An Anomaly Detection Framework for BGP

Abstract—**Abnormal events such as large scale power outages, misconfigurations, and worm attacks can affect the global routing infrastructure and consequently create regional or global internet service interruptions. As a result, early detection of abnormal events is of critical importance. In this study we present an anomaly detection framework based on data mining algorithms that is applied to anomaly detection on global routing infrastructure. To show the applicability of our framework, we conduct extensive experiments with a variety of abnormal events and classification algorithms. Our results demonstrate that when we train our system ample variety of abnormal events including worm attacks, power supply outages, submarine cable cuts, and misconfigurations such as network prefix hijacks or routing table leaks, we can detect a similar type of event that is happened later.**

Keywords: BGP, anomaly detection, Data mining, machine learning, IRF

I. INTRODUCTION

Current research [1] on Internet growth trends indicate several factors that will lead to project an even more unprecedented growth such as massive Internet penetration in markets belonging to developing countries, mainly driven by China and India and the substantial increase of 3-4G mobile broadband subscribers. These factors represent a challenge requiring interdisciplinary efforts that guarantee scalable and stable solutions to keep Internet services available. Detection of abnormalities in the routing framework of the internet is one of them and this is the focus of our study.

Routing in the Internet is divided into two well-defined domains. First domain consist of a fine-grained topological detail of connected segments of the Internet. Second domain includes wider scope interconnection of segments under the same routing policy, known as autonomous systems (ASs), managed by an inter-domain routing protocol, such as Border Gateway Protocol (BGP) [2]. BGP has allowed the Internet to effectively become a truly decentralized system. Routers send BGP messages between each other in order to exchange route and prefix visibility and reachability.

Abnormal events such as large scale power outages, misconfigurations, and worm attacks can affect the global routing infrastructure and consequently create regional or global internet service interruptions. As a result, early detection of abnormal events is of critical importance. In this study we present an anomaly detection framework based on data mining algorithms that is applied to BGP anomaly detection problem. The framework consists of two parts. First part is an advanced feature extraction system which extracts a wide range of features from a stream of BGP messages. Second part consists of learning and classification abnormal events using data mining / machine learning algorithms augmented by intelligent normalization and a sliding window approach. Although our anomaly detection framework is demonstrated on BGP anomaly detection, it can be used for detection of abnormalities on similar domains as long as aggregate features can be obtained from a stream of entities. Additionally, our framework can be used for real time detection and characterization of abnormal events. We report the results of experiments on a wide range of abnormal BGP events including worm attacks, power supply outages, submarine cable cuts, and misconfigurations such as routing table leaks.

II. BACKGROUND AND RELATED WORK

A. *Abnormal BGP Events*

There are several different types of abnormal events which affect interdomain routing infrastructure. Some of these events such as the Mediterranean Sea cable cut [3] and the slammer worm attack cause global service interruptions on internet. Events like the Moscow power blackout, or the Turk Telecom (TTNet) table leak, on the other hand, have limited regional influence. One of the most famous BGP anomalies that cause global service interruptions is Slammer worm attack. Slammer worm is a well understood event that has been subject to thorough analysis [4] in recent years. Slammer worm exhibited an aggressive propagation behavior which had a severe infrastructural impact. As can be seen in Figure 1, Slammer worm caused a massive spike of BGP traffic that lead to unreachable prefixes due to congestion in many of the main Internet backbone links. Nimda worm, however, had different properties as it aimed a longer control time over the infected hosts. Consequently, the BGP footprint of Nimda's design was reduced in comparison with Slammer's.

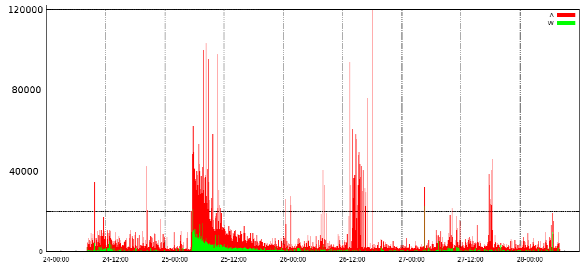

Figure 1. Announcement (A) and withdrawal (W) messages on BGP traffic during Slammer worm attack

As of 2010, the Internet routing is still built on top of a trust model in which network operators from Internet Service

Providers (ISP) and telecom companies do not always implement the currently available tools (RPKI, IRR filters, any of the BGP securing mechanisms [5]) to validate the legitimate ownership of network prefixes. Misconfigurations can take place either unwillingly or deliberately and come in different shape and forms. Amongst these routing table leaks and prefix hijacks are the most common examples network operators face in a more or less frequent basis. These types of events can be summarized as follows:

1) A "table leak" occurs when an ISP incorrectly announces a varying degree of prefixes learnt from other peers that it actually does not own. This brings general routing instability as well as unreachable prefixes. One of the examples of this type is AS9121 leak which is also known as TTnet. During this event at 24 December 2004, 100K+ routes leaked resulting misdirected or lost traffic for tens of thousands of networks.[1]

2) Prefix hijacks take place when a given organization injects a new BGP network prefix announcement impersonating the actual owner of that particular range of IP addresses, effectively becoming the destination of an address space not belonging to it.

The power outages and cable cuts also effect networking infrastructure. They are similar type of events since they both ceases network prefixes become abruptly withdrawn and physically isolated. This isolation is more visible in blackouts. This was the case of the Moscow blackout, where every member of Moscow Internet Exchange (MSK-IX) where disconnected, critically damaging the European and Asia Pacific region. Cable cuts [6], on the contrary, disable critical egress links of a given region and even though traffic can be rerouted, associated network load for such redistribution may create service disruptions.

The Internet has witnessed many abnormal events since it became openly interconnected and operational. It has only been during this last decade or so, though, that Regional Internet Registries have developed the tools and techniques to make comprehensive datasets publicly available for the research community. One example of such datasets is RIPE NCC's Routing Information Service [7]. The RIS collects BGP routing information messages in near real time from 600 peers from 16 collection boxes spread globally, called "Remote Route Collectors". Along with these updates, RRCs also take and store snapshots of their routing tables three times a day.

This work is based on the BGP updates supplied by the nearest RRC box to the abnormal BGP event being researched, based on the assumption that the attributes extracted per event are more genuine and reflect more reliably the behavior we are modeling as we gather data from the nearest possible location

B. BGP Anomaly Detection Systems

There has been an increasing interest in BGP anomaly detection and as a result several different types of anomaly detection systems are proposed for BGP traffic. A review of these systems can be found in [8]. There are several methods proposed for detection of abnormalities in BGP traffic including signature based and statistics based detection approaches, wavelets, k-means clustering [9,10,11]. In another approach, they characterized three abnormal BGP events (Slammer worm, Witty worm and 2003 East-coast blackout), by using attributes extracted from BGP messages in three-second bins [12]. Based on this work, in [13] an higher-order collective classifier is proposed. They compare their system with standard Naive Bayes technique for masquerade detection and report that their technique performs better. In [14] they used statistical pattern recognition techniques to detect instabilities and tested their system on several different types of events such as misconfiguration, node failures, and several worm attacks. One of the most related type detection systems is proposed in [15]. They employ decision tree, a data mining algorithm to detect internet routing anomalies by using attributes derived from BGP traffic. These attributes are basically the counts of different types of BGP messages divided into one minute bins.

Our approach differs from these works by learning to detect different types of abnormal events from normal traffic and the diversity of anomalous events we use in our experiments as well as the data mining algorithms we employ.

III. APPROACH

Our approach is similar to [15]. We use data mining algorithms to learn from labeled data (e.g. abnormal event traffic) and apply the learned model to the previously unseen BGP event traffic by using a sliding window approach. If the number of bins matching the model exceeds a certain threshold we raise an alarm. However, our framework is much more flexible so that we can use several different classification algorithms that are available in WEKA data mining software [16]. Furthermore, we use varying bin size and sliding window size to achieve best results. During our experiments we have trained our classifiers using one abnormal event and applied learned model to a different abnormal event. For all events we have extracted datasets that are consist of pre-event period bins and event period bins. We select 480 minutes before event and 240 minutes after the event started. In an event dataset, for an attribute, we normalized the values by dividing them the average of normal instances only. As a result attribute values of normal bins are close to one but for event bins they are much larger. This novel normalization approach allowed us to generalize classification models across wide variety of different events. This is important because these events occurs on different date and times in a ten year period and during this period BGP traffic has increased drastically. Additionally, event in a week period of BGP traffic exhibits different patterns in different hours and in different days of the week. Although BGP traffic pattern is dynamic and fluctuates, when an abnormal event occurs, the pattern changes significantly and we can see spikes on certain statistics. These strong differences in traffic allow our algorithms to detect them. We use decision tree (J48) [17], Naïve Bayes [18], and Support Vector Machines (SVM) [19] algorithms that are implemented in WEKA. Our results show that SVM performs best among these when used with polynomial kernel. On the other hand setting the kernel type to Radial Basis Function or Linear generated

[1] http://www.renesys.com/tech/presentations/pdf/renesys-nanog34.pdf

poor results. J48 performs well for some of the events but fail to detect some of the events that could be detected using SVM. Therefore for all the results reported hereafter, we use SVM as the classification algorithm.

Our anomaly detection implementation consists of two parts. First part is preprocessing which takes BGP update feeds as input and extract numerical features for a certain period of time. This certain amount of time is called a bin and it corresponds to an instance in a machine learning dataset. The bin size is parametric and it can be adjusted based on the needs of application domain. This preprocessing system need to process large number of messages that is in textual format therefore it is implemented in perl. This feature extraction part is further explained in the following section.

IV. EXPERIMENT SETUP

We have used wide variety of abnormal events which can be categorized in three types of events such as blackouts, cable cuts and worm attacks. A list of these events can be found in Table 1. We have extracted 30 second and 60 second (one minute) bins. Our experiments do not show significant performance differences between these different bin sizes therefore we choose 30 second bins for our system. This choice will allow us to detect events earlier and increase the applicability of our framework as a real time anomaly detection system. As sliding window size we choose ten bins and it moves forward four bins. These parameters are determined experimentally. Since we use 30 second bins our sliding window corresponds to five minutes of BGP traffic and it slides two minutes every time. If more than 60% of the bins are classified as abnormal event our system gives an alert. This corresponds to six instances in our sliding window. Theoretically, our system allows us to detect an abnormal event as early as four minutes.

TABLE I. LIST OF ABNORMAL EVENTS

Event name	*RRC*	*Date*
Nimda worm	rrc04, Geneva	Sept., 2001
Slammer worm	Routeviews, Oregon	Jan., 2003
East-coast blackout	Routeviews, Oregon	Aug., 2003
TTNET routing table leak	rrc05, Vienna	Dec., 2004
Moscow blackout	rrc05, Vienna	May, 2005
Luzon cable cut	rrc06, Otemachi	Dec., 2006
Mediterranean sea cable cuts	rrc10, Milan	Jan., 2008

A. Feature Extraction

We have developed a software solution in perl to pre-process the data and extract numerical features from BGP traffic. The methodology followed by our software consists of sequential parsing of BGP update feeds while certain statistical features are tracked and calculated. Per every time bin of a tunable length in seconds, the script extracts a fixed set of rich attributes that are meant to model the abnormal event and subsequently train the data mining algorithms. Some of these attributes belong to the classical set in [12,15]. On the other hand, some attributes introduced in this study such as concentration ratios and they are novel. To the best of our knowledge, no previous study has used them before. Table II provides a summary of the features.

TABLE II. LIST OF FEATURES

Id	*Definition*
1	Number of announcements
2	Number of withdrawals
3	Number of updates
4	Number of announced prefixes
5	Number of withdrawn prefixes
6	Max. announcements per prefix
7	Avg. announcements per prefix
8	Maximum AS path length
9	Average AS path length
10	Maximum unique AS path length
11	Average unique AS path length
12	Announcements to longer path
13	Announcements to shorter path
14	Concentration ratio
15	First order concentration ratio
16	Second order concentration ratio
15	Third order concentration ratio

For reasons detailed previously, the number of announcements and withdrawals exchanged by neighboring peers are an important feature always reproduced during instability periods. Hence, it is considered an effective yet simple feature that accurately spots bins differing from normal trends.

V. RESULTS

In our experiments we included both same type of events and different types of events, as long as the test event occurred after the training event. In the following figures X axis represents time and Y axis values represent the actual or detected number of abnormal events in a sliding window. Each point in the graph represents a window of 10 bins and as a result values in the Y axis range between 0 and 10 showing number of bins classified as event. Intuitively, we decide that if more than half of the bins in the window are classified as abnormal event we raise an alarm. This corresponds to six classifications or in other words 60% of the classifications in the context of window.

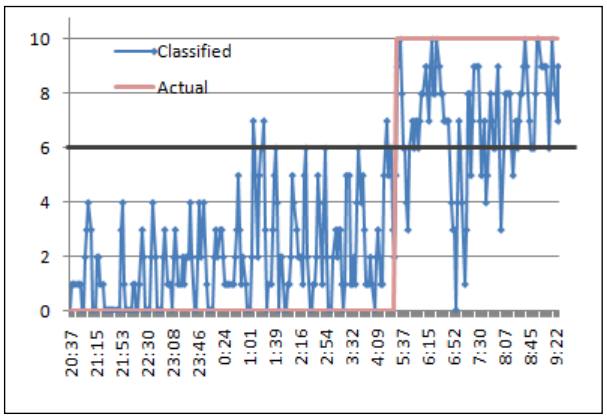

Figure 2. Training with Nimda worm data and testing on Slammer worm.

Figure 2 shows the output of our system when attempting to classify abnormal instances of BGP data during the Slammer period, based on a model built with training data belonging to the Nimda worm, circa 2001. Several experiments were performed in order to adjust the learning process of the algorithm. However, we concluded Slammer was a particularly aggressive event that can be detected on its own by observing the update message count, but, overall, worms are not easily detectable with our system as of today, for the reasons outlined in Section II.

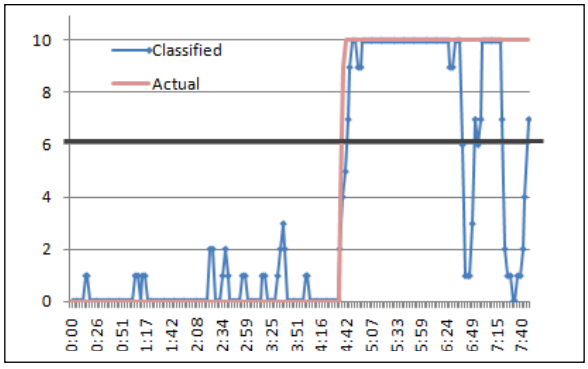

Figure 3. Training with East-coast blackout data and testing on Moscow blackout

Figure 3 provides evidence our research can yield accurate classification in near real time. Data from the East-coast blackout was used this time to train our system and help us detect a subsequent event of the same kind, as the Moscow blackout. The plot shows a clear and accurate detection of the abnormally classified bins, sustained over roughly two hours.

The next two plots (see Figure 4 and Figure 5) cross-tests our systems with data extracted from abnormal BGP events of very different nature, as it is the case of power outages and routing table leaks. Even though these events are implicitly different, they seem to be compatible using data mining models. This may be due to the similarities both events have in common, as blackouts cause a large amount of withdrawals to be exchanged between peers as routers power off and networks become unreachable. As explained in the introduction, the withdrawal of prefixes hastens a surge of prefix announcements analogous to the announcements seen when a table leak is broadcasted to neighboring BGP peers. In this case, and most of the times, the actual prefixes announced are not routable on the Internet itself as they are part of the leaking organization's internal structure, and hence not routable. As peers realize this fact prefixes start to be withdrawn and yet another surge of withdrawals is perceived. These two processes are similar enough for our system to detect. Our current suspicions are this detection can only happen if the routing table leak event large enough to quantify.

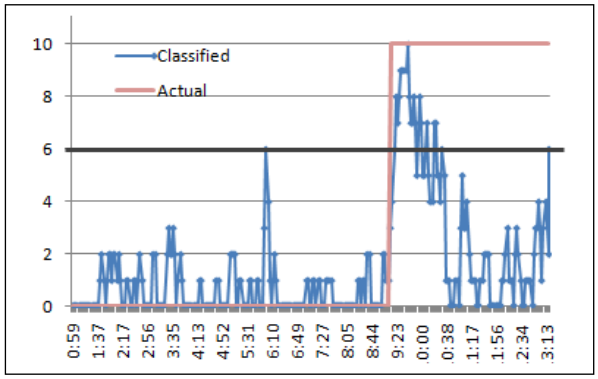

Figure 4. Training with East-coast blackout data and testing on TTNET routing table leak.

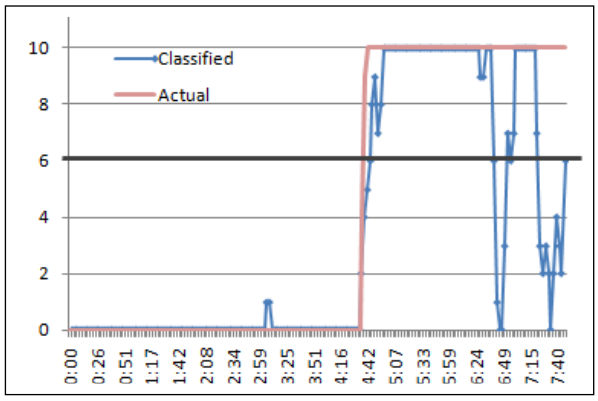

Figure 5. Training with TTNET routing table leak data, testing on Moscow blackout.

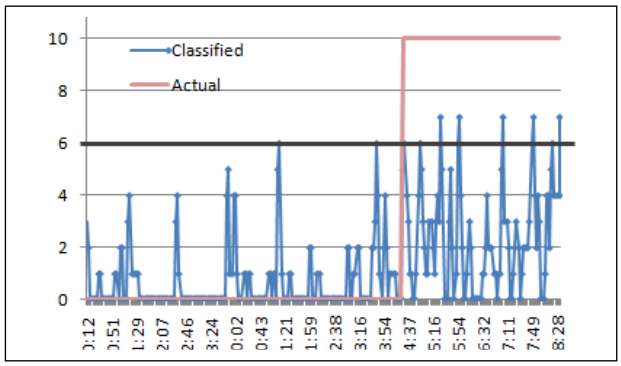

Figure 6. Training with Luzon cable cut data and testing on Mediterranean Sea cable cut.

The last figure (Figure 6) presents the results obtained when comparing two prominent submarine cable cuts. Contrary to our expectations, the system just raised five alerts during the first 240 minutes of the event, taking up to an hour to reach the alert threshold (six and above).

VI. Conclusion and Future Work

In this work we developed an anomaly detection framework combined with a feature extraction mechanism. We apply our framework to anomaly detection on BGP traffic. Our preprocessing module extracts several high level features from BGP traffic. Some of these features are inspired from the literature yet we also introduce several novel features and a novel normalization method that allow classification algorithms to perform better. Additionally, our framework allows us to use several different machine learning algorithms for training and detection of abnormal events. It is not limited to a single technique or algorithm. Combination of advanced features, intelligent normalization of feature values and a powerful classification algorithm, our framework can learn a model from an abnormal event that is happened in the past and use this model to detect a similar type of event in the future. To show the applicability of our framework, we conduct extensive experiments with a variety of abnormal events and classification algorithms. Our results demonstrate that when we train our system ample variety of abnormal BGP events including worm attacks, power supply outages, submarine cable cuts, and misconfigurations such as network prefix hijacks or routing table leaks, we can detect a similar type of event that is happened later.

We plan to explore and test our framework for real time detection and characterization of abnormal events. In this study our anomaly detection framework is demonstrated on BGP anomaly detection. However, it can be used for detection of abnormalities on similar domains as long as aggregate features can be obtained from a stream of entities. In the future we would like to apply our framework to different domains.

Acknowledgment

Large amount of this work is done when the first author was a visiting student in Dogus University, Computer Engineering department through Erasmus program.

References

[1] G. Huston “The changing Foundation of the Internet: confronting IPv4 Address Exhaustion,” in The Internet Protocol Journal, volume 11, number 3, pp. 19-36 September 2008.

[2] Y. Rekhter and T. Li. "A Border Gateway Protocol 4 (BGP-4)," .RFC 1771, IETF, Mar. 1995.

[3] A. Antony, L. Cittadini, D. Karrenberg, R. Kisteleki, T. Refice, T. Vest, R. Wilhelm. "Mediterranean fiber cable cut (January-February 2008) analysis of network dynamics", 2008, unpublished

[4] M. Lad, X. Zhao, B. Zhang, D. Massey, and L. Zhang. "Analysis of BGP update surge during slammer worm attack". In Proc. 5th International Workshop on Distributed Computing, 2003.

[5] K. Butler, T. Farley, P. Mcdaniel and J. Rexford. "A survey of BGP security issues and solutions" Technical report, AT&T Labs - Research, Florham Park, NJ, February 2004.

[6] Cable breaks:Ultimate test of network robustness. (2008). Retrieved March 4, 2011, from http://www.menog.net/menog-meetings/menog3/presentations/poppe-KuwaitMENOG3cablebreaksYPapr08.pdf

[7] RIPE Network Coordination Centre. http://www.ripe.net/data-tools/stats/ris/ris-raw-data

[8] K. Sriram, O. Borchert, O. Kim, P. Gleichmann, D. Montgomery. “A Comparative Analysis of BGP Anomaly Detection and Robustness Algorithms”. CATCH, 2009

[9] C. Kruegel, D. Mutz, W. Robertson and F. Valeur. "Topology-based detection of anomalous BGP messages," Proceedings of ACM Symposium on Recent Advances in Intrusion Detection, (28)20 pp. 17-35, Sept, 2003.

[10] K. Zhang, A. Yen, X. Zhao, D. Massey, S.F. Wu and L. Zhang. "On Detection of Anomalous Routing Dynamics in BGP". Networking 2004, 3042, pp. 259 – 270

[11] J. Zhang, J. Rexford, and J. Feigenbaum. "Learning-Based Anomaly Detection in BGP Updates". Proceeding of the 2005 ACM SIGCOMM Workshop on Mining Network Data, 2005 pp. 219 - 220

[12] M.C. Ganiz, S. Kanitkar,M.C. Chuah, W.M. Pottenger. “Detection of interdomain routing anomalies based on higher-order path analysis,” In: ICDM '06: Proceedings of the Sixth International Conference on Data Mining, Washington, DC, USA, IEEE Computer Society (2006) pp. 874-879

[13] V. Menon, W.M. Pottenger. “A Higher Order Collective Classifier for detecting and classifying network events,” IEEE International Conference on Intelligence and Security Informatics, 2009, pp. 125-130

[14] S. Deshpande, M. Thottan, T. K. Ho, B. Sikdar. "An Online Mechanism for BGP Instability Detection and Analysis," IEEE transactions on Computers, vol. 58, no.11, pp. 1470-1484, 2009

[15] J. Li, D. Dou, Z. Wu, S. Kim. "An Internet Routing Forensics Framework for Discovering Rules of Abnormal BGP Events". ACM Sigcomm, (35)5 pp. 55-66, Oct, 2005.

[16] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, I.H. Witten. "The WEKA Data Mining Software: An Update". SIGKDD Explorations, Volume 11, Issue 1. 2009, pp. 10-18

[17] R. Quinlan. "C4.5: Programs for Machine Learning". Morgan Kaufmann Publishers, San Mateo, CA. 1993

[18] G. H. John and P. Langley. "Estimating continuous distributions in Bayesian classifiers," in Eleventh Conference on Uncertainty in Artificial Intelligence, San Mateo, pp. 338-345, 1995.

[19] C.C. Chang and C.J. Lin, LIBSVM : a library for support vector machines, 2001.