Multilevel Resilience Analysis of Transportation and Communication Networks

Egemen K. Çetinkaya · Mohammed J.F. Alenazi · Andrew M. Peck · Justin P. Rohrer* · James P.G. Sterbenz

Received: date / Accepted: date

Abstract For many years the research community has attempted to model the Internet in order to better understand its behaviour and improve its performance. Since much of the structural complexity of the Internet is due to its multilevel operation, the Internet's multilevel nature is an important and non-trivial fea-

This research was supported in part by NSF FIND (Future Internet Design) Program under grant CNS-0626918 (Postmodern Internet Architecture), NSF grant CNS-1219028 (Resilient Network Design for Massive Failures and Attacks), by NSF grant CNS-1050226 (Multilayer Network Resilience Analysis and Experimentation on GENI), by the EU FP7 FIRE Programme ResumeNet project (grant agreement no. 224619), and by the Battelle Institute under contract number NFP00696666: Interdomain Resilience.

*Justin P. Rohrer is now with the Department of Computer Science, Naval Postgraduate School, Monterey, CA 93943– 5285 (e-mail: jprohrer@nps.edu); work performed primarily at The University of Kansas.

Egemen K. Çetinkaya

Department of Electrical Engineering and Computer Science, Information and Telecommunication Technology Center, The University of Kansas, Lawrence, KS 66045, USA ekc@ittc.ku.edu, +1 785 864 7290

Mohammed J.F. Alenazi malenazi@ittc.ku.edu

Andrew M. Peck apeck@ittc.ku.edu

Justin P. Rohrer rohrej@ittc.ku.edu

James P.G. Sterbenz

Department of Electrical Engineering and Computer Science, Information and Telecommunication Technology Center, The University of Kansas, Lawrence, KS 66045, USA http://www.ittc.ku.edu/resilinets jpgs@ittc.ku.edu, +1 508 944 3067 School of Computing and Communications, InfoLab21 Lancaster University, Lancaster, UK, jpgs@comp.lancs.ac.uk

ture that researchers must consider when developing appropriate models. In this paper, we compare the normalised Laplacian spectra of physical- and logical-level topologies of four commercial ISP and two research networks against the US freeway topology, and show analytically that physical level communication networks are structurally similar to the US freeway topology. We also generate synthetic Gabriel graphs of physical topologies and show that while these synthetic topologies capture the grid-like structure of actual topologies. they are more expensive than the actual physical level topologies based on a network cost model. Moreover, we introduce a distinction between geographic graphs that include degree-2 nodes needed to capture the geographic paths along which physical links follow, and structural graphs that eliminate these degree-2 nodes and capture only the interconnection properties of the physical graph and its multilevel relationship to logical graph overlays. Furthermore, we develop a multilevel graph evaluation framework and analyse the resilience of single and multilevel graphs using the flow robustness metric. We then confirm that dynamic routing performed over the lower levels helps to improve the performance of a higher level service, and that adaptive challenges more severely impact the performance of the higher levels than non-adaptive challenges.

Keywords Internet modelling, critical infrastructure, attack · graph spectrum, flow robustness, algebraic connectivity, spectral radius · network science, Gabriel graph, geographic graph, structural graph, multilevel graph, network cost · resilience, survivability, dependability, performability

1 Introduction and Motivation

The Internet has evolved to become a multilevel infrastructure critical to the functioning of society. The multilevel behaviour emerged in part due to the fact that protocols interact in multiple levels and in part because of the ways in which players operate, provide, and use the services of the Internet. Over the years, studies by the research community investigating the resilience of the Internet have suggested controversial findings [56], one being that an attack on a few central nodes could bring the entire Internet down. But this claim was dismissed by other researchers [11,29] based on the mesh-like structure of actual service-provider backbones. Therefore, *realistic* models are required to mathematically understand the properties of the Internet and improve its resilience.

The Internet can be examined at the physical, IP, router, PoP (point of presence), and AS (autonomous system) level from a topological point of view [28]. At the bottom is the physical topology consisting of elements such as fibre and copper cables, point-to-point wireless links, ADMs (add drop multiplexers), crossconnects, and layer-2 switches. The logical level consists of devices operating at the IP-layer. A PoP is a collection of routers in a geographic location, and PoPlevel topology can be seen as an aggregated view of the routers. At the AS-level, different provider networks peer with each other at IXPs (Internet exchange points) and private peering points [52]. Understanding the evolution of the Internet from a multilevel point of view is more realistic than examining its properties at individual levels. On the other hand, the primary focus of previous studies has been on the logical aspects of the topology, since tools have been developed to collect, measure, and analyse IP-level properties of the Internet (e.g. Rocketfuel [57]). However, given that physical networks provide the means of connecting nodes in the higher levels, the study of physical connectivity is an important area of research [30,37,44]. Furthermore, it is essential to model the impact of large-scale disasters and attacks against the physical infrastructure using the physical-level graph [21]. There are only a few studies that analyse graphs holistically from a multilevel point of view [45, 46, 51], but in very specific contexts.

In this paper, we begin our multilevel analysis of communication and transportation networks by expanding upon our previous work. First, we analyse the normalised Laplacian spectra of several commercial and research networks and find that *fibre* topologies are structurally similar to other critical infrastructures such as freeways. In our earlier work, we show this by examining only two commercial networks [19], in this study

we expand this to six networks. Second, we show that physical level topologies can be modelled by Gabriel graphs [34], since both are grid-like structures. Third, while Gabriel graphs capture the structure of the physical topologies, they are more expensive than the actual physical level topologies based on a network cost model. Fourth, we show that properties of *geographic* physical networks that capture information on the paths followed by links are dominated by degree-2 nodes. The removal of these degree-2 nodes does not change the structure significantly, but these *structural graphs* provide a more accurate representation of graph-theoretic properties such as degree distribution and clustering coefficient. Fifth, we develop a formal multilevel graph model and a framework to analyse flow robustness of a multilevel graph [22]. Furthermore, we extend this to include multiprovider graphs that captures logical IXP links. We analyse the flow robustness of a number of two-level graphs constructed from real-world communication networks and show that resilience of networks can be captured more realistically than a single level graph. Finally, as a side benefit of our work, we provide online adjacency matrices of these networks to the research community for future studies.

The rest of the paper is organised as follows: The communication and transportation topologies we use in this study are presented in Section 2. The evaluation of graphs using metrics is presented in Section 3. The correlation of transportation and communication networks using graph spectra is presented in Section 4. The structural properties of communication networks, as well as Gabriel graph models of physical networks is presented in Section 5. We evaluate multilevel transportation and communication graphs using flow robustness in Section 6. We analyse graphs and rank their resilience using flow robustness and spectral properties in Section 7. We provide a brief discussion in Section 8, followed by our conclusion as well as proposed future work in Section 9.

2 Topological Dataset

We study real networks (i.e. transportation and communication) that are geographically located within the continental United States. Therefore, we only include the 48 contiguous US states, the District of Columbia, and exclude Hawaii, Alaska, and other US territories. Furthermore, we have developed the KU-TopView (KU Topology Map Viewer) [59] using the Google Map API and JavaScript to visually present and assist in analysis of these topological maps. Unlike other visualisation tools, KU-TopView makes raw data conveniently available in the universal form of an adjacency matrix along



Fig. 1 Visual representation of transportation, and physical and logical level service provider networks in KU-TopView [10]

with the node coordinates and permits their manipulation. We have made these topologies publicly available [10].

2.1 Transportation Network

We have generated the freeway topology to represent the transportation network. Our starting point is the American Association of State Highway and Transportation Officials (AASHTO) data, which lists control cities and their sequential listing along each interstate highway. A control city is a major population center or destination on or near the interstate highway system determined by each state [13]. However, while generating the transportation topology, we realised that the existing list of control cities was not sufficient to represent the graph accurately. For example, there is no control city at some interchanges between interstate highways. Therefore, we add 6 additional $cities^1$ in those cases after verifying the crossing on Google Maps, as well as two that are needed to correspond to physical fiber junctions². There are also a few important newer freeways that are not listed in the 2001 AASHTO document that we add to reflect current connectivity³. This US freeway graph with 411 nodes, 553 links, and an average degree of 2.7 is shown in Figure 8a. We note that in a previous study of US interstate highway system,

the authors used GIS (geographic information system) databases from the year 2000 (unfortunately there is no reference to the source of data), and the resulting interstate freeway network consisted of 1337 links and 935 nodes with an average degree of 2.86 [35]. We note that the number of nodes and degree distribution in this *geographic graph* is highly dependent on the number of control cities used for geographic representation and that a number of cities are degree-2 vertices in between higher degree nodes at interchanges. We will discuss a uniform solution to this problem in Section 5.

2.2 Communication Networks

The Internet is a complex and large-scale network for which collective analysis is non-trivial. Therefore, we restrict this study to include physical fibre and logical level topologies. We note that throughout this paper we refer to IP router, PoP, and AS level graphs as *logical* level graphs or L3, whereas fibre level topologies as physical level topologies and denote them as L1. We use Rocketfuel-inferred AT&T, Level 3, and Sprint PoPlevel topologies [9,57] to study logical level topologies. We note that international links, as well as links crossing over Pacific and Atlantic Oceans, are removed intentionally to compare the logical level topologies against the US fibre deployments and freeway topologies.

We then use a US long-haul fibre-optic routes map data to generate physical topologies for AT&T, Sprint, and Level 3^4 [43]. In this map, US fibre-optic routes

¹ Benton Harbor MI, Country Club Hills IL, Effingham IL, Gary IN, Joilet IL, Lake Egypt IL

² Blaine WA, Hannibal MO

 $^{^3\,}$ I-335 Kansas Turnpike, I-86 East, I-97, I-68, I-495 in NY, and the important non-Interstate US-101 in California between Los Angeles and San Francisco

⁴ We also utilised the Level 3 network map in an effort to reflect the data as accurately as possible [4].

Topology	Star	Lin.	Tree	Ring	Grid	Tor.	Mesh	Star	Lin.	Tree	Ring	Grid	Tor.	Mesh
Nodes	10	10	10	10	10	10	10	100	100	100	100	100	100	100
Links	9	9	9	10	13	15	45	99	99	99	100	180	200	4950
Max. degree	9	2	3	2	3	3	9	99	2	3	2	4	4	99
Avg. degree	1.8	1.8	1.8	2	2.6	3	9	2	2	2	2	3.6	4	99
Deg. assort.	-1	-0.1	-0.5	1	0.3	1	1	-1	0	-0.3	1	0.6	1	1
Closeness	0.6	0.3	0.4	0.4	0.4	0.5	1	0.5	0	0.1	0	0.2	0.2	1
Clust. coeff.	0	0	0	0	0	0	1	0	0	0	0	0	0	1
Algeb. con.	1	0.1	0.2	0.4	0.4	1.4	10	1	0	0	0	0.1	0.4	100
Diameter	2	9	5	5	5	3	1	2	99	12	50	18	10	1
Radius	1	5	3	5	3	3	1	1	50	6	50	10	10	1
Hopcount	1.8	3.7	2.8	2.8	2.3	1.9	1	2	33.7	7.8	25.3	6.7	5	1
Max. Node betweenness	36	20	26	8	11	4	0	4851	2450	3068	1201	616	200	0
Max. Link betweenness	9	25	24	13	12	6	1	99	2500	2496	1250	341	200	1

Table 1 Topological characteristics of baseline networks

cross cities throughout the US and each ISP has a different coloured link. We project the cities to be physical node locations and connect them based on the map, which is sufficiently accurate on a national scale. We use this data to generate adjacency matrices for each individual ISP. To capture the geographic properties as well as the graph connectivity, cities are included as nodes even if they are merely a location along a link between fibre interconnection. As with the freeway graph, we will further discuss this in Section 5. Finally, we also make use of the publicly available Telia-Sonera⁵ network [6], Internet2 [3], and CORONET [8, 25] topologies. CORONET is a synthetic fibre topology designed to be representative of service provider fibre deployments, and this does not have a corresponding logical topology.

The physical and logical commercial service provider networks are shown in Figures 8b, 8c, 8d, and 8e. The Internet2 research network at the physical and logical level is shown in Figure 8f. Initial visual inspection suggests that the physical topologies are similar to the freeway topology. The relation of the physical level topology and other physical infrastructures has been stated before [59,62]; however, to best of our knowledge, we are not aware of any previous work that quantitatively demonstrates the correlation between these different infrastructures rigorously.

3 Properties of Networks

Although topology viewing is a powerful tool, it does not suffice for rigorous analysis of topologies [65]. We therefore calculate the graph metrics of regular networks (shown in Table 1) and critical infrastructures as shown in Table 2 using the Python NetworkX library [38].

3.1 Graph Metrics

Some of the well-known metrics provide insight on a variety of graph properties, including distance, degree of connectivity, and centrality. Network diameter, radius, and average hop count provide distance measures [37]. Eccentricity of a node is the longest shortest path from this node to every other node; the largest value of eccentricity among all nodes is the diameter and the smallest eccentricity is the radius. Betweenness is the number of shortest paths through a node or link and provides a centrality or importantness measure [49]. Clustering coefficient is a centrality measure of how well a node's neighbours are connected [37]. Closeness centrality is the inverse of the sum of shortest paths from a node to every other node [54,55]. Assortativity provides a measure of degree variance in a network [53]. Algebraic connectivity, a(G), is the second smallest eigenvalue of the Laplacian matrix [33]. For the graphs of the same order (number of vertices), algebraic connectivity provides a very good measure of how well the graph is connected and it indicates robustness of networks against node and link failures [40, 42, 48].

3.2 Graph Properties

Based on the metrics we describe above, we analyse regular graphs and real networks in this section.

3.2.1 Baseline Networks

We start our metrics-based analysis on seven regular graphs: star, linear, binary tree⁶, ring, grid, toroid, and

⁵ TeliaSonera physical graph has a link between Houston and Miami that appears to cross over the Gulf of Mexico. This is because TeliaSonera does not provide intermediate geographic path information.

 $^{^{6}}$ We note that not all leaves are binary as needed for a given order.

Network	Nodes	Links	Avg. Node Degree	Clust. Coeff.	Diam.	Rad.	Avg. Hop.	Close.	Max. Node Between.	Max. Link Between.
AT&T L1	383	488	2.6	0	39	20	14.1	0.1	17011	14466
AT&T L3	107	140	2.6	0.1	6	3	3.4	0.3	2168	661
Level 3 L1	99	130	2.6	0.1	19	10	7.7	0.1	1628	1046
Level 3 L3	38	376	19.8	0.8	3	2	1.5	0.7	59	37
Sprint L1	264	312	2.4	0	37	19	14.8	0.1	11275	9570
Sprint L3	28	76	5.4	0.4	4	2	2.2	0.5	100	27
TeliaSonera L1	21	25	2.4	0.2	9	6	4.1	0.3	75	61
TeliaSonera L3	16	29	3.6	0.5	4	2	2.1	0.5	34	17
Internet2 L1	57	65	2.3	0	14	8	6.7	0.2	630	521
Internet2 L3	9	13	2.9	0.4	4	2	2	0.5	9	11
CORONET L1	75	99	2.6	0	17	9	6.45	0.2	1090	704
US freeways	411	553	2.7	0.1	42	21	13.65	0.1	23872	19785

Table 2 Topological characteristics of communication and transportation networks

full mesh. We investigate the effect of an increase in the order (number of nodes) from n = 10 to n = 100for the baseline topologies as shown in Table 1. We note that the values are rounded to the nearest tenth decimal. The number of edges (links) are increased as necessary for each topology to scale to the number of nodes. Some metrics yield the same values for graphs of the same order (e.g. average degree for star, linear, tree), and others yield the same values for graphs of differing sizes and orders (e.g. same a(G) for 10 node linear and 100 node grid), therefore relying on a single metric for graph analysis is clearly not sufficient.

3.2.2 Real Networks

We investigate the graph-theoretic properties of the logical and the physical topologies of four commercial ISP networks (AT&T, Level 3, Sprint, TeliaSonera) and the Internet2 research network, as well as the fibre-link level of the CORONET synthetic topology. We also study the US Interstate Highway graph. Our results are shown in Table 2. In general, the metrics for the *logical* topologies differ from the *physical* topologies in that the physical topologies have more nodes and links compared to logical topologies.

The maximum degree of each provider's physical topology is less than that of its corresponding logical topology. This is due to the ability of logical topologies to arbitrarily overlay virtual links. The average degree of each provider's physical topology is less than that of its corresponding logical topology, in particular for the Level 3 topology in which the average degree for the logical level graph is a relatively highly meshed 19.8. Physical topologies have a higher value of network diameter, radii, and average hopcount than that of logical topologies. Betweenness values also differ for physical and logical topologies, showing a difference of one or two orders of magnitude higher for physical topologies. Clustering coefficient and closeness centrality metrics are also higher for the logical topologies compared to physical topologies.

From a distance metrics perspective, clearly physical topologies have higher values. We observe that the values of degree-based metrics also differ between physical and logical topologies. This can be attributed to the ease with which nodes can be connected in a logical topology as compared to the difficulty involved in connecting node in a physical topology, in which one must physically lay down fibre between nodes. Long links are added to logical topologies to reduce the forwarding overhead of multihop paths. From a centrality metrics perspective, we can see that physical topologies are not as clustered and have more homogeneous degree distributions.

We can also see that US freeway graph metrics are closer to those of the physical topologies. This is not surprising: both the US Interstate Highway system and the physical level of the Internet are physical infrastructures rather than logical overlays, and they frequently share the same paths since freeways (and railways) provide inexpensive right-of-way along which to lay fibre. Collective analysis of graph metrics provides a good indication of resilience of different topologies; however, it is difficult to infer sensible conclusions about the structure of a network or how similar two different networks are. Therefore, we redirect our attention to the spectra of these graphs.

4 Spectrum of Networks

In this section we first present a brief background on the spectra of graphs and then present our spectral analysis of the transportation and physical topologies.

4.1 Background on Spectra

Let G = (V, E) be an unweighted, undirected graph with *n* vertices and *m* edges. Let $V = \{v_0, v_1, \dots, v_{n-1}\}$ denote the vertex set and $E = \{e_0, e_1, \ldots, e_{m-1}\}$ denote the edge set. A graph can be represented by several methods including an adjacency matrix, incidence matrix, Laplacian matrix, and normalised Laplacian matrix [24,63]. A(G) is the symmetric adjacency matrix with no self-loops where $a_{ii} = 0$, $a_{ij} = a_{ji} = 1$ if there is a link between $\{v_i, v_j\}$, and $a_{ij} = a_{ji} = 0$ if there is no link between $\{v_i, v_j\}$. The Laplacian matrix of G is: L(G) = D(G) - A(G) where D(G) is the diagonal matrix of node degrees, $d_{ii} = \deg(v_i)$. Given degree of a node is $d_i = d(v_i)$, the normalised Laplacian matrix $\mathcal{L}(G)$ can be represented:

$$\mathcal{L}(G)(i,j) = \begin{cases} 1, & \text{if } i = j \text{ and } d_i \neq 0\\ -\frac{1}{\sqrt{d_i d_j}}, & \text{if } v_i \text{ and } v_j \text{ are adjacent}\\ 0, & \text{otherwise} \end{cases}$$

Let M be a symmetric matrix of order n and I be the identity matrix of order n. Then, eigenvalues (λ) and the eigenvector (**x**) of M satisfy $M\mathbf{x} = \lambda \mathbf{x}$ for $\mathbf{x} \neq 0$. In other words, eigenvalues are the roots of the characteristic polynomial, $det(M - \lambda I) = 0$. The set of eigenvalues $\{\lambda_1, \lambda_2, \ldots, \lambda_n\}$ together with their multiplicities (number of occurrences of an eigenvalue λ_i) define the *spectrum* of M. Spectral graph theory has been extensively covered in several monographs [15, 17, 24,26,63]. The spectrum of the AS-level topology of the Internet has been analysed based on the k largest values of the adjacency matrix [36]. The IP-level topology of the Internet has also been investigated and its Laplacian spectrum compared against synthetically generated topologies [41]. The normalised Laplacian spectrum of AS-level topologies has been shown to differ significantly from that of synthetically generated topologies [64]. Recently, a weighted spectral distribution metric has been proposed and has shown that synthetically generated graphs can be fine-tuned using spectral properties [32]. While previous studies utilised graph spectra to analyse *logical* level topologies, in this study we focus on *physical* networks and how they relate to each other structurally, as well as to their logical overlays.

4.2 Spectral Analysis of Networks

The normalised Laplacian spectrum provides insight into the structure of networks that are different in order (number of nodes) and size (number of links). The eigenvalues of the $\mathcal{L}(G)$ reside in the [0, 2] interval and take values $\{0 = \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n\}$. The algebraic multiplicity of $\lambda = 0$ indicates the number of connected components. Hence, there is always at least one eigenvalue equal to 0. Furthermore, matrices which resemble

one another may have similar eigenvalues and multiplicity. The spectrum of $\mathcal{L}(G)$ is quasi-symmetric⁷ around 1, which means a large algebraic multiplicity for the eigenvalue $\lambda = 1$ may indicate *duplications* in a network [14]. In other words, two separate nodes $\{u, v\}$ might have all or some of their neighbours being same. For example in a star graph with all other nodes connecting to the single central node, the leaves will all have the same neighbour, which is the central node. Likewise, while in a full mesh all nodes have the same neighbours, a partial mesh will have partial duplications. The presence of many small eigenvalue multiplicities may indicate that there are many components within a graph and these components are loosely connected to each other [14]. An eigenvalue of 2 indicates the graph is bipartite; eigenvalues close to 2 indicates the graph is nearly bipartite [14]. A bipartite graph is a graph in which its vertex set can be divided into two groups in such a way that there will be no edges between the vertices within each group. Once the discrete and deterministic eigenvalues are calculated for a given graph, the *relative frequency* of eigenvalues yield valuable information about the structure of a network. Moreover, spectra can be presented in *relative cumulative frequency* as well, and we describe our choice in the next section. For the rest of this paper we abbreviate relative frequency as RF and relative cumulative frequency as RCF.

4.2.1 Spectra of Baseline Networks

The RF (relative frequency) of the normalised Laplacian eigenvalues for baseline topologies (star, linear, ring, tree, grid, toroid, full mesh) of order n = 100 is shown in Figure 2. Since *most* of the eigenvalues have very small multiplicities, the RF of eigenvalues has a floor that is too noisy to be able to gather useful information. Because of the noisy floor in representing multiple RFs, we use the RCF (relative cumulative frequency) for the baseline graph analysis and for the rest of the paper. Furthermore, we note that while some researchers use RFs that they term density of eigenvalues to represent the spectra [14] and others use RCF that they term normalised index of eigenvalues to represent the spectra [64]. Since we show multiple curves in a plot to compare different graphs, our preference is to show spectra using RCFs since it is more informative.

The RCFs of the eigenvalues for these baseline topologies are shown in Figure 3. The star topology has

 $^{^{7}}$ We use the term quasi-symmetric to represent almost symmetric graph spectra. For example, a finite full-mesh graph is quasi-symmetric, since all eigenvalues except the first (which is equal to 0) are equal to a value close to 1. We will detail those graphs in the next section.



Fig. 2 Spectra of baseline topologies, n = 100



Fig. 3 Spectra of baseline topologies, n = 100



Fig. 4 Spectra of complete graphs

its eigenvalues fixed, independent of the graph order: $\{0 = \lambda_1 \leq 1 = \lambda_2 = \ldots = \lambda_{n-1} \leq \lambda_n = 2\}$. The spectrum of a 100 node full mesh looks similar to a star, except that it does not have an eigenvalue of 2 and the eigenvalues are fixed at 1.0101 (we comment on that later). An interesting observation is that the spectrum of these two baseline topologies look very similar. In-

deed, at a micro level we can think of each individual node in a mesh as a star motif. Furthermore, the algebraic connectivity of a star is 1 [33]. However, since node centrality measures are largest for a star topology, the central node in a star can be the target of an attack or the single point of failure from a network engineering perspective. An attack against the root node of a binary tree is also the worst case scenario, however, this partitions the network into two islands, in which nodes in each island can communicate with each other whereas this is impossible for a star topology. The spectrum of linear and ring topologies look almost identical, since a ring has an additional link compared to a linear topology, and both linear and ring topologies have the lowest algebraic connectivity values. Likewise, multiplicities of grid and toroid topologies look very similar, since a toroid has additional links to connect the nodes on the edge of a grid. We also observe that since a Manhattan grid is a combination of linear topologies, its spectrum looks similar to a linear topology. Multiplicities of a tree topology lie somewhere between the two extremes of mesh and linear.

We show the spectra of five different full-mesh complete graphs in Figure 4. The eigenvalues of a *n*-order complete graph are: $\{0 = \lambda_1 \leq \frac{n}{n-1} = \lambda_2 = \ldots = \lambda_n\}$. The multiplicity of the eigenvalue equal to n/(n-1) for complete graphs is n-1. Moreover, as the order of the graph approaches infinity, the eigenvalues will converge to a value of 1 since $\lim_{n\to\infty} \frac{n}{n-1} = 1$. However, eigenvalues λ_2 through λ_n are never exactly equal to 1 in a finite full mesh topology. Furthermore, the algebraic connectivity is equal to the order of a complete graph a(G) = n.

4.2.2 Spectra of Real Networks

We plot the RCFs of eigenvalues of US freeways against physical⁸ and logical level topologies in Figure 5 and Figure 6 respectively. Clearly, the spectra of the logical and physical topologies differ. Furthermore, the spectra of the physical topologies resemble the spectra of the US Interstate Highway graph as shown in Figure 5. This confirms our supposition that the properties of networks are similar since fibre is laid along right-of-ways, such as freeways. The spectra of logical level topologies along with the US Interstate Highway graph is shown in Figure 6. We intentionally include the transportation graph to compare it against the logical level topologies, which clearly shows the spectra do not match to that of freeways. The algebraic multiplicity for the eigenvalue

 $^{^{8}}$ These plots use the geographic version of the physical graphs; this will be explained in Section 5.

 $\lambda = 1$ is largest for the AT&T logical topology, indicating that this topology contains the largest number of node duplications. In other words, this topology has the most star-like components, as is evident by visually inspecting it on KU-TopView [10]. The largest eigenvalues indicate to what degree a graph is bipartite [14]. The largest eigenvalues of the physical topologies and the largest eigenvalues of the freeways graph are the eigenvalues closest to 2. Hence, the physical topologies and the freeways topology are the most nearly bipartite graphs.



Fig. 5 Spectra of geographical physical networks



Fig. 6 Spectra of logical networks

5 Structure of Communication Networks

In Section 4 we showed that the transportation graph closely matches the physical level topologies. In this section we analyse the structure of individual service provider networks and develop different models to represent physical level topologies. Furthermore, we calculate the fibre length of the physical level topologies.

5.1 Gabriel Graphs

In an attempt to understand the analytical models of physical level networks, we generate Gabriel graphs of the six service provider networks we study. The Gabriel graphs are useful in modelling graphs with geographic connectivity that resemble grids [34,50]. In a Gabriel graph, two nodes are connected directly if and only if there are no other nodes that fall inside the circle whose diameter is given by the line segment joining the two nodes. Since physical level communication networks resemble grids, we generate Gabriel graphs using the physical level network node locations. We note that the physical node locations we use in KU-TopView are in GPS (Global Positioning System) format, and we therefore convert node locations from GPS to ECEF (earth-centered, earth-fixed) format to accurately calculate the distance between any nodes. The number of nodes, links, and average degree of Gabriel graphs of six geographical physical networks (Gabriel-G) are shown in Table 3. In all cases, the synthetically generated Gabriel graph has more links than the corresponding geographical physical level topology. Note that since we use the same node locations, the number of nodes remain the same for Gabriel graphs.

5.2 Structural Graphs

As discussed in Section 3, the geographical physicallevel topologies consist of degree two intermediate nodes for accurate geographic representation. This is necessary to model area-based challenges on the network, such as power failures and severe weather [20,21]. However, these intermediate nodes artificially change the graph theoretic properties of the networks, in particular artificially skewing the degree distribution toward degree-2 nodes. Therefore, we modify the existing *geographical* physical level graphs by removing nodes with a degree of two, as long as there is not a logical level node at that location for which the physical node provides service to upper layers.

We show an example to emphasise the difference in representing geographic and structural physical level graphs in Figure 7. In this case, the map we have [43] has a path from Spokane, WA (dark coloured pin on upper left corner of Figure 7) to Billings, MT (cyan coloured pin on lower right corner of Figure 7) crossing five cities (Coeur d'Alene, ID; Thompson Falls, MT;

Provider	Graph	Nodes	Links	Avg. Degree	Tot. Link Length [km]	Avg. Link Length [km]	Max. Link Length [km]	Min. Link Length [km]
			100	Dogree		100	Longen [mm]	Longen [mm]
	Geographical	383	488	2.6	50,026	103	630	6
AT&T	Structural	130	191	2.9	37,489	196	1,195	6
11100.1	Gabriel-G	383	686	3.6	66,157	96	563	6
	Gabriel-S	130	218	3.4	36,459	167	695	6
	Geographical	99	130	2.6	28,538	220	1,063	10
Lovel 2	Structural	48	71	3.0	25,390	358	1,322	10
Level 3	Gabriel-G	99	170	3.4	33,991	200	736	10
	Gabriel-S	48	70	2.9	19,154	274	963	10
	Geographical	264	312	2.4	33,627	108	602	6
Sprint	Structural	52	73	2.8	25,190	345	1,299	11
Sprint	Gabriel-G	264	474	3.6	57,104	121	605	6
	Gabriel-S	52	73	2.8	19,853	272	1,064	11
	Geographical	21	25	2.4	14,190	568	1,592	27
TaliaSanara	Structural	18	21	2.3	14,040	669	1,592	41
Tenasonera	Gabriel-G	21	26	2.5	12,111	467	1,523	27
	Gabriel-S	18	23	2.6	12,482	543	1,523	41
	Geographical	57	65	2.3	19,050	293	910	41
Internet?	Structural	16	24	3.0	18,146	756	2,737	264
Internet2	Gabriel-G	57	94	3.3	28,786	306	871	41
	Gabriel-S	16	23	2.9	13,815	601	1,063	264
	Geographical	75	99	2.6	28,325	286	943	20
CODONET	Structural	39	63	3.2	27,579	438	1,195	44
CORONEI	Gabriel-G	75	127	3.4	33,265	262	758	20
	Gabriel-S	39	61	3.1	22,757	373	1,064	44

 Table 3 Structural properties of fibre topologies

Missoula, MT; Helena, MT; Bozeman, MT), forming a zigzag shaped path that captures geography of the path. This geographic physical graph is necessary to accurately study area-based challenges such as severe weather. On the other hand, this physical path from Spokane, WA to Billings, MT can be represented as a single link structurally. After all, it does not matter where the link is cut between Spokane, WA and Billings, MT since there is no logical PoP that is a traffic source or sink between these cities. We note that, the nodes that have a degree 3 or higher are kept to capture the physical layer structure even if there is not a logical PoP in these locations. Moreover, stub nodes are removed since we are interested in *backbone networks* and in the geographic representation these stub nodes might as well represent access networks. In this study, we primarily focus on node failures, and therefore we utilise structural physical level graphs.

corresponding physical level graph. However, with the exception of TeliaSonera, each structural graph has a larger average degree than its corresponding physical level graph. For example, the structural graph of Internet2 has 16 nodes, 24 links, and an average degree of 3 whereas the original Internet2 physical graph has 57 nodes, 65 links, and an average degree of 2.28. We believe that the structural graph of TeliaSonera has a smaller average degree than the original graph of Telia-Sonera due to the latter's small order and size. Furthermore, the TeliaSonera geographic graph was not obtained from the fiber map with significant detail of intermediate cities, but rather from a far more abstract map on TeliaSonera's Web site with only a few intermediate degree-2 nodes. However, we note that total fibre length of the structural graph (14,040 km) is close to that of the original physical graph (14,190 km).



Fig. 7 Geographical vs. structural graphs

The number of nodes, links, and average degree of these *structural* graphs are shown in Table 3. Each structural graph has fewer nodes and links than its Finally, we also generate synthetic Gabriel graphs using the node locations of the modified structural topologies. The number of nodes, links, average degree of nodes, and the total link length for these graphs (Gabriel-S) are shown in Table 3. An observation is that all these Gabriel graphs result in a total link length that is less than structural graphs. On the other hand, when we generate the Gabriel graphs of the geographical physical topologies, with the exception of TeliaSonera, these incur more cost than the geographical physical level topologies.



Fig. 8 Spectra of service provider networks

5.3 Spectral Comparison of Networks

Next, we study the spectra of logical level, physical level, Gabriel, and structural graphs as shown in Figure 8. We note that the CORONET synthetic topology does not include any logical level topology since CORO-NET is aimed at modelling a generic fibre topology. We observe that physical level, Gabriel, and structural graphs all share a similar spectrum due to the fact that all of them have a grid-like structure. The normalised Laplacian spectra of AT&T, Level 3, and Sprint logical level networks show an s-curve that corresponds to spectra of star- and mesh-like topologies (cf. Figure 3). On the other hand, the spectra of the TeliaSonera and Internet2 logical level topologies show a flatdiagonal curve that corresponds much more closely to spectra of grid-like topologies, as shown in Figure 1. Note that Internet2 is a research network that has a sparsely connected topology that is very similar to its physical topology.

We conclude that the logical and physical level topologies do indeed have different structures. The structural differences between the logical and physical level topologies, while impossible to determine using a *single* graph metric, are successfully captured by the normalised Laplacian spectrum. Moreover, the normalised Laplacian spectrum can be used to compare graphs of differing sizes and orders. Furthermore, while the spectra of the physical level topologies reveal grid-like structures, the spectra of the *large* logical level topologies reveal star- or mesh-like structures.

5.4 Cost of Networks

Structural properties impact the connectivity and *cost* of building networks. While at the logical level the cost is captured by the number of nodes and the capacity of each node (i.e. bandwidth and number of ports available in a router [11,29]), at the physical level, the length of the fibre dominates the cost. As we discussed, logical level links are arbitrarily overlaid links on top of the underlying physical links. Previously, we provided a network cost model as:

$$C_{i,j} = f + v \times d_{i,j} \tag{1}$$

where f is the fixed cost associated with link (including termination), v is the variable cost per unit distance for the link, and $d_{i,j}$ is the length of a link [39,58,59]. Based on the assumption that the variable cost dominates in long haul fibre networks, we ignore the fixed cost associated with links, and simplify network cost as:

$$C = \sum_{i} l_i \tag{2}$$

where l_i is the length of the *i*-th link [12,18]. With this simplified network cost model, we observe in Table 3

that, with the exception of TeliaSonera, each Gabriel graph incurs a cost that is approximately 20% larger than its corresponding original physical level graph. The Gabriel graph corresponding to TeliaSonera incurs a cost that is actually 15% smaller than that of its corresponding physical level graph. The structural topologies incur a slightly smaller cost than the geographical topologies because links are line-of-sight rather than following the geographic path of the actual fibre. We note that the number of nodes has changed from the original geographical physical to the modified structural graphs, since we only consider degree-3 and logical PoP nodes when constructing the structural graphs and delete degree-2 intermediate nodes needed for accurate geographic representation.

Physical Networks	Nodes	Links	Tot. $l \times 10^6$ [km]	Max. l [km]	Min. <i>l</i> [km]
AT&T	383	73, 153	116.8	4,400	6
Level 3	99	4,851	7.5	4,260	10
Sprint	264	34,716	57.8	4,330	6
TeliaSonera	21	210	0.4	4,079	27
Internet2	57	1,596	2.7	4,233	41
CORONET	75	2,775	4.6	4,260	20

Finally, for each physical level topology, we consider a baseline of the full-mesh topology whose vertex set is identical to that of the original topology. We then calculate the maximum link length, minimum link length, and total link length of each full-mesh topology as shown in Table 4. The maximum possible length of a fibre link is around 4,200 km for each topology corresponding to a diagonal of the continental US. Note that the total lengths are given in millions of km for a hypothetical full-mesh physical level topology, emphasising that real networks cannot have unlimited resilience due to cost constraints.

6 Multilevel Analysis

A holistic graph analysis of the Internet is non-trivial and does not exist to the best of our knowledge. Understanding the evolution of the Internet from a multilevel point of view is more realistic than examining its properties at individual levels. Therefore, we have developed a framework to analyse the *flow robustness* [54] of multilevel and multiprovider networks. When designing a resilient network, our main goal is providing service to the users in a cost-efficient manner. Hence, it is extremely important that we ensure connectivity between pairs of end systems. It is for this reason that we introduced the flow robustness metric, which quantifies resilience as the fraction of node pairs that remain connected in a network after it has been subjected to a number of node and link failures. Furthermore, we categorise networks in the following four groups:

- Single level, single provider: These networks consist of the physical or logical level of a single provider. Most previous studies analysed this type of graph [19, 21, 54].
- Single level, multiprovider: These networks consist of AS-level graphs that include several provider networks, but as a single adjacency matrix in which each provider is a single vertex of the graph. While several studies exist in the analysis of AS-level graphs [31], they treat multiprovider graphs at an abstract level (i.e. AS number), but they lack capturing the inter-AS level connectivity into the structure of the provider graphs.
- 3. Multilevel, single provider: These networks consist of multilevel graphs within a single provider. There are a few studies examining multilevel graphs for a single provider [45,46,51].
- 4. Multilevel, multiprovider: This type of model and analysis is the most realistic to capture the complexity of the Internet, and is the ultimate target of our research.

We begin our multilevel analysis of flow robustness of a 3-level graph and a 2-level graph in which the top two level graphs are the same. We show that the two multilevel graphs exhibit different performance and using fewer levels of graphs obscures accurate resilience evaluation of the top level of a multilevel graph. We then analyse the flow robustness of a number of twolevel graphs constructed from real-world communication networks. Next, we analyse a multiprovider graph, which is constructed by aggregating four different ISP networks into a single adjacency matrix. Our results confirm that it is difficult to partition the tier-1 ISP connectivity using attacks targetted at logical links.

6.1 Multilevel Graph Model

In an effort to further understand the structure of a number of communication networks, we employ a framework for studying multilevel graphs. A multilevel graph \mathcal{G} is a sequence of graphs, $\mathcal{G} = (G_{\ell_0}, G_{\ell_1}, ..., G_{\ell_{L-1}})$, ordered from lowest-level graph to highest-level graph where:

- 1. L is the number of levels
- 2. G_{ℓ_i} is the graph corresponding to level ℓ_i , where ℓ_i can be any desired label, given by $G_{\ell_i} = (V_{\ell_i}, E_{\ell_i})$



(a) Connected multilevel network



(b) Disconnected multilevel network (c) Partitioned multilevel network

- For all non-negative integers i and j such that i ≤ j, V_{ℓj} ⊆ V_{ℓi}
 For all non-negative integers i and j such that i ≤ j
- 4. For all non-negative integers i and j such that $i \leq j$ and all nodes u and v such that $u, v \in V_{\ell_j}$, if $\operatorname{conn}_{\ell_i}(u, v) = \text{false}$, then $\operatorname{conn}_{\ell_j}(u, v) = \text{false}$, where the function $\operatorname{conn}_{\ell_m}$ takes as its two parameters nodes in V_{ℓ_m} and returns true if the two nodes are connected in G_{ℓ_m} and false otherwise.

In other words, a multilevel graph consists of multiple graphs, one for each level, arranged such that for any pair of levels, the set of all nodes in the higher level is a subset of the set of all nodes in the lower level, and such that nodes that are not connected in a lower level are not connected in a higher level. In this paper, we only consider unweighted and undirected graphs. A connected multilevel graph is depicted in Figure 9a, and when a link is removed at the bottom level, this does not impact the higher level graphs if dynamic routing is utilised as shown in Figure 9b. Note that in Figure 9c, the removal of links (1, 6) and (3, 4) in the lowest level partitions the graph and necessitates the removal of all links between the disconnected clusters in the above levels as well.

A number of authors have discussed the importance of multilevel graphs as a means of further studying the resilience and survivability of the Internet [27, 45–47, 51, 61]. Some have even developed multilevel graph frameworks of their own [45, 46]. One study made use of a multilevel framework in order to study railway, peerto-peer, brain, and random graph topologies [46]. Each topology was subjected to random and loaded [45] link deletions, which were used to simulate errors and attacks, respectively. The robustness of each topology was then quantified in two different ways: as the fraction of logical link weight remaining and as the size of the largest connected component, both as a function of the number of link deletions. In our work, we study challenges [23] on multilevel networks by subjecting topologies to deletions drawn from a far more extensive group of graph metrics. Moreover, rather than treating robustness as the fraction of remaining logical link weight or as the size of the largest connected component, we consider the quantity *flow robustness*, which is defined as the fraction of node pairs that remained connected after a number of deletions [54].

We implement our model in Python. Our code takes as input a collection of adjacency matrices – one for each level – and stores them in a single multilevel graph data structure in memory, with the following requirements:

- 1. For any pair of levels, the set of all nodes in the level above are required to be a subset of the set of all nodes in the level below.
- 2. For any pair of levels, nodes that are disconnected from one another in the level below are also required to be disconnected from one another in the level above.

If the above requirements are met, we can then perform node and link deletions at any level and calculate any number of graph metrics with the help of the Python NetworkX library [38]. When node and link deletions are performed within a given level, the effects of the deletion are propagated to the higher levels to ensure that requirement 2 remains satisfied.

6.2 Multilevel Graph Analysis

We first employ our multilevel graph analysis framework to demonstrate the effect using multiple levels of graphs on the *service resilience* [60] at the top level. For this demonstrative analysis, we use a 3-level graph (US freeways, geographical physical, and logical-level topology of Internet2 research network) and a 2-level graph (physical- and logical-level topology of Internet2 research network) in which top two levels are identical

Fig. 9 Multilevel graph example

both for 3-level and 2-level graphs. We emphasise that the lowest graph in the 3-level graph is the freeways graph, and it does not provide a service in the conventional sense to the physical topology other than the provision of right-of-way. This is an example to show the impact of using multiple levels of graphs on evaluating the service resilience of the top level. For both the 3- and 2-level network, we perform random node and link deletions at the lowest level and observe how these deletions affect the highest level. Moreover, we consider the effects of these deletions under two separate scenarios – dynamic routing and static routing. Under perfect dynamic routing, we allow any pair of nodes in a given level to remain connected so long as there exists some path between them in the level below. Under static routing, which we show for worst-case baseline comparison, we immediately sever the connection between two nodes within a given level the moment that the shortest path between them in the level below is disrupted.



Fig. 10 Robustness of multilevel network for node deletions



Fig. 11 Robustness of multilevel network for link deletions

The results of this experiment are shown in Figure 10 for node deletions and in Figure 11 for link deletions. For both networks, the average flow robustness of the topmost level is plotted against the number of random deletions performed at the lowest level. For a given number of deletions, the average flow robustness was computed by averaging the flow robustness over 1000 failure sets, each of which was generated by performing the specified number of random deletions. For each value of average flow robustness on the curve, we also plot the 95% confidence interval. We note that the 3-level network has higher values of average flow robustness for any given number of deletions than the 2-level network. For example in Figure 10, when we delete 50 random nodes in the lowest topology of the 3-level graph (in the freeways graph), the flow robustness at the top level is approximately 0.55, whereas in a 2-level graph when we delete random 50 nodes in the lowest topology (in the physical topology), the flow robustness at the top level is approximately 0. This shows that adding multiple levels of graphs in resilience analysis impacts the outcome significantly. The difference when considering multiple levels is due to the fact that the bottom level graph has *nodes* that are a superset of the top 2 levels. We also note that if the US freeway topology was less connected (e.g. instead of a grid-like, it was linear) then the flow robustness would be lower. However, it is outside scope of this paper to analyse different connected graphs at the lower layers, and it will be part of our future work. Moreover, both the 3-level and 2-level network have higher values of average flow robustness under dynamic routing than under static routing. Finally as expected, average flow robustness diminishes more severely with node deletions than with link deletions since a single node deletion results in the deletion of all of its incident links.

Our framework can handle graphs with any number of levels. Part of the reason behind the experiment given above was to demonstrate the ability of our framework to handle multilevel graphs with more than two levels, in particular, the 3-level graph with the Internet2 physical and logical topologies in the two upper levels and the freeway right-of-way graph in the lowest level. We focus on 2-level communication networks for the rest of our multilevel analysis. To that end, we use the geographical physical and logical level adjacency matrices for each of AT&T, Level 3, Sprint, TeliaSonera, and Internet2 to create multilevel graphs for each network, and then perform node and link deletions within each multilevel graph at the physical level. Finally, we calculate the resulting flow robustness in the logical level for every failure set. The results of the experiments involving node deletions are shown in Figures 12 through 16,











Fig. 14 Sprint flow robustness for dynamic and static routing during adaptive and non-adaptive node deletions







Fig. 16 Internet2 flow robustness for dynamic and static routing during adaptive and non-adaptive node deletions

while the results of link deletions are shown in Figures 17 through 21.

As can be seen in Figures 12 through 21, in some of cases we delete nodes and links at random while in others we delete nodes and links with very specific properties. The former experiments serve as a baseline for comparison against the latter, which focus on those nodes and links with large values of certain forms of centrality – in particular, betweenness, closeness, degree, link betweenness, current-flow betweenness, and current-flow closeness. We discussed the first four metrics earlier in the paper; here, we define current-flow betweenness and current-flow closeness [16].

Unlike conventional betweenness and closeness that measure a node's centrality based on the shortest paths going through that node, current-flow betweenness and current-flow closeness are both ways of measuring a node's centrality based on information flow alone. To understand these two measures, we must first view the graph under consideration as an electrical network into which one unit of current enters from a node known as the *source* and from which one unit of current exits through another node known as the $sink^9$. The locations of the source and sink suffice to specify a unique current for each link in the network, as argued in Lemma 1 of [16]. Moreover, once each link is assigned a current, it is possible to assign absolute potentials to each node throughout the network, as argued in Lemma 2 of $[16]^{10}$.

The current-flow betweenness of a node in a graph is simply the average of the total current passing through that node (from all of its incident links) over all possible electrical networks resulting from different possible (source, sink) pairs. The current-flow closeness of a node in a graph is the inverse of the average over all other possible nodes of the potential difference between that node when it is treated as the source and the other node when it is treated as the source and the other node when it is treated as the sink. If we view "current" as information, then in essence, current-flow betweenness is a measure of the amount of information that can pass through a given node, while current-flow closeness is a measure of the ease with which information can be sent out from one node into the rest of the network.



(a) Dynamic, adaptive & non (b) Static, adaptive & nonFig. 17 AT&T flow robustness for link deletions



(a) Dynamic, adaptive & non (b) Static, adaptive & nonFig. 18 Level 3 flow robustness for link deletions



(a) Dynamic, adaptive & non(b) Static, adaptive & nonFig. 19 Sprint flow robustness for link deletions



(a) Dynamic, adaptive & non(b) Static, adaptive & nonFig. 20 TeliaSonera flow robustness for link deletions



(a) Dynamic, adaptive & non(b) Static, adaptive & nonFig. 21 Internet2 flow robustness for link deletions

 $^{^9}$ Note that the concept of an electrical network – and therefore the measures of current-flow closeness and current-flow betweenness – make sense only if the graph is simple and connected. That is why these measures, along with closeness, are employed only for *non-adaptive deletions*, explained in the subsequent paragraph.

¹⁰ In order to compute these potentials, we assign each link one unit of resistance. In other words, we employ the standard practice of assigning each link of an unweighted graph a length of one. This provides the ability to capture the link capacity in future analysis.

We use all of these measures (betweenness, closeness, degree, link betweenness, current-flow betweenness, and current-flow closeness) as a means to study what sorts of deletions at the physical level have the most disruptive effect at the logical level. Furthermore, we consider two different categories of deletions: adaptive deletions and non-adaptive deletions. A non-adaptive deletion is defined as a deletion performed based on the initial node or link centrality rankings that existed prior to the occurrence of any deletion. An adaptive deletion is defined as a deletion performed based on centrality rankings that are recomputed after the most recent deletion. This can result from an attacker that has real-time access to internal network management and operations information.

Finally, note that for centrality-based deletions we compute flow robustness, while for random deletions we compute *average* flow robustness in the same manner as before, that is by averaging the flow robustness over 1000 failure sets, each of which was generated by performing the number of random deletions. We also plot the 95% confidence intervals on each of the points located on the random curves.

As before, flow robustness diminishes more severely under static routing than under dynamic routing, and node deletions have a greater impact on flow robustness than link deletions. Furthermore, adaptive deletions have a more severe impact on the network than non-adaptive deletions. The reason for this should be clear: an adaptive deletion is always selecting from the pool of existing nodes or links the one with the highest centrality value, whereas a non-adaptive deletion will select from the pool of one that used to – but may no longer – have the highest centrality value. Hence, adaptive deletions have a far greater tendency to select the most important nodes or links than non-adaptive deletions, which results in a more severe impact on the flow robustness of the logical level.

Given a sufficiently small number of deletions, random deletions tend to have less effect on flow robustness than any other type of deletion. This is unsurprising, since deletions based on centrality metrics have a greater tendency to delete more "important" nodes and links than random deletions. What is surprising, however, is that, given a sufficient number of deletions, the flow robustness resulting from non-adaptive deletions based on closeness and current-flow closeness surpasses the average flow robustness resulting from random node deletions. This holds true for all five of the networks under study. For example in Figure 14b, with 40 random node deletions the flow robustness of the Sprint network is about 0.3, whereas the flow robustness for closeness is about 0.55. Similarly in Figure 14b, for 60 random node deletions the flow robustness is about 0.1 and for flow closeness the flow robustness is about 0.2. We speculate that since these are *non-adaptive* challenges, by the time network arrives in a state in which several nodes are deleted, initially calculated rankings are no longer accurate. However, why this happens *only* for closeness and current-flow closeness centrality metrics is not known. The reasons for the occurrence of this phenomenon, and an investigation into the types of multilevel graphs to which it is restricted, will be the subject of future work.

6.3 Multiprovider Graph Analysis

We introduce a new graph-theoretic model in which we define the concept of a *multiprovider graph*. Within our framework, a multiprovider graph is an ordered pair (G_{L3}, G_{AS}) , in which L3 represents PoP-level topology and AS represents the interprovider AS topology, where $G_{L3} = (V_{L3}, E_{L3})$ and $G_{AS} = (V_{AS}, E_{AS})$ are graphs such that:

- 1. the vertices in $V_{\rm AS}$ are mutually disjoint connected subgraphs of $G_{\rm L3}$ that, when taken together, contain all of the vertices in $V_{\rm L3}$. More specifically, if $V_{\rm AS} = \{v_1, v_2, ..., v_n\}$, then
 - (a) any two distinct vertices $v_i, v_j \in V_{AS}$ will be connected subgraphs of G_{L3} given by $v_i = (V_i, E_i)$ and $v_j = (V_j, E_j)$ such that $V_i \cap V_j = \emptyset$
 - (b) if we let $v_i = (V_i, E_i)$ for all integers i such that $1 \le i \le n$, then $\bigcup_{i=1}^n V_i = V_{L3}$.
- 2. there exists some function $f: E_{AS} \to 2^{E_{L3}}$ such that for any pair of distinct vertices $v_i, v_j \in V_{AS}$ given by $v_i = (V_i, E_i)$ and $v_j = (V_j, E_j)$, if $\{v_i, v_j\} \in E_{AS}$, then $f(\{v_i, v_j\}) = V_{ij} \cap E_{L3}$ where V_{ij} is the set of unordered pairs $\{u_i, u_j\}$ such that $\{u_i, u_j\} \in V_{ij}$ if and only if $u_i \in V_i$ and $u_j \in V_j$. More explicitly, the mapping f is used to identify edges between specific AS peer routers that serve to connect two ASes $v_i, v_j \in V_{AS}$ that share a given AS-edge $\{v_i, v_j\} \in E_{AS}$.

To study multiprovider graphs, first we combine the PoP-level topologies of four commercial ISPs (AT&T, Level 3, Sprint, TeliaSonera). We treat each ISP as a single AS, and the resulting AS-level abstract graph is a full-mesh with 4 nodes, in which each AS is connected to the other through a logical IXP (Internet exchange point) link. We select Atlanta NAP [1], Equinix [2], Terremark [7], and MAE-East [5] as the IXPs in which 4 ISPs are connected. The reason we select these 4 IXPs is that we analysed a number of IXP websites and found that these IXPs do provide service to the 4 commercial ISPs. We do not claim that this is an exhaustive list of IXPs, however, it was sufficient to generate a full-mesh AS-level graph for those tier-1 ISP providers. The 4 IXPs are distributed across the US in 17 different cities and there are 51 logical links that connected the four ISPs.



Fig. 22 Flow robustness of multiprovider network

In Figure 22, the flow robustness of a multiprovider graph is shown. In this case we delete all inter-AS IXP links in a *city*, ranked based on betweenness. As expected, adaptive attacks inflict more harm than nonadaptive attacks, which, in turn, inflict more harm than randomly-placed attacks. In Figure 22, the sharp reductions of flow robustness due to targetted attacks indicate the disconnection of an AS from the AS-level graph following such attacks. Note that several cities must be deleted in order to disconnect a single AS. In contrast, the flow robustness values in random scenarios decrease at a smoother rate because the flow robustness is averaged over 1000 failure sets. For example, the flow robustness values indicate that a very high percentage of the failure sets following the twelfth city deletion did not partition the network in any manner. Furthermore, our results indicate that it is very difficult to partition the tier-1 ISP connectivity, which is a full-mesh, given that it requires at least 9 cities and all the IXP links in a city to be destroyed. If we had included all IXPs in more than 17 cities, intuitively it would have been even more difficult to partition the AS-level graph.

Next, we analyse flow robustness of two provider graphs and their connections via IXPs. In this case, we investigate the flow robustness of provider duos as shown in Figure 23. Since these graphs are constructed by only two providers, the impact of random deletions and centrality-based attacks result in the same flow robustness. The connectivity between these provider duos



Fig. 23 Resilience of provider duos



Fig. 24 Resilience of Level 3, Sprint, TeliaSonera trio

breaks when the *n*th IXP link is broken. For example, AT&T and Level 3 peer with each other in 16 cities, and flow robustness remains at 1 until the 16th IXP link breaks, in which flow robustness drops to 0.

Finally, we analyse the flow robustness of a provider trio combination (Level 3, Sprint, TeliaSonera). The network performance when challenged by random failures and targeted attacks against IXP links are shown in Figure 24. The number of IXP connections reduced from 17 to 10 as shown in Figure 24 when we consider only three providers. As expected, the flow robustness depends on the number of IXP links between the provider trio.

7 Flow Robustness and Spectral Properties

Previously, we presented how physical communication topologies match the structure of the right-of-way of freeways using normalised Laplacian spectra. In this section, we compare the flow robustness against spec-

Network	Avg. Flow Robustness	FR Rank	a(G)	$a(G) \ {f Rank}$	$ ho(\mathcal{L})$	$egin{array}{c} ho(\mathcal{L}) \ { m Rank} \end{array}$
Level 3 L3	0.9413	1	0.9758	1	1.5037	1
Sprint L3	0.6503	2	0.6844	3	1.6361	2
TeliaSonera L3	0.5963	3	0.7669	2	1.7237	3
Internet2 L3	0.4779	4	0.4885	4	1.8091	4
AT&T L3	0.2996	5	0.1324	5	1.9127	5
TeliaSonera L1	0.1615	6	0.1178	6	1.9642	6
CORONET L1	0.0958	7	0.0401	7	1.9688	7
Level 3 L1	0.0721	8	0.0261	9	1.9811	9
Internet2 L1	0.0626	9	0.0386	8	1.9858	11
US freeways	0.0323	10	0.0055	10	1.9752	8
AT&T L1	0.0222	11	0.0055	11	1.9892	12
Sprint L1	0.0164	12	0.0053	12	1.9840	10

Table 5 Ranking of flow robustness and spectral properties

tral properties of the networks we study. The two important spectral properties we are interested are algebraic connectivity a(G) and the spectral radius $\rho(\mathcal{L})$. We present these metrics of five logical level topologies, six physical level network topologies, and US freeway graph in Table 5.

Flow robustness was defined in the previous section; here, we calculate its value for single level graphs for comparison. We approximate the average flow robustness of a given network by averaging the flow robustness over its 10,000 link-failure sets drawn uniformly and randomly from the pool of all of its link-failure sets. Next, we consider the algebraic connectivity a(G)of these topologies. Algebraic connectivity is the second smallest eigenvalue of the Laplacian matrix and is well-suited for measuring graph connectivity and for comparing the connectivities of graphs with the same order [33]. Finally, we consider the spectral radius of these 12 topologies. The spectral radius ρ is the absolute value of the maximum eigenvalue, $\rho = |\lambda_{\max}|$. Moreover, if $\rho(\mathcal{L}) = 2$, then the graph is bipartite, and the closer the spectral radii to 2, the closer the graph is to a bipartite. We calculate the spectral radius of the normalised Laplacian matrix $\rho(\mathcal{L})$ shown in column 6 of Table 5. We note that we previously studied spectral radii of Laplacian matrices $\rho(L)$ and adjacency matrices $\rho(A)$, but did not observe any pattern for $\rho(L)$ and $\rho(A)$ [19].

We rank the flow robustness of networks in descending order in columns 1 and 2. The logical topologies have higher values compared to the physical topologies and the US freeway graph. When we rank the topologies according to descending values of a(G), we observe a similar ranking order. In this case, only the rankings of the Internet2 and Level 3 physical topologies are swapped. Finally, we rank the spectral radii of these 12 topologies in ascending order. The ranking according to the spectral radii of the first seven topologies matches the rankings of the flow robustness and a(G). Our conclusion from this ranking comparison is that flow robustness, algebraic connectivity, and spectral radii are suitable metrics for the resilience analysis of networks.

8 Discussion

In previous sections we presented how communication and freeway right-of-way graphs relate to each other as well as an analysis of multilevel communication and transportation networks. In this section we discuss what these result mean for the network research community.

Existing models of the Internet often employ a single level perspective. However, as we demonstrated, multilevel graphs yield different performance measures than single level graphs under network perturbations. Moreover, given that the Internet is built on multiple levels structurally, synthetic topology generators as well as Internet models should consider its multilevel nature. Another important factor is the geography of the nodes. While statistical frameworks provide analytical models of the Internet, fibre routes are built on the right-of-way of the freeways, railways, and other critical infrastructures [59,62]. Moreover, freeways and other critical infrastructures are not built randomly or via the preferential attachment model-instead, they are built between adjacent major population centers in a way that will reduce the *cost* of building these infrastructures [39]. Eventually, this impacts how the fibre routes are laid out since placing fibre along a freeway is more feasible than placing it over mountains or other unreachable terrain even if it might result in a shorter physical path.

Another important point of discussion is the difference in graph-theoretic properties between the physical and logical topologies. While freeways and physical fibre route graphs share similar grid-like structural characteristics, logical overlays clearly differ from the physical underlays in terms of well-studied graph metrics, spectral properties, and flow robustness values. Physical topologies have higher distance metric values, whereas logical topologies have higher centrality metric values. Intuitively, one would expect that a richlyconnected underlay is needed to provide a resilient service to users. We have shown that while lower level graphs have grid-like structure and higher level graphs have mesh-like structure, the resilient service at the highest level depends on the protection of critical network elements that are common at all levels.

A third and final point of discussion is that, for a sufficiently large number of challenges, random failures can degrade the performance of a network more severely than *non-adaptive* attacks targeted at nodes based on *closeness* and *current-flow closeness* centrality metrics. This is an interesting result and we aim to investigate the reasons behind it in future work.

9 Conclusions and Future Work

Realistically modelling the Internet requires a collective and systemic analysis of all of its structural properties. Intuitively, fibre-level topologies are laid along right-ofway of the freeways, since it is less costly than lineof-sight installation. We analytically show structural similarities between these physical infrastructures by using the normalised Laplacian spectra. Moreover, we show that synthetically generated Gabriel graphs can model physical level topologies that are grid-like; however, with a greater cost incurred. We then show that geographical physical level graphs are dominated by degree-2 nodes, and removal of them provides more accurate structural metrics, particularly for degree distribution. Next, using the flow robustness metric we evaluated multilevel graphs and analysed combined communication and transport networks with our multilevel framework. We confirmed that dynamic routing helps alleviate the impact of perturbations and that adaptive challenges degrade multilevel network performance more than non-adaptive challenges. We also computed the flow robustness values of single level graphs and our results indicate that the rankings provided by important spectral properties are comparable to the rankings provided by the flow robustness metric.

Our future work will include the incorporation of other critical infrastructures such as railways, pipelines, and the power grid. We will also investigate the significance of the closeness and current-flow closeness centrality metrics in order to determine why – for a sufficient number of deletions – they have less impact on the network than random failures. We will evaluate multi-domain networks including tier-2 and tier-3 networks using our multilevel graph framework. Finally, we will generate graphs with specific properties and adding links and nodes to increase resilience.

Acknowledgements This is a significantly extended version and substantial revision of papers that appeared in the 4th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM) 2012 [19] and 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN) 2013 [22]. We would like to acknowledge members of the ResiliNets group for discussions on this work, as well as Robert Beverly, Maciej Kurant, Padma Krishnaswamy, and Kave Salamatian.

References

- 1. AtlantaNAP. http://www.atlantanap.com
- 2. Equinix. http://www.equinix.com
- 3. Internet2. http://www.internet2.edu
- 4. Level 3 network map. http://maps.level3.com
- 5. MAE-East. http://en.wikipedia.org/wiki/MAE-East
- 6. TeliaSonera. http://www.teliasoneraic.com
- 7. Terremark. http://www.terremark.com
- The Next Generation Core Optical Networks (CORO-NET). http://www.darpa.mil/Our_Work/STO/Programs/ Dynamic_Multi-Terabit_Core_Optical_Networks_ (CORONET).aspx
- Rocketfuel: An ISP topology mapping engine. http://www.cs.washington.edu/research/networking/ rocketfuel/interactive (2002)
- ResiliNets Topology Map Viewer. http://www.ittc.ku. edu/resilinets/maps/ (2011)
- Alderson, D., Li, L., Willinger, W., Doyle, J.C.: Understanding Internet Topology: Principles, Models, and Validation. IEEE/ACM Transactions on Networking 13(6), 1205–1218 (2005)
- Alenazi, M.J.F., Çetinkaya, E.K., Sterbenz, J.P.G.: Network Design and Optimisation Based on Cost and Algebraic Connectivity. In: Proceedings of the 5th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM). Almaty (2013)
- American Association of State Highway and Transportation Officials: Guidelines for the Selection of Supplemental Guide Signs for Traffic Generators Adjacent to Freeways. Washington, DC (2001)
- Banerjee, A., Jost, J.: Spectral characterization of network structures and dynamics. In: N. Ganguly, A. Deutsch, A. Mukherjee (eds.) Dynamics On and Of Complex Networks, Modeling and Simulation in Science, Engineering and Technology, pp. 117–132. Birkhäuser Boston (2009)
- Biggs, N.: Algebraic Graph Theory, 2nd edn. Cambridge University Press (1993)
- Brandes, U., Fleischer, D.: Centrality measures based on current flow. In: Proceedings of the 22nd Annual Conference on Theoretical Aspects of Computer Science (STACS), *LNCS*, vol. 3404, pp. 533–544. Springer Berlin / Heidelberg, Stuttgart (2005)
- Brouwer, A.E., Haemers, W.H.: Spectra of Graphs. Springer New York (2012)
- Cetinkaya, E.K., Alenazi, M.J.F., Cheng, Y., Peck, A.M., Sterbenz, J.P.G.: 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)

- Çetinkaya, E.K., Alenazi, M.J.F., Rohrer, J.P., Sterbenz, J.P.G.: Topology Connectivity Analysis of Internet Infrastructure Using Graph Spectra. In: Proceedings of the 4th IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), pp. 752–758. St. Petersburg (2012)
- Çetinkaya, E.K., Broyles, D., Dandekar, A., Srinivasan, S., Sterbenz, J.P.G.: A Comprehensive Framework to Simulate Network Attacks and Challenges. In: Proceedings of the 2nd IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), pp. 538–544. Moscow (2010)
- Çetinkaya, E.K., Broyles, D., Dandekar, A., Srinivasan, S., Sterbenz, J.P.G.: Modelling Communication Network Challenges for Future Internet Resilience, Survivability, and Disruption Tolerance: A Simulation-Based Approach. Telecommunication Systems 52(2), 751–766 (2013)
- Çetinkaya, E.K., Peck, A.M., Sterbenz, J.P.G.: Flow Robustness of Multilevel Networks. In: Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN), pp. 274– 281. Budapest (2013)
- Çetinkaya, E.K., Sterbenz, J.P.G.: A Taxonomy of Network Challenges. In: Proceedings of the 9th IEEE/IFIP International Conference on the Design of Reliable Communication Networks (DRCN), pp. 322–330. Budapest (2013)
- Chung, F.R.K.: Spectral Graph Theory. American Mathematical Society (1997)
- Clapp, G., Skoog, R.A., Von Lehmen, A.C., Wilson, B.: Management of Switched Systems at 100 Tbps: the DARPA CORONET Program. In: International Conference on Photonics in Switching (PS), pp. 1–4. Pisa (2009)
- Cvetković, D., Rowlinson, P., Simić, S.: An Introduction to the Theory of Graph Spectra. London Mathematical Society (2009)
- Demeester, P., Gryseels, M., Autenrieth, A., Brianza, C., Castagna, L., Signorelli, G., Clemenfe, R., Ravera, M., Jajszczyk, A., Janukowicz, D., Doorselaere, K.V., Harada, Y.: Resilience in multilayer networks. IEEE Communications Magazine **37**(8), 70–76 (1999)
- Donnet, B., Friedman, T.: Internet topology discovery: a survey. IEEE Communications Surveys & Tutorials 9(4), 56–69 (2007)
- Doyle, J.C., Alderson, D.L., Li, L., Low, S., Roughan, M., Shalunov, S., Tanaka, R., Willinger, W.: The 'robust yet fragile' nature of the Internet. Proceedings of the National Academy of Sciences of the United States of America 102(41), 14,497–14,502 (2005)
- Durairajan, R., Ghosh, S., Tang, X., Barford, P., Eriksson, B.: Internet Atlas: A Geographic Database of the Internet. In: Proceedings of the 5th ACM HotPlanet Workshop. Hong Kong (2013)
- Faloutsos, M., Faloutsos, P., Faloutsos, C.: On power-law relationships of the Internet topology. In: Proceedings of the ACM SIGCOMM, pp. 251–262. Cambridge, MA (1999)
- 32. Fay, D., Haddadi, H., Thomason, A., Moore, A., Mortier, R., Jamakovic, A., Uhlig, S., Rio, M.: Weighted Spectral Distribution for Internet Topology Analysis: Theory and Applications. IEEE/ACM Transactions on Networking 18(1), 164–176 (2010)

- Fiedler, M.: Algebraic connectivity of graphs. Czechoslovak Mathematical Journal 23(2), 298–305 (1973)
- Gabriel, K.R., Sokal, R.R.: A New Statistical Approach to Geographic Variation Analysis. Systematic Zoology 18(3), 259–278 (1969)
- Gastner, M.T., Newman, M.E.: The spatial structure of networks. The European Physical Journal B - Condensed Matter and Complex Systems 49(2), 247–252 (2006)
- Gkantsidis, C., Mihail, M., Zegura, E.: Spectral Analysis of Internet Topologies. In: Proceedings of the IEEE IN-FOCOM, vol. 1, pp. 364–374. San Francisco, CA (2003)
- Haddadi, H., Rio, M., Iannaccone, G., Moore, A., Mortier, R.: Network Topologies: Inference, Modeling, and Generation. IEEE Communications Surveys & Tutorials 10(2), 48–69 (2008)
- Hagberg, A.A., Schult, D.A., Swart, P.J.: Exploring Network Structure, Dynamics, and Function using NetworkX. In: 7th Python in Science Conference (SciPy), pp. 11–15. Pasadena, CA (2008)
- Hameed, M.A., Jabbar, A., Çetinkaya, E.K., Sterbenz, J.P.: Deriving Network Topologies from Real World Constraints. In: Proceedings of IEEE GLOBECOM Workshop on Complex and Communication Networks (CC-Net), pp. 400–404. Miami, FL (2010)
- 40. Jamaković, A., Uhlig, S.: On the relationship between the algebraic connectivity and graph's robustness to node and link failures. In: Proceedings of the 3rd EuroNGI Conference on Next Generation Internet Networks, pp. 96–102. Trondheim (2007)
- Jamaković, A., Van Mieghem, P.: The Laplacian Spectrum of Complex Networks. In: Proceedings of the European Conference on Complex Systems (ECCS). Oxford (2006)
- Jamaković, A., Van Mieghem, P.: On the Robustness of Complex Networks by Using the Algebraic Connectivity. In: Proceedings of the 7th International IFIP Networking Conference, *Lecture Notes in Computer Science*, vol. 4982, pp. 183–194. Singapore (2008)
- KMI Corporation: North American Fiberoptic Long-haul Routes Planned and in Place (1999)
- Krioukov, D., claffy, k., Fomenkov, M., Chung, F., Vespignani, A., Willinger, W.: The Workshop on Internet Topology (WIT) Report. ACM Comput. Commun. Rev. 37(1), 69–73 (2007)
- Kurant, M., Thiran, P.: Layered complex networks. Phys. Rev. Lett. 96, 138,701 (2006)
- Kurant, M., Thiran, P., Hagmann, P.: Error and attack tolerance of layered complex networks. Phys. Rev. E 76(2), 026,103 (2007)
- 47. Lehman, T., Yang, X., Ghani, N., Gu, F., Guok, C., Monga, I., Tierney, B.: Multilayer Networks: An Architecture Framework. IEEE Communications Magazine 49(5), 122–130 (2011)
- Liu, W., Sirisena, H., Pawlikowski, K., McInnes, A.: Utility of algebraic connectivity metric in topology design of survivable networks. In: Proceedings of the 7th IEEE International Workshop on Design of Reliable Communication Networks (DRCN), pp. 131–138. Washington, DC (2009)
- Mahadevan, P., Krioukov, D., Fomenkov, M., Dimitropoulos, X., claffy, k., Vahdat, A.: The Internet ASlevel topology: three data sources and one definitive metric. ACM Comput. Commun. Rev. 36(1), 17–26 (2006)
- 50. Matula, D.W., Sokal, R.R.: Properties of Gabriel Graphs Relevant to Geographic Variation Research and the Clustering of Points in the Plane. Geographical Analysis 12(3), 205–222 (1980)

- 51. Medhi, D., Tipper, D.: Multi-layered network survivability-models, analysis, architecture, framework and implementation: An overview. In: Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX), vol. 1, pp. 173–186. Hilton Head Island, SC (2000)
- 52. Metz, C.: Interconnecting ISP Networks. IEEE Internet Computing 5(2), 74–80 (2001)
- Newman, M.E.J.: Assortative mixing in networks. Phys. Rev. Lett. 89(20), 208,701 (2002)
- 54. Rohrer, J.P., Jabbar, A., Sterbenz, J.P.: Path Diversification for Future Internet End-to-End Resilience and Survivability. Springer Telecommunication Systems (2012)
- Rohrer, J.P., Sterbenz, J.P.G.: Predicting topology survivability using path diversity. In: Proceedings of the IEEE/IFIP International Workshop on Reliable Networks Design and Modeling (RNDM), pp. 95–101. Budapest (2011)
- Roughan, M., Willinger, W., Maennel, O., Perouli, D., Bush, R.: 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems. IEEE Journal on Selected Areas in Communications 29(9), 1810–1821 (2011)
- 57. Spring, N., Mahajan, R., Wetherall, D., Anderson, T.: Measuring ISP topologies with Rocketfuel. IEEE/ACM Transactions on Networking **12**(1), 2–16 (2004)
- Sterbenz, J.P., Çetinkaya, E.K., Hameed, M.A., Jabbar, A., Rohrer, J.P.: Modelling and analysis of network resilience (invited paper). In: Proceedings of the Third IEEE International Conference on Communication Systems and Networks (COMSNETS), pp. 1–10. Bangalore (2011)
- Sterbenz, J.P., Çetinkaya, E.K., Hameed, M.A., Jabbar, A., Shi, Q., Rohrer, J.P.: Evaluation of Network Resilience, Survivability, and Disruption Tolerance: Analysis, Topology Generation, Simulation, and Experimentation (invited paper). Telecommunication Systems 52(2), 705–736 (2013)
- 60. Sterbenz, J.P.G., Hutchison, D., Çetinkaya, E.K., Jabbar, A., Rohrer, J.P., Schöller, M., Smith, P.: Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. Computer Networks 54(8), 1245–1265 (2010)
- 61. Sterbenz, J.P.G., Hutchison, D., Çetinkaya, E.K., Jabbar, A., Rohrer, J.P., Schöller, M., Smith, P.: Redundancy, Diversity, and Connectivity to Achieve Multilevel Network Resilience, Survivability, and Disruption Tolerance (invited paper). Springer Telecommunication Systems (2012). (accepted April 2012)
- Strand, J., Chiu, A.L., Tkach, R.: Issues for routing in the optical layer. IEEE Communications Magazine 39(2), 81–87 (2001)
- Van Mieghem, P.: Graph Spectra for Complex Networks. Cambridge University Press (2011)
- 64. Vukadinović, D., Huang, P., Erlebach, T.: On the Spectrum and Structure of Internet Topology Graphs. In: Proceedings of the Second International Workshop on Innovative Internet Computing Systems (IICS), *Lecture Notes in Computer Science*, vol. 2346, pp. 83–95. Kühlungsborn (2002)
- Zegura, E.W., Calvert, K.L., Donahoo, M.J.: A Quantitative Comparison of Graph-Based Models for Internet Topology. IEEE/ACM Transactions on Networking 5(6), 770–783 (1997)

10 Author Biographies

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

Mohammed J.F. Alenazi: is a Ph.D. student in the department of Electrical Engineering and Computer Science at The University of Kansas. He received his B.S. and M.S. degrees in Computer Engineering from the University of Kansas in 2010 and 2012 respectively. He is a graduate research assistant in the ResiliNets research group at the KU Information & Telecommunication Technology Center (ITTC). His research interests are in resilient networks, particularly disruption-tolerant networks. He is a member of the IEEE and ACM.

Andrew M. Peck: is an M.S. student in the department of Electrical Engineering and Computer Science at The University of Kansas. He received his A.B. degree in Astrophysical Sciences from Princeton University in 2011. He is a graduate research assistant in the ResiliNets research group at the KU Information & Telecommunication Technology Center (ITTC). His research interests are in resilient networks, particularly multilevel network analysis.

Justin P. Rohrer: is currently a Research Associate of Computer Science at the Naval Postgraduate School (NPS) and an Adjunct Assistant Professor of Electrical Engineering and Computer Science at the KU Information & Telecommunication Technology Center (ITTC). He received his Ph.D in Electrical Engineering from the University of Kansas in 2011 with honors. He received his B.S. degree in Electrical Engineering from Rensselaer Polytechnic Institute, Troy, NY, in 2004. From 1999 to 2004, he was with the Adirondack Area Network, Castleton, NY as a network engineer. He was also an ITTC Graduate Fellow from 2004–2006. He received the best paper award at the International Telemetering Conference in 2008 and the best graduate student paper award at the same conference in 2011. His research focus is on resilient and survivable transport and routing protocols. Interests also include highly-dynamic mobile

networks, and simulating network disruptions. Previous research has included weather disruption-tolerant mesh networks and free-space optical metropolitan networks. He is a member of the IEEE Communications and Computer Societies, ACM SIGCOMM, Eta Kappa Nu, and was an officer of the Kansas City section of the IEEE Computer Society.

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