

Network Adaptability from Disaster Disruptions and Cascading Failures¹

Biswanath Mukherjee

Distinguished Professor of Computer Science

University of California, Davis

mukherjee@cs.ucdavis.edu

Abstract

Recent disasters (e.g., Hurricane Sandy) showed that our networks (and relevant IT infrastructure) are vulnerable to large-scale disasters. Today's network protection techniques are lacking because they do not take into account the correlated cascading nature of disaster failures. While researchers in climatology, geology, and environmental sciences have been studying how to predict disasters and assess disaster risks for certain regions, networking research could exploit this information to develop novel methods to prepare networks to handle disasters with the knowledge of risky regions and to better prepare them for a predicted disaster. The events during the aftermath of a disaster should also be considered. For instance, methods to re-arrange network resources and services on a partially damaged network should be developed and new algorithms to manage post-disaster traffic deluge and to relieve the rescue operations after a disaster, with the knowledge of the post-disaster failures, should be investigated. In today's networks, the "cloud" has become an important feature because it simplifies the process of content/service sharing and distribution. Therefore, there is a pressing need to protect not only network services, but also cloud services. In this study, how to prepare network and cloud services against disasters and adapt them from disaster disruptions and cascading failures are discussed in the light of findings of a recent project supported by the Defense Threat Reduction Agency (DTRA).

1. Introduction

Recent disasters such as Hurricane Sandy demonstrate that our network infrastructures need to be better prepared for disasters, and they require intelligent and efficient recovery methods during post-disaster events. Note that telecom backbone networks employ optical mesh structures to provide highly-scalable connectivity across large distances; and these networks along with their "higher-layer" (virtual) networks (e.g., IP, MPLS, SONET, Ethernet, ATM) are integral to our economic well-being and national security because they are widely deployed in commercial and defense sectors to support many aspects of our daily life, cloud computing, battlefield surveillance/backhaul, etc. Thus, the need for survivability against disasters is acute, given the scale and criticality of these networks.

Techniques exist (and are implemented in operational networks) to provide fast protection at the optical and other layers, but they are optimized for limited faults without addressing the extent of disasters. Typically, failures caused by disasters are correlated, and cascading. From the network point of view, a disaster may cause a set of multiple correlated cascading failures. The initial shock, when a disaster occurs in a region, typically causes multiple network element failures at the physical (optical) layer. But after the initial failure, some correlated incidents might cause other failures at the optical layer (which we call horizontal correlated cascading failures). The optical layer provides lightpaths (i.e., optical circuits) to form connectivity between nodes (switches, routers, etc.) and datacenters of the upper layers (virtual networks, IP, SONET, etc.), so lack of restoration of lightpaths can cause failures on upper layers and data loss (which we call vertical correlated cascading failures). Figure 1 shows an example of a disaster inducing horizontal and vertical correlated cascading failures.

Figure 2 shows the timeline before and after a disaster with an example of average offered and requested bandwidth utilization; and it summarizes research challenges that we address in this study. While proactive (pre-disaster) techniques tend to minimize loss in case of a disaster by analyzing disaster regions, they usually overprovision the network, i.e., by exploiting the excess capacity in a network (discussed below). During a disaster, some businesses supported by telecom backbone networks may be

¹ This project is supported by the Defense Threat Reduction Agency (DTRA); PI: Biswanath Mukherjee, UC Davis.

temporarily closed, and some of the network resources may suffer major failures which may decrease both offered and requested bandwidth. At this point, immediate relief techniques can provide connectivity for crucial services and avoid the regions with high risk of correlated cascading failures. Typically, there will be many inquires to/from the disaster zone, and this will cause a traffic flood which may cause blocking or congestion of services required for rescue operations. Novel traffic deluge management techniques, which differentiate urgent and delay-tolerant services, can provide connectivity for urgent services while delay-tolerant services may be redirected to a temporary facility (e.g., message center). After the post-disaster period, when network equipment is recovered, intelligent relief techniques can restore services while increasing the offered bandwidth gradually.

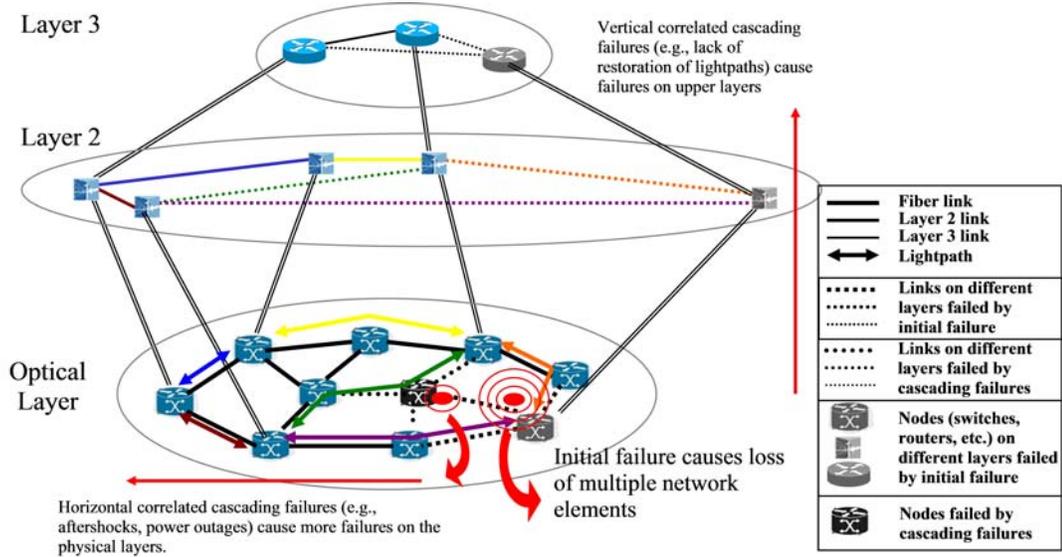


Figure 1 . Horizontal and vertical correlated cascading failures caused by a disaster.

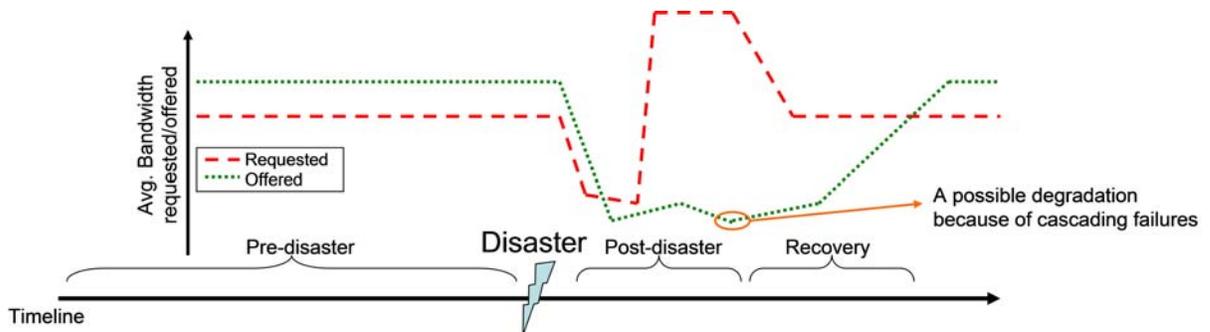


Figure 2 . Requested and offered bandwidth before and after a disaster.

2. Normal Preparedness

Network operators should proactively take necessary actions to minimize network disruptions and data loss in case of a disaster. Several studies consider regional/correlated failures caused by disasters (e.g., [1-3]) and aim to find the vulnerable parts of the network to disasters for analysis and/or design purposes. Knowledge of possible disaster zones (i.e., risk information) would help to utilize network resources and disseminate data accordingly. For instance, seismic hazard maps would be useful to determine vulnerable parts of the networks and to define the risk of traversing connections through these vulnerable parts in case of an earthquake. Similar methods can better prepare our military networks against their threats, e.g., weapons of mass destruction (WMD). By exploiting this information, we define a risk parameter which captures possible disaster zones, probability of disaster occurrence, and loss after a disaster; and we have developed a provisioning approach (called risk-aware provisioning) which minimizes the risk and reduces the loss in case of a disaster [4]. Figure 3(a) shows an example of US-wide topology with physical locations of fiber optic cables. This topology can be matched with seismic hazard maps (Fig. 3(b)) to find the vulnerable regions of the network to earthquakes where, for each

seismic level, there would be a different probability of damage. Same kind of matching can be done with tornado-activity map to obtain tornado zones (Fig. 3(c)).

The risk-aware approach encourages the network operator to choose less-risky regions during provisioning (i.e., *normal preparedness*). Figure 4 shows how risk-aware provisioning uses links which traverse less-risky regions more than others (the thicknesses of the links are proportional to the resource consumption on the links) for earthquake case where disaster zones are obtained by matching a typical nation-wide physical network topology with US seismic hazard map. Note that each zone has different probability of disaster occurrence.

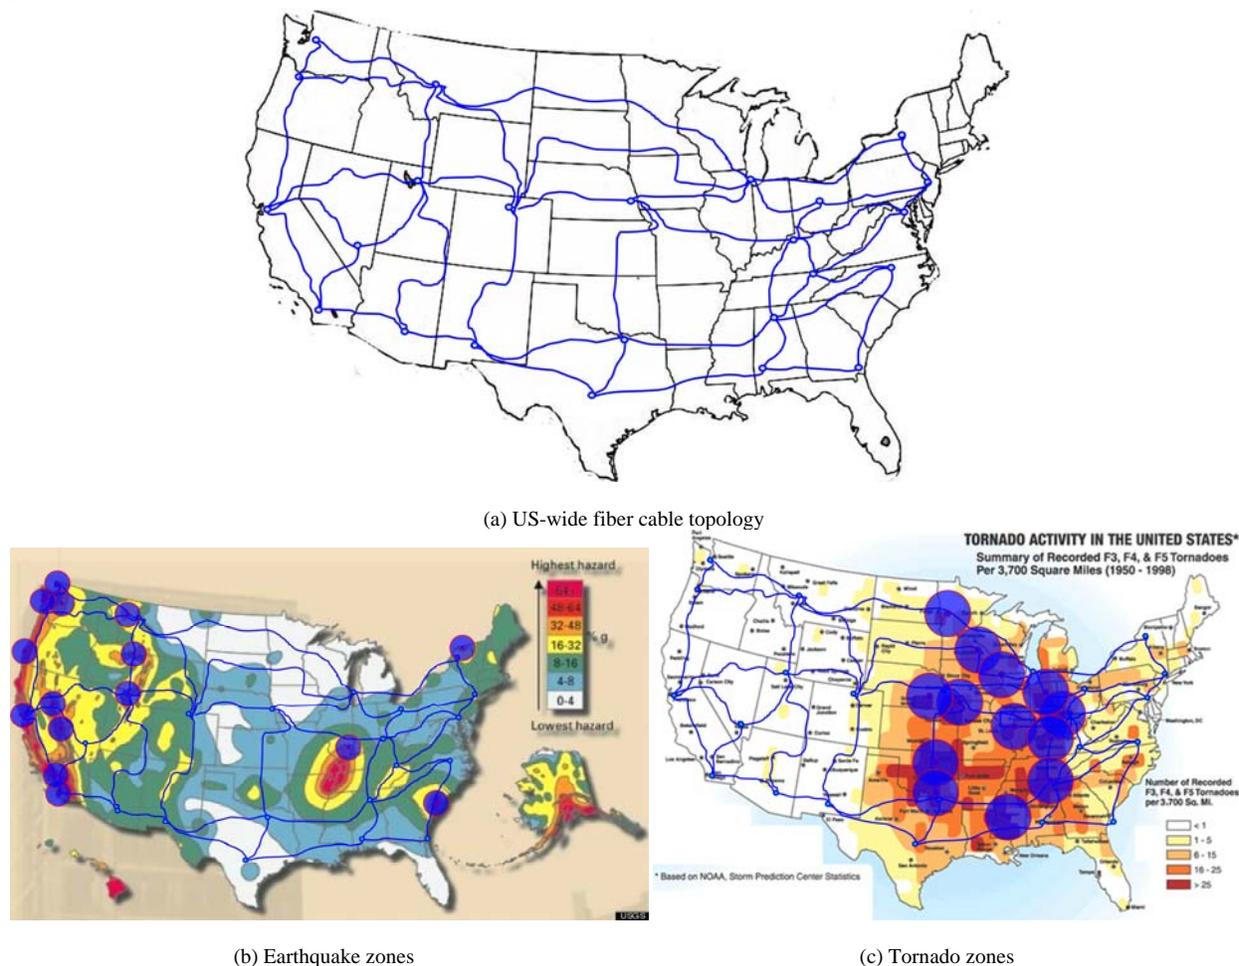


Figure 3 . (a) US-wide topology matched with (b) seismic-hazard map and (c) tornado-activity map to obtain disaster zones.

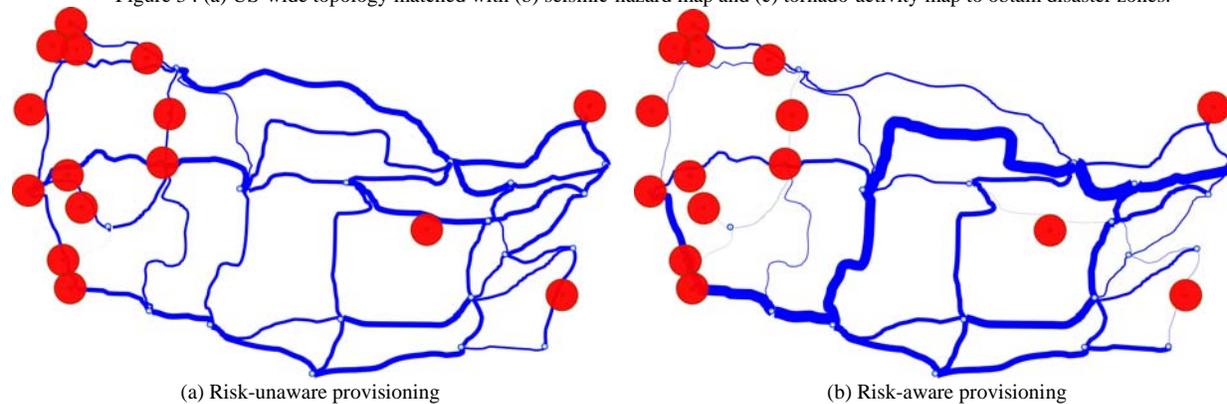


Figure 4 . Comparison between (a) risk-unaware and (b) risk-aware provisioning (the thicknesses of the links are proportional to their bandwidth utilization) shows higher utilization of links in less-risky regions in risk-aware provisioning.

Risk-aware provisioning requires more resources than risk-unaware provisioning. However, backbone networks usually have some unused capacity, called *excess capacity* (EC), to accommodate traffic fluctuations and to avoid capacity exhaustion. EC can be exploited to provide better protection against disasters (e.g., avoid disaster zones, whenever possible, by using the excess capacity of the links traversing safe zones), when there is a large amount of EC. When EC becomes low, traversing less-risky regions could be an acceptable risk to lower the resource consumption and to avoid capacity exhaustion. We studied how to exploit excess capacity to improve network robustness in several works [5-8].

Note that, while traditional approaches focus on protecting links and nodes (routers, switches, etc.) to provide “network connectivity”, the shifting paradigm towards cloud computing/storage require that we protect the data (or content). In today’s networks, more than 90% of traffic is due to content dissemination [9]. The advent of cloud computing has simplified the process of content sharing and distribution. In a cloud, a pool of configurable resources (e.g., content, services) is shared among multiple users with on-demand access [10]. So, protecting accessibility of content against any disaster is a fundamental problem. Therefore, we have also developed techniques for protection of datacenter networks (where clients request contents from datacenters) against disasters. Contents should be placed at datacenters so that they are accessible even after a disaster [11]. Figure 5 shows an example of data replication for disaster survivability, where there are two disaster zones (DZ1 and DZ2) and three datacenter locations at nodes C, D, and E. A specific data is replicated at datacenters C and E for disaster protection so that, for a request for this data, a primary path to datacenter C and a backup path to datacenter E can be provided. For instance, a request from node A, might have primary path (A-B-C) to datacenter C and a backup path (A-F-E) to a secondary datacenter at node E. If a disaster occurs (either in DZ1 or DZ2), the content will still be reachable.

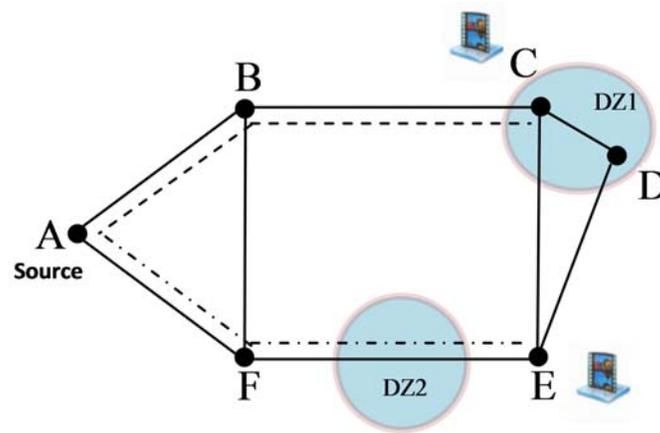


Figure 5 . Data/content protection against disasters.

Thus, we have defined a new concept, called *Content Connectivity*, which is the reachability of every content from any point of a network. It might be the case that a failure disconnects the network, but if a content is replicated in all disconnected portions of the network, the required service can still be provided. Maintaining content connectivity requires efficient content placement by exploiting the a priori knowledge of potential vulnerable locations in the network. Our preliminary investigation [12] has found the following: a) since maintaining network connectivity may not always be possible after failures, content connectivity can help us to provide / continue services in such scenarios; and b) ensuring content connectivity requires less network resources than network connectivity.

3. Enhanced (Better) Preparedness

If a disaster is predicted through scientific measurements and observations (e.g., the possible path and estimated time of arrival to main land of a hurricane occurring in the ocean can be known days in advance), the network can be better prepared by re-allocation of network resources and re-dissemination of data, and possibly by relocation of hardware resources also. In such scenarios, reactive measures can be taken to adapt to the changing risk level of attacks and to enhance the protection of the network. Risk-aware provisioning can be used using the changing risk inputs. Data/content protection can also be provided by replication data/content from a datacenter under high risk in case of an upcoming disaster to

a safe location. These pre-disaster actions, namely re-allocating network resources and replicating data, should be done considering cascading effects of the disaster (e.g., power outage). For instance, Fig. 6 shows the projection of Hurricane Sandy's path on October 29, 2012, with red lines and network elements that might be affected by the hurricane and/or its correlated cascading failures with green lines and circles. Any data in any datacenter which is located in these green circles should be replicated to safe nodes (blue circles) to avoid the loss of data due to failures by the hurricane. The connections traversing the links and nodes shown in green should also be reprovisioned by exploiting the excess capacity on links shown by blue lines.

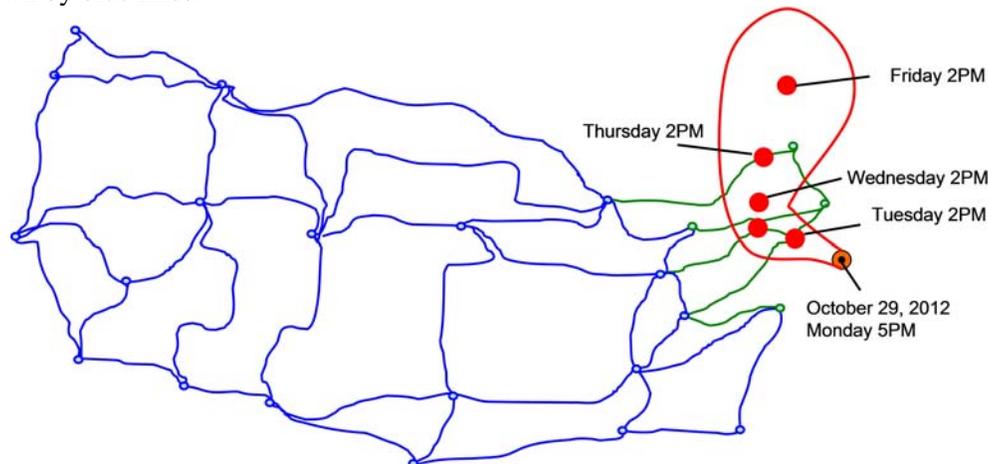


Figure 6 . Path of Hurricane Sandy projected on October 29, 2012, 5 PM (red) and network elements that might be affected by the hurricane and/or its correlated cascading failures (green).

4. Post-Disaster Events

After an attack, some traffic may be disrupted. This disrupted traffic can be reprovisioned in the network by using the excess capacity (EC) in the undamaged parts of the network. During the reprovisioning, cascading failures should also be taken into consideration, i.e., when a disaster occurs, a set of network elements fails simultaneously in the first phase; and, in a second phase, other failures in different parts of the physical network (and/or in the upper network layers) may occur due to the failure in the first phase. An important feature of these 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 network connectivity in view of such possible cascading failures. In [13], we have developed methods for reprovisioning of resources to recover at least the most crucial services in the immediate aftermath of the attack. We also developed multipath provisioning (i.e., a connection's full bandwidth is provided through multiple paths) approaches which may guarantee *degraded service* rather than full service where the offered bandwidth is less than requested bandwidth [13]. The classification of traffic may help us to decide which service (degraded vs. full) would be more appropriate for which kind of traffic.

Once correlated cascading failures are settled down (i.e., there is no expectation of more correlated cascading failures), the information on the network damage is more certain. The data on network recovery (e.g., through FCC's Disaster Information Reporting System (DIRS)) can give information on the network status, and this information can be exploited to better meet the service needs and current network state. A network recovery plan is prepared which schedules the repair time of network elements. With the knowledge of this plan, a series of intelligent reprovisioning actions (executed after recovery of each network element) could be developed with the option of degraded services to transition the network to a fully operational state. While the network elements are recovered, the network operator may aim to guarantee partial bandwidth which becomes 100% when the network is fully recovered.

5. Summary

Below is a summary of key concepts in our study.

- Exploiting excess capacity to improve network resilience
- Determination of disaster zones with risk (hazard) maps
- Risk-aware provisioning (for *normal preparedness*)

- Data replication
- Content connectivity
- Reprovisioning for *better preparedness* and *post-disaster events*
- Multipath provisioning for *degraded services* (reduced level of services vs. no services at all)

6. Conclusion

Methods to prepare the network for possible disasters, to better prepare for upcoming disasters, to provide some minimal level of services after a disaster to support critical operations, and to recover services, while the network is recovering, can significantly improve network resilience/robustness against disasters.

7. Acknowledgement

Many thanks to Mr. Ferhat Dikbiyik, PhD student at UC Davis, for his help in preparing this report.

8. References

- [1] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of fiber infrastructure to disasters," *IEEE/ACM Transactions on Networking*, vol. 19, pp. 1610-1623, Dec. 2011.
- [2] P. K. Agarwal, A. Efrat, S. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "The resilience of WDM networks to probabilistic geographical failures," in *Proc. IEEE INFOCOM*, Shanghai, China, Apr. 2011.
- [3] M. Rahnamay-Naeini, J. Pezoa, G. Azar, N. Ghani, and M. Hayat, "Modeling stochastic correlated network failures and assessing their effects on reliability," in *Proc. IEEE ICC*, Kyoto, Japan, June 2010.
- [4] F. Dikbiyik, M. De Leenheer, A. S. Reaz, and B. Mukherjee, "Minimizing disaster risk in optical telecom networks," *IEEE/Optical Fiber Communications (OFC) Conference*, Los Angeles, CA, March 2012.
- [5] F. Dikbiyik, L. Sahasrabuddhe, M. Tornatore, and B. Mukherjee, "Exploiting excess capacity to improve robustness of WDM mesh networks," *IEEE/ACM Transactions on Networking*, vol. 20, pp. 114-124, 2012.
- [6] F. Dikbiyik, M. Tornatore, L. Sahasrabuddhe, and B. Mukherjee, "Exploiting excess capacity Part II: Differentiated services under traffic growth," *IEEE/ACM Transactions on Networking* [in review], 2012.
- [7] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Exploiting excess capacity for survivable traffic grooming in optical WDM backbone networks," *IEEE Globecom*, Houston, TX, Dec. 2011.
- [8] F. Dikbiyik, M. Tornatore, and B. Mukherjee, "Optimal relocation of excess capacity in optical WDM networks," *IEEE Globecom*, Anaheim, CA, Dec. 2012.
- [9] Cisco Visual Networking Index: Forecast and Methodology, 2011-2016. White Paper, May 2012.
- [10] National Institute of Standards and Technology, "The NIST definition of cloud computing," <http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>
- [11] M. F. Habib, M. Tornatore, M. De Leenheer, F. Dikbiyik, and B. Mukherjee, "Design of disaster-resilient optical datacenter networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 30, no. 16, pp.2563-2573, 2012.
- [12] M. F. Habib, M. Tornatore, and B. Mukherjee, "Fault-tolerant virtual network mapping to provide content connectivity in optical networks," *IEEE/Optical Fiber Communications (OFC) Conference*, Anaheim, CA, Mar. 2013.
- [13] S. Huang, M. Xia, C. Martel, and B. Mukherjee, "A multistate multipath provisioning scheme for differentiated failures in telecom mesh networks," *IEEE/OSA Journal of Lightwave Technology*, vol. 28, no. 11, pp. 1585-1596, June 2010.