RAPTOR: a network tool for mitigating the impactof spatially correlated failures in infrastructure networks

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

Noname manuscript No.
(will be inserted by the editor)

RAPTOR: A Network Tool for Mitigating the Impact of Spatially Correlated Failures in Infrastructure Networks

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

Received: date / Accepted: date

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 geographic 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 RAPTOR, an advanced *Network Planning and Management Tool* 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 tool's capabilities, but also extend these capabilities to heterogeneous interdependent networks such as communication, power, water and satellite networks.

Keywords Spatially Correlated Faults · Region based Faults · Geographically Correlated Faults · Fault Tolerant Network Design · Network Robustness and Resilience · Network Tool

This work was supported in part by the NSF grant 1441214, and by grant HDTRA1-14-C-0015 from the U.S. Defense Threat Reduction Agency.

A. Das (✉), A. Sen
School of Computing, Informatics and Decision System Engineering, Arizona State University, Tempe, AZ 85287, USA
Tel.: +1-480-965-6153
E-mail: arun.das@asu.edu (✉), asen@asu.edu

C. Qiao
Department of Computer Science and Engineering, SUNY at Buffalo, Buffalo, NY 14201, USA,
E-mail: qiao@computer.org

N. Ghani
Department of Electrical Engineering, University of South Florida, Tampa, FL 33620, USA,
E-mail: nghani@usf.edu

N. Mitton
Inria, 40 Avenue Halley, 59650 Villeneuve D'ASCQ, France, E-mail: nathalie.mitton@inria.fr

1 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 by humans (nuclear attack or conventional weapon attack over a large geographical area). The characteristic of such *spatially correlated*, or *region based* faults is that they are *massive* but *localized* faults. As noted by the authors in [26], *network connectivity* [14] as a metric for evaluating the *fault-tolerance capability* of a network is inadequate for spatially correlated failures as it does not capture the characteristics of such failures. It may be noted that although the authors of [26] consider wireless sensor networks in their study, however, their reasoning for such an inadequacy is applicable for macro-scale nation-wide backbone networks as well. A primary point of difference between the two networks is that in wireless sensor networks a spatially correlated fault may not cause a "wireless" edge to fail independent of any of its incident nodes, i.e., for an edge to fail in a wireless setting at least one of its incident nodes must necessarily fail. On the other hand, in "wired" networks, such as in backbone networks, a fault can cause an edge to fail independent of its incident nodes, for instance, when a tropical thunderstorm causes damage to underground optical fibers without necessarily causing damage to the end-points of the fiber link. In spite of such differences, the justifications provided by the authors of [26] on the inadequacy of network connectivity as a metric for network fault-tolerance is equally applicable to large backbone networks. For instance, network connectivity as a metric ignores the *locality* of the fault, i.e., the faulty nodes (edges) may be close or far away from each other. Also, connectivity as a metric does not capture important structural properties of the network such as the *number*, or *size* of the *connected components* [14] into which a network disintegrates when the number of failed nodes (edges) *exceeds* the node (edge) connectivity of the network.

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*, 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 RAPTOR, an advanced network planning and management tool [19], 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 tool. The goal of this paper is to bring to the attention of the networking research community the existence of RAPTOR 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 region based faults. To this effect, this paper's contribution does not lie in new analytical findings, but in service to the networking research community.

The tool described in this paper is intended to support design and analysis of single layered and multi-layered interdependent heterogeneous networks. In this context, RAPTOR is particularly suitable for planning and design of critical infrastructures. For

example, from the single network layer perspective, RAPTOR can enable backbone communication network providers, such as AT&T, Sprint, Qwest and Level 3 Communications, to (i) identify the most vulnerable parts of their network against region based faults, and (ii) reinforce the network with least cost to reduce, or eliminate the threat of network disruption due to a region based fault. From a multi-layer perspective, RAPTOR can be used for design and analysis of smart cities, where heterogeneous networks such as power, communication, water, and gas distribution networks form a complex interdependent ecosystem where failures in one network may impact another. For instance, 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 RAPTOR can be used by utility companies and 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. It may be noted that the preliminary work performed in developing RAPTOR was presented in [12].

Several studies in the network research community have focused on different aspects of spatially correlated or region-based faults in networks [1,4,10,17,20,24,28], 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 tool RAPTOR is intended to fill that gap and be such a platform that can incorporate the outcomes developed in studies such as [1, 4, 9, 10, 17, 20, 24, 28] into executable modules that can be integrated into RAPTOR. This will allow network designers, planners and operators to use the results of these studies in their real world operational networks. It may be noted that although the current version of the tool includes several features to aid network designers, planners and operators to better design and manage their networks, however, the available features are in no respect comprehensive to include all user required functionality. In future iterations of the tool, it is our objective to closely work with researchers and industry professionals to include additional metrics, techniques, and functionality to make RAPTOR more usable and useful to its users. In future extensions, the tool is intended to include other network resiliency analysis techniques such as the ones proposed in [16, 22, 27] that are beyond the current scope of the tool.

The rest of the paper is organized as follows: In Section 2 we present an overview of the underlying concepts and theoretical results that the tool operates on. In Section 3 we outline the capabilities of the tool, and finally in Section 4 we conclude the paper.

2 Concepts, Metrics and Solution Techniques

In this section we give a brief overview of the underlying concepts, metrics and solution techniques that the RAPTOR tool utilizes to carry out its functions. The tool is built as a modular execution engine that can execute smaller reusable modules to perform desired operations on a network topology. A modular approach allows design, development and testing of the modules to be done independently and defers integra-

tion into the tool until a module meets its functional requirements. In the following sub-sections we give an overview of the analytical foundations of these modules.

2.1 Region-Based Fault Metrics Computation Module

As outlined in Section 1, network 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 other metrics and their computation techniques have been proposed by the research community. For a given network topology, in its current version, RAPTOR can compute the following metrics pertinent to region based faults:

2.1.1 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 [26], and (ii) Multiple Region Fault Model (mRFM) where faults are confined to k regions for some specified k [25].

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: Let $\{R_1,\ldots,R_k\}$ be 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 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.

Algorithms to compute region-based connectivity in sRFM was presented in [26] and RAPTOR provides an implementation of this algorithm for analysis of user selected network topologies.

2.1.2 Region-based Component Decomposition Number (RBCDN) Metric Computation

The *Region-Based Component Decomposition Number* (RBCDN) of graph $G=(V,E)$ with a specified definition of region R is defined as follows [5]: Let $\{R_1,\ldots,R_k\}$ be 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 graph G with region R is computed as follows:

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

RBCDN as a metric provides an insight into the worst case scenario on how fragmented a network can become in the presence of a region based fault. Techniques are proposed in [5] to compute the RBCDN, and RAPTOR provides an implementation of this algorithm for analysis of user selected network topologies.

2.1.3 Region-based Smallest/Largest Component Size Metric Computation

The *Region-Based Smallest* (*Largest*) *Component Size*, or RBSCS/RBLCS was proposed in [6], and is defined for a graph $G = (V, E)$ with a specified definition of region R, as follows: Let $\{R_1, \ldots, R_k\}$ be the set of all possible regions of 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:

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

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, network 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 RAPTOR tool allows the user to compute the RBLCS and RBSCS metrics for a chosen network topology.

2.2 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 that compute the set of regions, given a network and some fault specification. In [6], 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 [6] 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 RAPTOR tool is bundled with an implementation of the technique outlined in [6]. Given a network topology and a fault radius, RAPTOR can compute all distinct regions of the network which can then be used by other modules of the tool.

2.3 Region-disjoint Paths Computation Module

For a graph $G = (V,E)$, a set of region-disjoint paths $\mathscr{P}$ 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 $P_u \in \mathscr{P}$ contains a set of nodes and edges from G such that P_u forms a path from s to d, $\{s,d\} \in V$. Then, for every pair of paths $\{P_u,P_v\} \in \mathscr{P}, u \neq v$, P_u and 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 $i = 1,\ldots,k$:

$$|(P_u \cap R_i) \cap (P_v \cap R_i)| = 0, \forall \{P_u,P_v\} \in \mathscr{P}, u \neq v$$

Although region-disjoint path computation has been addressed in [28], the authors consider a model where faults do not cause edge failure unless a failed edge is associated with a failed node. 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 RAPTOR tool 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 tool 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. Future extensions of this module include computing more than two paths, and including other selection criteria such as minimizing the maximum path length.

2.4 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 [7, 8, 15]. 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 [23] have proposed an *Implicative Interdependency Model* (IIM) that is able to capture such complex interdependency.

This module will support multi-layer interdependent network modeling utilizing IIM, and implement the techniques proposed for problems studied under the IIM setting, such as (i) identification of the K most vulnerable nodes [23], (ii) root cause analysis of failures [11], (iii) the entity hardening problem [2], (iv) the smallest pseudo-target set identification problem [13], (v) the robustness analysis problem [3], and (vi)

progressive recovery from failure in multi-layer interdependent networks [18]. It may be noted that, as of writing this paper this module is currently under development and will be part of the tool upon its completion.

3 Architecture and System Capabilities

In this section we first outline the system architecture, and then discuss the different capabilities of RAPTOR.

3.1 System Architecture

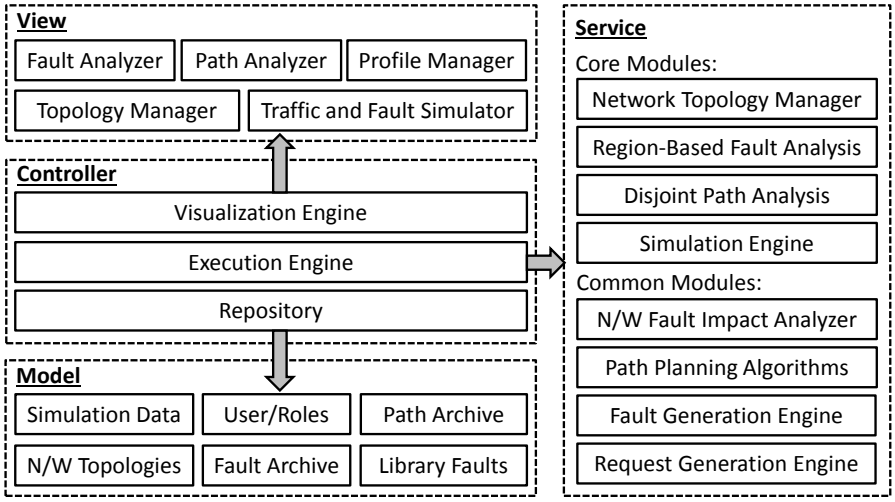

Fig. 1: High-Level Architecture of RAPTOR

RAPTOR 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 tier. The tool has been developed following the Model-View-Controller (MVC) design pattern. Fig. 1 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 the IBM ILOG CPLEX Optimization Studio (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 Ubuntu 14.04.

3.2 System Capabilities

The RAPTOR tool 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 tool are bundled around these three workflows. The following list enumerates the current high-level features of the tool and the corresponding workflows 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 is accessible from a tabbed interface and can be navigated to from any part of the application. In its current iteration, the tool supports user interaction and user input only through the graphical user interface, however, future extensions of the tool are expected to include other data interaction options such as RESTful Web Services. In the following subsections we discuss each of the tool's features and provide a brief functional overview.

3.2.1 Topology Management

The *Topology Manager* interface allows users to create, edit, save and delete network topologies. The user is presented with a geographical map interface that she can interact with to manage network topologies. The displayed map tiles are rendered from OpenStreetMap [21] and the OpenLayers API is used to support the interactive map. 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 to add the network entity. Capacities for each edge (in Gigabits per second), can also be specified during the edge creation process. Once a topology is created, it must be saved to be used for Network Analysis and Network Simulation. The topologies are saved on the server and can be retrieved to edit entities or attributes of the network.

Fig. 2 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 add nodes and edges, or can alternatively type in the latitude and longitude coordinates from the menu. The menu maintains the list of nodes and edges of the topology. Selecting an edge (node) from the list highlights the network entity on the map (in yellow), and the user can then remove or edit the entity attribute as necessary. The displayed map overlays can be toggled from a dropdown menu available on the map (in blue in Fig. 2). Options for saving, loading, and deleting topologies are available to the user directly below the displayed map. In the current

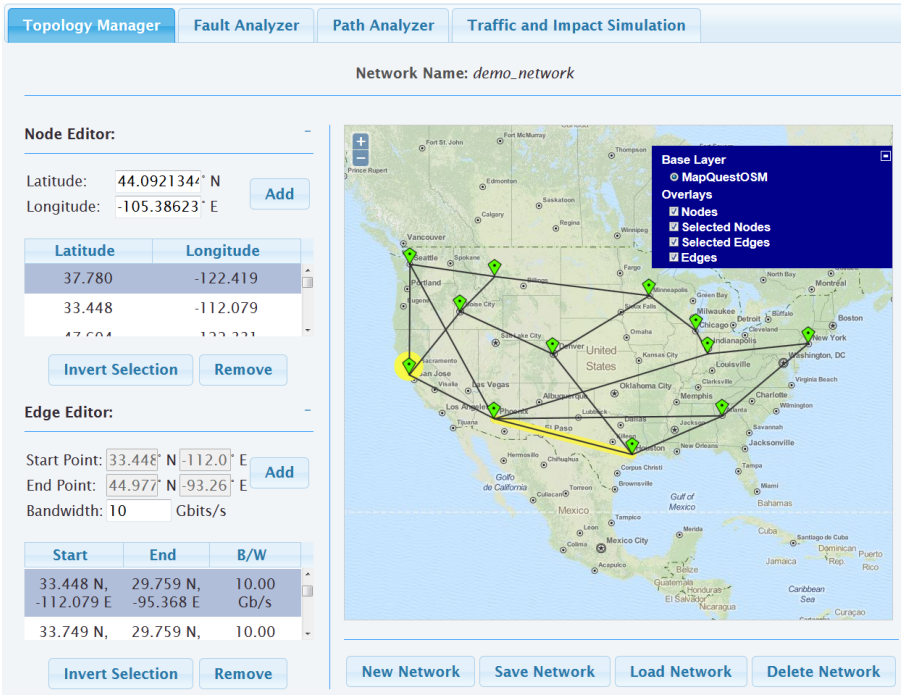

Fig. 2: Topology Manager – create, edit and manage network topologies

version of the tool random generation of topologies is not supported, however, this feature will be included in future releases for rapid topology instantiations.

3.2.2 Fault Analysis

The *Fault Analyzer* allows the user to analyze the created networks for their *resilience* in the presence of region based faults. To analyze network resilience, the tool can compute the metrics outlined in Section 2.1, and also allow the user to simulate the impact of custom region based faults on the created networks. For the purpose of this analysis, the tool assumes that any network entity (nodes/edges), that fall within the fault area are all rendered inoperable. 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 (described below).

Generic Fault Analysis: For the generic fault analysis the user can specify a fault feature, in this case circular faults with user specified fault radius r, and the tool can compute the individual metrics listed in Section 2.1.

As shown in Fig. 3, the user can specify the fault radius r from the left menu. The tool then performs the generic fault analysis by (i) computing all *distinct* regions with radius r using the techniques implemented in the module “Distinct Regions Com-

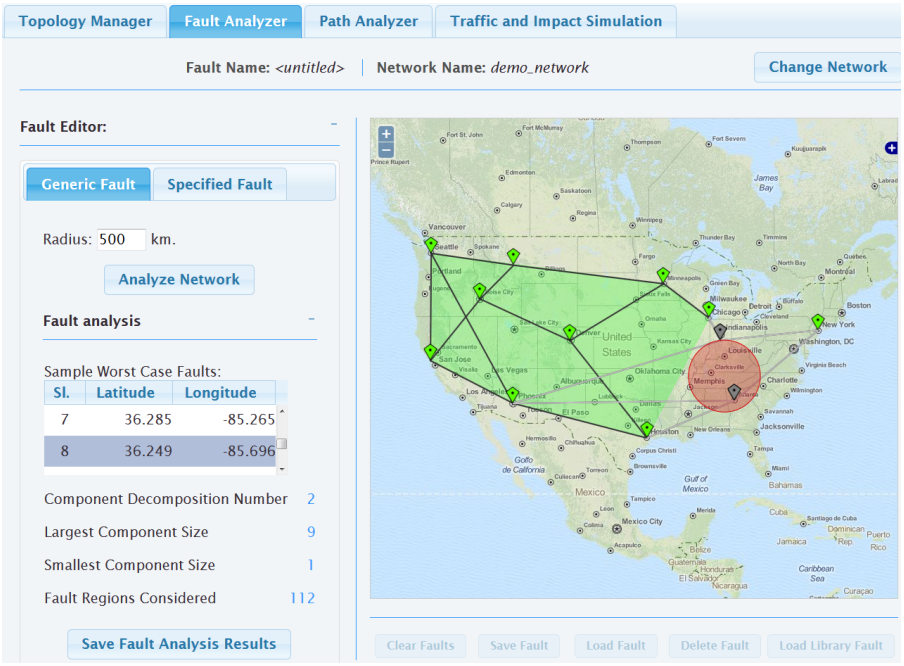

Fig. 3: Fault Analyzer – generic fault analysis, metric computations

putation Module" (Section 2.2), and (ii) computes the individual metrics using the techniques implemented in the module "Region-Based Fault Metrics Computation Module" (Section 2.1). The results are subsequently reported back to the user and are displayed on the left menu.

On the menu, 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. 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. 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 tool 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 tool is to provide the user with pre-defined fault scenarios based on known fault patterns, faults centered at a target of interest,

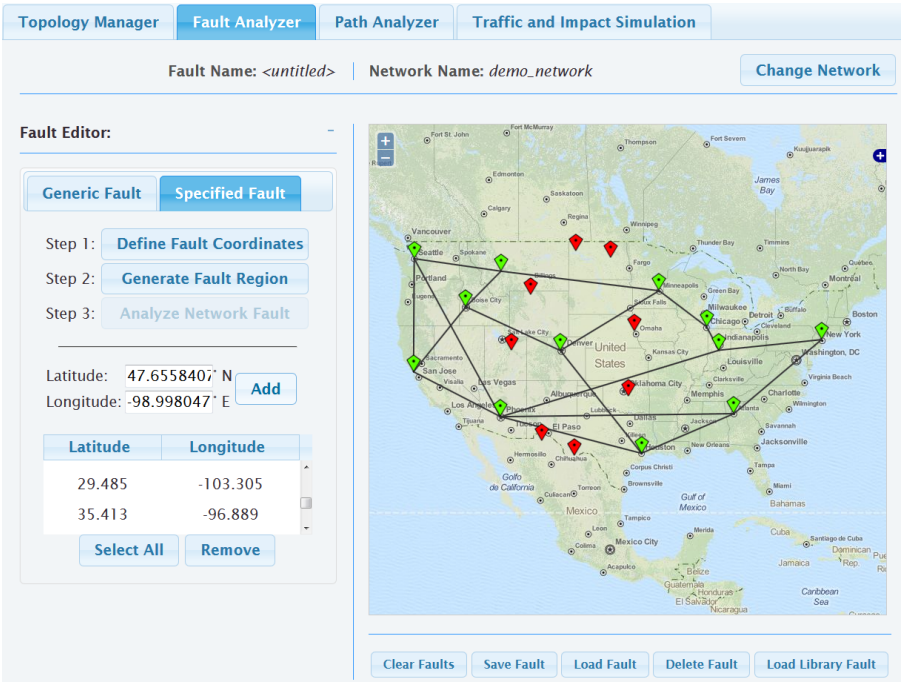

Fig. 4: Fault Analyzer – Specified Fault Analysis with user specified fault coordinates

or recorded faults, such as recorded fault impact zones of Category 4 hurricanes such as hurricane Katrina or hurricane Harvey.

As shown in Fig. 4, to specify the exact coordinates of the fault region the user can either type in the exact coordinates of the fault region, or can click on the map to add such coordinates. In Fig. 5, in addition to the user specified fault region of Fig. 4, 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 in Fig. 5 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. The menu displays impact statistics such as, the number of surviving nodes/edges and the number of connected components. The user has the option to save the analysis results, and the specified fault regions.

3.2.3 Path Analyzer

The *Path Analyzer* allows users to analyze a network by computing paths between source and destination nodes that provide protection against spatially correlated faults. In the current version of the tool, circular faults are considered for the path computation and the user can specify the fault radius r. The number of region based faults that may occur in the network is restricted to one, and the number of paths computed

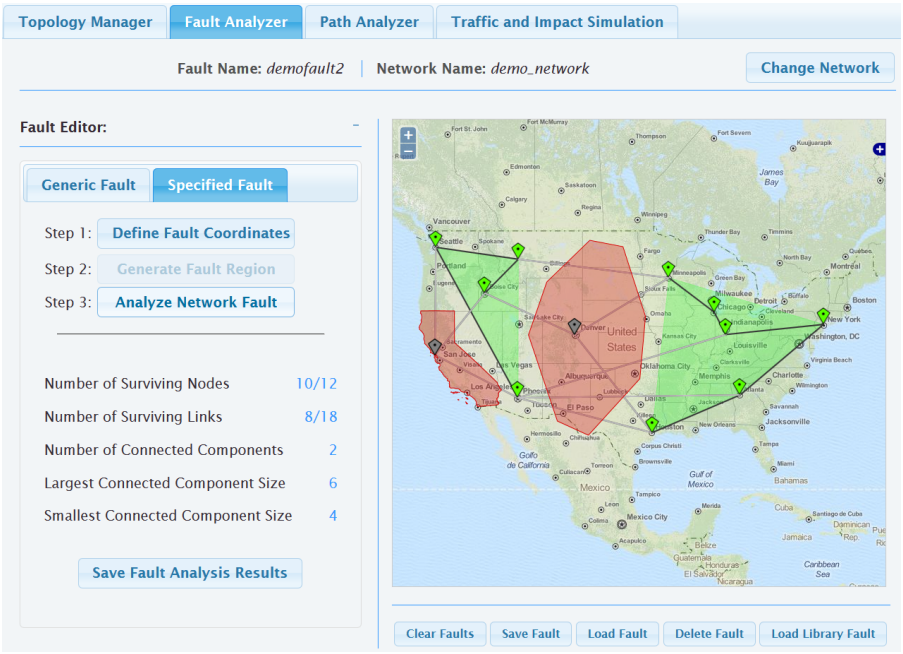

Fig. 5: Fault Analyzer – Fault impact of the user specified fault and an imported library fault (coordinates for the state of California, USA)

is two, i.e., RAPTOR computes two paths such that if a single circular fault with radius r occurs anywhere in the network, at least one of the computed paths will not be affected by the fault. The tool computes the region-disjoint paths such that the sum of lengths of the paths is minimum.

As there will always exist a fault of radius r that causes the source (destination) node to fail, to compute region-disjoint paths, the tool allows a circular *"No-Fault Zone"* centered at the source (destination) node. The user can specify a no-fault zone radius nf_r, and any network entities, or parts of a network entity (such as an edge segment), that fall within this no-fault zone are made immune to faults.

Fig. 6 and Fig. 7 show screen grabs of the path analyzer computation for different input values of fault radius r. The no-fault zone is 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 the 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 Fig. 6 and Fig. 7. In Fig. 6 when $r = 100$ km., the sum of the path lengths is 5793.24 km., however in Fig. 7 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.

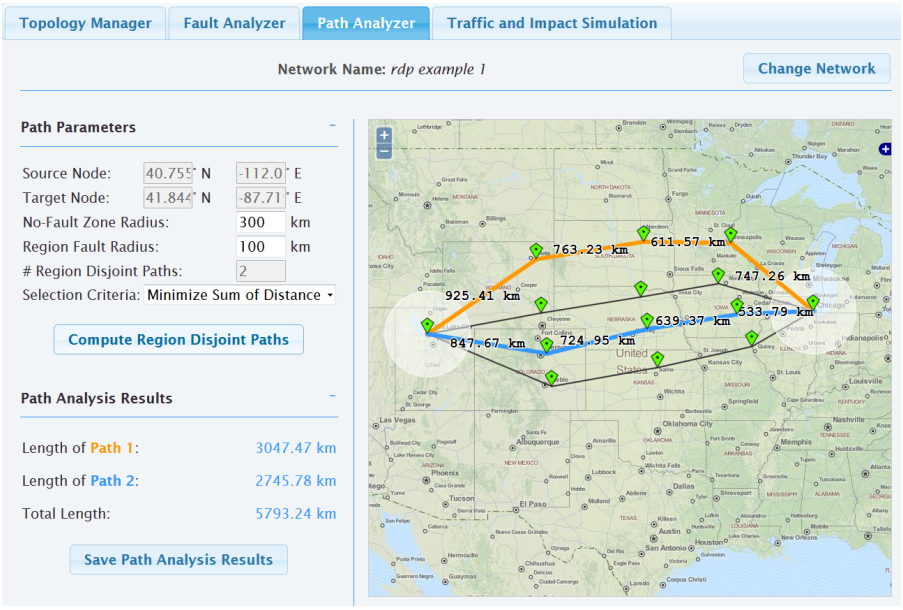

Fig. 6: Region-disjoint paths between a source and destination nodes for given fault radius $r = 100$ *km*. and *no-fault zone* radius $nf_r = 300$ *km*.

3.2.4 Traffic and Fault Impact Simulation

The *Traffic and Impact Simulator* is a discrete event simulation utility that allows users to generate traffic and faults to analyze the impact of faults on a load bearing network. To perform this analysis on a user selected network topology, a simulation schedule consisting of bandwidth requests and faults is generated by the tool using user provided simulation parameters. 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 fault coordinates, or introduce random circular faults for a specified fault radius. Time intervals of the faults can be user specified or generated randomly. Using these settings, the tool generates a discrete time stepped simulation schedule of requests and faults. The user can then select the routing algorithm to be used and run the simulation. As the simulation runs, the faults and requests are processed sequentially according to the schedule and the network state is visualized at every time interval. The impact of faults on the network state is dependent on the routing algorithm chosen by the user, the default algorithm is the shortest path routing algorithm without any bandwidth multiplexing. Future extensions of the tool will support other routing algorithms such as the ones proposed in [10, 27].

As shown in the screen grabs of Fig. 8 and Fig. 9, the left menu contains the fault and simulation parameters that can be used to generate the schedule and run the simulation. The tables below the map allow the user to edit the requests and faults

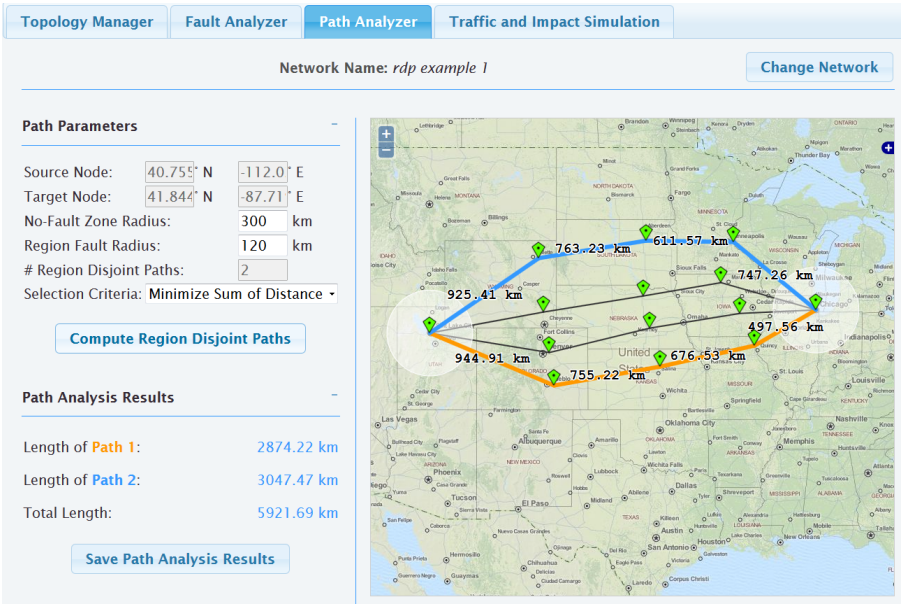

Fig. 7: Region-disjoint paths between a source and destination nodes for given fault radius $r = 120\ km$. and *no-fault zone* radius $nf_r = 300\ km$.

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 tool will iterate over all the time steps and update the map with the network state at each step. In Fig. 8 and Fig. 9 the impact of a fault and the corresponding response of the network is shown. In Fig. 8 the network is fault free, but in Fig. 9 a fault is introduced and an edge is rendered inoperable. It can be seen that the red and yellow flows of Fig. 8 are impacted by the fault, however, as bandwidth is available, in Fig. 9 the flows are rerouted in response to this fault.

3.3 Performance Analysis

For our preliminary performance analysis of the tool, we used a 40 node test network where each node had a degree of at least 3. Using two users, the Path Analysis, Fault Analysis and Traffic and Impact Simulator modules were executed concurrently on the test network using different input parameters. In these tests it was observed that for both users, computations were completed within a minute of the user request. The Path Analysis module, however, required variable computation time depending on the source and destination nodes chosen. This is due to the fact that the Path Analysis computation is executed as an ILP, however, in all our tests, the path computations took at most 3 minutes.

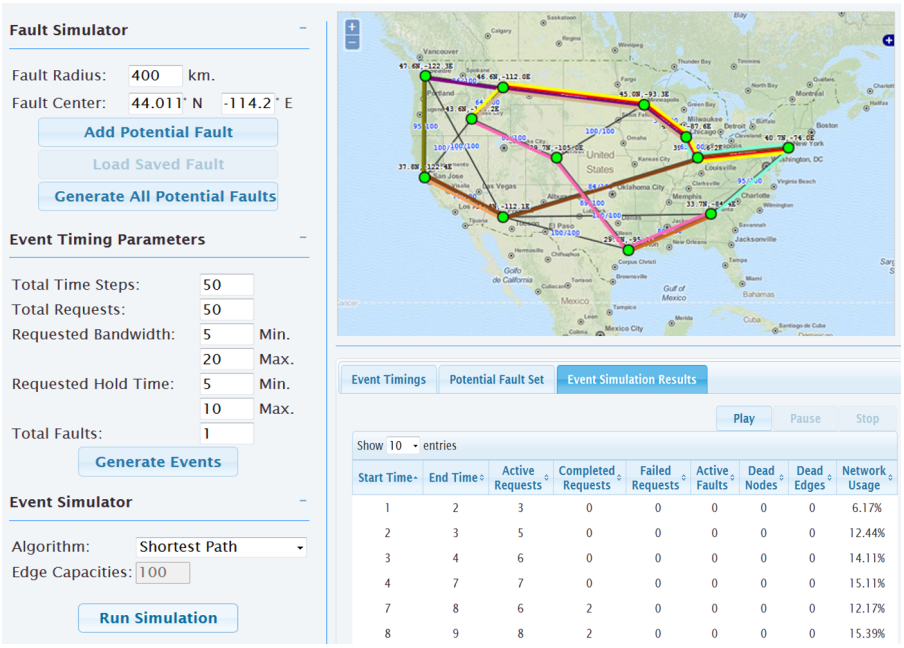

Fig. 8: Traffic and Fault Impact Simulator – Pre-Fault network state

4 Conclusion

In this paper we presented a summary of the work done towards developing RAPTOR, a network planning and management tool intended to support design and analysis of single layer and multi-layer networks in the presence of spatially correlated faults. We highlighted that RAPTOR 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 tool. The goal of this paper was to bring to the attention of the networking research community about the existence of RAPTOR 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. We also note that although the current version of the tool includes several features to aid network designers, planners and operators to better design and manage their networks, however, the available features are in no respect comprehensive to include all user required functionality. In future iterations of the tool, it is our objective to closely work with researchers and industry professionals to incorporate the outcomes of several studies and include additional metrics, techniques, and functionality into the tool to make RAPTOR more usable and useful to its users.

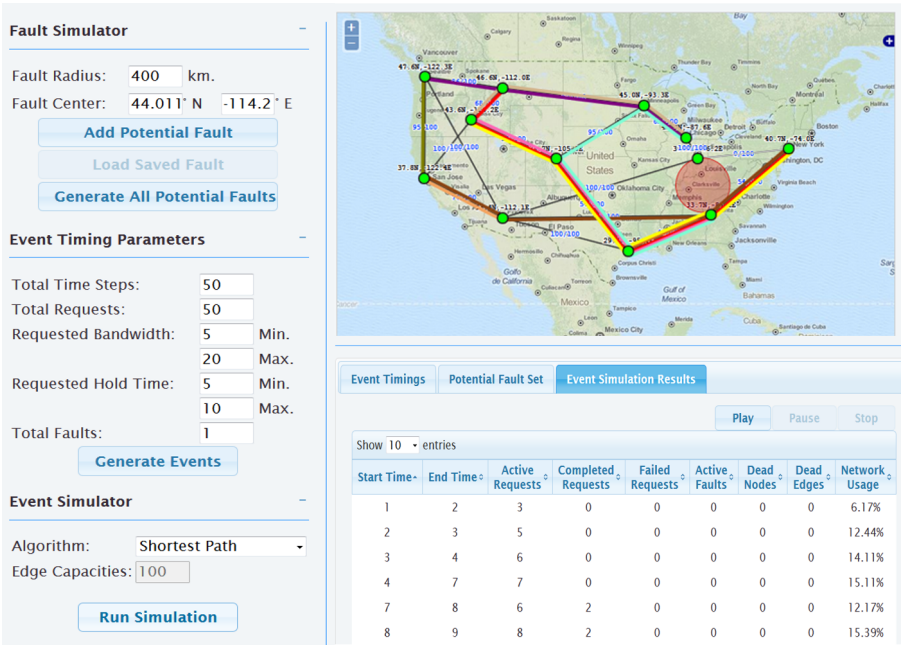

Fig. 9: Traffic and Fault Impact Simulator – Post-Fault network state, rerouted red and yellow flows

References

1. Agarwal, P., Efrat, A., Ganjugunte, S., Hay, D., Sankararaman, S., Zussman, G.: The resilience of wdm networks to probabilistic geographical failures. In: Proceedings of IEEE INFOCOM (2011)
2. Banerjee, J., Das, A., Zhou, C., Mazumder, A., Sen, A.: On the entity hardening problem in multi-layered interdependent networks. In: WIDN Workshop (INFOCOM WKSHPS), 2015 IEEE Conference on Computer Communications, pp. 648–653. IEEE (2015)
3. Banerjee, J., Zhou, C., Das, A., Sen, A.: On robustness in multilayer interdependent networks. In: International Conference on Critical Information Infrastructures Security, pp. 247–250. Springer (2015)
4. Banerjee, S., Das, A., Mazumder, A., Derakhshandeh, Z., Sen, A.: 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, pp. 78–82. IEEE (2014)
5. Banerjee, S., Shirazipourazad, S., Ghosh, P., Sen, A.: Beyond connectivity-new metrics to evaluate robustness of networks. In: High Performance Switching and Routing (HPSR), 2011 IEEE 12th International Conference on, pp. 171–177. IEEE (2011)
6. Banerjee, S., Shirazipourazad, S., Sen, A.: Design and analysis of networks with large components in presence of region-based faults. In: 2011 IEEE International conference on communications (ICC), pp. 1–6. IEEE (2011)
7. Buldyrev, S.V., Parshani, R., Paul, G., Stanley, H.E., Havlin, S.: Catastrophic cascade of failures in interdependent networks. Nature **464**(7291), 1025–1028 (2010)
8. Castet, J.F., Saleh, J.H.: Interdependent multi-layer networks: Modeling and survivability analysis with applications to space-based networks. PloS one **8**(4), e60,402 (2013)
9. Çetinkaya, E.K., Broyles, D., Dandekar, A., Srinivasan, S., Sterbenz, J.P.: Modelling communication network challenges for future internet resilience, survivability, and disruption tolerance: A simulation-based approach. Telecommunication Systems **52**(2), 751–766 (2013)
10. Cheng, Y., Gardner, M.T., Li, J., May, R., Medhi, D., Sterbenz, J.P.: Optimised heuristics for a geodiverse routing protocol. In: 10th International Conference on the Design of Reliable Communication Networks (DRCN), 2014, pp. 1–9 (2014)

11. Das, A., Banerjee, J., Sen, A.: Root cause analysis of failures in interdependent power-communication networks. In: Military Communications Conference (MILCOM), 2014 IEEE, pp. 910–915. IEEE (2014)
12. Das, A., Sen, A., Qiao, C., Ghani, N., Mitton, N.: A network planning and management tool for mitigating the impact of spatially correlated failures in infrastructure networks. In: Design of Reliable Communication Networks (DRCN), 2016 12th International Conference on the, pp. 71–78. IEEE (2016)
13. Das, A., Zhou, C., Banerjee, J., Sen, A., Greenwald, L.: On the smallest pseudo target set identification problem for targeted attack on interdependent power-communication networks. In: Military Communications Conference, MILCOM 2015-2015 IEEE, pp. 1015–1020. IEEE (2015)
14. Diestel, R.: Graph Theory. Springer (2005)
15. Gao, J., Buldyrev, S.V., Stanley, H.E., Havlin, S.: Networks formed from interdependent networks. Nature Physics **8**(1), 40–48 (2011)
16. Gardner, M.T., Beard, C., Medhi, D.: Using network measure to reduce state space enumeration in resilient networks. In: Design of Reliable Communication Networks (DRCN), 2013 9th International Conference on the, pp. 250–257. IEEE (2013)
17. Mazumder, A., Das, A., Zhou, C., Sen, A.: Region based fault-tolerant distributed file storage system design under budget constraint. In: Reliable Networks Design and Modeling (RNDM), 2014 6th International Workshop on, pp. 61–68. IEEE (2014)
18. Mazumder, A., Zhou, C., Das, A., Sen, A.: 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)
19. NetXT Lab, Arizona State University: Raptor: The Network Planning and Management Tool. URL `http://netsci.asu.edu/networktool/`
20. Neumayer, S., Modiano, E.: Network reliability with geographically correlated failures. In: INFOCOM, 2010 Proceedings IEEE, pp. 1–9. IEEE (2010)
21. OpenStreetMap Contributors: OpenStreetMap. URL `www.openstreetmap.org`
22. Rahnamay-Naeini, M., Pezoa, J.E., Azar, G., Ghani, N., Hayat, M.M.: Modeling stochastic correlated failures and their effects on network reliability. In: Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on, pp. 1–6. IEEE (2011)
23. Sen, A., Mazumder, A., Banerjee, J., Das, A., Compton, R.: 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, pp. 831–836. IEEE (2014)
24. Sen, A., Mazumder, A., Banerjee, S., Das, A., Zhou, C., Shirazipourazad, S.: Region-based fault-tolerant distributed file storage system design in networks. Networks (Wiley Online Library, 2015). DOI 10.1002/net.21655
25. Sen, A., Murthy, S., Banerjee, S.: Region-based connectivity-a new paradigm for design of fault-tolerant networks. In: High Performance Switching and Routing, 2009. HPSR 2009. International Conference on, pp. 1–7. IEEE (2009)
26. Sen, A., Shen, B.H., Zhou, L., Hao, B.: Fault-tolerance in Sensor Networks: A New Evaluation Metric. In: Proceedings of IEEE Infocom, pp. 1–12. Barcelona, Spain (2006)
27. Sterbenz, J.P., Çetinkaya, E.K., Hameed, M.A., Jabbar, A., Qian, S., Rohrer, J.P.: Evaluation of network resilience, survivability, and disruption tolerance: analysis, topology generation, simulation, and experimentation. Telecommunication systems **52**(2), 705–736 (2013)
28. Trajanovski, S., Kuipers, F., Ilic, A., Crowcroft, J., Van Mieghem, P.: Finding critical regions and region-disjoint paths in a network. IEEE/ACM Transactions on Networking **23**(3), 908–921 (2015)