Analysing GeoPath diversity and improving routing performance in optical networks

Yufei Cheng[a,*], M. Todd Gardner[c], Junyan Li[a], Rebecca May[c], Deep Medhi[c], James P.G. Sterbenz[a,b,c]

[a] *Information and Telecommunication Technology Center, The University of Kansas, Lawrence, KS 66045, USA*
[b] *School of Computing and Communications, Lancaster University, Lancaster LA1 4WA, UK*
[c] *Computing Department, The Hong Kong Polytechnic University, Hung Hom, Kowloon, Hong Kong*
[d] *University of Missouri-Kansas City, MO, USA*

Abstract

With the increasing frequency of natural disasters and intentional attacks that challenge telecommunication networks, vulnerability to cascading and regional-correlated challenges is escalating. Given the high complexity and large traffic load of optical networks, these correlated challenges cause substantial damage to reliable network communication. In this paper, we propose a network vulnerability identification mechanism and study different vulnerability scales using real-world optical network data. We further propose geographical diversity and incorporate it into a new graph resilience metric cTGGD (compensated Total Geographical Graph Diversity), which is capable of characterising and differentiating resiliency levels among different optical fibre networks. It is shown to be an effective resilience level indicator under regional network challenges or attacks. We further propose two heuristics for solving the path geodiverse problem (PGD) in which the calculation of a number of geographically separated paths is required. Geodiverse paths can be used to circumvent physical challenges such as large-scale disasters in telecommunication networks. We present the GeoDivRP routing protocol with two new routing heuristics implemented, which provides the end nodes with multiple geographically diverse paths. Our protocol demonstrates better performance compared to OSPF when the network is subject to area-based challenges. We have analysed the mechanism by which the attackers could use to maximise the attack impact with a limited budget and demonstrate the effectiveness of restoration plans.

Keywords:
network resilience and survivability; physical network topology diversity; graph-centrality targeted attacks; area-based network disasters; multipath geographic routing heuristics; cross-layer routing protocol;

1. Introduction and motivation

With the ever-increasing traffic demands for communication, optical fibre has been widely deployed due to its high-traffic carrying capacity. It relies on network components such as optical fibres, amplifiers, routers, and switches to maintain normal operations. With the increasing need for high network resilience, it is imperative to analyse and quantify the network robustness level in the process of network design. Networks are vulnerable to regional and correlated physical challenges including large-scale disasters such as earthquakes or hurricanes. Large-scale disasters can cause correlated failures within the affected area. For example, events such as earthquakes or hurricanes can have impact zones up to 800 km in diameter, and cable cuts potentially affect large geographic areas [1]. Unlike previous random failure analysis [2, 3], these challenges relate to the physical location of nodes and links in the network. We study the geographic path diversity characteristics of the network graph to understand the vulnerability level in different physical locations. We further propose mechanisms to optimise routing performance in optical-fibre networks so that they are robust to geographically correlated failures.

We propose a vulnerability identification mechanism using a probabilistic moving-circle challenge model [4]. Other models, such as scaling-circle and polygon challenges are similarly applicable and we plan to include such analysis in our future work. This mechanism captures the essence of physical challenges while maintaining simplicity and effectiveness. In this paper, we apply a similar probabilistic regional attack model [5, 6] in which network components adjacent to the attack centre fail with a high probability, while those away from the centre linearly decrease in the failure probability. We explain the mechanism in detail in Section 4.1. Based on the identified vulnerable areas, we extend the path-diversification metric [7] to consider the *geographic separation* of nodes and links for resiliency analysis. This is an extension to our previous mechanisms [7, 8] in order to represent graph resilience to geographically correlated failures, as opposed to only individual node or link outages. We present our GeoPath diversity metric: minimum distance d between any two nodes on alternate paths. Based on the geodiversity of different node pairs, we present *path geodiversitification* – a new mechanism (pro-

*Corresponding author.
Email addresses: yfcheng@ittc.ku.edu (Yufei Cheng), todd.gardner@faa.gov (M. Todd Gardner), junyan.li.0719@gmail.com (Junyan Li), rm5x8@mail.umkc.edu (Rebecca May), DMedhi@umkc.edu (Deep Medhi), jpgs@{ittc.ku.edu,comp.lancs.ac.uk,comp.polyu.edu.hk} (James P.G. Sterbenz)

posed in [7, 8]) to quantify the graph GeoPath diversity by selecting multiple geographically diverse paths between a given node pair using a quantified geodiversity measure to achieve high network survivability. We then apply this mechanism in the context of real-world optical-fibre networks to compare the relative robustness of these topologies. This mechanism allows future internetworking architectures to exploit naturally rich physical topologies to a far greater extent than is possible with only shortest-path routing or equal-cost load balancing.

We further propose a new distance d-separated resilient routing algorithm and incorporate it in the GeoPath Diverse Routing Protocol (GeoDivRP) that considers geographical diversity and provides multiple geodiverse paths that can be used to circumvent regional challenges given a threat model. It fits in the protocol stack as shown in Figure 1. Knobs $\mathbb{K}$ are used by higher layers to influence the lower layer operation while dials $\mathbb{D}$ are the mechanisms for lower layers to provide instrumentation to the layers above. The application passes a service specification (ss) and threat model (tm) down to the transport layer protocol ResTP (resilient transport protocol). ResTP then requests GeoDivRP to calculate geodiverse paths that meet the requirement tuple $(k, d, [h, t])$, where k is the total number of geodiverse paths requested, d is the distance separation criteria, $[h, t]$ is the optional path stretch (number of additional hops for diverse paths) and skew (delay difference across paths) t constraint. ResTP then establishes multiple transport flows and chooses among different reliability and error control modes (e.g., ARQ, HARQ, FEC) to meet the application service specification while taking advantage of the multiple geodiverse path set $P = P_0 ... P_{k-1}$ provided by GeoDivRP. We apply our multipath algorithm in the context of several real-world service provider networks to analyse the diversity gain and improvement in packet delivery ratio.

The remainder of this paper are organised as follows. Section 2 presents the background and related work. Section 3 provides an overview of the GeoPath diversity mechanism. Section 4 describes our area-based challenges and introduces the area scanning mechanism and further analyses targeted attacks with restoration suggestions. Section 5 introduces our two routing heuristics and the evaluation methodology. Section 6 presents the geodiverse routing algorithm, introduces our protocol implementation in ns-3, and provides routing simulation results. Section 7 concludes the paper and suggests future work.

2. Background and related work

Network challenge analysis is an increasingly popular research area with a number of works have considered only random link and non-correlated failures [9, 10, 11]. Mechanisms have been proposed to identify link- and PoP-disjoint paths for the same node pair in ISP networks [12]. However, most of these works only focus on random challenges using either random synthetic networks or IP-layer networks. Other works have shown that the geolocation of nodes has played a key role in optical fibre networks. This proves that different locations in the physical network have varying contribution to the overall network connectivity [13]. It has also been observed that

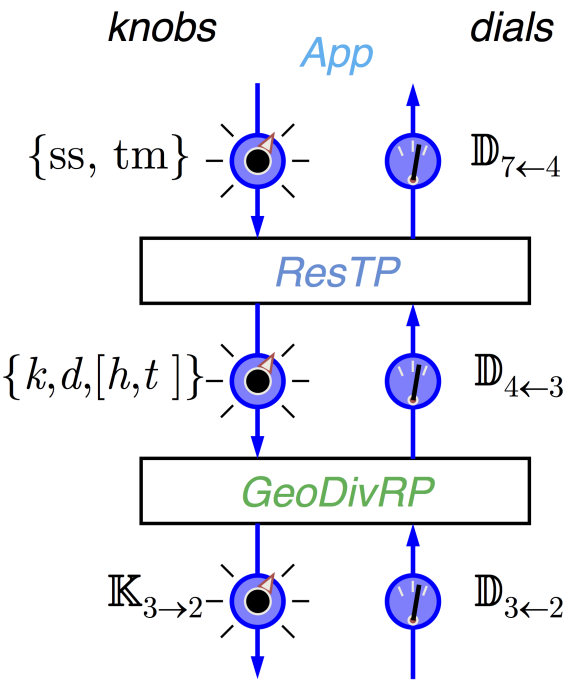

Figure 1: Layered block diagram of the GeoDivRP and ResTP

area-based challenges can cause a large number of failures in a geographical region and give rise to catastrophic damage to network communications [14]. Numerous events have demonstrated that area-based challenges can be modeled as a circular area with a certain challenge radius. For example, an earthquake or hurricane normally has a challenge radius from 0 to 500 miles [1]. There has been previous work on understanding the geographic vulnerabilities for certain topologies that proposed optimisation mechanisms to reduce the searching complexity [15, 16]; based on the vulnerable areas identified, optimisation mechanisms have been proposed to alleviate these impacts [17]. Another vulnerable network zone identification mechanism [5] divides the entire network area into a number of cells to identify the geographical distribution and size of the vulnerable network zones. Single location physical challenge scenarios have been analysed [18, 19, 20], while physical challenges of correlated failures and simultaneous challenges have been discussed [21]. A random line-cut mechanism has been used to assess the vulnerability to regional-based challenges [22]. Both correlated failures and targeted attacks with simulation results have been modeled [4]. However, none of those works incorporates the area-based challenge results into the routing algorithm and guides the process of making routing decisions.

The telecommunication network community has been studying link/node-disjoint paths and previous works have presented survivable network routing algorithms using disjoint paths over the past decade [23, 24, 25, 26]. However, those works did not consider the distance d-separated paths or the GeoPath diversity problem (PGD). As area-based challenges become more impor-

tant to analyse, an efficient algorithm is required to solve the PGD problem. We have proposed a path diversity mechanism for qualifying the network resilience [7, 8]. We further extended this work to the analysis of GeoPath diversity in optical network topologies and have proposed cTGGD (compensated Total Geographical Graph Diversity) to characterise the GeoPath diversity of different network topologies [14]. We have also proposed the GeoDivRP routing protocol that takes advantage of the GeoPath diversity in the physical network topology that achieves survivability by providing multiple geodiverse paths to different application scenarios [27].

The two-terminal and all-terminal reliability between two nodes or any component of the network is the probability that nodes remain connected after random independent links or node failures, respectively [28]. Flow robustness [7, 8] has similar definition and we use flow robustness as the performance metric in this work. cTGGD considers added geodiversity from newly calculated GeoPaths one at a time and proves to be an excellent indicator for network resilience in the face of area-based challenges, as explained in detail in the following section.

3. Path diversity overview

Most networked devices have access to multiple partial or complete physical-layer paths between endpoints, and many of these paths have a certain degree of diversity. However, we are currently unable to benefit from it since design decisions in the current Internet protocol stack assume unipath and shortest path routing. This dramatically decreases the ability to provide resilience under either targeted attacks or large-scale natural disasters. We can achieve improved performance and increased resilience with multiple geodiverse paths.

This paper presents a formal definition of the *GeoPath diversity* metric and its aggregate properties when applied to each node pair as well as to complete network graphs. It is an extension from link/node-disjoint diversity [7, 8, 29] by considering *geographical* diversity between different paths. We evaluate GeoPath diversity based on its ability to reflect the connectivity of the underlying graph, and the cost incurred in doing so in terms of *path stretch*. For this analysis, we have selected the well-connected Level 3 and Sprint physical networks, and the less-connected Internet2 and TeliaSonera physical networks [30, 31].

The definitions in the following sections provide background from our previous work on the mechanism to systematically understand geographical diversity and area-based challenges. We summarise the definitions in Table 1.

3.1. Alternative path mechanism

The primary concern of GeoDivRP is to select alternative distance d-separated paths to circumvent a challenged area when we have an accurate or estimated threat model. Based on the challenge characteristics, our protocol can quickly respond in terms of routing and path selection. Multiple paths along with ResTP end-to-end erasure coding can be used to significantly improve the packet delivery probability. We use these mechanisms for path selection based on GeoPath diversity:

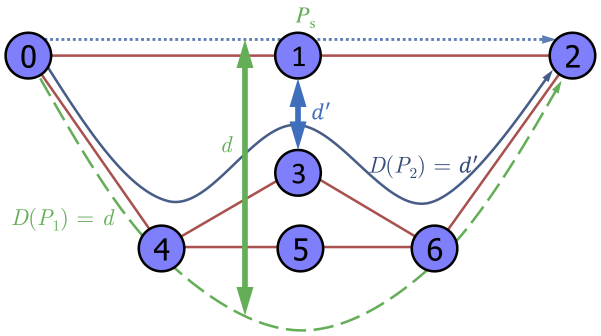

Figure 2: Geographical diversity: distance d

- **Path cache:** indexed by source–destination pairs and includes the unique identifiers for each node and link traversed.
- **Path diversity mode:** single path, k-diverse paths, or d-separated k-geodiverse paths.
- **Path selection:** choose paths that meet higher layer requirement and current network conditions.
- **Packet forwarding:** based on the source routes selected from the path cache.

3.2. Geographical diversity

We define geographical diversity as how much two paths are separated from each other in a geographical context. Note that we assume straight lines for network links for the analysis in this paper. However, when necessary, we add additional nodes to capture the geography of highly curved links. We list two useful definitions for the calculation of geographical diversity as follows.

Path is defined as a vector that contains all links (edges) L and intermediate nodes (vertices) N from a source node to a destination node

$$P = L \cup N \tag{1}$$

Geographical diversity $D(P_a)$ such that $D \geq d$ is defined as the minimum distance between any node members of vector P_a and that of the shortest path. Consider Figure 2 where node 0 is the source and node 2 is the destination node. The red dotted line shows the shortest path P_s consists of nodes 0–1–2. The green dashed line shows path P_1 and its diversity $D(P_1)$ equal d. The blue line shows path P_2 and its diversity $D(P_2)$ is d' since the minimum distance is d' between node 1 and node 3.

Based on the geographical diversity, we start the Effective Geographical Path Diversity (EGPD) metric calculation by taking weighted additional diversity from added paths similar to [8]:

$$\text{EGPD} = 1 - e^{-\lambda k_{sd}} \tag{2}$$

where λ is an experimentally determined constant that scales the impact of k_{sd} based on the utility of this added diversity,

Table 1: Path Diversity Metrics

Metric	Acronym	Definition
Flow Robustness	FR	# of reliable flows divided by # total flows
Effective Geographical Path Diversity	EGPD	weighted additional diversity from added paths
Total Geographical Graph Diversity	TGGD	average of EPGD for all node pairs
Compensated Total Geographical Graph Diversity	cTGGD	TGGD weighted by the total # links for a topology
Aggregated Remaining Flow	ARF	average FR after each challenge node set removed
Normalized Aggregated Remaining Flow	nARF	ARF normalized by total # of links
Cost Radius Relation	CRR	attack radius given budget and # of locations

while s is the source node and d is the destination node. k_{sd} is the sum of all non-zero diversity paths defined as:

$$k_{sd} = \sum_{i=1}^{m} D(P_i), \tag{3}$$

The range of EGPD is between $[0, 1]$ where 0 means that there is no diversity in the graph as there is no alternative path connecting any pair of nodes. When EGPD approaches 1, geographical diversity increases.

Path stretch is defined as the hop counts of L_{P_A} a given path P_A divided by the hop counts L_{P_s} of the shortest path P_s

$$S = L_{P_A}/L_{P_s}, \tag{4}$$

where we use the same definition from [29].

The Total Graph Geographical Diversity (TGGD) is simply the average of the EPGD value of all node pairs within that graph similar to [7, 8]. Therefore, this metric represents the overall resilience of the network topology in face of area-based challenges. Based on the TGGD calculated, we obtain the cTGGD value as follows:

$$\text{cTGGD} = e^{\text{TGGD}-1} \times (\frac{||G_M||}{||G||})^{-\rho}. \tag{5}$$

Here, $||G||$ is the total number of links in topology G, and $||G_M||$ is the total number of links for the largest network topology in consideration (in this case 488 links for AT&T). We weight the graph diversity based on the division result of $||G||$ and $||G_M||$. The purpose of the weight is two fold, first is to eliminate the penalty to a dense network for a given size of a physical region. This is because one dense network will have less geographical diversity for one node pair within a given area as the links are not able to be as separated geographically compared to one sparse network. Second, it is normalized by the number of links of the largest topology in the comparison topology group. ρ is experimentally chosen as 0.05. By weighting the TGGD metric by the link number of the largest topology in the set of topologies in consideration, the cTGGD metric indicates the relative resilience level of topologies against the largest topology.

cTGGD considers added geodiversity from newly calculated GeoPaths one at a time. Therefore, cTGGD represents the resilience level of a certain topology against area-based challenges through the incrementally added GeoPath. We present the metric comparison results in the next section and show cTGGD to be one good indicator for network resilience in face of area-based challenges. However, cTGGD has a couple of limitations. First, the metric requires weighing by the largest network topology in the comparison set of topologies. This means that when the set of network topologies are different, one given topology may not have a global unique cTGGD value. Second, the value ρ needs to be selected experimentally. We plan to explore these issues in future work.

4. Area-based challenges

Area-based challenges have drawn increasing attention in the network community due to their large-scale impact and their potential damage to optical fibre networks. In this section, we first propose an area scanning mechanism to identify vulnerable areas in a network and suggest mechanisms to prepare the current network for those types of challenges. Furthermore, we analyse how the network performs in the face of targeted attacks using centrality metrics.

4.1. Area scanning mechanism

We propose an area identification mechanism to better determine the vulnerable locations in optical networks. This is one endeavour to help optimise the design and maintenance of the normal operation of communication networks in face of challenges. Before explaining the scanning mechanisms in detail, we introduce the network performance metric used in this paper.

4.1.1. Flow Robustness (FR)

A flow is established between each node pair using a set of paths determined by the path geodiversification algorithm with a specified diversity threshold. Link and node removal, based on a fixed probability of failure, have been analysed [8]. We consider regional challenges, where one challenge removes nodes covered in the area and links connected to the challenged nodes. A flow is considered *reliable* if at least one path remains connected during the failure. We compute *flow robustness* to be the number of the reliable flows divided by the number of total flows that exist in the network. The algorithmic complexity depends on the time to find the number of components in a given graph, which makes the complexity as $O(|V| + |E|)$. The

limitation of flow robustness is that it is based on connectivity and does not take traffic into consideration. We plan to include traffic analysis in future work.

4.1.2. Scanning mechanism

The scanning mechanism starts by embedding the topology in the Euclidean plane based on each node's geolocation and operates greedily by scanning through the entire topology for possible vulnerable areas. The unit for both of the models is degrees in latitude and longitude. The model takes two input values from the user: the radius of the scanning circle and the degree that it forwards each step along the longitude scale. The nodes that fall into the circle in the Euclidean plane at any given step are defined as the challenged node set. Whenever an identified challenged node set has distinct node members compared to the previous set, we record it as a *possible challenge set*. The model starts from the network topology corner with the smallest longitude and latitude values. When the challenged node set is exhausted in the longitude level of a certain topology, the model moves forward along the latitude scale for one step and repeats the above process. By applying this mechanism in the optical fibre network, we are able to identify all the different challenge scenarios for geographically correlated challenges. We calculate the flow robustness of the network topology after each challenge node set has been removed from the topology. We can identify the relative contribution of the network components in each area to the overall resilience of the whole topology; in other words, we can identify the most vulnerable area in the network and offer better suggestions for network design and capacity planning. Based on the flow-robustness value for different challenge node sets, we plot different colour shades on the map to demonstrate the relative vulnerability of different geographical areas. This is one simple yet effective mechanism to scan through all the distinct node sets and identify possible vulnerable areas.

We further reduce the complexity of the mechanism by fixing the centre of the challenge circles at the nodes in the topology. This reduces the number of challenged node sets while at the same time captures the essential aspects of the topology as all the nodes are covered by at least one challenge circle. We employ a probabilistic challenge model with one inner circle and one outer one. The challenge probability of both inner and outer circles is tunable by the user, and defines what percentage of nodes in the circle fail. In this paper, we set the challenge probability for the inner circle as 1.0 and that of the outer circle as 0.5. The inner circle probability represents a deterministic challenge scenario while the outer circle represents a simple probabilistic challenge. Any nodes outside of the circles are not affected. We are able to identify vulnerable locations in the network realistically keeping the mechanism simple; this is a simplified challenge model from [32].

By averaging the list of flow robustness results obtained from the scanning mechanism, we can evaluate the relative resilience of different network topologies. We define ARF (Aggregated Remaining Flow) for a specific topology and its value is in the range $[0, 1)$.

Aggregated Remaining Flow is defined as the average flow robustness after each challenged node set has been removed

$$\frac{\sum_{i=1}^{n}(R_i/A_i)}{n} \tag{6}$$

where A_i is defined as all possible network flows within a given topology in challenge node set i and R_i is the remaining flow after each challenged node set i has been removed, n is the number of challenged node sets identified in one topology. The larger the ARF is, the more robust a certain topology is against area-based challenges. We further propose Normalized Aggregated Remaining Flow nARF as follows.

Normalized Aggregated Remaining Flow is defined as the remaining flow robustness after a list of nodes has been challenged, normalized by the total number of links

$$\text{nARF} = e^{\text{ARF}-1} \times \left(\frac{\|G_M\|}{\|G\|}\right)^{-\rho} \tag{7}$$

$\|G\|$ is the total number of links in topology G, and $\|G_M\|$ is the total number of links for the largest network topology in consideration (in this case 488 links for AT&T). The moving challenge circle model produces a limited number of challenged node sets by the degree difference in both longitude and latitude in a given network topology. Since both ARF and nARF are calculated after removing each challenged node set at a time, they are time-bounded.

We further introduce the notion of *challenge failure criteria*. When users provide a criteria value between zero and one, the scanning mechanism returns the list of challenged node sets that can reduce the flow robustness value below this criteria. The challenge radius is increased by one-degree increments if none of the challenged node sets meet the criteria. This is primarily useful when identifying the most vulnerable area in a specified network topology.

4.2. Scanning mechanism evaluation

As shown in Figures 3 and 4, we present the visual representation of the challenge locations plotted in the topology with different colour shades of challenged circles for both Sprint and Level 3 physical topologies.[1] The darker the shade, the more vulnerable the location is. The radius used in this mechanism is five degrees with one degree per step, and we use a fixed-size challenge circle in the scanning mechanism to find the most vulnerable geographical area for that given size. This is a fairly large range of challenges, and we use it here for easy visual representation. The dark colour circle in the Chicago area causes the flow-robustness to drop to 30%, which is the most vulnerable area inside this topology. The slightly lighter colour circle in the New York area reduces the flow-robustness to around 70%, while the even lighter circle in the San Diego area only drops it to around 90%. The probabilistic failing case introduced in Section 4.1 shows similar relative vulnerability levels and is not shown here.

[1] US network topologies are analysed in this paper; the same mechanisms can be applied to non-US, regional, and global networks.

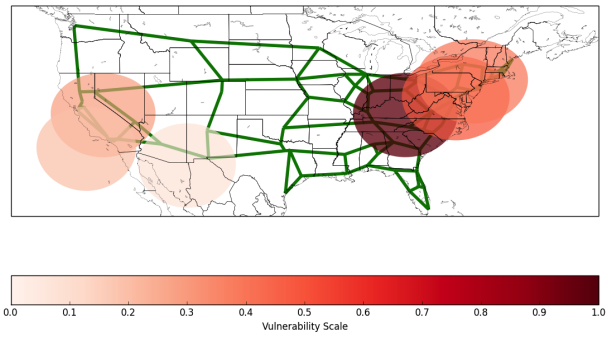

Figure 3: Sprint network with different vulnerable areas

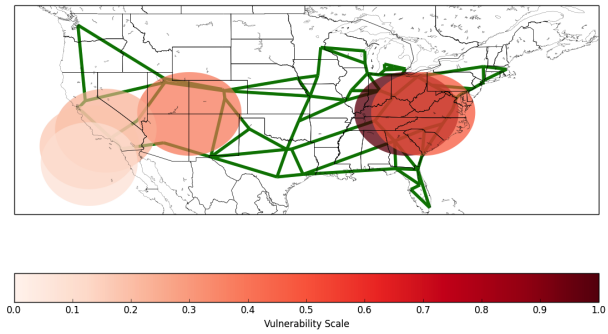

Figure 4: Level 3 network with different vulnerable areas

We further apply the scanning mechanisms to identify the most vulnerable area in the network topologies, with the challenge failure criteria set as 0.5. This means that the challenge radius increases until the flow robustness of any challenge circle drops the flow robustness below 50%. The reason we choose 0.5 as the failure criteria is because we consider it substantial damage to network connectivity, but other thresholds can be chosen as appropriate for a given scenario.

As we can observe from Table 2, the most vulnerable areas to the area-based challenge are around the Virginia and Pennsylvania locations. This is due to the fact that if the challenges happen in those locations, the communication for most of the northeast US will be disconnected from the rest of the network. For example, a challenge in the radius of four degrees can drop the flow robustness of the AT&T network to below 50%. Although challenges are rare at this large scale, we need to understand the consequences. If similar natural disasters happen in these vulnerable locations, the damage to the telecommunication network will be catastrophic, for which corresponding plans should be made.

We apply this mechanism to verify the path geodiversification mechanism. We present a comparison of our graph metrics in Table 3. As shown in this table, the difference among different topologies in terms of both TGGD and ARF is minimal. This is because all the topologies are designed with nodes and links separated, which provides survivability against area-based challenges. Although these optical networks have achieved similar TGGD and ARF, they have a different number of links. This analysis has penalises against dense networks with more links since when the geographical area is fixed, the geographical separation among links is limited. Both the cTGGD and nARF metrics demonstrate that after normalising the metrics based on the total number of links, we can compare the GeoPath diversity of different physical networks. cTGGD and nARF have shown comparable results by successfully distinguishing different geographical diversity levels among the optical fibre networks. The numbers in bold are the two topologies with the largest cTGGD and nARF values. nARF has a higher computational complexity since the number of different challenged node sets is $O(nm)$, with n as the step interval at the longitude scale and m at the latitude scale, which makes the complexity $O(nm(|V|+|E|))$. The cTGGD metric effectively indicates the resilience level of different topologies under regional challenges while at the same time having a substantially lower complexity compared to nARF.

4.3. Targeted attacks

In addition to understanding the challenges from natural disasters, we further explore how targeted attacks can affect physical layer networks. In this case, the attackers determine the exact attack location and the radius of the attack in order to do significant damage. An example of such attacks would be an EMP (electromagnetic pulse) weapon. From the standpoint of the attackers, we analyse the mechanisms they could apply to increase the damage with a given attack budget. We assume the cost to increase the attack area is proportional to the budget, which means that the radius of the attack corresponds to the square root of the budget. We use cost c to represent the cost to take down an area of A in the physical topologies, while the number of attack locations l corresponds to the number of challenges that share the total attack budget.

Cost Radius Relation is defined as the radius of each attack location, given the attack budget (c) and the number of attack locations (l)

$$r = \sqrt{\frac{c}{\pi l}} \tag{8}$$

We employ several best known centrality metrics: betweenness, closeness, eigenvector, load, and degree centrality [33, 34] to analyse different physical networks and provide a list of nodes sorted according to their different centrality values from high to low. Betweenness is defined as the number of the shortest paths that flow through a node; it signifies a node's importance in network communication [35]. Closeness is the inverse of the sum of the shortest paths from a node to every other node and indicates efficiency of a message's diffusion in a network [36]. Eigenvector centrality is a measure of the influence of a node in a network [37] and assigns relative scores to all nodes in the network based on the assumption that connections to high-scoring nodes contribute more than connections to low-scoring nodes. The load centrality of a node is the fraction of all shortest paths that pass through that node [35]. Degree centrality is the number of links affiliated to a node and can be viewed

Table 2: Physical topology vulnerable locations

Network	Number of Nodes	Number of Links	Flow Robustness	Challenge Locations	Challenge Coordinates	Challenge Radius	Number of Failed Nodes
AT&T	383	488	0.48	Pittsburgh, PA	40.44, -79.97	4	23
CORONET	75	99	0.44	Pittsburgh, PA	40.44, -79.98	5	11
Internet2	57	65	0.35	Nashville, TN	36.17, -86.78	6	3
Level 3	99	132	0.43	Pittsburgh, PA	40.44, -79.98	6	15
Sprint	77	114	0.45	Roanoke, VA	37.28, -79.96	5	15
TeliaSonera	18	21	0.43	Ashburn, VA	39.04, -77.48	6	6

Table 3: Network characteristics

Network	TGD	cTGD	TGGD	cTGGD	ARF	nARF
AT&T	0.90	0.06	0.99	**0.96**	0.86	**0.87**
CORONET	0.93	0.16	0.99	0.89	0.90	0.84
Internet2	0.88	0.26	0.94	0.81	0.88	0.80
Level 3	0.89	0.10	0.97	0.90	0.87	0.82
Sprint	0.91	0.08	0.98	**0.92**	0.87	**0.86**
TeliaSonera	0.75	0.15	0.87	0.75	0.87	0.75

as the relative importance of a node [36]. These graph centrality metrics have been used to study performance of networks against targeted attacks [38, 4].

The attack starts with a challenge area defined as the *Cost Radius Relation* (CRR) centered at the highest centrality node identified from the previous step. Given the fixed budget for the attack and assuming a certain budget defines a specific attack area, the attack can occur in one location or multiple locations each with a smaller attack radius. For simplicity of the analysis, we assume that the area-based attacks in different locations are divided equally in areas. For example, if the total challenge area is ten and the number of challenge locations is two, then each challenge location has an area of five. We present how the number of attack locations affects the overall flow robustness for the AT&T physical network in Figure 5. As the number of challenge locations increases, the flow robustness value decreases. For example, the degree centrality attack drops the flow robustness to below 40% when the number of locations is 16. Furthermore, after the challenge locations increase beyond four, the value of flow robustness stabilises. As it would be more complicated and costly to increase the number of attack locations, we conclude that by dividing the attacks into four locations and deploying them based on the higher degree centrality maximises the attack damage in the AT&T network.

Figure 6 shows similar results when the attacks happen in the Sprint physical network. Similar to the AT&T network, degree centrality still has the greatest impact to the flow robustness. However, the significant drop in flow robustness happens around eight challenge locations. This is partly due to the evenly distributed network nodes and links in the Sprint network. However, when the number of locations increases beyond eight, the flow robustness drops significantly, which is

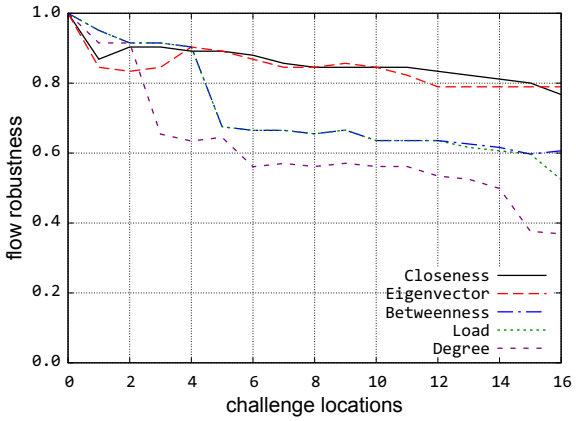

Figure 5: AT&T optical network under regional challenges

due to the fact that after the higher centrality nodes have been removed, the network is partitioned.

Figure 7 shows the Level 3 physical network under targeted attacks. It shows faster and more significant damage than the other networks. When the number of challenge locations increases to eight, the flow robustness drops below 20%. This targeted attack result demonstrates that with a certain amount of knowledge of the network topology and expertise to analyse it, attackers can cause a substantial amount of damage even with a small budget.

4.4. Suggestion for restoration and improvement

Network recovery time can vary from a few hundred seconds to days [39]. In this section, we analyse the effectiveness of network restoration schemes and provide network improvement

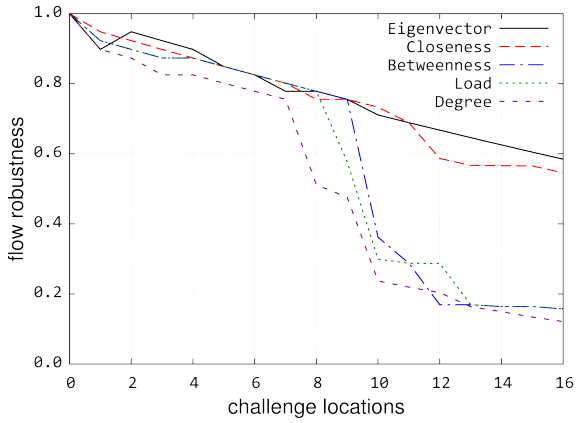

Figure 6: Sprint optical network under regional challenges

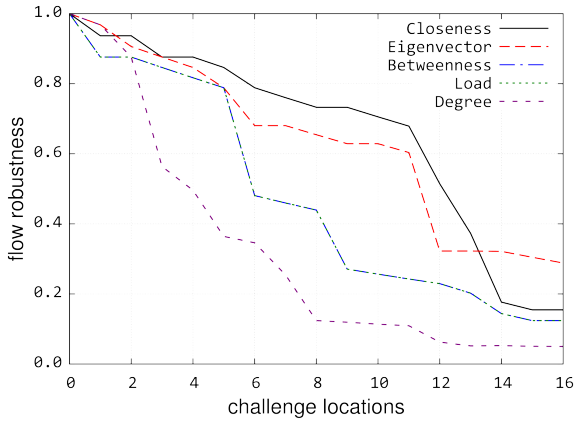

Figure 7: Level 3 optical network under regional challenges

suggestions. We present flow robustness results when the network has restoration plans and demonstrate the improvement for the overall network performance. This is one endeavour to better understand the challenge characteristics and suggest network design guidelines.

The result from the vulnerable area scanning mechanism in Section 4.1 reveals vulnerable network locations and can guide the improvement of overall network resilience. For example, adding physical protection for existing components in the vulnerable locations can mitigate large-scale physical attacks. Compared to analysing the overall resilience and global optimisation of networks, this is the local optimisation of the network based on the vulnerability level of each individual area.

We present the flow robustness improvement when a certain percentage of the challenged nodes have remained connected due to a particular restoration or protection plan.[2] The challenge locations come from the most vulnerable areas we have identified in Table 2. Due to the size of different networks, the number of challenged nodes in different locations varies. As we noticed from Figure 8, by protecting three nodes, all the physical networks increase to above 60% flow robustness. Protection can be done by shielding existing nodes or providing hot standby nodes. The 20% flow robustness improvement is valuable for network recovery as well as disaster recovery.

[2]Note that a specific restoration plan is not studied in this work.

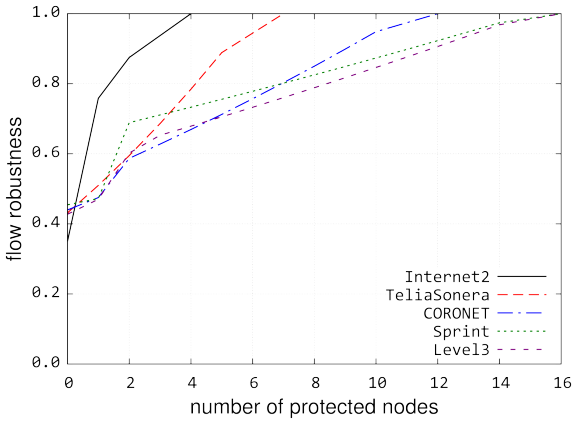

Figure 8: Protection plan improvement on different networks

5. Routing model description

Based on the previous regional challenge analysis, we design GeoPath Resilient Link-State Routing (GeoResLSR) and GeoPath Diverse Routing Protocol (GeoDivRP) with GeoPath diversity taken into consideration when routing traffic and nodes advertise their geocoordinates. Traditional routing protocols are designed to form the shortest path for each source–destination pair in the current Internet architecture for efficiency. However, this comes with the cost of not having the option to choose alternate paths when the current path is unavailable due to challenges or attacks. Fast IP reroute [40] is designed to protect the network from the event of a single failure, yet it lacks protection capacity for multiple simultaneous failure events. With a realistic regional-based challenge model, in order to quickly route around the challenged area, a new routing protocol is required to find multiple backup paths for every node pair in the network.

We formulate the Path Geodiverse Problem (PGD) for calculating k geographically separated paths and provide a two-step algorithm for solving it. This algorithm begins with the Suurballe's algorithm [23, 24] in which the shortest-path algorithm (SPA) is iteratively applied. After each iteration of the SPA, the weight of the edges from the constructed path is penalised by adding a penalty factor. Once the algorithm has identified n paths, it selects the path with distance d-separated by iteratively comparing the distance between each and every node pair from all the candidate paths. Based on this algorithm, we have designed the Geodiverse Resilient Link-State Routing (GeoResLSR). This mechanism guarantees choosing the best d-separated paths when assuming a large number of candidate paths [27]. However, as the SPA is applied n times for generating the candidate paths before selecting the qualified ones, its time complexity is $O(n|V|(|E| + |V|log|V|))$ [41] and the computation is slow. To reduce the complexity of the

two-step algorithm, we propose two heuristics for efficiently calculating geographically diverse paths. Specifically, we consider PGD that involves obtaining a set of paths that are distance d-separated from each and every node in different disjoint paths (d-separation). The proposed heuristics return a set of (S, D) paths from the graph $G = (V, E, w)$, where V is the vertex set, E is the edge set, and w is the link weight set. Dijkstra (G, n) is the standard Dijkstra algorithm we use to provide the shortest path. We list the graph notations used as follows:

- $G(V, E, w)$: input graph G with a set of vertices V, a set of edges E and weight of edges w
- S: source node.
- D: destination node.
- S_n: neighbor node chosen by source node.
- D_n: neighbor node chosen by destination node.
- k: number of geodiverse path requested.
- d: distance separation between each and every node in different disjoint paths.
- δ: delta distance when selecting waypoint node.

Our protocol is based on Open Shortest Path First (OSPF) [42]. It is a link-state interior gateway routing protocol that is widely used in a single autonomous system. It has become (along with IS-IS) the de facto interior gateway routing protocol. In OSPF, the Dijkstra algorithm is used to calculate the shortest path between a pair of nodes based on link-state information. Every node generates a link state advertisement (LSA) that carries the cost of all its links and floods throughout the network. To ensure liveness, each node sends HELLO packets to their neighbours over the *hello interval*, which is set at ten seconds by default. If a node does not receive its neighbours' HELLO packets after a HELLO interval, its adjacency no longer exists and will recalculate the shortest path. This means that if a challenge occurs, the network needs at least a HELLO interval to detect the challenge and to react in response.

5.1. Heuristics

In consideration of decreasing the complexity of the geodiverse path calculation, we propose two heuristics: iterative WayPoint Shortest Path (iWPSP) and Modified Link Weight (MLW) [27, 43]. As shown in Figure 9, for the case when $k = 3$, the iWPSP first selects neighbour nodes S_{k1} and D_{k2} that are d distance separated from source node S and destination node D, respectively (for simplicity in this presentation we assume that such nodes exist; otherwise, the nodes with the greatest distance will be chosen, iterating until nodes d apart are located). Assuming the straight line connecting S and D is L, the iWPSP selects waypoint nodes m' and m'' in the opposite direction that are distance $d + \delta$ apart from the middle node m in the shortest path, where the segment $m'mm''$ interleaves with the shortest path. Dijkstra's algorithm is performed

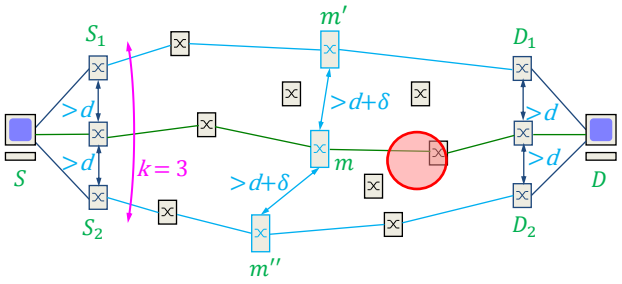

Figure 9: Iterative waypoint shortest path heuristic

for the two branches $S_{km'}$ and $D_{km'}$. By connecting the shortest path returned from the two branches, the heuristic obtains the first geodiverse path P_1. The same mechanism repeats for waypoint node m'' for the second geodiverse path. The variable d is a user-chosen parameter based on a threat model for a challenge of distance d, and δ is experimentally chosen for different network topologies to increase the probability of the heuristic successfully returning a d-separated path. The δ parameter is also introduced to prevent the edges in each of the two paths from interleaving and creating routing loops. By tweaking the value of δ, this heuristic can select a nearby waypoint node if the previous one fails running Dijkstra's algorithm. The pseudo code for the iWPSP is shown in Algorithm 1.

Our second heuristic MLW statistically modifies the link weights and performs Dijkstra's algorithm to calculate the geodiverse paths with the modified link weights in the network. The heuristic begins by increasing, linearly or squarely, the weight in one direction based on the perpendicular distance to the straight line L connecting source node S and destination node D. The weight increment ratio is inversely proportional to the distance from L. Dijkstra's algorithm is applied on the graph with modified link weights. The heuristic repeats the process for the other perpendicular direction to L. This way the heuristic can generate two paths that are geographically separated. If more diverse paths are required, the heuristic selects one of the geodiverse paths established as the starting line for modifying link weights and iteratively generates k geodiverse paths as shown in Figure 10.

We use a 5×5 grid network to demonstrate the d-separation paths calculated by the MLW. As shown in Figure 11, the MLW calculates two paths that are separated by distance d by statistically modifying link weights. Node 21 is the source and node 3 is the destination. The d value is set at twice the length of the edges in the grid. The iWPSP heuristic generates same results when using the heuristic mechanism shown in Figure 9 and Algorithm 1. The weight shown in different colours is used for calculating paths in its representative colours. For example, when the MLW is calculating the path shown in blue solid links (the first of the two weights before the slash), the link weight is statistically modified by decreasing towards the top right corner of the grid network. The other path shown is the red dashed links, corresponding to the second of the two weights after the slash. The detailed heuristic is presented in Algorithm 2.

Both of the heuristics have incorporated improvement mech-

Functions:
Calculate k paths from S to D separated by distance d
Input:
G_i:= input graph
S:= source node
D:= destination node
k:= number of requested geodiverse path
d:= separation distance between the paths
δ:= delta distance when selecting waypoint node
Output:
k number of geographically d distance separated paths
begin
segment L connecting S and D, with its middle point m;
choose neighbor node S_k, D_k that is at least d distance from S_{k-1}, D_{k-1}, respectively;
if *k is odd number* **then**
choose two nodes m_1 and m_2 that are separated by $d + \delta$ on each direction of L, where $m_1 m m_2$ is perpendicular bisector of L;
$P_1 = \text{SourceTree}_{DS} \leftarrow \text{Dijkstra}(D, S)$;
$k- = 3$;
else
choose two nodes m_1 and m_2 that are separated by $d/2 + \delta$ on each direction of L, where $m_1 m m_2$ is perpendicular bisector of L;
$k- = 2$;
end
$p_{m_1 S_1} = \text{SourceTree}_{S_1 m_1} \leftarrow \text{Dijkstra}(m_1, S_1)$;
$p_{m_2 S_2} = \text{SourceTree}_{S_2 m_2} \leftarrow \text{Dijkstra}(m_2, S_2)$;
$p_{m_1 D_1} = \text{SourceTree}_{D_1 m_1} \leftarrow \text{Dijkstra}(m_1, D_1)$;
$p_{m_2 D_2} = \text{SourceTree}_{D_2 m_2} \leftarrow \text{Dijkstra}(m_2, D_2)$;
while $k > 0$ **do**
segment L = newest established path;
choose one node m_k that is separated by distance $d + \delta$ from L on the farther direction from the absolute shortest path;
$p_{m_k S_k} = \text{SourceTree}_{m_k S_k} \leftarrow \text{Dijkstra}(m_k, S_k)$;
$p_{m_k D_k} = \text{SourceTree}_{m_k D_k} \leftarrow \text{Dijkstra}(m_k, D_k)$;
$k- = 1$;
end
if *k is odd number* **then**
$P_2 = p_{m_1 S_1} + p_{m_1 D_1}$;
$P_3 = p_{m_2 S_2} + p_{m_2 D_2}$;
...
$P_k = p_{m_{k-1} S_{k-1}} + p_{m_{k-1} D_{k-1}}$;
else
$P_1 = p_{m_1 S_1} + p_{m_1 D_1}$;
$P_2 = p_{m_2 S_2} + p_{m_2 D_2}$;
...
$P_k = p_{m_k S_k} + p_{m_k D_k}$;
end
return $(P_1, P_2, ..., P_k)$
end

Algorithm 1: Iterative waypoint shortest path heuristic

Functions:
cost(L):= cost function
Input:
G_i:= input graph
W_i:= link weights
S:= source node
D:= destination node
k:= number of diverse paths requested
buffer := distance buffer to increase link weight
Output:
k number of paths that are geographically separated by distance d
begin
straight line L connecting source S and destination D
if *k is odd number* **then**
$P_1 = \text{SourceTree}_{DS} \leftarrow \text{Dijkstra}(D, S)$;
modify link weight linearly or squarely on one direction perpendicular to line L until distance d;
$P_2 = \text{SourceTree}_{DS} \leftarrow \text{Dijkstra}(D, S)$;
repeat the process for the other direction;
buffer = d;
$k- = 3$;
else
modify link weight linearly or squarely on one direction perpendicular to line L until distance $d/2$;
$P_1 = \text{SourceTree}_{DS} \leftarrow \text{Dijkstra}(D, S)$;
repeat the process for the other direction;
buffer = $d/2$;
$k- = 2$;
end
while $k > 0$ **do**
buffer += d;
modify link weight linearly decreasing on one direction perpendicular to line l until buffer;
links beyond distance buffer, link weight = 1;
$P_{k-1} = \text{SourceTree}_{DS} \leftarrow \text{Dijkstra}(D, S)$;
repeat the process in the other direction;
$P_k = \text{SourceTree}_{DS} \leftarrow \text{Dijkstra}(D, S)$;
$k- = 1$;
end
return $(P_1, P_2, ..., P_k)$
end

Algorithm 2: Modified link weight shortest path heuristic

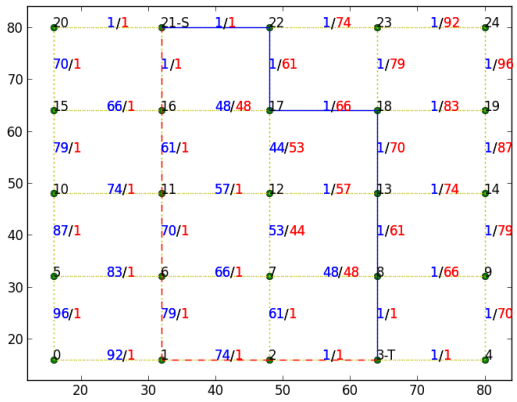

Figure 10: Geodiverse paths by MLW heuristic in grid network

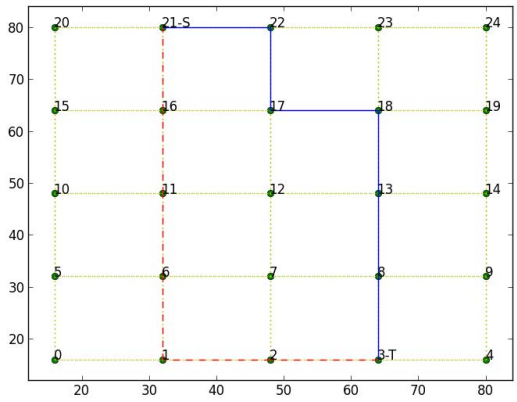

Figure 11: Grid network routing

anisms. When the calculated paths fail to qualify the d-separation criteria, the iWPSP will choose another waypoint that has a slightly larger δ distance; while the MLW will increase the link weight around the avoidance line. The heuristics initialise another iteration of Dijkstra's algorithm. The heuristics fall back to the two-step algorithm if the result still does not qualify for the d-separation criteria, which ensures that both of the heuristics can acquire the geodiverse path while not generating the worst case complexity. Another major component of both heuristics is loop detection. For example, the iWPSP algorithm can create routing loops when calculating paths for corner nodes in the topology. We use a loop detection algorithm so that if a previous node from one path is identified, the algorithm deletes that part.

Some may argue that geographical vulnerability should be fixed at the network planning phase instead of at the routing phase; however, this is not always the case for the following reasons. First, network planning with over-provisioning is a long term process; we still have to design our routing protocol to cope with regional challenges based on the current network topology. Second, although sophisticated network planning mechanisms can help reduce the impact to network traffic during area-based challenges, resilient routing is still needed to get around challenged areas quickly and be adaptive to traffic and congestion in the network.

5.2. Complexity analysis and evaluation

We analyse the complexity of the two heuristics compared to the two-step algorithm. For simplicity, we examine the complexity for obtaining two d-separated paths and assume the Fibonacci heap for Dijkstra's algorithm. The two-step algorithm starts by calculating k edge-disjoint paths using Suurballe's algorithm, which requires k iterations of Dijkstra's algorithm. Dijkstra's algorithm can be performed in time $O(m + n \log n)$ on a graph with n vertices and m edges. Therefore, the same time complexity applies to each path for the Suurballe's algorithm, which makes its complexity $O(km + kn \log n)$. After generating k disjoint paths, the two-step algorithm demands a choice of paths that qualify the distance separation criteria. This process requires n^2 time, which means the total complexity for the two-step algorithm is $m + n \log n + kn^2$, or $O(kn^2)$. The number of edge-disjoint paths k may be large to guarantee the quality of the paths calculated. For most application scenarios, k is chosen to be 1000 [44]. Therefore, for a network with vertices less than 1000, the complexity of the two-step algorithm goes up to $O(n^3)$.

iWPSP has a complexity of $2c^2n^2 \log n$, where c is the average number of neighbours for vertices; the complexity for choosing the waypoint node is $O(n)$, where n represents the number of nodes; and $2n \log n$ is for Dijkstra's algorithm to calculate the two shortest paths. Therefore, the worst case scenario is $O(n^2 \log n)$. Most of the physical topologies have an average degree below four [30]. This means that c in our complexity analysis is a small constant. This reduces the best case time complexity of the iWPSP to $O(n \log n)$. The complexity of the MLW is $O(2n \log n)$, which is the complexity for invoking Dijkstra's algorithm twice. The complexity for both of our heuristics is much better than that of the two-step algorithm, which is $O(n^3)$.

5.3. Path calculation validation

We use a 5×5 grid network to demonstrate the d-separated paths calculated by the algorithm. As shown in Figure 11, the algorithm demonstrates the calculation of two geographical diverse paths from Node 21 to Node 3. The d value is set as twice the length of the edges in the grid.

6. Real network results

In this section, we evaluate the proposed heuristics and compare their performance with the two-step algorithm [14]. We present the geodiverse paths calculated by our heuristics using the Nobel-EU (Pan-European Reference Network) with 28 nodes and 40 links [45]. We assume a challenge along the line from Amsterdam to Rome with a radius of 50 km. Nodes Strasbourg and Frankfurt are in the challenge circle. The result of the

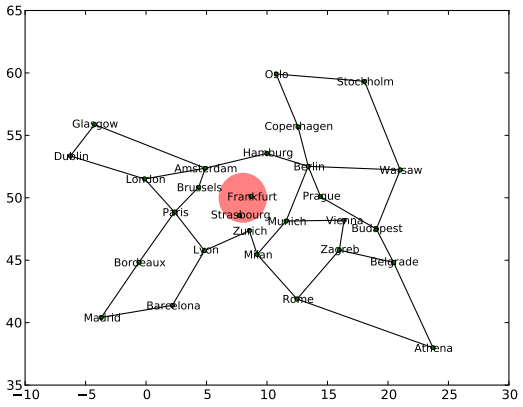

Figure 12: iWPSP heuristic in Nobel-EU network

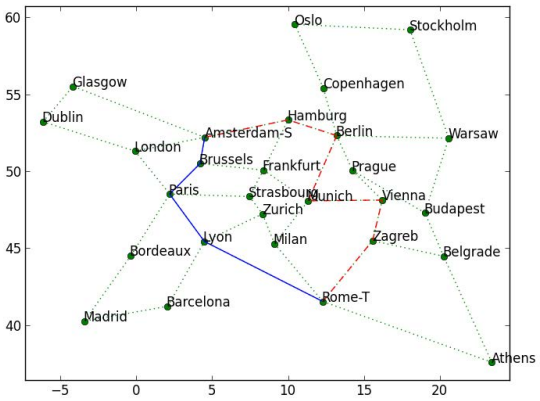

Figure 13: MLW heuristic in Nobel-EU network

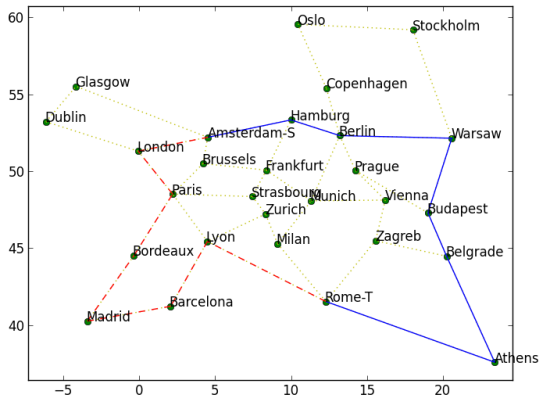

Figure 14: MLW heuristic in Nobel-EU network with large radius

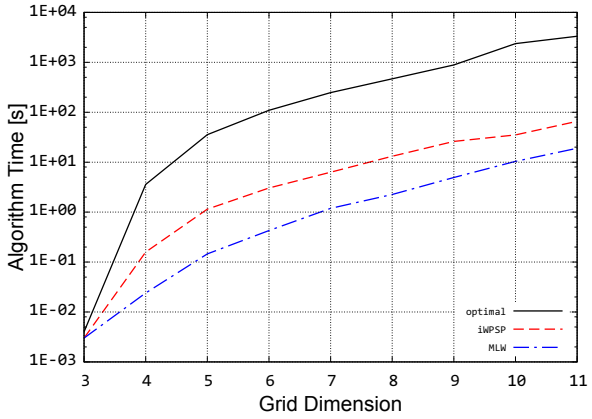

Figure 15: Complexity analysis and comparison

iWPSP is shown in Figure 12. The challenge area is shown in the red circle. The result of the MLW is shown in Figure 13 with its two paths. We only show the two paths from Amsterdam to Rome. The first path shown in red dashed link is Amsterdam–Hamburg–Berlin–Munich–Vienna–Zagreb–Rome, and the second path shown in blue solid link is Amsterdam–Brussels–Paris–Lyon–Rome. We present a large radius challenge case in Figure 14.

We present the execution time of the heuristics to demonstrate their effectiveness compared to the two-step algorithm in the case of calculating two d-separated paths. The evaluation is performed on a Linux machine with a 3.16 GHz Core 2 Duo CPU with 4 GB memory. We use different dimensions of grid networks to analyse the time complexity. The grid dimension ranges from 3×3 to 11×11, which means the number of nodes varies from 9 to 121. We present the time to calculate all the node pairs in the topology. When calculating only one path-pair that happens more often in real-world scenarios, the time is exponentially less. As shown in Figure 15, the x-axis is a grid dimension and the y-axis is the log-level algorithm execution time in seconds. Both the MLW and iWPSP algorithms show a better execution time compared to the two-step algorithm. For calculating all the paths in 11×11 grid, MLW takes 20 s, iWPSP takes 65 s, while the two-step algorithm takes more than 3000 s. We can observe that iWPSP has greater execution time compared to that of MLW. This is because of the extra time of Dijkstra's algorithm and selecting qualifying waypoint nodes. However, we observe that when calculating geodiverse paths in real-world topologies, iWPSP is more efficient in calculating the paths for the node pair around the topology boundary. This is because by selecting waypoints based on a distance and a delta value, iWPSP has more control over the distance separated from the two paths. One better algorithm might be combining the two heuristics in calculating one topology, and this will be analysed in future work.

6.1. GeoDivRP routing protocol

The implementation is done using ns-3 [46], a popular network simulator to analyse network protocols and network challenges. We base this protocol on the link state routing protocol methodology. At the beginning of the simulation, by obtaining node locations from the link state update messages that include node geolocation, we calculate the geodiverse paths and

store them in the path cache server. When the simulation begins, our protocol sends data traffic using the paths from the cache. When a challenge happens in the network, our protocol responds to the challenge faster than OSPF (Open Shortest Path First) [42] and calculates the paths according to the challenge estimation [14]. The distance value d is a user-provided value, and when challenges occur, network administers can modify d according to the different challenge characteristics and ensure the traffic quickly circumvents the challenged area. We have incorporated fallback mechanisms; when the generated d-separated paths do not satisfy the application requirement, OSPF is then used for further routing decisions.

We now present simulations using physical topologies including Sprint [30], Level 3 [47], Internet2 [48], and TeliaSonera [49]. We use CBR (constant bit rate) traffic, sending from each node to all the others at a data rate of one packet per second. We carry out the simulation once for each topology since there is no randomness because of CBR traffic. There are three area-based deterministic challenges we have simulated. From 20 to 40 s, the challenge occurs around Los Angeles, from 60 to 80 s in Kansas City, and the last challenge occurs in New York City from 100 to 120 s. The challenge locations come from the flow robustness analysis [14], and our challenge duration time is set as 20 s. We choose these different challenge areas so that the most vulnerable area is around Kansas City, due to its high betweenness as a major fibre exchange point in the US. The next damage area is around New York City. While it does not have many high-betweenness nodes, the network is dense and more nodes are challenged in a given radius. The least vulnerable area is around Los Angeles. The radii of the three challenge areas are 300 km. By assuming the correct estimation of the challenge radius and position, we compare our protocol's performance with standard OSPF in terms of the PDR (packet delivery ratio) as well as delay. Packet delivery ratio is the ratio of packets delivered divided by total packets sent, while delay is the time it takes for the data packet to travel end-to-end. We use the same challenge areas throughout all the topologies for ease of comparison. The iWPSP heuristic is used in the GeoDivRP for calculating the geodiverse paths. MLW achieves the same PDR and delay result as iWPSP when the links are carefully modified to guarantee the distance d-separation of the paths calculated. Since ns-3 is an event-driven network simulator and the algorithm execution time is not included in the simulation time, the delay in ns-3 for both the iWPSP and MLW is the same.

The Sprint physical network contains 77 nodes and 114 links. The PDR result for the Sprint network is shown in Figure 16a. We compare the performance of our GeoDivRP with standard OSPF. The second challenge in the Kansas City area occurs at 60 s and GeoDivRP shows substantial performance improvement compared to OSPF. The PDR of OSPF drops to 75% and it takes 10 s to converge while the time for the GeoDivRP is within one second and the PDR only drops by 2%. The paths calculated by the GeoDivRP to bypass the challenge is shown in Figure 17. The red circle shown in this figure is the challenge area. The last challenge occurs from 100 s to 120 s and the difference in the PDR between OSPF and GeoDivRP is small, only about 1%. This is because the challenge in New York City has little effect on the connectivity of the overall topology. The PDR for OSPF drops by about 1%, it takes 10 s to recover, and there is no noticeable PDR drop for our protocol. The first challenge happens at 20–40 s and there is no noticeable PDR drop for both of the protocols. This is due to the same reason as in New York City but the loss of the PDR to both the GeoDivRP and OSPF is even less.

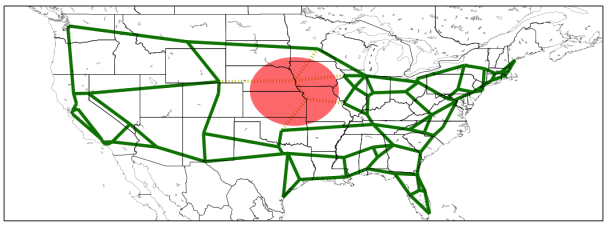

Figure 17: Sprint topology under regional challenges

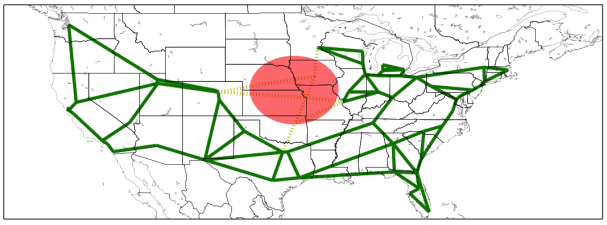

Figure 18: Level 3 topology under regional challenges

The delay analysis for the Sprint network is shown in Figure 16b. The reason that OSPF shows a lower delay when the network is under challenge compared to the GeoDivRP is because most of the data packets during the challenge have been dropped and the lost packets are not counted as delay. This is why there is a delay drop for OSPF before converging. Consider the first challenge in Figure 16b: the delay for OSPF drops from 20 to 30 s due to the packet losses, while the GeoDivRP converges and calculates geodiverse paths during that period of time and shows 1 s higher in delay. However, the extra delay is caused by an extra path stretch due to routing packets around the challenged area. We also notice a delay bump for OSPF right after the challenge is finished. For example, from 40 to 50 s, there is an increase in delay for OSPF. The same happens at 80–90 s, and 120–130 s. This is because OSPF needs to converge again after the topology has recovered from the challenge. In contrast, for our protocol, the convergence time is still one second and no noticeable delay increase is recorded.

The Level 3 physical network contains 99 nodes and 132 links. The PDR for the Level 3 network is shown in Figure 16c. Since Level 3 shares geographical similarities to the Sprint network, we observe a similar PDR result. The challenge in the Kansas City area reduces the PDR for OSPF significantly; it is even greater than for Sprint. This is because the Level 3 network lacks some of the nodes and links from Seattle to Chicago and the challenge around the Kansas City area causes more damage to the overall connectivity. As shown in Figure 18 using the Level 3 network, the similar challenge location as from

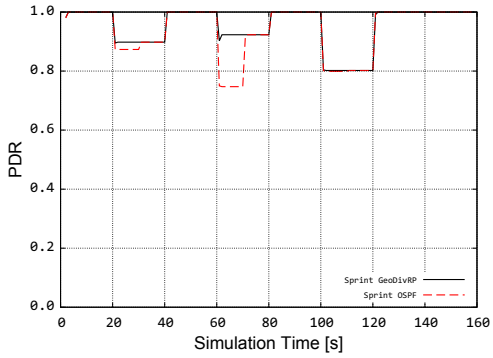

(a) Sprint PDR under area-based challenges

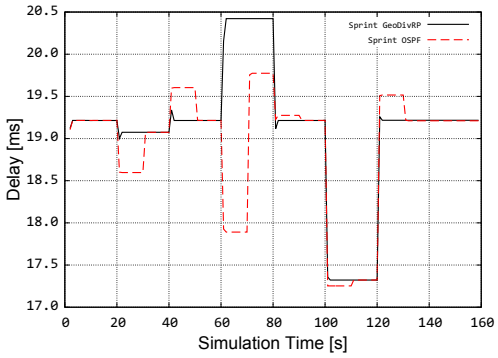

(b) Sprint delay under area-based challenges

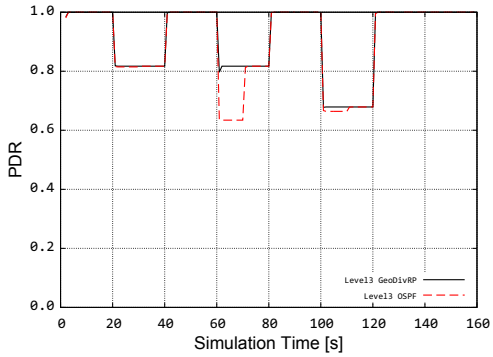

(c) Level 3 PDR under area-based challenges

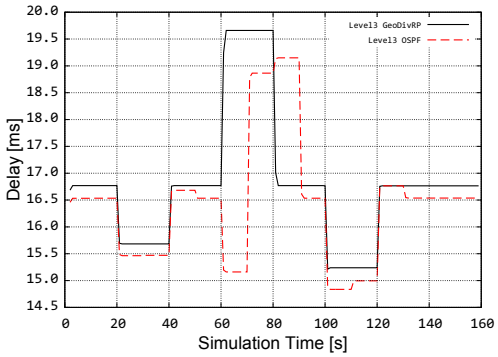

(d) Level 3 delay under area-based challenges

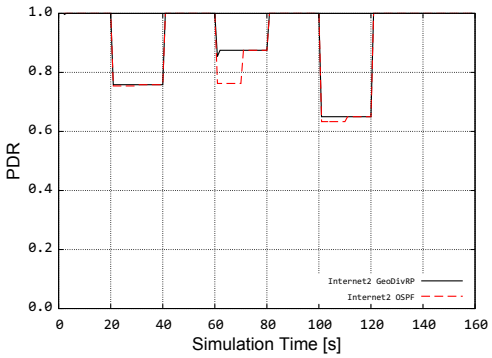

(e) Internet2 PDR under area-based challenges

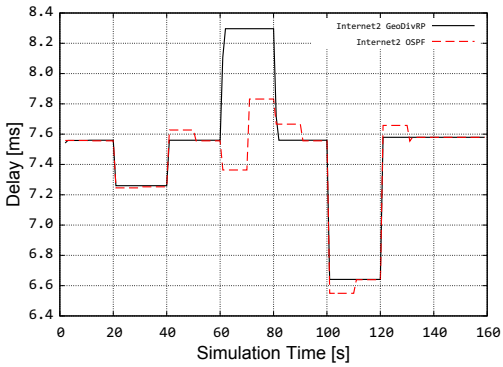

(f) Internet2 delay under area-based challenges

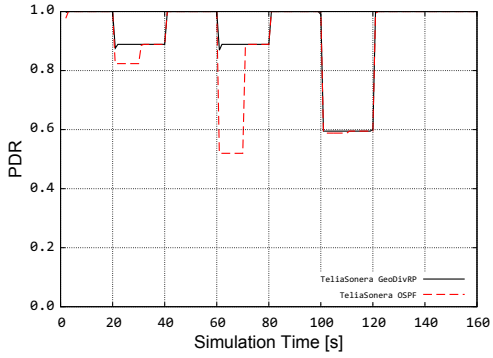

(g) TeliaSonera PDR under area-based challenges

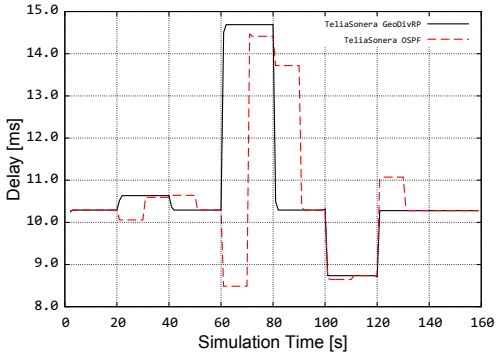

(h) TeliaSonera delay under area-based challenges

Figure 16: PDR and delay result under area-based challenges

the Sprint network has caused more nodes and links to fail. The delay case for the Level 3 network is similar to the Sprint network as shown in Figure 16d.

The Internet2 physical network is a smaller research network with only 16 nodes and 24 links. The PDR for the Internet2 network is shown in Figure 16e. The challenged PDR and delay show a similar trend. The first challenge does damage to the network connectivity and GeoDivRP converges within one second. The second challenge in Kansas City area causes OSPF to drop around 10% in the PDR and takes 10 s to converge and return the PDR to normal. The Los Angeles challenge has small impact on the network similar to the Sprint case. The delay analysis for the Internet2 network is shown in Figure 16f. For the same reason, OSPF shows a smaller delay compared to that of the GeoDivRP during challenges from 20 to 30 s, 60 to 70 s, and 100 to 110 s.

The TeliaSonera physical network contains 18 nodes and 21 links. The PDR for TeliaSonera is shown in Figure 16g. The second challenge at Kansas City area drops the PDR for OSPF to around 50%. This significant drop is caused by two reasons. First, the Kansas City node connects multiple nodes between the east and west coast. Second, the TeliaSonera network is very sparse so the damage from the Kansas City node is greater than that for the other networks. However, GeoDivRP recovers from the damage in only one second and limits the PDR drop within 1%. The PDR case for both the first and third challenges are similar. At the same time, OSPF drops about 1% of the total packets and recovers only after 10 s. The delay analysis is shown in Figure 16h. OSPF shows a smaller delay during challenges since the dropped packets are not counted for delay analysis. We notice that the delay increases after the challenge for OSPF at 80–90 s is larger than other challenge locations as well as the same challenge location in other topologies. This is because OSPF is using a path with more path stretch before convergence.

7. Conclusion and future work

We have proposed the path geodiversification mechanism and the global graph resilience metric cTGGD to characterise the geographical diversity for different physical topologies. We have verified its effectiveness in representing the geographical path diversity of a given topology. We have proposed a network vulnerability area identification mechanism and verified its effectiveness in identifying vulnerable areas in different physical network topologies. We have also implemented GeoDivRP that is capable of calculating and selecting single or multiple geographically diverse paths to meet the requirements from higher network layers. We have demonstrated its efficiency in routing around the challenged area and its improvement in both packet delivery ratio and delay compared to OSPF. We have also analysed how attackers could maximise the attack impact using a fixed budget with knowledge of the network structure and improve the effectiveness of restoration plans.

We have proposed two geodiversity heuristics to efficiently solve the path geodiverse problem (PGD): iWPSP (iterative WayPoint Shortest Path) and MLW (Modified Link Weight). We have implemented both of the heuristics in ns-3 and demonstrated the effectiveness of the heuristics in calculating and choosing different geographically diverse paths. It meets resilience requirements from the higher layers and shows better efficiency in routing data traffic around the challenged area. GeoDivRP shows significant improvement in both packet delivery ratio compared to OSPF with only 1 ms more delay on average. The two heuristics for GeoDivRP use different mechanisms to calculate geodiverse paths. By carefully modifying the link weights, the MLW is capable of providing one geodiverse path using one iteration of Dijkstra's algorithm. However, it is difficult to provide solutions when the paths required are for node pairs around a topology boundary, and the choice of a link weight needs to be carefully considered for different networks. On the other hand, the iWPSP requires one more run of Dijkstra's algorithm for each geodiverse path than the MLW; therefore, it takes a bit more time to execute and solve the PGD problem. However, by carefully selecting the waypoint node and the parameters d and δ, different topologies are similar and the iWPSP works better than the MLW when dealing with node pairs near topology boundaries.

For future work, we plan to implement these two heuristics in a testbed to emulate their effectiveness in real-world routers and examine the mechanisms to incorporate our geodiverse routing protocol into the current Internet. We will extend this work to analyse how the geographic multipath mechanism improves flow robustness. Furthermore, we will incorporate our protocol with ResTP (resilient transport protocol) to test the protocol stack and analyse the protocol performance with multiple geodiverse paths. We propose to fully analyse the relative benefits of the two heuristics and enable the GeoDivRP to automatically choose different heuristics in different network topologies. As for the resilience metric cTGGD, we plan to compare it against popular graph metrics in detail.

Acknowledgments

The authors would like to thank the members of the ResiliNets group for discussions that led to this work. This research was supported in part by US NSF Grant CNS-1219028 and CNS-1217736 (Resilient Network Design for Massive Failures and Attacks) at KU and UMKC.

References

[1] M. J. F. Denise M. B. Masi, Eric E. Smith, Understanding and Mitigating Catastrophic Disruption and Attack, Sigma Journal (September 2010) 16–22.

[2] D. Kuhn, Sources of failure in the public switched telephone network, Computer 30 (4) (1997) 31–36. doi:10.1109/2.585151.

[3] D. Oppenheimer, A. Ganapathi, D. A. Patterson, Why do internet services fail, and what can be done about it?, in: Proceedings of the 4th Conference on USENIX Symposium on Internet Technologies and Systems - Volume 4, USITS'03, USENIX Association, Berkeley, CA, USA, 2003, pp. 1–1. URL http://dl.acm.org/citation.cfm?id=1251460.1251461

[4] E. K. Çetinkaya, D. Broyles, A. Dandekar, S. Srinivasan, J. P. G. Sterbenz, Modelling Communication Network Challenges for Future Internet Resilience, Survivability, and Disruption Tolerance: A Simulation-Based Approach, Telecommunication Systems 52 (2) (2013) 751–766.

[5] J. Liu, X. Jiang, H. Nishiyama, N. Kato, Reliability Assessment for Wireless Mesh Networks Under Probabilistic Region Failure Model, Vehicular Technology, IEEE Transactions on 60 (5) (2011) 2253–2264. doi:10.1109/TVT.2011.2114684.
[6] H.-W. Lee, E. Modiano, K. Lee, Diverse routing in networks with probabilistic failures, IEEE/ACM Transactions on Networking 18 (6) (2010) 1895–1907. doi:10.1109/TNET.2010.2050490.
[7] J. P. Rohrer, A. Jabbar, J. P. G. Sterbenz, Path diversification: A multipath resilience mechanism, in: Proceedings of the IEEE 7th International Workshop on the Design of Reliable Communication Networks (DRCN), Washington, DC, 2009, pp. 343–351.
[8] J. P. Rohrer, A. Jabbar, J. P. Sterbenz, Path Diversification for Future Internet End-to-End Resilience and Survivability, Springer Telecommunication Systems 56 (1) (2014) 49–67.
[9] R. Cohen, K. Erez, D. ben Avraham, S. Havlin, Resilience of the Internet to Random Breakdowns, Phys. Rev. Lett. 85 (2000) 4626–4628. doi:10.1103/PhysRevLett.85.4626.
URL http://link.aps.org/doi/10.1103/PhysRevLett.85.4626
[10] D. Magoni, Tearing down the internet, IEEE J.Sel. A. Commun. 21 (6) (2006) 949–960. doi:10.1109/JSAC.2003.814364.
[11] Q. Zhou, L. Gao, R. Liu, S. Cui, Network Robustness under Large-Scale Attacks, 8th Edition, SpringerBriefs in Computer Science, 2013.
[12] R. Teixeira, K. Marzullo, S. Savage, G. M. Voelker, In Search of Path Diversity in ISP Networks, in: Proceedings of the 3rd ACM Internet Measurement Conference (IMC), Miami Beach, FL, 2003, pp. 313–318.
[13] A. Lakhina, J. Byers, M. Crovella, I. Matta, On the Geographic Location of Internet Resources, IEEE Journal on Selected Areas in Communications 21 (6) (2003) 934–948. doi:10.1109/JSAC.2003.814667.
[14] Y. Cheng, J. Li, J. P. G. Sterbenz, Path Geo-diversification: Design and Analysis, in: Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), Almaty, 2013.
[15] M. Gardner, C. Beard, Evaluating Geographic Vulnerabilities in Networks, in: IEEE International Workshop Technical Committee on Communications Quality and Reliability (CQR), 2011, pp. 1–6. doi:10.1109/CQR.2011.5996078.
[16] M. Gardner, C. Beard, D. Medhi, Using Network Measure to Reduce State Space Enumeration in Resilient Networks, in: Design of Reliable Communication Networks (DRCN), 2013 9th International Conference on the, 2013, pp. 250–257.
[17] M. Gardner, C. Beard, D. Medhi, Avoiding High Impacts of Geospatial Events in Mission Critical and Emergency Networks using Linear and Swarm Optimization, in: Cognitive Methods in Situation Awareness and Decision Support (CogSIMA), 2012 IEEE International Multi-Disciplinary Conference on, 2012, pp. 264–271. doi:10.1109/CogSIMA.2012.6188395.
[18] S. Neumayer, G. Zussman, R. Cohen, E. Modiano, Assessing the Impact of Geographically Correlated Network Failures, in: Military Communications Conference, 2008. MILCOM 2008. IEEE, 2008, pp. 1–6.
[19] S. Neumayer, E. Modiano, Network reliability with geographically correlated failures, in: Proc. of IEEE INFOCOM, 2010, pp. 1–9.
[20] W. Wu, B. Moran, J. Manton, M. Zukerman, Topology design of undersea cables considering survivability under major disasters, in: International Conference on WAINA, 2009, pp. 1154–1159.
[21] P. K. Agarwal, A. Efrat, S. Ganjugunte, D. Hay, S. Sankararaman, G. Zussman, Network Vulnerability to Single, Multiple, and Probabilistic Physical Attacks, in: MILITARY COMMUNICATIONS CONFERENCE, 2010 - MILCOM 2010, 2010, pp. 1824–1829.
[22] S. Neumayer, G. Zussman, R. Cohen, E. Modiano, Assessing the Vulnerability of the Fiber Infrastructure to Disasters, IEEE/ACM Transactions on Networking 19 (6) (2011) 1610–1623. doi:10.1109/TNET.2011.2128879.
[23] J. W. Suurballe, Disjoint paths in a network, Networks 4 (2) (1974) 125–145. doi:10.1002/net.3230040204.
URL http://dx.doi.org/10.1002/net.3230040204
[24] J. W. Suurballe, R. E. Tarjan, A quick method for finding shortest pairs of disjoint paths, Networks 14 (2) (1984) 325–336. doi:10.1002/net.3230140209.
URL http://dx.doi.org/10.1002/net.3230140209
[25] R. Bhandari, Survivable Networks: Algorithms for Diverse Routing, Kluwer Academic Publishers, Norwell, MA, USA, 1998.
[26] X. Long, D. Tipper, T. Gomes, Measuring the survivability of networks to geographic correlated failures, Optical Switching and Networking 14, Part 2 (0) (2014) 117–133. doi:http://dx.doi.org/10.1016/j.osn.2014.05.004.
[27] Y. Cheng, M. T. Gardner, J. Li, R. May, D. Medhi, J. P. Sterbenz, Optimised heuristics for a geodiverse routing protocol, in: Proceedings of the IEEE 10th International Workshop on the Design of Reliable Communication Networks (DRCN), Ghent, Belgium, 2014.
[28] C. J. Colbourn, D. D. Harms, Bounding all-terminal reliability in computer networks, Networks 18 (1) (1988) 1–12. doi:10.1002/net.3230180102.
[29] M. Motiwala, M. Elmore, N. Feamster, S. Vempala, Path Splicing, in: Proceedings of the ACM SIGCOMM, Seattle, WA, 2008, pp. 27–38.
[30] E. K. Çetinkaya, M. J. F. Alenazi, Y. Cheng, A. M. Peck, J. P. G. Sterbenz, On the Fitness of Geographic Graph Generators for Modelling Physical Level Topologies, in: Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), Almaty, 2013, pp. 38–45.
[31] J. P. Rohrer, M. J. F. Alenazi, J. P. G. Sterbenz, ResiliNets Topology Map Viewer (January 2011) [cited 2011-01-14].
URL http://www.ittc.ku.edu/resilinets/maps/
[32] P. K. Agarwal, A. Efrat, S. K. Ganjugunte, D. Hay, S. Sankararaman, G. Zussman, The resilience of WDM networks to probabilistic geographical failures, IEEE/ACM Transactions on Networking PP (99) (2013) 1. doi:10.1109/TNET.2012.2232111.
[33] S. P. Borgatti, Centrality and network flow, Social Networks 27 (1) (2005) 55–71. doi:http://dx.doi.org/10.1016/j.socnet.2004.11.008.
[34] S. P. Borgatti, M. G. Everett, A graph-theoretic perspective on centrality, Social Networks 28 (4) (2006) 466–484. doi:http://dx.doi.org/10.1016/j.socnet.2005.11.005.
[35] L. C. Freeman, A Set of Measures of Centrality Based on Betweenness, Sociometry 40 (1) (1977) 35–41.
[36] L. C. Freeman, Centrality in social networks conceptual clarification, Social Networks 1 (3) (1978–1979) 215–239.
[37] M. E. Newman, The mathematics of networks, The new palgrave encyclopedia of economics 2 (2008) 1–12.
[38] P. Holme, B. J. Kim, C. N. Yoon, S. K. Han, Attack vulnerability of complex networks, Phys. Rev. E 65 (2002) 056109.
[39] A. Sahoo, K. Kant, P. Mohapatra, Characterization of bgp recovery time under large-scale failures, in: IEEE ICC, Vol. 2, 2006, pp. 949–954. doi:10.1109/ICC.2006.254830.
[40] A. Atlas, A. Zinin, Basic Specification for IP Fast Reroute: Loop-Free Alternates, RFC 5286 (Proposed Standard) (Sep. 2008).
URL http://www.ietf.org/rfc/rfc5286.txt
[41] E. Bouillet, G. Ellinas, J.-F. Labourdette, R. Ramamurthy, Optical Networking, John Wiley & Sons, Ltd, 2007, pp. 1–23. doi:10.1002/9780470032985.ch1.
[42] J. Moy, OSPF specification, RFC 1131 (Proposed Standard), obsoleted by RFC 1247 (Oct. 1989).
URL http://www.ietf.org/rfc/rfc1131.txt
[43] M. Gardner, R. May, C. Beard, D. Medhi, Using multi-topology routing to improve routing during geographically correlated failures, in: Design of Reliable Communication Networks (DRCN), 2014 10th International Conference on the, 2014, pp. 1–8. doi:10.1109/DRCN.2014.6816134.
[44] V. Akgün, E. Erkut, R. Batta, On Finding Dissimilar Paths, European Journal of Operational Research 121 (2) (2000) 232–246.
[45] S. Orlowski, M. Pióro, A. Tomaszewski, R. Wessäly, SNDlib 1.0–Survivable Network Design Library, Networks 55 (3) (2010) 276–286.
[46] The ns-3 network simulator [online] (July 2009).
[47] Level 3 network map [online].
[48] Internet2 [online].
[49] TeliaSonera [online].

8. Author Biographies

Yufei Cheng is a Ph.D. student in the department of Electrical Engineering and Computer Science at the University of Kansas. He received his B.S. degree from Xidian University (China) in 2009, and his M.S. degree in Electrical Engineering from the University of Kansas in January 2014. He is a graduate research assistant at the KU Information & Telecommunication

Technology Center (ITTC). His research interests are survivable routing and multipath routing in both wired and wireless networks. He is a member of IEEE and ACM.

M. Todd Gardner is a Ph.D. Candidate in the School of Computing and Engineering at the University of Missouri, Kansas City. His M.S. and B.S. degrees in Electrical Engineering are from the University of Kansas and University of Missouri, Columbia, respectively in 2002 and 1990. He has spent over 20 years working for the U.S. Federal Aviation Administration on air traffic control systems as a communications and networking engineer. He is also a Registered Professional Engineer and an IEEE member. His research interests include high availability networks, resilient networks, and performance analysis.

Junyan Li graduated as a Master student in the department of Electrical Engineering and Computer Science at The University of Kansas. He received his B.S. degree in Information Security from Shanghai Jiao Tong University (China) in 2011. He is a graduate research assistant at the KU Information & Telecommunication Technology Center (ITTC). His research interest is resilient networks.

Rebecca "Becca" May is a senior undergraduate at the University of Missouri-Kansas City, USA majoring in Computer Science and Mathematics. She has been funded by the National Science Foundation – Research Experience for Undergraduates (NSF-REU) supplemental grant to work with Deep Medhi. Currently, she also works as a Junior SQL/C# Developer at Venture U.S. Enterprises in Kansas City, Missouri. Upon completion of her B.S. in Computer Science and Math, Becca hopes to advance her career in the field of computing and networking.

Deep Medhi is a Curators' Professor (and past Head) of the Computer Science & Electrical Engineering (CSEE) Department, School of Computing and Engineering (SCE) at the University of Missouri-Kansas City (UMKC). He received his B.Sc. (Hons.) degree in Mathematics from Cotton College, Gauhati University, India in 1981, and his M.S. degree in Mathematics from the University of Delhi, India in 1983. He then obtained his M.S. and Ph.D. degrees in Computer Sciences from the University of Wisconsin-Madison in 1985 and 1987, respectively. Prior to joining UMKC in 1989, he was a member of the technical staff in the traffic network routing and design department at the AT&T Bell Laboratories, Holmdel, New Jersey from 1987 to 1989. He is currently also an adjunct professor, Department of Computer Science & Engineering, Indian Institute of Technology-Guwahati (IIT-G), India.

James P.G. Sterbenz is a Professor of Electrical Engineering & Computer Science and on staff at the Information & Telecommunication Technology Center at The University of Kansas, a Visiting Professor of Computing and Communications in InfoLab 21 at Lancaster University in the UK, an Adjunct Professor of Computing at The Hong Kong Polytechnic University, and has been a Visiting Guest Professor at ETH Zürich. He received a doctorate in computer science from Washington University in St. Louis in 1991, with undergraduate degrees in Electrical Engineering, Computer Science, and Economics. He is director of the ResiliNets research group at KU, PI for the NSF-funded FIND Postmodern Internet Architecture project, PI for the NSF Multilayer Network Resilience Analysis and Experimentation on GENI project, lead PI for the GpENI (Great Plains Environment for Network Innovation) international GENI and FIRE testbed, co-I in the EU-funded FIRE ResumeNet project, and PI for the US DoD-funded highly-mobile airborne networking project. He has previously held senior staff and research management positions at BBN Technologies, GTE Laboratories, and IBM Research, where he has lead DARPA- and internally-funded research in mobile, wireless, active, and high-speed networks. He has been program chair for IEEE RNDM, GI, GBN, and HotI; IFIP IWSOS, PfHSN, and IWAN; and is on the editorial board of *IEEE Network*. He has been active in the Science and Engineering Fair organisation and judging in Massachusetts and Kansas for middle and high-school students. He is the principal author of the book High-Speed Networking: A Systematic Approach to High-Bandwidth Low-Latency Communication. He is a member of the IEEE, ACM, IET/IEE, and IEICE. His research interests include resilient, survivable, and disruption tolerant networking, future Internet architectures, active and programmable networks, and high-speed networking and systems.