June 27-29, 2004.

[XYWL04b] Bing Xiang, Hongfang Yu, Sheng Wang, and Lemin Li. A QoS-based Differentiated Protection Algorithm in WDM Mesh Networks. In *Proc. 2004 International Conference on Communications, Circuits and Systems ICCCAS 2004*, Chengdu, China, June 27-29, 2004.

[Yan05] X. Yang. Availability-Differentiated Service Provisioning in Free-space Optical Access Networks. *OSA Journal of Optical Networking*, 4(7):391–399, July 2005.

[YK97] T. Yahara and R. Kawamura. Virtual path self-healing scheme based on multi-reliability ATM network concept. In *Global Telecommunications Conference, 1997. GLOBECOM '97., IEEE*, volume 3, pages 1803–1807, 1997.

[YKO98] T. Yahara, R. Kawamura, and S. Ohta. Multi-Reliability Self-Healing Scheme that Guarantees Minimum Cell Rate. In *Proc. 1st International Workshop on the Design of Reliable Communication Networks DRCN 1998*, Brugge, Belgium, May 19-20, 1998.

[YR04] Wang Yao and B. Ramamurthy. Survivable traffic grooming with differentiated end-to-end availability guarantees in WDM mesh networks. In *Local and Metropolitan Area Networks, 2004. LANMAN 2004. The 13th IEEE Workshop on*, pages 87–90, 2004.

[YSJ$^+$04] M.-R. Yoon, J.-D. Shin, C.-H. Jeong, J.-M. Jo, O.-H. Kang, and S.-U. Kim. Optical-LSP Establishment and a QoS Maintenance Scheme Based on Differentiated Optical QoS Classes in OVPNs. *Photonic Network Communications*, 7(2):161–178, March 2004.

[ZD02] H. Zhang and A. Durresi. Differentiated multi-layer survivability in IP/WDM networks. In *IEEE/IFIP Network Operations and Management Symposium, (NOMS 2002)*, pages 681–694, Florence, Italy, April 15–19 2002.

[ZG05] Ling Zhou and W. D. Grover. A theory for setting the "safety margin" on availability guarantees in an SLA. In *Proceedings of International Workshop on Design of Reliable Communication Networks (DRCN)*, pages 403–409, 2005.

[ZZM03] J. Zhang, K. Zhu, and B. Mukherjee. Service Provisioning to Provide Per-Connection-based Availability Guarantee in WDM Mesh Networks. In *Proc. Optical Fiber Communication Conference and Exhibition OFC 2003*, Atlanta, GA, March 23-28, 2003.

[ZZZM03] Jing Zhang, Keyao Zhu, Hui Zang, and B. Mukherjee. A new provisioning framework to provide availability-guaranteed service in WDM mesh networks. In *IEEE International Conference on Communications (ICC)*, volume 2, pages 1484–1488, May 11-15 2003.