Enhancing Restoration Performance Using Service Relocation in PCE-based Resilient Optical Clouds

J. Ahmed[1], P. Monti[1], L. Wosinska[1], S. Spadaro[2]
[1]KTH Royal Institute of Technology, Stockholm, Sweden, [2]Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
{jawwad, pmonti, wosinska}@kth.se[1], spadaro@tsc.upc.edu[2]

Abstract: This paper investigates the benefits of dynamic restoration with service relocation in resilient optical clouds. Results from the proposed optimization model show that service availability can be significantly improved by allowing a few service relocations.

OCIS codes: (060.0060) Fiber optics and optical communications; (060.4250) Networks

1. Introduction

Grid and Cloud services leveraging on high speed optical transport networks (e.g., based on optical wavelength division multiplexing (WDM) technology) are gaining a lot of popularity. This is thanks to their ability to offer storage and computing services, i.e., Information Technology (IT) resources, by utilizing and jointly managing IT resources distributed over different data centers (DCs) connected across geographically spread areas. From a client point of view the actual location of the IT resources is (usually) not that important as long as the specific requirements of a cloud service (e.g., certain bandwidth demand, delay limit, reliability performance) are met. For this reason, a client can be potentially connected to any DC with sufficient IT resources. In other words, 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 infrastructure to select the most suitable destination DC node.

In optical clouds, failure recovery mechanisms are essential in guaranteeing high availability of the provisioned cloud services. In this respect, there is an extensive literature work addressing the benefits of utilizing protection-based survivability strategies, in particular the ones using shared path protection (SPP) schemes [1]. In the presence of anycast provisioning, some of these SPP protection strategies leverage on the concept of *service relocation,* (i.e., a protection path might use a DC that is different from the one the primary path was connected to) with obvious benefits in terms of the network resources usage [2][3]. Another way to improve resource efficiency is to consider alternative resilient schemes. It is well known that path protection strategies guarantee survivability against network failure(s) but they are, on the other hand, not as efficient (i.e., in terms of resource utilization) as restoration-based schemes [4][5]. The main advantage of restoration-based strategies is the efficient use of backup resources, which are dynamically provisioned upon failure, at the expense of a certain risk that the backup network resources are not available when needed. This problem can be potentially alleviated by using dynamic restoration together with the concept of service relocations. However, there is currently very little work presented in the literature investigating this option in optical cloud networks [6].

The work in this paper aims at assessing the benefits of using service relocation in a centralized (i.e., PCE-based [7]) cloud inter-networking scenario while trying to restore a number of failed cloud services as a consequence of a network failure. An ILP model specifically designed to solve this problem is proposed. The model tries to maximize the number of restored services while at the same time minimizing the total number of relocations needed. Results show that service availability can be significantly improved by relocating only a small fraction of the services.

2. ILP Model for Path Restoration with Relocation Capabilities

This section presents an ILP model that can be used to optimally solve the path restoration problem for a number of cloud services disrupted by a failure in a scenario where service relocation is allowed.

Let $G(N, E)$ represent a physical topology with $|N|$ nodes and $|E|$ fiber links, while N_{DC} is the set of network nodes connected to a DC**.** It is assumed that each network node can be connected to at most one DC, i.e., $N_{DC} \subseteq N$. Let VM_k and ST_k represent, respectively, the number of free virtual machines (VMs), and the number of free storage units of the DC connected to node $k \in N_{DC}$. W_{xy} is the number of free wavelengths on fiber link $(x, y) \in E$. Let D be the set of cloud services affected by a failure for which a specific path restoration problem needs to be solved, then $\lambda^{c}_{s_c,d_c}$ (with $\lambda^{c}_{s_c,d_c} \in D$) represents a unique cloud service c from source s_c to destination d_c requiring q_c VMs and t_c storage units. The variables used in the formulation are: p_{xy} , i.e., the total number of wavelengths used by the restoration paths on fiber link (x, y) ; p^{c}_{xy} , equal to 1 if the restoration path of service c uses fiber link (x, y) ; A_c , equal to 1 if service c is successfully restored; and A_c^k, equal to 1 if a service c is successfully restored using the DC at node $k \in N_{DC}$.

ILP Formulation for the Path Restoration with Relocation (ILP_ PR_RELOCATION) Problem

The following ILP model is used to solve the path restoration problem for a set of cloud services affected by a failure in the network ($\forall \lambda_c \in D$). The formulation optimizes the routes and selection of the destination DCs to be used to restore a given service. Two objective functions are proposed. Objective 1 focuses on the minimization of the number of services that cannot be restored and on the overall wavelength resources used in the restoration process. Objective 2, on the other hand tries also to minimize the number of services that are relocated to a different DC while being restored. The details of the model are presented next.

Objective 1:

Minimize $\alpha \cdot (|D| - \sum_{\forall \lambda_c, \forall k} A_c^k) + \beta \cdot \sum_{\forall (x,y)} p_{xy}$

Objective 2:

Minimize $\alpha \cdot (|D| - \sum_{\forall \lambda_c, \forall k} A_c^k) + \beta \cdot \text{Relocations} + \gamma . \sum_{\forall (x,y)} p_{xy}$

Constraints:

$$\sum_{\forall n \in N} p_{nj}^c - \sum_{\forall m \in N} p_{jm}^c = \begin{cases} -A_c, & j = s_c \\ A_c^j, & j \neq s_c \end{cases}, \forall \lambda_c, \forall j \in N \quad (2.1)$$

$$p_{xy} = \sum_{\forall \lambda_c} p_{xy}^c, \ \forall (x,y) \quad (2.2)$$

$$p_{xy} \leq W_{xy}, \ \forall (x,y) \quad (2.3)$$

$$\sum_{\forall k} A_c^k \leq 1, \ \forall \lambda_c \quad (2.4)$$

$$A_c^k = 0, \ \forall k \notin N_{DC}, \forall \lambda_c \quad (2.5)$$

$$A_c^k = 0, \ \forall k \in N_{DC} \mid (q_c > VM_k) \vee (t_c > ST_k), \forall \lambda_c \quad (2.6)$$

$$\sum_{\forall k \neq d_c, \forall \lambda_c} A_c^k = \text{Relocations} \quad (2.7)$$

The first term in Objective 1 accounts for the number of cloud services that cannot be successfully restored, while the second term measures the total number of wavelength resources used by all successful restored services. Objective 2 has an extra term (i.e., β·Relocations) that accounts for the number of relocations used while restoring the disrupted cloud services. Constraint (2.1) is used for flow conservation of the path of each restored service. Constraint (2.2) calculates the total number of (wavelength) resources used for restoration on each fiber link. Constraint (2.3) ensures that, on each link, the number of (wavelength) resources used for restoration does not exceed the number of free wavelengths. Constraint (2.4) ensures that each restored service is using at most one DC. Constraint (2.5) makes sure that a node without a data center cannot be selected as the destination of a restored service. Constraint (2.6) ensures that a DC cannot be selected to be the destination of a restored service if it has not sufficient IT resources (i.e., VMs and storage units). Constraint (2.7) is used to calculate the total number of relocations.

3. Simulation Setup and Numerical Results

This section presents performance assessment of the proposed ILP-based restoration strategy applied in a dynamic anycast provisioning scenario. More specifically, the intent is to evaluate the benefits of using a restoration strategy that has relocation capabilities compared to the one that doesn't have them. The NSF network [8] (Fig. 1(a)) is used as the topology of reference where nodes 3, 4, 10, and 11 are assumed to be DC nodes, picked because of their high connectivity. All fiber links in the network are assumed to be bidirectional, with one fiber in each direction, and with 16 wavelengths per fiber. The fiber links connecting DCs to their respective network nodes are, on the other hand, assumed to have unlimited bandwidth resources. The presented results are averaged over 16 different experiments with one million cloud service requests generated in each experiment. Only anycast services are considered during the provisioning phase, i.e., each service requires an all optical lightpath between a non-DC (client) node to a DC node with enough IT resources to satisfy the cloud service requirements. Each lightpath is assumed to require the entire capacity of a wavelength channel. In all the experiments, services are provisioned using a modified version of the weighted least congested routing (WLCR) heuristic [8], i.e., adapted to solve the anycast routing and wavelength assignment problem. Path provisioning operations are assumed to be coordinated by a PCE-based controller [7] specifically designed for concurrent optimization. Each cloud service holding time is assumed to be exponentially distributed with an average equal to 60 time-units while service request arrivals at the client nodes follow a Poisson process. The ILP model (ILP_PR) presented in [5] is used for benchmarking purposes, i.e., it represents an optimal restoration strategy without relocation capabilities. All ILP model instances are solved dynamically using the Java API provided by *Gurobi Optimizer 5.50* [9] and execution times are measured to be in few tens of milliseconds.

The assessment study presented in this work considers only fiber link failures. The time between the occurrences of two consecutive failures is assumed to be exponentially distributed, with a mean time to failure (*MTTF*) equal to

1000 time units. The link reparation time is uniformly distributed, with a mean time to repair (*MTTR*) equal to 10 time units. Link failures are uniformly distributed in the network. For the sake of simplicity, it is assumed that all DCs have the same capacity in terms of IT resource (i.e., 3,000 storage units and 150 VMs). The amount of storage units and virtual machines required by one cloud service are uniformly chosen in the (1-100) and in the (1-5) interval, respectively. The values for α, β and γ used in Objective 1 and 2 are: 10^5, 10^3 and 1, respectively.

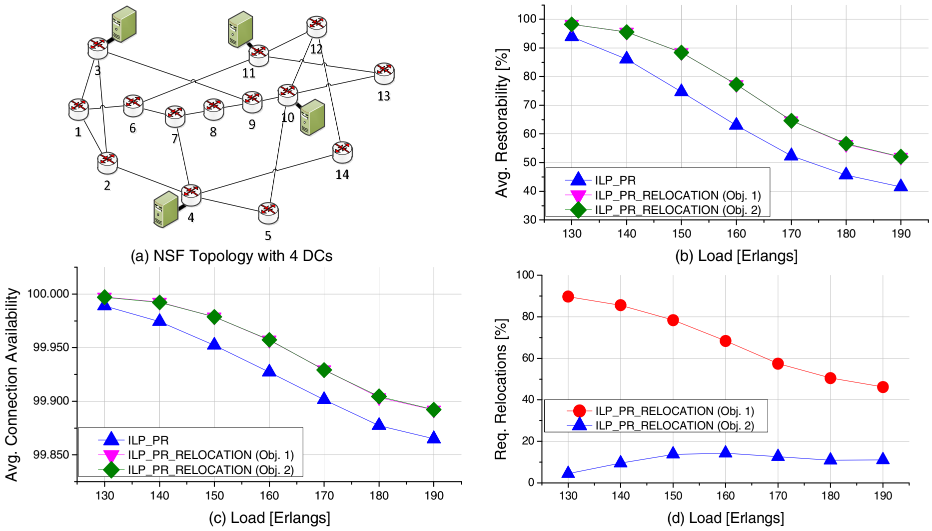

Fig.1. (a) NSF network topology (b) Avg. Restorability [%], (c) Avg. Connection Availability, (d) Required Relocations [%].

Fig. 1(b) shows the impact of service relocation on the average number of overall successful restorations over the total number of restoration attempts (i.e., the restorability). As expected, up to 20% higher restorability can be achieved when relocation capabilities are enabled in the network. This is because service relocation allows for a better use of both network and IT resources, i.e., bottlenecks are avoided as much as possible. A similar effect can also be appreciated in terms of average connection availability (i.e., ratio between the time a service is operative over the entire service holding time). More specifically, ILP_PR_RELOCATION (Fig. 1(c)) provides the highest availability performance where availability drops below "three 9s" only at the highest network load. Fig. 1(d), on the other hand, shows that the total number of service relocations required during restoration can be substantially reduced without losing anything in terms of restorability and connection availability performance (Fig. 1(b) and Fig. 1(c), respectively). In fact, with ILP_PR_RELOCATION (Obj. 2) at most 15% of the successfully restored services require relocation, thus the overhead typically associated with these operations (e.g., state transfer, and VM migration) [10] can be avoided.

4. Conclusion

This paper proposes to utilize the relocation concept while restoring services in PCE-based optical cloud networks. To assess the benefits of the proposed approach we formulated an ILP model providing an optimal solution for the dynamic path restoration problem with service relocation. It can be concluded that service relocation is indeed very beneficial in terms of both average restorability and average connection availability. In addition it was also found that these benefits can be achieved requiring only a relatively small fraction of the restored services to be relocated to a different data center.

References

[1] J. Buysse, et al., "Exploiting relocation to reduce network dimensions of resilient optical grids", DRCN 2009.
[2] J. Buysse, et al., "On the impact of relocation on network dimensions in resilient optical grids", ONDM 2010.
[3] B. Jaumard, et al., "Column Generation for Dimensioning Resilient Optical Grid Networks with Relocation", GLOBECOM 2010.
[4] S. Ramamurthy, et al., "Survivable WDM Mesh Networks", IEEE/OSA JLT. vol. 21, no. 4, pp. 870-883, April 2003.
[5] J. Ahmed, et al., "Hybrid Survivability Schemes Achieving High Connection Availability with Reduced Amount of Backup Resources", IEEE/OSA JOCN, Sept. 2013.
[6] Y. Chen, et al., "Strategies for fault tolerance in optical grid networks", IEEE ICNC 2012.
[7] J. Ahmed, et al., "A Dynamic Bulk Provisioning Framework for Concurrent Optimization in PCE-based WDM Networks", IEEE/OSA JLT, vol. 30, issue 14, pp. 2229-2239, 2012.
[8] X.-W. Chu, et al., "A dynamic RWA algorithm in a wavelength-routed all optical network with wavelength converters", INFOCOM 2003.
[9] Gurobi Optimizer, Online URL: http://www.gurobi.com/.
[10] T. Wood, et al., "CloudNet: Dynamic Pooling of Cloud Resources by Live WAN Migration of Virtual Machines", ACM SIGPLAN Notices, 2011.