

Feature Extraction and Visual Feature Fusion for the Detection of Concurrent Prefix Hijacks^{*}

Stavros Papadopoulos^{1,2}, Konstantinos Votis², Christos Alexakos³,
and Dimitrios Tzovaras²

¹ Department of Electrical and Electronic Engineering, Imperial College London,
SW7 2AZ, London, UK

s.papadopoulos11@imperial.ac.uk

² Information Technologies Institute, Centre for Research and Technology Hellas,
P.O. Box 361, 57001 Thessaloniki, Greece

{spap,kvotis,tzovaras}@iti.gr

³ Pattern Recognition Laboratory, Computer Engineering and Informatics,
University of Patras, Patras, Greece

alexakos@ceid.upatras.gr

Abstract. This paper presents a method for visualizing and analyzing Multiple Origin Autonomous System (MOAS) incidents on Border Gateway Protocol (BGP), for the purpose of detecting concurrent prefix hijack. Concurrent prefix hijacks happen when an unauthorized network originates prefixes that belong to multiple other networks. Towards the goal of accurately identifying such events, multiple features are extracted from the BGP records and visualized using parallel coordinates enhanced with visual querying capabilities. The proposed visual queries enable the analyst to select a significant subset of the initial dataset for further analysis, based on the values of multiple features. This procedure allows for the efficient visual fusion of the proposed features and the accurate identification of prefix hijacks. Most of the previous approaches on BGP hijack detection depend on static methods in order to fuse the information from multiple features and identify anomalies. The proposed visual feature fusion, however, allows the human operator to incorporate his expert knowledge into the analysis, so as to dynamically investigate the observed events, and accurately identify anomalies. The efficiency of the proposed approach is demonstrated on state-of-the-art BGP events.

1 Introduction

Multiple Origin Autonomous System (MOAS) [1] conflicts occur when a particular prefix appears to originate from more than one AS (Autonomous System).

^{*} This work has been partially supported by the European Commission through the project 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.

Although MOAS incidents occur often for legitimate reasons (e.g. multi-homing, exchange point addresses etc[1]), they might also be the result of a fault or an attack, where a Border Gateway Protocol (BGP) router falsely originates prefixes belonging to other organizations. The detection of such faults and attacks has been the main focus of the research community in the last years. In this paper, the main focus is the detection of concurrent prefix hijacks, in which an unauthorized network originates prefixes that belong to multiple other networks, usually due to router misconfigurations.

The proposed methods for prefix hijack identification are separated into two categories: authentication and detection. Authentication schemes [2][3] depend on the a priori knowledge of the prefix owners, while detection schemes make no such assumption but instead analyze the routing behavior in order to identify suspicious events. The approach presented in this paper falls into the second category.

The detection schemes depend on some type of feature extraction and feature fusion in order to detect BGP prefix hijacks. Although there are many well known techniques for algorithmic fusion [4], the result of such approaches is static and it does not capture the dynamic nature of the modern multivariate systems and the corresponding alterations of the features' cross-correlations. For this reason, visualization techniques can be utilized for the purpose of performing the feature fusion, and allow the user to incorporate his experience and knowledge into the analysis.

In this respect, this paper introduces a novel visual feature fusion method for the combination of different BGP features, extracted from the observed MOAS incidents. The ultimate goal of the proposed approach is the detection of concurrent prefix hijacks, in which an unauthorized network originates prefixes that belong to multiple other networks. In order to achieve this, visual queries are applied on the parallel coordinates visualization in order to enable the analyst to focus on a significant subset of the data for further analysis, based on the values of multiple features.

It should be noted that the BGP exhibits continuous alterations, with over 20,000 BGP messages exchanged every 5 minutes (in 2014). All these messages are analyzed using the proposed approach in order to firstly identify MOAS incidents, and afterwards extract meaningful features that will enable the anomaly detection procedure. The proposed approach has been applied on data collected for over 3 months with out any problems. But in this paper and without loss of generality, the demonstration is carried out with respect to data corresponding to a time period of 5 days.

The rest of the paper is organized as follows: Section 2 presents the related work. Section 3 presents the feature extraction methodology, while the proposed visual feature fusion method is detailed in section 4. Use case analysis takes place in section 5. Finally, the paper concludes in section 6.

2 Related Work

The internet uses BGP as the defacto protocol for the exchange of routing information between Autonomous Systems (AS). An AS represents a collection of Internet Protocol prefixes under the control of one network operator. The distributed nature of BGP and the lack of security mechanisms, render it vulnerable to various types of attacks, as for example prefix hijacking or Man- In-the-Middle attacks [5]. Towards this end, the research community has focused its efforts on the development of methods that could enable the detection of such anomalies in the BGP infrastructure.

Specifically, Deshpande et al. [6] introduced multiple features that represent various aspects of the BGP update messages. A Generalized Likelihood Ratio (GLR) based hypothesis test is utilized onto each feature in order to detect time periods of instabilities. Majority voting is afterwards utilized in order to correlate the detected instabilities, and generate alerts for which more than half of the features indicate an anomaly. Unlike the proposed visual feature fusion method, the disadvantage of this approach is that the analyst is not able to change the parameters of the anomaly detection algorithm in order to get better detection rates.

Li et al. [7] propose the use of a signature based classifier, that is trained to recognize certain types of behaviors specific to BGP anomalies. The classifier is comprised of a collection of IF-THEN threshold based rules, which are applied onto multiple BGP features, extracted in different time windows. As with the aforementioned approach the anomaly results are also static, as the analyst is left out of the detection procedure.

Al-Rousan et al. [8] introduce multiple features that characterize the BGP activity. Feature selection methods are applied onto this set of features in order to select a specific subset that better characterizes BGP anomalies. Afterwards, Naive Bayes classifiers are used for the detection of BGP anomalies, including worm attacks and router misconfigurations. In comparison with the proposed approach the user is not able to change the parameters of the classification in order to find alternative or better results.

Zhang et al. [9] propose the use of signature and a statistical based methods for detecting anomalies in BGP. The signature based methods are comprised of a standard set of patterns in the BGP update messages, which are specific to BGP anomalies. The statistical based methods are applied onto multiple BGP features, in order to acquire an anomaly score for each one of them. A linear weighted sum of the scores of each feature is afterwards utilized, so as to fuse the features and acquire a single anomaly score. This procedure is repeated for each time window under investigation, in order to detect time periods of instability. The results of this work were utilized by Teoh et al. [10], and combined with visualization methods. Specifically, the anomaly scores of each feature, as well as the global anomaly score are visualized along the BGP update messages selected by the analyst. Using the feedback provided by the visualization, the analyst is able to adjust the parameters of the statistical and signature based methods and increase the detection accuracy. The adjustable parameters, however, do not

refer to the feature fusion procedure but rather to the definitions of anomalous BGP events and the threshold of the fused anomaly metric. The approach of including the analyst into the analytical procedure is related to the present work, which also incorporates visual methods to aid the analysis procedure. But unlike [10], the present work utilizes the visualization to aid the visual feature fusion procedure.

3 MOAS Feature Extraction

This section presents the feature extraction methodology proposed in the context of this paper. The extracted features are able to quantify the degree of anomaly of each MOAS incident, while their fusion provides a complete view of the MOAS characteristics and enables the detection of concurrent prefix hijacks.

There are four features that are extracted from the MOAS events that occurred over a specific time period: 1) *Country MOAS Probability*, 2) *AS Relationship*, 3) *Ownership duration*, and 4) *MOAS duration*. Each one of them is described in detail in the paragraphs that follow.

The first feature is *Country MOAS Probability*, which represents the probability of occurrence of a MOAS event from the country of the new owner AS to the country of the old owner AS. This country based aggregation and analysis of the BGP activity, has been shown to capture the underlying geo-spatial coherence of the Internet routing information [11]. Although the probability of occurrence of a MOAS event between two ASes can be used directly for anomaly detection, the country aggregation takes into account all the MOAS events between the two countries, and thus holds additional semantic information regarding the geo-spatial distribution of the MOAS incidents. This fact leads to a more accurate identification of concurrent prefix hijacks, since most of them target countries that are usually not targeted in normal behavior.

The *AS Relationship* between the two ASes involved in the MOAS incident has been used as a heuristic for reducing the number of false positives in prefix hijack detection [12]. Under this consideration, it is also utilized in this context for the generation of a binary feature. This feature captures the AS pairs that have commercial relationships, and thus might use prefixes from the same set. The AS relationships were defined by Gao [13] and are comprised of Provider-To-Customer, Peer-To-Peer, and Sibling-To-Sibling relationships. If any type of relationship exists between two ASes, the *AS Relationship* feature has value 1, else it has value 0.

The *Ownership duration* feature captures the duration in hours that the new owner has been announcing the prefix involved in the MOAS incident. Since the new owner AS involved in the concurrent prefix hijacks is not the legitimate owner of the prefix, it is highly unlikely that it has been announcing the prefix for a long period. Thus, this feature captures this behavior and provides additional information for the accurate identification of concurrent prefix hijacks. It should be noted that absolute thresholds for anomaly detection using the aforementioned feature have been used in the literature [12], but in the proposed

Table 1. Overview of the MOAS features and metadata used for the identification of concurrent prefix hijacks

		Description
Features	Country MOAS Probability	The probability of a MOAS occurrence from the country of the Attacking AS to the country of the Victim AS
	AS Relationship	Captures the AS that have direct commercial relationships.
	Ownership duration	The duration in hours that the Attacking AS has been announcing the corresponding prefix
	MOAS duration	The duration in hours of the MOAS incident between the two ASes
Metadata	Attacking AS	The AS number of the new owner of the prefix
	Victim AS	The AS number of the old owner of the prefix
	Time	The exact date and time that the MOAS event occurred

approach combines the actual value of this feature with multiple other features, and allows the analyst to select appropriate thresholds dynamically, using visual queries.

The *MOAS duration* feature captures the duration in which each MOAS event is active. In the case of concurrent prefix hijacks, the MOAS incidents are highly unlikely to last for a long period. After the misconfiguration has been fixed, the prefixes return to their original owners. Although the MOAS duration can be a useful heuristic to distinguish between valid MOAS conflicts and invalid ones[1], such differentiation on this feature alone can not be accurate enough to be a solution prefix hijack identification. In this context, the visual feature fusion procedure proposed in the context of this paper combines the values of this feature with multiple other features, in order to provide a holistic view of the MOAS activity to the analyst.

In addition to the aforementioned features that characterize the degree of anomaly of each MOAS incident, the MOAS events are also characterized by metadata. These metadata provide additional semantic information about each event and are used for root cause analysis by the visualization. These metadata include the following: *Attacking AS*, i.e. the AS number of the new owner of the prefix, *Victim AS*, i.e. the AS number of the old owner of the prefix, and *Time*, i.e. the exact date and time that the MOAS event occurred.

An overview of the extracted features and the metadata that are utilized for visual feature fusion and prefix hijack detection is presented in Table 1.

4 Visual Feature Fusion Using Parallel Coordinates and Visual Queries

The proposed visual feature fusion scheme is detailed in this section. Specifically, the features and metadata describing each MOAS event, which were presented in section 3, are visualized using parallel coordinates. The parallel coordinates are enhanced with filtering capabilities in order to enable the analyst to detect significant MOAS incidents that might contain anomalies.

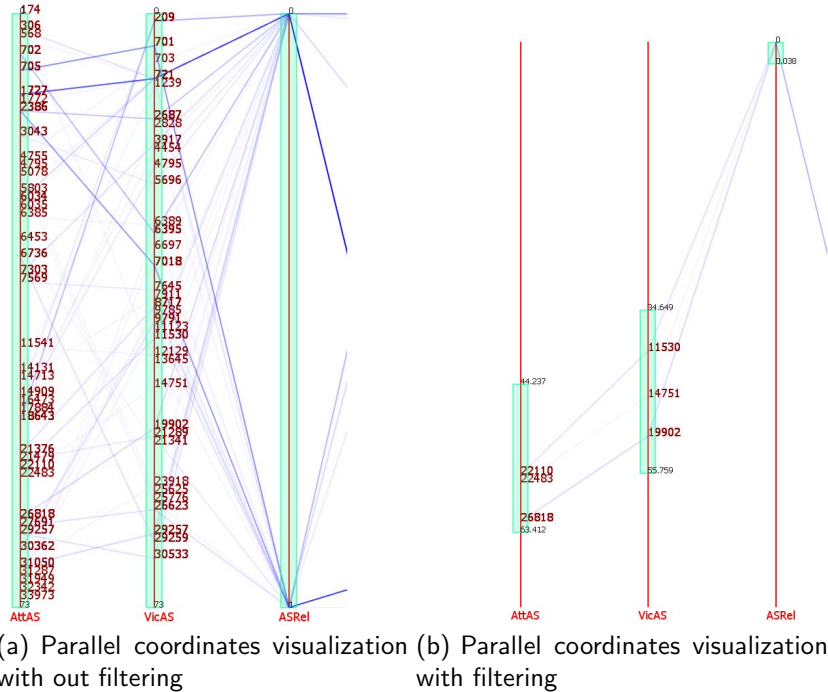


Fig. 1. The application of visual queries on the parallel coordinates utilized for visual feature fusion and the selection of significant data subsets

Specifically, bar sliders are attached onto each parallel coordinate that specify the allowed upper and lower bounds of each feature. These bar sliders can change size and position and are directly manipulated by the user of the visualization. This procedure allows the analyst to dynamically change the bounds of the features and as a result the MOAS events that are visualized. The direct feedback provided by the visualization with regards to the selection of MOAS events, enables the analyst to redefine the bounds, and find the best parameters in each case, according to the underling information.

Thus, for each feature and metadata f_i the bounds are defined as $B_i = \{f_i^{low}, f_i^{upper}\}$. Each MOAS event is characterized by its value in all the features and metadata. Furthermore, the following function is defined:

$$g(e_j) = \begin{cases} true, & \text{if } f_i^{low} \leq f_i(e_j) \leq f_i^{high}, \forall i \\ false, & \text{else} \end{cases} \quad (1)$$

where $f_i(e_j)$ is the value of the j_{th} event e_j in the i_{th} feature f_i . If $g(e_j)$ is *true* then the MOAS event e_j is visualized, while in the case that is *false* it is not.

This interactive procedure enables for the visual fusion of the values of multiple features, in order to focus on suspicious MOAS incidents and lead to the accurate identification of concurrent prefix hijacks.

5 Discovery and Analysis of Prefix Hijacks

This section demonstrates the use of the proposed feature extraction and visual feature fusion approach for the identification of concurrent prefix hijacks.




Fig. 2. Visual feature fusion for the detection of a router misconfiguration event which tool place on 08-April-2010[12]. The *AS Relationship* feature is set to 0 (no relationship exists), the *Country MOAS Probability* is set to be smaller than 6%, the *Ownership duration* feature is set to be smaller than 2 hours, and the *MOAS duration* feature is set to be smaller than 24 hours. After this filtering, AS-23724 is still involved in multiple MOAS events, occurring simultaneously (at around 16:00 GMT)

The BGP data were collected and analyzed from the RIPE routing monitoring service [14]. Specifically, the monitoring point of AS-3333 was utilized for the collection of the routing updates, and the analysis of the observed MOAS incidents.

It should be underlined that the monitoring period is larger than the period used for the visualization. The reason for this is that the values of the features are relevant to the total monitoring period, and small monitoring periods might not be enough for their accurate calculation. In both use cases examined in this section, the visualization period is set to one day, while the monitoring period is set to five days, two days before and two days after the visualization period.

The first use case involves a concurrent prefix hijack that occurred on 08-April-2010, around 16:00 GMT. In this incident, AS-23724 announced multiple prefixes it did not own, belonging to multiple other ASes[12]. The monitoring period is set from 06-April-2010 to 10-April-2010, while the visualization period is set to 08-April-2010. The result of the visual feature fusion is depicted in figure 2. The values of the *AS Relationship*, *Country MOAS Probability*, *Ownership duration*, and *MOAS duration* are set to be: 0 (no relationship exists), smaller than 6%, smaller than 2 hours, and smaller than 24 hours, respectively. After the application of these visual queries, there is a small set of MOAS incidents that remains. It is apparent that in the first parallel coordinate, which

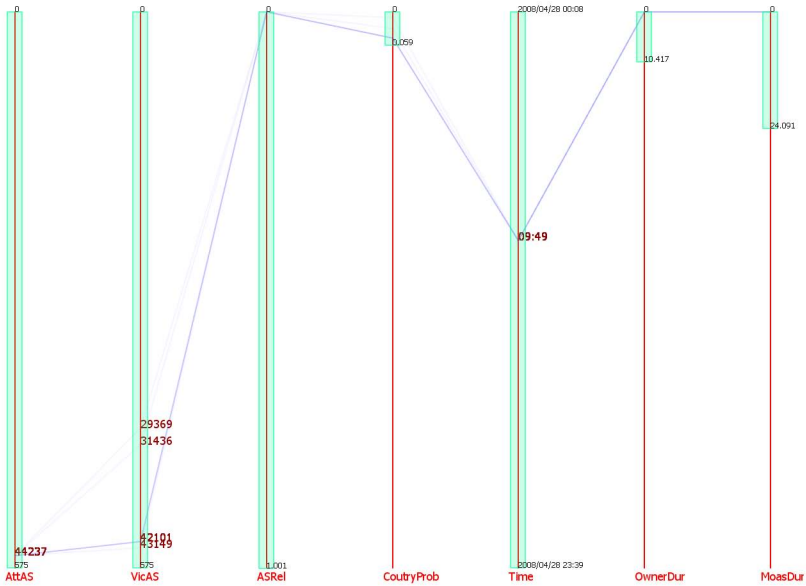


Fig. 3. Visual feature fusion for the detection of a router misconfiguration event which tool place on 28-April-2008[12]. The *Country MOAS Probability* is set to be smaller than 6%, the *Ownership duration* feature is set to be smaller than 10 hours, and the *MOAS duration* feature is set to be smaller than 24 hours. AS-44237 is involved in four suspicious MOAS events visible from the monitoring point of AS-3333.

represents the Attacking AS metadata, that AS-23724 is involved in multiple suspicious MOAS incidents, which last for a short period and have low probability of appearance. Furthermore, all these MOAS events are occurring almost simultaneously, around 16:00 GMT.

It should be noted that after the filtering procedure has taken place, there are two more MOAS incidents that remain. These incidents are not concurrent, since they do not concern multiple ASes. These events are short lived in the monitoring period and also have low probability of appearance. But due to the nature of BGP and the lack of ground truth, the answer to the legitimacy of these events is difficult and out of the scope of this paper, which focuses on the detection of concurrent prefix hijacks.

The second use case concerns a concurrent prefix hijack that occurred on 28-April-2008 by AS-44237[12]. The monitoring period is set from 26-April-2008 to 30-April-2008, while the visualization period is set to 28-April-2008. The result of the visual fusion is depicted in figure 3. The bounds of the features are set using visual queries to: *Country MOAS Probability* smaller than 6%, *Ownership duration* smaller than 10 hours, and *MOAS duration* feature smaller than 24 hours. It is apparent in this figure that AS-44237 is involved in four suspicious MOAS events, which are visible from the monitoring point of AS-3333. All these MOAS incidents occurred on 28-April-2008 around 9:49 GMT.

6 Conclusions

The present paper proposed a method for visualizing and exploring BGP MOAS incidents, in order to detect concurrent prefix hijacks. Multiple features are extracted from the detected MOAS incidents, which quantify their degree of anomaly. An important aspect of the proposed features is the fact that they present complex cross-correlations, and they have to be combined within the context of each other, as well as within the context of the other MOAS events. Towards this goal, this paper proposed the use of visual feature fusion of the aforementioned features, in order to allow for the accurate identification of concurrent prefix hijacks. To achieve this, a parallel coordinates view is proposed, which is further enhanced with filtering capabilities so as to discriminate between abnormal and normal events. This procedure allows the analyst to incorporate his/her expert knowledge and steer the analysis dynamically, so as to effectively detect prefix hijacks. The proposed was applied into a real-world BGP events, in order demonstrate its analytical potential.

It should be noted, that the proposed approach is general and can be applied in any set of relevant BGP features. Furthermore, it is also scalable, which renders it able to visualize and fuse a relatively large number of features by utilizing the power of parallel coordinates.

References

1. Zhao, X., Pei, D., Wang, L., Massey, D., Mankin, A., Wu, S.F., Zhang, L.: An Analysis of BGP Multiple Origin AS (MOAS) Conflicts. In: SIGCOMM Internet Measurement Workshop, p. 31. ACM Press (2001)
2. Zheng, C., Ji, L., Pei, D., Wang, J., Francis, P.: A light-weight distributed scheme for detecting ip prefix hijacks in real-time. *ACM SIGCOMM Computer Communication Review* 37(4), 277 (2007)
3. Lad, M., Massey, D., Pei, D., Wu, Y., Zhang, B., Zhang, L.: PHAS: A prefix hijack alert system. In: Proc. USENIX Security Symposium, vol. 2, pp. 153–166 (2006)
4. Chowdhury, P., Das, S., Samanta, S., Mangai, U.: A Survey of Decision Fusion and Feature Fusion Strategies for Pattern Classification. *IETE Technical Review* 27(4), 293–307 (2010)
5. Ballani, H., Francis, P., Zhang, X.: A study of prefix hijacking and interception in the internet. *ACM SIGCOMM Computer Communication Review* 37(4), 265 (2007)
6. Deshpande, S., Thottan, M., Ho, T.K., Sikdar, B.: An online mechanism for BGP instability detection and analysis. *IEEE Transactions on Computers* 58(11), 1470–1484 (2009)
7. Li, J., Dou, D., Wu, Z., Kim, S., Agarwal, V.: An Internet routing forensics framework for discovering rules of abnormal BGP events. *ACM SIGCOMM Computer Communication Review* 35(5), 55–66 (2005)
8. Al-Rousan, N.M., Haeri, S., Trajkovic, L.: Feature selection for classification of BGP anomalies using Bayesian models. In: ICMLC, pp. 140–147 (2012)
9. Zhang, K., Yen, A., Zhao, X., Massey, D., Wu, S.F., Zhang, L.: On detection of anomalous routing dynamics in BGP. In: Mitrou, N.M., Kontovasilis, K., Rouskas, G.N., Iliadis, I., Merakos, L. (eds.) NETWORKING 2004. LNCS, vol. 3042, pp. 259–270. Springer, Heidelberg (2004)
10. Teoh, S.T., Zhang, K., Tseng, S.-M., Ma, K.-L., Wu, S.F.: Combining visual and automated data mining for near-real-time anomaly detection and analysis in BGP. In: Proceedings of the 2004 ACM Workshop on Visualization and Data Mining for Computer Security, VizSECDMSEC 2004, p. 35 (2004)
11. Theodoridis, G., Tsigkas, O., Tzovaras, D.: A Novel Unsupervised Method for Securing BGP Against Routing Hijacks. In: Computer and Information Sciences III, pp. 21–29. Springer (2013)
12. Khare, V., Ju, Q., Zhang, B.: Concurrent prefix hijacks: Occurrence and impacts. In: Proceedings of the 2012 ACM Conference on Internet Measurement Conference, pp. 29–36. ACM (2012)
13. Gao, L.G.L.: On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking* 9(6), 733–745 (2001)
14. RIPE Network Coordination Centre, Routing Information Service project (RIS), <http://www.ripe.net>