A Network Planning and Management Tool for Mitigating the Impact of Spatially Correlated Failures in Infrastructure Networks

Arun Das, Arunabha Sen, Chunming Qiao, Nasir Ghani, Nathalie Mitton

▶ To cite this version:

HAL Id: hal-01254982

https://inria.hal.science/hal-01254982

Submitted on 21 Mar 2016

A Network Planning and Management Tool for Mitigating the Impact of Spatially Correlated Failures in Infrastructure Networks

Arun Das*, Arunabha Sen*, Chunming Qiao†, Nasir Ghani‡, Nathalie Mitton§

*School of Computing, Informatics and Decision System Engineering, Arizona State University, Tempe, Arizona 85287, USA
†Department of Computer Science and Engineering, SUNY at Buffalo, Buffalo, NY 14201, USA
‡Department of Electrical Engineering, University of South Florida, Tampa, FL 33620, USA
§Inria, 40 Avenue Halley, 59650 Villeneuve D'ASCQ, France
Email: arun.das@asu.edu, asen@asu.edu, qiao@computer.org, nghani@usf.edu, nathalie.mitton@inria.fr

Abstract—Current practices of fault-tolerant network design ignore the fact that most network infrastructure faults are *localized* or *spatially correlated* (i.e., confined to regions). Network operators require new tools to mitigate the impact of such *region based faults* on their infrastructures. Utilizing the support from the U.S. Department of Defense, and by consolidating a wide range of theories and solutions developed in the last few years, the authors of this paper have developed an advanced *Network Planning and Management Tool* (NPMT) that facilitates the design and provisioning of robust and resilient networks. The tool provides multi-faceted network design, evaluation and simulation capabilities for network planners. Future extensions of the tool currently being worked upon not only expand the tools capabilities, but also extend these capabilities to heterogeneous interdependent networks such as communication, power, water and satellite networks.

I. Introduction and Motivation

It is extremely important that planners for large wide area networks have the right tools to design robust and resilient networks that can effectively withstand large scale *geographically correlated failures* in their networks. Such failures can be triggered by nature (hurricane or earthquake) or humans (nuclear attack or conventional weapon attack over a large geographical area). With research support from the U.S. Defense Threat Reduction Agency, an agency whose mission is to protect the U.S. against *Weapons of Mass Destruction* (WMD), such as nuclear, biological or chemical attacks, the authors of this paper, over the last six years have developed a wide ranging set of concepts and techniques for enhancing network robustness against *spatially correlated* or *region based* faults. We have recently incorporated these concepts and techniques into a *Network Planning and Management Tool* (NPMT) [1] for the benefit of network designers, planners and operators.

In this paper, we first describe the novel concepts developed to design networks that are robust against region based faults, and then describe how these concepts have been incorporated into the NPMT. The goal of this paper is to bring to the attention of the networking research community, and the audience of the workshop on DRCN in particular, the existence of NPMT as a tool that consolidates a large body of work on spatially correlated failures, and as a tool that can be used by the community to meet the needs for robust network design against spatially correlated failures. In essence, this paper's contribution should not be measured in terms of new analytical findings, but in terms of service to the networking community.

We use the term *WMD attack* to imply a large scale *geographically correlated failure* such as failures caused by an earthquake, hurricane or nuclear attack. The characteristic of a WMD attack is *massive* but *localized* faults. The *connectivity* of a network [2] is generally accepted as a metric for evaluating the *fault-tolerance capability* of a network [3]. If a network's connectivity is $k + 1$, then the network can tolerate up to k faults, implying that the surviving network remains connected even after k failures. The connectivity metric, however, has no way of capturing *locality*, i.e., the faulty nodes/edges may be close or far away from each other. Thus, the connectivity metric cannot distinguish between faults that are geographically correlated (a WMD fault characteristic), and faults that are not. Connectivity as a metric also fails to capture other important structural properties of the network such as the *number* or *size* of the *connected components* [2] into which a network disintegrates when the number of failed nodes/edges *exceeds* the node/edge connectivity of the network.

Recognizing the limitations of connectivity as a metric for capturing the special characteristics of geographically correlated failures, the authors of [4] introduced the notion of *region-based connectivity*. A *region* may be defined either with reference to the network graph or to the network geometry (i.e., layout of the network in a two or three dimensional space). For example, a region may be defined as a subgraph with *diameter* d (where the diameter of a graph is defined as the maximum of the shortest path distance between a pair of nodes, taken over all source-destination node pairs). Or, a region may also be defined as a collection of nodes and edges in the network graph layout that is covered by a *circular* area in that layout. Figure 1(a) shows an example of a circular region-based fault.

The NPMT described in this paper is intended to support design and analysis of single layered and multi-layered interdependent heterogeneous networks. In essence, the NPMT is particularly suitable for planning and design of critical infrastructures. For example, from the single network layer perspective, the NPMT enables backbone communication network providers, such as, AT&T, Sprint, Qwest and Level 3 Communications, to (i) identify the most vulnerable parts of
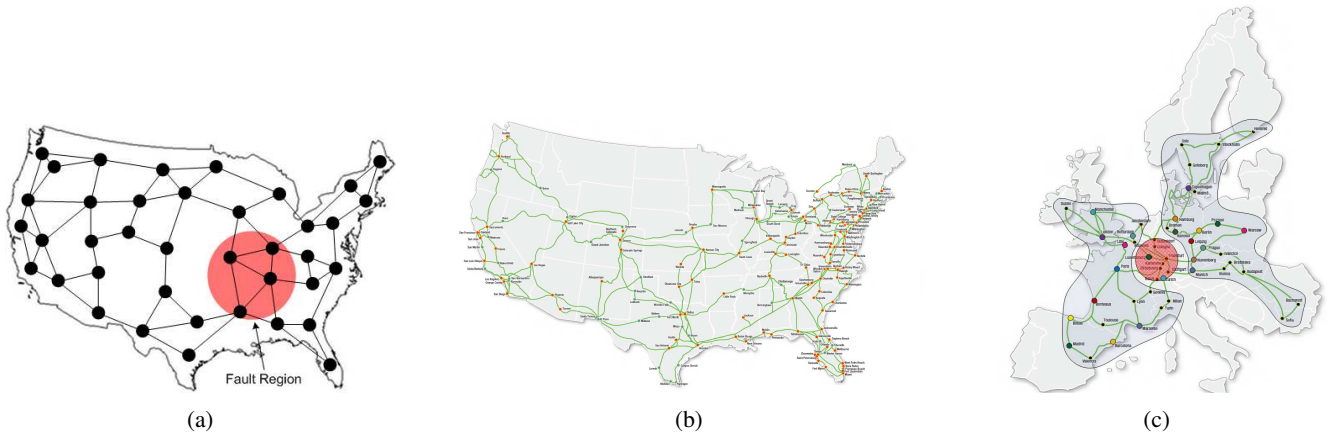
Fig. 1: (a) Network with circular fault region, (b) Optical fiber network of a major U.S. provider, (c) Optical Fiber network of a major European provider disrupted by a WMD attack

their network against a WMD attack, and (ii) reinforce the network with least cost to eliminate or significantly reduce the threat of network disruption due to a WMD attack. Figure 1(b) shows the backbone network of a major U.S. provider and Figure 1(c) shows how the backbone network of a major European provider can potentially be disrupted by a WMD attack. From a multi-layer perspective the NPMT can be used for design and analysis of smart cities, where heterogeneous networks ranging from disparate telecom networks (such as 2G, 3G, WiFi, Bluetooth, etc.) to water, electricity and gas distribution networks, form a complex interdependent ecosystem. Subsequently, failures in one network, for example a leak in the water distribution network, may deteriorate other nearby (spatially correlated) infrastructures such as gas or electricity whose pipes and cables may get affected due to the leak. In this context, a tool like NPMT can be an asset for utility companies and smart city planners to quickly perform (i) root cause analysis of failure, and (ii) forecast fault evolution, to direct repairs and maintenance towards specific network components and restrict fault propagation. To the best of our knowledge no such tool is available today that supports features for planning and designing of single layer and multi-layer interdependent networks in the presence of spatially correlated faults.

Several studies in the network research community have focused on different aspects of spatially correlated or region-based faults in networks [5-11], however, to the best of our knowledge there does not exist an executable platform that consolidates the findings and techniques of these studies into a readily usable tool. The NPMT is intended to fill that gap and be such a platform that can incorporate the outcomes developed in [5-11] into executable modules to be integrated into the NPMT. This will allow network designers, planners and operators to use the results of these studies in their real world operational networks.

The rest of the paper is organized as follows: In Section II we present an overview of the underlying concepts and theoretical results that the NPMT operates on. In Section III we outline the capabilities of the NPMT and finally Section IV concludes this paper.

II. CONCEPTS, METRICS AND SOLUTION TECHNIQUES

In this section we give a brief overview of the underlying concepts, metrics and solution techniques that the Network

Planning and Management Tool (NPMT) utilizes to carry out its functional operations. The NPMT is built as a modular execution engine that can execute smaller reusable modules to perform desired operations on a network topology. In this respect, the current version of the NPMT comprises of different modules that deal with both static and dynamic aspects of robust and resilient network design. The modular approach allows design, development and testing of these modules to be done independently and defers the integration into the NPMT until a module meets it's functional requirements. In the following sub-sections we give a brief overview of the analytical foundations of these modules. It may be noted that, as of writing this paper not all modules have been implemented and integrated into the NPMT. Accordingly, we highlight our ongoing work in the discussion below.

A. Region-Based Fault Metrics Computation Module

As outlined in Section I, connectivity as a metric fails to capture several characteristics of the network in presence of spatially correlated failures. For instance, the number or size of the connected components into which a network disintegrates in the presence of a spatially correlated fault is not captured by the traditional connectivity metric. In order to overcome these gaps and capture such network state characteristics, several metrics and their computation techniques have been proposed by the research community. For a given network topology, the NPMT can analyze the network and compute metrics pertinent to network state in the presence of spatially correlated faults. The following metrics are supported by the NPMT:

Region-based Connectivity Metric Computation

Region based connectivity can be considered under two fault models – (i) Single Region Fault Model (sRFM) where faults are confined to a single region [4], and (ii) Multiple Region Fault Model (mRFM) where faults are confined to k regions for some specified k [12].

Formally, in sRFM, the *single-region-based (node) connectivity* of graph G with a specified definition of region R, $s\kappa_R(G)$, is defined as follows: Suppose that $\{R_1, \ldots, R_k\}$ is the set of all possible regions of the graph G. Consider a k-dimensional vector T whose i-th entry, $T[i]$, indicates the number of nodes in region R_i whose failure will disconnect the graph G. If the graph G remains connected even after the

failure of all nodes of the region R_i then $T[i]$ is set equal to ∞. The *region-based connectivity* of a graph G with region R, is then computed as follows:

$$s\kappa_R(G) = \min_{1 \leq i \leq k} T[i]$$

In mRFM, the *multi-region-based (node) connectivity* of graph G with a specified definition of region R, $m\kappa_R(G)$, is defined as the minimum number of regions whose removal (i.e., removal of all nodes in the regions and edges incident on them) will disconnect the graph.

Polynomial time algorithms to compute region-based connectivity in sRFM was presented in [4]. The NPMT contains an implementation of this algorithm that can be used to compute the Region-based Connectivity for a given network topology.

Region-based Component Decomposition Number (RBCDN) Metric Computation

Proposed by the authors of [13], the Region-Based Component Decomposition Number, or RBCDN of graph $G = (V, E)$ with a specified definition of region R is defined the following way: Suppose that $\{R_1, \ldots, R_k\}$ is the set of all possible regions of the graph G. Consider a k-dimensional vector C whose i-th entry, $C[i]$, indicates the number of connected components in which G decomposes when all entities in R_i fails. RBCDN of a graph G with region R is defined as follows:

$$\delta_R(G) = \max_{1 \leq i \leq k} C[i]$$

RBCDN as a metric provides a insight into the worst case scenario on how fragmented a network can become in the presence of a spatially correlated fault. In [13] the authors propose techniques to compute the RBCDN and the NPMT provides an implementation of this algorithm that can be used on user selected network topologies.

Region-based Smallest/Largest Component Size Metric Computation

The Region-Based Smallest (Largest) Component Size, or RBSCS/RBLCS was proposed in [14], and is defined for a graph $G = (V, E)$ with a specified definition of region R, as follows: Suppose that $\{R_1, \ldots, R_k\}$ is the set of all possible regions of the graph G. Consider a k-dimensional vector C_S (C_L) whose i-th entry, $C_S[i]$ ($C_L[i]$), indicates the size of the smallest (largest) connected component in which G decomposes when all nodes in R_i fails. The RBSCS $\alpha_R(G)$ and RBLCS $\beta_R(G)$ of graph G with region R is defined as:

$$\alpha_R(G) = \min_{1 \leq i \leq k} C_S[i] \text{ and } \beta_R(G) = \min_{1 \leq i \leq k} C_L[i]$$

The RBLCS and RBSCS metrics provide insights on how well a network's performance degrades in the presence of region-based faults. Depending on the needs of graceful performance degradation, networks designers may choose to design networks that have a small value of RBCDN ($\delta_R(G)$) and a high value of either RBLCS ($\alpha_R(G)$) or RBSCS ($\beta_R(G)$). The NPMT allows the user to compute the RBLCS and RBSCS metrics for a chosen network topology.

B. Distinct Regions Computation Module

It may be noted that all the previously defined metrics operate on a given graph and a set of regions. Thus, there is a need for techniques to compute the set of regions, given a network and some fault specification. In [14], given a graph G's layout on a two-dimensional plane and a fault radius r, the authors provide a polynomial time algorithm to compute all *distinguishable* or *distinct* circular regions with radius r. Two fault regions are considered *indistinguishable* if they contain the same set of links and nodes. The authors considered both wired networks, where nodes and edges can be part of a failure region, and wireless networks, where only nodes can be part of a failure region. It was shown in [14] that the number of distinct regions in wireless and wired networks are $O(n^2)$ and $O(n^4)$ respectively, and that all distinct regions can be computed in $O(n^6)$ time, where n is the number of nodes.

The NPMT is bundled with an implementation of the technique outlined in [14]. Given a network topology and a fault radius, the NPMT can compute all distinct regions of the network which can then be used by other modules of the NPMT, such as the Metric Computation Module and the Region-disjoint Path Computation Module (discussed next).

C. Region-disjoint Paths Computation Module

For a graph $G = (V, E)$, a set of region-disjoint paths $\mathcal{P}$ between a source node s and destination node d with a specified definition of region R, is defined as follows: Suppose that $\{R_1, \ldots, R_k\}$ is the set of all possible regions of graph G and path $\mathcal{P}_u \in \mathcal{P}$ contains a set of nodes and edges from G such that $\mathcal{P}_u$ forms a path from s to d, $\{s, d\} \in V$. Then, for every pair of paths $\{\mathcal{P}_u, \mathcal{P}_v\} \in \mathcal{P}, u \neq v$, $\mathcal{P}_u$ and $\mathcal{P}_v$ are region-disjoint, i.e. there is no region in R that both the paths traverse. Formally, region-disjoint paths are defined as follows, for all $i = 1, \ldots, k$:

$$|(\mathcal{P}_u \cap R_i) \cap (\mathcal{P}_v \cap R_i)| = 0, \forall \{\mathcal{P}_u, \mathcal{P}_v\} \in \mathcal{P}, u \neq v$$

Although region-disjoint path computation has been addressed in [8], the authors consider a model where faults do not cause edges to fail unless a failed edge is associated with a failed node. In this model an edge cannot fail on it's own and can only fail when one of the nodes incident on the edge fails. This assumption is considerably restrictive and possibly unusable for designers of larger networks where spatially correlated faults can affect nodes and edges independently. In order to overcome this limitation the NPMT supports computation of region-disjoint paths in the presence of circular faults using an Integer Linear Program (ILP) that doesn't presuppose any such restrictions. The NPMT is capable of computing two region-disjoint paths from given source and destination nodes such that the sum of lengths of the two paths is minimum. Also, as the source (destination) node is part of a region that is traversed by both paths (as both paths have the same staring and ending points), no region disjoint path may exist. To accommodate this situation the NPMT accommodates the use of *no-fault zones* – a circular area around the source and destination nodes that is immune to faults. Future extensions of this module include computing more than two paths, and including other selection criteria such as minimizing the maximum path length.

D. Region-based Fault Tolerant Distributed File Storage Module

In the preceding discussions the importance of a node in keeping the network connected is emphasized, however,

individual nodes can also act as data stores of the network and the removal of a node from a network (due to a region-based fault), may not only cause connectivity losses, but also data losses. To address such data loss risks, distributed storage techniques are often employed that enhances data survivability in the presence of faults. One such technique is redundancy, such as by (i) storing multiple copies of the entire file, or (ii) storing different fragments of the same file at different nodes in the network. In the popular (N, K), $N \geq K$ file distribution scheme, from a file F of size $|F|$, N segments of size $|F|/K$ are created in such a way that it is possible to reconstruct the entire file by accessing any K segments. For such a reconstruction scheme to work, it is essential that the K segments of the file are stored in nodes that are connected to each other in the network. However, in the event of failures, the network may become disconnected (i.e., split into several connected components) and K segments may not be accessible in the residual network to reconstruct the file F.

From the context of data survivability in the presence of spatially correlated faults in networks, the NPMT supports a "Region-based Distributed File Storage Module" that implements an algorithm proposed in [11] that ensures that: (i) even when the network is fractured into disconnected components due to a region-based fault, at least one of the largest components will have access to at least K distinct file segments with which to reconstruct the entire file, and (ii) the total storage requirement is minimized. As of writing this paper, this module is currently under development and will be part of the NPMT upon its completion.

E. Robust Multi-layer Interdependent Network Design Module

In today's world, a multitude of heterogeneous interconnected networks form a symbiotic ecosystem that supports all of the economic, political and social aspects of human life. For example, the critical infrastructures of the nation such as the power grid and the communication network are highly interdependent on each other, and any adverse effects on one network can affect the other network. Thus, isolated network analysis is no longer sufficient to design and operate such interconnected and interdependent network systems.

Recognizing this need for a deeper understanding of the interdependency in such multi-layered network systems, significant efforts have been made by the research community in the last few years, and accordingly, a number of analytical models have been proposed to analyze such interdependencies [15-17]. However, most of these models are simplistic and fail to capture the complex interdependencies that may exist between entities of the power grid and communication networks. To overcome the limitations of existing models, the authors of [18] have proposed an *Implicative Interdependency Model* that is able to capture such complex interdependency. Utilizing this model, several problems on multi-layer interdependent networks have been studied, such as (i) identification of the $\mathcal{K}$ most vulnerable nodes [18], (ii) root cause analysis of failures [19], (iii) the entity hardening problem [20], (iv) the smallest pseudo-target set identification problem [21], and (v) the robustness analysis problem [22].

This module will support multi-layer network interdependency modeling using the *Implicative Interdependency Model*,

and analysis of multi-layer networks using the techniques proposed in [18-22]. The module is currently under development and will be part of the NPMT upon its completion.

F. Module for Progressive Recovery from Region-based Failures

With this module, the NPMT addresses post-fault recovery techniques in the aftermath of region-based faults on multi-layer interdependent networks. To restore an interdependent network system from a post-fault scenario to its pre-failure state, all the faulty network entities (nodes/edges) have to be repaired or replaced. However, resource limitations may prevent simultaneous restoration of all failed units of the network. Accordingly, the failed units have to be restored in a *sequenced* manner. As each network entity in its *operational* state adds some *utility value* to the interdependent network system, when a unit recovers from a *failed state* to an *operational* state, the unit starts providing some "benefit" to the system. Since different units have different utility values to the system, the sequence in which the failed units are restored is important as the recovery sequence determines the cumulative system utility during the recovery process.

As discussed in Section II-E, the *Implicative Interdependency Model* provides a powerful technique for modeling dependencies in multi-layer interdependent networks. Using this model the authors of [23] studied the progressive recovery problem in interdependent networks with the objective of maximizing system utility during the system recovery process. This module implements the progressive recovery algorithm of [23], and can be used to sequence recovery of network entities from a post-fault to a pre-fault network state that maximizes system utility during the recovery process. The module is currently under development and will be part of the NPMT upon its completion.

III. Architecture and System Capabilities

In this section we first outline the system architecture, and then discuss the different capabilities of the NPMT.
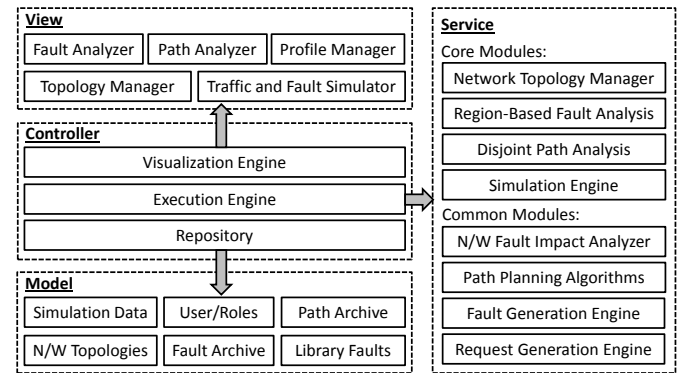
A. System Architecture

Fig. 2: The NPMT High-Level Architecture

The NPMT is implemented as a web-application that allows the user to remotely connect and operate the tool from a browser. The web-application follows the standard three-tier architecture and has a client tier, application tier, and database

(a)   (b)

Fig. 3: (a) Topology Manager – create, edit and manage network topologies, (b) Fault Analyzer – generic fault analysis, metric computations

tier. The tool has been developed following the Model-View-Controller (MVC) design pattern. Figure 2 outlines the high level architecture and some of the components of the tool.

The tool is currently accessible from Arizona State University's WAN, and runs from our testbed server. The tool's web-application is deployed on an Apache Tomcat 7 instance, and the repository used is MySQL. The application tier business logic for operations on network topologies, such as Region-Based Fault Analysis and Region Disjoint Path Analysis, are implemented in Java. Additional packages and libraries, such as IBM ILOG CPLEX Optimization Studio libraries (required for solving Integer Linear Programs), are setup and made available on the testbed server. Our testbed server is a 64-bit Intel Core 2 Quad Core (2.66 GHz) system with 8 GB of RAM running an Ubuntu 14.04 instance.

B. System Capabilities

The NPMT is designed to be used by following a three step workflow comprising of (i) Network Creation, (ii) Network Analysis, and (iii) Network Simulation. Accordingly, the individual features and the executable modules of the NPMT are bundled around these three workflow steps. The following list enumerates the current high-level features of the tool and the corresponding workflow step that each feature emulates:

1) Topology Management (Network Creation)
2) Fault Analysis (Network Analysis)
3) Path Analysis (Network Analysis)
4) Traffic and Fault Impact Simulation (Network Simulation)

Each of the above features are accessible from a tabbed interface of the tool and can be navigated to from any part of the web-application. In the following subsections we discuss each of the features and provide a brief functional overview.

Topology Management

Network Creation is the first step of the NPMT workflow and the *Topology Manager* interface allows the user to create, edit, save and delete network topologies. The Topology Manager presents the user with a geographical map interface that she can interact with to manage network topologies. The displayed

map tiles are rendered from OpenStreetMap [24]. The NPMT uses the OpenLayers API to support an user interactive map interface.

To create the topology and place nodes and edges on the map, the user can either point-and-click on the map itself, or can type in specific latitude and longitude coordinates and then proceed to add the network entity. Capacities for each edge (in Gigabits per second), can also be specified during the edge creation process. Once a network topology is created, the topology must be saved to be used for Network Analysis and Network Simulation. The topologies are saved on the NPMT server and can be loaded back into the Topology Manager to edit any entity or attribute of the network.

Figure 3(a) shows a screen grab of the Topology Manager. As seen in the figure, the map interface is on the right and the user interact-able menu is on the left. The user can click on the map to to add nodes and edges, or can alternatively type in the latitude and longitude coordinates in the input fields available on the menu. The menu also lists the nodes and edges that are part of the topology. Selecting an edge or node from these lists highlights the network entity on the map (in yellow), and the user can then proceed to remove the entity from the network if necessary. The displayed map overlays can be toggled from a dropdown menu available on the map (in blue in Figure 3(a)). Finally, as seen in Figure 3(a), options for saving, loading, and deleting topologies are available to the user directly below the displayed map's dimensions.

Fault Analysis

Once network topologies are created from the Topology Manager, the *Fault Analyzer* can be used to analyze the created networks for their *resilience* in the presence of spatially correlated faults. In the context of the NPMT, network resiliency is measured by how well the network performs when benchmarked against the metrics outlined in Section II-A. It may be noted that the metrics of Section II-A emphasize resilience from the aspect of connectivity in the presence of a spatially correlated fault. For example, the more number of disconnected components a network has due to a fault, the worse is the network's resilience (as captured by the metric RBCDN). It
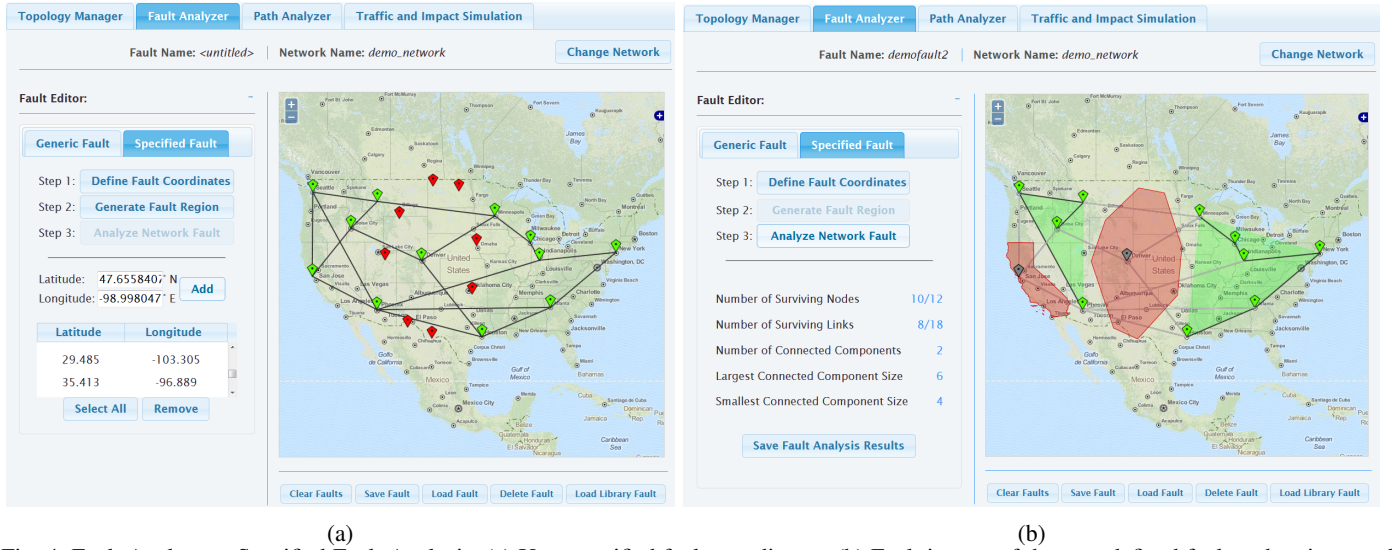
Fig. 4: Fault Analyzer - Specified Fault Analysis. (a) User specified fault coordinates, (b) Fault impact of the user defined fault and an imported library fault (coordinates for the state of California, USA)

may be noted that, for the purpose of this analysis the NPMT assumes that any network entity (nodes/edges), that fall within the fault area are all rendered inoperable, i.e. the fault model is deterministic and if a network entity falls within the fault region, it necessarily fails. To carry out this analysis, the user first selects a network topology and can then choose to either perform a *generic fault* analysis, or a *specified fault* analysis. These analyses are described below.

Generic Fault Analysis: In the generic fault analysis, for a selected network topology, the user specifies a fault feature and the tool computes the values of the individual metrics listed in Section II-A. The NPMT can generically analyzes circular faults, and the supported fault feature is the fault radius r.

As shown in Figure 3(b), the user can specify the fault radius r from the left menu. The tool then performs the generic fault analysis by (i) computing all the *distinct* regions with radius r using the techniques implemented in the module "Distinct Regions Computation Module" (Section II-B), and (ii) computes the individual metrics using the techniques implemented in the module "Region-Based Fault Metrics Computation Module" (Section II-A). The results are subsequently reported back to the user. For the network selected in Figure 3(b) and radius $r = 500$ km., the computed Region-based Component Decomposition Number (RBCDN) is 2, the Region-based Largest Component Size (RBLCS) is 9 and the Region-based Smallest Component Size (RBSCS) is 1. Finally, the number of distinct regions computed is 112.

As shown in Figure 3(b), the user is also presented with sample worst case fault scenarios where a distinct fault causes the network to fragment into the same number of components as the RBCDN. Selecting one of the listed faults updates the displayed network with the fault's impact. In Figure 3(b) the fault centered at $36.249°N, -85.696°E$ is selected. The nodes and edges rendered inoperable by the fault are grayed out, while the surviving nodes and edges are shown in green and black respectively. The connected components in the fragmented network are highlighted by a light-green region. In this example, the loss of the two grayed out nodes causes the network to fragment into two disconnected components: one

with 9 components, and the other with 1 component. Options for saving the analysis results are available from the menu.

Specified Fault Analysis: In the specified fault analysis, the user can provide the exact coordinates of one or more faults and visualize the impact of these faults on the selected network. The user has the option to save and load faults to visualize the impact of a fault on different networks. The NPMT also comes bundled with a set of *library faults* that the user can choose from to simulate fault impact on a network. The current set of library faults consist of the coordinates of the 50 states of the USA. The inclusion of a fault library in the NPMT is to provide the user with pre-defined fault scenarios based on known fault patterns, faults centered at a target of interest, or recorded faults. For example, fault impact zones of Level 4 hurricanes such as hurricane Katrina or hurricane Sandy.

As shown in Figure 4(a), to specify the exact coordinates of the fault region the user can either type in the exact coordinates of the fault region coordinates, or can click on the map to add such coordinates. The user also has the option for importing library faults. Once all the fault regions are defined, the NPMT can simulate the impact of the fault on the selected network. In Figure 4(b), apart from the user specified fault region, the boundary of the state of California has been imported from the fault library and the selected network has been analyzed for these two fault regions. The updated map shows the impacted nodes and edges in gray, while the operable nodes and edges are shown in green and black respectively. The connected components are shown with a green region. As seen in Figure 4(b) the menu displays impact statistics such as, the number of surviving nodes/edges and the number of connected components. The user is provided with the option to save the analysis results for later reference, and also the option to save the defined fault regions for later use.

Path Analyzer

The *Path Analyzer* allows the user to analyze the network by computing paths between source and destination nodes that provide protection against spatially correlated faults. As in the Fault Analyzer, the Path Analyzer allows the user to specify a fault feature, and the tool then proceeds to compute paths
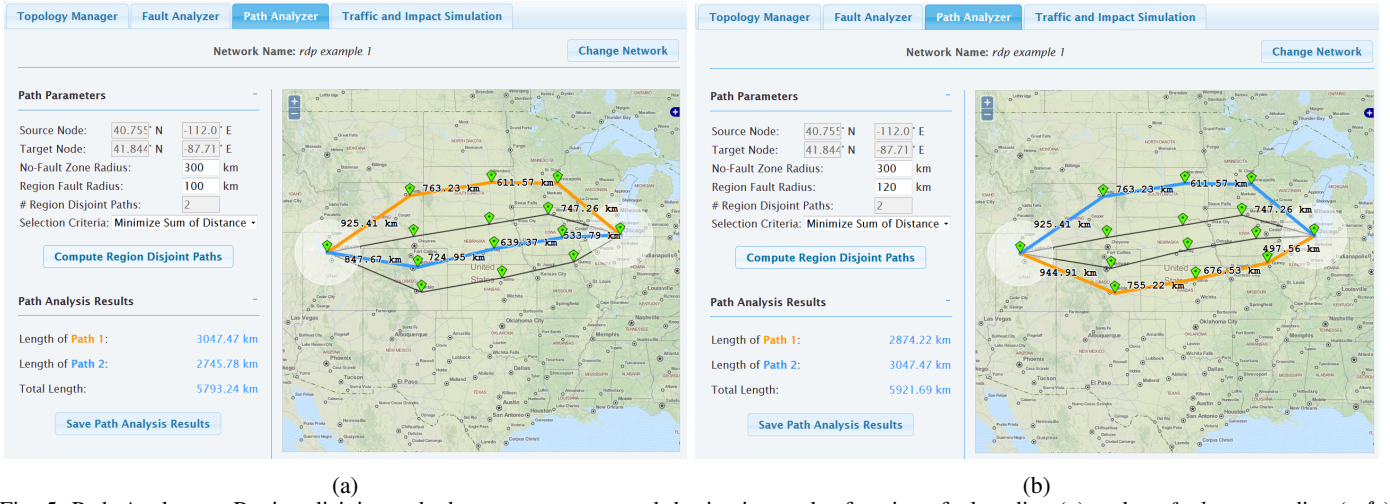
(a)   (b)

Fig. 5: Path Analyzer - Region disjoint paths between a source and destination nodes for given fault radius (r) and *no-fault zone* radius (nf_r) (a) $r = 100 \ km.$, $nf_r = 300 \ km.$ (b) $r = 120 \ km.$, $nf_r = 300 \ km.$

between a given source and destination node pair such that (i) at least one of the paths survive in the presence of one or more spatially correlated faults, and (ii) satisfy some other network resource constraint.

In the current version of the tool the faults considered are circular faults and the supported fault feature that can be specified by the user is the fault radius r. The number of spatially correlated faults considered for path analysis is one, and the number of paths computed is two, i.e. the NPMT computes two paths such that if a single circular fault with radius r occurs anywhere in the network, at least one of the two paths computed will not be affected by the fault. The network resource constraint supported is that the sum of lengths of the two paths computed must be minimum.

It may be noted that a single fault can also render inoperable either the source node, or the destination node, or both, and thus there always exists a fault region that affects all paths computed and no region disjoint paths can exist such that at least one path remains immune to the fault. To accommodate this case when the source and/or destination nodes themselves are part of the fault region, the NPMT supports a *"No-Fault Zone"* parameter. The user can specify a no-fault zone radius nf_r for both the source and destination nodes that reserves two circular areas with radius nf_r centered at the source and destination nodes such that network entities, or parts of a network entity (such as an edge segment), that falls within this no-fault zone are immune to faults.

Figures 5(a) and 5(b) show screen grabs of the path analyzer computation for different input values of fault radius (r). The no-fault zone set to a radius of $nf_r = 300$ km. and is shown as a white circular region centered at the source and destination nodes. The computed paths are shown in orange and blue, and the lengths of each of these to paths are reported in the left menu. The effect of the path selection criteria, i.e. the sum of the lengths of the two paths must be minimum, is also visible in Figures 5(a) and 5(b). In Figure 5(a) when $r = 100$ km., the sum of lengths of the two paths is 5793.24 km., however in Figure 5(b) increasing r to 120 km. the previously computed paths are no longer feasible as a region fault exists that can impact both these paths. Hence, new paths are computed and the sum of the new lengths is 5921.69 km.

Traffic and Fault Impact Simulation

For a selected network, the *Traffic and Impact Simulator* allows users to generate traffic and faults to analyze the impact of faults on a load bearing network. To perform this analysis, a simulation schedule consisting of bandwidth requests and faults is generated by the NPMT using user provided simulation parameters. Parameters such as total number of time steps in the schedule, total number of requests in the schedule, minimum/maximum request bandwidth and minimum/maximum request hold times can be specified by the user. The source and destination nodes for each request can be generated randomly, or can be user specified. For introducing faults in the schedule, the user can specify the number of faults to introduce and can either specify the exact fault coordinates, or introduce random circular faults from the set of all possible *distinct* circular faults for a specified fault radius. Time intervals of the faults can be user specified, or can be randomly generated by the NPMT. Using the request and fault settings, the NPMT then generates a time stepped simulation schedule of requests and faults. Once the schedule is finalized, the user can specify the algorithm to be used in the simulation to route requests from source and destination nodes, and then proceed to run the simulation.

As shown in the screen grabs of Figures 6(a) and 6(b), the left menu of the Traffic and Impact Simulator contains the fault and simulation parameter fields that can be used to generate the schedule and run the simulation. The tables below the map's dimensions allow the user fine grained control over the requests and faults that will be simulated. Once the simulation is complete, for each time interval the network state can be visualized from the "Event Simulation Results" table. The user can click on a row of this table to visualize the network state on the map for that specific time interval. The user can also "play" the simulation results and the NPMT will iterate over all the time steps and update the map with the network state at each step. In Figures 6(a) and 6(b) the impact of a fault and the corresponding response of the network is shown. In Figure 6(a) the network is fault free, but in Figure 6(b) a fault is introduced and an edge is rendered inoperable. It can be seen that the red and yellow flows of Figure 6(a) are impacted by the fault, however, as bandwidth is available, in Figure 6(b) the flows are rerouted in response to this fault.
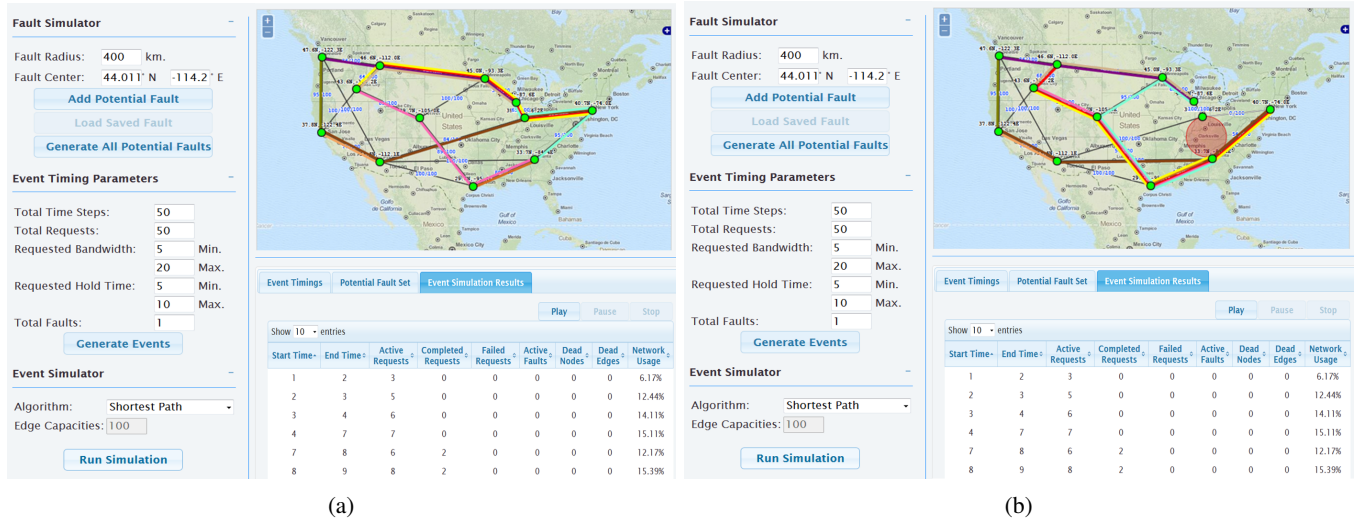
Fig. 6: Traffic and Fault Impact Simulator (a) Pre-Fault network state, (b) Post-Fault network state – rerouted red and yellow flows)

IV. CONCLUSION

In this paper we presented a summary of the work done towards developing a Network Planning and Management Tool (NPMT), intended to support design and analysis of single layer and multi-layer networks in the presence of spatially correlated faults. We highlighted that the NPMT is particularly suitable for planning and design of critical infrastructures. We described the underlying novel concepts that have been developed to enhance robustness of networks in presence of region based faults, and then described how those concepts have been incorporated into the NPMT. The goal of this paper was to bring to the attention of the networking research community, and to the audience of the workshop on DRCN in particular, about the existence of NPMT as a tool that consolidates a large body of work on spatially correlated faults. To the best of our knowledge no such tool is available today that supports planning and designing of single layer and multi-layer networks in the presence of spatially correlated faults.

REFERENCES

[1] NEXT Lab, Arizona State University. The Network Planning and Management Tool. [Online]. Available: http://netsci.asu.edu/networktool/
[2] R. Diestel, *Graph Theory*. Springer, 2005.
[3] E. Ganesan and D. K. Pradhan, "The Hyper-deBruijn Networks: Scalable Versatile Architecture," *IEEE Transactions on Parrallel and Distributed Systems*, vol. 4, no. 9, September 1993.
[4] A. Sen, B. H. Shen, L. Zhou, and B. Hao, "Fault-tolerance in Sensor Networks: A New Evaluation Metric," in *Proceedings of IEEE Infocom*, Barcelona, Spain, April 2006, pp. 1–12.
[5] S. Neumayer and E. Modiano, "Network reliability with geographically correlated failures," in *INFOCOM, 2010 Proceedings IEEE*. IEEE, 2010, pp. 1–9.
[6] Y. Cheng, M. T. Gardner, J. Li, R. May, D. Medhi, and J. P. Sterbenz, "Optimised heuristics for a geodiverse routing protocol," in *10th International Conference on the Design of Reliable Communication Networks (DRCN), 2014*, 2014, pp. 1–9.
[7] P. Agarwal, A. Efrat, S. Ganjugunte, D. Hay, S. Sankararaman, and G. Zussman, "The resilience of wdm networks to probabilistic geographical failures," in *Proceedings of IEEE INFOCOM*, 2011.
[8] S. Trajanovski, F. Kuipers, A. Ilic, J. Crowcroft, and P. Van Mieghem, "Finding critical regions and region-disjoint paths in a network," *IEEE/ACM Transactions on Networking*, vol. 23, no. 3, pp. 908–921, 2015.
[9] S. Banerjee, A. Das, A. Mazumder, Z. Derakhshandeh, and A. Sen, "On the impact of coding parameters on storage requirement of region-based fault tolerant distributed file system design," in *Computing, Networking and Communications (ICNC), International Conference on*. IEEE, 2014, pp. 78–82.
[10] A. Mazumder, A. Das, C. Zhou, and A. Sen, "Region based fault-tolerant distributed file storage system design under budget constraint," in *Reliable Networks Design and Modeling (RNDM), 2014 6th International Workshop on*. IEEE, 2014, pp. 61–68.
[11] A. Sen, A. Mazumder, S. Banerjee, A. Das, C. Zhou, and S. Shirazipourazad, "Region-based fault-tolerant distributed file storage system design in networks," *Networks*, Wiley Online Library, 2015.
[12] A. Sen, S. Murthy, and S. Banerjee, "Region-based connectivity-a new paradigm for design of fault-tolerant networks," in *High Performance Switching and Routing, 2009. HPSR 2009. International Conference on*. IEEE, 2009, pp. 1–7.
[13] S. Banerjee, S. Shirazipourazad, P. Ghosh, and A. Sen, "Beyond connectivity-new metrics to evaluate robustness of networks," in *High Performance Switching and Routing (HPSR), 2011 IEEE 12th International Conference on*. IEEE, 2011, pp. 171–177.
[14] S. Banerjee, S. Shirazipourazad, and A. Sen, "Design and analysis of networks with large components in presence of region-based faults," in *International Conference on Communications (ICC)*. IEEE, 2011.
[15] S. V. Buldyrev, R. Parshani, G. Paul, H. E. Stanley, and S. Havlin, "Catastrophic cascade of failures in interdependent networks," *Nature*, vol. 464, no. 7291, pp. 1025–1028, 2010.
[16] J. Gao, S. V. Buldyrev, H. E. Stanley, and S. Havlin, "Networks formed from interdependent networks," *Nature Physics*, vol. 8, no. 1, pp. 40–48, 2011.
[17] J.-F. Castet and J. H. Saleh, "Interdependent multi-layer networks: Modeling and survivability analysis with applications to space-based networks," *PloS one*, vol. 8, no. 4, p. e60402, 2013.
[18] A. Sen, A. Mazumder, J. Banerjee, A. Das, and R. Compton, "Identification of k most vulnerable nodes in multi-layered network using a new model of interdependency," in *NetSciCom Workshop (INFOCOM WKSHPS), Conference on Computer Communications*. IEEE, 2014, pp. 831–836.
[19] A. Das, J. Banerjee, and A. Sen, "Root cause analysis of failures in interdependent power-communication networks," in *Military Communications Conference (MILCOM), 2014 IEEE*. IEEE, 2014, pp. 910–915.
[20] J. Banerjee, A. Das, C. Zhou, A. Mazumder, and A. Sen, "On the entity hardening problem in multi-layered interdependent networks," in *WIDN Workshop (INFOCOM WKSHPS), 2015 IEEE Conference on Computer Communications*. IEEE, 2015, pp. 648–653.
[21] A. Das, C. Zhou, J. Banerjee, and A. Sen, "On the smallest pseudo target set identification problem for targeted attack on interdependent power-communication networks," in *Military Communications Conference (MILCOM), IEEE (To appear)*. IEEE, 2015.
[22] J. Banerjee, C. Zhou, A. Das, and A. Sen, "On robustness in multi-layer interdependent networks," in *Conference on Critical Information Infrastructures Security (CRITIS) (To appear)*. Springer, 2015.
[23] A. Mazumder, C. Zhou, A. Das, and A. Sen, "Progressive recovery from failure in multi-layered interdependent network using a new model of interdependency," in *Conference on Critical Information Infrastructures Security (CRITIS)*. Springer, 2014.
[24] OpenStreetMap Contributors. OpenStreetMap. [Online]. Available: www.openstreetmap.org