1-2019

SDN Testbed for Evaluation of Large Exo-Atmospheric EMP Attacks

Diogo Oliveira

Nasir Ghani

Majeed M. Hayat

Jorge Crichigno

Elias Bou-Harb

SDN Testbed for Evaluation of Large Exo-Atmospheric EMP Attacks

Diogo Oliveira
School of Information, Florida State University
Nasir Ghani
Department of Electrical Engineering, University of South Florida
Majeed Hayat
Electrical and Computer Engineering Department, Marquette University.
Jorge Crichigno
Department of Integrated Information Technology, College of Engineering and Computing, University of South Carolina
Elias Bou-Harb
Computer Science Department at Florida Atlantic University

Abstract:

Large-scale nuclear electromagnetic pulse (EMP) attacks and natural disasters can cause extensive network failures across wide geographic regions. Although operational networks are designed to handle most single or dual faults, recent efforts have also focused on more capable multi-failure disaster recovery schemes. Concurrently, advances in software-defined networking (SDN) technologies have delivered highly-adaptable frameworks for implementing new and improved service provisioning and recovery paradigms in real-world settings. Hence this study leverages these new innovations to develop a robust disaster recovery (counter-EMP) framework for large backbone networks. Detailed findings from an experimental testbed study are also presented.

Network recovery from catastrophic disaster events is a major concern. In particular there is renewed concern about large-scale electromagnetic pulse (EMP) attacks caused by nuclear detonations in space by hostile state-based actors.

Introduction

Network recovery from catastrophic disaster events is a major concern. In particular there is renewed concern about large-scale electromagnetic pulse (EMP) attacks caused by nuclear detonations in space by hostile state-based actors. These exo-atmospheric attacks can occur at varying height of bursts (HoB), ranging from 50–100 miles (low-altitude) to 500 miles (high-altitude) and will emit large amounts of gamma radiation. In turn, these emissions will interact with the Earth's atmosphere and magnetic fields to release electrons [1]. Specifically, three types of EMP pulses will be generated across a broad range, called E1, E2 and E3 (see Fig. 1). The initial E1 component is a high-frequency short-duration pulse (10 MHz-1 GHz, ns-ms) with very high electric fields, over 10 kV/m. Unshielded electronics are particularly vulnerable to this effect. Meanwhile, the E2 component is a lower frequency longer-lasting pulse generated by secondary gamma scattering (100 kHz-10 MHz, ms-sec), However, this component is similar to lightning strikes, and hence existing shielding methods can suffice. Finally, the E3 component has much longer durations and lower frequencies (10–100 Hz, seconds-min-utes) and can severely damage electrical grids [2]. Note that E3 pulses can also be generated by geomagnetic storms from solar flares, and smaller non-nuclear EMP (NN-EMP) devices can also emit intense emissions across the E1-E3 spectrum (albeit with much smaller footprints).

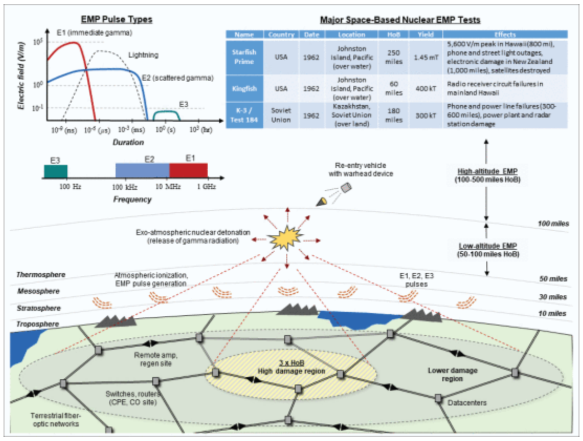

Name	Country	Date	Location	HoB	Yield	Effects
Starfish Prime	USA	1962	Johnston Island, Pacific (over water)	250 miles	1.45 mT	5,600 V/m peak in Hawaii (800 mi), phone and street light outages, electronic damage in New Zealand (1,000 miles), satellites destroyed
Kingfish	USA	1962	Johnston Island, Pacific (over water)	60 miles	400 kT	Radio receiver circuit failures in mainland Hawaii
K-3 / Test 184	Soviet Union	1962	Kazakhstan, Soviet Union (over land)	180 miles	300 kT	Phone and power line failures (300-600 miles), power plant and radar station damage

Figure 1. Effects of exo-atmospheric (space-based) nuclear EMP strike.

Overall, nuclear EMP attacks can cause near-instantaneous failure of unshielded electronics depending on warhead yield and epicenter location. These events are commonly called stressors and can damage critical infrastructures across a wide region, for example, telecommunications, electricity, oil, gas, water, transportation, aviation, agriculture, and so on. Systematic interdependencies between many of these infrastructures can trigger delayed cascades, further exacerbating fallouts [3]. Now several live space-based EMP tests have been conducted in the past, with empirical measurements showing notable damage at extended distances, for example, Starfish Prime, Kingfish, and K-3 (Fig. 1). These tests yielded vital empirical data and also helped uncover deficiencies with earlier EMP models (which largely underestimated the E1 component). Hence more accurate formulations were developed based on the pioneering work of theoretical physicist Dr. Conrad Longmire at Los Alamos National Laboratory, New Mexico, USA. Based on these models, a potential 1.5 megaton detonation over the central United States today (akin to the Starfish Prime test) will damage over 350 large transformers and cut power to almost 40 percent of the population [2]. Moreover, it will likely take years to rebuild and reinstall damaged power grids. Given these sobering facts, in 2011 the United States Congress established a commission to assess the EMP threat, the recommendations of which stressed the need for resilient communications to ensure effective recovery/response [2]. In particular, contingency and military planners are very concerned about shorter timescales in the immediate aftermath of an attack, that is, seconds to minutes.

Now most large backbones, including military, use fiber-optic systems for ultra-fast transmission over single-mode fiber. In some cases wireless links are also used to reach inaccessible regions or provide backup, that is, satellite and radio. Overall, the fiber medium is immune to EMP radiation as it does not conduct E1-E3 pulses. However, fiber-based transmission still relies on many opto-electronic systems that remain vulnerable to EMP effects, for example, routers, optical cross-connects, multiplexers, remote link regeneration/amplifier units, and so on. Many of these systems are located at dispersed sites such as central offices (CO), customer premises (CP), remote amplifier huts, and so on [2]. For the most part, these locations (in commercial settings) provide minimum shielding against weaker effects, for example, interference, electrostatic discharges, and lightning. Hence a powerful EMP device will cause rapid and extensive damage to backbone infrastructures. Delayed cascading power outages are also likely, but their impacts may not be immediate as most networking facilities have 48–72 hours of backup (batteries, fuel, and generators).

In light of the above, rapid post-attack recovery is a critical concern. However, most networks are only designed to handle single or dual node and link failures and provide little protection against large EMP attacks (natural disasters) causing multiple near-instantaneous failures with high spatial and temporal correlation [4], [5]. To address this concern, the Defense Threat Reduction Agency (DTRA) has funded various studies on disaster recovery, including single domain networks [5], [6], multi-domain networks [7], interdependent pow-er-communication grids [8], and robust infrastructure design/pre-planning [9]. These contributions leverage definitions/concepts of pre-fault probabilistic risk regions and have yielded key theoretical findings on failure modeling and multi-failure recovery. Nevertheless, the effectiveness of these solutions in realistic settings remains to be seen, and this forms the key motivation for this work.

Deploying advanced counter-EMP schemes in real-world communication networks is very challenging due to the prevalence of vendor-pro-prietary systems and tightly-coupled data-con-trol planes. These dependencies mandate niche expertise and limit the range of recovery options. However, emerging software defined networking (SDN) paradigms have redefined legacy boundaries and offer much better service programmability and agility [10]. Still, SDN only provides a control framework and relegates detailed service provisioning and recovery algorithm design to operators themselves. Despite some notable testbed studies [11]–[12][13], there are no known works on SDN-based disaster recovery. Hence, this effort focuses on translating advanced count-er-EMP service recovery schemes (only studied in research literature) into real-world settings by developing a resilient SDN-based management and orchestration (MANO) framework. Namely, these methods incorporate the definition of pre-fault probabilistic failure regions. A testbed facility is also built to evaluate the solution for large optical backbones with full conversion and regeneration.

This article is organized as follows. The following section presents a review of recent work in multi-failure network recovery and SDN testbeds, motivating the need for disaster recovery efforts. We then detail the resilient SDN MANO architecture, and present findings from the testbed study. Conclusions and future directions are outlined in the final section.

Existing Work

Various studies have addressed multi-failure network recovery. For example, some have looked at connection protection for multiple independent failures [4]. However, this approach does not accurately model large disasters or EMP attacks with high spatio-temporal fault correlation. Hence, [5] introduces a realistic probabilistic shared risk link group (p-SRLG) model to define a-priori stressor (risk) regions, that is, occurrence probabilities, conditional link/node failure rates, and so on. New optimization and heuristic methods are then proposed to minimize primary/backup path failure probabilities for pre-defined stressors. However, since these schemes focus on risk minimization, they yield longer (less efficient) routes. Hence, [6] outlines a load balancing heuristic to improve the reliability-efficiency tradeoff. Finally, others have also studied multi-failure recovery in multi-domain networks, for example, two-stage optimization models to minimize failures at the intra-domain/inter-domain levels, adhoc post-fault re-routing, etc [7].

Despite these contributions, few researchers have addressed multi-failure recovery under realworld conditions. Here, the complex vendor-pro-prietary nature of legacy control setups makes it extremely difficult (if not impossible) to implement tailored schemes. For example, most setups feature single-vendor network management systems (NMS) with limited recovery options for single/dual node and link failures [6]. It is here that emerging SDN frameworks [10] provide a much-needed avenue for building more capable (counter-EMP) recovery solutions. Namely, SDN decouples the data and control planes and relegates all path computation and provisioning to controllers operating with “bird's eye” network views. These entities are responsible for generating flow routes/rules and pushing them onto network nodes, which in turn are streamlined for data forwarding. This programmable open architecture allows operators to achieve flow-selective routing and can support a wide range of applications/policies, for example, load balancing, survivability, security, and so on.

Now various SDN protocols have emerged, including OpenFlow, OpenSwitch, OpenLight, NETCONF, others [10], [13]. Numerous testbeds have also deployed these SDN-based solutions for a range of scenarios, for example, datacen-ter, enterprise, metro-core, mobile, and so on. Although a detailed survey is out of scope here, some recent efforts are summarized. For example, [11] extends SDN control for a multi-do-main testbed running the *Floodlight* and *Open vSwitch* protocols. Novel methods are proposed for “east-west” inter-controller exchange and domain virtualization. Meanwhile, [12] proposes a resilient SDN architecture for connecting high-performance zones in a higher-education network. The authors also use OpenFlow devices to perform failover switching for isolated link outages. Finally, [13] studies controller failures and compares the scalability and reliability of multi-controller Open Network Operating System (ONOS) and OpenOaylight (ODL) setups, that is, for rule installation, network database synchronization, and failover mechanisms. Nevertheless, despite these (and other) contributions, there are no known SDN studies on large-scale disaster recovery. A detailed framework is now presented to address this need.

Architecture Design and Development

As noted earlier, several SDN protocols have been developed. However, *OpenFlow* is chosen here as it is the most widely-used solution. Several network operating systems have also been proposed to interface with the southbound (SDN protocol) application programmer interfaces (API) and automate service provisioning, for example, *ONOS*(by ON.lab) and the vendor-based *ODL* solution. Again, *ONOS* is chosen as it is designed for use with *OpenFlow.* Nevertheless, the Open-Flow and ONOS combination still lacks critical features for multi-failure recovery as it does not provide any path computation element (PCE) for route generation. Therefore, a novel resilient SDN system MANO framework is developed for disaster recovery support. This setup is specifically designed to handle large multi-failure (e.g., EMP-type) events and allows operators to pre-specify vulnerable regions and stressor likelihoods. The framework is shown in Fig. 2 and is comprised of three key entities, that is, an Abstraction Module, a Resilient Path Computation (RPC) Module, and an ONOS Driver. The solution leverages the *OpenFlow* and *ONOS* frameworks and adds critical features for multi-failure recovery, that is, state visibility, pre-fault/post-fault schemes, and so on.

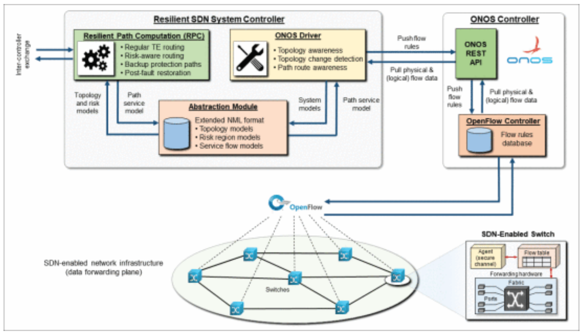

Figure 2. Resilient SDN MANO system architecture for multi-failure network recovery.

Overall, the proposed setup assumes that all network nodes are generic bandwidth/circuit provisioning entities, that is, opaque/full-conversion optical nodes or backbone routers. However, future extensions can be added to handle all-optical or partial-conversion nodes requiring wavelength continuity. Also, an SDN controller can represent a single point-of-failure and bottleneck. Hence in practice, multi-controller setups are required to improve *control plane* redundancy, and these designs can leverage from related work in [13]. However, this is left for future study as the focus is on data plane recovery. Consider the details.

Abstraction Module

This module implements a database to allow the resilient SDN-based MANO system to visualize the network, its risk vulnerabilities, and active flows. This overall state is then used to drive the other key system components, for example, the ONOS Driver module. In particular, the network markup language (NML) standard [14] is used to describe the network topology using abstract models based on classes (also called resources), attributes, relations and parameters. These entities are then encoded in appropriate schema files. For example, NML defines base schema constructs for physical resources at switches and links using topology-related resources and properties. These constructs include inbound and outbound link ports and neighbors. Fig. 3 shows a fraction of the ontology needed to define such a topology.

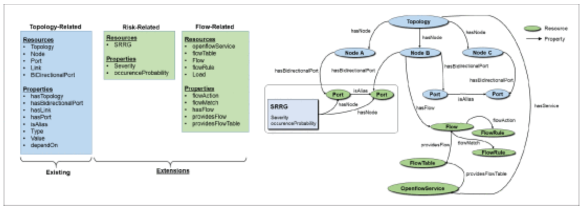

Figure 3. Extended network markup language (E-NML) schema.

Nevertheless, the base NML schema [14] only details static topologies. Hence, new class and attribute definitions are also introduced for critical risk vulnerability and path route information to support multi-failure disaster recovery, called extended NML (E-NML) (see Fig. 3). Foremost, a shared risk resource group (SRRG) service class is defined to group nodes and links with common a-priori risks. Here a SRRG resource identifies a group of nodes and ports with specific severity and failure probability values, respectively, given by the *severity* and *occurenceProbability* properties. Note that the shared risk (link) group concept has been studied by researchers [5]–[6][7], but its adaptation in operational SDN settings for multi-failure recovery is lacking.

Next, a path service model is introduced to enhance the NML base schema and define control plane flow rules to detail logical relations. Namely, the nodes, ports and flow rules associated with a source-destination pair are specified by a new flow service class, called *openFlowService* (see Fig. 3). This class defines several other flow-related classes, relations, attributes and parameters. Foremost, *openFlow-Service* resources are associated with a topology, which in turn supports SDN nodes with flow tables, that is, *providesFlowTable* property. A flow table is also defined by the *flowTable* class and provides a set of flows (*Flow* resources) through the *providesFlow* relation property. Ideally, a flow class is created to match a flow, and hence the *flowMatch* relation defines a *flowRule* traffic criteria, that is, source/des-tination media access control (MAC) addresses, source/destination IP addresses, inbound/ outbound ports, virtual LAN (VLAN) identifiers, and so on. Furthermore, a *flowRule* is also associated with a *flow* by a *flowAction* relation to determine control plane decision-making for specific traffic data, that is, forward frame through outbound port *N.* Finally, a set of NML *flow* objects (instantiated *flow* classes) are also defined for each source-destination path pair, and each path is also assigned a unique identification number (UUID).

The Resilient Path Computation Module retrieves state information from the Abstraction Module and performs constrained route computation between requesting source-destination nodes. Topology information is used to build a graph-based network view comprised of vertices and *edges*.

Resilient Path Computation (RPC) Module

This module retrieves state information from the Abstraction Module and performs constrained route computation between requesting source-destination nodes. Topology information is used to build a graph-based network view comprised of vertices and edges. Meanwhile, input requests are defined via a new path request service model which specifies a source, destination, and desired bandwidth capacity (see Fig. 3). Now as noted above, several polynomial-time graph-based algorithms have been developed for multi-failure disaster recovery [5]–[6][7]. Hence, some of these (and some other baseline) solutions are adapted here (Fig. 4).

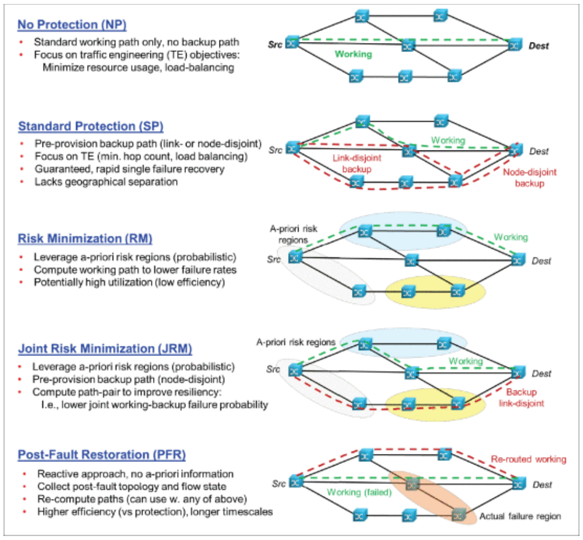

No Protection (NP)

- Standard working path only, no backup path
- Focus on traffic engineering (TE) objectives: Minimize resource usage, load-balancing

Standard Protection (SP)

- Pre-provision backup path (link- or node-disjoint)
- Focus on TE (min. hop count, load balancing)
- Guaranteed, rapid single failure recovery
- Lacks geographical separation

Risk Minimization (RM)

- Leverage a-priori risk regions (probabilistic)
- Compute working path to lower failure rates
- Potentially high utilization (low efficiency)

Joint Risk Minimization (JRM)

- Leverage a-priori risk regions (probabilistic)
- Pre-provision backup path (node-disjoint)
- Compute path-pair to improve resiliency: I.e., lower joint working-backup failure probability

Post-Fault Restoration (PFR)

- Reactive approach, no a-priori information
- Collect post-fault topology and flow state
- Re-compute paths (can use w. any of above)
- Higher efficiency (vs protection), longer timescales

Figure 4. Network service recovery mechanisms (pre- and post-fault).

No Protection (NP)

This baseline scheme does not provide any backup. A primary path is provisioned for a request using Dijkstra's shortest path algorithm based upon a desired network traffic engineering (TE) objective, for example, hop count minimization, load balancing, and so on.

Standard Protection (SP)

This scheme computes a backup path, that is, a primary-backup path pair (Fig. 4). The algorithm starts by computing the *k* shortest-paths between the source and destination. Node-disjoint protection paths are then derived for each candidate path, and the path pair with the lowest combined TE cost is chosen based on hop count or load cost. This joint computation strategy is more efficient than a greedy *sequential* approach (which first computes a primary path, prunes related links/nodes, and then computes a backup route).

Risk Minimization (RM)

This scheme computes a primary working path in a "risk-averse" manner to increase reliability based upon pre-specified risk regions defined by operators, that is, occurrence probabilities and conditional node/link failure rates (p-SRLG models in the Abstraction Module). A constrained path is computed to minimize risks over a modified topology with link weights set as logarithmic functions of the risk probabilities [5], [6]. Expectedly, path reliability will depend upon the accuracy of the a-priori SRLG models.

Joint Risk Minimization (JRM)

This scheme extends the RM scheme to compute a "risk-averse" backup path. Namely, the *k* shortest-paths are computed over a modified graph using the same logarithmic link weights. Node-disjoint backup paths are then computed for each possible primary path, and the working-backup path pair with the minimum TE cost (using hop count or load balancing) is chosen (see [6]). Hence, this scheme tries to achieve a balance between cost and risk, but for small *k* values (k≤3) will still favor the latter.

Post-Fault Restoration (PFR)

This method focuses on *post-fault* recovery and ca*n* be used in conjunction with any of the above schemes. The RPC module waits for a pre-defined hold-off time to allow the Abstraction Module to collect updated topological and path service state. New routes are then re-computed for affected paths experiencing one/more link failures to re-estab-lish connectivity. The affected paths are randomly selected for recovery, and path re-computation is done using the baseline NP scheme.

After computation, the path is translated/ stored in a path service model in the Abstraction Module database (see Fig. 2). Note that the unprotected NP and JRM schemes only store one main path, whereas the protected SP and JRM schemes store *an*additional backup path (albeit only one is active, as detailed next).

ONOS Driver

ONOS provides a framework for managing *Open-Flow-capable* devices. Leveraging this, the ONOS Driver module in Fig. 2 implements the main interface with the underlying ONOS controller via the ONOS REST representational state transfer (REST) API. Several key directional transfers are implemented. Foremost, topology state is pulled in the *northbound* direction from SDN-enabled devices, that is, to populate the Abstraction Module database when prompted by an incoming request. Meanwhile, path route information is pushed in the *southbound* direction to instantiate end-to-end routes. Namely, the ONOS Driver retrieves the relevant path request service model from the Abstraction Module to obtain the flow rules, that is, frame processing/forwarding. This model is then parsed and encoded into JavaScript object notation (JSON) format and sent to the ONOS controller to distribute to all path nodes. The OpenFlow flow identification numbers (FID) assigned to each deployed flow are also retrieved and passed to the Abstraction Module database to track the deployed flow rules, that is, northbound transfer.

Now only one path service model can be active at a given time for a demand. Hence, primary path rule sets are established after a request is processed. However, for protected demands (SP, JRM schemes), backup path switchover/ reversion is only performed after link/node failures are detected. Here the ONOS Driver module retrieves **all** currently-active primary FID routes from the Abstraction Module database and signals the ONOS controller to remove affected flows with failed link(s). Flow rules for all corresponding backup paths (unaffected by failures) are then pushed to the SDN controller to re-establish routes. Now if path protection is not provisioned (NP, SP, or RM schemes) and/or backup protection paths are also affected by failures, then the PFR scheme can (optionally) be used to recover failed routes. Here the ONOS Driver removes the SDN flow table entries for affected paths by invoking a procedure to retrieve failure-affected paths. For each such path, the PFR algorithm searches for a recovery path and allocates bandwidth on its links before performing switchover. Akin to the pre-computation schemes, this approach relies on a service request model with a source and destination pair. Once the shortest path is computed, a service model is defined and stored in the database.

Note that larger multi-area networks will require multiple (SDN) controllers, mandating distributed inter-domain path computation. Specifically, multiple ONOS Driver instances will have to be instantiated to communicate (push, pull) required topology information between peer SDN controllers. However, inter-domain path computation is complicated by the lack of global visibility, *and* many different schemes have been proposed.

Note that larger *multi-area* networks will require multiple (SDN) controllers, mandating distributed *inter-domain* path computation. Specifically, multiple ONOS Driver modules will have to be instantiated to communicate (push, pull) required topology information between peer SDN controllers. However, inter-domain path computation is complicated by the lack of global visibility, and many different schemes have been proposed (see [7], [13]). Owing to these complexities, such features are left for future study.

Testbed Study

The resilient SDN-based MANO framework is implemented in an advanced testbed and evaluated for a range of nuclear EMP threats. The impact of such attacks is best gauged for a large conti-nental-sized backbone, and hence the testbed uses a realistic 75 node/99 link Continental US (CONUS) topology from the Defense Advanced Research Project Agency (DARPA) CORONET project [15]. This network is implemented using a combination of real-world physical and soft-ware-emulated network nodes and links. Specifically, five *OpenFlow-enabled iwNetworks Pronto 3290* bare-metal switches are used to build part of the CONUS network at the University of Maryland Mid-Atlantic Crossroads (MAX) facility. The remaining nodes and links are emulated in *Mininet* on a quad-core Intel Xeon E3-1200 processor with 8 gigabytes of memory. **All** key system modules (Fig. 3) are separately implemented i*n* a quad-core Intel Xeon E3-1200 processor with 32 gigabytes of memory.

A range of potential EMP attacks are developed for the central United States. This region is of key strategic interest (vulnerability) since its nodes/links serve in critical *transit* roles for cross-country routes. Note that several stressors are also tested over the East and West Coast regions, but these give low recovery rates since a large percentage of the failed nodes are source/ destination end-points, that is, they need physical repair. Now as detailed i*n* [1], low-altitude EMP attacks (HoB in 50–300 miles) exhibit “smi-ley face” geometries (due to Earth's magnetic fields) with three failure sub-regions (see Fig. 5a). The (red) sub-region closest to the epicenter experiences near complete failure of unshielded nodes and links. The wider surrounding (yel-low) sub-region also sees relatively high failures, that is, 75 percent range. However, the outer (green) sub-region has notably lower outages, averaging 20 percent. Based on this, three EMP attack types are defined, as shown superimposed over the CORONET test network in Fig. 5. These stressors include a 50 mile HoB with 700 mile diameter (Stressor 1, Fig. 5a), an 80 mile HoB with 1,120 mile diameter (Stressor 2, Fig. 5b), and a 100 mile HoB with 1,440 mile diameter (Stressor 3, Fig. 5b). The smaller Stressor 1 tests are further evaluated for five epicenters to gauge geographic sensitivity, that is, Central (C), West Central (WC), East Central (EC), South Central (SC), and North Central (NC) (see Fig. 5a). Also, three random permutations are averaged for each stressor site to evaluate different node/ link failure combinations i*n* the yellow and green regions. Meanwhile, the larger Stressor 2 and 3 scenarios are only tested for one epicenter as they span most of the central United States, that is, almost cause network partitioning. Still, three random permutations are tested. Note that these test cases include link failures between surviving/working end-points, that is, to model optical amolifier/link failures.

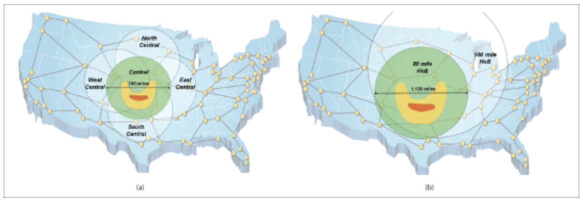

Figure 5. Nuclear EMP stressor footprints over 75-node CONUS topology: a) Stressor 1 (50 mile HoB); b) Stressors 2, 3 (80, 100 mile).

The scenarios are evaluated for medium-to-heavy load with 100 cross-country connections between random East and West Coast locations (requesting 5 percent link capacity). In general, these test cases represent *an* optical backbone carrying wavelength lightpaths with full/ adequate conversion and regeneration. Now the actual "attacks" are triggered by running scripts to fail the impacted nodes/links specified i*n*the test cases. As noted earlier, the "risk-aware" (RM, JRM) schemes require a-priori SRRG regions for path computation and hence three types are defined/inputted, that is, East Coast, Central, and West Coast. As the Central region is most relevant here, its SRRG footprint is purposely set different from any of the test cases (Stressors 1–3 in Fig. 5). This choice injects a high degree of realism as actual and predicted failures will rarely match. A subset of tests are also done using a minimum level of EMP hardening. Twenty percent of nodes and links (opto-electronic regeneration units) in the central region of the CONUS network are assumed to be immune to E1 and E2 effects, for example, via MIL 188–125 or TEMPEST shielding.

Overall post-attack failure rates are shown i*n* Fig. 6, that is, the number of failed connections after recovery attempts. A tiered recovery strategy is also implemented where the postfault re-routing (PFR) scheme is done in case of backup route failure. Foremost, the results for Stressor 1 attacks in Fig. 6a indicate very high outages for regular unprotected demands, that is, over 70 percent of failures for the NP scheme in non-hardened networks. Sensitivity to epicenter location is also evident here, with attacks in the SC and WC regions failing over 90 percent of unprotected East-West connection paths. However, these findings also confirm notable improvement with the survivability schemes. For example, the basic node-disjoint protection (SP scheme) gives 10–30 percent fewer failures depending upon Stressor 1 epicenters, and these gains increase with higher initial failures. More importantly, the "risk-aware" schemes give even better results. For example, using a-priori information to compute more resilient routes (RM scheme) yields 20 percent fewer failures versus regular TE routing (NP scheme) for attacks in the WC region and over 50 percent i*n* the SC region. "Risk-aware" route protection (JRM scheme) is even more effective, with post-fault failures dropping to under 30 percent in most cases. The improvement over basic protection (SP scheme) is also quite significant, with 20 to 40 percent fewer route failures. Finally, active post-fault re-routing (PFR scheme) recovers all failed demands regardless, as Stressor 1 footprints leave sufficient post-fault working link capacity.

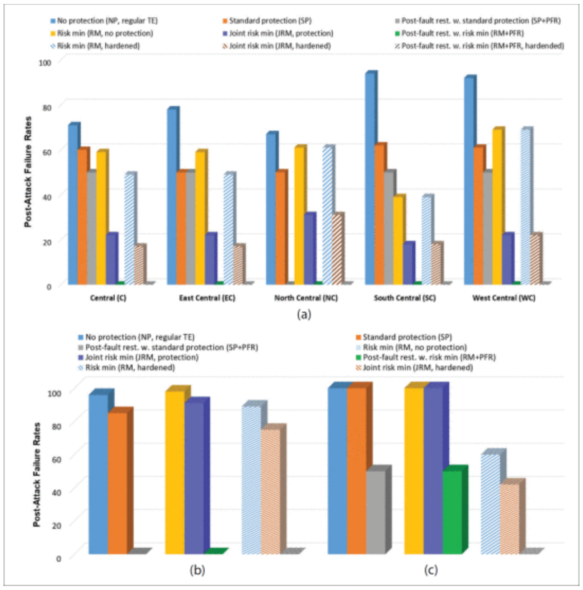

Figure 6. Post-attack connection failure rates: a) Stressor 1, b) Stressor 2, c) Stressor 3.

However, results for larger EMP attacks show much higher outages, that is, Fig. 6b(Stressor 2) and Fig. 6c (Stressor 3). For example, unprotected TE-routed demands experience near total failure in all cases. Furthermore, basic protection (SP scheme) only gives nominal recovery, under 10 percent. More importantly, a-priori risk information (RM, JRM schemes) provides minimal failure reduction for both stressors, that is, since larger failure regions dwarf the smaller a-pri-ori central failure region. In other words, the realistic discrepancy between predicted and actual failure footprints reduces the effectiveness of a-priori methods. By contrast, theoretical studies show notably better recovery since they assume ideal matching failure regions [4], [5]. Nevertheless, post-fault re-routing (PFR scheme) is still very effective for larger attacks, that is, recovering all Stressor 2 outages (Fig. 6b) and lowering Stressor 3 failures to below 50 percent (Fig. 6c). However this scheme still cannot provide full recovery in the latter case since only one or two working links are left to prevent network partitioning (bandwidth congestion when re-routing large numbers of affected demands).

As noted earlier, the “risk-aware” methods (RM, JRM, and PFR schemes) are also evaluated for moderate levels of network EMP hardening (the shaded bars in Fig. 6). These results indicate that the “risk-aware” RM and JRM schemes generally give moderate (non-negligible) recovery gains in hardened networks for smaller Stressor 1 and 2 type attacks, i*n* the 0–20 percent range (Figs. 6a and 6b). However, failure reduction with network hardening is much more pronounced for larger Stressor 3 events, that is, i*n* the 40 to 50 percent range (Fig. 6c). Furthermore, post-fault re-routing (PFR scheme) even recovers all failed routes for larger Stressor 3 attacks, that is, since hardened networks generally provide more capacity *on* critical transit links. These findings indicate that operators can reduce, that is, tradeoff, their investments i*n*network hardening to an extent by deploying more capable SDN-based recovery solutions.

Average route length (hop count) is also measured to gauge resource utilization (see Table 1). These results confirm notably longer pre-fault routes with backup protection, about 40 percent. However, “risk-aware” routing only gives a slight increase i*n* usage versus regular TE provisioning, that is, 1 to 2 percent higher average hop counts with the RM scheme (JRM scheme) versus the TE scheme (SP scheme). Meanwhile, post-fault PFR recovery shows sensitivity to stressor size and location. For example, the larger Stressor 2 and 3 attacks yield longer routes, averaging 13 to 14 hops (non-hardened). At the same time, some smaller Stressor 1 epicen-ters also give higher utilization, *on* par with larger attacks. These findings also show a slight decrease in post-fault utilization with network hardening, that is, less detoured routes.

A-priori risk information (RM, JRM schemes) provides minimal failure reduction for both stressors, that is, since larger failure regions dwarf the smaller a-priori central failure region. In other words, the realistic discrepancy between predicted and actual failure footprints reduces the effectiveness of a-priori methods. By contrast, theoretical studies show notably better recovery since they assume ideal matching failure regions.

Finally, in terms of delay, it takes about 2 minutes (1.25 minutes) to compute a primary (protection) path. However, since all routes are pre-computed and stored beforehand, post-at-tack recovery is not impacted. Meanwhile, average connection setup times are about 2 minutes during working conditions. Conversely, average post-fault recovery delays with the PFR scheme are generally higher, as shown in Table 1, includi*ng* cumulative switchover times from failed protection paths to re-computed working routes (SP→PFR or JRM→PFR). These findings confirm that "risk-aware" provisioning gives faster recovery as there are fewer failed routes to recover, as seen by comparing SP→PFR with IRM→PFR times (non-hardened). Moreover, hardening gives equivalent or lower averages as it yields fewer failures.

This study presents a novel resilient SDN control framework for disaster recovery, with a focus on large space-based nuclear EMP attacks. The solution introduces new survivability-related schema definitions *and* implements a range of pre-fault and post-fault strategies.

Table 1. Average path lengths (hops) and post-fault switchover times

Pre-fault path lengths (hops)				
Scheme TE		RM	SP	JRM
Stressor 1–3 11.46		11.73	15.76	15.9
Post-fault path lengths (hops)				
Non-		ardened		Hardened
Scheme ⍰		SP ⍰ PFR	JRM ⍰ PFR	JRM ⍰ PFR
Stressor 1	C	10.98	11.87	11.48
	NC	11.71	12.42	12.44
	SC	10.93	11.27	11.00
	EC	11.10	12.40	11.40
	WC	11.25	12.01	11.59
Stressor 2		12.04	13.88	12.27
Stressor 3		13.35	14.38	13.00
Post-fault switchover time (s)				
Non-		ardened		Hardened
Scheme ⍰		SP ⍰ PFR	JRM ⍰ PFR	JRM ⍰ PFR
Stressor 1	C	8.5	3.3	2.6
	NC	9.2	4.7	4.7
	SC	5.9	2.7	2.7
	EC	8.9	3.3	2.6
	WC	10.4	3.3	3.3
Stressor 2		14.7	13.7	13.4
Stressor 3		15.0	15.0	9.0

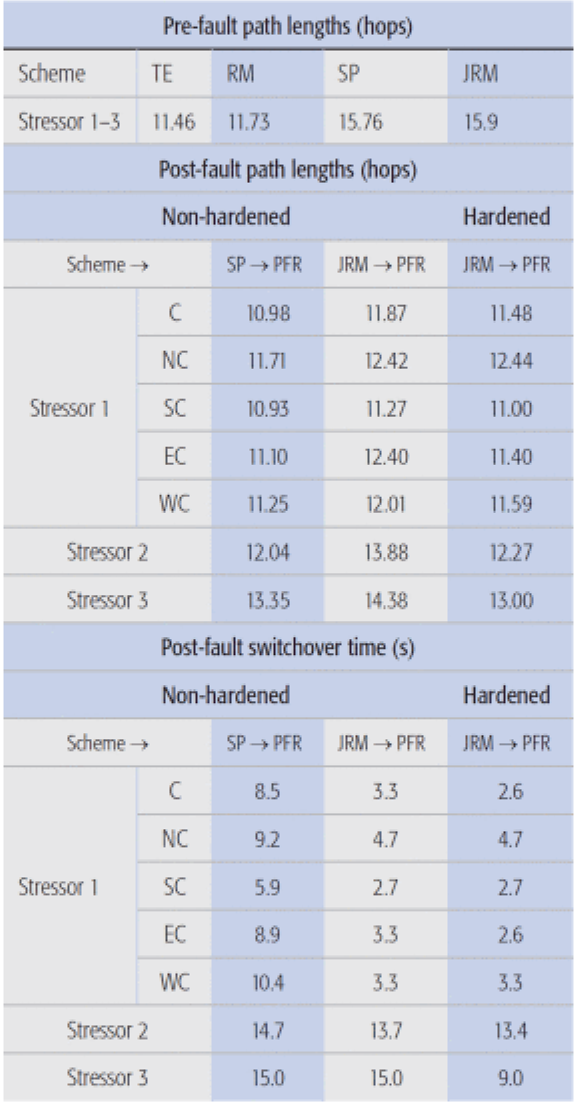

Pre-fault path lengths (hops)				
Scheme	TE	RM	SP	JRM
Stressor 1–3	11.46	11.73	15.76	15.9
Post-fault path lengths (hops)				
Non-hardened				Hardened
Scheme →		SP → PFR	JRM → PFR	JRM → PFR
Stressor 1	C	10.98	11.87	11.48
	NC	11.71	12.42	12.44
	SC	10.93	11.27	11.00
	EC	11.10	12.40	11.40
	WC	11.25	12.01	11.59
Stressor 2		12.04	13.88	12.27
Stressor 3		13.35	14.38	13.00
Post-fault switchover time (s)				
Non-hardened				Hardened
Scheme →		SP → PFR	JRM → PFR	JRM → PFR
Stressor 1	C	8.5	3.3	2.6
	NC	9.2	4.7	4.7
	SC	5.9	2.7	2.7
	EC	8.9	3.3	2.6
	WC	10.4	3.3	3.3
Stressor 2		14.7	13.7	13.4
Stressor 3		15.0	15.0	9.0

Conclusions and Future Efforts

This study presents a novel resilient SDN control framework for disaster recovery, with a focus *on* large space-based nuclear EMP attacks. The solution introduces new survivability-related schema definitions and implements a range of pre-fault and post-fault strategies. The framework is deployed i*n* an advanced SDN testbed running the *ONOS* framework to manage *OpenFlow* switches. Various scenarios are evaluated, and the findings validate results from earlier theoretical studies *on* the benefit of using a-priori risk information and fast re-routing recovery. More importantly, this effort demonstrates the effectiveness of using SDN protocols to implement counter-EMP network recovery. As such, this testbed provides a vital operational facility to develop and evaluate further pre-fault/ post-fault EMP mitigation schemes. These efforts ca*n*include modified path recovery schemes, delayed cascading failure methods, expanded multi-domain disaster recovery methods, static network pre-design (EMP hardening), and post-fault repair strategies. The latter efforts ca*n* include rapid deployment of backup cellular and/or satellite communication assets to re-establish connectivity.

ACKNOWLEDGMENT

This work has been funded by the Defense Threat Reduction Agency (DTRA) under Grant HDTRA1-13-C-0027. The authors are very grateful for this support.

References

1. S. Gladstone, P. Dolan, The Effects of Nuclear Weapons, United States Department of Defense and the Energy Research and Development Administration, 1977.

2. S. Gladstone, P. Dolan, "Critical National Infrastructures" in Report of the Commission to Assess the Threat to the US from Electromagnetic Pulse (EMP) Attack, April 2008.

3. M. Naeini et al., "Modeling Stochastic Correlated Failures and Their Effects on Network Reliability", *Proc. IEEE ICCCN 2011*, Aug. 2011.

4. Q. She, X. Huang, J. Jue, "Maximum Survivability Under Multiple Failures", *IEEE/OSA Optical Fiber Conf. (OFC) 2006*, Mar. 2006.

5. H. Lee, E. Modiano, K. Lee, "Diverse Routing in Networks with Probabilistic Failures", *IEEE/ACM Trans. Networking*, vol. 18, no. 6, pp. 1895-1907, Dec. 2010.

6. O. Diaz et al., "Network Survivability for Multiple Probabilistic Failures", *IEEE Commun. Lett.*, vol. 16, no. 8, pp. 1320-23, 2012.

7. N. Ghani, M. Peng, A. Rayes, N. Anto-niades, G. Ellinas, I. Roudas, "Service Provisioning and Survivability in Multi-Domain Optical Networks" in Design and Engineering of WDM Systems and Networks, Springer, pp. 481-519, 2011.

8. S. Buldyrev et al., "Catastrophic Cascade of Failures in Interdependent Networks", *Nature*, vol. 464, pp. 1025-28, Apr. 2010.

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

10. B. Nunes et al., "A Survey of Software-Defined Networking: Past Present and Future of Programmable Networks" in IEEE Commun. Surveys & Tutorials, pp. 1617-34, 3rd Quarter, Feb. 2014.

11. P. Lin et al., "A West-East Bridge Based SDN Inter-Domain Testbed", *IEEE Commun. Mag.*, vol. 53, no. 2, pp. 190-97, Feb. 2015.

12. M. Menth et al., "Resilient Integration of Distributed High-Performance Zones into the BelWue Network Using OpenFlow", *IEEE Commun. Mag.*, vol. 55, no. 4, pp. 94-99, Apr. 2017.

13. D. Suh et al., "Toward Highly Available and Scalable Software Defined Networks for Service Providers", *IEEE Commun. Mag.*, vol. 55, no. 4, pp. 100-07, Apr. 2017.

14. J. Ham et al., "Network Markup Language Base Schema Version 1" in Open Grid Forum, 2013.

15. G. Clapp et al., "Management of Switched Systems at 100 Tbps: The DARPA CORONET Program", *Proc. Intl Conf. Photonics in Switching 2009 (PS 2009)*, Sept. 2009.