Detecting BGP Instability Using Recurrence Quantification Analysis (RQA)

Bahaa Al-Musawi[1,2], Philip Branch[1], and Grenville Armitage[1]

[1]CAIA, Swinburne University of Technology, Melbourne, Australia
[2]University of Kufa, Al-Najaf, Iraq
{balmusawi,pbranch,garmitage}@swin.edu.au

***Abstract*—The Border Gateway Protocol (BGP) is the default Internet routing protocol that manages connectivity among Autonomous Systems (ASes). Although BGP disruptions are rare, when they occur the consequences can be very damaging. Consequently there has been considerable effort aimed at understanding what is normal and abnormal BGP traffic and, in so doing, enable potentially disruptive anomalous traffic to be identified quickly. In this paper, we make two contributions. We show that over time BGP messages from BGP speakers have deterministic, recurrence and non-linear properties, then build on this insight to introduce the idea of using Recurrence Quantification Analysis (RQA) to detect BGP instability. RQA can be used to provide rapid identification of traffic anomalies that can lead to BGP instability. Furthermore, RQA is able to detect abnormal behaviours that may pass without observation.**

I. Introduction

The Border Gateway Protocol (BGP) is the current inter-domain routing protocol that maintains and exchanges network reachability information between Autonomous Systems (ASes). BGP was developed at a time when information provided by an AS could be assumed to be accurate. Consequently, it includes few security mechanisms [1]. Although propagation of inaccurate information via BGP is fortunately rare, when an AS, either deliberately or accidentally propagates incorrect information, the consequences can be very serious. The process of detecting abnormal data in a series of BGP update messages represents a challenge for researchers and operators, especially during unstable periods, as routing data is complex, noisy, and voluminous. To compound the challenge, single events such as link failure can produce multiple update messages, affect routing decisions, and erroneously propagate incorrect destination prefixes [2].

Routing instability has been defined as a fluctuation in topology information and network reachability [3]. BGP routing instability appears as fluctuations in the number of BGP updates and/or path length for an AS [4], [5]. BGP instability can be a result of different types of disruptions such as hardware failure, misconfiguration, hijacking, software bugs, faulty equipment, and Denial of Service (DoS) attacks. Instability affects performance, processing load, and distribution balance of traffic load for BGP speakers [4]. It is worth noting that it is not just direct attacks on BGP that can affect its stability. Although malware such as Nimda and Slammer were directed at web servers, BGP routing stability was also affected during these attacks [6].

In the years since it was deployed, many types of disruptions have threatened BGP stability such as panix.com domain hijack, TTNet misconfiguration, and Moscow blackout [6]. In addition to many reported events, other types of events remain unreported or even unnoticed [7]. Shi et al. in [8] present statistics and trends of BGP anomalies during a period of 1 year monitoring from May 2012. During this period, around 40k bogus routes were detected. Among these bogus routes, there were 193 BGP hijacks and 27 misconfigurations. Furthermore, about 20% of the hijacking and misconfigurations last less than 10 minutes but with ability to pollute 90% of the Internet in less than 2 minutes. These statistics demonstrate the need for a real-time detection of BGP instability caused by different types of disruptions.

Considerable research has been carried out into BGP stability. Generally, research work can be classified as improving stability during route changes caused by link failures [9], [10] and detecting different types of anomaly which indicate the likely onset of instability [2]–[4], [6], [8]. Our interest is in the latter. In particular, we are interested in the rapid detection of BGP disruptions indicative of routing instability at the AS level.

BGP is an incremental protocol. BGP updates should reflect significant network engineering decisions by the AS to add or remove a network or as a result of a major reachability issue. However, real-world BGP update traffic is of a substantial volume that is much larger than might be expected. There is a substantial background traffic consisting of route announcements followed soon after by withdrawals that do not appear related to underlying network management decisions or events.

The source of this oscillatory behaviour has attracted great attention. In [11] the authors showed that the widely used multi discriminator (MED) attribute can lead to persistent oscillatory behaviour in BGP. In [3] Labovitz et al. examined BGP data and found that it was overwhelmingly what they described as 'pathological updates' made up of duplicate withdrawal announcements, oscillating reachability announcements and duplicate path announcements. Varadhan et al in [12] observed that interdomain routing is prone to persistent route

978-1-4673-8590-9/15/$31.00 ©2015 IEEE

oscillations. Others have also identified the phenomenon of oscillatory behaviour of BGP updates [13].

Regardless of the cause, much of the background traffic of BGP is made of oscillations from different ASes of different frequencies. Detecting anomalous behaviour in this environment is challenging. Indeed defining what is and is not anomalous is difficult. All traffic that does not reflect underlying network changes can be thought of as anomalous. Nevertheless, such traffic does not necessarily represent dangerous instability that can lead to disruption. What is needed is a method that can detect changes to existing patterns of behaviour. In this paper, we describe the use of Recurrence Quantification Analysis (RQA) to detect such changes.

We begin by examining traffic volume and message path length from the ten most active BGP speakers and show that each generates approximately periodic updates but with different frequencies. Aggregated, unsynchronised traffic of this kind has the characteristics of a deterministic, stable system. RQA is an established technique that has been used successfully to model the dynamic behaviour of such systems found in physiology, physics and mechanical engineering. In this paper, we describe its application to modelling BGP traffic. We demonstrate that background BGP traffic is well modelled as a deterministic stable system. We also analyse a recent significant BGP instability and show that RQA is able to quickly identify the onset of instability and is able to show otherwise hidden information about the event.

The rest of this paper is structured as follows. Section II explores related work in detection of BGP instability. Section III presents a brief statistical analysis of BGP traffic from the ten most active ASes and shows that they generate approximately periodic traffic, but of different frequencies. Section IV outlines our approach using RQA. Section V is our results section where we demonstrate the effectiveness of RQA in analysing BGP traffic. Finally, section VI is our Conclusion where we summarise the paper and outline future directions for this work.

II. Previous Work

In this section we briefly cover BGP research related to our topic. This has been in two areas: measuring and modeling BGP behaviour [3], [14], [15] and BGP instability detection [2], [4], [6].

BGP measurements such as those reported in [3] and [14] focus on characteristics of BGP updates while modeling studies such as [15], [16] focus on analysis of BGP 'churn' in term of its evolution and causes; that is the announcement and withdrawal of routing updates. One of the earliest efforts at characterising BGP updates and identifying instability was by Labovitz et al. [3]. Labovitz observed that much BGP activity at the core of the Internet was 'churn'. To detect BGP instability, the authors applied the Fast Fourier Transform (FFT) to routing update rates and demonstrated that rapid changes in routing updates are correlated with instability [3]. In [14], the authors use Wavelet Transforms and Median Filtering to identify patterns in BGP updates. They observed BGP traffic as being self-similar and bursty. In addition, they detected other phenomena such as prolonged oscillations of updates ranging in duration from hours to months.

Detecting BGP instability has attracted a great deal of attention [2]–[4], [6], [14]. Huang et al. [2] introduced a technique to detect BGP node, link, and peer failure. Their technique uses Principal Component Analysis (PCA) to detect the failure type. They used BGP update volume as a single BGP feature extracted every 10 minutes with window size of 200 minutes. Although the approach is able to detect and identify BGP node, link, and peer failure, it requires routers configuration information and is slow, typically taking 9 to 96 minutes.

GLRT is a standard statistical technique used in hypothesis testing. Deshpande et al. adopted it as an instability detection technique [6]. Their approach was based on statistical pattern recognition incorporating the technique. The approach was evaluated against data collected from notable disruptions including the Moscow blackout, the Nimda worm, and the Panix domain hijack. However, once again this detection approach is slow, typically requiring around one hour to detect instability.

Haeberlen et al. [4] presented a prototype to detect BGP faults at the AS level called NetReview. This prototype uses one year of BGP data to detect BGP faults, where BGP faults include BGP router and link failure, misconfiguration, policy violations, and attacks. Although NetReview can quickly detect different types of BGP faults and identify their source cause, it requires information about each ASes policy configuration and requires the storing and processing of large log files, leading to scalability issues especially for large Internet Service Providers (ISPs).

In our work we present evidence that BGP updates are dominated by unsynchronised oscillations and can be modelled very effectively as a dynamical system with RQA. We demonstrate that RQA is able to quickly detect subtle changes in the underlying background BGP traffic without needing to store large amounts of historical data. It is also able to detect hidden information such as changes to update patterns which might otherwise pass without observation.

III. Modeling BGP as a dynamical system

In this section, we examine BGP traffic using online repositories to access BGP data, and model a BGP speaker as a dynamical system based on BGP updates it sends. We show that BGP speakers show stable, deterministic, and non-linear behaviour. We also analyse BGP updates sent from most active ASes and show that ASes are approximately periodic with unsynchronised oscillations.

A. BGP Dataset

We use publicly available BGP control plane datasets[1]to model BGP speakers and detect instability. We refer to BGP control plane's traffic as BGP traffic. The RouteViews project [17] and Réseaux IP Européens (RIPE) Network Coordinate

[1]Route Information Base (RIB) and BGP update messages

Centre (NNC) [18] are the two most well known repositories of BGP control plane data. Each of these repositories has multiple collectors which run BGP sessions with several routers, referred to as monitors, in many networks.

Figure 1 show an example for BGP topology of the Route-Views project at collector 4 which was peered with 40 peers on 26th of July 2015. In this example, AS24516 and AS58511 represent monitoring ponits, AS22059 is a source AS, and AS2764, AS7545, AS6939, AS174, AS11404 are intermediate ASes. When AS22059 sends a BGP update, AS24526 may receive multiple copies of this update with different path lengths such as (6939 22059), (174 11404 22059), and (2764 7545 6939 22059).

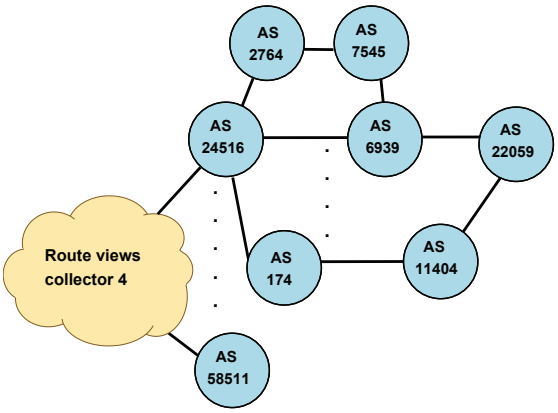

Fig. 1: Simple BGP topology

BGP traffic has been characterised as noisy and bursty [14], [19]. The substantial background traffic for this behaviour consists of route announcements followed soon after by withdrawals or updates that do not appear related to underlying network management decisions or events. It is not clear where this behaviour comes from, but there has been some speculation that it is caused by misconfiguration [14], [19].

B. Modeling

In this section, we model a BGP speaker as a dynamical system sending BGP updates and path lengths depending on BGP messages received from neighbors and local routing policies. When a BGP speaker receives a BGP message that changes its routing table it will propagate that message to all or a group of its neighbor based on its local policies. Otherwise, the message will be terminated.

In a dynamical system, the state changes over time in different ways. The study of dynamics for different systems is an important task in many disciplines. It provides a way of understanding these systems and predicting their behaviour over time [20]. Formally, a dynamical system is defined by a phase space, a time evolution law, and continuous or discrete time. In phase space, all possible states of a system are represented where each possible state of the system corresponds to a unique point in phase space. For example, state of a system at a fixed time t can be specified by d variables to form a vector $x(t)$ in d-dimensional phase space [20]. This vector can be described as

$$\vec{x}(t) = (x_1(t), x_2(t), \ldots, x_d(t))^T \tag{1}$$

The time evolution law allows determination of the state of the system at time t based on previous states. The time evolution for a continuous time system can be described by a set of differential equations.

$$\dot{\vec{x}}(t) = \frac{d\vec{x}(t)}{dt} = \vec{F}(\vec{x}(t)), \qquad F: \mathbb{R}^d \to \mathbb{R}^d \tag{2}$$

Where vector $\vec{x}(t)$ is a trajectory in phase space [20]. Experimentally, not all components are known or can be measured. A scalar and discrete time series (u_i) can be an alternative option; $u_i = u(i\triangle t)$, where $i = 1, \ldots, N$ and $\triangle t$ is a sample rate. In this case, phase space can be reconstructed using the time embedding method by

$$\hat{\vec{x}}_i = \sum_{j=1}^{m} u_i + (j-1)_\tau \vec{e}_j, \tag{3}$$

Where m is the embedding dimension, τ is the time delay, $\vec{e}_j$ are the unit vectors.

For the analysis of time series, the phase space parameters are represented by embedding dimension and time delay. These parameters have to be selected carefully. Different approaches to estimate the smallest sufficient value of embedding dimension are available such as the false nearest-neighbor algorithm while auto-correlation function and the mutual information function are used to estimate time delay [21].

The concept of phase space is powerful in modelling deterministic systems. For a purely deterministic system all future states can be determined when its current state is known. Phase space is also useful for understanding non-deterministic systems when they are described as a set of states that specify system transition. To that end, we analyse a BGP speaker as a dynamical system in terms of type of motion, determinism, and linearity. These properties help to understand system behaviour and can be used to detect different transitions that identify system instability.

C. Type of motion

In dynamical systems, there are different types of motion such as stable, where a system's behaviour appears stable around a point in the phase space, and noise, where the behaviour is fully random. Identifying the type of motion for a dynamical system can help to understand system behaviour.

While estimating the type of motion in a dynamical system is comparatively easy if the equation of motion in the phase space is available, it is a difficult task when only a series of data is available. With lack of knowledge about the underlying dynamics, maximal Lyapunov exponent is a good measurement to estimate the type of motion in dynamical systems [22], [23].

Different methods have been proposed to find the maximum Lyapunov exponents, the most well-known methods are described in [22] and [23]. While the method described in [23] does not depend on the correct embedding dimension, the method in [22] does. We use TISEAN [21], a software package for analysis of time series with methods based on the theory of nonlinear deterministic dynamical systems, to estimate maximum Lyapunov exponents based on methods described in [22] and [23] respectively. If s_{n1} and s_{n2} are two points in phase space with distance $s_{n1} - s_{n2} = \triangle_0 \ll 1$, distance after a time Δl is $\delta_{\Delta l} = s_{n1+\Delta l} - s_{n2+\Delta l}$. The maximal Lyapunov exponents represents the slope of the case defined in (4), where a positive value is an indication of a chaotic system and a zero slope corresponds to a stable fixed point system.

$$\delta(\epsilon, m, t) = \left\langle \ln \left(\frac{1}{\mid \upsilon_n \mid} \sum_{s_{n2} \epsilon \upsilon_n} \mid s_{n1+t} - s_{n2+t} \mid \right) \right\rangle_n \tag{4}$$

where m is the embedding dimension, t is time delay, υ_n is the $\epsilon - neighborhood$ of s_n [21].

We estimate the maximal Lyapunov exponents for multiple ASes on different dates and found the slope is zero which we interpret as a stable type of motion for BGP speakers. Figure 2 shows values of Lyapunov exponents for multiple dimensions on the AS10102 where $\delta(\epsilon, m, t)$ exhibits a flat line which indicates possible stable behaviour.

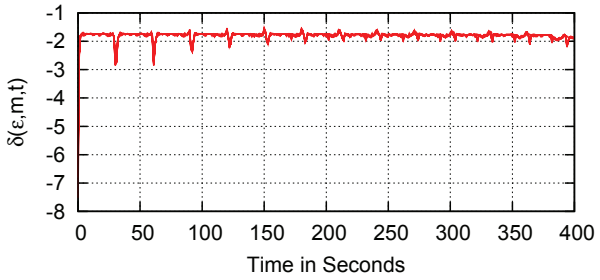

Fig. 2: Lyapunov exponents estimation for AS10102

D. BGP Periodicity

The phase plane behaviour of BGP can be attributed to underlying unsynchronised periodic behaviour of highly active BGP speakers. One possible source of periodicity is the Minimal Route Advertisement Interval (MRAI). The MRAI refers to the minimum amount of time between two subsequent advertisements to a particular destination [24], the default value in Cisco routers is 30 seconds while in Juniper routers is 0 second.

Periodicity can also be seen in the most active ASes. These ASes show reasonably periodic behaviour in terms of sending BGP updates. [25] provides a weekly BGP instability report for the 50 most active ASes. We analysed the periodicity of the most active noisy ASes as reported by [25] during the period 20th to 27th July 2015 and found these ASes

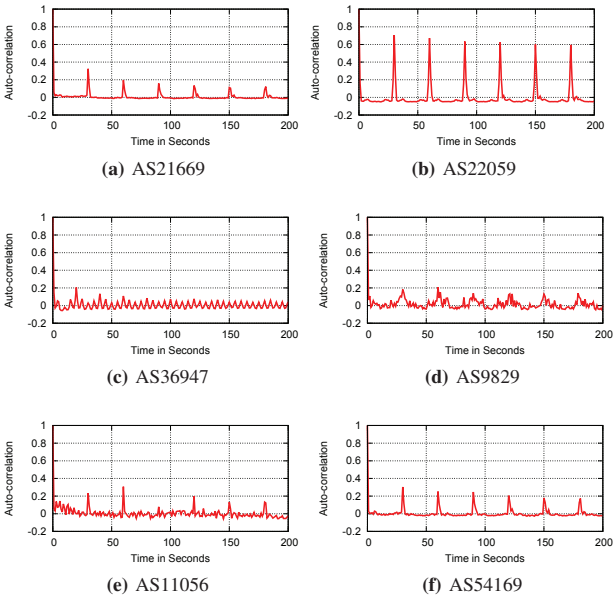

(a) AS21669 **(b)** AS22059

(c) AS36947 **(d)** AS9829

(e) AS11056 **(f)** AS54169

Fig. 3: periodicity of unstable ASes

periodically sent different numbers of BGP updates with different time intervals. Figure 3 shows periodicity for the six most active ASes. For example, AS21669 and AS22059 show periodicity of 30 seconds, and AS36947 shows periodicity of 5 seconds as observed from monitoring point AS24516 at route-views4.routeviews.org which was peered with 40 peers on 26th of July 2015. Figure 4 shows the aggregated BGP updates from the ten most unstable ASes. The unsynchronised aggregation of different periodic updates leads to recurrence behaviour in the underlying system.

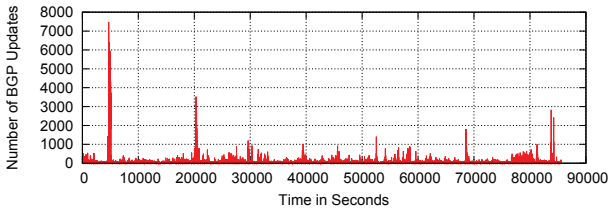

Fig. 4: Unsynchronised aggregation of different periodic updates

E. Determinism and non-linearity

Detecting determinism and linearity properties in a system helps to select an appropriate method to predict system behaviour. Different methods for detecting the existence of determinism and/or non-linearity in time series are available such as [26] and [27]. We use delay vector variance (DVV) described in [27], a method based on the concepts of false nearest neighbor and Kaplan's method [26], to examine BGP data for determinism and non-linearity. The DVV uses an

approach for comparing the characteristics of time series based on its predictability against those obtained for linearised versions of the signal.

The DVV requires the proper selection of time delay and embedding dimension. The examination of determinism and non-linearity can be interpreted using a DVV plot and DVV scatter diagram respectively. The examination of determinism can be observed in a DDV plot by observing DVV plots converging to unity while non-linearity can be examined in DVV scatter by deviation from the bisector line. We estimate determinism and linearity for multiple ASes on different dates and have found that most BGP speakers show properties of determinisim and non-linearity. In Figure 5a, we can see the variance converges to a value of 1 which indicates determinism while in Figure 5b shows DVV scatter where we can observe a deviation from the bisector line as an indication of non-linearity.

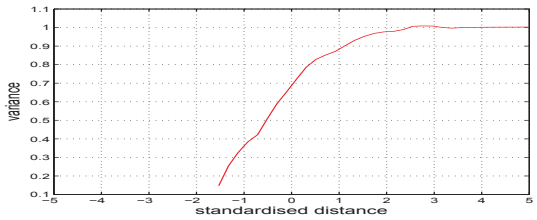

(a) DVV plot

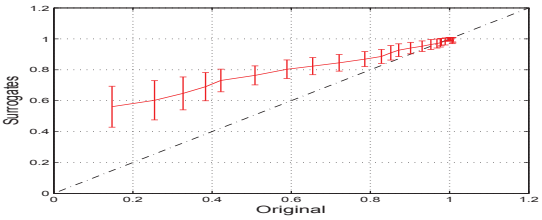

(b) DVV scatter

Fig. 5: Estimation of determinism and non-linearity

IV. A non-linear approach to detecting BGP instability

As outlined in the previous section, BGP messages sent from BGP speakers have been characterized as deterministic, stable and non-linear. While recurrence is a fundamental property in many dynamical systems, it is the most important feature of stable systems.

Different methods and models for analysis of time series and forecasting are available. The Fourier transform (FT) and the autoregressive integrated moving average (ARIMA) model are the most well-known for analysis and forecasting time series. The FT and the ARIMA have some limitation for analysis of time series in general and for BGP data in particular. The limitations of the Fourier transform are time and frequency positions, non-stationarity or abrupt changes in a signal can spread out for whole signal as well as resolution. These drawbacks limit the application of the FT in detecting short periods of BGP instability. The ARIMA model has two significant limitations: (1) future values are assumed to be a linear function of past values and (2) a large amount of historical data is required to obtain reliable predictions. BGP data has been characterised in section III as non-linear and stable. In [14], [15] it is described as a bursty and self-similarity which motivate us to look for another approach.

Here, we use Recurrence Quantification Analysis (RQA), a non-linear technique based on a phase plane trajectory, to detect BGP instability based on calculating RQA measurements for observed BGP data sent from a BGP speaker. RQA was introduced to quantify the important aspects revealed by Recurrence Plot (RP). RQA has been developed to find different transitions between regular, laminar, non-stationarity, and stable behaviours in complex systems such as those found in astrophysics and geosciences. This approach has been introduced to find these transitions using short data series where most nonlinear techniques are ineffective or require a long data series [20]. RQA has shown its ability to discover time correlation between data that do not rely on linear or non-linear assumptions and are not distinguishable through using a direct study of one dimension time series.

A. Recurrence Plot

Recurrence Plot (RP) is an advanced nonlinear analysis technique introduced by Eckmann et al. [28]. The RP was initially produced to graphically display recurring patterns and non-stationarity in time series. RP was introduced as a tool to visualise the time dependent behaviour of the dynamics of a system as a square matrix where each element corresponds to a point in time states. With enough data, structural patterns in the RP can reveal information about the time evolution of the phase space. RP is not limited to long data sets. It can be used for short data sets, noisy data, and non-stationary data [29].

RPs can be formally expressed by the matrix R

$$\mathrm{R_{i,j}}(\varepsilon) \quad = \quad \Theta\Big(\varepsilon - \|\overrightarrow{x_i} - \overrightarrow{x_j}\|\Big), \qquad i, j = 1, \ldots, N, \quad (5)$$

where $R_{i,j}$ is an element of the recurrence matrix R, N is the number of measure points, ε is a threshold distance, $\Theta(\cdot)$ the Heaviside function and $(\|\cdot\|)$ is a normalization operation.

In RP, three parameters have to be selected carefully: time delay (τ), embedding dimension (m), and threshold (ε). We use the auto-correlation function to estimate time delay and false nearest neighbor (FNN) to estimate the embedding dimension [29]. Finally, threshold value is selected based on recommendations in [20] where the value of the threshold is 10% of the maximum phase space diameter.

Figure 6 shows a recurrent plot and underlying time series for aggregated BGP updates sent from the most active ASes with embedding dimension=3, threshold=1.8 and time delay=30 seconds.

From the RP structure, we can infer the characteristics of large and small scale patterns. At the large scale in Figure 6, we can see periodic recurrent structures indicating an oscillating system. At the small scale, single isolated points indicate large fluctuations of stochastic behaviour while vertical and

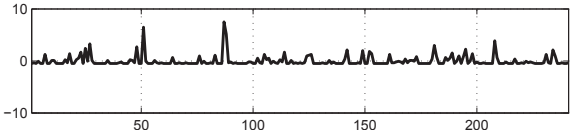

(a) Underlying time series

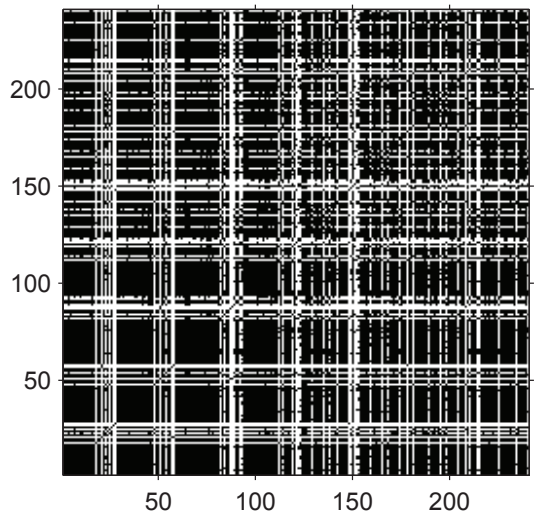

(b) Recurrence Plot

Fig. 6: Recurrence plot for top-ten noisy ASes

horizontal lines forming rectangles indicate some states that do not change or change slowly. This behaviour is expected as BGP speakers do not continually send BGP updates based on MRAI timer values.

The large and small scale patterns from Figure 6 support our outlines in Section III-C and Section III-D, where BGP speakers show recurrence behaviour caused by unsynchronised periodicity. However, the visual interpretation of RP requires some experience; therefore, the quantification of RP was introduced to offer a more objective way of evaluating the system under investigation.

B. Recurrence Quantification Analysis

To reduce the difficulty of RP interpretation, Recurrence Quantification Analysis (RQA) was introduced by Zbilut and Webber [30] to provide objective quantification of important aspects revealed by RPs. In RQA, the density of *recurrence* points as well as the histograms of the length of the diagonal and vertical lines in the RP is quantified.

RQA provides several measures of complexity. The first measurement is Recurrence Rate (RR) which measures the percentage of recurrent points in the phase space. The second measurement is determinism (DET) which can be interpreted as the probability that two closely evolving segments of the phase space trajectory will remain close in the next time step. The third measurement is laminarity which measures the probability that a state will not change (within the ε error) for the next time step. The fourth measurement is Trapping Time (TT) which contains information about the vertical structures in the RP. It can be used to measure how long the system remains in a specific state [20].

V. Instability Detection with RQA

As noted earlier, we define BGP instability as fluctuations in the number of BGP updates and/or path lengths for an AS [4], [5]. We have adopted these two factors in our approach as an input to measure their fluctuation over time.

The strength of RQA applied to this approach is in its ability to rapidly distinguish between fluctuations that are part of normal behavior and fluctuations that indicate instability. Furthermore, RQA is able to detect behaviour that cannot be detected with other techniques.

One of the advantage of this approach is that systems based upon it can be deployed in different scenarios. For example, it can be used to monitor instability on a particular BGP router through analysing its BGP updates or monitor unstable behaviour caused by its neighbors through analysing BGP messages received from them. It also can be used to remotely monitor a set of ASes by monitoring BGP messages belonging to these ASes from different monitoring points.

A. Method

Our approach is based on instability behaviour observed from a BGP speaker in terms of the number of BGP updates and AS-PATH lengths to detect BGP instability. It is comprised of two main components. The first is extraction and analysis of BGP features while the second is calculating RQA measurements for the BGP features to detect variations from these features. We use BGP volume (V) and average length of AS-Path (AL) as BGP features. BGP volume refers to total number of BGP updates per second while average AS-Path length refers to the average length of AS-PATH for BGP announcements per second. These two features can be calculated as follows:

$$V = A + W \tag{6}$$

$$AL = \left[\frac{TA}{A}\right] \tag{7}$$

where A is number of announcements, W number of withdrawals, TA is total AS-PATH length for announcements, and [] is the nearest integer function.

We have found TT and RR (IV-B) to be very effective variables for detecting BGP instability. We have adopted the approach described in [31] and [32] that calculates the RQA variables TT and RR for both BGP features. These measurements are calculated every second with a window size of 300 seconds. There is a tradeoff in window size. A large window size may fail to identify some transitions in system behaviour while choosing too small a window can generate spurious fluctuations in RQA measurements. We have tested different window sizes such as 100, 180, 300, 600, 900, and 1200 seconds and have found the window size of 300 seconds is a good choice as a sliding window.

B. Results and Discussion

One of the most recent incidents of BGP instability was observed on the 12th of July 2015 by Telekom Malaysia (TMnet) which caused significant network problems for the global routing system. TMnet (AS4788) accidentally announced approximately 179,000 prefixes to Level3, the global crossing AS, leading to significant packet loss and slow Internet service around the world [33].

We apply our approach based on RQA to detect the effect of this incident on multiple ASes, in particular on many peers that are connected to route-views4 in the routeviews project such as AS10102, AS1299, AS197264, AS3267, and AS58511 [17]. RQA shows its ability to detect BGP instability caused by high volume of BGP updates as well as hidden abnormal behaviour using only 300 seconds as a BGP history.

As outlined in Section (IV-B), TT and RR variables measure fluctuations and recurrence ratio respectively for a given dataset, where high value of TT indicates a high fluctuation and low value of RR indicates low recurrence behaviour for a given dataset.

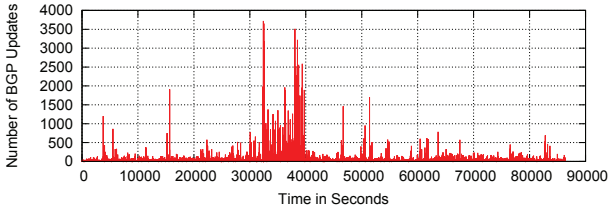

(a) BGP Volume

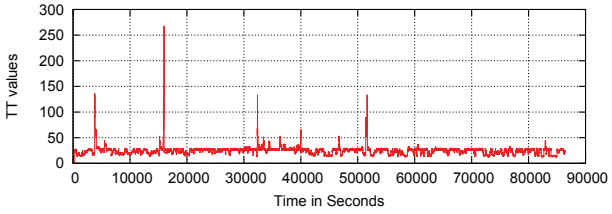

(b) RQA with TT Measurement

Fig. 7: Instability detection for BGP volume feature

Figure 7 shows the number of BGP updates sent from AS10102 and TT measurement during the TMnet event. RQA can detect multiple periods of instability caused by high volume of BGP updates sent during TMnet. Figure 8 shows the ability of RQA to quickly detect BGP instability where TT value flagged after one second of rises in BGP volume.

Furthermore, RQA can detect multiple periods of hidden abnormal behaviour with of RR measurement such as (10000-20000) and (30000-34000) seconds as shown in Figure 9. The detection was not caused by a notable changing in AS-PATH length, but it is related to anomalous behaviour for MRAI timers. For example, during the period (32200-33200) seconds the AS10102 sent multiple BGP updates during some seconds,

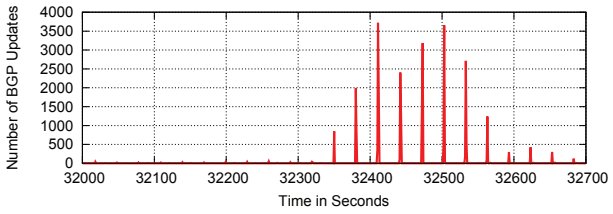

(a) BGP Volume

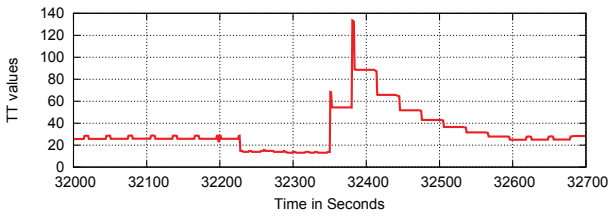

(b) RQA with TT Measurement

Fig. 8: Rapid detection for instability with RQA

which is abnormal behaviour in term of MRAI, instead of 30 seconds as shown in Figure 10.

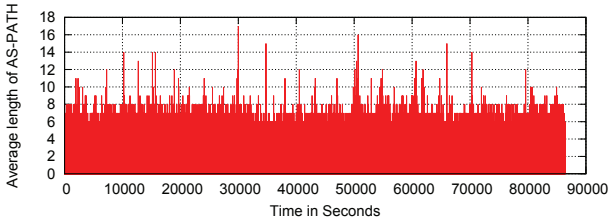

(a) Average AS-PATH length

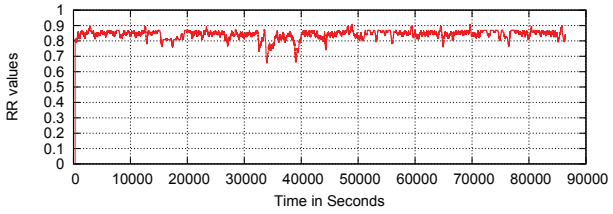

(b) RQA with RR Measurement

Fig. 9: Instability detection for average AS-PATH length feature

VI. Conclusions

In this paper, we model a BGP speaker as a dynamical system sending different values of BGP updates and AS-PATH lengths depending on updates received from neighbors and local routing policies. The analysis of BGP updates sent from a BGP speaker shows stable, deterministic, and non-linear behaviour. These characteristics help to explain BGP behaviour and detect instability which may arise from different sources such as hardware failure, misconfiguration, and

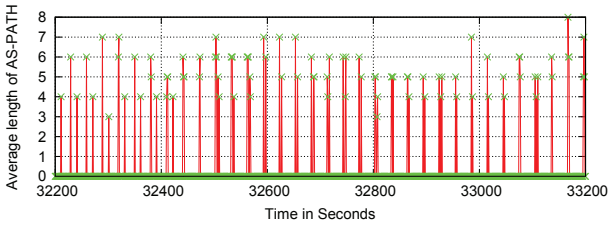

Fig. 10: Anomalous behaviour for MRAI in AS10102

software bugs. We also speculate on possible sources for this recurrence behaviour, including MRAI and the unsynchronised aggregation updates for the most active ASes.

The outlined characteristics of our model motivate us to introduce a new approach to detect BGP instability based on RQA. RQA is a tool to extract hidden information from statistics of dynamic non-linear systems. We have shown that RQA can rapidly identify anomalous fluctuations in the number of BGP updates and AS-PATH lengths without the need for a long BGP history. It is also able to detect hidden anomalous behaviour such as changes to update patterns which might otherwise pass without observation.

Our work suggests that the approach has the potential to not just identify the fact of BGP instability but also the type of disruptions that cause instability. To that end we will be investigating in more detail different types of disruptions and RQA measurements using a controlled testbed.

References

[1] G. Huston, M. Rossi, and G. Armitage, "Securing BGP-A Literature Survey," *Communications Surveys Tutorials, IEEE*, vol. 13, no. 2, pp. 199–222, Second 2011.

[2] Y. Huang, N. Feamster, A. Lakhina, and J. J. Xu, "Diagnosing Network Disruptions with Network-wide Analysis," *SIGMETRICS Perform. Eval. Rev.*, vol. 35, no. 1, pp. 61–72, Jun. 2007.

[3] C. Labovitz, G. R. Malan, and F. Jahanian, "Internet Routing Instability," *IEEE/ACM Trans. Netw.*, vol. 6, no. 5, pp. 515–528, Oct. 1998.

[4] A. Haeberlen, I. Avramopoulos, J. Rexford, and P. Druschel, "NetReview: Detecting when Interdomain Routing Goes Wrong," in *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, ser. NSDI'09. Berkeley, CA, USA: USENIX Association, 2009, pp. 437–452.

[5] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP Routing Stability of Popular Destinations," in *Proceedings of the 2Nd ACM SIGCOMM Workshop on Internet Measurment*, ser. IMW '02. New York, NY, USA: ACM, 2002, pp. 197–202.

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

[7] M. Wählisch, O. Maennel, and T. C. Schmidt, "Towards Detecting BGP Route Hijacking Using the RPKI," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 4, pp. 103–104, Aug. 2012.

[8] X. Shi, Y. Xiang, Z. Wang, X. Yin, and J. Wu, "Detecting Prefix Hijackings in the Internet with Argus," in *Proceedings of the 2012 ACM Conference on Internet Measurement Conference*, ser. IMC '12. New York, NY, USA: ACM, 2012, pp. 15–28.

[9] J. Chandrashekar, Z. Duan, Z.-L. Zhang, and J. Krasky, "Limiting path exploration in BGP," in *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, vol. 4, March 2005, pp. 2337–2348 vol. 4.

[10] G. Huston, M. Rossi, and G. Armitage, "A Technique for Reducing BGP Update Announcements through Path Exploration Damping," *Selected Areas in Communications, IEEE Journal on*, vol. 28, no. 8, pp. 1271–1286, October 2010.

[11] T. G. Griffin and G. Wilfong, "Analysis of the MED Oscillation Problem in BGP," in *Network Protocols, 2002. Proceedings. 10th IEEE International Conference on*. IEEE, 2002, pp. 90–99.

[12] K. Varadhan, R. Govindan, and D. Estrin, "Persistent route oscillations in inter-domain routing," *Computer networks*, vol. 32, no. 1, pp. 1–16, 2000.

[13] M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush, "10 lessons from 10 years of measuring and modeling the internet's autonomous systems," *Selected Areas in Communications, IEEE Journal on*, vol. 29, no. 9, pp. 1810–1821, 2011.

[14] B. A. Prakash, N. Valler, D. Andersen, M. Faloutsos, and C. Faloutsos, "BGP-lens: Patterns and Anomalies in Internet Routing Updates," in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '09. New York, NY, USA: ACM, 2009, pp. 1315–1324.

[15] N. C. Valler, M. Butkiewicz, B. A. Prakash, M. Faloutsos, and C. Faloutsos, "Non-binary information propagation: Modeling BGP routing churn," in *Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on*. IEEE, 2011, pp. 900–905.

[16] G. Huston and G. Armitage, "Projecting future IPv4 router requirements from trends in dynamic BGP behaviour," in *Proc. of ATNAC*, 2006.

[17] University of Oregon, "University of Oregon Route Views Project." [Online]. Available: http://www.routeviews.org/

[18] Reseaux IP Europeens Network Coordination Center. [Online]. Available: http://www.ripe.net/

[19] A. Elmokashfi, A. Kvalbein, and C. Dovrolis, "BGP churn evolution: a perspective from the core," *Networking, IEEE/ACM Transactions on*, vol. 20, no. 2, pp. 571–584, 2012.

[20] N. Marwan, M. C. Romano, M. Thiel, and J. Kurths, "Recurrence plots for the analysis of complex systems," *Physics Reports*, vol. 438, no. 5, pp. 237–329, 2007.

[21] R. Hegger, H. Kantz, and T. Schreiber, "Practical implementation of nonlinear time series methods: The TISEAN package," *Chaos: An Interdisciplinary Journal of Nonlinear Science*, vol. 9, no. 2, pp. 413–435, 1999.

[22] M. T. Rosenstein, J. J. Collins, and C. J. De Luca, "A practical method for calculating largest Lyapunov exponents from small data sets," *Physica D: Nonlinear Phenomena*, vol. 65, no. 1, pp. 117–134, 1993.

[23] H. Kantz, "A robust method to estimate the maximal Lyapunov exponent of a time series," *Physics letters A*, vol. 185, no. 1, pp. 77–87, 1994.

[24] Y. Rekhter, T. Li, and S. Hares, "RFC 4271: A Border Gateway Protocol 4 (BGP-4)," RFC 4271 (Proposed Standard), Internet Engineering Task Force, January 2006. [Online]. Available: http://tools.ietf.org/html/rfc4271

[25] G. Huston, "The BGP Instability Report," July 2015. [Online]. Available: http://bgpupdates.potaroo.net/instability/bgpupd.html

[26] D. T. Kaplan, "Exceptional events as evidence for determinism," *Physica D: Nonlinear Phenomena*, vol. 73, no. 1, pp. 38–48, 1994.

[27] T. Gautama, D. P. Mandic, and M. M. Van Hulle, "The delay vector variance method for detecting determinism and nonlinearity in time series," *Physica D: Nonlinear Phenomena*, vol. 190, no. 3, pp. 167–176, 2004.

[28] J.-P. Eckmann, S. O. Kamphorst, and D. Ruelle, "Recurrence plots of dynamical systems," *Europhys. Lett*, vol. 4, no. 9, pp. 973–977, 1987.

[29] N. Marwan and J. Webber, CharlesL., "Mathematical and Computational Foundations of Recurrence Quantifications," in *Recurrence Quantification Analysis*, ser. Understanding Complex Systems, C. L. Webber, Jr. and N. Marwan, Eds. Springer International Publishing, 2015, pp. 3–43. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-07155-8_1

[30] J. P. Zbilut and C. L. Webber, "Embeddings and delays as derived from quantification of recurrence plots," *Physics Letters A*, vol. 171, no. 3, pp. 199 – 203, 1992. [Online]. Available: http://www.sciencedirect.com/science/article/pii/037596019290426M

[31] N. Marwan, "Commandline Recurrence Plots," September 2013. [Online]. Available: http://tocsy.pik-potsdam.de/commandline-rp.php

[32] N. Marwan, "CROSS RECURRENCE PLOT TOOLBOX 5.18 (R29.3)," July 2015. [Online]. Available: http://tocsy.pik-potsdam.de/CRPtoolbox/

[33] A. Toonk, "Massive route leak causes Internet slowdown," BGPMON, June 2015. [Online]. Available: http://www.bgpmon.net/massive-route-leak-cause-internet-slowdown/