

Disaster-survivable cloud-network mapping

Carlos Colman-Meixner · Ferhat Dikbiyik ·
M. Farhan Habib · Massimo Tornatore ·
Chen-Nee Chuah · Biswanath Mukherjee

Received: 21 February 2014 / Accepted: 3 April 2014 / Published online: 29 April 2014
© Springer Science+Business Media New York 2014

Abstract Cloud-computing services are provided to consumers through a network of servers and network equipment. Cloud-network (CN) providers virtualize resources [e.g., virtual machine (VM) and virtual network (VN)] for efficient and secure resource allocation. Disasters are one of the worst threats for CNs as they can cause massive disruptions and CN disconnection. A disaster may also induce post-disaster correlated, cascading failures which can disconnect more CNs. Survivable virtual-network embedding (SVNE) approaches have been studied to protect VNs against single physical-link/-node and dual physical-link failures in communication infrastructure, but massive disruptions due to a disaster and their consequences can make SVNE approaches insufficient

to guarantee cloud-computing survivability. In this work, we study the problem of survivable CN mapping from disaster. We consider risk assessment, VM backup location, and post-disaster survivability to reduce the risk of failure and probability of CN disconnection and the penalty paid by operators due to loss of capacity. We formulate the proposed approach as an integer linear program and study two scenarios: a natural disaster, e.g., earthquake and a human-made disaster, e.g., weapons-of-mass-destruction attack. Our illustrative examples show that our approach reduces the risk of CN disconnection and penalty up to 90 % compared with a baseline CN mapping approach and increases the CN survivability up to 100 % in both scenarios.

This work has been supported by Defense Threat Reduction Agency (DTRA) Grant No. HDTRA1-10-1-0011. A preliminary version of this work was presented in [1].

Keywords Cloud computing · Disaster survivability · Cloud-network mapping · Virtual-network mapping · Virtual machine

C. Colman-Meixner (✉) · M. F. Habib · C.-N. Chuah ·
B. Mukherjee · M. Tornatore
University of California, Davis, CA, USA
e-mail: cecolmanmeixner@ucdavis.edu

M. F. Habib
e-mail: mfhabib@ucdavis.edu

C.-N. Chuah
e-mail: chuah@ucdavis.edu

B. Mukherjee
e-mail: bmukherjee@ucdavis.edu

M. Tornatore
e-mail: mtornatore@ucdavis.edu

M. Tornatore
Politecnico di Milano, Milan, Italy
e-mail: tornator@elet.polimi.it

F. Dikbiyik
Sakarya University, Sakarya, Turkey
e-mail: fdikbiyik@sakarya.edu.tr

1 Introduction

Reliable provisioning of cloud-computing services depends on robust resource allocation over a common physical infrastructure, formed by datacenters and communication networks [2–4]. Physical infrastructure is often abstracted as “infrastructure as a service (IaaS)” layer which provides computational and communication resources to the upper service layers (e.g., platform as a service (PaaS) and software as a service (SaaS)) of the cloud-computing framework [5,6]. Cloud-network (CN) mapping is the combination of virtual-network (VN) mapping and virtual-machines (VMs) allocation (i.e., network and server virtualization) over a physical infrastructure. CN survivability is crucial for computational resource allocation in a consistent and secure environment for cloud-computing services [4,6,7].




Fig. 1 Cloud-network mapping into a cloud-infrastructure provider (IaaS)

Figure 1 presents an example of two CNs consisting of interconnected VMs mapped over an optical network that interconnects datacenters (DC) of a cloud-infrastructure provider. Failures in the physical infrastructure can reduce the available resources (optical network and DCs) and disconnect multiple CNs. This may severely affect the upper-layer services [8]. CN survivability for a small number of failures in the physical infrastructure has been modeled as a survivable virtual-network embedding (SVNE) problem defined as the resilient VN mapping over the physical infrastructure to avoid disconnection due to failures [9]. Most SVNE studies considered single and multiple physical-link (-node) failures (e.g., datacenter and shared-risk group (SRG)), and a regional failure that may or may not be a disaster [9–13].

Disaster failure is a special case of SRG failure which may produce multiple failures in cascade, i.e., when a disaster occurs, some network elements may fail simultaneously in the first phase, and, later, other failures in different parts of the physical network (and upper layers) may occur (e.g., power outage and aftershocks after an earthquake). An important feature of cascading failures is that they tend to be more predictable from the damage and location of the initial failure, and this prediction can be used to reorganize the network to reduce disruptions [14].

An example of a disaster failure is the 2012 Hurricane Sandy, where post-disaster cascading failures (caused by flooding and power blackouts) shut down many datacenters and network nodes in the New York area [15], and caused disruption in communication services in the northeastern US [16]. Given the scale of their impact in CNs, network operators should take measures to protect cloud-computing services from disaster and post-disaster failures despite their rare occurrences.

In this study, we consider a disaster-survivable CN mapping approach using risk assessment (similar to [17]), virtual-machine (VM) backup location, and post-disaster survivability constraints to substantially reduce risk of failure, penalty, and probability of CN disconnection in case of disaster and post-disaster failures. In this work, to the best of our knowledge, we study for the first time:

- Integration of disaster and post-disaster survivable CN mapping with a risk assessment model to reduce the risk of CN disconnection.
- Use of a virtual-backup-node approach that can relocate VMs (i.e., VM backup location) to increase the cloud-computing survivability in case of disasters.

The rest of this study is organized as follows. Section 2 presents a brief review on cloud-network protection schemes and related works. Section 3 presents the survivable CN mapping problem. Section 4 describes our approach with an example. Section 5 introduces the variables and symbols and the ILP formulations of the baseline approach with risk minimization objective function. Section 6 introduces the ILP formulation of the proposed approach including VM backup location and post-disaster survivability constraints. An illustrative example is presented in Sect. 7, and our study concludes in Sect. 8.

2 Background and related works

A survey on network virtualization highlighting the importance of survivable virtual-network embedding (SVNE) is presented in [18]. Ref. [14] surveyed works on disaster survivability and pointed out works on disaster SVNE combined with VM location for datacenter networks.

Most studies on the SVNE problem suggested protection or restoration (e.g., reactive) approaches to deal with single physical-link (-node) failure. To deal with single physical-link failure, Ref. [19] proposed a fast-rerouting approach to recover failed VN, and Ref. [20] suggested to mix protection and restoration with backup capacity sharing to maximize revenue. Ref. [21] studied the SVNE problem for IP-over-WDM optical networks considering single and dual-link failures, introducing cut disjoint as a survivability constraint and a routing metric MINCUT. Cut-disjoint constraint avoids the mapping of two virtual links on the same physical resource if failures on both links disconnect the virtual topology (i.e., a cut of the topology). Ref. [22] used dedicated-path-protection and cut-disjoint approaches to increase the survivability. Ref. [23] showed the advantage of cut-disjoint approach over path-disjoint approach to provide protection in VN.

Refs. [12,24] proposed two versions of SVNE approach for physical-node failures (i.e., a datacenter failure in a regional failure) by adding backup node: 1-backup node (one backup node for each VN) and k-backup nodes (1+1 node protection). Ref. [25] presented an extension of these approaches, considering the network flow perspective to increase survivability.

Ref. [26] studied the SVNE problem in the context of grid- and cloud-computing survivability over optical networks, highlighting the importance of the survivable CN mapping

(SCNM) problem which combines the SVNE problem and VM survivability. In this regard, the study in [13] suggested server capacity relocation and lightpath re-provisioning for virtualized datacenters to offer survivability. Ref. [10] presented a model that helps to reduce the disaster failure in cloud services (i.e., cloud contents) provisioned over optical datacenter networks using a SRG-disjoint approach. Refs. [27,28] studied the SCNM problem combining with anycast routing, where VN mapping and anycast routing are optimized together to provide CN survivability. Ref. [11] studied disaster survivability in CN mapping, suggesting a disaster disjoint combined with non-survivable mapping to maximize revenue.

In this work, we address the SCNM problem for disaster failures using risk minimization, cut-disjoint constraint, virtual-machine (VM) backup location, and post-disaster survivability approaches.

3 Survivable CN mapping (SCNM)

The survivable CN mapping (SCNM) problem combines SVNE and VM resiliency. To address this problem, we consider a baseline SCNM approach to provide CN resiliency for any single physical-link failure while minimizing resources (Min-Res). To extend the baseline approach for disaster survivability, we also consider minimization of the risk of damage given the occurrence of a disaster (Min-Risk).

3.1 SCNM problem statement

Inputs:

- CN mapping requests and VM allocation requests with required communication and processing capacity.
- Physical network with link and node capacity (i.e., data-center capacity).

Output:

- Single physical-link failure survivable CN mapping.

Goal:

Minimize the communication resources used (i.e., wave-length channels).

3.2 Survivable mapping constraint

The survivable mapping constraint guarantees a survivable CN mapping for any single physical-link failure by enforcing cut-disjoint mapping as studied in [21–23]. This constraint ensures that virtual links of the same cut (i.e., set of links

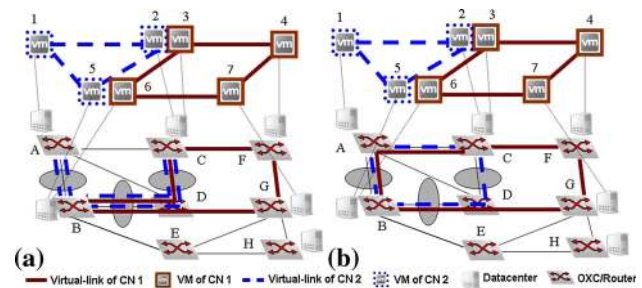


Fig. 2 a Non-survivable and b survivable CN mapping over a WDM optical network

whose simultaneous failures disconnects the virtual topology) do not share the same physical link. A simple example of SCNM approach is shown in Fig. 2. Two CNs are considered: CN 1 = {3, 4, 6, 7} and CN 2 = {1, 2, 5} mapped over an optical network with physical nodes (i.e., optical cross-connects (OXC/Routers) {A, B, C, D, E, F, G, H}, where some physical nodes {A, B, C, F, G, H} connect datacenters. Each virtual link is mapped using a lightpath. Figure 2a shows a non-survivable mapping where, if any of the physical links (shown in circles) fails (C–D or B–D or A–B), one or both CNs will be disconnected. Figure 2b shows an example of SCNM where no single physical-link failure will disconnect a CN.

3.3 Resource minimization (Min-Res)

The baseline objective is to minimize resource usage (Min-Res):

$$\min \sum_{\gamma \in \Gamma} (\text{Resources used by } \gamma) \tag{1}$$

where γ represents a CN request and Γ is the set of requests.

3.4 Disaster-survivable CN mapping with risk minimization (Min-Risk-DS)

The disaster-survivable CN mapping with risk minimization approach (Min-Risk-DS) extends Min-Res by including a disconnection constraint. Risk minimization offers two important advantages for the case of disaster survivability. The first advantage is the reduction in capacity (for backup) usage. The second advantage is the feasibility of the mapping in disaster zones (DZs) where the SRG-disjoint approach will not give a feasible mapping without additional resources for backup.

3.4.1 Risk assessment

Risk is defined as the expected value of an outcome seen as undesirable. In this work, we analyze the risk of CN based on damage/loss caused by a disaster [17], as shown below:

$$\min \sum_{n \in N} \sum_{\gamma \in \Gamma} (\text{Loss of } \gamma \text{ due to disaster } n) p_n \quad (2)$$

where the loss of CN γ ($\gamma \in \Gamma$) represents the sum of two values: (1) the penalty for CN disconnection which is the sum of the total disconnection penalty which represents capacity lost from the CN (i.e., total bandwidth) multiplied by a CN disconnection coefficient (i.e., value defined in the service-level agreement (SLA) which indicates the additional cost paid by the network provider to the customer or tenant when their CN is disconnected) and (2) the penalty of virtual-link disconnection in term of capacity lost. Finally, the risk is calculated by multiplying the resulting loss (i.e., total penalty) of γ by the probability p_n that disaster n can occur in the given disaster zone from the set of N possible disasters. Disasters are defined according to the approach used in [17] where the probability of a disaster and probability of damage are calculated based on hazard maps (see Sect. 7).

3.4.2 Example of risk minimization in CN mapping

To illustrate the impact of a disaster failure in CNs and the advantage of the Min-Risk-DS approach, we compare the mapping using Min-Res (Fig. 2b) with the mapping using Min-Risk-DS (Fig. 3b). Two disaster zones are included in Fig. 3, DZ1 and DZ2, with probability of occurrences (p_n) 0.3 and 0.5, respectively. Since DZ1 affects an entire node C, a SRG-disjoint approach will demand more resources for backup. To compare the two mappings, we calculate the total risk of CN disconnection using Eq. (2), assuming the bandwidth of each virtual link is 10 Mbps and a CN disconnection coefficient of 10 (we assume a value between 1 and 10). The risk of CNs mapped in Fig. 3a into the physical infrastructure in DZ1 is: Penalty for CN 1 disconnection, 40 Mbps (4 virtual links of 10 Mbps each) \times 10 = 400+ Penalty for CN 2 disconnection, 30 Mbps \times 10 = 300. The total risk of CN disconnection is $700 \times 0.3(p_1) = 210$. DZ2 does not disconnect any CN; hence, only 20 Mbps is affected, 20 Mbps (i.e., penalty for virtual-link disconnection) \times 0.5(p_2) = 10. Then, the total risk will be 220.




Fig. 3 a Min-Res approach (SCNM), b disaster-survivable CN mapping with risk minimization (Min-Risk-DS), with two DZs

Similarly, we can calculate the risk of CN mapping in Fig. 3b which is 210. The mappings of Fig. 3a, b use the same amount of resources (i.e., 120 Mbps each). However, the risk minimization can force the use of more resources in case of having more DZs. Hence, in this example, we confirm the necessity of VM backup location for further reduction in the risk of CN disconnection which is introduced in Sect. 4.

4 Disaster and post-disaster survivable CN mapping with risk minimization (Min-Risk-D-PDS)

Min-Risk-D-PDS extends Min-Risk-DS by adding two new functions to increase the disaster and post-disaster survivability of CNs. Note that, in the mapping of Fig. 3b, the risk is reduced by 10 units only and a disaster in DZ1 can still disconnect both CNs. To reduce the risk and increase CN survivability for case of disaster failures, Min-Risk-D-PDS introduces the concept of VM backup location (VBL) and post-disaster survivability (PDS).

4.1 Virtual backup node for VM backup location (VBL)

VBL maps one or more virtual backup node to relocate VMs of a CN, following three main steps: selection, connection, and sharing. For comparative purpose, we use the CN 1 nodes (3, 4, 6, 7) already used in Fig. 3 with one and two VM backup location (Fig. 4). These three steps are the main novelty and advantages of our proposed VBL approach over previous works in [11, 12, 25], in which risk of disaster and post-disaster survivability are not considered.

4.1.1 Selection of datacenter for VM backup location

The physical node (i.e., datacenter) selected as backup must not only have enough excess processing capacity but also should be located in a safer place to lower the risk of disconnection.

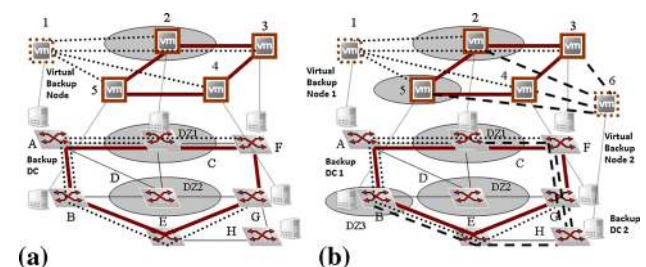


Fig. 4 Virtual backup node for VM backup location: a one VM backup location per CN, b two VM backup locations per CN

4.1.2 Connectivity of VM backup location

Every virtual backup node has to be connected using one virtual link to a set of working VMs in its own CN (Fig. 4a). The virtual links which connect the CN with its backup VM have 50% of the bandwidth of the working virtual link.

4.1.3 Physical node (i.e., datacenter) sharing for VM backup location

The selected physical node to provide VM backup location for one CN can be shared by another CN as working VM location and/or VM backup location. To increase the survivability to post-disaster failures, this approach will not allow to share the same physical node if both CNs can be disconnected by the same disaster. VBL has the flexibility to choose more than one physical node to relocate VMs based on the demand (Fig. 4b).

4.1.4 Example of VM backup location

By adding VBL into Min-Risk-DS approach (Fig. 4a), the risk of disconnection of CN 1 (Fig. 3b) is reduced from 120 (note that, we assume a penalty of disconnection of 400 and a p_n is equal to 0.3, so $120 = 400 \times 0.3$) to 9 (30 of penalty $\times 0.3$). Thanks to our approach, the CN does not get disconnected, so the risk of CN disconnection is reduced by 92% with an additional capacity of 30 Mbps (assuming 5 Mbps for each backup-virtual link).

As an example of two VM backup locations, in Fig. 4b, we add a third disaster zone, DZ3, with $p_3 = 0.5$, which increases the risk to 210 in the mapping of Fig. 4a. Then, we map a second virtual backup node which reduces the risk to 28 or 91.4% because only independent virtual links can be affected by disaster and the CN may remain connected. Also, the CN may survive if a disaster and post-disaster disconnect two VMs and create additional physical-link failures.

4.2 Post-disaster survivability (PDS)

However, if a disaster in DZ1 occurs, a post-disaster-correlated cascading failure of the physical link A–B will still disconnect the CN of Fig. 4a. Additionally, a post-disaster failure of physical links A–B and F–G will disconnect the CN of Fig. 4b. Hence, post-disaster survivability (PDS) constraint is added in our model to increase the survivability during recovery periods, given the vulnerability of CNs to post-disaster failures [14, 16]. Our (PDS) approach consists of two functions: cut extension and a survivability constraint.

4.2.1 Cut extension

We implement a new algorithm called ExCuts, which is an extension of the approach proposed in [22]. ExCuts extends

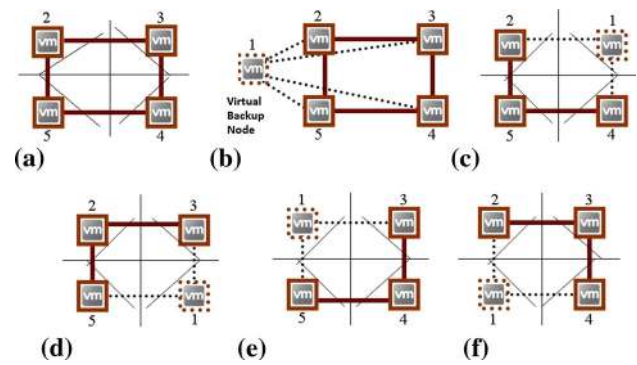


Fig. 5 Basic cuts, post-disaster cuts, and one VM backup location per CN. **a** CN with basic cuts, **b** CN with one VM backup location and **c–f** extended cuts for any replacement

Table 1 Example of basic and extended cuts

Basic cuts in Fig. 5a	Extended cuts in VM 1 replaces VM 3 (Fig. 5c)
(2–3)(2–5)	(2–1)(2–5)
(2–5)(5–4)	(2–5)(5–4)
(5–4)(4–3)	(5–4)(4–1)
(3–2)(4–3)	(1–2)(4–1)
(2–5)(4–3)	(2–5)(4–1)
(3–2)(4–5)	(1–2)(4–5)
(3–2)(4–5)(2–5)	(1–2)(4–5)(2–5)
(3–4)(2–5)(5–4)	(1–4)(2–5)(5–4)
(3–4)(2–5)(5–4)(2–3)	(1–4)(2–5)(5–4)(2–1)

the basic cuts of the CN 1 topology in three steps. To describe the steps, we use CN 1 (Fig. 5a) and one possible replacement of VM 3 by VM 1 (i.e., as virtual backup node).

Step i: ExCuts replaces the working VM 3 for VM 1 as possible relocation and builds a new topology (Fig. 5c).

Step ii: ExCuts renumbers the basic cuts with virtual links of the resulting topology of Fig. 5c. In Table 1, we show the basic and extended cuts of the resulting topology when VM 3 is disconnected and replaced by VM 1.

Step iii: ExCuts eliminates redundant cuts and repeats the three steps for each possible VM relocation of Fig. 5(c–f).

In this example, we consider only one datacenter for VM backup location. However, ExCuts will generate new cuts considering all possible VM relocation given a disaster failure.

4.2.2 Survivability constraint

The extended cuts are input to the novel survivability constraint which enforces survivable mapping against any post-disaster single physical-link failure. The constraint applies the concept of cut-disjoint approach introduced in Sect. 3 but considering post-failure cuts to increase the post-disaster




Fig. 6 Post-disaster cuts for two VM backup locations per CN. **a** CN with two VM backup locations, and **b–g** extended cuts for the replacement of the two failed VMs

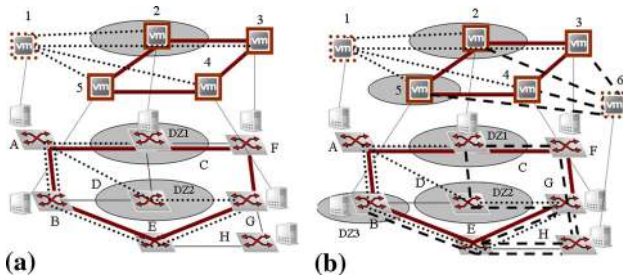


Fig. 7 Resulting mapping by Min-Risk-D-PDS with **a** one and **b** two VM backup locations

survivability. Figure 6 presents the cut extension of Fig. 5 for two VM backup locations.

4.3 Example of Min-Risk-D-PDS approach

In the mapping of Fig. 4a, if a disaster, e.g., in DZ1, occurs, the physical node C and its physical links will fail, but the CN will not be disconnected, because the failed VM in node 2 will be relocated into physical node A (VM in node 1). However, a post-disaster failure in physical link A–B will disconnect the CN, because virtual links 1–5 and 1–4 will be disconnected. Similarly, failure of any of physical links B–E, F–G, and E–G may disconnect the CN.

Min-Risk-D-PDS obtains the mapping in Fig. 7a, where the CN will not be disconnected by any single physical-link failure, disaster failure, or post-disaster single physical-link failure, and the expected loss of bandwidth and processing capacity will be reduced.

5 ILP formulation of Min-Risk-DS

In this section, we present the ILP formulation of the baseline approach Min-Risk-DS which has three elements: Min-Risk formulation, CN mapping, and survivability constraints.

Before we describe the formulation, we introduce the parameters and variables of the problem.

5.1 Variables and symbols

Given

- $G(V, E)$: Physical topology, where V is the set of physical nodes and E is the set of physical links.
- \hat{V} : Set of VM datacenter locations, $\hat{V} \subset V$.
- $G_\gamma(V_\gamma, E_\gamma)$: Topology of CN γ where V_γ is the set of working VM locations (virtual nodes, $V_\gamma \subset \hat{V}$), and E_γ the set of virtual links of CN.
- C_γ : Set of basic cuts of CN topology γ .
- \hat{E}_γ : Set of virtual links including the links in E_γ and virtual links from each node in V_γ to each node in $\{\hat{V} - V_\gamma\}$.
- \hat{C}_γ : Set of extended cuts of CN topology γ formed by a possible relocation of working VM of V_γ to a physical node b with free processing capacity in $\{\hat{V} - V_\gamma\}$.
- $\Gamma = \{\gamma = \langle V_\gamma, E_\gamma, C_\gamma, \hat{E}_\gamma, \hat{C}_\gamma \rangle\}$: Set of cloud networks (CNs).
- $s_{i,j}^n$: 1 if the physical link $\{i, j\}$ is disconnected by disaster n , zero otherwise.
- $S_n = \{s_{i,j}^n\}$, $S_n \subset E$.
- p_n : Probability of occurrence of disaster n .
- $N = \{\langle S_n, p_n \rangle\}$: Set of disasters zones (i.e., DZs).
- P_u^γ : Processing capacity required to allocate VM u used by CN γ ($u \in V_\gamma$).
- P_{free}^v : Excess processing capacity in physical node v .
- $F_{i,j}$: capacity of physical link (i, j) .
- d : CN disconnection coefficient ($1 \geq d \leq 10$).
- b_e : Bandwidth requirement of virtual link e .
- b_c : Total capacity that can be lost if the links of the cut c are disconnected (i.e., the CN is disconnected).
- m_c : Number of virtual links in cut c .

Binary variables

- D_e^n : 1 if virtual link e is disconnected by disaster n .
- $M_{i,j}^e$: 1 if virtual link e is mapped on physical link (i, j) .
- $K_{u,v}^{\gamma,e}$: 1 if virtual link e from node u to v in γ .
- Y_b^γ : 1 if b is assigned as virtual backup node of γ .
- Q_c^n : 1 if virtual links of the cut c is disconnected by disaster n .
- X_γ^n : 1 if CN γ may be disconnected by disaster n , 0 otherwise.¹
- $T_{g,h}^n$: is an auxiliary variable.
- $Z_{u,b}^\gamma$: 1 if VM u can be relocated to datacenter b , $b \in \hat{V}$.

¹ Since we are using a probabilistic model, this variable only indicates if a cloud network can be affected by a disaster or not. The actual probability of disconnection will depend on the disaster intensity.

5.2 Min-Risk formulation and constraints

5.2.1 Objective function

The objective is to minimize the total capacity that can be lost if a disaster occurs. The risk as defined in Sect. 4 is the total penalty for capacity loss multiplied by the probability of occurrence. The total penalty for capacity lost is the sum of penalty for CN and virtual links' disconnections. The penalty for CN disconnection is calculated by $\sum_{c \in C_\gamma} d Q_c^n b_c$ which is the sum of capacity b_c that is lost if a CN is disconnected by disaster n multiplied by a CN disconnection coefficient d . The penalty for virtual-link disconnection is calculated by $\sum_{e \in \hat{E}_\gamma} D_e^n b_e$ which is the sum of capacity b_e that is lost when virtual links e is disconnected by disaster n . Finally, the objective function is:

$$\min \sum_{n \in N} \sum_{\gamma \in \Gamma} \left(\sum_{c \in C_\gamma} d Q_c^n b_c + \sum_{e \in \hat{E}_\gamma} D_e^n b_e \right) p_n + \left(\epsilon \times \sum_{(i,j) \in E} \sum_{\gamma \in \Gamma} \sum_{e \in \hat{E}_\gamma} M_{i,j}^e \times b_e \right) \quad (3)$$

To avoid the mapping of virtual links over long lightpaths, a resource minimization formula is added with a coefficient ϵ . A very small value of ϵ will give more importance for risk minimization in the mapping over resources used.

5.2.2 Constraint to determine whether a virtual link is affected by a disaster

$$D_e^n \geq \frac{1}{M} \sum_{(i,j) \in E} s_{i,j}^n M_{i,j}^e, \forall e \in \hat{E}_\gamma, \gamma \in \Gamma, n \in N \quad (4a)$$

$$D_e^n \leq \sum_{(i,j) \in E} s_{i,j}^n M_{i,j}^e, \forall e \in \hat{E}_\gamma, \gamma \in \Gamma, n \in N \quad (4b)$$

where M is a large number.

5.2.3 Constraint to determine a CN disconnection (i.e., cut failure) due to a disaster

$$Q_c^n \leq \frac{\sum_{e \in E_c} D_e^n}{m_c}, \forall c \in C_\gamma, \gamma \in \Gamma, n \in N \quad (5a)$$

$$Q_c^n \geq \sum_{e \in E_c} D_e^n - m_c + 1, \forall c \in C_\gamma, \gamma \in \Gamma, n \in N \quad (5b)$$

The CN is disconnected when the value of Q_c^n is 1, i.e., disaster n disconnects all the virtual links e (D_e^n) belonging to a cut c .

5.3 CN mapping constraints

The basic constraints used in the mapping are for virtual-link mapping, flow conservation and, physical-link (i.e., optical link) capacity in number of wavelengths available for mapping.

5.3.1 Virtual-link mapping constraint

$$K_{u,v}^{\gamma,e} = 1, \forall u, v \in V_\gamma, u \neq v, \gamma \in \Gamma, e \in \hat{E} \quad (6)$$

This constraint maps the CN γ , connecting the VMs u and v .

5.3.2 Flow-conservation constraints

$$\sum_{(i,s_e) \in E} M_{i,s_e}^e - \sum_{(s_e,j) \in E} M_{s_e,j}^e = -K_{s_e,d_e}^{\gamma,e} \quad (7a)$$

$$\sum_{(i,d_e) \in E} M_{i,d_e}^e - \sum_{(d_e,j) \in E} M_{d_e,j}^e = K_{s_e,d_e}^{\gamma,e} \quad (7b)$$

$$\sum_{(k,j) \in E} M_{k,j}^e - \sum_{(i,k) \in E} M_{i,k}^e = 0, \forall e \in \hat{E}_\gamma, \gamma \in \Gamma, k \in \hat{V} - \{s_e, d_e\} \quad (7c)$$

These constraints ensure that each virtual link is mapped on a lightpath, and it does not pass the same physical node more than once.

5.3.3 Physical-link capacity constraint

$$\sum_{e \in \hat{E}_\gamma} M_{i,j}^e \leq F_{i,j}, \forall (i,j) \in E, \gamma \in \Gamma \quad (8)$$

5.4 Survivability constraint

The survivability constraint uses the basic cuts of the CN topology C_γ . The constraint enforces that all links (m_c) of the cut c do not use the same physical link.

$$\sum_{e \in E_c} M_{i,j}^e \leq m_c - 1, \forall c \in C_\gamma, \gamma \in \Gamma, (i,j) \in E \quad (9)$$

6 ILP formulation of Min-Risk-D-PDS

Min-Risk-D-PDS is our comprehensive approach which extends the ILP formulation of the baseline approach Min-Risk-DS by adding the VM backup location (VBL) and post-disaster survivability (PDS) constraints.

6.1 VBL constraints

6.1.1 Disaster-disjoint VM backup location constraint

This set of constraints enforce that two or more CNs do not share the same physical node as VM backup location if the CNs are affected by the same disaster [Eqs. (10), (11), and (12)]. Equation (10) identifies which disaster n disconnects the CN γ , giving value 1 to X_γ^n , 0 otherwise.

$$X_\gamma^n \geq \frac{1}{M} \sum_{c \in C_\gamma} Q_c^n, X_\gamma^n \leq \sum_{c \in C_\gamma} Q_c^n, \forall \gamma \in \Gamma, n \in N \quad (10)$$

Equation (11) uses the value of X_γ^n and an auxiliary variable $T_{g,h}^n$ to identify the disaster which disconnect CNs h and g .

$$T_{g,h}^n \leq X_g^n, T_{g,h}^n \leq X_h^n, \forall g, h \in \Gamma, g \neq h, n \in N \quad (11a)$$

$$T_{g,h}^n \geq X_g^n + X_h^n - 1, \forall g, h \in \Gamma, g \neq h, n \in N \quad (11b)$$

Equation (12) restricts two CNs (g and h) to share the same physical node (b) for VM backup location if both CNs are disconnected by the same disaster.

$$Y_b^g + Y_b^h \leq 2 - T_{g,h}^n, \forall g, h \in \Gamma, g \neq h, n \in N, \\ b \in [\hat{V} - (V_g \cup V_h)] \quad (12)$$

6.1.2 Mapping of VM backup location constraint

This constraint gives the bound for the number of VM backup location per CN. It has two set of equations: VM backup location selection and bound on number of VM location per CN. Equation (13) chooses the less risky VM backup location b for each CN γ . Equation (13a) ensures that the VM backup location b will not be chosen from the working VM V_γ of CN γ .

$$Y_b^\gamma = 0, \forall b \in V_\gamma, \gamma \in \Gamma \quad (13a)$$

$$Y_b^\gamma \geq \sum_{(u \in V_\gamma)} \frac{Z_{u,b}^\gamma}{M}, \forall b \in (\hat{V} - V_\gamma), \gamma \in \Gamma, u \in E \quad (13b)$$

$$Y_b^\gamma \leq \sum_{u \in V_\gamma} Z_{u,b}^\gamma, \forall b \in (\hat{V} - V_\gamma), \gamma \in \Gamma \quad (13c)$$

Equation (14) bounds the number of VM backup location between 2 and certain maximum number.

$$\sum_{b \in (\hat{V} - V_\gamma)} Y_b^\gamma \geq 2, \sum_{b \in (\hat{V} - V_\gamma)} Y_b^\gamma \leq |V_\gamma|, \forall \gamma \in \Gamma \quad (14)$$

6.1.3 Connecting the VM backup node for relocation

When VM backup location is selected, virtual links connect it to working VMs [Eq. (15)]. The connection follows two conditions:

- (i) When one or more VMs choose a VM backup location. In this regard, $Z_{v,b}^\gamma$ is 1, meaning that working VM used by CN in physical node v chose to be relocated to physical node b . As a result, the variable $K_{v,b}^{\gamma,e}$ will be 1, forcing the mapping of virtual link e into the physical network.
- (ii) When the VM backup location mapped in b is already connected to v , ($K_{v,b}^{\gamma,e} = 1$), and the VM in physical node u is neighbor of v . Hence, a virtual link connects one working VM u with a VM backup location b of the same CN ($K_{u,b}^{\gamma,e} = 1$).

$$K_{u,b}^{\gamma,e} \leq Z_{v,b}^\gamma, K_{u,b}^{\gamma,e} \leq K_{v,u}^{\gamma,e}, K_{u,b}^{\gamma,e} \geq Z_{v,b}^\gamma + K_{v,u}^{\gamma,e} - 1 \quad (15a)$$

$$K_{u,b}^{\gamma,e} = Z_{v,b}^\gamma, \forall v, u \in V_\gamma, (b \in (\hat{V} - V_\gamma), \gamma \in \Gamma) \quad (15b)$$

6.2 Processing capacity required for VM backup location

This constraint manages the free capacity of each physical node used for VM backup location. If P_{free}^b is zero, the physical node ($Y_b^\gamma = 0$) cannot be used (e.g., the required capacity of the CN (P_u^γ) is higher or the free capacity (P_{free}^b) is not enough).

$$P_{\text{free}}^b - \sum_{u \in V_\gamma} P_u^\gamma \cdot Y_b^\gamma \geq 0 \quad (16)$$

6.3 PDS constraint

PDS uses the same formulation presented in Eq. (9) with the extended cut \hat{C}_b^γ as additional input.

7 Illustrative examples

7.1 Experimental setup

We test our approaches on a 24-node US mesh opaque WDM optical network (Fig. 8b) with 32 wavelengths per link. Two types of disasters are considered: natural disasters (earthquake) and human-made disasters (weapons-of-mass-destruction (WMD) attacks), originally modeled in [17] and shown in Fig. 8b. For earthquakes, the probability of occurrence and damage is obtained with seismic hazard maps. And for WMD attacks, the probability of attack and damage is based on cities population and importance [17].

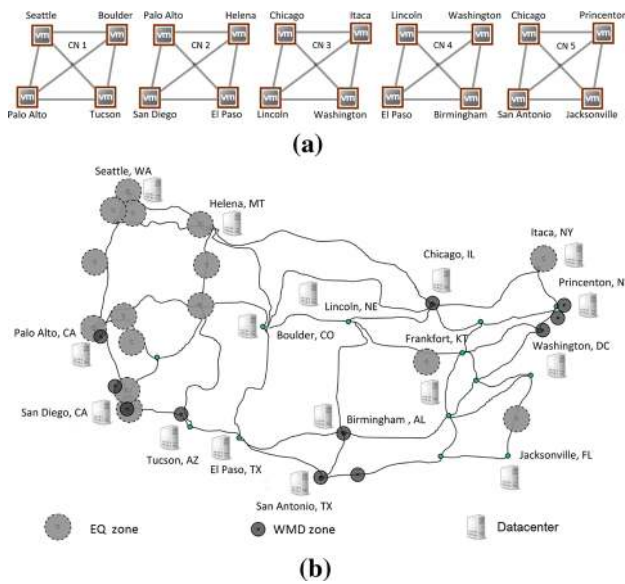


Fig. 8 **a** CNs studied and **b** physical topology with disaster zones for earthquake and potential WMD attacks [17], and datacenter locations

We consider five full-mesh cloud networks (CNs), each consisting of four virtual nodes (i.e., VMs) distributed over 16 datacenters (Fig. 8a). We assume that each virtual link requires a full lightpath (i.e., wavelength channels), and each datacenter has enough processing capacity.

7.2 Survivable CN mapping approaches

We tested eight approaches: four minimizing resources (Min-Res) and four minimizing risk (Min-Risk). All approaches use a set of baseline survivability constraints (SC). Some of them use a disaster survivable mapping (DS), disaster and post-disaster survivable constraints (D-PDS), and VM backup location (VBL) with number of backup location: one (1L) or two (2L). Min-Res-DS-1L indicates minimization of resources, disaster-survivable mapping with 1 VM backup location which we call RESA-1L. The list of approaches is presented in Table 2 including our proposed approaches.

Table 2 Approaches used in illustrative examples

Name	Approach	PDS	VBL	Cuts
RESA	Min-Res			Basic
RISKA	Min-Risk-DS			Basic
RESA-1L	Min-Res-DS-1L		1L	Basic
RISKA-1L	Min-Risk-DS-1L		1L	Basic
RESA-PDS	Min-Res-D-PDS	X	1L	Extended
RISKA-PDS	Min-Risk-D-PDS	X	1L	Extended
RESA-2L	Min-Res-D-PDS-2L	X	2L	Extended
RISKA-2L	Min-Risk-D-PDS-2L	X	2L	Extended

7.3 Evaluation and comparative methodologies

Our examples are evaluated using risk and penalty, disaster and post-disaster survivability, and resource usage analysis.

7.3.1 Risk and penalty

The risk of CN disconnection is evaluated using the first part of Eq. (3). The penalty for capacity loss is the total capacity that can be lost due to a disaster.

7.3.2 Disaster and post-disaster survivability analysis

The second analysis is the evaluation of the probability of CN disconnection (PoD). The PoD is calculated by an algorithm called cloud-network resiliency test algorithm (CNRT) which tests the vulnerability of the CN to all possible combinations of disaster and post-disaster failures. CNRT gets the mapping of each CN and simulates disaster damage over the physical infrastructure based on given disaster scenarios (Table 3). Then, the algorithm tests the connectivity of every VM and counts the number of possible failure scenarios caused by a disaster in which the CN is disconnected. With these numbers, CNRT obtains one PoD for each CN and type of failure using Eq. (17).

$$PoD = \frac{\text{Total number CN disconnection}}{\text{Total number of possible failures}} \tag{17}$$

7.4 Numerical analysis

To study the risk and penalty, we use the mapping of the five CNs presented in Fig. 8a. However, we select CN 1 for earthquake and CN 3 for WMD to study the disaster and post-disaster scenarios, as these two CNs are more affected by the disasters.

7.4.1 Risk and penalty analysis

Figure 9 compares the expected risk of CN disconnection of different approaches. In Fig. 9 we observe that:

Table 3 Simulated failures

Symbols	Description	Disaster	Post-disaster failures	
			Physical link/s	Disaster
DF	Any single disaster occurs	Single	–	–
DSLDF	One physical link fails after a disaster	Single	Single	–
DDLDF	Two physical links fail after a disaster	Single	Dual	–
DFDF	Second disaster occurs after a disaster	Single	–	Single

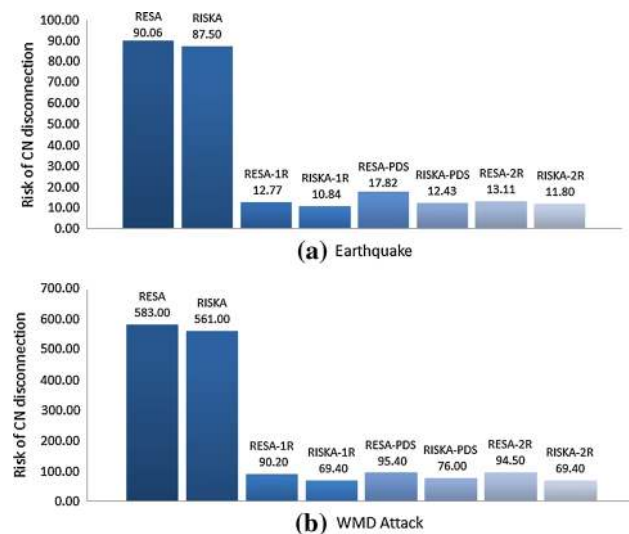


Fig. 9 Risk of CN disconnection **a** earthquake and **b** weapon of mass destruction (WMD)

- (i) RISK A approach reduces the risk of CN disconnection and penalty by 2.75–3.77%. These results show a low risk reduction without VBL constraint, and the limitation of SVN M-based approaches to deal with disaster and post-disaster failures.
- (ii) By adding the VM backup location (VBL), RISK A-1L approach reduces the risk of CN disconnection and penalty up to 87% for earthquake, and up to 88% for WMD. Also, RESA-1L approach reduces risk up to 85% for earthquake, and up to 87% for WMD. It confirms that VBL approach reduces considerably the CN disconnection and penalty for capacity loss. However, VBL works better with RISK A (i.e., larger risk and penalty reduction by 10–30%).
- (iii) PDS constraint slightly increases the risk because the extended cuts force virtual links to be mapped in longer lightpaths. However, PDS constraint increases survivability against post-disaster failures by 60–100% (Table 4).
- (iv) The combination of PDS and VBL with two VM backup locations per CN obtains more reduction in risk and penalty. However, the risk and penalty reduction tend to be lower in earthquake case and higher for WMD for one VM backup location per CN.

7.4.2 Disaster and post-disaster survivability study

After risk and penalty analysis, we study the probability of disconnection (PoD) due to a disaster failure and three kinds of post-disaster failures presented in Table 3.

Table 4 presents the PoD of CN 1 and CN 3. We observe that:

Table 4 Probability of disconnection (PoD)

Approach	CN 1—Earthquake		CN 3—WMD attack	
	DF	DSL F	DF	DSL F
RESA	0.27	0.45	0.18	0.38
RISK A	0.27	0.35	0.09	0.38
RESA-1L	0	0.30	0	0.29
RISK A-1L	0	0.26	0	0.20
RESA-PDS	0	0	0	0.14
RISK A-PDS	0	0	0	0
RESA-2L	0	0	0	0
RISK A-2L	0	0	0	0
	DDL F	DFDF	DDL F	DFDF
RESA	0.50	0.52	0.42	0.35
RISK A	0.42	0.49	0.41	0.22
RESA-1L	0.38	0.19	0.35	0.04
RISK A-1L	0.35	0.11	0.24	0.02
RESA-PDS	0.35	0.13	0.17	0
RISK A-PDS	0.20	0.13	0.15	0
RESA-2L	0.23	0.01	0.17	0
RISK A-2L	0.20	0	0.15	0

- (i) DF: CNs with VBL will completely survive any failure as any VM can be relocated from one datacenter to another, i.e., PoD = 0. In addition, RISK A approach increases the survivability by 50% in WMD case compared with RESA approach.
- (ii) DSL F: RISK A approach reduces PoD by 0–22% compared with RESA approach. And, RISK A-1L (i.e., with VBL) increases the survivability by 37–100% compared with RESKA-based approaches. PDS constraint increases the survivability to 100% independent of the number of VM backup locations and the objective function (RISK A or RESA).
- (iii) DDL F: RISK A achieves a reduction in PoD by 2.3% in WMD case and 16% in earthquake case compared with RESA. However, when VBL is used, the reduction in PoD is higher (between 24 and 64%). PDS constraint has positive impact, because the reduction is higher for RISK A-PDS compared to other approaches without PDS constraints.
- (iv) DFDF: VBL reduces the PoD remarkably by 78–100%. Also, including PDS constraint with RISK A-based approach does not enhance the performance significantly. However, RESA-based approaches with PDS achieve an important reduction of 33% in PoD.

7.4.3 Resource consumption analysis

In this analysis, we study the resources used to provide reduction in risk, penalty for capacity loss, and PoD. From the

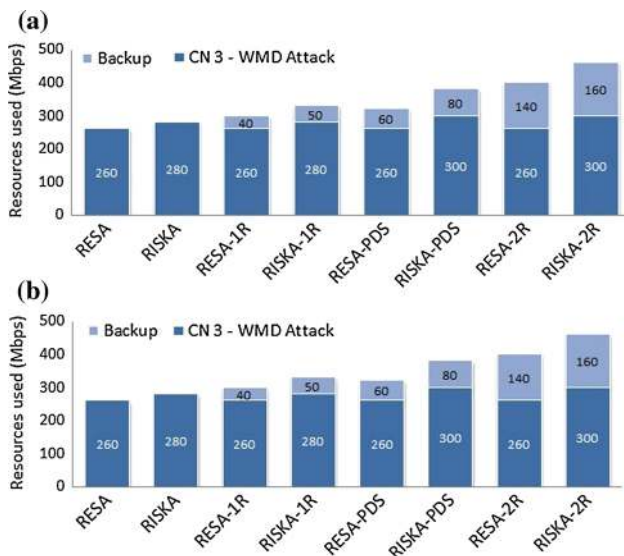


Fig. 10 Resources used (in Mbps) by the mapping of **a** CN 1 in earthquake case **b** CN 3 in WMD case

previous analysis and the results of Fig. 10, we observe that:

- (i) RISK A-based approaches require additional resources by 7.8–16 % to reduce the risk and penalty and PoD. RISK A with VBL constraints increases resource usage by 16–37 % for one VM backup location (RISK A-1L) to provide risk and penalty reduction by 85–87 %, and a reduction in the PoD by 24–100 % (i.e., increasing the survivability by 24–100 %). This result confirms that SVN M cannot deal with disasters and their consequences.
- (ii) PDS constraint with RISK A and VM backup location (RISK A-PDS) increases the resources by 25–50 % in CN 1 (earthquake) and by 23–38 % for CN 3 (WMD). However, the risk and penalty are reduced up to 88 %, and the survivability increases up to 100 % in cases of disaster and post-disaster failures.
- (iii) Two VM backup locations require more resources, but increase the survivability for more severe disaster scenarios which may disconnect two VMs.

8 Conclusion

We studied the disaster and post-disaster survivable cloud-network (CN) mapping problem. We proposed a CN mapping approach Min-Risk-D-PDS using (i) VM backup location for each CN (VBL) and (ii) post-disaster survivability constraint (PDS), which offer an economically sustainable disaster and post-disaster survivable CN mapping approach.

We formulated the Min-Risk-D-PDS as an integer linear program. We compared our approach with seven different approaches characterized by different combinations of VBL

and PDS constraints with risk and resources minimization as objective function.

Results on a case study formed by five CNs mapped over a US network and two disaster cases (earthquake and WMD) showed that Min-Risk-D-PDS (RISK A-PDS) reduces the risk of CN disconnections and penalty for capacity loss by 85–90 %. As a consequence, our approach increases the CN survivability by 60 and 100 % against three kinds of post-disaster failures with the cost of 23–50 % of additional resources usage.

Hence, our illustrative examples confirm the importance of VM backup location and post-disaster survivability constraints for CN survivability against any disaster and post-disaster correlated, cascading failures that may occur in the network.

As future work, we are exploring the use of heuristic approaches to increase the scalability of Min-Risk-D-PDS for dynamic scenarios and to extend our disaster-resiliency study by adding new comparative metrics (e.g., blocking probability of dynamic CN mapping request).

References

- [1] Meixner, C.C., Dikbiyik, F., Tornatore, M., Chuah, C., Mukherjee, B.: Disaster-resilient virtual-network mapping and adaptation in optical networks. In: 17th International Conference on Optical Network Design and Modeling (ONDM), Brest, France (2013)
- [2] Devellder, C., De Leenheer, M., Dhoedt, B., Pickavet, M., Colle, D., De Turck, F., Demeester, P.: Optical networks for grid and cloud computing applications. *Proc. IEEE* **100**(5), 1149–1167 (2012)
- [3] Contreras, L., Lopez, V., De Dios, O., Tovar, A., Munoz, F., Azanon, A., Fernandez-Palacios, J., Folgueira, J.: Toward cloud-ready transport networks. *IEEE Commun. Mag.* **50**(9), 48–55 (2012)
- [4] Mogul, J.C., Popa, L.: What we talk about when we talk about cloud network performance. *SIGCOMM Comput. Commun. Rev.* **42**(5), 44–48 (2012)
- [5] Rimal, B.P., Choi, E., Lumb, I.: A taxonomy and survey of cloud computing systems. In: Proceedings of the IEEE International Joint Conference on INC, IMS and IDC, Washington, DC, USA (2009)
- [6] Abbadi, I.: Clouds infrastructure taxonomy, properties, and management services. *Adv. Comput. Commun.* **193**, 406–420 (Jun. 2011)
- [7] Sun, G., Yu, H., Anand, V., Li, L., Di, H.: Optimal provisioning for virtual network request in cloud-based data centers. *Photonic Netw. Commun.* **24**(2), 118–131 (2012)
- [8] Kounev, S., Reinecke, P., Brosig, F., Bradley, J.T., Joshi, K., Babka, V., Stefanek, A., Gilmore, S.: Providing dependability and resilience in the cloud: challenges and opportunities, chap. 4. In: Wolter, K., Avritzer, A., Vieira, M., van Moorsel, A. (eds.) *Resilience Assessment and Evaluation of Computing Systems*, pp. 65–81. Springer Berlin Heidelberg (2012)
- [9] Chowdhury, N., Rahman, M., Boutaba, R.: Virtual network embedding with coordinated node and link mapping. In: Proceedings of the IEEE International Conference on Computer Communications (INFOCOM), Rio de Janeiro, Brazil (2009)

- [10] Habib, M., Tornatore, M., De Leenheer, M., Dikbiyik, F., Mukherjee, B.: Design of disaster-resilient optical datacenter networks. *IEEE/OSA J. Lightw. Technol.* **30**(16), 2563–2573 (2012)
- [11] Gu, F., Alazemi, H., Rayes, A., Ghani, N.: Survivable cloud networking services. In: *Proceedings of the IEEE International Conference on Computing, Networking and Communications (ICNC)*, San Diego, USA (2013)
- [12] Yu, H., Anand, V., Qiao, C.: Virtual infrastructure design for surviving physical link failures. *Comput. J.* **55**(8), 965–978 (2012)
- [13] Xu, J., Tang, J., Kwiat, K., Zhang, W., Xue, G.: Survivable virtual infrastructure mapping in virtualized data centers. In: *Proceedings of the IEEE Cloud Computing Conference (CLOUD)*, Honolulu, Hawaii, USA (2012)
- [14] Habib, M.F., Tornatore, M., Dikbiyik, F., Mukherjee, B.: Disaster survivability in optical communication networks. *Comput. Commun.* **36**(6), 630–644 (2013)
- [15] Carew, S.: Hurricane Sandy disrupts Northeast U.S. Telecom Networks. Reuters, [Online]. <http://uk.reuters.com/article/2012/10/30/us-storm-sandy-telecommunications-idUKBRE89T0YU20121030> (2012)
- [16] Henderson, N.: Noise filter: hurricane sandy floods NYC data center, impacts hosts, colocation providers. *WebHost Ind. Rev.* [Online]. <http://www.thewhir.com/web-hosting-news/noise-filter-hurricane-sandy-floods-nyc-data-center-impacts-hosts> (2012)
- [17] Dikbiyik, F., Leenheer, M.D., Reaz, A., Mukherjee, B.: Minimizing the disaster risk in optical telecom networks. In: *Proceedings of the IEEE/OSA Optical Fiber Communication Conference (OFC)* (2012)
- [18] Chowdhury, N., Boutaba, R.: A survey of network virtualization. *Comput. Netw.* **54**(5), 862–876 (2010)
- [19] Rahman, M., Aib, I., Boutaba, R.: Survivable virtual network embedding. In: *Crovella, M., Feeney, L., Rubenstein, D., Raghavan, S. (eds.) NETWORKING 2010*, ser. Lecture Notes in Computer Science, vol. 6091, pp. 40–52. Springer, Berlin (2010)
- [20] Guo, T., Wang, N., Moessner, K., Tafazolli, R.: Shared backup network provision for virtual network embedding. In: *Proceedings of IEEE International Conference on Communications (ICC)*, Kyoto, Japan (2011)
- [21] Lee, K., Modiano, E., Lee, H.: Cross-layer survivability in WDM based networks. *IEEE/ACM Trans. Netw.* **19**(6), 1000–1013 (2011)
- [22] Vadrevu, C.S., Tornatore, M.: Survivable IP topology design with re-use of backup wavelength capacity in optical backbone networks. *Opt. Switch. Netw.* **7**(4), 196–205 (2010)
- [23] Jaumard, B., Hoang, A., Bui, M.: Path vs. cutset approaches for the design of logical survivable topologies. In: *Proceedings of IEEE International Conference on Communications (ICC)*, Ottawa, Canada (2012)
- [24] Yu, H., Anand, V., Qiao, C., Sun, G.: Cost efficient design of survivable virtual infrastructure to recover from facility node failures. In: *Proceedings of IEEE International Conference on Communications (ICC)*, Kyoto, Japan (2011)
- [25] Hu, Q., Wang, Y., Cao, X.: Survivable network virtualization for single facility node failure: a network flow perspective. *Opt. Switch. Netw.* **10**(4), 406–415 (2013)
- [26] Develder, C., Buysse, J., Shaikh, A., Jaumard, B., De Leenheer, M., Dhoedt, B.: Survivable optical grid dimensioning: anycast routing with server and network failure protection. In: *Proceedings of IEEE International Conference on Communications (ICC)*, Kyoto, Japan (2011)
- [27] Bui, M., Jaumard, B., Develder, C.: Anycast end-to-end resilience for cloud services over virtual optical networks (invited). In: *Proceedings of 15th International Conferent Transparent Optical Networks (ICTON)*, Cartagena, Spain (2013)
- [28] Barla, I., Schupke, D., Hoffmann, M., Carle, G.: Optimal design of virtual networks for resilient cloud services. In: *Proceedings of*

9th International Conference on the Design of Reliable Communication Networks (DRCN), Budapest, Hungary (2013)



Carlos Colman-Meixner is Ph.D. student in the Department of Electrical and Computer Engineering at University of California, Davis. He received his M.Sc. in Systems Engineering at National University of Asuncion, Paraguay in 2008. He received his B.Sc. degree from National University of the East, Ciudad del Este, Paraguay, in 2000. He has worked as a Networking Instructor and Professor in the Computer Engineering undergrad program of University of Cone Sur of the Americas (UCSA), Paraguay between 2001 and 2011. He was co-founder of the UCSA Technological Center one of the first information technology, educational and, entrepreneurial center of Paraguay. His research interest includes robustness and disaster survivability in virtualized networks and machines, optical networks and, cloud computing. Additionally, he has research interests in multi-objective optimization in next-generation networks and cloud-computing design problems.



especially excess capacity management, survivability against disasters, and network upgrade.



survivability, datacenter network, network modeling, and optimization.

M. Farhan Habib received the B.Sc. (Eng.) degree in Computer Science and Engineering from Bangladesh University of Engineering and Technology, Dhaka, Bangladesh and the M.Sc. degree in Computer Science from University of California, Riverside, in 2004 and 2010, respectively. He is currently a Ph.D. candidate in the Department of Computer Science at the University of California, Davis. His research interests include network



Massimo Tornatore is currently working as an assistant professor in the Department of Electronics, Information and Bioengineering at Politecnico di Milano, where he received a Ph.D. degree in Information Engineering in 2006 and a Laurea (M.Sc. equivalent) degree in October 2001. He also holds an appointment as adjunct associate professor in the Department of Computer Science at the University of California, Davis, where he served as a postdoc researcher in 2008 and 2009. He started his carrier as a researcher in 2001, as an intern at CoreCom Optical Network Laboratory. During 2002, he was granted a research scholarship and he worked in collaboration with PSTS (Pirelli Submarine Telecom Systems) and TILAB (Telecom Italia

Labs). From May to December 2004 he visited the Networks Lab in U.C. Davis, under the supervision of Prof. Biswanath Mukherjee. He was also visiting Ph.D. student in the Optical Networking group of CTTC laboratories. He is author of more than 170 conference and journal papers and his research interests include design, protection and energy efficiency in optical transport and access networks and group communication security. He was a co-recipient of 6 Best Paper Awards from IEEE conferences.



Chen-Nee Chuah is a Professor in Electrical and Computer Engineering at the University of California, Davis. She received her Ph.D. degree in Electrical Engineering and Computer Sciences from the University of California, Berkeley. Her research interests include Internet measurements, network management, anomaly detection, online social networks, and vehicular ad hoc networks. Chuah is an ACM Distinguished Scientist. She received the NSF CAREER Award in 2003, and the Outstanding Junior Faculty Award from the UC Davis College of Engineering in 2004. In 2008, she was named a Chancellor's Fellow of UC Davis. She has served on the executive/technical program committee of several ACM and IEEE conferences. She was an Associate Editor for IEEE/ACM Transactions on Networking from 2008 to 2013.



Biswanath Mukherjee is Distinguished Professor at University of California, Davis, where he was Chairman of Computer Science during 1997–2000. He received the B.Tech. degree from Indian Institute of Technology, Kharagpur (1980) and Ph.D. from University of Washington, Seattle (1987). He was General Co-Chair of the IEEE/OSA Optical Fiber Communications (OFC) Conference 2011, Technical Program Co-Chair of OFC'2009, and Technical Program Chair of the IEEE INFOCOM'96 conference. He is Editor of Springer's Optical Networks Book Series. He has served on eight journal editorial boards, most notably IEEE/ACM Transactions on Networking and IEEE Network. In addition, he has Guest Edited Special Issues of Proceedings of the IEEE, IEEE/OSA Journal of Lightwave Technology, IEEE Journal on Selected Areas in Communications, and IEEE Communications. He has supervised 56 Ph.D.s to completion and currently mentors 17 advisees, mainly Ph.D. students. He is co-winner of Optical Networking Symposium Best Paper Awards at IEEE Globecom 2007 and 2008. He is author of the graduate-level textbook *Optical WDM Networks* (Springer, January 2006). He served a 5-year term on Board of Directors of IPLocks, a Silicon Valley startup company. He has served on Technical Advisory Board of several startup companies, including Teknovus (acquired by Broadcom). He is an IEEE Fellow.