How Resilient are Individual ASes against AS-Level Link Failures?

Wenping Deng[1], Peidong Zhu[1], Naixue Xiong[2], Yang Xiao[3], Xiaofeng Hu[1]
[1]National University of Defense Technology, China
[2]Georgia State University, USA
[3]University of Alabama, USA
{wpdeng,pdzhu,xfhu}@nudt.edu.cn, nxiong@gsu.edu, yangxiao@cs.ua.edu

***Abstract*—Most ISPs rely on their upstream providers in achieving global reachability, thus they tend to make backup upstream links for the purpose of improving their own resilience. Yet, the Internet AS topology does not entail as much redundancy as it is widely thought, e.g., in our observation, 23.6% of all ASes only have single provider. Moreover, even though an AS has multiple upstream links, if all its upstream paths share a common link, the failure on the shared link may destroy its global reachability. This raises the question, how resilient are individual ASes against AS-level link failures? In this paper, we investigate the connectivity resilience for individual ASes from a global perspective. We identify the main factor that affects the resilience of individual ASes–the number of edge-disjoint uphill paths to Tier-1 they hold. To this end, we propose algorithms to quantify the resilience for individual ASes against AS-level link failures. We apply our methods on large-scale AS relationships data sets in order to study the resilience properties of the actual Internet. Our findings reveal that a considerable proportion (29.9%) of ASes are vulnerable to even one link failure in their upstream.**

***Index Terms*—BGP, inter-domain routing, autonomous system, connectivity, resilience**

I. Introduction

The Internet is composed of tens of thousands of Autonomous Systems (AS), where an AS is a domain operating one or more networks with common routing policies under the same administration entity, such as an Internet Service Provider (ISP) or a university. On top of physical connections, ASes use the Border Gateway Protocol (BGP) [1] to exchange routing information and forward traffic on global scale. Yet, this critical infrastructure is not as robust as it is widely thought. At AS level, link failures could be caused by contractual reasons (as the Tier-1 depeering between Cogent and Level 3 [2]), misconfigurations or physical damages [3], [4]. Such incidents have resulted in large-scale disruptions of services–many involved ISPs partially or completely lost their global reachability in a considerable period.

Indeed, Internet-wide reachability requires both considering *topological connectivity* and *policy compliance*. With respect to the former, the structure of the Internet has been of some interest for a variety of reasons; most commonly, because its topology plays a significant role in determining the performance of the Internet, although pure scientific interest has also played a substantial role in these investigations. To achieve global reachability, ASes exchange routing information via the Border Gateway Protocol (BGP) [1]. Much research, e.g., [5], [6], [7], [8], [9] has gone into studying Internet topology at different levels of granularity, including AS-level, PoP(point of presence)-level, and router-level.

Generally, most ISPs rely on their upstream ISPs in achieving their global reachability. For the purpose of improving their own resilience, most ISPs tend to make multiple upstream links to different providers. However, on the one hand, the Internet AS-level topology does not entail as much redundancy as it is widely thought. In our findings, 23.6% of all the ASes only have single provider, i.e., for such an AS, failure to its unique upstream link may destroy its global reachability. On the other hand, even though an AS has multiple upstream links, at the worst, if all its upstream paths share a common link, the failure on the shared link may disconnect this AS from its upstream thus destroy its global reachability. This raises the question: how resilient are individual ASes against link failures in their upstream?

Our goal is to quantify the connectivity resilience of individual ASes based on the Internet AS graph. In light of Menger's Theorem [10] and the inherent hierarchy [11], [12] of Internet AS-level topology, we identify the main factor that affects the resilience of individual ASes against AS-level link failures–the number of edge-disjoint uphill paths to Tier-1 they hold. Hence, for any given AS, its resilience is quantified by the number of its edge-disjoint uphill paths to Tier-1. We then illustrate that the problem can be solved by the Max-Flow techniques. We apply our method on large-scale AS relationships data sets from AquaLab [13] and CAIDA [14] in order to study the resilience properties of the actual Internet. Our findings reveal that a considerable proportion of ASes are vulnerable to even one link failure in their upstream–29.9% for AquaLab data set, yet 47.7% for CAIDA data set.

The remainder of this paper is structured as follows. Section II describes the background. Section III presents our algorithms for measuring the resilience from individual ASes, followed by the statistical analysis in Section IV. We conclude the paper in Section V.

II. Background

The Border Gateway Protocol (BGP) [1] is a *policy-based* routing protocol. The need for routing policies stems from the economic structure of the Internet.

AS relationships are the most wide-spread of all policy models that have been suggested in the past. AS relationships classify any directed link of an AS graph into one of the following types: customer-provider (c2p), peer-to-peer (p2p), or sibling (s2s). While in a c2p relationship the customer pays the provider to obtain transit through the provider's network, p2p assumes that two peering ASes share the deployment and maintenance cost for the connecting links. Siblings are peering ASes that have a mutual transit agreement, e.g., merging ISPs. Drawing further on AS relationships, the Internet AS topology can be denoted by a *Type-of-Relationship* (ToR) graph $G = (V, E, R)$ [15]: the nodes V are ASs, the edges E reflect AS-level peerings, and the edges are annotated with AS relationships $R = \{p2c, c2p, p2p, s2s\}$–each edge $e =< u, v >$ in E is assigned by a unique relationship of R, e.g., $R(e) = c2p$ implies that u is a customer of v.

Under the AS relationship policy model, a common assumption is that the valley-free property [16] holds: ASes want to avoid being used as a transit. For this reason, routes learned from provider and peer neighbors are not propagated to other provider or peer ASes. Formally, this can be expressed as follows: Let p be an AS path $p = (v_1, v_2, \ldots, v_k)$ from v_1 to v_k. If (v_i, v_{i+1}) $(1 \leq i < k)$ is a *p2c* edge or a *p2p* edge, then for any j $(i < j < k)$ the edge (v_j, v_{j+1}) must be of type *p2c* or *s2s*. Gao et al. [19] characterize a path as *downhill* (*uphill*) if it only contains p2c or s2s links (c2p or s2s links) and therefore any valid (valley-free) path *must* match one of the following *patterns*: (1) **Pattern 1**: an uphill path; (2) **Pattern 2**: a downhill path; (3) **Pattern 3**: an uphill segment followed by a downhill segment; (4) **Pattern 4**: an uphill segment followed by a p2p link; (5) **Pattern 5**: a p2p link followed by a downhill segment; (6) **Pattern 6**: an uphill segment followed by a p2p link, followed by a downhill segment.

Drawing further on AS relationships, Gao and Rexford [21] state that the inter-domain topology of the Internet is inherently hierarchical: a customer AS is at a lower level in the hierarchy than its provider ASes. At the top of the hierarchy there are around 10 core ASes as called Tier-1 ASes. These Tier-1 ASes have no provider and form a full mesh of AS-level peerings by p2p links [11]. Consequently, there should not be any directed cycle of p2c links, i.e., AS A cannot be provider of AS B, if B is provider of AS C and C is provider of A.

Ge *et al* [12] present a method to precisely construct the hierarchy of the ToR graph. Their method mainly consists of the following two steps: 1) An AS is classified as Tier-1 if it does not have a provider; 2) An AS node belongs to $level_{i+1}$ if it has a provider that belongs to $Level_i$. Each node is put at a level below all its provider ASes. Therefore the Internet ToR graph $G = (V, E, R)$ can eventually be represented by a hierarchical ToR graph $G_H = (V, E, R, H)$, where H is the hierarchy mapping from nodes to levels. Each node in V is assigned by a unique level number $H(v)$. Nodes in the m-th level can be represented by $V_m = \{v | v \in V, H(v) = m\}$. In particular, V_1 is the set of the Tier-1 ASes, non-Tier-1 ASes are the ASes at level 2 to M. In this hierarchical structure, nodes at Tier-1 are connected by p2p links. Nodes at the same level cannot be connected via p2c links, while p2p links may exist between nodes that are at the same level or nodes that are at different levels.

Based on the hierarchy, all ASes are generally classified into three categories: *core AS*, *transit AS*, and *stub AS*. While core ASes are the Tier-1 ASes (in the following sections, we will use the two notions interchangeably), stub ASes are the domains that have no customers and are located at the edge of the Internet. Finally, we refer to "transit ASes" as ASes that have both upstream providers and downstream customers, therefore, they are located in the middle of the hierarchy. Throughout this paper, we will sometimes use the term **non-Tier-1 AS** to denote the combined set of stub ASes and transit ASes.

Existing data sets [13], [14] have revealed that s2s relationships are rare in the AS topology. The literature generally simplifies the connectivity structure by neglecting s2s links [2], or treating s2s links as p2p links [18]. In this paper, we adopt the latter approach and treat s2s links as p2p links in the ToR graph. Consequently, a downhill path is simply defined as a sequence of links that exclusively consists of p2c links while uphill paths only include c2p links.

III. Measuring the Resilience for Individual ASes

In this section, we present details of our method for measuring the connectivity resilience from individual ASes. In Section III-A we analyze inherent properties in AS connectivity. Section III-B and Section III-C describe the main algorithm to quantify connectivity resilience of individual ASes.

A. Reachability Analysis

In this section, we first restate and disclose some properties that can be derived from the AS hierarchical structure and the valley-free property. Then we analyze the AS reachability under scenarios with different kinds of link failures.

(1) Tier-1 full mesh.

Constrained by the valley-free property, two Tier-1 ASes are reachable by each other if and only if they have a direct p2p relationship[1]. Any other scenario results in contradiction to the valley-free property: a) If they have a p2c (or c2p) relationship, then one should be a provider of the other, and this contradicts with that Tier-1 ASes are provider-free; 2) Otherwise, they should be connected by at least one middle node x. No matter whether x is a Tier-1 AS or not, $u - x - v$ inevitably violates the valley-free property. Hence, we tacitly assume that Tier-1 ASes form a full mesh by p2p relationships.

(2) Tier-1 uphill path.

We have demonstrated that an AS path connecting an AS pair s and t is considered as a valid path only if it follows any pattern (Pattern 1∼6) described in Section II. Examining the reachability of an AS pair turns to check whether there

[1] In the real AS-level topology, Tier-1 ASes are widely thought as a full mesh [11].

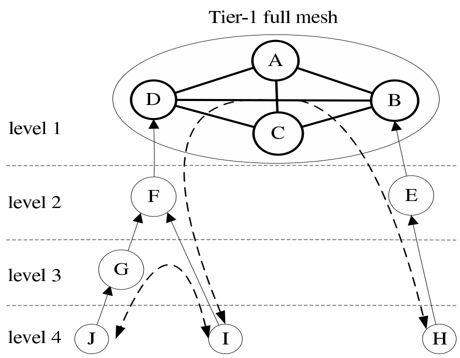

Fig. 1. Reachability between AS pairs that have Tier-1 uphill paths.

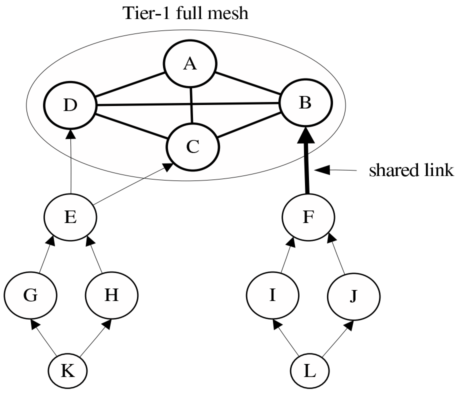

Fig. 2. Nodes with multiple upstream links (Although node L has two upstream links $< L, I >$ and $< L, J >$, while it may lose all its Tier-1 uphill paths under failure of the link $< F, B >$).

is a valley-free AS path connecting the two nodes. To this end, we introduce the term "Tier-1 uphill path" based on the notion "uphill path" (consisting of a sequence of c2p links), and demonstrate the important role it plays in AS reachability.

Definition 1 (Tier-1 uphill path)
Given a non-Tier-1 AS n_0 and a c2p link sequence $\{< n_0, n_1 >, < n_1, n_2 >, \cdots, < n_{i-1}, n_i >\}$, if n_i is the unique Tier-1 AS in $\{n_0, n_1, \cdots, n_{i-1}, n_i\}$, then $n_0 - n_1 - \cdots - n_{i-1} - n_i$ is called a ***Tier-1 uphill path****.*

Claim 1 *Given $G = (V, E, R)$ is the ToR graph of the Internet AS topology, where all Tier-1 nodes are fully connected by p2p links, any pair of nodes m and n in G ($m \neq n$) that have Tier-1 uphill paths are reachable by each other.*

Proof: According to the definition of Tier-1 uphill path, there is a Tier-1 uphill path p_m=$m - m_1 - \cdots - m_{i-1} - m_i$ for m and a Tier-1 uphill path p_n=$n - n_1 - \cdots - n_{j-1} - n_j$ for n respectively, where $H(m_i)$=$H(n_j)$=1. If p_m and p_n cover a common node, there is a valid path between m and n consisting of an uphill sub-path and a downhill sub-path following the Pattern 5 as described in Section II (as the reachability between nodes I and J in Figure 1); otherwise, there is a valid path between m and n formed by an uphill sub-path, followed by a p2p link ($< m_i, n_j >$) and a downhill sub-path following the Pattern 6 as described in Section II (as the reachability between nodes I and H in Figure 1). ■

(3) Reachability under failures.
For any two non-Tier-1 ASes, each of them holding a Tier-1 uphill path is not a necessary but sufficient condition for the reachability between the two ASes. As the case with nodes I and J in Figure 1, both of the two nodes have a Tier-1 uphill path, yet they can achieve their reachability without going through Tier-1. We now consider two typical scenarios of link failures in the graph.

•**Tier-1 depeering.** Tier-1 ASes and the full mesh formed by them construct the core of the Internet. Indeed, depeering over a Tier-1 p2p link can cause significant impact on the Internet as it disrupts the communication between their respective customers. It could be caused by contractual reasons, misconfigurations or physical damages, as pinpointed in [2] and as evidenced by contractual disputes between Cogent and Level 3 [2]. However, we argue that the Tier-1 full mesh can still be maintained when such a failure occurs. For instance, given the Tier-1 full mesh formed by the Tier-1 AS set $\{n_1, n_2, \cdots, n_{i-1}, n_i\}$, if the failure happens on the link between n_1 and n_2, the Tier-1 mesh will not be a full mesh any more. However, a new Tier-1 full mesh can be reconstructed by removing either n_1 or n_2 from the Tier-1 AS set.

•**Upstream link teardown.** c2p links connect the ASes in different tiers of the Internet, which contribute to the major connectivity. Now we study the case with a c2p (upstream) link tearsdown. When there is an upstream link teardown, for the ASes with single Tier-1 uphill paths, if their Tier-1 uphill paths must traverse this upstream link, they may inevitably lose their Tier-1 uphill paths. However, we argue that the Tier-1 full mesh is absolutely resilient to such failures, and most of the non-Tier-1 ASes can still maintain a Tier-1 uphill path.

In summary, the majority of the non-Tier-1 ASes can still maintain a Tier-1 uphill path for both of the two scenarios. Thus all AS pairs of the Tier-1 ASes and those who still have a Tier-1 uphill path are reachable by each other. We thus call the Tier-1 ASes and ASes that still have a Tier-1 uphill path the **main component**. While the ASes that lose all their Tier-1 uphill paths are only reachable by limited part of remained ASes, hence they are classified into the **disconnected component**.

B. Resilience Measurements

Now we have already demonstrated that Tier-1 uphill path plays a critical role in the global reachability of individual ASes. In light of this, ISPs expect to enhance their connectivity resilience by making backup upstream links. Unfortunately, as shown in Figure 2, even though an AS has multiple upstream links, it can still lose all its Tier-1 uphill paths in case of single link failures.

In classic graph theories, the connectivity resilience of a

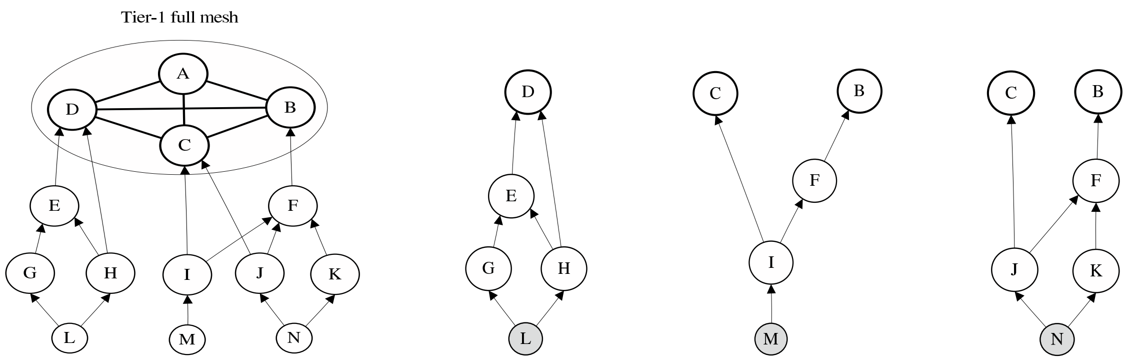

(a) The original ToR graph. (b) USG(L). (c) USG(M). (d) USG(N).

Fig. 3. Examples for extracting USGs.

node pair to risk of node (link) failures is measured by the minimum vertex (edge) cut between them. It is well known by Menger's theorem [10] that in directed or undirected graphs the maximum number of edge-disjoint s-t-paths is equal to the size of a minimum s-t-edge-cut. These measures can be computed efficiently using network flow techniques upon unit capacity networks–by assigning unit capacity to each edges [22]. To this end, we firstly present the definition for edge-disjoint Tier-1 uphill path as follows.

Definition 2 (Edge-disjoint Tier-1 uphill paths)
*For a given non-Tier-1 AS n_0, its two Tier-1 uphill paths p_1=n_0−n_1−$\cdots$−n_{i-1}−n_i and p_2=n_0−m_1−$\cdots$−m_{j-1}−m_j. Let $E_I = \{< n_0, n_1 >, < n_1, n_2 >, \cdots, < n_{i-1}, n_i >\}$ be the c2p edge set of p_1, $E_J = \{< n_0, m_1 >, < m_1, m_2 >, \cdots, < m_{j-1}, m_j >\}$ be the c2p edge set of p_2, if $E_I \bigcap E_J$=$\emptyset$, then p_1 and p_2 are two **edge-disjoint Tier-1 uphill paths** of n_0.*

In light of Menger's Theorem and the edge-disjoint Tier-1 uphill path, we derive the metrics for quantifying the connectivity resilience of individual ASes as follows.

Definition 3 (N_{edp})
Given a non-Tier-1 AS n, $N_{edp}(n)$ is defined as the number of $n's$ edge-disjoint Tier-1 uphill paths.

For any given non-Tier-1 AS, its resilience to AS-level link failures depends on the number of edge-disjoint Tier-1 uphill paths it holds. Hence an AS with k edge-disjoint Tier-1 uphill paths is resilient to arbitrary k-1 link failures.

C. Computing Disjoint Tier-1 Uphill Paths

We now focus on counting the number of edge-disjoint Tier-1 uphill paths for individual ASes based on given ToR graphs. Firstly we propose the definition of **uphill spanning graph (USG)** to capture the full upstream view of any given non-Tier-1 AS. Drawing further on this, we then demonstrate how to transform an uphill spanning graph to a unit capacity flow network, thus the disjoint paths problem can be equally changed into the maximum flow problem.

Algorithm 1 Constructing the flow network from USG

Input:
n_0, the given source node;
$USG(n_0) = (V_U, E_U)$, the directed uphill spanning graph of n_0;

Output:
$G_F(n_0) = (V_F, E_F)$ annotated by capacity values for all edges;

$V_F \leftarrow V_U$, $E_F \leftarrow E_U$;
for each e in E_F **do**
  $C(e) \leftarrow 1$; //assigning a unit capacity to each edge
end for
$V_F \leftarrow V_F \cup \{T\}$; //adding T as the common sink node
for each Tier-1 nodes v in V_F **do**
  $E_F \leftarrow E_F \cup \{< v, T >\}$; //adding the directed edge $< v, T >$ to E_F
  $C(< v, T >) \leftarrow \infty$; //assigning an infinity capacity to this extra edge
end for

Definition 4 (Uphill spanning graph, USG)
For a given non-Tier-1 AS n_0 in the ToR graph $G = (V, E, R)$, let $USG(n_0) = (V_U, E_U)$ be a directed graph, where V_U is the full set of nodes that are covered by at least one of $n_0's$ uphill paths, and E_U is the full set of c2p edges that are covered by at least one of $n_0's$ uphill paths (here a c2p edge $< u, v >$ is treated as a directed edge from u to v).

According to Definition 4, we know that $USG(n_0)$ consists of all nodes and c2p links that are covered by $n_0's$ uphill paths, hence all $n_0's$ Tier-1 uphill paths are inevitably involved in $USG(n_0)$. Figure 3 gives examples for extracting USGs from a given ToR graph.

Former literatures [22] have demonstrated that the number of edge-disjoint paths can be computed efficiently using maximum flow techniques upon unit capacity networks–by assigning unit capacity to each edges. In light of this, we firstly demonstrate how to transform a uphill spanning graph of a given AS to a unit capacity flow network. Then we leverage maximum flow techniques to compute the number of edge-

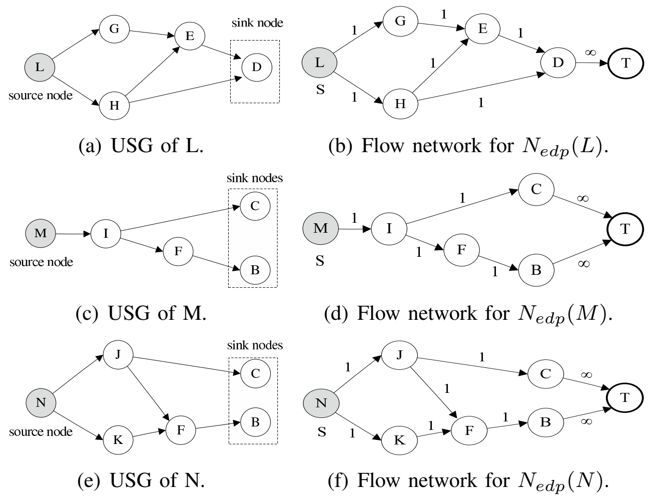

(a) USG of L. (b) Flow network for $N_{edp}(L)$.

(c) USG of M. (d) Flow network for $N_{edp}(M)$.

(e) USG of N. (f) Flow network for $N_{edp}(N)$.

Fig. 4. Examples for constructing flow networks.

disjoint Tier-1 uphill paths.

We start by constructing the unit capacity flow networks for edge-disjoint Tier-1 uphill paths computation. For a given node n_0 and its corresponding uphill spanning graph $USG(n_0)$, the flows start from the source node n_0 and are expected to be ended in Tier-1 AS nodes in $USG(n_0)$. Hence the flow networks are with single source node but multiple sink nodes (Tier-1 AS nodes in the uphill spanning graphs). We can transform the single-source multi-sink flow problem to maximum flow problem by adding an extra sink node T with infinity capacity edges connecting from all Tier-1 AS nodes in $USG(n_0)$. Algorithm 1 shows constructing the flow network from a given uphill spanning graph for N_{edp} computation. The problem has been transformed to a classic Max-Flow problem in a directed graph, hence it can be solved by Max-Flow algorithms [23]. Examples in Figure 4 are given to show construction of flow networks for N_{edp} computation.

IV. Evaluation

In this section, we start by explaining the data sets in Section IV-A. Then, in Section IV-B, we apply our methods to quantify connectivity resilience for individual ASes.

A. Data Sources

CAIDA [14] and AquaLab [13] both provide inferred AS relationships data sets for public use. The two data sets differ in the way how they were measured. While the CAIDA set relies on BGP snapshots from RouteViews [20], the AquaLab set leverages active measurements performed by an extension to a popular P2P system. For both data sets, AS relationships are inferred by using standard or improved techniques based on Gao's algorithms in [16].

Table I summarizes the comparison between the two data sets (Private ASes or relationships evolving such ASes are excluded from our statistics). It reveals that CAIDA observes more ASes than AquaLab, yet AquaLab provides a more comprehensive view of AS-level links. To mitigate the bias that might be caused by incompleteness of AS relationships, we perform all our analysis from both the data sets of AquaLab and CAIDA.

TABLE I
Comparison of the data sets of CAIDA and AquaLab.

Source	# of ASes	# of AS relationships	
		p2c	p2p(including s2s)
CAIDA (30/08/2009)	32,381	66,498	6,321
AquaLab (28/08/2009)	31,554	93,620	48,216

We point out again that any measured Internet map is inherently limited in terms of its visibility [24]. Nonetheless, the data set is sufficient for the purpose of performing a case study to demonstrate the feasibility and efficiency of our methods, and helpful for understanding to what extent current ASes are resilient to AS-level link failures.

For the purpose of extracting all Tier-1 ASes, relying on a "full mesh of p2p links between all provider-free ASes" as criterion turns out to be too strict. After all, our view of the Internet is limited and likely to miss actually existing AS-level edges. For this reason, we first select nodes with a degree higher than a threshold D as Tier-1 candidates, e.g., $D = 500$. Nodes having p2p relationships with a "large enough" proportion (P) of other candidates (e.g., P=80%, rather than 100%) are then classified as Tier-1 ASes. Combining the two sets of inferred Tier-1 ASes, we totally get 13 Tier-1 ASes, including AS 174 (Cogent), 209 (Quest), 701 (UUnet), 1239 (Sprintlink), 2914 (NTT-Com.), 3356 (Level3), 3549 (Gblx), 7018 (AT&T WorldNet), etc. We point out that our inferred list of Tier-1 ASes is highly consistent with CAIDA, Wikipedia, or other sources.

B. Quantifying Resilience of Individual ASes

With the the Tier-1 ASes, we seek to quantify the connectivity resilience against AS-level link failures by computing the numbers of edge-disjoint Tier-1 uphill paths of individual ASes. Meanwhile, we also show the correlations between the number of upstream links and the number of edge-disjoint Tier-1 uphill paths. Table II provides statistics on the number of upstream links and the number of edge-disjoint Tier-1 uphill paths for non-Tier-1 ASes. According to the statistics in Table II, we summarize our findings as follows.

TABLE II
Upstream links vs. edge-disjoint Tier-1 uphill paths (TUP).

m	% of non-Tier-1 ASes with $\geq m$ upstream links		% of non-Tier-1 ASes with $\geq m$ edge-disjoint TUPs	
	AquaLab	CAIDA	AquaLab	CAIDA
1	99.74%	99.64%	98.90%	97.40%
2	76.37%	60.05%	70.11%	52.29%
3	35.56%	18.54%	28.33%	12.96%
4	18.44%	9.01%	11.95%	4.78%
5	11.02%	5.39%	5.78%	2.04%

• From both of the two data sets, we observed that there are several dozens of non-Tier-1 ASes with no upstream link or tier-1 uphill path, e.g., even from AquaLab data set, we observed 82 ASes with no upstream link and 347 ASes with no Tier-1 uphill path. We argue that this is mainly caused by the incompleteness of the data sets.

• A big proportion of ASes tend to have multiple upstream links to different upstream providers. In AquaLab data set, 76.37% of all non-Tier-1 ASes have at least two upstream links. In CAIDA data set, there are 60.05% of all non-Tier-1 ASes having at least two upstream links.

• A considerable proportion of ASes only have single Tier-1 uphill path thus they are vulnerable to one link failure in their upstream. From AquaLab data set, there are totally 70.11% of all non-Tier-1 ASes having at least two edge-disjoint tier-1 uphill paths, i.e., 29.89% of all non-Tier-1 ASes only have single edge-disjoint tier-1 uphill path, hence they are vulnerable to one link failure in their upstream. From CAIDA, there are totally 52.29% of all non-Tier-1 ASes having at least two edge-disjoint tier-1 uphill paths, hence 47.71% of non-Tier-1 ASes have single edge-disjoint tier-1 uphill path.

• For any given AS, possessing multiple upstream links does not guarantee multiple edge-disjoint Tier-1 uphill paths. In our observation, (1) from AquaLab, about 6% of non-Tier-1 ASes have multiple upstream links, yet they only have single edge-disjoint Tier-1 uphill paths; (2) about 8% of non-Tier-1 ASes have multiple upstream links, yet they only have single edge-disjoint Tier-1 uphill paths. These numbers clearly point out that not all redundant UP links can improve the resilience for individual ASes from a global perspective.

V. Conclusions

This paper investigated the connectivity robustness of inter-domain reachability, focusing on individual ASes, but at a global scale. We demonstrate that the resilience of any given AS is not purely decided by the number of upstream links it has, but depends on the number of its edge-disjoint uphill paths to Tier-1. We then leverage Menger's theorem and the Max-Flow techniques to quantify connectivity resilience of individual ASes, by computing their edge-disjoint Tier-1 uphill paths. Our statistics reveals that only 70.1% of all non-Tier-1 ASes have more than one edge-disjoint Tier-1 uphill path, i.e., 29.9% of ASes are vulnerable to even one AS-level link failure in its upstream. Such a number clearly points out that the connectivity of individual ASes is not as resilient as it is widely thought. We hope that our study can help to increase the awareness of network operators about Internet-wide resilience for individual ASes.

However, we also see that the proposed methods and the analysis performed in this paper rely on the abstraction of an AS graph and ignore the underlying physical topology. Indeed, each AS can consist of multiple routers and each AS-level link may actually correspond to multiple physical peering links at different geographic sites [25]. Evidently, such redundancy improves the resilience of the Internet. However, at such a global scale, discovering the topology and understanding the resilience at physical level would be more challenging.

Acknowledgment

We would like to thank Merkouris Karaliopoulos and Wolfgang Mühlbauer for their constructive suggestions. This work is supported by the National Science and Technology Fund of China (NSF-60873214 and NSF-61070199).

References

[1] Y. Rekhter, T. Li, S. Hares, A Border Gateway Protocol 4 (BGP-4), IETF RFC4271, January 2006.
[2] J. Wu, Y. Zhang, Z. Mao, K. Shin, Internet Routing Resilience to Failures: Analysis and Implications, in: Proc. ACM CoNEXT, 2007.
[3] J. Cowie, A. Ogielski, B. Premore, E. Smith, T. Underwood, Impact of the 2003 Blackouts on Internet Communications, in: Technical report, Renesys Corporation, 2004.
[4] Y. Kitamura, Y. Lee, R. Sakiyama, K. Okamura, Experience with Restoration of Asia Pacific Network Failures from Taiwan Earthquake, IEICE Transactions 90-B (11):3095–3103, 2007.
[5] K. Chen, D. Choffnes, R. Potharaju, Y. Chen, F. Bustamante, D. Pei, Y. Zhao, Where the Sidewalk Ends, in: Proc. ACM CoNEXT, 2009.
[6] N. Spring, R. Mahajan, D. Wetherall, Measuring ISP Topologies with Rocketfuel, in: Proc. ACM SIGCOMM, 2002.
[7] R. Sherwood, A. Bender, N. Spring, DisCarte: A Disjunctive Internet Cartographer, in: Proc. ACM SIGCOMM, 2008.
[8] G. Siganos, M. Faloutsos, P. Faloutsos, C. Faloutsos, Powerlaws and the AS-level Internet Topology, IEEE/ACM Trans. Networking 11 (4):514–524, 2003.
[9] L. Li, D. Alderson, W. Willinger, J. Doyle, A first-principles approach to understanding the internet's router-level topology, in: Proc. ACM SIGCOMM, 2004.
[10] K. Menger, Zur allgemeinem Kurventheorie, Fund. Math. 10:96–115, 1927.
[11] J. R. L. Subramanian, S. Agarwal, R. Katz, Characterizing the Internet Hierarchy form Multiple Vantage Points, in: Proc. IEEE INFOCOM, 2002.
[12] Z. Ge, D. Figueiredo, S. Jaiwal, L. Gao, On the Hierarchical Structure of the Logical Internet Graph, in: Prof. of SPIE ITCOM, 2001.
[13] AquaLab, http://aqualab.cs.northwestern.edu/projects/SidewalkEnds.html.
[14] AS commercial relationship data, http://as-rank.caida.org/data/.
[15] G. Battista, T. Erlebach, A. Hall, M. Patrignani, M. Pizzonia, T. Schank, Computing the Types of the Relationships between Autonomous Systems, IEEE/ACM Trans. Networking 15 (2):267–280, 2007.
[16] F. Wang, L. Gao, Inferring and Characterizing Internet Routing Policies, in: Proc. ACM IMC, 2003.
[17] M. Caesar, J. Rexford, BGP Routing Policies in ISP Networks, IEEE Network Magazine 19 (6):5–11, 2005.
[18] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, S. Uhlig, Building an AS-Topology Model that Captures Route Diversity, in: Proc. ACM SIGCOMM, 2006.
[19] L. Gao, On Inferring Autonomous System Relationships in the Internet, IEEE/ACM Trans. Networking 9 (6):733–743, 2001.
[20] RouteViews, http://www.routeviews.org.
[21] L. Gao, J. Rexford, Stable Internet Routing Without Global Coordination, IEEE/ACM Trans. Networking 9 (6):681–692, 2001.
[22] J. Kleinberg, E. Tardos, Algorithm Design, Addison Wesley, 2006.
[23] R. Lestor, F. Ford,D. R. Fulkerson, Flows in Networks, Princeton University Press, 1962.
[24] R. Bush, O. Maennel, M. Roughan, S. Uhlig, Internet Optometry: Assessing the Broken Glasses in Internet Reachability, in: Proc. ACM IMC, 2009.
[25] P. Mérindol, V. V. Schrieck, B. Donnet, O. Bonaventure, J. Pansiot, Quantifying ASes Multiconnectivity Using Multicast Information, in: Proc. ACM IMC, 2009.