

Dynamic Analysis of the Autonomous System Graph*

Marco Gaertler

University of Karlsruhe, Faculty of Informatics

E-mail: gaertler@ira.uka.de

Maurizio Patrignani

Università di Roma Tre

E-mail: patrigna@dia.uniroma3.it

Abstract

In this paper we investigate to what extent the information provided by BGP routing tables about the graph of the Autonomous Systems (ASes) can be used to understand dynamic phenomena occurring in the network. First, we classify the time scales at which such an analysis can be performed and, consequently, the kinds of phenomena that could be anticipated. Second, we improve cutting-edge technologies used to analyze the structure of the network, most notably spectral methods for graph clustering, in order to be able to analyze a whole sequence of consecutive snapshots that capture the temporal evolution of the network. Finally, we use such tools to analyze the data collected by the Oregon RouteViews project [20] during the last few years. We confirm stable properties of the AS graph, find major trends and notice that events occurring on a smaller time-frame, like worm-attacks, misconfigurations, outages, DDoS attacks, etc. seem to have a very diverse degree of impact on the AS graph structure, which suggests that these techniques could be used to distinguish some of them.

1 Introduction

Several researchers have started studying the Internet and its properties. In fact, such a rich and dynamic environment can be analyzed with respect to both the adjacency relationships between the entities composing it and the complex information flowing through its links and stored in its nodes. Often, this field of research produces intriguing results. For example, self-similarity laws, which are usually found when studying natural or biological phenom-

ena, were used to explain the apparently chaotic behavior of traffic loads or the uneven distribution of the number of adjacencies between the network nodes.

Recently, the general attention focused on the Autonomous Systems (ASes), the independent organizations whose cooperation guarantees the delivery of the packets through the network. Since each AS exchanges traffic flows with some neighbors (BGP peers), and local and remote adjacencies can be gathered from the logs of the BGP conversations between peers, plenty of current and historic data is potentially available about the ASes, their peerings, and the IP prefixes that are contained inside them. This kind of information is used to produce reliable statistics on the scalability of Internet technology and on the state of health of global routing [17].

Usually, the AS graph is reconstructed by merging information collected by a number of repositories managed by private and public research organizations around the world. Several studies have the purpose of analyzing the static properties of the AS graph, including the kind of relationships between its nodes, the distribution of their degrees, or other graph-theoretic measures [14, 3, 10, 12, 25, 22].

In this paper we investigate to what extent the information available from the BGP repositories can be used to understand the dynamic phenomena taking place in the Internet. Since, to our knowledge, this is the first attempts to systematically study dynamic properties of the AS graph, we begin by classifying the timescales at which such observations can be carried out and, consequently, the kinds of phenomena which may be anticipated, namely, stability, major trends, and spot events. Second, we consider the cutting-edge technologies described in the literature to analyze the structure of the AS graph, most notably spectral methods for graph clustering, and we advance and modify them in order to be able to analyze not only a single snapshot, but a sequence of consecutive AS graphs capturing the temporal evolution of the network. Finally, we use such tools to analyze the data collected by the Oregon RouteViews project [20] during the last few years, confirming known results and sensible hypotheses about the stable properties of the Internet, measuring how the network is

*Work partially supported by European Commission - Fet Open project COSIN - COevolution and Self-organisation In dynamical Networks - IST-2001-33555, by European Commission - Fet Open project DELIS - Dynamically Evolving Large Scale Information Systems - Contract no 001907, by "Progetto ALINWEB: Algoritmica per Internet e per il Web", MIUR Programmi di Ricerca Scientifica di Rilevante Interesse Nazionale, and by "The Multichannel Adaptive Information Systems (MAIS) Project", MIUR Fondo per gli Investimenti della Ricerca di Base.

evolving, and showing how this approach may turn out to be a valuable tool for monitoring the network against some types of malicious attacks. In fact, the various kinds of transient phenomena affecting the Internet seem to have a very different magnitude of impact on our measures, from none at all (worm attacks, misconfigurations, etc.) to a significant one (DDoS against DNS servers), offering a way to distinguish them.

The paper is organized as follows: Section 2 describes the motivations that prompted our research. Section 3 provides the necessary background and describes our methodological contribution to this domain of research. Section 4 contains the results of our short- and long-term analysis of the last few years of networks evolution and incidents. The conclusions complete the paper.

2 Motivations

Data retrieved from BGP repositories is generally analyzed to determine the structure and properties of the network of the Autonomous Systems at a specific time [13, 14, 3, 10, 12, 25]. However, a more sophisticated investigation can only be accomplished by considering a whole sequence of snapshots taken at subsequent moments. This is particularly true when the properties to be analyzed are related to temporal aspects, like stability, growth, or the dynamics of transitory phenomena. This kind of analysis can be performed with different time granularity, addressing different needs:

Long-term analysis Typically, the measures of the AS graph are not constant, but heavily affected by noise, rapid fluctuations, and sudden changes. A more stable view, obtained by considering a sufficiently long time window, would allow us to decide whether a particular snapshot actually fits into the ‘average’ network behavior, to understand if a peculiar measure captures a persistent property of the network or if it must be ascribed to fortuitous fluctuations, and to spot slow phenomena that take place over a period of months or even of years. Combining acquired knowledge about long-term temporal behavior allows to identify current trends and helps to ensure that the network continues to work properly.

Short-term analysis: The accessibility of detailed BGP logs offers the opportunity to perform short-term analysis, that is to investigate changes occurring in a short time period (and sometimes in a small neighborhood) with the purpose of classifying them as noise, incidents, outages, misconfigurations or malicious attacks. Of course, an effective classification would allow us to detect anomalies that occurred in the past, but the objective of this research may be much more

ambitious: devise techniques to spot anomalies as they occur, by means of a preemptive monitoring of the real-time and near-real-time behavior of the network.

3 Methodology

In this section we give an overview of the terminology we use, the data collection process, and our techniques.

3.1 Terminology

As usual, we build the graph of the Autonomous Systems starting from a set P of AS paths gathered from BGP logs. Each AS path is associated with a prefix (an interval of IP addresses) and consists of the sequence of the numbers of the Autonomous Systems traversed by the traffic to be delivered to the associated prefix. The first number in the sequence is called *vantage point* and is the source that recorded the AS path, while the last number is the *announcing AS* hosting the prefix IP numbers. In order to preserve information about vantage points and their relative importance we enrich the AS graph $G(V, E)$ with a mapping $\phi: V \times E \rightarrow [0, \infty[$, where:

- each AS is represented by a node,
- there exists an undirected edge between two nodes, if there exists a path in P in which the corresponding AS numbers are consecutive, and
- the number $\phi(v, e)$ equals the number of different paths in P that have v as vantage point and contain e .

This model allows us to keep track of the edge set seen by each vantage point. AS graphs are very difficult to compare, even if built from data sets taken at short intervals of time. In fact, the very vantage points collecting data are often temporarily unavailable, inducing significant changes in the observed portion of the network. Therefore, it is sometimes desirable being able to consider the stable, common part of two subsequent AS graphs. We define the *commonly observed AS graph* of two AS graphs G_1 and G_2 (with attached mappings ϕ_1 and ϕ_2) as the graph G_c where an edge $\{u, v\}$ is introduced if it has a value different from zero for both $\phi_1(v, \{u, v\})$ and $\phi_2(v, \{u, v\})$ for some node v . If needed, the value of $\phi_c(v, \{u, v\})$ is chosen to be the minimum of the two values above. Isolated nodes are removed from G_c .

3.2 Data Collection

Data on the actual routing can be collected by logging the BGP routing tables. Unfortunately, some well-known testbeds available on the web, e.g. [2], do not provide the

needed granularity with respect to time. In order to have the needed width and density, we used data collected by the *Oregon RouteViews Project* ([20]). This project was especially founded to grant real-time information about the global routing system from the perspectives of several different backbones and locations around the Internet. Its data is widely used by other network researchers and operators. It started to operate on April 20th, 2001 and provides snapshots every two hours. The provided data is not purified and contains misleading information, like repetitions of paths, loops inside paths, or prepending of ASes. Further, other issues can alter data, like truncated data files, time-outs, lost connections, or connections closed too early. We therefore filtered data according to the following simple rules:

1. discard any data set that is truncated or shows signs of technical problems, like the IP space is too limited,
2. enumerate all paths only once, and
3. delete loops and prepending ASes.

3.3 Observed Indices

A common approach is to define or use indices that provide a succinct measure of the properties of the graph which are to be monitored. We consider the number of nodes, edges, observing peers and paths. The first two are the classical graph theoretic measures and give a rough measure of the size and the density of the network. The other two are specific to this domain. They are needed to reflect the various influences on the extracted network view. So far, the total number of BGP-peering between ASes currently present in the Internet is unknown. It is only known that a larger set of vantage points often implies a larger set of paths. More paths often also result in more edges. Unfortunately, the peer and path sets are affected by various fluctuations themselves. Such changes can be roughly classified as: Problems due to the collecting software's faults, connectivity problems that affected the collection process, and actual changes in the routing graph. The first two may lead to false conclusions and therefore need to be removed. Most of these events can be eliminated by an online cleaning process. Other indices that capture structural properties are the maximum core level and the size of the induced core graph (both will be introduced in Section 3.4.1). Besides this static indices, we also considered the similarity of consecutive snapshots, by considering the commonly observed AS graph and its size.

3.4 Abstract Views

In order to cope with the amount of data, we had to compute a very concise description. We used different kinds

of simplification and clustering techniques, which offer the opportunity to get highlevel views of the network structure and behavior.

3.4.1 Cleaning the Graph

When dealing with large amounts of data, a usual technique is filtering out irrelevant information. Several authors attempted to investigate relevance concepts for the AS graph. The most common is the degree of a node, used for example by Govindan and Reddy [16], Gao [14] and Tauro et. al. [25]. We used the concept of coreness which is strictly related to the degree of nodes and was introduced in [23] and [5]. The k -core of a graph is defined as the unique subgraph obtained by recursively removing all nodes of a degree less than k . A node has *coreness* value ℓ , if it belongs to the ℓ -core but not to the $(\ell + 1)$ -core. Coreness can be used to filter out peripheral ASes. A suitable value of k can be chosen, for example, based on the number of the remaining ASes. There are no evidences that the scale-free property of the AS graphs might influence this coreness-based cleaning. The coreness concept can be used to generate a reduced model of the graph. We define the *core graph* $G_{\text{core}}(H) = (V', E')$ of a graph $H = (V, E)$ as follows. Let $E_{i,j}$ be the subset of all edges (in H) incident to nodes of coreness i and j . In $G_{\text{core}}(H)$ there is a node for each distinct coreness value, and two nodes i and j are connected, if $E_{i,j}$ is not empty. An edge-weighting function w of H can be transformed into an edge weight w' of G_{core} by setting $w'(i, j)$ to the total weight of all edges in $E_{i,j}$.

The AS graph has a very special structure, i.e. it is the result of merging paths. This composition reflects fundamental properties and thus a good filtering technique should respect the inherent path properties. The quality of a filtering technique could be measured by the number of AS paths that remain connected in the reduced graph. For the k -core filtering we define $\mu_k(i)$ to be the fraction of paths that have i connected components in the k -core. This number depends on the input parameter k . In order to avoid such a dependency and increase the insight into its compatibility, we will consider lower bounds for $\mu_k(1)$. Such a bound is given by the fraction of paths that remain connected in all k -cores and is denoted by $\mu(1)$. On the other hand, we also like to judge the degree of fragmentation. Therefore we count the fraction of paths that would be cut into i components in a k -core for some value k and $i > 1$. This value, which is denoted by $\mu(i)$, gives an upper bound on $\mu_k(i)$. An experimental validation on the coreness is presented in Section 4.1.1.

3.4.2 Clustering

Clustering is a well-known technique to explore the inner structure of relationships and the roles of the participating

entities. It was already used for this purpose in research areas such as social networks [6, 11], data mining [18], and VLSI [4]. Gkantsidis et. al. [15] performed spectral analysis and clustering on the AS graph. As part of their preprocessing they delete all nodes of degree one or two, which is equivalent to consider the 3-core. Other operations involved normalization or third-party information. Their main analysis is based on spectral information and involves eigenvectors. In this special scenario eigenvectors are mappings from the node set to the real numbers and proximity in these values often corresponds to dense subgraphs. They used a common heuristic to partition, which is based on selecting connected subgraphs that share the same sign in a specific eigenvector. By suitably choosing a subset of the eigenvectors, they were able to produce clusters characterized by strong geographic or business relationships. Good eigenvectors for this task are among those that correspond to the largest eigenvalues.

We applied a clustering method, called Geometric MST Clustering (GMC), that embodies the same intuition of the proximity provided by the eigenvectors and that was introduced in [8]. Although it involves more complex algorithmic steps, it can be fully automated, it explores the eigenvector structure to a larger extent, and requires only the graph (and optionally an edge weighting). The general idea is to interpret several eigenvectors, that are associated with the largest eigenvalues of the normalized adjacency matrix, as embedding for the graph and search within this geometric object for dense parts. The search itself is performed by a Minimum-Spanning-Tree (MST) computation. By considering only the span of tree edges that have (geometric) distance smaller than a given threshold, the MST represents a hierarchy of clusterings. Evaluating this hierarchy with different clustering indices allows to choose a good clustering as well as to incorporate several user-defined aspects.

The AS graph is quite inhomogeneous with respect to the density. The major part of the nodes are only loosely connected ([24]). In order to overcome this property and concentrate on relevant ASes, we choose the minimum k -core that has at most 200 elements. This choice is based on a rough estimate on the number of important ASes and experiments. As objective function for the clustering we used $\text{quality}(\mathcal{C}) = \sqrt[4]{\text{coverage}(\mathcal{C}) \cdot \text{performance}(\mathcal{C})^3}$, where coverage is the ratio of numbers of edges within all clusters and the total number of edges; and performance is the number of all adjacent node pairs lying in the same cluster plus the number of all non-adjacent node pairs lying in a different cluster divided by the total number of node pairs. The coverage score is an indicator for the global density of a clustering. While performance rates the clustering with respect to the disjoint-clique model. We chose this objective function to focus more on structural issues than on pure density. Again experiments verified the usability

of this objective function. Some properties of the resulting clusterings are presented in Section 4.1.2.

The integration of dynamic aspects, like the clustering of a sequence of graphs, entails lots of problems. First of all, the result format is not clear. Is a single partition of the node set still required or for each point of time? What is the importance ratio between edge set and time horizon? Second, how can static algorithms be adapted? And finally what should be the relation between static and dynamic clusterings? In the few instances such problems already arose, the following two methods were used to reduce the dynamic case to the static one, thus always calculating a single partition of the node set. The first one collapsed all graphs into one. Edges are present if their frequency is large enough. This corresponds to consider the 'average' graph over time and cluster it. The second approach clusters each graph individually and combines the results afterwards to a single partition. In this case, the 'average' clustering is output. Both versions have their disadvantages; the major problem is to find suitable thresholds for the combination step. We will use the second approach because it allows us to observe node movement more easily. However, statistical disturbances are the major source of false conclusions for both methods, especially when the given time frame is very short.

4 Experiments and Results


Basic properties and the quality of our auxiliary reduction techniques are presented first. This is necessary to increase the understanding of the impact of these utilities and to avoid wrong conclusions. It is followed by the major elements of dynamic analysis – baselines, trends and anomalies.

4.1 Evaluation of Reduction Techniques


The reduction techniques consist of cleaning and filtering on the one side and clustering on the other. The major difference is that the first keeps atomic elements – nodes still represent single ASes – while clustering produces many non-singleton groups. In this first part we present some considerations that justify the use of the reduction techniques in the later sections.

4.1.1 Validation of Cleaning and Filtering

It can be observed that $\mu(1)$ always exceeded 0.9, thus more than 90% of all AS paths remained connected in every k -core. Furthermore, only $\mu(2)$ and $\mu(3)$, which are the upper bounds on the fraction of paths that are split into two and three components, have values that are significantly larger than zero. Figure 1 shows the temporal evolution.



(a) Lower bound $\mu(1)$



(b) Temporal evolution of $\mu(i)$ for $i = 1, 2, 3$

Figure 1. Fraction of paths that are guaranteed to be connected and are split into two or three components over time

These observations are good indicators that the core structure of the graph is compatible with the path structure in general and therefore may be used for filtering purposes without altering the graph structure too much.

4.1.2 Significance of Clustering

We selected several random snapshots and clustered them individually. The observed features were averaged, used as hypotheses for common properties, and verified against other randomly selected samples. Thus, we found the following characteristics:


- two to six clusters have more than four elements; these clusters cover together more than 70% of all elements

- the remaining clusters often contained only one element
- clusters reflect geographical issues, i.e. the major clusters separated US, Europe and Asia

The Cooperative Association for Internet Data Analysis (CAIDA) provides limited geographic positions for several ASes. These may include coordinates or affiliation to certain states. The relationship between clusters and this data is shown in Figure 2 and Table 1. Thus, clusters reflect



(a) United States of America



(b) Europe

Figure 2. Geographical positions of some ASes. Different shapes represent different clusters, i.e. crosses for cluster 1 (33/86 positions available), squares for cluster 2 (10/17 positions available), and sharps for cluster 3 (4/6 positions available).

actual coherences of ASes as well as geographical issues. Therefore it might be useful to detect and classify changes in the network structure.


Country	Clusters			
	1.	2.	3.	remaining
US	75	3	1	22
Europe	3	10	4	24
Africa	2	1	1	2
Asia	5	1	-	2
unknown	1	2	-	2

Table 1. Distribution of countries on clusters


4.2 Baselines and Trends

Determining the standard undisturbed behavior is one fundamental task of dynamic analysis. As an initial step we investigate the evolution of static indices. The temporal evolution is shown in Figure 3(a) and 3(b). The number of nodes and edges seem to be locally stable over time. In order to verify this observation, we calculated the standard deviations of time frames of different length. Both indices seem to be non-constant, therefore a larger time window results in a larger deviation. In Figure 4 the time frame length is plotted versus the average standard deviation (ASD). As expected the ASDs are isotonic in the frame length, furthermore their increases can be expressed by linear function, thus verifying the local stability. The slope for the number of edges is three times larger than for the number of nodes. This can be explained by the fact, that edges are subject to more consistent changes and depend on the number of paths which in turn depend on the number of participating peers. For example, number of edges and paths as well as number of paths and peers are highly correlated (≈ 0.977 and ≈ 0.897 , respectively). In this special scenario, we can utilize the path structure to normalize the views by using the commonly observed AS graph of two consecutive AS graphs. The size of the node set and the edge set will be denoted by $|V'|$ and $|E'|$. This enables us to decide whether a change in the number of nodes (or edges) is only due to a change in the vantage points or not. Figure 5(a) displays the evolution of the number of nodes and edges for the AS graphs and the commonly observed AS graphs. Utilizing the detailed view of Figure 5(b) we conclude that the first two increases (2am and 10am) of peers led to a larger view (new edges were added) while the third increase (2pm) did not effect it (only redundant information was added). However, we have to be careful judging decreases, since the absence of vantage points can have many reasons. Some of these causes are independent of the (global) network status, like internal reorganization or connection failures during the collection process.

Another approach to judge the undisturbed behavior of the network involves a more structural and high-level view: the core graph (Section 3.4.1). We extend the concept by



(a) Number of nodes and edges



(b) Number of AS paths and peers

Figure 3. Graph theoretic and domain specific measures. The x-axis represents time, starting in May 2001 till September 2003.

adding a node-weighting function that represents the fraction of nodes (in the original graph) having a certain coreness. We observed that most nodes have a very low coreness score (one, two, or three). When edges are uniformly weighted, there are two different types of edges with large weight: edges that connect cores with a small index and edges that connect low-index cores with high-index cores (*transition edges*). In the case where edges are weighted according to the number of paths that use them, transition edges and edges connecting cores with large index are present. Figure 6 shows such a core graph with the two different edge weightings. This reflects the general intuition, that many links are needed to connect the customer to their provider and that the links that occur more frequently in paths are links to or within the 'backbone' and not in-




Figure 4. Local stability of number of nodes and edges. The time frame length is plotted versus the average standard deviations.

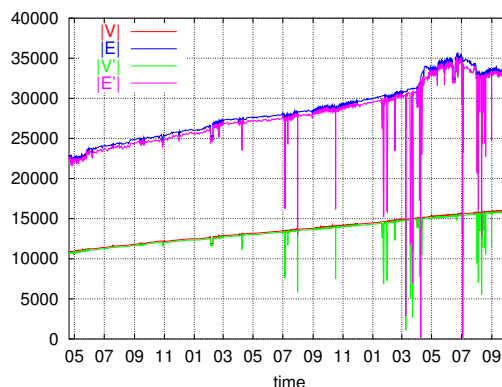
side the periphery. This description is stable with respect to time. In Figure 7 the relative distribution of coreness values are drawn; from the diagram it is evident that fractions of nodes with the lowest and the highest coreness are stable.

Long-term Patterns


Using these observed phenomena, we recognized two major patterns. The first one is the linear growth of the number of nodes and edges. Although it seems to be a trivial observation (Figure 3(a)), it is not at all expected. Many researchers have reported an exponential growth in the time frame of 1997 till 2000. Because both indices grow in a linear fashion, the ratio of observed edges to the total number of possible edges decreases. The second pattern is the distribution of growth. ASes are separated according to their importance in the core graph. By inspecting the evolution of the core graph, we observed that most nodes enter the graph with a very low coreness score, while edges (in the path-weighted version) connect low-score nodes with high-score nodes or only high-score nodes with each other. This verifies the general intuition, that more customers than (high-level) providers enter the system and that more connections are established to connect the small providers with the backbone or enlarge the backbone.

4.3 Short-Term Analysis – Anomalies

In this section we describe the results of short-term analysis. Namely, we measure the impact on our measurements of dramatic short-lived events that occurred in the network in the last few years. This section also offers the opportunity to test the effectiveness of the techniques described in Section 3.4.2 to overcome the difficulties implied by the




(a) Evolution over time of the different indices




(b) Detailed view (2001 May from 24th till 27th).

Figure 5. Sizes of AS graphs and commonly observed AS graphs.

use of the clustering approach in a dynamic scenario on real-world data. As already mentioned, events that last very short are problematic. In these cases a higher granularity of the data would be needed for better observations, but is often absent. However, our measurements suggest that events occurring on a small time-window may be effectively classified with our clustered analysis. In particular, we show in this section that some kind of sporadic phenomena, like worm attacks or misconfigurations, even if they may have a dramatic effect from the user's point of view, do not seem to have a measurable impact on the AS graph structure. On the other hand, some very specific events, notably DDoS attacks against DNS root servers, happen to cause a significant change of the structure of the clusters. Thus, this kind of analysis may be a promising tool, to be used as a litmus-paper to test if one of these specific events is occurring in



(a) Uniform edge weight



(b) Edge weight according to used paths

Figure 6. Core graph of AS graph (May 1st, 2001 0:00). The area of the nodes is proportional to the number of nodes having that coreness score. The thickness of the edges is proportional to their weight. (Note, that edges with very little weight are omitted.)

the network.

BGP storm due to worm attack. On July 19th 2001, the Code Red II worm was spreading on the Internet ex-





Figure 7. Distribution of coreness values (relative to the number of nodes in the graph)

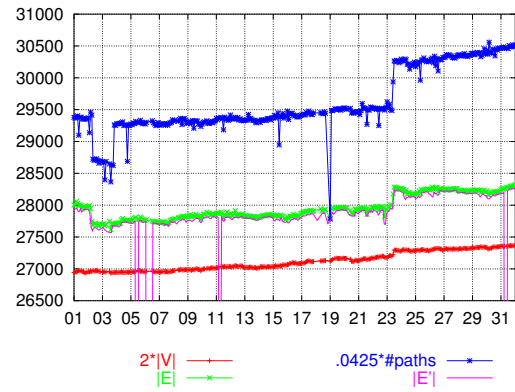
ploiting the indexing service vulnerability in the Microsoft Internet Information Server MS-IIS ([9]). Although an exponentially growth in the advertisement rate was observed by [7], we could not assess a significant change in the number of nodes and edges. In fact, Figure 8 shows that the number of nodes and edges is quite stable while the drop in the number of peers accounts for the loss of paths. These missing peers - AS715 and AS19092 - usually contribute with more than 15,000 paths each. Using the clustering technique, we could not spot any significant changes. In fact, about 75% of all nodes were either perfectly stable or switched between two clusters, the other nodes blinked in and out of our observed view or could not be associated to one cluster. The number and sizes of the clusters were stable. A similar behavior could be observed when on September 18th 2001, a extremely virulent worm, called W32.NIMDA, spread throughout the Internet using multiple methods to infect both Windows servers and user machines ([27]).

Shutdown of KPNQwest/Ebone and RIPE NCC Test Traffic measures a 50% increase of alarms.

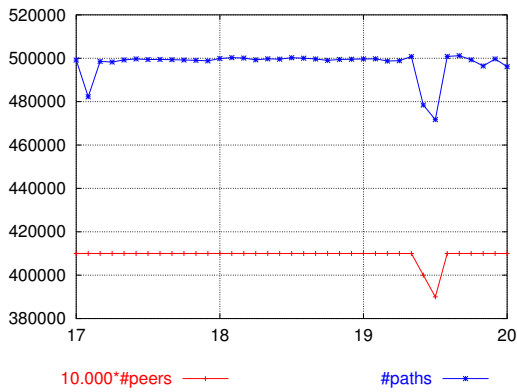
In July 2002, KPNQwest, one of Europe’s largest Internet backbone provider had a time-out. The company went offline on the 3rd and returned on the 25th ([19]). At the beginning of this event, the RIPE NCC measures a 50% increase of alarms on their TTMs (Test Traffic Measurement Boxes, [26]). We observed a drop in the number of peers, paths and edges. While the first two lasted only a few days, the last one was present during the whole period. Figure 9 shows this evolution. The clustering is rather stable when we considered the three time periods (before, during and after the shutdown) separately. However, we could observe temporal migrations (small subsets formed new clusters or were swallowed up), new ASes ‘entered’ the system and




(a) Number of nodes and edges



(a) Number of nodes, edges and paths



(b) Number of peers and paths



(b) Sizes of clusters (sorted according to size)

Figure 8. Evolution around July 19th 2001


Figure 9. Evolution during July 2002

old ASes 'left' during the transitions. The few drops in the number of edges of the commonly observed AS graph (see Figure 9(a)) are probably due to technical issues and do not reflect changes.

Misconfiguration in UUNet. UUNet (WorldCom) stopped talking to the rest of the Internet on October 3rd 2002 between 12am and 9pm (UTC), due to misconfiguration in some of their routers. This caused increases of respond times, delays, and packet losses ([21]). It affected many other ASes, because UUNet carried half of the world's Internet traffic. Although the number of edges seems to be affected, we could not observe a significant change (see Figure 10).

DDoS Attacks against 13 Internet Root Servers. On October 21st 2002, a series of well-coordinated, simulta-

neous DDoS attacks were launched from various points around the world, against each of the 13 Root Servers that are used for the Internet's Domain Name System (DNS) ([21]). The attack, which disabled nine of the 13 Root Servers, started at 8:45pm (UTC) and lasted approximately two hours. We could only partially observe the event, because *Oregon Routeview* had a 6-hours blackout, due to connection time-outs, starting at October 22nd 2am (UTC) and the peer AS701 disappeared, causing a loss of $\approx 17,600$ paths earlier on the 21st (8am UTC). It never returned as a vantage point. Thus, we can only analyze the structure of the network before and during the attack, but not immediately afterwards. Also, the loss of paths limits the viewpoint. Similar to the previous worm attacks, the number of nodes and edges is stable (see Figure 11). The clusterings are very stable, except for the point in time of 21st 8pm. In that instant a kind of splitting occurred. A subset of 29 elements and a single node emerged from the



(a) Number of nodes and edges

Figure 10. Evolution around October 3rd 2002

largest cluster and built their own clusters. For simplicity we denote the remaining part of the largest cluster by $1'$ and the other two with $1''$ and $1'''$ respectively, according to their size. Table 2 contains the geographic distribution of the ASes located in those clusters. Most of the elements in $1''$ and $1'''$ were perfectly stable while belonging to cluster 1 during the other time frames. The four Asian ASes – AS2518, AS4143, AS4728, and AS4766 – are kind of exchange servers. The cluster $1''$ contains also several well-know ASes like AT&T, Cable and Wireless, former Genuity, Globalcrossing, Sprintlink, UUnet, Verio, Microsoft (three times) and Google. A more fascinating fact is that in each cluster $1'$ and $1''$ an unaffected or a less affected DNS root server (AS297 and AS3557) is present. This is a strong indicator that the clustering was affected. However, the degree of interaction remains open. The current granularity is not high enough for a further and more detailed analysis.

countries	clusters			
	1	1'	1''	1'''
US	74	48	26	-
Europe	9	9	-	-
Asia	11	7	3	1
Africa	1	1	-	-
Australia	1	1	-	-
unknown	1	1	-	-

Table 2. The distribution of countries of the cluster and its segmentation.

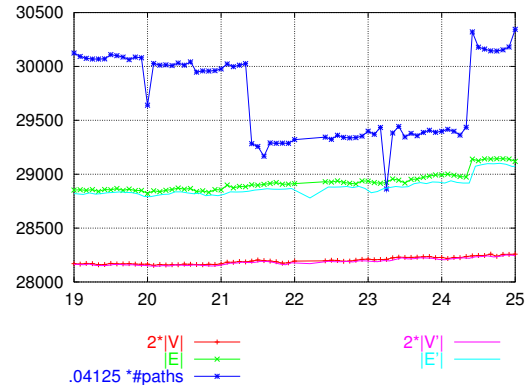
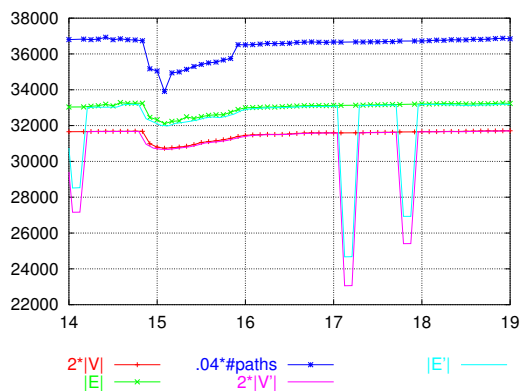


Figure 11. Evolution around October 21st 2002


Sapphire worm attack. In the morning of 25th January 2003, the Sapphire worm (also known as SQL slammer) was released on the Internet. Exploiting a vulnerability in Microsoft SQL server it multiplied itself rapidly and soon spread out over networks worldwide. From news headlines and activity on mailing lists it was clear the attack had an impact on the Internet's performance. Similar to the *DDoS Attack* our data source *Oregon Routeview* had a blackout from 8am till 12am due to connection timeouts. Also two large peers - AS7660 and AS8297 - were absent on 26th. Each contributed with more than 16,000 paths. Thus our view point is very limited and can hardly be used to distinguish between worm attack and peer loss.

DDoS Attack on the RIPE NCC. Starting from 2pm (UTC) February 27th, 2003, the RIPE NCC network suffered a large DDoS attack (a distributed ICMP echo attack). Their network structure was affected not only ICMP traffic. Network condition returned back to normal the same day at 4:30pm UTC. Shortly before (2am till 2pm) the peer AS5459 did not contribute, causing a loss of $\approx 7,100$ paths. Nothing unusual could be observed.

Blackouts. During August and September 2003, several blackouts happened in different geographic regions. On August 14th, around 6pm (UTC) started the cascading chain of events that lead to a gigantic blackout in the north-eastern part of America and Canada (refer to [1] for details). A 40-minute blackout took place in London on August 28th (around 6 pm). The last one struck parts of France, Italy and Switzerland on September 28th. It started around 2:30 am and power returned during the afternoon. The event in London could not be observed, since it was too short. The other two exhibited a similar pattern: a



(a) US and Canada



(b) France, Italy and Switzerland

Figure 12. Evolution of different blackouts

large drop in the number of paths occurred at the start of the event. This drop caused also a drop in the static indices (number of nodes and number of edges). The recovery phase was very short and the indices quickly got back to their old scores. The drops in the number of nodes and edges of the commonly observed AS graph (Figure 12(a)) is not due to the blackout itself. They are results of the disruption in the collection process. The reduced data set was relatively stable. Most node were present more than 80% of the time. Only a small set (≈ 20 nodes) appeared only once or twice. The clustering itself is relatively stable. Figure 13 shows the temporal development of the sizes of the clusters. Smaller migrations, splittings or merging phenomena could be observed. Although there is no clear pattern or regularity apparent, it is a fact that these changes occurred more frequently during the recovery phase, thus giving a good indicator that clustering reflects aspects of the event. Similar observations could be made during the




Figure 13. Sizes of the different clusters during the blackout in the US and Canada

blackout in Europe. However, the impact on the structure was not comparable with respect to size and the recovery time much shorter.

5 Conclusions

Clustering and reduction techniques are well-known and widely-used approaches for analyzing large amounts of data. We adapted these tools to the domain of the dynamic analysis of the Internet graph at the AS level. Several experiments were performed to deepen the knowledge of short- and long-term phenomena. The first target was to describe the baseline behavior of the network in the absence of pathologic phenomena. Although the analysis was relatively plain, we could verify several intuitional properties. A second effort went in the direction of measuring the impact of short-lived events, such as viruses, misconfigurations and blackouts. We found, that viruses and worms had no measurable effect on the network structure, even if Internet services might have been disrupted. In contrast, blackouts and DDoS showed measurable consequences. Obviously, the analysis showed indisputable and precise evidences of blackouts and power outages, but the symptoms registered for DDoS are perhaps more promising from a practical point of view.

Although hampered by lots of technical problems, as the continuous changing of the available vantage points or the uncertain identification of the ever-changing clusters through time, we think that the dynamic analysis of the AS-graph is a promising field of research that could have many useful outcomes in the future. In this paper we made a first step towards the definition of a practical methodical approach to this intriguing problem.

References

- [1] A timeline of the 2003 blackout. <http://www.cnn.com/2003/US/08/16/blackout.chron.ap/index.html>.
- [2] Characterizing the internet hierarchy from multiple vantage points. <http://www.cs.berkeley.edu/~sagarwal/research/BGP-hierarchy/>.
- [3] S. Agarwal, L. Subramanian, J. Rexford, and R. Katz. Characterizing the internet hierarchy from multiple vantage points. In *Proceedings of IEEE Infocom 2002*, 2002.
- [4] Charles J. Alpert and Andrew B. Kahng. Recent directions in netlist partitioning: A survey. *Integration: The VLSI Journal*, 19:1–81, 1995.
- [5] V. Batagelj and M. Zaveršnik. Generalized cores. Preprint 799, Universtiy of Ljibljana, 2002.
- [6] Vladimir Batagelj, Anuska Ferligoj, and Patrick Doreian. Generalized blockmodeling. *Informatica (Slovenia)*, 42(4), 1999.
- [7] Code Red II and Nimda Worms and BGP Instability. http://www.renesys.com/projects/bgp_instability.
- [8] Ulrik Brandes, Marco Gaertler, and Dorothea Wagner. Experiments on graph clustering algorithms. In *Proceedings of the 11th Annual European Symposium on Algorithms (ESA'03)*, Lecture Notes in Computer Science. Springer-Verlag, 2003. To appear.
- [9] The Code Red Worm. <http://www.ciac.org/ciac/bulletins/1-117.shtml>.
- [10] G. Di Battista, M. Patrignani, and M. Pizzonia. Computing the types of the relationships between autonomous systems. In *IEEE INFOCOM 2003*, 2003.
- [11] Patrick Doreian, Vladimir Batagelj, and Anuska Ferligoj. Symmetric-acyclic decompositions of networks. *Journal of Classification*, 17(1):3–28, 2000.
- [12] T. Erlebach, A. Hall, and T. Schank. Classifying customer-provider relationships in the internet. In *Proceedings of the IASTED International Conference on Communications and Computer Networks*, pages 538–545, 2002.
- [13] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM*, pages 251–262, 1999.
- [14] Lixin Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, 2001.
- [15] Christos Gkantsidi, Milena Mihail, and Ellen Zegura. Spectral analysis of internet topologies. In *IEEE Infocom 2003*, 2003.
- [16] R. Govindan and R. Reddy. An analysis of internet inter-domain topology and route stability. In *IEEE INFOCOM 1997*, 1997.
- [17] Geoff Huston. BGP table statistics. <http://bgp.potaroo.net>.
- [18] A. K. Jain, M. N. Murty, and P. J. Flynn. Data clustering: A review. *ACM Computing Surveys (CSUR)*, 31(3):264–323, 1999.
- [19] KPNQwest limps back after shutdown. http://news.com.com/2104-1033_3-946262.html.
- [20] David Meyer. University of oregon route views project. <http://www.routeviews.org/>.
- [21] The internet outage and attacks of october 2002. <http://www.isoc-chicago.org/internetoutage.pdf>.
- [22] M. Rimondini, M. Pizzonia, G. Di Battista, and M. Patrignani. Algorithms for the inference of the commercial relationships between autonomous systems: Results analysis and model validation. submitted to this workshop, 2004.
- [23] S. B. Seidman. Network structure and minimum degree. *Social Networks*, 5(5):269–287, 1983.
- [24] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos. Power laws and the as-level internet topology. In *IEEE/ACM Transactions on Networking (TON)*, volume 11, pages 514–524. ACM Press, 2003.
- [25] S. Tauro, C. Plamer, G. Siganos, and M. Faloutsos. A simple conceptual model for the internet topology. In *IEEE Globalcom 2001*, volume 3, pages 1667–1671, 2001.
- [26] Test Traffic Analysis Update. <http://www.ripe.net/ripe/meetings/archive/ripe-43/presentations/ripe43-tt-analysis/>.
- [27] The W32.nimda Worm. <http://www.ciac.org/ciac/bulletins/1-144.shtml>.