UNIVERSIDADE FEDERAL DO PARÁ
INSTITUTO DE TECNOLOGIA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

TESE DE DOUTORADO

EXPLORING SERVICE RELOCATION AND DIFFERENTIATION TO IMPROVE SURVIVABILITY ON OPERATION OF RESILIENT OPTICAL CLOUD NETWORKS

CARLOS NATALINO DA SILVA

TD 09/2016

UFPA / ITEC / PPGEE
BELÉM/PA
2016

UNIVERSIDADE FEDERAL DO PARÁ
INSTITUTO DE TECNOLOGIA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

CARLOS NATALINO DA SILVA

EXPLORING SERVICE RELOCATION AND DIFFERENTIATION TO IMPROVE SURVIVABILITY ON OPERATION OF RESILIENT OPTICAL CLOUD NETWORKS

UFPA / ITEC / PPGEE
BELÉM/PA
2016

UNIVERSIDADE FEDERAL DO PARÁ
INSTITUTO DE TECNOLOGIA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

CARLOS NATALINO DA SILVA

EXPLORING SERVICE RELOCATION AND DIFFERENTIATION TO IMPROVE SURVIVABILITY ON OPERATION OF RESILIENT OPTICAL CLOUD NETWORKS

Tese de doutorado submetida à Banca Examinadora do Programa de Pós-Graduação em Engenharia Elétrica da UFPA para a obtenção do Grau de Mestre/Doutor em Engenharia Elétrica na área de Computação Aplicada.

UFPA / ITEC / PPGEE
BELÉM/PA
2016

Dados Internacionais de Catalogação - na – Publicação (CIP) Sistema de Bibliotecas da UFPA

Silva, Carlos Natalino da, 1987-

Exploration service relocation and differentiation to improve survivability on operation of resilient optical cloud networks. / Carlos Natalino da Silva. - 2016.

Orientador: Carlos Renato Lisboa Francês

Tese (Doutorado) - Universidade Federal do Pará, Instituto de Tecnologia, Programa de Pós-Graduação em Engenharia Elétrica , Belém, 2016.

1. Computação em nuvem. 2. Comunicações óticas. 3. Recursos de redes de computadores. 4. Serviços da web. I. Título.

CDD 23. ed. 004.6782

UNIVERSIDADE FEDERAL DO PARÁ
INSTITUTO DE TECNOLOGIA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

"EXPLORING SERVICE RELOCATION AND DIFFERENTIATION TO IMPROVE SURVIVABILITY ON OPERATION OF RESILIENT OPTICAL CLOUD NETWORKS"

AUTOR: **CARLOS NATALINO DA SILVA**

TESE DE DOUTORADO SUBMETIDA À BANCA EXAMINADORA APROVADA PELO COLEGIADO DO PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA, SENDO JULGADA ADEQUADA PARA A OBTENÇÃO DO GRAU DE DOUTOR EM ENGENHARIA ELÉTRICA NA ÁREA DE COMPUTAÇÃO APLICADA.

APROVADA EM: 17/03/2016

BANCA EXAMINADORA:

Prof. Dr. Carlos Renato Lisboa Francês
(Orientador – PPGEE/UFPA)

Prof. Dr. João Crisóstomo Weyl Albuquerque Costa
(Avaliador Interno – PPGEE/UFPA)

Prof. Dr. Marcelo Eduardo Vieira Segatto
(Avaliador Externo – UFES)

Prof. Dr. Nandamudi Lankalapalli Vijaykumar
(Avaliador Externo – INPE)

Prof. Dr. Solon Venâncio De Carvalho
(Avaliador Externo – INPE)

VISTO:

Prof. Dr. Evaldo Gonçalves Pelaes
(Coordenador do PPGEE/ITEC/UFPA)

Contents

List of Figures

List of Tables

List of abbreviations and acronyms

ASON Automatically Switched Optical Network

C-RAN Cloud Radio Access Networks or Centralized Radio Access Networks

CCC Command and Control Center

CSP Cloud Service Provider

CPU Central Processing Unit

DC DataCenter

DPP Dedicated Path Protection

GMPLS Generalized Multiprotocol Label Switching

HRP Heuristic for Relocation with Priorities

IaaS Infrastructure as a Service

ICT Information and Communication Technology

ILP Integer Linear Programming

IRP ILP for Relocation with Priorities

ISP Internet Service Provider

IT Information Technology

JVM Java Virtual Machine

MILP Mixed Integer Linear Programming

MTTF Mean Time To Failure

MTTR Mean Time To Repair

NCP Network Control Plane

OS Operating System

PaaS Platform as a Service

PCE Path Computation Element

PR-SRD Path Restoration with Service Relocation and Differentiation

QoE Quality of Experience

QoS Quality of Service

RAM Random Access Memory

RNG Random Number Generator

RWA Routing and Wavelength Assignment

SDN Software Defined Network

SLA Service Level Agreement

SML Service Middleware Layer

SPP Shared Path Protection

SRG Shared Risk Group

SLRG Shared Link Risk Group

VM Virtual Machine

WAN Wide Area Network

VON Virtual Optical Network

WDM Wavelength Division Multiplexing

List of Papers

List of publications included in this thesis:

I. **Natalino, Carlos**; Ahmed, Jawwad; Monti, Paolo; Wosinska, Lena; Francês, Renato. A relocation-based heuristic for restoring optical cloud services. In: *2014 13th International Conference on Optical Communications and Networks (ICOCN)*, p. 1-4, 2014, Suzhou, China.

II. **Natalino, Carlos**; Wosinska, Lena; Spadaro, Salvatore; Costa, João C. W. A.; Francês, Carlos Renato L.; Monti, Paolo;. Restoration in Optical Cloud Networks with Relocation and Services Differentiation. *Journal of Optical Communications and Networking*, vol. 8, n. 2, p. 100-111, 2016.

List of publications not included in this thesis:

1. Jailton, José; Carvalho, Tássio; Valente, Warley; **Natalino, Carlos**; Francês, Renato; Dias, Kelvin. A quality of experience handover architecture for heterogeneous mobile wireless multimedia networks. *IEEE Communications Magazine*, v. 51, p. 1, 2013.
2. Carvalho, Tássio; Jailton, José; Valente, Warley; **Natalino, Carlos**; Francês, Renato; Lopes, Kelvin. A Mobile WiMAX Mesh Network with Routing Techniques and Quality of Service Mechanisms. In: *Dr. Gianni Pasolini. (Org.). Selected Topics in WiMAX*. 1ed.: InTech, 2013.
3. Almeida, Luciana Abdon; Monteiro, Flávia Pessoa; Santana, Ádamo Lima; **Natalino, Carlos**; Francês, Carlos Renato; Cardoso, Diego Lisboa. Green-fuzzy - An intelligent user allocation model for macro-femto co-channel networks. In: *2013 SBMO/IEEE MTTS International Microwave and Optoelectronics Conference (IMOC)*, 2013, Rio de Janeiro.
4. **Natalino, Carlos**; Monti, Paolo; França, Luis; Furdek, Marija; Wosinska, Lena; Francês, Carlos R. L.; Costa, João C. W. A. Dimensioning Optical Clouds with Shared-Path Shared-Computing (SPSC) Protection. *2015 IEEE 16th International Conference on High Performance Switching and Routing (HPSR)*, 2015, Budapest, Hungary.
5. **Natalino, Carlos**; Chiaraviglio, Luca; Idzikowski, Filip; Monti, Paolo; Listanti, Marco; Francês, Carlos R. L.; Wosinska, Lena. Lifetime-Aware Provisioning in Green Optical Backbone Networks. *Optical Fiber Communication Conference and Exhibition (OFC)*, 2016, Anaheim, California, USA.

Resumo

Nuvens ópticas são a junção de redes ópticas de transporte com computação em nuvem, permitindo o gerenciamento integrado de ambas as infraestruturas por um único controlador. Este gerenciamento integrado permite o surgimento de novos modelos de serviço como por exemplo demandas *anycast* formadas por uma demanda para a rede óptica (requisitando recursos de rede) e uma demanda de serviço para a nuvem (requisitando recursos de processamento e armazenamento) no mesmo ciclo de vida. Este novo paradigma permite a migração de serviços (habilidade de migrar um serviço que está em um *datacenter* para outro), o qual deve ser utilizado com moderação pois possuem execução complexa e demorada, podendo causar sobrecarga na rede ou *datacenters* envolvidos na operação. Atualmente trabalhos da literatura consideram um único modelos de demanda de serviço, o que pode não representar corretamente a variedade de demandas esperadas para serviços em nuvem. Neste contexto, propriedades e características de serviços afetados por falhas, como por exemplo a prioridade ou o tempo restante do serviço, podem ser estudados visando a melhoria da sobrevivência de tais serviços durante o processo de restauração de conexões. Esta tese propõe o estudo de nuvens ópticas e os impactos causados pelo uso de migração de serviços na sobrevivência destas nuvens. Para avaliar estes impactos, um cenário de simulação de nuvens ópticas é proposto, visando aumentar a heterogeneidade dos serviços considerados. O processo de restauração de serviços afetados por falhas na rede é estudado e novas estratégias para restauração de serviços capazes de migrar serviços entre *datacenters* são propostas visando aumentar a resiliência destas nuvens a falhas. As propostas são avaliadas em um conjunto de diferentes cenários, e comparadas com estratégias da literatura. Resultados mostram uma considerável melhoria na resiliência de nuvens ópticas quando usando as estratégias propostas.

Palavras-chaves: Comunicações óticas. Computação em nuvem. Nuvens óticas. Operação de nuvens óticas. Sobrevivência de redes óticas. Recuperação de falhas. Migração de serviços.

Abstract

Optical clouds are the combination of optical transport networks and cloud computing, which allows the integrated management of both infrastructures in one controller element. In this paradigm, cloud services can be provisioned in an *anycast* fashion, i.e., only the source node asking for a service and the amount of IT resources are specified, while it is up to the cloud control/management system to select the most suitable destination datacenter (DC) node. During the cloud service provisioning process resiliency is crucial in order to guarantee continuous network operations also in the presence of failures. On one hand, a survivability strategy needs to be able to meet the availability requirements of each specific cloud service, while on the other hand it must be efficient in using backup resources. Service relocation (i.e., the ability to live re-allocate one provisioned service to another DC) is one of the new features that can be used in this new paradigm, but needs to be applied carefully given its associated overhead may overload the network and DCs. Current works in the literature consider a single service model that lack representation of the heterogeneity expected for cloud services. In this context, some disrupted connection properties can be considered to improve the network survivability during the restoration process, e.g., priorities and service remaining holding time. This thesis proposes a restoration-based survivability strategy, which combines the benefits of both cloud *service relocation* and *service differentiation* concepts. The former is used to enhance the restorability performance (i.e., the percentage of successfully restored cloud services) offered by restoration, while the latter ensures that critical services are given the proper consideration while backup resources are assigned. An ILP and a heuristic are presented in order to solve the proposed survivability strategy. The proposed strategies are evaluated considering a variety of different simulation scenarios. Results show that performance achieved by the proposed ILP and heuristic are close to the ones achieved when using protection strategies, but with the inherent benefits in terms of efficient use of resources offered by restoration-based approaches.

Key-words: Cloud computing. Optical clouds. Optical clouds operation. Network survivability. Failure recovery. Service relocation. Service Availability. Service Differentiation.

1 Introduction

The combination of high capacity DataCenters (DCs) with high capacity optical networks interconnecting them allows the emergence of optical clouds concept (DEVELDER et al., 2012). Advances on both cloud computing (e.g., elasticity, scalability, customization) and optical networks (e.g., network virtualization, software defined networks) are fundamental parts for building an efficient optical cloud.

The “cloud” denotes a computing infrastructure based on on-demand model, in which businesses and individual users can have access to applications from any part of the world at anytime (VOORSLUYS; BROBERG; BUYYA, 2011). Processing, storage and networking capacities are the main components of the aforementioned infrastructure. Processing and storage are spread over DCs interconnected by high-speed optical networks. Figure 1 shows an overview of the usual network architecture used to provide internal communication between DCs located at the same Cloud Service Provider (CSP) and external communication between users and the DCs that host the services. From the user point of view, cloud is seen as a single entity accessible by its ISP connection, but Fig. 1 shows that this cloud can be understood as a “cloud of clouds”, with each provider (e.g., ISP of CSP) owning its own cloud.

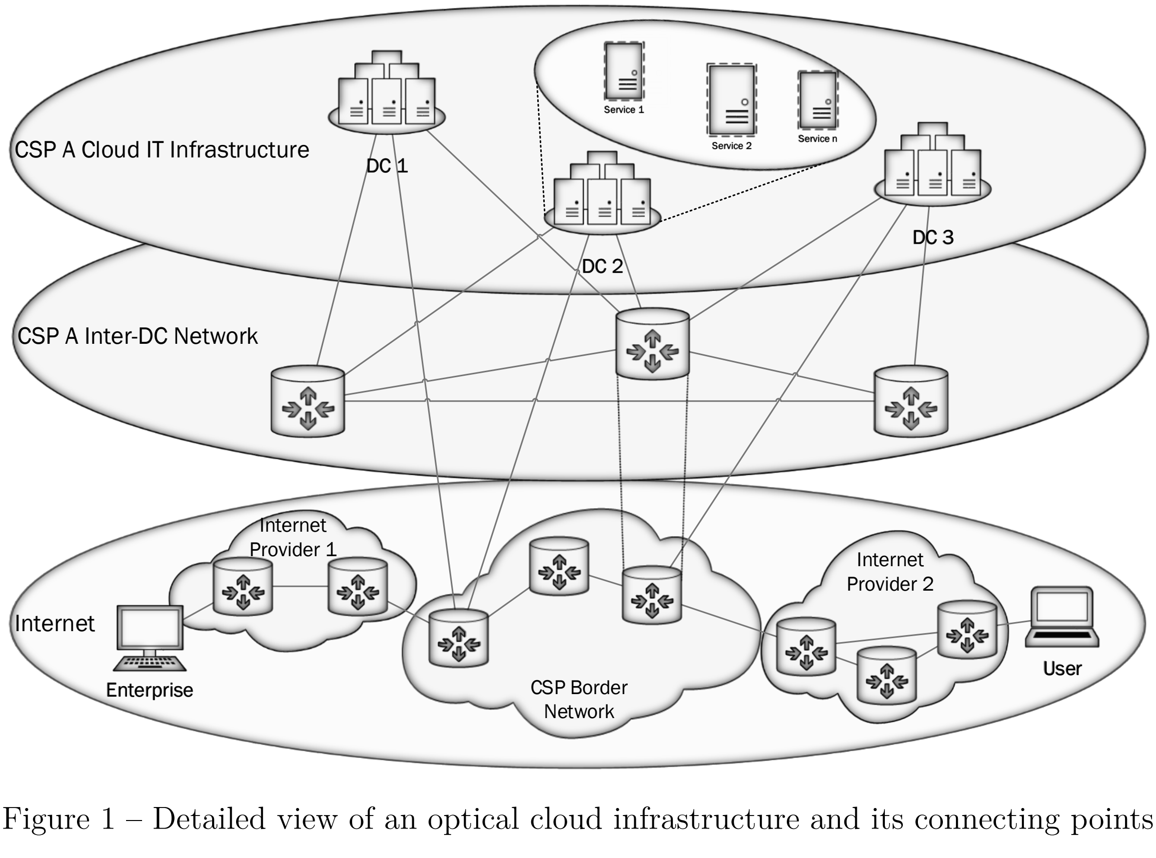

Figure 1 – Detailed view of an optical cloud infrastructure and its connecting points

This set of resources available in optical clouds can be analyzed from different points of view, i.e., user, enterprise and cloud service provider. From the user point of view, i.e., application user, an optical cloud is the source of services that are available through an Internet connection. From the enterprise point of view, i.e., business side, an optical cloud is an infrastructure in which services can be deployed and made available for users inside or outside the companies (e.g., customers). From the CSP point of view, an optical cloud is its main product, an owned infrastructure that is used to attend the enterprises (also referred as CSPs customers, in this case) and help them to deploy services to its users or costumers (BILAL et al., 2014).

When a cloud service is deployed, the client makes a request to the CSP for Information Technology (IT) and network resources. Usually IT resources are composed mainly by storage and processing, with their respective quantities, independent from each other. The type of application to be installed defines the amount of resources. Network resources are usually composed by transport units connecting the IT resources to the CSP network border.

The transport units can be represented as optical network lightpaths with one wavelength capacity connecting the physical DC resources to the CSP border, where the CSP network connects to other network operators. The physical DC resources are servers with storage and processing capacities that are responsible for handling requests coming through the network (DEVELDER et al., 2013).

Given that IT resources are spread over different DC locations, the CSP can respond to a cloud service request in an *anycast* fashion, which means that only the source node is defined by the request and the requested IT resources need to be previously known in the provisioning phase. Thus, any DC that meets the IT and networking resource requirements can be assigned to respond this cloud service request. From this point, the cloud service is composed by IT and networking resources, representing a single entity from the management point of view. Therefore, all the operations made on this entity will affect both networking and IT infrastructure associated with it (LANDI et al., 2012; DEVELDER et al., 2012).

Optical clouds create a set of new possibilities that were never experienced before, as well as new concerns related to the management of this set of resources. The combination of different quantities of each resource and configuration levels can lead us to new business models and service types. Quality of Service (QoS) and Service Level Agreement (SLA) requirements now comprise both networking and DC resources, and QoE highlights as a concurrent measure for cloud services (VOORSLUYS; BROBERG; BUYYA, 2011; LANDI et al., 2012). Therefore, integrated optical cloud management has to provide strategies to ensure survivability requirements fulfillment over the entire cloud system.

Dedicated protection figures as the most effective survivability strategy at providing dedicated backup for each resource. For example, Dedicated Path Protection (DPP) resilient against single link failures reserves, for each demand, a dedicated primary and a dedicated backup paths (i.e., backup paths cannot be shared among demands) (AHMED et al., 2013). Thus, the protected path has full survivability against single link failures. However, dedicated protection uses excessive amount of resources and such behavior becomes more pronounced on optical clouds, where there are networking, storage and processing resources to be protected.

To reduce the amount of resources used for backup, Shared Path Protection (SPP) allows to share a backup wavelength among demands with disjoint primary paths. Nevertheless, SPP still using a considerable amount of backup resources, which can still lead to an increased cost for the CSP.

In some cases, it is possible to decrease protection reserved resources by allowing the use of unprotected services, thus reducing the amount of resources needed for services and lowering the costs. These unprotected services do not use any previously reserved resource, which improves the resource usage. However, the restoration strategies used plays an important role to improve the survivability performance. Usually, the restoration process of an unprotected connection relies on the remaining unused infrastructure resources to establish a new path replacing the failed one.

Service relocation can also be used during the restoration process of a disrupted cloud service (AHMED et al., 2014). By using service relocation, the entire content of a service can be moved from one DC to another based on some criteria (e.g., unable to establish a new path to the DC). This process can be used to enhance survivability performance for the optical clouds.

1.1 Proposal

Given the actual scenario, some opportunities brought by optical clouds are still unexplored by survivability strategies. Its impacts on the design and operation of optical clouds can possibly enhance survivability performance using the same amount of resources, or reduce the need of spare resources for the same level of protection. Evaluation of the drawbacks and advantages of using the new opportunities are important to define if optical clouds are worth, since enabling these operations will require planning and resources not considered before.

This work relies on the hypothesis that service relocation can be used in the restoration process in order to improve survivability performance of optical clouds, and the achieved improvements can reach higher levels in the presence of a heterogeneous scenario. Even with the ability of relocating services, only a few relocation operations are required. Furthermore, such modifications on the restoration process do not impact significantly on the capacity for accepting requests.

To validate the proposed hypothesis this work aims to create restoration strategies able to relocate cloud services between DCs, as well as adapt existing strategies to consider service relocation, and analyze the survivability performance of such strategies on the operation of optical cloud scenarios. The achieved performance is benchmarked against literature strategies in order to get a better sense of the advantages of adopting the proposed strategies using different performance metrics.

Modeling heterogeneity and dynamics of optical clouds operation is also a challenge, which needs to be addressed properly so that the simulation scenarios reflect a more realistic scenario. Such heterogeneity is given by the different requirements that the optical cloud must meet from its users (CSP, businesses and users), and the dynamics inherent of users behavior. Once this scenario is modeled, simulation techniques can be applied to evaluate the performance of each tested strategy and to define each one is more suitable for a given scenario.

1.2 Document Organization

The remaining of this document is organized as follows: The chapter 2 details the optical clouds, its main concepts and components, mainly the ones whose differentiate the optical clouds from the traditional optical network and cloud computing. The new operations possible in an optical cloud service are also explained. Finally, management operations important for this work are also detailed.

The literature review of the related works is in the Chapter 3, where some of the recent publications on survivability strategies on optical networks and optical clouds are listed. Some works related to important subjects needed for the implementation of optical clouds, but not central for this work, are also commented. A review of survivability performance metrics is also presented to support this thesis proposal and results.

Chapter 4 details the proposed work of this thesis. The proposal scope delimits the central areas of this work, and also the other important areas needed for the successful implementation of this proposal. The contributions of this work are also highlighted.

Performance assessment of the proposed strategies is shown in Chapter 5, where published results from this research are presented, mentioning the accomplished contributions.

The final remarks in Chapter 6 summarizes the presented results in this document. Its first section highlights the contributions accomplished for this thesis. Finally, the second section gives some insights about future works that can be derived from this work.

2 Optical Clouds

Optical clouds concept has been strengthened due to recent advances on cloud computing and optical network fields. Among the characteristics of optical clouds, the *anycast* provisioning fashion and IT resources distributed over geographically spread DC locations interconnected via high-speed optical links (e.g., Wavelength Division Multiplexing (WDM) links) are the most prominent.

As stated in the Introduction, the “cloud” denotes a computing infrastructure which hosts applications to be available through an Internet connection. There are some requirements expected to be addressed and fulfilled by any infrastructure that wants to be considered cloud. Scalability, flexibility and elasticity must be strong characteristics to allow for the deployed applications a wide range of possibilities. Virtualization of resources is a key point as well, since it enhances the granularity of resources use and management, and allows the physical infrastructure to be shared among a significant number of applications. Automatized management offers dynamic and on-demand setup to easily (re)allocate needed resources (VOORSLUYS; BROBERG; BUYYA, 2011). These requirements can be considered as “cloud requirements”. Current DCs address most of the requirements, and are commonly referred to as “cloud computing”.

Recent researches have shed light on optical networks to address some of the cloud requirements by enhancing virtualization of optical equipment and resources, developing reconfigurable devices, proposing new management protocols, etc. Thus, the combination of cloud computing DCs with highly dynamic optical networks creates the concept know as optical clouds.

Storage, processing and networking are the fundamental resources provided by an optical cloud, and often used as main resource requirements for applications deployed in optical clouds. Storage and processing capacities are provided by DC layer, whereas network layer provides data transmission capabilities. Fostered by the usual high transmission rates required by actual applications, network layer uses high-capacity optical fibers (DEVELDER et al., 2013).

These 3 fundamental resources can be combined to offer different infrastructure functionalities “as a service”. Figure 2, adapted from (DEVELDER et al., 2012) and (VOORSLUYS; BROBERG; BUYYA, 2011), shows the usual commercial different levels of service classes offered by a cloud, its main access and management tools, and some existing commercial examples (VOORSLUYS; BROBERG; BUYYA, 2011; DEVELDER et al., 2013). On the Infrastructure as a Service (IaaS) class, virtual resources are deployed over physical infrastructure in order to offer basic infrastructure services, usually allowing access to its basic configuration as, for instance, administrative access to Virtual Machines (VMs). For the Platform as a Service (PaaS) class, environments for development and/or execution of applications are offered (e.g., Java Virtual Machine (JVM) up and running for deploying JVM-based applications), reducing needs of infrastructure administrative knowledge. For the Software as a Service case, the offered services are accessed by its final users, which means all the infrastructure and software required for the service is provided by the CSP (VOORSLUYS; BROBERG; BUYYA, 2011).

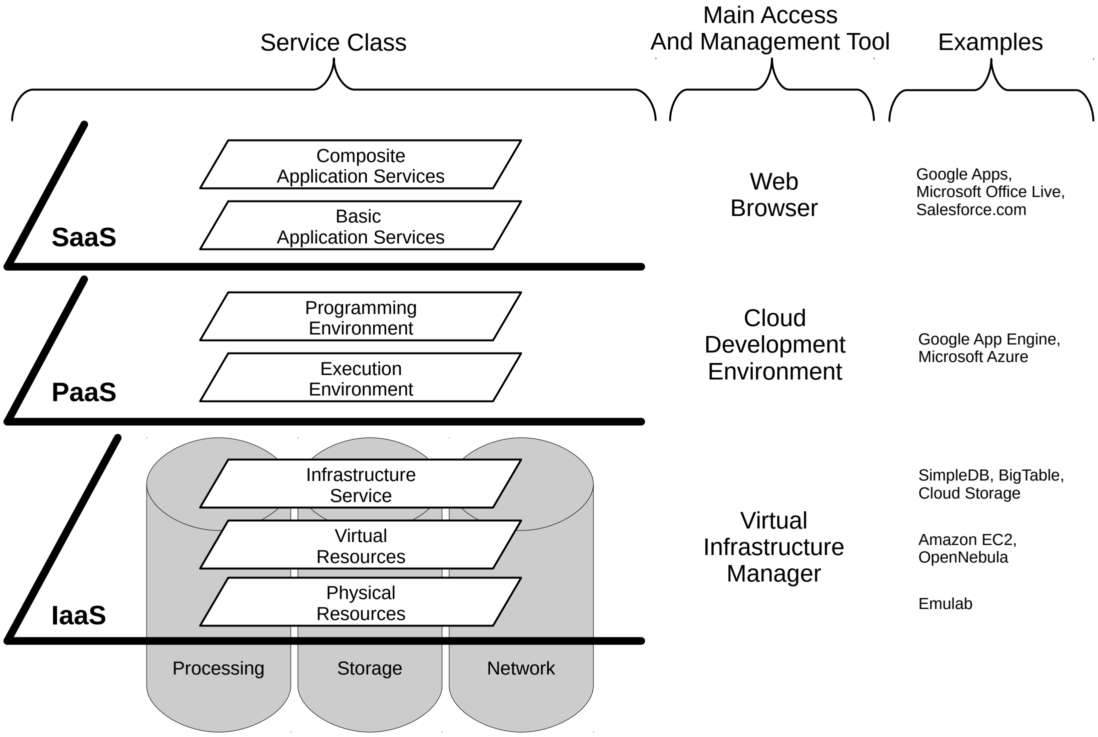

Figure 2 – Cloud computing stack and the main access and management tools

For all the different service classes an optical cloud can offer, advanced, efficient and high capacity communication networks are always required. Thus, optical networks figures as the most suitable networking technology for such task, and it plays an important role on implementing optical clouds (DEVELDER et al., 2013). Given the flexibility associated to the clouds, applications can require a different amount of each fundamental resource.

From this perspective, academic (also referred to as scientific), enterprise (business) or nonprofessional (consumer) applications can be deployed in an optical cloud (DEVELDER et al., 2013). Such applications may share the same physical infrastructure (BILAL et al., 2014), as long as virtualization ensures the isolation of resources assigned to each application. To use all the potential of an optical cloud, the application must be correctly integrated with cloud functionalities (VOORSLUYS; BROBERG; BUYYA, 2011). An application deployed on an optical cloud and prepared to use its functionalities is commonly called cloud service, and this denomination is used throughout this document.

Table 1, adapted from (DEVELDER et al., 2012), summarizes the main requirements an application can require from the cloud: required storage, processing and network capacities for different types of applications. Characteristics as scalability and elasticity are also demonstrated. The amount of required resources by one application to each resource is independent, i.e. the amount can vary independently. In this sense, CSP must provide a high level of granularity to its customers, allowing for the allocation of a quantity of resources very close to what is effectively needed. This characteristic also reduces the amount of resources paid to the CSP and eventually not used (VOORSLUYS; BROBERG; BUYYA, 2011).

Table 1 – Requirements for Novel Applications (adapted from Develder et al. (2012))

Application	**Storage***	**Processing***	**Network***	**Scale**$^\circ$	**Elasticity**†
Scientific					
Scientific computing	+	+++	o/+/++	−	+
Data computing	+++	++	++	−	+
Sensor applications	++(+)	+	+(+)	-	-
Business					
Virtual meeting	-	+	+	+	-
Colaborative framework	o	+	+	+	-
Multimedia processing & editing	+	++	+++	-/o	++
Multimedia storage & retrieval	++(+)	+	++	-/o	++
Data mining	+	++	+++	-	-
Transactional systems	+	+	+	-	-
Consumer					
Multimedia storage, editing & processing	++	+	++	-/o	++
Action games	-	+(+)	+(+)	++	+
Interactive TV	-	+(+)	++	+	+
Augmented reality	-	+	++	++	+
Virtual tourism	-	+	+(+)	+	+

*: Qualitative measures: - Low, o Neutral, + High

$^\circ$: Number of entities for a single application instance

†: Requirement for amount of resources of a single application instance: +Fluctuating, -Constant

The independence relation among the resources is very important for the understanding of optical clouds. For example, scientific applications (e.g., numerical simulations, solution of optimization models) can require a short amount of storage, but a huge processing capacity. Its scalability requirement is very low, as its resource requirements usually do not vary over time, but the networking needs can vary depending on a series of application characteristics. On the other hand, multimedia storage and retrieval applications can require very high amount of storage and a short amount of processing capabilities, with a moderate need for networking.

There are some common requirements for cloud applications not included in Table 1. Such applications rely on on-demand instantiation of required IT and network resources. Mechanisms for setup, allocation, configuration and exposing the necessary resources dynamically are key features in this environment (DEVELDER et al., 2012).

Resilience and reliability figures also as key concerns. Although there is an extensive literature addressing this problem, each one of the domains (network and IT domains) have been analyzed separately. The successful deployment of optical clouds calls for the inter-working of both IT and network resources, and their interdependency needs to be addressed to meet the cloud applications requirements (DEVELDER et al., 2012). Despite the challenge to consider such complex environment, new alternatives and approaches can be derived from this interdependence, and this work aims at exploring such alternatives.

Optical clouds share characteristics with optical networks as well as cloud computing (e.g., service or demand lifecycle). But as stated before, it has some characteristics which differentiate it from previous optical networks or cloud computing concepts (e.g., anycast traffic), as well as allows for some new operations that can be executed in the infrastructure not allowed before (e.g., service relocation). Understanding the mentioned characteristics and operations is mandatory for this work, and it is explained in details in the next Sections of this Chapter.

2.1 Cloud Service Lifecycle

The lifecycle of a service in an optical cloud scenario is similar to the one found in an optical network. As this work focuses on operation of optical clouds, a clear definition of the service lifecycle is important. Figure 3 shows a diagram which represents the service lifecycle considered in this work.

Upon the arrival of a service request in the system, the connection provisioning (or assignment) algorithm decides if the request is accepted or blocked. In optical networks, it usually uses Routing and Wavelength Assignment (RWA) algorithms to define the route to follow, the wavelength to be used and the protection strategy (e.g., dedicated or shared path protection, unprotected) used for this request defines the amount of resources to reserve. If, for some reason, the provisioning algorithm decides to drop the connection, it does not receive any resource, and accounts for the blocked connections. If the request is accepted, the required resources are reserved and provided for the service (e.g., networking and IT resources), which starts its operation period. An important performance metric for

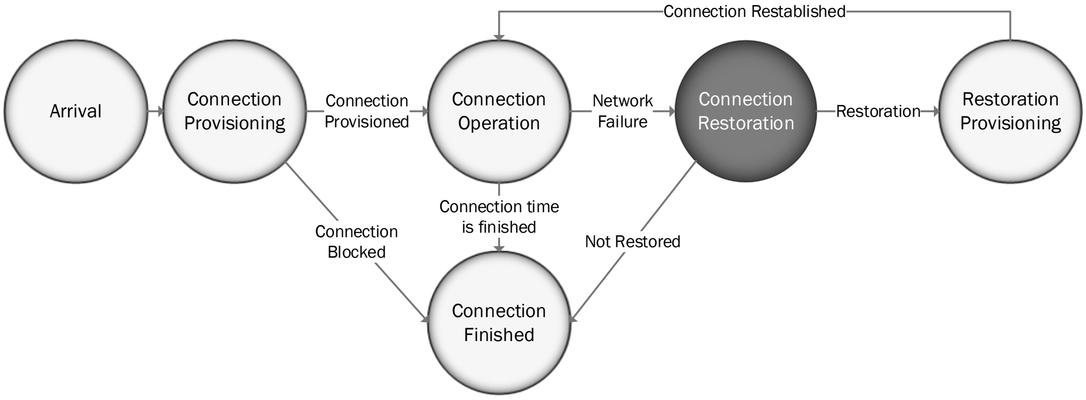

Figure 3 – Cloud service lifecycle

provisioning algorithms is the blocking probability, which measures the algorithm's ability to allocate requests, and is expected to reach as low a number as possible. In simulation scenarios, it is defined as the relation between the total number of blocked connections by the total number of arrivals. Other mechanisms in other connection phases can also affect the blocking probability (e.g., restoration strategies). The provisioning strategies can consider also other impacts, e.g., energy efficiency (BUYSSE et al., 2013).

Is important to notice that in this work long-lived cloud services are considered. For this type of service, during the operation phase end-user requests access such services during a certain time. Once the request finishes, the processing and/or storage resources needed to respond to the request are released for the Operating System (OS) running in the VM. However, both networking and DC resources will keep running to be able to serve any existing or future end-user request. The same service model is described in (MANDAL et al., 2013).

Once the service is established, the operation is responsible for keeping the service working until its holding time finishes and it leaves the infrastructure, releasing all the resources reserved for this service. During this period, if some failure occurs, the restoration process is responsible to restore the disrupted service to its working conditions. If the service has protection resources reserved and offers resilience against the experienced failure, the restoration process needs to switch its state for the resources: from the failed state to the working state resources, and then reestablish the service. If the protection strategy is unable to recover from the experienced failure or if the service has no protection resources, the restoration process must try to find available resources able to restore the service. If there are no available resources to restore the service, it is dropped and accounts as unrestored service.

Restoration is the phase where the strategies proposed in this work can be applied. In this phase, to recover the connection, several strategies can be used. One is to create a new network connection between the source node and its DC is one option already existing in optical networks. However, as optical clouds offer higher flexibility degree, other operations (e.g., service relocation) can be used to restore the service.

In cases where a new set of resources is provided to restore the disrupted service, the restoration provisioning releases the previous reserved resources and reserves the new set. Once the service is fully restored, it returns to the operation phase, and stays until either its holding time finishes or a new failure occurs.

2.2 Anycast Traffic

The *anycast traffic* is the point-to-point flow of packets between a single client and the “nearest” destination server (METZ, 2002). The reasoning behind anycast traffic is that the user wants to send and receive data to/from any one of several possible servers since such server has the content or service required by the user, but the user does not have any requirement regarding which server should be used or where the server should be placed. Then, the “nearest” destination server aforementioned becomes the one which better fits the constraints regarding content or service required by the user, availability of resources to be provisioned for the user and specific service requirements.

Optical clouds are based on high speed optical networks and provide for its users high capacity communication and reinforces the geographical independence of where the services provided for users are currently deployed. Thus, from the user point of view the real location of the IT resources is not important as long as the service meets the user requirements (e.g., as listed on Table 1). This characteristic allows for the optical clouds allocate each service based on parameters other than user location. Then, anycast traffic is used to better represent the geographical independence of cloud services. From this perspective, a service demand initially does not know which server will be used to serve it.

Consider the topology showed in Fig. 4, where nodes 1 and 4 are source of demands and nodes 2, 5, and 6 are DCs with a given amount of IT resources, interconnected by 5 links. In traditional optical networks, at the provisioning phase (as explained on the previous section) the destination nodes must be provided for the provisioning strategy in order to calculate the resources to be used. In optical clouds only the source node needs to be known at the provisioning phase, and the destination (DC) is chosen by the provisioning strategy. Thus, all the DCs are initially candidates to host the two services.

Considering that all the DCs have enough resources to receive all the demands, the closest DC (i.e., one with the least number of hops) as criteria to assign the DC, node 2 would receive demand from node 1 and node 6 would receive demand from node 4. Considering the farthest DC, node 5 would receive both the demands (LANDI et al., 2012). Thereafter, anycast traffic enhances the flexibility in the provisioning phase, and

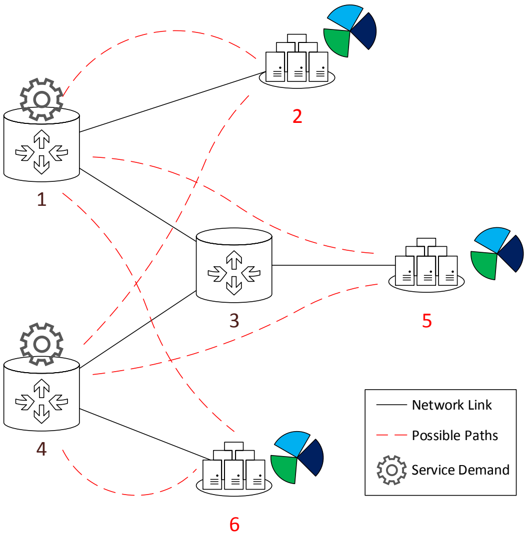

Figure 4 – Anycast traffic example

creates new possibilities for a smarter use of resources on the network.

2.3 Service Relocation

Traditional applications deployed on DCs remain at the same DC during their entire lifetime. Their change to another DC requires a manual operation, which becomes costly both time wise and financially. The cloud computing concept and the introduction of infrastructure virtualization allowed the complete automation of moving one application from one DC to another at low or zero downtime (VOORSLUYS; BROBERG; BUYYA, 2011), originating the process formally called migration or relocation. In this thesis we apply the relocation process not only to a VM, but also acting at the networking used by this VM (i.e., the lightpath used to establish the communication) (COLMAN-MEIXNER et al., 2016).

This possibility enabled the creation of a new set of application and services. Multi-DC backups or replicas, content distribution networks, geographically distributed concurrently working instances and dynamic load balancing, traffic offload and quick service replication are some of the innovations enabled or enhanced by the relocation (VOORSLUYS; BROBERG; BUYYA, 2011). Service relocation can be also used to enhance resilience of cloud infrastructure (DEVELDER et al., 2013), as primary and backup resources can be assigned to different DCs, which can be located on different geographical regions to ensure the service survives up to wide area disasters (HABIB et al., 2011; HABIB et al., 2012). More detailed comments about works related to service relocation can be seen in the next Chapter.

From the DC management point of view, each service is a VM hosted at a physical infrastructure. The main components that differ from one VM to another is its Random

Access Memory (RAM) and file storage. Thus, RAM and file storage represent the VM's current state that must be transported from one infrastructure to another during the service relocation process. The relocation procedure can be executed mainly using two strategies: live relocation and continual migration. Both enable the service relocation, but differ in some aspects explained next.

2.3.1 Live Relocation

Live service migration (or relocation) is a fundamental procedure across DCs, since it is used in many higher level DC procedures. IT consolidation, placement of services closer to demands, etc. (MANDAL et al., 2013) are some examples of procedures that rely on the live relocation. In traditional infrastructures the migration would incur in some service downtime, since the VM needs to be turned off to transport its state and reconfigure the infrastructure.

Recent studies and industry solutions have overcome these limitation by proposing seamless live VM migration over a Wide Area Network (WAN) (WOOD et al., 2014), but this process still generates a significant amount of network traffic and complex operations on both source and destination DCs (MANDAL et al., 2013). The live migration is done in several phases, to synchronize all the data during the VM execution. However, if we consider the live migration of a disrupted cloud service, the VM can be turned off, as there are no end-users accessing it. Thus, the migration process can be done in only one phase.

For example, consider a cloud service with 1.7 GB of RAM and 6.6 GB of storage used out of a 10 GB disk (WOOD et al., 2014). The total amount of data needed to be transported is 8.3 GB. From this point, the inter-DC network needs to transport this amount of data from the original DC to the new one (LIN et al., 2013). If the inter-DC network offers 1 Gbps capacity, it will take 66.4 seconds to transfer the VM state (i.e., 66.4 Gb) over ideal link conditions. If the inter-DC network offers 100 Mbps capacity, then the migration will need 67,993.6 seconds to complete in ideal link conditions. Table 2 summarizes all the analyzed data for migration time.

Table 2 – Live relocation VM migration time example

RAM	1.7 GB
Storage disk	10 GB
Storage space used	6.6 GB
Total VM size	8.3 GB
Time for migrate over 1 Gbps	66.4 sec
Time for migrate over 100 Mbps	18.89 hours

Other aspect of live migration that must be considered is that this process can be done only if the VM is under working conditions and can be reached, i.e., there is network connectivity and storage and RAM are working. If either the VM is not reachable by the network or the storage is under failure conditions, the live migration process becomes impossible.

2.3.2 Continual Migration

To circumvent the drawbacks of live migration, continual migration proposes the automatic and continuous propagation of VM state to a backup infrastructure using live migration techniques. To transfer the data from one DC to another, efficient inter-DC networks can be used for it (PERELLO; PAVON-MARINO; SPADARO, 2013; LIN et al., 2013; PENG et al., 2014). By using this strategy, the relocation process can be applied to VMs under failure conditions with minimal downtime (CUI et al., 2009).

The work in (COUTO et al., 2015) analyses the impacts of data rate and distance between DCs in the continual migration, used to create replication of VMs. Some requirements for continual migration, as the latency, are also mitigated. The study shows that the continual migration can be successfully applied in WANs.

As the VM state is continuously synchronized, if the backup DC is different from the primary one, the cloud service can recover from link, server, storage, RAM or even when the DC completely fails (e.g., all the links connecting the DC failed, disaster failures, etc.), given a consistent VM state already exists at the backup DC. The requirement for local storage backup can decrease as it is already backed up at other DC. As the state transfer is done continuously, only the changed RAM and storage are transferred at each synchronization, reducing drastically the transfer time.

For continual migration, a cloud service needs to have a dedicated storage backup at the backup DC. Some performance degradation can also be experienced by the VM, mainly because the storage access of continual migration mechanism is concurrent with the normal OS storage access (CUI et al., 2009).

2.3.3 Restoration with Relocation

The assignment of backup spare resources during the provisioning phase enhances the survivability of services, but at the cost of higher blocking probability for the same load level (AHMED et al., 2014). Thus, unprotected services require no spare resources, but no survivability against failures. Restoration strategies are designed to recover services from failures, improving resource utilization and reducing blocking probability.

Optical clouds allow for the service relocation, and this operation can be combined to restoration strategies to improve survivability performance. Then, the use of service relocation in the restoration process improves network survivability by increasing the restoration possibilities upon the occurrence of a failure (AHMED et al., 2014). In this scenario, a service disrupted by a failure can be recovered not only by establishing a new lightpath from its source to the destination, but can also be relocated to a different DC. Figure 5 shows an example of how could it work using Fig. 4 as a starting point.

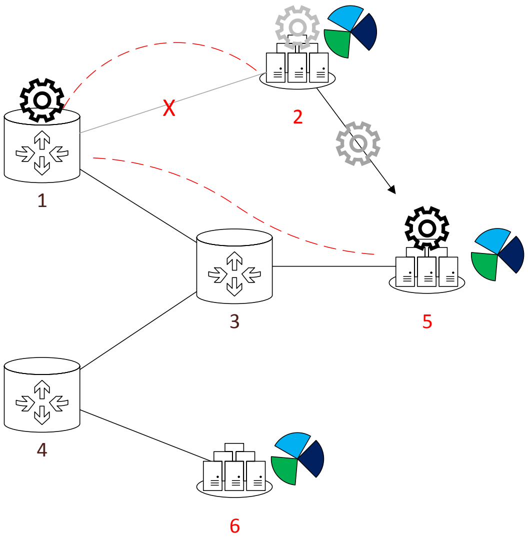

Figure 5 – Restoration with relocation example

In this case, consider the cloud service demand from node 1 is assigned, at first, to DC 2 using network link (1,2), and is called service 1. Consider that link (1,2) suffers a fiber failure, and then becomes unavailable, hence service 1 is disrupted by the failure. On traditional restoration approaches, the only way to restore the disrupted service is establishing a new path from the source to the destination, which is not possible considering the topology. In such conditions the service is dropped.

Now consider the service relocation is possible during the restoration process. As there are no working paths between source and assigned DC, the restoration process can relocate (i.e., migrate) the cloud service from DC 2 to DC 5 or 6, since both DC 5 and DC 6 have enough IT and network resources. Figure 5 considers that the service is relocated to DC 5, and a new path is established (1,3,5), allowing for the cloud service to be alive. For this case, the relocation must be done using an inter-DC network, which belongs to other network layer (i.e., inter-DC network) not considered in Fig. 5, but explained in Fig. 1.

2.4 Optical Clouds Management

The convergence of IT and optical network services is an important step to fulfill the requirements of new emerging applications as showed in Table 1. To provide the required flexibility, scalability and other cloud requirements, the cloud services as well as each physical and virtual resources used by this service need to be managed from one only point of view (LANDI et al., 2012), as commented in Section 2.1.

Some extensions for standard protocols as Automatically Switched Optical Network (ASON), Generalized Multiprotocol Label Switching (GMPLS) and Path Computation Element (PCE) solutions to manage virtual optical cloud infrastructure has been proposed. To enable such integration, the DC controllers (e.g., Service Middleware Layer (SML)) and Network Control Plane (NCP) must share responsibilities and data about each infrastructure. The main goal is to create a management entity capable of controlling an entire optical cloud infrastructure, as in Fig. 1. Thus, SML is aware of the network resources as well as NCP is aware of the IT resources. The integrated management is crucial for providing functionalities as restoration with service relocation, the provisioning of anycast demands, etc (LANDI et al., 2012).

2.4.1 Performance Metrics

Simulation of dynamic systems involves the modeling of a given system and usually the evaluation of some performance metrics in order to analyze which strategies can be used to improve the performance. In computational systems, simulation is a powerful tool to analyze performance of systems that may not yet exist, or to evaluate impacts of changes in the systems. The use of simulation also reduces the cost of creating all the infrastructure for testing.

Consider, for a given optical cloud simulation environment, the following variables that define the scenario of a simulation execution:

- $MTTF$ defines the mean time to failure for a given network link;
- $MTTR$ defines the mean time to repair for a given disrupted network link;
- CA is the set of $|CA|$ connection arrivals experienced by the system during its entire simulation time, each one (CA_i) requesting CA_i^{tu} time unit duration;
- CA_i^E is equal to 1 when the request CA_i is successfully provisioned in the provisioning phase (i.e., the request is accepted by the system), and 0 if the request is blocked;
- CA_i^{down} represents the number of time units the service CA_i experienced interruptions during its lifetime, and 0 if has not been disrupted;
- DS is the total number of disrupted services for the fiber failures experienced during the simulation;
- RS represents the number of disrupted services restored (among the DS) during the restoration phase;
- RE is the total number of service relocation performed among all the restored services RS during the restoration phase.

2.4.1.1 Blocking probability

Blocking probability metric is used to evaluate provisioning (also known as assignment) strategies. This metric is also used for evaluate if some strategy placed in other lifecycle phase is impacting on the system capacity for establish request successfully. In simulation environments, the resulting blocking probability (BP) represents the total number of established services (i.e., $CA_i^E = 1$) over the total number of connection arrivals (CA), as define in Eq. (2.1). Thus, as lower is the number as better if for the overall network performance.

$$BP = \frac{\sum CA_i \left(CA_i^E\right)}{|CA|} \tag{2.1}$$

2.4.1.2 Restorability

In optical networks, an important parameter related to the effectiveness of a survivability strategy is the resulting network restorability (NR) (IANNONE, 2011). In simulation environments, it can be measured as the relation between the total number of successful restored services (RS) over the total number of disrupted services (DS), as defined on Eq. (2.2). Thus, as higher is the number as better is the restoration strategy.

$$NR = \frac{RS}{DS} \tag{2.2}$$

2.4.1.3 Availability

Network availability is also an important performance parameter to evaluate the effectiveness of a survivability strategy (IANNONE, 2011). It represents the probability of a given service operates for a certain time with no interruptions. The availability can be also calculated for a given link (LA) in the optical network topology by knowing in advance its $MTTR$ and $MTTF$, as defined in (2.3), which represents the probability of founding the link working looking at it in a random instant.

$$LA = \frac{MTTF}{MTTF + MTTR} \tag{2.3}$$

For example, consider a link has $MTTF$ equal to 1000 time units, and the $MTTR$ is equal to 10 time units. Equation (2.4) describes the calculation of the link availability. In this case, the result is an availability of 99%. Notice that this availability considers a long-term time frame, and during a certain time frame the experienced availability can vary.

$$\begin{aligned} LA &= \frac{MTTF}{MTTF + MTTR} \\ &= \frac{1000}{1000 + 10} \\ &= 0.990099 \equiv 99\% \end{aligned} \tag{2.4}$$

In many cases, the link availability by itself is not sufficient for a complete QoS analysis, and more detailed parameters are needed. The path availability can be useful for analyzing the impacts of failures in a given path. This can also be used for calculate the availability of an unprotected service. Let S_i denotes all the link components along some link i and a_i be the availability of link i (CAVDAR; TORNATORE; MUKHERJEE, 2007). The path is available when all of its links are available, so the path availability (A_i) can be calculated following Eq. (2.5).

$$A_i = \prod_{j \in S_i} a_j \tag{2.5}$$

However, the analytic definition of Eq. (2.5) is used for long-term analysis, which may not be accurate for a single case. To calculate the exact experienced availability (CSA or Cloud Service Availability) of a given service i, Eq. (2.6) is used. In this case, the availability is calculated after the service leave the system.

$$CSA_i = \frac{CA_i^{tu} - CA_i^{down}}{CA_i^{tu}} \tag{2.6}$$

For a simulation scenario, the system receives a certain amount of requests, and the resulting network availability (NA) is the total uptime over the total requested time, for all the requests accepted in the system. It means if any service has been dropped during its operation, the network availability drops below 100%, as defined in (2.6).

$$NA = \frac{\sum\limits_{CA_i | CA_i^E = 1} \left(CA_i^{tu} - CA_i^{down} \right)}{\sum\limits_{CA_i | CA_i^E = 1} \left(CA_i^{tu} \right)} \tag{2.7}$$

Unprotected services rely on the restoration strategies ability to recover from failures to survive. However, such restoration strategies have limited effectiveness under multiple failure conditions. Thus, the analytic multiple failure probability is an important parameter to define if a survivability strategy is enough to ensure the required QoS for a given system.

To calculate the link failure probability ($p(x)$) in a topology, considering each link has the same analytical unavailability, the binomial probability can be used. Lets consider

n as the number of elements of this topology, and ua as the unavailability of each element (i.e., network link). The probability to find x failures happening at the same time in the topology is given according to Eq. (2.8).

$$p(x) = \binom{n}{x} \times ua^x \times (1 - ua)^{n-x} \tag{2.8}$$

Consider unavailability of a link is the inverse of the availability, i.e., $ua = 1 - LA$, and $\binom{n}{x} = \left(\frac{n!}{x!(n-x)!}\right)$. Notice that the number of links in the network affects the probability, since as more links the network has, higher is the probability of multiple links under failure at the same time.

Consider a network with 21 links, and each link has the same availability calculated in Eq. (2.4). The probability to find the network in full operating conditions (without any link failure) at a random time is calculated in Eq. (2.9). The probability to find a single link failure in the network at a random time is calculated in Eq. (2.10). The probability to find a double link failure in the network at a random time is calculated in Eq. (2.11). The probability to find the network with more than 3 failed links at a random time is 0.113% as demonstrated in Eq. (2.12).

$$\mathbf{p(0)} = \left(\frac{21!}{0!(21-0)!}\right) \times (1 - 0.990099)^0 \times 0.990099^{21-0} = 0.81143 \equiv \mathbf{81{,}143\%} \tag{2.9}$$

$$\mathbf{p(1)} = \left(\frac{21!}{1!(21-1)!}\right) \times (1 - 0.990099)^1 \times 0.990099^{21-1} = 0.1704 \equiv \mathbf{17{,}04\%} \tag{2.10}$$

$$\mathbf{p(2)} = \left(\frac{21!}{2!(21-2)!}\right) \times (1 - 0.990099)^2 \times 0.990099^{21-2} = 0.01704 \equiv \mathbf{1{,}704\%} \tag{2.11}$$

$$p(x \geq 3) = 0.113\% \tag{2.12}$$

For a more complete assessment on the link failures and link reliability, we refer to (WIATR; MONTI; WOSINSKA, 2012) and (CHIARAVIGLIO et al., 2015).

2.4.1.4 Relocations

For optical clouds, the performance of a restoration strategy can be measured also by its capability of retaining the number of required relocation. Reducing the number of required relocations is mandatory given the overhead related to the operation, and described in Section 2.3.

The percentage of relocations can be calculated following the Eq. (2.13), i.e., the relation between the number of relocations and the number of disrupted services. The work presented in (AHMED et al., 2014) uses this parameter to compare the ability to retain the number of relocations.

$$RF[\%] = \frac{RE}{DS} \tag{2.13}$$

However, this strategy to calculate the relocation percentage can be affected by the network restorability, since the same number of relocations can result in a higher relocation percentage case the restoration strategy does not achieve the same restorability performance.

2.5 Applicability in Smart Cities

Cities are expecting unprecedented growth, and it puts significant stress on city infrastructure as demand outpaces supply for water, energy, transportation, health care, education and safety. To address these issues, and make the city smarter and more attractive, reducing costs and increasing efficiency, the Information and Communication Technology (ICT) are required to bring innovation in planning, deployment, management and operations in a city (NAPHADE et al., 2011). To deploy a smart city, a complex ICT ecosystem composed by applications, user and/or citizen data, connectivity and energy efficiency figure as challenges that have been highlighted in recent researches (WALRAVENS; BALLON, 2013).

The smart city concept also means a continuous improvement process over a city environment, processing a huge amount of data from different sources to generate usable information, used by different stakeholders during the decision process through the Command and Control Center (CCC). Hence, the ICT infrastructure is the basement of a CCC operation, supporting its processes. This ICT infrastructure is composed mainly by data transporting, storing and processing. The effective management of these 3 main components can be decisive role in reducing the ICT cost in a smart city deployment. Actually, optical networks are the best option to create high capacity data transporting networks, and cloud computing is the best option in high capacity data storing and processing.

In developing countries, the need for improvements in ICT infrastructure is a key point in a smart city deployment, and can be decisive in a CCC deployment. In this case, the planning and management of ICT is even more important to a successful deployment and long-term operation. So, define the needed infrastructure is not always the most important task, but develop ways to a smart usage of available ICT resources on cities in which there are scarce resources.

2.5.1 The ICT Infrastructure on Attractive Cities

In a smart city scenario, there are different actors that will communicate sharing data and information. Figure 6 shows a view of these actors and the ICT infrastructure that will provide the basic transport, storing and processing capacities to support the CCC operation. The ICT infrastructure is composed by telecommunication (CT) and IT layers. In the case of cloud computing layer, the infrastructure can be divided in storage and processing infrastructures, which can be combined to process or attend the requests. The IT layer can be geographically distributed over different data centers (places containing a certain amount of IT resources) interconnected by the telecommunication layer.

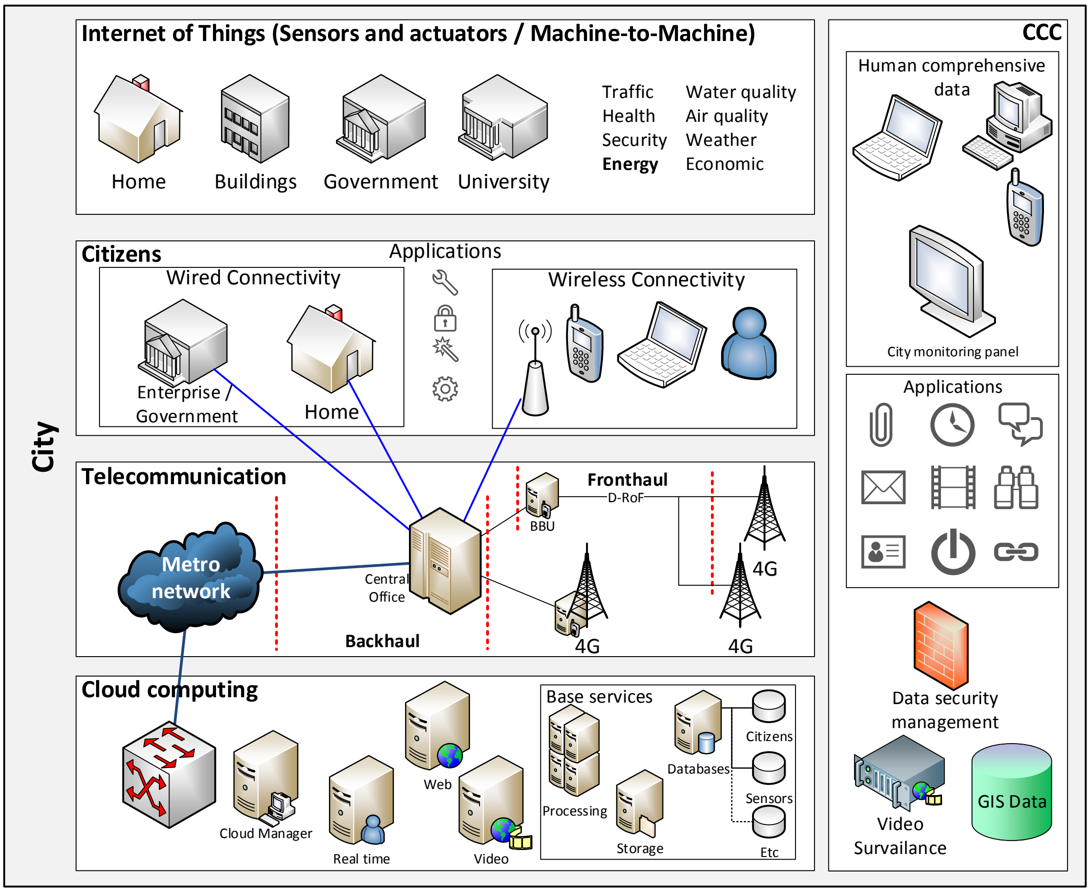

Figure 6 – Smart city ICT scenario

In the initial step, the telecommunication layer will transport all the data from the different sources around the city, which can come from sensors and/or citizens to an eligible DC (a DC with enough available IT resources). All the collected data can be stored on the cloud computing layer using storage resources. In the second step, the CCC requests information that can be generated using the data stored on the storage infrastructure. The information generation can use computer intensive processes, for example, data mining, map-reduce, image processing, 3D mapping, etc. This computer intensive processing will be done by the processing infrastructure, also placed on cloud computing layer.

2.5.2 Modeling Details

In this section, the components of optical clouds and its relation with the CCC operation will be detailed for a more complete understanding. It is important to specify all the technical components and its roles for the CCC.

2.5.2.1 Optical network

Optical networks are high capacity networks used to transport huge amounts of data using high speed in large distances. Recently, optical networks have been used to create metropolitan area networks, interconnecting public sector and educational buildings in a city. This new high capacity interconnection makes easy the creation of applications that share data among the different public units. So, by using optical networks, the CCC deployment will be able to integrate and share data in easier and faster way. The recent advances in both deployment and management on optical networks have reducing the costs of installation and maintenance on these networks, making this the main option when high capacity is needed.

2.5.2.2 Data centers

The public institutions usually have IT infrastructure in which all the internal applications can be installed and from where these applications can be accessed. In a city, there are many buildings in which each public institution install and manage its own IT infrastructure, keeping each one isolated from each other. This isolation does not allow one institution to use infrastructure from another institution, which could help the public administration to reduce the amount of infrastructure in each institution, and also improve the use of the installed infrastructure. The need of infrastructure improvement is also a huge requirement in all the institutions, and the IT integration can reduce this need by make another IT resources available when needed.

With the creation of the metropolitan optical networks, new possibilities of management and use of the IT infrastructure need to be analyzed. The integrated management of different data centers has been used by the industry for some years, usually called cloud computing, showing good results in terms of costs and specifically in energy consumption. This innovative management strategy can create a layer that manage IT resources geographically distributed, and provides access to them in a transparent way, improving the convergence of the different sector in a city. All these characteristics show that the use of cloud computing can improve the CCC performance.

So, to study the benefits of an integrated IT infrastructure management in a smart city scenario can make the deployment less expensive, using the pre-installed infrastructure as much as possible, making new acquisitions only when extremely needed.

2.5.2.3 Cloud service

Cloud service is a concept of application that can use all the benefits of innovative ICT infrastructure (optical networks and DC. In this kind of application, different data sources (possibly geographically distributed) can be used while the application processes these data. Using the cloud computing, the processing of this application can be assigned to any DC that has enough IT resources (storage and processing) and available networking resources to interconnect the data sources to the processing DC.

A cloud service can be any application or algorithm designed to use storage and processing capacity, and in a smart city scenario, there are many examples of high capacity applications that will need to be processed in the ICT infrastructure. One example is the use of forecast weather models that use a huge amount of IT resources while processing the data to generate the forecast. Applications that use image processing also need high processing capacity. Air pollution models are also an example, once they usually need to use a huge amount of data from different sources, and also need a large processing capacity. To generate citizen-based applications, using a human sensor concept, the multitude of virtual sensors and data also will need a smart ICT infrastructure management to use the least number of resources generating the best results.

2.5.2.4 Infrastructure resilience

To reduce the costs and changes in the ICT infrastructure, resilience strategies have an important role in the management of ICT. These strategies can reduce the amount of the needed backup resources (reducing costs) by using effective failure recovery strategies. So, in scenarios of low cost-or-changes scenarios, the resilience strategies play an important role to ensure the CCC continuity.

3 Literature Review

This chapter aims to show the related works and the state-of-art of survivability topics in optical networks and optical clouds subject. The benefits achieved by adding awareness of the service heterogeneity is commented, as well as some works not related to survivability, but important for optical clouds. At the end, a short review of performance metrics used in this work is provided.

3.1 Survivability on Optical Networks

The actual optical WDM networks carry a huge amount of data from many different connections and a single link failure can cause huge data and profit losses and a multitude of higher-level protocol failures (GHANI et al., 2008)(HABIB et al., 2013). The growing needs for high-survivability in optical networks is reducing the amount of time that connections can become offline, making SLA contracts stricter in relation to downtime on the networks (XIA et al., 2011), and increasing the needs for enhancements on provisioning/restoring strategies.

Ahmed et al. (2013) present a framework with a set of survivability strategies that aim to improve the network survivability in the presence of double-link failures, using hybrid provisioning and reprovisioning schemes to both primary and backup lightpaths. Simulation scenarios are used to evaluate the performance of the proposed strategy and benchmark it against previous strategies from literature during the operation of optical networks. Results show the advantages of the restoration-based schemes compared to simple protection-based schemes in terms of blocking probability, and also show the reduction on resource usage by the restoration schemes, with slightly drop on connection availability.

Paper (PEREIRA; Camillo Penna, 2014) proposes a heuristic for design survivable WDM networks using SPP. The design strategy computes the resource capacity required for a given demand and availability requirement considering protection against multiple link failures. The method works, for each demand, in two steps: the first selects out of a set of precomputed paths which are the selected paths for the demand, and; the second step calculates the availability of the availability based on the chosen paths to ensure the availability requirements are met.

The frequent occurrence of natural disasters, which can affect large areas, and the rising hazard of intentional attacks increase the risk of damaging large portions of optical networks with just a single event raises the need for protection to large scale network damages. Thus, network vulnerabilities to multiple cascading, correlated and collocated failures have become a major concern. The authors in Habib et al. (2013) present a comprehensive survey showing all the concepts related to disaster resilience for optical networks. Modeling Shared Risk Groups (SRGs) capable of representing the risks of disasters in the network is a key issue among the pointed open issues and challenges for preparing optical networks to survive upon the occurrence of such large failures.

The authors in Mukherjee, Habib e Dikbiyik (2014a) raise the need for novel robust survivability methods to mitigate the effects of disaster on telecommunication networks. The use of data from climate, geology and environmental science on how to predict disasters and assess risk of disaster in regions must be used for prepare networks to handle such events. Self organized networks and content distributions are powerful functionalities that can be used to improve the networks disaster survivability.

The advantages of considering degraded services under multiple correlated failures in the network is investigated by (HUANG; MARTEL; MUKHERJEE, 2009). In this work node failures are considered and different protection strategies are benchmarked, showing that with a small extra cost the dropping probability due to a node failure can be significantly reduced.

3.2 Status Awareness

To address the growing quality requirements of optical connections, many works consider some request and/or network parameters to improve the network performance. The work in (XIA et al., 2011) presents a strategy of provisioning lightpaths with the least associated risk on the violation of SLAs. Results show that is possible to reduce the SLAs violation by calculating the risks associated with each available lightpath.

The requirements to support a wide variety of different application types and is increasing for the Internet Service Providers (ISPs), established through SLA contracts. The work of (HARLE et al., 2007) proposes a framework to be included in SLA contracts in order to provide the required level of survivability. Present the different recovery techniques and how it implies on the survivability performance for the customers is a key concern for ISPs. Table 3, adapted from (HARLE et al., 2007), shows how different application requirements can be translated into survivability strategies and requirements.

The current heterogeneity of services that can use optical networks is represented in Table 3. The higher protection level implies in a higher cost, but reduces considerably the recovery time. Restoration schemes impact in the best effort application, which has no restorability or availability guarantee. The cost reduction for best effort applications is achieved by reducing the spare capacity required for the service.

Applications		Survivability Scheme	Restorability	Availability Under Normal Operation	Availability Under Failure Operation	Spare Capacity	Recovery Time	Price
Non-preempted	Protected	Dedicated protection (1+n)	Guaranteed recovery	100% guaranteed availability	100% guaranteed	Highest	Fastest	Expensive
		Dedicated protection (1:1, 1:2)			Premium availability	High	Fast	
		Shared protection (1:n)			High availability	Medium	Medium	
		Pre-allocated restoration			High availability	Low	Slow	
	Best effort	Restoration	No guaranteed recovery	100% guaranteed availability	Availability not guaranteed	None	Slowest	
	Unprotected	None	No recovery	100% guaranteed availability	No availability	None	N/A	
Pre-empted	None	No recovery	Simple	Availability is not guaranteed	No availability	N/A	N/A	Cheap

Table 3 – Application identification and characteristics

Thus, as one only fiber cable can carry many different wavelengths supporting heterogeneous connections, each connection needs to be provisioned and managed according to its own quality requirements. In order to improve the network performance, a variety of works have proposed the use of connection properties. The work in (CAVDAR; TORNATORE; MUKHERJEE, 2007) uses the holding time of a connection during the provisioning phase to ensure the availability requirements in a service-differentiated environment. The connection protection level is chosen by analyzing the required priority and holding time, taking into account also the networks and link availability.

A similar approach is used in (SONG; ZHANG; MUKHERJEE, 2007), where the authors propose a dynamic provisioning strategy to route and protect (if necessary) each connection. The algorithm calculates the theoretical connection availability of each candidate path and protection strategy, defining the most suitable option to use for each connection. The proposed strategy also reduces the resource overbuild in the system, which can impact in the price reduction of the connection.

3.3 Survivability on Optical Clouds

More recently, due to constant evolution of optical networks, some characteristics from cloud computing have been used to add flexibility to optical networks. In the same manner it brings advantages to the optical networks, some survivability issues need to be revisited in order to update the strategies integrating the new set of functionalities from cloud computing.

In order to protect optical clouds against single link failures, Buysse et al. (2009) and Buysse et al. (2010) propose models for dimensioning the network accounting the reduction achieved when enabling service relocation. In the first, it is shown that relocation can reduce around 20% of the required network capacity, while it can require 25% more DC resources than when relocation is not considered. In the second, a heuristic for dimensioning optical clouds is proposed in order to reduce scalability issues (e.g., for solving larger problem instances) faced by the Integer Linear Programming (ILP) models, and its performance is compared with the optimal strategy.

A model to dimension both network and DC is proposed in Develder et al. (2011), which evaluates the resulting optical cloud for 3 different case studies: single link failures without relocation, single link failures with relocation, and single link or server site (also called here as DC) failure with relocation, modeled through Shared Link Risk Groups (SLRGs). Results showed that enabling relocation incurs an increase in required DC capacity, but reduces the network dimensions. A similar work, (DEVELDER et al., 2012), introduces a relation $1 : N$ to increase network and DC dimensions, improving reliability on cases where relocation is not used; and the same relation is used in (DEVELDER et al., 2013).

In (DEVELDER et al., 2013) it is shown, as expected, that protecting services against single link and server failures without relocation and with a $1 : 1$ protection strategy, requires twice the server resources. Considering only link failures, the required server capacity is the same as the number of demands (one demand requires only one server), given that services are not protected against server or DC failures. The dedicated server protection ($1 : 1$ protection) at the same DC is used to protect against single link and server failures, but is unable to protect against DC failures, and can also require more network capacity than when using relocation. Thus, protecting services using backup servers at the same DC requires more network capacity, reducing the protection level or doubling DC capacity.

For protecting cloud services against disaster failures (e.g., multiple cascading, correlated and collocated failures), paper (HABIB et al., 2011) proposes an ILP model for resource assignment and content placement which explore the use of relocation. The disaster risks are represented as SRGs, and used for the ILP to distribute the content and define the paths between client and its content in such a way that the service is resilient to the disasters comprised in the SRG set. The work in (HABIB et al., 2012) extends the previous work by proposing an ILP relaxation capable of reducing the complexity of the model, as well as a heuristic which achieves a good performance. Once a disaster event occurs, there are few alternatives to avoid data losses. Such alternatives are investigated in (FERDOUSI et al., 2015).

The work in (ZHANI; BOUTABA; HELLERSTEIN, 2014) proposes a provisioning approach that tries to find the best match between the specific QoS requirements of the cloud service and the available IT resources. This matching takes place during the cloud service provisioning phase, and it is used to make sure that critical cloud services are given precedence over non-critical applications. The work in (YI; DING; RAMAMURTHY, 2016) proposes strategies to allocated inter-DC resources optimizing budget and reducing blocking probability.

Regarding operation of optical clouds, the paper of (AHMED et al., 2014) proposes an ILP model to restore concurrently disrupted cloud services upon the occurrence of a failure, analyzing the survivability performance improvements achieved by enabling the live relocation of failed services to another DC. Two variants of the ILP are used to evaluate the impacts of contain the number of relocation: the first does not consider the number of relocations in the objective function; the second tries to minimize the number of relocations by adding a term for penalize the model for each required relocation. The results show significant improvements on survivability performance for both ILP variants, represented by total network availability and restorability, without significant impacts on the overall blocking probability. The second ILP, which tries to minimize the number of relocations, does not impact on the survivability or blocking probability performance, and halves the number of relocations.

The use of ILP for solving restoration problems allows for the concurrent optimization and overall minimization of resource usage, however is complex and difficult to predict the time it will need to solve the problem, which can have a high impact on the restoration performance (e.g., increasing downtime between failure and restoration). To overcome it the authors in (SILVA et al., 2014), which is **Paper I** of this thesis, propose a heuristic capable of calculate the restoration with relocation problem sequentially and using a set of pre-computed paths, which reduced the average solution time in two orders of magnitude. The results of this research will be described in details in Chapter 5.

3.4 Complementary Works

The study and deployment of optical clouds rely on the synergy of a multitude of components. This work studies the contributions service relocation can do to improve survivability by enhancing restoration strategies, but there are some other components also important for the successful implementation of optical clouds. These other components are not the main focus of this work, however need to be analyzed in order to expand the view of optical clouds.

With the growing concern about energy consumption in cloud infrastructure and the increasing number and scale of DCs, the authors in Mandal et al. (2013) investigate strategies capable of use VM relocation to replace the non-renewable energy used in DCs for renewable energy. Similar techniques to those presented in 2.3 are used to improve flexibility of DCs and overcome predictability issues in renewable energy (e.g., intermittency and volatility), which disable its use with full potential. Results show that in spite of the slightly increase in overall energy consumption, the non-renewable energy use can be reduced in 30%.

A set of procedures needed to improve the reliability of continual migration is proposed in (CUI et al., 2009), showing the performance of different use cases. In (CHANCHIO; THAENKAEW, 2014) and (WOOD et al., 2014) strategies are proposed to reduce the overhead caused by live migration on both DC and network.

Cloud service providers need to interconnect the IT infrastructure spread over geographically distributed DCs. The work in (LIN et al., 2013) proposes a model for efficient inter-DC wide area networks with static traffic scenario for unicast and multicast connection requests. An ILP and a heuristic are presented in order to dimension the required resources to connect DCs. The work in (PERELLO; PAVON-MARINO; SPADARO, 2013) proses a Mixed Integer Linear Programming (MILP) model to introduce flexibility in the virtual-physical node mapping in order to reduce the usage of physical resources on the interconnectivity among DCs, allowing the Virtual Optical Network (VON).

The RWA problem and alternative solutions are analyzed in the book chapter (SARBAZI-AZAD; ZOMAYA, 2014). Example with full conversion capabilities or without conversion capability are presented. The chapter proposes a new model to represent a network with partial conversion capability as a full conversion capability. At the end the algorithms are evaluated showing the blocking probability, average start time and scheduling overhead performance.

4 Thesis Proposal

Just to recall, the requirements for flexibility, scalability and elasticity are common issues in innovative ICT applications, as demonstrated in Table 1. Such requirements result in the individualization of resources for each application deployed on the optical cloud instead of previous applications which could be deployed in a shared infrastructure. This leads to the infrastructure virtualization in order to provide the required isolation in a shared environment. Thus, optical clouds need, for each application, to identify resources needed. The enhancement of individualization and granularity is expected to reduce the gap between the requirements of an application and the amount of resources allocated to this application.

Cloud computing (i.e., IT resources) currently offer a high level of granularity for each basic resource. However, analyzing Table 1 and Table 3 it is possible to notice that there is a gap between application requirements and network services offered by CSPs. Furthermore, IT and networking infrastructure need to be offered in a coherent manner. Once individual needs for each application are identified, the CSP business model can adapt the costs according to resources needed for the application with the billing costs expected for the customers, or adapt resources available for the applications according to the costs customers have previously agreed.

This work relies on the hypothesis that service relocation and differentiation can be used in the restoration process in order to improve survivability performance of optical clouds, without increasing the need for backup resources.

The first step towards the objective of this work is to analyze the impacts of enabling optical clouds for service relocation. Work in (AHMED et al., 2014) analyzes such impacts providing valuable performance analysis. However, strategies used for restoration must meet time and complexity constraints of optical clouds environment. Thus, **Paper I** provides a simpler but powerful strategy capable of meeting time and complexity constraints, and its results will be detailed in the next Chapter.

Nevertheless, previous works consider a homogeneous type of service, not representing the heterogeneity of optical clouds environments. Table 4, adapted from (HARLE et al., 2007), and Table 3 increase the existing options for provisioning connectivity to optical cloud applications. In this sense, best effort applications figures as the most suitable option for deployment of low-cost but reasonable-availability applications. The main proposal of Table 4 is to separate such connections into two different priorities, allowing them to be considered differently on the restoration process. The implementation of this behavior allows for the proof of the second hypothesis statement of this thesis, which

Applications		Survivability Scheme	Restorability	Availability Under Normal Operation	Availability Under Failure Operation	Spare Capacity	Recovery Time	Price
Non-preempted	Protected	Dedicated protection (1+n)	Guaranteed recovery	100% guaranteed availability	100% guaranteed	Highest	Fastest	Expensive
		Dedicated protection (1:1, 1:2)			Premium availability	High	Fast	↑
		Shared protection (1:n)			High availability	Medium	Medium	
		Pre-allocated restoration			High availability	Low	Slow	
	Best effort	Restoration	No guaranteed recovery with high priority	100% guaranteed availability	Availability not guaranteed with high priority	None	Slowest	
			No guaranteed recovery with low priority		No guaranteed recovery with low priority			
	Unprotected	None	No recovery	100% guaranteed availability	No availability	None	N/A	
Pre-empted	None	No recovery	Simple	Availability is not guaranteed	No availability	N/A	N/A	Cheap

Table 4 – Proposal for classifying the different application types with two-level best effort applications

states the benefits of relocation can be increased in heterogeneous scenarios.

To represent real scenarios, the simulation must consider industry-based fiber capacities. Current works consider only 16 network units in each fiber. To represent an average industry capacity, this work aims to consider 80 units (e.g., 80 wavelengths per fiber). In addition to the scenario representation, it allows the analysis of survivability performance in higher number of connections at each failure (i.e., the maximum number of disrupted connections by a single failure is equal to the number of network units in the disrupted fiber).

Ahmed et al. (2014) proposes an optimal model considering only the dynamics of the network part of optical clouds. The model receives a pre-computed decision if a given DC a is able to receive the relocation of a given cloud service s previously allocated to DC b. Lets consider that DC a is not able to receive the cloud service s at the moment of the disruption. The drawback of this approach is that during the solution of the model, a cloud service assigned to DC a can be relocated to other DC, changing the available resources in DC a and allowing cloud service s to be relocated to it. In other words, the changes happening at the DCs during the solution of the model are not represented in the model. This behavior can be improved to reflect the changes of DCs during the model solution.

Survivability performance may be improved if cloud service differentiation is considered in the restoration process. Restorability is the main performance metric addressed by previous works, but availability is more important than restorability once it is used as a major parameter in SLA agreements. An alternative way to improve both restorability and availability is to consider each disrupted service not only as an unit (as it is considered in the previous works), but also to consider the remaining time of the service to sort the set of disrupted connections. This strategy has proven its effectiveness, as commented in

Section 3.2.

4.1 Scope Definition and Assumptions

Figure 7, adapted from (HABIB et al., 2013), shows a classification of research works in network survivability field. This thesis aims to study the performance improvements achieved by enabling optical clouds to perform service relocation under link failure. We consider the occurrence of single link failures in the optical cloud topology and all the disrupted connections are submitted to the restoration strategy.

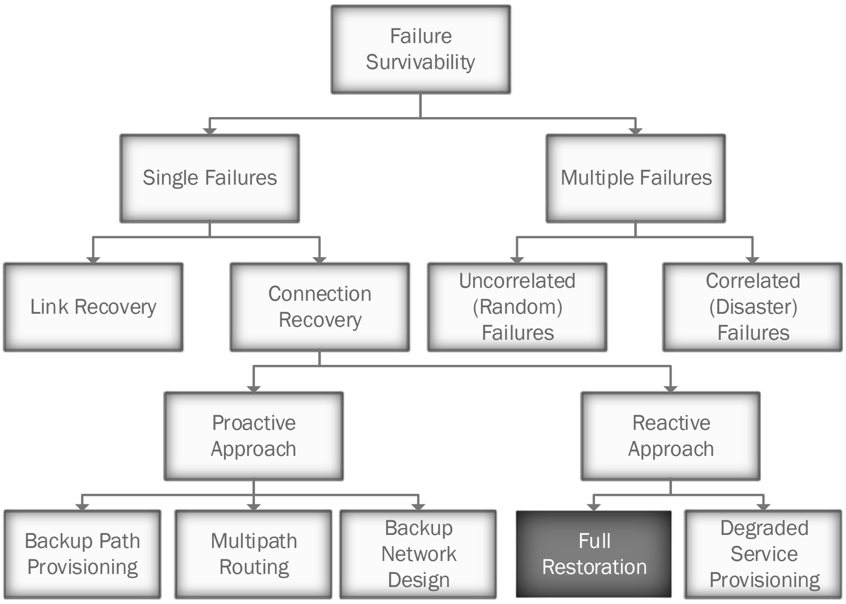

Figure 7 – Classification of research works in network survivability

Our restoration strategy works in a reactive approach. The restoration strategy is performed upon the occurrence of a link failure. The restoration under degraded conditions (i.e., when a service is restored with less resources than assigned or requested initially) is not considered (refer to (SAVAS et al., 2014) for details). The restoration strategies proposed by this work cannot be used under DC failures (e.g., storage failures) that disable the live relocation process for the cloud service.

This work does not propose or evaluate strategies for phases of cloud service life-cycle (Fig. 3) other than the restoration. Thus, impacts of the assignment strategies are not studied or analyzed. The assignment strategy called DC_CLOSEST, proposed by (LANDI et al., 2012), is used in all the experiments, allowing for the analysis of the real impacts of our proposal for the restoration phase.

As indicated in Chapter 1, the cloud can be analyzed from several points of view. For this study, the optical cloud is analyzed from the CSP point of view. It means that all the considered resources are managed from one only management entity.

Long-lived cloud services are considered for the arrivals in the system. Upper part of Fig. 8 represents this model in which the optical cloud receives requests its customers. In this example the first request that arrives demands a certain amount of resources and this request stays in the system until the service reaches its service time or a failure disrupts it. The second request then arrives and consumes more resources than the first, but remains for a shorter period of time. The total assigned resources (i.e., reserved resources or resources in use) is the sum of all the provisioned resources currently on the optical cloud. During the operational time the cloud service is deployed by the CSP customer in the cloud and remains listening to the network during its entire lifetime.

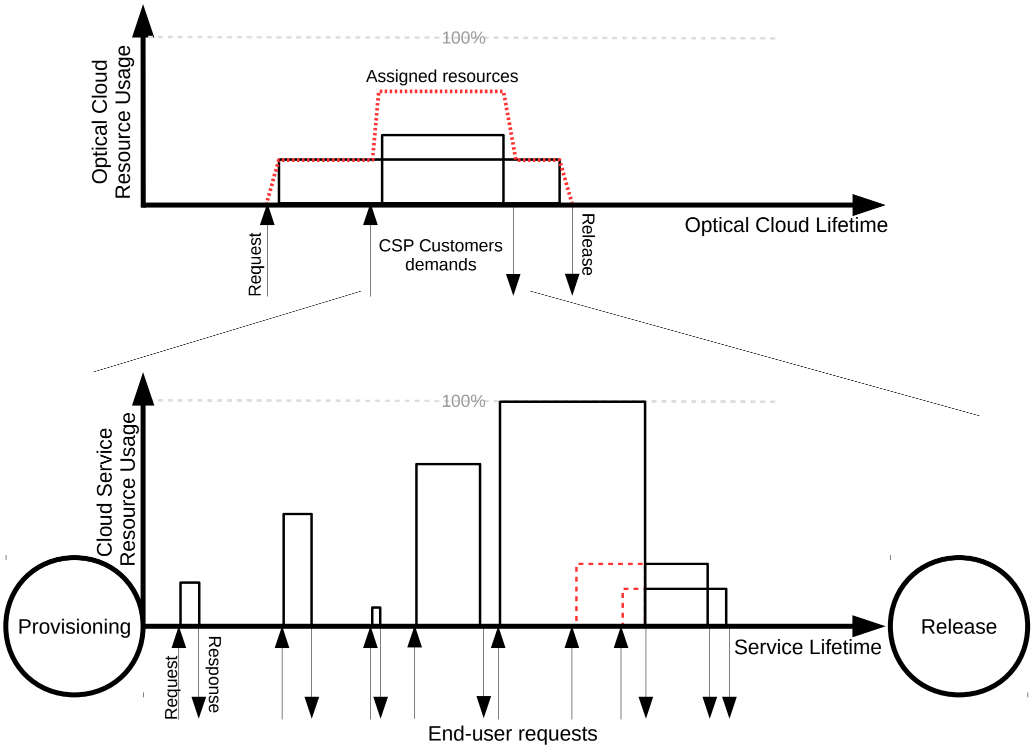

Figure 8 – Representation of the lifetime of a optical cloud and its services

The lower part of Fig. 8 details the lifetime of a cloud service. End-users can access the service by sending requests to the service and waiting for the response. After sending the response to the end-user the service remains executing and listening the network, waiting for new user requests. Usually the applications are capable of receiving simultaneous requests from different end-users. To process an end-user request the application uses a certain amount of IT resources (i.e., processing and storage) and usually the time needed to process a request is defined by the amount of resources used for the application and the amount of resources effectively assigned to the application . If the application needs to use the entire available resources for a certain period to process a request, incoming requests during this period may need to wait until the previous request to be processed.

This work considers the cloud service lifecycle represented by upper part of Fig. 8, comprising requests and release of optical clouds and its resource management (i.e., networking, processing and storage resources). End-user requests, cloud service resource usage and performance are not analyzed in this work.

In the topologies considered in this work, once a link failure occurs, such link is fixed according to the Mean Time To Repair (MTTR) defined in the simulation. Given that a failure disrupts a cloud service there are two options to keep it running: either the connection is reestablished through a new networking path or it is relocated to other DC.

Due to the use of unprotected services in the provisioning phase, the financial cost of deployment is expected to be low preventing backup in secondary path and DCs. To allow the relocation of services from a DCs, we consider the existence of a inter-DC network layer able to relocate services and capable of preserving DC connectivity in cases where the user network layer is unable to reach the DC. It results in a scenario where the relocation can be performed in any optical cloud network failure scenario.

This work considers all the intermediate nodes are capable of full network unit conversion (e.g., full wavelength conversion capability). Thus, all the paths calculated in this work do not consider the wavelength-continuity constraint.

4.2 Path Restoration with Service Relocation and Differentiation Problem

This section formally introduces the Path Restoration with Service Relocation and Differentiation (PR-SRD) problem for optical cloud services. In addition, the section also proposes two methods to solve the PR-SRD problem, one based on an ILP formulation, referred to as ILP for Relocation with Priorities (IRP), and the second one on a heuristic called Heuristic for Relocation with Priorities (HRP).

The problem considers a dynamic provisioning scenario where cloud services require continuously bandwidth between a client and a DC node. Moreover, for the sake of simplicity at most one fiber link can be down at any point in time, i.e., single fiber link failure scenario is assumed. However, the proposed solutions can be easily applied to multiple fiber link failures.

It is assumed that cloud services are divided in traffic classes, each one with a different priority value in order to reflect their importance (e.g., platinum, gold, silver, and bronze services). When competing for the same spare resources, cloud services belonging to a traffic class with a higher priority value are given precedence over the ones with a lower priority. If a cloud service cannot be successfully restored upon the occurrence of a failure, the cloud service is dropped and its remaining service time is accounted as downtime. In the problem just described, the objective of the PR-SRD solution is to minimize the average downtime of all cloud services (i.e., defined as the portion of the service holding time during which a cloud services is not available).

Let $G(N, E)$ be a graph representing the optical transport network after a fiber link failure, i.e., $G(N, E)$ does not include the failed fiber link. $G(N, E)$ consists of $|N|$ network nodes and $|E|$ fiber links. Let N_{DC} be the set of DC nodes ($N_{DC} \subseteq N$) where each $DC_k \in N_{DC}$ has DC_k^{st} available storage units, and DC_k^{pu} available processing units. Let Q be the set of cloud services disrupted by a failure that need to be restored. Each $Q_i \in Q$ requires Q_i^{st} storage, and Q_i^{pu} processing units, with Q_i^{st}, $Q_i^{pu} \in \mathbb{Z}_+^*$. The arrival and holding time values of each cloud service are represented by Q_i^{at} and Q_i^{ht}, respectively, with Q_i^{at}, $Q_i^{ht} \in \mathbb{R}_+^*$. The source node of Q_i, i.e., the client node at which Q_i originates, is Q_i^{src}, while Q_i^{dst} represents the DC node that was serving Q_i before the failure. Finally, let $\alpha_i = 1, 2, ..., M$ with $M \in \mathbb{Z}_+$, be the set of priority values (one for each traffic class) that can be assigned to the cloud service Q_i. If $\alpha_i > \alpha_j$, then Q_i has a priority value higher than Q_j. Thus, the higher is the value of α_i, the higher is the importance of the cloud service.

At the occurrence of a failure, for each $Q_i \in Q$ we define a quantity called remaining service time (Q_i^{rt}) defined as:

$$Q_i^{rt} = Q_i^{ht} - (CT - Q_i^{at}), \tag{4.1}$$

where CT represents the time at which the failure occurs (expressed in time units). If Q_i cannot be restored, Q_i^{rt} becomes equal to the downtime value. Figure 9 shows an example of how the value Q_i^{rt} can be calculated using (4.1), assuming a failure happening in the network at $CT = 40$ time units. In the figure, cloud services 2 and 4 have different arrival and holding time values, but they have the same value of Q^{rt}. On the other hand, the value of Q^{rt} of cloud services 1 and 3 is lower than the one of cloud service 4, even if the value of their holding time is higher. In this way, cloud services 2 and 4 have the higher impact in the average availability and the restoration strategy should consider it during its process.

Given that the objective of the PR-SRD problem is to minimize the average downtime of the cloud services, the value of Q_i^{rt} plays a central role in the solution. More specifically it is used to decide which cloud services should be given precedence when the restoration solution is calculated. For example, looking at Fig. 9 and assuming that all four cloud services belong to the same traffic class, it would make sense to try to restore cloud services 2 and 4 first, as they would make the largest contribution to the average service downtime. The same intuition is used in the ILP formulation and in the heuristic that are presented in the next sections.

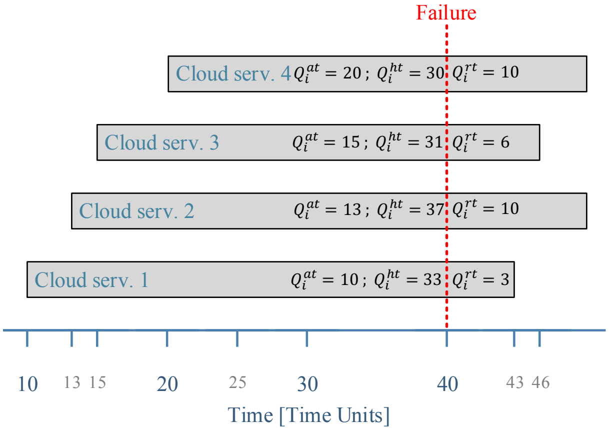

Figure 9 – Example of how the value of the remaining (service) time (Q_i^{rt}) can be calculated for a set of four cloud services disrupted by a fiber link failure happening at $CT = 40$ time units.

4.2.1 ILP for the PR-SDR Problem

The ILP formulation proposed in this section, referred to as IRP, solves the PR-SRD problem for a set of cloud services disrupted by a single fiber link failure. Given the set Q, the solution of the IRP formulation provides a set of restoration paths (i.e., one for each restored cloud service) and, when necessary, the identity of the DC nodes to which a restored cloud service has been relocated to. The ILP model presented in this section works under the assumption that each cloud service requires a full networking unit capacity. The disrupted cloud services that cannot be assigned any restoration path are dropped.

While being included in the IRP formulation the value of remaining service time (Q_i^{rt}) of each disrupted cloud service, i.e., (Q_i^{rt}) defined in (4.1), is normalized as follows:

$$t_i = \left\lceil 100 \times \frac{Q_i^{rt}}{RT} \right\rceil, \quad RT = max(Q_i^{rt}), \forall Q_i \in Q. \tag{4.2}$$

In this way the remaining service time has a value t_i that is always within the [1,100] interval, making it easy to be used in a multi objective cost function, as explained later in this section.

When a cloud service is relocated it is important to update the values of t_i to account for the time spent during the relocation process. For a given cloud service i relocated to DC k (i.e., $Q_{i,k}$) the *relocation downtime* $Q_{i,k}^{rd}$ is proportional to: *(i)* the number of storage units to be relocated (Q_i^{st}), *(ii)* the rate at which storage units are transmitted (i.e., Δ, expressed in number of storage units per time unit), *(iii)* the distance between the DC node where Q_i was being served before the failure and the DC node k where Q_i is relocated (i.e., $d_{Q_i^{dst},k}$, expressed in km), and *(iv)* the propagation time (i.e., Θ, the speed of light in fiber, expressed in km per time units). More formally $Q_{i,k}^{rd}$ can be expressed as:

$$Q_{i,k}^{rd} = \frac{Q_i^{st}}{\Delta} + \frac{d_{Q_i^{dst},k}}{\Theta}, \forall Q_i \in Q, k \in N_{DC}. \tag{4.3}$$

When a cloud service is relocated the value of t_i defined in (4.2) is modified in order to take into account the impact of the relocation downtime, as:

$$t_{i,k} = \begin{cases} \left\lceil 100 \times \frac{(Q_i^{rt} - Q_{i,k}^{rd})}{RT} \right\rceil & \text{, if } Q_i^{rt} - Q_{i,k}^{rd} > 0 \\ 0 & \text{, otherwise} \end{cases}, RT = max(Q_i^{rt}), \forall Q_i \in Q, k \in N_{DC}. \tag{4.4}$$

The proposed IRP model relies on the following inputs and variables.

Inputs:

- W_{xy}: the number of free wavelengths on fiber link $(x, y) \in E$;
- $\alpha_i \in A$: the set of priority values of $Q_i \in Q$;
- t_i: normalized value of Q_i^{rt}, when $Q_i \in Q$ is restored but not relocated, calculated according to (4.2);
- $t_{i,k}$: normalized value of Q_i^{rt}, when $Q_i \in Q$ is relocated to DC $k \in N_{DC}$, calculated according to (4.4);
- Q_i^{dst}: the node serving $Q_i \in Q$ before the failure;
- Q_i^{st}: storage units required by $Q_i \in Q$;
- Q_i^{pu}: processing units required by $Q_i \in Q$.

Variables:

- wl_{xy}: the total number of wavelengths used by the restoration paths on fiber link $(x, y) \in E$;
- $wl_{xy}^i \in \{0, 1\}$: equal to 1 if the restoration path of cloud service Q_i traverses fiber link (x, y), 0 otherwise;
- $A_i \in \{0, 1\}$: equal to 1 if the service Q_i is successfully restored, 0 otherwise;
- $A_{i,k} \in \{0, 1\}$: equal to 1 if cloud service Q_i is successfully restored using the DC at node $k \in N_{DC}$.

The formulation of the PR-SRD problem is presented next.

Objective function:

$$min \sum_{Q_i \in Q} \left[\alpha_i \left(t_i - \sum_{k \in N_{DC}} t_{i,k} \times A_{i,k} \right) \right] + \beta \sum_{Q_i \in Q} \sum_{k \in N_{DC} | k \neq Q_i^{dst}} A_{i,k} + \gamma \sum_{(x,y)} wl_{x,y} \quad (4.5)$$

Subject to:

$$\sum_{\forall n \in N} wl^i_{nj} - \sum_{\forall m \in N} wl^i_{jm} = \begin{cases} -A_i \text{ , if } j = Q_i^{src} \\ A_{i,j}, \text{ if } j \in N_{DC},\ \forall Q_i \in Q, \forall j \in N \\ 0 \text{ , otherwise} \end{cases} \quad (4.6)$$

$$wl_{xy} = \sum_{\forall Q_i} wl^i_{xy},\ \forall (x,y) \in E \quad (4.7)$$

$$wl_{xy} \leq W_{xy},\ \forall (x,y) \in E \quad (4.8)$$

$$\sum_{\forall k} A_{i,k} \leq 1,\ \forall Q_i \in Q \quad (4.9)$$

$$A_{i,k} = 0,\ \forall k \in N_{DC}, \forall Q_i \in Q | k \neq Q_i^{dst} \wedge Q_i^{rt} \leq Q_i^{rd} \quad (4.10)$$

$$\sum_{\forall Q_i \in Q} \left(Q_i^{st} \times A_{i,k} \right) \leq DC_k^{st}, \forall k \in N_{DC} \quad (4.11)$$

$$\sum_{\forall Q_i \in Q} (Q_i^{pu} \times A_{i,k}) \leq DC_k^{pu}, \forall k \in N_{DC} \quad (4.12)$$

$$t_{i,k} = \begin{cases} t_i & \text{, if } k = Q_i^{dst} \\ \left\lceil 100 \times \frac{(Q_i^{rt} - Q_{i,k}^{rd})}{RT} \right\rceil & \text{, if } Q_i^{rt} - Q_{i,k}^{rd} > 0, \\ 0 & \text{, otherwise} \end{cases} \quad (4.13)$$
$$RT = max(Q_i^{rt}), \forall Q_i \in Q, k \in N_{DC}.$$

Equation (4.5) describes the objective function consisting of three terms. The first term is the sum of the downtime of all the cloud services (each one weighted by its priority value) that cannot be restored. By minimizing this term, the model minimizes the amount of downtime caused by the failure in the network. The second term counts the number of cloud services that require relocation while being restored. By minimizing this term the model avoids unnecessary relocation and its associated overhead. The third term accounts for the number of networking units used by all the restoration paths. By minimizing this term the model ensures to use the minimum amount of resources while calculating the solution, and prevents the model and the restoration process to have a larger impact on the blocking probability of future arriving requests.

The role of the β and γ parameters is to make sure that the results obtained after solving the ILP formulation represent a good trade-off among the metrics considered in this work: overall cloud service downtime, the wavelength resource usage, and the number of cloud service relocations. For example setting a value of β similar to γ encourages a higher number of relocations, if relocating a cloud service helps in saving networking resources (i.e., by having a shorter restoration path). When the value of β gets closer to the value of α, the model might prefer solutions where cloud services are not restored (i.e., dropped) in order to reduce the number of relocations.

Constraint (4.6) guarantees the flow conservation of each restored cloud service. Constraint (4.7) computes the total number of networking units used on each fiber link for restoration purposes. Constraint (4.8) ensures that the number of networking units used on each fiber link does not exceed the actual number of available resources. Constraint (4.9) checks that each relocated cloud service uses at most one DC. Constraint (4.10) ensures that a cloud service cannot be relocated to DC node k if the relocation downtime of moving the service to that DC node is larger than the remaining service time. Constraints (4.11) and (4.12) ensure that a DC node cannot be used to relocate a cloud service if it has not enough IT resources (either storage or processing units) to accommodate it. Finally, constraint (4.13) computes the normalized value of the remaining service time of $Q_i \in Q$, which is used in the objective function.

4.2.2 Heuristic for the PR-SDR Problem

This section describes a heuristic, referred to as Heuristic for Relocation with Priorities (HRP), which can be used to solve the PR-SRD problem as an alternative to the IRP model. The objective of the heuristic is the same as of the IRP, i.e., the minimization of the following three metrics: *(i)* the average downtime value of all unrestored cloud services, *(ii)* the number of successfully restored cloud services that needed to be relocated, and *(iii)* the number of networking units used in the restoration process. The HRP heuristic is described in Algorithm 1 and it works as follows.

First the heuristic defines a sorted set Q' obtained from Q after applying a weight function Q_i^{wt} to each $Qi \in Q$. The role of Q_i^{wt} (defined as $Q_i^{wt} = \omega \times Q_i^{rt}$) is to make sure that cloud services belonging to higher priority classes are considered first in the restoration process. Based on this rationale high priority cloud services will be assigned a value of ω that is higher than the one used for low priority services. The heuristic then

Algorithm 1 HRP

```
$Q$' = $Q$ sorted by $Q_i^{wt}$
for all $Q_i \in Q$' do
    curRoute = selRoute = selDC = NULL
    if shortestPath($Q_i^{src}$, $Q_i^{dst}$) $\neq NULL$ then
        selRoute = shortestPath($Q_i^{src}$,$Q_i^{dst}$)
        restorePath($Q_i$, selRoute)
    else
        for all $DC_k \in N_{DC} \mid Q_i^{rt} > Q_{i,k}^{rd} \wedge DC_k^{st} \geq Q_i^{st} \wedge DC_k^{pu} \geq Q_i^{pu}$ do
            curRoute = shortestPath($Q_i^{src}$,$Q_i^{dst}$)
            if hopCount(curRoute) < hopCount(selRoute) then
                selRoute = curRoute
            end if
        end for
        if selRoute $\neq$ NULL then
            relocateAndRestorePath($Q_i$, selRoute)
        else
            dropService($Q_i$)
        end if
    end if
end for
```

tries to restore each cloud service $Q_i \in Q'$ sequentially as described next.

For each $Q_i \in Q'$, the heuristic first checks if there is a path with enough networking resources in $G(N, E)$ from Q_i^{src} to the DC node already in use (i.e., Q_i^{dst}). This step is done in order to reduce the number of unnecessary cloud service relocations. If such a path exists, then a new networking path from Q_i^{src} to Q_i^{dst} is established. If it does not, the heuristic tries to check if it is possible to connect Q_i^{src} with a different DC with enough storage (Q_i^{st}) and processing units (Q_i^{pu}) to accommodate Q_i. While checking the presence of free storage and compute resources the heuristic also makes sure that the value of the relocation downtime, i.e., $Q_{i,k}^{rd}$, for the candidate DC node under evaluation does not exceed the value of the remaining service time of Q_i. If more than one DC node fulfilling the requirements are reachable, the heuristic chooses the one that is the closest in terms of hops to Q_i^{src} (this is done to minimize the number of networking unit resources used to restore Q_i). Once this new DC and route are selected, Q_i is relocated and a networking path from Q_i^{src} to the new DC is established. In case neither an available path nor an alternate DC can be found, the cloud service is dropped.

The heuristic uses external functions to calculate shortest path and hop count, described as follow:

shortestPath(**a**,**k**)

This function uses a pre-computed set of k-shortest-paths, generated using a *k*-

shortest-path algorithm applied to the $G(N, E)$ graph which creates a set of paths between each node pair, and returns the shortest path available (i.e., a path with at least one networking unit available at each segment or link) path between two given nodes (i.e., a and k, usually a source node and a DC, respectively). If either no such paths exist or there are no available networking resources on any of the pre-computed paths, this function returns $NULL$;

hopCount(r)

This function returns the number of hops of a given route r, or MAX_VALUE when received argument is $NULL$;

restorePath(i,r)

This function is responsible to establish a new networking path for service i following route r and update network state accordingly;

relocateAndRestorePath(i,r)

This function is responsible to relocate service i to the destination of route r, establish a new networking path for service i following route r and update network state accordingly;

dropService(i)

This function is responsible to drop the service i and release all the associated resources, updating the network state, i.e., $G(N, E)$ accordingly.

5 Performance Assessment

This chapter investigates the performance of the IRP and HRP approaches described in the previous chapter. The first section describes the assumptions used in the performance evaluation work. The second section presents and discusses the simulation results.

5.1 Simulation Setup and Metrics

The IRP and HRP strategies are evaluated considering cloud services with and without traffic classes differentiation, and assuming two network topologies. The purpose of considering two topologies is to assess how different network connectivity characteristics (i.e., average nodal degree) and sizes may affect the performance of IRP and HRP. With this purpose, three scenarios are examined, i.e., Scenario A, B, and C. The first two scenarios use the NSF network with 14 nodes and 21 links showed in Fig. 10a as the reference topology (AHMED et al., 2014). In this network nodes 3, 4, 10, and 11 are assumed to host DataCenters (DCs) because they are among the most connected nodes in the network. Thus, nodes 3, 4, 10 and 11 are referred to as DC nodes. Each DC node in the NSF topology is equipped with 15000 storage units and 900 processing units.

The third scenario (i.e., Scenario C) uses the Italian network with 32 nodes and 56 links, showed in Fig. 10b, as the reference topology. In this network nodes 4, 9, 12, 13 and 14 are assumed to be DC nodes, because of their high connectivity. In the Italian network, each DC node is equipped with 21000 storage units and 1300 processing units. In both topologies DC nodes are located in the network nodes with the highest nodal degree. The rationale is the following: the higher is the nodal degree the higher are the chances that the DC storage and processing resources can be reached by the other nodes in the network (i.e., there will be more networking resources available to reach the DC node resources) (DEVELDER et al., 2013). As a result, in the NSF network DCs are connected to network nodes that on average have nodal degree 3.5, while in the Italian network DC are connected to network nodes that have nodal degree 5. This is the reason behind the choice of having in the Italian network 1.4 times more storage and processing units than in the NSF network.

The positioning of DCs, i.e., the choice of which network node to deploy a DC, impacts the reachability of such a DC. The positioning can be based on several criteria, such as reliability, electricity availability, etc. Nevertheless, the positioning of DCs in this work follows a straight forward approach. Thus, all the experiments use the same DC positioning, which will not affect the results of this work.

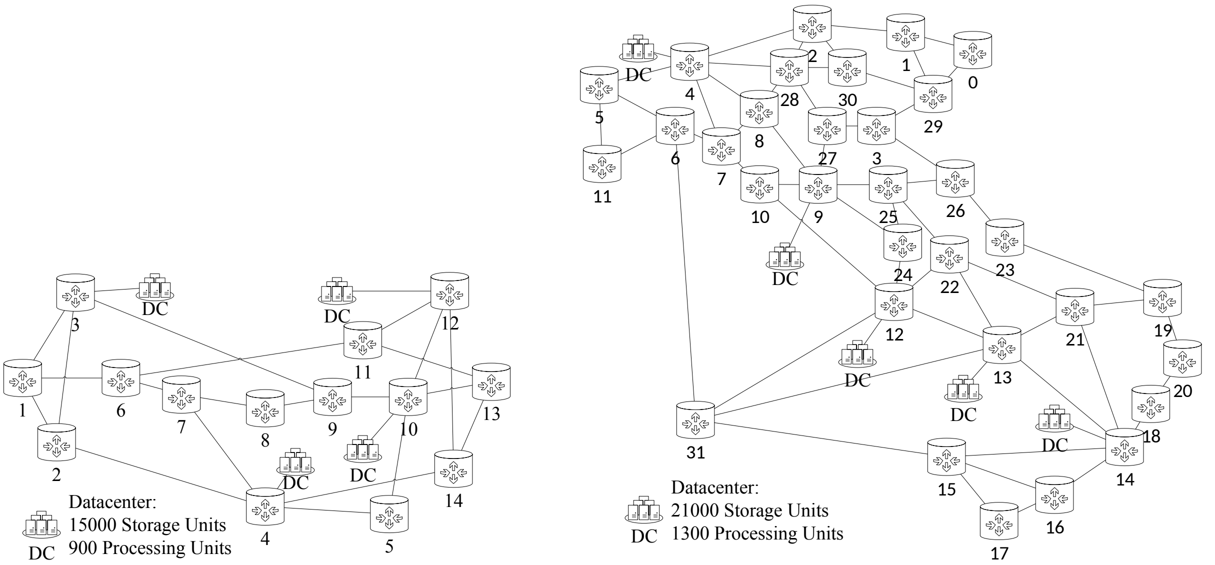

(a) NSF topology with 14 nodes and 21 links

(b) Italian topology with 32 nodes and 56 links

Figure 10 – Network topologies.

To ensure that the dimensioning of all the DCs is coherent to the network connectivity, the number of IT resources (i.e., processing and storage units) is defined mainly by the average node degree of the nodes connected to the DCs, the number of networking units existing at each link connecting the DC and the average number of IT resources required by each cloud service. The reasoning is that, in average, the DCs would be capable of hosting all the cloud services that the network can connect to it. Equations (5.2) and (5.1) show how to calculate approximately the number of storage and processing units, respectively, where $W_{(x,k)}$ is the number of networking units connecting DC k to any other node.

$$DC_k^{st} \approx \sum_{(x,k)\in L} W_{(x,k)} \times AVG(Q_i^{st}), \forall k \in N_{DC} \tag{5.1}$$

$$DC_k^{pu} \approx \sum_{(x,k)\in L} W_{(x,k)} \times AVG(Q_i^{pu}), \forall k \in N_{DC} \tag{5.2}$$

For instance, considering NSF topology, the DCs have in average 3.5 links connecting them to other network nodes, and 80 networking units at each link. The services have in average 51 storage and 3 processing storage. This scenario gives us the following version of Eq. (5.1):

$$DC_k^{st} = 3.5 \times 80 \times 51 = 14,280.$$

In order to ensure a more readable quantity, and as the addition of some IT resources will not impact significantly in the overall optical cloud cost neither in the performance results, we decide to round this number up, adopting 15000 storage units in all the DCs.

In Scenario A, all cloud services belong to the same traffic class (i.e., they all have the same priority value $\alpha_i = \alpha, \forall Q_i \in Q$). In Scenario B and C, cloud services belong to two different traffic classes: one with high priority, and the other with low priority. The probability distribution to choose the class of a service is uniform. It is assumed that the cloud services in each experiment have 20% probability to be in high priority, while the remaining part, i.e., 80%, has low priority. The notion of priority is used only during the restoration phase, and it has no influence on how cloud services are initially provisioned in the network, i.e., during normal network operations.

Regardless of the topology, all fiber links in the network are bidirectional, with 80 networking units in each direction. DCs are assumed to be co-located with the network nodes to which they are connected. For this reason, the fiber links connecting DCs to their respective network nodes are assumed to have always enough capacity to cater for the traffic to and from the DC (i.e., they are not the bottleneck of the system). Additionally, the links connecting the DC to its respective network node are not considered to fail at any time. All fiber links in both the NSF and the Italian topology are assumed to have the same length, while all network nodes have full wavelength conversion capability.

Each simulation experiment consists in establishing one million cloud services, each one to be provisioned from a client (i.e., non-DC) to a DC node that has enough storage and processing resources to accommodate the cloud service requirements. The amount of storage and processing units required by each cloud service are chosen uniformly in the intervals [1,100] and [1,5] respectively (AHMED et al., 2014). Connecting a client node to a DC node requires the establishment of a networking path with a capacity equal the capacity of one networking channel. The holding time of each cloud service is exponentially distributed with an average value of 60 time units. The arrival rate of the cloud services follows a Poisson distribution, where the mean time between arrivals that varies according to the load value chosen for the specific experiment. The client node at which a cloud service originates is uniformly selected among all non-DC network nodes. Unless otherwise stated the value of Δ (i.e., the rate at which storage units are transmitted) is equal to 100 [storage_units/s]. This assumption refers to a scenario where transponders in the optical networks work at 100 Gbps, with storage units of approximatively 1.3 Gb in size. These values are in line with the assumptions presented in (GIFRE et al., 2015; WOOD et al., 2014). The value of Θ is assumed to be equal to 2×10^5 [km/s], while the hop length for all the links in the network in the NSF and Italian topology is set to 1086 and 224 km, respectively.

The simulation study presented in the thesis considers a single fiber link failure scenario. All fiber links in the network have the same probability to fail. The time between two consecutive fiber link failures is exponentially distributed, with a mean value (i.e., MTTF) equal to 1000 time units, whereas the fiber link reparation time is exponentially distributed with a Mean Time To Repair (MTTR) equal to 10 time units, which gives availability of 99% for each link (see Eq. (2.4) for more details).

In Scenario A (i.e., where cloud services belong to the same traffic class) α, β, and γ are equal to 10^5, 10^4, and 1, respectively, while $Q_i^{wt} = Q_i^{rt}$, for all cloud services. It means that in the HRP strategy decisions on which cloud service should be first restored are based only on their respective value of the remaining service time. In Scenarios B and C (i.e., where cloud services have different priority values), the values of β and γ are still 10^4, and 1, respectively. The value of α_i, on the other hand, is different and varies with the cloud service type. For high priority services $\alpha_i = 10^5$, while for low priority services $\alpha_i = 10^4$. Finally, Q_i^{wt} is defined as follows:

$$Q_i^{wt} = \begin{cases} 8 \times Q_i^{rt} & ,\forall Q_i \in Q \mid \alpha_i = 10^5 \\ Q_i^{rt} & ,\forall Q_i \in Q \mid \alpha_i = 10^4 \end{cases}. \tag{5.3}$$

Equation (5.3) is used in the HRP strategies to make sure that cloud services with high priority ($\alpha_i = 10^5$) are considered first during the sequential restoration process.

Note that the choice for the values of the α, β, and γ parameters is the result of a number of tests aimed at finding which combination is able to guarantee a good trade-off among the performance parameters defined in (4.5).

In normal operating conditions (i.e., in the absence of a failure) each cloud service is provisioned in the network upon request using the DC_CLOSEST heuristic (LANDI et al., 2012), which chooses the DC node with enough storage and computing resources, and that is the closest to the client node at which a cloud service originates. Upon the occurrence of a failure, the disrupted cloud services are restored using either one of the evaluated strategies (i.e., IRP or HRP). The process is not revertive, i.e., once a failure is repaired each restored cloud service will not be switched back to its original networking path and/or DC node. The function *shortestPath()* used by HRP heuristic uses a set of $k = 10$ pre-computed paths. Cloud service provisioning and restoration operations are assumed to be coordinated by a Path Computation Element (PCE) (AHMED et al., 2012) or Software Defined Network (SDN) based controller specifically designed for concurrent optimization of IT and transport resources. For benchmarking purposes, an additional restoration strategy that does not allow for relocation is also considered in the study. In this way, the assessment of the benefits introduced by the relocation procedure during the restoration process can be highlighted. This benchmarking approach is based on the solution of an ILP formulation and it is referred to as ILP_PR (AHMED et al., 2013).

All the results presented in the next section are the average of 100 experiments using different Random Number Generator (RNG). The confidence interval of these results never exceeds 5%, and it has been calculated assuming a confidence level of 95%. Experiments are carried out using a Java-based discrete event driven simulator (AHMED; NAWAZ, 2004) adapted to the scenario of this thesis. The ILP formulation is solved using the Gurobi Optimizer (OPTIMIZATION, 2016). A Debian Linux workstation with an Intel Xeon Central Processing Units (CPUs) (6 cores per CPU) clocked at 2.2 GHz and with 32 GB of RAM is used for the simulations.

5.2 Results

The performance of the IRP and HRP strategies is evaluated in terms of the following metrics: *(1) blocking probability*, defined as the ratio between the number of cloud services that could not be successfully provisioned in the network and the total number of service requests; *(2) average availability*, defined as the ratio between the sum of the uptime of all the provisioned cloud services and the sum of the their service holding time values; *(3) average restorability*, defined as the ratio between the number of cloud services that were successfully restored and the number of cloud services disrupted by a failure), and *(4) average relocations*, defined as the ratio between the number of restored cloud services that required relocation and the number of cloud services that were successfully restored.

Figure 11 presents the performance results for Scenario A. The proposed strategies (i.e., both IRP and HRP) show a slightly worse performance in terms of blocking probability than the benchmark ILP_PR i.e., Fig. 11a. It can be explained by the better performance in terms of restorability (Fig. 11c) that, in turns, leads to less resources available for provisioning future traffic. Figure 11b presents the average availability values, where the two red dashed lines on the top of the figure represent the five 9s and four 9s availability values. The figure confirms the intuition that allowing for the relocation of cloud services to alternative DC nodes during the restoration process has a beneficial effect on the average value of the cloud service availability. More specifically, both IRP and HRP show availability performance exceeding four 9s at low and medium load conditions, while ILP_PR is not able to guarantee the same availability performance.

As already mentioned, IRP and HRP show also very good performance improvement in terms of average restorability values, i.e., Fig. 11c, when compared to ILP_PR. The IRP strategy shows up to 48% better restorability performance compared to ILP_PR, while with HRP the improvement versus the benchmark strategy is at most 42%.

Figure 11d shows the average number of relocations needed to restore the cloud services in relation to the number of restored services. The lower this number is, the better, because relocations are costly in terms of Virtual Machine (VM) state transfer overhead. It can be noticed from the figure that in the best case (i.e., with the IRP strategy) only

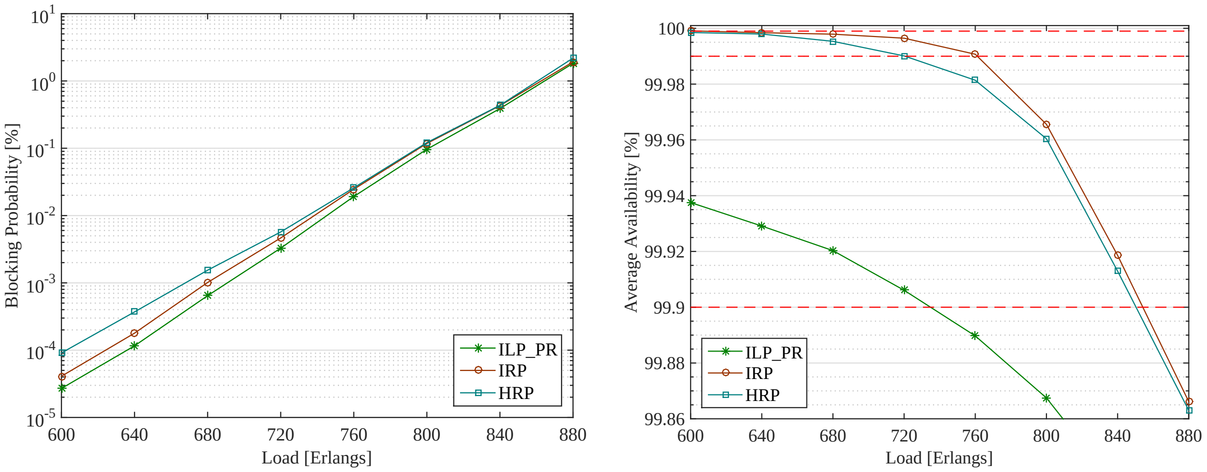

(a) Average blocking probability (b) Average availability

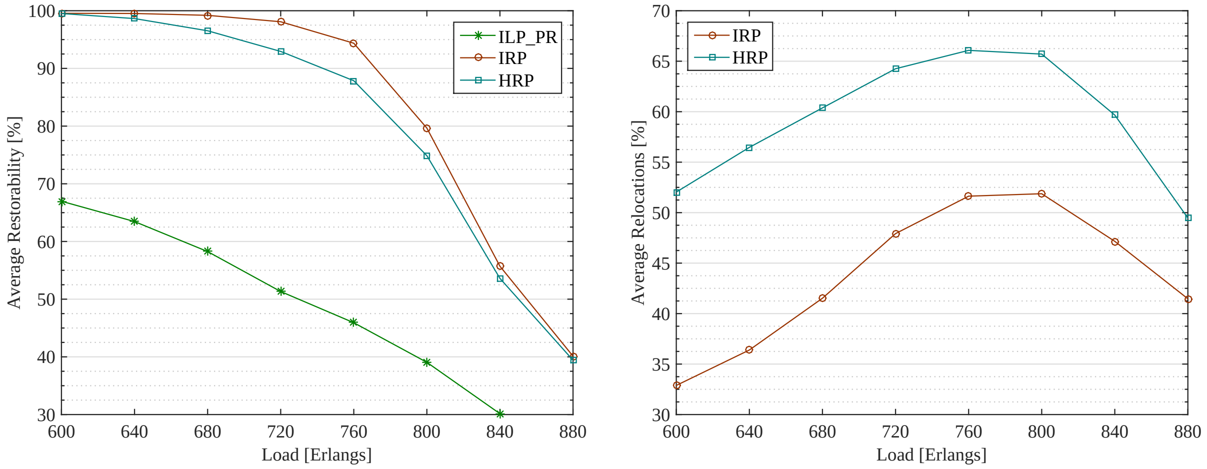

(c) Average restorability (d) Average percentage of relocations

Figure 11 – Simulation results for Scenario A.

52.5% of the successfully restored cloud services needed to be relocated.

Finally, the performance results presented in Fig. 11 allows to estimate how close the results from the HRP are from the optimum, i.e., the results of the IRP. In almost all the cases HRP behaves closely to IRP. The only higher difference is in the number of relocations. This behavior can be explained by the sequential nature of the heuristic that restores one cloud service at a time, not having the benefits that a concurrent approach such as IRP has.

Figure 12 presents the performance results for Scenario B (i.e., NSF topology with different traffic classes). In the figure the "_P1" and "_P2" notation refers to the results specific to the cloud services belonging to the traffic class with high and low priority, respectively. Since the notion of priority is used only during the restoration process, the results for the blocking probability of all the proposed strategies, i.e., Fig. 12a, are the same as in Scenario A. This is not the case for the other metrics. Figure 12b presents the availability performance of the IRP and the HRP. The figure presents curves for both

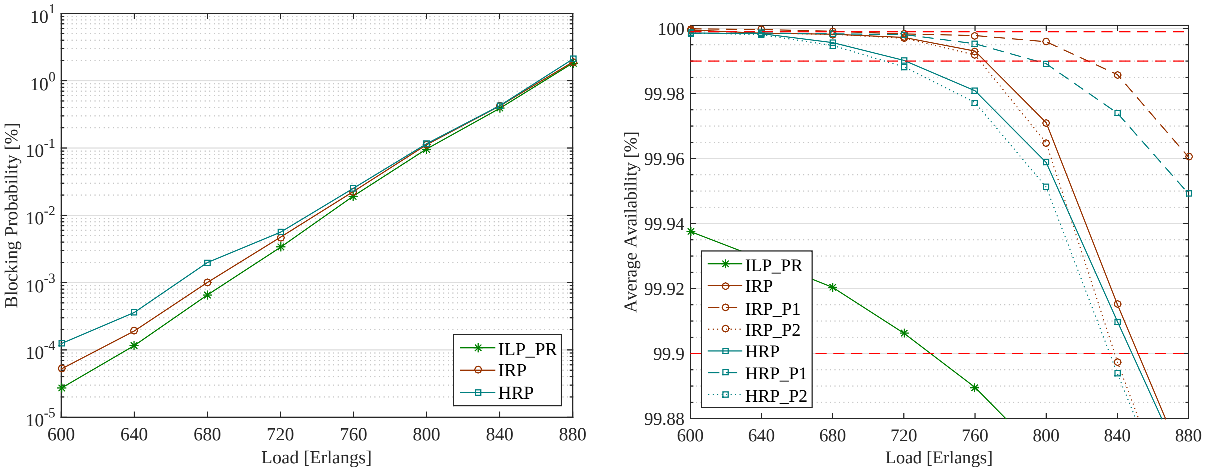

(a) Average blocking probability (b) Average availability

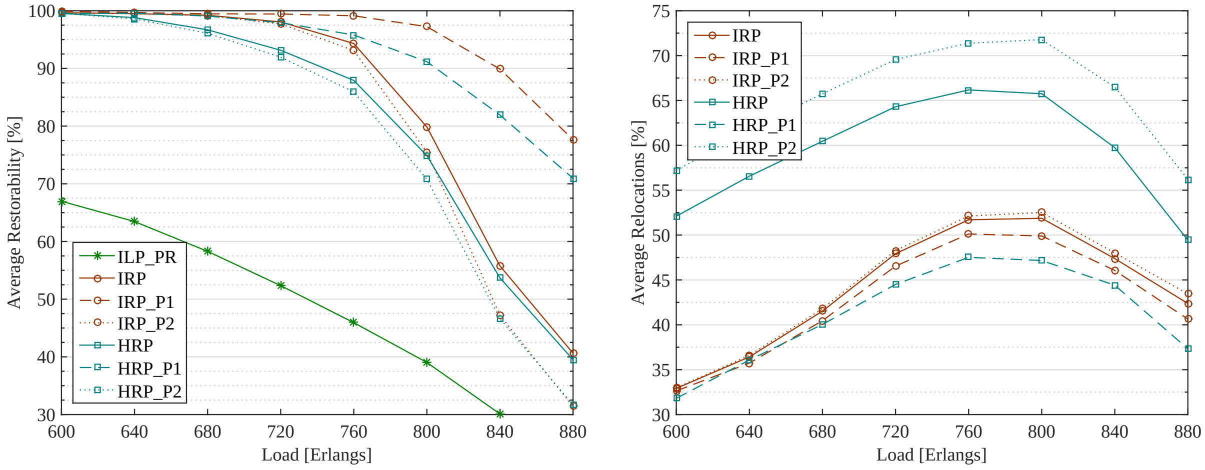

(c) Average restorability (d) Average percentage of relocations

Figure 12 – Simulation results for Scenario B.

the average (IRP curve) and the per-class (i.e., IRP_P1 and IRP_P2 curves, specific for each priority) performance values. The average availability performance (IRP curve) is close to the one of Scenario A, i.e., Fig. 12b, leading to the conclusion that IRP and HRP present the same benefits over the performance of ILP_PR. In other words, the presence of traffic classes does not have a negative effect on the general performance of the proposed restoration strategies.

On the other hand, the good availability performance for the cloud services in the high priority class is achieved at the expense of a slight degradation of availability values of cloud services in the low priority class. This reduction in performance is not dramatic since it is close to the average value (IRP curve). The HRP approach offers availability results that are very close to IRP.

Figure 12c presents results for the restorability performance. The conclusions that can be drawn are similar to the one just discussed for the availability results. The average performance (i.e., the IRP and HRP curves in the figure) are very close to the ones of

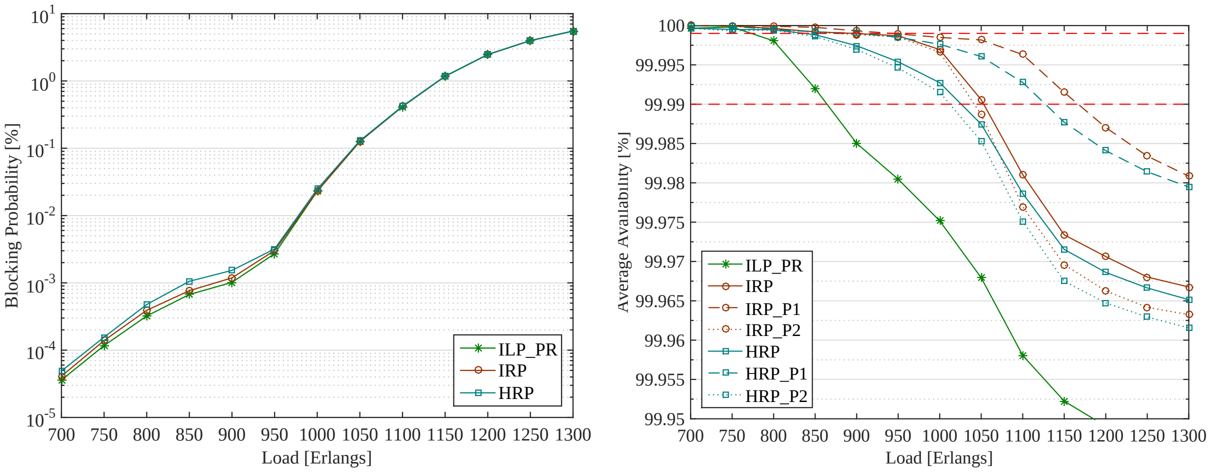

(a) Average blocking probability

(b) Average availability

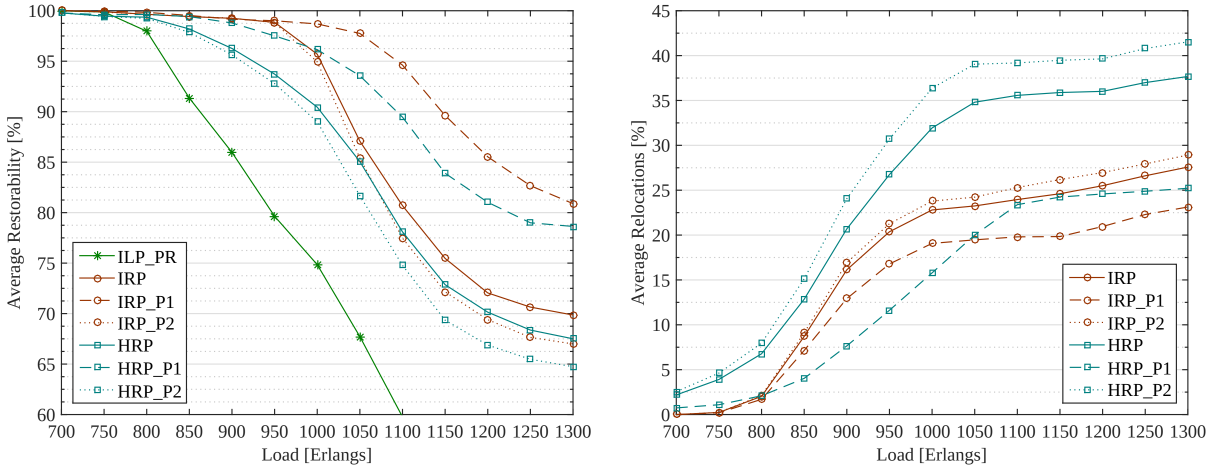

(c) Average restorability

(d) Average percentage of relocations

Figure 13 – Simulation results for Scenario C.

Scenario A, i.e., Fig. 11c, while in the presence of cloud services belonging to different traffic classes it is possible to accommodate the needs of the high priority cloud services without compromising much the restorability performance of the cloud services with low priority. The low priority services experience a decrease of at most 9% in restorability compared to the average case (i.e., the IRP curve), a value that is still a good improvement over the ILP_PR curve that does not allow neither for relocation nor for handling cloud services with priority.

Figure 12d presents the average relocation performance. Moreover, in this case the results show that the introduction of priorities in the system does not have a large influence on the average number of relocations, i.e., when comparing these results with the ones in Fig. 12d. All three curves for the IRP case show a similar behavior, with a maximum number of relocations of 52.5% for the low priority case. The HRP shows average results (HRP curve) close to the optimal, i.e., presented in Fig. 12d, but with a number of relocations for the low priority cloud services that exceed 70% at medium to high load conditions.

Figure 13 presents the performance results for Scenario C (i.e., Italian topology with different traffic classes). The load values change compared to the ones presented in Scenarios A and B in order to have comparable values in terms of blocking probability. From the results presented in Fig. 13 it can be noticed that in general both the IRP and the HRP present the same advantages in terms of average availability and restorability that were highlighted in Scenario B, confirming the general validity of the conclusions drawn so far. One difference that is worth noting is about the absolute values of the various metrics. As it can be seen in Fig. 13, both IRP and HRP can achieve on average higher availability and restorability values, with a lower required number of relocations. It can be explained by the nature of the Italian topology that is on average more connected than the NSF.

Figure 14 presents the performance results of the IRP formulation in Scenario A as a function of different values of Δ. Three cases are considered. Each case is used to represent a transmission rate of optical transport networks currently deployed (i.e., $\Delta = 10$ and/or $\Delta = 100$), and of optical transport networks that will most probably be deployed in the short term future ($\Delta = 400$). As it was explained in Section 4.2.1 the value of Δ has an impact on the relocation downtime. For this reason, it is interesting to understand under which conditions the relocation downtime plays a crucial role in the restorability performance of cloud services.

As it can be expected, the figure shows that in the scenario under consideration the higher is the value of Δ the lower is the impact on the relocation downtime. In general, the different values of Δ do not impact significantly in the blocking probability. The same behavior can be found for high loads in the availability and restorability (i.e., Figs. 14b and 14c, respectively). When considering $\Delta = 10$ for low and medium loads, the availability and restorability are substantially impacted. However, considering the upgrade of Δ from 100 to 400 do not significantly improve the availability and restorability performance, which means that considering $\Delta = 100$ have a good compromise between current and future network technologies.

Table 5 presents the average processing time results for each restoration attempt for ILP_PR, IRP and HRP in each one of the presented scenarios. Three load values (i.e., low, medium and high) are considered. The processing time of HRP is always below 1 millisecond in all load conditions, whereas IRP requires processing times that are orders of magnitude higher. The processing time is an important metric especially in scenarios in which failures can affect a high number of cloud services. Given the relative small difference in availability and restorability performance of the HRP compared to IRP, it can be concluded that HRP is a very good compromise between performance and complexity.

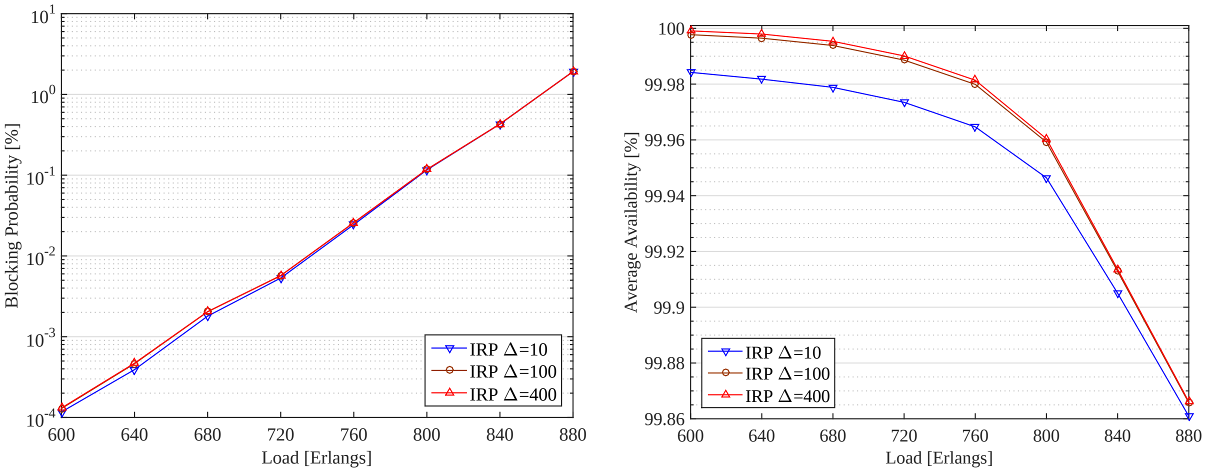

(a) Average blocking probability (b) Average availability

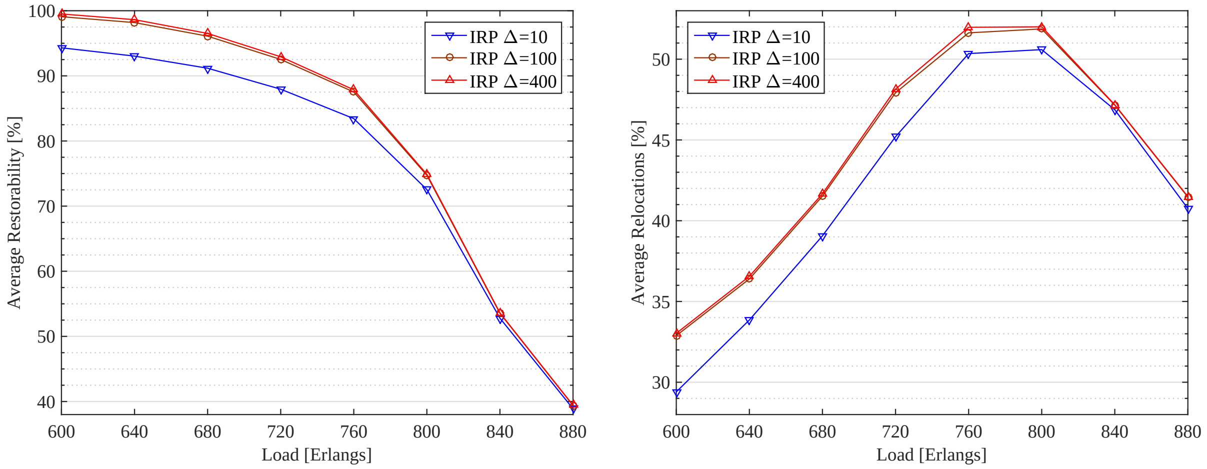

(c) Average restorability (d) Average percentage of relocations

Figure 14 – Simulation results for Scenario A for different Δ values.

Table 5 – Average processing time for each disrupted cloud service restoration attempt

		Processing Time (ms)		
	Load (Erlangs)	**HRP**	**IRP**	**ILP_PR**
Scenario A	600	0.089	138.66	129.98
	760	0.104	148.82	134.23
	880	0.199	201.86	160.27
Scenario B	600	0.085	138.47	-
	760	0.095	142.99	-
	880	0.188	200.79	-
Scenario C	700	0.039	225.81	-
	1000	0.074	276.49	-
	1300	0.131	306.12	-

6 Conclusions and Final Remarks

This thesis proposes a restoration-based survivability strategy that can be used to recover cloud services disrupted by fiber link failures. The intuition behind the proposed strategy is to combine the service relocation and the service differentiation concepts. The former is used to enhance the average service availability and restorability performance offered by a restoration-based approach. The latter is leveraged to make sure that when cloud services with different priorities compete for the same backup resources, critical services are given precedence over non-critical ones.

We proposed both a solution based on ILP formulation referred to as ILP for Relocation with Priorities (IRP) and on heuristic referred to as HRP. It is shown that HRP provides results very close to IRP, but with a significantly lower processing time. The performance of both IRP and HRP were assessed in a dynamic provisioning simulation, considering two different optical cloud network scenarios. The results from the performance assessment study show that both IRP and HRP are able to improve the average service availability and restorability performance with a limited number of cloud service relocations when compared to conventional restoration-based techniques. In addition, thanks to service differentiation, the availability and restorability performance of critical cloud services are very close to the one achievable with a protection-based strategy, but with the inherent benefits in terms of efficient resources usage deriving from a restoration-based approach.

6.1 Contributions

This thesis identifies a set of survivability research challenges on the operation of optical clouds. Section 3.3 presented works related to the survivability of optical clouds, however the evaluation of survivability performance on the operation of clouds remained an unexplored research topic.

The scenario definition and problem statement for optical clouds is a strong contribution of this work. The joint definition of optical networks and cloud computing infrastructures is an important step towards more complete simulation environments. The optical cloud environment is heterogeneous, and need to be defined considering the characteristics of each component.

Application requirements are described in the literature. However, a more detailed requirement definition for applications can be used to simulate a more real system. Translate high level QoE and system requirements (e.g., maximum response time, elasticity, flexibility and reliability) into optical cloud QoS requirements (e.g., amount of IT resources and service priority), as well as the independent variation of each resource requirement, are key points to provide the expected quality for customers and end-users.

Based on the scenario generated by the previous contributions, the restoration strategies can be evaluated in many different aspects. The evaluation of the impacts of using different inter-DC data rates is one important contribution of this work and shows the impacts of inter-DC capacities in the survivability of cloud services which rely on the restoration strategies.

This thesis has also contributions that goes beyond scientific work. During the development of this thesis, the already established relation with researchers from KTH in Sweden could be strengthened. The author received scholarship from Science without Borders program and could stay one year working closely with KTH researchers, and such joint work continues with publication and submission of more works related to the optical networks field. The author was part also of a partnership between Brazilian government and SAAB, which allowed the author to contribute with a program to develop a plan for Brazilian smart cities deployment.

During the development of this thesis, the author participate of more than 10 research and development projects. The projects involved the development of software solutions to be employed in regional and national private and public sector.

6.2 Future Works and Trends

This work presented solutions to enhance the performance evaluation of restoration techniques in optical clouds. However, the optical cloud subject cover a multitude of components, each one having is own requirements. The research of each component can enhance the comprehension and improve the modeling details of optical clouds.

The complexity analysis of proposed strategies (e.g., ILP and heuristic) can increase the knowledge about the models, and provides information on how to improve the performance of the models. In this way, a more detailed study and comprehension of algorithm complexity is a good contribution to the field.

The lower part of Fig. 8 shows the interaction of end-users with cloud services. Understand the characteristics of this interaction can help Cloud Service Providers (CSPs) and Internet Service Providers (ISPs) improve the perceived quality, QoS and QoE of its services.

In the last years disaster survivability has gained attention from scientific community (MUKHERJEE; HABIB; DIKBIYIK, 2014b). Study how to protect data and services against disaster failures is a crucial topic today (FERDOUSI et al., 2015). Dimensioning and provisioning disaster survivable optical clouds are important to reduce network dimensions or blocking probability while increase network survivability (GU et al., 2015; SAVAS et al., 2014).

Optical clouds are becoming important also for the wireless communications. The virtualization of wireless network components and the respective assignment from antenna equipment to optical clouds helps to lower antenna prices, as well as allows the flexibility of allocate resources where there are end-users needing them. The increasing interest in Cloud Radio Access Networks or Centralized Radio Access Networks (C-RAN) architectures and its relation to the anycast traffic model can be an interesting subject for research nowadays (WANG et al., 2015; FIORANI et al., 2015a).

For the next decade the fifth generation of wireless communication will be an important point for the innovation in the wireless communication. To provide the required QoS and QoE combined with high flexibility and scalability optical clouds will play an important role. The Information and Communication Technology (ICT) infrastructure will need to meet all the user application and networking demands. Some works (e.g., (FIORANI et al., 2015b)) present studies in this direction.

Bibliography

AHMED, J. et al. A dynamic bulk provisioning framework for concurrent optimization in PCE-based WDM networks. *Lightwave Technology, Journal of*, v. 30, n. 14, p. 2229–2239, July 2012. ISSN 0733-8724. Disponível em: <http://dx.doi.org/10.1109/JLT.2012.2195296>. Cited in page 54.

AHMED, J. et al. Hybrid Survivability Schemes Achieving High Connection Availability With a Reduced Amount of Backup Resources [Invited]. *Journal of Optical Communications and Networking*, v. 5, n. 10, p. A152, set. 2013. ISSN 1943-0620. Disponível em: <http://www.opticsinfobase.org/abstract.cfm?URI=jocn-5-10-A152>. Cited 3 times in pages 13, 33, and 54.

AHMED, J. et al. Enhancing Restoration Performance Using Service Relocation in PCE-based Resilient Optical Clouds. In: *Optical Fiber Communication Conference*. Washington, D.C.: OSA, 2014. p. Th3B.3. ISBN 978-1-55752-993-0. Disponível em: <http://www.opticsinfobase.org/abstract.cfm?URI=OFC-2014-Th3B.3>. Cited 8 times in pages 13, 23, 29, 37, 39, 40, 51, and 53.

AHMED, J.; NAWAZ, S. Pose: A new discrete event optical simulator for the java platform. In: CITESEER. *National Conference on Emerging Technologies*. [S.l.], 2004. p. 184. Cited in page 55.

BILAL, K. et al. Trends and challenges in cloud datacenters. *IEEE Cloud Computing*, v. 1, n. 1, p. 10–20, maio 2014. ISSN 2325-6095. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6848725>. Cited 2 times in pages 12 and 16.

BUYSSE, J. et al. Exploiting relocation to reduce network dimensions of resilient optical grids. In: *2009 7th International Workshop on Design of Reliable Communication Networks*. IEEE, 2009. p. 100–106. ISBN 978-1-4244-5047-3. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5340020>. Cited in page 35.

BUYSSE, J. et al. On the impact of relocation on network dimensions in resilient optical Grids. In: *2010 14th Conference on Optical Network Design and Modeling (ONDM)*. IEEE, 2010. p. 1–6. ISBN 978-1-4244-5973-5. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5431591>. Cited in page 35.

BUYSSE, J. et al. Energy-Efficient Resource-Provisioning Algorithms for Optical Clouds. *Journal of Optical Communications and Networking*, v. 5, n. 3, p. 226, fev. 2013. ISSN 1943-0620. Disponível em: <http://www.opticsinfobase.org/abstract.cfm?URI=jocn-5-3-226>. Cited in page 19.

CAVDAR, C.; TORNATORE, M.; MUKHERJEE, B. Holding-time-aware and availability-guaranteed connection provisioning in optical WDM mesh networks. In: *2007 International Symposium on High Capacity Optical Networks and Enabling Technologies*. IEEE, 2007. p. 1–7. ISBN 978-1-4244-1828-2. Disponível em:

<http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4600278>. Cited 2 times in pages 27 and 35.

CHANCHIO, K.; THAENKAEW, P. Time-Bound, Thread-Based Live Migration of Virtual Machines. In: *2014 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*. IEEE, 2014. p. 364–373. ISBN 978-1-4799-2784-5. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6846472>. Cited in page 38.

CHIARAVIGLIO, L. et al. Is green networking beneficial in terms of device lifetime? *IEEE Communications Magazine*, v. 53, n. 5, p. 232–240, May 2015. ISSN 0163-6804. Cited in page 28.

COLMAN-MEIXNER, C. et al. A survey on resiliency techniques in cloud computing infrastructures and applications. *IEEE Communications Surveys Tutorials*, PP, n. 99, p. 1–1, 2016. ISSN 1553-877X. Disponível em: <http://dx.doi.org/10.1109/COMST.2016.2531104>. Cited in page 21.

COUTO, R. S. et al. Server placement with shared backups for disaster-resilient clouds. v. 93, p. 423–434, dec 2015. ISSN 13891286. Disponível em: <http://www.sciencedirect.com/science/article/pii/S1389128615003643>. Cited in page 23.

CUI, W. et al. Enhancing Reliability for Virtual Machines via Continual Migration. In: *2009 15th International Conference on Parallel and Distributed Systems*. IEEE, 2009. p. 937–942. ISBN 978-1-4244-5788-5. ISSN 1521-9097. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5395298>. Cited 2 times in pages 23 and 38.

DEVELDER, C. et al. Resilient network dimensioning for optical grid/clouds using relocation. In: *2012 IEEE International Conference on Communications (ICC)*. IEEE, 2012. p. 6262–6267. ISBN 978-1-4577-2053-6. ISSN 1550-3607. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6364981>. Cited in page 36.

DEVELDER, C. et al. Joint Dimensioning of Server and Network Infrastructure for Resilient Optical Grids/Clouds. *IEEE/ACM Transactions on Networking*, PP, n. 99, p. 1–1, 2013. ISSN 1063-6692. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6642129>. Cited 6 times in pages 12, 15, 16, 21, 36, and 51.

DEVELDER, C. et al. Survivable Optical Grid Dimensioning: Anycast Routing with Server and Network Failure Protection. In: *2011 IEEE International Conference on Communications (ICC)*. IEEE, 2011. p. 1–5. ISBN 978-1-61284-232-5. ISSN 1550-3607. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5963385>. Cited in page 36.

DEVELDER, C. et al. Optical Networks for Grid and Cloud Computing Applications. *Proceedings of the IEEE*, v. 100, n. 5, p. 1149–1167, maio 2012. ISSN 0018-9219. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6143975>. Cited 6 times in pages 5, 11, 12, 15, 17, and 18.

FERDOUSI, S. et al. Rapid data evacuation for large-scale disasters in optical cloud networks [invited]. *IEEE/OSA Journal of Optical Communications and Networking*, v. 7, n. 12, p. B163–B172, Dec 2015. ISSN 1943-0620. Cited 2 times in pages 36 and 62.

FIORANI, M. et al. Transport abstraction models for an sdn-controlled centralized ran. *IEEE Communications Letters*, v. 19, n. 8, p. 1406–1409, Aug 2015. ISSN 1089-7798. Cited in page 63.

FIORANI, M. et al. Transport abstraction models for an sdn-controlled centralized ran. *IEEE Communications Letters*, v. 19, n. 8, p. 1406–1409, Aug 2015. ISSN 1089-7798. Disponível em: <http://dx.doi.org/10.1109/LCOMM.2015.2446480>. Cited in page 63.

GHANI, N. et al. Control plane design in multidomain/multilayer optical networks. *IEEE Communications Magazine*, v. 46, n. 6, p. 78–87, jun. 2008. ISSN 0163-6804. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4539470>. Cited in page 33.

GIFRE, L. et al. ABNO-driven content distribution in the telecom cloud. *Optical Switching and Networking*, jul. 2015. ISSN 15734277. Disponível em: <http://www.sciencedirect.com/science/article/pii/S1573427715000478>. Cited in page 53.

GU, F. et al. Survivable cloud network mapping for disaster recovery support. *IEEE Transactions on Computers*, v. 64, n. 8, p. 2353–2366, Aug 2015. ISSN 0018-9340. Disponível em: <http://dx.doi.org/10.1109/TC.2014.2360542>. Cited in page 63.

HABIB, M. F. et al. A disaster-resilient multi-content optical datacenter network architecture. In: *2011 13th International Conference on Transparent Optical Networks*. IEEE, 2011. p. 1–4. ISBN 978-1-4577-0881-7. ISSN 2161-2056. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5970928>. Cited 2 times in pages 21 and 36.

HABIB, M. F. et al. Design of Disaster-Resilient Optical Datacenter Networks. *Journal of Lightwave Technology*, v. 30, n. 16, p. 2563–2573, ago. 2012. ISSN 0733-8724. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6208805>. Cited 2 times in pages 21 and 36.

HABIB, M. F. et al. Disaster survivability in optical communication networks. *Computer Communications*, v. 36, n. 6, p. 630–644, mar. 2013. ISSN 01403664. Disponível em: <http://www.sciencedirect.com/science/article/pii/S0140366413000224>. Cited 3 times in pages 33, 34, and 41.

HARLE, D. et al. Service Level Agreement Framework for Differentiated Survivability in GMPLS-based IP-over-Optical Networks. In: *2007 IEEE International Conference on Communications*. IEEE, 2007. p. 2249–2256. ISBN 1-4244-0353-7. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4289049>. Cited 2 times in pages 34 and 39.

HUANG, S.; MARTEL, C.; MUKHERJEE, B. A multistate multipath provisioning scheme for combating node failures in telecom mesh networks. In: *MILCOM 2009 - 2009 IEEE Military Communications Conference*. IEEE, 2009. p. 1–5. ISBN 978-1-4244-5238-5. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5379920>. Cited in page 34.

IANNONE, E. *Telecommunication Networks*. Taylor & Francis, 2011. (Devices, Circuits, and Systems). ISBN 9781439846377. Disponível em: <https://books.google.com.br/books?id=KHLRBQAAQBAJ>. Cited in page 26.

LANDI, G. et al. A Network Control Plane architecture for on-demand co-provisioning of optical network and IT services. In: *2012 Future Network & Mobile Summit (FutureNetw)*. Berlin: [s.n.], 2012. p. 1–8. Disponível em: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=6294216>. Cited 6 times in pages 12, 20, 24, 25, 41, and 54.

LIN, R. et al. Design of light-tree based optical inter-datacenter networks. *IEEE/OSA Journal of Optical Communications and Networking*, v. 5, n. 12, p. 1443–1455, Dec 2013. ISSN 1943-0620. Disponível em: <http://dx.doi.org/10.1364/JOCN.5.001443>. Cited 3 times in pages 22, 23, and 38.

MANDAL, U. et al. Greening the cloud using renewable-energy-aware service migration. *IEEE Network*, v. 27, n. 6, p. 36–43, nov. 2013. ISSN 0890-8044. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6678925>. Cited 3 times in pages 19, 22, and 37.

METZ, C. Ip anycast point-to-(any) point communication. *IEEE Internet Computing*, v. 6, n. 2, p. 94–98, Mar 2002. ISSN 1089-7801. Disponível em: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=991450>. Cited in page 20.

MUKHERJEE, B.; HABIB, M.; DIKBIYIK, F. Network adaptability from disaster disruptions and cascading failures. *IEEE Communications Magazine*, v. 52, n. 5, p. 230–238, maio 2014. ISSN 0163-6804. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6815917>. Cited in page 34.

MUKHERJEE, B.; HABIB, M.; DIKBIYIK, F. Network adaptability from disaster disruptions and cascading failures. *IEEE Communications Magazine*, v. 52, n. 5, p. 230–238, May 2014. ISSN 0163-6804. Disponível em: <http://dx.doi.org/10.1109/MCOM.2014.6815917>. Cited in page 62.

NAPHADE, M. et al. Smarter Cities and Their Innovation Challenges. *Computer*, v. 44, n. 6, p. 32–39, jun. 2011. ISSN 0018-9162. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5875937>. Cited in page 29.

OPTIMIZATION, I. G. *Gurobi Optimizer Reference Manual*. 2016. Disponível em: <http://www.gurobi.com>. Cited in page 55.

PENG, S. et al. A novel SDN enabled hybrid optical packet/circuit switched data centre network: The LIGHTNESS approach. In: *2014 European Conference on Networks and Communications (EuCNC)*. IEEE, 2014. p. 1–5. ISBN 978-1-4799-5280-9. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6882622>. Cited in page 23.

PEREIRA, D.; Camillo Penna, M. A new algorithm for dimensioning resilient optical networks for shared-mesh protection against multiple link failures. *Optical Switching and Networking*, v. 13, p. 158–172, jul. 2014. ISSN 15734277. Disponível em: <http://www.sciencedirect.com/science/article/pii/S1573427714000320>. Cited in page 33.

PERELLO, J.; PAVON-MARINO, P.; SPADARO, S. Cost-Efficient Virtual Optical Network Embedding for Manageable Inter-Data-Center Connectivity. *ETRI Journal*, Electronics and Telecommunications Research Institute, v. 35, n. 1, p. 142–145, fev. 2013. ISSN 1225-6463. Disponível em: <http://etrij.etri.re.kr/etrij/journal/article/article.do?volume=35&issue=1&page=142>. Cited 2 times in pages 23 and 38.

SARBAZI-AZAD, H.; ZOMAYA, A. Routing and wavelength assignment in optical networks. In: SARBAZI-AZAD, H.; ZOMAYA, A. (Ed.). *Large Scale Network-Centric Distributed Systems*. Wiley-IEEE Press, 2014. ISBN 9781118640708. Disponível em: <http://ieeexplore.ieee.org/servlet/opac?bknumber=6670810>. Cited in page 38.

SAVAS, S. S. et al. Network adaptability to disaster disruptions by exploiting degraded-service tolerance. *IEEE Communications Magazine*, v. 52, n. 12, p. 58–65, December 2014. ISSN 0163-6804. Disponível em: <http://dx.doi.org/10.1109/MCOM.2014.6979953>. Cited 2 times in pages 41 and 63.

SILVA, A. P. da et al. Congestion control in disruption-tolerant networks. In: *Proceedings of the 9th ACM MobiCom workshop on Challenged networks - CHANTS '14*. New York, New York, USA: ACM Press, 2014. p. 65–68. ISBN 9781450330718. Disponível em: <http://dl.acm.org/citation.cfm?id=2645672.2645684>. Cited in page 37.

SONG, L.; ZHANG, J.; MUKHERJEE, B. Dynamic provisioning with availability guarantee for differentiated services in survivable mesh networks. *IEEE Journal on Selected Areas in Communications*, v. 25, n. 3, p. 35–43, abr. 2007. ISSN 0733-8716. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4202537>. Cited in page 35.

VOORSLUYS, W.; BROBERG, J.; BUYYA, R. Introduction to Cloud Computing. In: BUYYA, R.; BROBERG, J. T. U. o. M.; GOSCINSKI, A. (Ed.). *CLOUD COMPUTING: Principles and Paradigms*. [S.l.]: John Wiley & Sons, Inc., 2011. cap. 1, p. 40. ISBN 978-0-470-88799-8. Cited 6 times in pages 11, 12, 15, 16, 17, and 21.

WALRAVENS, N.; BALLON, P. Platform business models for smart cities: from control and value to governance and public value. *IEEE Communications Magazine*, v. 51, n. 6, p. 72–79, jun. 2013. ISSN 0163-6804. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6525598>. Cited in page 29.

WANG, X. et al. Handover reduction via mobility-prediction-based vpon formation in optical-access-enabled cloud-ran. In: *Optical Fiber Communications Conference and Exhibition (OFC), 2015*. [s.n.], 2015. p. 1–3. Disponível em: <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=7121809>. Cited in page 63.

WIATR, P.; MONTI, P.; WOSINSKA, L. Power savings versus network performance in dynamically provisioned wdm networks. *IEEE Communications Magazine*, v. 50, n. 5, p. 48–55, May 2012. ISSN 0163-6804. Disponível em: <http://dx.doi.org/10.1109/MCOM.2012.6194382>. Cited in page 28.

WOOD, T. et al. CloudNet: Dynamic Pooling of Cloud Resources by Live WAN Migration of Virtual Machines. *IEEE/ACM Transactions on Networking*, PP, n. 99, p. 1–1, 2014. ISSN 1063-6692. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6880406>. Cited 3 times in pages 22, 38, and 53.

XIA, M. et al. Risk-Aware Provisioning for Optical WDM Mesh Networks. *IEEE/ACM Transactions on Networking*, v. 19, n. 3, p. 921–931, jun. 2011. ISSN 1063-6692. Disponível em: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5671517>. Cited 2 times in pages 33 and 34.

YI, P.; DING, H.; RAMAMURTHY, B. Budget-optimized network-aware joint resource allocation in grids/clouds over optical networks. *Journal of Lightwave Technology*, PP, n. 99, p. 1–1, 2016. ISSN 0733-8724. Disponível em: <http://dx.doi.org/10.1109/JLT.2016.2521722>. Cited in page 36.

ZHANI, M. F.; BOUTABA, R.; HELLERSTEIN, J. L. Dynamic heterogeneity-aware resource provisioning in the cloud. *IEEE Transactions on Cloud Computing*, v. 2, n. 1, p. 14–28, jan. 2014. ISSN 2168-7161. Cited in page 36.