Flow Robustness of Multilevel Networks

Egemen K. Çetinkaya, Andrew M. Peck, and James P.G. Sterbenz
Information and Telecommunication Technology Center
Department of Electrical Engineering and Computer Science
The University of Kansas, Lawrence, KS, 66045, USA
{ekc, apeck, jpgs}@ittc.ku.edu
www.ittc.ku.edu/resilinets

***Abstract*—The Internet topology can be viewed at the physical-, router-, PoP-, and AS-level. Intuitively, a richly connected lower level can improve the resilience of a service at higher levels. Understanding the structure of the Internet from a multilevel point of view is more realistic than examining its properties at individual levels. We have developed a framework to analyse the robustness of multilevel and multiprovider networks. We show that multilevel graphs exhibit different performance from single level graphs. Moreover, our framework validates the robustness of the Internet core due to its rich connectivity.**

***Keywords*-Internet, topology modelling, multilevel network, multiprovider network, resilience, flow robustness**

I. Introduction and Motivation

The Global Internet is formed by the interconnection of ASes (autonomous systems). BGP (Border Gateway Protocol) enables the communication among these ASes. It has a hierarchical structure in which tier-1 ISP (Internet Service Provider) networks reside on top of this hierarchy [1]. Although traffic measurements indicate an evolution towards a flatter topology, structurally a loose hierarchy remains [2].

The Internet topology can be described by functional levels [3]. The levels are as follows: physical-, router-, PoP- (point of presence), and AS-level. An abstract view of different levels of the Internet is shown in Figure 1. At the bottom is the physical topology consisting of components such as fibre and copper cables, switches, and ADMs (add drop multiplexers). The router level consists of devices operating at the IP-layer. A PoP is a collection of routers in a geographic location, and PoP-level topology can be seen as an aggregated view of the routers. At the AS-level, different provider networks peer with each other at the IXPs (Internet eXchange Points) and private peering points [4].

A holistic graph analysis that systematically analyses the Internet as a multilevel infrastructure 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 *robustness* [6] of multilevel and multiprovider networks. We categorise networks in the following four groups:

1) **single level, single provider:** These networks consist of the router- or PoP-level of a single provider. Most studies analysed this type of graph [6], [7].

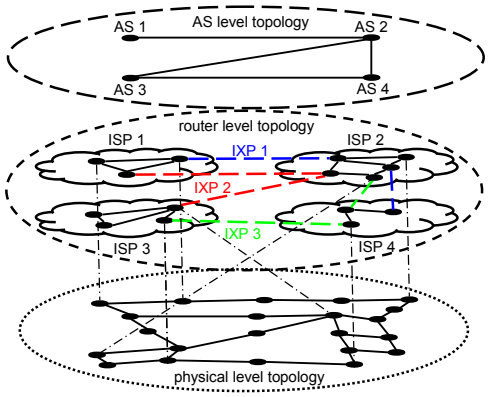

Fig. 1. An abstract view of multilevel and multiprovider Internet graph

2) **multilevel, single provider:** These networks consist of multilevel graphs within a single provider. There are a few studies examining multilevel graphs [8]–[10].
3) **single level, multiprovider:** These networks consist of AS-level graphs that include several provider networks, but as a single adjacency matrix. Several studies analyse AS-level graphs, but fail to capture inter-AS relationships via IXP links [11]. In this paper, we show multiprovider analysis that captures inter-AS relations via IXP links.
4) **multilevel, multiprovider:** This type of model and analysis is the most realistic. To the best of our knowledge, this type of model has not been studied. It will be part of our future work.

In this paper, we study real-world communication and transportation networks. We begin our multilevel analysis of flow robustness of 1-, 2-, and 3-level graphs, and show that the single level and multilevel graphs exhibit different performance. We then analyse the flow robustness of a number of two-level 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 indicate that it is difficult to partition the tier-1 ISP connectivity using attacks targetted at logical links.

The rest of the paper is organised as follows: The topological dataset we use in this study is presented in Section II. The properties of graphs we analyse are presented in Section III.

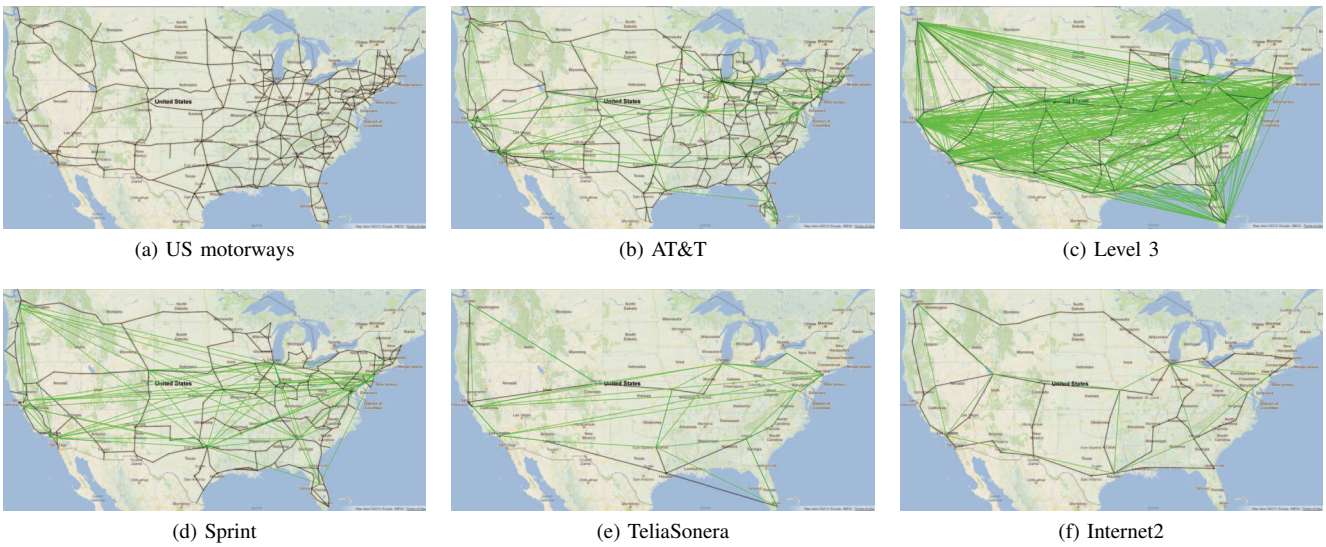

(a) US motorways (b) AT&T (c) Level 3

(d) Sprint (e) TeliaSonera (f) Internet2

Fig. 2. Visual representation of US motorways graph, and physical and logical level service provider networks in KU-TopView [5]

We present the robustness of multilevel and multiprovider graphs in Section IV. Finally, we summarise our findings as well as propose future work in Section V.

II. Topological Dataset

We study real networks (i.e. transportation and communication) that are geographically located within the continental United States. We have made these topologies publicly available [5]. We note that the focus of the paper is to present the applicability of our framework, and a comprehensive study of other networks remains part of our future work.

A. Transportation Network

We have generated the interstate highway topology to represent the transportation network based on the American Association of State Highway and Transportation Officials (AASHTO) data [12]. We presented the details of how we constructed the US highways topology in our earlier work [7]. The visual representation of this topology is shown in Figure 2a.

B. Communication Networks

We restrict this study to include PoP-level and physical fibre topologies. We use Rocketfuel-inferred AT&T, Level 3, and Sprint PoP-level topologies [13]. We also use the publicly available TeliaSonera network [14] and Internet2 [15] topologies. We note that international links, as well as discontiguous US links, are removed intentionally to compare the PoP-level topologies against the US fibre deployments and motorways topologies.

We then use US *long-haul fibre-optic routes* map data to generate geographic physical topologies for AT&T, Sprint, and Level 3 [16], [17]. In this map, US fibre-optic routes 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. The physical and logical commercial service provider networks are shown in Figures 2b, 2c, 2d, and 2e. The Internet2 research network at the physical and logical level is shown in Figure 2f.

III. Properties of Networks

We calculate the graph metrics of critical infrastructures as shown in Table I using the Python NetworkX library [18].

A. 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 [19]. Clustering coefficient measures how well a node's neighbours are connected [19]. Closeness centrality is the inverse of the sum of shortest paths from a node to every other node [6]. Betweenness centrality is the number of shortest paths through a node or link [20].

B. Graph Properties

We investigate the graph-theoretic properties of the US motorways graph, and the PoP-level (L3) and physical level (L1) topologies of four commercial ISP networks (AT&T, Level 3, Sprint, TeliaSonera) and the Internet2 research network. 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, as shown in Table I. The US motorways graph metrics are closer to those of the physical topologies. This is not surprising: since both the US highway system and the physical level of the Internet are physical infrastructures rather than logical overlays, they frequently share the same right-of-way [7].

From a distance metrics perspective, clearly physical topologies have higher values. This is expected since physical topologies are grid-like and have more nodes with low degree. We observe that the values of degree-based metrics also differ

TABLE I
TOPOLOGICAL CHARACTERISTICS OF TRANSPORTATION AND COMMUNICATION NETWORKS

Network	Nodes	Links	Avg. Node Degree	Clustering Coefficient	Diameter	Radius	Hopcount	Closeness	Max. Node Betweenness	Max. Link Betweenness
US motorways	411	553	2.69	0.05	42	21	13.65	0.07	23872	19785
AT&T L1	383	488	2.55	0.04	39	20	14.13	0.07	17011	14466
AT&T L3	107	140	2.62	0.09	6	3	3.38	0.30	2168	661
Level 3 L1	99	130	2.63	0.07	19	10	7.65	0.14	1628	1046
Level 3 L3	38	376	19.80	0.80	3	2	1.50	0.69	59	37
Sprint L1	264	312	2.36	0.03	37	19	14.75	0.07	11275	9570
Sprint L3	28	76	5.43	0.41	4	2	2.19	0.48	100	27
TeliaSonera L1	21	25	2.38	0.21	9	6	4.06	0.25	75	61
TeliaSonera L3	16	29	3.63	0.51	4	2	2.08	0.49	34	17
Internet2 L1	57	65	2.28	0.00	14	8	6.69	0.15	630	521
Internet2 L3	9	13	2.89	0.44	4	2	2.03	0.50	9	11

between physical and logical topologies. This can be attributed to the ease with which nodes can be connected in a logical topology (i.e. logical topologies are mesh-like) as compared to the difficulty involved in connecting node in a physical topology, where one must physically lay down fibre between nodes. From a centrality metrics perspective, we can see that physical topologies are not as clustered.

IV. MULTILEVEL ANALYSIS

In this section we present our framework as well as analyses of multilevel and multiprovider networks.

A. Multilevel Graph Model

In order to better 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 $G_{\ell_i} = (V_{\ell_i}, E_{\ell_i})$
3) For all non-negative integers i and j such that $i \leq j$, $V_{\ell_j} \subseteq V_{\ell_i}$
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 $\mathrm{conn}_{\ell_i}(u, v) =$ false, then $\mathrm{conn}_{\ell_j}(u, v) =$ false, where the function conn_{ℓ_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 3a, 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 3b. Note that in Figure 3c, 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.

Some authors have discussed the importance of multilevel graphs for studying the resilience and survivability of the Internet [8]–[10], [21]–[23]. Some have developed a multilevel graph framework [8], [9] and used it to analyse railway, peer-to-peer, brain, and random graph topologies [9]. Each topology was subjected to random and loaded [8] 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 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 remain connected after a number of deletions [6].

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, provided the following requirements are met:

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 using the Python NetworkX library [18]. 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.

B. Multilevel Graph Analysis

We first employ our multilevel framework to demonstrate the effect of using multiple levels of graphs on the *service*

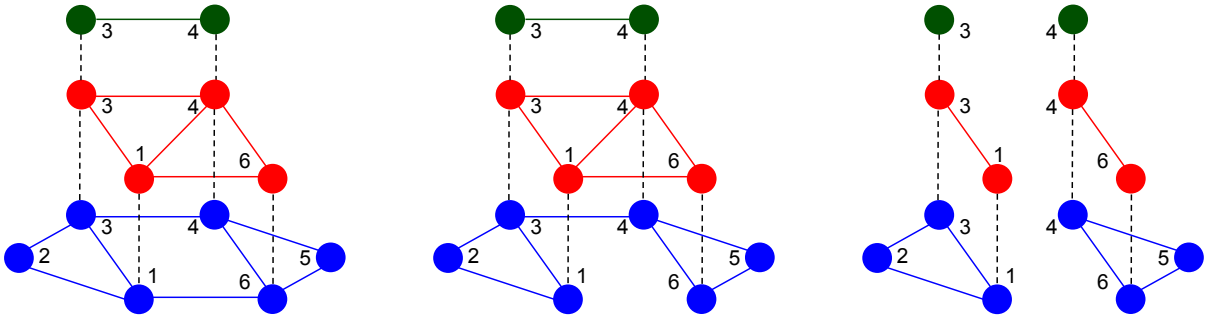

(a) Connected multilevel network (b) Disconnected multilevel network (c) Partitioned multilevel network

Fig. 3. Multilevel graph example

resilience at the top level. To demonstrate this, we use a 3-level graph (US freeways, physical, and logical-level topology of Internet2), a 2-level graph (physical- and logical-level topology of Internet2), and a single level graph (logical-level topology of Internet2). Note that the top level is identical in all cases. We acknowledge that the US motorways topology does not provide service to the physical level topology of Internet2 other than as right-of-way. We merely use it in order to construct a 3-level topology to better demonstrate how the resilience at Internet2's logical level behaves when we add more levels. For both the single and multilevel graphs, 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 sever the connection between two nodes within a given level the moment that the shortest path between them in the level below is disrupted.

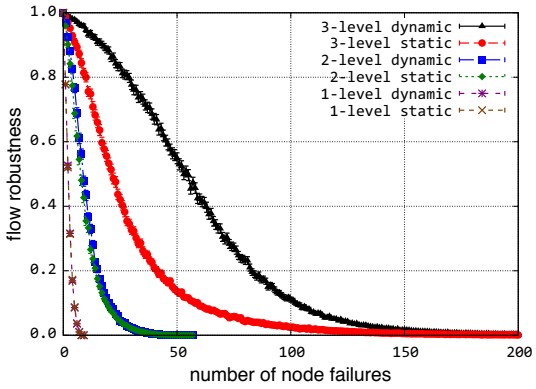

Fig. 4. Robustness of multilevel network for node deletions

The results of this experiment are shown in Figure 4 for node deletions and in Figure 5 for link deletions. For all 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

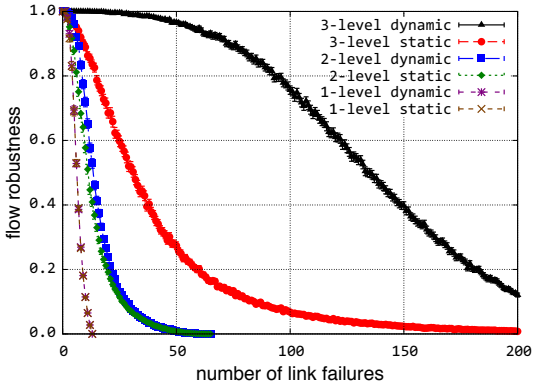

Fig. 5. Robustness of multilevel network for link deletions

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 and 1-level networks. For example, in Figure 4, when we delete 50 nodes in the lowest topology of the 3-level 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, 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 highway topology were less connected (e.g. instead of a grid-like, it was linear) then the flow robustness would be smaller. 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

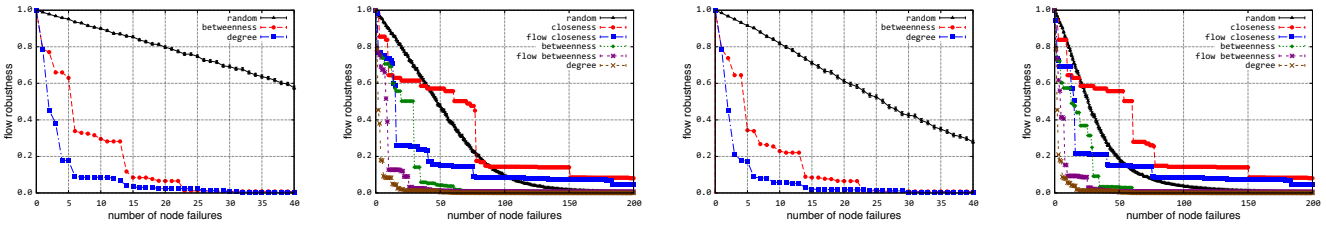

(a) Dynamic, adaptive (b) Dynamic, non-adaptive (c) Static, adaptive (d) Static, non-adaptive

Fig. 6. AT&T flow robustness for dynamic and static routing during adaptive and non-adaptive node deletions

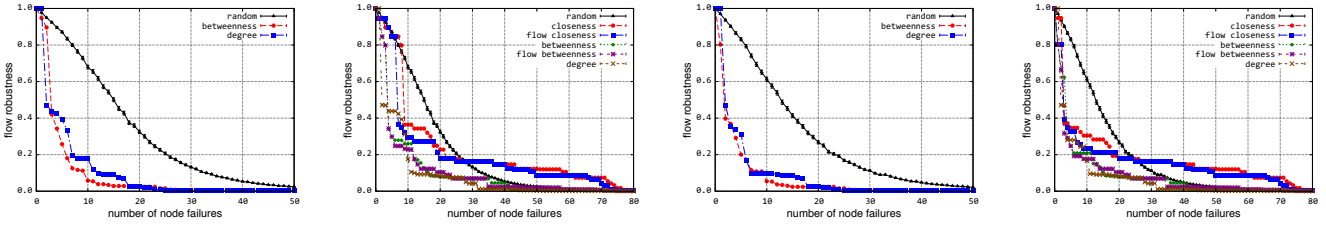

(a) Dynamic, adaptive (b) Dynamic, non-adaptive (c) Static, adaptive (d) Static, non-adaptive

Fig. 7. Level 3 flow robustness for dynamic and static routing during adaptive and non-adaptive node deletions

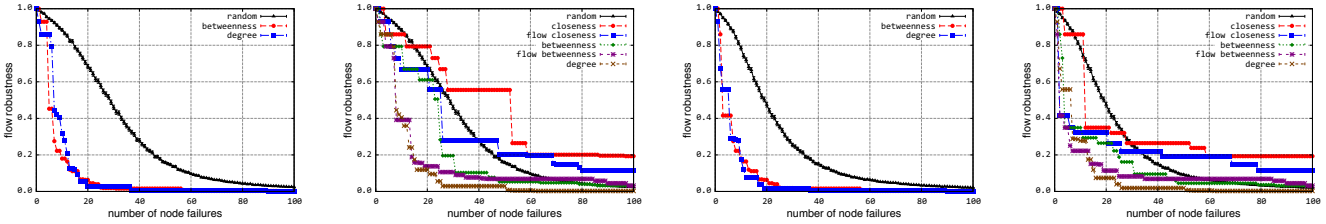

(a) Dynamic, adaptive (b) Dynamic, non-adaptive (c) Static, adaptive (d) Static, non-adaptive

Fig. 8. Sprint flow robustness for dynamic and static routing during adaptive and non-adaptive node deletions

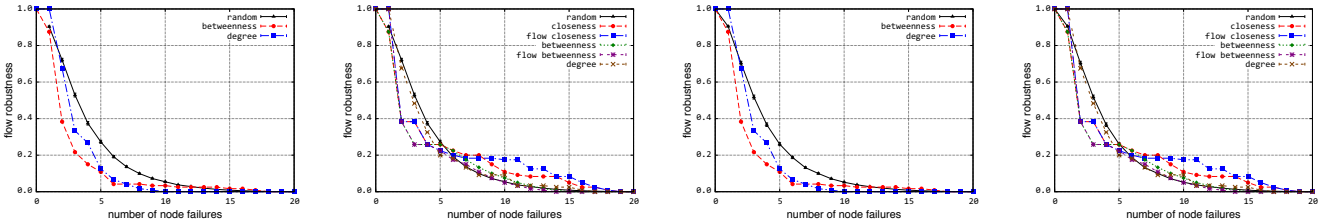

(a) Dynamic, adaptive (b) Dynamic, non-adaptive (c) Static, adaptive (d) Static, non-adaptive

Fig. 9. TeliaSonera flow robustness for dynamic and static routing during adaptive and non-adaptive node deletions

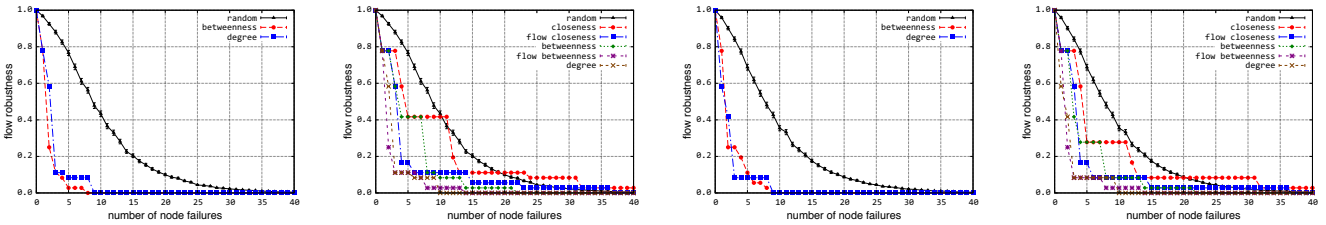

(a) Dynamic, adaptive (b) Dynamic, non-adaptive (c) Static, adaptive (d) Static, non-adaptive

Fig. 10. Internet2 flow robustness for dynamic and static routing during adaptive and non-adaptive node deletions

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 right-of-way motorways graph in the lowest level. However, considering that the motorways graph does not in actuality provide a service to the physical level of any given communication network, we focus on 2-level communication networks for the rest of our multilevel analysis. To that end, we use physical and logical level adjacency matrices for each of AT&T, Level 3, Sprint, TeliaSonera, and Internet2, use them 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 our experiments involving node deletions are shown in Figures 6 through 10, while the results of our experiments involving link deletions are shown in Figures 11 through 15.

As can be seen in Figures 6 through 15, in some cases we delete nodes and links at random while in others we delete nodes and links with 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: betweenness, closeness, degree, link betweenness, current-flow betweenness, and current-flow closeness. Next, we explain current-flow betweenness and current-flow closeness [24].

Current-flow betweenness and current-flow closeness are both ways of measuring a node's centrality based on information flow[1]. Consider 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*. 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 [24]. 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 [24][2].

The current-flow betweenness of a node is 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 and sink pairs. The current-flow closeness of a node 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 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.

[1]These two measures and closeness centrality are only applicable to simple and connected graphs. That is why they are employed only for *non-adaptive deletions*, as explained in the subsequent paragraph.

[2]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.

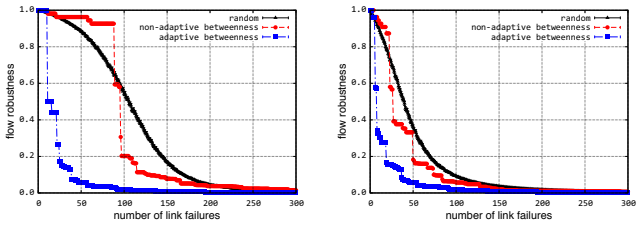

(a) Dynamic, adaptive & non-adaptive (b) Static, adaptive & non-adaptive

Fig. 11. AT&T flow robustness for link deletions

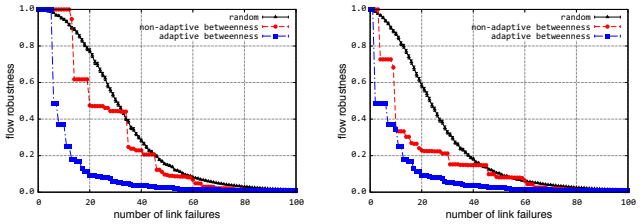

(a) Dynamic, adaptive & non-adaptive (b) Static, adaptive & non-adaptive

Fig. 12. Level 3 flow robustness for link deletions

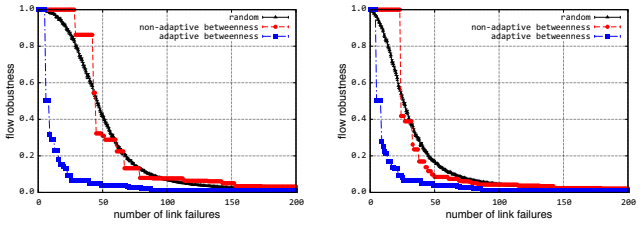

(a) Dynamic, adaptive & non-adaptive (b) Static, adaptive & non-adaptive

Fig. 13. Sprint flow robustness for link deletions

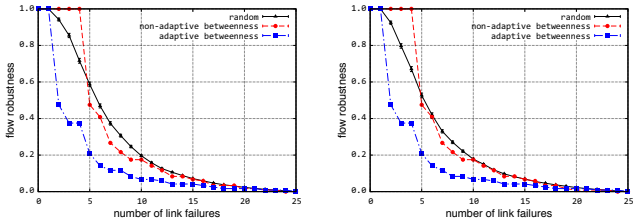

(a) Dynamic, adaptive & non-adaptive (b) Static, adaptive & non-adaptive

Fig. 14. TeliaSonera flow robustness for link deletions

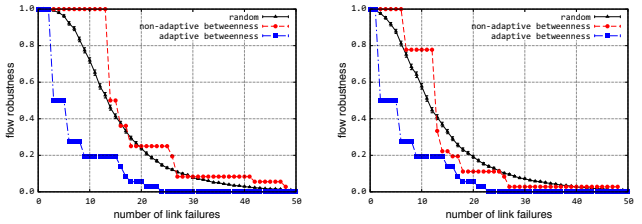

(a) Dynamic, adaptive & non-adaptive (b) Static, adaptive & non-adaptive

Fig. 15. Internet2 flow robustness for link deletions

We use all of these measures (betweenness, closeness, degree, link betweenness, current-flow betweenness, and current-flow closeness) to study what types 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 were recomputed after the most recent deletion. This could result from an attacker that had 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 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 8b, 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 8b, 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 the network arrives in a state in which several nodes are deleted, the 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 will be the subject of future work.

C. 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 $(G_{\mathrm{L3}}, G_{\mathrm{AS}})$, where $G_{\mathrm{L3}} = (V_{\mathrm{L3}}, E_{\mathrm{L3}})$ and $G_{\mathrm{AS}} = (V_{\mathrm{AS}}, E_{\mathrm{AS}})$ are PoP- and AS-level graphs that obey the following properties:

1) the vertices in V_{AS} are mutually disjoint connected subgraphs of G_{L3} that, when taken together, contain all of the vertices in V_{L3}. More specifically,
   if $V_{\mathrm{AS}} = \{v_1, v_2, ..., v_n\}$, then
   a) any two distinct vertices $v_i, v_j \in V_{\mathrm{AS}}$ 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 $1 \leq i \leq n$, then $\bigcup_{i=1}^{n} V_i = V_{\mathrm{L3}}$.
2) there exists some function $f : E_{\mathrm{AS}} \rightarrow 2^{E_{\mathrm{L3}}}$ such that for any pair of distinct vertices $v_i, v_j \in V_{\mathrm{AS}}$ given by $v_i = (V_i, E_i)$ and $v_j = (V_j, E_j)$, if $\{v_i, v_j\} \in E_{\mathrm{AS}}$, then $f(\{v_i, v_j\}) = V_{ij} \cap E_{\mathrm{L3}}$ 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 $v_i, v_j \in V_{\mathrm{AS}}$ that share a given AS-edge $\{v_i, v_j\} \in E_{\mathrm{AS}}$.

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 [25], Equinix [26], Terremark [27], and MAE-East [28] 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 have 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.

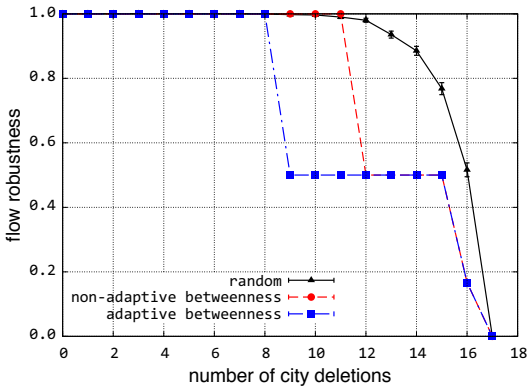

Fig. 16. Flow robustness of multiprovider network

In Figure 16, 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 non-adaptive attacks, which, in turn, inflict more harm than randomly-placed attacks. In Figure 16, 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.

V. Conclusions and Future Work

Realistically modelling the Internet requires a collective analysis of all of its structural properties. We evaluated multilevel graphs using the flow robustness metric and analysed combined communication and transport networks with our multilevel framework. We showed that dynamic routing helps alleviate the impact of perturbations and that adaptive challenges degrade multilevel network performance more than non-adaptive challenges.

Our future work will include investigating 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. Additionally, we are working on incorporating the *policy* aspect of a multiprovider network into our model by considering relationships between pairs of ASes.

Acknowledgments

We would like to acknowledge Mohammed J.F. Alenazi and other members of the ResiliNets group, as well as Padma Krishnaswamy for discussions on this work. 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 NFP0069666: Interdomain Resilience.

References

[1] D. Alderson, L. Li, W. Willinger, and J. C. Doyle, "Understanding Internet Topology: Principles, Models, and Validation," *IEEE/ACM Transactions on Networking*, vol. 13, no. 6, pp. 1205–1218, 2005.

[2] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *Proceedings of the ACM SIGCOMM*, (New Delhi, India), pp. 75–86, 2010.

[3] B. Donnet and T. Friedman, "Internet topology discovery: a survey," *IEEE Communications Surveys & Tutorials*, vol. 9, no. 4, pp. 56–69, 2007.

[4] C. Metz, "Interconnecting ISP Networks," *IEEE Internet Computing*, vol. 5, no. 2, pp. 74–80, 2001.

[5] "ResiliNets Topology Map Viewer." http://www.ittc.ku.edu/resilinets/maps/, January 2011.

[6] J. P. Rohrer, A. Jabbar, and J. P. Sterbenz, "Path Diversification for Future Internet End-to-End Resilience and Survivability," *Springer Telecommunication Systems*, 2012.

[7] E. K. Çetinkaya, M. J. Alenazi, J. P. Rohrer, and J. P. G. Sterbenz, "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)*, (St. Petersburg), October 2012.

[8] M. Kurant and P. Thiran, "Layered complex networks," *Phys. Rev. Lett.*, vol. 96, p. 138701, April 2006.

[9] M. Kurant, P. Thiran, and P. Hagmann, "Error and attack tolerance of layered complex networks," *Phys. Rev. E*, vol. 76, p. 026103, August 2007.

[10] D. Medhi and D. Tipper, "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, (Hilton Head Island, SC), pp. 173–186, January 2000.

[11] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *Proceedings of the ACM SIGCOMM*, (Cambridge, MA), pp. 251–262, 1999.

[12] 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.

[13] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP topologies with Rocketfuel," *IEEE/ACM Transactions on Networking*, vol. 12, no. 1, pp. 2–16, 2004.

[14] "TeliaSonera." http://www.teliasoneraic.com.

[15] "Internet2." http://www.internet2.edu.

[16] "Level 3 network map." http://maps.level3.com.

[17] KMI Corporation, "North American Fiberoptic Long-haul Routes Planned and in Place," 1999.

[18] A. A. Hagberg, D. A. Schult, and P. J. Swart, "Exploring Network Structure, Dynamics, and Function using NetworkX," in *7th Python in Science Conference (SciPy)*, (Pasadena, CA), pp. 11–15, August 2008.

[19] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network Topologies: Inference, Modeling, and Generation," *IEEE Communications Surveys & Tutorials*, vol. 10, no. 2, pp. 48–69, 2008.

[20] P. Mahadevan, D. Krioukov, M. Fomenkov, X. Dimitropoulos, k. claffy, and A. Vahdat, "The Internet AS-level topology: three data sources and one definitive metric," *ACM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 17–26, 2006.

[21] J. P. G. Sterbenz, D. Hutchison, E. K. Çetinkaya, A. Jabbar, J. P. Rohrer, M. Schöller, and P. Smith, "Redundancy, Diversity, and Connectivity to Achieve Multilevel Network Resilience, Survivability, and Disruption Tolerance (invited paper)," *Springer Telecommunication Systems*, 2012. (accepted April 2012).

[22] P. Demeester, M. Gryseels, A. Autenrieth, C. Brianza, L. Castagna, G. Signorelli, R. Clemenfe, M. Ravera, A. Jajszczyk, D. Janukowicz, K. V. Doorselaere, and Y. Harada, "Resilience in multilayer networks," *IEEE Communications Magazine*, vol. 37, pp. 70–76, August 1999.

[23] T. Lehman, X. Yang, N. Ghani, F. Gu, C. Guok, I. Monga, and B. Tierney, "Multilayer Networks: An Architecture Framework," *IEEE Communications Magazine*, vol. 49, no. 5, pp. 122–130, 2011.

[24] U. Brandes and D. Fleischer, "Centrality measures based on current flow," in *Proceedings of the 22nd Annual Conference on Theoretical Aspects of Computer Science (STACS)*, vol. 3404 of *LNCS*, pp. 533–544, Stuttgart: Springer Berlin / Heidelberg, February 2005.

[25] "AtlantaNAP." http://www.atlantanap.com.

[26] "Equinix." http://www.equinix.com.

[27] "Terremark." http://www.terremark.com.

[28] "MAE-East." http://en.wikipedia.org/wiki/MAE-East.