# Modelling and Analysis of Network Resilience

*Invited Paper*

James P.G. Sterbenz*†, Egemen K. Çetinkaya*, Mahmood A. Hameed*, Abdul Jabbar* and Justin P. Rohrer*
*Department of Electrical Engineering and Computer Science, Information and Telecommunication Technology Center
The University of Kansas, Lawrence, KS, 66045, USA
†Computing Department, InfoLab21, Lancaster University, Lancaster, LA1 4WA, UK
{jpgs|ekc|hameed|jabbar|rohrej}@ittc.ku.edu, jpgs@comp.lancs.ac.uk
http://www.ittc.ku.edu/resilinets

***Abstract*—As the Internet becomes increasingly important to all aspects of society, the consequences of disruption become increasingly severe. Thus it is critical to increase the resilience and survivability of the future network. We define *resilience* as the ability of the network to provide desired service even when challenged by attacks, large-scale disasters, and other failures. This paper describes a comprehensive methodology to evaluate network resilience using a combination of analytical and simulation techniques with the goal of improving the resilience and survivability of the Future Internet.**

***Keywords*- Future Internet architecture, resilience, survivability, performability, dependability, topology, population, attack, disaster, challenge, metrics, generation, simulation, modelling**

## I. Introduction and Motivation

The increasing importance of the Global Internet has lead to it becoming one of the critical infrastructures [1] on which almost every aspect of our lives depend. Thus it is essential that the Internet be *resilient*, which we define as the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation [2]. It is generally recognised that the current Internet is not as resilient, survivable, dependable, and secure as needed given its increasingly central role in society [3]–[7]. Thus, we need to ensure that *resilience is a fundamental design property of the Future Internet*, and seek ways to increase the resilience of the current and future Internet. This requires an understanding of vulnerabilities of the current Internet, and a methodology to test alternative proposals to increase resilience. In particular, we are interested in understanding, modelling, and analysing the properties of *dependability* that quantifies the reliance that can be placed on the service delivered by a system including reliability and availability [8] and *performability* that quantifies the level of performance [9]. This notion of resilience subsumes *survivability*, which is the ability to tolerate the correlated failures that result from attacks and large-scale disasters [10], [11].

This paper describes a comprehensive approach to evaluate network resilience through analysis and simulation, with a brief discussion on experimentation, and is organised as follows: Section II reviews an architectural framework for network resilience based on a two-phase strategy $D^2R^2$+DR. Section III presents the problem of generating realistic topologies that can be used to evaluate network resilience, and introduces the Ku-LoCGen topology generator. Section IV describes an analytical formulation of resilience as the trajectory through a multilevel two-dimensional state space with operational and service dimensions. Section V describes a simulation methodology to evaluate the resilience of alternative network architectures with emphasis on attacks and area-based challenges using the KU-CSM challenge simulation module. Finally, Section VI summarises the main points of the the paper.

## II. Resilience Framework and Strategy

There have been several systematic resilience strategies, including ANSA [12], T1A1.2 [13], CMU-CERT [10], and SUMOWIN [11]. This section briefly reviews the *ResiliNets* resilience framework and strategy [2], based in part on these previous frameworks, which provides the basis for the resilience evaluation methodology described in the rest of the paper. The framework begins with a set of four axioms that motivate the strategy: 1) *Faults are inevitable*; it is not practical nor possible to construct perfect systems, nor is it possible to prevent challenges and threats. 2) *Understanding normal operation* is necessary, including the environment and application demands. It is only by understanding normal operation that we have any hope of determining when the network is challenged or threatened. 3) *Expectation and preparation for adverse events and conditions* is necessary, so that defences and detection of challenges that disrupt normal operations can occur. These challenges are inevitable. 4) *Response to adverse events and conditions* is required for resilience, by remediation ensuring correct operation and graceful degradation, restoration to normal operation, diagnosis of root cause faults, and refinement of future responses.

The ResiliNets strategy consists of two phases $D^2R^2$+DR, as shown in Figure 1. The first strategy phase $\mathbf{D^2R^2}$ (defend, detect, remediate, recover) consists of a passive core and a cycle of four steps that are performed in real time throughout the network and are directly involved in network operation and service provision:

***Defend.*** The basis for a resilient network is a set of defences that reduce the probability of a fault leading to a failure (fault-tolerance) and reduce the impact of an adverse event on network service delivery. These defences are identified by developing and analysing threat models, and

978-1-4244-8953-4/11/$26.00 

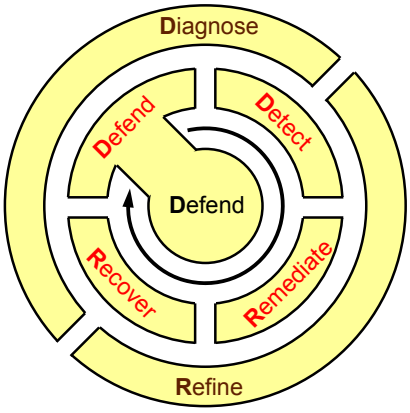


Fig. 1. ResiliNets strategy

consist of a passive and active component. Passive defences are primarily structural, consisting of geographically diverse redundant paths using alternative technologies such as simultaneous wired and wireless links, so that a challenge to part of the network permits communication to be routed around the failure [14]. Trust boundaries between network realms allow active defences consisting of self-protection mechanisms operating in the network that defend against challenges, such as firewalls that filter traffic for anomalies and known attack signatures, and the eventual connectivity paradigm that permits communication to occur even when stable end-to-end paths cannot be maintained [11], which is the basis for the discipline of disruption-tolerant networking. Clearly, defences will not always prevent challenges from penetrating the network, which leads to the next strategy step.

***Detect.*** The second step for the network as a distributed system, as well as individual components such as routers, is to detect challenges and to understand when the defences have been penetrated. There are three main ways to determine if the network is challenged. The first of these involves understanding the service requirements and the normal operational behaviour of a system and detecting deviations from it – *anomaly detection* based on metrics (described in Sec. IV). The second approach involves detecting when errors occur in a system that could lead to a service failure. Finally, a system should detect service failures; an essential facet of this is understanding service requirements. An important aspect of detecting a challenge is determining its nature, which requires context awareness. This along with detection informs the appropriate remediation strategy in the next step.

***Remediate***. The third step is to remediate the effects of the detected adverse event or condition to minimise the effect on service delivery. The goal is to do the best possible after an adverse event and during an adverse condition. This requires adaptation and autonomic behaviour so that corrective action can be taken at all levels without direct human intervention, to minimise the impact of service failure, including correct operation with graceful degradation of performance. A common example of remediation is for dynamic routing protocols to reroute around failures (e.g. [15]) and for adaptive applications and congestion control algorithms to degrade gracefully from acceptable to impaired service (Section IV). There may be a number of strategies that can be used to remediate against a given challenge.

***Recover.*** Finally, once the challenge is over after an adverse event or the end of an adverse condition, the network may remain in a sub-optimal state as an artifact of remediation (such as longer-path routing around links that have been restored). Thus, the network should recover to its original optimal normal operation, since continued remediation activities may incur an additional resource cost.

The second phase **DR** (diagnose, recover) consists of two background operations that observe and modify the behaviour of the $D^2R^2$ cycle: diagnosis of faults and refinement of future behaviour. While currently these activities generally have a significant human involvement, a future goal is for autonomic systems to automate diagnosis and refinement.

***Diagnose.*** While it is not possible to directly *detect* faults, we may be able *diagnose* the fault that caused an observable error, using techniques such as root-cause analysis. The goal is to either remove the fault (generally a design flaw as opposed to an intentional design compromise) or add redundancy for fault-tolerance so that service failures are avoided in the future.

***Refine.*** The final aspect of the strategy is to refine behaviour for the future based on past $D^2R^2$ cycles. The goal is to learn and reflect on how the system has defended, detected, remediated, and recovered so that all of these can be improved to continuously increase the resilience of the network. This is an ongoing process that requires that the network infrastructure, protocols, and resilience mechanisms be evolvable.

## III. Topology Generation

A key aspect of understanding and analysing network resilience is to accurately represent the topology of the existing network, as well as to be able to generate representative alternative topologies to evaluate resilience properties, and to be the basis of comparing candidate mechanisms.

The majority of topology modelling is based on *logical* topologies focusing on the generation of either router-level or AS-level topologies [16], motivated by the desire to study Internet layer-3 connectivity and protocols such IP, BGP, and IGPs, constrained by the fact that the majority of inference mechanisms [17] are only able to collect data on the the router-level connectivity of commercial networks. However, a router-level topology is an abstraction of the underlying physical topology and not an exact representation. Links visible to layer 3 are logical interconnections consisting of multiple physical links between layer 2 and layer 1 components such as switches, multiplexers, regenerators, and optical amplifiers. Furthermore, layer 3 topologies are frequently not representative of the underlying infrastructure due to layer 2.5 technologies such as MPLS, SONET, and Metro Ethernet that permit rearrangement of paths for traffic engineering, policy, and restoration. Thus it is possible for two distinct IP paths to share the physical same link, making it difficult

to understand and engineer the resilience of the network by assuming that IP links correspond to physical links. If we can not understand the geographic location of physical network nodes and links we will not know if they share fate, as was the case in the Baltimore tunnel fire [18] in which many logically distinct links failed at the same time. The next step is to model the overlay logical topologies, with a key challenge of understanding the relationships of each topology level, in part to avoid the problems of shared fate; this is a generalisation of the concept of shared-risk link groups (SRLGs) [19].

We argue that resilience evaluation of a network must begin with the *physical topology* because it ultimately determines the network's ability to survive infrastructure failures. The service and overall network dependability and performability in the face of failures is highly dependent on physical structure.

Thus, a key piece of our resilience evaluation strategy is to have a flexible and realistic topology generation tool that reflects the hierarchical structure of the networks including the differing topological characteristics at each level, with the ability to *geographically* place nodes, constrained by cost, population density and technology penetration, and availability of network infrastructure such as the fibre optic plant. We have implemented a topology generator KU-LoCGen (KU Location and Cost-Constrained Topology Generator), described in the following subsections.

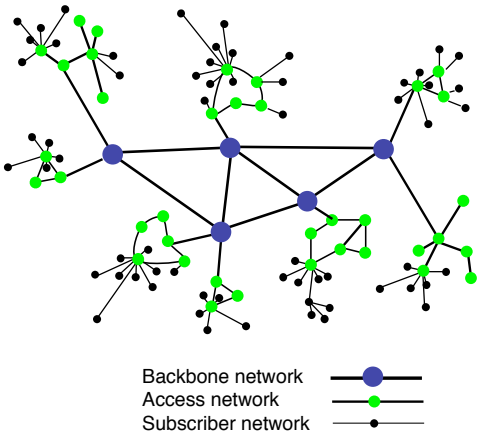


Fig. 2. Hierarchical topology model

### A. Hierarchical Topology Model

A topology generation model should be representative of the actual network structure and evolution process. While analytical models such as power law [20] are important for understanding mathematical properties of the network graph, they are not necessarily representative of real-world hierarchical deployments [21], [22], and diferent modelling methods are needed to incorporate real-world optimisations and tradeoffs (e.g. [23]–[25]). Our goal is to provide a flexible framework that allows for a hierarchical, modular structure with level-specific graphs, constrained by cost, population, and infrastructure location. The graph models used at each level vary from closed-form general-purpose mesh models (e.g. Waxman [26]) typical in backbones, to pre-structured models such as rings and trees typical in access networks based on particular technologies such as SONET rings and HFC (hybrid-fibre coax) trees, to a modified power-law preferential attachment of subscribers to access networks. Figure 2 shows an example of a topology modelled by KU-LoCGen representing a mesh backbone at level 1, various access topologies at level 2, and preferential attachment of subscribers at level 3.

### B. Location Constraints

The physical topology of networks is highly constrained by the *geographic location* of its components. It has also been observed that the router-level topology shows a very high correlation to the population density [27]. Moreover, the probability of link deployment is strongly related to the distance between the nodes. It has been shown that the geographic distance-based models such as Waxman accurately model the link distribution when considering location constraints [27].

Furthermore, the ability to model area-based challenges such as large-scale disasters depends on geographic node placement rather than the random placement of traditional topology generators. Examples of applying area-based challenges to geographic topology models will be shown in Section V. Our ultimate goal is to understand the graph-theoretic properties that relate to network resilience, including spatial diversity that requires node geolocation information.

We define a basic measure of diversity that quantifies the degree to which alternate paths share the same nodes and links [14], [28], and are enhancing it to incorporate geographic diversity, measured as the minimum distance between any pair of nodes along alternate paths, and as the area inside a polygon defined by a pair of alternate paths. Thus, it should be possible to specify diverse paths among a set of candidates with a given degree of sharing and distance metric (*effective path diversity*) constrained by stretch, as well as to measure the diversity inherent in a graph across all paths (*total graph diversity*).

Generating topologies with location constraints can be done in two ways. We can use the known location of existing infrastructure to geographically place nodes (for example Rocketfuel [29] for backbone node placement that generally corresponds to PoP locations). In this case we can synthetically generate links under cost constraints, as described later. Alternatively, we can use population density to drive node placement, as described next.

### C. Population Constraints

The physical topologies of networks are highly constrained by the geographic location of its components, which in turn are determined by two factors. The location of nodes is determined primarily by the population centres that links connect. The paths of links are further constrained by topographic features the minimise the deployment cost of fibre-optic cables; long-distance runs are typically laid along railways and motorways.

One of the goals of our geographically-constrained topology generator is to use realistic constraints to *deduce* node placement. This can be used either to compare the resilience

of existing networks to alternatives in developed areas such as the US and Europe, or to predict where new infrastructure should be deployed in developing nations.

We use the *k-means* clustering algorithm on the 1 km$^2$ gridded population density data sets from CIESIN [30] to determine optimal locations to place a backbone PoP [31].

The two inputs to our algorithm are the population data and the number of cluster centers. From the inferred topologies obtained from Rocketfuel [29], we note the number of PoPs for various ISPs and geographical areas. For example, the Sprint backbone has 27 nodes spatially distributed across the USA. We use this number as the input to our algorithm and generate an equal number of population centers. We consider multiple ISPs so that we can aggregate across major tier-1 providers, so as to not neglect certain parts of a country that may not be serviced by a specific ISP. Figure 3 shows a comparison of 112 PoPs generated using our population based model with the existing combined ISP PoPs of Sprint, AT&T, and Level3.

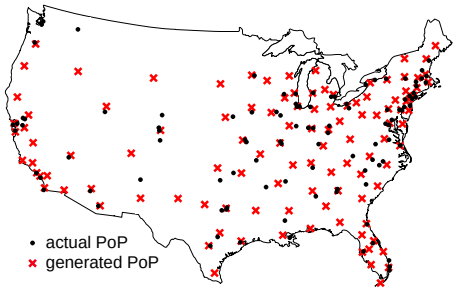


Fig. 3. Relative node locations for combined ISPs in USA

We quantify the distance between inferred PoP locations and population based cluster centers as the *offset distance* for a pair of nodes. The complementary cumulative distribution function (CCDF) of the offset distance for individual and combined ISPs is shown in Figure 4. We note that when we combine ISPs, almost 90% of the nodes generated by our algorithm are within 50 km offset distance

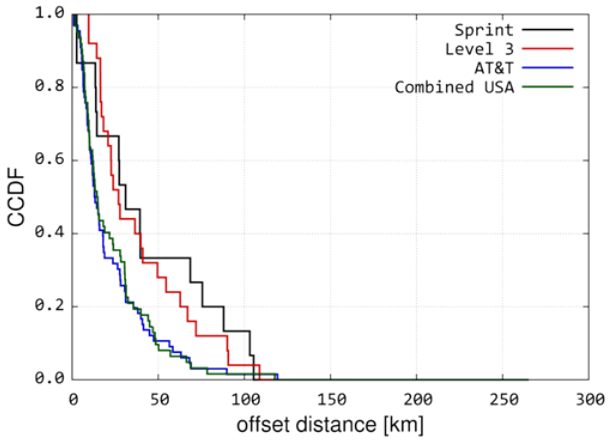


Fig. 4. CCDF of offset distance

*1) Technology Penetration:* The other fundamental aspect governing the location of the PoPs is technology penetration. The location of the backbone PoPs is highly dependent on the number of Internet users in a given area. We denote the technology penetration factor as $\gamma$, defined as the fraction of *Internet* users to the total population in a particular area. We assume this factor is uniform for a developed regions such as the US and Europe, for which we consider $\gamma$=1. This factor particularly has significant influence on a developing country such as India, where technology penetration is not homogeneous in all areas. We incorporate technology penetration into our model by weighting the population of each grid in an area by corresponding $\gamma$ and then clustering the resulting data set. Figure 5 shows the impact of technology penetration factor on the VSNL network in India [29], and the improvement over pure-population clustering is clearly visible for Mumbai, Kolkata, and Hyderabad.

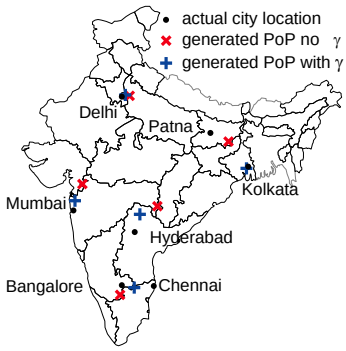


Fig. 5. Illustration of $\gamma$ factor for India

*2) Link Path Constraints:* Given the prediction of major nodes determined by population distributions, actual node placement should be further influenced by the location of existing network infrastructure, including fibre routes. To model this infrastructure and potential new deployment opportunities, we are currently adding existing fibre paths, railway mainlines, and Interstate freeways to our US adjacency matrices. This will permit us to add the additional step of "snapping-to-grid" nodes to infrastructure, and should improve the accuracy of node placement. For example, in Figure 3 there are a number of nodes in the sparsely-populated Western US that would snap to larger cities at fibre junctions and be located even more closely to existing PoP cities.

### *D. Cost-Constrained Connectivity*

Given a set of node locations, either based on existing networks or predicted as discussed in the previous subsection, we want to explore the resilience of alternative interconnection topologies. This only makes sense under realistic cost constraints, otherwise all networks would be full meshes – maximum resilience can be obtained with unlimited cost, but this is not practical. Therefore, our model uses a cost constrained connectivity models to generate feasible topologies.

Given the impracticability of a universal cost function, we use modular cost functions that are highly tunable and

allow network designers to select as well as define new cost functions based on fundamental variables such as fixed and variable costs per link and per node. Our baseline model assumes that the cost of all nodes in the backbone network is the same $C_{\mathrm{b}}$. The link cost $C_{i,j}$ of a link $i, j$ is calculated as $C_{i,j} = f_{i,j} + v_{i,j} \times d_{i,j}$ where $f_{i,j}$ is the fixed cost associated with terminating the link, $v_{i,j}$ is the variable cost per unit distance for link deployment, and $d_{i,j}$ is the length of the link. For simplicity we choose $v_{i,j} = \bar{d} \times v_{i,j}$ where $\bar{d}$ is the average link length of the network. The level-1 nodes in our model are connected using a cost-constrained Waxman model, which is a reasonable representation of link connectivity in a backbone network [27]. Figure 6 shows an example level-2 topology generated by our model using the 27-node topology (equal to the number of Sprint PoPs) with population-based node clustering and random node placement about the PoPs for the 2nd level. The objective is to be able to generate alternative realistic topologies to compare their resilience with one another as well as against existing network deployment. This motivates the need to for metrics and a methodology to *quantify* resilience, described in the next section.

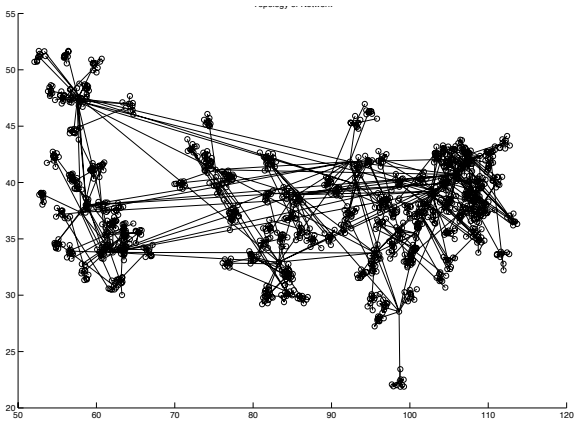


Fig. 6. Sample 2-level topology using 27 nodes

## IV. Analytical Resilience Framework

This section describes a new analytical framework to evaluate network resilience based on a two-dimensional state space. There have been only a few attempts to quantify resilience (or survivability) in particular contexts, such as for large-scale disasters [32] and for distributed denial of service (DDoS) attacks [33]. The T1A1.2 working group defines survivability [34] based on availability and network performance, which has then been used to quantify survivability using multidimensional Markov chains [35].

### A. Metrics Framework

Recall that we define resilience at the ability of the network to provide a certain service under the presence of challenges, manifest as adverse events and conditions to normal operations [2]. The major complexity in resilience evaluation comes from the varied nature of services that the network provides, the numerous layers and their parameters over which these services depend, and the plethora of adverse events that present as challenges to the network as a whole. This complexity renders an exhaustive cross-product analysis of all aspects of the resilience analysis intractable. Therefore, our approach abstracts and simplifies the resilience evaluation process in three ways: First, we isolate the impact of challenges at each level by evaluating resilience at each layer boundary (physical, link, topology, network path, end-to-end transport, and application) thereby avoiding a continually increasing parameter set of the network protocol stack as a whole. Secondly, we quantify resilience as a change in service corresponding to a change in the operational state at any given layer [36], [37]. Finally we choose representative scenarios of applications and network deployments. Therefore resilience is characterised as a mapping between the operations and service, wherein the operations are affected by the challenges, which in turn result in degradation of the service at that layer boundary. In other words, instead of evaluating the impact of each challenge or attack separately, which leads to intractable number of cases, we focus on quantifying the service given perturbations in the underlying operational conditions.

*1) Operational State:* The first step in our framework is to quantify the operational state at any layer given a set of *operational metrics*. For a given system $\mathcal{S}$, where the system refers to the network at an arbitrary level, let the $\ell$ operational metrics be represented as $N_{\mathcal{S}} = \{N_1, \ldots, N_\ell\}$. Each operational metric $N_i, 1 \leq i \leq \ell$, is in itself a set of $m$ values, representing all possible settings of the particular operational metric, $N_i = \{n_{i,1}, \ldots, n_{i,m}\}$. The *operational state space* of $\mathcal{S}$ is $\mathcal{N}_{\mathcal{S}} = \times_i N_i$ ($\times$ is the cross product). Thus, the operational state space consists of all possible combinations of the operational metrics. We define an *operational state*, $\mathbb{N}$ as a subset of the complete state space $\mathcal{N}_{\mathcal{S}}$.

*2) Service State:* The second step is to characterise the service provided at a given network layer. The *service parameters* represent the requirements of the service that is being provided across the service interface. Let the $\ell$ service parameters of system $\mathcal{S}$ be represented by $P_{\mathcal{S}} = \{P_1, \ldots, P_\ell\}$. Each service parameter $P_i, 1 \leq i \leq \ell$, is in itself a set of $m$ values (representing all possible values of the particular service parameter), $P_i = \{p_{i,1}, \ldots, p_{i,m}\}$. The *service state space* of $\mathcal{S}$ is $\mathcal{P}_{\mathcal{S}} = \times_i P_i$. Therefore, the service state space consists of all possible combinations of the service parameters. We define *service state*, $\mathbb{P}$, as a subset of the complete state space $\mathcal{P}_{\mathcal{S}}$.

*3) Network State:* The operational and service states described above represent the state of the network at any given time. Therefore, we define the overall *state* $S_{\mathcal{S}}$ of the system $\mathcal{S}$, as a tuple of operational state and service state: $(\mathbb{N}, \mathbb{P})$. Therefore the $k^{\text{th}}$ network state $S_k = (\mathbb{N}_k, \mathbb{P}_k)$. The network state represents a mapping between the operational state space $\mathcal{N}_{\mathcal{S}}$ and service state space $\mathcal{P}_{\mathcal{S}}$. Furthermore, this mapping is an onto mapping, meaning that for every service state there is an operational state.

Note that both the operational and the service state spaces are multivariate. In order to visualise this state space on a two

dimensional state space, we project both the operational state space and service state space on to one dimension each. This projection is achieved via objective functions that are applied to both the state spaces. The specific function used depends on the scenario. For example, it may be a linear combination with normalised weights or logical functions (e.g., AND, OR), for example a minimum goodput $\wedge$ maximum delay.

*4) Resilience Quantification:* We formulate that challenges in the form of adverse events transform the network from one *state* to another based on the severity of the event. Network resilience can be evaluated in terms of the various network states transitions under the presence of challenges. Resilience $\mathbb{R}_{ij}$ is defined at the boundary $B_{ij}$ between any two adjacent layers $L_i, L_j$. Resilience $\mathbb{R}_{ij}$ at the boundary $B_{ij}$ is then evaluated as the transition of the network through this state space. The goal is to derive the $\mathbb{R}_{ij}$ as a function of $\mathbb{N}$ and $\mathbb{P}$, measured as the area under the state-space trajectory. The operational and service space is covered fully by its states and can be decomposed in a fixed set of large states termed *regions*: The network operational space is divided into *normal*, *partially degraded*, and *severely degraded* regions as shown in Figure 8. Similarly, the service space is divided into *acceptable*, *impaired*, and *unacceptable* regions.

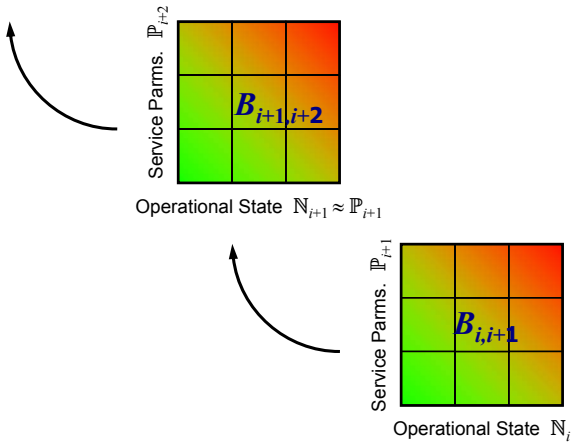


Fig. 7. Resilience across multiple levels

*5) Multilevel Resilience:* In the multilevel analysis, shown in Figure 7, the service parameters at the boundary $B_{ij}$ become the operation metrics at boundary $B_{i+1,j+1}$. In other words, the service provided by a given layer becomes the operational state of the layer above, which has a new set of service parameters characterising its service to the layer above.

### B. Relationship to the Strategy

The relationship of the the state-space formulation to the ResiliNets strategy described in Section II is shown in Figure 8 with the inner $D^2R^2$ loop trajectory shown. Defence prevents the system from leaving its initial state $S_0$. If a challenge causes the state to change significantly, this is detected by a change in the operational or service parameters when the state goes to a challenged state $S_c$. Remediation improves the

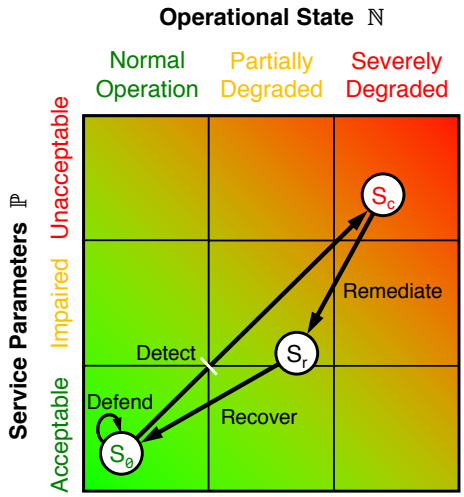


Fig. 8. Resilience state space and strategy inner loop

situation to $S_r$, and recovery finally returns the system to its original state $S_0$.

The outer control loop reduces the impact of a given challenge in the future, as shown in Figure 9, in which the challenged state $S_c\prime$ is not as bad as the previous $S_c$, and remediation performs better with $S_r\prime$ resulting in a smaller area under the trajectory (shaded) and better overall resilience.

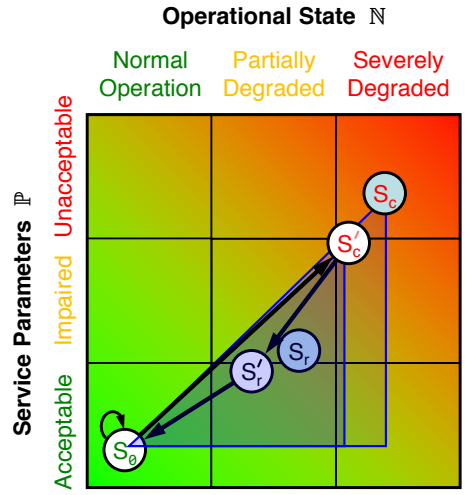


Fig. 9. $D^2R^2$ strategy and resilience state space

### C. Example Analysis of Topology

In order to demonstrate the application of this framework, we conduct the resilience analysis of a few example ISP networks at the topology layer $(3t)$ wherein the objective is to study the impact of node and link failures on the topology. In this case, a set of vertices $V$ and edges $E$ and link failures $f$ characterise the operational state of the network. The service provided by this layer is *topological connectivity*. Since we

consider only link failures, we choose a single operational metric $n_1$ to represent the number of link failures. We define the *topology service* by selecting two parameters: the relative size of the largest connected component $p_1$ that represents the reachability of the graph, and clustering coefficient $p_2$ representing the local richness of the topology. We conducted simulations in MATLAB to evaluate the impact of link failures on the service parameters, averaged over 100 runs. We use three service provider backbone network topologies: Sprint (US), AT&T (US), (both inferred from [38]) and GÉANT2 [39].

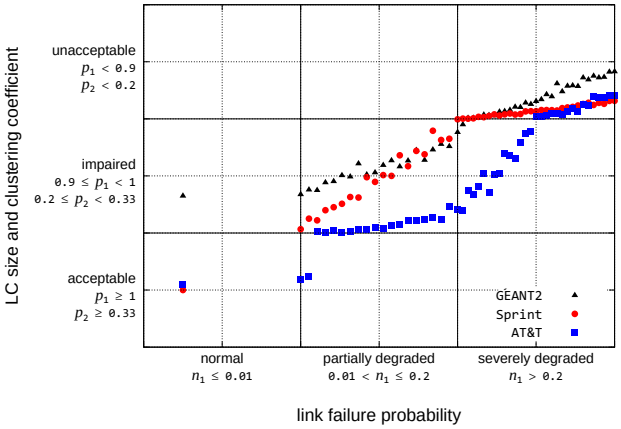


Fig. 10. Comparing resilience of ISP topologies

We plot the steady state resilience of three networks to link failures on a piece-wise linear axis as shown in Figure 10. The AT&T and Sprint topologies are in the acceptable service region under normal operating conditions with no link failures. Depending upon the connectivity of the network, the service may remain acceptable even when the network starts degrading. However, as link failures increase the topology moves toward impaired service on the vertical axis. As the network operational conditions are severely degraded the service transitions from impaired to unacceptable regions. In order to get the aggregate measure of resilience for each topology, we calculate the area under the curve formed by linear interpolation between the states. In order to get a normalised value of resilience, we define resilience $\mathbb{R} = 1 -$ normalised area, where normalised area is the total area divided by the span of the $x$-axis. For the plot shown in Figure 10, the resilience $\mathbb{R}$ for AT&T is 0.6338, Sprint is 0.5410, and GÉANT2 is 0.4721. In this case, we observe that AT&T has better resilience compared to Sprint (comparable US backbones), as well as the GÉANT2 topology. We note that the GÉANT2 Euopean research topology has a very low clustering coefficient and is not biconnected, thus this result confirms expectations.

## V. Simulation Methodology

The previous section presented an analytical framework for resilience analysis, and presented an example of its application to analysis at the link/topology level boundary. At upper levels of the network and for more complex challenge scenarios, simulation is necessary to obtain tractable results. The results of these simulations can still be applied to the state space framework [37], with the operational dimension as simulation parameters and the service dimension as the simulation output. This section describes a simulation framework and methodology that can be used to understand the impact of challenges more complex than random link and node failures, given a variety of simulated protocols and application traffic.

### A. Simulation Framework

The goal of KU-CSM (KU Challenge Simulation Module) [40] is to provide a modular framework that can be used to investigate the resilience of a number of network scenarios $n$ to a variety of challenges $c$. Traditionally, the type of challenge has been part of a monolithic simulation model, making it difficult to test the effects of different challenges. In the worst case this results in $n \times c$ models to explore the entire cross-product of network and challenge scenarios.

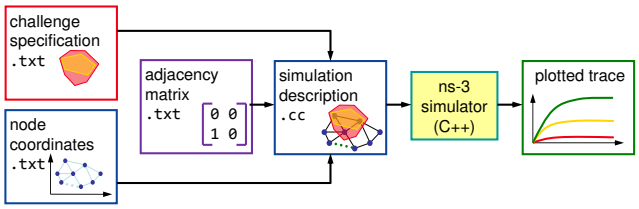


Fig. 11. Framework flow diagram

KU-CSM decouples the network topology from the challenge type, as shown in Figure 11, reducing the number of models to $n + c$ distinct network models and challenges. We use the ns-3 [41] network simulator since it is open source, flexible, provides mixed wired and wireless capability (unlike ns-2), and is modular in its C++ implementation.

The first step is to provide a challenge specification that includes the type of the challenge and specifics of the challenge type. The second step is to provide a description of the network topology, consisting of node geographical or logical coordinates and an adjacency matrix. This can either be based on an actual network deployment, or be the output of the KU-LoCGen topology generator described in Section III. The third step is the automated generation of ns-3 simulation code based on the topology and challenge descriptor. Finally, we run the simulations and analyse the network performance under challenge scenarios given protocols and application traffic.

### B. Challenge Characteristics

Understanding network disruptions and their cause is crucial for planning and designing the networks. The redundancy and diversity that increase resilience add to the cost of the network. Therefore, we need to understand the types of challenges and their impact on network operation and the service delivered to users, in order to understand which of the alternative structures and mechanisms will actually improve resilience.

A *challenge* is an event that impacts normal operation [2]. A challenge triggers *faults*, which are the hypothesised cause of errors. Eventually, a fault may manifest itself as an error. If the

error propagates it may cause the delivered services to fail [8]. Challenges to the normal operation of networks include unintentional misconfiguration or operational mistakes, malicious attacks, large-scale natural disasters, and environmental challenges [2], [11]. Network challenges can be categorised based on intent, scope, and domain they impact [40].

We model the challenges based on the intent as non-malicious or malicious. Non-malicious challenges can be due to incompetence of an operator or designer. These random events affect node and link availability, and result in the majority of the failures observed [42], [43]. On the other hand, malicious attacks, orchestrated by an intelligent adversary, target *specific* parts of a network and can have significant impact if critical elements of the network fail.

The scope of a challenge can be further categorised based on nodes, links, or network elements affected within a geographic area. Hurricanes, earthquakes, and solar storms are examples of natural disasters that can impact the network at a large scale [44]. Furthermore, geographically correlated failures can result from dependency among critical infrastructures, as experienced in the 2003 Northeast US blackout.

The challenges that are inherent in the wireless domain include weakly connected channels, mobility of nodes in an ad-hoc network, and unpredictably long delays [11]. These are the natural result of noise, interference, and other effects of RF propagation such as scattering and multipath, as well as the mobility of untethered nodes. Furthermore, weather events such as rain and snow can cause the signals to attenuate the wireless communication network [15]. Malicious nodes may jam the signal of legitimate users to impair communication in the open wireless medium.

While the above-mentioned challenge models are orthogonal to each other, challenge scenarios are a combination of challenge sub-categories. For example, a failure due to natural aging of a component can be categorised as a non-malicious, wired (or wireless), node failure.

### C. Example Simulation Analysis

In this section, we apply our challenge framework and evaluation methodology to an example topology to demonstrate the utility of this approach. We use the inferred Sprint backbone network topology of 27 nodes and 68 links [29], shown in Figure 13. A full explanation of the challenge specifications, as well as details of simulation parameters and further example results are presented in [40].

*1) Non-malicious and Malicious Challenges:* First, we evaluate the performance of the Sprint topology (Figure 13) under the presence of malicious and non-malicious failures of up to 10 nodes or links, with the packet delivery ratio (PDR) shown in Figure 12. We measure the instantaneous PDR at the steady-state condition during the challenges for each point.

The top curve in Figure 12 shows the PDR with random link failures. In this case for 10 random link failures averaged over 100 runs, the PDR drops to 87%. The second curve from the top shows the PDR for link attacks. In this case, we first calculate the betweenness for each link in the topology,

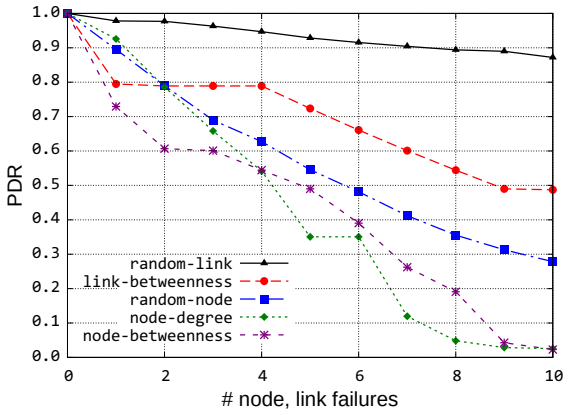


Fig. 12. PDR during non-malicious and malicious challenges

and provide the challenge file as the list of the links to be brought down. As can be seen, link attacks have more degrading impact than the random failures, 50% PDR for highest ranked 10 links. The middle curve shows random node failures, worse than link attacks or failures, since each node failure is equivalent of the failure of all links incident to that node. The bottom two curves show the PDR during node attacks based on degree of connectivity and betweenness; these are the most damaging attacks to the network. The primary difference between the two attack scenarios is that an attack based on betweenness can be more damaging for the few highest ranked nodes. When the highest betweenness two nodes in rank are attacked, PDR is reduced to 60%, while an attack based on degree of connectivity only reduces the PDR to 80%. This example confirms the intuition that attacks launched with knowledge of the network topology can cause the most severe damage.

*2) Area-based Challenges:* Recently, the research community has recognised the importance of understanding the impact of geographically correlated failures on the networks [40], [45]–[48]. Our framework uses circles or polygons to model geographically correlated failures representative of large-scale disasters needed for network survivability [10], [11]. Next, we present the results of three scenarios that demonstrate area-based challenges that evolve spatially and temporally using the same Sprint topology, as shown in Figure 13. Application traffic is generated from 2 to 29 sec. and challenge scenarios were applied from 10 until 22 sec. for the plots as shown in Figure 14, which verify the impact of the example challenges.

To demonstrate a scaling circle area-based challenge scenario, we simulate a circle centered at in New York, USA as shown in Figure 13a, with a radius of approximately 111 km. We choose the scenario to be representative of an electromagnetic pulse (EMP) attack [49]. The PDR is shown in Figure 14a. We choose the simulation parameters such that the radius doubles in every 4 sec. As can be seen, the PDR reduces as the circular area doubles. The PDR drop depends on how many nodes and links are covered in each step.

Next, we demonstrate an area-based scenario that can evolve

(a) Scaling circle (b) Moving circle (c) Scaling polygon

Fig. 13. Area-based challenge scenarios

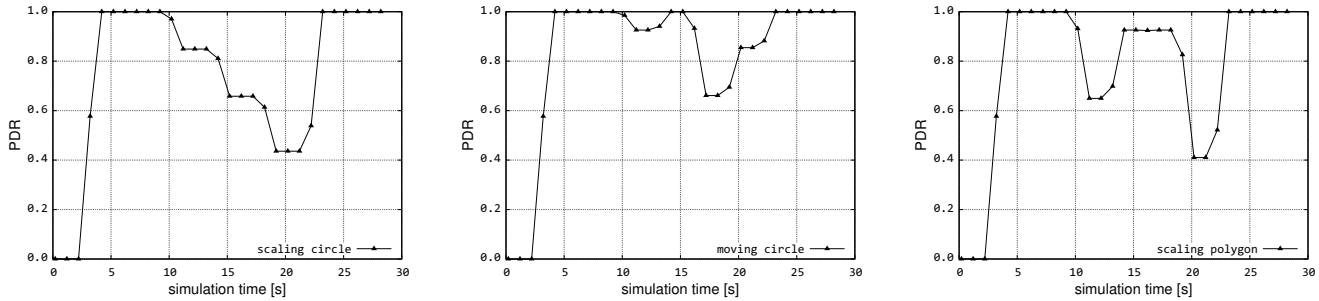


(a) Scaling circle PDR (b) Moving circle PDR (c) Scaling polygon PDR

Fig. 14. PDR during area-based challenges

spatially and temporally. We simulate a moving circle in a trajectory from Orlando to New York that might model a severe hurricane, but with rapid restoration of links as the challenge moves out of a particular area. Three snapshots of the evolving challenge are shown in Figure 13b. The radius of the circle is kept at approximately 222 km. We choose the simulation parameters for illustration such that the circle reaches NY in seven seconds (to constrain simulation time), with route recomputation every 3 sec.

As shown in Figure 14b PDR reduces to 93% as the challenge starts only covering the node in Orlando at 10 sec. As the challenge moves towards NY in its trajectory, the PDR reaches one at the 13 sec. In this case, the challenge area includes only the link between Orlando and NY, but since there are multiple paths a single link failure does not affect the PDR, showing that *geographic diversity for survivability* is crucial [2]. As the challenge moves into the northeast US region at 16 sec., the PDR drops to 66% as the challenge covers several nodes and links. The simulation shows that as the circle moves out of the more crowded region of the network, the PDR improves

Polygons are useful to model specific geographic challenges such as power failures. For a scaling polygon example, we show a 6-sided irregular polygon in the Midwest region of the US, roughly representative of the North American Electric Reliability Corporation (NERC) Midwest region [49], as shown in Figure 13c.

The PDR throughout the simulation is shown in Figure 14c. In this scenario, the edges of the irregular polygon increase 1.8 times every three sec. At 10 sec. the challenge affects 16 links, which causes the PDR to drop to 65%. The PDR then increases to 93%, even though more links and nodes are affected at 13 sec. because of route reconvergence. As the polygon increases in size, the PDR drops to as low as 41%, because the challenge area partitions the network at 21 sec. This type of scenario can be used either to understand the relationship between the area of a challenge and network performability, or to model a temporally evolving challenge, such as a cascading power failure that increases in scope over time.

## VI. Summary

Resilience is an essential property of the Future Internet, including performability, dependability, and survivability. This requires metrics to quantify resilience, and a methodology to evaluate the resilience of current networks as well as alternative topologies and mechanisms that are candidates for deployment in the Future Internet. This paper has described a comprehensive framework consisting of a resilience strategy, metrics for quantifying resilience, and evaluation techniques, and provided example results from our ongoing research. We believe that we have shown the potential for these techniques to help gain insight on the resilience analysis of current networks, and to evaluate the benefits of proposed architectures, mechanisms, and protocols for the Future Internet. Furthermore, we plan to extend the methodology to emulation using the international GpENI [50] programmable testbed to cross-verify with the analytical and simulation models, and to evaluate real network resilience at scale [51].

## ACKNOWLEDGMENTS

The authors would like to thank members of the ResiliNets research group at the University of Kansas and Lancaster University, as well as members of the EU ResumeNet project for discussions on aspects of this work. In particular we acknowledge Shi Qian at KU, David Hutchison and Paul Smith at Lancaster, and Marcus Schöller of NEC Laboratories. This research was supported in part by the the National Science Foundation FIND (Future Internet Design) Program under grant CNS-0626918 (Postmodern Internet Architecture), by NSF grant CNS-1050226 (Multilayer Network Resilience Analysis and Experimentation on GENI), and the European Commission under grant FP7-224619 (ResumeNet). We mourn the recent passing of Jean-Claude Laprie, whose seminal work in dependability is an important foundation for this work.

## REFERENCES

[1] "Protecting America's infrastructures," President's Commission on Critical Infrastructure Protection, Report, October 1997.

[2] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines," *Computer Networks*, vol. 54, no. 8, pp. 1245–1265, June 2010.

[3] "A roadmap for cybersecurity research," Department of Homeland Security (DHS), Technical Report, November 2009.

[4] S. Goodman and H. Lin, *Toward a Safer and More Secure Cyberspace*. National Academies Press, 2007.

[5] F. Schneider, *Trust in Cyberspace*. National Academies Press, 1999.

[6] (2010, February) UK resilience homepage. http://www.cabinetoffice.gov.uk/ukresilience.aspx.

[7] (2010, February) European information society. http://ec.europa.eu/information_society/policy/nis/strategy/activities/ciip/index_en.htm.

[8] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, "Basic concepts and taxonomy of dependable and secure computing," *IEEE Trans. on Dependable and Secure Computing*, vol. 1, no. 1, pp. 11–33, 2004.

[9] J. Meyer, "On Evaluating the Performability of Degradable Computing Systems," *IEEE Trans. on Comp.*, vol. 100, no. 29, pp. 720–731, 1980.

[10] R. J. Ellison, *et al.*, "Survivable network systems: An emerging discipline," CMU, CMU/SEI-97-TR-013, 1999.

[11] J. P. G. Sterbenz, *et al.*, "Survivable mobile wireless networks: issues, challenges, and research directions," in *ACM WiSe*, 2002, pp. 31–40.

[12] N. Edwards, "Building dependable distributed systems," ANSA, Technical report APM.1144.00.02, February 1994.

[13] T1A1.2 Working Group, "Network survivability performance," ATIS, Technical Report T1A1.2/93-001R3, November 1993.

[14] J. P. Rohrer, A. Jabbar, and J. P. G. Sterbenz, "Path diversification: A multipath resilience mechanism," in *IEEE DRCN*, October 2009, pp. 343–351.

[15] A. Jabbar, J. P. Rohrer, A. Oberthaler, E. K. Çetinkaya, V. Frost, and J. P. G. Sterbenz, "Performance comparison of weather disruption-tolerant cross-layer routing algorithms," in *IEEE INFOCOM*, April 2009, pp. 1143–1151.

[16] A. Medina, I. Matta, and J. Byers, "On the origin of power laws in Internet topologies," *SIGCOMM CCR*, vol. 30, no. 2, pp. 18–28, 2000.

[17] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network topologies: inference, modeling, and generation," *IEEE Communications Surveys and Tutorials*, vol. 10, no. 2, pp. 48–69, 2008.

[18] H. C. Styron, "CSX tunnel fire: Baltimore, MD," Federal Emergency Management Administration, Emmitsburg, MD, US Fire Administration Technical Report USFA-TR-140, 2001.

[19] S. Chaudhuri, G. Hjalmtysson, and J. Yates, "Control of lightpaths in an optical network," Optical Internetworking Forum OIC2000.04, IETF Internet Draft, January 2000.

[20] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos, "Power laws and the AS-level Internet topology," *IEEE/ACM ToN*, vol. 11, no. 4, pp. 514–524, Aug. 2003.

[21] K. Calvert, M. Doar, and E. Zegura, "Modeling Internet topology," *IEEE Communications Magazine*, vol. 35, no. 6, pp. 160–163, Jun 1997.

[22] J. M. Carlson and J. Doyle, "Highly optimized tolerance: A mechanism for power laws in designed systems," *Phys. Rev. E*, vol. 60, no. 2, pp. 1412–1427, Aug 1999.

[23] D. Alderson, J. Doyle, R. Govindan, and W. Willinger, "Toward an optimization-driven framework for designing and generating realistic Internet topologies," *SIGCOMM CCR*, vol. 33, no. 1, pp. 41–46, 2003.

[24] J. C. Doyle, *et al.*, "The "robust yet fragile" nature of the internet," *PNAS*, vol. 102, no. 41, pp. 14 497–14 502, 2005.

[25] C. Wang and J. W. Byers, "Generating representative ISP topologies from first-principles," in *ACM SIGMETRICS*, 2007, pp. 365–366.

[26] B. Waxman, "Routing of multipoint connections," *IEEE JSAC*, vol. 6, no. 9, pp. 1617–1622, Dec 1988.

[27] A. Lakhina, *et al.*, "On the geographic location of Internet resources," *IEEE JSAC*, vol. 21, no. 6, pp. 934–948, 2003.

[28] J. P. Rohrer, R. Naidu, and J. P. G. Sterbenz, "Multipath at the transport layer: An End-to-End resilience mechanism," in *IEEE RNDM*, St. Petersburg , Oct 2009, pp. 1–7.

[29] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP Topologies with Rocketfuel," *IEEE/ACM ToN*, vol. 12, no. 1, pp. 2–16, Feb 2004.

[30] CIESIN, Gridded Population of the World Version 3 (GPWv3): Population Density Grids. SEDAC, Columbia University. [Online]. Available: http://sedac.ciesin.columbia.edu/gpw

[31] M. A. Hameed, A. Jabbar, E. K. Çetinkaya, and J. P. Sterbenz, "Deriving network topologies from real world constraints," in *IEEE GLOBECOM CCNet Workshop*, 2010, to appear.

[32] S. Liew and K. Lu, "A framework for network survivability characterization," in *IEEE ICC*, 1992, pp. 405–410.

[33] S. Hariri, *et al.*, "Impact analysis of faults and attacks in large-scale networks," *IEEE Security and Privacy*, vol. 01, no. 5, pp. 49–54, 2003.

[34] T1A1.2 Working Group, "Enhanced network survivability performance," ATIS, Technical Report T1.TR.68-2001, February 2001.

[35] K. Trivedi, D. Kim, A. Roy, and D. Medhi, "Dependability and security models," in *IEEE DRCN*, 2009, pp. 11–20.

[36] A. J. Mohammad, D. Hutchison, and J. P. G. Sterbenz, "Towards quantifying metrics for resilient and survivable networks," in *IEEE ICNP*, November 2006, pp. 17–18.

[37] A. Jabbar, "A framework to quantify network resilience and survivability," Ph.D. dissertation, The University of Kansas, May 2010.

[38] (2008, September) Rocketfuel: An ISP topology mapping engine.

[39] (2009, December) GÉANT2. http://www.geant2.net/.

[40] E. K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, and J. P. Sterbenz, "A Comprehensive Framework to Simulate Network Attacks and Challenges," in *IEEE RNDM*, Moscow, October 2010.

[41] (2009, July) The ns-3 network simulator. http://www.nsnam.org.

[42] D. Kuhn, "Sources of failure in the public switched telephone network," *Computer*, vol. 30, no. 4, pp. 31–36, April 1997.

[43] D. Oppenheimer, A. Ganapathi, and D. A. Patterson, "Why do Internet services fail, and what can be done about it?" in *Proc. of USENIX USITS*, 2003, pp. 1–16.

[44] Y. Kitamura, Y. Lee, R. Sakiyama, and K. Okamura, "Experience with restoration of asia pacific network failures from taiwan earthquake," *IEICE Trans. on Comm.*, vol. E90-B, no. 11, pp. 3095–3103, 2007.

[45] R. Mahmood, "Simulating challenges to communication networks for evaluation of resilience," MS thesis, The University of Kansas, Aug. 2009.

[46] B. Bassiri and S. S. Heydari, "Network survivability in large-scale regional failure scenarios," in *ACM C3S2E*, 2009, pp. 83–87.

[47] S. Neumayer, G. Zussman, R. Cohen, and E. Modiano, "Assessing the vulnerability of the fiber infrastructure to disasters," in *IEEE INFOCOM*, 2009, pp. 1566–1574.

[48] S. Neumayer and E. Modiano, "Network reliability with geographically correlated failures," in *IEEE INFOCOM*, 2010, pp. 1–9.

[49] "Report of the commission to assess the threat to the United States from electromagnetic pulse (EMP) attack," Critical National Infrastructures, Report, 2004.

[50] J. P. G. Sterbenz, *et al.*, "The Great plains Environment for Network Innovation (GpENI): A programmable testbed for Future Internet architecture research," in *TridentCom*, Berlin, May 2010.

[51] J. P. Sterbenz, J. P. Rohrer, and E. K. Çetinkaya, "Multilayer network resilience analysis and experimentation on GENI," KU, ITTC Technical Report ITTC-FY2011-TR-61349-01, July 2010.