Network Adaptability to Disaster Disruptions by Exploiting Degraded-Service Tolerance

S. Sedef Savas, M. Farhan Habib, Massimo Tornatore, Ferhat Dikbiyik, and Biswanath Mukherjee

Introduction

Telecom networks are exposed to many threats such as malicious attacks, equipment failures, human errors (e.g., misconfigurations), and large-scale disasters, both human-made (e.g., due to weapons of mass destruction, WMD, attacks) and natural. Disasters represent a challenging threat for our networks as they affect large geographical areas, and can trigger correlated and/or cascading failures of multiple network elements, resulting in huge data loss and disruptions in network connectivity. The 2011 Japan Earthquake/Tsunami and 2012 Hurricane Sandy [1] are two recent disasters that have deprived people of essential network services and severely hindered rescue operations for weeks. Thus, disaster-aware provisioning schemes (e.g., routing around risky disaster areas) have been proposed. The emergence of bandwidth-hungry applications has led to rapid growth in the volume of the data traffic in our networks, making cost-effective survivability methods against disasters even more crucial.

Today's networks support diverse services, from cloud computing and video streaming to traditional ones (HTTP, VoIP, etc.). Services have different requirements (e.g., delay/latency tolerance and bandwidth) and characteristics (e.g., importance and revenue generation). With such heterogeneity, using the same protection policies for all services can result in suboptimal solutions. Thus, we consider the different tolerances of services (i.e., degraded-service tolerance) to develop fault-tolerant (survivable) methods that can sustain an acceptable level of service even when disasters occur. Some services are sensitive to the amount of capacity provided, while others (e.g., video streaming or file transfers) can operate with reduced bandwidth. *Degraded service* refers to a reduced amount of resource allocation for a service vs. its normal operational requirement. Thus, even with degraded service, some services can still achieve lower but acceptable quality.

The degraded-service concept has been investigated for survivable service provisioning schemes against large-scale disasters [2]. Providing 100 percent protection against disasters (by routing them via primary and backup paths) would require massive and economically unsustainable bandwidth overprovisioning, as disasters are difficult to predict and statistically rare, and may create large-scale failures. Some researchers have shown that providing protection for a portion of the requested bandwidth, or *partial protection*, can alleviate the extensive resource usage of full protection schemes [2]. The fraction of the requested bandwidth that will be guaranteed, even under failures, is determined by the degraded-service tolerance of services, generally stated in a service level agreement (SLA).

A limitation of prior studies exploiting degraded service is that they do not adapt their resource allocation based on the network state, which leads to suboptimal network utilization in the presence of disasters. They exploit degraded-

S. Sedef Savas, M. Farhan Habib, and

Biswanath Mukherjee are with the University of California, Davis.

Massimo Tornatore is with the University of California, Davis, and Politecnico di Milano.

Ferhat Dikbiyik is with Sakarya University.

This work has been sup-ported by the Defense Threat Reduction Agen-cy (DTRA) Grant HDTRA1-10-1-0011.

A preliminary version of this work was presented at OFC 2014. This extended work includes detailed analysis specific to disaster scenarios.

service tolerance of connections only during admission to provide them partial protection so that when a disaster occurs, the connection will continue to operate at its minimum required service level [1, 3]; hence, when a disaster occurs, already accepted (and unaffected) connections are not considered by these studies for network resource optimization.

Networks may experience *resource crunch*, that is, an undesired reduction of network capacity due to disaster failures. Since disasters are rare, resource crunch is an unusual situation, so different measures should be taken to get through disasters with minimal damage. There are two major contributions of this article besides exploiting degraded-service tolerance as partial protection schemes do. First, our disaster-aware provisioning considers the decision process in the aftermath of disasters separate from the provisioning step, as it also considers the degraded-service tolerance of incoming connections. Second, during disasters, we allow provisioned connections to also be part of the resource allocation optimization. To offer the affected connections acceptable levels of service, unaffected connections can be degraded, rerouted, or even halted depending on their importance or profitability characteristics.

In this article, we first present an overview of existing disaster-aware service provisioning schemes that exploit degraded-service tolerance. Then we describe our disaster-aware adaptive resource allocation scheme. Finally, we illustrate and quantitatively compare these mechanisms using a case study in wavelength-division multiplexing (WDM) mesh networks. Note that our techniques are applicable to general mesh networks as well.

Overview of Disaster-Aware Provisioning Schemes

Disaster repercussions (e.g., disconnections, data loss, and service disruptions) can be minimized using protection schemes (preconfigured before disasters) and restoration schemes (reactive, after disasters).

Disaster-Aware Provisioning Schemes

Some studies propose survivable provisioning to proactively alleviate the impact of disasters. They determine possible disaster zones in the network, such as *risk (hazard) maps* to highlight its vulnerable regions using interdisciplinary contributions from climatology, geology, and so on. (Figure 1a shows risky zones of the United States with a heat map against some natural disasters [4]). By utilizing risk maps, we can estimate the probability of occurrence of a disaster and probability of a network device getting damaged by this disaster. These two parameters give us the risk levels of disaster zones. Figure 1b shows a seismic-hazard map of the United States with its risk levels. Locations of high-risk earthquake zones (Fig. 1c), which have different probabilities of failures, can be determined by matching a network topology with the seismic hazard map. Using these maps, network planners can develop systems that select less risky regions for routing connections; hence, the expected loss will be minimized and the network becomes better prepared to handle a disaster.

The set of links or nodes that are vulnerable to a common failure (e.g., a disaster) can be represented as a shared risk group (SRG) [1]. The most prevalent protection strategy against disasters is to route connections over disaster-zone-disjoint (i.e., SRG-disjoint) primary and backup paths (or using multiple primary paths, e.g., multi-path provisioning). However, fully protecting primary paths with backups requires extensive resource usage, especially for multiple failures (as in disasters).

Some services may accept a reduced level of bandwidth during failures, depending on their characteristics. For services that can tolerate reduced bandwidth, network operators may offer partial protection, possibly at lower cost. The partial protection guarantee is determined by the connection's degraded-service tolerance.

Multipath routing (i.e., multiplexing a connection over multiple paths) is another scheme for providing partial protection [6]. For a multi-path-routed service, even if some paths are down or overloaded, other paths may provide the required degraded service. Thus, some SLAs for partial protection can be satisfied without any redundant resource allocation.

The shifting paradigm toward cloud computing is creating new opportunities for optimizing disaster-aware network design. Contents in cloud systems can be replicated at multiple servers/data centers from which users can be served. New service models are introduced such as *anycasting* (providing service from any of the data centers that host the requested service) and *manycasting* (providing service from a subset of the data centers). These models can be exploited, for example, for file transfer and media streaming to enhance the resilience of cloud services. Resilience against destination node failures is very crucial due to cloud services and data centers hosting content. These schemes (shown in Fig. 2) are resilient against destination-node failures (since the paths connect to disaster-zone-disjoint destination nodes, so the service will not be disrupted if such a node fails due to a disaster).

Figure 2 shows examples of a degraded-service-aware single-primary path using anycasting and multipath provisioning using manycasting [2], both of which guarantee a minimum tolerable bandwidth of 60 percent of the required bandwidth under normal operation even under disaster scenarios. Backup-path protection using anycasting is shown in Fig. 2a, where a 1 Gb/s connection is partially protected by a backup path with 0.6 Gb/s capacity. Reference [7] used inverse multiplexing over multiple paths (the least risky ones) to provision bandwidth for services distributed over multiple servers/data centers with manycasting (Fig. 2b). Also, it ensures degraded service (vs. no service at all) after a failure without using extra resources since it uses multipath routing. For instance, in Fig. 2b, during normal operation, the customer receives 1 Gb/s service, which is multiplexed over three paths (one with 0.4 Gb/s and two with 0.3 Gb/s) destined to different data centers; and any pre-

For a multi-path-routed service, even if some paths are down or overloaded, other paths may pro-vide the required degraded service. Thus, some SLAs for partial protection can be satisfied without any redundant resource allocation.

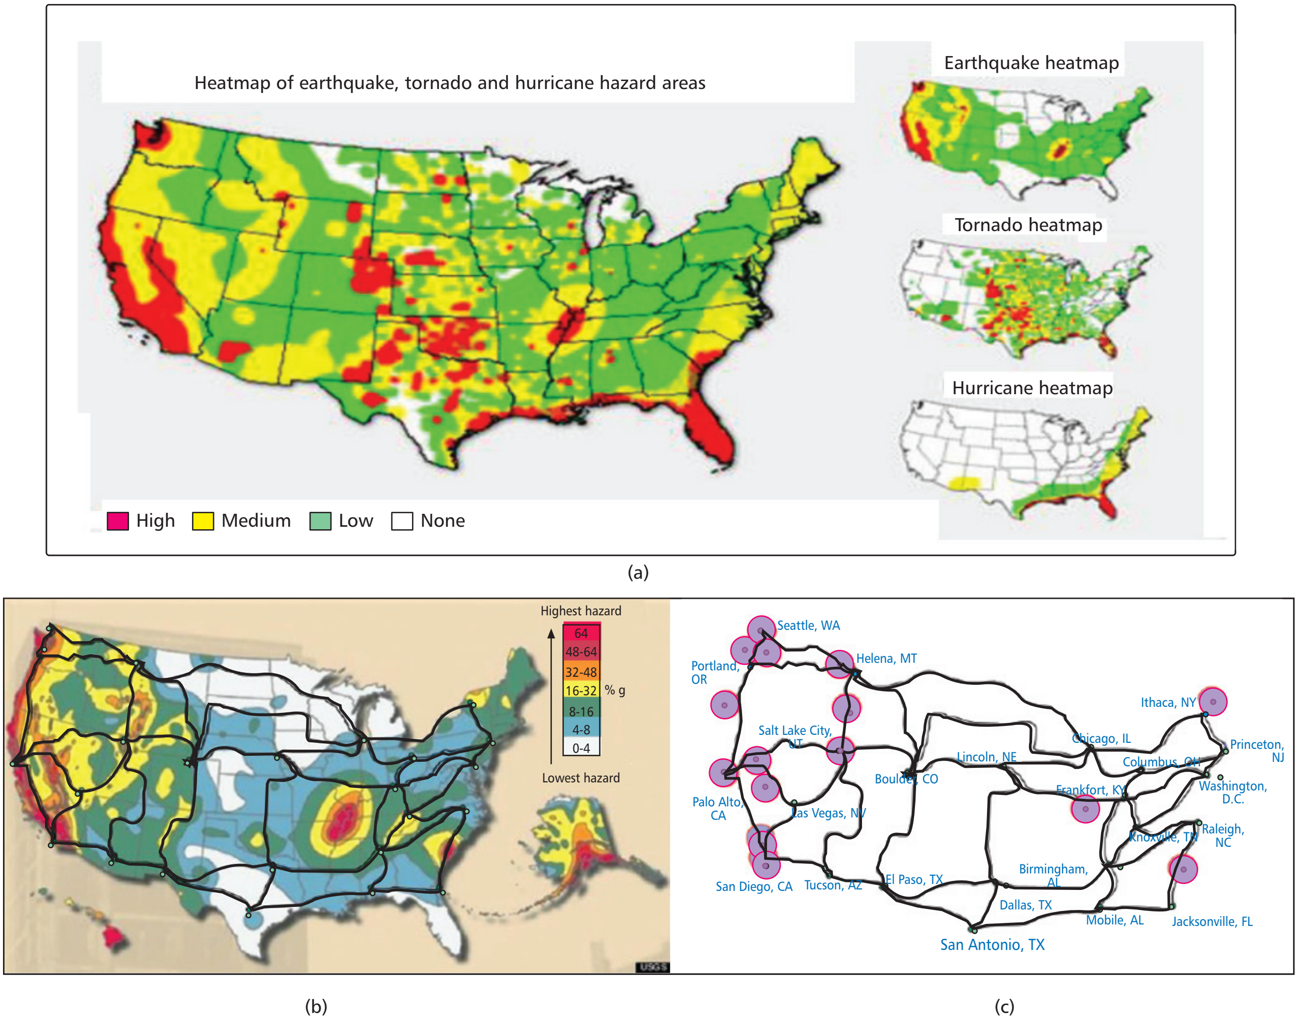

Figure 1. Exploitation of hazard maps to determine disaster zones: a) natural disaster risk map (Credit: U.S. Geological Survey); b) earthquake risk map; c) earthquake zones (shown in circles).

dicted disaster in the figure affects only one path, so the guaranteed bandwidth is at least 0.6 Gb/s.

Restoration/Reprovisioning Schemes

Despite the above measures, it is not always possible to avoid all disaster zones and provide protection to all services for all disaster scenarios; moreover, unforeseen attacks and disasters may occur. Therefore, taking actions in the aftermath of a disaster should also be considered. Reprovisioning is a reactive approach where network resources are re-allocated for existing connections. If an unpredicted disaster occurs, restoration schemes can be used to preserve the targeted quality of service (QoS) level or to ensure graceful degradation using the remaining resources in the undamaged parts of the network. Since usually only some parts of the network are damaged, the unused capacity available in the network's remaining parts can be used to reprovision the disrupted connections. Note that during restoration, secondary failures such as aftershocks following an earthquake should be considered as they are predictable with good accuracy after the primary failure. Reprovisioning schemes are robust against different types of failures as they adapt the network according to its current state but they do not give restoration guarantee for the disrupted services as provisioning schemes do. Full-service restoration schemes are as costly as full protection schemes, and reduce restoration chances due to high bandwidth requirement in a limited-resource environment. Partial-bandwidth restoration [1] (exploiting degraded-service tolerance) may be a good option as it requires fewer resources and hence increases restoration probability. Although resource consumption is low, restoration schemes are not favored for live traffic (e.g., rerouting primary paths), as they cause service disruptions due to reprovisioning of network resources.

To handle the unusual resource crunch problem caused by disaster failures, besides exploiting degraded-service tolerance just for partial-protection purposes as the above works do, we propose to exploit it further to perform degraded-service-aware resource re/allocation in case of disasters. Our solution is applicable for

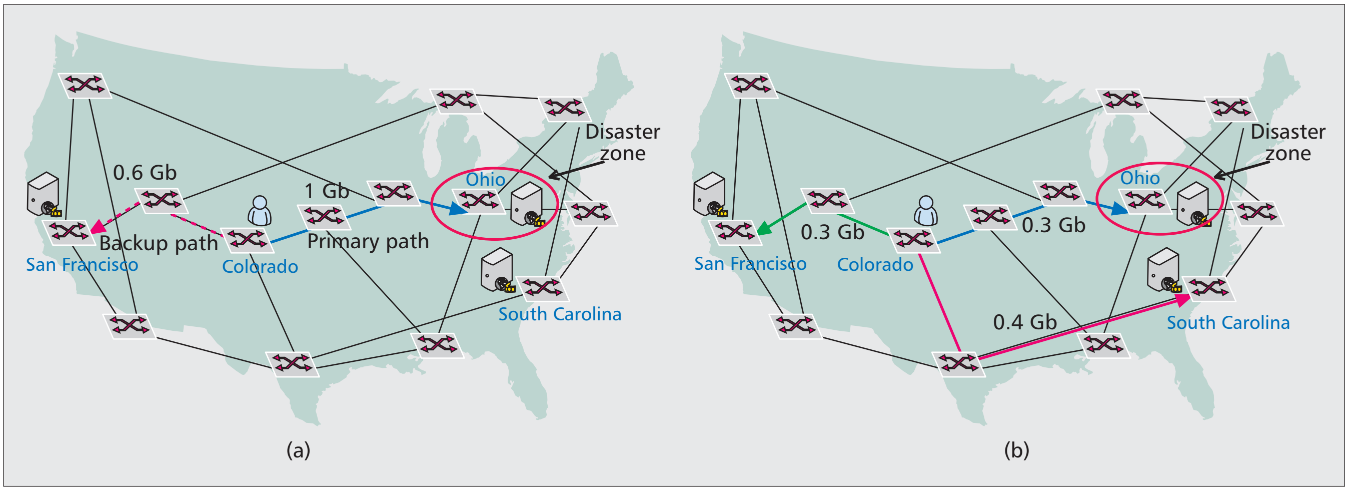

Figure 2. Disaster-aware provisioning schemes in cloud networks: a) anycasting with partial protection; b) manycasting (multipath) with partial protection.

all networks with service differentiation functionality, such as optical WDM networks, IntServ, DiffServ over IP networks, multiprotocol label switching (MPLS) on layer 2 networks, and software defined networking (SDN)-enabled networks [8].

The Proposed Approach: Degradation-Aware Adaptive Resource Allocation

The above solutions perform static resource allocation; for example, once admitted, no further actions are taken on a connection unless there is a fault that affects it. Such non-adaptive approaches result in suboptimal solutions.

We propose and study the characteristics of a novel resource allocation framework (service provisioning and reprovisioning) to enable the network to be robust to disasters by adaptively responding to changes in the network state. This is a new approach to (re)distribute resources among existing connections. Our work can coexist with prior disaster-aware provisioning schemes. We aim to minimize blocking rate (unadmitted connection ratio over all requests), dropping rate (losing a connection), and disruption rate (rate of reprovisioning of connections) during disasters. Also, we aim to provide the best service possible to connections with remaining network resources by rearranging the resource allocation if needed.

We exploit the degraded-service concept to combat disasters as follows:

1. Accept degraded services during call admission to increase service acceptance rate, and if necessary, degrade existing connections
2. Degrade existing connections to reduce dropping rate
3. Apply an upgrade process to restore degraded connections to full service whenever possible

Figure 3 shows our proposal in three steps: *provisioning*, *recovery*, and *upgrade*. The *provisioning step* is applied only when a disaster occurs; otherwise, disaster-aware provisioning with full service is used. Specifically, we deal with the limitation of available resources caused by a disaster during the admission process by accepting connections with degraded service and/or adapting the network by degrading existing connections to release some resources for incoming connections. *The recovery step* kicks in when some connections get disrupted due to a disaster failure; then we try to release some capacity to avoid them being dropped. Finally, the *upgrade step* is triggered when new resources are reactivated after the repair of some network elements, to provide the best service to connections within available network resources.

Our approach may increase management complexity under faulty network conditions. This increase might be considered acceptable as it represents the cost to reduce the connection dropping rate and provide the best possible service with the residual resources while not violating SLAs. Recent progress in automatic control/management solutions (e.g., SDN) can provide technological support to harness this increase in management complexity. SDN eases the network control, monitoring, and management, thus enabling more dynamic approaches such as our proposed solution. Modifying end-to-end service levels of existing connections that are not affected by the disaster requires a multi-tenant, multilayer management scheme, which will be enabled by SDN. This multi-layer optimization together with a holistic view of the network may increase the chances of discovery of resources, which is crucial in a stressed scenario such as a disaster. This feature also improves disaster recovery time as post-disaster convergence of the network may take a large amount of time (seconds) if addressed only at a higher layer of the network such as the IP layer.

With evolving cloud and video applications, traffic patterns are becoming more dynamic. Therefore, even without disasters, the overload from sudden bursts of traffic can lead to resource crunch. To cost-effectively handle growing/bursty traffic, networks may need to be run at high uti-

The decision of what to degrade and how much to degrade will depend on the opti-mization objective. This decision affects the number of con-nections that are reprovisioned, and the degradation pro-cess can be per-formed by allowing either global reprovi-sioning or essential (local) reprovisioning.

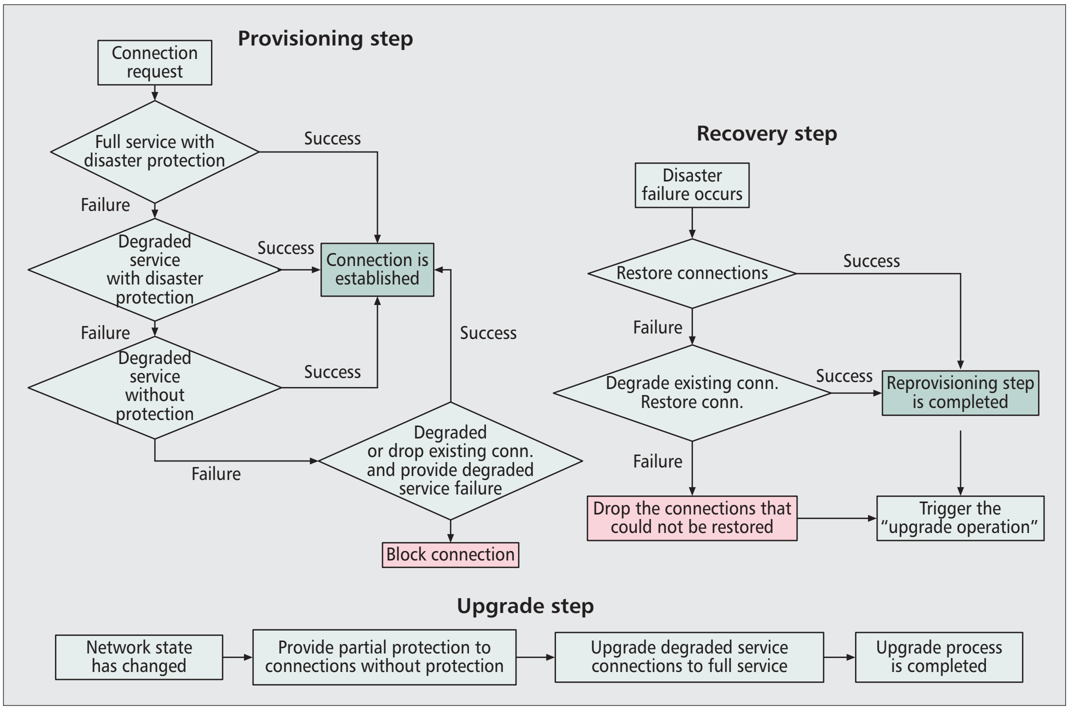
Provisioning step
Connection request
Full service with disaster protection
Success
Failure
Degraded service with disaster protection
Success
Connection is established
Failure
Degraded service without protection
Success
Success
Failure
Degraded or drop existing conn. and provide degraded service failure
Block connection
Recovery step
Disaster failure occurs
Restore connections
Success
Failure
Degrade existing conn. Restore conn.
Success
Reprovisioning step is completed
Failure
Drop the connections that could not be restored
Trigger the "upgrade operation"
Upgrade step
Network state has changed
Provide partial protection to connections without protection
Upgrade degraded service connections to full service
Upgrade process is completed

Figure 3. Operational steps.

lization [9]. These highly utilized networks are at risk of resource crunch, so our adaptive resource allocation scheme can apply to them as well.

What and How to Degrade

Degraded service is exploited to achieve effective traffic engineering with high network utilization in disasters. The decision of what to degrade and how much to degrade will depend on the optimization objective. This decision affects the number of connections that are reprovisioned, and the degradation process can be performed by allowing either global reprovisioning or essential (local) reprovisioning. For instance, as a result of a global reprovisioning approach (no restriction on the number of reprovisioned connections), better resource utilization may result in better network performance in terms of recovered connections and blocking probability. However, a high number of reprovisioning actions may cause excessive disruptions, which in turn could cause data loss due to the switching time required to reprovision the connections. There is a trade-off between better network utilization and providing uninterrupted service. Some optimization objectives to determine the level of flexibility between these two are as follows (and note that only the first objective below is used in our current study):

- *Maximize profit*: Prioritizing the low-revenue-generating services in the degradation process can maximize profit by degrading the minimum number of connections that generate low income.
- *Minimize penalties for SLA violations*: Some penalty can be applied to a service when the service level goes below its agreed QoS in SLA. There is a minimum service level that a customer can accept, and this must be satisfied.
- *Customer satisfaction*: Customers can tolerate degraded service under unusual circumstances. For better customer satisfaction, the full bandwidth requested must be provided. To minimize dissatisfaction from degraded service, the fewest connections should be degraded whenever possible.
- *Load balancing*: If some resources need to be released, while selecting connections to be degraded we need to consider their physical locations as well. As the number of degraded-service connections increases in a region, the chance of upgrading them to full service will be less than a scheme where connections to be degraded are selected by region (i.e., regions with a high density of degraded services will not be selected).

Also, the optimization objective can be a mix of some of the above objectives; this is a topic for further investigation.

Other Opportunities to Exploit Service Differentiation Characteristics

Services have different characteristics (e.g., revenue generation, importance) and different requirements (e.g., delay tolerance, availability, and bandwidth) [10]. These metrics should be exploited in degraded-service provisioning to determine which services should have reduced capacity (and by how much) under resource crunch. Below, we list some possible service differentiations:

- *Scalability*: Scalability of services is an important factor for the degradation process. Services may not operate at every bit rate between a requested bit rate and the minimum tolerable bit rate. For instance, if we assume that a streaming video is encoded at two different qualities (high and low),

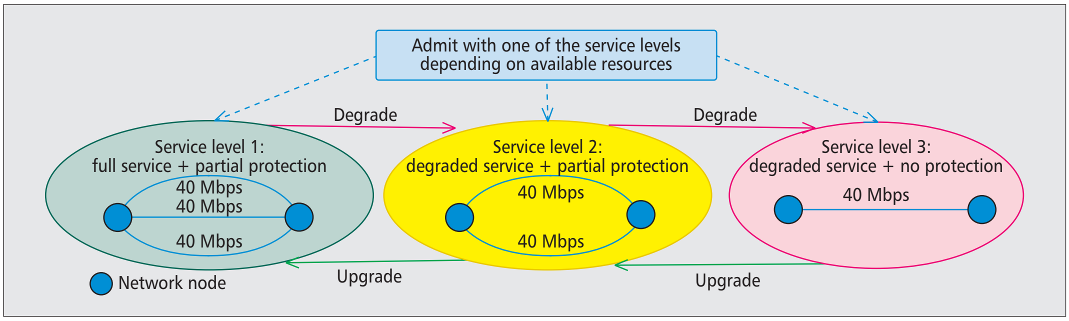
Admit with one of the service levels depending on available resources
Degrade
Degrade
Service level 1:
full service + partial protection
40 Mbps
40 Mbps
40 Mbps
Service level 2:
degraded service + partial protection
40 Mbps
40 Mbps
Service level 3:
degraded service + no protection
40 Mbps
Network node
Upgrade
Upgrade

Figure 4. Possible service levels provided during admission. (The service level of a connection may be degraded during provisioning and recovery steps to free some resources for incoming or disrupted connections. The service level is only upgraded during the upgrade step.)

Without causing any disruption, by tearing down (degrading) or adding some paths (upgrading) of/to a connection, we can adapt the network according to its cur-rent state. Our proposed solution is not restricted by multi-path provisioning and can be applied to any existing provi-sioning scheme.

if the bit rate allocated to this service is between low and high bit rates, low-quality video will be served and the remaining bit rate will be wasted. Thus, this aspect should be considered.

- *Mission-critical vs. non-mission-critical*: Some services are critical such as military services. Service availability is increasingly becoming a requirement for certain mission-critical applications. Attention must be paid for these connections to reduce their disruption rate.
- *Real-time vs. data traffic*: Traffic can be classified as:
  –Real-time (delay-sensitive, e.g., VoIP, video conferencing, web search) and
  –Data (non-real-time traffic, which is not delay-sensitive, e.g., file transfer)

Our solution can exploit these QoS differences. While we do not reprovision delay-sensitive connections as reprovisioning causes delay, non-real-time traffic can be reprovisioned to release resources that can be utilized by other connections to recover from degradations to their normal operational state.

Degraded-Service-Aware Provisioning: A Case Study

We present a case study to show how our proposed framework can be applied to enhance resilience, resource utilization, and QoS. Our study covers two differentiated service types:

- Degraded-service-tolerant and -intolerant services
- Mission-critical and regular services

A dynamic scenario is considered where traffic/service requests arrive, hold for a while, and then depart. Also, disasters randomly hit the network, and affected resources remain unavailable until they get repaired. Disaster protection is granted to connections by multiplexing the connection over SRG-disjoint paths (multipath provisioning). In case of a multi-path-routed service, even if a path is down or overloaded, the other paths may provide the required degraded service. Thus, SLAs can be satisfied without additional resource usage or re-allocation of bandwidth among competing resources. In our study, each service level provides partial protection, except degraded-service-intolerant services (which always require full bandwidth, i.e., full protection).

We consider three service levels which can be provided to a connection during its lifetime. These three service levels can be seen in Fig. 4 (service levels 1–3 in descending order of quality) for a connection request with 120 Mb/s full service and 40 Mb/s degraded service requirement:

- *Service level 1 (full service with partial protection*), which provides full bandwidth during normal operation and degraded service during failures
- *Service level 2 (degraded service with partial protection)*, which provides the requested degraded service during normal operation and partial protection after failures
- *Service level 3 (degraded service without protection)*, which provides the requested degraded service during normal operation and does not give any guarantee after failures

Without causing any disruption, by tearing down (degrading) or adding some paths (upgrading) of/to a connection, we can adapt the network according to its current state. Our proposed solution is not restricted by multipath provisioning and can be applied to any existing provisioning scheme.

We analyze the steps (provisioning, recovery, and upgrade) depicted in Fig. 3 by comparing three different approaches explained below. The *degraded-service* and *extreme-degraded-service schemes* are the proposed approaches, whereas the *full-service scheme* is the traditional approach. These schemes can work with any number of service levels. For illustration, we consider the above-mentioned three service levels.

Full-service scheme (FS): In FS, degraded service is not exploited during resource crunch; that is, no reallocation of network resources is allowed, and connections are only accepted with full service. This is the traditional admission approach where an incoming connection is denied if the network cannot provide full service. In the recovery step, disrupted connections are attempted for recovery to full service by reprovisioning. If not successful, they are dropped.

Degraded-service scheme (DS): At admission, connections under DS can be accepted with the best possible service level between the ranges of degraded service to full service (service levels 1–3). During recovery, we reprovision disrupted connections by gradually degrading the service level until it is satisfied. In this scheme, reallocation of network resources among existing connections to release resources is not allowed.

Extreme-degraded-service scheme (EDS): At both admission and reprovisioning, this scheme allows rearrangement in the network by relocating or degrading existing connections to release resources for incoming and disrupted connection after trying the steps in the degraded-service scheme. This scheme includes all steps shown in Fig. 3.

The upgrade step in Fig. 3 can be applied to all of the above schemes. The upgrade step is triggered after each network state change to provide the best possible service according to current network state. The network state changes whenever:

- A new request arrives.
- An existing connection terminates.
- A network failure occurs.
- A failed link/node is repaired.

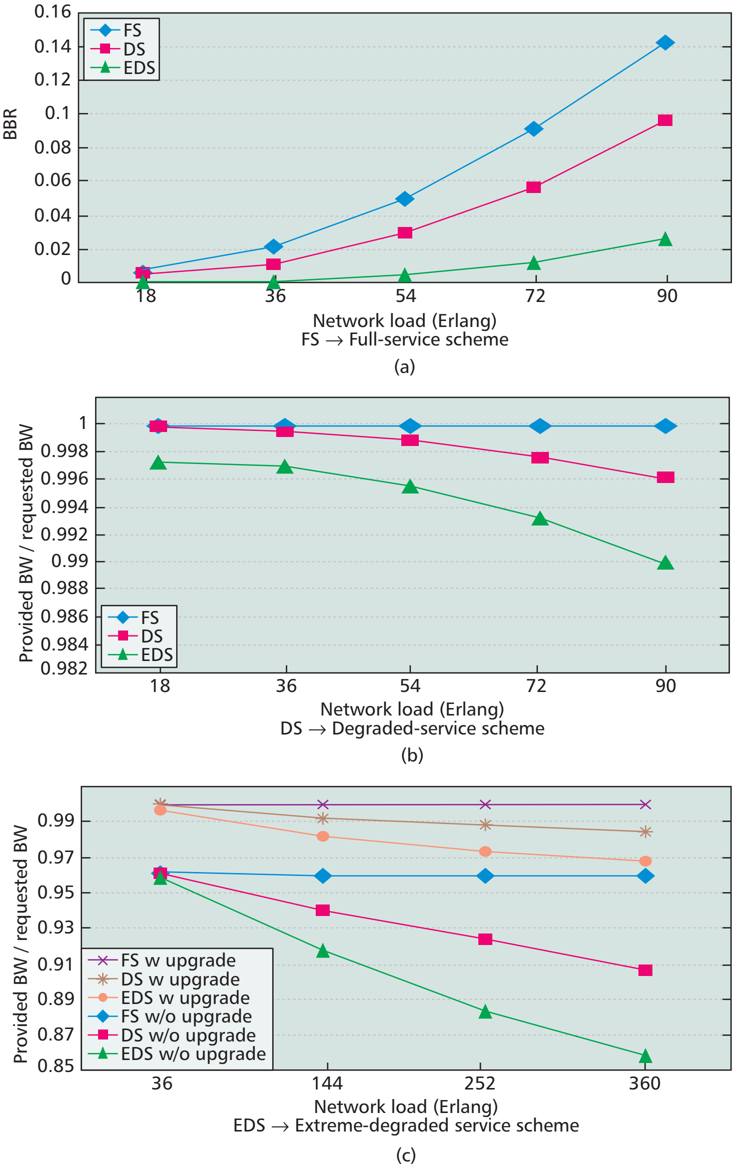

FS → Full-service scheme

(a)

DS → Degraded-service scheme

(b)

EDS → Extreme-degraded service scheme

(c)

Figure 5. Comparison of the proposed schemes (DS and EDS) vs. a non-adaptive provisioning scheme (FS): a) bandwidth blocking probability (BBR); b) average bandwidth provided; c) effect of upgrade operation.

In the upgrade step, first, partial protection is provided to connections that do not have protection. Then connections that receive degraded service are upgraded to full service if possible. Even though our degraded-service schemes increase acceptance rate (at the cost of slightly decreasing average bandwidth provided), the upgrade operation alleviates the decrease in average bandwidth, which stems from the fact that some existing connections' service levels are degraded to accommodate new ones. During the upgrade step, to reduce disruptions, we only reprovision a connection while upgrading it from service level 3 or 2 to service level 1.

Case Study: Results and Evaluation

To evaluate the proposed approaches, we simulated a realistic dynamic environment. The connection arrival process is Poisson with arrival rates 0.3, 0.6, 0.9, 1.2, and 1.5 requests/min for Figs. 5a–5b; and 0.6, 2.4, 4.2, and 6 requests/min for Fig. 5c. Connection holding time follows a negative exponential distribution with a mean of 60 min. We simulate wavelength-division-multiplexed (WDM) optical networks. Each link has 10 wavelengths, and each wavelength has 10 Gb/s of capacity. The bandwidths of the connection requests follow the realistic distribution 150M: 600M: 2.5G: 10G = 40: 30: 20: 10 [2]. Degraded-service requirements of connections are uniformly distributed between 40 and 100 percent of their requested bandwidth (on average 70 percent), and connections that need the whole requested capacity at all times are intolerant to degraded service. We simulated 100,000 connection requests on a sample 24-node U.S. topology with the shown disaster scenario, with occurrence rate 2×10^{-5} disasters/min (Fig. 1c). In this study, connections arrive with the information of their connection type (either mission-critical, which is 15 percent of all connections, or regular) and degraded-service tolerance. We prioritize mission-critical connections at every decision step (e.g., we upgrade mission-critical connections first).

Figure 5a compares bandwidth-blocking ratio (BBR) (i.e., the amount of rejected bandwidth over the total requested bandwidth) of the FS scheme with our degraded-service approaches. At low loads, there is no significant difference as acceptance ratio is high, but at high loads, our degraded-service schemes outperform the traditional scheme by decreasing BBR significantly. Since our schemes are much more flexible in terms of resource allocation, we can serve more connections with the same amount of bandwidth as the traditional scheme.

Since connections may receive degraded service in DS and EDS, the average bandwidth provided by these schemes is slightly lower than in FS (Fig. 5b). Figure 5c shows the benefit of upgrade; and to observe it, the connection arrival rates for this figure are higher than the others as at low load, disrupted connections can get reprovisioned with full service. At high network loads (e.g., 360 Erlang in Fig. 5c), the proposed degraded service schemes affect average provided bandwidth excessively, as much as 15 percent decrease, when no special action is taken. Our upgrade mechanism successfully remedies this problem by reducing the decrease in provided bandwidth from 15 percent in EDS without upgrade to 3 percent with upgrade and from 10 percent in DS without upgrade to 1 percent with upgrade (when load is 360 Erlang in Fig. 5c). Also, we observe that the acceptance and recovery rates of mission-critical connections are higher than regular connections as critical connections are prioritized.

Conclusion

Recent disasters have shown that current survivability schemes in our networks are lacking, and enhanced schemes need to be considered. By exploiting service differentiation and the degraded-service concept, we can improve the network's adaptability against disasters. We propose a method that accepts service degradation not only during failures but also during the admission process to increase service acceptance and/or availability. Also, in our proposal, some additional resources can be released by downgrading some existing connections to avoid dropping some other connections that are affected by the disaster. A case study where our proposal methods are applied to a U.S.-wide network topology shows our method's advantageous properties.

By exploiting service differentiation and the degraded-service concept, we can improve the net-work's adaptability against disasters. We proposed a method that accepts service degradation not only during failures but also during admis-sion process to increase service acceptance and/or availability.

References

[1] M. F. Habib *et al.*, "Disaster Survivability in Optical Communication Networks," *Computer Commun.*, vol. 36, no. 6, Mar. 2013, pp. 630–44.

[2] S. Huang *et al.*, "A Multistate Multipath Provisioning Scheme for Differentiated Failures in Telecom Mesh Networks," *J. Lightwave Tech.*, vol. 28, no. 11, June 2010, pp. 1585–96.

[3] H. Chang, "A Multipath Routing Algorithm for Degraded-Bandwidth Services Under Availability Constraint in WDM Networks," WAINA, 2012.

[4] P. Agarwal *et al.*, "The Resilience of WDM Networks to Probabilistic Geographical Failures," *IEEE/ACM Trans. Networking*, vol. 21, no. 5, 2013, pp. 1525–38.

[5] T. L. Weems, "How Far is Far Enough," *Disaster Recovery J.*, vol. 16, no. 2, Spring 2003.

[6] W. Zhang *et al.*, "Reliable Adaptive Multipath Provisioning with Bandwidth and Differential Delay Constraints," *IEEE INFOCOM*, Mar. 2010.

[7] S. S. Savas *et al.*, "Disaster-Aware Service Provisioning by Exploiting Multipath with Manycasting in Telecom Networks," *IEEE ANTS*, Chennai, India, Dec. 2013.

[8] N. Bitar, S. Gringeri, and T. J. Xia, "Technologies and Protocols for Data Center and Cloud Networking," *IEEE Commun. Mag.*, vol. 51, no. 9, pp. 24-31, Sept. 2013.

[9] S. Jain *et al.*, "B4: Experience with A Globally-Deployed Software Defined WAN," *ACM SIGCOMM*, Aug. 2013.

[10] S. Oueslati and J. Roberts, "Method and a Device for Implicit Differentiation of Quality of Service in a Network," U.S. Patent No. 7,646,715, 12 Jan. 2010.