

MARIUSZ KAMOLA
PIOTR ARABAS

**NETWORK RESILIENCE ANALYSIS:
REVIEW OF CONCEPTS
AND A COUNTRY-LEVEL.
CASE STUDY**

Abstract

This paper presents the rationale behind performing an analysis of Internet resilience in the sense of maintaining a connection of autonomous systems in the presence of failures or attacks — on a level of a single country. Next, the graph of a network is constructed that represents interconnections between autonomous systems. The connectivity of the graph is examined for cases of link or node failure. Resilience metrics are proposed, focusing on a single autonomous system or on overall network reliability. The process of geographic location of networking infrastructure is presented, leading to an analysis of network resilience in the case of a joint failure of neighboring autonomous systems.

Keywords

network resilience, valley-free path, autonomous system geolocation

1. Introduction

The concept of disaster-related network vulnerability analysis results from the increasing impact that networking services have on everyday life — not only in the sense of human living conditions, but also in the sense of human life itself. Our dependence on networked equipment (not so obvious as in the case of, say, social networking or online banking) has crept into areas that are critical to human survival. Medical monitoring appliances, smart power grids, and air traffic control are just some applications which directly influence human safety; they all rely (or will rely) heavily on a network, however, which is not directly operated by humans.

We should first clarify the meaning of the very term ‘resilience’, providing the concepts and ideas encompassed by it. Here we focus on resilience defined as the vulnerability of a network of autonomous systems (AS) located in a single country to link or node failures. In this paper, we try to present practical solutions to the problem of identifying Internet weak spots; i.e., parts of the infrastructure whose failure could be really harmful. Solving problems of this type usually requires many computing resources, unless extra assumptions are made. A handful of relatively-simple technologies and publicly-available data does the work, letting one rank the importance of network links as well as simulate the results of a disaster affecting some area.

Intense research activity related to network resilience is going on, which can be observed by the number of scientific publications and collaborative research projects supported by governmental and international programs. The activities differ in their focus on key concepts linked to the very term ‘network resilience’:

- *evaluation* — how resilience should be understood; what metrics should be picked up to properly assess resilience; how measurement results should be interpreted in practice;
- *structure* — what specific following aspects of network operation should be addressed, and in how much detail, in order to evaluate resilience properly: ISO/OSI layers, systems hierarchy, physical location of devices, natural traffic, and topology variability in time;
- *scenario* — what kinds of network threats and responses to consider, and how to model them;
- *economy* — what can be the cost of the desired resilience level, and how it relates to the income model?

Let us address each of those concepts, respectively, by reviewing representative — although not all — research contributions.

1.1. Resilience evaluation

The intuitive definition of resilience is how well it withstands challenges placed on it, against infrastructure disruption, or traffic surge. More formal quantitative measures place the requirement to maintain network graph connectivity or cohesion in case of

a link or node failure, respectively [21]. The number of simultaneous failures resulting in service disruption is the resilience metric in either case. Alternatively, one may use the probabilistic approach, calculating service degradation upon the given probabilistic models of traffic load and link failure. Such a probabilistic metric of traffic loss due to low-level failures is proposed and evaluated in [21] for typical networking topologies. The analysis is usually computationally demanding, as it is applied in multidimensional space; therefore, improvements are proposed, as in [22], to effectively consider only the most probable failure scenarios. Here, the authors also consider link capacities, which result in extra metrics defined as the link overload.

To address the inherent multi-criterial nature of resilience, aggregate measures are proposed [29] in order to generalize the problem and provide a single performance grade to the network operator. The authors argue that a single measure (like average node connectivity) may be deceiving, so they propose using simple weighted means of measurement with individual tolerance intervals. The work has been carried within EC's FP7 project ResumeNet [1], whose other remarkable paper [28] provides a vast overview of previous works as well as sound classification of resilience-related issues. The term resilience gets divided into challenge tolerance and trustworthiness. The former means the ability to withstand: failures for internal reasons (termed "survivability"), failures for external factors ("disruption tolerance"), and unusual traffic patterns ("tolerance"). The latter provides users with dependability (so network services are reliable), security of the data being sent, and a contracted quality of service (QoS).

Remarkably, the report also addresses the question of classification and handling of any resilience measurements, suggesting to bin them into three classes (with appropriate action assigned to each of them). The authors propose an overall resilience management framework on a general level.

1.2. Network structure

Analysis of network resilience is rarely done on the level of abstracted graph concepts — even if this is the intention of Internet Protocol authors. To model correlated challenges (cf. also Sec. 1.3), one needs to reveal the underlying shared infrastructure. On the other hand, to model application-level resilience or even inter-AS resilience, the IP graph needs to be aggregated. There is economical rationale behind modeling a traffic interchange between autonomous systems; however, the specifics of that interaction adds extra complexity to the IP-level model.

The physical proximity of networking appliances is a special form of interaction; network topology (in the sense of its geographical layout) plays a crucial role in case of accidental damage as well as coordinated attacks affecting a specific area. Much effort has been put in order of getting a precise estimation of network geographical location. The seminal work [20] presents a number of specific technologies and datasets to establish a mapping an IP address, both to a region (e.g., a city) and to an AS. An important observation is made on the correlation of router and population

densities, which has been confirmed and elaborated further in [19]. Contrary to such a passive approach, [30] suggests the deployment of active probes and application of a precise-location algorithm based on triangle inequality, where distance is represented by packet delay. The algorithm operation is based on a set of nodes with a precisely-known geographical location (a.k.a. landmarks). This idea gets further developed in [13], where a set of constraints is imposed on the distance vectors, and a maximum likelihood approach is applied. The proposed algorithm is verified on the known Akamai node locations in the US. In another example of statistical approach to node geolocation [16], the authors use machine-learning techniques, considering delay and its statistical properties, hop count, and population density as relevant model inputs. Once having found the location of networking infrastructure, one is able to perform an analysis of its resistance to natural disasters or geographically-extensive attacks, as in [24] (where a bipartite graph approach is used to model network response to an electromagnetic pulse attack).

Contemporary research on geolocation has specialized, as in [25, 23], on the location of certain network resources as content delivery networks (CDN). The complementary task to infrastructure location is to locate mass users using that infrastructure as well — i.e., the so-called “eyeball networks.” Such focused research definitely links the definition of resilience with the cost of resilience, which is addressed in Sec. 1.4.

The geographical location of IP resources and AS-level network modeling are not tasks to be addressed disjointly; one can find a good survey of relevant techniques in [12]. Important issues in AS relationship are that they are asymmetric and non-transitive: the mere fact that AS Y is located between X and Z does not imply that traffic from X to Z goes through Y . One can examine BGP prefixes — but they are usually incomplete and outdated and do not reveal the economy of relationships (which may silently and efficiently prevent the use of sequence $X \rightarrow Y \rightarrow Z$). The proper categorization of network nodes — and whole AS’s, depending on modeling accuracy — as customers or providers is a necessary preprocessing stage in network analysis. A number of approaches exist: the classification of AS to multi-tier hierarchy levels based on BGP routing tables is presented in [15], while [18] uses social metrics of network nodes to achieve the same goal. Resilience metrics for an inter-AS graph with valley-free property is proposed and studied in [10]; practical suggestions are given there about the most effective investments in resilience improvement on the AS level.

Revealing the actual and precise structure of the Internet is as important as learning its statistical properties. In the latter case, the information is used to *generate* artificial networks with properties analogous to the real ones, and to appropriately analyze their resilience — as in [8], where a model of Internet growth on the AS level is constructed on the basis of real data. The proposed approach extends the Highly Optimized Tolerance (HOT) model (used earlier) in which the new node’s choice about which AS to join is based on AS centrality and proximity. The contribution allows us to model an AS as a set of points of presence (POP), which may be joined. The

main aim of yet another generator of inter-AS topology [6] is to reflect “local dense cores”; i.e., clusters of nodes correlated with population density — and to model AS hierarchy creation in such conditions. The approach needs extra arbitrary information about region boundaries to make the model accurate.

As it turns out in [27], metropolis or administrative units are not the only natural phenomena to “attract” the location of networking equipment. A suite of tools for physical location of network devices used in that research shows a clear correlation of network density with overall technology penetration in the region, represented by road or railway densities. Moreover, certain data paths are aligned with major communication routes. This is an important observation, as a collocation of goods, energy, and data transportation corridors adds up to network vulnerability to geographically-extensive attacks.

A good summary of technologies for geographical location, inter-AS relationship, and topology generation can be found in [17].

1.3. Network scenarios

The effects that various events have on the network are generally modeled in two non-exclusive ways: i) as a graph of resource or risk dependencies that indicate event impact on higher-level services; and ii) as a game played between the challenging and network response forces. The former can be modeled in the form of a Bayesian belief network [21]. A set of separate entities in higher layers that use common underlying device (like VPNs running on a physical trunk line) is termed a shared resource group (SRG). Upward propagation of risks is given adequate attention in previously-mentioned works [22, 11], with the latter emphasizing that small local incidents can develop into global malfunction, analogous to the butterfly effect in meteorology.

A good example of a complex, game-theoretic approach to model a challenge as a sequence of actions and counter-measures can be found in [7]. The authors develop an appropriate formalism for describing the interaction between the challenging and defending parties. The resulting Markov graph of network states can then be used in a quantitative assessment of systems resilience. Also, a similar but much-simplified approach is applied in [26] for the specific case of interdomain routing with path costs specified.

1.4. Economy of resilience

The issue of an economy model for providing resilience globally is rather omitted in the literature. Some authors [10] postulate that it should be addressed adequately. The most complete assessment of the current situation is contained in an ENISA report on resilience [9]. This vast document is a comprehensive overview of all of the aspects of resilience presented here thus far, aggregating important developments in the domains of network modeling, resource localization, network dynamics, network generation, and so on. As far as economical drivers for inter-AS resilience are concerned, virtually none exist (except for strong ISP competition for demanding customers). In general,

ISPs at all levels have no incentive to improve resilience individually: redundant links or multihoming raise operational costs that are not compensated by additional income.

However, in the long term, network reliability matters for end users. It is one of the reasons for migration — and its role will grow for those users sensitive to service disruption. In general, this means all of us as we depend on the web more and more. But it seems crucial for ISPs to select customers particularly keen on high-quality networking — and to find ways to make them pay for resilience-related investments and technology. As stated in [14], eyeball networks (i.e., broadband access networks) connect users ready to pay for the rich content they receive. In this sense, they are valuable equally to the ISP as well as the content provider. Thus, the point is revenue sharing — if not in general, then in what relates to increased resilience, because this guarantees undisturbed revenue for both business partners.

2. Network model

As much as everyone agrees upon using a graph for network modeling, everyone argues about what data to use for graph construction. There exist several commonly-used sources:

- Database of AS's managed by RIPE [2]. The data consists of information declared when registering an AS. They are updated relatively rarely. Of most value are BGP export and import records, which allows us to infer neighboring AS's and their relationships; however, only a small part of that information is usually implemented in BGP paths. The data also contains country information about each AS and — sometimes but rarely — the street address of the AS legal site.
- BGP paths gathered by available Polish *looking glasses* together with the paths collected by CAIDA [3] probes. This data, although live, is also usually superfluous, while the set of paths in everyday use is usually much smaller.
- Team Cymru database [4]. In this work, IP to AS number mapping was used to aggregate sites of interest to autonomous systems; also, a country-location database helps us to complete information retrieved from RIPE and CAIDA.

Let us observe that a discrepancy between current inter-AS routing paths and the possible relationships declared by AS's are of the utmost importance in traffic engineering or short-term failure prevention. For resilience analysis, all paths that can be *potentially* set up in case of a breakdown are almost as important as the ones currently used. That is why we stick with RIPE&CAIDA records in network graph construction.

2.1. Network graph limited to a single country

We also have decided to model the network consisting of AS's whose locations are limited to a single country. Obviously, such simplification deprives the graph of AS's located abroad (along with links to them), thus weakening the overall connectivity. However, there are sound reasons for doing that. Firstly, it allows the graph to remain

reasonably small, allowing the application of an exact graph search method without any heuristics. Even when applying methods that formally run in polynomial time (but with high ‘big O’ exponent) to a large graph, it renders analysis impossible. Secondly, the discarded foreign AS usually belongs to tier-1 or tier-2 operators, which, in general, maintain very good mutual connectivity and do not affect the overall network resilience much [10]. Thirdly, network resilience is usually under the care of an appropriate national authority; maintaining people connected in case of a disaster can, therefore, be achieved by legal, social, and economical measures — all of them being executed on the national stage.

Table 1

BGP statistics at border routers for three different AS’s.

BGP paths	AS X	AS Y	AS Z
a) total	186784	272905	987729
b) terminating in the country	4738	6835	12314
c) <i>not</i> contained in the country	826	271	108
d) ...but with alternative in the country	794	254	63
avg. number of foreign AS’s traversed in case c)	1.94	2.48	3.67

A quick check shows how much BGP paths connecting AS’s in a single country rely on foreign infrastructure. Let us consider here (and throughout the rest of the paper) the case of Poland. Such a choice was determined by the authors’ ability to access relatively-much traffic-related data that is up to date. Table 1 presents how many BGP paths that connect pairs of AS’s located in the same country traverse AS’s that are located abroad. Path counts in rows a) and b) roughly correspond to the size of X, Y, and Z (in the sense of exchanged traffic volume). Note that the percentage of paths that have destinations in the same country but ‘detour’ the national core in favor of providers abroad is not big; it ranges here from 17.4 to just 0.88 percent (row c). Interestingly, the majority of such paths have their pure national alternatives (row d), thus bringing a fraction of detouring paths to the edge of the error rate in observed original datasets.

2.2. AS’s reachable via uphill paths

As mentioned in Sec. 1, the edges in a graph of AS interconnections should be attributed to the proper type of inter-AS relationship, in order to correctly find feasible paths. There exists four types of relationships that an AS may have:

- P2C *provider to customer* — the AS acts as provider; it must forward all traffic originated by the client AS (and its customers as well);
- C2P *customer to provider* — the AS may forward all its own traffic, and the traffic of its customers, upwards to its provider (this relationship is complementary to P2C);
- P2P *peer to peer* — two AS’s may exchange their own traffic (and the traffic of their customers), but may not forward it using C2P links;

S2S *sibling to sibling* — two AS's act as a single unit, merely maintaining separate AS numbers.

Let us now consider a graph consisting only of P2C/C2P relationships (an extension to a more general case will follow). A valid path connecting AS1 to AS2 must contain a sequence of only C2P-type links, followed by a sequence of only P2C-links. No more link-type alternation may happen. Observe also that the AS1-AS2 relationship is symmetric; i.e., the reverse path traversing the same sequence of AS's is also valid. Therefore, let us construct a graph $G(V, E)$ consisting of vertices V corresponding to AS's and *directed* edges E , corresponding to all C2P links. For a given $v_i \in V$, the set of vertices v_i it is weakly connected with is defined as $R(v_i) = \{v_k : \text{path from } v_i \text{ to } v_k \text{ exists in } G(V, E)\}$. Let us call $R(v_i)$ a set of vertices *reachable from* v_i . For another vertex, v'_i , the reachable set is $R(v'_i)$; if the intersection $R(v_i) \cap R(v'_i)$ is not empty, then at least one valid path connecting v_i and v'_i exists. The intersection of reachable sets contains the AS's where the switch from C2P to P2C on the connecting path happens.

We can measure network resilience by a calculation of reachable sets in case of a single link or node failure. Consequently,

$$R(v_i, e_j) = \{v_k : \text{path from } v_i \text{ to } v_k \text{ exists in } G(V, E \setminus e_j)\}, \quad (1)$$

$$R(v_i, v_j) = \{v_k : \text{path from } v_i \text{ to } v_k \text{ exists in } G(V \setminus v_j, E_{v_j}^-)\} \quad (2)$$

denote reachable sets in situation of e_j or v_j failure, respectively. ($E_{v_j}^-$ is the original edge set with links to/from v_j removed from it.)

2.3. Vulnerability measures

Computation of $R(\cdot)$ for all (v_i, e_j) or (v_i, v_j) pairs is affordable for medium-sized systems, as it involves only several graph searches (which can be perfectly done in parallel). With such statistics at hand, we can go for calculating some ranks for individual links. Let us devise the following metrics:

$$u_E(v_i) = \frac{\sum_{j \in E} \text{card}(R(v_i) \setminus R(v_i, e_j))}{\text{card}(R(v_i))}, \quad (3)$$

$$u_V(v_i) = \frac{\sum_{j \in V, j \neq i} \text{card}(R(v_i) \setminus R(v_i, v_j))}{\text{card}(R(v_i))} \quad (4)$$

for assessing a relative decrease in the number of AS's reachable from v_i via C2P paths — summed over all possible link (3) or node (4) failures. For AS's poorly connected with the Internet core, many failure scenarios cause a severe decrease of available AS's; hence, u_E and u_V will tend to be high.

For a more-accurate check of overall network connectedness (w.r.t., a given link or vertex failure), one needs to check the resulting uphill paths *against* the downhill paths — for all possible pairs of AS's. If a set,

$$S(e_j) = \{(v_i, v'_i), i \neq j : R(v_i, e_j) \cap R(v'_i, e_j) = \emptyset\}, \quad (5)$$

contains all pairs of AS's disconnected by e_j failure, then we can use cardinality of $S(\cdot)$ for ranking edges, finding those whose failures deteriorate network connectivity most. This idea can be further developed by incorporating any extra information that could be available (e.g., traffic matrix, AS importance) into the rank. The complementary measure for node importance is $\text{card}(S(v_j))$, with:

$$S(v_j) = \{(v_i, v'_i), i \neq j : R(v_i, v_j) \cap R(v'_i, v_j) = \emptyset\}, \quad (6)$$

Sibling relationship can be handled by defining extra C2P links in both directions between AS's. However, a trick is needed to handle the peering case without a change in the apparatus presented above. We propose to represent every P2P link between AS1 and AS2 by *two* C2P links to an artificial AS acting as a provider. Such a proposition makes some BGP paths longer than in reality, but it does not affect the connectivity measure definition at all.

3. Assessing resilience in a real case

The specific graph that has undergone analysis was constructed from RIPE&CAIDA data [2, 3], from which Polish AS's were extracted. It consists of 1868 AS's linked by 5390 edges of P2C/C2P type. There are also extra 24 AS's connected to the rest of the graph by P2P links only; the total number of P2P links in the graph amounts to 835. The results of that specific graph analysis are presented below.

3.1. AS-oriented analysis

Metrics proposed in (3,4) provide a comprehensive overview of the connectivity of a given AS with the rest of the network. One can see in histograms in Figure 1 that the vast majority of AS's have $u_E(v_i)$ measure value close to 1. This means that a link failure on the uphill path deprives them on average of connectivity to only *one* AS — the one that was connected via the faulty link. If the upward connectivity to other AS's remains intact in such a case, we can hope that connectivity to only one of potentially many top-level providers is deficient, while there are alternatives to maintaining connectivity with the rest of the network. The prevalent value of 2 for $u_V(v_i)$ can be explained in a similar way, for the fact that a single AS failure leaves this AS unavailable by definition, plus the AS's that are located directly behind it. However, there exist nodes much more vulnerable to failure: four cases having $u_V(v_i) > 5$; five more having $u_V(v_i) > 4$, and as many as 130 others with parameters bigger than three — and analogous to u_E .

Quite expectedly, the most vulnerable AS's (w.r.t. both metrics) are exactly the same ones. The top- (or rather bottom-) ranked AS belongs to a small local ISP located in the countryside. For that specific AS, we have calculated the number of unreachable AS's along C2P paths — in the case of just a single link or AS failure. The failures causing biggest losses are listed in Table 2. The cases are sorted by disconnection severity; remarkable is the fact that as many as three individual link or




Figure 1. Histograms of AS's overall vulnerability to a single link or node failure — $u_E(v_i)$ on left , $u_V(v_i)$ on right. Wide bars refer to a complete network graph; black peaks refer to a graph with P2P relationships removed.

AS failures can cause a major disruption with that AS service. A further investigation of this case should involve BGP paths and topology analysis (which lie beyond the scope of this paper).

Table 2

Worst potential failure cases, ranked in terms of missing connections along C2P paths experienced by a specific AS. The specific AS is as v_i — cf (3,4).

$\text{card}(R(v_i)) - \text{card}(R(v_i, e_j))$ descending order	$\text{card}(R(v_i)) - \text{card}(R(v_i, v_j))$ descending order
136, 135, 134, 26, 20, 5,	136, 135, 134, 27, 26, 20, 12, 12,
4, 3, 3, 3, 3, 2, 2, 2, 2, 2,	9, 6, 5, 5, 4, 3, 3, 3, 3, 3, 3, 3, 3,
2, 2, 2, 2, 2, 2, 2, 2, 2, 1,...	2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 2, 1,...

Histograms in Figure 1 cover both cases of: a complete graph; and a graph from which P2P links have been removed. Please note little importance of P2P connections manifested in bar plots: forgetting them does not perceptibly change the statistics. This feature results from the highly-hierarchical nature of the contemporary Internet, where replicated transit agreements effectuating at the top of the hierarchy count more for network resilience than peering agreements elsewhere.

3.2. Network-oriented analysis

When it comes to examining end-to-end connectivity measures defined in terms of the number of AS pairs disconnected by a single link or single AS failure — cf. (5,6) — the relevant histograms are given in Figure 2. Surprisingly, there is not a single completely redundant link in the graph; a failure of any of them will result in some disconnection of AS's. Links tend to cluster into two groups; w.r.t., the effect of their failure on the network, which are not so diverse after all (observe $\text{card}S(e)$). Regarding AS failures,

their failures have a much heavier impact on the network than link failures, reaching $3.6 \cdot 10^4$ disconnected pairs of AS's, which is about 2% of all AS's pairs possible.




Figure 2. Histograms of number of all disconnected AS's pairs in effect of a single link or node failure — $\text{card}(S(e_i))$ on left , $\text{card}(S(v_i))$ on right.

We can conclude here that single AS's exist that are vulnerable to single-link disruption, thus failing to reach any of the top-level interconnection points (cf. cases presented in Figure 1). However, from the point of overall network connectivity (in the sense of connected AS pairs), the influence of a single AS failure is small, while the influence of a single link failure is next to nothing. Such observations are in force when infrastructure failures happen rarely and independently.

3.3. AS geographical location

In practice, links, interconnection points, and even whole AS's happen to be collocated. Tracing intricacies of infrastructure collocation is as valuable as it is difficult. Layered network structure and ISP policies to hide physical topology make the task of proper Internet resource geolocation practically impossible. As with many other authors, we present here the results of our own trials to find geographical location of infrastructure. We focus here on determining autonomous system location. Being aware that many AS's are, in fact, geographically extensive, we claim that considering small AS's as confined to a small region is justified and useful in overall resilience analysis.

Our approach assumes that infrastructure location in the case of small AS's is identical to their legal address. Such an address is sometimes provided for RIPE records, but it can be outdated if not absent altogether, particularly for small AS's. Fortunately, many national telecommunication authorities run their registries of telcos. The data provided there is always accurate and up-to-date. In the presented case, the national authority provides data for all 6456 issued telecommunication concessions [5]. A fuzzy match of 'clean' data was done to RIPE records, with some name regularization preprocessing, and the application of Levenshtein distance metrics of two to four characters. Successful matching of over 30% of AS's was done. Next, street

addresses were mapped to geographical coordinates using Google Maps online API. Eventually, 601 AS's have been attributed with their geographical location, as shown in Figure 3 (available on request from the authors). As observed by other authors, AS's density correlates well with population density (not shown), although with significant exceptions (note the sparsely populated NE corner of the map where numerous AS's emerge in response to absence of adequate infrastructure of the incumbent telco).

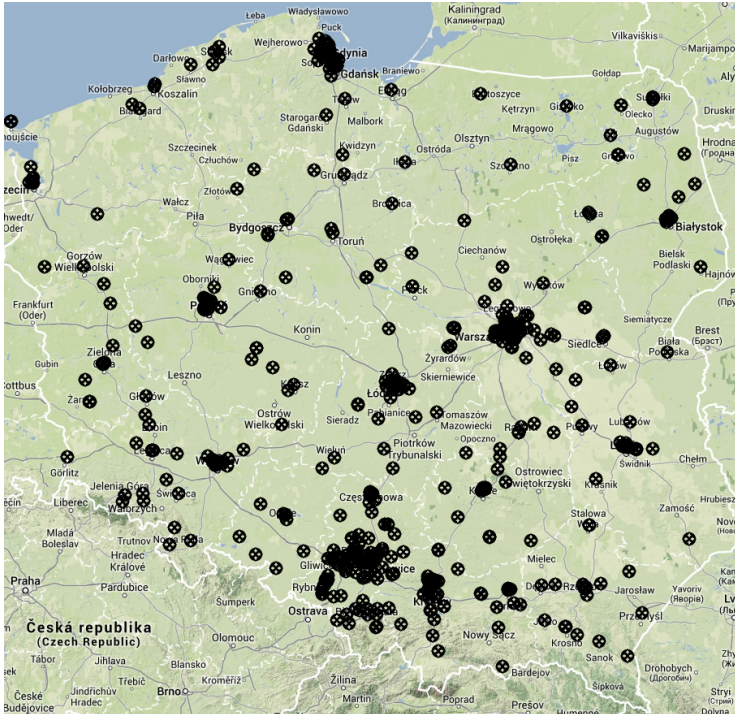


Figure 3. Google map with AS geographical locations. Map Data © Basarsoft, GeoBasis-DE/BKG (©2009), Google, basado en BCN IGN España.

3.4. Analysis of joint failures of neighboring AS's

With AS location found, one may perform an analysis of the effects of joint failures of AS's located nearby. Let us define a cluster of radius r as a set of AS's located geographically within a circle of that radius. Such a definition may be quite useful in terms of disaster effects: a weapon or a natural disaster with destruction range of r could destroy all AS's from the cluster.

For the simplest case, $r = 0$, we have 42 clusters on the graph. In this special case, they are in fact collocated AS's, in groups of two, three, or four. An analysis of the effects of their failure follows the definitions (2,4,6) but with more than one vertex removed from the graph. Failures of collocated AS's may cause from 13,000 to 55,000

AS pairs to stay disconnected; the latter number is quite significant as compared to single-AS failure repercussions: see Figure 2 (right).




Figure 4. Bivariate histogram of number AS pairs disconnected by faulty clusters of varying radii.

Allowing bigger $r = 100, 200, 400, 800$ meters and, finally, 1.6 km yields many more clusters whose failure consequences are obviously far more reaching. Naturally, many of these clusters overlap. The effects, in the sense of the number of AS pairs disconnected by cluster failures are presented in Figure 4. The failure cases have been grouped in logarithmic bins; w.r.t., r and the number of disconnected AS pairs. The most straightforward way to read the graph is to start with the assumed r ; then all possible damage scenarios with r are classified according to their effect on the network, with bar heights representing the number of corresponding failure cases.

We can observe that choosing r to be 0, 100 or 200 m does not really impact the overall damage to the network. This is because AS's are either collocated exactly in the same building or located far apart; in the other case, there is little chance to find any neighboring AS's within a radius of 100 to 200 meters. However, for larger r , we denote rapid growth of disconnected AS pairs — 220,000 in the worst case; i.e., 10% of all AS-to-AS relations. One must keep in mind that the location of only 30% of the AS's was done; with all equipment properly localized, the overall service deterioration would probably manifest at a much smaller r .

4. Conclusion

This study is focused deliberately on resilience analysis of a single-country-sized network. Statistics from several BGP routers show that relatively few paths connecting

AS's located in the same country would go through other countries. Thus, the resulting AS graph (constructed from RIPE&CAIDA datasets) is of manageable size and allows for the application of a standard graph-analysis method. We utilize the concept of a set of reachable AS's via customer-to-provider links (uplinks), determined for each autonomous system — as well for a complete network graph as for a graph with missing (or failed) links or AS's. Here, the connectivity between pairs of AS's is the main metric for the network as a whole. We find that the overall network connectivity in the considered case is rather sound; w.r.t., single element (AS or link) failures. However, there exist peripheral AS's without redundant links altogether. We also find the importance of P2P links to be marginal.

In order to assess the impact of geographically-extensive failures, we reached for well-established national registers to obtain street addresses of organizations that own AS's. A cursory expert overview of the geolocation results gives a good impression: for small AS's (but not only), the infrastructure is usually collocated with that organization's legal address. However, this part of the analysis deserves a more-objective evaluation. Performing a complete geographical location of resources is definitely a laborious but prospective task.

An analysis of joint failures of collocated AS's shows that it may have a much-greater impact on overall connectivity than individual ones. Expansion of a 'fire radius' that would destroy AS's has noticeable effects only when large areas are affected (800 meters or more). Alternatively, one can consider a vulnerability analysis for joint failures of any number of links (or AS's), without any assumption as to their geographical location. In such a case, the measure (1) becomes $R(v_i, \Psi) = \{v_k : \text{path from } v_i \text{ to } v_k \text{ exists in } G(V, E \setminus \Psi)\}$, where Ψ is a set containing faulty links. Measure (2) and the following metrics can be defined analogously. Specifically, in case of just two links failing simultaneously, in order to calculate e.g., (3), one has to iterate over all possible failure scenarios. While it is a fairly laborious task, we checked the results for the most-vulnerable AS in the test network, covered in Sec. 3.1 and in Table 2. To make a vulnerability comparison between single-link and double-link failures, (3) must get normalized by the total number of links or link pairs, respectively. For such a vulnerability measure, we get the value of $9.2 \cdot 10^{-4}$ for single-link and $4.7 \cdot 10^{-3}$ for double-link failures. Those values are not proper probabilities yet, but still provide a useful comparison. Quite expectedly, joint failures are on order of magnitude more dangerous for network connectivity. Especially if the "failures" are, in fact, orchestrated attacks.

References

- [1] URL: <http://www.resumenet.eu>.
- [2] URL: <http://www.ripe.net/data-tools/db>.
- [3] URL: <http://www.caida.org/data/>.
- [4] URL: <https://www.team-cymru.org/Services/ip-to-asn.html>.
- [5] URL: <http://www.uke.gov.pl/marta/index.php>.

- [6] Bar S., Gonen M., Wool A.: A geographic directed preferential internet topology model. *Computer Networks*, vol. 51(14), pp. 4174–4188, 2007. ISSN 1389-1286. <http://dx.doi.org/10.1016/j.comnet.2007.04.021>.
- [7] Bursztein E., Goubault-Larrecq J.: A logical framework for evaluating network resilience against faults and attacks. In: *Advances in Computer Science–ASIAN 2007. Computer and Network Security*, pp. 212–227. Springer, 2007.
- [8] Chang H., Jamin S., Willinger W.: *Internet connectivity at the AS-level: an optimization-driven modeling approach*. In: *Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research*, Mo-MeTools '03, pp. 33–46. ACM, New York, NY, USA, 2003. ISBN 1-58113-748-6. <http://dx.doi.org/10.1145/944773.944780>.
- [9] Chris Hall Richard Clayton R. A. E. O.: *Inter-alpha: Resilience of the Internet Interconnection Ecosystem*. Tech. rep., European Network and Information Security Agency, 2011.
- [10] Deng W., Karaliopoulos M., Mhlbauer W., Zhu P., Lu X., Plattner B.: k -Fault tolerance of the Internet {AS} graph. *Computer Networks*, vol. 55(10), pp. 2492–2503, 2011. ISSN 1389-1286. <http://dx.doi.org/10.1016/j.comnet.2011.04.009>.
- [11] Doerr C., Hernandez J.: A Computational Approach to Multi-level Analysis of Network Resilience. In: *Dependability (DEPEND), 2010 Third International Conference on*, pp. 125–132. 2010. <http://dx.doi.org/10.1109/DEPEND.2010.27>.
- [12] Donnet B., Friedman T.: Internet topology discovery: a survey. *Communications Surveys Tutorials, IEEE*, vol. 9(4), pp. 56–69, 2007. ISSN 1553-877X. <http://dx.doi.org/10.1109/COMST.2007.4444750>.
- [13] Eriksson B., Barford P., Maggs B., Nowak R.: Posit: a lightweight approach for IP geolocation. *SIGMETRICS Perform. Eval. Rev.*, vol. 40(2), pp. 2–11, 2012. ISSN 0163-5999. <http://dx.doi.org/10.1145/2381056.2381058>.
- [14] Faratin P., Clark D. D., Bauer S., Lehr W., Gilmore P. W., Berger A.: The growing complexity of Internet interconnection. *Communications & strategies*, (72), p. 51, 2008.
- [15] Gao L.: On inferring autonomous system relationships in the internet. *IEEE/ACM Trans. Netw.*, vol. 9(6), pp. 733–745, 2001. ISSN 1063-6692. <http://dx.doi.org/10.1109/90.974527>.
- [16] Gkantsidis C., Mihail M., Zegura E.: Spectral analysis of Internet topologies. In: *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications*. *IEEE Societies*, vol. 1, pp. 364–374 vol.1. 2003. ISSN 0743-166X. <http://dx.doi.org/10.1109/INFCOM.2003.1208688>.
- [17] Haddadi H., Rio M., Iannaccone G., Moore A., Mortier R.: Network topologies: inference, modeling, and generation. *Communications Surveys Tutorials, IEEE*, vol. 10(2), pp. 48–69, 2008. ISSN 1553-877X. <http://dx.doi.org/10.1109/COMST.2008.4564479>.

- [18] Jaiswal S., Rosenberg A., Towsley D.: Comparing the structure of power-law graphs and the Internet AS graph. In: *Network Protocols, 2004. ICNP 2004. Proceedings of the 12th IEEE International Conference on*, pp. 294–303, 2004. ISSN 1092-1648. <http://dx.doi.org/10.1109/ICNP.2004.1348119>.
- [19] Kim D., Son S. W., Jeong H.: Demographic studies of Internet routers. *Journal of the Korean Physical Society*, vol. 60(4), pp. 585–589, 2012. ISSN 0374-4884. <http://dx.doi.org/10.3938/jkps.60.585>.
- [20] Lakhina A., Byers J. W., Crovella M., Matta I.: On the geographic location of Internet resources. *IEEE J. Sel. A. Commun.*, vol. 21(6), pp. 934–948, 2006. ISSN 0733-8716. <http://dx.doi.org/10.1109/JSAC.2003.814667>.
- [21] Liu G., Ji C.: Scalability of Network-Failure Resilience: Analysis Using Multi-Layer Probabilistic Graphical Models. *Networking, IEEE/ACM Transactions on*, vol. 17(1), pp. 319–331, 2009. ISSN 1063-6692. <http://dx.doi.org/10.1109/TNET.2008.925944>.
- [22] Menth M., Duelli M., Martin R., Milbrandt J.: Resilience analysis of packet-switched communication networks. *IEEE/ACM Trans. Netw.*, vol. 17(6), pp. 1950–1963, 2009. ISSN 1063-6692. <http://dx.doi.org/10.1109/TNET.2009.2020981>.
- [23] Mátray P., Hága P., Laki S., Vattay G., Csabai I.: On the spatial properties of internet routes. *Computer Networks*, vol. 56(9), pp. 2237–2248, 2012. ISSN 1389-1286. <http://dx.doi.org/10.1016/j.comnet.2012.03.005>.
- [24] Neumayer S., Zussman G., Cohen R., Modiano E.: Assessing the Vulnerability of the Fiber Infrastructure to Disasters. *Networking, IEEE/ACM Transactions on*, vol. 19(6), pp. 1610–1623, 2011. ISSN 1063-6692. <http://dx.doi.org/10.1109/TNET.2011.2128879>.
- [25] Rasti A. H., Magharei N., Rejaie R., Willinger W.: Eyeball ASes: from geography to connectivity. In: *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement, IMC '10*, pp. 192–198. ACM, New York, NY, USA, 2010. ISBN 978-1-4503-0483-2. <http://dx.doi.org/10.1145/1879141.1879165>.
- [26] Secci S., Liu K., Rao G., Jabbari B.: Resilient Traffic Engineering in a Transit-Edge Separated Internet Routing. In: *Communications (ICC), 2011 IEEE International Conference on*, pp. 1–6, 2011. ISSN 1550-3607. <http://dx.doi.org/10.1109/icc.2011.5963439>.
- [27] Sterbenz J. P., Cetinkaya E. K., Hameed M. A., Jabbar A., Qian S., Rohrer J. P.: *Evaluation of network resilience, survivability, and disruption tolerance: analysis, topology generation, simulation, and experimentation. Telecommunication Systems*, pp. 1–32, 2011.
- [28] Sterbenz J. P., Hutchison D., Cetinkaya 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*, vol. 54(8), pp. 1245–1265, 2010. ISSN 1389-1286. <http://dx.doi.org/10.1016/j.comnet.2010.03.005>.

- [29] Van Mieghem P., Doerr C., Wang H., Hernandez J. M., Hutchison D., Karaliopoulos M., Kooij R.: *A framework for computing topological network robustness*. Delft University of Technology, Report20101218, 2010.
- [30] Ziviani A., Fdida S., de Rezende J. F., Duarte O. C. M.: Improving the accuracy of measurement-based geographic location of Internet hosts. *Computer Networks*, vol. 47(4), pp. 503–523, 2005. ISSN 1389-1286.
<http://dx.doi.org/10.1016/j.comnet.2004.08.013>.

Affiliations

Mariusz Kamola

1. Naukowa i Akademicka Sieć Komputerowa – Instytut Badawczy, ul. Wąwozowa 18, 02-796 Warszawa, Mariusz.Kamola@nask.pl,
2. Politechnika Warszawska, Instytut Automatyki i Informatyki Stosowanej, ul. Nowowiejska 15/19, 00-665 Warszawa

Piotr Arabas

1. Naukowa i Akademicka Sieć Komputerowa – Instytut Badawczy, ul. Wąwozowa 18, 02-796 Warszawa, Piotr.Arabas@nask.pl,
2. Politechnika Warszawska, Instytut Automatyki i Informatyki Stosowanej, ul. Nowowiejska 15/19, 00-665 Warszawa

Received: 14.02.2014

Revised: 22.05.2014

Accepted: 22.05.2014