

Learning-Based Anomaly Detection in BGP Updates*

Jian Zhang[†]
Computer Science Dept.
Yale University
New Haven, CT 06520
zhang-jian@cs.yale.edu

Jennifer Rexford
Computer Science Dept.
Princeton University
Princeton, NJ 08544
jrex@cs.princeton.edu

Joan Feigenbaum[‡]
Computer Science Dept.
Yale University
New Haven, CT 06520
feigenbaum@cs.yale.edu

ABSTRACT

Detecting anomalous BGP-route advertisements is crucial for improving the security and robustness of the Internet’s interdomain-routing system. In this paper, we propose an instance-learning framework that identifies anomalies based on deviations from the “normal” BGP-update dynamics for a given destination prefix and across prefixes. We employ wavelets for a systematic, multi-scaled analysis that avoids the “magic numbers” (*e.g.*, for grouping related update messages) needed in previous approaches to BGP-anomaly detection. Our preliminary results show that the update dynamics are generally consistent across prefixes and time. Only a few prefixes differ from the majority, and most prefixes exhibit similar behavior across time. This small set of abnormal prefixes and time intervals may be further examined to determine the source of anomalous behavior. In particular, we observe that many of the unusual prefixes are unstable prefixes that experience frequent routing changes.

Categories and Subject Descriptors

C.2.3 [Computer-Communication Networks]: Network Operations—*network monitoring*

General Terms

Algorithms, Management, Measurement

Keywords

Anomaly Detection, Wavelets, Instance-Based Learning

1. INTRODUCTION

The stability of BGP affects the stability, availability, and efficiency of the Internet. It is thus of great importance to understand the behavior of BGP. In recent years, a lot of

*This work was supported by the DoD University Research Initiative (URI) administered by the Office of Naval Research. A full version of the paper is available as Yale CS department technical report YALEU/DCS/TR-1318.

[†]Supported by NSF grant number 0331548.

[‡]Supported in part by NSF grants 0219018, 0331548, and 0428422 and by ONR grants N00014-01-1-0795 and N00014-04-1-0725.

research has been done on BGP instability [2, 3, 4, 6, 9]. It is understood that not all route changes cause instability problems for the network. It is the “abnormal” route changes (*e.g.*, frequent updates due to flaky equipment, protocol oscillation, route hijacking, *etc.*) that require the network operator’s attention.

Previously, statistics-based anomaly detection[8, 12] has been used to detect abnormal route changes. In a statistics-based system for detecting BGP-route anomalies, the behaviors of BGP updates are normally represented by simple aggregates and their statistics. Thus, such a system is simple and may be easily deployed and run with high efficiency. However, the very simplicity of its representation also makes it unable to capture complex features that may be important for a better analysis of BGP behavior. Furthermore, the representations in many such systems have the “magic-number” problem. That is, they use parameters, set either arbitrarily or according to statistics, for controlling granularity of the analysis or for the threshold that determines when a burst of messages ends. For example, a prefix may be determined to have converged to a stable route if there have been no updates for that prefix in the last T minutes. Clearly, it is hard to determine a good value for such a parameter T (although it is easy to give a very loose upper bound for T , *e.g.*, T should be smaller than 50,000.)

These problems motivate our search for new representations and new frameworks that are independent of “magic numbers” and more powerful in characterizing BGP updates. Towards this end, we propose a framework that uses instance-based learning. In our framework, BGP-update behaviors are represented by a vector of quantified features. Such a representation maps a particular behavior to a point in a multidimensional vector space. The set of normal behaviors maps to a set of points whose neighborhood defines the location of “good” behaviors. A behavior represented by an outlier point that is far away from this location raises suspicion and may require the network operator’s attention.

In this preliminary study, we use the features of BGP-update dynamics, such as burst duration and inter-burst intervals, to construct the representation vectors. As discovered in [5], different update types have different convergence times and numbers of update messages. The convergence and update-message numbers also differ from ISP to ISP. Hence, the configuration of the underlying network affects the way the new paths are explored, which in turn can lead to different timing and message numbers in different systems. If we assume that the underlying network does not change configuration often, the temporal dynamics of

the updates should also be similar. Thus, unusual dynamics may indicate anomalous updates, as the analysis of BGP updates during certain worm attacks have shown [7]. Indeed, several recent works [12, 10] have examined the BGP-update dynamics for signs of anomaly. We use wavelet transformations to construct our representation and use clustering to discover the locations of normal behaviors and the outlier points. By extending the representation to a vector of quantified features, our framework is much more expressive in characterizing BGP-update behaviors.

Our representation avoids the “magic-number” problem by employing a multi-scale transformation. Using a set of different time scales, the wavelet transformation provides a systematic, multi-granularity view of the structures and patterns of BGP updates. Therefore, it avoids the “magic-number” problem not by magically getting rid of all the parameters but by systematically examining the intervals in which the parameters may lie. In this preliminary study, the intervals are upper-bounded by 24 hours. That is, daily BGP updates are examined at different time scales. We remark that the upper bound can be greatly extended in a full system. Another feature of our representation is shift invariance. This feature prevents misclassifications that are often caused by shift-intolerant representations.

2. DETECTION SYSTEM AND PRELIMINARY RESULTS

A sequence of BGP-update messages can be viewed as a signal along time. An update-message burst within this signal can be viewed as a high-frequency signal (the individual updates) modulated by a low-frequency signal (the burst). Because wavelet transformation is a powerful tool for revealing such temporal structures in signals, we use it to extract features from BGP-update dynamics. We use the Haar wavelet because of its simplicity. We intend to experiment with other forms of wavelets in a full-scale study. Let $\Psi(\cdot)$ be the Haar wavelet. The discretized transformation is defined to be $\gamma(\delta, \tau) = \sum_x S(x) \cdot \frac{1}{\sqrt{\delta}} \Psi^*\left(\frac{x-\tau}{\delta}\right)$, where τ is the time translation, and δ is the time scale. Our analysis employs a set of values for τ and δ . The function value $\gamma(\delta, \tau)$ then defines a surface, and we are interested in the peak values, *i.e.*, local maxima, of $\gamma(\delta, \tau)$. This is because a burst of length t will give a large value of $\gamma(\delta, \tau)$ at the scale δ closest to t and at the time τ when the burst happens. Our representation consists of several histograms. We construct one histogram for the peak values of $\gamma(\delta, \tau)$ for each value of δ . To include time information about the bursts in our representation, we also use a histogram for the time intervals between each pair of consecutive peaks. Note that these two types of histograms are shift invariant. Therefore, so is our representation. (A detailed description of the representation can be found in our Technical Report [11].)

However, the above construction does not tell us what kind of features indicate anomalies in BGP updates. This is done by the learning component in our framework. We use clustering to discover the structures in the BGP-update dynamics. Clustering groups the points representing the dynamics into categories that can then be examined and labeled by network operators. Therefore, clustering helps to learn the positions in our representation space that correspond to anomalies in BGP updates. These positions can be used later in identifying anomalies in future BGP up-

dates. In this study, as an efficient way to test our framework, we use k-means clustering. (This clustering method, however, still requires a predefined parameter k . We plan to use hierarchical clustering that does not require any predefined parameter in a full-scale investigation.) We cluster the update dynamics of a single prefix as well as the update dynamics across prefixes over a view.

We experimented with a preliminary implementation of our framework, investigating BGP-update behaviors for six months (06/04-11/04) using data from RouteViews [1]. Focusing on each prefix in isolation, we show that, for most prefixes, update dynamics are similar 80% of the time. Furthermore, on a single day, there are several behaviors with the property that the prefixes displaying these behaviors constitute the majority of the prefixes being updated. Only a few prefixes exhibit behaviors that are quite different from the majority. This small set of prefixes or daily behaviors can be further examined for anomaly detection. In particular, we observe that most prefixes whose update dynamics deviate from the majority are unstable prefixes with frequent routing changes.

As a preliminary study, the system given in this paper still lacks sophistication and fine tuning. Furthermore, it is still an unsubstantiated hypothesis at this point that the outliers discovered by our system are “anomalous.” We plan to explore the “abnormal prefixes” and the “abnormal clusters” (for the same prefix) to understand the relationship between the structures revealed by our system and the operation of BGP. More importantly, we plan to validate, via data from ISPs and known Internet outages/anomalies, whether the “anomalies” raised by our system are truly anomalous.

3. REFERENCES

- [1] Route Views Project. <http://www.routeviews.org>.
- [2] M. Caesar, L. Subramanian, and R. H. Katz. Towards localizing root causes of BGP dynamics. *Technical Report, CSD-3-1292, UC Berkeley*, 2003.
- [3] D.-F. Chang, R. Govindan, and J. Heidemann. The temporal and topological characteristics of BGP path changes. In *Proc. IEEE ICNP*, 2003.
- [4] A. Feldmann, O. Maennel, Z. Mao, A. Berger, and B. Maggs. Locating internet routing instabilities. In *Proc. ACM SIGCOMM*, 2004.
- [5] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed internet routing convergence. *IEEE/ACM Transactions on Networking*, 9(3):293–306, 2001.
- [6] M. Lad, A. Nanavati, D. Massey, and L. Zhang. An algorithmic approach to identifying link failures. In *Proc. Pacific Rim Dependable Computing*, 2004.
- [7] 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.
- [8] S. T. Teoh, K. Zhang, S.-M. Tseng, K.-L. Ma, and S. F. Wu. Combining visual and automated data mining for near-real-time anomaly detection and analysis in BGP. In *Proc. VizSEC/DMSEC*, 2004.
- [9] J. Wu, Z. M. Mao, J. Rexford, and J. Wang. Finding a needle in a haystack: Pinpointing significant BGP routing changes in an IP network. In *Proc. Networked Systems Design and Implementation*, 2005.
- [10] K. Xu, J. Chandrashekar, and Z. Zhang. A first step toward understanding inter-domain routing dynamics. In *Sigcomm Workshop on Mining Network Data*, 2005.
- [11] J. Zhang, J. Rexford, and J. Feigenbaum. Learning-based anomaly detection in BGP updates. *Yale University Technical Report, YALEU/DCS/TR-1318*, 2005.
- [12] K. Zhang, A. Yen, X. Zhao, D. Massey, S. F. Wu, and L. Zhang. On detection of anomalous routing dynamics in BGP. In *Proc. NETWORKING LNCS 3042*, pages 259–270, 2004.