IET Information Security

Research Article

Border gateway protocol graph: detecting and visualising internet routing anomalies

ISSN 1751-8709
Received on 24th October 2014
Revised on 24th April 2015
Accepted on 14th July 2015
doi: 10.1049/iet-ifs.2014.0525
www.ietdl.org

Stavros Papadopoulos[1] ✉, *Konstantinos Moustakas*[2], *Anastasios Drosou*[3], *Dimitrios Tzovaras*[3]

[1]*Electrical and Electronic Engineering Department, Imperial College London, SW7 2AZ, London, UK*
[2]*Electrical and Computer Engineering Department, University of Patras, 26504, Rio Campus, Patras, Greece*
[3]*Information Technologies Institute, Centre of Research and Technology – Hellas, 6th km Xarilaou – Thermi 57001, Thessaloniki, Greece*
✉ *E-mail: spap@iti.gr*

Abstract: Border gateway protocol (BGP) is the main protocol used on the Internet today, for the exchange of routing information between different networks. The lack of authentication mechanisms in BGP, render it vulnerable to prefix hijacking attacks, which raise serious security concerns regarding both service availability and data privacy. To address these issues, this study presents BGPGraph, a scheme for detecting and visualising Internet routing anomalies. In particular, BGPGraph introduces a novel BGP anomaly metric that quantifies the degree of anomaly on the BGP activity, and enables the analyst to obtain an overview of the BGP status. The analyst, is afterwards able to focus on significant time windows for further analysis, by using a hierarchical graph visualisation scheme. Furthermore, BGPGraph uses a novel method for the quantification of information visualisation that allows for the evaluation, and optimal selection of parameters, in case of the corresponding visual analytics algorithms. As a result, by utilising the proposed approach, four new BGP anomalies were able to be identified. Experimental demonstration in known BGP events, illustrates the significant analytics potential of the proposed approach in terms of identifying prefix hijacks and performing root cause analysis.

1 Introduction

Today's rapidly expanding Internet provides data delivery and communication service to millions of end users. At the same time, the border gateway protocol (BGP) is responsible for exchanging external routing information amongst the autonomous systems (AS) that comprise the Internet.

Due to the lack of authentication mechanisms in BGP, an unauthorised network can originate prefixes [A prefix defines a set of consecutive IP addresses] owned by other networks [1]. This procedure is called prefix hijacking, and raises serious security issues concerning both service availability and data privacy. In general, there are two types of prefix hijacks [2]: hijack of prefix ownership and AS-path hijack. Because of these security issues, the BGP protocol is vulnerable to either intentional or unintentional attacks from ASes. One of the most common cases of prefix hijack is router misconfiguration [1, 3], which causes the announcement of multiple unauthorised prefixes, leading to large service availability problems in other networks as well (e.g. AS-9121 incident [4]).

Thus far, previous efforts for providing solutions on prefix hijacking are presented from two aspects [5]: hijack prevention and hijack detection. Hijack prevention solutions include cryptographic based authentications [6, 7], where BGP routers sign and verify the origin AS of each prefix. On the other hand, hijack detection mechanisms [1, 8, 9] are provided when a prefix hijack has already happened and needs to be detected.

This paper focuses on prefix hijack detection, with the use of anomaly metrics and visualisation. Specifically, two types of anomalies are addressed: isolated hijacks, and distributed anomalies. Isolated hijacks concern a prefix ownership hijacking between a specific pair or a small number of networks, in which one network originates prefixes belonging to this small set of other networks. Distributed anomalies concern cases of large routing deviation, such as when a specific network originates prefixes that belong to multiple other networks, or a network that withdraws a lot of prefixes due to internal problems. These two categories of anomalies will be referred to, for the rest of the paper as *small scale* and *large scale* anomalies, respectively.

2 Related work

Quite a number of methods have been proposed in the literature for BGP prefix hijack detection. Deshpande *et al.* [8] proposed the use of multiple features that characterise the BGP activity. A generalised likelihood test is applied on each feature separately, to detect time windows that deviate more than a specified threshold, regarding to the normal activity, and thus, are rendered as candidates for including anomalies. The alerts created for each feature are afterwards correlated in time. In addition, Zhang *et al.* [10] proposed the use of wavelets applied on multiple feature vectors extracted from BGP data. This approach is able to detect outliers for each feature separately, which represent possible BGP anomalies. Moreover, Al-Rousan *et al.* [9] proposed the generation of multiple features extracted from the BGP messages, that characterise the BGP activity at different time periods. Afterwards, the authors used feature selection methods to select the most descriptive features, which are fed to a naive Bayes classifier for the identification of anomalous time periods. Each of the aforementioned approaches, are able to detect time windows of instability, but further require further processing to identify the ASes responsible for the detected anomalies. To overcome this issue, Khare *et al.* [1] proposed a method that identifies suspicious routing announcements, by taking each AS past routing announcements into account. Based on the assumption that the prefixes are usually announced by a stable set of ASes during their lifetime, the authors proposed a metric that captures the ASes that affect multiple other networks simultaneously. The approach proposed in this paper, also addresses the root cause analysis problem by utilising visualisation methods, which enable the analyst to identify the ASes responsible for the anomalies. In addition, unlike previous approaches which target the identification of specific types of anomalies, this paper addresses the identification of both *small scale* and *large scale* anomalies in a common framework.

There are many network visualisation systems available today [11], but very few of them deal with BGP routing changes. BGPlay [12] allows Internet Service Providers to monitor the reachability of a specified prefix from the perspective of a given border router, while incorporating animation to highlight its routing changes. Cortese *et al.* [13] used the idea of topographic maps to enhance BGPlay visualisations by positioning the ASes to different areas on the map according to their ranking. Wong, *et al.* [14] proposed TAMP, a system that uses statistical methods to aggregate BGP data, to visualise and diagnose BGP anomalies. AS are displayed as nodes while connectivity is illustrated by using links. More attributes such as size and colour are used to encode auxiliary information. The LinkRank visualisation approach [4], which is the closest work to the one presented on this paper, provides a high level view of Internet routing changes. Each node in the graph represents an AS, while the size of the nodes illustrates the amount of IP ownership change. All the aforementioned approaches do not take into account the information content of the visualisations, and the information loss, caused by the mapping of the input data to the visualisation. This information loss can cause the analyst to misinterpret the patterns in the data and conduct wrong conclusions. To address this issue, this paper utilises a novel method for the selection of the appropriate parameters in the visualisation so as to minimise the information loss, and as a result maximise the information that is presented to the analyst using the proposed visualisation scheme.

3 Motivation-contribution

Most of the previously referred approaches for prefix hijack detection, focus on the identification of time windows that include anomalies, but require further processing for root cause analysis, that is, the identification of the ASes responsible for the detected anomalies. Towards this end, the proposed BGPGraph approach on this paper, utilises visualisation methods for root cause analysis, in which the analyst is presented with a time series of anomaly scores, and is able to focus on significant time periods for further analysis. In addition, unlike previous approaches which target the identification of specific types of anomalies, this paper addresses the identification of both *small scale* and *large scale* anomalies in a common framework.

Recently, it has been suggested [15, 16] that it might be possible to employ concepts from the theories of data communication for the purpose of evaluating and improving the effectiveness of information visualisation techniques. Information theory can be used to measure the information content of a specific visualisation approach, and the information loss that is inevitably caused by the components of the visualisation pipeline [15]. Such an approach could enable quantitative comparisons of visualisation approaches, with respect to their information content. Moreover, it could provide sufficient means to improve the way information is visualised and, thus, lead to automated visualisation optimisation processes. Towards this direction, the present work presents the use of entropy measures to capture the information loss, caused by the visualisation of the input dataset. The proposed metrics are afterwards utilised to recognise the definition of the parameters of the visualisation, so as to maximise its information content. This optimisation scheme is applied in a highly challenging analytics problem, namely the Internet routing visualisation, where approaches are challenged in, both terms of analytics potential and scalability.

The contributions of BGPGraph are summarised as follows:

- Propose multiple descriptive metrics for anomaly detection in BGP, taking into account both prefix hijacks and routing information.
- Identify both *small scale* and *large scale* anomalies using a common framework.
- Perform root cause analysis in addition to anomaly detection, using visualisation methods that are optimised with respect to their information content.

4 Analysis of small scale anomalies

This section describes the analysis procedure followed to define multiple metrics, capable of capturing *small scale* anomalies. In the context of this paper, the term *small scale* anomalies refers to an isolated or a small set of prefix hijacks caused by the same network. For the identification of isolated prefix hijacks the multiple origin AS (MOAS) events are taken into account. Specifically, a MOAS event occurs when the same prefix appears to belong to at least two different ASes, that is, the same prefix is announced by two different ASes. The analysis utilised, is based on the methodology presented in [17], applied in the case of MOAS events.

Specifically, four descriptive metrics are extracted and assigned to every MOAS event based on the country of origin of the ASes involved in these events (more details can be found in [17]):

i. *CAP*: Probability of appearance of each MOAS event, as calculated for a specific target country.
ii. *CAPZ*: Z-score of the probability of appearance *CAP* of each MOAS event, as calculated for a specific target country.
iii. *CGL*: The geographic length between the two countries involved in the MOAS incident.
iv. *CGLZ*: The z-score of the geographic length *CGL* between the two countries involved in the MOAS incident.

These four metrics characterise the MOAS behaviour, but the detection of MOAS anomalies without the combination of these metrics into a common framework is still difficult. The work presented in [17] does not combine these metrics, but instead uses scatterplots to show that they are able to discriminate between normal and abnormal cases. Towards this end, BGPGraph combines these small scale metrics into a single anomaly metric using a weighted sum approach. The generated metric is able to capture *small scale* anomalies in an efficient manner utilising the descriptive power of the small scale metrics.

5 Hierarchical visualisation

This section presents the visualisation method that is used in the context of the BGPGraph approach for root cause analysis. Specifically, a graph metaphor is utilised to visualise the BGP routing changes on per AS level. The graph representation is commonly used in the visualisation domain [18]. The reason for this is that it is a very intuitive way to present the connections between objects, as it directly targets the human perception system [18]. The objects are represented by vertices, while their connections/relationships by edges.

It must be underlined that the graph representation of the AS inter-relationships is too large (more than 50,000 AS and 100,000 connections) to be visualised on a limited display size. For this reason, a hierarchical clustering scheme is utilised, to reduce the size of the visualised graphs and enable scalable analysis.

The AS-graph is a graph $G(V, E)$ whose set of vertices $V = \{v_i | i \in [1, N]\}$, where N is the total number of ASes, is the set of all the ASes and each vertex represents an AS. The set of edges is $E = \{e_i(v_j, v_k) | v_j, v_k \in V\}$, where each edge represents the existence of a physical connection between the ASes. The graph is created from the AS-paths of the BGP announcements.

The proposed graph visualisation approach, is used to visualise the difference on traffic between consecutive time windows, measured in the number of IP addresses which changed paths. The reason for using the difference of IPs as weights of the edges and the vertices, lies in the nature of the Internet routing protocol. When a source AS sends information to a destination AS, this information follows a path consisting of ASes and links. As it was mentioned

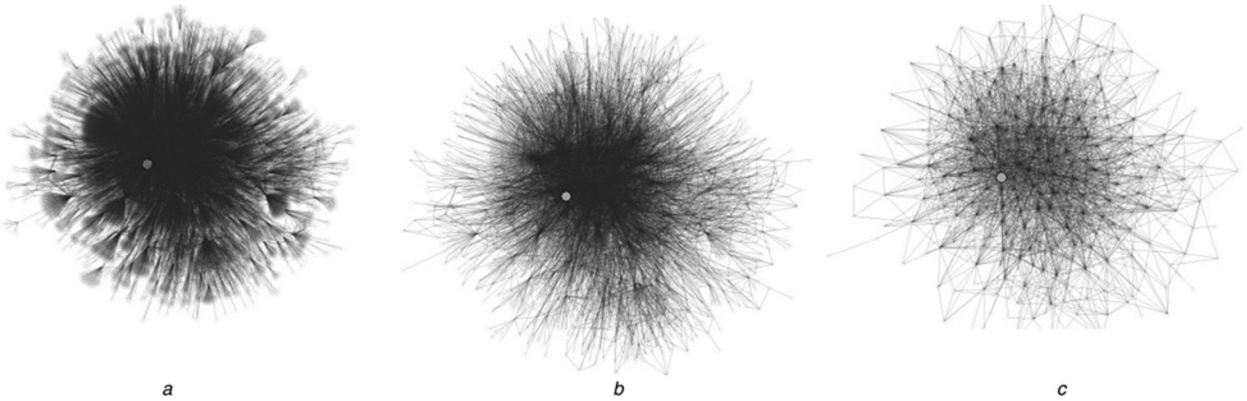

Fig. 1 *Cluster Hierarchy of the AS-Graph. As it is shown the number of vertices is reduced as the level increases, but the topology of the original graph (a), is preserved. The monitoring point is the orange node*

a level-0
b level-3
c level-7

earlier, if a link or an AS fails (for various reasons such as misconfiguration, hardware problem etc.), according to the inherent Internet dense architecture, the router searches for another path to reach its destination. Thus, there is a transfer of BGP routes from one path to another. Another example of route transfer is the case of prefix hijacks. This routing transition is captured well, using the proposed IPs different metrics. Thus, the visualisation shows these events in a practical and meaningful way. Red colour in the visualisations represents negative weights, while green colour represents positive weights.

The magnitude of the corresponding weight is mapped to the width of the edges or the radius of the vertices, which are represented by two sets, one that defines the quantisation levels of the vertex radii and another that defines the quantisation levels of the edge widths:

$$R = \{r_i | r_{\min} \le r_i \le r_{\max},\ i \in [0, L_r] \text{ and } i \in \mathbb{N}\} r_i \neq r_j,\ \text{for } i \neq j,\ r_0 = 0 \text{ and } r_i < r_{i+1} \quad (1)$$

$$W = \{w_i | w_{\min} \le w_i \le w_{\max},\ i \in [0, L_w] \text{ and } i \in \mathbb{N}\} w_i \neq w_j,\ \text{for } i \neq j,\ w_0 = 0 \text{ and } w_i < w_{i+1} \quad (2)$$

where r_i is the radius of the i_{th} circle in visual degrees, $r_{\min}$ is the minimum radius and $r_{\max}$ is the maximum radius. L_r is the number of different radius values or alternatively the radius quantisation levels. The same notation is used for the edges width set W. Furthermore, it is assumed that the sets are ordered, which is expressed through the relation: $r_i < r_{i+1}$. To define the sets R and W, the visual acuity of the human visual system needs, to be initially measured and quantified. There are many studies and metrics on the acuity of the human visual system [18], but the most practical one is the Weber's Law or Just-Notable Difference [19], which is the one utilised here. It is practical because it is directly applicable to the task of defining the sets R and W due to the fact that it calculates the minimum required difference in stimuli in order for a human to understand that the inputs are different. Other alternatives for perceiving visual stimuli such as Ricco's law and Stevens' power law [20] do not measure the minimum required difference, but instead measure the perceived intensity with respect to the actual intensity. In addition, the Weber's Law, has been widely utilised in multiple research studies addressing perception issues in visualisation [21–23].

Since the AS-Graph is too large to be visualised at once, a clustering method is utilised to produce a hierarchy of coarse to fine graphs that are easier to visualise and perceive in a hierarchical manner, while still maintaining the significant information. Each cluster represents a collection of ASes. For the clustering creation the approach proposed by Gansner *et al.* [24] is utilised. In this case the notation $G^l(V^l, E^l)$ represents the graph of the l level of the hierarchy. The superscript is used in general to represent the level. V^l is the set of vertices and E^l is the set of edges of the lth level. In this case $G^0(V^0, E^0) \equiv G(V, E)$. The results of the application of the clustering algorithm on the AS-Graph of 2004 are shown in Fig. 1.

It should be noted that the utilised clustering approach [24] allows the analyst to select parts of the graph, that must be visualised with higher granularity than the rest of the graph, and thus, create hybrid graphs comprised of multiple levels.

6 Quantification of the information

In this section the information content of the input data and proposed visualisation method is quantified using entropy metrics. The input dataset will be referred to as the input signal, while the visualisation will be referred to as the output signal. These metrics are used for the calculation of the information loss caused by the mapping of the input data to the visual attributes of the visualisation.

6.1 Entropy of the input signal

The input signal represents the input dataset that is used by the proposed approach for visualisation purposes. In the respective case, the input signal is composed of the weights that directly reflect the routing changes between consecutive time windows.

The entropy metric is used to calculate the information content of the entire graph that is currently visualised. The notation $G(V^c, E^c)$ represents the graph that is currently visualised by the system, along with all the respective weights. The graph $G(V^c, E^c)$ can represent either a level of the hierarchy $c = l$, or a hybrid graph comprised of vertices end edges from different levels. This way, the entropy of the edge weights, of the entire graph, that is currently visualised is defined as:

$$H_G^{\text{in}}(E^c) = -\sum_{i=1}^{Y} \frac{y_i}{y_{\text{total}}} \log\left(\frac{y_i}{y_{\text{total}}}\right) \quad (3)$$

where Y is the number of different edge weight instances that are visualised, y_i is the number of occurrences of the ith weight and $y_{\text{total}} = \sum_{i=1}^{Y} y_i$ the total number of weight occurrences, but with respect to the edge weights of the entire graph. The same procedure is repeated to find the entropy $H_G^{\text{in}}(V^c)$ of the vertex weights of the entire graph.

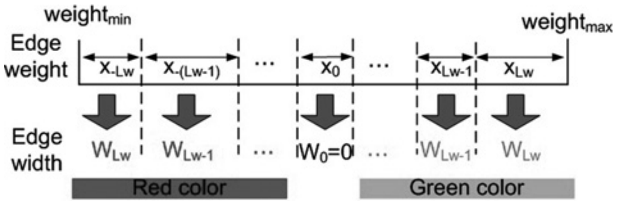

Fig. 2 *Mapping function $F: \Re^n \rightarrow V^m$ from the edge weight to the edge width. Green colour represents positive weight values and red negative*

6.2 Entropy of the output signal

In general, a visualisation system uses a mapping function to map the input data to the visualisation features: $F: \Re^n \rightarrow V^m$, where $\Re^n$ is the input signal in the space of real numbers and has n features, and V^m is the visualisation space (scatterplots, graphs, glyphs etc.) that has m features. In the proposed scheme, the output signal is the visualised graph. Because of various factors, such as the display capacity, visual clutter and the limitations of the human vision system, a transfer of all the information of the input to the output signal, is usually not possible. As explained in Section 5, in the context of the proposed approach, the edge weights are mapped to the width and colour of the edges and the vertex weights are mapped to the radius and colour of the vertices. Hence the visualisation space V^m has four features ($m = 4$): the width and colour of edges and the radius and colour of vertices.

Furthermore, in (1) and (2), the set of the vertex radius values R and the set of edge width values W were defined. In addition, $|W| = L_w$ quantisation levels were defined for the edges width and $|R| = L_r$ quantisation levels were defined for the vertices radius. The role of the mapping function F is to map the edge and vertex weights to the elements of the sets W and R. Fig. 2 shows the mapping function $F: \Re^n \rightarrow V^m$, that maps the edge weight to the edge width. A similar mapping function is used for the mapping of the vertex weight to the vertex radius.

To formulate the mapping function F, a variable X_i is initially defined as follows:

$$X_i = \left\{ \begin{array}{ll} 0, & i < -L_w \\ \sum_{j=-L_w}^{i} x_j, & -L_w \leq i \leq L_w \\ \sum_{j=-L_w}^{L_w} x_j, & i > L_w \end{array} \right\} \quad (4)$$

where x_i represents the size of the section of the input weight that is mapped to the width w_i or radius r_i, as depicted in Fig. 2. The mapping function F is defined as:

$$F(e_{\text{weight}}) = |i|, \quad \text{if} \quad X_{i-1} \leq e_{\text{weight}} < X_i \quad (5)$$

where e_{weight} is the corresponding edge weight that is mapped to width $w_k \in W$, for $k = F(e_{\text{weight}})$.

To find the entropy of the output signal, (3) is used, but the difference is that the weights are first mapped to the elements of the sets W and R through the mapping function F. This way the corresponding entropies $H_G^{\text{out}}(E^c, F)$ (entropy of the edges weights of the entire graph), $H_G^{\text{out}}(V^c, F)$ (entropy of the vertices weights of the entire graph) are calculated. The superscript 'out' represents the output signal. It is apparent that the value of the output entropy depends on the mapping function $F: \Re^n \rightarrow V^m$ that inevitably induces information loss.

7 BGPGraph framework

This section is comprised of two subsections. Section 7.1 utilises entropy metrics to define the mapping function F from the input edge and vertex weights to the visual attributes of the graph, so as to minimise the information loss caused by their visualisation. Section 7.2 presents the proposed BGP anomaly metric, which is capable of capturing both *small* and *large scale* anomalies.

7.1 Visualisation optimisation

In this section, the proposed visualisation method is optimised with respect to its information content. Due to limited display size, and in many cases, limited number of visual attributes (e.g. limited number of available edge widths or vertex sizes), the transfer of all the information to the visualisation is usually not possible. The amount of information that is not visualised represents the amount of information loss. Large information loss reflects poor visualisation quality, since most of the data are not visualised and important patterns are lost (e.g. difference in the size of two vertices/edges might not be visible). On the other hand, low information loss reflects the fact that most of the data are visualised, and insights about the data are easier by the analyst. Thus, low information loss represents good quality visualisations. The optimisation procedure results in visualisations which convey the larger amount information (e.g. reveals a difference in the size of two vertices/edges not seen previously), and thus, render the analytical procedure easier.

The entropy metrics defined in Section 6 are used to optimise the visualisation result. In particular, the method proposed, estimates the mapping function $F: \Re^n \rightarrow V^m$, which maximises the entropy of the output signal and as a consequence, to minimise the information loss, caused by the traversal of information through the visualisation pipeline [15]. This procedure enables the analyst to make decisions guided through the acquired information, while minimising the amount of false conclusions that might arise due to large information loss. The proposed optimisation procedure highlights parts of the graphs that have high information content, and as a result are relevant to the analysis.

In Section 6.2, is explained that the role of the mapping function F is to map the edge and vertex weights to the elements of the sets R and W ((1) and (2)). Without loss of generality only the edge width case will be analysed in this section.

The mapping function $F: \Re^n \rightarrow V^m$ is defined in such a way, so that the entropy of the output signal is maximum, according to the following optimisation problem:

$$\boldsymbol{x} = \arg\max_{\bar{x}} \left\{ H_G^{\text{out}}\left(E^c, F(\bar{x})\right) \right\} \quad (6)$$

where $\boldsymbol{x}$ is a vector of edge weight quantisation ranges $\boldsymbol{x} = (x_{-L_w}, x_{-(L_w-1)}, \ldots x_0, \ldots x_{(L_w-1)}, x_{L_w})$. The mapping function F is defined in (5). Furthermore, H_G^{out} is the output entropy of the edge weights of the currently visualised graph.

The result is the definition of the mapping function from the input data to the sizes of the nodes and edges of the graph, so as to reduce the number of values that are mapped to the same sizes, and enable the analyst detect otherwise not visible small differences between different values. Results of the minimisation of (6) on the visualisation are illustrated in Figs. 5–7.

The downhill simplex method [25] is utilised for the optimisation procedure. Alternative optimisation schemes have also been considered and tested (such as gradient descent, simulated annealing, and compass search [25]). However, it is highly underlined that the proposed optimisation problem has many local minimum values, and it is very difficult to identify the exact global minimum in such problems. The downhill simplex method was selected due to the fact that it provided better results in the majority of the cases, since it was able to identify a local minimum closer to the global minimum.

7.2 BGP anomaly metric

This section presents the procedure followed for the definition of the *small* and *large scale* anomaly metrics, as well as their subsequent fusion for the definition of a single anomaly metric, capable of capturing both *small* and *large scale* anomalies.

7.2.1 Large scale anomaly metric: In Section 6, two entropy metrics where proposed so as to quantify the information of the input signal (i.e. routing changes) at each time instance. The main characteristic of the *large scale* anomalies is that they represent large routing deviation between ASes and links. The large number of routing deviations is captured well using the proposed entropy metrics, and thus, they are used for the definition of the *large scale* anomaly metric. Specifically, the entropy of the edge weights of the entire graph $H_G^{\text{in}}(E^c)$, and the entropy of the vertex weights of the entire graph $H_G^{\text{in}}(V^c)$, are utilised for this calculation.

It should be pointed out that the entropy is not a perceptually linear metric. Specifically, the difference of one (measured in bits) indicates twice as much information. To make the entropy metric perceptually consistent, the entropies are transformed as follows: $H_G^{\text{in}'}(E^c) = 2^{H_G^{\text{in}}(E^c)}$ and $H_G^{\text{in}'}(V^c) = 2^{H_G^{\text{in}}(V^c)}$.

The *large scale* anomaly metric is defined as the fusion of the entropies of the edges and vertices of the current graph, and is defined as:

$$A_E = w_1{}^* H_G^{\text{in}'}(E^c) + w_2{}^* H_G^{\text{in}'}(V^c) \tag{7}$$

where w_1 and w_2 are the weights of the edges and vertices entropy, respectively, and $H_G^{\text{in}'}(E^c) = 2^{H_G^{\text{in}}(E^c)}$, $H_G^{\text{in}'}(V^c) = 2^{H_G^{\text{in}}(V^c)}$. The higher the value of the *large scale* anomaly metric is, the higher the probability of a *large scale* BGP anomaly occurrence.

The calculation of the *large scale* anomaly metric depends on two parameters: the selected time step ΔT, and the value of *c*, which can represent either a level of the hierarchy $c = l$, or a hybrid graph comprised of vertices end edges from different levels. As it was mentioned in Section 5, the clustering algorithm was designed in such a way that the higher levels of the clustering hierarchy contain the important information in an abstract form, while irrelevant information is filtered. Thus, by using high levels of the clustering hierarchy for the calculation of the entropy anomaly metric, the clutter induced by the irrelevant information is reduced. The value of $c = 11$ was used for the experiments on this paper. The value of the second parameter, the time step ΔT, depends on the size of the time period *T* that is examined for anomalies. For large time periods, larger time steps ΔT must be selected, to reduce the number of iterations needed. In general the algorithm needs $r = T/\Delta T$ iterations. On the other hand, for small time periods the value of the time step ΔT must be relatively small, so as to achieve finer granularity.

There are multiple research works which utilise ΔT parameter for aggregation and feature extraction purposes. In all of the cases the value of ΔT is set manually to a period of 1 to 10 minutes [8, 9, 26–28]. In the context of this paper, similarly to previous approaches, the value of the ΔT parameter is manually selected. Extensive experimentation on multiple known prefix hijack events, where the algorithm was applied on single day periods ($T = 1$ day), revealed that a value of $\Delta T = 30$ s is sufficient for their successful identification, while still keeping the number of algorithm iterations/runtime small. Specifically, Fig. 3 shows the effect that the different values of the ΔT parameter has on: (i) The average running time of the algorithm, (ii) The number of algorithm iterations (in logarithmic scale), (iii) The number of small scale detections, and (iv) The number of large scale detections. The data represent a total period of $T = 1$ day. The number of small scale detections is static (equal to 15), due to the fact that the maximum operator is utilised within the time window ΔT (see (9)). On the contrary, the number of large scale anomalies that are detected by the algorithm drops from the maximum which is 5 for $\Delta T \leq 300$ s, down to 3 for $\Delta T = 600$ s, and 1 for $\Delta T > 600$ s. This means that $\Delta T = 300$ s is a threshold for the sufficient identification of the large scale anomalies. In this paper, $\Delta T = 30$ s is chosen due to the fact that it has a sufficient distance from the threshold of 300 s, while the running time of the algorithm is still relatively small (around 10 minutes).

7.2.2 Small scale anomaly metric: In Section 4 four metrics were defined that are capable of capturing *small scale* anomalies.

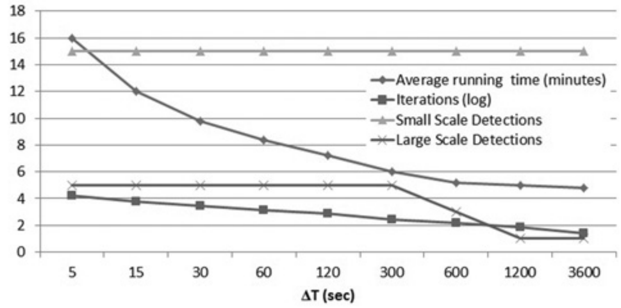

Fig. 3 *Effect of the* ΔT *parameter has on: (1) The average running time of the algorithm, (2) The number of algorithm iterations (in logarithmic scale), (3) The number of small scale detections, and (4) The number of large scale detections. The data represent a total period of T = 1* day

In this section, these four metrics are fused into one metric, the *small scale* anomaly metric. This metric is used to discover *small scale* events that deviate from normal behaviour, and as a result are considered as anomalies. First of all, the four metrics defined in Section 4 are normalised to [0, 1] and afterwards the fusion procedure is defined as follows:

$$A'_M = w_3(1 - \text{CAP}) + w_4(1 - \text{CAPZ}) + w_5\text{CGL} + w_6\text{CGLZ} \tag{8}$$

where w_3, w_4, w_5, and w_6 are the weights of each of the four small scale metrics.

It should be noted that the anomalous BGP MOAS incidents are characterised by low *CAP* and *CAPZ* values, while also demonstrating high *CGL* and *CGLZ* values. Thus, by using (1 − *CAP*), (1 − *CAPZ*), (*CGL*) and (*CGLZ*) in the calculation of the *small scale* anomaly metric, high values of this metric point to possible BGP MOAS anomalies.

Unlike the case of the *large scale* anomaly metric in which the calculation is performed in predefined time windows ΔT, the *small scale* anomaly metric is defined for each MOAS incident detected. To facilitate the fusion procedure, the two metrics must be defined in a common time frame. Thus, the *small scale* anomaly metric A'_M is registered in specific time windows ΔT as follows:

$$A_M = \max(A'_M), \text{ for each time step } \Delta T \tag{9}$$

where ΔT is the time step that is also used for the *large scale* anomaly metric. In other words, the maximum *small scale* anomaly score of all the MOAS events that occurred within a specified time window, is chosen as a representative for this time window.

7.2.3 Global anomaly metric: The global anomaly metric is defined as the weighted sum of the normalised values of the *large scale anomaly metric* and *small scale anomaly metric*:

$$A = w_E{}^* A_E + w_M{}^* A_M \tag{10}$$

Thus, the calculation of the BGP anomaly metric depends on the values of the parameters of the following vector: [*c*, ΔT, w_1, w_2, w_3, w_4, w_5, w_6, w_E, w_M]. The value of this vector is set to [11, 30 *sec*, 0.5, 0.5, 0.25, 0.25, 0.25, 0.25, 0.5, 0.5] so as to take all the metrics equally into account.

By utilising the weighted sum approach, large importance metrics provide large contribution to the value of the final anomaly metric. In other words, the weights characterise the degree of contribution of each metric to the final anomaly score.

The BGP anomaly metric *A* is used to construct plots that summarise the anomaly score of the BGP activity over a specific period of time, called anomaly plots. These plots are bar plots, in which the length of each bar represents the magnitude of the corresponding anomaly score for the specific time window ΔT.

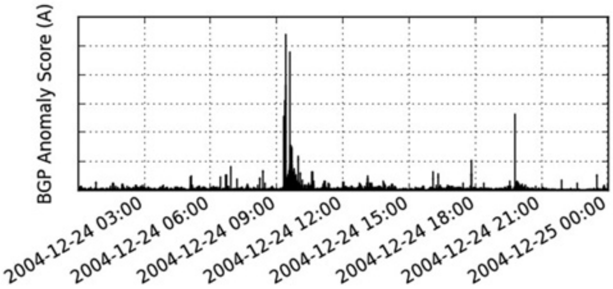

Fig. 4 *Anomaly plot for the 24-Dec-2004. The time step is set to* $\Delta T = 30$ *s. Two spikes are clearly observed, one around 9:24 GMT, and one around 19:48 GMT. The monitoring point is AS-3549*

Using the anomaly plots, the analyst can select to focus on specific periods of time were the anomaly score is high and as a result something significant might be happening. Afterwards, the analyst can examine the BGPGraph visualisation of the selected period, so as to perform the analysis on a greater depth.

8 Event analysis and experimental results

In this section the BGPGraph framework is used to explore state of the art prefix hijack events, and demonstrate the analytical potential of the proposed approach. The BGP announcement data in this paper are collected from the RIPE [29] BGP monitoring project, and specifically from the monitoring points of AS-3333, AS-3549, AS-4608, AS-4777, and AS-7018. A list of 12 known anomalous BGP events which are visible from these monitoring points is used to evaluate the BGPGraph approach, while additional new events are being detected, by utilising the proposed method are also presented.

8.1 Event analysis – router misconfiguration

In this section, the BGPGraph framework is applied on the detection and analysis of a BGP router misconfiguration event that occurred on 24th December 2004 [4]. The observations in this section are made through the monitoring point of AS-3549.

Fig. 4 depicts the anomaly plot of 24th December 2004. Observing this overview plot, a pattern emerges. At 9.24 GMT the anomaly score suddenly increased its value by a factor of five.

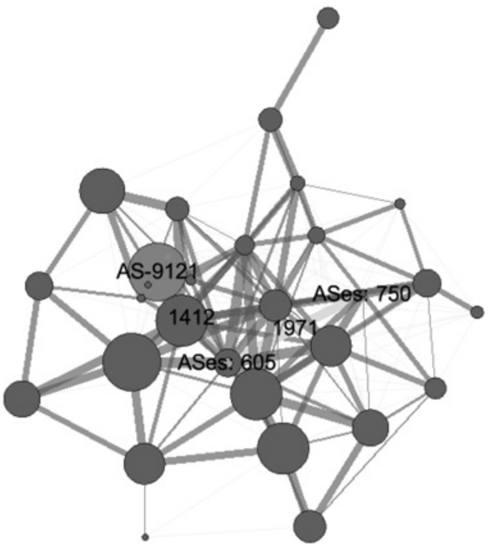

Fig. 6 *Entropy optimised view of level 11 after increasing the granularity of the green cluster of Fig. 5. It is apparent that AS-9121 is the one that hijacks a big portion of Internet routes*

This is an indicator that something might be happening. As a result, the user selects a small time window around 9.24 GMT and the AS-Graph of the level-11 of the clustering hierarchy is then visualised (Fig. 5*a*), thus allowing a more focused and in depth analysis.

Observing Fig. 5*a*, it is apparent that one cluster of ASes gains weight (gains IPs), and is represented with green colour, while all the neighbouring clusters lose weight (lose IPs), a fact which is shown by using red colour. This behaviour is also observed in the case of the edges, which all lose weight, except for one path having as starting point the green cluster.

After increasing the granularity of the cluster that is gaining weight (Fig. 6) it is obvious that AS-9121 is responsible for this event, since it gains weight (announces a lot of IPs) and all its neighbour vertices lose weight (lose IPs). This event concerns a clear case of router misconfiguration, in which AS-9121 originated many prefixes that it did not own, thus making a route through AS-9121 more preferable than some of the longer but genuine routes. Although, this event was caused by an error and not a malicious act, the consequences had a global impact, as for several

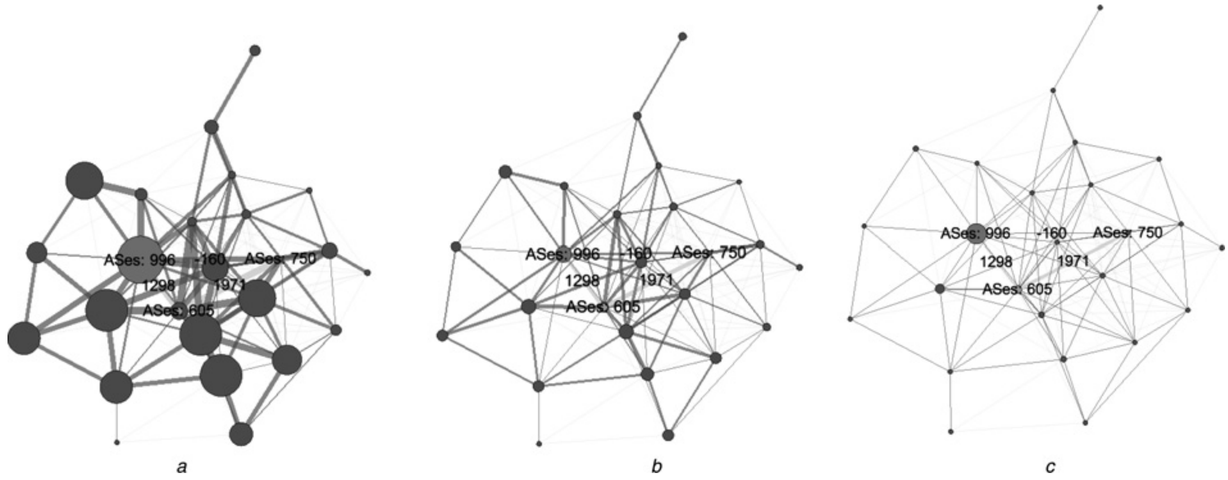

Fig. 5 *Abstract view of level 11 of the clustering hierarchy, visualising the event that took place on Dec 24, 2004 concerning AS-9121. Subfigure (a) is the entropy optimised version*

a $H_G^{\text{out}}(E^c, F') = 3.31$ $H_G^{\text{out}}(V^c, F') = 3.97$
b $H_G^{\text{out}}(E^c, F') = 1.93$ $H_G^{\text{out}}(V^c, F') = 2.62$
c $H_G^{\text{out}}(E^c, F') = 0.67$ $H_G^{\text{out}}(V^c, F') = 0.87$

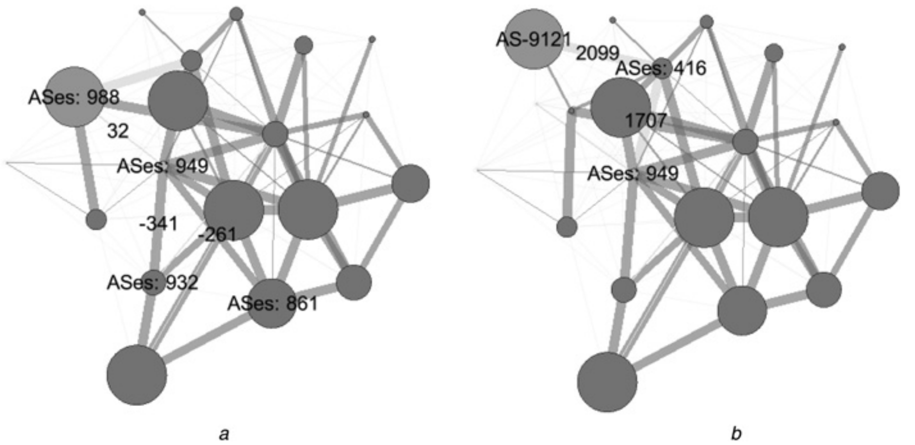

Fig. 7 *Entropy optimised views of level 11 of the clustering hierarchy, visualising the AS-Graph around 19:48 GMT*

a The Green cluster comprised of 988 Ases gains a lot of IPs and paths
b Increasing the granularity of the green cluster reveals AS-9121 as the main actor of this event

hours several of Internet users were unable to reach a large number of Internet sites.

The entropy of the visualisation in Fig. 5*b*, Fig. 5*c* is lower compared with 5*a*. It is apparent from these figures that the lower the entropy, the lower the information content of the visualisation. For example in Fig. 5*c* the user cannot discriminate between many of the edge as well as the vertex weights. On the contrary, in the optimised version of Fig. 5*a* the differences are more apparent and insights about the data are more easily obtained by the user and as a consequence the analytic potential is increased. Even if the event is visible in both the high and the medium entropies, in the high entropy view, the impact of the event to the neighbouring ASes is much clearer.

In addition to the aforementioned misconfiguration event, it is obvious from the anomaly plot depicted in Fig. 4 that there is a second event, captured by the spike around 19:48 GMT. Level-11 of the clustering hierarchy for this time period is depicted in Fig. 7. The green cluster in Fig. 7*a* that is comprised of 988 ASes gains IPs, while all the neighbouring clusters lose IPs. After increasing the granularity of this cluster, AS-9121 is revealed as the only vertex that gains IPs, as depicted in Fig. 7*b*. This means that AS-9121 after almost ten hours of the previous misconfiguration event, announces once again prefixes it does not own. This event is clearly visible using BGPGraph even if it has, compared with the first, a much smaller duration and lower anomaly score, as depicted in the anomaly plot.

8.2 List of detected BGP anomalies

Table 1 provides a list of *small* and *large scale* anomalies, which have been detected with use of the BGPGraph approach and other methods in the literature. This table provides several attributes for each anomaly, such as the time of the event, the responsible ASN, the monitoring points from which the anomalies were detected, and the status of the anomaly (i.e. if the anomaly is confirmed in the literature or detected using the proposed approach). The last column represents the methods that were able to identify the anomaly represented by each row of the table.

As shown in Table 1, concerning the *small scale* events, the proposed approach was able to efficiently detect eleven known prefix hijacking events [1, 4], including the youtube-Pakistan incident [30] which is captured only by BGPGraph. In addition, four new anomalies detected, using BGPGraph and were not previously reported in the literature are also presented in this table.

Furthermore, the analysis of the *large scale* anomalies in Table 1 identified large routing changes, most of which concern prefix hijacks, also detected with the *small scale* analysis. In addition to prefix hijack detection, however, the *large scale* analysis revealed an event or routing change that did not concern prefix hijacking. Specifically, on 21st October 2005 around 06:09 GMT, AS-3356 experienced internal problems, which resulted in losing a lot of prefixes that it previously owned [4], and caused a large number

Table 1 List of small and large scale anomalies detected using the proposed approach

Date (yyyy-MM-dd)	Start time (hh: mm)	Responsible ASN	Monitoring point(s)	Alarm type	Anomaly status	Detection method
2010-04-08	15:54	23,724	3333, 3549, 4608, 4777, 7018	small scale, large scale	confirmed [1]	BGPGraph, [1, 8, 9]
2009-12-15	09:53	39,386	3333, 3549, 4608, 4777	small scale	confirmed [1]	BGPGraph, [1]
2009-08-13	00:34	4800	3333, 3549, 4608, 4777, 7018	small scale	confirmed [1]	BGPGraph, [1]
2009-05-05	16:34	10,834	3333, 3549, 4608, 4777, 7018	small scale	confirmed [1]	BGPGraph, [1]
2009-02-14	11:09	8895	3333, 4608, 4777, 7018	small scale	confirmed [1]	BGPGraph, [1]
2008-12-31	07:30	6849	3333, 3549, 4608, 4777, 7018	small scale, large scale	confirmed [1]	BGPGraph, [1, 8, 9]
2008-08-26	07:18	24,739	3333, 4608, 4777, 7018	small scale	confirmed [1]	BGPGraph, [1]
2008-06-17	14:28	8953	4608, 7018	small scale, large scale	confirmed [1]	BGPGraph, [1]
2008-04-28	09:49	44,237	3333, 7018	small scale	confirmed [1]	BGPGraph, [1]
2008-02-24	19:47	17,557	3549, 3333	small scale	confirmed [30]	BGPGraph
2005-10-21	06:09	3356	333, 3549, 4608, 4777, 7018	large scale	confirmed [4]	BGPGraph, [8, 9]
2004-12-24	09:20	9121	3333, 3549, 4608, 4777, 7018	small scale, large scale	confirmed [4]	BGPGraph, [1, 8, 9]
2010-09-10	04:05	9260	7018	small scale	new anomaly	BGPGraph
2010-05-19	21:15	29,550	3549	small scale	new anomaly	BGPGraph
2009-05-05	00:49	8001	3333	small scale	new anomaly	BGPGraph
2008-09-22	18:06	45,433	4777	small scale	new anomaly	BGPGraph

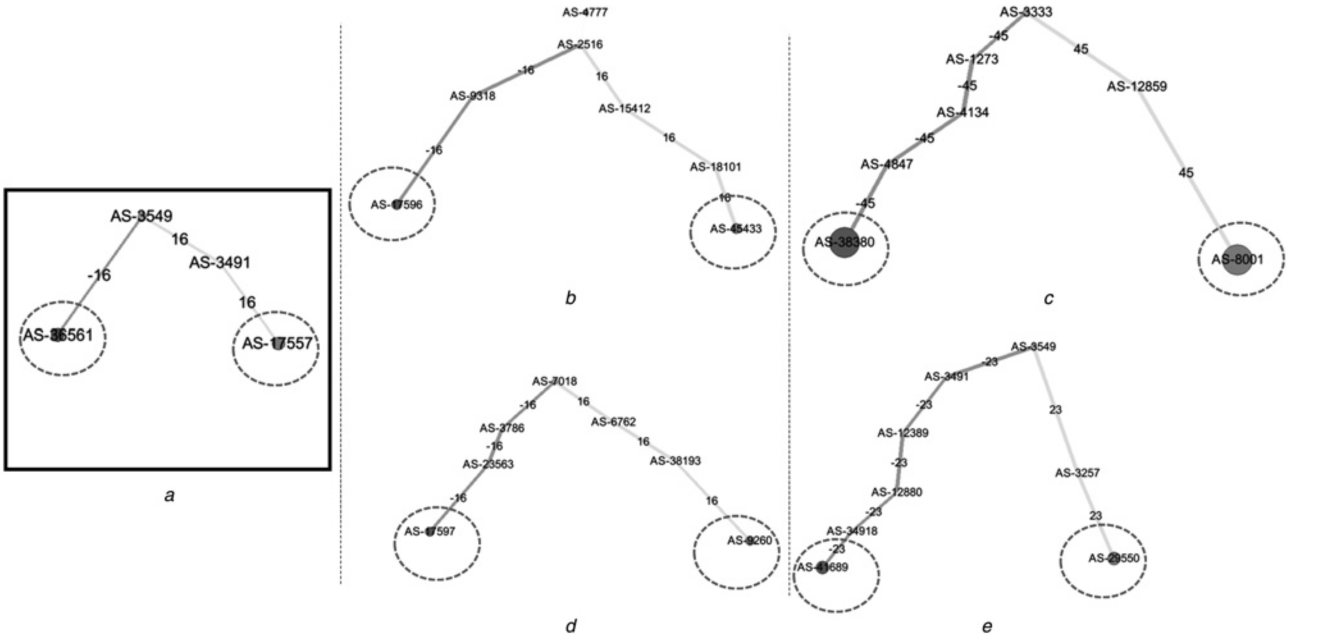

Fig. 8 *Root cause analysis of a known anomaly, and all the new anomalies detected using the proposed BGPGraph approach*

a Youtube-Pakistan incident [30], where AS-17557 (Pakistan) announces prefixes previously belonging to AS-36531 (Youtube)
b New anomaly detected using the proposed approach on 2008-09-22 around 18:06, where AS-45433 announces prefixes previously belonging to AS-17596
c New anomaly detected using the proposed approach on 2009-05-05 around 00:49, where AS-8001 announces prefixes previously belonging to AS-38380
d New anomaly detected using the proposed approach on 2010-09-10 around 04:05, where AS-9260 announces prefixes previously belonging to AS-17597
e New anomaly detected using the proposed approach on 2010-05-19 around 21:15, where AS-29550 announces prefixes previously belonging to AS-41689. The ASes involved in the anomalies are manually highlighted in the red ellipsoids

of BGP withdrawal messages. This event concerns a case that the *small scale* analysis is not able to detect, since there are no MOAS incidents involved.

For the detection of new anomalies, the same procedure presented in Section 8.1 was followed. The anomaly plots for the specific dates were created, and large anomaly scores were identified in the corresponding time instances. This fact indicated the existence of anomalies in these time periods. Further analysis using the proposed methods revealed the responsible ASes for these anomalies. Specifically, Fig. 8 shows multiple anomalies, a known anomaly, and all the new anomalies detected using the proposed BGPGraph approach. Fig. 8*a* illustrates the Youtube-Pakistan incident [30], where AS-17557 (Pakistan) announces prefixes previously belonging to AS-36531 (Youtube). Figs. 8*b*–*e* illustrate the set of new anomalies detected using the proposed approach. The ASes involved in the anomalies are manually highlighted in the red ellipsoids. When compared with the known Youtube-Pakistan, in all the cases the value of the anomaly score is high, while there is also a transfer of traffic from one path to the other, which shows that the responsible ASes announce prefixes previously belonging to other ASes.

Table 2 presents a comparison of the proposed approach with three other recent BGP anomaly detection approaches proposed in the literature: Khare *et al.* [1], Deshpande *et al.* [8], and Al-Rousan *et al.* [9]. The comparison takes place with regards to the (i) Percentage of known *large scale* anomalies detected, and (ii) Percentage of known *small scale* anomalies detected. As shown in this table, the proposed BGPGraph approach is able to identify more anomalies in all the cases.

Table 2 Comparison of the different anomaly detection methods

Method	BGPGraph	Khare *et al.* [1]	Deshpande *et al.* [8]	Al-Rousan *et al.* [9]
Percentage of known *large scale* anomalies detected	100%	80%	80%	80%
Percentage of known *small scale* anomalies detected	100%	90.1%	–	–

9 Conclusions

This paper presented a novel BGP anomaly metric, capable of capturing both *large* and *small scale* BGP anomalies. Anomaly plots of this metric were utilised, so as to provide an overview of the BGP activity over a specific period of time, and help the analyst focus on interesting time windows to perform their analysis by using the BGPGraph visualisation approach.

The BGPGraph utilised a hierarchical graph visualisation scheme to enable the exploration of BGP routing changes. The proposed approach also introduced an information theoretic metric for the quantification and optimisation of the generated visualisations. The hierarchical visualisation provides high level overview of the input graph, while the analyst can further investigate abnormalities in more detail.

The framework is seen to be very efficient in the analysis of BGP routing changes and aids the analyst through the automated estimation of parameters and thresholds that are usually manually selected.

10 Acknowledgments

This work has been partially supported by the European Commission through projects FP7-ICT-257495-VIS-SENSE and FP7-ICT-317888-NEMESYS funded by the 7th framework program. The opinions expressed in this paper are those of the authors and do not necessarily reflect the views of the European Commission.

11 References

1 Khare, V., Ju, Q., Zhang, B.: ‘Concurrent prefix hijacks: Occurrence and impacts’. Proc. of the 2012 ACM Conf. on Internet Measurement Conf., 2012, pp. 29–36

2 Ballani, H., Francis, P., Zhang, X.: ‘A study of prefix hijacking and interception in the internet’. ACM SIGCOMM Computer Communication Review, 2007, vol. 37, no. 4, p. 265

3 Mahajan, R., Wetherall, D., Anderson, T.: ‘Understanding BGP misconfiguration’. ACM SIGCOMM Computer Communication Review, 2002, vol. 32, pp. 3–16

4 Lad, M., Massey, D., Zhang, L.: ‘Visualizing Internet routing changes’, *IEEE Trans. Vis. Comput. Graphics*, 2006, **12**, (6), pp. 1450–1460

5 Butler, K., Farley, T.R., McDaniel, P., *et al.*: ‘A survey of BGP security issues and solutions’, *Proc. IEEE*, 2010, **98**, (1), pp. 100–122

6 Subramanian, L., Roth, V., Stoica, I., *et al.*: ‘Listen and whisper: Security mechanisms for BGP’. Proc. First Symp. on Networked Systems Design and Implementation (NSDI), 2004

7 Kent, S., Lynn, C., Seo, K.: ‘Secure border gateway protocol (S-BGP)’, *IEEE J. Sel. Areas Commun.*, 2000, **18**, (4), pp. 582–592

8 Deshpande, S., Thottan, M., Ho, T.K., *et al.*: ‘An online mechanism for BGP instability detection and analysis’, *IEEE Trans. Comput.*, 2009, **58**, (11), pp. 1470–1484

9 Al-Rousan, N.M., Haeri, S., Trajkovic, L.: ‘Feature selection for classification of BGP anomalies using Bayesian models’. ICMLC, 2012, pp. 140–147

10 Zhang, J., Rexford, J., Feigenbaum, J.: ‘Learning-based anomaly detection in BGP updates’. Proc. of the 2005 ACM SIGCOMM Workshop on Mining Network Data, 2005, pp. 219–220

11 Shiravi, H., Shiravi, A., Ghorbani, A.A.: ‘A survey of visualization systems for network security’, *IEEE Trans. Vis. Comput. Graphics*, 2011, **1**, (1), pp. 1–19

12 Colitti, L., Di, G., Federico, B.: ‘Visualizing interdomain routing with BGPlay’, *J. Graph Algorithms Appl.*, 2005, **9**, (1), pp. 117–148

13 Cortese, P.F., Di Battista, G., Moneta, A., *et al.*: ‘Topographic visualization of prefix propagation in the internet’, *IEEE Trans. Vis. Comput. Graphics*, 2006, **12**, (5), pp. 725–732

14 Wong, T., Jacobson, V., Alaettinoglu, C.: ‘Internet routing anomaly detection and visualization’. 2005 Int. Conf. on Dependable Systems and Networks (DSN’05), 2005, pp. 172–181

15 Chen, M.C.M., Jaenicke, H.: ‘An information-theoretic framework for visualization’, *IEEE Trans. Vis. Comput. Graphics*, 2010, **16**, (6), pp. 1206–1215

16 Yang-Peláez, J., Flowers, W.C.: ‘Information content measures of visual displays’. IEEE InfoVis, 2000, pp. 99–104

17 Theodoridis, G., Tsigkas, O., Tzovaras, D.: ‘A novel unsupervised method for securing BGP against routing hijacks’. Computer and Information Sciences III, 2013, pp. 21–29

18 Ware, C.: ‘Information visualization: Perception for design, vol. 22 of The Morgan Kaufmann series in interactive technologies / ed. Stuart Card. – San Francisco: Morgan Kaufmann’ (Morgan Kaufmann, 2004)

19 Smeets, J.B.J., Brenner, E.: ‘Grasping Weber’s law’, 2008

20 Van Hateren, J.H.: ‘Spatiotemporal contrast sensitivity of early vision’, *Vis. Res.*, 1993, **33**, (2), pp. 257–267

21 Cheng, I., Shen, R., Yang, X.-D., *et al.*: ‘Perceptual analysis of level-of-detail: The jnd approach’. Eighth IEEE Int. Symp. on Multimedia, 2006. ISM’06., 2006, pp. 533–540

22 Cheng, I., Boulanger, P.: ‘A 3D perceptual metric using just-noticeable-difference’. Proc. of Eurographics, 2005, pp. 97–100

23 Green, M.: ‘Toward a perceptual science of multidimensional data visualization: Bertin and beyond’, *ERGO/GERO Hum. Factors Sc.*, 1998, **8**

24 Gansner, E.R., Koren, Y., North, S.C.: ‘Topological fisheye views for visualizing large graphs’, *IEEE Trans. Vis. Comput. Graphics*, 2005, **11**, (4), pp. 457–468

25 Ziegel, E., Press, W., Flannery, B., *et al.*: ‘Numerical recipes in C: The art of scientific computing’ (Cambridge University Press, 1987), vol. 29

26 Li, Y., Xing, H.-J., Hua, Q., *et al.*: ‘Classification of BGP anomalies using decision trees and fuzzy rough sets’. 2014 IEEE Int. Conf. on Systems, Man and Cybernetics (SMC), 2014, pp. 1312–1317

27 Li, J., Brooks, S.: ‘I-seismograph: Observing and measuring Internet earthquakes’. 2011 Proc. IEEE INFOCOM, , 2011, pp. 2624–2632

28 Wang, Y., Wang, Z., Zhang, L., *et al.*: ‘Situation assessment model for inter-domain routing system’, *IET Softw.*, 2014, **8**, (2), pp. 53–61

29 RIPE Network Coordination Centre: ‘Routing information service project (RIS)’. Available at http://www.ripe.net

30 RIPE: ‘YouTube Hijacking’, 2014. Available at http://www.ripe.net/internet-coordination/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study