# A Framework to Quantify Network Resilience and Survivability

*Abdul Jabbar*

Submitted to the graduate degree program in Electrical Engineering & Computer Science and the Graduate Faculty of the University of Kansas School of Engineering in partial fulfillment of the requirements for the degree of Doctor of Philosophy.

**Committee:**

Dr. James P.G. Sterbenz (chair)

Dr. Tyrone Duncan

Dr. Victor S. Frost

Dr. David Hutchison

Dr. Gary Minden

Dr. Caterina Scoglio

Dr. Alexander M. Wyglinski

Date Defended

The dissertation Committee for Abdul Jabbar certifies that this is the approved version of the following thesis:

**A Framework to Quantify Network Resilience and Survivability**

**Committee:**

______________________________

(Chairperson)

______________________________

______________________________

______________________________

______________________________

______________________________

______________________________

Date Approved:________________________

Page left intentionally blank.

# Abstract

The significance of resilient communication networks in the modern society is well established. Resilience and survivability mechanisms in current networks are limited and domain specific. Subsequently, the evaluation methods are either qualitative assessments or context-specific metrics. There is a need for rigorous quantitative evaluation of network resilience. We propose a service oriented framework to characterize resilience of networks to a number of faults and challenges at any abstraction level. This dissertation presents methods to quantify the operational state and the expected service of a network using functional metrics. We formalize resilience as transitions of the network state in a two-dimensional state space quantifying network operational characteristics and service parameters. One dimension represents the network as normally operating, partially degraded, or severely degraded. The other dimension represents network service as acceptable, impaired, or unacceptable. Our goal is to initially understand how to characterize network resilience, and ultimately how to guide network design and engineering toward increased resilience. We apply the proposed framework to evaluate the resilience of three ISP backbone network topologies and the topologies generated using a new realistic topology generator. Furthermore, we quantify the resilience of MANETs at multiple layer boundaries. We develop new predictive routing algorithms for weather disruption-tolerant networks, which are shown to be more resilient than the existing protocols. Lastly, we develop a new routing algorithm for highly-mobile ad hoc airborne networking.

Behind this dissertation are two women to whom I dedicate this work.
My mom gave me the inspiration through her sacrifice.
My wife, who is my best friend, gave me the support though her love.

# Acknowledgments

I would like to thank my advisor Dr. James P.G. Sterbenz for guiding me in my research and providing immense support throughout my graduate work. Thanks to Dr. Gary J. Minden for helping me with the formalization of the metrics framework. I would also like to thank Dr. Victor S. Frost for his technical help and guidance in my work on weather disruption tolerant networks. Thanks are due to Dr. Alexander M. Wyglinski, Dr. Tyrone Duncan, Dr. David Hutchison, and Dr. Caterina Scoglio for their valuable comments on my dissertation.

I would like to thank my colleagues in ResiliNets group with whom I collaborated on various aspects of my graduate research. Particularly, I would like to acknowledge Justin P. Rohrer for joint work on path diversification and WDTN, Egemen K. Çetinkaya, Mahmood Abdul Hameed, and Qian Shi for the collaboration on topology generation, Hemanth Narra for help with ns-3 simulations, and David Hutchison, Marcus Schöller, and Paul Smith for the collaboration on ResiliNets architecture.

Lastly, I would like to thank all the staff at the Information and Telecommunication Technology Center for their support through the years.

Page left intentionally blank.

# Contents

# List of Figures

# List of Tables

Page left intentionally blank.

# Chapter 1

# Introduction and Motivation

Society increasingly relies on computer networks as essential for individuals, businesses, and governments. These networks include the Global Internet, PSTN (public switched telephone network – wired and mobile), SCADA networks (supervisory control and data acquisition), and emerging sensor and mobile ad hoc networks. They have developed into large scale systems with increasing complexity both in terms of physical infrastructure as well as the operational protocols and user applications. Essential services are provided by distributed networked systems in the sectors of energy, finance, banking, education, health care, defense, transport, and communication. More recently, personal communication and sensor networks have exploded in to the mainstream market. A number of new services have emerged that depend on the dependability of these wireless networks.

The consequences of disruption to either legacy or emerging networks are thus increasingly severe, and threaten the lives of individuals, the financial health of business and other organizations, as well as the economic stability and security of nations and the world. Canonical examples include the dependence of military operations on the Global Information Grid [3] and interdependency of the Internet and the electrical grid for power [4]. With this increasing importance of, and reliance on the networks, so follows their increasing attractiveness as a target from bad guys: recreational and professional crackers, terrorists, and from information warfare. The U.S commission on critical infrastructure concluded that network attacks have the potential to be catastrophic [5]. The need for resilience in communication infrastructure has been widely recognized as a priority [6–10].

However, malicious attacks are not the only challenge that post-modern internetworks face. Given the diversity and heterogeneity in the infrastructure, technology, and applications there is a complete set of challenges that a systematic resilience approach must consider. For example, emerging mobile and wireless networks are not only susceptible to malicious attacks but also face environmental challenges due to the open nature of their communication channels. In the following section, we attempt to characterize various challenges faced by communication networks.

## 1.1 Challenges to Communication Networks

Let us consider the various challenges affecting both traditional and evolving networks. We characterize these challenges in to the following five categories:

1. **Environmental:** The challenges to the communication environment include high-mobility of nodes and subnetworks; weak, asymmetric, and episodic connectivity of wireless channels; and unpredictably-long-delay paths either due to length (e.g. satellite) or as a result of episodic connectivity.

2. **Malicious:** These include attacks against the network hardware, software, or protocol infrastructure from recreational crackers, industrial espionage, terrorism, or warfare.

3. **Non-malicious:** Failures due to misconfiguration or operational errors; unusual but legitimate traffic load (e.g. flash crowds); accidents leading to component failures.

4. **Large scale disasters:** Large-scale natural disasters such as hurricanes, earthquakes, ice storms, tsunami, and floods as well as man made disasters such as large scale power failures.

5. **Lower level failures:** The service failure at a lower level (or layer) is a challenge to the higher layers. For example, the failure of paths (at link layers) is a challenge to the network layer. A detailed discussion of this challenge is presented in Section 4.1.1

A detailed explanation of each of these challenge categories along with examples of past failures is presented in the [11]. Besides the above mentioned challenges, the network is also subjected to inherent faults in the system such as design flaws, defective components, and software bugs. These faults manifest as dormant faults in the network and activate only when triggered either through challenges or normal operation of the system. Given the fact that the networks are constantly facing a variety of challenges, we investigate *how to measure the ability of the network to tolerate these challenges*.

## 1.2 Goodness Measure

In order to study the impact of the above mentioned challenges on various systems, several disciplines have evolved over time. These include *fault tolerance*, *survivability*, *disruption tolerance*, and *traffic tolerance*. We present the formal definitions below:

- **Fault Tolerance:** "The ability of a functional entity to mask or mitigate the impact of faults on its specified operation" [12]. In other words, it is the ability of a system to tolerate component faults such that service failures do not result. Fault tolerance generally covers single or at most a few faults, and is thus a subset of survivability. It is generally used in the context of component failures.

- **Survivability:** "The capability of a system to fulfill its mission, in a timely manner, in the presence of attacks, failures, or accidents." [13, 14]. It is often agreed that survivability handles a larger set of faults and failures as compared to other measures (in particular, fault tolerance).

- **Disruption Tolerance:** "The ability of a system to tolerate disruptions in connectivity among its components, consisting of the environmental and tolerance of energy (or power) challenges" [11].

- **Traffic Tolerance:** "The ability of a system to tolerate unpredictable offered load without a significant drop in carried load (including congestion collapse), as well as to isolate the effects from cross traffic, other flows, and other nodes" [11].

In order to differentiate networks and services in terms of their ability to meet the objectives of the above mentioned disciplines, there is a need for some sort of *goodness measure*. There have been several such measures proposed in the literature and many of these are often applied in industry as well. Depending upon the context of service being offered, these measures characterize and quantify different aspects of a system. Furthermore, many of the measures defined in the literature overlap with one another. Some of the well-known measures and their definitions are:

- **Reliability:** "The probability that an entity (unit) will complete its intended mission (i.e. perform a required function) as required over a specified period of time in its intended environment (or stated conditions)" [12, 15, 16].

- **Availability:** "The proportion of the operating time in which an entity meets its in-service functional and performance requirements in its intended environment" [12, 15, 16].

- **Dependability:** "Dependability is that property of a system such that reliance can justifiably be placed on the service it delivers" [15]. There are different facets of dependability. The various attributes of dependability include availability (readiness for usage), reliability (continuity of service), correctness of service and maintainability [16, 17].

- **Performability:** Performability is the probability that the system will stay above a certain accomplishment level over a fixed period of time. It is often described by the QoS (quality of service) measures for a given set of operational conditions [18].

We use the term *resilience* to include *all* of these measures, defined as [2, 11]:

**Definition 1.2.1.** *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.*

Resilience includes the ability for users and applications to access information when needed (e.g. Web browsing and sensor monitoring), the maintenance of end-to-end communication association (e.g. tele- and video-conferences), and the operation of distributed processing and networked storage in the presence of challenges discussed in Section 1.1.

Henceforth, in this document, we shall refer to *resilience* as the property that is superset to all the other properties mentioned in this section.

## 1.3 Problem Statement

Researchers agree that due to lack of consistency in evaluating network resilience, it is difficult to guarantee that the networks being designed and developed would satisfy the requirements of both the end users and their applications [19]. Without standard metrics to measure the relative effectiveness of resilience mechanisms, it is difficult to identify potential solutions that lead to resilient networks. While it is clear that a number of new and innovative solutions are needed to provide network resilience, a key problem is *how to measure and specify resilience*. We need a methodology to *measure* the resilience (or lack thereof) of current and proposed networks and *evaluate the benefit* of architectures, designs, and mechanisms. This methodology needs to be both rigorous in capturing service parameters and operational metrics, as well as tractable so that it is useful in practice. The challenge is to bring order to a fundamentally complex problem; we do not underestimate the difficulty in this task and note that the QoS (quality of service) community has struggled with a related problem for some years. The development of resilience measures is further complicated by the heterogeneity of communication networks. In other words, a resilience scheme that applies well to a specific network scenario may not work as well on a different network scenario.

Secondly, the resilience strategy must be *multilevel*. As we will argue in this dissertation, it is necessary to improve the resilience of a communication network at each layer in the protocol stack. Hence the proposed resilience evaluation methodology must support the multilevel approach. With this discussion, we arrive at our *thesis statement*:

> A multilevel two-dimensional state-space framework can be used to quantify the resilience of networked systems, and be the basis of understanding the resilience of current networks as well as evaluating the ability of new mechanisms to improve future network resilience, survivability, and disruption tolerance.

The goal of this dissertation is to develop such a framework and utilize it to develop and evaluate new resilience mechanisms for challenged networks, especially at the topology and routing sub-layers.

## 1.4 Proposed Solution

We propose a *new multilevel framework to measure and analyze network resilience* at a given layer boundary. Resilience is effectively quantified as robustness, which measures service degradation in the presence of challenges (perturbations) to the operational state of the network. Hence, the network can be viewed (at any layer) as consisting of two orthogonal dimensions as shown in Figure 1.1: one is the operational state of the network, which consists of its physical infrastructure and their protocols; the second dimension is the service being provided by the network and its requirements [20]. We

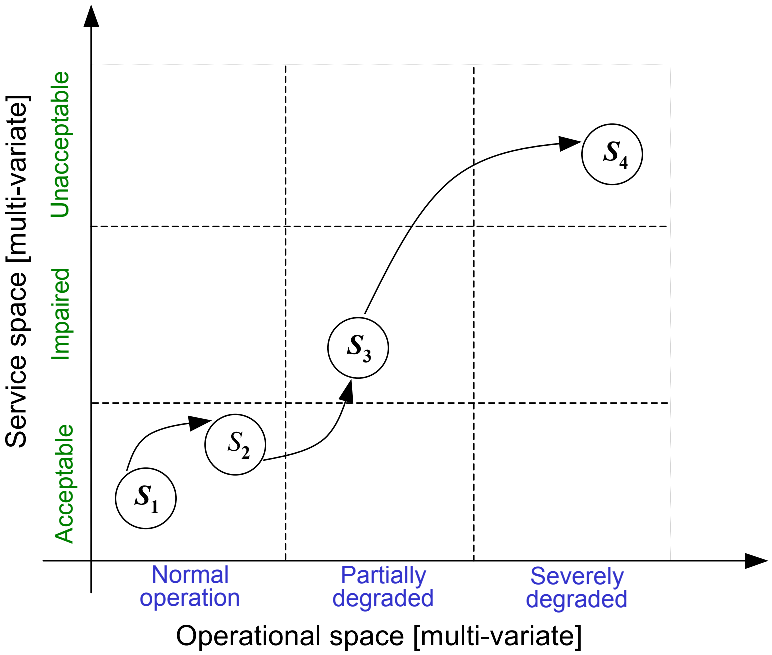

Figure 1.1: Network state space

characterize these two dimensions using *operational metrics* and *service parameters* respectively. Note that both of these dimensions are multi-variate and there is a clear mapping between the operational metrics and service parameters. Simply put, for a given set of operational conditions, the network provides a certain level of service for a given application. Thus, the *state* of a network is an aggregate of several points in this two dimensional space, represented as $S_i$ in Figure 1.1.

In order to limit the number of states, the operational and service space of the network may be divided into three regions each as shown in Figure 1.1. The network operational space is divided into *normal*, *partially degraded*, and *severely degraded* regions. Similarly, the service space is divided into *acceptable*, *impaired*, and *unacceptable* regions. While an arbitrary number of such regions is possible, one of the primary goals of this work is to achieve tractable yet useful solutions, and this set of nine $(3 \times 3)$ regions provides the necessary abstraction while limiting the number total regions. Each region may contain multiple states if the service demands such a fine granularity. In the limiting case, each region represents just one state.

Adverse events (challenges) are manifest as degradations in the operational condition of the network. When such events degrade the operational state of the network, the level of service being provided degrades as well resulting in *state transitions*, e.g., $S_1 \to S_2$ and $S_2 \to S_3$. We then evaluate *resilience* as the trajectory of the network through this state space and measure resilience as the area under this trajectory. The result is that based on a characterization of metrics for network operation and service requirement parameters, we are able to analyze and quantify the resilience of a network in the face of various faults, challenges, and attacks. We believe that this is a critical part of understanding, designing, and evaluating the resilience, survivability, and trustworthiness of networks.

Lastly, we utilize the proposed framework to devise a resilience strategy for defensive networks as well as for resilient topology and routing in case studies involving millimeter-wave networks and highly-dynamic airborne MANETs.

## 1.5 Contributions

The main contribution of this dissertation are as follows:

1. A multilevel resilience framework to quantify network resilience based on state transitions. We propose metrics to quantify the operational state of the network at various levels. Similarly, we define service parameters to quantify topology, routing, transport service.

2. A methodology to analyze the steady state resilience at any layer boundary. This enables the user to compare and contrast competing mechanisms in terms of their resilience benefits. For example, two different network topologies can be compared to see which one provides better resilience by evaluating the area under the state space trajectory.

3. A methodology for transient analysis of resilience strategy. The instantaneous states of the network shifts farther away from the origin as it experiences various challenges (failures, attacks, etc). On the other hand, the state transitions towards the origin due to remediation and recovery mechanisms. An analysis of these transition quantifies the overall resilience of a network in defending against attacks.

4. Resilient routing mechanisms in millimeter-wave mesh and aeronautical telemetry networks. Using the proposed resilience framework, we designed and evaluated resilient routing algorithm for two scenarios: a

cross-layered predictive weather-assisted routing protocol (P-WARP) to provide reliable communication in the presence of weather disruptions to millimeter-wave mesh networks and a location-aware highly-adaptive routing protocol (AeroRP) to provide acceptable service in a highly-dynamic airborne network.

5. Resilience topologies: We developed a model to generate realistic network topologies and used the resilience framework to evaluate the impact of graph properties on the resilience of the topologies for several applications. As a part of this work, we developed a location and cost constrained network topology generator (KU-LoCGen).

6. Resilience strategy: We developed formal methods to characterize the resilience strategy: $\mathrm{D}^2\mathrm{R}^2 + \mathrm{DR}$. We present case studies to show the impact of each phase of resilience strategy on the overall resilience of the network.

7. Evaluation of various scenarios in MATLAB and ns-3: We conducted simulations to demonstrate the use of the propose metrics framework in different scenarios. A simulation based analysis was conducted to quantify resilience at a number of levels (topology, routing, transport), for a number of different networks (ISP topologies, MANETs, and mesh networks).

# Chapter 2

# Related Work

While the generic concept of reliability has been in literature for a very long time, specific disciplines such as fault-tolerance, reliability, availability, dependability, vulnerability, and survivability have been rigorously studied over the past several decades. As discussed in Chapter 1, we use the term *resilience* as the superset of the above disciplines and properties. There is also a varied set of targeted systems over which resilience studies have been conducted. For example, resilience studies have been conducted on computer systems, operating systems, software development, control systems, and mechanical systems. More recently, telecommunication systems such as the PSTN (public switched telephone network) and Internet have become one of the focus areas for resilience analysis. Due to the vast nature of the subject area and the differences in the resilience analysis for each area, we will primarily focus our discussion on the previous research efforts in the field of telecommunication networks.

The earliest works in fault tolerance that were published include the 1965 Moore and Shannon paper [21] on reliable circuits and the Peirce [22] and Avizienis [23] papers on fault tolerant computing. The initial work on reliability and fault tolerance was focused on the design of computing systems [24]. In 1974, one of first resilience works in communication networks was presented by Frank [25, 26] 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 realized in late 1970's in the context of fault-tolerant computing systems [27–29]. Hence the concept of *degradable systems* was introduced in which the system has at least some degraded performance under the presence of faults with out failing completely. Markov models are used to evaluate the performance and reliability of the degradable system by Huslende [30], Meyer [31], Gay [32], and others. Meyer [31] first coined the term *performability* as the probability that the system will stay above a certain accomplishment level over a fixed period of time [18]. Until then, reliability and performance of communication networks were treated separately. Huslende [30] defined performance as the second dimension of the classical reliability, thereby defining reliability of a degradable system as the probability that the system will operate with a performance measure above certain threshold. The system is then represented with continuous time discrete states and the Markov model is used to calculate the reliability and availability based on the probability that the system will remain in states such that aggregate performance measure of all the visited states is above a certain threshold.

Often communication systems were considered a special class of distributed real-time systems. Hence, the same techniques that are used to evaluate distributed systems were used to evaluate the resilience of communication networks. Lately, rigorous analytical definitions of survivability, reliability and availability [13,17,33] were developed that apply to all systems in general.

In the following sections we present detailed computational methods to evaluate the resilience characteristics of telecommunication networks.

## 2.1 Resilience Computation Methods

We now further narrow our survey of literature to delve into the specific frameworks developed to compute and quantify the various aspects of resilience such as reliability, availability, vulnerability, and survivability. While considering papers in related areas, we focus primarily on communication network related frameworks. In order to provide structure to this large body of research, we divide it in to the following six categories: *topology based evaluation*, *survivability and reliability frameworks*, *security and vulnerability analysis*, *fault injection and visualization schemes*, *service oriented resilience metrics*, and *passive monitoring.*

### 2.1.1 Topology Based Evaluation

Research efforts have focussed on the two most common failure events: node and link failures. In other words, the connectivity of the underlying topology is a primary measure of resilience. Techniques commonly used for topology analysis include graph theory and linear algebra. This method is well suited to analyze the resilience of transport networks in which the reliability of the physical links is vital.

Liew and Lu suggest a survivability characterization schemes specifically in the presence of large-scale disasters [34, 35]. They propose a network survivability function that can be used to derive the actual survivability of the network with respect to the service metric of interest that measures the "goodness" of the network. The primary failures of the network considered are the node and link failures. The survivability function, $S$ is defined as:

$$P[S = s] = \sum_{e:S_e=s} P_e \tag{2.1}$$

where the $s$ is the fraction of nodes still connected and $S_e$ is the fraction of nodes that remain connected when a sample node failure event $e$ occurs with a probability of $P_e$. The *expected survivability* $E[S]$ and worst case survivability $s^0$, are then calculated as:

$$E[S] = \sum_{s} sP[S = s] \tag{2.2}$$

$$s^0 = \min_{P[S=s]>0} s \tag{2.3}$$

Similarly, the $r$-percentile survivability and the probability of zero survivability can be calculated.

Antonopoulos proposed a metrication framework to assess the performance of ring-based transport networks [36]. In this work, network unavailability due to link failures is the primary concern. Metrics based on path unavailability are used to determine the resilience of the network. A worst case path unavailability, $\mathrm{PU}_w$, defines the unavailability of the longest path in the network. The average unavailability, $\mathrm{PU}_a$, of a network with $n$ paths is defined as:

$$\mathrm{PU}_a = \frac{\sum_{i=1}^{n} \mathrm{PU}_i}{n} \tag{2.4}$$

Finally, the deviation of the network unavailabilities , $\Delta\mathrm{PU}$, is calculated as

$$\Delta\mathrm{PU} = \frac{\sum_{i=1}^{n} |\mathrm{PU}_i - \mathrm{PU}_a|}{n\mathrm{PU}_a} \tag{2.5}$$

where the unavailability $U_i$, of any path $i$ is given as

$$U_i = 1 - A_i = \frac{\mathrm{MTTR}}{\mathrm{MTBF} + \mathrm{MTTR}} \tag{2.6}$$

where MTTR is the mean time to recover and MTBF is the mean time between failures. Hence a measure of goodness of the network is given by $\mathrm{PU}_a$ and the distribution is given by $\Delta\mathrm{PU}$.

### 2.1.2 Survivability and Reliability Frameworks

In 2004, Westmark [37] surveyed an impressive 107 papers out of a selected 270 computational survivability papers from research databases such as ACM, IEEE and SEI. The objective of the author was to summarize the standard method to compute survivability. However, it was concluded that there is no consensus among the research community on the definition or evaluation methods of survivability. Furthermore, it was noticed that individual areas have unique interpretation of the idea. The author proposes a composite template to characterize survivability such that it is applicable to all the research areas in general:

> Survivability = the ability of a given system with a given intended usage to provide a pre-specified minimum level of service in the event of one or more pre-specified threats [37, pg. 7]

The author claims that survivability must be context-specific. Hence a survivability characterization must include the system environment, the usage service, minimum level of service and the specific threats.

As opposed to purely topological frameworks, Gan and Helvik [38] recognize that dependability models should also consider the dynamic and behavioral characteristics of the networks such as changes in topology, routes and traffic loads. Since the underlying network model has to consider above mentioned properties in addition to the connectivity, the state space increases dramatically subsequently leading to increased complexity. Stochastic activity networks (SAN) are used to model both the faults of components as well as the dynamics of the network. In order to tackle the problem of state-space explosion, the authors propose limiting state space at the model level by limiting the state generation to most probable states.

Another approach to compute resilience is to derive a closed form survivability function [39]. A survivability function is derived for random occurrence of failure (ROF) of nodes and links while considering the effects of routing protocol and traffic demand. Network survivability is defined as “the probability function of the percentage of total data delivered in the event of a failure” [39]. Furthermore, *survivability attributes* are used as a measure of disaster-based routing performance. The network is modeled as a directed graph $\Gamma(\boldsymbol{N}, A)$, where $\boldsymbol{N}$ is the set of nodes $|\boldsymbol{N}| = N$ and $\boldsymbol{A}$ is the set of directed arcs $|\boldsymbol{A}| = M$. The network topology is described by the incidence matrix $\boldsymbol{A}$ and the traffic demand between the nodes is specified by demand matrix $\mathbf{R}$. Each arch $e_m \in \boldsymbol{A}$ has a capacity given by $c_m$. A failure scenario $\zeta$ is defined a set of component failures, e.g. $N_D$ nodes out of $N$ nodes or $M_D$ links out of $M$ link. Assuming that each failure, $\zeta$ occurs with a specific probability $P(\zeta)$, and is independent of all other failures, the survivability function is then defined as:

$$S(x) = \sum_{\zeta : X(\zeta) = x} P(\zeta) \tag{2.7}$$

where the random variable $X(\zeta)$ represents the percentage of flow delivered after the failure $\zeta$. The "survivability function $S(x)$ is the sum of probabilities of all failure scenarios in which $x$ percent of total flow is still delivered". In order to simplify the calculations, uniform link failure distribution is assumed. Hence the probability of each link failure in a network with $M$ edges is $M^{-1}$. The probability of a failure scenario $P(\zeta)$ with $M_D$ faulty (failed) links is given by:

$$P(\zeta) = \prod_{l=1}^{M_D} \frac{1}{M} \prod_{k=1}^{M-M_D} \frac{M-1}{M} \tag{2.8}$$

The above probability function $P(\zeta)$ decreases very quickly with the increase in $M_D$. Therefore a cumulative probability mass function (CPM) is used to achieve a fixed threshold (e.g. 99%). Furthermore, the number of different scenarios (combinations) that lead to $M_D$ faults is given by:

$$L_{M_D} = \left( \begin{array}{c} M \\ M_D \end{array} \right) = \frac{M!}{M_D!(M-M_D)!} \tag{2.9}$$

The increase in the CPM due to the failure scenario $\zeta$ is given as:

$$v_{M_D} = L_{M_D} \times P(\zeta) \tag{2.10}$$

Hence the overall failure probability is obtained by adding the probability of a scenario with zero link failures, one link failure, two link failures, all the way up to all link failures. However, due to decreasing contribution from higher link failures, we can achieve a 99% CPM with relatively small number of scenarios. Finally, the authors also calculate the effect of shortest path routing given the overall failure probability. A linear programming formulation is used to derive the optimal feasible multicommodity flows for a given network. The survivability of the network is then calculated as the fraction of the total flow that is delivered considering the overall failure probability for a given traffic load. The traffic load is specified as a percent of the optimal feasible multicommodity flows. One such analysis on a Polish backbone network consisting of 12 nodes and 34 links has been conducted [39]. It was observed that the effect of failures is worst in case of higher traffic loads and node failures.

One rigorous survivability evaluation method is described by Knight [33], in which a system is said to be survivable if it satisfies or complies with the survivability specification. Therefore, much emphasis is placed on survivability specification, which is quoted as follows: [33]

**Definition 2.1.1.** *A survivability specification is a four-tuple,* $\{E, R, P, M\}$ *[33, 40] where:*

- *$E$ is a statement of the assumed operating environment for the system. It includes details of the various hazards to which the system might be exposed together with all of the external operating parameters. To the extent possible, it must include any anticipated changes that might occur in the environment.*

- *$R$ is a set of specifications each of which is a complete statement of a tolerable form of service that the system must provide. This set will include one distinguished element that is the normal or preferred specification, i.e. the specification that provides the greatest value to the user and with which the system is expected to comply most of the time.*

- *$P$ is a probability mass function across the set of specifications, $R$. A probability is associated with each member of the set $R$ with the sum of these probabilities being one. The probability associated with the preferred specification defines the fraction of operating time during which the preferred specification must be operational.*

- *$M$ is a finite-state machine denoted by the four-tuple $\{S, s_0, V, T\}$ with the following meanings:*

  - *$S$: A finite set of states each of which has a unique label which is one of the specifications defined in $R$.*
  - *$s_0$: $s_0 \in S$ is the initial or preferred state for the machine.*
  - *$V$: A finite set of customer values.*
  - *$T$: A state transition matrix.*

This definition of survivability is frequently used by a researchers developing methods to quantitatively evaluate survivability [40, 41].

### 2.1.3 Security and Vulnerability Analysis

Research has also been conducted to evaluate the resilience of networks with emphasis on the security aspects. Here the vulnerability of the networks to attacks and failures is used as a measure of the goodness of the network. More recently, methods developed to evaluate dependability such as state-based model checking and combinatorial techniques are being leveraged to evaluate the security of the system [42].

A framework to analyze the network as well as service vulnerabilities under the presence of attacks, specifically the distributed denial of service (DDOS) attacks has been developed [43, 44]. A three step method is described: 1) develop metrics to identify network vulnerabilities, 2) characterize the state of the network in to three states, and 3) evaluate a application or network vulnerability index (VI) based on the vulnerability metrics. The vulnerability index function is used to quantify the network services under attacks and failures.

The authors propose a set of global metrics to evaluate the effect of faults on the network systems. In order to monitor these metrics, an agent module is placed at each network entity that collects data and recognizes individual attacks as well as correlation effects due to one or more attacks. The data gathered from these agents is used to derive the vulnerability index with respect to each metric. Simulation conducted using three subnets with DDoS attacks showed that the VI correctly predicts the potential vulnerabilities in the network.

Similarly, Guirguis [45] developed metrics to quantify the damage cause due to reduction of quality (RoQ) type attacks and to measure the potency of the attack.

### 2.1.4 Fault Injection and Visualization Schemes

Various resilience evaluation schemes have been developed based on the fault injection method. The idea is to model a network and inject various expected and unexpected faults in the network. The fraction of the service delivered after the failure is the measure of resilience. This research is particularly helpful during the design and planning state of the network.

Jha and Wing present a survivability analysis framework for networked systems based on fault injection and visualization techniques [46,47]. The general method outlined in their research involves modeling the network using state machines. Faults are injected in to the network as node and link failures. The survivability metrics include the physical connectivity of the network and some measure of system service. The analysis method consists of generating the sequence of states the network traverses upon the introduction of faults. The service measure of the final state determines the survivability of the network for the given fault that was injected. In order to evaluate the reliability, the random occurring faults are considered. Bayesian networks are used to specify the probabilities of individual faults as well as correlated and conditional faults. Constrained Markov decision processes are used to evaluate the reliability of the network.

### 2.1.5 Service Oriented Resilience Metrics

The resilience of the network can be evaluated specifically in terms of the service delivered to the *end user*. At every abstraction level, the service of the network to higher entity is defined and statistical averages are used to quantify the its resilience against various faults and challenges.

The effects of network disruptions to end users has been widely studied. Zolfaghari and Kaudel evaluated the outage in the network as a $(U, D, W)$ triple, where the unservability (U) is the percentage of unsuccessful attempts at obtaining a service; duration (D) is the length of time for which the service is unavailable; and weight (W) is a measure of the impact of the service failure such as number of people and area affected [1]. The author presents a multi-layer approach, in which each layer serves is responsible to provide service to upper layers. Furthermore, the link capacity of a higher layer generates traffic for the layer below. Due to the multilayer approach, failures in given layer can be guarded against at that layer or at a higher layer. The authors also introduce two methods of survivability analysis based on the statistical nature of the failures. The first approach is based on random occurrence of failures (ROF), in which failures are assumed to adhere to a given probability distribution. The second method is a conditional approach, in which the survivability of the network is evaluated for a *given* failure. Furthermore, this paper introduces the concept of using different measures to quantify the survivability at different layers. A step-by-step procedure to calculate the survivability measures for both ROF and GOF methods is also presented. This procedure involves obtaining the failure occurrence data (random or given), defining a survivability measurement and then listing all the combination of events that can occur for the failure under consideration. For each event, the survivability measure is calculated.

An example of this approach is shown in Figure 2.1, where a survivability measure $x$, is being captured under a failure scenario. For a given survivability measure, the following survivability attributes can be defined [1].

$S_a$: fraction of $x$ that remains after failure (before restoration)$(1 - U)$

$S_u$: fraction of $x$ unservable after failure $(U)$

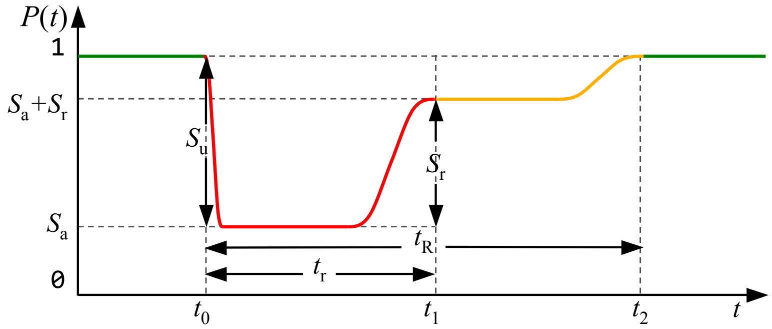

Figure 2.1: Example of survivability attributes (adapted from [1])

$S_r$: the fraction of $x$ that is restored at $t_1$

$t_0$: the time the failure occurs

$t_r$: the duration required until fraction $S_a + S_r$ of $x$ is restored

$t_R$: the duration required until all of $x$ ($t_2$) is restored (D)

Assuming that $S_a, S_u, S_r$ are random variables, the *survivability function* is the probability distribution function $f(y)$. Based on this approach, the expected values of the survivability attributes can be obtained as:

$$E[S_a] = \sum_y yP[S_a = y] \tag{2.11}$$

$$E[S_r] = \sum_y yP[S_r = y] \tag{2.12}$$

$$E[t_r] = \sum_y yP[t_r = y] \tag{2.13}$$

The above measures can be used to estimate the survivability of any given service layer against failures based on a given survivability measure.

The ATIS/ANSI T1A1.2 working group on network survivability performance have produced a number of documents on evaluation of network resilience and have recommended standard metrics to be used by the telecommunications industry. A framework to quantify the severity of failure events in terms of user service outage has been presented [12,19]. Based on the popular distinction in the industry between network supplier, service provider and end user, the metrics framework is divided in to the following two domains [19]:

- **Network service provider domain** consists of metrics that characterize the end users' service reliability requirements with respect to the cost of maintenance. For IP networks, these include the service downtime, service denial, loss of data probability, service failure, maintenance downtime, billing downtime, network upgrade cost and spares inventory cost. These various downtime metrics are specified as the percentage of the total operation time.

- **Network supplier's domain** consists of metrics to quantify the resilience of the network solutions such that the network service provider's requirements are met. These metrics deal with the network infrastructure and the technologies implemented. Sample metrics are service outage failure rate, interruption failure rate, service outage downtime, fault isolation, repairability, link and node restoration time, network failure containment.

The objective of the above metrics is to optimize the network infrastructure costs while maintaining the end user service reliability requirements. These metrics can be applied at various network abstraction levels.

Similarly, Grover [17] has established resilience measures such as *user-lost erlangs*, $(U, D, E)$ triples, and other restorability indexes for specific networks.

### 2.1.6 Passive Monitoring

Passive monitoring is a fairly simple method of evaluating network resilience. The network administrators monitor key places in a given network and gather data using dedicated hardware systems or software snooping programs. The gathered data is then post-processed to generate performance metrics. In the event of faults and challenges, post-processing can be used to determine the resilience of the network. Many Internet service providers (ISPs), business organizations and universities employ this method since it is economical and can be done at any stage of network operation.

## 2.2 Summary

Based on this literature survey, it is obvious that network resilience has been evaluated in several different contexts and levels. In summary, resilience is evaluated as node and link connectivity based on topology; as point-to-point or overall capacity of the network; as an ability to survive potential attacks; or as the reliability seen by the end user. Previous research efforts have targeted different levels of the network. While some focus on the logical layer (transport network connectivity), others have focused on application layer (user metrics). While these methods capture some aspects of resilience in specific scenarios and layers, they do not present a *comprehensive view of resilience at all levels*. Resilience evaluation at any level must consider both aspects of the communication network: the operational condition of the network and the service delivered. Hence we arrive at the question: *At what abstraction level should resilience be evaluated and what is the best evaluation methodology*?

Page left intentionally blank.

# Chapter 3

# ResiliNets Architecture

This chapter presents a multi-layer resilience architecture from the perspective of metrics. The metrics framework proposed in this dissertation is specifically designed to support the evaluation of strategy and mechanisms of this architecture. First we define the fundamental axioms on which the architecture is based. Then we present a unique resilience strategy followed by a set of principles with which the architecture can be deployed.

There are several existing frameworks that take a systematic approach towards various sub-categories or aspects of resilience. The ANSA (Advanced Networked Systems Architecture) project [48] framework characterizes dependability using the two-dimensional space of value and time. The ATIS (Alliance for Telecommunications Industry Solutions) workgroup on network survivability has developed a framework to quantify network survivability in a service oriented fashion. This multilevel framework quantifies survivability over physical, system, logical, and service levels using the notion of unservability [12, 19] as discussed in Section 2.1.5. The CMU architecture specifies a 4-step strategy to resilience consisting of resistance, recognition, recovery, and adaptation and evolution [49]. The architecture presented in this chapter focusses on the resilience of the network at any arbitrary level, which we defined earlier to be subsuming of all related disciplines [1].

## 3.1 Axioms

The ResiliNets architecture is based on the following four fundamental axioms [2].

### Axiom A0. Inevitability of Faults

Faults are inevitable. It is not possible to construct perfect fault-free systems, nor is it possible to prevent challenges and threats. A fault is a flaw in the system that can cause an error [15, 50]. Internal faults are a result of imperfections in the design of the system and external faults are triggered by challenges. It is not possible to design fault free systems due to the difficulty in designing perfect systems and predicting all possible challenges nor practical under cost constraints. A metrics framework to quantify resilience should be able to quantify the impact of these faults on the system when activated by challenges.

[1] This work is a part of the **ResiliNets** initiative at the University of Kansas and Lancaster University [2], to which the author of this thesis provided substantial contributions [11]. It provides a fundamental approach to research and build resilient networks.

### Axiom A1. Understand Normal Operations

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 Every network is designed for a set of parameters that are derived from the anticipated scenarios. These parameters characterize the normal operations of the network when there are no challenges or attacks to the system. In order to understand how faults and challenges impact a system, it is necessary to first understand the *normal* behavior. Hence the metrics framework should be able to characterize the normal operations of the network as well as the expected service.

### Axiom A2. Expect Adverse Events and Conditions

Expectation and preparation for adverse events and conditions is necessary, so that that defenses and detection of challenges that disrupt normal operations can occur. These challenges are inevitable. Adverse events are challenges to normal operations of the network. These can be either anticipated events based on past experience or unanticipated events that cannot be predicted. However, in both cases, resilience warrants the system to be prepared to respond. One way this can be achieved is to understand the impact of any adverse event on the network operations. While there are infinitely many potential challenge scenarios, there is a finite range over which the network parameters vary. The resilience framework proposed in Chapter 4 addresses this issue by characterizing the expected service as the operational parameters of interest vary over their respective ranges.

#### Axiom A3. Respond to Adverse Events and Conditions

It is clear that 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. One of the key distinguishing feature of a resilient network is its ability to respond to adverse events in a manner that preserves the service being delivered. There are several phases involved in the life cycle of an attack over a resilient network. The ResiliNets architecture proposes a specific six stage strategy as discussed in the following section. Furthermore, the metrics framework proposed by this dissertation describes how to quantify the effectiveness of this strategy.

## 3.2 Strategy

There is a need for a comprehensive and rigorous strategy towards resilient networks. Give the fundamental axioms of Section 3.1, it is clear that the resilience of the network must be addressed before, during, and after an adverse event challenges the normal operations. This approach is consistent with the previously published literature on degradable systems. In this section, we present a six-step, two phase resilience strategy, formalized as $\mathrm{D}^2\mathrm{R}^2 + \mathrm{DR}$, under the presence of faults and challenges [11]. Note that this strategy as shown in Figure 3.1 motivates and determines the requirements for the metrics framework discussed in Chapter 4. We first present the strategy and then discuss the its implications on the metrics framework.

Before discussing the details of the strategy, we consider how challenges affect a system. In a communication network, a challenge can be viewed in two orthogonal dimensions. First, the challenge can impact the network operations by causing an active fault that may lead to an failures of the network components. This causes the network to deviate from its normal operations. Secondly, this deviation in the operating conditions of the network may result in the perturbation (degradation) in the service being provided by the network.

### 3.2.1 First Phase

The first phase of the strategy is the real-time active loop $\mathrm{D}^2\mathrm{R}^2$ with a passive core as shown in Figure 3.1. This consists of a cycle of four steps: *defend*, *detect*, *remediate*, and *recover* that are involved with the network operations and service simultaneously. Furthermore, more than one instance of this cycle may be active in the network, corresponding to individual network modules or associations, and responds to any network challenge or attack in real time.

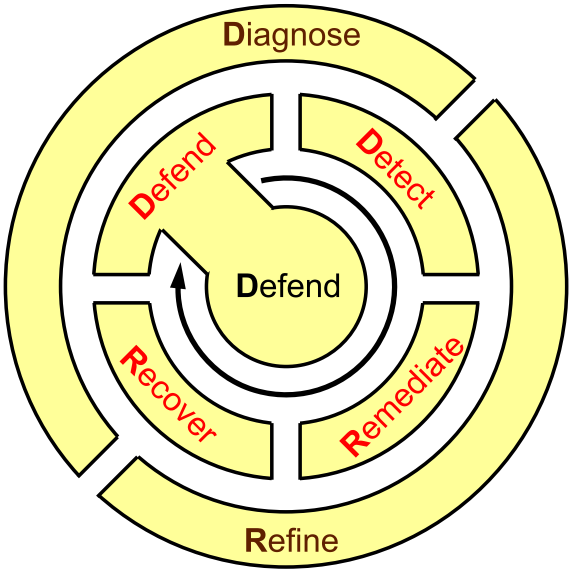

Figure 3.1: ResiliNets strategy

- **Defend** against challenges and threats to normal operation. The first step in any resilient system is to defend against challenges so that faults do not result in observable failures. Hence the network defenses first attempt to prevent the challenges from triggering failures resulting in deviation of the normal operations. Secondly, the defense also tries to isolate these perturbations to the network operational from the service being provided by the network.

  The first aspect of defense is referred commonly as fault tolerance and survivability and can be achieved by using passive mechanisms such as redundancy and diversity. The second aspect of the defense is to prevent the service from being affected even if the challenges occur by using active mechanisms of self-protection, e.g. filtering, firewalls. Examples include use of connectivity paradigms that enable the end-to-end communication in the face of unstable paths. However, it is not always possible to be completely defend against challenges, which is when the system must *detect* these failures.

- **Detect** when an adverse event or condition has occurred. As discussed above, challenges can and will overwhelm network defenses leading to failures in network operations and service. In this case, it is necessary for the network to detect this event.

  As with defense, detection can occur in two dimensions. First, the detection mechanism can detect the changes in operational conditions. This requires the system to understand normal operations (Axiom 1) and detect deviations from it. Secondly, the system can detect changes in the service parameters. In order to achieve this, the system must understand the primary service requirements.

  In order to support the first mode of detection, we need to characterize the normal behavior of the network. This requirement provides the first guideline for the metrics framework (Chapter 4): the need for metrics to characterize the operational state of the network and a means to distinguish normal state from errored state. Similarly, in order to support the second mode of detection, the metrics framework must provide a measure to characterize the service as well as a mechanism to distinguish the quality of service being provided.

- **Remediate** the effects of the adverse event or condition to minimize the impact.

  When the defense of the network fails and this is detected, the next best thing is to remediate the effect of the challenge (perturbation in operational conditions) on the service delivery. In other words, the system must take measures to minimize the impact of the failure. For example, when a link fails, the system must reroute the traffic around failed link. Depending upon the severity of the event, the remediation mechanism may not be able to keep the service levels to the ones prior to the adverse event, in which case it must gracefully degrade. This requires the system to be adaptable and autonomic.

  This step of the strategy specifies the second major guideline for the metrics framework: *the ability to characterize performability.* Not only should the framework characterize normal operations and service properties, it must also quantify the degradation in the service with respect to degradation in operations.

- **Recover** to original and normal operations.

  Once the adverse event passes over, the network must recover to normal operations. This requires the network to detect the end of the specific event and the recovery mechanism should restore the network to its original state. For example, after the failed link is repaired, the network should reroute on the original path (assuming it was the best available path). The recovery mechanism returns the network operations to the normal state and subsequently the service quality should return to its original value. Metrics help both to determine when the challenge is over and when the network has been restored to its original state.

In order to analyze the resilience of a network under the presence of adverse events, it is necessary to characterize each of the these phases. Hence the metrics framework must support representation of each step individually as well as the entire cycle collectively. In Chapter 4, we present exactly such an evaluation framework.

### 3.2.2 Second Phase

The second phase consists of two steps: *diagnose* and *refine* (DR). These background processes observe and modify the behavior of the $D^2R^2$ cycle as shown in Figure 3.1.

- **Diagnose** the fault that has been the root cause of an error or failure.

  Faults cannot be observed or detected directly, they can only be detected only when they manifest as errors or failures. Hence following an adverse event that was successful in causing a failure, it is necessary to diagnose the fault that was the root cause of the failures. For example, lack of sufficient redundancy in paths (a design fault) can be diagnosed after a link cut (challenge or attack) results in a route failure.

- **Refine** behavior for the future based on past $D^2R^2$ cycles.

  After observing how the $D^2R^2$ cycle performed in the presence of adverse events, the next step is to refine, enhance, and evolve the process so that the network is more resilient to future challenges. In the above example, a redundant link may be added as a hot standby so that future link cuts do not result in route failures. Metrics measure the effectiveness of $D^2R^2$ inner loop to help analyze how to refine.

In order to determine which refinement mechanisms yield the highest gain in resilience, it is critical to develop a metrics framework that facilitates such an evaluation. We carry the requirements presented by this strategy to the following chapter in designing a metrics framework for quantitative resilience evaluation.

## 3.3 Principles

Based on the axioms presented in Section 3.1, synthesis of the strategy presented in Section 3.2 and past experience of resilience disciplines, a detailed set of guiding principles has been developed for resilient network [11]. While the entire set of principles are summarized in Table 3.1, in this section, we discuss those specific principles that are either directly impacted by or set requirements for the metrics framework.

### 3.3.1 Prerequisites

The following are the prerequisites necessary to build resilient networks:

P1. **Service requirements** of applications need to be determined to understand the level of resilience the system should provide.

P2. **Normal behaviour** of the network is a combination of design and engineering specification, along with monitoring while unchallenged to learn the networks normal operational parameters.

P3. **Threat and challenge models** are essential to understanding and detecting potential adverse events and conditions.

P4. **Metrics** quantifying the service requirements and operational state are needed to measure the operational state and service state to detect and remediate and quantify resilience to refine future behavior and thus are critical for quantifying P1 – P3 above.

In order to build resilient systems the measurement, evaluation, and understanding of both the operations (normal behaviour) and service (requirements) aspect of the network, along with the challenges, is a prerequisite. Without a framework to quantify these aspects of the network, one cannot successfully move towards resilience. Therefore, we propose a metrics framework that facilitates these computations in Chapter 4.

### 3.3.2 Enablers

The ResiliNets framework identifies a set of enablers that can help fundamentally build resilient systems. These include:

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.

P11. **Redundancy** in space, time, and information increases resilience against faults and some challenges if defenses are penetrated. This includes spatial, temporal, informational, operational and implementational redundancies.

P12. **Diversity** in space, time, medium, and mechanism increases resilience against challenges to particular choices. Diversity consists of providing alternatives so that even when challenges impact particular alternatives, other alternatives prevent degradation from normal operations.

P13. **Multilevel resilience** Multilevel resilience 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 Internetwork including attached end systems.

All the *enablers* specified above attempt to improve resilience a network. However, for a given scenario, it is difficult to evaluate which one or combination of enables yields the best resilience. Another way to look at the problem is: given a finite fixed cost, which enablers should be applied for a specific scenario. Our approach to this very complex question is to quantify the resilience of the network based on a set of operational metrics, thereby determining which parameters have the highest sensitivity to the resilience of that particular network and choose enablers to optimize those parameters. A more detailed analysis of this method will be presented in Chapter 4

### 3.3.3 Behaviour

The ResiliNets architecture proposes that amongst other behavior one of the key resilience property is the self-organizing and autonomic behavior as discussed below:

P16. **Self-organizing and autonomic behaviour** is necessary for network resilience that is highly reactive with minimal human intervention. The phases of autonomic networking consist of initialization, autoconfiguration, self-organization, self-managing, self-optimizing, self-diagnosing, and self-repair.

In order to evaluate this autonomic behaviour, we need to evaluate the resilience of the network in real-time as the network experiences adverse events.

While the prerequisite and enabling principles require a steady-state resilience evaluation, the behavioral principles require an analysis of transient resilience. In the following chapter we will introduce a metrics framework than can achieve both of these objectives.

Table 3.1: Summary of resilience principles [2]

| | |
|---|---|
| P1 | Service requirements determine the need for network resilience |
| P2 | Normal behavior must be specified, verified, and refined through monitoring to understand normal operations |
| P3 | Threat and Challenge Models are essential to understanding and detecting potential adverse events and condition |
| P4 | Metrics are needed to measure and engineer network resilience |
| P5 | Heterogeneity in mechanism, trust, and policy must be addressed |
| P6 | Resource tradeoffs determine the deployment of resilience mechanisms |
| P7 | Complexity of the network in general, and resilience in particular, must be reduced to maximize overall resilience |
| P8 | Multilevel resilience is needed with respect to protocol layer, protocol plane, and hierarchical network organization |
| P9 | Translucency is needed to control the degree of abstraction vs. the visibility between levels |
| P10 | Heterogeneity in mechanism, trust, and policy among different network realms is a reality of emerging multi-provider networks; resilient mechanisms must admit this heterogeneity |
| P11 | Redundancy in space and time increases resilience against faults and some challenges |
| P12 | Diversity in space, time, medium, and mechanism increases resilience against challenges to particular choices. |
| P13 | Self-organizing and autonomic behavior is necessary for network resilience that is highly reactive with minimal human intervention |
| P14 | Security and self-protection are essential properties of entities to defend against challenges in a resilient network |
| P15 | State management is an essential aspect of networks in general, and resilience mechanisms in particular; the alternatives of how to distribute and manage this state are critical to resilience |
| P16 | Connectivity and association among communicating entities should be maintained when possible, but information flow should still take place even when a stable end-to-end path does not exit |
| P17 | Context awareness is necessary for network components to operate autonomously to detect challenges |
| P18 | Adaptability to the network environment is essential for a node in a resilient network to detect, remediate, and recover from challenges |
| P18 | Evolvability is needed to refine future behavior to improve the response to challenges, as well as for the network architecture and protocols to respond to emerging threats and application demands |

Page left intentionally blank.

# Chapter 4

# Metrics Framework

We start this chapter by revisiting our thesis statement. *A multilevel two-dimensional state-space framework can be used to quantify the resilience of networked systems, and be the basis of understanding the resilience of current networks as well as evaluating the ability of new mechanisms to improve future network resilience, survivability, and disruption tolerance.*

Resilience of a communication network is conventionally specified as a change in specific performance measures under the presence of individual faults and challenges. For example, resilience may be specified as the percentage of traffic delivered (performance measure) following an edge router failure. With this approach, for a complete appraisal of resilience, one must evaluate the performance measure in all the possible challenge scenarios. However, it is neither feasible to foresee the unique challenges that a network might experience, nor practical to characterize and model every challenge. The only commonality across all challenges is that they manifest as degradations in the operational condition of the network. This implies that there is a direct mapping between the challenges and the operational state of the network in that the challenges cause the network to deviate from its normal behavior.

Therefore, we propose that resilience should be specified as the change in level of service delivered under degrading operational state of the network. Resilience must be quantitatively specified as a function of the operational condition of the network and the level of service to applications. We present a new metrics framework based on this basic understanding of network characterization and challenge modeling.

This chapter is organised as follows: First we present an overview of the proposed framework (Section 4.1) followed by formulation of metrics state space (Section 4.2). A resilience evaluation using this state space is presented in Section 4.3. A detailed step-by step methodology to apply the proposed resilience framework for a given level as well as across multiple levels is illustrated with a numerical examples in Section 4.4. Lastly, we present the application of this metrics framework in ResiliNets architecture of Chapter 4.

## 4.1 Overview

We propose a new approach to measure and quantify the resilience of the network against various challenges and attacks using functional metrics.[1]

[1] As part of this characterization, we will also determine a way in which conventional security mechanisms (such as authentication and integrity) map into these metrics.

For simplicity, assume that we are interested in evaluating the resilience of the overall network. In other words, we are considering the resilience at the application layer interface to the network. A more detailed analysis of resilience at multiple layers will be presented in Section 4.1.1.

Our approach is a three step process. First, we represent the operational condition of the network using metrics derived from the fundamental characteristics of the network. These are termed as *operational metrics* since they define the operational state of network parameters such as link utilization. Secondly, the level of service being provided by the network is quantified using representative functions based on application requirements such as goodput and delay; these are termed as *service parameters*. Probability distributions are used to represent metrics in cases where a single mean value does not adequately capture the dynamics and distribution of a particular metric.

Hence, the network can be viewed (at any layer) as consisting of two orthogonal dimensions as shown in Figure 4.1: one is the operational state of the network, which consists of its physical infrastructure and their protocols; the second dimension is the service being provided by the network and its requirements.

The full representation of the network state, thus requires a knowledge of both the operational metrics and service parameters at any given instant of time. Therefore, the third step involves aggregating operational metrics and their corresponding service parameters into discrete states that we call *network state* represented by the circles in Figure 4.1. Due to the time-varying nature of these metrics, especially in dynamic networks, a continuous representation gets increasingly complex with the number of such metrics. Hence, we choose a discrete representation that scales well with the number of metrics and service parameters.

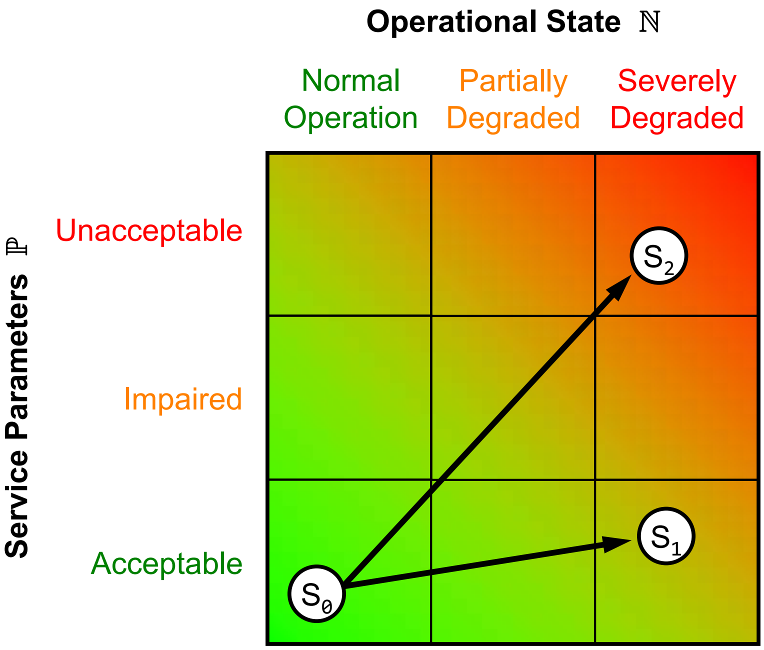

Figure 4.1: Resilience state space

In order to quantify the resilience of the system, 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. Hence, network resilience can be evaluated in terms of the various network states that can be supported with a given network infrastructure (e.g. technology and topology) and their transitions under the presence of challenges. Evaluating network resilience in this way effectively quantifies it as a measure of service degradation in the presence of challenges (perturbations) to the operational state of the network. Therefore, a comprehensive view of resilience requires the knowledge of quantitative performance of the network in all the states that it may visit under normal or adverse conditions.

In order to provide a second level of granularity, the operational and service space of the network may be divided into three regions each as shown in Figure 4.1. This purpose of this set of coarse grained regions in which the states reside is to simplify the resilience analysis. The network operational space is divided into *normal*, *partially degraded*, and *severely degraded* regions. Similarly, the service space is divided into *acceptable*, *impaired*, and *unacceptable* regions. While an arbitrary number of such regions is possible, one of the primary goals of this work is to achieve tractable yet useful solutions, and this set of nine ($3 \times 3$) regions provides the necessary abstraction while limiting the number total regions. Each region may contain multiple states if the service demands such a fine granularity. In the limiting case, each region represents just one state.

When an adverse event degrades the operational state of the network, the level of service being provided degrades as well resulting in *state transitions.* For example, Figure 4.1 shows the sample trajectory $S_0 \rightarrow S_1$ that an arbitrary application may take through the network if a malicious attack were to occur. The resilience is then evaluated as the transition of the network through this state space. In the simplest case, resilience is the area under the curve obtained by plotting operational metrics versus service parameters on a multivariate piecewise axis. For example, when comparing two services over a given network, the service with a smaller slope ($S_0 \rightarrow S_1$) is considered more resilient than one with a steeper slope as ($S_0 \rightarrow S_2$) shown in Figure 4.1.

### 4.1.1 Multilevel Approach

In this section, we discuss the different levels at which resilience is evaluated. In the literature review, we have seen resilience evaluated at various layers starting from the physical layer connectivity to the application layer service outage indexes. We propose that in order to be useful, *a resilience framework should facilitate the evaluation of resilience at any level of abstraction*, i.e. it should have the capability to evaluate resilience at any layer. In our proposed resilience framework, the service interface can be defined at any layer of architecture and we consider multilevel resilience an important aspect of our work. Consider the two scenarios shown in Figure 4.2.

In the first scenario 4.2(a), the service interface is defined at the network path routing layer[2] and the service being provided by the network layer (possibly to the higher transport layer) is to *find and establish paths* that meet specific quality of service requirements. The network layer itself and the layers below

[2] We divide the traditional network layer functionality into two sub-layers: path routing and topology. The reason for this will become apparent later in the resilience analysis.

contribute to the operational state. The ability of a network layer entity, such as a routing protocol, to find the end-to-end paths is based on the mechanisms it employs and the connectivity of the topology provided by the lower layers. Resilience is evaluated as the ability to find quality paths (service parameter) under the presence of challenges that affect the link connectivity and routing mechanisms (operational metrics).

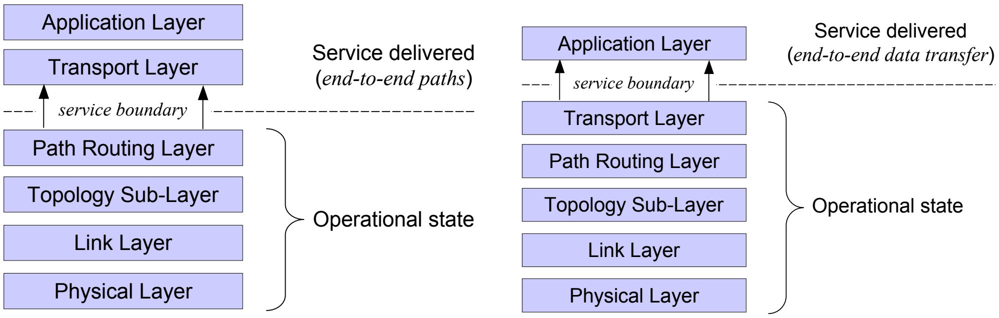

(a) Resilience at path routing layer

(b) Resilience at transport layer

Figure 4.2: Service interface at architectural layers

This represents one of the challenges categories discussed in Section 1.1: the service failure at a lower layer is a challenge to the higher layers.

In the second scenario, as shown in Figure 4.2(b), we are interested in evaluation of the resilience at the transport layer. Hence the service interface is defined at the transport layer. The service being provided is the *end-to-end transfer of data segments*. The operational state consists of the transport layer and all the layers below. For example, the quality of paths provided by the network layer and the state of transport protocols constitute the operational state. Hence resilience is evaluated as the ability to deliver data segments (service parameter) in the presence of adverse conditions that affect the stability and quality of the end-to-end paths and transport mechanisms (operational metrics).

Finally, the overall network resilience is evaluated from the user perspective. In this case all the layers contribute to the operational condition of the network and the user requirements define the service expected. While this scenario is particularly useful in evaluating the resilience in terms of end user requirements, it is equally important to be able to determine resilience at different levels of abstraction. In fact, such a capability is critical to development of multilevel resilience mechanisms.

In summary, the multilevel approach consists of defining the operational state of the network, the service being delivered, and the challenges experienced at any abstraction level. Then resilience is evaluated as *how well the service is provided under the presence of challenges that affect the operational state of the network*. Note that for the remainder of this document, resilience shall refer to the overall resilience from end-user perspective unless specifically stated otherwise.

## 4.2 Metrics State Space

In this sections, we present the formulation of operational and service state spaces using metrics as well as overall network states and the mathematical relationships between them. Finally, we discuss the impact of challenges on network states in terms of state transitions.

### 4.2.1 Operational State Space

Operational metrics capture the operational state of the network at any arbitrary service boundary. Let the system $\mathcal{S}$ (network at an arbitrary level) be represented by $\ell$ operational metrics, $N_{\mathcal{S}} = \{N_1, \ldots, N_\ell\}$. Each operational metric $N_i, 1 \leq i \leq \ell$, is in itself a set of $m$ values, representing all possible settings of the particular operational metric, $N_i = \{n_{i,1}, \ldots, n_{i,m}\}$. For example, at the physical layer of an ISP network, the number of link failures and link capacities could be two operational metrics.

*Special Case i*: If $N_i$ is numeric and ordered, then it is a set of $\ell$ values, $N_i = \{\underline{n}_i, \ldots, \overline{n}_i\}$, where $\underline{n}_i$ and $\overline{n}_i$ represent the lower and upper limit of the $i^{\text{th}}$ operational metric, respectively.

*Special Case ii*: If $N_i$ is numeric, ordered, and continuous then it is a set of all real values bounded by $[\underline{n}_i, \overline{n}_i]$.

The *operational state space* of $\mathcal{S}$ is $\mathcal{N}_{\mathcal{S}} = \times_i N_i$ where × represents the cross product operator. Therefore, the operational state space consists of all possible combinations of the operational metrics.

---

**Example 1:** Consider a system $\mathcal{S}$ with two operational metrics, $N_{\mathcal{S}} = \{N_1, N_2\}$. Assume that each operational metric has three possible settings (range). Therefore, $N_1 = \{n_{1,1}, n_{1,2}, n_{1,3}\}$ and $N_2 = \{n_{2,1}, n_{2,2}, n_{2,3}\}$. Then the operational state space is given as:

$$\mathcal{N}_{\mathcal{S}} = \{(n_{1,1}, n_{2,1}), (n_{1,1}, n_{2,2}), (n_{1,1}, n_{2,3}), (n_{1,2}, n_{2,1}), (n_{1,2}, n_{2,2}), \\ (n_{1,2}, n_{2,3}), (n_{1,3}, n_{2,1}), (n_{1,3}, n_{2,2}), (n_{1,3}, n_{2,3})\}$$

---

We now define an *operational state*, $\mathbb{N}$ as a subset of the complete state space $\mathcal{N}_{\mathcal{S}}$. 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, $\mathbb{N}_{\mathcal{S}} = \{\mathbb{N}_1, \ldots, \mathbb{N}_k\}$. $\mathbb{N}_{\mathcal{S}}$ is valid if $\mathbb{N}_{\mathcal{S}}$ is a partition of $\mathcal{N}_{\mathcal{S}}$. That is $\mathbb{N}_i \cap \mathbb{N}_j = \varnothing, \mathbb{N}_i, \mathbb{N}_j \in \mathbb{N}_{\mathcal{S}}$ and $i \neq j$ and $\cup_i \mathbb{N}_i = \mathbb{N}_{\mathcal{S}}$ where $\cup$ represents the union operator. Hence, in the generic case, an operational state is defined as a subset of $\mathcal{N}_{\mathcal{S}}$.

Special Case *i*: If $N_i$ is numeric and ordered $\forall i$ such that $N_i \in N_{\mathcal{S}}$, then the $k^{\text{th}}$ 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 $N_{\mathcal{S}}$. Therefore, $\mathbb{N}_k = \{N_{1k}, \ldots, N_{ik}, \ldots, N_{\ell k}\}$. A member $N_{ik}$ in the set $\mathbb{N}_k$ is in itself a set of valid values bounded by $[\underline{n}_{ik}, \overline{n}_{ik}]$, representing the lower and upper limit of the $i^{\text{th}}$ operational metric. We can now define $N_{ik} \equiv \{\underline{n}_{ik}, \ldots, \overline{n}_{ik}\}$. Thus $N_{ik}$ represents the set of $i^{\text{th}}$ operational metric values that correspond to the operational state $\mathbb{N}_k$. However, note that irrespective of the way in which the individual states are defined, an operational state $\mathbb{N}_k$ is always a partition of the state space $\mathcal{N}_{\mathcal{S}}$.

**Definition A.** *If the $i^{th}$ operational metric of a network at a given instant of time $t$ is $n_i(t)$, then the necessary condition for the network to be in operational state $\mathbb{N}_k$ is $\forall \{i : N_{ik} \in \mathbb{N}_k\}$, $n_i(t) \in N_{ik}$. In the special case of continuous $N_{ik}$, the necessary condition can be stated as $\forall \{i : N_{ik} \in \mathbb{N}_k\}$, $\underline{n}_{ik} \leq n_i(t) \leq \overline{n}_{ik}$.*

The network properties that are used in deriving operational metrics depend upon the type of network and the specific layer at which resilience is being categorized. Later in this chapter we will present an example of how operational metrics are obtained for a mobile wireless ad hoc network. Lastly, various security concerns can be captured using operational metrics. For example, malicious attacks that affect the physical state or behavior of network components can be modeled using metrics. The challenge lies in developing a compact set of metrics that are easy to understand and practical in current networks.

### 4.2.2 Service State Space

We now present the service state space which is orthogonal to the operational state space. The service parameters capture the requirement of the service that is being provided across the service interface. For example, the service from the transport layer to the application layer can be quantified using end-to-end delay in case of a voice application (in which latency affects the quality of service of the voice chat). Let the the system $\mathcal{S}$ (network at an arbitrary level) be represented by $\ell$ service parameters, $P_{\mathcal{S}} = \{P_1, \ldots, P_\ell\}$. Each service parameter $P_i, 1 \leq i \leq \ell$, is in itself a set of $m$ values (representing all possible values of the particular service parameter), $P_i = \{p_{i,1}, \ldots, p_{i,m}\}$. For example, service parameters metrics such as largest connected component and clustering coefficient may be used to characterize the topology service. *Special Case i*: If $P_i$ is numeric and ordered, then it is a set of $\ell$ values,

$P_i = \{\underline{p}_i, \ldots, \overline{p}_i\}$, where $\underline{p}_i$ and $\overline{p}_i$ represent the lower and upper limit of the $i^{\text{th}}$ service parameter, respectively.

Special Case ii: If $P_i$ is numeric, ordered, and continuous then it is a set of all real values bounded by $[\underline{p}_i, \overline{p}_i]$.

The *service state space* of $\mathcal{S}$ is $\mathcal{P}_{\mathcal{S}} = \times_i P_i$. Therefore, the service state space consists of all possible combinations of the service parameters.

**Example 2:** Consider a system $\mathcal{S}$ with two service parameters, $P_{\mathcal{S}} = \{P_1, P_2\}$. Assume that each operational metric has three possible settings (range). Therefore, $P_1 = \{p_{1,1}, p_{1,2}, p_{1,3}\}$ and $P_2 = \{p_{2,1}, p_{2,2}, p_{2,3}\}$. Then the operational state space is given as:

$$\begin{aligned}\mathcal{P}_{\mathcal{S}} = \{&(p_{1,1}, p_{2,1}), (p_{1,1}, p_{2,2}), (p_{1,1}, p_{2,3}), (p_{1,2}, p_{2,1}), (p_{1,2}, p_{2,2}),\\ &(p_{1,2}, p_{2,3}), (p_{1,3}, p_{2,1}), (p_{1,3}, p_{2,2}), (p_{1,3}, p_{2,3})\}\end{aligned}$$

We now define a *service state*, $\mathbb{P}$, as a subset of the complete state space $\mathcal{P}_{\mathcal{S}}$. 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, $\mathbb{P}_{\mathcal{S}} = \{\mathbb{P}_1, \ldots, \mathbb{P}_k\}$. $\mathbb{P}_{\mathcal{S}}$ is valid if $\mathbb{P}_{\mathcal{S}}$ is a partition of $\mathcal{P}_{\mathcal{S}}$. That is, $\mathbb{P}_i \cap \mathbb{P}_j = \varnothing, \mathbb{P}_i, \mathbb{P}_j \in \mathbb{P}_{\mathcal{S}}$ and $i \neq j$ and $\cup_i \mathbb{P}_i = \mathbb{P}_{\mathcal{S}}$. In a generic case, service states are specified as partitions of the complete service state space.

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

**Definition B.** *If the $i^{th}$ service parameter of a network at a given time instant $t$ is $p_i(t)$, then the necessary condition for the network to be in service state $\mathbb{P}_k$ is $\forall \{i : P_{ik} \in \mathbb{P}_k\}$, $p_i(t) \in P_{ik}$. In the special case of continuous $P_{ik}$, the necessary condition can be stated as $\forall \{i : P_{ik} \in \mathbb{P}_k\}$, $\underline{p}_{ik} \leq p_i(t) \leq \overline{p}_{ik}$.*

The service parameters invariably depend upon the service and application being supported. Hence the resilience of the network must be evaluated in terms of the particular service metric that is critical for the application. Given this framework, it is also possible for new and emerging application to define new metrics. Furthermore, the some of the data security issues such as confidentiality, integrity and authentication are modeled as service requirements. That is, an application requests the security services that it deems necessary.

### 4.2.3 Network State

As discussed earlier, in order to characterize a network at a service boundary we need to define both operational state and service state of the network. Hence, we define the overall *state* $S_{\mathcal{S}}$ of the system $\mathcal{S}$, (also termed as *network state*) as a tuple of operational state and service state: $(\mathbb{N}, \mathbb{P})$. Therefore the $k^{\text{th}}$ network state $S_k = (\mathbb{N}_k, \mathbb{P}_k)$

This overall state of the system $S_{\mathcal{S}}$ represents a mapping between the operational state space $\mathcal{N}_{\mathcal{S}}$ and service state space $\mathcal{P}_{\mathcal{S}}$. Furthermore, this mapping is an onto mapping, meaning that for every service state there is an operational state. There are no service states without a corresponding operational state. In other words, all service states are derived from the system.

In a deterministic system, the mapping of $\mathcal{N}_{\mathcal{S}}$ to $\mathcal{P}_{\mathcal{S}}$ 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. This is particularly true of the Monte-Carlo simulations and analysis presented in Chapter 5 and 6.

On order to eliminate the stochastic nature of the $\mathcal{N}_{\mathcal{S}}$ to $\mathcal{P}_{\mathcal{S}}$ mapping, in our analysis, we present the $\mathbb{N}_{\mathcal{S}}$ to $\mathbb{P}_{\mathcal{S}}$ 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 link failure probability (operational metric) to the largest component size (service parameter), we focus on the mapping of normal operating range of the link failure probability (operational state) to acceptable region of the largest connected component size (service state).

Lastly note that a single service state can occur (or be derived) from multiple operational states. Also note that every operational state must map to a service state and therefore we have a complete description of the behavior of the system in terms of service states. In both the operational metrics and service parameters are numeric and ordered, we can define the following:

**Proposition 1.** *For a given service boundary, if the $i^{th}$ metric of a network at an instant $t$ is $n_i(t)$ and the $j^{th}$ service parameter of a network at an instant $t$ is $p_j(t)$, then the network is said to be in a state $S_k$ if and only if* $\forall \{i : N_{ik} \in \mathbb{N}_k\}, n_i(t) \in N_{ik}$ and $\forall \{j : P_{jk} \in \mathbb{P}_k\}, p_j(t) \in P_{jk}$.

Proposition 1 suggests that every state of the network is defined by a unique set of operational metrics $N_k$ corresponding to the operational state $\mathbb{N}_k$ and a set of service parameters $P_k$ corresponding to the service state $\mathbb{P}_k$. However, each element of the set $N_k$ and $P_k$ is in itself a range of values thereby providing some region for the network to move around within a given state. We term these as *sub-states*. When the stimulus to the network drives the network beyond the range of the current state both in terms of operating region and network performance, the network is said to be in a different state altogether.

For example, consider the state space shown in the Figure 4.3. There are two states $S_1$ and $S_2$, in which each state is represented by three operational metrics $N_{1k}, N_{2k}, N_{3k}$ and two service parameters $P_{1k}, P_{2k}$. The internal points such as $S_1(t)$ within each state represent sub-states based on the instantaneous values of operational metrics and service parameters. Since each sub-state refers to a single point value of the each of the operational metrics and service parameters, there are a very large number of these instantaneous sub-states within a given state. However, as long as the operational metrics and service parameters are in the range of $\mathbb{N}_1$ and $\mathbb{P}_1$, the network remains in state $S_1$. When a perturbation of larger magnitude results in a change that exceeds the range of $P_{11}$ or $P_{21}$, the network moves to a different state $S_2$.

The challenge in this formulation is to limit the number of states to manageable value while achieving the required granularity is the services. However, the number of states are limited by the granularity of service required. For

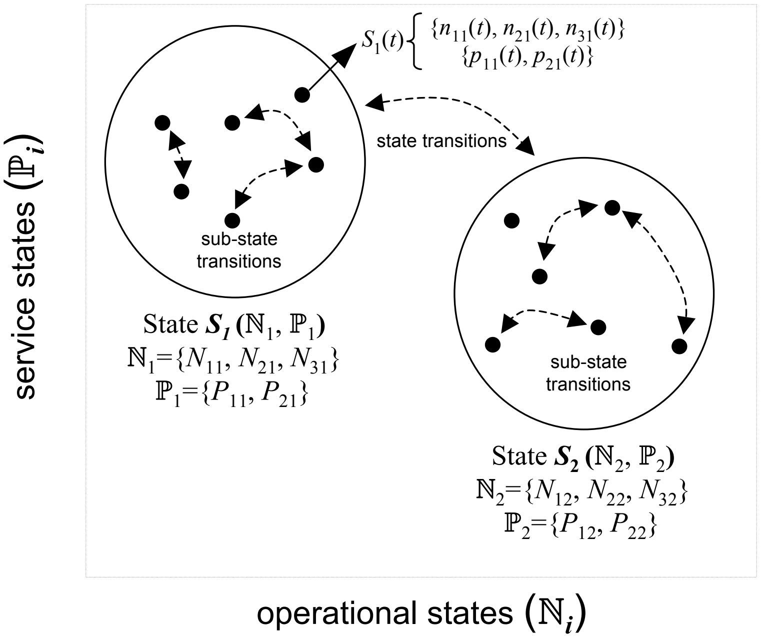
$S_1(t)$
$\{n_{11}(t), n_{21}(t), n_{31}(t)\}$
$\{p_{11}(t), p_{21}(t)\}$
state transitions
sub-state transitions
State $\boldsymbol{S_1}$ $(\mathbb{N}_1, \mathbb{P}_1)$
$\mathbb{N}_1=\{N_{11}, N_{21}, N_{31}\}$
$\mathbb{P}_1=\{P_{11}, P_{21}\}$
sub-state transitions
State $\boldsymbol{S_2}$ $(\mathbb{N}_2, \mathbb{P}_2)$
$\mathbb{N}_2=\{N_{12}, N_{22}, N_{32}\}$
$\mathbb{P}_2=\{P_{12}, P_{22}\}$
service states $(\mathbb{P}_i)$
operational states $(\mathbb{N}_i)$

Figure 4.3: Composition of network states and sub-states

example, in a voice application the number of states are dictated by the levels of service offered. Secondly, in order to provide a second coarser level of granularity, we divide the operational and the service space of the network in to *regions* based on network design and application supported, as discussed in Section 4.3.1. This further simplifies the analysis because of the limited number of regions. Further details on regions along with examples are presented in Section 4.3.

This approach of characterizing networks based on the fundamental network properties and expected performance can be used in different network research areas such as network resilience, adaptive routing, design and evaluation of security solutions.

### 4.2.4 Projected State Space

The operational state space $\mathcal{N}_{\mathcal{S}}$ and the service state space $\mathcal{P}_{\mathcal{S}}$ are both multivariate. As shown in Example 1, each element of the operational state space is a set with $\ell$ elements. Similarly, each element of the service state space is also a set with $\ell$ elements. In order to visualize this state space on a two dimensional state space, we project both the operational state space and service state space on to one dimension. This projection is achieved via an objective function that is applied in the both the state spaces. This is only possible if all operational metrics $N_i$ and service parameters $P_i$ are numeric and ordered.

Let $\mathcal{N}_{\mathcal{S}}^*$ be the projected operational state space of the original state space $\mathcal{N}_{\mathcal{S}}$. This is achieved via an objective function $f$ such that ${\mathcal{N}_{\mathcal{S}}}^* = f(\times_i N_i)$. This means that for each set in the $\mathcal{N}_{\mathcal{S}}$, we apply a objective function on its $\ell$ member elements. This objective function may be a linear combination with normalized weights or logical functions (e.g., AND, OR).

Similarly let ${\mathcal{P}_{\mathcal{S}}}^*$ be the projected service state space of the original service state space $\mathcal{P}_{\mathcal{S}}$. This is achieved via an objective function $f$ such that ${\mathcal{P}_{\mathcal{S}}}^* = f(\times_i P_i)$. Therefore, for each set in the $\mathcal{P}_{\mathcal{S}}$, we apply a objective function on its $\ell$ member elements. This objective function could be a linear combination with normalized weights or logical functions (e.g., AND, OR).

In the case of numeric, ordered, and continuous operational metrics and service parameters, the individual operational $\mathbb{N}_i$ and service $\mathbb{P}_i$ states with the range of their respective members. When these are projected over two dimension, we represent them as $\mathbb{N}_i^* = f(\mathbb{N}_i)$ and $\mathbb{P}_i^* = f(\mathbb{P}_i)$. When states are defined over the projected operational and service states, we can represent these states on a piece-wise linear axis.

---

**Example 3:** Lets consider the example that we have been building throughout this section. So far, we have defined the operational state space $N_{\mathcal{S}}$ and the service state space $P_{\mathcal{S}}$ for a given system $\mathcal{S}$. The operational state is defined as:
$N_1 = \{n_{1,1}, n_{1,2}, n_{1,3}\}$
$N_2 = \{n_{2,1}, n_{2,2}, n_{2,3}\}$
$\mathcal{N}_{\mathcal{S}} = \{(n_{1,1}, n_{2,1}), (n_{1,1}, n_{2,2}), (n_{1,1}, n_{2,3}), (n_{1,2}, n_{2,1}), (n_{1,2}, n_{2,2}),$
$(n_{1,2}, n_{2,3}), (n_{1,3}, n_{2,1}), (n_{1,3}, n_{2,2}), (n_{1,3}, n_{2,3})\}$
$P_1 = \{p_{1,1}, p_{1,2}, p_{1,3}\}$
$P_2 = \{p_{2,1}, p_{2,2}, p_{2,3}\}$
$\mathcal{P}_{\mathcal{S}} = \{(p_{1,1}, p_{2,1}), (p_{1,1}, p_{2,2}), (p_{1,1}, p_{2,3}), (p_{1,2}, p_{2,1}), (p_{1,2}, p_{2,2}),$
$(p_{1,2}, p_{2,3}), (p_{1,3}, p_{2,1}), (p_{1,3}, p_{2,2}), (p_{1,3}, p_{2,3})\}$

In order to project these multivariate spaces on to a two dimensional plot, we apply a objective function to obtain $N_{\mathcal{S}}^*$ and $N_{\mathcal{S}}^*$. Suppose the objective function is a linear combination of the operational metrics and service parameters:
$f = \alpha y_1 + (1 - \alpha) y_2$
$\mathcal{N}_{\mathcal{S}}{}^* = f(\mathcal{N}_{\mathcal{S}}) = \{(\alpha n_{1,1} + (1 - \alpha) n_{2,1}), (\alpha n_{1,1} + (1 - \alpha) n_{2,2}), \ldots\}$
$\mathcal{P}_{\mathcal{S}}{}^* = f(\mathcal{P}_{\mathcal{S}}) = \{(\beta n_{1,1} + (1 - \beta) n_{2,1}), (\beta n_{1,1} + (1 - \beta) n_{2,2}), \ldots\}$

Now both $\mathcal{N}_{\mathcal{S}}{}^*$ and $\mathcal{P}_{\mathcal{S}}{}^*$ are uni-dimensional and can be represented on two dimensional state space plots. Furthermore, the individual states can be defined as a subset of these new projected state space instead of the original state space.

---

### 4.2.5 State Transitions

There are two types of network transitions: sub-state transitions and state transitions. The stimuli that triggers these transitions include normal operational conditions such as traffic dynamics as well as various challenges and attacks (Section 1.1). The sub-state transitions reflect the instantaneous changes (of lesser magnitude) in the operational metrics with time due to dynamic nature of the network and traffic, especially in large networks. These transients sub-states are represented as internal points with in a state, as shown in Figure 4.3.

Sub-state transitions aside, as long as the operational metrics and service parameters do not violate the state boundaries, the network remains in its current state and only sub-state transitions are possible. However, events of large magnitude (often due to an external challenge or attack) result in state transitions. The range of operational metrics and service parameters for a given state is determined by the specific scenario. For example, a voice application may require two states based on the service metric – end-to-end delay: one state in which the delay is less than 200 msec and the other state for delays greater than 200 msec. On the other hand, data applications may require more number of states to differentiate service requirements of HTTP, P2P, and FTP traffic. The boundaries of each of the states and the number of states are both determined by the application being supported and the expected service.

Note that the sub-state transitions do not impact the resilience evaluations directly. Hence in the remainder of this dissertation we focus on state transitions alone.

## 4.3 Resilience Evaluation

In this section, we present a generic framework to *evaluate* the network resilience by tracing the movement of the network through various states.

### 4.3.1 Resilience State Space

Given that the network is characterized with network states and state transitions, we now consider the following two questions: *How do we evaluate resilience and what should be the resolution between states?* We propose a new approach to limit the number of states as well as facilitate an easy evaluation of resilience. We propose that the operational and service space of the network be divided into *regions*. The operational space of the network is divided into $r_o$ regions based on the physical infrastructure and $r_s$ regions based on the application and end user requirements. We present this approach with a simple case where $r_o = r_s = 3$.

We divide the operating region of the network in to three regions: normal operating, partially degraded, and severely degraded. Similarly, the performance of the network is also classified in to three regions: acceptable, impaired, and unacceptable. The stimuli, in this case, are various adverse conditions discussed in Section 1. Note that there may be more than one state in each region. Let us consider the transition of the network between different states as shown in Figure 4.4.

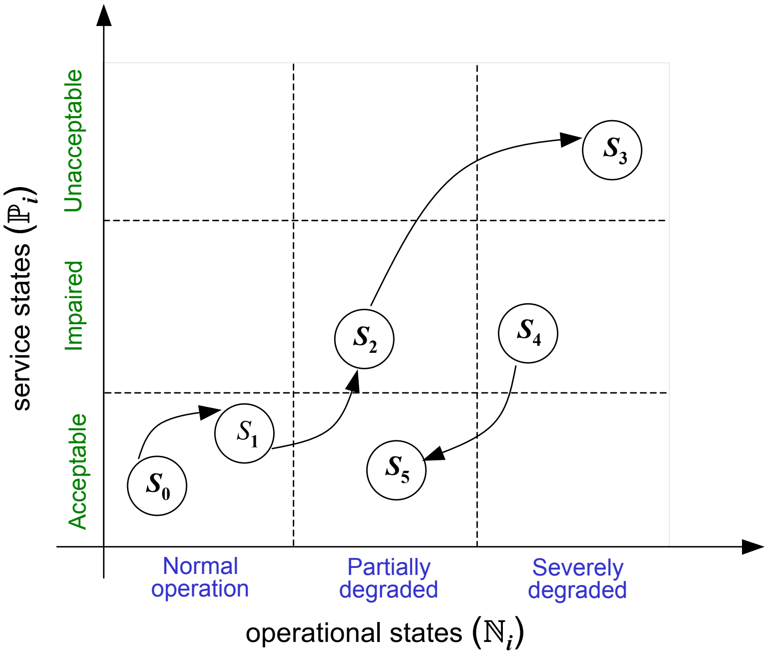

Figure 4.4: Network states in the given operating region

Let the current state of the network be $S_0$ as shown in Figure 4.4. Following the occurrence of an adverse event, the network stays in its current state if the change in the $i^{\text{th}}$ metric, $n_{i0}$, does not exceed the allowed range $[\underline{n}_{i0}, \overline{n}_{i0}]$ and the service parameters remain with in limits $[\underline{p}_{j0}, \overline{p}_{j0}]$. If an adverse event does result in one or more metrics exceeding their range in the current state, the network proceeds to a different state. Assume that following an adverse event, the network moves to state $S_1$ in which service parameters are in the limits $[\underline{p}_{j1}, \overline{p}_{j1}]$. The network may be engineered so that for a given application, both $S_0$ and $S_1$ lie in the normal operating region, in which the service is acceptable. On the other hand, adverse events of greater impact may drive the network to a state $S_2$ in the partially degraded region with impaired service, or to a state $S_3$ in severely degraded region with unacceptable service parameters. The range of network operational metrics for which the network will remain in each state is clearly quantified along with the expected service in that state. The unique advantage of this method is that it results in manageable number of states.

In summary, there are two potential issues with this approach: (1) the number of metrics in each dimension may be very large and (2) there may be state explosion due to the number of quantifiable network states can be avoided in this context. First, the number of metrics can be limited to those that affect the service under consideration. Secondly, the number of states are limited by the granularity of the service differentiation required as illustrated above. Finally, all transients in network operation (originating from network dynamics) that result in a predetermined service level are aggregated into one state. This approach significantly reduces the number of states.

### 4.3.2 Effects of Security Mechanisms on State Space

Security plays a very important role in the overall resilience of the network. We propose to map various security aspects of end-to-end applications on the state space of the network. The threats that effect the physical state of the network, such as jamming, hijacking, or corrupting a node can be captured in the operational metrics. On the other hand, data confidentiality and integrity are treated as service requirements that are quantified using service parameters. Separating the operational and service aspects of security has the added benefit of reducing the cost since service requirements can be met on a per application basis. For example, a bank transaction may require data confidentiality that is provided at increased cost by encryption, whereas a news feed with open access can be serviced at lower costs.

**Detecting attacks using the state space:** Consider an instant of time when the expected operational metrics belong to a state $S_k$, $\mathbb{N}(t) \in S_k$. However, the observed performance expressed in service parameters does not belong to state $k$, $\mathbb{P}(t) \notin S_k$. Assuming the correctness of measurements, this is an indication of a hidden malicious attack. Furthermore, the degradation in the network due to the attack can be evaluated from the state in which the observed service parameters are placed.

For example, using our formulation, the RoQ attacks can be modeled as changes in the traffic distribution, resulting in reduced performance. When an RoQ attack occurs, the observable traffic load places the network in state $S_1$ in the normal region, however, the observed service parameters place the network in a different state $S_2$ in the impaired or unacceptable region. Mathematically, $\mathbb{N}(t) \in S_1$ but $\mathbb{P}(t) \in S_2$ and since $S_1 \neq S_2$, an attack is detected. In summary, since the expected service under various network conditions is quantified, an anomaly in the service is immediately visible.

### 4.3.3 Resilience Measure

Given the formulation discussed in the above section, we now *establish a measure of resilience*. In this framework, resilience is the ability to stay in the acceptable service region or degrade gracefully under the presence of attacks. Resilient networks remain either in the acceptable service region or move slowly in to the impaired region with increasing degradation in the network. Furthermore, two applications with different service requirements may follow two completely different trajectories through the state space. For example in the Figure 4.5, a resilient application $X$ may follow the state trajectory $S_0 \rightarrow S_1 \rightarrow S_2 \rightarrow S_3$, whereas another resource intensive application $Y$ may follow the state trajectory $S'_0 \rightarrow S'_1 \rightarrow S'_2 \rightarrow S'_3$.

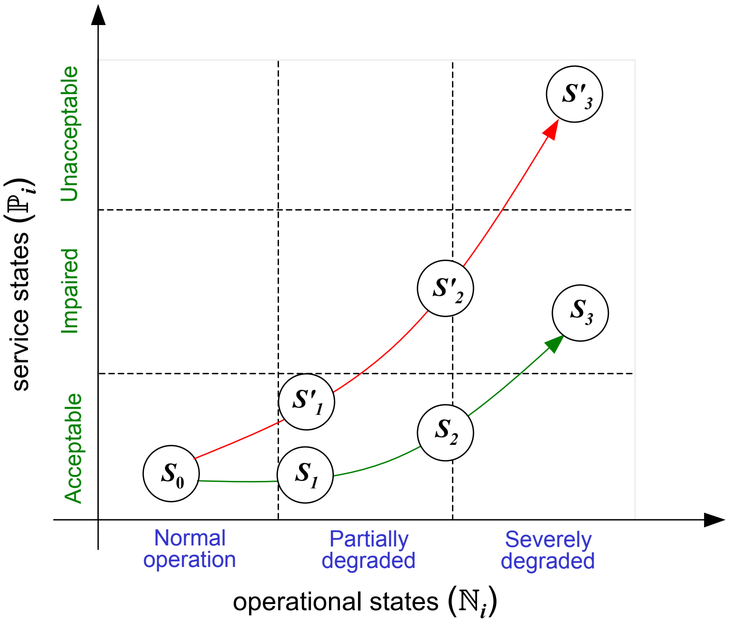

Figure 4.5: State trajectory of two applications

As is evident from the figure, application $Y$ deteriorates faster than application $X$ when the network degrades. Observe that this method provides the knowledge of system behavior under attacks before hand. This is not achieved by an exhaustive analysis of all the possible attacks, but rather by analyzing the manifestation of such attacks in network characteristics (operational metrics) and service parameters.

### 4.3.4 Steady State versus Transient Analysis

The proposed methodology supports both steady state analysis and transient analysis. In the steady-state analysis we evaluate the long term view of the network resilience by conducting either theoretical calculations or network simulations to understand the impact of perturbations in the operational state (due to challenges and attacks) on the service parameters. This leads to best, worst, and average case resilience measures. In the transient analysis, we observe the instantaneous state of the network and evaluate the state transitions in real time as the network challenges are countered by detection, defense, remediation, and recovery mechanisms. In this case, resilience is characterized by state transitions occuring in real time.

## 4.4 Methodology

In this section, we present the methodology to evaluate resilience of networks using the metrics framework presented in the previous sections. First, we show a step-by-step method to build the state-space using a numerical example. Secondly, we show how the framework applies to multilevel resilience.

### 4.4.1 Building the State-Space

In this section, we evaluate a specific scenario to demonstrate the steps involved in building the two-dimensional state-space using the proposed framework. Consider a generic wireless metropolitan area network consisting of a number of mobile wireless nodes uniformly distributed. Some modes are information sources (servers) and others are information sinks (clients). For instance, servers may have access to the Internet and other mobile nodes access the Web through the server nodes. The wireless client nodes also communicate amongst themselves (analogous to ad-hoc mode in 802.11). Our objective is to build the state-space at the application layer, meaning that we are interested in the overall resilience of the network when subject to a specific set of challenges.

**Step 1: Characterize the network operational condition:**

The first step is to characterize the network using operational metrics. The set of metrics used depend on the specific challenge being considered. In this example, we evaluate the impact of traffic and degree of network connectivity on the delivered service. For simplicity we consider the following two operational metrics for this scenario, assuming that the other metrics remain at a constant value:

$n_1 =$ traffic load normalized to $\rho_0$

$n_2 =$ average node degree

The normalized traffic is used as a metric to characterize the traffic on the network, given that $0.75\rho_0$ is the normal load for which the network is engineered. For simplicity, the burstiness of the traffic is ignored in this example. Node degree is used as a metric to quantify the connectivity of the network. Hence the *operational state* is represented as: $\mathbb{N} = \{N_1, N_2\}$ where $N_1$, $N_2$ are a set of values of $n_1$ and $n_2$ respectively.

Note that in this example we are not considering the numerous other operational metrics from all the layers between physical layer and application layer. The resilience analysis across multiple layers is presented in Section 4.4.2

**Step 2: Characterize the service:**

Service is characterized based on the application and the performance parameters of interest. Consider an interactive application such as Web browsing where the objective of the user is to retrieve a web page. Again, we consider two specific service metrics to specify user requirements.

$p_1 =$ response time expressed as: $2 \times$ end-to-end delay

$p_2 =$ service availability expressed as connection success rate (CSR)

The response time determines the time user has to wait for a response in cases where the service is available and the number of failed connection attempts characterizes the service availability. Hence the *service state* is represented as: $\mathbb{P} = \{P_1, P_2\}$ where $P_1$, $P_2$ are a set of values of $p_1$ and $p_2$ respectively.

**Assumptions:** In order to simplify the example scenario, we make the following non-realistic assumption. The objective of this example is to demonstrate the use of framework for hypothetical network scenario with two operational metrics and two service parameters:

- the dominant factor in the end-to-end delay is the queuing time assuming the transmission, propagation, and processing time to be negligible
- other operational conditions of the network remain constant and do not impact the variations in the service parameters

**Step 3: Observable network states:**

Given that the operational condition of the network and the expected service of the network is characterizes, we can now define the state as:

$$S = (\mathbb{N}, \mathbb{P}) \text{ where}$$

$$\mathbb{N} = \{N_1, N_2\} \text{ and}$$

$$\mathbb{P} = \{P_1, P_2\}.$$

We make a simplifying assumption that the user is interested in five network states $(S_1, S_2, S_3, S_4, S_5)$. This implies that we evaluate the resilience of the network in terms of its state transitions between these five states. The corresponding operational metrics and service parameters are given in Table 4.1.

Table 4.1: Example of network states

| State | Operational Metrics | | | | Service Parameters | | | |
|---|---|---|---|---|---|---|---|---|
| $S_k$ | traffic load $N_{1k}$ | | node degree $N_{2k}$ | | response time $P_{1k}$ | | success % $P_{2k}$ | |
| | $\underline{n}_{1k}$ | $\overline{n}_{1k}$ | $\underline{n}_{2k}$ | $\overline{n}_{2k}$ | $\underline{p}_{1k}$ | $\overline{p}_{1k}$ | $\underline{p}_{2k}$ | $\overline{p}_{2k}$ |
| $k = 1$ | 0 | $0.5\rho_0$ | 5 | 6 | 0 | 0.5 | 0.95 | 1 |
| $k = 2$ | $0.5\rho_0$ | $0.75\rho_0$ | 6 | 8 | 0 | 1 | 0.90 | 0.95 |
| $k = 3$ | 0 | $0.6\rho_0$ | 3 | 5 | 1 | 5 | 0.75 | 0.85 |
| $k = 4$ | $0.75\rho_0$ | $0.85\rho_0$ | 5 | 7 | 5 | 10 | 0.75 | 0.80 |
| $k = 5$ | 0 | $0.6\rho_0$ | 0 | 3 | 0 | 1 | 0.5 | 0.6 |

**Step 4: Formulate the state-space:**

Given the application and the observed states, the next step involves developing the state-space diagram of the network. In order to further simplify the resilience analysis, we divide the operational space and service space in to smaller *regions*. In this particular example, we divide the operational space and service space in to three regions each based on the design of the network and user expectations of the service provided. For the interactive application, currently under consideration, Table 4.2 shows the classification of regions in service space.

The boundaries of the service regions are based on the generic utility curve for interactive applications [51]. The user does not perceive a change if the response time is less than 300 ms and delays up to 1 s are barely notice-

Table 4.2: Regions of service space

| Region | Service Parameters | |
|---|---|---|
| | response time $p_1$ | % successful attempts $p_1$ |
| Acceptable | $p_1 \leq 1$ s | $p_2 \geq 0.90$ |
| Impaired | $1 < p_1 \leq 10$ s | $p_2 \geq 0.75$ |
| Unacceptable | $p_1 > 10$ s | $p_2 < 0.75$ |

Table 4.3: Regions of operational space

| Region | Operational metrics | |
|---|---|---|
| | traffic load $n_1$ | avg. node degree $n_1$ |
| Normal | $n_1 \leq 0.75\rho_0$ | $n_2 \geq 5$ |
| Partially degraded | $0.75\rho_0 < n_1 \leq 0.9\rho_0$ | $n_2 \geq 3$ |
| Severely degraded | $n_1 > 0.9\rho_0$ | $n_2 < 3$ |

able. Hence, for the application under consideration, the acceptable region of performance extends for all sub-second response times. The user begins to notice a delay in system response starting from 1 s up to 10 s. This is classified as impaired performance. For delays beyond 10 s, the user gives up and the service is unacceptable. The boundary conditions for second service parameters, i.e. percentage of successful connection attempts, are derived arbitrarily for this example.

Similarly, the operational space is divided in to 3 regions as shown in Table 4.3. The boundaries for each of the region is determined by the original design of the network. In this particular example, the network is engineered for optimal performance under a traffic load of $\rho_0$. Furthermore, based on the steady state mobility patterns the network is *expected* to have an minimum node degree of 3.

**Step 5: Assign observed/derived states to respective regions:**

In the final step, we assign the available states (observed or derived) to the regions of the state-space. For a state to be present in a given operational or service region, it must satisfy all the constraints of the regions as given in Tables 4.2 and 4.3. Furthermore, if a given network state satisfies the constraints of more than one region, it is placed in the best region. Matching the states in Table 4.1 with the region constraints of Tables 4.2 and 4.3, we obtain the state-space diagram shown in Figure 4.6. Thus, we obtain the state-space formulation for a given resilience scenarios. Note that in this case, the objective functions $\mathcal{N}_{\mathcal{S}}^{*}$ are the logical $\mathsf{AND}$ of the individual metrics.

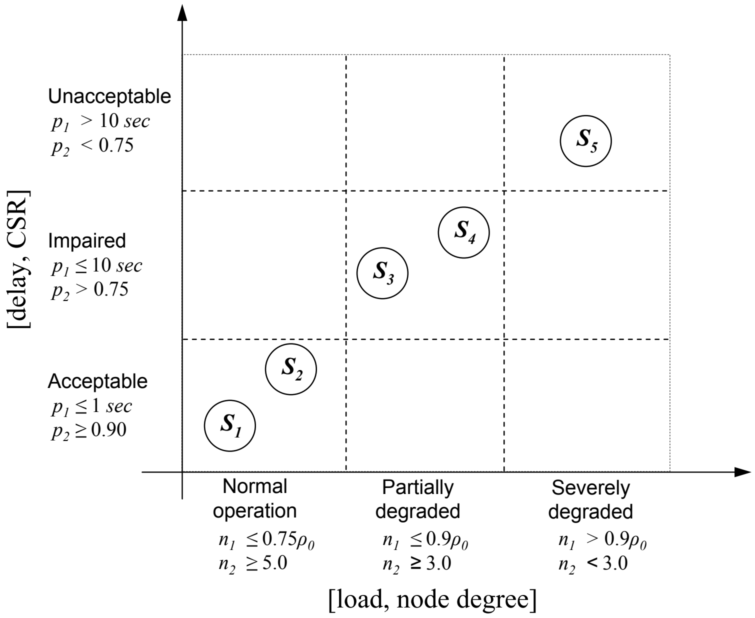

Figure 4.6: Final state-space diagram for example 3.1

In this example, we demonstrated the methodology to derive the state-space at a given level. The application of this methodology over multiple levels follows.

### 4.4.2 Multilevel Resilience Evaluation

In this section, we discuss the multilevel aspect of the metrics framework. Furthermore, we use a mobile adhoc network (MANET) example to demonstrate how resilience propagates across layer boundaries.

The multilevel principle [P8] of ResiliNets architecture suggests that resilience be addressed at all levels, in the sense that each layer does the best it can, given practical constraints. These constraints are often based on the cost of resilience. *Therefore, resilience must be analyzed at each layer individually as well as for the network as a whole.* For this purpose, the metrics framework supports multilevel resilience evaluation. Formally, resilience $\mathbb{R}_{ij}$ is defined at the boundary $B_{ij}$ between any two adjacent layers $L_i$, $L_j$. Based on the formulation of Section 4.2, let there be a set of $k$ *operational metrics* $\mathbb{N} = \{N_1, N_2, \ldots, N_k\}$ that characterize the state of the network below the boundary $B_{ij}$,. Similarly, let there be a set of $l$ *service parameters* $\mathbb{P} = \{P_1, P_2, \ldots, P_l\}$ that characterize the service from layer $i$ to layer $j$. Resilience $\mathbb{R}_{ij}$ at the boundary $B_{ij}$ is then evaluated as the transition of the network through this state space. The goal is to derive the $\mathbb{R}_{ij}$ as a function of $\mathbb{N}$ and $\mathbb{P}$. In the simplest case $\mathbb{R}_{ij}$ is the area under the curve obtained by plotting $\mathbb{P}$ vs. $\mathbb{N}$ on a multivariate piecewise axis. We will revisit the calculation of $\mathbb{R}_{ij}$ based on a given set of state transitions in Chapter 5.

In the multilevel analysis, 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. This process is shown in Figure 4.7.

We model different options at all layers under consideration, e.g., topology, routing and transport. In this process, we will evaluate how a specific mechanism at a given layer (say path routing) performs by using a fixed standard setting at all other levels. In other words, we isolate the resilience of each

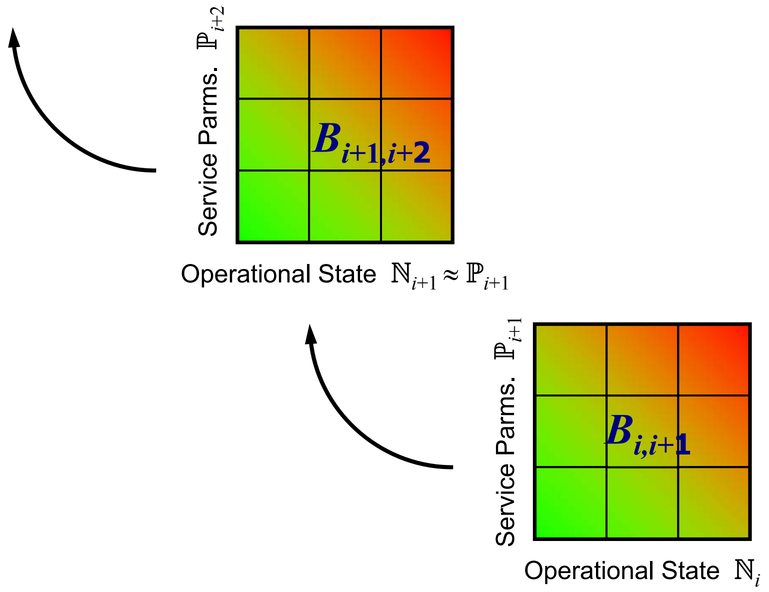
Service Parms. $\mathbb{P}_{i+2}$
$B_{i+1,i+2}$
Operational State $\mathbb{N}_{i+1} \approx \mathbb{P}_{i+1}$
Service Parms. $\mathbb{P}_{i+1}$
$B_{i,i+1}$
Operational State $\mathbb{N}_i$

Figure 4.7: Resilience across multiple levels

layer and evaluate the effectiveness of the resilience mechanisms and principles discussed in Chapter 3.

### MANET Example

In this section, we evaluate a mobile ad hoc network with the objective of characterizing the network at each layer from bottom up. In this example, we will consider which metrics can be used at each layer and how they propagate upwards. Figure 4.8 shows the operational metrics and service parameters at each layer boundary. Note that these levels do not directly correspond to OSI layers.

At the lowest boundary $B_{1,2}$ between the physical and the link+MAC layer, the operational metrics are the radio characteristics, the transmit power, and

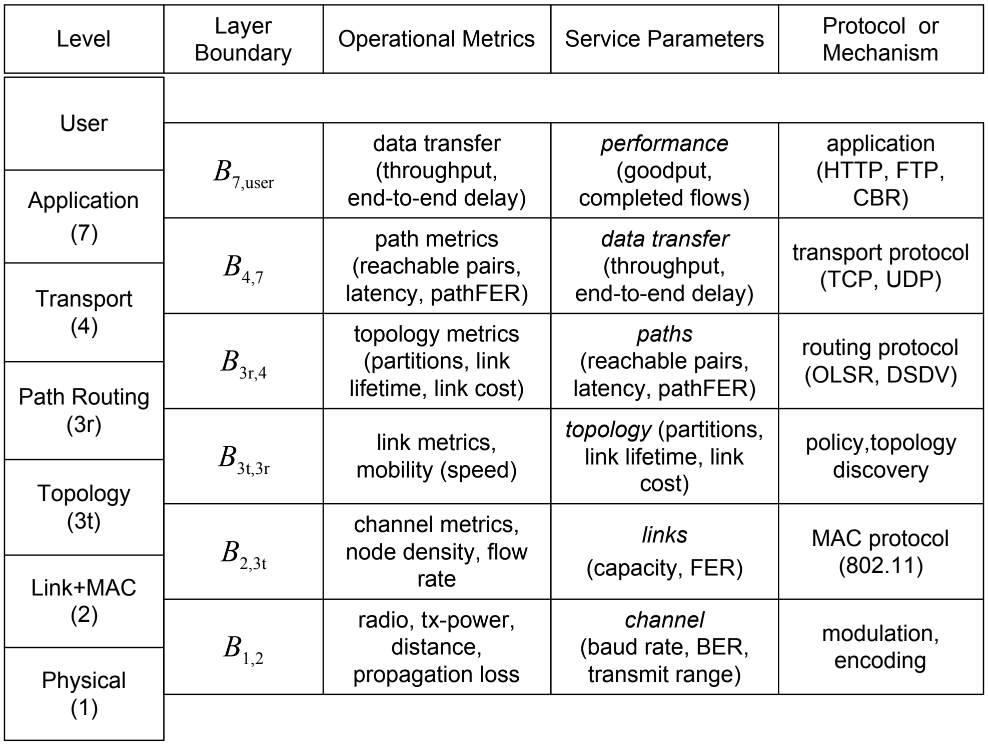

| Level | Layer Boundary | Operational Metrics | Service Parameters | Protocol or Mechanism |
|---|---|---|---|---|
| User | | | | |
| | $B_{7,\mathrm{user}}$ | data transfer (throughput, end-to-end delay) | *performance* (goodput, completed flows) | application (HTTP, FTP, CBR) |
| Application (7) | | | | |
| | $B_{4,7}$ | path metrics (reachable pairs, latency, pathFER) | *data transfer* (throughput, end-to-end delay) | transport protocol (TCP, UDP) |
| Transport (4) | | | | |
| | $B_{3r,4}$ | topology metrics (partitions, link lifetime, link cost) | *paths* (reachable pairs, latency, pathFER) | routing protocol (OLSR, DSDV) |
| Path Routing (3r) | | | | |
| | $B_{3t,3r}$ | link metrics, mobility (speed) | *topology* (partitions, link lifetime, link cost) | policy,topology discovery |
| Topology (3t) | | | | |
| | $B_{2,3t}$ | channel metrics, node density, flow rate | *links*<br>(capacity, FER) | MAC protocol (802.11) |
| Link+MAC (2) | | | | |
| | $B_{1,2}$ | radio, tx-power, distance, propagation loss | *channel* (baud rate, BER, transmit range) | modulation, encoding |
| Physical (1) | | | | |

Figure 4.8: Multilevel resilience evaluation of a MANET

channel conditions. The service provided across this boundary (from physical layer to the link+MAC) is a communication channel which is measured by the baud rate, BER, and the transmit range. Therefore, $\mathbb{N}_1$ = {tx-power, propagation loss, node distance} and $\mathbb{P}_2$={baud rate, BER, transmit range}. We evaluate the resilience the physical layer by measuring these service parameters (representing the quality of the channel) under the presence of perturbations to the operational conditions. For example, if a challenge such as a weather storm were to increase the propagation losses of a millimeter-wave network, the ability of the radio to overcome this increased losses determines its resilience. Therefore $\mathbb{R}_{1,2} = f(\mathbb{N}_1, \mathbb{P}_2)$. In an ideal (but impractical) case, the radio would be able to survive this challenge without affecting the service being provided.

At the boundary $B_{2,3\mathrm{t}}$ we can evaluate the resilience of the link+MAC protocol. In this case the service parameters from lower boundary $B_{1,2}$ become the operational metrics. In addition, other factors that are included in the operational state of the network at this boundary are the node density and flow rates because both of those parameters affect the performance of 802.11 MAC algorithm. The service provided across this boundary is *links*. Hence, $\mathbb{N}_2$ = {$\mathbb{P}_2$, node density, flow rate} and $\mathbb{P}_{3\mathrm{t}}$={link capacity, FER}. The resilience at this level is evaluated as the ability of the link+MAC protocol to provide stable, high capacity, and low error links in the presence of perturbations to the its operational metrics, i.e. $\mathbb{R}_{2,3\mathrm{t}} = f(\mathbb{N}_2, \mathbb{P}_{3t})$.

At the next boundary $B_{3\mathrm{t},3\mathrm{r}}$, we evaluate the resilience of the topology sub layer. Note that we divide the traditional network layer into a topology (3t) and path routing (3r)sub-layers so as to isolate the resilience of a given topology from the ability of the routing protocol to find paths over that topology. At this level, the service parameters from the level below become the operational metrics for this level. Furthermore, the mobility is another metric that gets added to the operational state at this level. Assuming a fixed mobility model, we only consider the node speed in this example. Here we define $\mathbb{N}_{\mathrm{3t}} = \{\mathbb{P}_{\mathrm{3t}}, \text{node speed}\}$. The service provided by this sub-layer is the *topology*, therefore, the service metrics relate to quantifying certain aspect of the network topology. In this example we choose three such parameters: partitions, average link duration, and link costs. Hence, $\mathbb{P}_{\mathrm{3rr}}$={partitions, avg. link durations, link costs}. Resilience is then the ability of this layer to provide connected topologies in the presence of perturbations to the links.

At the boundary $B_{\mathrm{3r},4}$, we measure the resilience of the routing protocol. In this case the operational conditions consists of the service parameters from the layer below and the service being provided is *paths*. The operational metrics are simply the service parameters from the layer below. $\mathbb{N}_{\mathrm{3r}}$=$\mathbb{P}_{\mathrm{3t}}$={partitions, avg. link durations, link costs}. The service parameters reflect the ability of the routing to find optimal paths given a topology, $\mathbb{P}_4$={% reachable pairs, path latency, and path FER}. The resilience at this level is the ability of the network to find optimal paths in the presence of perturbations to the underlying topology, $\mathbb{R}_{\mathrm{3r},4} = f(\mathbb{N}_{\mathrm{3r}}, \mathbb{P}_4)$.

Going a level above, at the $B_{4,7}$ we evaluate the resilience of the transport protocol under the presence of challenges. Note that we are ignoring the session (5) and presentation (6) layers. Mapping the service parameters from the level below, we arrive at the operational metrics at this boundary, $\mathbb{N}_4$=$\mathbb{P}_{3\mathrm{r}}$={% reachable pairs, path latency, and path FER}. The service provided by the transport is *end-to-end data transfer*. In order to measure this service, we define two metrics: $\mathbb{P}_4$={throughput, end-to-end delay}.

Lastly, at the boundary $B_{7,8}$ we evaluate the resilience of the application (e.g. HTTP, FTP). For the sake of generality, we are representing the end user as layer 8. The operational metrics as derived from the service parameters from layer below are: $\mathbb{N}_7$=$\mathbb{P}_4$={throughput, end-to-end delay}. The service provided by the application is simply measured as its performance to the user. In this example we define two service parameters: $\mathbb{P}_7$={goodput, percentage of completed flows or transactions}. Hence resilience at this boundary is evaluated as the ability to provide acceptable performance in the presence of perturbations to the service provided by the transport layer.

Thus, the proposed framework can evaluate the resilience of a network at each level, thereby providing a means to evaluate the multilevel resilience principle of the ResiliNets architecture. Furthermore, it allows us to determine the weakest link in the protocol stack from a resilience perspective. Finally, it is also possible to evaluate the resilience between any two arbitrary boundaries $B_{ij}$ where $j > i$ and $j \neq i + 1$. This is especially useful in simulation based studies in which the simulation model may abstract certain levels. We have conducted simulation studies to evaluate the resilience of the MANET at several boundaries. These results are presented later Chapter 5.

## 4.5 Metrics in ResiliNets Architecture

In this section, we present how the metrics framework fits in the overall resilience architecture. Furthermore, we also discuss how various resilience strategies and mechanisms influence the evaluation of resilience using the proposed framework. Figure 4.9 shows an integrated view of the network behavior at *any* abstraction level (layer).

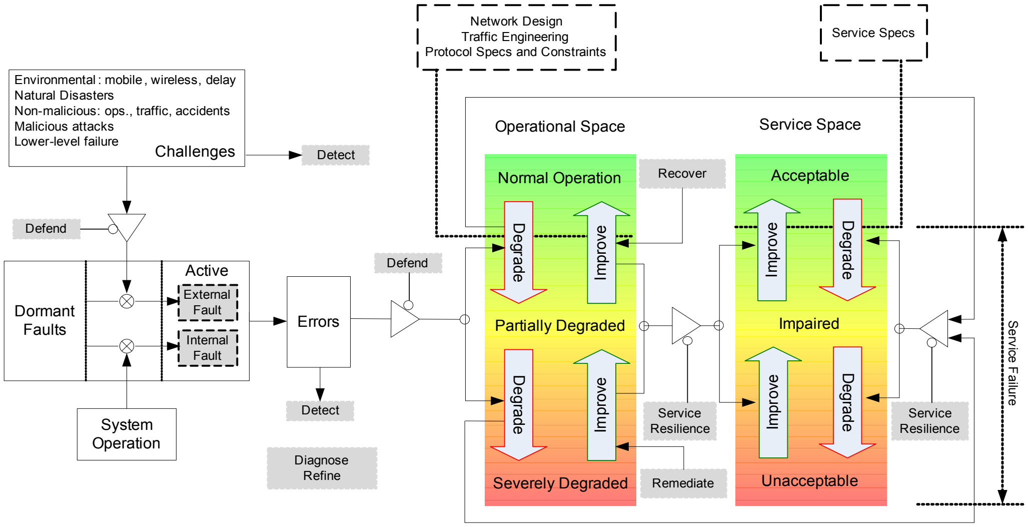

Figure 4.9: Instrumenting resilience using proposed framework

This block diagram shows how the errors develop in a network and how they affect the service being delivered with respect to the $D^2R^2 + DR$ strategy.

The first section of the block diagram is reproduced for clarity in Figure 4.10 and shows the evolution of the errors as explained below.

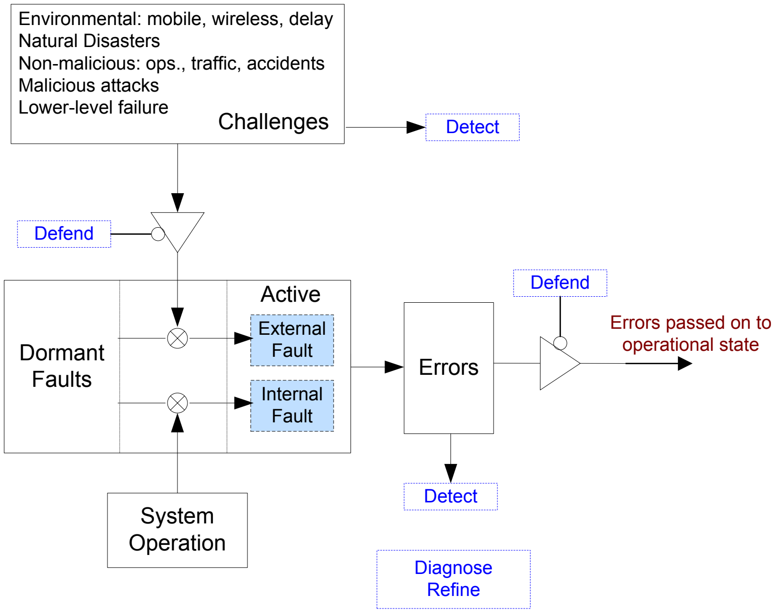

Figure 4.10: Evolution of errors in communication systems

As discussed in Chapter 3, dormant faults are inherent to all systems. Careful network design can eliminate faults to some extent. However, it is not possible to build absolutely fault free systems. This (dormant fault) vulnerability in the system is exploited by adverse events such as malicious attacks and environmental challenges in the absence of defense mechanisms to cause (trigger) external faults. On the other hand, internal faults are caused by those dormant faults that are triggered merely by system operation (for example execution of buggy code). Both internal and external faults are called active faults [15]. Active faults cause errors that if not detected and corrected or compensated directly affect the operational state of the network. The process by which such errors affect the system is reproduced in Figure 4.11.

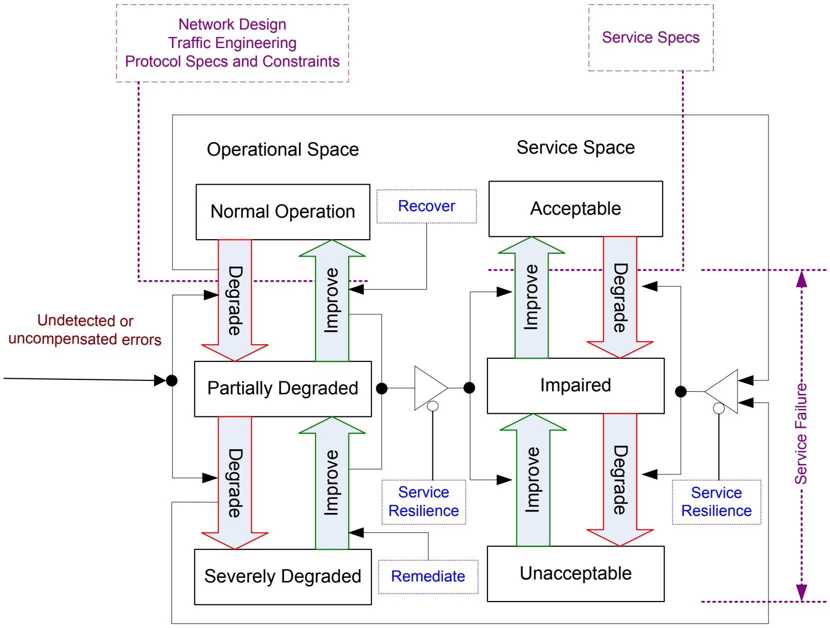

Figure 4.11: Effect of errors on the communication systems

The extent to which the network degrades depends on the severity of the error and the design of the system. Note that this degradation of the operational state is captured by operational metrics in the proposed framework. The effect of the network degradation on the users is determined by the resilience of the service and is captured by service parameters in our framework. Depending upon the extent of resiliency, the service parameters may either remain acceptable or degrade to some lower value. If the service does degrade, the remediation and recovery mechanisms detect this degradation in network state and repair the network leading to improvement in the services.

**Example 3.2:** An example of these relationships for the physical layer is shown in Figure 4.12. It is seen that for every service interface, it is possible

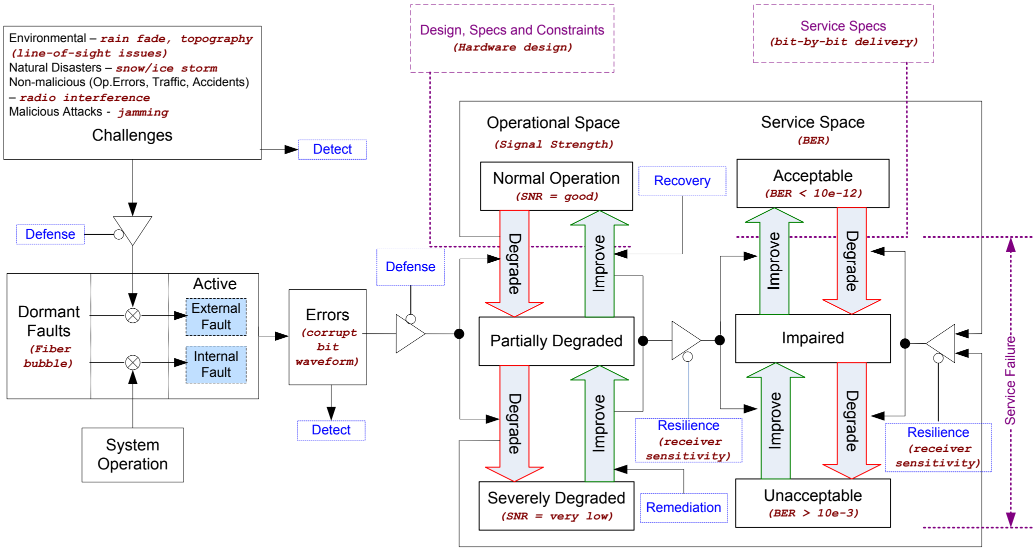

Figure 4.12: Example of resilience at physical layer

to identify the elements leading to errors. Also, the effect of these errors on the operational space and subsequently on the service can also be determined if the rate of errors and service resilience is known. We could then use the proposed framework to analyze the state-space of the network and derive resilience functions. Referring back to the Figure 4.4, the state transition from $S_1$ to $S_2$ shows a degradation in the operational condition of the network that leads to impaired service. On the other hand, the state transition from $S_4$ to $S_5$ shows the recovery by the network from impaired service to acceptable.

### 4.5.1 Resilience Strategy Revisited

Figure 4.9 shows the relationship of the $D^2R^2$+DR strategy to the 2- dimensional state space. The left part shows the fault → error → failure chain with the defense resisting the activation of faults and propagation of errors to service failures. Challenges and errors can be detected to initiate remediation. The right side of the Figure shows the operational dimension $\mathbb{N}$ and service dimension $\mathbb{P}$ as ranges that are degraded when errors propagate to operational state and when the effects of the operational state impact the service state ($S_0 \rightarrow S_c$ trajectory in Figure 4.13). Remediation mechanisms help drive the operational state towards improvement, and service resilience

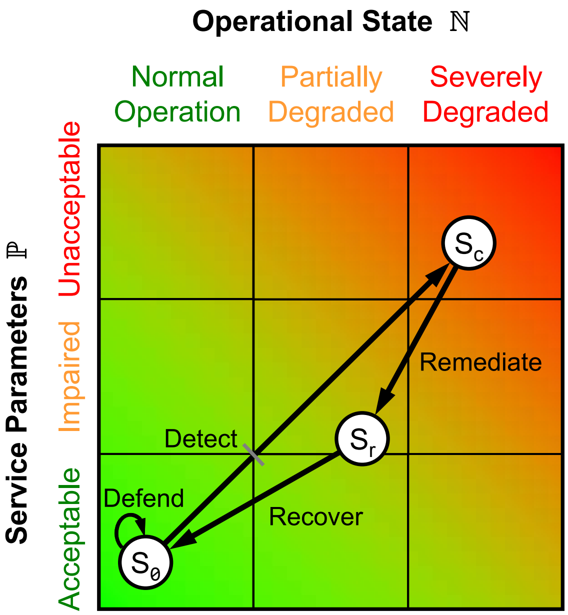

Figure 4.13: Resilience strategy evaluation using metrics framework

resists the effects of degraded operational state from impairing the service ($S_{\mathrm{c}} \rightarrow S_{\mathrm{r}}$ trajectory). Recovery moves the operational state back to normal operation at the end of the inner strategy loop. Diagnosis and refinement are the outer loop, used to improve this entire process.

In summary, an analysis of the process by which errors are produced will provide a quantitative measure of how often and to what extent will the operational state of the network degrade. Also, the proposed state-space framework provides a clear idea of the service behavior for different regions of the operational space. By combining the evaluation of errors with the proposed framework, it is possible to determine steady state resilience function of the system at any service interface.

# Chapter 5

# Network Topology Resilience

In this chapter, we present resilience evaluation of network topologies using the metrics framework proposed in Chapter 4. Furthermore, we present mechanisms to improve the resilience of the network at the topology and routing sub-layers. Since the resilience of a network depends strongly on the actual physical topology, we need to generate realistic physical topologies. Hence we first introduce a realistic topology model based on location and cost constraints. Secondly, we apply the proposed metrics framework to evaluate the resilience of the generated topologies. Simulations are conducted to evaluate the resilience of topologies in the presence of link failures. One of the resilience measures against link failures is to use path redundancy and diversity as suggested by principles P11 and P12 of the ResiliNets architecture in Chapter 3. Therefore, using the proposed metrics framework, we evaluate the effectiveness of one such mechanism [52] that exploits path-diversity to improve the resilience of the network against link and node failures.

The rest of the chapter is organised as follows: first, we present a realistic physical topology generation model in Section 5.1 followed by resilience evaluation of various topologies in Section 5.2. Lastly, we present the resilience analysis of a multi-path routing protocol in the presence of link failures in Section 5.3.

## 5.1 Generating Physical Topologies

Realistic network topology models are crucial for network research and are commonly used for the analysis, simulation, and evaluation of various mechanisms and protocols. The vast majority of work from this perspective is aimed at recreating router- or AS-level graphs that are representative of the inferred Internet topology. In this section, we consider network topology models in the context of physical topologies. In order to evaluate the resilience and survivability of networks, it is necessary to generate realistic *physical* topologies that are governed by the infrastructure as opposed to only logical topologies that are governed by policy or higher-layer abstractions. We argue that the dominant factor that influences the actual physical topologies is the geographic node location and distribution, which in turn is based on population distribution. Hence we explore location and cost constraints to generate realistic physical topologies. Furthermore, given the tiered nature of real networks and the diversity of topology at each tier, we propose the use of hierarchical models that utilize distinct graph generation methods at each level. Based on these principles, we present a generation model and discuss the resulting topologies.

### 5.1.1 Logical versus Physical topologies

The majority of the existing body of research is based on logical topology models focusing on the generation of either router-level [53] or AS-level topologies [53, 54]. In the router-level graph, each router is represented as a vertex and a logical (IP) link between a pair of routers that forms the edge between the vertices. However, an edge between a pair of vertices does not necessarily imply a direct physical link without any intermediary lower layer nodes. In an AS-level topology, each AS (autonomous system) is represented as a single node and the BGP (border gateway protocol) connectivity between the ASes represents the graph edge connectivity. Hence, *neither the router-level nor the AS-level graph represents the actual physical connection between nodes.* Layer 3 links are logical connections consisting of multiple physical links between Layer 2 and layer 1 components such as switches, multiplexers, regenerators, and optical amplifiers. Furthermore, Layer 3 topologies are frequently not representative of the underlying infrastructure due to Layer 2.5 technologies such as MPLS, SONET, and Metro Ethernet that permit dynamic rearrangement of paths for traffic engineering, policy, and restoration. Thus it is possible for two distinct IP paths to share the same link, making it difficult to understand and engineer the resilience of the network by assuming that IP links correspond to physical links. If we do not understand the geographic location of physical network nodes and links we do not know if they share fate, as was the case in the Baltimore tunnel fire [55] in which many logically distinct links failed at the same time.

Recently, there has been increased interest in the resilience and survivability aspects of communication networks. All currently evolving networks consider survivability in network planning and design phase to some extent. Various studies on future networks [56, 57] place tremendous emphasis on the resilience and survivability of the networks. For example, one of the goals of the Future Internet Design (FIND) [58] Postmodern Internetwork Architecture [59] project is to design networks with high availability. In this case, it is not good enough to merely look at logical topologies, because IP-level connectivity while being fundamental to numerous aspects of the network is not fundamental to the survivability of the network. In other words, while logical topologies are valuable from the protocol perspective, they do not form a very good basis for resilience evaluation in the presence of physical challenges to the network. Furthermore, evaluating resilience from a service perspective requires an end-to-end network model considering the end-hosts as opposed to just the core network.

Realistic physical topology models and generators are required to permit rigorous resilience studies including formal analysis, detailed simulations, and experimentation. This process is a step further from the use of simplistic graph properties such as clustering coefficient and betweenness as an approximate measure of resilience.

### 5.1.2 Approach

The main thrust of this work is to model and generate realistic physical-layer topologies. Therefore, the generation model should be representative of the actual network structure and evolution process. Some of the well known, yet fundamental aspects of real physical topologies are: a) hierarchical nature with significantly varying structure at each level, b) modular with level-specific graphs, and c) location and cost constrained.

We propose a generation model that is hierarchical with $n$-levels [60]. Furthermore, the graph models used at each level differ significantly. In other words, there is not a single overarching graph model that can generate realistic graphs at each level. Instead the graph models used at each level vary from closed form general-purpose models (e.g. Waxman [61]) to pre-structured subgraphs (e.g. trees, rings). It is known that the physical topology of real networks is highly structured and follows to a large extent the population distribution. So instead of using random node positioning, we use well-defined geographic models for node positioning at each level. These range from using exact geocoordinates from known existing networks[1] to probabilistic distributions (e.g. heavy tailed). Lastly, real physical topologies, especially at the backbone level, are highly cost constrained. High resilience can be achieved at unacceptably high costs, however unrealistic. If this were the case, net-

[1] While it is very difficult to infer data on the physical links of existing networks, the node locations, often coinciding with major population and commercial hubs, are easier to obtain.

works would be full meshes. Hence realistic generators must have the ability to *produce feasible topologies at finite cost.*

In this section, we present a synthesis of principles on which real networks are designed and apply then to topology generation. Furthermore, we discuss their implementation in a sample model followed by the analysis of the proposed model.

### Hierarchy

Hierarchical models were studied very early on in the context of logical topologies (router-level and AS-level). It has been observed that Internet exhibits loose hierarchy [62, 63]. We now revisit the hierarchy in the context of physical network topology. Based on *real* networks, the proposed model is hierarchical and this hierarchy permeates through various aspects of the generator – from the graph models, to design constraints, to optimizations desired, to location and cost constraints.

However, we note that when compared to the usage of hierarchy in logical topologies, there is a fundamental difference in which we apply hierarchy to physical layer topology. Hierarchy in the proposed approach is not limited to the way it shapes the connectivity between different levels, but also enables the use of distinct graph evolution mechanisms at each level independently as well as for different parts of a given level, e.g. separate access networks.

### Level-Specific Models

Real networks tend to be modular, i.e. different levels of hierarchy in the network use different graph models. For example, while backbone topologies are tend to be generally connected in a mesh or mesh-like graph [64], access networks on the other hand tend to use ring, star, and tree-like graphs. Lastly subscriber networks are connected directly to the access networks [64] forming a star like topology. Hence, we use level-specific graph generation models for each level of the hierarchy. Furthermore, some levels of the network use a set of discrete pre-structured graphs instead of a single homogenous graph model, in particular access networks are modeled using a combination of ring (e.g. SONET), tree (e.g. HFC – hybrid fiber coax), and mesh.

### Location Constraints

The physical topology of networks is highly constrained by the geographic location of its components. Several works in the past (e.g. [65,66]) have shown that the router-level topology shows a very high correlation to the population density. Moreover, the probability of links is strongly related to the distance between the nodes. It has been shown that the geographic distance-based models such as Waxman accurately model the link distribution when considering location constraints [65]. However, we are not aware of any existing models that apply these fundamental principles to physical network topologies. 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 growth model for each level. While the position and distribution of the level-1, also known as backbone PoPs (point of presence), nodes is based on the population distribution (or other location constraints such as existing PoP locations in a given region or presence of NAPs – network access points), the distribution of the access networks and access network nodes requires further research to determine the distributions that model it accurately.

We generate level-1 node locations based on the population distribution of the United States and compared it with the existing ISP node locations. Figure 5.1 shows a map of 104 nodes generated by finding clustering points on the US population dataset. This map also shows the actual PoP locations

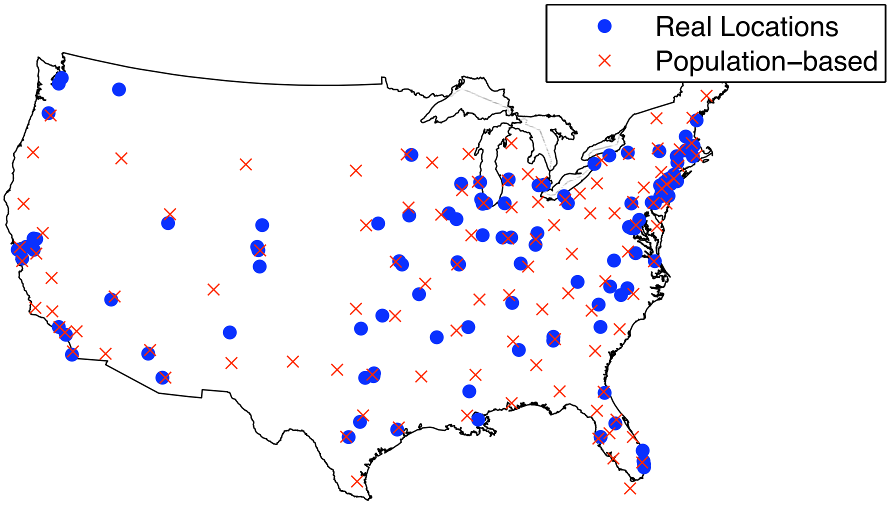

Figure 5.1: Comparing population based and real PoP locations

for four major US ISPs (Sprint, Level 3, and AT&T) from the Rocketfuel database [67]. In order to compare the error between the points generated purely based on population clusters and the known locations, we show the CCDF (complementary cumulative distribution function) of the error (offset distance) in Figure 5.2.

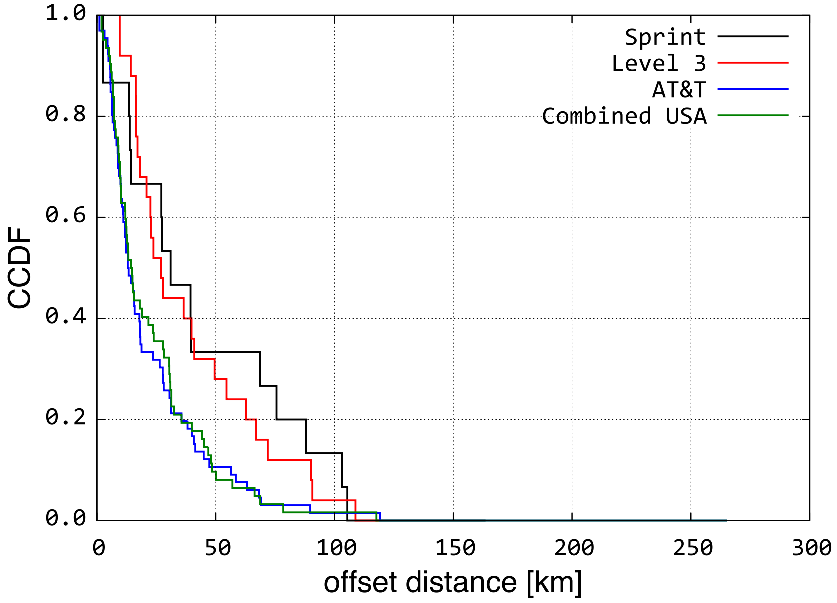

Figure 5.2: Offset distance between real and generated node locations

The offset distance is defined as the absolute distance between a generated point and the closest real PoP. The offset distance is plotted when comparing the points against a particular ISP network as well as against the combined PoPs of all three networks. The most meaningful conclusions can be derived when comparing the node locations generated using the population dataset against the combined PoP locations of all networks. In this case, we observe that 98% of nodes generated by KU-LocGen are accurate within an offset distance of 60 km.

#### Cost Constraints

Cost constraints significantly impact practical network design and evolution. Economic factors shape physical level infrastructure [68]. 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 propose the use of modular cost functions: the model should support multiple 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. We present one such model in our preliminary implementation.

### 5.1.3 Generation Model

In this section, we present a 3-level hierarchical model. These hierarchical levels represent the backbone, access, and subscriber networks whose geographic distribution is represented by $\Psi_p$, $\Psi_n$, $\Psi_s$ 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 $\Psi_p$ supports three different location constraints including fixed geographic positions based on known point-of-presence locations (of existing networks), user defined, and a random distribution as discussed below. The number of level-2 access networks $N(i)$ 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, thus $\Psi_n = N[\mu_n, \sigma_n^2]$, where $\mu_n$ represents the PoP location and $\sigma_n^2$ is the variance. The subscriber networks are distributed normally $\Psi_s = N[\mu_s, \sigma_s^2]$, where $\mu_s$ represents the access network location and $\sigma_n^2$ 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 $M(i) = \frac{\max(d_t);t=1,2,\ldots N(j)}{d_i} \times M_{\min}$, 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; $S(i) = \frac{M(i)}{\max(M(j)); j=1,2,...N} \times S_{\max}$ where $S_{\max}$ is the predefined limit on the maximum number of subscribers per access network.

The location constraints are implemented at any given level using a simple cost model based on node and link costs. For simplicity, we assume 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 link, $v_{i,j}$ is the variable cost per unit distance for the link and $d_{i,j}$ is the length of the link. For simplicity we choose $v_{i,j} = \bar{d} \times v_{i,j}$ where $\bar{d}$ is the average link length of the network.

The link generation is based on a number of different models depending upon the level of the network. The backbone nodes are connected using a cost-constrained Waxman model. The Waxman model accurately represents the link connectivity in backbone networks [65][2]. On the other hand, the access networks are connected using pre-structured topologies that reflect real deployments such as ring (SONET, Metro Ethernet), tree (HFC), mesh and star. For each access network, one of these topologies is chosen randomly. Finally, subscriber networks are connected to the geographically closest access network node. Figure 5.3 shows the components of a hypothetical topology

[2] While it is generally agreed that backbone networks are mesh-like, there is some contention as to exact relationship between link probability and its distance. While some works claim that this is exponential [65], others claim that this is linear [66].

and illustrating the hierarchy, location, and level-specific modeling of our approach. For the sake of simplicity, in this example we do not consider the individual (optical) spans that make up links. Hence, link are treated as direct entities between PoP nodes. However, the model itself is flexible and does not impose any limitations on the use of fiber spans.

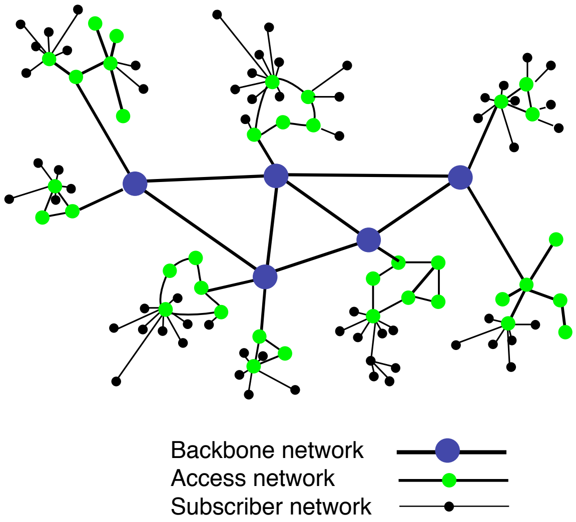

Figure 5.3: Illustrating the hierarchical topology model

Lastly, it should be noted that since the overall generator is modular, it allows the users to implement new models to be plugged in at the appropriate levels. The distributions and graph models used in our preliminary implementation provide an example of a typical physical network topology model. Further research is required to determine the optimal models that would yield highly representative plots. The tradeoff between abstraction and fidelity is supported by letting the network designer specify arbitrary detail.

We generated several topologies based on the location constraints imposed by existing infrastructure. For example, we used the Sprint and GÉANT PoP locations to constrain the distribution of the level-1 backbone nodes of the proposed model. Figure 5.4 shows one of the sample topologies generated using the proposed model. In this Figure, the access network nodes are aggregated for visual purpose. Figure 5.5 shows the detailed 2-level topology along with the access network node connectivity. We point out that visual inspection of topologies does not provide any rigorous analysis. The purpose of these figures is to simply provide an illustration of the effect of geographic location constraints and discrete graph models. As for the degree distributions, we observe that the degree-frequency relationship does obey power law. However, the generated graphs do not strictly obey *all* power laws in the literature.

## 5.2 Evaluating Network Topologies

Several measures have been proposed in the literature [69] to evaluate the ability of the network to survive link and node failures as well as various attacks. In this section, we apply the proposed resilience to wired network topologies. We conduct this study at the $B_{\mathrm{3t,3r}}$ boundary between the topology sub-level (3t) and the path routing sub-level (3r). In this case, a set of vertices $V$ and edges $E$ and link failures $f$ characterize the operational state of the network. The service across this boundary is *topological connectivity.*

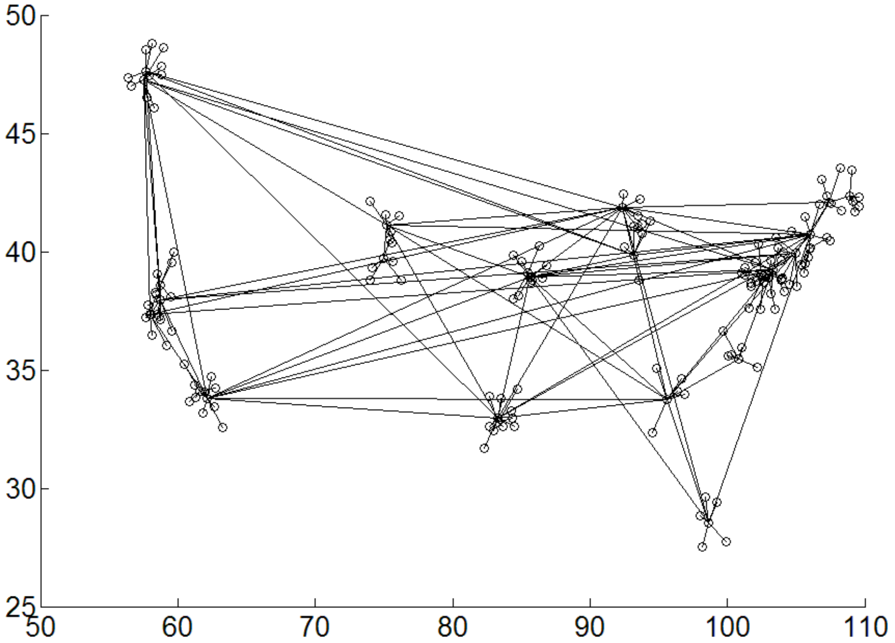

Figure 5.4: 2-level topology based on Sprint PoP coordinates

For the purpose of this study we selected three service provider backbone network topologies, Sprint (US), AT&T (US) (Rocketfuel database [67]), and GÉANT2 [70]. These topologies are shown in Figures 5.6, 5.7, and 5.8 and their statistics are shown in Table 5.1.

Table 5.1: Network statistics

| | Sprint | AT&T | GÉANT2 |
|---|---|---|---|
| nodes | 27 | 25 | 34 |
| links | 136 | 92 | 102 |

Since we consider only link failures, we chose a single operational metric $n_1$ to represent the number of link failures. Therefore $\mathbb{N}_{3\mathrm{r}} = \{n_1\}$. In order

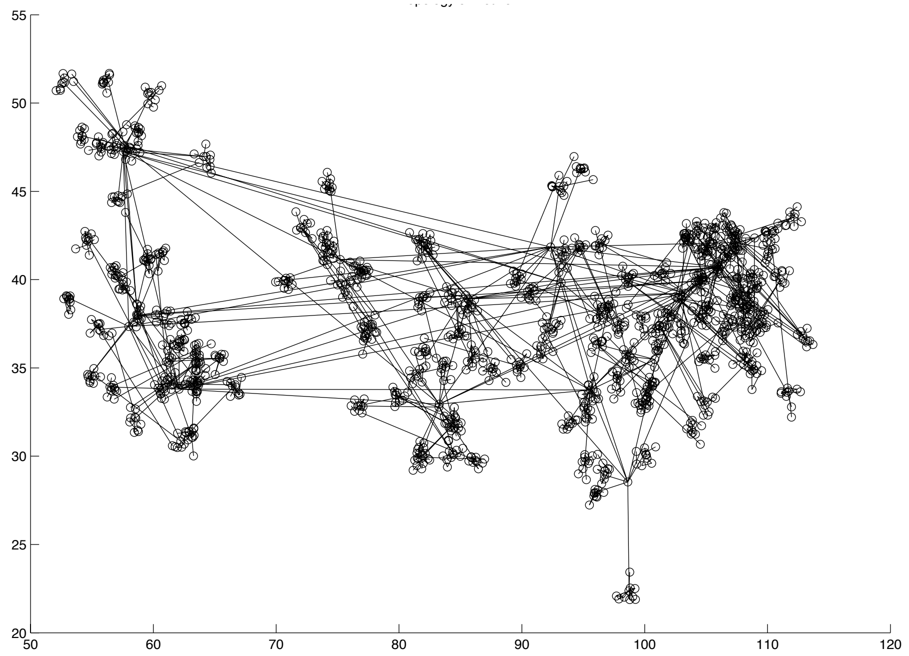

Figure 5.5: 2-level topology with access networks expanded

to characterize this service, we choose two service parameters as the relative size of the largest connected component $p_1$ and the clustering coefficient $p_2$. Therefore, $\mathbb{P}_{3\mathrm{t}} = \{p_1, p_2\}$. We conducted simulations in MATLAB to evaluate the impact of link failures on the service parameters at this boundary. We explore the operational metric (link failure probability) over the range of [0, 0.5] in intervals of 0.01. The simulations results are averaged over 100 runs. Figures 5.9, 5.10, 5.11, 5.12 show the impact of link failures on four graph metrics: number of partitions, relative size of the largest connected component, average degree, and clustering coefficient for the AT&T topology.

For all plots, 95% confidence intervals of the mean are shown. As expected

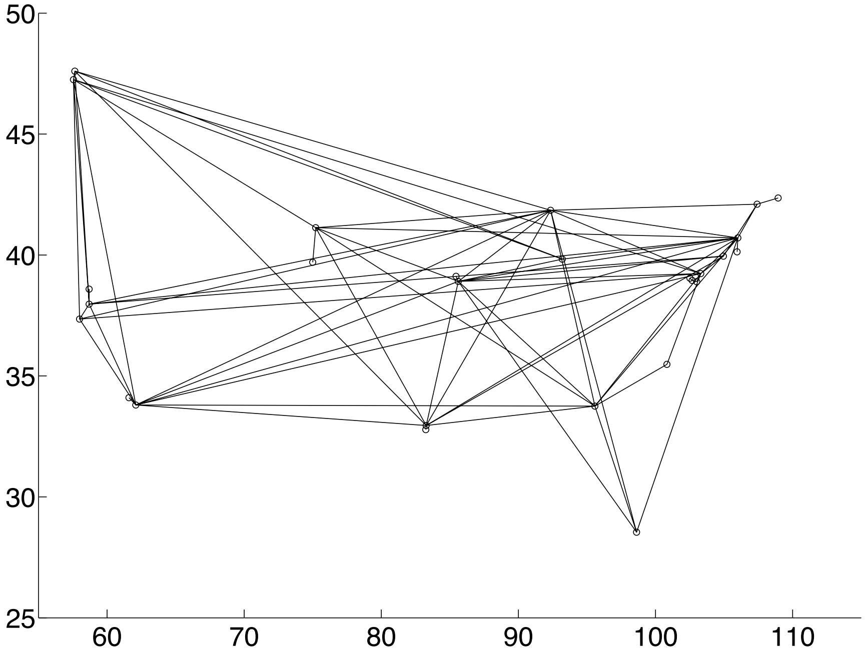

Figure 5.6: Sprint topology

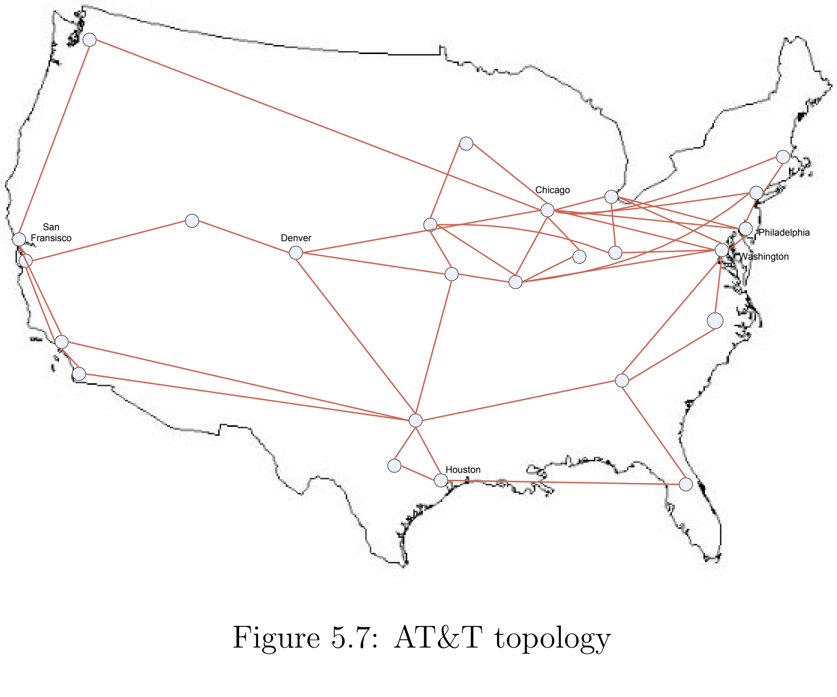

Figure 5.7: AT&T topology

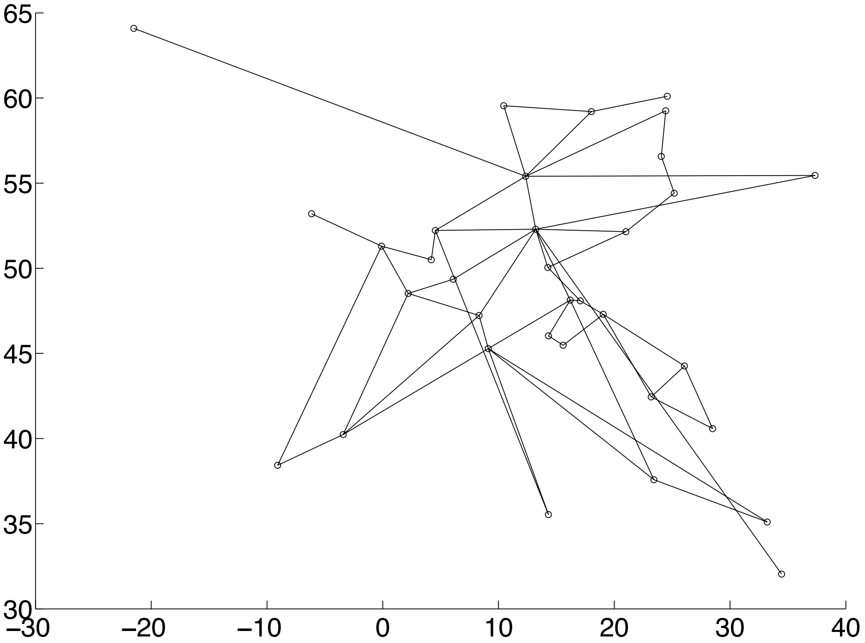

Figure 5.8: GÉANT2 topology

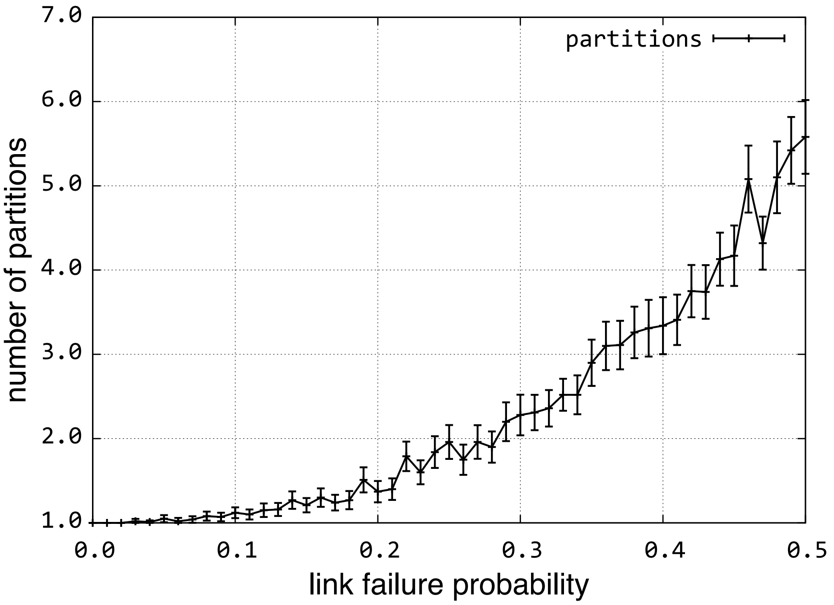

Figure 5.9: Effect of link failures on AT&T partitions

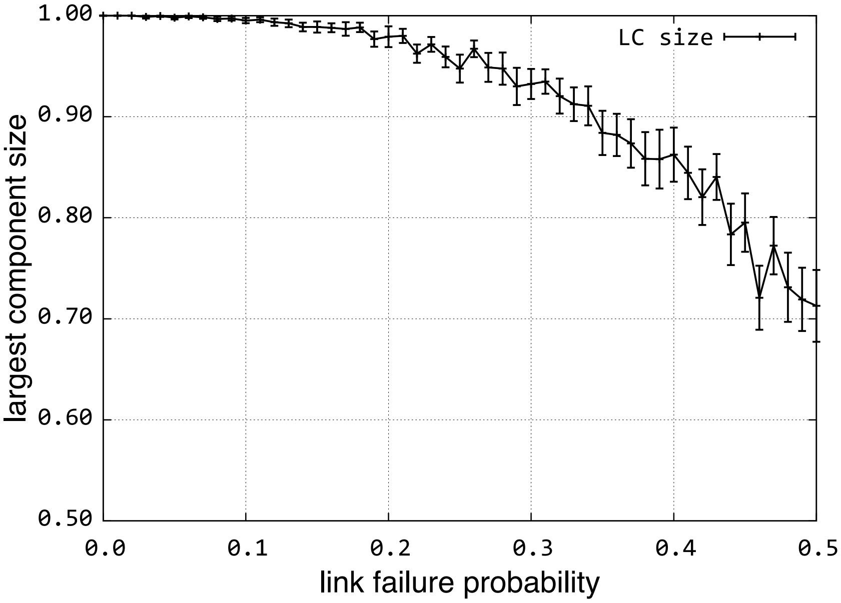

Figure 5.10: Effect of link failures on AT&T connectivity

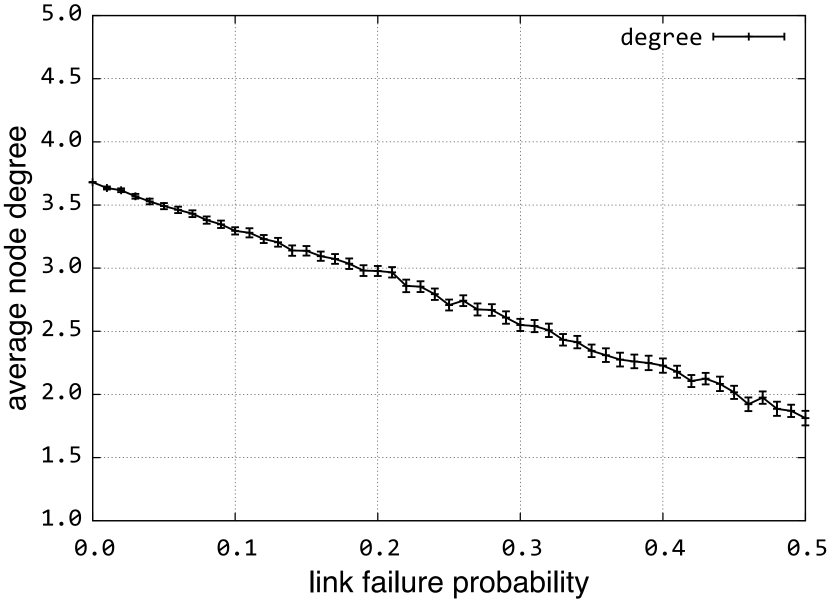

Figure 5.11: Effect of link failures on AT&T partitions

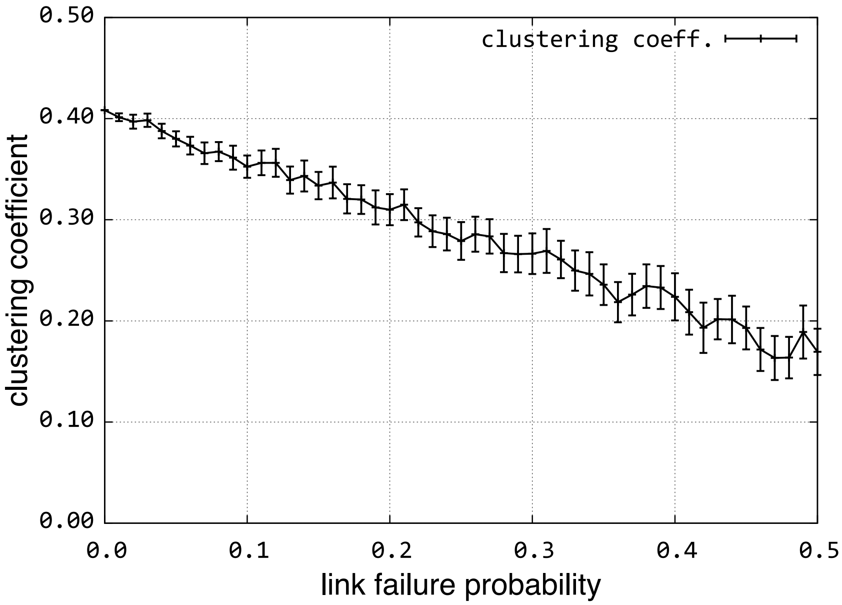

Figure 5.12: Effect of link failures on AT&T partitions

almost all the metrics decline with increasing probability of link failures. The slope of this decline varies with the metric as well as the topology being evaluated. Though traditional metrics like these give partial insight into the topology, in order to evaluate resilience, we need to look at a particular service that is delivered by the network at the given level boundary. In this example, we define the topology service by selecting two service parameters: the relative size of the largest connected component, which represents the reachability of the graph and clustering coefficient and thus represents 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. As an example, note that the clustering coefficient of a grid network (manhattan grid) is zero and a single link failure would result in significantly higher path lengths.

The regions for operational metrics (link failure probability) are defined in Table 5.2 and service parameters are defined in Table 5.3. Note that these are arbitrarily chosen boundaries for the purpose of analysis. However, the bounds are kept *reasonable* in a general sense. Depending upon the scenario, the framework can be used with a different set of bounds to evaluate resilience in that particular context.

Table 5.2: Operational regions at $B_{\mathrm{3t,3r}}$

| Region | Operational metrics<br>link failure probability $n_1$ |
|---|---|
| Normal | $n_1 \leq 0.01$ |
| Partially degraded | $0.01 < n_1 \leq 0.20$ |
| Severely degraded | $n_1 > 0.20$ |

Table 5.3: Service regions at $B_{\mathrm{3t,3r}}$

| Region | Service Parameters | |
|---|---|---|
| | LC size $p_1$ | clustering coefficient $p_2$ |
| Acceptable | $p_1 = 1$ | $p_2 \geq 0.33$ |
| Impaired | $0.90 \leq p_1 < 1$ | $0.25 \leq p_2 < 0.33$ |
| Unacceptable | $p_1 < 0.90$ | $p_2 < 0.25$ |

Given these operational and service regions, we plot the simulation results on a piece-wise linear axis. Figure 5.13 shows the steady state resilience of the AT&T network to link failures as degradation in the service from acceptable to unacceptable region. The region boundaries in both the operational and service dimensions are arbitrarily chosen. 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.

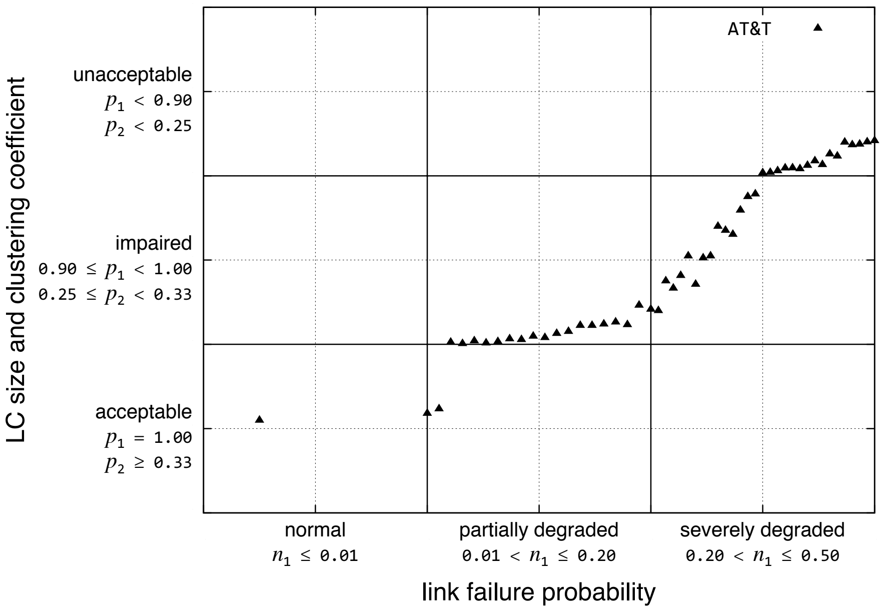

Figure 5.13: Resilience of AT&T topology to link failures

We see that the network lies in the acceptable service region under normal operating conditions. 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 impaired to unacceptable regions. In this example, 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 network remains acceptable for all operational conditions, the area under the curve will be zero. This would represent perfect resilience. In order to get a normalized value of resilience, we define resilience $\mathbb{R} = 1 -$ normalized area, where normalized area is the total area divided by the span of the $x$-axis. For the plot shown in Figure 5.13, $\mathbb{R}$ is calculated to be 0.6338. This area represents an aggregate measure of the resilience at this boundary [3].

Figure 5.14 compares the AT&T, Sprint, and GÉANT2 topologies. In this case, we observe that AT&T has better resilience compared to Sprint and GÉANT2 topology. The resilience $\mathbb{R}$ for AT&T is 0.6338 and that for sprint is 0.5410 and for GÉANT2 is 0.4721. We observe that due to fewer number of links, the GÉANT2 topology has very low clustering coefficient and the topology service performs poorly even in the normal operational regions.

Next, we apply the metrics framework to the synthetic topologies generated using the proposed generation model. We evaluate the resilience of a 104

[3] While this method provides a single value for resilience, it should be noted that to gain a complete understanding of the resilience, it is necessary to inspect the state-space plot as shown in Figure 5.13

node level-1 topology generated using the KU-LoCGen. Figure 5.15 shows the state space at the topology–routing boundary with varying value of $a$, the link probability factor of the Waxman link generation model.

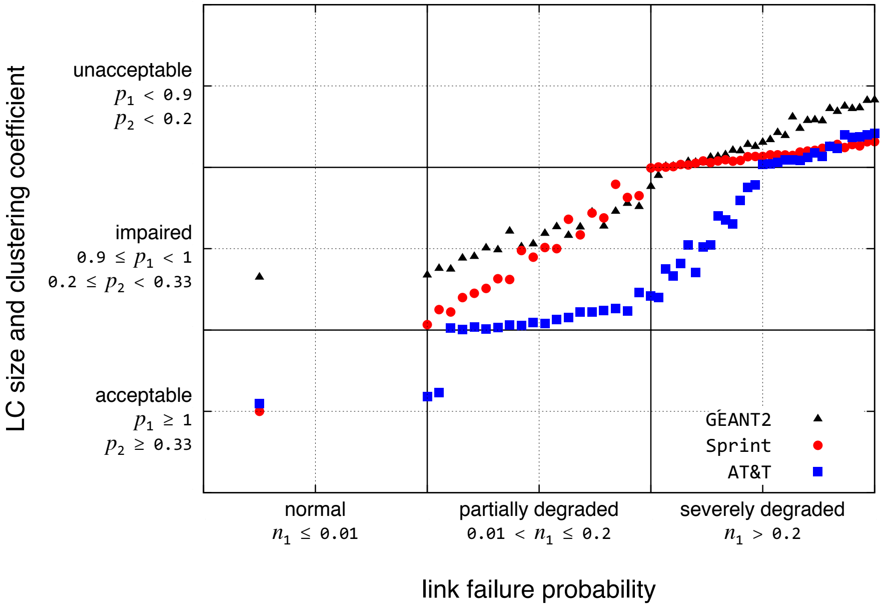

Figure 5.14: Comparing resilience of AT&T, Sprint, and GÉANT2

We observe that as $a$ increases, the resilience of the network increases for a given cost. While this particular graph is only an example, it demonstrates how the proposed framework can be used to evaluate the impact of topology model parameters on the resilience of the resulting graph.

## 5.3 Evaluating Path Routing

In this section we discuss how to evaluate the resilience of a network at the boundary $B_{3\mathrm{r},4}$ between topology and the transport boundary. In order to

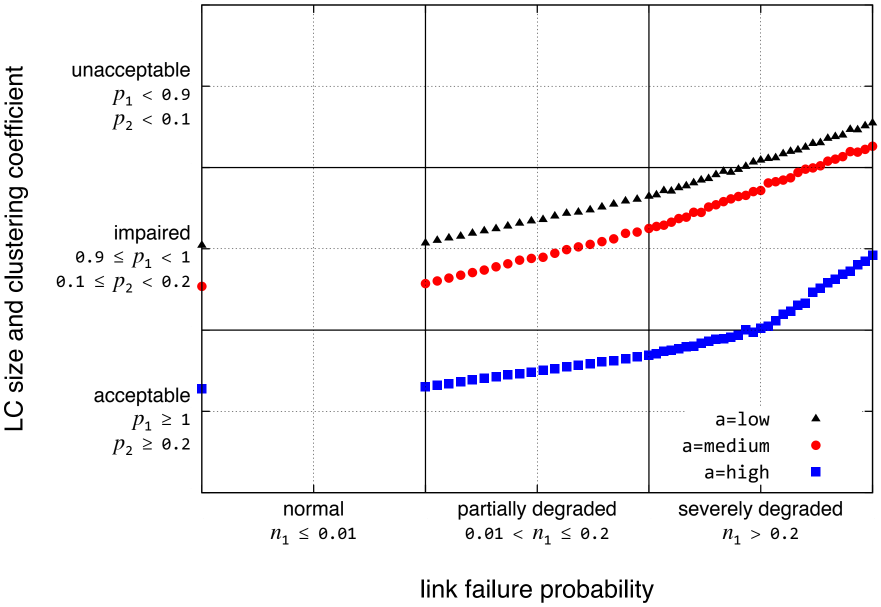

Figure 5.15: Resilience of topologies generated with KU-LoCGen

demonstrate the resilience at this level, we consider a multi-path routing mechanism [52] that exploits the diversity of improve the resilience of the paths. A brief summary of this algorithm is presented below.

### 5.3.1 Path Diversification

The primary motivation for path diversification mechanism is to increase resilience by choosing paths such that they will not experience correlated failures. For this purpose a diversity measure is defined that quantifies the degree to which alternate paths share the same nodes and links. The definitions for a *path* and *path diversity* are:

**Path:** 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:

$$P = L \cup N \tag{5.1}$$

and the length of this path $|P|$ is the combined total number of elements in $L$ and $N$.

**Path diversity:** Let the shortest path between a given ($s$, $d$) pair be $P_0$. 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:

$$D(P_k) = 1 - \frac{|P_k \cap P_0|}{|P_0|} \tag{5.2}$$

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:

$$D(P_b, P_a) = 1 - \frac{|P_b \cap P_a|}{|P_a|} \tag{5.3}$$

where $|P_a| \leq |P_b|$

The paths selection depends on the mode of operation. In this dissertation, we evaluate the resilience of multipath routing when using *k-path* mode in which the objective 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. In this mode, the multipath mechanism uses these paths in hot standby mode, in which if one of the paths fail, it seamlessly switches to using the other paths associated with the node pair. Hence a node pair is reachable as long as one of the $k$ paths chosen is still available.

### 5.3.2 Multipath Resilience

We now evaluate the resilience of this mechanism over the topologies discussed in the previous section. For this purpose, we conduct the resilience analysis at the topology and routing sublayer boundary $B_{\text{3r,4}}$. As mentioned in the Section 4.4.2, the service parameters from the level below become the operational metrics at the current level. In this case, the operational metrics $\mathbb{N}_{\text{3r}}$ at boundary $B_{\text{3r,4}}$ are the service parameters $\mathbb{P}_{\text{3t}}$ from the resilience boundary $B_{\text{3t,3r}}$. Hence, $\mathbb{N}_{\text{3r}} = \{n_1, n_2\} = \{\text{LC size, clustering coeff.}\}$.

Table 5.4: Operational regions at $B_{\text{3r,4}}$

| Region | Operational Metrics | |
|---|---|---|
| | LC size $n_1$ | clustering coefficient $n_2$ |
| Normal | $n_1 = 1$ | $n_2 \geq 0.33$ |
| Partially degraded | $0.90 \leq n_1 < 1$ | $0.25 \leq n_2 < 0.33$ |
| Severely degraded | $n_1 < 0.90$ | $n_2 < 0.25$ |

The service provided by the path routing (sub) level to the level above (transport) is *paths.* In order to characterize this service, we define two service metrics: path robustness and stretch. We compute path robustness as the number of reliable flows, divided by the total number of flows in the network. A flow is considered reliable if at least one of its paths remains unbroken by the link failures. Hence $\mathbb{P}_4 = \{p_1, p_2\} = \{\text{path robustness}, \text{stretch}\}$. The normal, partially degraded, and severely degraded regions at this boundary are given in Table 5.4. Similarly, the acceptable, impaired, and unacceptable service regions are defined in Table 5.5.

Table 5.5: Service regions at $B_{\text{3r},4}$

| Region | Service Parameters | |
|---|---|---|
| | path robustness $p_1$ | stretch $p_2$ |
| Acceptable | $p_1 \geq 0.99$ | $p_2 \leq 1.10$ |
| Impaired | $0.75 \leq p_1 < 0.99$ | $1.10 \leq p_2 < 1.50$ |
| Unacceptable | $p_1 < 0.75$ | $p_2 > 1.50$ |

Figure 5.16 shows the resilience of multi-path mechanism when using 1, 2 and 3 paths over the AT&T topology.

We observe that as the number of paths used $k$ increase, 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[4]. The resilience state space plots for Sprint and GÉANT2 networks when using multi-path mechanism are shown in Figures 5.17 and 5.18. Note that the GÉANT

[4] We are not considering active rerouting effects in this example.

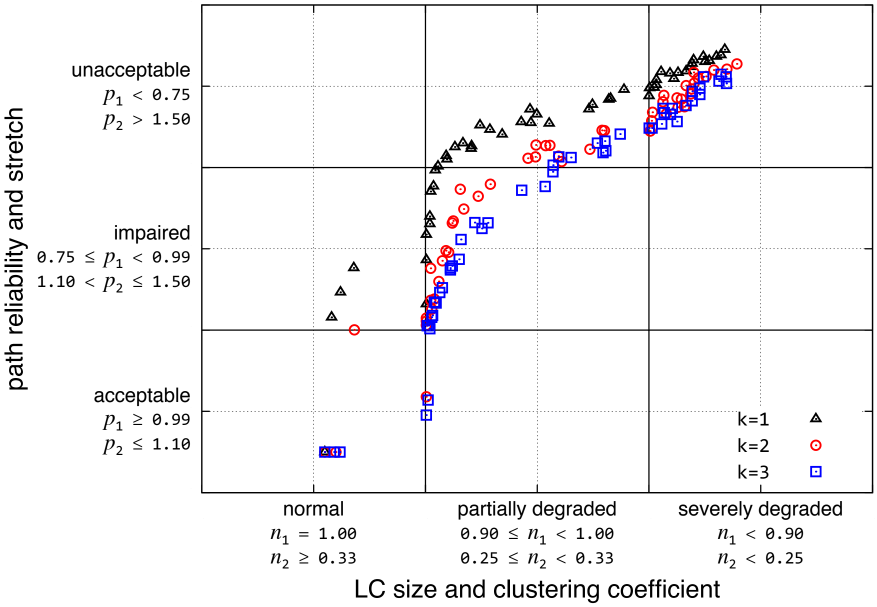

Figure 5.16: Resilience of multi-path routing over AT&T topology

topology is a sparse graph with low clustering coefficient and did not have any samples in the normal operating regions with the boundaries defined in this example.

In order to get a single measure of the resilience at this service boundary, we show the aggregate resilience $\mathbb{R}$ for the three networks in Table 5.6.

## 5.4 Summary

In this chapter we have shown how the proposed metrics framework can be applied to evaluate the resilience of networks at multiple levels. Specifically, we showed the evaluation of resilience at the topology–path routing bound-

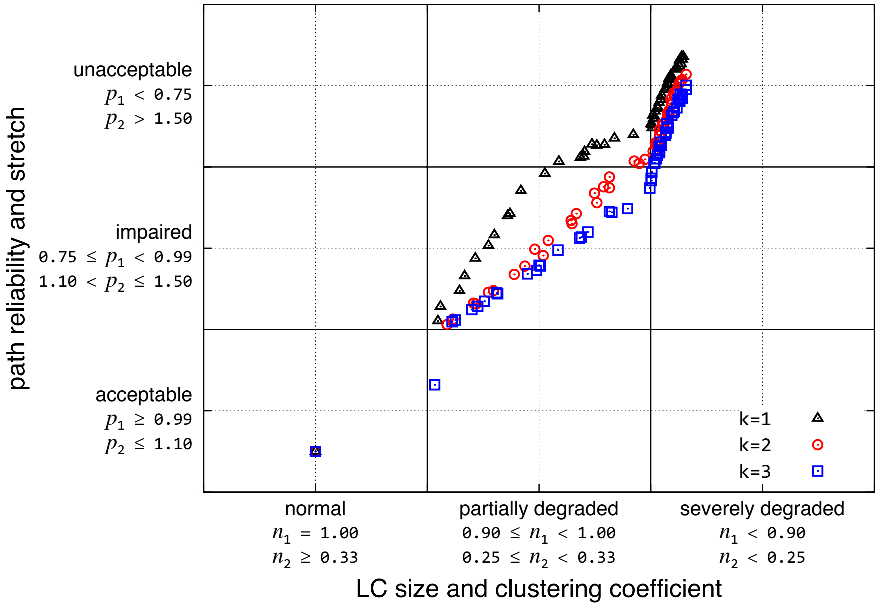

Figure 5.17: Resilience of multi-path routing over Sprint topology

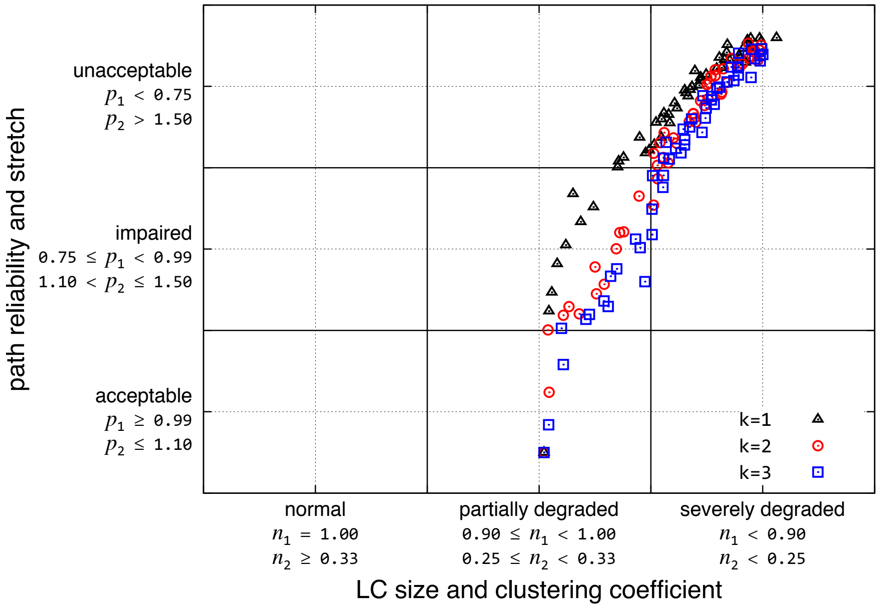

Figure 5.18: Resilience of multi-path routing over GÉANT2 topology

Table 5.6: Aggregate resilience at $B_{3\mathrm{r},4}$

| Number of paths | Aggregate Resilience $\mathbb{R}$ | | |
|---|---|---|---|
| | $k = 1$ | $k = 2$ | $k = 3$ |
| AT&T | 0.5885 | 0.6379 | 0.6854 |
| Sprint | 0.7259 | 0.7596 | 0.7830 |
| GÉANT2 | 0.4909 | 0.5967 | 0.6195 |

ary and path routing–transport boundary. Simulations were conducted to evaluate the resilience of the topology against link failures and the resilience of a multi-path mechanism against perturbations in the underlying topology. Note that while we presented our analysis based on a specific numerical examples, the same analysis could be conducted with different boundaries for the state space regions or a different service definition. We presented a model to generate realistic hierarchical network topologies using real world constraints. Lastly, we also presented a path diversification mechanism that exploits the multiple paths to improve the resilience of the routing service in the presence of challenges to the underlying network.

Page left intentionally blank.

# Chapter 6

# Resilience of MANETs

Wireless mobile ad hoc networks (MANETs) present an interesting case for resilience analysis. MANETs are inherently challenged due their mobile wireless environment and are therefore prone to service failures. They are the target of several mechanisms whose objective is to enhance the survivability. We discussed in the previous chapter (Section 4.4.2) how the proposed metrics framework could be applied to evaluate the resilience of the a MANET at multiple levels. In this chapter, we apply our framework to simulation-based studies of MANETs to evaluate their resilience in the presence of various challenges to normal operations.

The rest of the chapter is organised as follows: Section 6.1 describes the simulation setup. The resilience of MANET at three boundaries: topology–routing, routing–transport, and transport–application is presented in Section 6.2, Section 6.3, and Section 6.4 respectively. This is followed by design and evaluation of resilient routing protocols for two different scenarios:

millimeter-wave mesh networks in Section 6.5 and aeronautical networks in Section 6.6. Lastly, a summary of chapter is presented in Section 6.7.

## 6.1 Simulation Setup

We used the ns-3 simulator [71] for conducting the simulations presented in this chapter. While considering all the factors that affect the simulations, we choose parameters such that they cover a wide operational range from normal to severely degraded operations in order to evaluate the resilience of different layers in the presence of challenges. The simulation setup consists of 25 nodes in a 1000 × 1000 meter region. The random way point mobility model [] with zero pause times was used. The wireless simulation uses the 802.11 PHY module in infrastructure mode; the physical channel uses Friis propagation model. The simulation parameters that are varied over the course of the simulation runs include node speed, transmission power (and hence range), and the network load. A detailed list of simulation parameters is given in Table 6.1. All simulations are averaged over 10 runs and 95% confidence intervals are shown as appropriate.

Figure 6.1, reproduces the table of metrics in the multilevel resilience analysis of a MANET. Now, we will evaluate the resilience of the network at various level boundaries.

Table 6.1: Simulation parameters

| Parameter | Category | Value |
|---|---|---|
| number of nodes | fixed | 25 |
| simulation region | fixed | 1000 × 1000 metes |
| transmit range [meters] | variable | 100 – 800 |
| node speed | variable | 5, 10, 20 , 50, 100 |
| propagation mode | fixed | Friis propagation |
| PHY | fixed | YANS wifi Phy |
| MAC | fixed | 802.11 b |
| routing protocol | variable | DSDV, OLSR |
| transport protocol | fixed | UDP |
| data rate | variable | 0.25, 0.5, 1, 2, 4, 8, 16 Kbps |
| packet size | fixed | 1000 Bytes |
| application | fixed | CBR |
| traffic model | fixed | $n(n-1)$ flows = 600 flows |

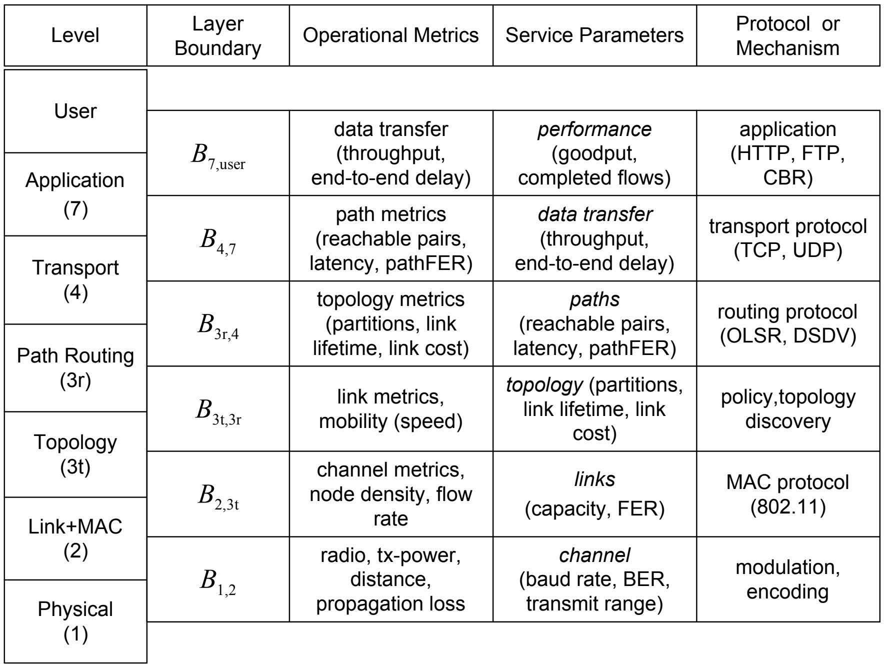

| Level | Layer Boundary | Operational Metrics | Service Parameters | Protocol or Mechanism |
|---|---|---|---|---|
| User | | | | |
| | $B_{7,\text{user}}$ | data transfer (throughput, end-to-end delay) | *performance* (goodput, completed flows) | application (HTTP, FTP, CBR) |
| Application (7) | | | | |
| | $B_{4,7}$ | path metrics (reachable pairs, latency, pathFER) | *data transfer* (throughput, end-to-end delay) | transport protocol (TCP, UDP) |
| Transport (4) | | | | |
| | $B_{3\text{r},4}$ | topology metrics (partitions, link lifetime, link cost) | *paths* (reachable pairs, latency, pathFER) | routing protocol (OLSR, DSDV) |
| Path Routing (3r) | | | | |
| | $B_{3\text{t},3\text{r}}$ | link metrics, mobility (speed) | *topology* (partitions, link lifetime, link cost) | policy,topology discovery |
| Topology (3t) | | | | |
| | $B_{2,3\text{t}}$ | channel metrics, node density, flow rate | *links* (capacity, FER) | MAC protocol (802.11) |
| Link+MAC (2) | | | | |
| | $B_{1,2}$ | radio, tx-power, distance, propagation loss | *channel* (baud rate, BER, transmit range) | modulation, encoding |
| Physical (1) | | | | |

Figure 6.1: Multilevel resilience evaluation of a MANET

## 6.2 Resilience at Topology – Routing

At this boundary $B_{3t,3r}$ (reference Figure 6.1 ). we will evaluate the resilience of the the MANET topology under the presence of challenges to its normal operations. The operational metrics to represent the operational state of the network at this layer are node speed and the transmission range. Secondly, the service parameters that are relevant at this boundary are the relative size of the largest connected component and the average link duration.

### 6.2.1 Variation in parameters

First, we show the variation of the first operational metric using standard two dimensional plots. Figure 6.2 shows the variation of the average link durations with transmit range and speed. The plot shows that the average link duration is significantly affected by both the parameters and varies over a wide range. Since the ability of the routing protocol to find paths depends on the churn in topology, this is a crucial metric to characterize the topology service. The confidence intervals for these measurements are shown in Figure 6.3 which shows relatively high confidence in the values.

Figure 6.4 shows the variation of second operational metric, the relative size of the largest connected component (LC size). As can be seen from the plot, the LC size depends primarily on the transmit range and does not vary much with node speed. This is because the connectivity of the network is heavily dominated by the transmit range especially at the higher end of the

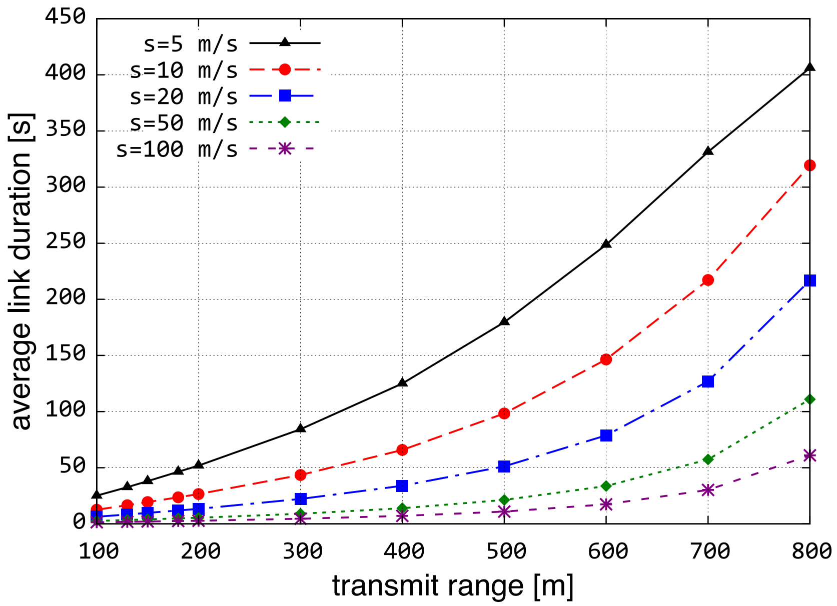

Figure 6.2: Variation of link durations

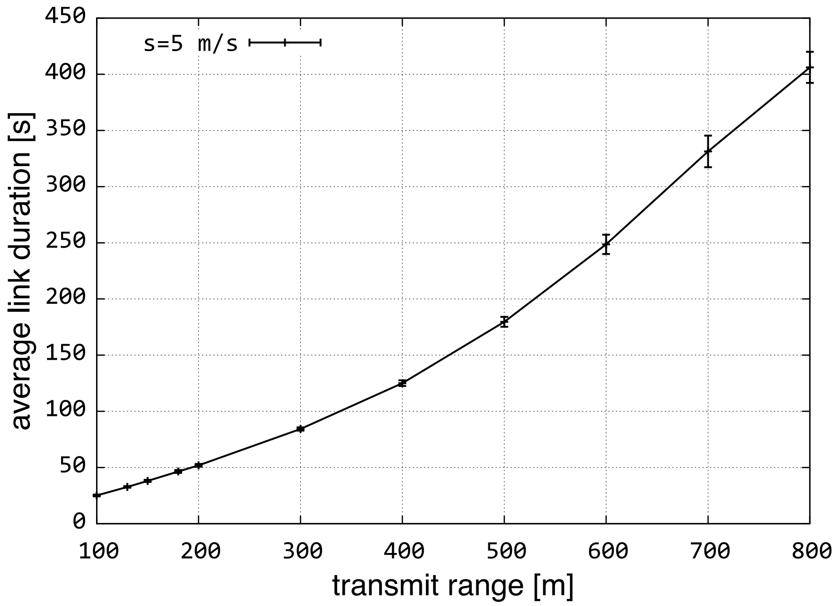

Figure 6.3: Confidence intervals for link duration measurements

transmit range. Again, the 95% confidence intervals in Figure 6.5 show high confidence in the averaged results. Finally, we look at the variation in the number of partitions in Figure 6.6 and observe that this is the inverse of the LC size as expected.

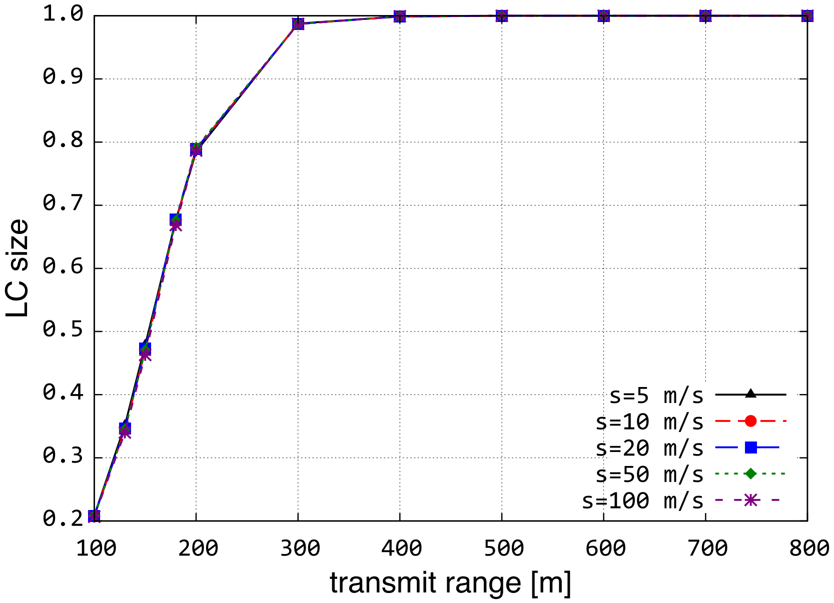

Figure 6.4: Variation of LC size

### 6.2.2 State Space Computations

We now generate the state space representation at this boundary. For all values of the the operational metrics and service parameters, we need to calculate the corresponding projected values, $N^*$ and $P^*$ of the state space region. In order to get a single $N^*$ or $P^*$ value from a set of operational metrics and

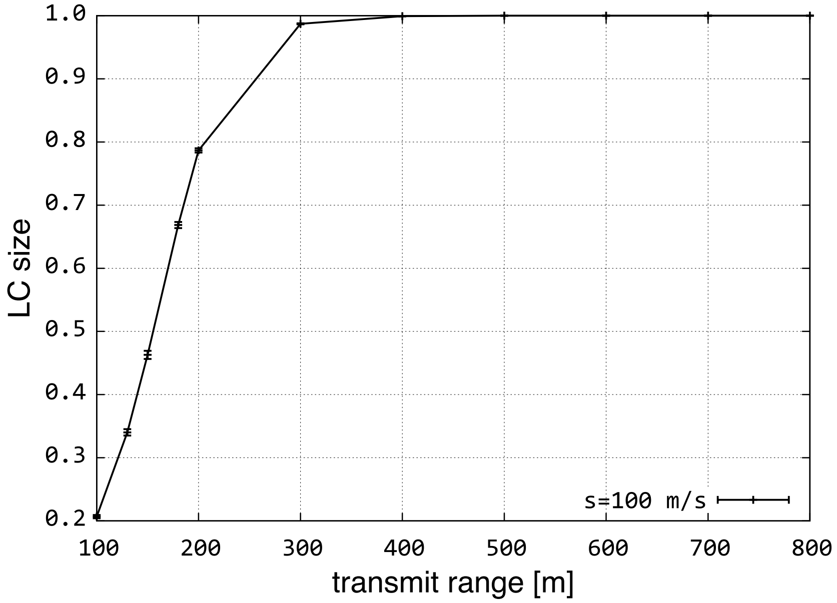

Figure 6.5: Confidence intervals for LC size measurements

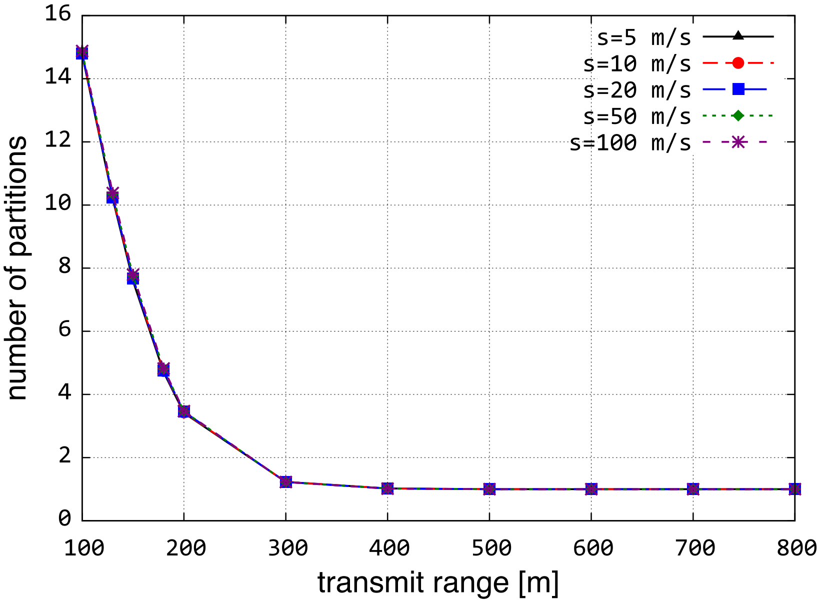

Figure 6.6: Variation of partitions

service parameters, we define an objective function. The methodology used to perform these calculations is as follows:

When operational state and service space is represented more than one non-independent metrics, we need to apply an objective function to obtain the operational metrics and service parameters. Lets say there are two operational state of the network $N_{\mathcal{S}}$ be represented by two metrics $N_{\mathcal{S}} = \{N_1, N_2\}$. In order to project these metrics on to a two dimensional state space, we calculate the projected state $N^*_{\mathcal{S}} = f(N_1, N_2)$ Secondly, we define the boundaries for the normal, partially degraded and severely degraded regions. When the regions are defined this way, they are are simply three states $\mathbb{N}_1, \mathbb{N}_2, \mathbb{N}_3$ with their respective ranges in the projected space $N^*_{\mathcal{S}}$.

In order to calculate the x-axis value (say, $n^*$) for pair of instantaneous values of the operational metrics $n_1, n_2$, we calculate the $n^*_1$ corresponding to $n_1$ and $n^*_2$ corresponding to $n_2$ on a piecewise linear scale. So if $n_1$ lies in the range of the range of the $i_{th}$ regions, then $n^*_1 = \frac{n_1}{\overline{n}_{1i} - \underline{n}_{1i}}$. Once we calculate $n^*_1$ and $n^*_2$ independently, we apply an objective function such that $n^* = \alpha n^*_1 + (1-\alpha) n^*_2$. The values of $\alpha$ is determined by the service specification from the layer above. Furthermore, the framework also supports logical objective functions of AND and OR. These are treated as the special cases and the program written to compute states supports this mode via special flags. The same procedure is repeated for deriving the $p^*$ value from a set of selected service parameters and the region boundaries. The objective function used is either logical (AND, OR) or a linear function: $p^* = \beta p^*_1 + (1-\beta) p^*_2$.

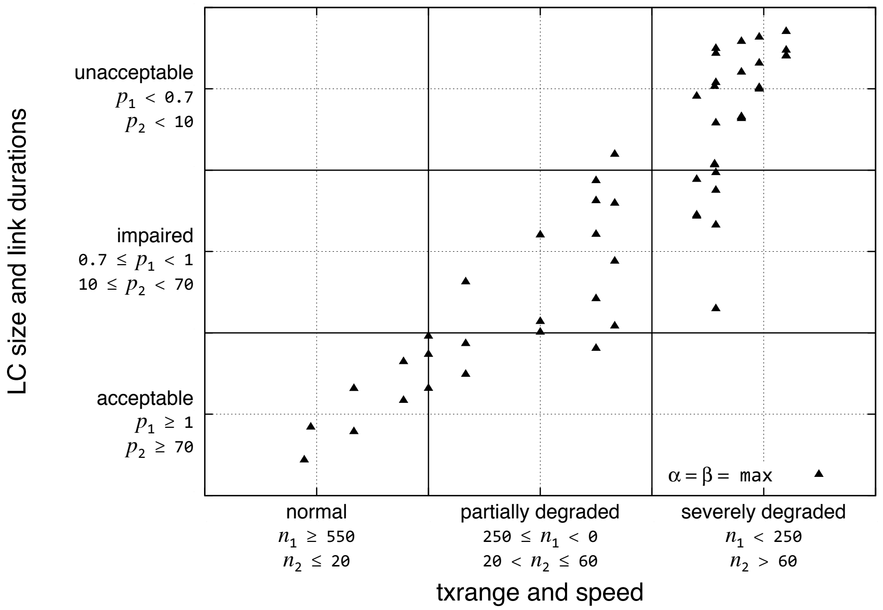

Figure 6.7: Resilience state space at $B_{3\text{t},3\text{r}}$ when $\alpha$ =max and $\beta$ =max

Figure 6.7, 6.8, and 6.9 shows the state space transitions for varying values of $\alpha$ and $\beta$. In these Figures, the keyword *"max"* is used to indicate the logical AND condition. These plots show that the MANET provides an acceptable topology as long as the operational metrics at the level, transmit range and node speed remain normal. However, as the operations degrade, the service degrades in a near linear fashion. The slope of this resilience curve depends on the objective function chosen. In the next section, we evaluate the impact of this objective function and if an upper and lower limit on the resilience can be found.

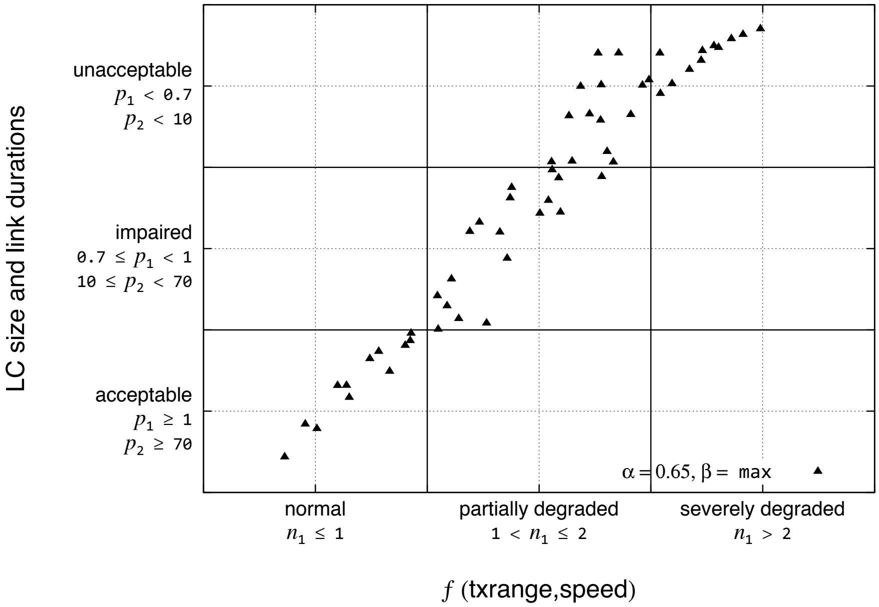

Figure 6.8: Resilience state space at $B_{3\text{t},3\text{r}}$ when $\alpha = 0.65$ and $\beta$ =max

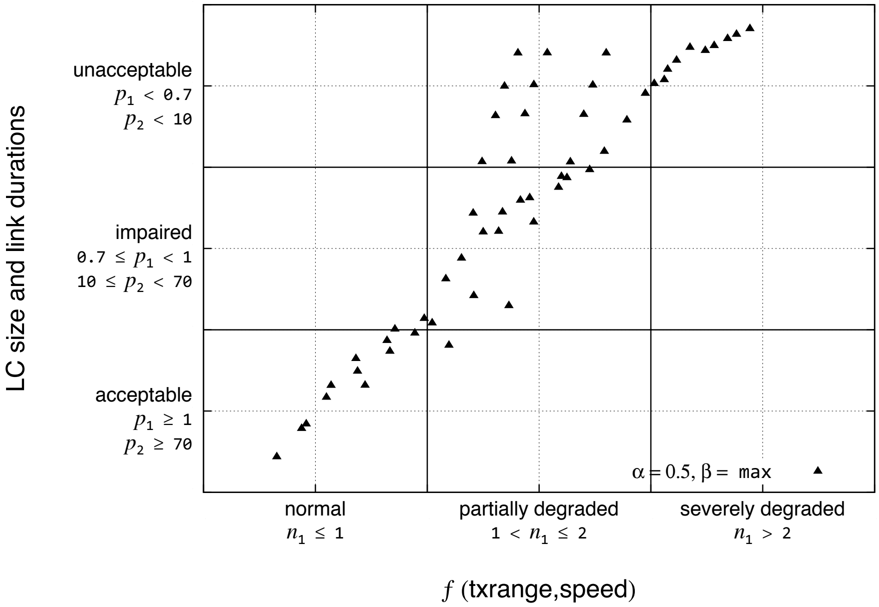

Figure 6.9: Resilience state space at $B_{3\text{t},3\text{r}}$ when $\alpha = 0.5$ and $\beta$ =max

## 6.3 Resilience at Routing – Transport

In this section, we evaluate the resilience of routing protocols (e.g. OLSR, DSDV). We define the service at this boundary as the ability to provide reachable paths to the transport layer in the presence of disruptions or perturbations to the underlying topology. Therefore, we characterize this service using one parameter: *path availability*, which is defined as percentage of time the network is able to find valid path between a pair of nodes, averaged over all node pairs. Therefore, $\mathbb{P} = \{P_1\} = \{\text{path reliability}\}$. Secondly, the operational metrics at this level are nothing but the service parameters from the layer below. Therefore, $\mathbb{N} = \{N_1, N_2\} = \{\text{LC size, link durations}\}$.

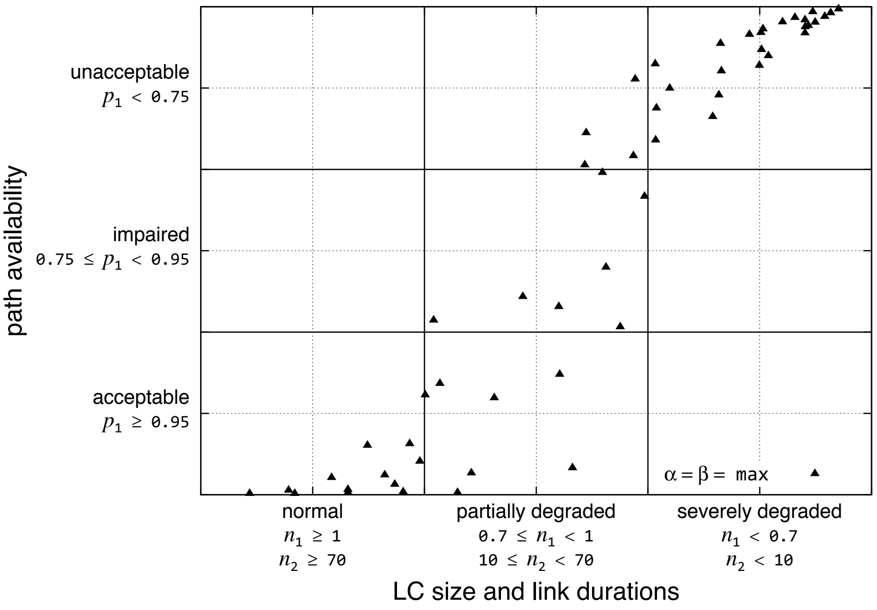

Figure 6.10: Routing state space with OLSR

Based on the simulation states we compute the state transitions using the 3×3 regions modeled as three states each in the operational space and service space. We conducted simulations using two different routing protocols: OLSR and DSDV. Figure 6.10 shows the state space for the OLSR routing protocol when using $\alpha = \beta = \max$, meaning that logical AND is being used to derive $x$ values from the two operational metrics. For the same values of $\alpha$ and $\beta$, we show the state space for DSDV protocol in Figure 6.11. Comparing these two protocols (as in Figure 6.12), we see that OLSR is more resilient to perturbations in the normal operating conditions compared to DSDV. While the absolute location of points varies with different values of $\alpha$, $\beta$ as shown in Figure 6.13, in generally OLSR has a better resilience profile than DSDV.

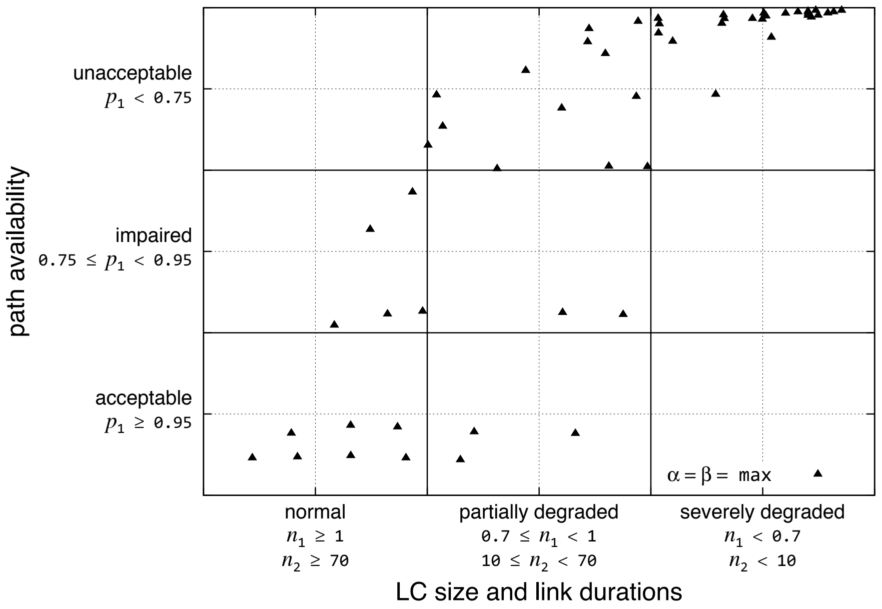

Figure 6.11: Routing state space with DSDV

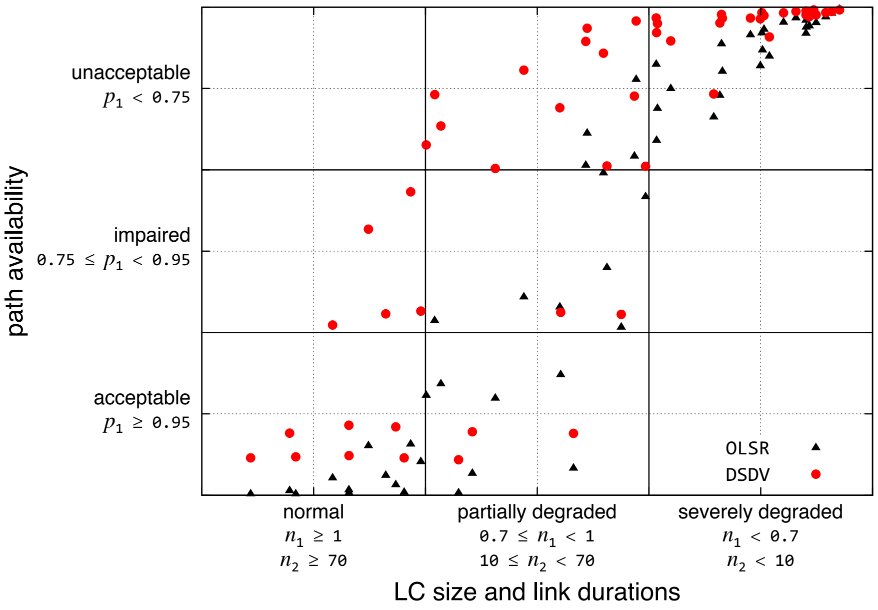

Figure 6.12: Comparing the resilience of OLSR and DSDV with $\alpha = \text{max}$

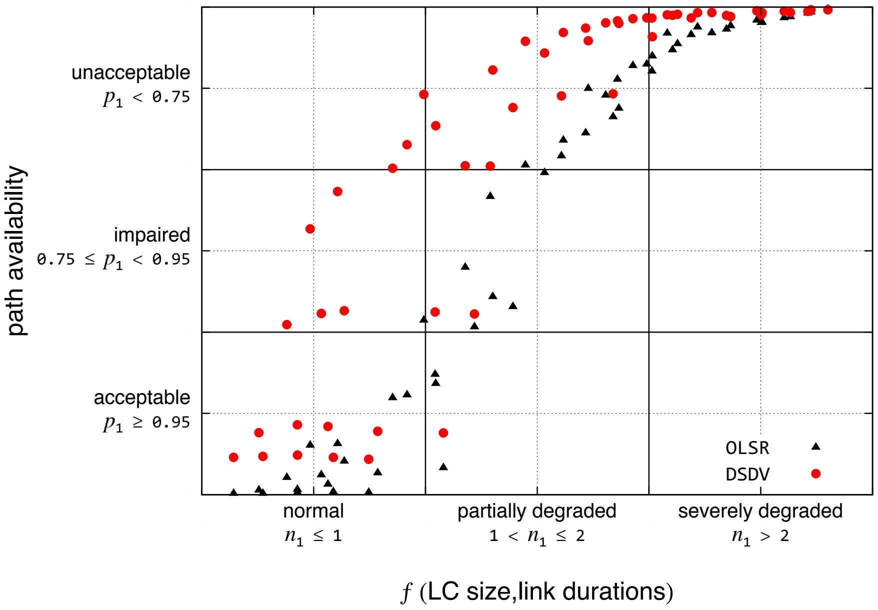

Figure 6.13: Comparing the resilience of OLSR and DSDV with $\alpha = 0.35$

### 6.3.1 Impact of Objective Function Parameters

In order to determine the impact of the objective function in the analysis of the resilience, we explore the full range of $\alpha$ and $\beta$ in this section. Figure 6.14 shows the state space plot when the value if $\alpha$ is varied from 0 to 1 in increments of 0.01. Note that since there is only one service metric, $\beta = 1$ for all runs.

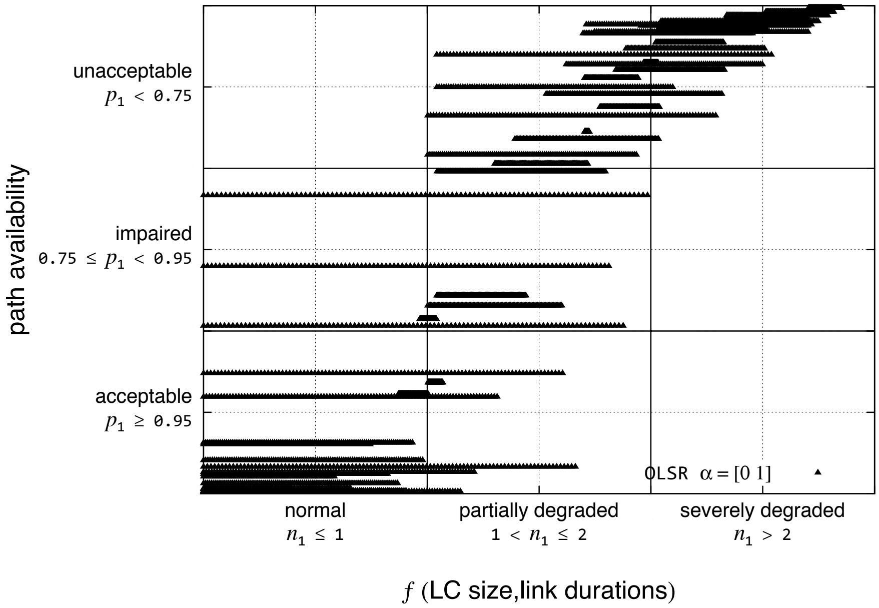

Figure 6.14: Impact of $\alpha$ on OLSR state space

From this Figure, we observe that if we probe the entire space of the objective function, we get a *envelope of the resilience.* In other words, we get an empirical bound on the resilience. Similarly, Figure 6.15 shows the DSDV states when varying $\alpha$ over the range of 0 to 1. Comparing the two in

Figure 6.16, we see that the envelope of the DSDV tends to lean more towards the impaired and unacceptable service region when compared to OLSR.

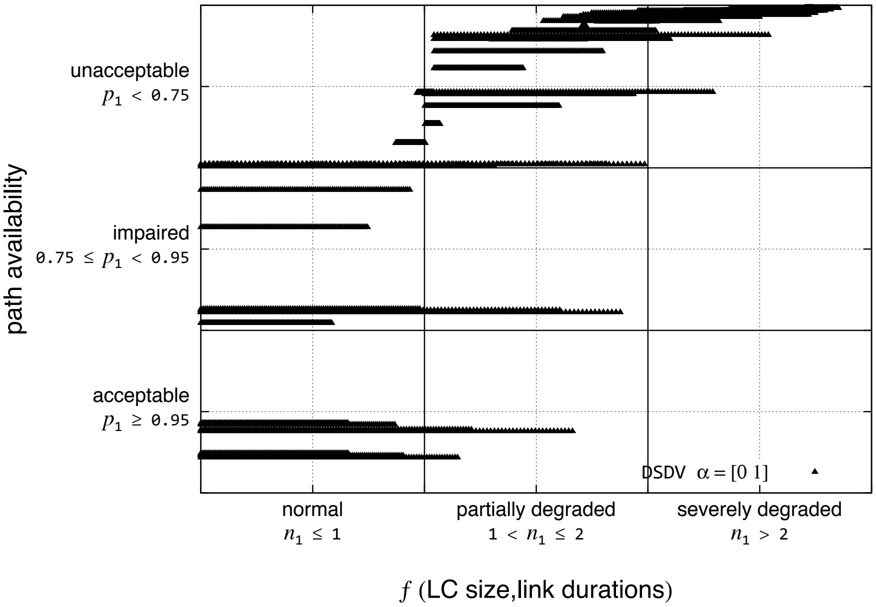

Figure 6.15: Impact of $\alpha$ on DSDV state space

## 6.4 Resilience at Transport – Application

In this section, we evaluate the resilience of the transport protocol under the presence of challenges. The service provided by the transport protocol to the application is end-to-end data transfer in the presence of perturbation in the paths provided by the underlying transport layer. The service parameters at this level are the packet delivery ratio (PDR) and end-to-end delay. Therefore, $\mathbb{P} = \{P_1, P_2\} = \{\text{PDR, delay}\}$. In addition to the the service parameters from the layer below, the operational metrics at this level include the relative

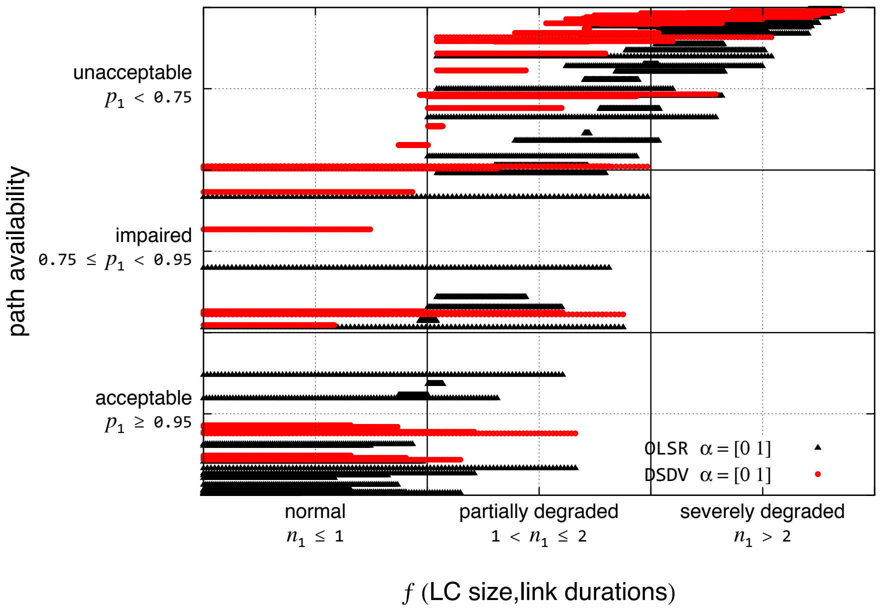

Figure 6.16: Impact of $\alpha$ on routing state space

traffic load that is dependent on the rate of the transport protocol. Relative load in this example is arbitrarily calculated as 1 for a sending rate of 0.54 Mb/s. Therefore, $\mathbb{N} = \{N_1, N_2\} = \{\text{path availability, relative load}\}$.

Using the state computation procedure detailed in Section 6.2.2, we plot the state transitions of the UDP protocol on a 3 × 3 region for two different data sets in Figure 6.17 and Figure 6.18. Looking at all the state instances, we clearly observe a specific pattern, an envelope that characterizes the resilience of the protocol. We observe that the service degrades almost linearly with respect to degradations in the operational state. Furthermore, the service sharply declines as the operations become severely degraded

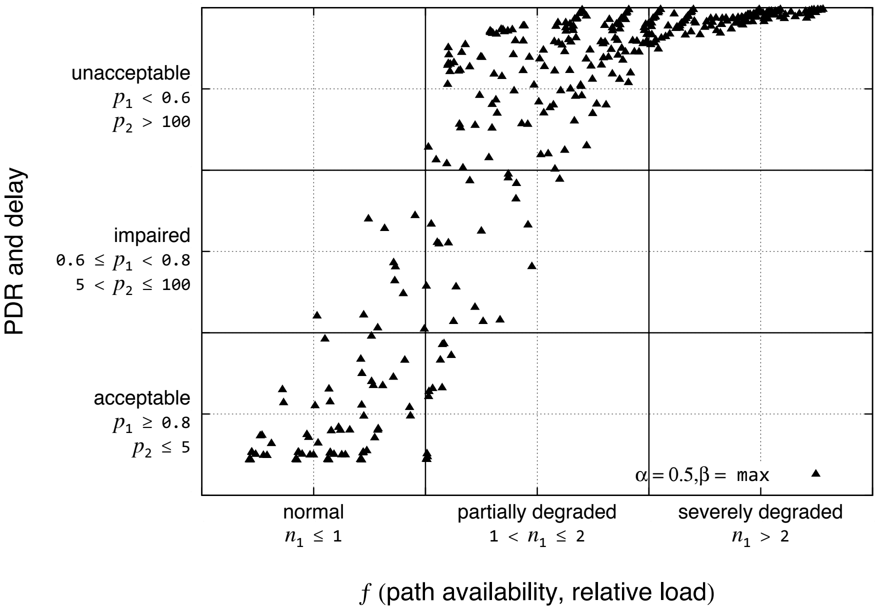

Figure 6.17: Transport state space: UDP resilience, run 1

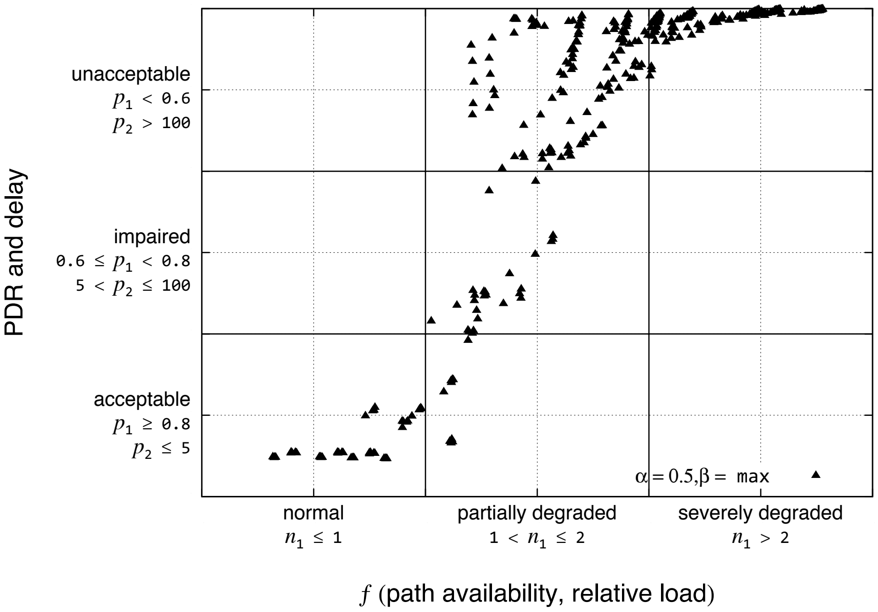

Figure 6.18: Transport state space: UDP resilience, run 2

## 6.5 Resilient Routing in WDTNs

In this section we consider several mechanisms as dictated by the principles presented in Section 3.3 to improve the resilience of weather disruption tolerant networks (WDTNs) which employ millimeter-wave links. In the previous sections of this chapter, we evaluated the resilience of some of the ISP-level topologies against random link failures. However, in the case of millimeter-wave mesh network, the main challenge observed is due to weather disruptions results in correlated link failures. Furthermore, due to the nature of these transmissions in high frequency bands, the physical layer can only provide a certain level of resilience in terms of link robustness. The topology of mesh networks is fixed: they form a basic grid network, therefore the resilience of the topology sublayer is also fixed. This implies that additional resilience against failures that penetrate through the defenses of physical and topological levels must be addressed at path routing level. In Section 5.3.1, we presented multipath as a mechanism to improve the resilience of the routing in the presence of perturbations to the network topology. In this section, we present predictive weather-assisted routing to enhance the resilience of millimeter-wave mesh network in the presence of environmental challenges. Specifically, we apply the principles of translucency [P9], diversity [P12], self-organizing and autonomic behavior [P13], and adaptability [P18] to design a resilient routing algorithm. Furthermore, we evaluate the resilience of the proposed algorithm to weather disruptions using the metrics framework.

While we present the resilient routing mechanisms in this section, a rigorous approach to resilient millimeter-wave mesh networks based on experimentation, modeling, and simulation is given in Appendix A.

### 6.5.1 Millimeter-Wave Mesh Networks

With increasing demand for high-bandwidth data access to end users, especially in metropolitan areas where the infrastructure cost of laying in new fiber is prohibitive, wireless broadband access technologies provide a viable alternative. One such alternative for 3G and potentially 4G service provider networks is the millimeter-wave mesh network that operates in the 71 – 86 GHz frequency band and have been proposed as a cost-effective high-speed alternative for fixed wireless mesh networks [72–74]. Existing commercial radios in this band can deliver data rates as high as 1 Gb/s with the potential for higher speeds on the order of 10 Gb/s using advanced modulation schemes [75]. Additionally, line-of-sight beams enable spatial reuse with mesh topologies.

However, their frequency of operation is highly susceptible to weather disruptions. In particular, millimeter-wave transmissions suffer heavy attenuation due to precipitation [72, 76–78]. As a result, link availability and reliability is significantly impaired during rain storms [72]. We conducted experiments to quantify the impact of weather disruption on millimeter-wave links. These experiments are presented in Appendix A. Based on the data collected, an empirical analysis of millimeter-wave links and mesh topologies using these

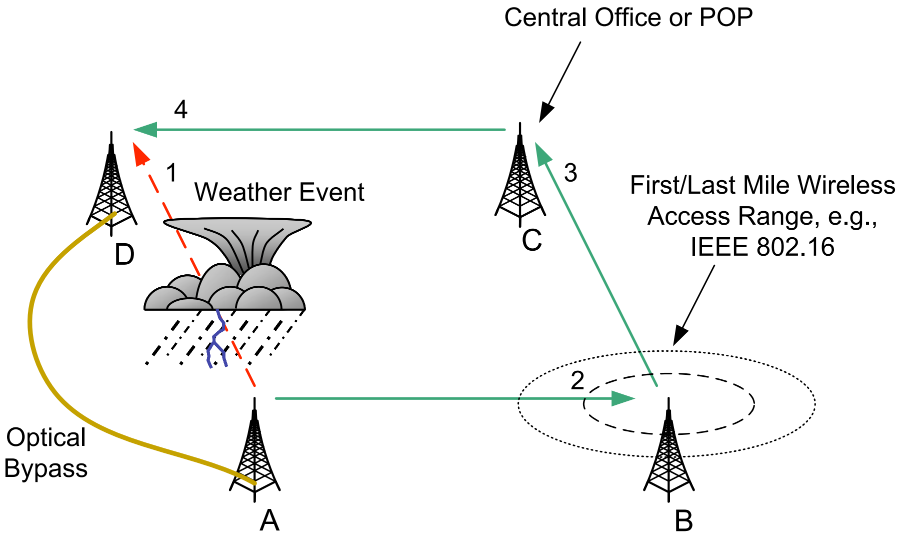

Figure 6.19: Schematic of an MWMN with weather system.

links is presented in Appendix B. In order to provide dependable paths in such a network, it is essential to design and engineer disruption tolerance into the network. In a mesh topology, as shown in Figure 6.19, this can be achieved through adaptive routing at the network layer. The following sections present a network-layer approach to overcome link instability by routing around failures within a mesh topology.

### 6.5.2 Mesh Routing Algorithms

This section presents two routing mechanisms (proactive and predictive) that exploit cross-layer mechanisms to leverage physical-layer information at the network layer. First, we present XL-OSPF, a cross-layered version of the well-known OSPF [79] routing protocol. This proactive approach uses link-cost metrics derived from physical-layer information to maximize the performance of OSPF in the presence of degraded links without the usual penalty of frame loss detection. Secondly, we present a *predictive weather assisted routing protocol* (P-WARP), that utilises short-term weather forecasts to reroute *ahead* of an impending link failure. We evaluate the performance of both approaches in terms of efficiency and survivability under weather disruptions.

### 6.5.3 XL-OSF: Cross-Layered OSPF

A cross-layer approach to link metrics could significantly improve the performance of dynamic link state algorithms [80, 81]. We choose OSPF (Open Shortest-Path First) [79] as the link-state routing protocol due to its wide deployment, use in research, and applicability to fixed networks. OSPF relies on two basic mechanisms to determine link state. One is the link state advertisements (LSAs) generated by each node that carry the status of all its links along with their costs. These are flooded throughout the network. Secondly, hello packets are used to determine if the link to a given neighbor is still alive. A *dead interval* based on the *hello interval* is used to detect dead links. The routes are proactively updated after the LSAs propagate through the network. With rapidly varying link quality, the only mechanism through which OSPF can detect link degradations is when four consecutive hello packets are dropped, in its default configuration. Since the size of hello packets is much smaller than data packets, a BER that results in four consecutive hello drops will correspond to a significantly higher data packet drop rate.

The first mechanism that can be used to improve the performance of OSPF is a cost metric that is proportional to the bit error rate of the link. However, this is a difficult proposition given the lack of information exchange between the physical (and MAC) layer that sees the actual packet losses and the network layer which determines the routes. Several mechanisms have been proposed in the literature [82, 83] that use in-band (packet header) or out-of-band (probe packets) signaling to determine the actual packet error rate.

For the purpose of analysis, we assume such a mechanism that informs the end hosts of the effective packet error rate. We define a cost metric that is proportional the the effective packet error rate. Assuming uniform distribution of the bit errors, the cost of a link between two nodes $i$ and $j$ is calculated as:

$$C_{ij} = P \times \mathrm{BER}_{i,j} \times \gamma, \tag{6.1}$$

where $P$ is the average packet size on the network, $\mathrm{BER}_{ij}$ is the bit error rate observed on the link, and $\gamma$ is the scale factor. The scale factor determines the sensitivity of the link cost with respect to change in BER and is set to 1000 in our simulations. A $\mathrm{BER}_{\mathrm{thresh}}$ of $10^{-8}$ is used to define the minimum observable change in BER. Further, hysteresis is used with a $H_{\mathrm{thresh}}$ of 10% to avoid excessive route flaps in the network. Finally, the value of cost is bounded in the range of $[1, 1000]$ which determines the maximum number of hops a packet can traverse in order to avoid an error-prone or lossy link. Since, the primary objective in the MWMN is to avoid disrupted links at all cost, we set this range to 1000. The performance of the modified XL-OSPF with this cost metric is discussed in Appendix C.

Even with the error based cost metrics, OSPF remains a proactive protocol that requires a finite amount of time before it adapts to changes in link state. If the application or service demands a highly-reliable service, proactive protocols must have a *very* short update interval on the order of milliseconds. But this adds an unacceptable level of overhead, even for broadband networks. In the following section, we discuss a predictive routing scheme that is intended to overcome this problem.

### 6.5.4 P-WARP: Predictive Weather-Assisted Routing Protocol

As discussed above, proactive algorithms may not be able to meet stringent service requirements (e.g. 50-ms restoration for circuit emulation and CBR traffic) in MWMNs during weather disruptions. Furthermore, it is difficult to measure effective BER or FER at end hosts without an explicit signaling mechanism. In this section, we investigate the use of information *external* to the network in order to predict the state of links over the next time epoch or several epochs ahead.

The proposed predictive routing algorithm is a link-state algorithm that utilises weather radar data to forecast the *future* condition of the link. In contrast to the XL-OSPF discussed above, the primary difference is the mechanism through which the link costs are obtained. While XL-OSPF depends on BER measurement from errored packets, we propose P-WARP (predictive weather-assisted routing protocol), in which BER of each link is calculated from weather radar reflectivity data modeled in real-time using the methodology discussed in Section A.3. This processing is done at a either a single *core node* or a small subset of core nodes which are connected to the external Internet and are capable of receiving weather radar data. In either case, multiple-path connectivity into the mesh is necessary for high-availability of the radar data[1]. The topology and physical locations of the fixed network nodes are pre-programmed in to the software module that performs the link BER calculation as well as the PER (packet error rate) for a predefined average-packet size. The cost metric for individual links is based on the effective link BER similar to XL-OSPF. While XL-OSPF *proactively* derives costs based on measured BER and propagates updates with conventional LSAs, P-WARP uses short term weather forecast to *predict* link costs. Thus the link cost is calculated using Equation (6.1) with the same thresholds described in Section 6.5.3.

### Link Status Updates and Route Computation

The weather-based link-status *updates* (WLSUs) in P-WARP are slightly different from the conventional link-state advertisements (LSAs) of OSPF.

[1] Note that this connectivity can be provided by low rate links such as lower frequency radio links that are much susceptible to weather

WLSUs are generated from the core nodes and contain the costs of *all* links in the network based on their predicted quality. These weather-based updates are flooded throughout the network. When an individual node receives a WLSU, it recomputes routes using the shortest-path first algorithm. However, unlike OSPF, individual nodes do not generate separate LSAs for the links to their neighbors. This approach significantly reduces the protocol overhead because only one update is generated for all the links and updates are generated only when a change in one or more link costs is predicted. Thus, the network reroutes traffic *ahead* of the incoming storm thereby minimizing, and perhaps eliminating packet loss. It is important to note that while we are using weather predictions to alter the network state, the time scale of the weather predictions are on the order of tens of seconds, a short but accurate interval in weather time, however a very long interval in network time, sufficient for predictive routing.

### Route Sensitivity

It is clearly evident that the effective BER on each link will vary continuously over the duration of the storm. In order to avoid route flaps and false alarms, we use thresholds along with hysteresis. A minimum noticeable change $\text{BER}_{\text{thresh}}$ is defined below which all BER changes are ignored. Further, a hysteresis percentage $H_{\text{thresh}}$ determines the minimum change in the cost of a link for an update to be generated. The various steps in the operation of the P-WARP are enumerated in Algorithm 1. Table 6.2 shows a brief comparison of the proposed routing protocols XL-OSPF and P-WARP with a standard OSPF implementation.

Table 6.2: Comparison of routing protocols

| Metric | Std. OSPF | XL-OSPF | P-WARP |
|---|---|---|---|
| cross-layered | no | yes | yes |
| link cost $C_{ij}$ | typically 1 | $\propto$ $\text{BER}_{ij}$ | $\propto$ $\text{BER}_{ij}$ |
| link failure detection | 40 s | 10 s | 0 s |
| data rerouting | proactive | fast proactive | predictive |
| control packet | LSA | LSA | WLSU |
| update rate | periodic 10s | periodic 10s | aperiodic |

### 6.5.5 Resilience Analysis

In order to evaluate the resilience of XL-OSPF and P-WARP, we conducted simulations using actual storm patterns observed during our experimentation phase. Details of experimentation setup, field data collection and simulations are presented in Appendix A. We evaluate the resilience of the millimeter-wave network at the routing – transport layer. At this layer boundary, the service the end-to-end paths provided by the routing protocol. Therefore, we characterize this service using one parameter: *path availability*, which is defined as percentage of time the network is able to find valid path between a pair of nodes, averaged over all node pairs. Therefore, a service state is represented as: $\mathbb{P} = \{P_1\} = \{\text{path reliability}\}$. The operational state is characterized by the underlying topology and the perturbations due to

**Algorithm 1** Predictive Weather-Assisted Routing

```
Step 1: Generate WLSU at the central node
Input: vertices $V$, edges $E$, radar reflectivity data RRD, geographic node
    positions, forecast window $\delta_t$
At time $t$, receive predicted weather for time epoch $t + \delta_t$
update $\leftarrow 0$
for all $(i, j) \in E$ do
    $\text{GSM}(t + \delta_t) \leftarrow f(\text{RRD}(t + \delta_t)$ {calculate geometric storm model}
    $A_{ij}(t + \delta_t) \leftarrow g(\text{GSM}(t + \delta_t)$ {calculate link attenuation based on link
    and storm overlap}
    $\text{BER}_{ij}(t + \delta_t) \leftarrow h(A_{ij}(t + \delta_t))$ {calculate link BER as a function of its
    attenuation and radio characteristics}
    if $\text{BER}_{ij}(t + \delta_t) - \text{BER}_{ij}(t) > \text{BER}_{\text{thresh}}$ then
        $C_{ij}(t + \delta_t) \leftarrow \text{BER}_{ij} \times P \times \gamma$ {calculate predicted link cost}
        if $C_{ij}(t + \delta_t) - C_{ij}(t) > H_{\text{thresh}}$ then
            update $\leftarrow 1$
        end if
    end if
end for
if update $= 1$ then
    generate $\text{WLSU}(t + \delta_t)$ consisting of $C_{ij}(t + \delta_t),\ \forall (i, j) \in E$
end if
Step 2: Recompute routes at each node
Input: $\text{WLSU}(t + \delta_t)$
update local cost matrix $C$ such that $C_{ij} \leftarrow C_{ij}(t + \delta_t),\ \forall (i, j) \in E$
compute least cost paths to obtain $R(t + \delta_t)$
schedule $R \leftarrow R(t + \delta_t)$ at time $t = t + \delta_t$
```

weather disruptions and is quantified by the packet error rate averaged over all network links. Therefore, $\mathbb{N} = \{N_1\} = \{\text{averaged PER}\}$.

Table 6.3: Operational regions: routing in millimeter-wave networks

| Region | Operational metrics<br>averaged PER $n_1$ |
|---|---|
| Normal | $n_1 \leq 0.05$ |
| Partially degraded | $0.05 < n_1 \leq 0.50$ |
| Severely degraded | $n_1 > 0.50$ |

In order to quantify resilience, we evaluate the state transitions over the $3 \times 3$ state space when the network is subjected to rain storms. The regions for operational metrics (avg PER) are defined in Table 6.3 and service parameters are defined in Table 6.4.

Table 6.4: Service regions: routing in millimeter-wave networks

| Region | Service Parameters<br>path availability $p_1$ |
|---|---|
| Acceptable | $p_1 \geq 0.99$ |
| Impaired | $0.75 \leq p_1 < 0.99$ |
| Unacceptable | $p_1 < 0.75$ |

Figure 6.20 compares the resilience of the P-WARP against traditional OSPF and static routing when the network is subjected to the first observed storm. For details of this storm, please refer to Appendix A.2. As shown in figure,

it is observed that P-WARP is significantly more resilient when compared to the other two algorithms. With P-WARP, the service remains acceptable under normal operating conditions as well as to some extent in partially degraded conditions. Furthermore, as the network degraded severely, P-WARP outperforms the other two in terms of resilience. Given that this is a millimeter-wave network which is significantly affected by rain as well as humidity, the normal operating conditions are defined to be the case with average PER up to 5%. This is because, even under normal operations, the link observes non-negligible errors. The path availability with both Static and OSPF routing degrades sharply with increasing PER.

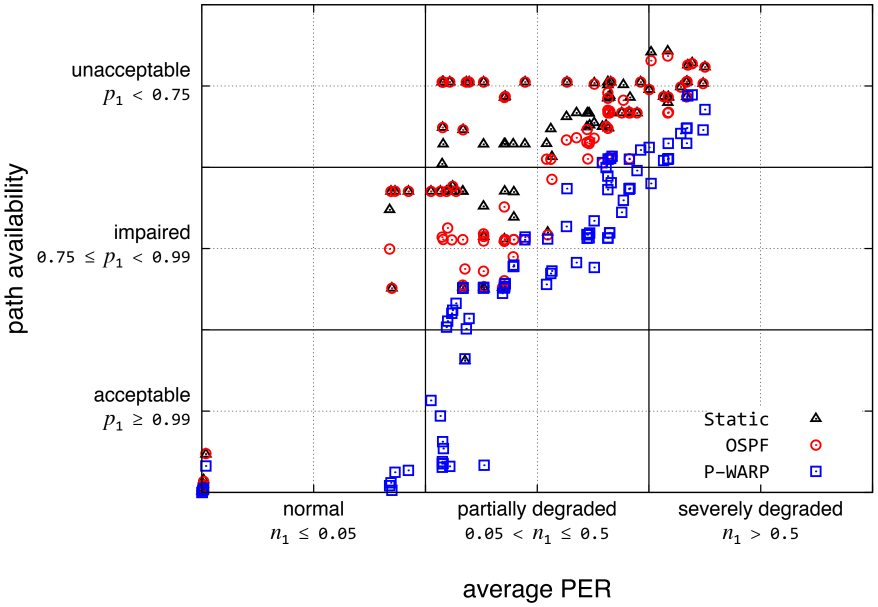

Figure 6.20: Comparing routing resilience in presence of first rain storm

Secondly, we compare the resilience of XL-OSPF versus P-WARP for the same storm. As shown in Figure 6.21 We observe that the state space transitions for these two mechanisms are very similar, with P-WARP being slightly more resilient. This is due to the predictive nature of P-WARP wherein the network reroutes around links before they degrade using weather prediction. Contrast this to of XL-OSPF, for which the algorithm actively measures the observed PER and quickly routes around links that are experiencing weather disruptions. However, for a brief duration between detection and rerouting, several paths are unavailable thereby leading to reduced path availability.

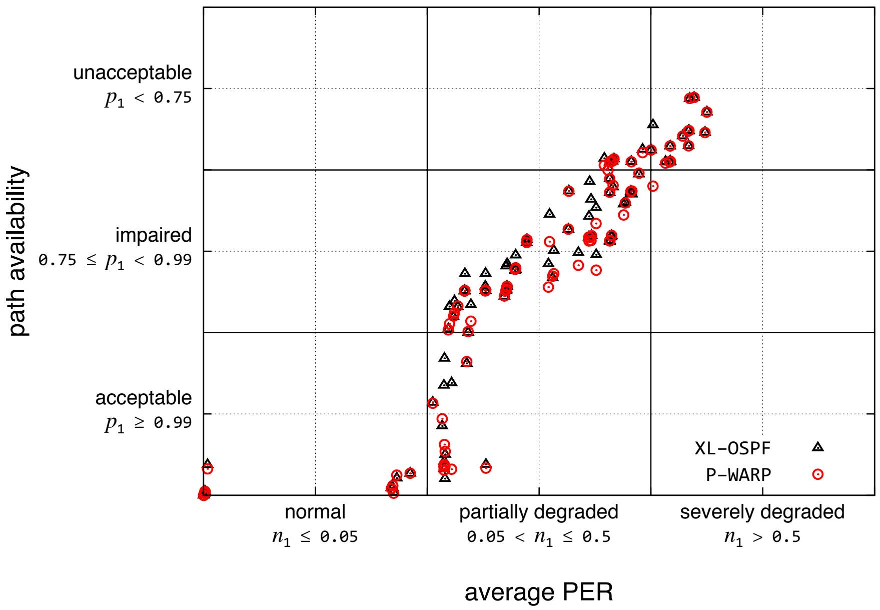

Figure 6.21: Comparing XL-OSPF vs P-WARP, first rain storm

Figures 6.22 and 6.23 show the resilience comparison of various protocols when the network is subjected to a second rain storm. As compared to the first storm, this storm was more intense, covered a larger geographical

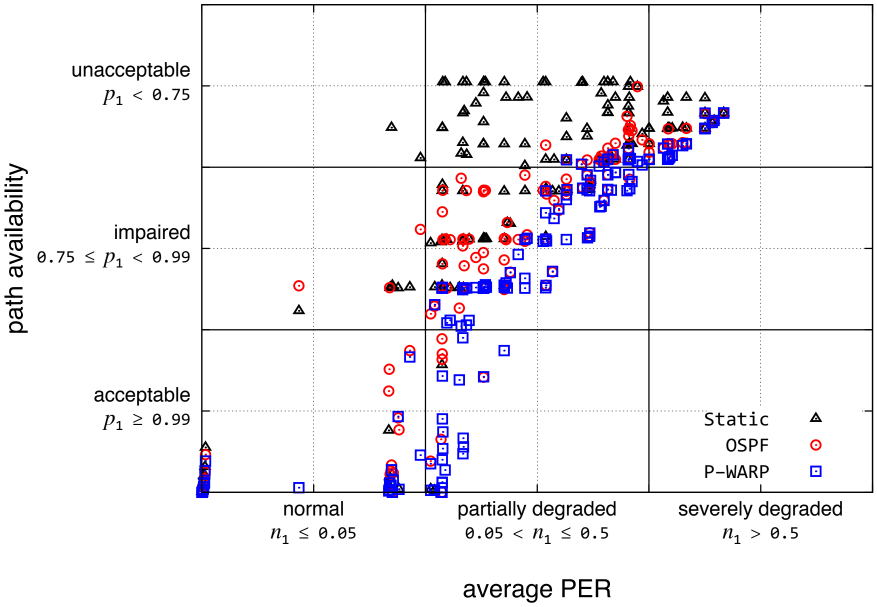

Figure 6.22: Comparing routing resilience in presence of second rain storm

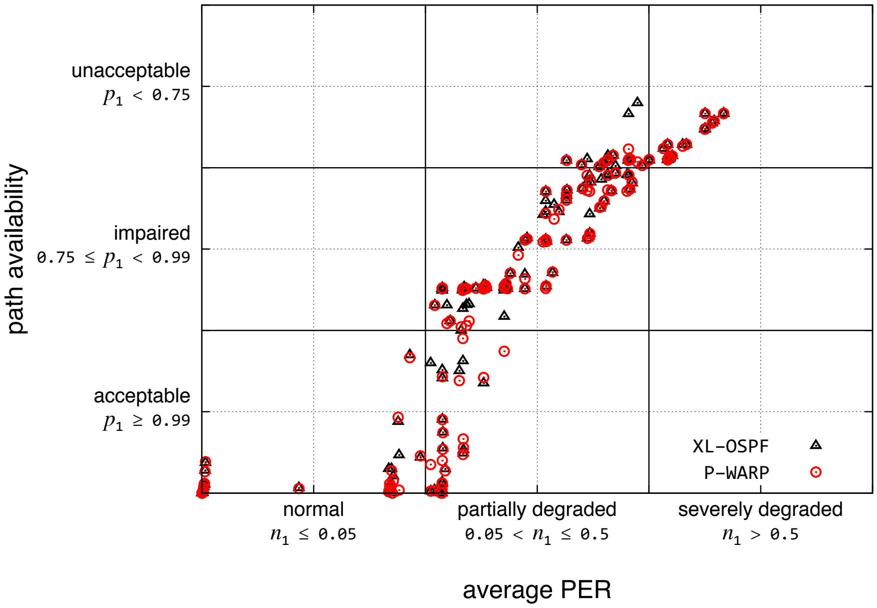

Figure 6.23: Comparing XL-OSPF vs P-WARP, second rain storm

area and lasted for a longer duration (Appendix A.2). We again calculated the path availability of each protocol while the storm system passed over the simulated network. The results indicate that P-WARP and XL-OSPF show similar resilience while outperforming traditional protocols such as static and OSPF.

## 6.6 Resilient Routing in Airborne Networks

In the Section 6.5, we considered resilient routing in millimeter-wave mesh networks in which the challenge was weather disruptions. We presented new mechanisms that are resilient to perturbations in link quality. This section considers yet another scenario, aeronautical telemetry networks in which the primary challenge is the highly dynamic nature of the network arising from the very high node speeds [84]. The objective is to develop resilient routing mechanisms that can provide acceptable service in the presence of high node mobility or extremely short contact durations. Again, we apply the principles of Section 3.3 to design resilient routing for highly dynamic mobile ad hoc networks.

### 6.6.1 Aeronautical Telemetry Networks

Aeronautical telemetry networks, as the one discussed in [85] primarily consist of airborne nodes traveling at extremely high speeds. When the node speeds are in excess of Mach 3, the contract durations can be extremely short leading to a highly dynamic network with constantly changing topologies. Other challenges to communication in this environment includes bandwidth constraints due to the limited spectrum availability, and limited transmission range due to power and weight constraints. For a typical scenario [84], the contact duration can be less than few seconds. Furthermore, due to the sparse nature of the telemetry network, the network is not always connected resulting in intermittent connectivity.

Given the various challenges, resilient routing in this environment demands exploitation of cross-layer information, especially the geo-location and trajectory of nodes to assist in routing. However, most of the existing routing protocols are not suitable for this environment because of slow route convergence and lack of cross-layering support. On the other hand, cross-layer aware protocols such as location aware and beaconless routing are not designed for extremely high speed operation.

### 6.6.2 AeroRP: Location Aware Routing

The aeronautical routing protocol (AeroRP) [86] is a location-aware highly adaptive routing algorithm that is specifically designed to be resilient to speed induced disruptions. It utilizes a number of resilience mechanisms as discussed below:

1. **Proactive behavior:** AeroRP is a fundamentally proactive routing protocol, but with limited updates thereby lowering protocol overheard.

2. **Exploits cross-layer controls:** AeroRP is designed to exploit the explicit cross-layering support provided by AeroNP and the geographic node location and trajectory information available at nodes.

3. **Per-hop behavior:** Unlike existing protocols, AeroRP forwards data per-hop based on partial local information and routes thereby avoiding the necessity for global convergence, making it especially suitable for highly-dynamic environments.

4. **Multi-modal:** Military applications present a high level of variation in their operational parameters. For example, based on the security requirements of the test application, the geolocation of the nodes may or may not be available. In order to support these dynamics in operation, policies, and constraints, AeroRP provides multiple modes of operation.

The basic operation of AeroRP consists of two phases. In the first phase, each AN learns and makes a list of available neighbors at any given point in time. It utilizes a number of different mechanisms to facilitate neighbor discovery, including snooping, hello beacons, and periodic updates. The second phase of the algorithm is to find the appropriate next hop to forward the data packets. In order to forward the packets toward a specific destination, additional information such as location data or route updates is required. AeroRP utilizes the node location and trajectory, when available, in making forwarding decisions. Based on the set of discovered neighbors, each node decides the best course for each packet such that with every transmission the packet ends up being physically closer to its destination. For each of these two phases, AeroRP defines a number of different mechanisms to choose from. The particular choice of mechanism to be used is dependent upon the mode of operation because the mode of operation determines which of the mechanisms are feasible given the scenario specific restrictions (e.g. security restrictions). Simulations results have shown that AeroRP is significantly more resilient when compared to other protocols [2].

## 6.7 Summary

In this chapter, we applied the proposed metrics framework to MANETs in order to evaluate multilevel resilience. We presented how state space computation can be performed with the help of an objective function. The impact of the parameters of the objective function was also explored. The resilience analysis was conducted for 4 successive levels or 3 level boundaries: topology → routing → transport → application. We quantified the resilience of the topology and compared the OLSR and DSDV in terms of their ability to survive perturbations in the topology. Secondly, applying the principles of ResiliNets architecture, we developed resilient routing mechanisms: XL-OSPF and P-WARP for millimeter-wave networks and AeroRP for aeronautical telemetry networks. The resilience of the proposed mechanisms was quantified using the metrics framework.

[2] Details of the each phase and the simulation results are presented in [86].

Page left intentionally blank.

# Chapter 7

# Conclusions and Future Work

This dissertation presents a comprehensive approach to resilient networks with the focus on quantitative evaluation of resilience. To this end, we have introduced a new metrics framework and resilience strategy both of which were applied to a number of scenarios. This chapter presents a summary of the dissertation, major contributions, conclusions drawn, and future work.

## 7.1 Conclusions

Modern society depends on the networks in general and the Internet in particular for their critical services. These services range from every day convenience to economic stability and national security. Besides this dependence on legacy networks, there are a growing number of services that are based on emerging networks such as wireless, mobile ad hoc, and sensor networks. The challenges to the normal operation of these networks and services come from a variety of sources. While some are inherent to the environment (e.g., wireless) others are generated from forces external to the network (e.g. natural disasters). While some challenges are on small scale (e.g., link cuts), others happen on massive scale (e.g., power failures). And while some are unintentional (e.g., component failures), others are intentional attacks (e.g., viruses). We termed the ability of a network to provide and maintain an acceptable level of service in the face of these challenges as *resilience*. Similar measures have been studied in literature in the fields of fault tolerance, disruption tolerance, and survivability. Some of the well known measures include reliability, availability, dependability, and performability. Existing methodologies to evaluate resilience-like properties of the system are often domain, failure, and level specific. While these provide valuable information regarding systems behavior, they do not provide a comprehensive view of resilience at any arbitrary level.

The contributions of this dissertation are three-fold: first, motivated by the multilevel resilience principle and the resilience strategy $\mathrm{D}^2\mathrm{R}^2 + \mathrm{DR}$, we present a multilevel metrics framework. Secondly, we present a comprehensive rigorous treatise towards network topologies including a model to generate realistic topologies and a rigorous evaluation of resilience using the proposed framework. Thirdly, we present unique resilience mechanisms at the routing layer including path diversification, predictive weather-assisted routing, and location-aware highly adaptive routing and characterize their resilience in the presence of perturbations to the underlying network.

In this dissertation, we define resilience at a given service boundary, generally referred to as a level boundary $B_{i,j}$. This boundary can be defined between any two successive or arbitrary layers. The communication network at any given level is viewed in two dimensions: operational space and service space. The operational space $\mathbb{N}$ represents the state of all elements of the network below the service boundary and are quantified using a set of metrics intuitively termed as *operational metrics*. The other dimension $\mathbb{P}$ represents the service that is provided to the network above this boundary and is quantified by a set of metrics termed as *service parameters*. The state of the network is then represented by the tuple of operational metrics and service parameters. Furthermore, we formulate that the challenges to the network manifest as perturbations to the operational conditions of the network leading to possible (but not necessarily) service degradations. Given that networks cannot be infinitely resilient, a "good" network is one which degrades gracefully in the presence of increasing challenges. Therefore, we evaluate the resilience $\mathbb{R}_{ij}$ as a function of these transitions in the network state-space and quantify as the area under the curve. At any given boundary, the network is said to be resilient if it prevents degradation in the operational condition from leading to degradations in service. Lastly, the proposed approach is inherently multilevel. The objective is to quantify the network resilience at a given layer so that resilience mechanisms can be implemented and subsequently evaluated. This approach is based on the ResiliNets architecture and its goal to provide the best resilience at any level given practical constraints. To this end, we show the service parameters at a given level map to the operational metrics of the level above, thereby providing a seamless method to conduct multilevel resilience analysis. We demonstrated this approach with numerical examples at the link–topology, topology–routing, routing–transport, and transport–application boundaries.

The second major contribution of this dissertation is a model to generate realistic network topologies based on practical constraints as well as evaluation of topological resilience using the proposed framework. Network topology research plays a key role in evaluation of various protocols and mechanisms. However, the existing body of research primarily focusses on generating logical topologies based on link connectivity models. While these are crucial to evaluate the performance of protocols such as BGP that operate over the logical topologies, they do not accurately model the resilience of the actual topologies. For this purpose, we developed a location and cost constrained topology generation model that focusses on first realistic node location and secondly cost constrained link connectivity. We evaluated the topological resilience against random link failures for ISP-level topologies of AT&T, Sprint, and GÉANT. We define service at the topology–path routing boundary as the *topological connectivity* and presented an aggregate measure of resilience $\mathbb{R}$ over the range of [0 1] where, a value of 1 represents perfect resilience and a value of zero indicates no resilience. We also present a path diversification mechanism which exploits the diversity in the graph to improve path robustness. We show that by using multiple paths, the resilience at the routing–transport boundary can be increased to varying extent depending upon the number of paths used.

The third major contribution of this dissertation is a new resilience mechanism at the routing–transport layer to improve the resilience of millimeter-wave mesh networks in the presence of weather disruptions. We characterized the challenge posed to millimeter-wave mesh networks based on actual experimental data along with a respective geometric model. We proposed two resilient routing protocols XL-OSPF and P-WARP based on cross-layering between the physical and network layers. While XL-OSPF uses cost metrics proportional to bit errors in order to select optimal paths, P-WARP uses a predictive mechanism based on weather forecasts to find optimal paths thereby providing paths that are resilient to weather disruptions. Performance analysis has shown that the proposed mechanisms outperform traditional routing. Lastly, we developed a resilient routing protocol, AeroRP for aeronautical telemetry networks that exploits location and trajectory information to provide acceptable service in the presence of highly dynamic and intermittent connectivity.

## 7.2 Future Work

The metrics framework presented in this dissertation provides a versatile basis to evaluate resilience at any arbitrary boundary. We showed resilience analysis at specific boundaries in our analysis. Evaluating resilience at other levels over as well as over different network types can be done as a part of future work. Further research is needed to derive a single multilevel resilience metric $\mathcal{R} = f(\mathbb{R}_{ij})\ \forall ij$, if feasible. We presented simple aggregate measure based on the resilience states to get a single resilience metric. Additional methods could be explored to derive very specific metric from the state-transitions. The metrics framework can be applied both to steady state as well as transient analysis. We have considered cases for steady state resilience analysis. A transient analysis of networks can be conducted to determine how a network reacts to challenges in general and malicious attacks in particular. The instantaneous state transitions as the network experiences a challenge can quantify the resilience in terms of the $\mathrm{D}^2\mathrm{R}^2+\mathrm{DR}$ cycles. Specifically, the time taken to revert to original state and the path taken through the state space (say the area under the curve) determine the resilience of the network.

In this dissertation, we presented a location and cost constrained hierarchical model to generate realistic physical network topologies. The emphasis of the current model is on location of the nodes and the reproducing the hierarchy of the real networks. Further research can be conducted to determine realistic link models based on actual fiber paths when subjected to area-based challenges such as natural disasters [87]. Even though the logical topologies are heavily influenced by the overlay topologies and policy decisions, they are still shaped by the underlying physical topology that the presented model generates. Further research can be conducted to model logical topologies given an underlying physical topology.

# Appendix A

# Millimeter-Wave Experiments

In order to understand the impact of actual weather events on millimeter-wave transmission, two millimeter-wave links were deployed and test data was collected for a period of one year long. For the same time duration, radar reflectivity data of the region in which the links were deployed was also collected. This appendix describes the setup for the link measurements and the radar measurements.

## A.1 Link Measurements

To collect link-performance data, two millimeter wave radio links were deployed in Lawrence, Kansas. In order to observe actual signal attenuation and the resulting frame errors, we used one radio with no error correction and to evaluate the effectiveness of error correction coding in this specific case, we used a second radio with forward error correction (FEC). As shown

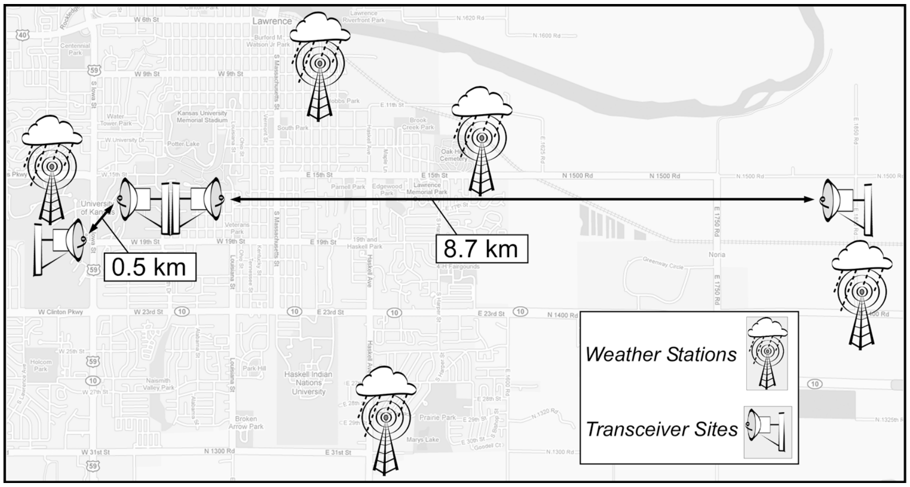

Figure A.1: Location of weather stations and radio links

in Figure 1, the transmission path was configured with an intermediate hop; however the primary performance effects arose from the 8.7 km link. To collect weather data for this research project, weather stations using Vaisala WXT510 [88] instruments were deployed at five locations, each equipped with sensors capable of recording temperature, humidity, wind, and most importantly, precipitation [88]. Figure A.1 shows the locations of the weather stations and the radio links. Since the purpose of this study is to evaluate real world scenarios, two off-the-shelf radio systems were chosen. Even though the two radios differ slightly in their design parameters, as shown in Table A.1, their sensitivity to weather-induced attenuation is similar because of the identical BER vs. SNR curves of the OOK (on-off keying) and BFSK (binary frequency shift keying) modulation schemes. Furthermore, one radio implements a Reed-Solomon (204,188) FEC while the other has no FEC.

Table A.1: Radio parameters

| Parameter | Without FEC | With FEC |
|---|---|---|
| TX / RX Frequency (FDD) | 73.5 GHz / 83.5 GHz | 72.5 GHz / 82.5 GHz |
| TX Power | 17 dBm | 17 dBm |
| Antenna Dia. / Gain / Beamwidth | 63 cm / 51 dBi / 0.4° | 31 cm / 43 dBi / 0.8° |
| RX Noise Figure | 6 dB | 7 dB |
| Modulation | OOK | BFSK |
| Data Rate | 1250 Mb/s | 1000 Mb/s |
| FEC Type | None | RS (204, 188) |

The objective is to measure the impact of low intensity precipitation on the radios as well as evaluate the effectiveness of FEC in overcoming bit errors.

The link was set up in a loop-back configuration and a GigE tester was used to measure FER along with a number of other performance metrics. To provide input for mesh networking and resilient routing algorithm studies, the link performance metric collected was the FER. Here the FER is the percentage of data frames that are lost during every 30 second interval. The tester transmits on average 3,348,000 512-byte data frames every 30 seconds (0.5 Gb/s). These transmissions, after being looped back at the other end of the line, are recovered at the transmitting end. The FER is then calculated as the percentage of frames lost or in error during the 30 second interval. Given the total number of frames transmitted, a FER lower than $10^{-7}$ is not observable. Since it takes 10 seconds to process and record this data, samples were collected every 40 seconds. All observations reported here are from 01 October 2007 through 30 September 2008. Over the course of a single day, around 14,400 (2,880 per station) weather samples and approximately 4,320 (2,160 per link) link performance samples were recorded. Later, we will present an analysis of the data collected to derive statistical conclusions regarding the impact of weather events on millimeter-wave links.

## A.2 Radar Measurements

In order to determine the impact of actual rain storms, we collected radar reflectivity data from the National Weather Service for the Midwest US. The effect of a weather disruption on millimeter-wave network mainly depends on two factors: rain rate and the geographic footprint of the storm [89]. Both of these parameters vary from one geographic region to the other. For example, the analysis of weather data from southern Great Plains region of the US [90] shows that approximately 78% of all storms are smaller than 25 km in diameter and account for only 1.0% of the precipitation. Furthermore, only 1% of the storms are large (over 40 km diameter) and account for 85% of precipitation. The remaining 20% are medium sized storms (20–40 km diameter) accounting for 14% of precipitation. We draw two conclusions [89] from this study: First, the majority of the storms are small enough for a metropolitan size mesh network (approx. 1000 km$^2$) to reroute traffic around the storm. Secondly, even moderate sized storms are likely to have small size heavy intensity regions since they don't account for a significant percentage of rainfall. This is consistent with the our measurements presented in Section B.1.1 (See Figure B.2). It is important to note that similar studies will be needed for other geographic regions with significantly different weather patterns, such as the Pacific Northwest US and Europe.

We selected specific weather-events to conduct a in-depth analysis on the impact of individual events on single links as well as mesh networks. In order to get a diverse set of weather patterns in this study, we specifically choose eight storms that were topologically different with significantly different characteristics, e.g. small and large cells, multiple cells on a front line, and intense front line. Their duration over a 1000 km$^2$ mesh varied from just under an hour to several hours[1]. Figures A.2 and A.3 each show an instance of two of the eight selected storms.[2,3] These match the previously described common Midwest US characteristic of distinct cells with high intensity.

## A.3 Modeling Rain Storms

In order to evaluate the effective attenuation on each link in a MWMN, we use the weather-radar echo intensity to determine the overlap of a storm cell with a given link. This is done through a geometric model as discussed below:

[1] MWMN grid Node 0 is geographically located at 38.8621N, 95.3793W

[2] Distribution 1 observed at 20:39:26 Z on 30 September 2008

[3] Distribution 2 observed at 05:04:11 Z on 22 April 2008

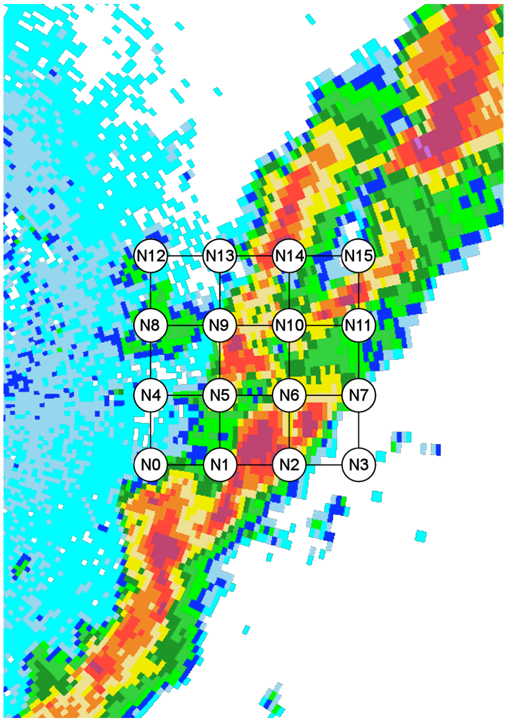

Figure A.2: Rain distribution 1 over MWMN links

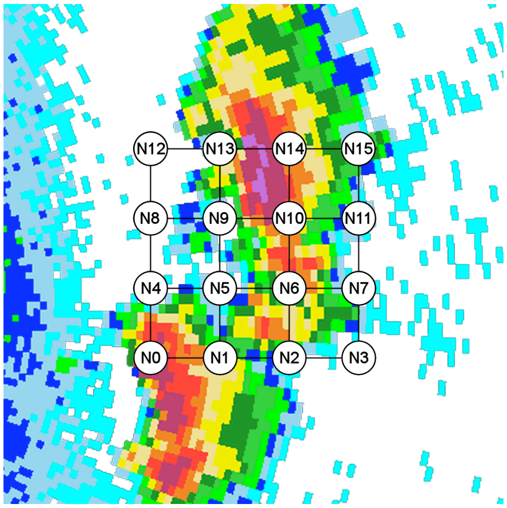

Figure A.3: Rain distribution 2 over MWMN links

## A.4 Geometric Storm Model (GSM)

In this model, each region in a radar reflectivity map shown in Figure A.4 is modeled as an ellipse ($R_1$–$R_5$) as shown in Figure A.5. While the resolution of the radar reflectivity differentiates many levels of precipitation, the proposed model uses only three levels indicating high (red), low (yellow), and little or no (green) precipitation so that the method remains tractable. The exact value of the rain rate corresponding to a given region depends on the location and the precipitation characteristics. For example, the simulations in our study use a value of 2mm/hr and 5 mm/hr to represent low and high intensity regions respectively. Note that there could be several regions of the same rain intensity in a given storm and the rain rates corresponding to each color vary depending upon the geographical location. Finally, individual links are modeled as line segments. The procedure to calculate the effective attenuation for a given storm as it passes over a fixed mesh network is discussed below.

For a given storm, we first generate the geometrical model of the storm before it enters the geographic region of the network. Secondly, as the storm moves across the network, we generate snap shots of the storm pattern (using ellipsoids) at different time intervals to capture both the regular progress as well as key points when the storm changes direction and the splitting and merging of cells occur[4]. Then, with the help of an interpolation program we

[4] This process currently relies to some extent on manual visualization of the radar data.

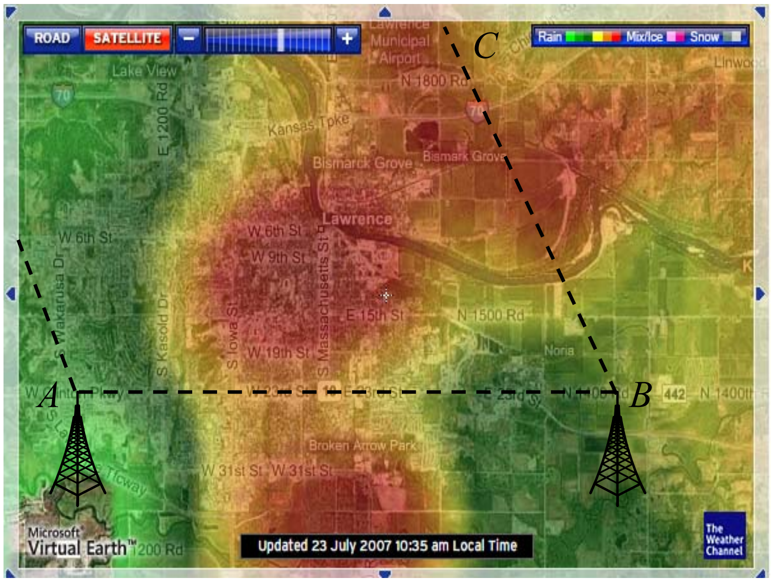

Figure A.4: Example radar image

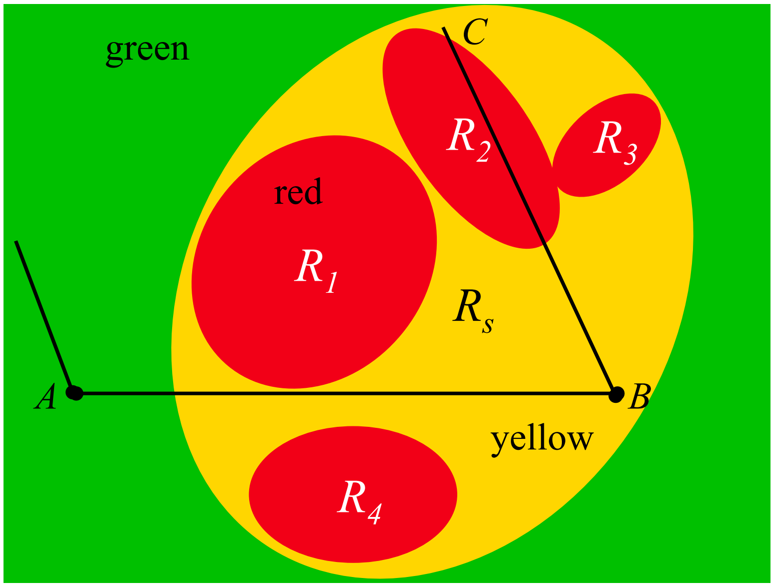

Figure A.5: Storm model corresponding to Fig. A.4

generated the movement of these ellipsoids over a fixed time step (e.g. one second).

Given that the position of storm ellipses at any time instant $t$ for the entire duration of the simulation, along with coordinates of the nodes and hence the links, we calculate the intersection between the links and the storm ellipses so as to determine the attenuation experienced be each link. Using simple geometry, we calculate the intersection of links with an ellipsoid to determine the length of the link $l_i$ affected by a given storm region $i$. Using ITU-R P.838-3 model [91], the attenuation of link due to this storm is calculated as $\gamma_R \times l_i$, where $\gamma_R$ is the rain rate associated with the storm region. The total attenuation on a given link is the summation of the attenuation resulting from all the individual regions that intersect or encompass the line segment[5]. Subsequently, the effective BER is obtained from the attenuation based on the specific radio design. Note that this model is considers only the attenuation effects due to rain. While other precipitation factors such as humidity contribute significantly lower in the signal attenuation, additional models are needed to capture their impact.

This example shows that while certain links (e.g. $\overline{\mathsf{BC}}$) are severely degraded

Future plans include a complete data processing module to automate the process, which has the added benefit of increased rain rate resolution as discussed later in this section.

[5] Note that the quantization of the rain rates in a storm to 3-levels and assumption of constant rain rate inside a region are two simplifying measures in this model. However, the impact of this simplification is conservative estimate of link quality. Future work involves automating this process which would permit the integration of the point attenuation values along the length of the link using the actual value of the rain rate from the radar data as opposed to the quantized values of this 3-level model.

due to the heavy rain, it is possible to re-route traffic on other adjacent links (e.g. $\overline{\mathsf{BA}}$) that are not in an intense rain region. We argue in the following section that this pattern of rain intensity distribution is typical for a majority of storms in the US Great Plains. While Figure A.4 shows the storm at one time instant, it can be said without loss of generality that an actual rain event is a series of such snapshots that change over time depending upon the velocity and evolution of the storm. In order to model a storm in real time, we continuously calculate the attenuation for all links as the storm moves through the network. The changes in the link attenuation are dependent on the dynamics of the storm relative to the network.

# Appendix B

# Empirical Analysis of Millimeter-Wave Links

This appendix presents an empirical analysis of the impact of weather events on millimeter-wave links using the data collected over test links.

## B.1 Impact on Millimeter-Wave Links

In this section, we evaluate the impact of the weather events on a single millimeter-wave link.

### B.1.1 Effect of Precipitation on FER

The ITU-R P.837 recommendation [92] provides an *estimate* of the precipitation intensities based on *year long* rain rate statistics for various locations. However, due to the significant annual variations, accurate modeling of millimeter-wave links requires precipitation measurements that are time correlated to the link test. As mentioned before, we measured rain intensities at five different locations in and around the link span. During our measurements, we recorded precipitation for approximately 1.6% of the time. Figure B.1 shows the probability distribution of rain rates for the observed events. The cumulative distribution function for the precipitation is given in Figure B.2. It is observed that 95% of the overall precipitation had a rain rate less than 25 mm/hr. The maximum observed rain rate was 160 mm/hr, but occurs with very low probability. As seen in Figure B.2, rain rates greater than 60 mm/hr contribute less than 1% towards the overall precipitation. This understanding of the probability and duration of high-intensity events is necessary to build effective routing algorithms that are resilient to link disruptions.

### Rain

Figure B.3 shows the variation of effective FER during a heavy precipitation event for the system with FEC. It is seen that FEC overcomes degradations to low intensity rain (e.g. below 5 mm/hr). However, for the short-lived, heavy-intensity rain events, the FER can be as high as one leading to complete link failures. For the same event, it was observed that the first radio link, lacking FEC, performed poorly with significant frame errors before and after the rain event due to the latent moisture in the atmosphere.

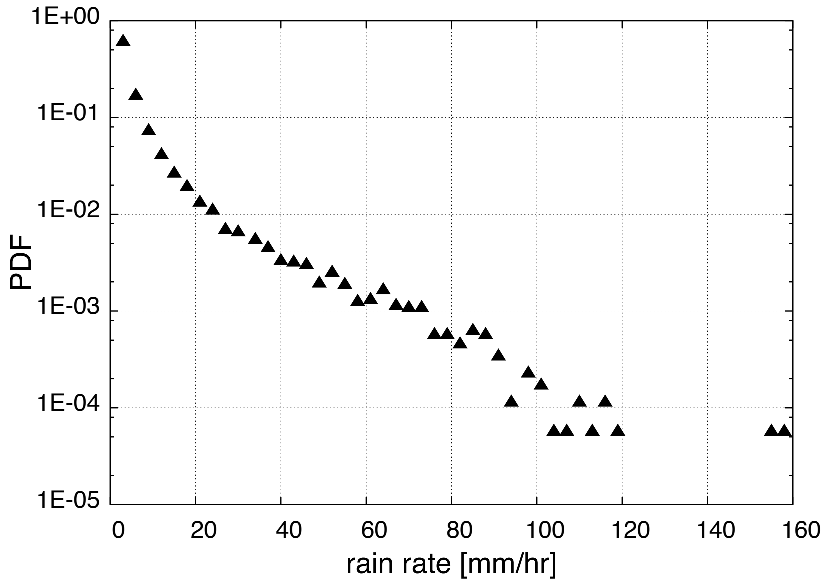

Figure B.1: Rain rate distribution

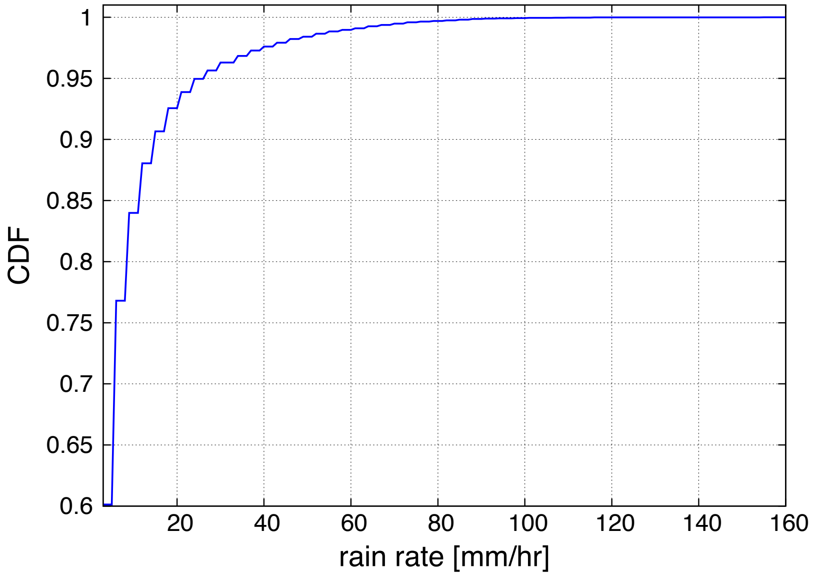

Figure B.2: Cumulative distribution function of recorded weather events

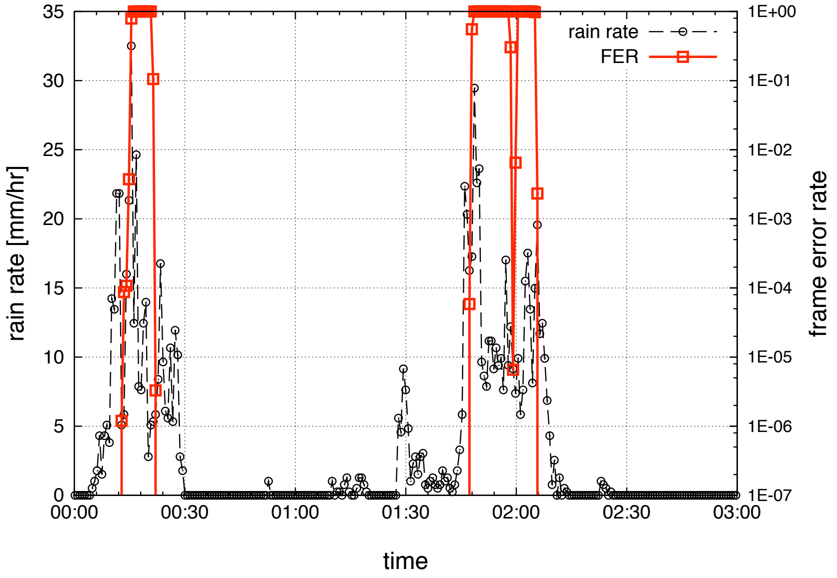

Figure B.3: FER due to a typical rain event on FEC enabled radio link

### Humidity

It is well known that high relative humidity increases the path loss by one dB/km or more depending on the specific conditions [93]. Figure B.4 shows variation in humidity as a function of time over a period of five days (03–08 August 2007), together with the associated FER for first radio link. As expected, the FER shows a strong correlation to humidity, tracking its diurnal characteristic. A relative humidity of over 60% is required to observe a noticeable change in the FER. As the humidity increases above 60%, FER increases exponentially. However, an absolute value of $2 \times 10^{-4}$ FER for a relative humidity of 90% is still small enough to be overcome by Reed-Solomon (204,188) FEC, as implemented in second radio in which no losses were observed.

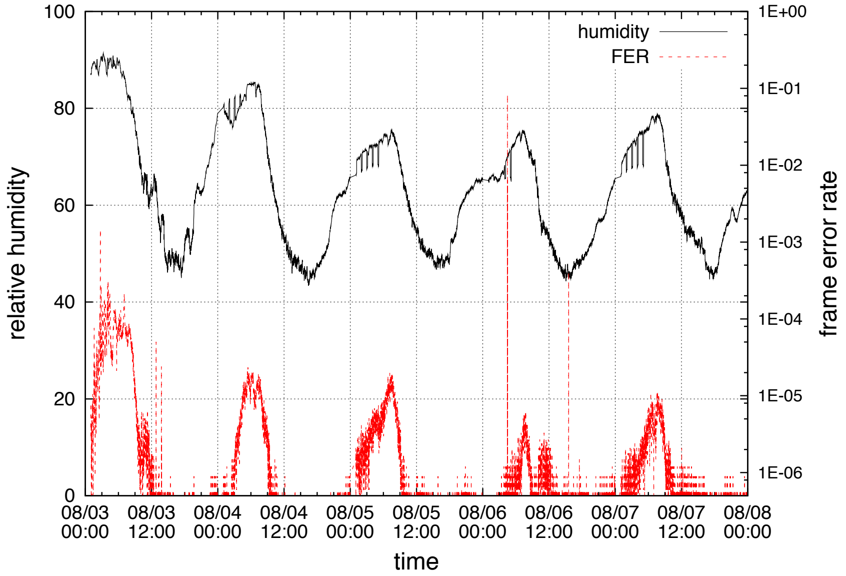

Figure B.4: Correlation between humidity and FER

### Snow

Figure B.5 shows FER as a function of time for a snow event on 29 January 2008 for first radio link which does not employ FEC. Light to heavy snow was observed between 10:00 and 14:00 hours. As with humidity, the system with FEC did not show sensitivity to snow events, while for the link without FEC the system availability showed significant degradation.

## B.1.2 Link Availability Analysis

We show the FER distribution on the two test links in Figure B.6. Note that these distributions are based on those samples in which FER > 0; in other words, we calculate the PDF and CDF for *non-zero* error samples. These

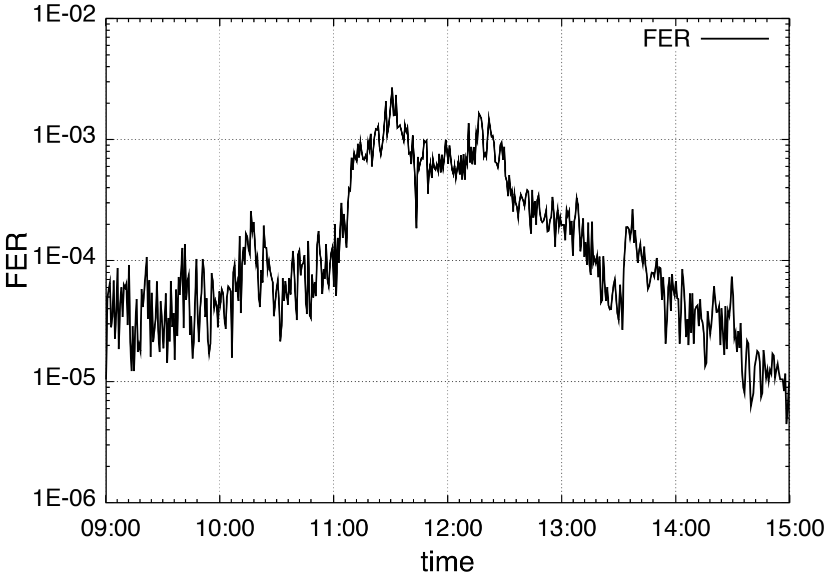

Figure B.5: FER due to a snow event on Jan 29, 2008

cumulative distribution function of the FER shows that the extremes (low and high FERs) contribute more in the error distribution. Comparing the two radios, we see that the CDF of the radio without FEC increases faster in the low FER region when compared to that of radio with FEC. However, the CDF shows that the radio with FEC performs significantly better than its counterpart because the low error rates are well masked by the error correction codes.

In this section, we evaluate the availability of the millimeter-wave links based on the data presented above. We define *link availability* as the percentage of time that the link has a FER less than a specified threshold. Each sample of FER is measured over a 40 sec interval transmission as described above.

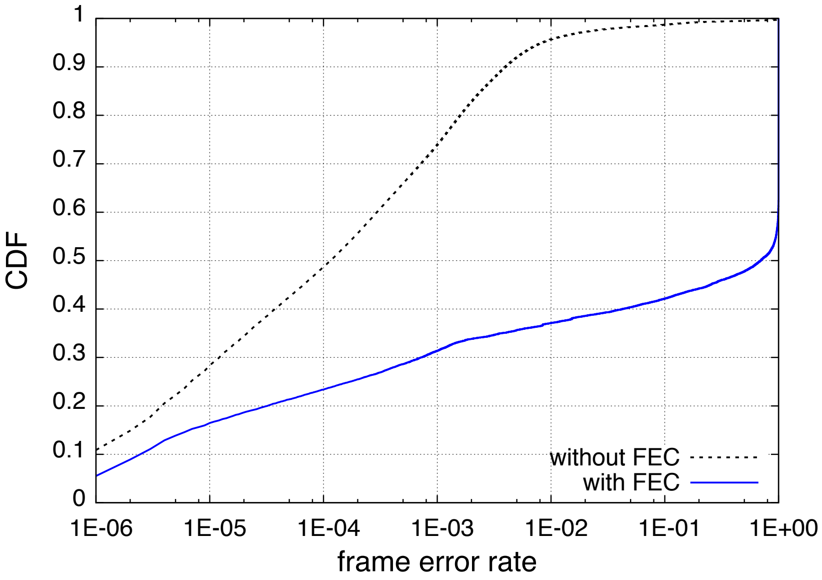

Figure B.6: Cumulative distribution function of observed FER

These individual samples are then used to determine link availability over longer time intervals. Availability as a function of FER threshold is given in Figure B.7. As expected, these results show the benefit of FEC, however, they also provide an example of the relatively high availability (∼95%) possible for a millimeter-wave link over a long (8.7 km) span.

The availability presented in Figure B.7 is specific to the radios used in the experiments. In order to derive a more general measure of link availability under weather disruption, we apply the same analysis to the statistical precipitation data presented in previous section (Section B.1.1). Figure B.8 shows the predicted availability as a function of rain rate tolerable by a given radio. For example, if the link budget of a given radio can overcome the at-

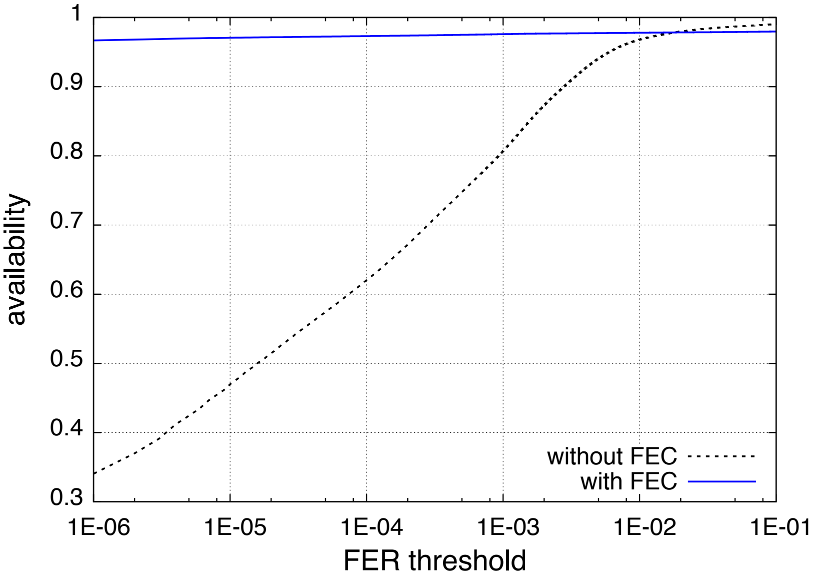

Figure B.7: Availability as a function of FER threshold

tenuation caused by rain rate of up to 10 mm/hr, the predicted availability is 0.997[1].

In order to characterize the link from a service perspective, we define three link states: strong, weak and disconnected; each represents a particular range of FER. The thresholds chosen roughly correspond to the current industry standards on the service requirements of cellular backhaul links:

**State 1:** strongly connected if FER $\leq 5 \times 10^{-5}$;

**State 2:** weakly connected if $5 \times 10^{-5} <$ FER $\leq 5 \times 10^{-2}$;

**State 3:** disconnected if $5 \times 10^{-2} <$ FER.

[1] Rain rate statistics vary depending upon the geographic location.

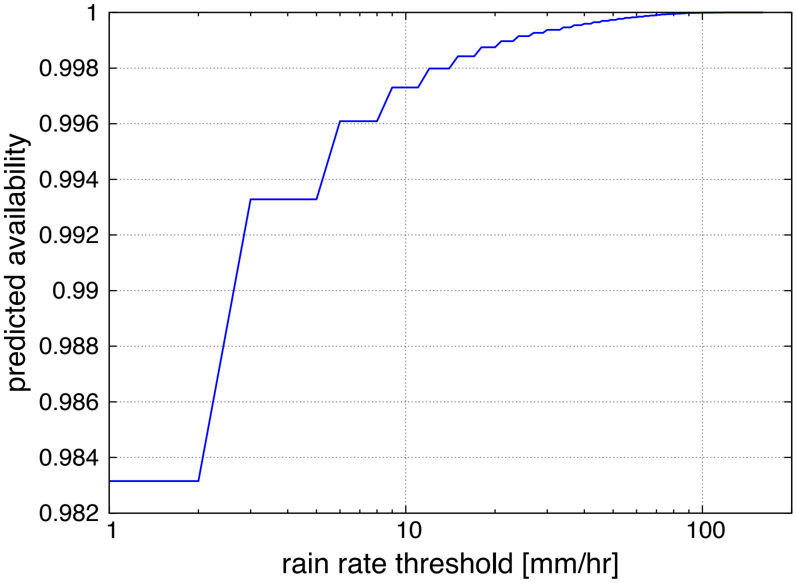

Figure B.8: Availability as a function of rain rate threshold

Table B.1: State probabilities over the duration of the experiments

| State | Radio 1 probability | Radio 2 probability |
|---|---:|---:|
| strongly connected | 0.3864 | 0.9469 |
| weakly connected | 0.5203 | 0.0255 |
| disconnected | 0.0833 | 0.0275 |

The objective is to quantify link state probabilities based on actual observed weather events (including rain, humidity, and snow) at a given geographical location. Table B.1 shows the observed state probabilities for the the two radios.These results were used to drive simulation scenarios for the routing protocol evaluation in Section C.

This shows that, while the RS(204,188) FEC was sufficient to overcome disruptions due to humidity, snow, and low intensity rain, reliable use of millimeter-wave links in the presence of heavy rain warrants the need for solutions above the physical (radio) and link layers.

## B.2 Impact on Mesh Topology

In this section we examine the effects of the previously-mentioned eight rain storms moving across a MWMN consisting of a 4×4 grid topology, with 16 nodes and 24 links. The individual link lengths are 10 km and the network spans a region of approximately 1000 $\text{km}^2$ representing a metropolitan-area mesh network. We analyzed the attenuation and BER experienced by all the links in the network during the duration of the storm. The duration of the storm is defined as the time difference between the instance when the first link is affected by the storm and the instant last link recovers from the storm. We then characterize the effect of the individual storms on a per-link basis as well as for the entire network. Finally, we aggregate the results across all storms to get average statistics.[2]

### B.2.1 Channel Error Rate

As a specific rain event moves across the grid, it affects a number of the links that are in its path. Individual links suffer attenuation to varying degrees depending upon the geographical distribution of the high-intensity regions in the storm. The disruptive effect of a given storm on all the MWMN

[2] Due to space constraints, we show illustrative results for per-storm analysis from a rain event that was observed on 9 July 2008 in Lawrence, KS, USA

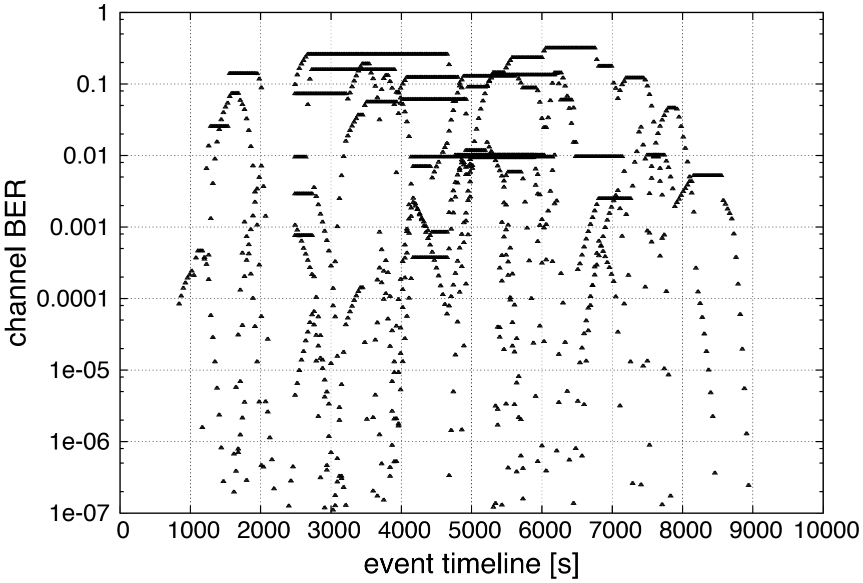

Figure B.9: Storm effect on link error rates at each time interval

links is presented using a using a scatter plot of each link BER at every time interval during the storm duration, shown in Figure B.9. The time interval used for polling was 10 seconds. The plot shows only those links that suffer a BER greater than $1 \times 10^{-7}$; all other links were error free. This distribution of BER values indicates that while a number of links were severely degraded, a significant number were either partially degraded or remained normal throughout the duration of the event.

### B.2.2 Mesh Availability Analysis

As shown in Section B, the sensitivity of the millimeter wave transmission to precipitation and humidity can be compensated using FEC. Therefore, we

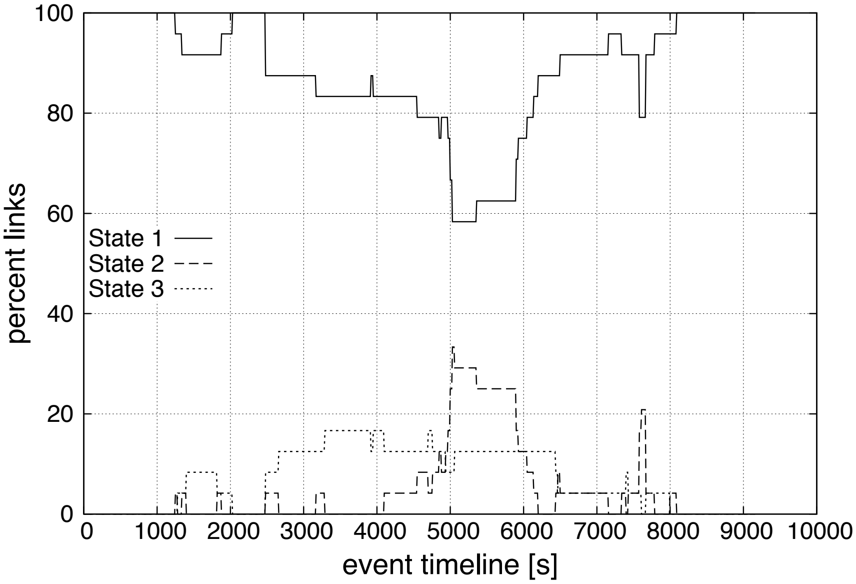

Figure B.10: Percentage of links in each state for each time interval

calculate the link availability based on the effective BER after the applying the FEC gain from the Reed Solomon (204,188) code. We then quantize the effective BER range into three levels representing the state of the links as *normal*, *partially degraded*, and *severely degraded*. We define *normal* as the state in which the effective BER of a link is less than the threshold of $5 \times 10^{-8}$. Links with BER greater than $5 \times 10^{-8}$ but less than $5 \times 10^{-5}$ are defined to be *partially degraded*. Finally, links with BER greater than the threshold of $5 \times 10^{-5}$ are *severely degraded*. These BER thresholds correspond to the FER thresholds mentioned previously. Using these thresholds we determine the percentage of links which fall into a given region and any time during the rain event as shown in Figure B.10.

We observe that just before the storm event, 100% of the links in the mesh are in the normal state (state 1), but as the storm moves over the grid a number of links begin transitioning to the partially and severely degraded states. As the storm moves out of the region the links return to the normal state.

Page left intentionally blank.

# Appendix C

# Mesh Network Performance

In this section, we present the methodology used to model and simulate the storms used in this study. Furthermore, we present the parameters for XL-OSPF and P-WARP used in the simulations. Finally, we quantify the disruptive effect of storms on the MWMN, and compare the disruption tolerance of the proposed mechanisms.

## C.1 Simulation Setup

We conducted simulations using storm modeling software that we developed in MATLAB and the ns-2 simulator [94]. The simulated topology consists of 16 nodes connected in a square mesh as shown in Figure C.1. The millimeter-wave links between each pair of nodes are 10 km long. The nodes remain fixed at their locations throughout the simulation.

To model a cellular backhaul network, nodes 0 and 15 are connected to the external network (Internet) and hence all traffic passes through one of these

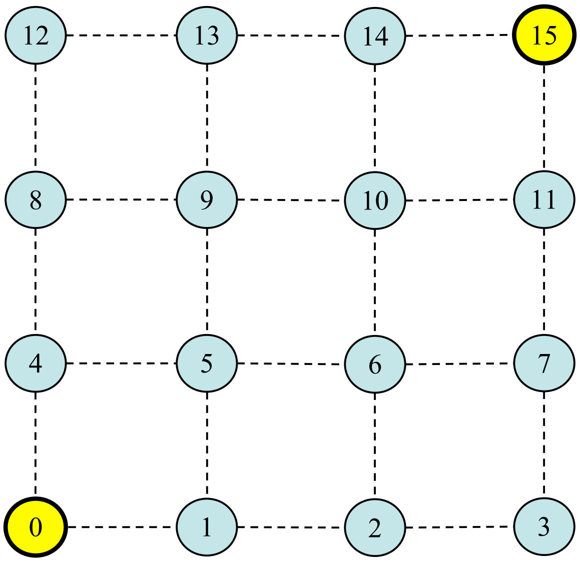

Figure C.1: Simulation topology

nodes. The remaining 14 nodes generate traffic at a rate of 2.4 Mb/s. The generated traffic is CBR over UDP with a packet size of 1000 bytes. We use eight different storms that took place in the Topeka – Lawrence – Kansas City corridor in 2007 and 2008 to evaluate the disruption tolerance of the proposed mechanisms. These storms, which are modeled using the procedure described in Section A.3, consist of an outer ellipse (partially degraded) varying between 30–200 km in diameter and inner ellipses (severely degraded) with a diameter varying between 5–30 km. We have selected two of these storms for illustrative purposes in this section. The first of these occurred 30 September 2007 in Lawrence, KS, and the second on 22 April 2008. We evaluate the packet delivery ratio and the service availability of the network for four different routing mechanisms. .

## C.2 Simulation Results

We compare the performance of P-WARP and XL-OSPF to a baseline of conventional OSPF (reactive to link failure but with no radar-data input) and static routing that provides a worst-case lower bound on performance. The metrics of interest are packet delivery ratio, delay, and overhead.

### C.2.1 Packet Delivery Ratio

The packet delivery ratios averaged over a window of two seconds are shown in Figures C.2 and C.3 for all four routing protocols. This plot shows the instantaneous response of the network to the first simulated storm. As the individual links fail due the storm, the delivery ratio of the network falls rapidly.

The time taken by the network to recover from link failures depends on the routing protocol. Static routing, as expected, performs very poorly and we show it as a lower performance bound reference. OSPF without any modification performs better than static because it can sense link outages from the loss of four consecutive hello packets. However, the delay in detection and route re-computation results in significant packet loss. XL-OSPF performs better than static routing and standard OSPF because it can detect degrading links in a shorter time from the their cost as advertised in the link state updates. Since the cost metric is directly proportional to the link error rate, high error paths are avoided whenever possible. P-WARP outperforms all

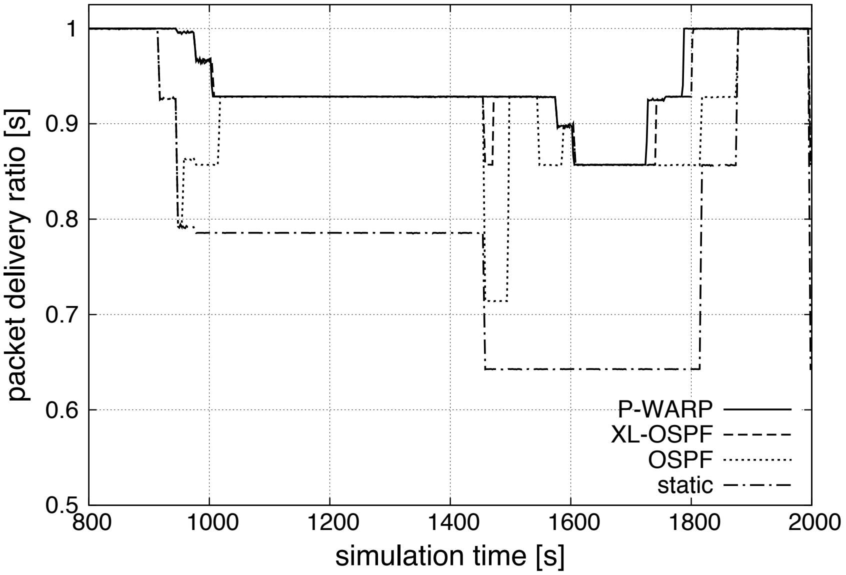

Figure C.2: Windowed average of received packets: first storm

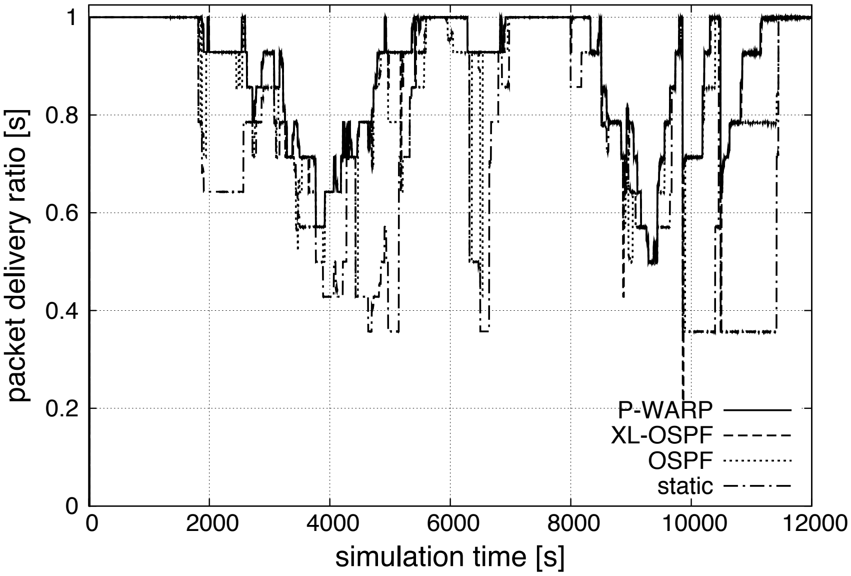

Figure C.3: Windowed average of received packets: second storm

three protocols because it can predict an upcoming link failure from weather updates and reroutes traffic ahead of the disruption.

For example, consider the packet delivery ratios at $t = 1400$ sec in Figure C.2. At $t = 1420$ sec, the storm disrupts additional links (given that the network is already degraded to 93% PDR) causing severe packet loss in case of static routing as the PDR drops to 65% and does not recover until the storm has passed at $t = 1800$ sec. On the other hand, OSPF detects failed links and recovers back to the starting PDR at $t = 1460$ sec, indicating a 40s recovery time that corresponds to the dead interval. XL-OSPF recovers much more quickly than OSPF and static routing. It takes approximately 10s (at $t = 1430$) to recover to maximum delivery ratio. Finally, P-WARP maintains the maximum possible delivery ratio indicating a negative reaction (predictive) time. Accurately predicting the impending disruption, P-WARP preemptively routes data on stable paths, thereby avoiding the failed links completely. The reason the maximum possible delivery ratio is not always one is that an intense storm cell located directly on top of a node may affect *all* the outbound links from a particular node causing all packets sourced from and destined to that node to be dropped irrespective of the routing protocol used. The only way to mitigate this effect is to provide alternative paths from a given node: either a fiber connection to the network when practical, or a lower-frequency, lower-bandwidth, but less weather susceptible link such as in the 23 GHz range which also can be used for weather reflectivity data and WLSU dissemination.

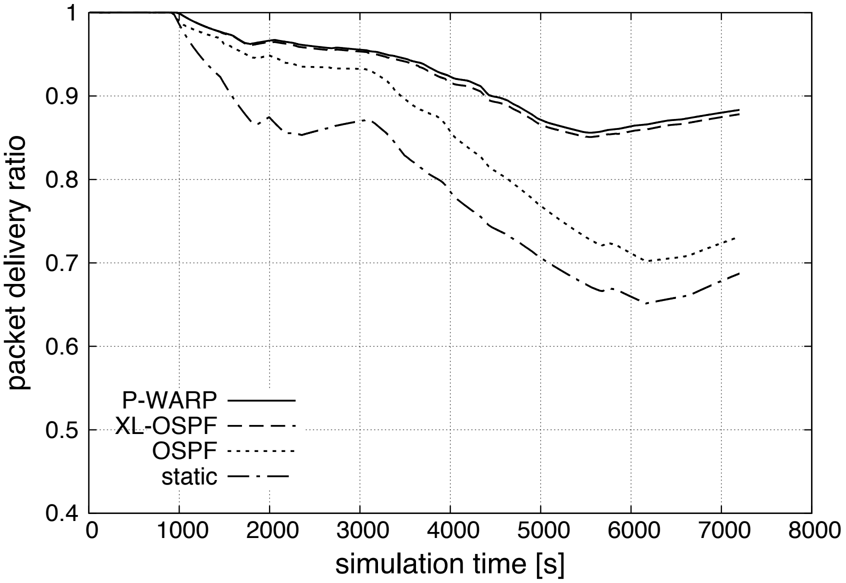

Figure C.4: Cumulative average of received packets: first storm

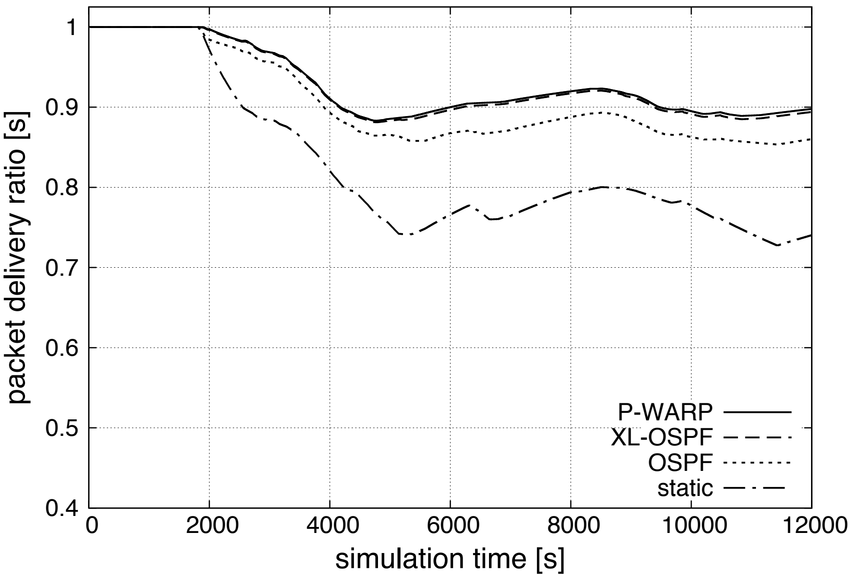

Figure C.5: Cumulative average of received packets: second storm

In order to compare the aggregate performance of the protocols with respect to each other, the cumulative average of the packet delivery ratio is shown in Figures C.4 and C.5. The cumulative average of packets delivered by XL-OSPF is very close to that of P-WARP in each case, and both outperform conventional OSPF. In the case of XL-OSPF, the frequency of link state updates determine the reaction time of the network; strict restoration times would require a very small value of update intervals. Because of its predictive nature, P-WARP has two distinct advantages: first, it has *negative* reaction time that might be necessary if stringent service requirements (such as 50ms restoration time frequently advertised by network service providers) are to be met; second, the frequency of weather updates does not scale with the restoration times leading to lower protocol overhead.

### C.2.2 Delay

While the absolute value of the end-to-end delays in the wireless-mesh topology are negligible, it is worth noting that P-WARP and XL-OSPF can cause higher average delays as they routes packets around link outages, while OSPF and static routing lose more packets during link outages and thus do not incur these delays. Figures C.6 and C.7 show the end-to-end delays averaged over a two second interval.

We have conducted a number of additional simulations with varying update intervals for OSPF and XL-OSPF. We have observed that while the general relationship between the OSPF, XL-OSPF and P-WARP routing remains

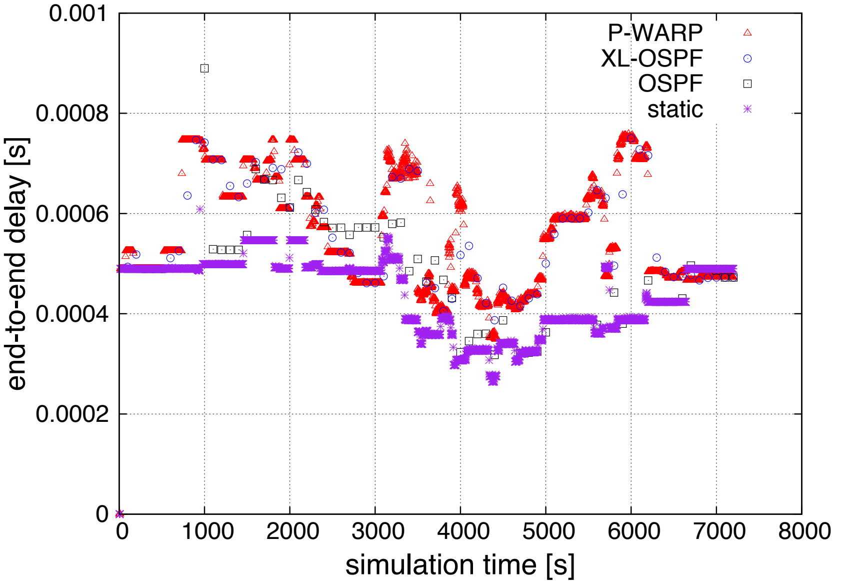

Figure C.6: End-to-end delay: first storm

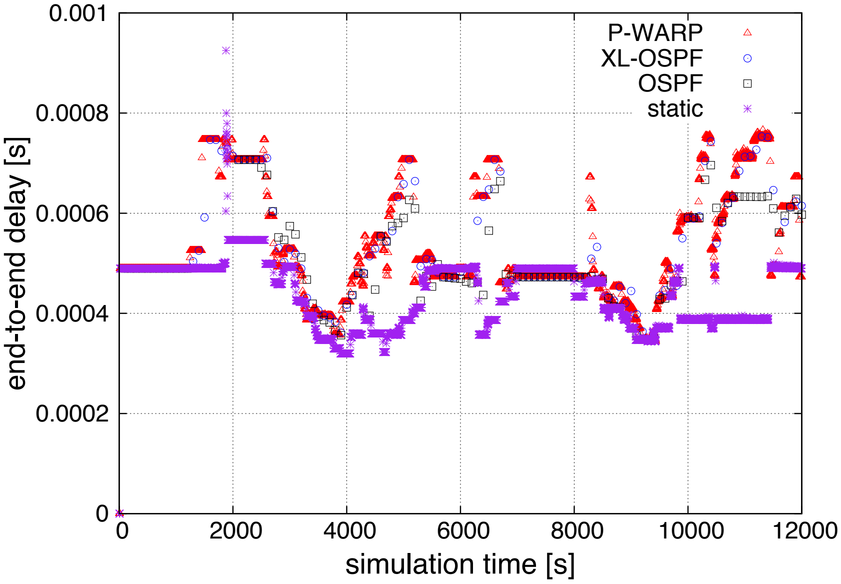

Figure C.7: End-to-end delay: second storm

the same, the performance gap between XL-OSPF and P-WARP decreases with increasing LSA update frequency, as expected.

### C.2.3 Overhead

Static routing does not generate any overhead traffic. Conventional OSPF generates periodic hello messages as well as LSAs. However, the frequency of the LSA does not affect the performance because the quality of the link is not reflected in the its cost metric. On the other hand, XL-OSPF uses a cost metric that is proportional to link BER and therefore generates frequent updates that are flooded in the network. As discussed above, in order to react quickly to weather disruptions, XL-OSPF must generate LSAs at a higher rate. This leads to significant increase in the overhead. The number of updates generated in P-WARP are comparatively lower for two reasons. First, a single WLSU update carries the predicted costs for all the links. Hence, individual nodes do not generate link state updates. Second, an update is generated only when there is a change in the predicted BER of one or more links. Since, the performance of the P-WARP is not dependent on the arbitrary update frequency, WLSUs are rate limited to 30 seconds to bound the overhead.

Page left intentionally blank.

# Bibliography

[1] Ali Zolfaghari and Fed J. Kaudel. Framework for network survivability performance. *IEEE Journal on Selected Areas in Communications*, 12(1):46–51, 1994.

[2] James P.G. Sterbenz and David Hutchison. ResiliNets: Multilevel resilient and survivable networking initiative web page. http://www.ittc.ku.edu/resilinets/index.html, 2007.

[3] Department of defense global information grid architectural vision: Vision for a net-centric, service-oriented DoD enterprise. Unclassified, DoD CIO, June 2007.

[4] K. Stouffer, J. Falco, and K. Kent. Guide to supervisory control and data acquisition (SCADA) and industrial control systems security. Special Publication NIST-SP-800-82-2006, National Insitute of Standards and Technology (NIST), September 2006.

[5] James Ellis, David Fisher, Thomas Longstaff, Linda Pesante, and Richard Pethia. Report to the president's commission on critical infrastructure protection, 1997.

[6] A roadmap for cybersecurity research. Technical report, Department of Homeland Security (DHS), November 2009.

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

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

[9] UK resilience homepage. `http://www.cabinetoffice.gov.uk/ukresilience.aspx`, February 2010.

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

[11] James P.G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. 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, June 2010.

[12] 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), June 2004.

[13] R. J. Ellison, D. A. Fisher, R. C. Linger, H. F. Lipson, T. Longstaff, and N. R. Mead. Survivable network systems: An emerging discipline. Technical Report CMU/SEI-97-TR-013, Carnegie-Mellon Sorfware Engineering Institute, PA, 1999.

[14] James P. G. Sterbenz, Rajesh Krishnan, Regina Rosales Hain, Alden W. Jackson, David Levin, Ram Ramanathan, and John Zao. Survivable mobile wireless networks: issues, challenges, and research directions. In *WiSE '02: Proceedings of the 3rd ACM workshop on Wireless security*, pages 31–40, New York, NY, USA, 2002. ACM Press.

[15] Jean-Claude Laprie. Dependability: Basic concepts and terminology. Draft, IFIP Working Group 10.4 – Dependable Computing and Fault Tolerance, August 1994.

[16] Algirdas Avizienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr. Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable Secure Computing*, 1(1):11–33, January 2004.

[17] Wayne D. Grover. *Mesh-Based Survivable Networks*, chapter 3. Prentice-Hall PTR Pearson, Upper Saddle River NJ, 2004.

[18] Alexander A. Hagin. Performability, reliability, and survivability of communicationnetworks: system of methods and models for evaluation. In

*Proceedings of the 14th International Conference on Distributed Computing Systems*, pages 562–573, June 1994.

[19] T1A1.2 Working Group. Enhanced network survivability performance. Technical Report T1.TR.68-2001, Alliance for Telecommunications Industry Solutions (ATIS), February 2001.

[20] Abdul Jabbar Mohammad, David Hutchison, and James P.G. Sterbenz. Towards quantifying metrics for resilient and survivable networks. In *Proceedings of the 14th IEEE International Conference on Network Protocols (ICNP)*, pages 17–18, November 2006.

[21] E.F. Moore and C.E. Shannon. Reliable circuits using less reliable relays. *Journal of the Franklin Institute*, 262(3):191–208, 1956.

[22] W.H. Pierce. *Failure-tolerant Computer Design*. Academic Press, 1965.

[23] A. Avizienis. Design of Fault-Tolerant Computers. *Fall Joint Computer Conference*, 31:733–743, 1967.

[24] RE Lyons and W. Vanderkulk. The use of triple-modular redundancy to improve computer reliability. *IBM Journal of Research and Development*, 6(2):200–209, 1962.

[25] H. Frank. Survivability Analysis of Command and Control Communications Networks–Part I. *Communications, IEEE Transactions on [legacy, pre-1988]*, 22(5):589–595, 1974.

[26] H. Frank. Survivability Analysis of Command and Control Communications Networks–Part II. *Communications, IEEE Transactions on [legacy, pre-1988]*, 22(5):596–605, 1974.

[27] M.D Beaudry. Performance-related reliability measures for computing systems. *IEEE Transactions on Computers*, 27:540–547, 1978.

[28] YW Ng and A. Avizienis. A Reliability Model for Gracefully Degrading and Repairable Fault-tolerant Systems. In *Proceedings of 7th International Symposium on Fault-Tolerant Computing*, pages 29–34. IEEE Computing Society Publications, 1977.

[29] J. Losq. Effects of Failures on Gracefully Degradable Systems. In *Proc. of 7th Fault-Tolerant Computing Symposium*, pages 29–34, 1977.

[30] Ragnar Huslende. A combined evaluation of performance and reliability for degradable systems. In *SIGMETRICS '81: Proceedings of the 1981 ACM SIGMETRICS conference on Measurement and modeling of computer systems*, pages 157–164, New York, NY, USA, 1981. ACM Press.

[31] JF Meyer. On Evaluating the Performability of Degradable Computing Systems. *Computers, IEEE Transactions on*, 100(29):720–731, 1980.

[32] FA Gay and ML Ketelsen. Performance evaluation for gracefully degrading systems. In *Proc. of the 9th Annual Int. Symp. on Fault Tolerant Computing*, pages 51–58, June 1979.

[33] John C. Knight, Elisabeth A. Strunk, and Kevin J. Sullivan. Towards a rigorous definition of information system survivability. In *Proceedings of the DARPA Information Survivability Conference and Exposition DISCEX III*, pages 78–89, Washington DC, April 2003.

[34] SC Liew and KW Lu. A framework for network survivability characterization. In *SUPERCOMM/ICC'92: Proceedings of IEEE International Conference on Communications, 1992. ICC 92, Conference record, /, Discovering a New World of Communications.*, pages 405–410, 1992.

[35] SC Liew and KW Lu. A framework for characterizing disaster-based network survivability. *Selected Areas in Communications, IEEE Journal on*, 12(1):52–58, 1994.

[36] Andreas Antonopoulos. Metrication and performance analysis on resilience of ring-basedtransport network solutions. In *GLOBECOM'99: Proc. of the Global Telecommunications Conference*, volume 2, pages 1551–1555, 1999.

[37] Vickie R. Westmark. A definition for information system survivability. In *HICSS '04: Proceedings of the 37th Annual Hawaii International Conference on System Sciences (HICSS'04) - Track 9*, page 90303.1, Washington, DC, USA, 2004. IEEE Computer Society.

[38] Q. Gan and B.E. Helvik. 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*, April 2006.

[39] W. Molisz. Survivability function-a measure of disaster-based routing performance. *IEEE Journal on Selected Areas in Communications*, 22(9):1876–1883, 2004.

[40] Dongyan Chen, Sachin Garg, and Kishor S. Trivedi. Network survivability performance evaluation:: a quantitative approach with applications in wireless ad-hoc networks. In *MSWiM '02: Proceedings of the 5th ACM international workshop on Modeling analysis and simulation of wireless and mobile systems*, pages 61–68, New York, NY, USA, 2002. ACM Press.

[41] Y. Liu and K.S. Trivedi. A general framework for network survivability quantification. In *Proc. of the 12th GI/ITG Conf. Measuring, Modelling and Evaluation of Computer and Comm. Systems*, 2004.

[42] David M. Nicol, William H. Sanders, and Kishor S. Trivedi. Model-based evaluation: From dependability to security. *IEEE Transactions on Dependable and Secure Computing*, 01(1):48–65, 2004.

[43] G. Qu, R. Jayaprakash, S. Hariri, and CS Raghavendra. A Framework for Network Vulnerability Analysis. In *CT '02: Proceedings of the 1st IASTED International Conference on Communications, Internet, Information Technology*, pages 289–298, St. Thomas, Virgin Islands, USA, Sep-Oct 2002.

[44] Salim Hariri, Guangzhi Qu, Tushneem Dharmagadda, Modukuri Ramkishore, and Cauligi S. Raghavendra. Impact analysis of faults and attacks in large-scale networks. *IEEE Security and Privacy*, 01(5):49–54, 2003.

[45] Mina Guirguis, Azer Bestavros, Ibrahim Matta, and Yuting Zhang. Reduction of quality (RoQ) attacks on internet end-systems. In *INFOCOM 2005: Proceedings of 24th Annual Joint Conference of the IEEE Computer and Communications Societies*, pages 1362–1367, Miami, Florida, march 2005.

[46] S. Jha, J. Wing, Richard C. Linger, and Thomas A. Longstaff. Survivability analysis of network specifications. In *DSN '00: Proceedings of the 2000 International Conference on Dependable Systems and Networks (formerly FTCS-30 and DCCA-8)*, pages 613–622, Washington, DC, USA, 2000. IEEE Computer Society.

[47] Somesh Jha and Jeannette M. Wing. Survivability analysis of networked systems. In *ICSE '01: Proceedings of the 23rd International Conference on Software Engineering*, pages 307–317, Washington, DC, USA, 2001. IEEE Computer Society.

[48] Nigel Edwards. Building dependable distributed systems. Technical report APM.1144.00.02, ANSA, February 1994.

[49] R.J. Ellison, D. Fisher, R.C. Linger, H.F. Lipson, T. Longstaff, and N.R. Mead. Survivable network systems: An emerging discipline. Technical Report CMU/SEI-97-TR-013, Software Engineering Institute, Carnegie Mellon University, 1997.

[50] M. Steinder and A. Sethi. A survey of fault localization techniques in computer networks. *Science of Computer Programming*, 53(2):165–194, November 2004.

[51] James P.G. Sterbenz and Joseph D. Touch. *High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication.* Wiley, 1st edition, May 2001.

[52] Justin P. Rohrer, Abdul Jabbar, and James P.G. Sterbenz. Path diversification: A multipath resilience mechanism. In *Proceedings of the 7th International Workshop on the Design of Reliable Communication Networks (DRCN)*, Washington, DC, USA, October 2009.

[53] Alberto Medina, Ibrahim Matta, and John Byers. On the origin of power laws in internet topologies. *SIGCOMM Comput. Commun. Rev.*, 30(2):18–28, 2000.

[54] Jared Winick and Sugih Jamin. Inet-3.0: Internet topology generator. Technical Report UM-CSE-TR-456-02, EECS, University of Michigan, 2002.

[55] Hilary C. Styron. Csx tunnel fire: Baltimore, md. US Fire Administration Technical Report USFA-TR-140, Federal Emergency Management Administration, Emmitsburg, MD, 2001.

[56] Autonomic Network Architecture. http://www.ana-project.org/, July 2009.

[57] ResumeNet. http://www.resumenet.eu/, July 2009.

[58] Future Internet Design (FIND). NSF NeTS research program, http://www.nets-find.net/, July 2009.

[59] Bobby Bhattacharjee, Ken Calvert, Jim Griffioen, Neil Spring, and James P.G. Sterbenz. Postmodern Internetwork Architecture. Technical Report ITTC-FY2006-TR-45030-01, Information Telecommunication and Technology Center, University of Kansas, Lawrence, KS, 2006.

[60] Abdul Jabbar, Qian Shi, Egemen Çetinkaya, and James P.G. Sterbenz. KU-LocGen: Location and cost-constrained network topology generator. ITTC Technical Report ITTC-FY2009-TR-45030-01, The University of Kansas, Lawrence, KS, December 2008.

[61] B.M. Waxman. Routing of multipoint connections. *Selected Areas in Communications, IEEE Journal on*, 6(9):1617–1622, Dec 1988.

[62] K.I. Calvert, M.B. Doar, and E.W. Zegura. Modeling internet topology. *Communications Magazine, IEEE*, 35(6):160–163, Jun 1997.

[63] Hongsuda Tangmunarunkit, Ramesh Govindan, Sugih Jamin, Scott Shenker, and Walter Willinger. Network topology generators: degree-based vs. structural. In *SIGCOMM '02: Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 147–159, New York, NY, USA, 2002. ACM.

[64] J.C. Doyle, D.L. Alderson, L. Li, S. Low, M. Roughan, S. Shalunov, R. Tanaka, and W. Willinger. The "robust yet fragile" nature of the internet. *Proceedings of the National Academy of Sciences*, 102(41):14497, 2005.

[65] A. Lakhina, JW Byers, M. Crovella, and I. Matta. On the geographic location of Internet resources. *IEEE Journal on Selected Areas in Communications*, 21(6):934–948, 2003.

[66] S.H. Yook, H. Jeong, and A.L Barabasi. Modeling the internet's large-scale topology. *Proceedings of the National Academy of Sciences*, 99(21):13382–13386, 2002.

[67] Rocketfuel: An ISP topology mapping engine, September 2008.

[68] David Alderson, Lun Li, Walter Willinger, and John C. Doyle. Understanding internet topology: principles, models, and validation. *IEEE/ACM Trans. Netw.*, 13(6):1205–1218, 2005.

[69] A. Sydney, C. Scoglio, P. Schumm, and R. E. Kooij. Elasticity: topological characterization of robustness in complex networks. In *BIONETICS '08: Proceedings of the 3rd International Conference on Bio-Inspired Models of Network, Information and Computing Sytems*, pages 1–8, ICST, Brussels, Belgium, Belgium, 2008. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering).

[70] GÉANT2. `http://www.geant2.net/`, December 2009.

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

[72] ITU-R F.1704. Characteristics of multipoint-to-multipoint fixed wireless systems with mesh network topology operating in frequency bands above about 17 GHz. ITU-R Recommendation F.1704, 2005.

[73] J. A. Khan and H. M. Alnuweiri. Traffic engineering with distributed dynamic channel allocation in BFWA mesh networks at millimeter wave band. In *Proceedings of the 14th IEEE Workshop on Local and Metropolitan Area Networks*, pages 1–6, Chania, Greece, September 2005.

[74] K. Ohata, K. Maruhashi, M. Ito, and T. Nishiumi. Millimeter-wave broadband transceivers. *NEC Journal of Advanced Technology*, 2(3):211–216, 2005.

[75] E. Torkildson, B. Ananthasubramaniam, U. Madhow, and M. Rodwell. Millimeter-wave MIMO: Wireless Links at Optical Speeds. In *Proceedings of the 44th Allerton Conference on Communication, Control and Computing*, Monticello, Illinois, USA, September 2006.

[76] H. Izadpanah. A millimeter-wave broadband wireless access technology demonstrator for the next-generation internet network reach extension. *IEEE Communications Magazine*, pages 140–145, September 2001.

[77] G. Hendrantoro, Indrabayu, T. Suryani, and A. Mauludiyanto. A multivariate autoregressive model of rain attenuation on multiple short radio links. *IEEE Antennas and Propagation Letters*, 5:54–57, 2006.

[78] K. S. Paulson and C. J. Gibbins. Rain models for the prediction of fade durations at millimetre wavelengths. *IEE Proceedings - Microwaves, Antennas and Propagation*, 147(6):431–436, December 2000.

[79] J. Moy. OSPF version 2. RFC 2328 (Standard), April 1998.

[80] L. Iannone, R. Khalili, K. Salamatian, and S. Fdida. Cross-layer routing in wireless mesh networks. *Wireless Communication Systems, 2004. 1st International Symposium on*, pages 319–323, 2004.

[81] Guangyu Pei, Phillip A. Spagnolo, Sang Bae, Thomas R. Henderson, and Jae H. Kim. Performance improvements of ospf manet extensions: A cross layer approach. *Military Communications Conference, 2007. MILCOM 2007. IEEE*, pages 1–7, Oct. 2007.

[82] M.E.M. Campista, P.M. Esposito, I.M. Moraes, L.H.M. Costa, O.C.M. Duarte, D.G. Passos, C.V.N. de Albuquerque, D.C.M. Saade, and M.G. Rubinstein. Routing metrics and protocols for wireless mesh networks. *Network, IEEE*, 22(1):6–12, Jan.-Feb. 2008.

[83] Hung Quoc Vo, Young Yig Yoon, and Choong Seon Hong. Multi-path routing protocol using cross-layer congestion-awareness in wireless mesh network. In *ICUIMC '08: Proceedings of the 2nd international conference on Ubiquitous information management and communication*, pages 486–490, New York, NY, USA, 2008. ACM.

[84] Justin P. Rohrer, Abdul Jabbar, Erik Perrins, and James P.G. Sterbenz. Cross-layer architectural framework for highly-mobile multihop airborne telemetry networks. In *Proceedings of the IEEE Military Communications Conference (MILCOM)*, San Diego, CA, USA, November 2008.

[85] Abdul Jabbar, Erik Perrins, and James P.G. Sterbenz. A cross-layered protocol architecture for highly-dynamic multihop airborne telemetry networks. In *Proceedings of the International Telemetering Conference (ITC)*, San Diego, CA, October 27–30 2008.

[86] Abdul Jabbar and James P.G. Sterbenz. AeroRP: A geolocation assisted aeronautical routing protocol for highly dynamic telemetry environments. In *Proceedings of the International Telemetering Conference*, Las Vegas, NV, October 26–29 2009.

[87] Egemen K. Çetinkaya, Dan Broyles, Amit Dandekar, Sripriya Srinivasan, and James P.G. Sterbenz. A Comprehensive Framework to Simulate Network Attacks and Challenges. In *RNDM'10 - Second International Workshop on Reliable Networks Design and Modeling*, Moscow, Russia, October 20010. to appear.

[88] Vaisala weather transmitter WXT510: The most essential of weather. Vaisala Instruments Catalog, 2008.

[89] Abdul Jabbar, Bharatwajan Raman, Victor S. Frost, and James P. G. Sterbenz. Weather disruption-tolerant self-optimising millimeter mesh networks. In *Proceedings of IWSOS : Third International IFIP/IEEE Workshop on Self-Organizing Systems*, volume 5343 of *Lecture Notes in Computer Science*, pages 242–255. Springer, 2008.

[90] Donna F. Tucker and Xingong Li. Characteristics of warm season precipitating storms in the Arkansas-Red river basin. *Journal of Geophysical Research*, 2009. Submitted.

[91] ITU-R P.838. Specific attenuation model for rain for use in prediction methods. ITU-R Recommendation P.838-3, 2005.

[92] ITU-R P.837. Characteristics of precipitation for propagation modelling. ITU-R Recommendation P.837-4, 2003.

[93] Hans J. Liebe. An updated model for millimeter wave propagation in moist air. *Radio Science*, 20:1069–1089, 1985.

[94] The network simulator: ns-2. `http://www.isi.edu/nsnam/ns/`, December 2007.