Experimental Observations and Analysis of BGP Route Flapping Dynamics

Karina Cereceda*, Jason But† and Philip Branch‡
School of Science, Computing and Engineering Technologies
Swinburne University of Technology
Melbourne, Australia
*ncerecedacastellani@swin.edu.au, †jbut@swin.edu.au, ‡pbranch@swin.edu.au

***Abstract*—The Border Gateway Protocol (BGP) is responsible for updating inter-domain routing information. One of the poorly understood aspects of BGP behaviour that has a direct impact on Internet stability is route flapping, where a route is repeatedly advertised but then withdrawn. This behaviour leads to slow convergence times resulting in Autonomous Systems (ASes) having inconsistent views of the Internet. Route flapping is an important matter that needs to be addressed to ensure the stability of the Internet as it continues to grow. However, there is currently no satisfactory solution to prevent or mitigate the effects of route flapping. Before a satisfactory solution can be developed, a better understanding of how route flapping originates and propagates is needed. In this paper, we used our large network testbed to construct a network with twenty ASes. We were able to generate route flapping and demonstrated that it can be caused by a congested link between ASes. We demonstrated that the frequency of route flapping is proportional to the level of congestion. We also showed that route flaps propagate across the entire network. As part of this research, we developed an automation tool to construct medium-to-large BGP testbeds where all interconnected ASes are placed under a common out-of-band management domain, providing a single vantage point to observe the behaviour and propagation of routing flaps across affected BGP peers. As far as we are aware this is the first time route flapping has been generated and analysed using a large scale testbed of commercial equipment. However, others have carried out research using simulation and modelling. We found our results consistent with theirs.**

***Index Terms*—BGP, BGP oscillations, route flapping, BGP testbeds, link congestion.**

I. Introduction

The Border Gateway Protocol is the defacto routing protocol for inter-domain connectivity among the more than 75,000[1] Autonomous Systems (ASes) actively exchanging traffic today. BGP responds to topology changes by propagating updated reachability information from the origin point of the event, and across the Internet. The efficiency of BGP mechanisms, including path selection and routing updates, have a global effect on Internet stability. Moreover, we believe that the scalability of these processes is a point of concern given the growing nature of the Internet. For these reasons, BGP dynamics is an important area of ongoing research.

Wang et al. [1] note that the two most prominent issues in BGP dynamics are the rapidly increasing rate of BGP updates (*BGP churn*) due to increasing route table entries, and slow convergence (*path exploration*) inherent to the path vector nature of BGP. Several active and reactive solutions have been researched to either reduce churn or accelerate path exploration. However, only two approaches have been implemented by router vendors, Route Flap Damping (RFD) and Minimum Route Advertisement Interval (MRAI) [1].

RFD was designed to mitigate the effects of "route flapping", a phenomenon where routes are repeatedly withdrawn and announced in a short time scale [2]. RFD helps reduce BGP churn derived from route flapping by actively limiting the propagation of unstable routes. MRAI, on the other hand, sets a minimum time between advertisements and withdrawal of routes, regardless of their stability [3], thus rate-limiting BGP updates overall.

In this paper we study the behaviour and propagation of oscillating routes using an experimental approach. To this end, we deployed medium-to-large scale BGP testbeds, using Cisco Routers, and generated traffic congestion scenarios that resulted in route flapping. In addition, all BGP peers shared a common out-of-band (OOB) management network, facilitating real time collection of synchronised BGP update data throughout. This testbed design allowed us to observe BGP dynamics across multiple ASes under a single management domain, giving us a rare vantage point that could not be easily replicated on production networks. We were able to deliberately generate route flapping events in our controlled testbed environment and explore the propagation parameters of BGP oscillations. Also, for congestion-induced oscillations, we were able to demonstrate a direct relationship between congestion ratio and frequency of route flapping.

The rest of this paper is organised as follows. Section II establishes the relevance of further research into BGP route flapping, as well as its causation and mitigation. Section III describes our experimental approach to observing BGP dynamics under link congestion scenarios. Section IV presents the experimental results and observations. Section V concludes this paper and discusses future work.

II. Background and Related Work

The increasing rate of BGP updates, known as BGP churn, has been characterised as one of the major points of concern when it comes to BGP performance [1]. BGP anomalous behaviours often trigger route oscillations, which not only

[1]"CIDR Report," www.cidr-report.org. https://www.cidr-report.org/as2.0/

cause reachability issues, but also contribute to increasing BGP churn. Furthermore, because oscillations can propagate far beyond the point of origin, affected parties several hops away are unable to correlate oscillations to a root cause event, and are therefore unable to take corrective actions.

BGP route oscillations are a common occurrence across the Internet. However, as noted, only two of many researched mitigation mechanisms have been implemented by router vendors, i.e. RFD and MRAI [1]. Unfortunately, both these mechanisms delay BGP convergence. Consequently, the RIPE routing working group recommended against using RFD in 2006 [4], with a 2020 network tomography study showing that only 9% of measured ASes enable RFD [5]. In addition, a 2014 survey showed that most operators disable MRAI [3]. For these reasons, we contend that route flapping is a largely unaddressed problem. As the Internet continues to grow, this problem raises questions of scalability and stability, hence it needs to be better understood as a step towards developing suitable mitigation techniques.

In the remainder of this section, we describe a taxonomy of BGP anomalies to contextualise the specific type of anomalous behaviour we have chosen to explore. We also discuss how link congestion affects the lifetime of BGP sessions, potentially resulting in frequent route flapping. Finally, we discuss related work and observe that we still do not have suitable techniques to mitigate this problem.

A. BGP Anomalies Classification

Al-Musawi et al. [4] classified BGP anomalies into four categories. Anomalies derived from malicious events are classified as *direct intended* anomalies or *indirect* anomalies. Direct intended anomalies are the result of BGP interactions purposely staged to trigger undesired routing behaviors. Indirect anomalies result from malicious activity targeting Internet components, causing collateral BGP instability. Two additional categories are defined to classify anomalies resulting from non-malicious events, *direct unintended* anomalies and *link failures*. Direct unintended anomalies refer to BGP misconfigurations or incompatibilities between policies. Link failures refer to peering link disconnections.

Link failures are often associated with cable breakages, blackouts or natural disasters [4], [6]. However, in the early stages of our experimental work, we observed that link congestion led to peering disconnections that trigger BGP oscillations. Given the similaraities in behaviour, we believe congestion can also be classified as a type of link failure anomaly. As congestion events are a common occurrence in production networks, we found it highly relevant to study route flapping in the vicinity of congested links and beyond.

B. BGP Session Lifetime Under UDP Congestion

BGP peers maintain Keepalive and Hold timers to manage the session status [7]. KEEPALIVE messages are sent and received at a fixed time interval and, upon receiving a KEEPALIVE message, peers reset the Hold timer and the session remains active. However, if enough consecutive KEEPALIVE messages are lost so as to cause the Hold timer to expire, the peers will reset the BGP session. BGP session resets are followed by several processes that can potentially destabilize the routing plane, which in turn can result in harmful behaviours, including route flapping.

Xiao et al. [7] used empirical studies and approximate models to characterise the packet drop probability under different queuing mechanisms, as well as under different congestion scenarios. They also suggest that the probability of successfully receiving a KEEPALIVE message is inversely related to the packet drop probability. As the probability of dropping KEEPALIVE messages increases, so does the number of session resets, Therefore, the packet drop probability on the peering link largely determines the lifetime of a BGP session. We further suggest that the packet drop probability also determines route flapping probability and frequency pattern.

In this paper, we will quantitatively correlate our results to Xiao's findings. More specifically, we expect to observe through our results that, as suggested in [7], as the UDP link saturation increases, the packet drop probability also increases and the lifetime of BGP sessions tend to shorten. We propose that the latter also means that at high congestion ratios, resulting in very high packet drop probability, all BGP KEEPALIVE messages are dropped causing periodic route oscillations. Furthermore, we contend the uptime of these periodically oscillating routes is determined by the value of the MRAI and Hold timers.

C. Other Related Work

As stated above, BGP recurrent route oscillations are known to be a common occurrence across the Internet, and a contributing factor to increasing BGP churn, slow convergence and "noisy" BGP background traffic. Several studies support the latter, including a 2017 analysis of two well-known BGP data repositories [8]. Al-Musawi et al. [8] concluded that most BGP traffic consists of updates and withdrawals that do not relate to network management goals. Instead, most BGP traffic is the result of recurrent oscillations, at different frequencies, from different ASes.

The literature suggests that an effective technique to reduce BGP oscillations is yet to be developed. Despite extensive research [1], only two mechanisms have been implemented by router vendors, RFD and MRAI. However, both techniques failed to meet their design goals, and are known to slow BGP convergence [9], [10]. Further research has been done into RFD and MRAI following their commercial implementations in an attempt to find and propose optimised implementations, [1], [11], [12]. Regardless, recent studies [3], [5] show that these mechanisms are rarely used.

We contend that in order to develop a practical solution to this problem, we first must gain a better understanding of how BGP oscillations originate, behave and propagate. Furthermore, we propose that an experimental approach to studying route flapping can close an existing research gap, as most studies in this area take a simulation or modelling

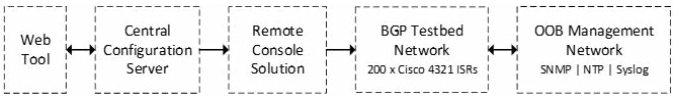

Fig. 1. Automated BGP Testbed System Modules

approach, or a combination of both rather than an experimental approach using commercial hardware.

III. BGP Testbeds Under Link Congestion

For the purpose of studying BGP dynamics, we developed an automated system that allows us to deploy medium-to-large scale testing environments, with up to 200 interconnected ASes, managed from a central configuration server. We developed a web tool to input topology, BGP and internal gateway specifications. Based on these specifications, the central configuration server auto-generates configuration files for all BGP peers, and deploys them through a commercial remote console solution. Additionally, all testbed devices have an out-of-band connection to a network management module that enables the real time collection of BGP data across all ASes in a synchronised fashion. Fig. 1 depicts the different modules in the automated BGP testbed system.

The remainder of this section describes a network topology built with the automated BGP testbed system to observe the behaviour and propagation of BGP oscillations under link congestion scenarios. We will also describe the network management services and congestion generation mechanisms implemented during our experiments.

A. BGP Testbed Network Topology

As a first step in studying BGP oscillations, we built a medium-size testbed, with a twenty-AS span, where BGP peers interconnect via point-to-point links in a physical topology pattern that repeats every four ASes as observed in Fig. 2. We will refer to a four-ASes grouping as a *pod*. Each pod is comprised of four Cisco 4321 Integrated Services Routers, interconnected using 2 Mbps serial links as shown in Fig. 2. Individual pods are interconnected using 1 Gbps Ethernet links.

For all experiments, RFD was not enabled and the BGP Keepalive, Hold and MRAI timers were kept at Cisco default values as per Table I. Each AS was configured to originate one or more internal gateway networks and advertise them to all BGP neighbours. Further, default path attributes were maintained, hence the shortest AS path available was always preferred.

TABLE I
Cisco Default BGP Timers

BGP Timer	Cisco Default Value[a]
Hold Timer	180 secs
Keepalive Timer	60 secs
MRAI Timer	30 secs

[a]Default values maintained for all BGP adjacencies.

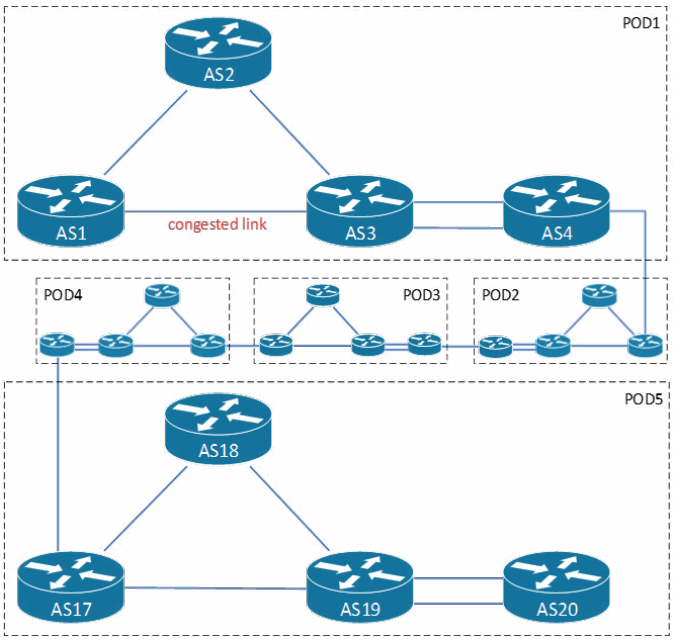

Fig. 2. BGP testbed network using five pods for a span of twenty ASes

Our system generated configuration files that assigned numbered router hostnames, AS numbers and BGP router IDs in increments of one, starting from user-entered initial values, and in the order represented by the AS labels in Fig. 2. The initial values set for this BGP testbed network are shown in Table II and should be used as reference when reading topology information. For example, the BGP peer at AS4, has its hostname set to *Router4* and its local AS number set to *65004*. Also, when the BGP topology is fully converged, AS4 will have two AS paths in its BGP table for a route originated from AS1, *65003 65001* and *65003 65002 65001*.

TABLE II
BGP Peers Parameters

BGP Peer Parameter	Value for ASx
Hostname	`RouterX`
AS Number	`6500X`
Router ID	`1.1.1.X`
Internal Gateway Network	`172.16.0.X/29`
Point-to-Point Peering Networks	`10.0.0.X/31`
Out-of-band Management IP	`192.168.0.10X/24`

Note: IP addressing is not relevant to the discussions in this paper, but included for documentation purposes.

B. OOB Management Network

As previously mentioned, in order to observe and collect synchronised real time BGP data from a single vantage point, all network devices in our BGP testbed connect to a management network module. We used 1 Gbps out-of-band Ethernet interfaces to establish these management connections. The components in the OOB management network module are listed below along with a brief description of how they were used during our experiments.

- Kiwi Syslog Server[2] : a Syslog monitoring tool used to centrally view and maintain system logs generated by BGP testbed network devices.
- PRTG Monitor[3] : an SNMP-based network manager used to gather and plot interface bandwidth utilization and detect interface disconnections.
- Meinberg Time Server[4] : an NTP-based tool used to ensure that the system time is tightly synchronised across all BGP testbed components.
- Cisco EEM[5] : an IOS-embedded applet used to generate custom system logs triggered by route change events as a mechanism to track BGP oscillations.

C. Link Congestion Mechanism

Once the BGP testbed network was deployed, fully converged and monitored under a single management domain, the next step was to implement a link congestion mechanism to stress the bandwidth between AS1 and AS3 (see Fig. 2). We deployed the open source tool *ntttcp*[6] to generate a controlled stream of traffic from the internal gateway network on AS1 to the internal gateway network on AS3. Traffic was generated at increasing bit rates to the point where BGP oscillations were observed at AS4 and beyond.

Multiple experiments were carried out at varying link congestion ratios by modifying both the serial peering link capacity by adjusting the clock rate, and the traffic forwarding rate by adjusting the traffic generator parameters. We will discuss results drawn from three datasets, collected under varying traffic conditions as described in Table III.

TABLE III
LINK CONGESTION SCENARIOS

Dataset ID number	Total Time	*ntttcp* UDP Traffic Range [Mbps]	*ntttcp* UDP Traffic Steps [Mbps]
202409131100	16 hrs.	5-20	+1 hourly
202409160700	16 hrs.	20-35	+1 hourly
202409181750	7 hrs.	1-64	1, 6, 8, 12, 20, 36, 64 (increased hourly)

Note1: AS1 to AS3 link bandwidth at 2 Mbps for all datasets.
Note2: UDP traffic with avg. frame size of 128 bytes for all datasets.

IV. RESULTS

We performed fourteen experiments where we instigated varying UDP congestion ratios on the link between AS1 and AS3, while tracking route changes for AS1's internal gateway network. Route change events were captured on all ASes beyond the congested link, i.e. AS4 through to AS20. It was evident from the very early stages that, above a certain congestion ratio, we could successfully originate a BGP oscillation that propagated across all ASes beyond the origin point. As our experimental work continued, we observed oscillation behaviours consistent with expectations derived from the literature. We believe that this validates the assertion that our BGP testbed system produces meaningful datasets for the purpose of studying the behaviour of route oscillations.

[2]https://www.solarwinds.com/kiwi-syslog-server
[3]https://www.paessler.com/prtg
[4]https://www.meinbergglobal.com/english/info/ntp.htm
[5]https://www.cisco.com/c/en/us/products/ios-nx-os-software/ios-embedded-event-manager-eem/index.html
[6]https://am.net/lib/tools/Microsoft/Tools/NTttcp/ntttcp.htm

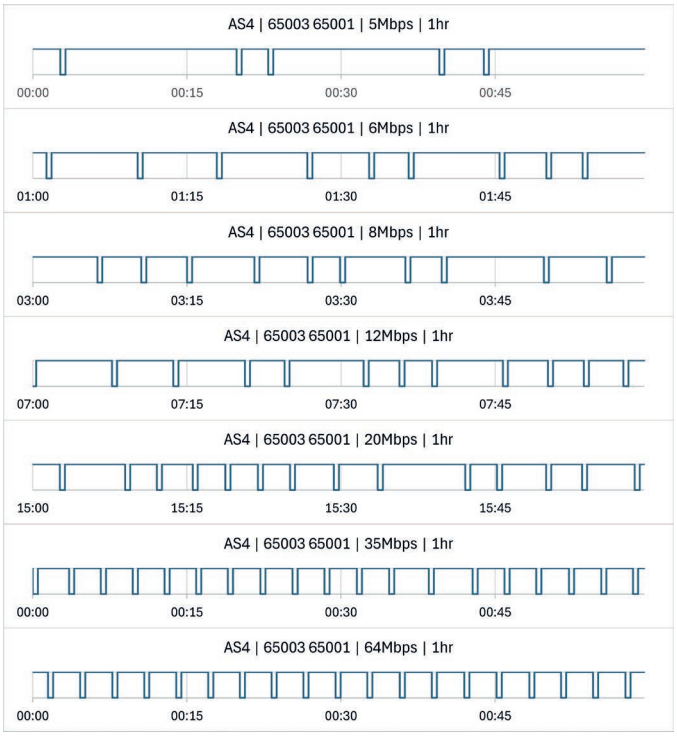

Fig. 3. BGP oscillations observed on AS4 at increasing congestion ratios

In the following three subsections we present our results. In subsection A, we discuss how the oscillation pattern and average route uptime respond to increasing congestion ratios. We then discuss the oscillation propagation span in subsection B. Finally, in subsection C, we discuss the propagation delay at increasing AS distances.

A. BGP Oscillation and Congestion Ratio

To observe the response of a BGP oscillation to increasing congestion ratios, we analysed data from all three datasets listed and described in Table III. More specifically, we studied route change events at AS4 related to the internal gateway network advertised by AS1. Fig. 3 plots the oscillating behavior of this route via the shortest AS path (*65003 65001*) as observed at AS4. Each graph plots the behaviour for different levels of congestion traffic.

Additionally, we studied the length of uptime periods in the oscillating route to find the correlation between this characteristic and the congestion ratio. The uptime period length for a single route flap cycle was calculated as the time differential between a route "down" event and the previous route "up" event. We calculated the uptime period length for all full-cycle oscillations within 1-hour periods at increasing

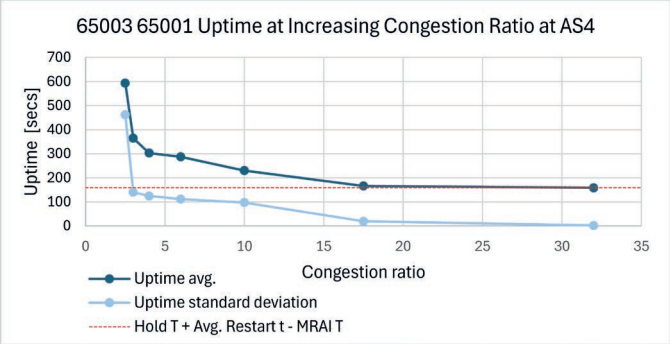

Fig. 4. Oscillation uptime observed on AS4 at increasing congestion ratios

congestion ratio steps. We then calculated the average uptime period length for each hourly step.

Finally, in order to study the response of the route flapping pattern to increasing congestion ratios, we calculated the standard deviation of uptime period lengths at each hourly step. A lower standard deviation implies greater consistency in the measured route uptimes. As the standard deviation approaches zero, the length of the uptime period tends to be the same for all oscillation cycles, and the route flap pattern becomes consistent.

Fig 4 plots both the average uptime length and uptime length standard deviation of the oscillating route observed at AS4 as the congestion ratio increases. The results shown in Fig. 3 and Fig 4 are qualitatively consistent with Xiao's findings [7] and other expectations derived from our literature review.

Firstly, according to [7], we should expect that as the congestion ratio on the link between AS3 and AS1 increases, the KEEPALIVE message drop probability also increases, resulting in more frequent BGP session resets. We observed that these BGP resets between AS3 and AS1 do occur, and originate a BGP oscillation at AS4 that we see in Fig. 3. Moreover, Fig. 3 suggests that the oscillation period of these congestion-induced route flaps become shorter and more consistent as the congestion ratio increases.

Secondly, in [7] Xiao et al. characterised the lifetime of BGP sessions under UDP saturation, concluding that as the packet drop probability increases with the congestion ratio, the lifetime of BGP sessions decrease faster than exponentially. Fig 4 is qualitatively consistent with Xiao's conclusion. To make this assertion, we must consider that the uptime period length of a route flap observed at AS4 is directly related to the lifetime of the BGP session between AS3 and AS1. This is because, neglecting processing and transmission times, the route "down" event happens at the BGP session reset time, and the route "up" event happens after the fixed MRAI time. The latter is true when the BGP session is re-established before the MRAI timer elapses, which was always the case during our experiments.

Finally, our results support that UDP link saturation at high ratios induces a very high packet drop probability, causing all KEEPALIVE messages to be lost as suggested in [7]. Furthermore, it was our contention that for this scenario we would find a direct correlation between the oscillation pattern and BGP MRAI and Hold Timers. This correlation has been demonstrated in Fig. 4, where uptime periods of oscillating routes converge to a persistent length that can be approximated as per Eq (1), where $Hold_T$ is the value of the Hold timer that must elapse for the BGP session between AS1 and AS3 to reset and withdraw the route, $AvgRestart_t$ is the average time that it takes for the BGP session to re-establish, and $MRAI_T$ is the time interval that AS3 must wait after the session has been re-established and before advertising the previously withdrawn route.

$$Hold_T + AvgRestart_t - MRAI_T \tag{1}$$

Additionally, the consistency of the oscillation pattern at high congestion ratios is demonstrated by the standard uptime deviation shown in Fig. 4, which approaches zero as the congestion ratio increases. The latter is also supported by the periodically consistent oscillation observed in Fig. 3, for the last congestion traffic step.

B. BGP oscillation propagation span

To observe the propagation behaviour of BGP oscillations, we studied route change events for AS1's internal gateway network on all ASes beyond the congested link, and at the same *ntttcp* traffic step. This information was retrieved from dataset ID 202409160700 (see Table III) and plotted in Fig. 5. Our results show that a congestion-induced BGP oscillation propagates from the origin point across the entire span of the BGP testbed network, with the same oscillation pattern being observed at every AS. The same outcomes were observed for all configured congestion ratios.

If we consider that both persistent and transient link congestion scenarios are ordinary events across the Internet, these results suggest that congestion-induced BGP oscillations are common occurrences under seamlessly harmless conditions, contributing to BGP background traffic that, as described by [4], does not relate to network management goals.

C. BGP oscillation propagation delay

To further understand the behaviour of flapping BGP routes, we studied the propagation delay of an oscillating route across the full propagation span. We calculated the time differential of route "up" events at increasing AS distances, where AS4 is at 1-AS distance from the origin point and is used as the zero time reference. We considered all route "up" events for AS1's internal gateway network detected within 1-hour periods for all *ntttcp* traffic steps in dataset ID 202409181750 (see Table III). The results of this study are summarised in Fig. 6 and are discussed below.

As shown in Fig. 6, the mean propagation delay increases with AS distance. This is expected considering that each BGP peer along the path introduces transmission and processing delays when forwarding BGP updates. The non-linearity of the upwards trend is attributed to the fact that BGP peer links were

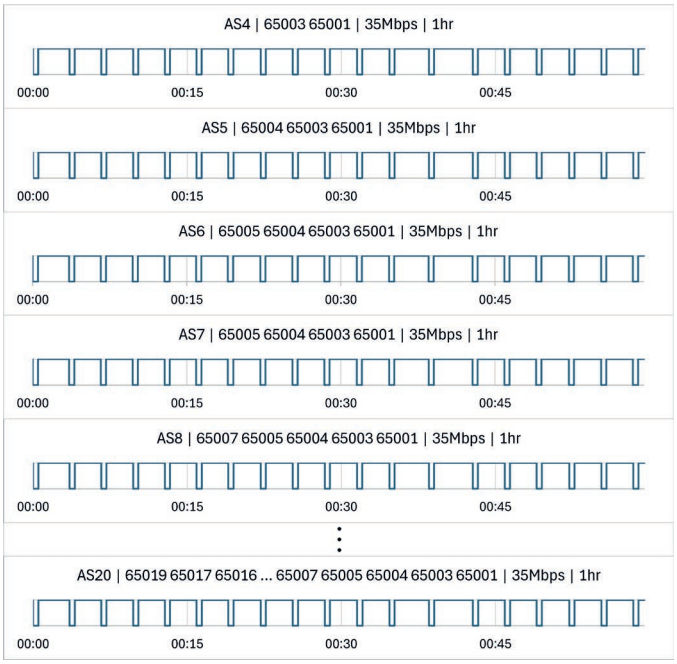

Fig. 5. Singular BGP oscillation observed on multiple ASes

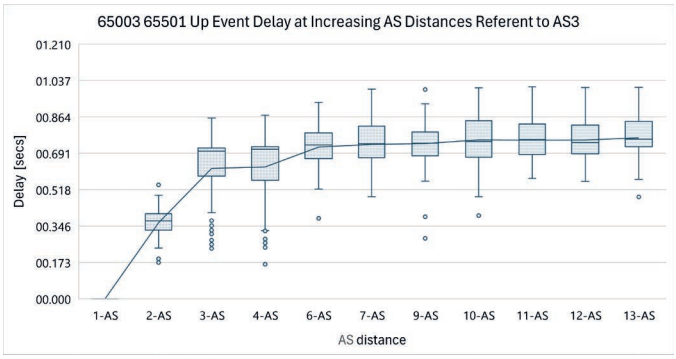

Fig. 6. Route "up" event delay at increasing AS distances referent to AS3

of varying speeds, hence introducing varying transmission delays.

V. Conclusions and Future Work

We have successfully generated congestion-induced oscillations, in an automated BGP testbed environment, and at varying UDP saturation conditions. Also, with the implementation of a common out-of-band management network, we were able to collect meaningful BGP traffic datasets across twenty ASes, from a single vantage point, and in a synchronised manner. The analysis of our experimental data generates results consistent with those from modelling and simulation analyses reported elsewhere in the literature.

In particular, our results demonstrate that BGP oscillations caused by UDP saturation on peering links are consistent with Xiao's findings [7]. Also, as we expected based on [7], at high congestion ratios oscillations reach a persistent, periodic pattern that directly correlates with the value of the MRAI and Hold Timers. In terms of propagation, we found that the propagation delay of a route oscillation increases with distance, however the periodic pattern does not change. Finally, our results show that congestion-induced oscillations propagate across all ASes beyond the origin point, suggesting that this phenomenon contributes to the high BGP churn and slow convergence documented in the literature.

Future work includes reconciling experimental data with our theoretical understanding of how BGP oscillations occur, in order to characterise the findings presented in this paper as it relates to the propagation delay of flapping routes, their response to varying UDP saturation, and the correlation between BGP timers and oscillation patterns. To this end, designing an improved route change tracking mechanism must also be considered for future work, as the current mechanism introduces an unpredictable processing delay, hence introducing an non-quantifiable error to propagation delay calculations in the order of milliseconds.

To conclude, we would like to highlight that our BGP testbed system has generated datasets that are consistent with current knowledge. We propose that using this methodology for further experimental research can yield new knowledge and a better understanding of BGP oscillations and BGP dynamics in general.

References

[1] X. Wang, O. Bonaventure, and P. Zhu, "Stabilizing bgp routing without harming convergence," in *2011 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2011, pp. 840–845.

[2] W. Lijun, W. Jianping, and X. Ke, "Modified flap damping mechanism to improve inter-domain routing convergence," *Computer communications*, vol. 30, no. 7, pp. 1588–1599, 2007.

[3] A. García-Martínez, P. R. Torres Jr, and M. Bagnulo, "Bgp convergence in an mrai-free internet," *Computer Networks*, vol. 240, p. 110183, 2024.

[4] B. Al-Musawi, P. Branch, and G. Armitage, "Bgp anomaly detection techniques: A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 377–396, 2016.

[5] C. Gray, C. Mosig, R. Bush, C. Pelsser, M. Roughan, T. C. Schmidt, and M. Wahlisch, "Bgp beacons, network tomography, and bayesian computation to locate route flap damping," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 492–505.

[6] T. B. Paiva, Y. Siqueira, D. M. Batista, R. Hirata, and R. Terada, "Bgp anomalies classification using features based on as relationship graphs," in *2021 IEEE Latin-American Conference on Communications (LATINCOM)*. IEEE, 2021, pp. 1–6.

[7] L. Xiao, G. He, and K. Nahrstedt, "Understanding bgp session robustness in bandwidth saturation regime," 10 2004.

[8] B. Al-Musawi, P. Branch, and G. Armitage, "Recurrence behaviour of bgp traffic," in *2017 27th International Telecommunication Networks and Applications Conference (ITNAC)*, 2017, pp. 1–7.

[9] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz, "Route flap damping exacerbates internet routing convergence," in *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, 2002, pp. 221–233.

[10] R. Gill, R. Paul, and L. Trajković, "Effect of mrai timers and routing policies on bgp convergence times," in *2012 IEEE 31st International Performance Computing and Communications Conference (IPCCC)*, 2012, pp. 314–323.

[11] E. A. Alabdulkreem, H. S. Al-Raweshidy, and M. F. Abbod, "Mrai optimization for bgp convergence time reduction without increasing the number of advertisement messages," *Procedia Computer Science*, vol. 62, pp. 419–426, 2015.

[12] C. Pelsser, O. Maennel, P. Mohapatra, R. Bush, and K. Patel, "Route flap damping made usable," in *Passive and Active Measurement: 12th International Conference, PAM 2011, Atlanta, GA, USA, March 20-22, 2011. Proceedings 12*. Springer, 2011, pp. 143–152.