

Evaluation of Network Resilience, Survivability, and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation

invited paper

James P.G. Sterbenz · Egemen K. Çetinkaya · Mahmood A. Hameed ·
Abdul Jabbar · Shi Qian · Justin P. Rohrer

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 future networks. 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 topology generation, analytical, simulation, and experimental emulation techniques with the goal of improving the resilience and survivability of the Future Internet.

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 FIRE (Future Internet Research and Experimentation Programme) under grant FP7-224619 (ResumeNet).

James P.G. Sterbenz
Electrical Engineering and Computer Science,
Information and Telecommunication Technology Center,
The University of Kansas, Lawrence, Kansas, USA
<http://www.ittc.ku.edu/resilinet>
jpgs@ittc.ku.edu, +1 785 864 7890
Computing Department, InfoLab21
Lancaster University, Lancaster, UK,
jpgs@comp.lancs.ac.uk

Egemen K. Çetinkaya
ekc@ittc.ku.edu

Mahmood A. Hameed
hameed@ittc.ku.edu

Abdul Jabbar
jabbar@ge.com

Shi Qian
shiqian@ittc.ku.edu

Justin P. Rohrer
rohrej@ittc.ku.edu

Keywords resilient survivable disruption-tolerant network · dependability performability · diverse topology generation · network analysis experimentation · ns-3 simulation methodology

1 Introduction and Motivation

The increasing importance of the Global Internet has lead to it becoming one of the critical infrastructures [2] 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 [133, 132]. 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 [14, 70, 126, 20, 17, 148]. 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, as well as 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 including reliability and availability [89] and *performability* that quantifies the level of performance [101] when the network is challenged. This notion of resilience subsumes *survivability* that is the ability to tolerate the correlated failures that result from attacks and large-scale disasters [134, 105, 59, 66] and *disruption-tolerance* that is the ability to communicate even when stable end-to-end paths may not exist due to weak channel connectivity, mobility, unpredictable delay, and energy constraints [134, 62].

This paper describes a comprehensive approach to evaluate network resilience through analysis, simulation, and experimentation, and is organised as follows: Section 2 reviews the ResiliNets architectural framework for network resilience based on a two-phase strategy and principles for resilient network design. Section 3 presents the problem of generating realistic topologies that can be used to evaluate network resilience, introduces the KU-LoCGen topology generator and KU-TopView, and discusses the issues of multi-level topology overlays. Section 4 describes an analytical formulation of resilience as the trajectory through a multilevel two-dimensional state space with operational and service dimensions and presents a few examples of this analysis. Section 5 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 with example simulation results. Section 6 briefly describes how the GpENI large-scale programmable testbed infrastructure can be used to experimentally validate and cross-verify with analytical and simulation-based resilience analysis. Finally, Section 7 summarises the main points of the paper and discusses future research in this area.

2 Resilience Framework, Strategy, Principles

This section reviews the ResiliNets framework for resilient, survivable, and disruption-tolerant network architecture and design [133,132]. First, a two-phase resilience strategy is described that provides the basis of the metrics framework presented in Section 4. Then, a set of design principles is presented with emphasis on heterogeneity, redundancy, and diversity that are used in the topology analysis in Section 3 and simulation methodology in Section 5.

2.1 ResiliNets Strategy

There have been several systematic resilience strategies, including ANSA [56], T1A1.2 [140], CMU-CERT [58], and SUMOWIN [134]. This ResiliNets resilience framework and strategy [133,132] are based in part on these previous frameworks and provides the basis for the resilience evaluation methodology described in the rest of the paper. More recently, the policy aspects of resilience mechanisms are being studied [130,125]. The framework begins with a set of four axioms that motivate the strategy:




Fig. 1 ResiliNets strategy

0. Faults are *inevitable*; it is not possible to construct perfect systems, nor is it possible to prevent challenges and threats¹.
1. **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.
2. **Expectation** and preparation for adverse events and conditions is necessary, so that that defences and detection of challenges that disrupt normal operations can occur. These challenges are inevitable.
3. **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. At the core are passive structural defences. The first active phase D^2R^2 : *defend, detect, remediate, recover*, is the inner control loop and describes a set of activities that are undertaken in order for a system to rapidly adapt to challenges and attacks and maintain an acceptable level of service. The second active phase DR : *diagnose, refine*, is the outer loop that enables longer-term evolution of the system in order to enhance the approaches to the activities of the inner loop. The following sections briefly describe the steps in this strategy.

2.1.1 D^2R^2 Inner Loop

The first strategy phase consists of a passive core and a cycle of four steps that are performed in real time

¹ The strict usage of fault, error, and failure terminology is based on [28] and fully explained in the ResiliNets context in [133].

and are directly involved in network operation and service provision. In fact, there is not just one of these cycles, but many operating simultaneously throughout and across the network for each resilient subsystem, triggered whenever an adverse event or condition is detected.

Defend against challenges and threats to normal operation. 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 consist of a passive and active component.

Passive defences are primarily structural, suggesting the use of trust boundaries, redundancy, diversity, and heterogeneity. The main network techniques are to provide geographically diverse redundant paths and 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 [123].

Active defences consist 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. Clearly, defences will not always prevent challenges from penetrating the network, which leads to the next strategy step: detect.

Detect when an adverse event or condition has occurred. The second step is for the network as a distributed system, as well as individual components such as routers, to detect challenges and to understand when the defence mechanisms have failed [67]. There are three main ways to determine if the network is challenged. The first of these involves understanding the service requirements and normal operational behaviour of a system and detecting deviations from it – *anomaly detection* based on metrics (described in Section 4). The second approach involves detecting when errors occur in a system, for example, by calculating CRCs (cyclic-redundancy checks) to determine the existence of bit errors that could lead to a service failure. Finally, a system should detect service failures; an essential facet of this is an understanding of service requirements. An important aspect of detecting a challenge is determining its nature, which requires context awareness.

Remediate the effects of the adverse event or condition. The next step is to remediate the effects of the detected adverse event or condition to minimise the impact on service delivery. The goal is to do the best possible at all levels after an adverse event and dur-

ing 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. [80]) and for adaptive applications and congestion control algorithms to degrade gracefully from acceptable to impaired service (Section 4).

Recover to original and normal operations. Once the challenge is over after an adverse event or the end of an adverse condition, the network may remain in a degraded state (Section 4). When the end of a challenge has been detected (e.g., a storm has passed, which restores wireless connectivity), the system must recover to its original optimal normal operation, since the network is likely not to be in an ideal state, and continued remediation activities may incur an additional resource cost.

2.1.2 D+R Outer Loop

The second phase 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 the fault that was the root cause. While it is not possible to directly *detect* faults, we may be able *diagnose* the fault that caused an observable error. In some cases this may be automated, but more generally it is an offline process of 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. An example of network-based fault diagnosis is the analysis of packet traces to determine a protocol vulnerability that can then be fixed.

Refine behaviour for the future based on past D^2R^2 cycles. 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 using the evaluation techniques described in this paper. This is an ongoing process that requires that the network infrastructure, protocols, and resilience mechanisms be evolvable. This is a significant challenge given the current Internet hourglass waist of

IPv4, BGP (border gateway protocol), and DNS (domain name system), as well as other mechanisms (e.g. NAT – network address translation) and protocol architectures (e.g. TCP and HTTP) that are entrenched and resist innovation.

2.2 Resilience Design Principles

The D²R²+DR strategy leads to a set of *principles* for the design of resilient networks and systems, developed as part of the ResiliNets framework [133,127] that provides detailed explanations with their derivation from the strategy and inter-relationships, as well as more extensive background references. This section provides a brief summary of the principles, shown in Figure 2, and describes the way in which they relate to the evaluation of network resilience that is the subject of the rest of this paper.

2.2.1 Prerequisites

The first set of five principles span the domain of prerequisites necessary to build a resilient system. Three of these are essential for the evaluation and analysis of resilience: *service requirements*, *normal behaviour*, and *metrics*.

P1. Service requirements of applications need to be determined to understand the level of resilience the system should provide. Service parameters are the vertical axis \mathbb{P} of the metrics state space described in Section 4. The resilience requirements at a particular service level, consisting of a set of parameters P , define *acceptable*, *impaired*, and *unacceptable service*.

P2. Normal behaviour of the network is a combination of design and engineering specification, along with monitoring while unchallenged to learn the network's normal operational parameters [129]. Operational parameters are the horizontal axis \mathbb{N} of the metrics state space described in Section 4. The resilience of the underlying system when challenged, consisting of a set of parameters N , define *normal*, *partially degraded*, and *severely degraded operation*. Understanding normal behaviour is a fundamental requirement for detecting challenges to normal operation.

P3. Threat and challenge models are essential to understanding and detecting potential adverse events and conditions. It is not possible to understand, define, and implement mechanisms for resilience that defend against, detect, and remediate challenges without such a model.

P4. Metrics quantifying the service requirements and operational state are needed to measure the operational state \mathbb{N} (in the range normal \leftrightarrow partially-degraded \leftrightarrow

severely-degraded) and service state \mathbb{P} (in the range acceptable \leftrightarrow impaired \leftrightarrow unacceptable) to detect, remediate, and quantify resilience, as well as to refine future behaviour. The set of parameters (N, P) and the way in which they are combined as objective functions to determine the scales (\mathbb{N}, \mathbb{P}) of the two dimensional state space are the fundamental basis for the measurement of resilience \mathbb{R} at a particular service level, leading to the multi-level composition into overall network resilience \mathbb{R} described in Section 4.

P5. Heterogeneity in mechanism, trust, and policy are the realities that no single technology is appropriate for all scenarios, and choices change as time progresses. The emerging Future Internet will be a collection of realms [46] of disparate technologies [34], which also define trust and policy boundaries across which there is *tussle* [47]. Resilience mechanisms must not only deal with this heterogeneity, but can also *exploit* it by using diversity in mechanism as a defence, and by providing self-protection mechanisms at realm boundaries.

2.2.2 Design Tradeoffs

The second set of principles describe fundamental trade-offs that must be made while developing and analysing a resilient system.

P6. Resource tradeoffs determine the deployment of resilience mechanisms. The relative composition and placement of these resources must be balanced to optimise resilience and cost. Resources to be traded against one-another include bandwidth, memory [116], processing, latency [136], energy, and monetary cost. These can either be viewed as resources contributing to the operational state \mathbb{N} or as constraints that define the service state \mathbb{P} . Of particular note is that maximum resilience can be obtained with unlimited cost, consisting in part of a full mesh of hardened overprovisioned links, but there are cost constraints that limit the use of enablers such as redundancy and diversity.

P7. Complexity of the network results due to the interaction of systems at multiple levels of hardware and software, and is related to scalability. While many of the resilience principles and mechanisms increase this complexity, complexity itself makes systems difficult to understand and manage, and thereby threatens resilience. The degree of complexity must be carefully balanced in terms of cost vs. benefit, and unnecessary complexity should be eliminated.

P8. State management is an essential part of any large complex system. It is related to resilience in two ways: First, the choice of state management impacts the resilience of the network. Second, resilience mech-




Fig. 2 ResiliNets principles

anisms themselves require state and it is important that they achieve their goal in increasing overall resilience by the way in which they manage state, requiring a tradeoff among the design choices. Resilience tends to favour *soft, distributed, inconsistency-tolerant* state rather than hard, centralised, consistent state, but careful choices must be made in every case, and it is the measurement of resilience \mathcal{R} that helps determine the proper state management decisions.

2.2.3 Enablers

Seven principles are enablers of resilience that guide network design and engineering. These are implemented as resilience mechanisms at each level of the network architecture, and come with the cost of implementation and deployment. The cost–benefit analysis of these mechanisms using the techniques described in the rest of the paper determine the applicability and degree to which each should be used.

P9. Self-protection and security are essential properties of entities to defend against challenges in a resilient network. Self-protection is implemented by a number of mechanisms, including but not limited to mutual suspicion, the AAA mechanisms of authentication, authorisation, and accounting, as well as the additional conventional security mechanisms of confidentiality, integrity, and nonrepudiation.

P10. Connectivity and association among communicating entities should be maintained when possible based on eventual stability, but information flow should still take place even when a stable end-to-end path does not exist based on the eventual connectivity model [134], using DTN (disruption-tolerant networking) techniques such as partial paths, store-and-forward with custody transfer, and store-and-haul (store-carry-forward).

P11. Redundancy in space, time, and information increases resilience against faults and some challenges if defences are penetrated. Redundancy refers to the replication of entities in the network, generally to provide

fault-tolerance. In the case that a fault is activated and results in an error, redundant components are able to operate and prevent a service failure. It is important to note that redundancy does not inherently prevent the redundant components from sharing the same fate, motivating the need for diversity.




Fig. 3 Diversity in path and mechanism

P12. Diversity is closely related to redundancy, but has the key goal to avoid *fate sharing*. Diversity in space, time, medium, and mechanism increases resilience against challenges to particular choices, and consists of providing alternatives so that even when challenges impact particular alternatives, other alternatives prevent degradation from normal operations. Diverse alternatives can either be simultaneously operational, in which case they defend against challenges [124], or they may be available for use as needed to remediate. The provision and analysis of diversity for resilience and the relationship between logical topologies and physical diversity is discussed in considerably more detail in Sections 3–5. Figure 6 shows an example of several kinds of diversity. Communicating subscribers are multihomed to service providers that are diverse in both geography and mechanism. Protection against a fibre cut is provided by the wireless access network; protection against wireless disruptions such as weather or jamming is provided by the fibre connection.

P13. Multilevel resilience [99] is needed in three orthogonal dimensions: *Protocol layers* in which resilience at each layer provides a foundation for the next layer

above; *planes*: data, control, and management; and *network architecture* inside-out from fault-tolerant components, through survivable subnetwork and network topologies, to the Global Internet including attached end systems and applications. The multilevel aspect of resilience analysis is discussed in Section 4.2.6.

P14. Context awareness is needed for resilient nodes to monitor the network environment (channel conditions, link state, operational state of network components, etc.) and detect adverse events or conditions (e.g. [80]). Remediation mechanisms must take the current context of system operation into account.

P15. Translucency is needed to control the degree of abstraction vs. the visibility between levels (layer, plane, and system organisation). An opaque level boundary can hide too much and result in suboptimal and improper behaviour based on incorrect implicit assumptions about the adjacent level [136,122]. Thus it is important that level boundaries be *translucent* in which cross-layer control loops allow selected state to be explicitly visible across levels; *dials* expose state and behaviour from below; *knobs* influence behaviour from above [34].

2.2.4 Behaviour needed for Resilience

The last group of three principles encompass the behaviours and properties a resilient system should possess. These properties are inherently complex, and their cost and benefit to resilience is measured by the analysis techniques described in the rest of this paper.

P16. Self-organising and autonomic behaviour [53, 35] is necessary for network resilience that is highly reactive with minimal human intervention. A resilient network must initialise and operate itself with minimal human configuration and operational management. Ideally human intervention should be limited to that desired based on high-level operational policy.

P17. Adaptability to the network environment is essential for a node in a resilient network to detect, remediate, and recover from challenges. Resilient network components need to adapt their behaviour based on dynamic network conditions, in particular to remediate from adverse events or conditions, as well as to recover to normal operations. At the network level, programmable and active network techniques enable adaptability [38,82].

P18. Evolvability is needed to refine future behaviour to improve the response to challenges, as well as for the network architecture and protocols to respond to emerging threats and application demands. Refinement of future behaviour is based on reflection on the inner strategy loop D²R²: the defence against, detection, and




Fig. 4 Rain intensity in millimeter wireless mesh network

remediation of adverse events or conditions and recovery to normal operation. Furthermore, it is essential that the system can cope with the evolution and extension of the network architecture and protocols over time, in response to long term changes in user and application service requirements, including new and emerging applications and technology trends, as resource trade-offs change, and as attack strategies and threat models evolve.

Weather disruption-tolerant networking [80] provides an example of the application of these principles to increase network resilience. In this domain, precipitation such as thunderstorms challenge areas of a millimeter-wave mesh network. In this case the main *challenge model* is area-based attenuation due to precipitation. *Context-awareness* of the precipitation as measured by radar echo intensity (Figure 4) is used by *translucent* cross-layer controls to allow predictive routing to *adapt* such that flows are not disrupted by the challenge. This is enabled by the *redundancy* and spatial *diversity* of the mesh network.

3 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. These alternative topologies may be based on a particular existing network, for example exploring alternative link interconnections among existing nodes or augmenting with additional components to increase resilience. Alternatively, we may wish to explore the resilience of entirely new network deployments, but

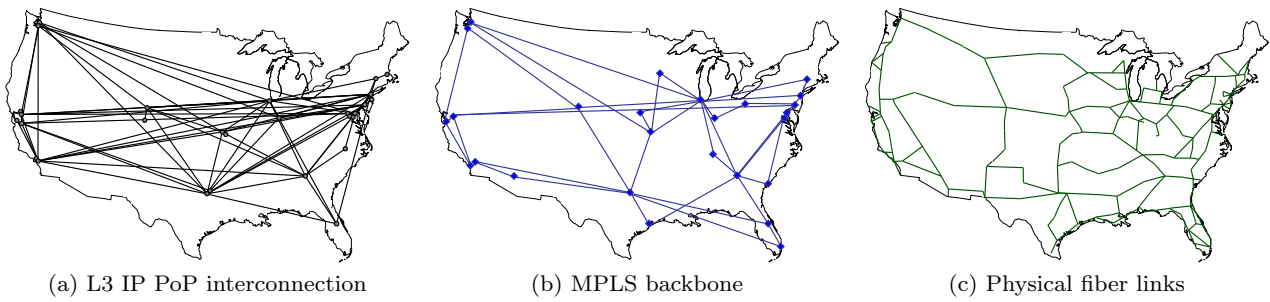


Fig. 5 Comparison of logical overlay and physical topologies for Sprint

grounded in a understanding of the structure of real networks constrained by cost, location, and practicality of infrastructure deployment. The rest of this section will explore these issues and introduce the topology generator KU LoCGen (The University of Kansas Location and Cost-Constrained Topology Generator) [81] and the topology viewer KU-TopView (The University of Kansas Topology Viewer and Combiner) [22].

3.1 Logical vs. Physical Topologies

The majority of the existing body of research is based on logical topology models focusing on the generation of either AS-level [149, 100] or router-level [100] topologies. In an AS-level topology, each Internet autonomous system (AS) is represented as a single node and the BGP (border gateway protocol) connectivity between the ASes represents the graph edge connectivity; this models the highest-level service provider structure of the Internet. While most of the approaches aggregate the intra-AS topology to a single node, some do consider the complexity or structure within a given AS as a mesh.

In the router-level L3 (layer-3) graph, each IP router is represented as a vertex and a logical IP link between a pair of routers forms the edge between the vertices. An example of the Rocketfuel-inferred [131] Sprint L3 topology is shown in Figure 5a.

One of motivating factors for the study of logical topologies is that the L3 protocols such as IP, IGPs (interior gateway protocols), and BGP only see L3 connectivity of the Internet. Furthermore, the majority of inference mechanisms [73] are only able to collect data on the the router-level connectivity of commercial networks. To date, results from topology modelling have been used for evaluating various aspects of networks [36] including security, performance, traffic modeling and engineering, protocol development and analysis, as well as evaluation of numerous other algorithms.

However, an edge between a pair of vertices almost never corresponds to a direct physical link without any intermediary lower layer nodes due to the underlying structures that provide IP connectivity, including L2.5-traffic-engineering underlays such as MPLS (multiprotocol label switching – Figure 5b [19]), L2 structure such as SONET/SDH (synchronous optical network / synchronous digital hierarchy) rings including cross connects and ADMs (add-drop multiplexors), and fibre links interconnected by regenerators and amplifiers (Figure 5c [85]). Hence, neither the AS-level nor router-level graphs represents the actual physical connection between nodes, as can be seen in Figure 5c for the Sprint network. In this example, the San Jose – Kansas City IP interconnection might go through the Stockton or Anaheim – Ft. Worth MPLS nodes, which follows a geographic fiber path significantly different from that implied by either the L3 or L2.5 graphs.

While L3 topologies are useful for modelling the the resilience of L3 services such as BGP and IGP routing, they are not sufficient to understand the resilience of the physical infrastructure to a number of challenges, including large-scale disasters and attacks against the infrastructure, explored in Section 5. Furthermore, since it is possible for two distinct IP paths of different service providers to share the same physical conduit, it is difficult to understand and engineer the resilience of the network by assuming that IP links correspond to physical links. Without an understanding of the *geographic location* of physical network nodes and links and their correspondence to logical links, it is not possible to know if the logical components share fate, as was the case in the Baltimore tunnel fire [138] in which many logically distinct links failed at the same time when all the fibre running through the tunnel melted.

Therefore, we argue that resilience evaluation of a network must begin with the *physical* topology and geography because it is the physical topology that ultimately determines the ability to survive infrastructure failures. Service and network dependability and performability in the face of failures is highly dependent

on the physical topology. For example, in the case of recovery after a large-scale disaster, it becomes the surviving physical infrastructure that drives traffic management decisions. This leads to the need for realistic topology models and generators that reflect the physical and geographic structure of the network, as well as the logical topology overlays.

We note that one of the reasons for the previous lack of interest in physical layer topologies may in part be the abstraction of protection mechanisms. For example, link-level protection such as SONET/SDH automatic protection switching (APS) [57], p -cycles [72], and f -cycles [106] provide fault-tolerant masking of uncorrelated failures, and shared-link risk groups (SLRGs) [137] provide topological diversity, but not necessarily *geographic* diversity. These mechanisms do not solve the fate-sharing problem nor provide resilience against correlated failures and attacks.

One of the fundamental challenges in developing a physical topology model is the lack of real data for validation of the models. The physical topology of commercial networks including the Internet are not readily available. Previous research has considered several inference mechanisms to determine geographic node locations and physical link distances [64, 113, 88], but despite these efforts, the inference of physical topologies remains an open problem. There are, however, a few educational and research networks such as GÉANT2, NLR (National LambdaRail), and Internet2 for which the physical topology is available for validation, but unfortunately research networks are generally significantly smaller than large commercial ISPs. It should be noted that physical topology generation and analysis is fundamentally an intra-domain issue. Hence, we can independently validate a model against an individual ISP physical topology.

3.2 Path Diversity

As described in Section 2.2.3, a key enabler to resilience is *diversity* [124, 123, 120] such that when challenges impact part of the network, other parts do not share fate and are able to continue communicating. In the case of topological resilience, it is important that diverse physical paths exist, and that end-to-end communication is able to exploit this capability and choose paths that are unlikely to experience correlated failures. To this end, we define a measure of diversity (introduced in [124]; further developed in [123]) that quantifies the degree to which alternate paths share the same nodes and links. Note that in the WAN (wide-area network) context in which we are concerned with events and connections on a large geographic scale, a node may be thought of

as representing an entire PoP (point-of-presence) area, and a link as the physical bundle of fibers buried in a given right-of-way. This distinction between WAN and LAN (local-area network) component identifiers affects only the population of the path database, not the usage of the diversity metric.

3.2.1 Diversity Metric

Given a (source s , destination d) node pair, a path P between them is a vector containing all links L and all intermediate nodes N traversed by that path $P = L \cup N$ and the length of this path $|P| = |L| + |N|$ is the combined total number of elements in L and N .

Let the shortest path between a given (s, d) pair be P_0 , L_0 be a vector containing the set of links traversed by P_0 , and N_0 be a vector containing the nodes which lie on P_0 . Then, for any other path P_k between the same source and destination, we define the diversity function $D(x)$ with respect to P_0 as:

$$D(P_k) = 1 - \frac{|P_k \cap P_0|}{|P_0|}$$

The path diversity has a value of 1 if P_k and P_0 are completely disjoint and a value of 0 if P_k and P_0 are identical. For two arbitrary paths P_a and P_b the path diversity is given as:

$$D(P_b, P_a) = 1 - \frac{|P_b \cap P_a|}{|P_a|}$$

where $|P_a| \leq |P_b|$.




Fig. 6 Shortest path P_0 and alternatives P_1 and P_2

It has been claimed [109] that measuring diversity (referred to as novelty) with respect to *either* nodes *or* links is sufficient, however we assert that this is not the case. Figure 6 shows the shortest path, P_0 , along with the alternate paths P_1 and P_2 both of which have a (link) novelty of 1. However, given a failure on node 1, both P_0 and P_2 will fail. In our approach, $D(P_2) = \frac{2}{3}$, which reflects this vulnerability. P_1 on the other hand has both a novelty of 1 and a diversity of 1, and does not share any common point of failure with P_0 . Similarly, the wavelengths or fibres from multiple nodes may in fact be shared by a single physical conduit such as was the case in the Baltimore tunnel fire [138], resulting in a single point of failure, thus illustrating the need

for including both nodes and links into the diversity measure.

3.2.2 Effective Path Diversity

Effective path diversity (EPD) is an aggregation of path diversities for a selected set of paths between a given node-pair (s, d) . To calculate EPD we use the exponential function

$$\text{EPD} = 1 - e^{-\lambda k_{sd}}$$

where k_{sd} is a measure of the added diversity defined as

$$k_{sd} = \sum_{i=1}^k D_{\min}(P_i)$$

where $D_{\min}(P_i)$ is the minimum diversity of path i when evaluated against all previously selected paths for that pair of nodes. λ is an experimentally determined constant that scales the impact of k_{sd} based on the utility of this added diversity. A high value of λ (> 1) indicates lower marginal utility for additional paths, while a low value of λ indicates a higher marginal utility for additional paths. Using EPD allows us both to bound the diversity measurement on the range $[0,1]$ (an EPD of 1 would indicate an infinite diversity) and also reflect the decreasing marginal utility provided by additional paths in real networks. This property is based on the aggregate diversity of the paths connecting the two nodes.




Fig. 7 Total graph diversity vs. number of paths selected

3.2.3 Path Selection

We identify three different criteria for choosing a set of diverse paths for a given node pair: number of paths, diversity threshold, and stretch limit. The objective in the

first mode is to find k maximally diverse paths. We first find the shortest fully disjoint paths, and if additional paths are required we continue finding paths that add maximum diversity as calculated using the equation for k_{sd} . The second mode selects as many maximally diverse paths as are required to achieve the requested EPD. Finally, the third mode selects all maximally diverse paths with stretch less than the stretch limit. In all modes, the set of maximally diverse paths are found using the Floyd-Warshall algorithm with modified edge weights [33]. In this algorithm, only those paths are used that increase the EPD for the node pair in question. Recall that only paths with one or more disjoint elements (links+nodes) will result in non-zero D_{\min} and consequently increase EPD.

3.2.4 Measuring Graph Diversity

The total graph diversity (TGD) is simply the average of the EPD values of all node pairs within that graph. This allows us to quantify the diversity that can be achieved for a particular topology, not just for a particular flow. For example a star or tree topology will always have a TGD of 0, while a ring topology will have a TGD of 0.6 given a λ of 1. In Figure 7 we compare three different real network topology TGD plots with those of four regular topologies (full-mesh, Manhattan grid, ring, and star). The Sprint and AT&T topologies are inferred from Rocketfuel [4]; GÉANT2 [9] nodes are the actual location.

Table 1 Network statistics

	Nodes	Links	Avg. deg	TGD
full-mesh	20	190	19.00	0.99
grid	25	40	3.20	0.75
AT&T	25	92	7.36	0.71
Sprint	27	136	10.00	0.70
GÉANT2	34	102	6.00	0.58
ring	25	25	2.00	0.39
star	25	24	1.92	0.00

Table 1 shows the number of nodes and links, average node degree, and TGD with $k = 10$ and $\lambda = 0.5$ for each network [123]. Of importance here is that the average node degree alone is not sufficient to indicate the diversity of a network in real-world cases, although it may be used to rank regular synthetic topologies. We see that while Sprint has a higher average node degree than AT&T, AT&T is slightly more diverse.

We have produced the Web-based KU-TopView network mapping tool [22] to visualize these topologies with screen-shots of the Sprint physical and logical topolo-

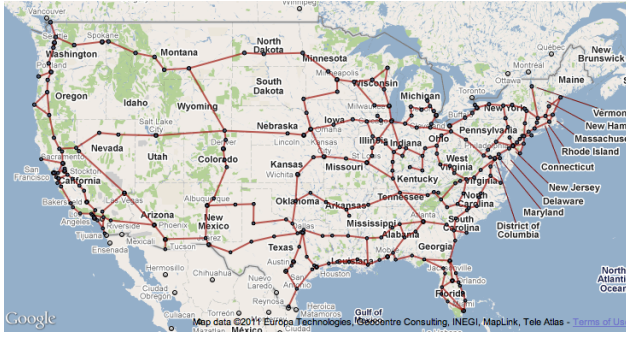


Fig. 8 Sprint fiber map in KU-TopView

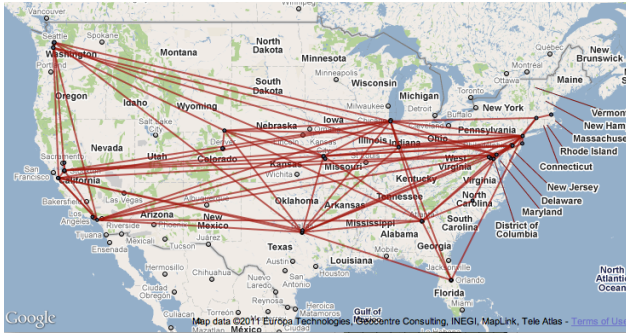


Fig. 9 Sprint layer-3 map in KU-TopView

gies shown in Figures 8 and 9. KU-TopView also provides adjacency matrices to be used by analysis tools such as KU-CSM described in Section 5.

Here we note that for diversity to make sense in the graph context it should be computed considering only path components (nodes and links) at the level of network hierarchy for which the diversity value is desired. For example, in computing the diversity of a service provider's backbone, only core nodes should be considered, otherwise the comparatively large number of subscriber nodes (typically stubs) will artificially reduce the calculated diversity. We also note here that the diversity measure is designed such that it does not penalize longer paths in favor of shorter paths, meaning that graph diameter and average path lengths are independent metrics that should be considered in addition to the diversity metric.




Fig. 10 Geographic diversity: distance d and area A

3.2.5 Geographic Diversity

The previous EPD and TGD measures consider the sharing of components, but do not capture the geographic characteristics necessary for area-based challenges such as large-scale disasters (simulated in Section 5) or to prevent the geographic fate sharing of distinct links in the same conduit as in the Baltimore tunnel fire. Therefore we are augmenting the diversity measures with a minimum distance between any pair of nodes along alternate paths, and as the area inside a polygon or set of polygons, the borders of which are defined by a pair of alternate paths, as shown in Figure 10. Thus, it should be possible to specify diverse paths among a set of candidates with a given degree of sharing and distance metric $EPD(d)$ constrained by stretch, and measure the geographic area between the paths $EPD(A)$ as well as to measure the diversity inherent in a graph across all paths $TGD(d, A)$.

3.3 Hierarchical Topology Model

Even though the majority of existing research deals with logical topologies, there is a significant overlap in generation models with the physical topology models needed for the analysis of resilience. Furthermore, we can draw upon the lessons learned through the evolution of topology research [73].

3.3.1 History of Models

The field of topology analysis and generation goes as far back as 1950s [60], and has been studied in various fields including computer science, mathematics, and physics [43]. Pre-power-law studies include random models such as Waxman [147] and hierarchical models such as Tiers [37] and transit-stub [151], in which the focus was on recreating the structure of networks.

Later, it was observed that the L3 degree-distribution in the Internet follows three power laws [63], followed by work that enhanced these power laws, theorised the underlying causes [30, 128], and developed models to generate graphs that faithfully reproduced these degree-based properties. Since structure-based models such as hierarchical did not strictly produce these properties in the graphs, they were discarded as not being representative of Internet topology. Hence power law in the degree distribution was considered a necessary, and in some studies sufficient, condition for the representativeness of the graphs.

In the post power-law era, further research and analysis was conducted to better understand observed properties in inferred topologies as well as the limitations

of inference mechanisms and the factors that lead to such properties [39, 44, 61]. The latest studies have emphasised the need to model the process of network development instead of replicating properties in random graphs [90, 25].

Emphasis is increasingly placed on reconciling the differences between purely analytical models and practical design principles. It was observed that Internet does not have an “achilles heel” [55, 25] of a purely power-law graph and that the processes that led to Internet properties are quite complex involving various optimisations that are characteristics of real networks [24, 39]. The latest understanding of the research community is that the purely analytical models are not very representative of actual networks and are hence not capable of producing realistic topology models. Several heuristic methods that incorporate real world optimisations and tradeoffs have been proposed [55, 146]. Specifically, the method of highly-optimised tolerance [39, 24, 25] proposes network topology as a result of resilience optimisation with non-generic, highly engineered configurations.




Fig. 11 KU-LoCGen hierarchical topology model

3.3.2 Hierarchy in KU-LoCGen

The goal of KU-LoCGen (The University of Kansas Location and Cost-Constrained Topology Generator) is to provide a flexible framework that allows for an n -level hierarchical, modular structure with level-specific graphs, constrained by cost, population, and infrastructure location. Furthermore, the graph models used at each level differ significantly. These vary from closed-form general-purpose mesh models (e.g. Waxman [147]) typical in backbones, to pre-structured models such as rings and trees typical in access networks based on particular technologies such as SONET/SDH rings and

HFC (hybrid-fibre coax) trees, to a modified power-law preferential attachment of subscribers to access networks. Figure 11 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.

The current KU-LoCGen implementation generates three levels representing the backbone, access, and subscriber networks whose geographic distribution is represented by Ψ_p , Ψ_n , Ψ_s respectively. Furthermore, the number and geographic spread of nodes at any given level is strongly correlated to the higher level nodes in the hierarchy as discussed below. The backbone node (level-1) distribution model Ψ_p supports three different location constraints including fixed geographic positions based on known point-of-presence geolocations of existing networks, user defined location, and a random distribution as discussed below.

The number of access networks (level-2) $N(i)$ are chosen based on a uniformly distributed random variable. $N = U(n_{\min}, n_{\max})$, where n_{\min} and n_{\max} are the lower and upper limits on the number of access networks per backbone node. The $N(i)$ access networks are distributed around a given PoP using a Gaussian distribution: $\Psi_n = N[\mu_n, \sigma_n^2]$, where μ_n represents the PoP location and σ_n^2 is the variance. The subscriber networks are distributed normally: $\Psi_s = N[\mu_s, \sigma_s^2]$, where μ_s represents the access network location and σ_s^2 is the variance. Obviously, the variance determines the geographical extent and the spread of the subscribers. Additionally, the variance of each access network may vary according to the size and location of the access network as well as the PoP to which the access network is connected: $\sigma_s^2(i) \propto \frac{1}{N(i)}$

The number of access nodes $M(i)$ in the i^{th} access network of the j^{th} backbone node is based on the distance of the access network from the backbone node relative to the other access networks connected to the same backbone node. The number of nodes in the access network is given as

$$M(i) = \frac{\max(d_t); t = 1, 2, \dots, N(j)}{d_i} \times M_{\min}$$

where d_i is the distance of the i^{th} access network and M_{\min} is the minimum number of access networks defined per PoP. Furthermore $M(i)$ is also the upper bound to a predefined maximum value of M_{\max} . The access network nodes are then uniformly distributed in a circular region of radius r around the first access node. Therefore $\Psi_m = U(0, r)$. The number of subscribers in an access network is directly proportional to the size of

the access network;

$$S(i) = \frac{M(i)}{\max(M(j)); j = 1, 2, \dots, N} \times S_{\max}$$

where S_{\max} is the predefined limit on the maximum number of subscribers per access network.

3.4 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 [88]. Moreover, the probability of link deployment is strongly related to the distance between the nodes. Geographic distance-based models such as Waxman accurately model the link distribution when considering location constraints [88].

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 5. Our ultimate goal is to understand the graph-theoretic properties that relate to network resilience [139], including spatial diversity that requires node geolocation information. As described in Section 3.2.5, nodes may be located such that total graph diversity meets a constraint $TGD(d)$.

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 [131] for backbone node placement that generally corresponds to PoP locations). In this case we synthetically generate links under cost constraints, as described later. Alternatively, we can use population density to drive node placement, as described next, additionally constrained to meet graph theoretic properties such as clustering coefficient and $TGD(d)$.

We consider both the absolute distribution of the nodes as relating to population density and the distribution of nodes with respect to each other. The use of a hierarchical model enables us to achieve this by defining a separate structure or growth model for each level. While the position and distribution of the level-1 (backbone PoP) nodes is based on the population distribution (or other location constraints such as existing PoPs or fibre infrastructure), the distribution of the access networks and access network nodes requires further research to determine the distributions that model it accurately.

3.4.1 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 that minimise the deployment cost of fibre-optic cables; long-distance runs are typically laid along railways, motorways, pipelines, and transmission lines.

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 km² gridded population density data sets from CIESIN [42] to determine optimal locations for backbone PoP placement [75]. *K-means* is an iterative clustering method that works in two phases. The goal is to minimise the sum of the distances between all data points to cluster centres for all clusters. The initial selection of the cluster centres is random. The first *batch* phase recomputes the cluster centres by re-associating each data point to its nearest cluster centre. This phase provides an approximate but fast computation of cluster centres. The second *on-line* phase uses the output of the batch phase as the initial cluster centres and re-associates points to a different cluster only if doing so reduces the sum of distances. Cluster centres are recomputed after each re-association. Each iteration requires one pass through all data points. This is computationally complex and time consuming phase, especially for such large data sets.




Fig. 12 Relative node locations for combined ISPs in USA

The two inputs to the algorithm are population data and the number of cluster centres. In this example, we consider multiple ISPs to aggregate across tier-1

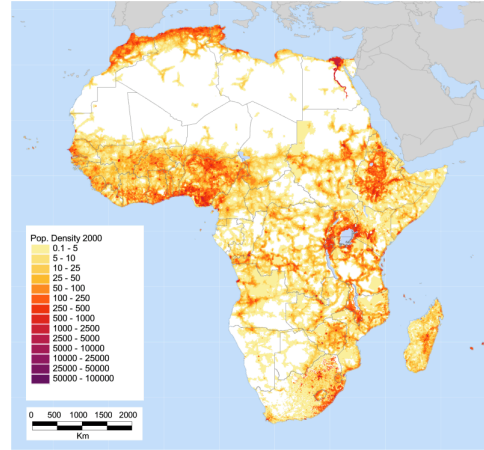
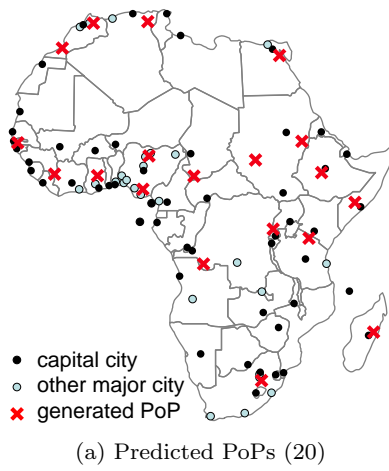


Fig. 13 Cluster centres for Africa

providers, so as to not neglect certain parts of the US that may be under-served by a particular ISP. Figure 12 shows a comparison of 112 PoPs generated using our population based model with the existing 112 combined Rocketfuel-inferred [131] L3 PoP cities of Sprint, AT&T, and Level3.




Fig. 14 CCDF of offset distance

We quantify the distance between inferred PoP locations and population based cluster centres 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 14. Note that with combined ISPs, almost 90% of the nodes generated are within 50 km offset distance; only a very small percentage of nodes are outliers.

Next we consider an under-developed area that does not yet have significant network infrastructure. We gen-

erate the optimal location of backbone PoPs that could be used by an ISP desiring to have a continent-wide topology. Figure 13a shows the predicted location of 20 PoPs for Africa, next to a population-density map (Figure 13b) [23] for visual comparison. Since there is no continent-wide ISP in Africa, we cannot compare predicted node locations with existing infrastructure.

3.4.2 Technology Penetration

The other fundamental aspect governing the location of the PoPs is technology penetration. The location of backbone PoPs is highly dependent on the number of Internet users in a given area. We denote the technology penetration factor as γ , defined as the fraction of *Internet* users to the total population in a particular area, and assume this factor is uniform for a developed regions such as the US and Europe: $\gamma=1$. This factor has particularly significant influence on a developing country such as India, where technology penetration is not homogeneous across all areas. Hence, placing network resources solely based on the population density data set would not lead to a realistic network deployment.

India is highly populated in the northern belt of the river Ganges. However, the number of Internet users for this region is small compared to the absolute population. We consider the inferred topology of the VSNL network in India [131], which has only five PoPs located in Delhi, Mumbai, Hyderabad, Bangalore, and Chennai. The clustering algorithm is run on both on the absolute population data set as well as the effective γ weighted data. Figure 16 shows that four of the predicted PoPs match VSNL closely. However, instead of a PoP near Chennai, it is placed near Patna for two reasons: Patna

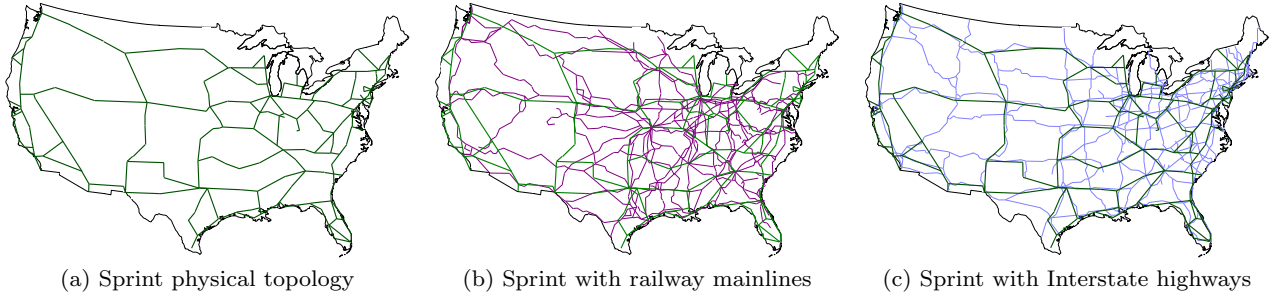


Fig. 15 Comparison of Sprint fiber topology and main transportation routes

is much denser in population than Chennai and the PoP placed near Bangalore is close enough to Chennai for our algorithm to place another PoP.

The quarterly report released by *Telecom Regulatory Authority of India* [143] is used to get the state-wide list of broadband subscribers in India. Technology penetration is incorporated into our model by weighting the population of each grid in an area by the corresponding γ and then clustering the resulting data set. After incorporating γ , the 4 PoPs which matched earlier get closer to the real locations, while the one in Patna moves to Kolkata as it is one of the metropolitan areas with a high number of Internet users.

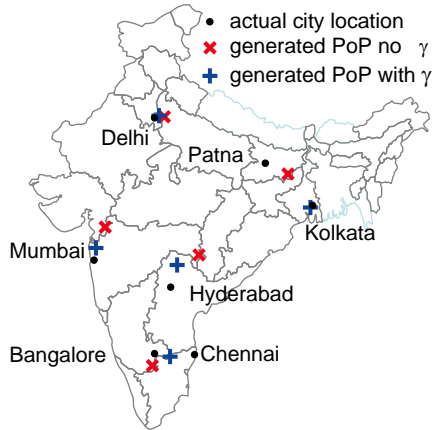


Fig. 16 Illustration of γ factor for India

3.4.3 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 ad-

ditional step of “snapping-to-grid” nodes to infrastructure, and should improve the accuracy of node placement over purely population based. For example, in Figure 12 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. Figure 15 shows the relationship of the Sprint physical fibre topology to railway mainlines (based on [121]) and Interstate freeways (based on [21]) in the US.

3.5 Cost-Constrained Connectivity

Given a set of node locations, either based on existing networks or predicted as discussed previously, 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 cost constrained connectivity models to generate feasible topologies.

Economic factors shape physical level infrastructure [25]. The resilience and survivability of networks is almost always limited by the cost, therefore, realistic models must incorporate cost constraints. However, this poses a significant challenge due to the lack of standard cost functions for network infrastructure. Furthermore, the cost function is not only location and time dependent, but also depends on the level within the network hierarchy.

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_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 generally 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 [88]. While it is generally agreed that backbone networks are mesh-like [55], there is some contention as to exact relationship between link probability and its distance; some works claim that this is exponential [88], but others claim that it is linear [150].

According to the Waxman model [147] the probability that two nodes u and v have a link between them is given by

$$P(u, v) = \beta e^{-\frac{d(u,v)}{L\alpha}}$$

where $0 < \alpha, \beta \leq 1$ and L is the maximum distance between any two nodes. The Waxman parameters α and β are controlled by the cost. A high value of α corresponds to a high fraction of short to long links and β is directly proportional to the link density; d is the Euclidian distance between the two nodes. The Waxman model as traditionally applied begins with uniform node distribution, but we use the location constrained node locations for a realistic backbone topology model.

Figure 17 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 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 Section 4.




Fig. 17 Sample 2-level topology using 27 nodes

Table 2 Offset distance with existing PoPs in km

Network (PoPs)	Mean	σ	Min.	Max.
Sprint (27)	54.2	45.3	2.6	163.6
AT&T (106)	26.5	37.3	1.1	265.2
Level 3 (38)	43.4	31.7	9.6	118.6
GÉANT2 (34)	101.5	54.1	20.2	252.3
Ebone (27)	56.3	27.9	17.7	131.1
Tiscali (47)	34.8	22.3	2.47	80.6
VSNL (5)	26.7	34.9	2.6	265.1

Table 2 is a summary of our results pertaining to the locations of PoPs for various ISPs in the US, Europe, and India [75]. Noted that all of them are Rocketfuel-inferred topologies except for the European GÉANT2 [9] research network. Our predictions match very well with ISPs with large infrastructure. For example, in the case of AT&T, the mean separation between real and clustered nodes is 26.5 km and the closest match is with an offset of 1.1 km.

3.6 Example Synthetic Network Graph

In this section, we demonstrate the ability to generate a realistic 27 node topology based on US population density data set. We use a cost-constrained Waxman model to connect the backbone nodes. The objective is to go from realistic node locations to understanding realistic topologies and evaluate resilience of synthetic graphs. Figure 18 shows the backbone topology generated by our model.



Fig. 18 Synthetic topology for 27 nodes

Note that we do not claim that this topology is resilient, but we evaluate the graphs by comparing various metrics in Table 3 and show that the graph metrics of synthetically generated topologies are comparable to realistic topologies, if not better.

We compute betweenness, average node degree and clustering coefficient metrics of the graph. The betweenness metric is the number of shortest paths through a

Table 3 Topological characteristics of sample networks

Network Topology	Sprint PoPs	Synthetic
Number of nodes	27	27
Number of edges	68	71
Maximum degree	12	14
Average degree	5.04	5.23
Clustering coeff.	0.43	0.28
Network diameter	6	6
Average hop count	2.44	2.22
Node betweenness (max/min/avg)	144/28/72	124/1/32
Link betweenness (max/min/avg)	72/2/12.6	35.1/2.9/11.3

particular node or link [97]. A high value of betweenness (for example, average node betweenness of 72 for Sprint topology) indicates a high stress level. This indicates presence of more critical nodes in a graph. An average node betweenness value of 32 for the generated synthetic topology implies uniform stress levels for most nodes. Similar reasoning applies for link or edge betweenness metric. A higher average node degree value generally indicates that a graph is better connected and is more robust [97]. We observe that average node degrees for both graphs are almost equal. Clustering coefficient, almost same for both topologies, is the measure of how well neighbors of a node are connected. The other metrics for both topologies compare well. The ability to *generate graphs with specified resilience properties* such as TGD(d) is a key part of our future work.

4 Analytical Resilience Framework

This section describes a new analytical framework to evaluate network resilience based on a two-dimensional state space: the horizontal axis representing the operational state of the network and the vertical axis the service delivered. The resilience of a network is quantified as the trajectory through the state space as the network is challenged by failures, attacks, or large-scale disasters. We show that for particular scenarios, and at particular service levels, we can indeed characterise the resilience with a single number given by the area under the resilience trajectory.

4.1 Background

The earliest works in fault tolerance include the 1956 Moore and Shannon paper [108] on reliable circuits and the Peirce [118] and Avizienis [27] publications on fault-tolerant computing. The initial work on reliability and fault tolerance was focused on the design of computing

systems [96]. In 1974, one of first resilience works in communication networks was presented [65] as the survivability analysis of command and control networks in the context of military systems. The inability to design systems with sufficient redundancy to overcome all failures was realised in late 1970s in the context of fault-tolerant computing systems [32, 114, 95]. Hence the concept of *degradable systems* was introduced in which the system has at least some degraded performance under the presence of challenges without failing completely. Markov models are used to evaluate the performance and reliability of the degradable system [78, 101, 69]. Meyer [101] first coined the term *performability* as the probability that the system will stay above a certain accomplishment level over a fixed period of time [74]. Until then, reliability and performance of communication networks were treated separately. Huslende [78] defined performance as the second dimension of the classical reliability, thereby defining the reliability of a degradable system as the probability that the system will operate with a performance measure above certain threshold. Since then there have been a number of rigorous definitions of survivability, reliability and availability [59, 86, 71, 105]. Existing research on fault-tolerance measures such as reliability and availability targets single (or at most several) instances of random faults, such as topology based analysis considering node and link failures [91, 92, 26, 83].

Frameworks to characterise survivability were also developed in specific contexts, such as for large-scale disasters [91, 92], vulnerabilities under the presence of DDoS attacks [119, 76], and in conjunction with network dynamics [86, 68, 104]. The T1A1.2 working group defines survivability [141, 142] based on availability and network performance models; later approaches have used this approach to quantify survivability [94, 77, 144]. In this paper, we quantify network resilience as a measure of service degradation in the presence of challenges that are perturbations to the operational state of the network.

4.2 Metrics Framework

In this section, we present a framework to quantify network resilience in the presence of challenges using functional metrics [79, 103], beginning with a brief overview of our approach. Recall that we define resilience 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. 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 de-

pend, and the plethora of adverse events and conditions that present as challenges to the network as a whole. This complexity renders an exhaustive resilience analysis intractable. Our approach simplifies the resilience evaluation process by using two novel methods. First, we isolate the impact of challenges at each layer in the network by evaluating resilience at each service-layer boundary, thereby avoiding a continually increasing parameter set as we move up the network layers. Secondly, we quantify resilience as a (negative) change in service corresponding to a (negative) change in the operating conditions at any given layer [103]. Therefore resilience is characterised as a mapping between the network operation and service, wherein the operation is affected by challenges, which in turn may result in degradation of the service at the service-layer boundary. In other words, instead of evaluating the impact of each challenge or attack separately, which leads to an intractable number of cases, we focus on quantifying the service to varying operational conditions. Given the right set of metrics, the operations can always be defined such that most challenges manifest as perturbations in these operational metrics. We now present the formal framework.




Fig. 19 Resilience state space $\mathbb{N} \times \mathbb{P}$

4.2.1 Operational State \mathbb{N}

The first step in our framework is to quantify the operational state (the givens) at any layer using a set of metrics termed intuitively as *operational metrics*. For a given system \mathcal{S} , in which the system refers to the network at an arbitrary level, let the ℓ operational metrics be represented as $\mathcal{N}_{\mathcal{S}} = \{N_1, \dots, N_{\ell}\}$. 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, $N_i = \{n_{i,1}, \dots, n_{i,m}\}$. For example, at the physical layer of an ISP network, the number of link failures and link capacities could be two operational metrics. The *operational state space* of \mathcal{S} is $\mathcal{N}_{\mathcal{S}} = \mathbf{X}_i N_i$ where \mathbf{X} represents the cross product operator. Therefore, the operational state space consists of all possible combinations of the operational metrics.

We now define an *operational state*, \mathbb{N} as a subset of the complete state space $\mathcal{N}_{\mathcal{S}}$, represented as the horizontal axis in Figure 19. Therefore, \mathbb{N} is an operational state if $\mathbb{N} \subseteq \mathcal{N}_{\mathcal{S}}$. Let $\mathbb{N}_{\mathcal{S}}$ be a set of operational states, $\mathbb{N}_{\mathcal{S}} = \{\mathbb{N}_1, \dots, \mathbb{N}_k\}$. $\mathbb{N}_{\mathcal{S}}$ is valid if $\mathbb{N}_{\mathcal{S}}$ is a partition of $\mathcal{N}_{\mathcal{S}}$. That is $\mathbb{N}_i \cap \mathbb{N}_j = \emptyset, \mathbb{N}_i, \mathbb{N}_j \in \mathbb{N}_{\mathcal{S}}$ and $i \neq j$ and $\cup_i \mathbb{N}_i = \mathcal{N}_{\mathcal{S}}$ where \cup represents the union operator. Hence, in the generic case, an operational state is defined as a subset of $\mathcal{N}_{\mathcal{S}}$.

If N_i is numeric and ordered $\forall i$ such that $N_i \in \mathcal{N}_{\mathcal{S}}$, then the k^{th} operational state \mathbb{N}_k can be defined using the same notation used to define the complete state space instead of specifying it as a subset of $\mathcal{N}_{\mathcal{S}}$. Therefore, $\mathbb{N}_k = \{N_{1k}, \dots, N_{ik}, \dots, N_{\ell k}\}$. A member N_{ik} in the set \mathbb{N}_k is in itself a set of valid values bounded by $[n_{ik}, \bar{n}_{ik}]$, representing the lower and upper limit of the i^{th} operational metric. We can now define $N_{ik} \equiv \{n_{ik}, \dots, \bar{n}_{ik}\}$. Thus N_{ik} represents the set of i^{th} operational metric values that correspond to the operational state \mathbb{N}_k . We divide the operational state into three regions: *normal operation*, *partially degraded*, and *severely degraded*.

4.2.2 Service State \mathbb{P}

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. For example, the service provided by the routing layer (to the transport layer is) discovery of end-to-end paths. Let the ℓ service parameters of system \mathcal{S} be represented by $\mathcal{P}_{\mathcal{S}} = \{P_1, \dots, P_{\ell}\}$. 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), $P_i = \{p_{i,1}, \dots, p_{i,m}\}$. The *service state space* of \mathcal{S} is $\mathcal{P}_{\mathcal{S}} = \mathbf{X}_i P_i$. Therefore, the service state space consists of all possible combinations of the service parameters.

Similar to an operational state, we define *service state*, \mathbb{P} , as a subset of the complete state space $\mathcal{P}_{\mathcal{S}}$, represented as the vertical axis in Figure 19. Therefore, \mathbb{P} is a service state if $\mathbb{P} \subseteq \mathcal{P}_{\mathcal{S}}$. Let $\mathbb{P}_{\mathcal{S}}$ be a set of service states, $\mathbb{P}_{\mathcal{S}} = \{\mathbb{P}_1, \dots, \mathbb{P}_k\}$. $\mathbb{P}_{\mathcal{S}}$ is valid if $\mathbb{P}_{\mathcal{S}}$ is a partition of $\mathcal{P}_{\mathcal{S}}$. That is, $\mathbb{P}_i \cap \mathbb{P}_j = \emptyset, \mathbb{P}_i, \mathbb{P}_j \in \mathbb{P}_{\mathcal{S}}$

and $i \neq j$ and $\cup_i \mathbb{P}_i = \mathbb{P}_S$. Note that a union of all service states forms the complete service state space. In other words, service states are simply partitions of the complete service space.

If P_i is numeric and ordered, then the k^{th} service state can be represented as $\mathbb{P}_k = \{P_{1k}, \dots, P_{ik}, \dots, P_{\ell k}\}$. A member P_{ik} in the set \mathbb{P}_k is in itself a set of values bounded by $[p_{ik}, \bar{p}_{ik}]$, representing the lower and upper limit of the i^{th} service metric. We can define $P_{ik} \equiv \{p_{ik}, \dots, \bar{p}_{ik}\}$. Thus, P_{ik} represents the set i^{th} service parameter values that correspond to the service state \mathbb{P}_k .

4.2.3 Network State S

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

In a deterministic system, the mapping of \mathcal{N}_S to \mathcal{P}_S is functional, meaning that for each operational state there is one and only one service state. However, if the system is stochastic then this mapping is also stochastic in which one operational state maps to multiple service states based on the randomness in the execution of the system. In order to eliminate the stochastic nature of the \mathcal{N}_S to \mathcal{P}_S mapping in our analysis, we present the \mathbb{N}_S to \mathbb{P}_S mapping, thereby focussing on the mapping of *aggregates* rather than individual operational or service states. In other words, instead of looking at the mapping of a instantaneous value of an operational metric to a service parameter, we focus on the mapping of the operational state to the service state.

Note that both the operational and the service state spaces are multi-variate. In order to visualise this state space on a two dimensional state space as in Figure 19, 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 service and operational parameters. 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)

Figure 19 shows the system in an initial state S_0 for which acceptable service is delivered during normal operations. As the network is challenged, its operational state may be degraded, represented by the

states labelled S_c, S_a . Depending on the service specification and resilience, various trajectories are possible. The lower S_a is preferable since the service remains acceptable even when the network degrades. Next, we describe how this can be quantified.




Fig. 20 Resilience \mathbb{R} measured in state space

4.2.4 Resilience Evaluation \mathbb{R}

Under normal conditions, the network continues to operate in a given state corresponding to normal operational and service states. When a challenge causes a large perturbation in the operational state, the service may also be impaired below the acceptable service specification. A significant change in either dimension leads to a *network state transition*. 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 state 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 N and \mathbb{P} . The operational and service space is covered fully by its states and can be decomposed in a fixed set of large states which we term as regions. For simplicity, the network operational space N is divided into *normal operation*, *partially degraded*, and *severely degraded* regions as shown in Figure 20. Similarly, the service space \mathbb{P} is divided into *acceptable*, *impaired*, and *unacceptable* regions.

We then quantify the resilience \mathbb{R}_{ij} for a particular scenario at a particular layer boundary B_{ij} as the area under the resilience trajectory, shown by the shaded triangular area under the $S_0 \rightarrow S_c$ trajectory in Figure 20. This results in a *static resilience analysis* [79] that does not consider the temporal aspects of the state space trajectory.

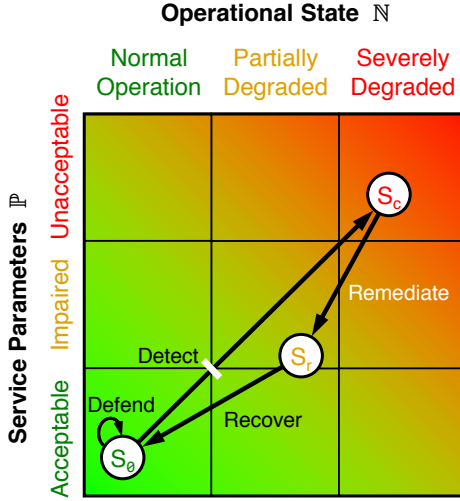


Fig. 21 Resilience state space

4.2.5 Relationship to the Strategy

The relationship of the the state-space formulation to the ResiliNets strategy described in Section 2 is shown in Figure 21. The inner D^2R^2 loop trajectory is 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 situation to S_r , and recovery finally returns the system to its original state S_0 .




Fig. 22 Temporal aspects of resilience

To measure the benefits of remediation mechanisms, a *temporal resilience analysis* should not only consider the area under the trajectory, but factor in the time that the system is in challenged and remediated states, as shown in Figure 22 (based in part on on [140]). At time t_0 a challenge lowers performability by P_u (fraction unserved) corresponding to the $S_0 \rightarrow S_c$ state transition, but then remediation mechanisms at t_1 increase performability by P_r (fraction remediated) corresponding to the $S_c \rightarrow S_r$ state transition. Eventually recovery at t_2 returns the network to its original normalised performability of $P = 1$. Clearly the shorter the time to remediate t_r and time to recover t_R , the more resilient the network.




Fig. 23 Resilience state space

The outer control loop reduces the impact of a given challenge in the future, as shown in Figure 23, in which the challenged state S'_c is not as bad as the previous S_c , and remediation performs better with S'_r resulting in a smaller area \mathbb{R}' and better overall resilience. Temporal analysis considers the improvement not only for the static condition of $\mathbb{R}' < \mathbb{R}$, but also for reduced remediation and recovery times: $t'_r \leq t_r$ and $t'_R \leq t_R$.

4.2.6 Multilevel Multiscenario Resilience \mathbb{R}

In the multilevel analysis, as shown in Figure 24, 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 characterizing its service to the layer above.




Fig. 24 Resilience across multiple levels

By beginning at the bottom level and progressing up the service layers, an overall multilevel resilience value can be computed [79], and by composing these across all scenarios of interest for a given network architecture, it may be possible to derive a single resilience value \mathcal{R} .

4.3 Example Analyses

In order to demonstrate the application of this framework, we conduct a static resilience analysis of example ISP networks at the topology-service layer (3t) wherein the objective is to study the impact of node and link failures on the topology, followed by the routing-service layer (3r) in which the objective is to construct a path given a (hopefully connected) topology. Note that DTNs (disruption tolerant networks) are primarily concerned with the case in which layer 3t is unable to deliver stable topologies over which layer 3r can create persistent end-to-end paths. We divide the traditional network layer 3 into topology and routing sublayer services.

In the layer 3t 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. Therefore $\mathcal{N}_{3t} = \{N_1\}$. In this example, we define the topology service by selecting two service 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. While reachability directly affects the number of pairs that are reachable, the clustering coefficient affects how the local paths will be affected by link failures. Therefore, $\mathcal{P}_{3t} = \{P_1, P_2\}$. We conduct simulations in MATLAB to evaluate the

impact of link failures on the service parameters at this boundary. We explore the operational link-failure-probability metric over the range of $[0.0, 0.5]$ to commercial US ISPs and a European research network, the latter significantly smaller and less connected. The simulation results are averaged over 100 runs. The purpose of this example is to show how the proposed metrics framework can be applied at a service boundary given a certain set of service constraints expressed in terms of what is acceptable, impaired, and unacceptable.

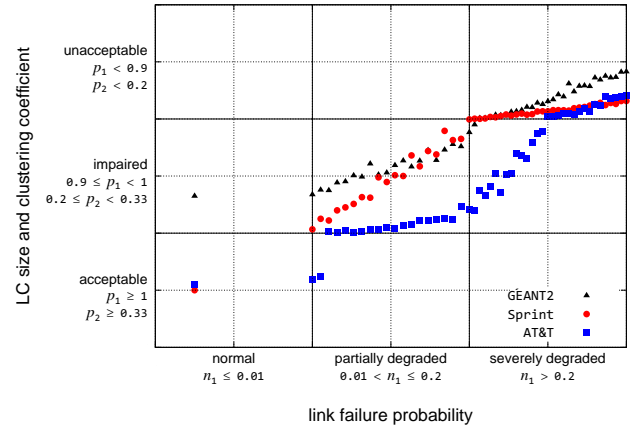


Fig. 25 Comparing resilience of ISP topologies

Given these operational and service regions, we plot the simulation results on a piece-wise linear axis. Figure 25 compares the steady-state resilience of the Rocketfuel-inferred [4] AT&T, Sprint L3, and actual GÉANT2 [9] network topologies to link failures as degradation in the service from the acceptable to unacceptable region. The region boundaries in both the operational and service dimensions are arbitrarily chosen based on operational targets and service requirements.

We see that the AT&T and Sprint networks lie in the acceptable service region under normal operating conditions; the research GÉANT2 network does not because it is not richly enough connected to meet this service specification even in normal operation. Given the rich connectivity of the AT&T network, the service remains acceptable even when the network starts degrading. However, as the failures continue the network eventually moves to impaired service. As the network operational conditions are severely degraded, the service transitions from the impaired to unacceptable region. The Sprint network provides unacceptable service in the presence of a single link failure. In order to get an aggregate measure of resilience, we calculate the area under the curve formed by linear interpolation between the states. The smaller the area, the better is the

resilience; in the limiting case, if the service remains acceptable for all operational conditions, the area under the curve will be zero, representing perfect resilience $\mathbb{R} = 1$. In order to get a normalised value of resilience, we define resilience $\mathbb{R} = 1 - \text{normalised area}$, where *normalised area* is the total area divided by the span of the x -axis.

The resilience \mathbb{R} for AT&T is 0.63, for sprint 0.54 and for GÉANT2 0.47. We observe that due to a fewer number of links, the GÉANT2 topology has very low clustering coefficient and the topology service performs poorly even in the normal operational regions.




Fig. 26 Resilience of multi-path routing on AT&T topology

Now we go up a level in the analysis in which $\mathbb{N}_{3r} = \mathbb{P}_{3t}$ to analyse the resilience of the routing service given a topology. Figure 26 shows the resilience of the path-diversity mechanism explained in Section 3.2, when using 1, 2, and 3 paths over the AT&T topology. We observe that as the number of paths selected (k) increases, the service remains longer in the acceptable region and degrades more gracefully. Note that $k = 1$ represents unipath routing in which even a single link failure will result in failure of certain paths even if the network is connected. As expected, multipath routing is more resilient to poorly connected topologies. A significantly more detailed analysis with more examples (including MANET resilience) is presented in [79]; dynamic temporal analysis remains in future work. A related analytical approach computes the robustness R -value [145] and uses GraphExplorer [54].

5 Simulation Methodology

This section describes our simulation framework and methodology for understanding the resilience of networks and the impact of challenges. First the types of

challenges are presented. Then, the KU-CSM (The University of Kansas Challenge Simulation Module) is described, followed by a few example simulation results. More details are presented in [41, 40].

5.1 Simulation Framework

Simulation via abstraction is one of the techniques to analyse networks in a cost-effective manner. We have chosen the ns-3 [110] network simulator since it is open source, flexible, provides mixed wired and wireless capability (unlike ns-2), and the models can be extended. Unfortunately, the simulation model space increases multiplicatively with the different number of challenges and network topologies being simulated. Hence, for n different topologies subjected to c different challenges, $n \times c$ models have to be generated and simulated. KU-CSM decouples the challenge generation from topologies by providing a comprehensive challenge specification framework, thereby reducing the simulation model space to $n + c$ consisting of c challenges applied to n network scenarios. We have created an automated simulation model generator that arbitrarily combines network topologies and challenge specifications, thus increasing the efficiency of the simulation model generation process. Our simulation framework consists of four distinct steps as shown in Figure 27.

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. The third step is the automated generation of ns-3 simulation code based on the topology and challenge descriptor. Finally, we run the simulations with traffic sources and protocols as appropriate, and analyse the network performance under challenge scenarios. Additionally, the simulation framework can also be enabled to generate ns-3 network animator (NetAnim) traces for visualisation purposes.

5.2 Challenge Modelling

A *challenge* is an event that impacts normal operation [133]. A *threat* is a potential challenge that might exploit a *vulnerability*. 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* [28]. Challenges to the normal operation of networks include unintentional misconfiguration or operational mistakes, malicious attacks, large-scale disasters, and environmental




Fig. 27 Framework flow diagram

challenges [133,134]. Network challenges can be categorised based on *intent*, *scope*, and *domain* they impact [41]. Considerably more detail on challenge modelling is presented in a companion paper [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 [87, 117, 83]. 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 [84]. Furthermore, geographically correlated failures can result from dependency among critical infrastructures, as experienced in the 2003 Northeast US blackout [93, 51].

A *domain* of a challenge is *wired* or *wireless*; some challenges can affect both, particularly area-based challenges. Challenges to the wired domain include fibre-optic cable cuts and failure of switching nodes and transmission equipment. Challenges that are inherent in the wireless domain include weakly connected channels, mobility of nodes in an ad-hoc network, and unpredictably long delays [134]. 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 in wireless communication networks [80]. Malicious nodes may jam the signal of legitimate users to impair communication in the open wireless medium.

While these challenge model categories are orthogonal to one other, particular 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.

5.3 Implementation of Challenge Models

Modelling and simulating network performance under challenge conditions is non-trivial [115]. There have been several studies that analyse different aspects of networks under challenges (see [40]). Here we briefly describe the way in which challenges are implemented in KU-CSM.

5.3.1 Non-malicious challenges

In the case of wired domain challenges in this category, the number of nodes or links k and challenge period is specified in the challenge specification file. This type of challenge models failures that are uncorrelated with respect to topology and geography, that is, random node and link failures.

5.3.2 Malicious attacks

We use topological properties of the graph in order to determine the *critical* elements in the network based on properties such as the degree of connectivity of nodes and betweenness of nodes and links [97, 107]). The critical nodes or links are shut down for the duration of the challenge period to simulate an attacker with knowledge of the network structure.

5.3.3 Large-scale disasters

The challenge specification for area-based challenges is an n -sided polygon with vertices located at a particular set of geographic coordinates or a circle centered at a user specified coordinates with radius r . The simulation framework then determines the nodes and links that are encompassed by the polygon or circle, and disables them during the challenge interval using the Computational Geometry Algorithms Library (CGAL) [1]. We also implement dynamic area-based challenges, in which the challenge area can evolve in shape over time: scale (expand or contract), rotate, and move on a trajectory during the simulation. Examples of the need

to simulate arbitrary polygons are to model large-scale power blackouts [51, 29, 48] and large-scale natural disasters such as hurricanes and earthquakes [50, 52, 49]. Circles are useful to model solar coronal mass ejections (CME) [5] and electromagnetic pulse (EMP) weapons [3].

5.3.4 Wireless challenges

To simulate challenges in the wireless domain, we create a new propagation loss model that includes a mobility model parameter and range of influence [41, 40]. Using these parameters, the user can specify where the loss takes place and how it moves over time. In this way, we model a realistic challenge instead of relying solely upon statistical methods. Unlike signal loss due to scattering and line-of-sight obstacles, jammers can cause radio channel interference that increase channel noise and reduce the signal to noise ratio that is critical to a receiver's ability to discern the data bits correctly. We implement a jammer module that sends high power signals with high data rate packets continuously on the same channel.

5.4 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 [131], shown in Figure 28 and Sprint fiber-optic topology [85] in Figure 5c. The traffic matrix for logical topologies consists of every node pair. For the fibre topology the traffic matrix only consists of MPLS PoP nodes as shown in Figure 29, since these are the only nodes that can inject or extract traffic, eliminating cities that only house amplifiers and regenerator nodes. The physical topology has 245 cities of which 90 MPLS PoP locations match the cities on the physical topology. Since not all cities are traffic source or sinks, the statistical failure scenarios would not be useful determining the performability of the network; therefore we do not do random node or link failures on the physical topology. A full explanation of the challenge specifications, as well as details of simulation parameters and further example results are presented in [41, 40].

5.4.1 Non-malicious and Malicious Challenges

First, we evaluate the performance of the Sprint topology (Figure 28) 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 30. We measure the instantaneous PDR at the steady-state



Fig. 28 Sprint inferred topology




Fig. 29 Sprint MPLS PoP locations

condition during which the challenge is causing a particular set of nodes and links to fail for a given time interval; we are not concerned here with route-convergence effects.

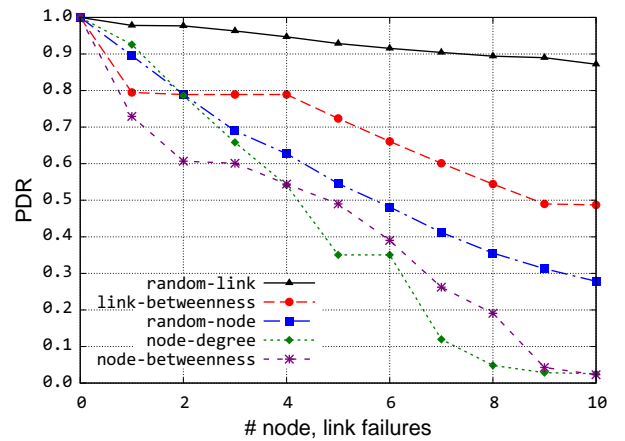


Fig. 30 PDR during non-malicious and malicious challenges

The top curve in Figure 30 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

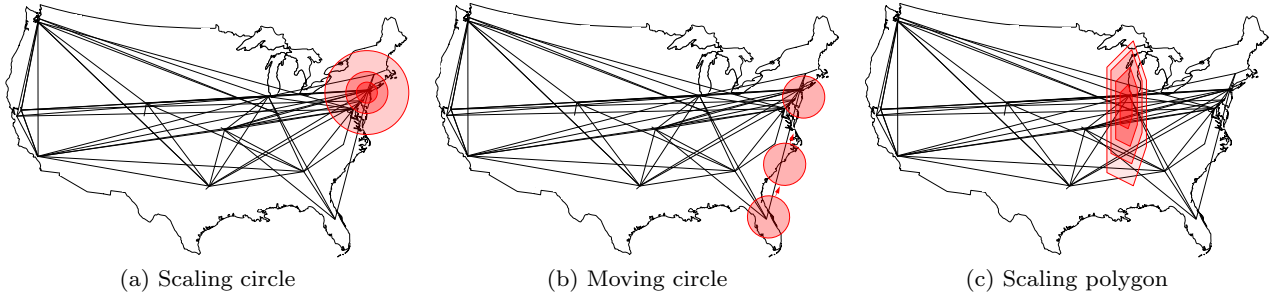


Fig. 31 Area-based challenge scenarios for Sprint PoP topology

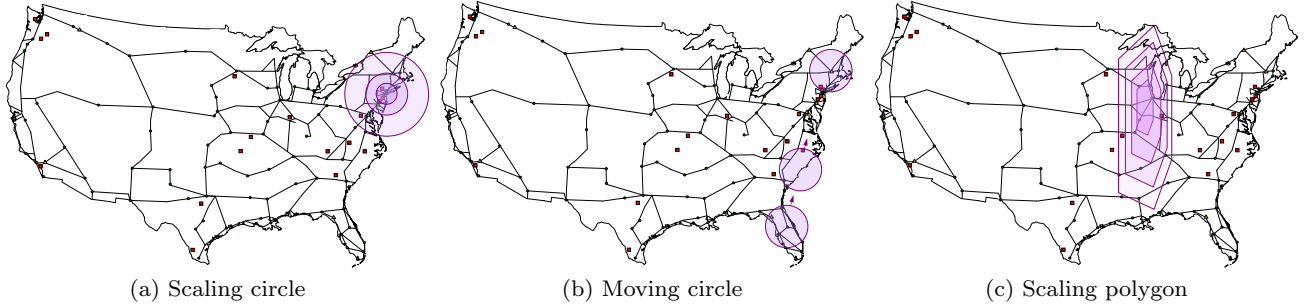


Fig. 32 Area-based challenge scenarios for Sprint physical topology




Fig. 33 PDR during area-based challenges

for each link in the topology, 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.

5.4.2 Area-based Challenges

Recently, the research community has recognised the importance of understanding the impact of geographically correlated failures on networks [98, 41, 31, 111, 112]. Our framework uses circles or polygons to model geographically correlated failures representative of large-scale disasters needed for network survivability [134, 59] analysis. Next, we present the results of three scenarios that demonstrate area-based challenges that evolve spatially and temporally using the Sprint logical and physical topologies, as shown in Figure 31 and Fig-

ure 32. 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 33, 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 as shown in Figure 31a and in Figure 32a for inferred and physical topologies respectively, with a radius of approximately 111 km. We choose the scenario to be representative of an electromagnetic pulse (EMP) attack [3]. The PDR is shown in Figure 33a. 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 drops depending on how many nodes and links are covered in each step for both physical and logical topologies.

Next, we demonstrate an area-based scenario that can evolve 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 31b and Figure 32b. 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 33b PDR reduces to 93% for the logical topology as the challenge starts only covering the node in Orlando at 10 sec and 82% for the physical topologies, since there are several PoPs being affected. As the challenge moves towards New York in its trajectory, the PDR reaches 1.0 at the 13 sec. In this case, the challenge area includes only the link between Orlando and New York, but since there are multiple paths for the Rocketfuel-inferred topology a single link failure does not affect the PDR, showing that *geographic diversity for survivability* is crucial [133]. On the other hand, for the physical topology in the same instance, the traffic source in Fairfax, South Carolina resides in the challenge area, therefore the cumulative PDR does not reach 100%. As the challenge moves into the Northeast US region, cumulative PDR values depend on the number of traffic sources being affected by the challenge area.

Polygons are useful to model specific geographic challenges such as power failures, earthquakes, and floods. For a scaling polygon example, we show a 6-sided irregular polygon in the Midwest US, roughly representative of the North American Electric Reliability Corporation (NERC) Midwest region [3], as shown in Figure 31c and in Figure 32c for logical and physical topologies respectively.

The PDR throughout the simulation is shown in Figure 33c. In this scenario, the edges of the irregular polygon increase 1.8 times every three sec. The characteristics of the network performability in this scenario is similar for both physical and logical topologies. The PDR drops as the challenge area affects Chicago in the smallest area at 10 sec. Despite the increase in the challenge area at 13 sec., the PDR improves due to completion of route reconvergence. As the area increases, the PDR drops as low as 40% since the network is partitioned. 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.




Fig. 34 South central area-based challenge scenario

Next, we demonstrate an area-based scenario representative of a hurricane hitting the South-central US as shown in Figure 34. In the smallest area are the physical nodes in New Orleans and Biloxi of which only New Orleans node is a MPLS PoP node generating and sinking traffic. In the second circular area challenge, the physical nodes are: New Orleans, Baton Rouge, Lafayette, Biloxi, and Mobile, in which 4 out of the 5 affected nodes are MPLS PoP nodes. In the largest affected area there are total of 10 physical nodes, 6 of which are MPLS PoP nodes. However none of the three circular challenge areas cover any logical links or nodes on the map in Figure 28, permitting us to investigate the differences between logical and physical topologies.

The network performance of physical and logical topologies when the South-central US region is challenged is shown in Figure 35. Since there are no nodes or links in the logical topology impacted, the PDR is 100%. On the other hand, the PDR of the physical topology drops to 98%, 91%, and 86% respectively as the challenge area covers more physical nodes and links. This demonstrates that it is imperative to study the impact of area-based challenges on the *physical* topologies.




Fig. 35 South central US challenge PDR

Traditional layer-3 logical topologies are insufficient to understand the impact of physical challenges against the network infrastructure.

6 Experimental Evaluation

This section describes experimental evaluation and cross-verification of resilience using a large-scale programmable testbed: GpENI.

6.1 GpENI Overview

The Great Plains Environment for Network Innovation – GpENI is an international programmable network testbed centered on a regional optical network in the Midwest US, providing flexible infrastructure across the entire protocol stack [135, 45]. The goal of GpENI is to build a collaborative research infrastructure enabling the research community to conduct experiments in Future Internet architecture. GpENI is funded in part by the US National Science Foundation GENI (Global Environments for Network Innovation) and GENI experimentation programs and by the EU FIRE (Future Internet Research and Experimentation) programme, and is affiliated with a project funded by the NSF FIND (Future Internet Design) program. Two of the key characteristics of GpENI needed for experimental evaluation of resilience are programmability at all levels and a large-scale flexible topology.

6.1.1 Programmability and Flexibility

The defining characteristic of GpENI is programmability of *all* layers, as shown in Table 4, implemented on a *node cluster* of general- and special-purpose processors. At the top layer Gush [10] provides experiment control and Raven [13] distributes code; both are software

Table 4 GpENI programmability layers

GpENI Layer	Programmability
experiment	Gush, Raven
7 application	PlanetLab
4 end-to-end	
3 router	Quagga, XORP, Click
topology	VINI
2 VLAN	DCN
lightpath	
1 photonics	site-specific

developed as part of the GENI program. Layer 7 and 4 programmability are provided by the GENIwrapper version of PlanetLab [11]. At layer 3, programmable routers are implemented in Quagga [12], XORP [16], and Click [6], supplemented by any other technology GpENI institutions should choose to deploy. Flexible network-layer topologies are provided by VINI [15]. At layer 2, dynamic VLAN configurations are provided by DCN-enabled managed Gigabit-Ethernet switches at the center of each GpENI node cluster. GpENI institutions directly connected to the optical backbone use DCN-enabled [7] Ciena switches to provide dynamic lightpath and wavelength configuration. At layer 1, the architecture even permits programmability at the photonic layer for switches that provide such support. Furthermore, each GpENI institution can connect site specific networking testbeds; plans include wireless, sensor, and cognitive radio testbeds (e.g. KUAR [102]). External users in the broader research community may request GpENI accounts with which to run network experiments.

6.1.2 Topology

GpENI is built around the core GpENI optical backbone centered in the Midwest, shown in the centre of Figure 36, among the principal institutions of KU (The University of Kansas), KSU (Kansas State University), UMKC (University of Missouri – Kansas City), and UNL (University of Nebraska – Lincoln), with connectivity to other Midwest US universities including the GMOC (GENI Meta-Operations Center). The optical backbone consists of a fibre-optic run from KSU to KU to the Internet2 PoP in Kansas City, interconnected with dedicated wavelengths to UMKC and UNL.

Each of the four core institutions has a node cluster that includes optical switching capabilities provided by a Ciena CoreDirector, with the ultimate goal of permitting flexible spectrum, wavelength, and lightpath configurations.

GpENI is extended to Europe across Internet2 to GÉANT2 and NORDUnet and then to regional or na-

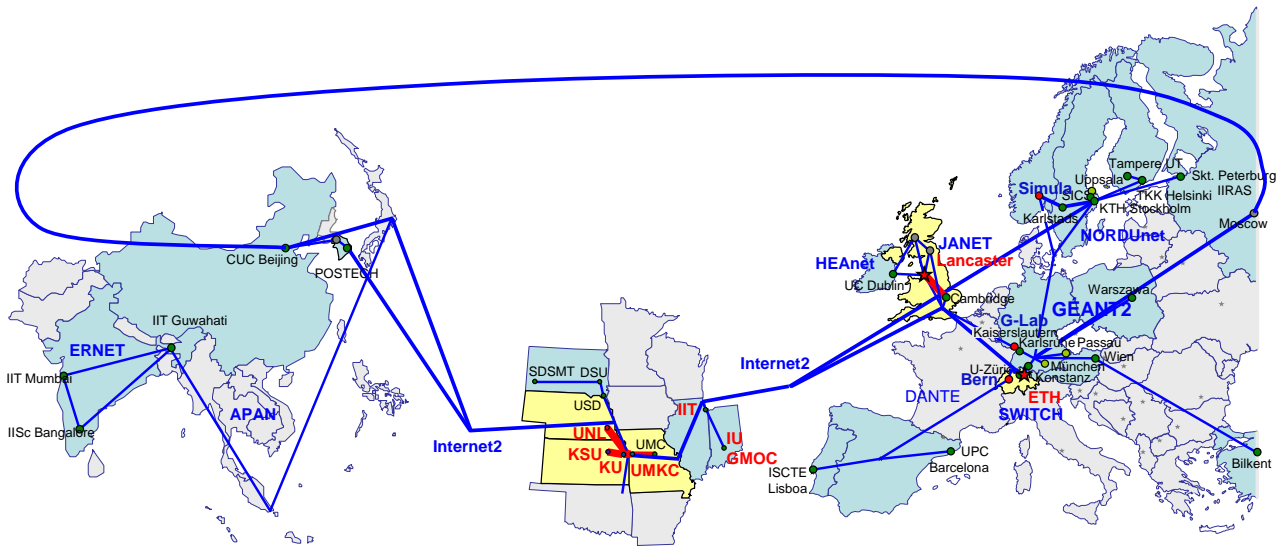


Fig. 36 GpENI topology

tional networks, as shown in Figure 36. Currently, connectivity is achieved using L2TPv3 and IP tunnels. A direct fibre link over JANET is deployed between Lancaster and Cambridge Universities. The principal European GpENI institutions are Lancaster University in the UK and ETH Zürich in Switzerland, with additional core nodes at Universität Bern in Switzerland, G-Lab Kaiserslautern in Germany, and Simula Research Laboratory in Norway. Similarly, GpENI is extended to Asia across Internet2 to APAN, then to national research network infrastructure including ERNET in India. Furthermore, GpENI is interconnected to the Emulab-based ProtoGENI cluster [18] in Utah, and is deploying several small ProtoGENI clusters of its own. Thus GpENI provides a large scale, rich topology on which to perform resilience experiments.

6.2 Experimental Evaluation of Resilience

Resilient topologies generated by KU-LoCGen using GpENI node geographic coordinates and analysed by KU-CSM can be used to generate layer-2 topologies that configure the topology of GpENI experiments. We can then evaluate performance when GpENI slices are challenged by correlated failures of nodes and links, measuring connectivity, packet delivery ratio, goodput, and delay, when subject to CBR (constant bit rate), bulk data transfer, and transactional (HTTP) traffic. We can also characterise the packet-loss probability of wireless links at the Utah Emulab [8], and the capabilities for emulating jamming and misbehaving nodes within the Emulab-federated CMU wireless emulator.

The goal is to cross-verify identical configurations of the simulated topologies and protocols discussed in

Sections 3 and 5 to GpENI experiments. GpENI experiments will have the advantage of incorporating real networking not easy to emulate in a simulation, albeit still at smaller scale than large simulation topologies.

7 Summary and Future Outlook

Resilience is an essential property of the Future Internet, including performability, dependability, and survivability. While a number of aspects of resilience have been an active area of research for a half-century, it is generally recognised that the Global Internet lacks resilience and is vulnerable to attacks and disasters. It is critical to make progress toward evaluating proposals for alternative topologies, protocols, and mechanisms that are candidates for deployment in the Future Internet.

However, we have lacked a comprehensive framework to evaluate the resilience of current and proposed network architectures, in part due to the complexity of the problem. This requires metrics to quantify resilience, and a tractable methodology to evaluate network resilience using appropriate abstractions in analysis, simulation, and emulation permitting cross-verification among these techniques.

This paper aims to contribute to this task by describing a comprehensive framework consisting of a resilience strategy, metrics for quantifying resilience, and evaluation techniques. The key to a tractable solution is multilevel composition of scenario-based evaluation of the resilience \mathbb{R} at each level, measured as the normalised inverse of the area under the trajectory through the \mathbb{N}, \mathbb{P} state space. Complex scenarios are simulated using the KU-LoCGen topology generator and KU-CSM

challenge simulation module in ns-3, which permit realistic challenge, topology, and protocol simulations, whose results can be mapped onto the state space for analysis. Much future work remains to be done to further refine the methodology, as well as to understand the properties of a network that make it resilient, and apply them to design and engineer the Future Resilient Internet.

Acknowledgements The authors would like to thank members of the ResiliNets research group at the University of Kansas, Lancaster University, as well as members of the EU ResumeNet project for discussions on, and contributions to aspects of this work. In particular we acknowledge David Hutchison and Paul Smith at Lancaster, and Marcus Schöller of NEC Laboratories, and Bernhard Plattner of ETH Zürich. We would like to thank Jacek Rak and Jon Crowcroft for inviting the RNDM keynote address and COMSNETS paper, respectively, that form the basis of this paper. We mourn the recent passing of Jean-Claude Laprie, whose seminal work in dependability is an important foundation for this work.

References

1. CGAL, Computational Geometry Algorithms Library. URL <http://www.cgal.org>
2. Protecting America's infrastructures. Report, President's Commission on Critical Infrastructure Protection (1997)
3. Report of the commission to assess the threat to the United States from electromagnetic pulse (EMP) attack. Report, Critical National Infrastructures (2004)
4. Rocketfuel: An ISP topology mapping engine (2008)
5. Severe space weather events: Understanding societal and economic impacts. Workshop report, National Research Council (2008)
6. The click modular router project. <http://read.cs.ucla.edu/click/> (2009)
7. Dynamic resource allocation via GMPLS optical network. <http://dragon.maxgigapop.net/> (2009)
8. Emulab: Network emulation testbed. <http://www.emulab.net/> (2009)
9. GÉANT2. <http://www.geant2.net/> (2009)
10. Gush: GENI user shell. <http://gush.cs.williams.edu/trac/gush> (2009)
11. PlanetLab. <http://www.planet-lab.org/> (2009)
12. Quagga routing suite. <http://www.quagga.net/> (2009)
13. Raven provisioning service. <http://raven.cs.arizona.edu/> (2009)
14. A roadmap for cybersecurity research. Technical report, Department of Homeland Security (DHS) (2009)
15. VINI: A virtual network infrastructure. <http://www.vini-veritas.net/> (2009)
16. XORP: Extensible open-source routing platform. <http://www.xorp.org/> (2009)
17. European information society. http://ec.europa.eu/information_society/policy/nis/strategy/activities/ciip/index_en.htm (2010)
18. ProtoGENI wiki. <http://www.protopeni.net/trac/protopeni> (2010)
19. Sprint network maps (2010). URL https://www.sprint.net/network_maps.php
20. UK resilience homepage. <http://www.cabinetoffice.gov.uk/ukresilience.aspx> (2010)
21. Map of Current Interstates. http://en.wikipedia.org/wiki/File:Map_of_current_Interstates.svg (2011)
22. Resilinet topology map viewer. <http://www.ittc.ku.edu/resilinet/maps/> (2011)
23. Africa: Population Density, 2000. Center for International Earth Science Information Network (CIESIN), Columbia University; and Centro Internacional de Agricultura Tropical (CIAT). 2005. Gridded Population of the World Version 3 (GPWv3). Palisades, NY: Center for International Earth Science Information Network (CIESIN), Columbia University. Available at <http://sedac.ciesin.columbia.edu/gpw>
24. Alderson, D., Doyle, J., Govindan, R., Willinger, W.: Toward an optimization-driven framework for designing and generating realistic Internet topologies. *SIGCOMM Computer Communication Review* **33**(1), 41–46 (2003)
25. Alderson, D., Li, L., Willinger, W., Doyle, J.C.: Understanding Internet topology: principles, models, and validation. *IEEE/ACM Trans. Netw.* **13**(6), 1205–1218 (2005)
26. Antonopoulos, A.: Metrication and performance analysis on resilience of ring-based transport network solutions. In: *GLOBECOM'99: Global Telecommunications Conference*, vol. 2, pp. 1551–1555 (1999)
27. Avizienis, A.: Design of fault-tolerant computers. In: 1967 Fall Joint Computer Conf., *AFIPS Conf. Proc.*, vol. 31, pp. 733–743. Thompson Books (1967)
28. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. *Transactions on Dependable and Secure Computing* **1**(1), 11–33 (January–March 2004)
29. Bacher, R., Näf, U.: Report on the blackout in Italy on 28 September 2003. report, Swiss Federal Office of Energy (SFOE) (2003)
30. Barabasi, A.L., Albert, R.: Emergence of Scaling in Random Networks. *Science* **286**(5439), 509–512 (1999). URL <http://www.sciencemag.org/cgi/content/abstract/286/5439/509>
31. Bassiri, B., Heydari, S.S.: Network survivability in large-scale regional failure scenarios. In: *Proceedings of the 2nd Canadian Conference on Computer Science and Software Engineering (C3S2E)*, pp. 83–87. ACM, New York, NY, USA (2009)
32. Beaudry, M.: Performance-related reliability measures for computing systems. *IEEE Transactions on Computers* **27**, 540–547 (1978)
33. Bhandari, R.: *Survivable Networks: Algorithms for Diverse Routing*. Kluwer Academic Publishers, Norwell, MA, USA (1998)
34. Bhattacharjee, B., Calvert, K., Griffioen, J., Spring, N., Sterbenz, J.P.G.: Postmodern internetwork architecture. Technical Report ITTC-FY2006-TR-45030-01, Information and Telecommunication Center, 2335 Irving Hill Road, Lawrence, KS 66045-7612 (2006)
35. Bouabene, G., Jelger, C., Tschudin, C., Schmid, S., Keller, A., May, M.: The autonomic network architecture (ANA). *IEEE Journal on Selected Areas in Communications (JSAC)* **28**(1), 4–14 (2010)
36. CAIDA: Cooperative Association for Internet Data Analysis (caida). <http://www.caida.org/home/> (2009). URL <http://www.caida.org/home/>
37. Calvert, K., Doar, M., Zegura, E.: Modeling Internet topology. *IEEE Communications Magazine* **35**(6), 160–163 (1997)
38. Calvert, K.L., Bhattacharjee, S., Zegura, E.W., Sterbenz, J.P.G.: Directions in active networks. *IEEE Communications Magazine* **36**(10), 72–78 (1998)
39. Carlson, J.M., Doyle, J.: Highly optimized tolerance: A mechanism for power laws in designed systems. *Phys. Rev. E* **60**(2), 1412–1427 (1999)
40. Çetinkaya, E.K., Broyles, D., Dandekar, A., Srinivasan, S., Sterbenz, J.P.G.: Modelling Communication Network Challenges for Future Internet Resilience, Survivability, and Disruption Tolerance: A Simulation-Based Approach. *Telecommunication Systems* (in this issue)
41. Çetinkaya, E.K., Broyles, D., Dandekar, A., Srinivasan, S., Sterbenz, J.P.G.: A comprehensive framework to simulate network attacks and challenges. In: *Proceedings of the 2nd IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM)*, pp. 538–544. Moscow, Russia (2010)
42. Center for International Earth Science Information Network (CIESIN), Columbia University; and Centro Internacional de Agricultura Tropical (CIAT). 2005. Gridded Population of the World Version 3 (GPWv3): Population Density Grids. Palisades, NY: Socioeconomic Data and Applications Center (SEDAC), Columbia University. URL <http://sedac.ciesin.columbia.edu/gpw>
43. Chakrabarti, D., Faloutsos, C.: Graph mining: Laws, generators, and algorithms. *ACM Comput. Surv.* **38**(1), 2 (2006)

44. Chen, Q., Chang, H., Govindan, R., Jamin, S.: The origin of power laws in Internet topologies revisited. In: *Proceedings of the 21st Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, vol. 2, pp. 608–617 (2002)
45. Cherukuri, R., Liu, X., Bavier, A., Sterbenz, J., Medhi, D.: Network virtualization in GpENI: Framework, implementation & integration experience. In: *IEEE/IFIP ManFI*. Dublin, Ireland (2011). (to appear)
46. Clark, D., Sollins, K., Wroclawski, J., Katabi, D., Kulik, J., Yang, X., Braden, R., Faber, T., Falk, A., Pingali, V., Handley, M., Chiappa, N.: New arch: Future generation Internet architecture. Technical report, DARPA, MIT, ISI (2003)
47. Clark, D.D., Wroclawski, J., Sollins, K.R., Braden, R.: Tussle in cyberspace: Defining tomorrow's Internet. *IEEE/ACM Transactions on Networking* **13**(3), 462–475 (2005)
48. Cowie, J.: Lights Out in Rio. <http://www.renesys.com/blog/2009/11/lights-out-in-rio.shtml> (2009)
49. Cowie, J.: Japan Quake. <http://www.renesys.com/blog/2011/03/japan-quake.shtml> (2011)
50. Cowie, J., Popescu, A., Underwood, T.: Impact of hurricane Katrina on Internet infrastructure. report, Renesys (2005)
51. Cowie, J.H., Ogielski, A.T., Premore, B., Smith, E.A., Underwood, T.: Impact of the 2003 blackouts on Internet communications. Preliminary report, Renesys Corporation (2003). (updated March 1, 2004)
52. Davis, T., Rogers, H., Shays, C., Others: A failure of initiative: The final report of the select bipartisan committee to investigate the preparation for and response to hurricane Katrina. Congressional Report H.Rpt. 109-377, US House of Representatives, Washington, DC (2006)
53. Dobson, S., Denazis, S., Fernández, A., Găiti, D., Gelenbe, E., Massacci, F., Nixon, P., Saffre, F., Schmidt, N., Zambonelli, F.: A survey of autonomic communications. *ACM Transactions on Autonomous and Adaptive Systems* **1**(2), 223–259 (2006)
54. Doerr, C., Hernandez, J.M.: A Computational Approach to Multi-level Analysis of Network Resilience. In: *Proceedings of the Third International Conference on Dependability (DEPEND)*, pp. 125–132 (2010)
55. Doyle, J.C., Alderson, D.L., Li, L., Low, S., Roughan, M., Shalunov, S., Tanaka, R., Willinger, W.: The “robust yet fragile” nature of the Internet. *Proceedings of the National Academy of Sciences of the United States of America* **102**(41), 14,497–14,502 (2005)
56. Edwards, N.: Building dependable distributed systems. Technical report APM.1144.00.02, ANSA (1994)
57. Ellinas, G., Stern, T.: Automatic protection switching for link failures in optical networks with bi-directional links. In: *Proceedings of the Global Telecommunications Conference (GLOBECOM)*, vol. 1, pp. 152–156 (1996)
58. Ellison, R., Fisher, D., Linger, R., Lipson, H., Longstaff, T., Mead, N.: Survivable network systems: An emerging discipline. Tech. Rep. CMU/SEI-97-TR-013, Software Engineering Institute, Carnegie Mellon University (1997)
59. Ellison, R.J., Fisher, D.A., Linger, R.C., Lipson, H.F., Longstaff, T., Mead, N.R.: Survivable network systems: An emerging discipline. Tech. Rep. CMU/SEI-97-TR-013, Software Engineering Institute, Carnegie Mellon University, PA (1999)
60. Erdős, P., Rényi, A.: On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci* **5**, 17–61 (1960)
61. Fabrikant, A., Koutsoupias, E., Papadimitriou, C.: Heuristically optimized trade-offs: A new paradigm for power laws in the Internet. *Lecture notes in computer science* pp. 110–122 (2002)
62. Fall, K.: A delay-tolerant network architecture for challenged internets. In: *SIGCOMM '03: Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 27–34. ACM, New York, NY, USA (2003)
63. Faloutsos, M., Faloutsos, P., Faloutsos, C.: On power-law relationships of the Internet topology. In: *SIGCOMM '99: Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communication*, pp. 251–262. ACM, New York, NY, USA (1999)
64. Francis, P., Jamin, S., Jin, C., Jin, Y., Raz, D., Shavitt, Y., Zhang, L.: IDMaps: a global Internet host distance estimation service. *IEEE/ACM Transactions on Networking* **9**(5), 525–540 (2001)
65. Frank, H.: Survivability Analysis of Command and Control Communications Networks—Part I. *IEEE Transactions on Communications*, [legacy, pre-1988] **22**(5), 589–595 (1974)
66. Frank, H., Frisch, I.: Analysis and Design of Survivable Networks. *IEEE Transactions on Communication Technology* **18**(5), 501–519 (1970)
67. Fry, M., Fischer, M., Karaliopoulos, M., Smith, P., Hutchison, D.: Challenge identification for network resilience. In: *Proc. of the IEEE 6th EURO-NF Conference on Next Generation Internet (NGI)*, pp. 1–8 (2010)
68. Gan, Q., Helvik, B.: Dependability modelling and analysis of networks as taking routing and traffic into account. In: *NGI '06: Proceedings of the Conference on Next Generation Internet Design and Engineering* (2006)
69. Gay, F., Ketelsen, M.: Performance evaluation for gracefully degrading systems. In: *Proc. of the 9th Annual Int. Symp. on Fault Tolerant Computing*, pp. 51–58 (1979)
70. Goodman, S., Lin, H.: *Toward a Safer and More Secure Cyberspace*. National Academies Press (2007)
71. Grover, W.D.: *Mesh-Based Survivable Networks*. Prentice Hall PTR Pearson, Upper Saddle River, New Jersey (2004)
72. Grover, W.D., Stamatelakis, D.: Cycle-oriented distributed preconfiguration: Ring-like speed with mesh-like capacity for self-planning network restoration. In: *Proceeding of the IEEE International Conference on Communications (ICC'98)*, vol. 1, pp. 537–543 (1998)
73. Haddadi, H., Rio, M., Iannaccone, G., Moore, A., Mortier, R.: Network topologies: inference, modeling, and generation. *Communications Surveys and Tutorials*, *IEEE* **10**(2), 48–69 (2008)
74. Hagin, A.A.: Performability, reliability, and survivability of communication networks: system of methods and models for evaluation. In: *Proceedings of the 14th International Conference on Distributed Computing Systems*, pp. 562–573 (1994)
75. Hameed, M.A., Jabbar, A., Çetinkaya, E.K., Sterbenz, J.P.: Deriving network topologies from real world constraints. In: *Proceedings of IEEE GLOBECOM Workshop on Complex and Communication Networks (CC-Net)*, pp. 415–419 (2010)
76. Hariri, S., Qu, G., Dharmagadda, T., Ramkishore, M., Raghavendra, C.S.: Impact analysis of faults and attacks in large-scale networks. *IEEE Security and Privacy* **01**(5), 49–54 (2003)
77. Heegaard, P.E., Trivedi, K.S.: Network survivability modeling. *Computer Networks* **53**(8), 1215–1234 (2009).

- Performance Modeling of Computer Networks: Special Issue in Memory of Dr. Gunter Bolch
78. Huslende, R.: A combined evaluation of performance and reliability for degradable systems. In: *Proceedings of the ACM Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, pp. 157–164. ACM Press, New York, NY, USA (1981)
 79. Jabbar, A.: A framework to quantify network resilience and survivability. Ph.D. thesis, The University of Kansas, Lawrence, KS (2010)
 80. Jabbar, A., Rohrer, J.P., Oberthaler, A., Çetinkaya, E.K., Frost, V., Sterbenz, J.P.G.: Performance comparison of weather disruption-tolerant cross-layer routing algorithms. In: *Proc. IEEE INFOCOM 2009. The 28th Conference on Computer Communications*, pp. 1143–1151 (2009)
 81. Jabbar, A., Shi, Q., Hameed, M., Çetinkaya, E.K., Sterbenz, J.P.: ResiliNets topology modelling. https://wiki.ittc.ku.edu/resilinets/Topology_Modeling (2011)
 82. Jackson, A.W., Sterbenz, J.P.G., Condell, M.N., Hain, R.R.: Active network monitoring and control: The SENCOMM architecture and implementation. In: *DARPA Active Networks Conference and Exposition (DANCE)*, pp. 379–393. IEEE Computer Society, Los Alamitos, CA, USA (2002)
 83. Jamakovic, A., Uhlig, S.: Influence of the network structure on robustness. In: *IEEE International Conference on Networks (ICON)*, pp. 278–283 (2007)
 84. Kitamura, Y., Lee, Y., Sakiyama, R., Okamura, K.: Experience with restoration of asia pacific network failures from taiwan earthquake. *IEICE Transactions on Communications* **E90-B**(11), 3095–3103 (2007)
 85. KMI Corporation: North American Fiberoptic Long-haul Routes Planned and in Place (1999)
 86. Knight, J.C., Strunk, E.A., Sullivan, K.J.: Towards a rigorous definition of information system survivability. In: *Proceedings of the DARPA Information Survivability Conference and Exposition DISCEX III*, pp. 78–89. Washington DC (2003)
 87. Kuhn, D.: Sources of failure in the public switched telephone network. *Computer* **30**(4), 31–36 (1997)
 88. Lakhina, A., Byers, J., Crovella, M., Matta, I.: On the geographic location of Internet resources. *IEEE Journal on Selected Areas in Communications* **21**(6), 934–948 (2003)
 89. Laprie, J.C.: Dependability: Basic concepts and terminology. Draft, IFIP Working Group 10.4 – Dependable Computing and Fault Tolerance (1994)
 90. Li, L., Alderson, D., Willinger, W., Doyle, J.: A first-principles approach to understanding the Internet's router-level topology. *SIGCOMM Computer Communication Review* **34**(4), 3–14 (2004)
 91. Liew, S., Lu, K.: A framework for network survivability characterization. In: *SUPERCOMM/ICC'92: Proceedings of IEEE International Conference on Communications (ICC)*, pp. 405–410 (1992)
 92. Liew, S., Lu, K.: A framework for characterizing disaster-based network survivability. *IEEE Journal on Selected Areas in Communications* **12**(1), 52–58 (1994)
 93. Liscouski, B., Elliot, W.J.: Final report on the august 14, 2003 blackout in the United States and Canada: Causes and recommendations. Tech. rep., U.S. – Canada Power System Outage Task Force (2004)
 94. Liu, Y., Mendiratta, V., Trivedi, K.: Survivability analysis of telephone access network. In: *Proceedings of the 15th International Symposium on Software Reliability Engineering*, pp. 367–378. IEEE Computer Society Washington, DC, USA (2004)
 95. Losq, J.: Effects of Failures on Gracefully Degradable Systems. In: *Proc. of 7th Fault-Tolerant Computing Symposium*, pp. 29–34 (1977)
 96. Lyons, R., Vanderkulk, W.: The use of triple-modular redundancy to improve computer reliability. *IBM Journal of Research and Development* **6**(2), 200–209 (1962)
 97. Mahadevan, P., Krioukov, D., Fomenkov, M., Dimitropoulos, X., Claffy, K.C., Vahdat, A.: The Internet AS-level topology: three data sources and one definitive metric. *ACM SIGCOMM CCR* **36**(1), 17–26 (2006)
 98. Mahmood, R.A.: Simulating challenges to communication networks for evaluation of resilience. Master's thesis, The University of Kansas, Lawrence, KS (2009)
 99. Medhi, D., Tipper, D.: Multi-layered network survivability-models, analysis, architecture, framework and implementation: An overview. In: *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX)*, vol. 1, pp. 173–186 (2000)
 100. Medina, A., Matta, I., Byers, J.: On the origin of power laws in Internet topologies. *SIGCOMM Computer Communication Review* **30**(2), 18–28 (2000)
 101. Meyer, J.: On Evaluating the Performability of Degradable Computing Systems. *IEEE Transactions on Computers* **100**(29), 720–731 (1980)
 102. Minden, G., Evans, J., Searl, L., DePardo, D., Petty, V., Rajbanshi, R., Newman, T., Chen, Q., Weidling, F., Guffey, J., Datla, D., Barker, B., Peck, M., Cordill, B., Wyglinski, A., Agah, A.: KUAR: A flexible software-defined radio development platform. In: *2nd IEEE International Symposium on New Frontiers in Dynamic Spectrum Access Networks (DySPAN)*, pp. 428–439 (2007)
 103. Mohammad, A.J., Hutchison, D., Sterbenz, J.P.G.: Towards quantifying metrics for resilient and survivable networks. In: *Proceedings of the 14th IEEE International Conference on Network Protocols (ICNP)*, pp. 17–18 (2006)
 104. Molisz, W.: Survivability function—a measure of disaster-based routing performance. *IEEE Journal on Selected Areas in Communications* **22**(9), 1876–1883 (2004)
 105. Molisz, W., Rak, J.: End-to-end service survivability under attacks on networks. *Journal of Telecommunications and Information Technology* **3**, 19–26 (2006)
 106. Molisz, W., Rak, J.: f-cycles—a new approach to providing fast service recovery at low backup capacity overhead. In: *10th Anniversary International Conference on Transparent Optical Networks (ICTON)*, vol. 3, pp. 59–62 (2008)
 107. Molisz, W., Rak, J.: Impact of WDM network topology characteristics on the extent of failure losses. In: *12th International Conference on Transparent Optical Networks (ICTON)*, pp. 1–4 (2010)
 108. Moore, E., Shannon, C.: Reliable circuits using less reliable relays. *Journal of the Franklin Institute* **262**(3), 191–208 (1956)
 109. Motiwala, M., Elmore, M., Feamster, N., Vempala, S.: Path splicing. In: *Proceedings of the ACM SIGCOMM conference on data communication*, pp. 27–38. ACM, New York, NY, USA (2008)
 110. The ns-3 network simulator. <http://www.nsnam.org> (2009)
 111. Neumayer, S., Modiano, E.: Network reliability with geographically correlated failures. In: *Proc. of IEEE INFOCOM*, pp. 1–9 (2010)

112. Neumayer, S., Zussman, G., Cohen, R., Modiano, E.: Assessing the vulnerability of the fiber infrastructure to disasters. In: Proc. of IEEE INFOCOM, pp. 1566–1574 (2009)
113. Ng, T.S.E., Zhang, H.: Predicting Internet network distance with coordinates-based approaches. In: INFOCOM, pp. 170–179 (2001)
114. Ng, Y., Avizienis, A.: A Reliability Model for Gracefully Degrading and Repairable Fault-tolerant Systems. In: Proceedings of 7th International Symposium on Fault-Tolerant Computing, pp. 29–34. IEEE Computing Society Publications (1977)
115. Nicol, D.M., Sanders, W.H., Trivedi, K.S.: Model-based evaluation: From dependability to security. IEEE Transactions on Dependable and Secure Computing **01**(1), 48–65 (2004)
116. Nussbaumer, J., Patel, B.V., Schaffa, F., Sterbenz, J.P.G.: Networking requirements for interactive video on demand. IEEE Journal on Selected Areas in Communications **13**, 779–787 (1995)
117. Oppenheimer, D., Ganapathi, A., Patterson, D.A.: Why do Internet services fail, and what can be done about it? In: Proc. of USENIX USITS, pp. 1–16 (2003)
118. Pierce, W.: Failure-tolerant Computer Design. Academic Press (1965)
119. Qu, G., Jayaprakash, R., Hariri, S., Raghavendra, C.: A framework for network vulnerability analysis. In: CT '02: Proceedings of the 1st IASTED International Conference on Communications, Internet, Information Technology, pp. 289–298. St. Thomas, Virgin Islands, USA (2002)
120. Rak, J.: k-penalty: a novel approach to find k-disjoint paths with differentiated path costs. IEEE Communications Letters **14**(4), 354–356 (2010)
121. Richards, C.W.: Map of the Month: Mainline Tonnage 1980 / 2005 (2007)
122. Rohrer, J.P., Jabbar, A., Perrins, E., Sterbenz, J.P.G.: Cross-layer architectural framework for highly-mobile multihop airborne telemetry networks. In: Proceedings of the IEEE Military Communications Conference (MILCOM), pp. 1–9. San Diego, CA, USA (2008)
123. Rohrer, J.P., Jabbar, A., Sterbenz, J.P.G.: Path diversification: A multipath resilience mechanism. In: Proceedings of the IEEE 7th International Workshop on the Design of Reliable Communication Networks (DRCN), pp. 343–351. Washington, DC, USA (2009)
124. Rohrer, J.P., Naidu, R., Sterbenz, J.P.G.: Multipath at the transport layer: An end-to-end resilience mechanism. In: Proceedings of the IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), pp. 1–7. St. Petersburg, Russia (2009)
125. Schaeffer-Filho, A., Smith, P., Mauthe, A.: Policy-driven Network Simulation: a Resilience Case Study. In: 26th ACM Symposium on Applied Computing (SAC). Taichung, Taiwan (2011)
126. Schneider, F.: Trust in Cyberspace. National Academies Press (1999)
127. Schöller, M., Smith, P., Rohner, C., Karaliopoulos, M., Jabbar, A., Sterbenz, J., Hutchison, D.: On realising a strategy for resilience in opportunistic networks. In: Future Network and Mobile Summit, pp. 1–8 (2010)
128. Siganos, G., Faloutsos, M., Faloutsos, P., Faloutsos, C.: Power laws and the AS-level Internet topology. IEEE/ACM Transactions on Networking **11**(4), 514–524 (2003)
129. Smith, P., Hutchison, D., Banfield, M., Leopold, H.: On understanding normal protocol behaviour to monitor and mitigate the abnormal. In: Proceedings of the IEEE/IST Workshop on Monitoring, Attack Detection and Mitigation (MonAM), pp. 105–107. Tuebingen, Germany (2006)
130. Smith, P., Schaeffer-Filho, A., Ali, A., Schöller, M., Kheir, N., Mauthe, A., Hutchison, D.: Strategies for Network Resilience: Capitalising on Policies. In: B. Stiller, F. De Turck (eds.) Mechanisms for Autonomous Management of Networks and Services, *Lecture Notes in Computer Science*, vol. 6155, pp. 118–122. Springer Berlin / Heidelberg (2010)
131. Spring, N., Mahajan, R., Wetherall, D.: Measuring ISP Topologies with Rocketfuel. IEEE/ACM TON **12**(1), 2–16 (2004)
132. Sterbenz, J.P.G., Hutchison, D.: Resilinet: Multilevel resilient and survivable networking initiative wiki. <http://wiki.ittc.ku.edu/resilinet> (2008)
133. Sterbenz, J.P.G., Hutchison, D., Çetinkaya, E.K., Jabbar, A., Rohrer, J.P., Schöller, M., Smith, P.: Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. Computer Networks: Special Issue on Resilient and Survivable Networks (COMNET) **54**(8), 1245–1265 (2010)
134. Sterbenz, J.P.G., Krishnan, R., Hain, R.R., Jackson, A.W., Levin, D., Ramanathan, R., Zao, J.: Survivable mobile wireless networks: issues, challenges, and research directions. In: WiSE '02: Proceedings of the 3rd ACM workshop on Wireless security, pp. 31–40. ACM Press, New York, NY, USA (2002)
135. Sterbenz, J.P.G., Medhi, D., Ramamurthy, B., Scoglio, C., Hutchison, D., Plattner, B., Anjali, T., Scott, A., Buffington, C., Monaco, G.E., Gruenbacher, D., McMullen, R., Rohrer, J.P., Sherrell, J., Angu, P., Cherukuri, R., Qian, H., Tare, N.: The Great plains Environment for Network Innovation (GpENI): A programmable testbed for future internet architecture research. In: Proceedings of the 6th International Conference on Testbeds and Research Infrastructures for the Development of Networks & Communities (Trident-Com), pp. 428–441. Berlin, Germany (2010)
136. Sterbenz, J.P.G., Touch, J.D.: High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication, 1st edn. Wiley (2001)
137. Strand, J., Chiu, A., Tkach, R.: Issues for routing in the optical layer. IEEE Communications Magazine **39**(2), 81–87 (2001)
138. Styron, H.C.: CSX tunnel fire: Baltimore, MD. US Fire Administration Technical Report USFA-TR-140, Federal Emergency Management Administration, Emmitsburg, MD (2001)
139. Sydney, A., Scoglio, C., Youssef, M., Schumm, P.: Characterising the robustness of complex networks. International Journal of Internet Technology and Secured Transactions **2**, 291–320 (2010)
140. T1A1.2 Working Group: Network survivability performance. Technical Report T1A1.2/93-001R3, Alliance for Telecommunications Industry Solutions (ATIS) (1993)
141. T1A1.2 Working Group: Enhanced network survivability performance. Technical Report T1.TR.68-2001, Alliance for Telecommunications Industry Solutions (ATIS) (2001)
142. T1A1.2 Working Group: Reliability-related metrics and terminology for network elements in evolving communications networks. American National Standard

- for Telecommunications T1.TR.524-2004, Alliance for Telecommunications Industry Solutions (ATIS) (2004)
143. Telecom Regulatory Authority of India: The Indian Telecom Services Performance Indicators April–June 2008. Quarterly press release **109** (2008)
 144. Trivedi, K., Kim, D., Roy, A., Medhi, D.: Dependability and security models. In: Proceedings of the International Workshop of Design of Reliable Communication Networks (DRCN), pp. 11–20. IEEE (2009)
 145. Van Mieghem, P., Doerr, C., Wang, H., Hernandez, J., Hutchison, D., Karaliopoulos, M., Kooij, R.: A Framework for Computing Topological Network Robustness. Tech. rep., Delft University of Technology (2010). URL http://www.nas.ewi.tudelft.nl/people/Piet/papers/RobustnessRmodel_TUDreport20101218.pdf
 146. Wang, C., Byers, J.W.: Generating representative ISP topologies from first-principles. In: Proceedings of the ACM international conference on measurement and modeling of computer systems (SIGMETRICS), pp. 365–366. ACM, New York, NY, USA (2007)
 147. Waxman, B.: Routing of multipoint connections. Selected Areas in Communications, IEEE Journal on **6**(9), 1617–1622 (1988)
 148. Winick, J., Jamin, S.: Information Security: Computer Hacker Information Available on the Internet. Tech. Rep. T-AIMD-96-108, United States General Accounting Office (1996). URL <http://www.fas.org/irp/gao/aimd-96-108.htm>
 149. Winick, J., Jamin, S.: Inet-3.0: Internet topology generator. Tech. Rep. UM-CSE-TR-456-02, EECS, University of Michigan (2002). URL <http://topology.eecs.umich.edu/inet/inet-3.0.pdf>
 150. Yook, S., Jeong, H., Barabasi, A.: Modeling the Internet's large-scale topology. Proceedings of the National Academy of Sciences **99**(21), 13,382–13,386 (2002)
 151. Zegura, E., Calvert, K., Bhattacharjee, S.: How to model an internetwork. In: Proceedings of the 15th Annual Joint Conference of the IEEE Computer Societies (INFOCOM), vol. 2, pp. 594–602 (1996)

Bios

Dr. James P.G. Sterbenz: is Associate Professor of Electrical Engineering & Computer Science and on staff at the Information & Telecommunication Technology Center at The University of Kansas, and is a Visiting Professor of Computing in InfoLab 21 at Lancaster University in the UK. He received a doctorate in computer science from Washington University in St. Louis in 1991, with undergraduate degrees in electrical engineering, computer science, and economics. He is director of the ResiliNets research group at KU, PI for the NSF-funded FIND Postmodern Internet Architecture project, PI for the NSF Multilayer Network Resilience Analysis and Experimentation on GENI project, lead PI for the GpENI (Great Plains Environment for Network Innovation) international GENI and FIRE testbed, co-I in the EU-funded FIRE ResumeNet project, and PI for the US DoD-funded highly-mobile airborne networking project. He has previously held senior staff and

research management positions at BBN Technologies, GTE Laboratories, and IBM Research, where he has lead DARPA- and internally-funded research in mobile, wireless, active, and high-speed networks. He has been program chair for IEEE GI, GBN, and HotI; IFIP IWSOS, PfHSN, and IWAN; and is on the editorial board of IEEE Network. He has been active in Science and Engineering Fair organisation and judging in Massachusetts and Kansas for middle and high-school students. He is principal author of the book High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication. He is a member of the IEEE, ACM, IET/IEE, and IEICE. His research interests include resilient, survivable, and disruption tolerant networking, future Internet architectures, active and programmable networks, and high-speed networking and systems.

Egemen K. Çetinkaya: is a Ph.D. candidate in the department of Electrical Engineering and Computer Science at The University of Kansas. He received the B.S. degree in Electronics Engineering from Uludağ University (Bursa, Turkey) in 1999 and the M.S. degree in Electrical Engineering from University of Missouri–Rolla in 2001. He held various positions at Sprint as a support, system, design engineer from 2001 until 2008. He is a graduate research assistant in the ResiliNets research group at the KU Information & Telecommunication Technology Center (ITTC). His research interests are in resilient networks. He is a member of the IEEE Communications Society, ACM SIGCOMM, and Sigma Xi.

Mahmood A. Hameed: received the B.S. degree in Electronics and Communications Engineering from Osmania University (Hyderabad, India) in 2005 and the M.S. degree in Electrical Engineering from the University of Kansas in 2008. He is currently working toward the Ph.D. degree in Electrical Engineering at the University of Kansas, where he is engaged in research in energy efficient optical technologies for switching in future networks. His research interests include nonlinear optics, advanced optical modulation formats as well as signal processing.

Dr. Abdul Jabbar: is currently a Research Engineer in the Advanced Communication Systems Lab at GE Global Research in Nikayuna, NY. He received his Ph.D in Electrical Engineering from The University of Kansas in 2010 with honors. He received his M.S. degree in Electrical Engineering from KU in 2004 and B.S. degree in Electrical Engineering from Osmania University, India in 2001. He also holds an Adjunct Research Associate position at the KU Information & Telecommunication Technology Center. His interests include resilience and survivability, network algorithms, design

and analysis of network architectures, topologies, and protocols, highly dynamic networks, wireless access, and future networks. Abdul is the recipient of Moore award for best M.S. thesis and is a member of IEEE Communications Society, IEEE Computer Society, and the ACM Special Interest Group on Data Communications.

Shi Qian: received the M.S. degree in Electrical Engineering from The University of Kansas in 2009. He was a graduate student at the Information & Telecommunication Technology Center (ITTC).

Justin P. Rohrer: is a Ph.D. candidate in the department of Electrical Engineering and Computer Science at The University of Kansas. He received the B.S. degree in Electrical Engineering from Rensselaer Polytechnic Institute, Troy, NY, in 2004. From 1999 to 2004, he was with the Adirondack Area Network, Castleton, NY as a network engineer. He is currently a graduate research assistant at the KU Information & Telecommunication Technology Center (ITTC) and an IFT Graduate Fellow, and was an ITTC Graduate Fellow from 2004–2006. He received the best paper award at the International Telemetering Conference in 2008. His research focus is on resilient and survivable transport protocols. His interests also include highly-dynamic mobile networks, simulating network disruptions, and developing the GpENI network testbed for the GENI program. Previous research has included weather disruption-tolerant mesh networks and free-space-optical metropolitan networks. He is a member of the IEEE Communications and Computer Societies, ACM SIGCOMM, Eta Kappa Nu, and is an officer of the Kansas City section of the IEEE Computer Society.