Restoration in Optical Cloud Networks with Relocation and Services Differentiation

Carlos Natalino, Lena Wosinska, Salvatore Spadaro, João C. W. A. Costa, Carlos R. L. Francês and Paolo Monti

Abstract—Optical cloud networks allow for the integrated management of both optical and IT resources. In this paradigm, cloud services can be provisioned in an *anycast* fashion, i.e., only the source node asking for a service is specified, while it is up to the cloud control/management system to select the most suitable destination data center (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 the 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. This paper 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. The paper proposes both an ILP formulation which guarantees optimal results and a heuristic, which trades the optimality of the solution achieved by the ILP for faster processing times. Simulation results show that the average service availability and restorability performance obtained by both the ILP and the heuristic are very close to the one achievable using a protection-based strategy, but with the inherent benefit, in terms of efficient use of resources, offered by a restoration-based approach.

Index Terms—Optical cloud; wavelength division multiplexing (WDM); failure recovery; service relocation; service differentiation; network survivability; service availability.

I. INTRODUCTION

WITH an ever-growing demand for cloud-based services, the *optical cloud* concept is becoming increasingly popular [1]. The term refers to a paradigm in which client nodes require access to *Information Technology (IT) resources* (i.e., mostly storage and computation services) offered at selected data center (DC) nodes located in different geographical areas. In the optical cloud paradigm, DC nodes are interconnected and made accessible to client nodes via high-speed optical wavelength division multiplexing (WDM) fiber links, referred to as *optical transport resources*.

Differently from conventional network connectivity services, which require the provisioning of transport resources only, the establishment of a cloud service is the result of a

Carlos Natalino da Silva, João C. W. A. Costa and Carlos Renato L. Francês are with the Post-Graduation Program in Electrical Engineering, Institute of Technology of Federal University of Pará, Pará, Brazil (corresponding authors to provide e-mail: cns, jweyl, rfrances@ufpa.br).

Paolo Monti and Lena Wosinska are with the KTH Royal Institute of Technology, School of Information and Communication Technology, Electrum 229, 164 40 Kista, Sweden (e-mail: pmonti, wosinska@kth.se).

Salvatore Spadaro is with Polytechnic University of Catalonia, Department of Signal Theory and Communications (TSC).

joint optimization process that consider the availability of both transport and IT resources [1], [2]. For example, the network architecture proposed in [2], offers the possibility to manage jointly both IT and transport resources in an optical cloud network, and allows for the integrated coordination of all aspects related to cloud service provisioning, management, and tear down.

Cloud services are also anycast in nature. More specifically, from the client point of view, the location of the IT resources assigned to a cloud service is not important as long as the specific quality of service (QoS) requirements of the cloud service are met (e.g., in terms of bandwidth, latency, and availability). As a result, a client node can be connected to potentially any DC node equipped with sufficient storage and computing resources. It is then up to the cloud control/management infrastructure to select the DC node that is the most suitable to serve a specific cloud service.

A crucial aspect to consider while provisioning optical cloud services is resiliency. In fact, at the occurrence of a failure, the cloud control/management infrastructure needs to be able to promptly react to guarantee the required availability levels for the cloud services already provisioned. One possibility to guarantee survivability in optical cloud networks is to use proactive *protection-based techniques*. These strategies pre-reserve backup resources during the provisioning phase so that they can be made promptly available to recover a cloud service disrupted by a failure. As a result, a cloud service is guaranteed 100% survivability against the failure scenario(s) considered by the protection strategy [3]. This very good resiliency performance comes at the expense of large resource consumption since backup resources are pre-reserved and cannot be used to provision other cloud services. The design of survivable optical cloud networks using protection-based techniques translates in additional transport and IT resources to be deployed [4], while, in a dynamic provisioning scenario, pre-reserving backup resource negatively affects the blocking probability for service requests [5].

Another option to provide resiliency for cloud services is based on *restoration-based techniques* [3]. These strategies are reactive and they do not rely on pre-reserved backup resources to recover the cloud services disrupted by a failure. Restoration strategies are more flexible than protection techniques in adapting to failure scenarios, and they allow for lower deployment costs and/or better network blocking probability performance. On the other hand restoration cannot guarantee 100% survivability against failures (i.e., backup resources might not be available when they are needed to recover one or more cloud services), an aspect that might be crucial for some critical cloud service applications.

The restorability performance (i.e., defined as the portion of disrupted cloud services that can be successfully restored

[6]) of restoration-based techniques can be improved by leveraging on the anycast nature of cloud services. If the original DC node where the cloud service is running is not reachable anymore after a failure (i.e., not enough spare transport resources can be found to provision a backup lightpath), the cloud service can be migrated to a different DC. This operation is known as *cloud service relocation* and it has already shown good potential in improving the restorability performance of restoration-based techniques in optical cloud networks [5], [6]. On the other hand, the migration process is not instantaneous and need to be used carefully to not significantly affect the cloud service downtime [7]. In addition, despite the benefits of using cloud service migration, a restoration strategy may still not be able to recover all the cloud services disrupted by a failure because of a lack of transport resources. In such a case, a selection of cloud services that should be given priority when using the existing spare wavelengths needs to be made. One possibility is to differentiate the cloud services based on their survivability requirements [8]. For example, critical applications (e.g., transactional systems, security applications, infrastructure as a service) may require service availability levels close to 99.999%, while less critical cloud services (e.g., multimedia offline processing, preemptible virtual machines), may be provisioned with less stringent availability figures (i.e., between four or three 9s). Normally these requirements are specified in the Service Level Agreement (SLA) [9], [10].

This paper aims at investigating the benefits of applying both the service relocation and the service differentiation concepts while restoring optical cloud services. The work considers a dynamic provisioning scenario where cloud services are divided in service classes (i.e., based on their availability requirements). Single fiber link failure scenarios are considered. Two approaches are presented. The first one is based on an Integer Linear Programming (ILP) formulation, referred to as *IRP (ILP for Relocation with Priorities)*. The IRP strategy targets the minimization of both the average service downtime (defined as the portion of the holding time during which a cloud services is not available) and the number of cloud services that are relocated. The second approach, i.e., *HRP (Heuristic for Relocation with Priorities)*, is developed to provide performance results very close to IRP, but with a significantly lower processing time.

The results from the performance assessment study show that the proposed restoration strategy is 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 benefit in terms of efficient resources usage.

II. RELATED WORK

The service relocation concept has been addressed in a number of recent research works. The authors in [11], [12], [13], and [14], base their works on an intuition similar to the one explored in this paper. The main difference is that they investigate the benefits of service relocation while designing an optical cloud network that uses protection-based techniques. These works show that service relocation

helps in reducing the planned number of IT and transport resources needed to accommodate a given number of cloud services.

When looking at scenarios where optical cloud networks are in operation, there are very few studies that address the impact of service relocation on the performance of a resilience strategy in the presence of dynamic traffic. The authors in [6], and [5] propose to apply the service relocation concept together with restoration-based techniques. They demonstrate that service relocation is indeed beneficial in improving the restoration performance. In addition it was also found that these benefits could be achieved requiring only a relatively small fraction of the restored services to be relocated to a different data center. However, these studies do not take into account and exploit the fact that different cloud services may have different survivability requirements.

Adapting the performance of a resilience strategy to the specific service availability requirements is not a new concept. Service differentiation has been already explored in the past for conventional connectivity services in order to improve the resources usage efficiency by minimizing backup resource overprovisioning [15], [16]. Regarding the provisioning of cloud services, the work in [17] proposes an 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 application of similar techniques could be helpful while restoring cloud services, more specifically when resource contentions take place among different cloud services with different availability requirements. Nevertheless, to the best of our knowledge, the benefits of cloud service restoration with service differentiation have not been studied and quantified yet.

III. THE PATH RESTORATION WITH SERVICE RELOCATION AND DIFFERENTIATION (PR-SRD) 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 IRP (ILP for Relocation with Priorities), and the second one on a heuristic called HRP (Heuristic for Relocation with Priorities).

The paper considers a dynamic provisioning scenario where cloud services require continuously bandwidth between a client and a data center (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 extended to consider 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 high priority value are given precedence over the ones with a low priority. If a cloud service cannot be successfully restored upon the occurrence of a failure, the cloud service is dropped. In the scenario just described, the objective of the PR-SRD problem is to minimize the average downtime of all cloud
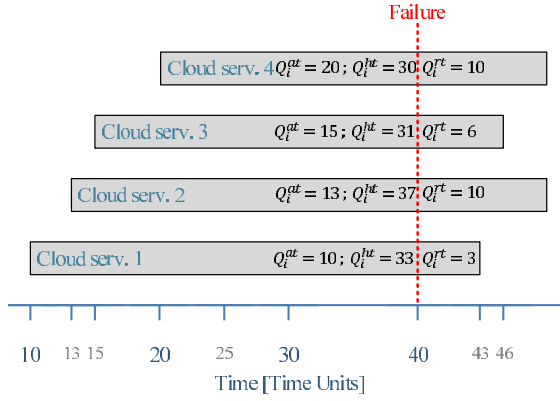
Fig. 1. Example of how the value of the remaining (service) time (Q_i^{rt}) can be calculated for a set of four cloud services disrupted by a fiber link failure happening at $CT = 40$ time units.

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 data center (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. This prioritization strategy is used to mimic a scenario where cloud services should be treated differently depending on their importance (e.g., platinum, gold, silver, and bronze services). 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}):

$$Q_i^{rt} = Q_i^{ht} - (CT - Q_i^{at}), \tag{1}$$

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 1 shows an example of how the value Q_i^{rt} can be calculated using (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. 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 of the PR-

SRD problem. More specifically it is used to decide which cloud services should be given precedence when restoration resources are assigned. For example, looking at Fig. 1 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. It is because 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.

A. The ILP for Relocation with Priorities (IRP) Model

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 wavelength capacity. As already mentioned, 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 (1), is normalized as follows:

$$t_i = \left\lceil 100 \times \frac{Q_i^{rt}}{RT} \right\rceil, \quad RT = max(Q_i^{rt}), \forall Q_i \in Q. \tag{2}$$

In this way the remaining service 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: the number of storage units to be relocated (Q_i^{st}), the rate at which storage units are transmitted (i.e., Δ, expressed in number of storage units per time unit), 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 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:

$$Q_{i,k}^{rd} = \frac{Q_i^{st}}{\Delta} + \frac{d_{Q_i^{dst}, k}}{\Theta}, \forall Q_i \in Q, k \in N_{DC}. \tag{3}$$

When a cloud service is relocated the value of t_i defined in (2) is modified as follows, in order to take into account the impact of the relocation downtime:

$$t_{i,k} = \begin{cases} \left\lceil 100 \times \frac{(Q_i^{rt} - Q_{i,k}^{rd})}{RT} \right\rceil & , \text{if } Q_i^{rt} - Q_{i,k}^{rd} > 0, \\ 0 & , \text{otherwise} \end{cases}$$
$$RT = max(Q_i^{rt}), \forall Q_i \in Q, k \in N_{DC}. \tag{4}$$

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 (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);
- 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:

$$
\begin{aligned}
min \sum_{Q_i \in Q} & \left[\alpha_i \left(t_i - \sum_{k \in N_{DC}} t_{i,k} \times A_{i,k} \right) \right] \\
& + \beta \sum_{Q_i \in Q} \sum_{k \in N_{DC} | k \neq Q_i^{dst}} A_{i,k} + \gamma \sum_{(x,y)} wl_{x,y}
\end{aligned} \tag{5}
$$

Subject to:

$$
\sum_{\forall n \in N} wl_{nj}^i - \sum_{\forall m \in N} wl_{jm}^i = \begin{cases} -A_i & , \text{if } j = Q_i^{src} \\ A_{i,j} & , \text{if } j \in N_{DC}, \\ 0 & , \text{otherwise} \end{cases}
$$
$$
\forall Q_i \in Q, \forall j \in N \tag{6}
$$

$$
wl_{xy} = \sum_{\forall Q_i} wl_{xy}^i, \forall (x,y) \in E \tag{7}
$$

$$
wl_{xy} \leq W_{xy}, \forall (x,y) \in E \tag{8}
$$

$$
\sum_{\forall k} A_{i,k} \leq 1, \forall Q_i \in Q \tag{9}
$$

$$
A_{i,k} = 0, \forall k \in N_{DC}, \forall Q_i \in Q | k \neq Q_i^{dst} \wedge Q_i^{rt} \leq Q_i^{rd} \tag{10}
$$

$$
\sum_{\forall Q_i \in Q} \left(Q_i^{st} \times A_{i,k} \right) \leq DC_k^{st}, \forall k \in N_{DC} \tag{11}
$$

$$
\sum_{\forall Q_i \in Q} \left(Q_i^{pu} \times A_{i,k} \right) \leq DC_k^{pu}, \forall k \in N_{DC} \tag{12}
$$

$$
t_{i,k} = \begin{cases} t_i & , \text{if } k = Q_i^{dst} \\ \left\lceil 100 \times \frac{(Q_i^{rt} - Q_{i,k}^{rd})}{RT} \right\rceil & , \text{if } Q_i^{rt} - Q_{i,k}^{rd} > 0, \\ 0 & , \text{otherwise} \end{cases}
$$
$$
RT = max(Q_i^{rt}), \forall Q_i \in Q, k \in N_{DC}. \tag{13}
$$

Equation (5) describes the objective function consisting of three terms. The first one is the sum of the downtime of all the cloud services (each one weighted by its priority value) that cannot be restored. The second term counts the number of cloud services that needed relocation while being restored. The third term accounts for the number of wavelength links used by all the restoration paths. The role of the β and γ parameters is to make sure that the results obtained after solving the ILP formulation represent a good tradeoff 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 high number of relocations, if relocating a cloud service helps in saving wavelength 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 (6) guarantees the flow conservation of each restored cloud service. Constraint (7) computes the total number of wavelength used on each fiber link for restoration purposes. Constraint (8) ensures that the number of wavelengths used on each fiber link does not exceed the actual number of available wavelengths. Constraint (9) checks that each relocated cloud service uses at most one DC. Constraint (10) ensures that a cloud service cannot be relocated to DC node k if the relocation downtime to that DC node is larger than the remaining service time. Constraints (11) and (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 (13) computes the normalized value of the remaining service time of $Q_i \in Q$.

B. Heuristic for Relocation with Priorities (HRP)

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 approach. 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 wavelength links 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 a high priority class 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 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 wavelength resources in $G(N,E)$ from Q_i^{src} to the DC node already in use (i.e., Q_i^{dst}). This is done to reduce the number of unnecessary cloud service relocations. If such a path exists, then a new lightpath 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 availability of storage and compute resources the heuristic makes also sure that the value of the relocation downtime, i.e., $Q_{i,k}^{rd}$, for the candidate DC node under exam does not exceed the value of the remaining

Algorithm 1 HRP

1: $Q' = Q$ sorted by Q_i^{wt}
2: **for all** $Q_i \in Q'$ **do**
3:     curRoute = selRoute = selDC = NULL
4:     **if** $shortestPath(Q_i^{src}, Q_i^{dst}) \neq NULL$ **then**
5:         $selRoute = shortestPath(Q_i^{src}, Q_i^{dst})$
6:         $restorePath(Q_i, selRoute)$
7:     **else**
8:         **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**
9:             curRoute = $shortestPath(Q_i^{src}, Q_i^{dst})$
10:            **if** $hopCount($curRoute$) < hopCount($selRoute$)$ **then**
11:                selRoute = curRoute
12:            **end if**
13:         **end for**
14:         **if** selRoute $\neq$ NULL **then**
15:            $relocateAndRestorePath(Q_i, selRoute)$
16:         **else**
17:            $dropService(Q_i)$
18:         **end if**
19:     **end if**
20: **end for**

service time of Q_i. If more than one DC node with these characteristics is 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 wavelength link resources used to restore Q_i). Once this new DC and route are selected, Q_i is relocated and a lightpath 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 functions used in the heuristic are described next. The function *shortestPath(N_i , N_j)* checks a pre-computed set of k-shortest-paths between N_i and N_j and returns the shortest path with available wavelength resources between these two nodes. If no such path can be found the function returns *NULL*. The function *hopCount()* returns either the number of hops of a given route, or it returns infinite (i.e., ∞) when the received argument is *NULL*. The functions *restorePath()*, *relocateAndRestorePath()*, and *dropService()* are responsible for restoring, restoring plus relocating, and dropping a given cloud service, respectively. After each one of these actions the network state, i.e., $G(N, E)$, is updated accordingly.

IV. PERFORMANCE ASSESSMENT

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

A. Simulation Setup

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 IPR 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 (Fig. 2(a)) as the reference topology [6].
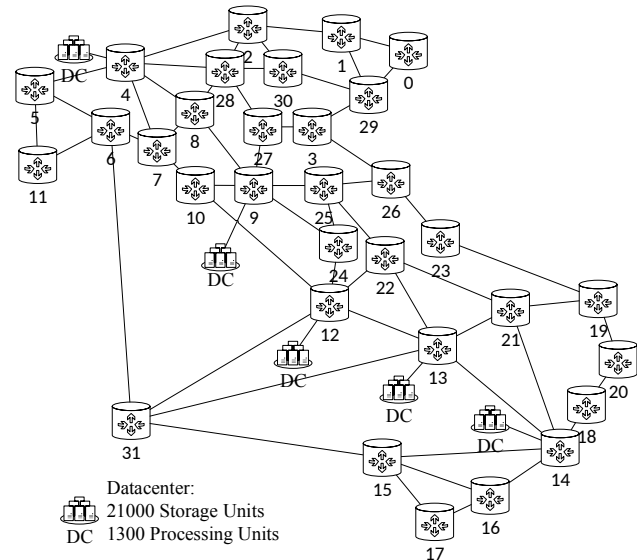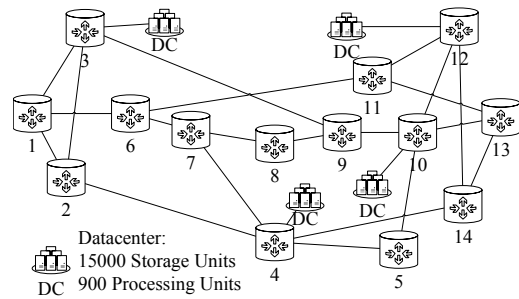

(a)

(b)

Fig. 2. Network topologies: (a) NSFNET Network Topology and; (b) Italian Network Topology.

In this network nodes 3, 4, 10, and 11 are assumed to host 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 (Fig. 2(b)) 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 wavelength resources available to reach the DC node resources) [11]. 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 DCs 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.

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
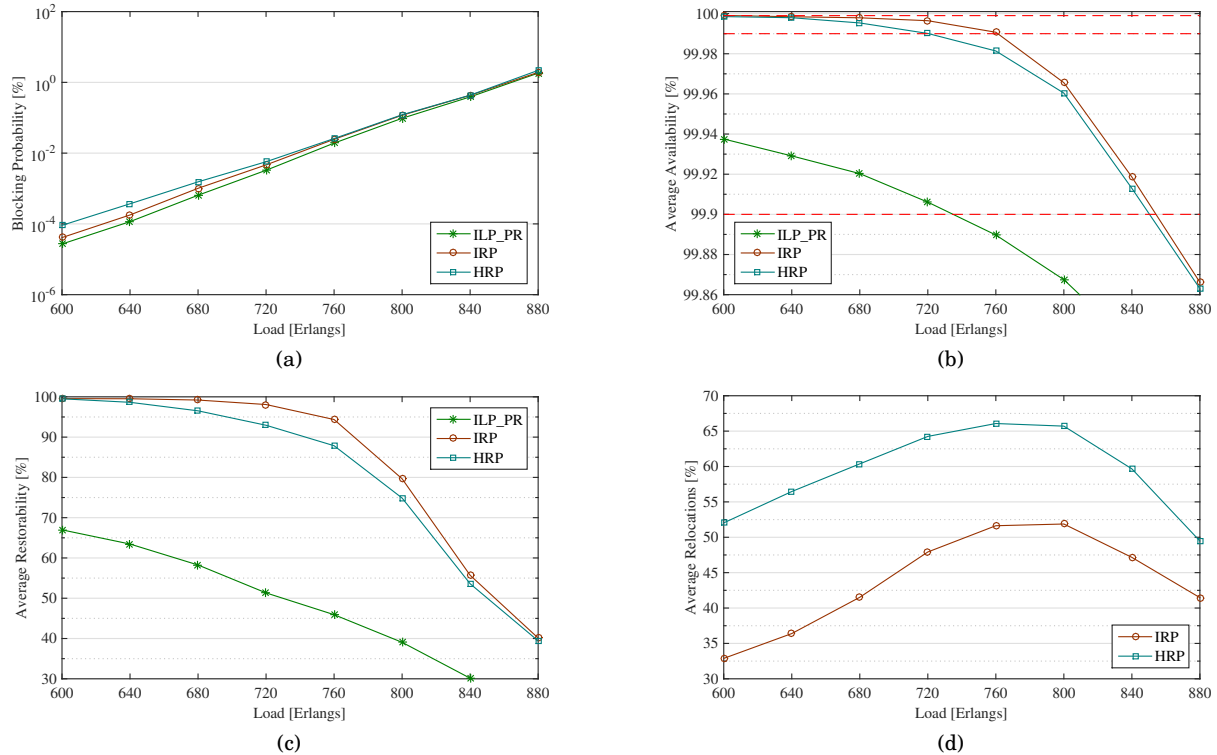
Fig. 3. Simulation results for Scenario A: (a) Blocking probability; (b) Availability; (c) Restorability and; (d) Percentage of service relocations.

other with low priority. It is assumed that up to 20% of the cloud services in each experiment have high priority, while the remaining part 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 wavelengths 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). 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 interval [1,100] and [1,5] respectively [6]. Connecting a client node to a DC node requires the establishment of a lightpath with a capacity equal the capacity of one wavelength 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 [18], [19]. The value of Θ is assumed to be equal to 2×10^5 [km/s], while the hop length in the NSF and Italian topology is set to 1086 and 224 km, respectively.

The simulation study presented in the paper considers a single fiber link failure scenario. All fiber links in the network have the same probability to be down. 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.

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:

$$Q_i^{wt} = \begin{cases} 8 \times Q_i^{rt} & , \forall Q_i \in Q \mid \alpha_i = 10^5 \\ Q_i^{rt} & , \forall Q_i \in Q \mid \alpha_i = 10^4 \end{cases} . \quad (14)$$
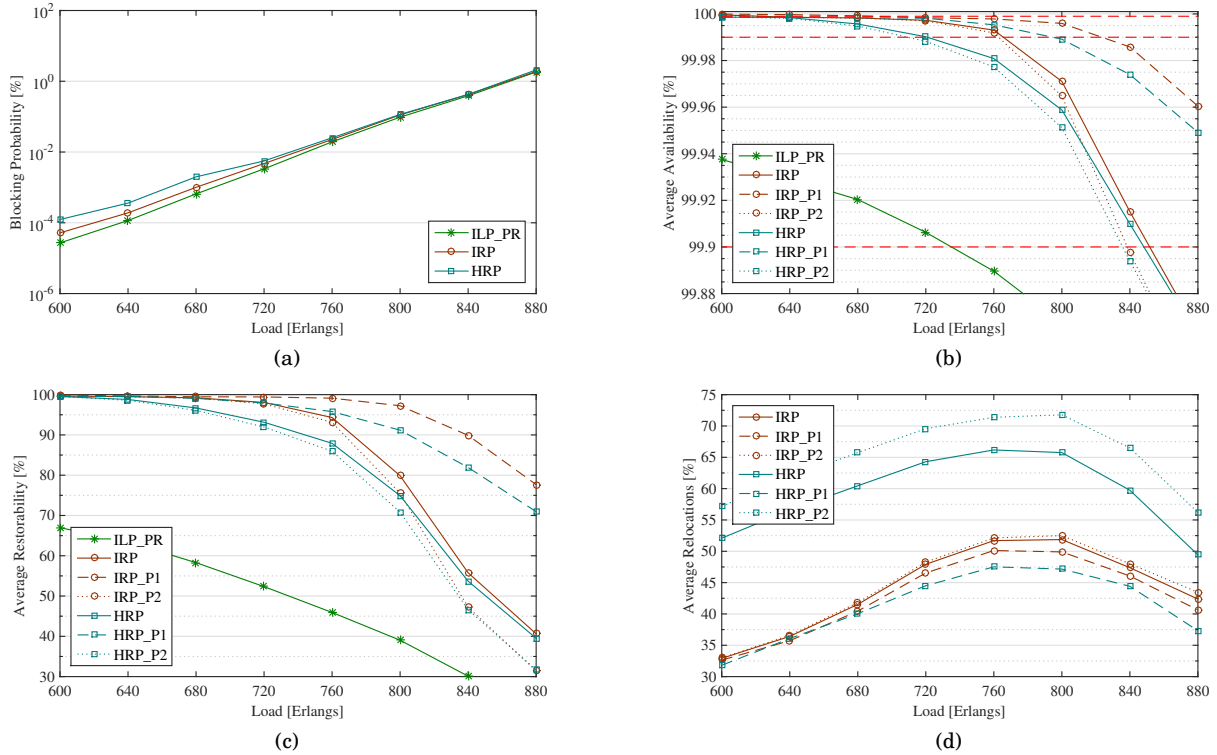
Fig. 4. Simulation results for Scenario B: (a) Blocking probability; (b) Availability; (c) Restorability and; (d) Percentage of service relocations.

Equation (14) 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 tradeoff among the performance parameters defined in (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 [2], which chooses the DC node, with enough storage and computing resources, that is the closest to the client node at which a cloud service originates. Upon the occurrence of a failure, the 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 lightpath 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 PCE-based controller [20] 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. This way, the assessment of the benefits of introducing the relocation option during the restoration process is performed. This benchmarking approach is based on the solution of an ILP formulation and it is referred to as ILP_PR [3].

All the results presented in the next section are the average of 100 different experiments. 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. The

ILP formulation is solved using the Gurobi Optimizer [21]. A Debian Linux workstation with 2 Intel Xeon CPUs (6 cores per CPU) clocked at 2.2 GHz and with 32 GB of RAM is used for the simulations.

B. 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 3 presents the performance results for Scenario A. The proposed strategy (i.e., both IRP and HRP) shows a slightly worse performance in terms of blocking probability than the benchmark ILP_PR i.e., Fig. 3(a). This is because of the better performance in terms of restorability (Fig. 3(c)) that, in turns, leads to less resources available for provisioning future traffic. Figure 3(b) 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 thresholds, respectively. The figure confirms the intuition that allowing for the relocation of cloud services to alternative DC nodes during the restoration process has a beneficial
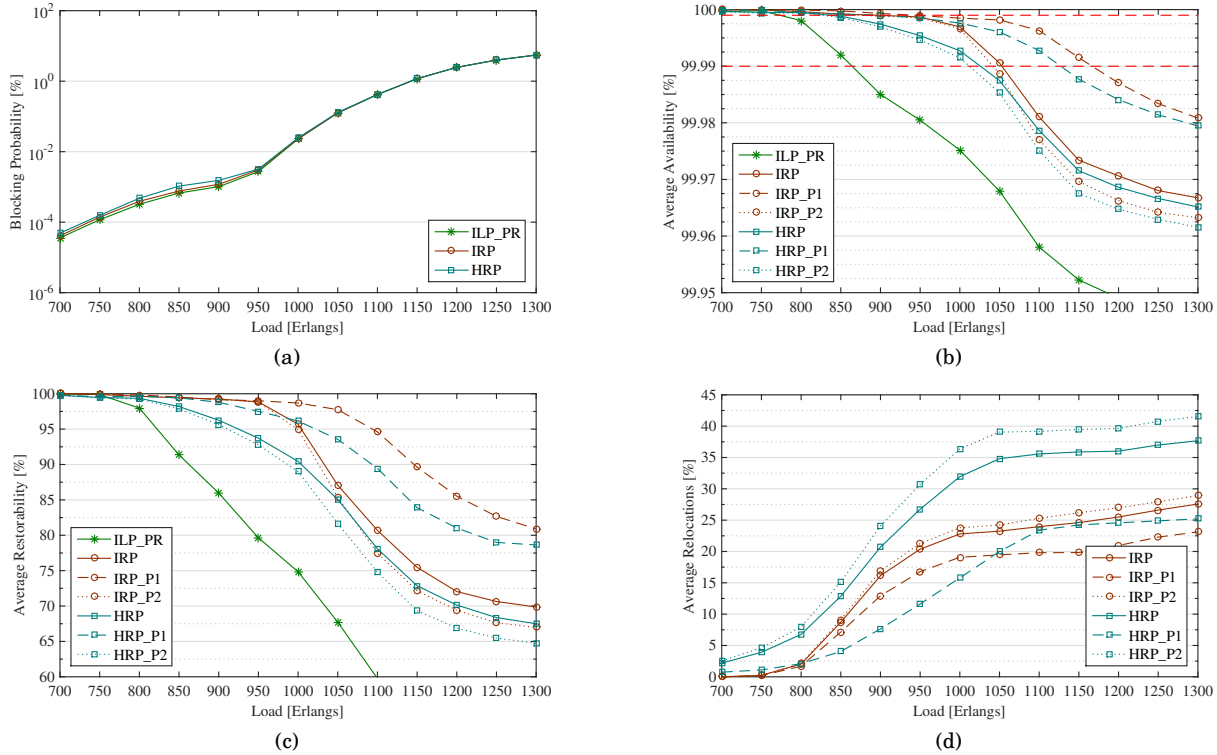
Fig. 5.   Simulation results for Scenario C: (a) Blocking probability; (b) Availability; (c) Restorability and; (d) Percentage of service relocations.

effect on the average value of the cloud service availability. More specifically both IRP and HRP show availability figures exceeding four 9s in 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 improvements in terms of average restorability values, i.e., Fig. 3(c), 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 3(d) shows the average number of relocations needed to restore the cloud services. The lower this number the better because relocations are costly in terms of virtual machine state transfer overhead. It can be noticed from the figure that in the best case (i.e., with the IRP strategy) only 52.5% of the successfully restored cloud services needed to be relocated. Finally, the performance results presented in Fig. 3 allows to estimate how close the results from the HRP are from the optimum, i.e., the results of the IRP. It can be noticed that in almost all the cases HPR behaves closely to IRP. The only difference is in the number of relocations. This is due to the sequential nature of the heuristic that restores one cloud service at a time, loosing all the benefits that a concurrent approach such as IRP has.

Figure 4 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. 4(a), are the same as in Scenario A. This is not the case

for the other metrics. Figure 4(b) presents the availability performance of the IRP and the HRP. The figure presents curves for both the average (IRP curve) and the per-class (i.e., IPR_P1 and IPR_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. 3(b), leading to the conclusion that IRP and HRP present the same benefits over the performance of ILP_PR. This means that 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 quite close to the average value (IRP curve). The HRP approach offers availability results that are very close to IRP. Figure 4(c) presents results for the restorability performance. The conclusions that can be drawn are very 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 Scenario A, i.e., Fig. 3(c), 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 too much the restorability performance of the cloud services with low priority. They 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 4(d) presents the average relocation performance. Moreover, in
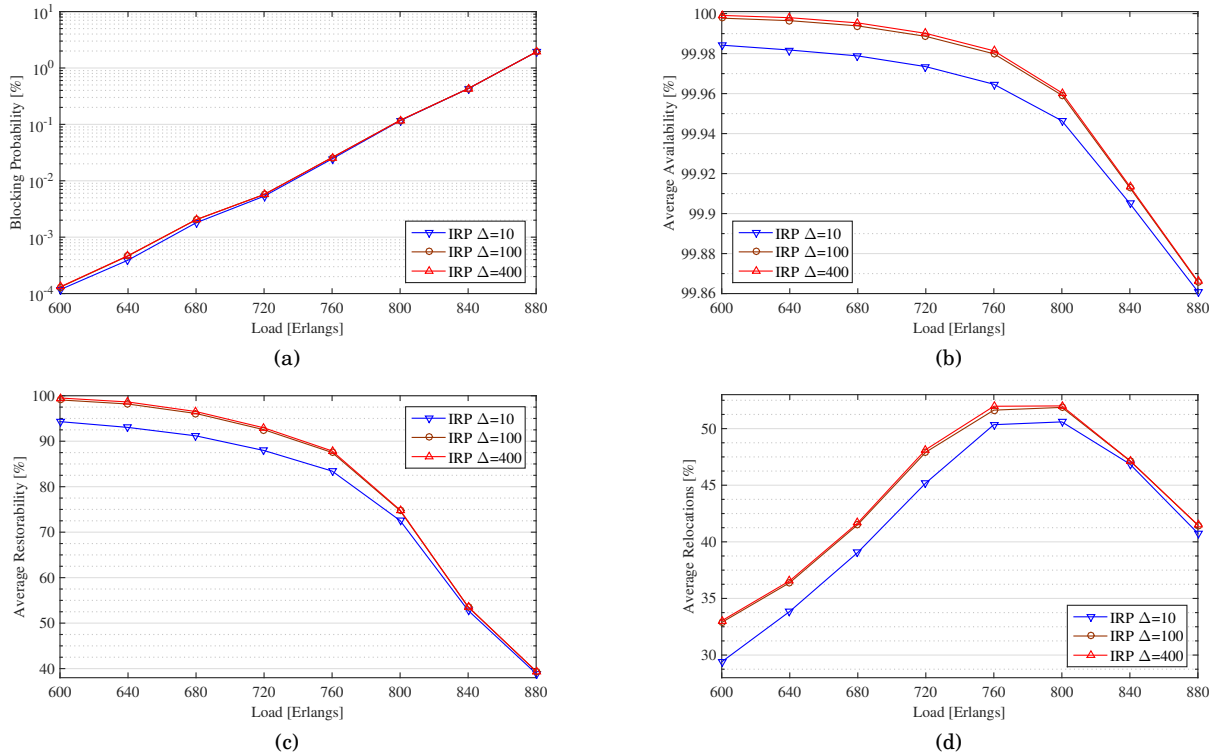
Fig. 6. Simulation results for Scenario A for different Δ values: (a) Blocking probability; (b) Availability; (c) Restorability and; (d) Percentage of service relocations.

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. 3(d). 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. 3(d), but with a number of relocations for the low priority cloud services that exceed 70% in medium to high load conditions.

Figure 5 presents the performance results for Scenario C (i.e., Italian topology with different traffic classes). The load values have been changed compared to the ones used in Scenarios A and B in order to have a fair comparison in terms of blocking probability. From the results presented in Fig. 5 it can be noticed that in general both the IPR 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. The only difference that is worth noting is about the absolute values of the various metrics. As it can be seen in Fig. 5, both IRP and HRP can achieve on average higher availability and restorability values, with a lower required number of relocations. This is motivated by the nature of the Italian topology that is on average more connected than the NSF.

Figure 6 presents the performance results of the IRP formulation in Scenario A as a function of different values of Δ. Three cases are considered. They are used to represent the 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

TABLE I
AVERAGE PROCESSING TIME FOR EACH DISRUPTED CLOUD
SERVICE RESTORATION ATTEMPT

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

the short term future ($\Delta = 400$). As it was explained in Section III 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.

Table I 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 milliseconds 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.

V. FINAL REMARKS

This paper proposes a restoration-based survivability strategy that can be used to recover cloud services disrupted by fiber link failures. The intuition behind this 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 IRP (ILP for Relocation with Priorities) and on heuristic referred to as HRP (Heuristic for Relocation with Priorities). It is shown that HPR provides results very close to IRP, but with a significantly lower processing time. The performance of both IRP and HPR were assessed in a dynamic provisioning paradigm, 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.

ACKNOWLEDGMENT

REFERENCES

[1] C. Develder, M. De Leenheer, B. Dhoedt, M. Pickavet, D. Colle, F. De Turck, and P. Demeester, "Optical networks for grid and cloud computing applications," *Proceedings of the IEEE*, vol. 100, no. 5, pp. 1149–1167, May 2012.

[2] G. Landi, N. Ciulli, J. Buysse, K. Georgakilas, M. Anastasopoulos, A. Tzanakaki, C. Develder, E. Escalona, D. Parniewicz, A. Binczewski, and B. Belter, "A network control plane architecture for on-demand co-provisioning of optical network and it services," in *2012 Future Network & Mobile Summit (FutureNetw)*, Berlin, 2012, pp. 1–8.

[3] J. Ahmed, C. Cavdar, P. Monti, and L. Wosinska, "Hybrid survivability schemes achieving high connection availability with a reduced amount of backup resources [invited]," *Journal of Optical Communications and Networking*, vol. 5, no. 10, p. A152, Sep. 2013.

[4] C. Ou, K. Zhu, B. Mukherjee, H. Zang, and L. Sahasrabuddhe, "Survivable traffic grooming in WDM mesh networks," in *Optical Fiber Communications Conference, 2003. OFC 2003*, March 2003, pp. 624–625 vol.2.

[5] C. Natalino, J. Ahmed, L. Wosinska, P. Monti, and R. Francês, "A relocation-based heuristic for restoring optical cloud services," in *13th International Conference on Optical Communication and Networks, 2014. ICOCN 2014*, November 2014.

[6] J. Ahmed, P. Monti, L. Wosinska, and S. Spadaro, "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.

[7] U. Mandal, M. Habib, S. Zhang, B. Mukherjee, and M. Tornatore, "Greening the cloud using renewable-energy-aware service migration," *IEEE Network*, vol. 27, no. 6, pp. 36–43, Nov. 2013.

[8] K. Bilal, S. U. R. Malik, S. U. Khan, and A. Y. Zomaya, "Trends and challenges in cloud datacenters," *IEEE Cloud Computing*, vol. 1, no. 1, pp. 10–20, May 2014.

[9] D. Harle, S. Albarrak, F. Ali, A. Urra, E. Calle, and J. L. Marzo, "Service level agreement framework for differentiated survivability in GMPLS-based IP-over-Optical networks," in *2007 IEEE International Conference on Communications*. IEEE, Jun. 2007, pp. 2249–2256.

[10] C. Cavdar, M. Tornatore, and F. Buzluca, "Availability-guaranteed connection provisioning with delay tolerance in optical WDM mesh networks," in *Conference on Optical Fiber Communication*, San Diego, CA, 2009, pp. 1–3.

[11] C. Develder, J. Buysse, B. Dhoedt, and B. Jaumard, "Joint dimensioning of server and network infrastructure for resilient optical grids/clouds," *IEEE/ACM Transactions on Networking*, vol. PP, no. 99, pp. 1–1, 2013.

[12] J. Buysse, M. De Leenheer, B. Dhoedt, and C. Develder, "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, pp. 1–6.

[13] C. Develder, J. Buysse, A. Shaikh, B. Jaumard, M. De Leenheer, and B. Dhoedt, "Survivable optical grid dimensioning: Anycast routing with server and network failure protection," in *2011 IEEE International Conference on Communications (ICC)*. IEEE, Jun. 2011, pp. 1–5.

[14] C. Develder, J. Buysse, M. De Leenheer, B. Jaumard, and B. Dhoedt, "Resilient network dimensioning for optical grid/clouds using relocation," in *2012 IEEE International Conference on Communications (ICC)*. IEEE, Jun. 2012, pp. 6262–6267.

[15] C. Cavdar, M. Tornatore, and B. Mukherjee, "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, Nov. 2007, pp. 1–7.

[16] P. Monti, M. Tacca, and A. Fumagalli, "Resource-efficient path-protection schemes and online selection of routes in reliable wdm networks," *J. Opt. Netw.*, vol. 3, no. 4, pp. 188–203, Apr 2004.

[17] M. F. Zhani, R. Boutaba, and J. L. Hellerstein, "Dynamic heterogeneity-aware resource provisioning in the cloud," *IEEE Transactions on Cloud Computing*, vol. 2, no. 1, pp. 14–28, Jan. 2014.

[18] L. Gifre, M. Tornatore, L. M. Contreras, B. Mukherjee, and L. Velasco, "ABNO-driven content distribution in the telecom cloud," *Optical Switching and Networking*, Jul. 2015.

[19] T. Wood, K. K. Ramakrishnan, P. Shenoy, J. Van der Merwe, J. Hwang, G. Liu, and L. Chaufournier, "CloudNet: Dynamic pooling of cloud resources by live WAN migration of virtual machines," *IEEE/ACM Transactions on Networking*, vol. PP, no. 99, pp. 1–1, 2014.

[20] J. Ahmed, C. Cavdar, P. Monti, and L. Wosinska, "A dynamic bulk provisioning framework for concurrent optimization in PCE-based WDM networks," *Lightwave Technology, Journal of*, vol. 30, no. 14, pp. 2229–2239, July 2012.

[21] I. Gurobi Optimization, "Gurobi optimizer reference manual," 2015. [Online]. Available: http://www.gurobi.com