Department of Information and Communication Engineering
Graduate School of Information Science and Technology
THE UNIVERSITY OF TOKYO

Master Thesis

Detecting BGP zombies and inferring their source

(BGP zombies の発見及びその原因に関する研究)

Yutaro Nomura
野村 祐太朗

Supervisor: Professor Hiroshi Esaki
Associate professor Hideya Ochiai

January 2019

Abstract

Border Gateway Protocol (BGP) is used to exchange reachability information across Autonomous Systems (ASes). When an IP prefix is withdrawn, withdraw messages are sent from an origin AS and all BGP routers remove the entry of the prefix in the routing table. However there are cases where BGP routers manage routes to withdrawn IP prefixes being reported. We refer to this problem as BGP zombies and examine their characteristics. We use route information of peer ASes collected by RIS or Route-Views projects. There is a very strong possibility that at least one peer AS manages a zombie route when a prefix is withdrawn (more than 90% for IPv4). But the impact of zombie outbreaks is limited. We observe 75% of peer ASes have zombie entries for less than 2.3% of zombie prefix withdraws. We also discover the source of zombie outbreaks and show the benefits for network operators. We compare the number of paths withdrawn as expected (normal) and the number of paths not withdrawn unusually (zombie). When a transit AS become a zombie AS, the AS causes a lot of zombie paths and has an large influence on its downstream ASes.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Background

The Border Gateway Protocol (BGP) is the protocol that manages inter-domain routing across the Internet. In the world of BGP, packets are directed between an Autonomous System (AS). Each BGP router maintains a BGP routing table used to send packets to the Internet. BGP makes routing decisions based on paths. When the network undergoes changes, BGP routers exchange update messages toward other ASes. It is important to announce or withdraw address space collectlly and forward the information toward the Internet.

When a prefix is withdrawn in an origin AS, withdraw messages of the prefix are forwarded to its neighbour ASes. The neighbour ASes that receive a withdrawal also propagate it to their neighbors. An AS does not always propagate withdraw messages to its neighbors. When the AS withdraws the best path but there are still alternative paths in the routing table, its neighbors do not receive withdrawals but receive the best alternative path. This process, called path hunting, causes several BGP path changes in a short period of time before all paths about a prefix is completely withdrawn. Theoretically this withdrawal process ends with the prefix completely withdrawn from all BGP speakers, as announcements and withdrawals propagate through the entire Internet similarly. In practice this sometimes fails, a phenomenon known by network operators as *stuck routes* or *zombie routes*. In this case some BGP routers still have route information of the prefix in their routing tables though the prefix is actually withdrawn. We can easily detect these inconsistents with route collector systems like RIS or Route-Views [1, 2].

1.2 Objective

We are motivated by the operational confusions caused by missing withdrawals. We have observed some cases where zombie routes caused confusions about the state of the withdrawn address space. It is hard for network operators to detect the source of zombie routes and clean them. BGP zombies are not familiar to network operators and generally not well understood. This work can help network operators for troubleshooting or resolving zombie routes.

1.3 Approach

We classify prefixes into two category, beacon prefixes and wild prefixes. A BGP beacon is a BGP speaker that announces and withdraws a particular prefix at predetermined time intervals.

We construct a detection method by using beacons and apply the algorizm to wild prefixes. We collect route information from RIS or Route-Views collectors and examine AS paths to detect zombie routes.

1.4 Contributions

We can show characteristics of zombie routes and the detection method of sources of zombies. Our experiments reveal that BGP zombies are not rare phenomenons. We can observe zombie routes every day in our dataset but the number of affected ASes is usually limited. In addition, we can distinguish transit zombie ASes that cause zombie routes to its downstream ASes from a peer zombie AS that does not forward routes obtained from other network to its neighbor ASes. Some transit zombie ASes are likely to be observed many times in our dataset. We can report failures of routers used in different ASes that we detected as sources of zombie routes.

1.5 Constitution of this thesis

In Chapter 2, we explain basic knowledge of BGP. In Chapter 3, We introduce phenomenons that are related to BGP zombies and works about anomaly detection about BGP. In Chapter 4, we explain about BGP zombies caused by missing withdrawals. In Chapter 5, we show characteristics of BGP zombies and zombie routes. In Chapter 6, we show case studies about detecting of sources of zombies.

Chapter 2

BGP

2.1 The Internet

The Internet is a globally connected network system. To send data packets from one node to other networks, the router in the network the node belongs to has to know paths to other networks. Each router obtain the routing information by routing protocols. Routing protocols enable routers to get path information to other networks dynamically, which is neccessary to keep the Internet.

2.2 BGP overview

2.2.1 IGP and EGP

Routing protocols can be classified into 2 groups by their purpose: IGP (Interior Gateway Protocol) and EGP (External Gateway Protocol). An autonomous system (AS) is a collection of routers under a common administration such as a company or an organization. IGP is used for routing within an AS. It is also referred to as intra-AS routing. Companies, organizations, and even service providers use an IGP on their internal networks. IGP includes RIP, EIGRP, OSPF, and IS-IS. EGP is used for routing between ASes. It is also referred to as a inter-AS routing. Service providers and large companies can interconnect using an EGP. The Border Gateway Protocol(BGP) is the only currently viable EGP and is the official routing protocol used by the Internet.

2.2.2 Autonomous system number(ASN)

A public AS has a globally unique number, an AS Number, associated with the AS. The number is used both in the exchange of exterior routing information and as an identifier of the AS itself. ASN ranges from 1 to 64,511. When an ASN is needed, the next highest unused number is assigned. The American Registory for Internet Numbers [3] manages IP address allocations and assignments. It is also the authority for assigning and tracking ASNs. There are two types of AS Numbers: Public AS Numbers and Private AS Numbers.

2.2.3 BGP session

Each individual AS establishes BGP peering sessions to other AS to exchange routing information. A BGP peering session is a TCP session established between two routers, each one in a

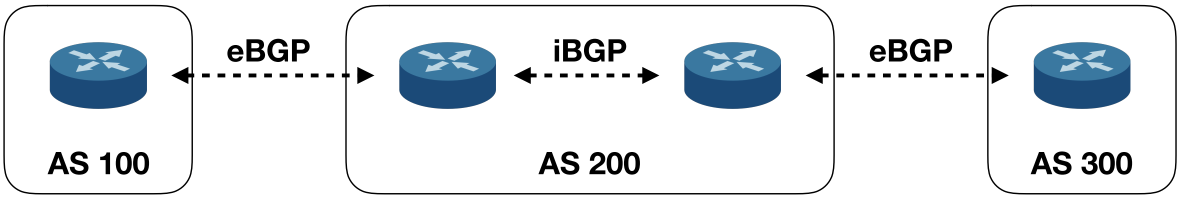

Fig. 2.1. Routers in different ASs communicate using eBGP or iBGP.

particular autonnomous system. This BGP peering session is a TCP session established between Ethernet interface between those routers. The routing infromation contains an IP address prefix and a subnet mask. This translates which IP address are associated with ASN (AS origin). Routing information propagates across these autonomous systems based upon policies that individual networks define. There are two main types of relationships between autonomous systems today: Transit and Peering. Transit is where an autonomous system will pay an upstream network for the ability to forward traffic towards them who will forward that traffic further. Peering is where an autonomous system will connect to another autonomous system and agree to exchange traffic with each other of their own networks and any costomers they have.

2.2.4 iBGP and eBGP

If a BGP session is established between two neighbors in different autonomous systems, the session is external BGP (eBGP), and if the session is established between two neighbors in the same AS, the session is internal BGP (iBGP). Multiple routers usually exist whithin an AS, so iBGP is necessary whenever BGP advertised information must be passed within a given AS. Figure 2.1 illustrates this concepts. In this figure, iBGP runs between routers in AS200 and eBGP runs between a router in AS100 and one in AS200 and between one in AS200 and one in AS300.

2.2.5 AS classification

ASes can be classified into Single-homed(Stub) or Multi-homed ASes. When an AS sends packets toward outside networks with only one exist point, the AS is a Single-homed AS. A single-homed ASes can be referred to as stub ASes. Stub ASes can depend on a default route to manage all traffic toward outside netwroks so BGP is not necessarily for stub ASes. Figure 2.2a shows an exmaple of a single-homed AS (AS10). When an AS have more than one connection to the Internet to increase the reliability of network, the AS is a Mulit-homed AS. One of the benefits of Multi-homing is that an Mutli-homed AS can use alternate path to the Internet in case of the primary link failure because of equipment failure, fiber cuts, maintenance windows and DDos attacks. If a mutli-homed AS does not allow pass on traffics that has a source outside network to different outside network, the AS is a non-transit AS. A non-transit AS advertises only routes information to the ISPs it connects to and do not advertise routes it obtained from one ISP to another. Figure 2.2b shows that AS10 does not transit routes from ISP2 to ISP1. ISP1 can not use AS10 to reach destinations that belong to ISP2 and ISP2 can not use AS10 to send packets to destinations that belong to ISP1. On the other hand a multi-homed transit AS with more than one

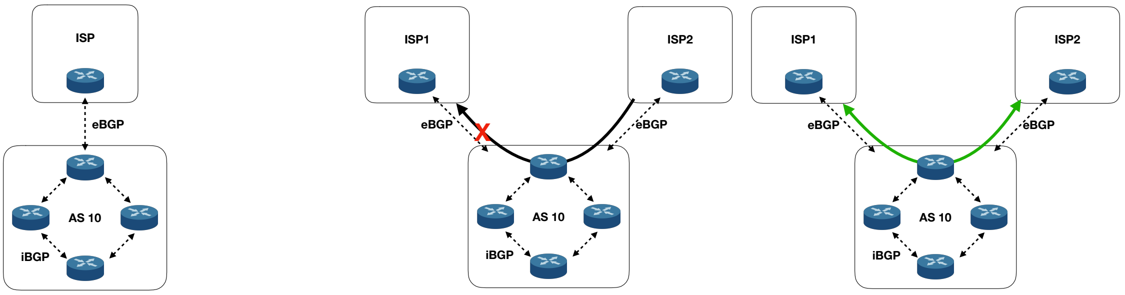

(a) Single-homed AS.

(b) Multi-homed non-transit AS.

(c) Multi-homed transit AS.

Fig. 2.2. AS path length from 2017/4/1 to 2017/4/30

connection to outside network can transit traffic sent from other ISPs in figure 2.2c.

2.2.6 BGP messages

The BGP communication uses four message types. After a TCP connection is established between two BGP systems, they exchange BGP open messages to create a BGP connection between them. Once the connection is established, the two systems can exchange BGP messages and data traffic. BGP systems send update messages to exchange network reachability information. BGP systems use this information to construct a graph that describes the relationships among all known ASes. Update messages consist of the BGP header plus the following optional fields:

Withdrawn Route Length this field shows the length of the Withdrawn Routes field in bytes. When it is set to 0, there are no routes withdrawn.

Withdrawn Routes this field shows IP address prefixes for the routes being withdrawn from service because they are no longer deemed reachable.

Total Path Attribute length Length of the path attributes field: it lists the path attribtues for a feasible route to a destination.

Path Attributes the BGP attributes for the prefix are stored here, for example: origin, as_path, next_hop, med, local preference, etc.

2.2.7 Path attributes

BGP routers usually receive multiple paths to the same destination, so the best path to each destination needed to be selected. IGPs select the path with the lowest metric. For example, RIP selects the path with the lowest hop count. OSPF selects the path with the lowest cost. BGP, however, selects the best path based on path attributes. There are four categories of path attribute.

Well-known mandatory an update packet must include this attribute. Missing this attribute creates a notification error and the BGP session is closed.

Well-known discretionary This attribute must be recognized by all BGP implementations but does not have to be included in each BGP UPDATE message.

Optional transitive This attribute is not required to be recognized by all BGP implementations. It is allowd to be sent to other peers.

Optional non-transitive This attribute is not required to be recognized by all BGP implementations. It is not allowed to be sent to other peers.

Following attributes are used for the best path selection.

1. ORIGIN Well-known mangatory.
   This attribute informs an AS about the originator of the route. There are three defferent types, IGP, EGP and Incomplete. IGP has higher priority than EGP and Incomplete is lowest priority of them.
2. AS_PATH Well-known mandatory.
   This attribute identifies the ASes through which the update message has passed.
3. NEXT_HOP Well-known mandatory.
   This attribute refers to the IP address take to reach the destination.
4. MULTI_EXIT_DISC (MED) Optional non-transitive.
   It is used to select the preferred entry path to an AS when more than one entry points are available for a network. It is an optional attribute, meaning that it is not always sent in the BGP announcements.
5. LOCAL_PREF Well-known discretionary.
   This attribute indicates the degree of preference for one BGP route over the other BGP routes. The BGP route with the highest local preference value is preferred. Local preference attribute is never advertised to external BGP peers.
6. ATOMIC_AGGREGATE Well-known discretinoary.
   This attribute informs that some information have been lost due to the route aggregation process.
7. AGGREGATOR Optional transitive,
   This attribute may be included in update messages sent by routers performing aggregation.
8. COMMUNITY Optional transitive.
   This attribute is assigned to a specific prefix and advertised to other neighbor peers. Neighbors that receive updates messages that have COMMUNITY attributes examine the community value and take proper action whether it is filtering or modifying other attributes. The BGP has 4 well known communities, Internet (advertise routes to all neighbors), Local-as (prevent forwarding routes outside the local AS within the confederation), No-Advertise (do not advertise this route to any peer, internal or external), No-Export (do not advertise this route to external BGP peers).
9. ORIGINATOR_ID Optional non-transitive
   This attribute contains the BGP identifier of the router that originated the route.
10. Cluster List Optional non-transitive
    This attribute contains all IDs of the clusters that a route has propagated through.

The algorithm for determing the best path is as follows:

1. Verify the reachability to the next hop.
2. Select the path with highest local preference value.
3. Prefer the path that originates locally.
4. Prefer the path with the shortest AS path value.
5. Choose the path with the lowest origin type.

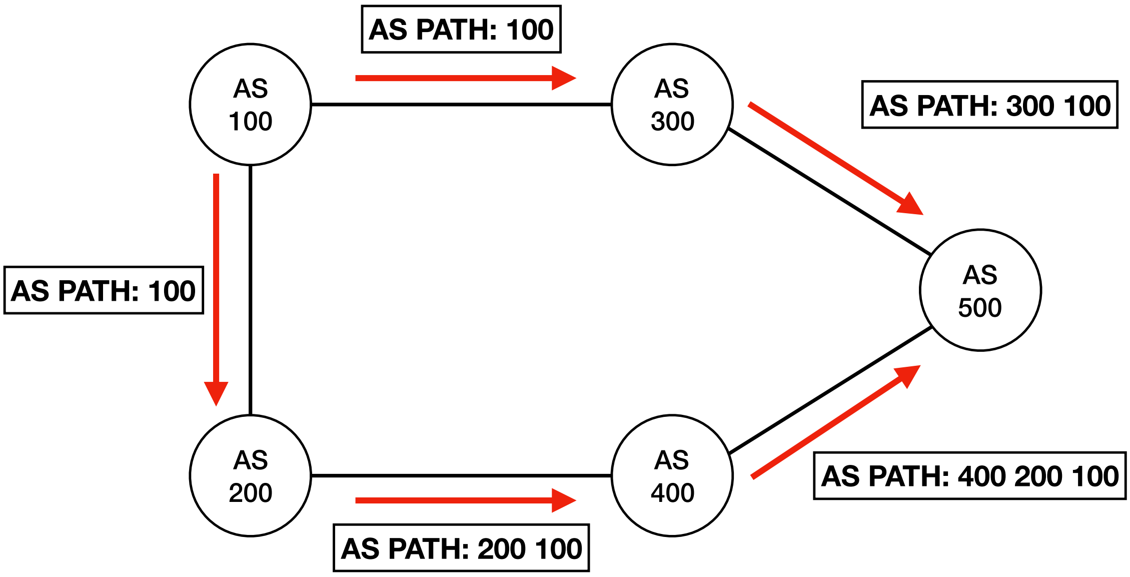

Fig. 2.3. BGP messages for a new prefix propagation..

6. The lowest value of MED is preferred.
7. Prefer eBGP over iBGP.
8. Prefer the path with lowest IGP value to the next hop.
9. The path that is older or was received first is preferred.
10. Prefer the path that originates from a router with the lowest router ID.
11. The path that have the lowest neighbor address is preferred.

Figure 2.3 depicts a propagation of BGP announcements. For example, AS100 obtains a new IP prefix, 1.1.1.0/24, it sends announcements about a new prefix 1.1.1.0/24 in AS 100 to its peering ASes. These announcements mean that the path to the 1.1.1.0/24 is AS 100. AS200 and AS300, peering with AS100, receive the announcements from AS100 and update their routing tables. AS200 and AS300 add their own AS number (200 or 300) to the path information and send this new route information about 1.1.1.0/24 to their peers. AS400 receive the announcement with path: 200 100 and AS500 receive the one with path: 300 100. AS400 updates its routing table and sends new announcements. Finally, AS500 obtains two different routes to the new prefix 1.1.1.0/24, 300 100 and 400 200 100. If AS500 uses the AS path length in the best path selection, AS500 selects the AS path 300 100.

2.2.8 BGP Routing Table

Any BGP speaker receives routing updates from other peers, processes the information for local use and then advertises selected routes to different peers based on predetefined policies. BGP routers store route information in a special type of database called the BGP Routing Information Base(RIB). BGP Routing Information Base consists of three parts as explained below.

The-Adj-RIBs-In This one stores BGP routing information received from different peers. The stored information is used as an input to BGP decision process.

The Local RIB This one stores the resulted information from processing the RIBs-In database's information. These are the routes that are used locally after applying BGP policies and decision process.

The Adj-RIBs-Out This one stores the routing information that was selected by the local BGP router to advertise to its peers through BGP update messages.

Chapter 3

Related Work

3.1 Phenomenon

In this section, we introduce a phenomenon that followed by withdrawing prefixes.

3.1.1 Path Hunting

A withdraw message for one prefix can cause to increase the number of announcements sharply[4]. Path hunting behavor means that when a prefix become unreachable, BGP will explore longer and longer paths before the prefix disappears from routing tables everywhere. In the case of reachablity failure, the notifying AS will generate a BGP withdrawal that is intended to propagate through the network. This propagation of a withdrawal sometimes trigger consequent announcements. When a BGP speaker believes that it has a choice to switch to a less preferred, but still valid path upon receipt of the withdrawal. In this case the BGP speaker involved does not propagate the withdrawal explicity, but will propagate an announcement for the remaining valid path to its BGP peers. This announcement is appropriately interpreted as an implicit withdrawal of its previous announcement and an announcement of a new best path.

Figure 3.1 shows an example of path hunting. When the link between AS1 and AS2 fails, AS2 sends withdrawal W to AS3 and AS5. When AS5 gets withdrawal from AS2, AS5 removes the path {5,2,1} and installs another path {5,3,2,1} since it has three paths in its BGP routing table. AS5 advertizes the new path {5,3,2,1} to all other BGP neighbors. After that, AS3 forwards withdrawals to AS4 and AS5. AS5 propagate the new path {5,4,3,2,1} to neighbor peers. Finally, AS4 sends withdrawals to AS5. AS5 remove the path {5,4,3,2,1} and there is no path in the routing table. AS5 recognizes that AS1 is unreachable and sends withdrawals to its peers at first. During this time, AS5 has propagated announcements of AS1 though AS1 is actually unreachable. Figure 3.2 shows when peers receive announcements and withdrawals of 84.205.64.0/24 from 2017/4/1 to 2017/4/3. 84.205.64.0/24 is one of the beacon prefixes with announcements at 00:00, 04:00, 08:00, 12:00, 16:00, 20:00 (UTC) and withdrawals at 02:00, 06:00, 10:00, 14:00, 18:00, 22:00 (UTC). The vertical axis shows routes that manage the prefix 84.205.64.0/24 and the horizontal axis shows the time. We exclude duplicated messages. Most peers receive update messages at collect time. However, many peers receive announcements at the time they should receive withdrawls in theory. We can guess that path hunting causes this phenomenon. Before all route are completely withdrawn, alternative paths that become the best path because the previous best path is withdrawn are propagated to the Internet.

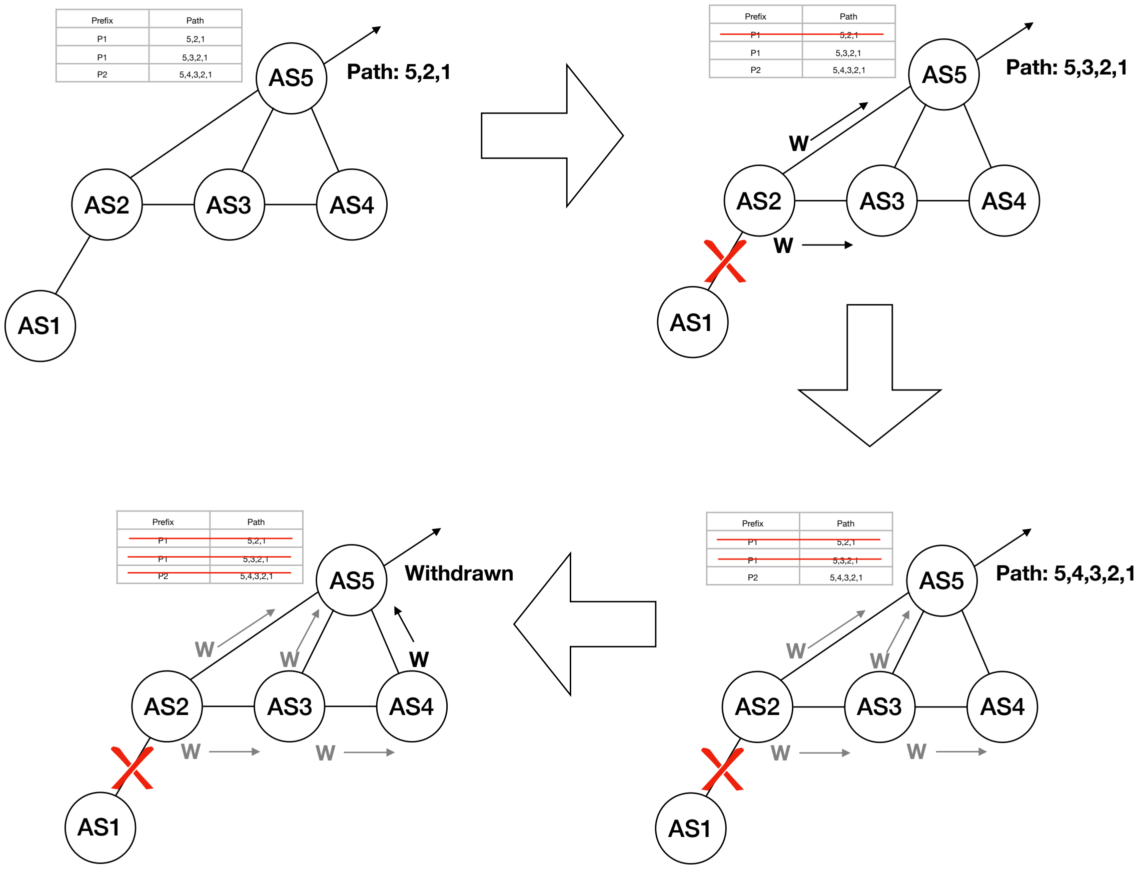

Fig. 3.1. An example of path hunting

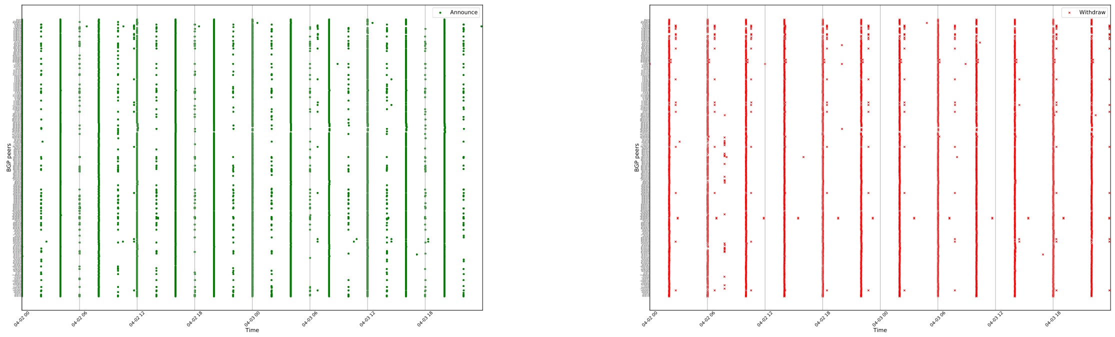

(a) Announcement.

(b) Withdrawals.

Fig. 3.2. update messages for from 2017/4/2 to 2017/4/4

3.2 BGP anomalies

In this section, we introduce BGP anomalies that prevent routers from managing route information and show detection or solution techniques.

3.2.1 Route Flap and damping

BGP Route flap [5] occurs when routes to a destination are advertised and withdrawn in a rapid sequence by a BGP router. A router also advertises an available route and then available one again. Route Flapping is usually caused by hardware or software errors on the router or configulation errors. The continuous advertisements by a BGP router causes its neighbor BGP routers to update their routing tables as well. Affected routers need to recalculate of topology many times and use more CPU and memory, leading to packet loss and performance degradation for traffic traversing the affected networks. Route flap damping is generally assumed by network operators to be widely deployed in today's infrastructure [6].

Route flap damping is a mechanism designed to selectively limit the propagation of unstable routing information [7]. A BGP router maintains a route penalty associated with every prefix announced by its BGP neighbors. This route penalty increments by some fixed value whenever the state of the route changes and exponentially decays with time. The penalty measures the instability of the route. The BGP router decide when to supprss the route on the basis of the locally configured threshold.

However route flap damping can actually exacerbate the convergence of relatively stable routing information (i.e., path hunting). A single withdrawal of a route cause path hunting and route penalty to get beyond the threshold and lead to the suppressed route. A single announcement can cause a suppressed route in the so-called focus [8] topology. In the focus topology, a router(node A) send BGP packets to its neighbors and all neighbors forward packets to one router(node B). Route announcements from node A arrvie at node 7 at different times and path length of the announcements is 2. Node B will also announce to its neighbor peers these different announcements in the order they are received. When neighbor peers receive these announcements in sequence, it can suppresses the route to node A. We can consider a simple solution for these surpression of stable routes caused by a announcement or a withdrawal. One way to find routes changes by path exploration is to avoid penalizing successive routes with non-decreasing path lengths. When the path length of a new route is the same or longer thant the existing route, router does not increment the flap penalty. On the other hand, it is recommended that the needs of BGP flap damping is no longer a major concern beacause of increasement of the power of routers [9]. In fact the cure has become worse than the disease.

3.2.2 Route Oscillations

The route oscillation problem for iBGP was first reported in a Field Notice from Cisco Systems [10]. This problem was reported for both route reflection configurations [11] and confederation configurations [12]. The use fo the Multi-Exit-Discriminator (MED) for route comparison is mainly related to this problem. The MED attribute value is used in configurations where multiple links connect the same AS pair and a router prefer a route with lower MED. Because a router use MED when it select the best path from routes that pass through the same neighboring AS, the presence of absence of a route can change the relative ranking of a different router and cause oscillations. One solution to this problem is a modification to iBGP and the modified protocol provably converges [13].

3.2.3 BGP configuration faults

Network operators configure BGP rotuers to manage the routes that are announced and withdrawn. Configuring a network of BGP routers is very complex because of configuration being distributed across the routers. BGP configuration faults can cause forwarding loops, packet loss and unintended paths between ASes. 0.2%-1.0% of the BGP table entries suffer misconfiguration events each day and misconfigurations increased the route update load by at least 10% for 2% of time [14]. There are many causes of misconfigurations, not only involuntary slips by network operators, errors by router initialization bugs, or a poor understanding of configuration semantics on some operators.

3.3 Anomaly detection

The BGP is the primary protocol for internet service providers to exchange packets between each other's networks. The BGP protocol was formaized in the 1980s and the major revision to this protocol has never made since 1995. So detecting BGP anomalous behavior is important for improving the security and robustness of the Internet. The Internet routing have many problems today, including instability [15], convergence delay. We introduce some monitoring system for exchanges of BGP route information related with BGP zombies.

RIPEstat [16] is a web-based interface that provides everything about IP address spaces, ASes and related information for hostnames and countries in one place. RIPE stat provides a lot of different data from different sources. We can see analyses of global Internet events and compare query results for multiple resources across multiple widgets. RIPEstat looking glass allows us to directly query their route collectors and we can detect zombie outbreaks. We can also accesses the raw data from the RIS archive using BGP stream [17] and check that withdraw messages are indeed missing in the raw traces. NLNOG looking glass [18] is also useful to detect BGP zombies.

Chapter 4

BGP zombies

4.1 BGP zombies

A **BGP zombie** refers to an active Routing Information Base(RIB) entry for a prefix that has been withdrawn by its origin network, hence it is not reachable anymore. In this paper we also refer to **zombie ASes** for ASes whose routers have BGP zombies, Moreover we refer zombie ASes with peering with RIS or Route-Views collectors for **zombie peers**. In addition we refer to all zombies that corresspond to the same prefix as a **zombie outbreak**, the outbreak size is the number of zombie ASes.

4.1.1 Beacon prefix

In order to observe BGP zombies we need to obtain when one prefix is withdrawn and find ASes whose router still have BGP route information about the prefix. However, it is difficult to know about prefix withdraws beforehand and it is hard to ask network operators about its routing policy. So before driving into the detailed analysis of BGP zombies in global network, we conduct experiments by using RIPE's Routing Information Serivice (RIS) BGP beacons [19]. The RIS BGP beacons are a set of IPv4 and IPv6 prefixes that are used solely for studying Internet inter-domain routing. These IP prefixes are announced and withdrawn at predetermined times intervals. Namely, RIS BGP beacons are announced every day at 00:00, 04:00, 08:00, 12:00, 16:00, and 20:00 UTC, and they are withdrawn two hours after the announcements (i.e. at 02:00, 06:00, 10:00, 14:00, 18:00, and 22:00 UTC). We refer to RIS BGP beacons for **beacon prefixs** and refer to not beacon prefixes for **wild prefixes**.

We conduct such controlled experiments with the help of BGP beacons, RIS BGP data registory, and The Oregon Route-Views project. RIS also archives RIB and BGP update messages collected at diverse places on the Internet. RIS collectors (named rrc00, rrc01, etc ...) are mainly locatd at Internet exchange Points (IXP) and peer with hundreds different ASes. Route-Views provides a real-time view of the Internet routing table from perspective of the ASes that provide their routing tables via BGP. Cisco BGP RIBs collected from route-views.routeviews.org (the collector script was written by Sean Mccreary) are available on http://archive.routeviews.org/. There are two data formats; those collected from the Cisco and those from the Zebra routing software. The Cisco format is collected on two hours intervals starting at 00:00. There Zebra files(RIBs and UPDATEs) have different intervals. RIBs are snapshots and are collected every 2 hours. UPDATEs are ongoing files that are rotated every 15 minutes.

Start	End	IPv4 beacons		IPv6 beacons	
		#beacons	#outbreaks	#beacons	#outbreaks
2017-04-01	2017-04-30	14	2488	13	2328
2017-09-01	2017-09-30	16	2642	17	2665
2018-04-01	2018-04-30	16	1434	17	2043
2018-09-01	2018-09-30	16	2852	17	2609

Table 4.1. Measurement periods, number of becons and number of detected zombie outbreaks for beacons.

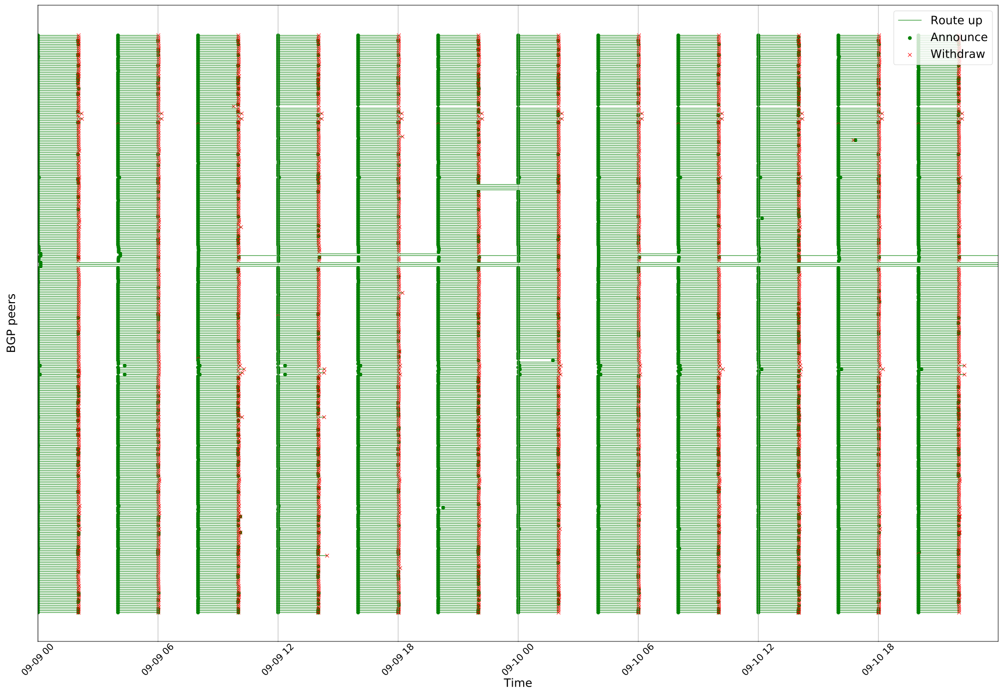

Fig. 4.1. Visibility for 84.205.64.0/24 from all RIS and Route-Views collectors on September 9th and 10th, 2018.

For beacon prefixes, the detection of zombies in RIS and Route-Views peers is straightforward. We keep track of the visibility of beacons for all RIS and Route-Views peers and report a zombie for each RIB entry that is still active while the prefix is withdrawn.

We conducted experiments for beacons during the four periods of time listed in Table 4.1. All beacon prefixes are in the same AS12674 but the probability of zombie outbreaks varies depending on prefix versions, 99% for IPv4 and 85% for IPv6.

4.1.2 Example

Figure 4.1 illustrates the visibility for beacon prefix (81.205.64.0/24) from all RIS and Route-Views peers on September 9th and 10th, 2018. Peers are represented on the y axis and time is represented by the x axis. At 2017/09/09 00:00 UTC RIS and Route-Views peers announced the avalability of the beacon prefix. Most peers withdrawed the prefix at 2017-09-09 02:00 UTC as expected. However, 3 peers did not withdraw the beacon prefix although this prefix was not reachable at that time. Similar zombie outbreaks appeared from 2018/09/09 11:00 to 2017-09-10 23:00 and from 2017-09-10 06:00 for the same peers. Other peers also did not withdraw the beacon prefix from 2017-09-09 23:00 to 2017-09-10 00:00. On the other hand, all peers behaved as expected from 2017-09-09 06:00 to 2017-09-09 08:00. Three zombie peers that failed to withdraw 81.205.64.0/24 from 2017-09-09 10:00 to 2017-09-10 00:00 also failed to withdraw the beacon prefix, 81.205.71.0/24. These observations imply us that the relationship between zombie outbreaks during the same period for different prefixes.

4.2 Detecting Zombie outbreaks

We can detect beacons zombie outbreaks easily because when beacon prefixes are withdrawn is already known. We just examine whether each peer AS obtain a withdraw message when the beacon prefix is withdrawn. However, it is difficult to detect zombie peers of wild prefixes because we cannnot know when wild prefixes are withdrawn beforehand. Regarding wild prefixes, we need to know from when to when the wild prefix is withdrawn to detect zombie outbreaks. If a prefix is withdrawn and everything goes well, all peers receive withdraw messages of the prefix. Montioring the change in the number of peers that have route to the prefix over time can help us to detect the prefix withdraw. However, update messages can reach ASes late and there are ASes that have route to the withdrawn prefix, zombie ASes. A prefix withdraw does not always mean that no peer has a reachability to the prefix. So we set the threshold of the number of peers and when the number of peers that have route informaion of the prefix get less than the threshold, we judge the prefix is withdrawn.

We also need to take into account of how long it takes for an update message to reach each peer to detect withdrawn time accurately. We determine beacons update (announce and withdraw) messages propagation time. Figure 4.2a illustrates the distribution of propagation time of update messages in April 2018. We already know when the beacon prefix is announced and withdrawn, so we can observe the time from the start of the update message propagation to the time each peer receive the update message. Most announcements propagate in the Internet in less than a few minutes for both IPv4 beacon prefixes and IPv6 beacon prefixes. For withdrawals to propagate, it takes longer than announcements because of the path hunting. But most peers withdraw the prefix less than 20 minutes after the prefix is withdranw in its origin AS. Also, the delay of the propagation of some update messages can happen. This delay can be seen in figure 4.2a from 2017-09-09 10:00 to 2017-09-09 15:00. So we set the threshold of the number of peers that can see the prefix. and we determine the period when the number of peers that do not receive the withdrawal gets less than the threshold as the prefix withdraw. The max rate of the number of zombie peers for beacon prefixes is about 20%. So we set the threshold as 25%. Figure 4.2b show this detecting process. Because we set update messages propagation time longer than the actual

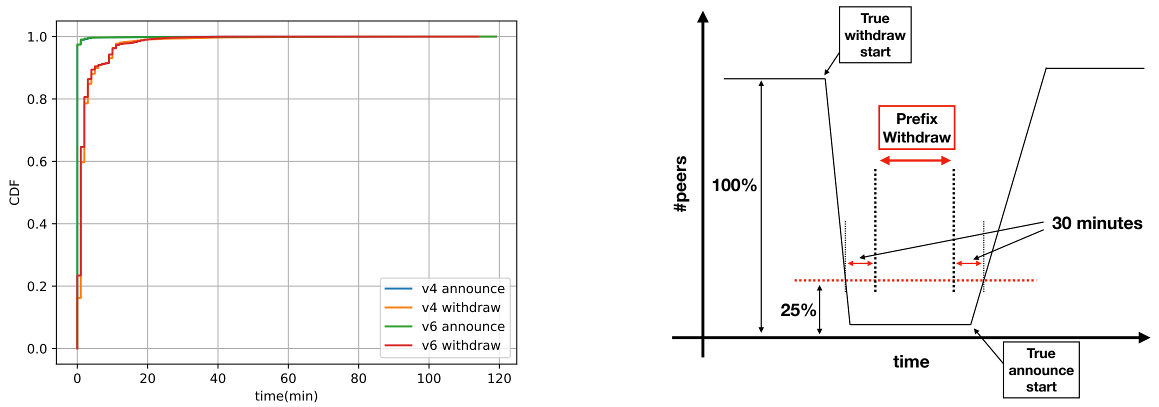

(a) Propagation time of updates messages of beacons.

(b) Change in #peers over time help us to detect the prefix withdraw.

Fig. 4.2. Zombie peers detection.

time, we observe shorter time of a prefix withdraw than in reality and ignore prefix withdraws that last less than one hour.

To detect the wild prefix withdraws, we keep track of the change of the number of peers that have routes to the prefix. When we find the rate of peers that have routes get less than 25%, we judge the time as the withdraw start of the prefix temporarily. We continue monitoring the change and we also determine the end of the withdraw when the rate get more than 25% temporarily. We need to take into account the update messages propagation time. We determine the time 30 minutes after the temporal withdraw start as the withdraw start of the prefix. We also determine the time 30 minutes before the temporal withdraw end as the withdraw end of the prefix.

4.3 Hunting zombies

In this section we show that the withdrawn and zombie AS paths collected by RIS and Route-Views also enable us to infer zombie ASes beyond RIS and Route-Views peers and estimate the scope of outbreaks. With the simple zombie detection technique described above we observe zombies only in ASes that are peering with RIS and Route-Views collectors. We classify AS paths into two categories: **the normal path** and **the zombie path**. **The normal path** is the last path the peer AS has in the RIB before it receives the withdrawal. **The zombie path** is the path that is not removed from the RIB when the prefix is withdrawn.

For each outbreak we retrieve the AS path of zombie entries and the last valid path for peers that have correctly withdrawn the prefix. A path alone provides little information, but putting together them reveals topological similarities that we consider as evidences for locations of zombies.

Figure 4.3 depicts AS paths for the zombie outbreak from 2017-04-12 00:52:00 to 2017-04-12 02:46:05 for 103.16.24.0/23. Each node represents an AS and consecutive ASes in the AS paths are connected by an edge. The green nodes represent RIS and Route-Views peers that have been correctly withdrawn the prefix before 2017-04-12 00:52:00. Thr red nodes represent zombie peers observed from 2017-04-12 00:52:00 to 2017-04-12 02:46:05. The gray nodes represent ASes that are not peering with RIS or Route-Views collectors, hence we have no direct observations for these

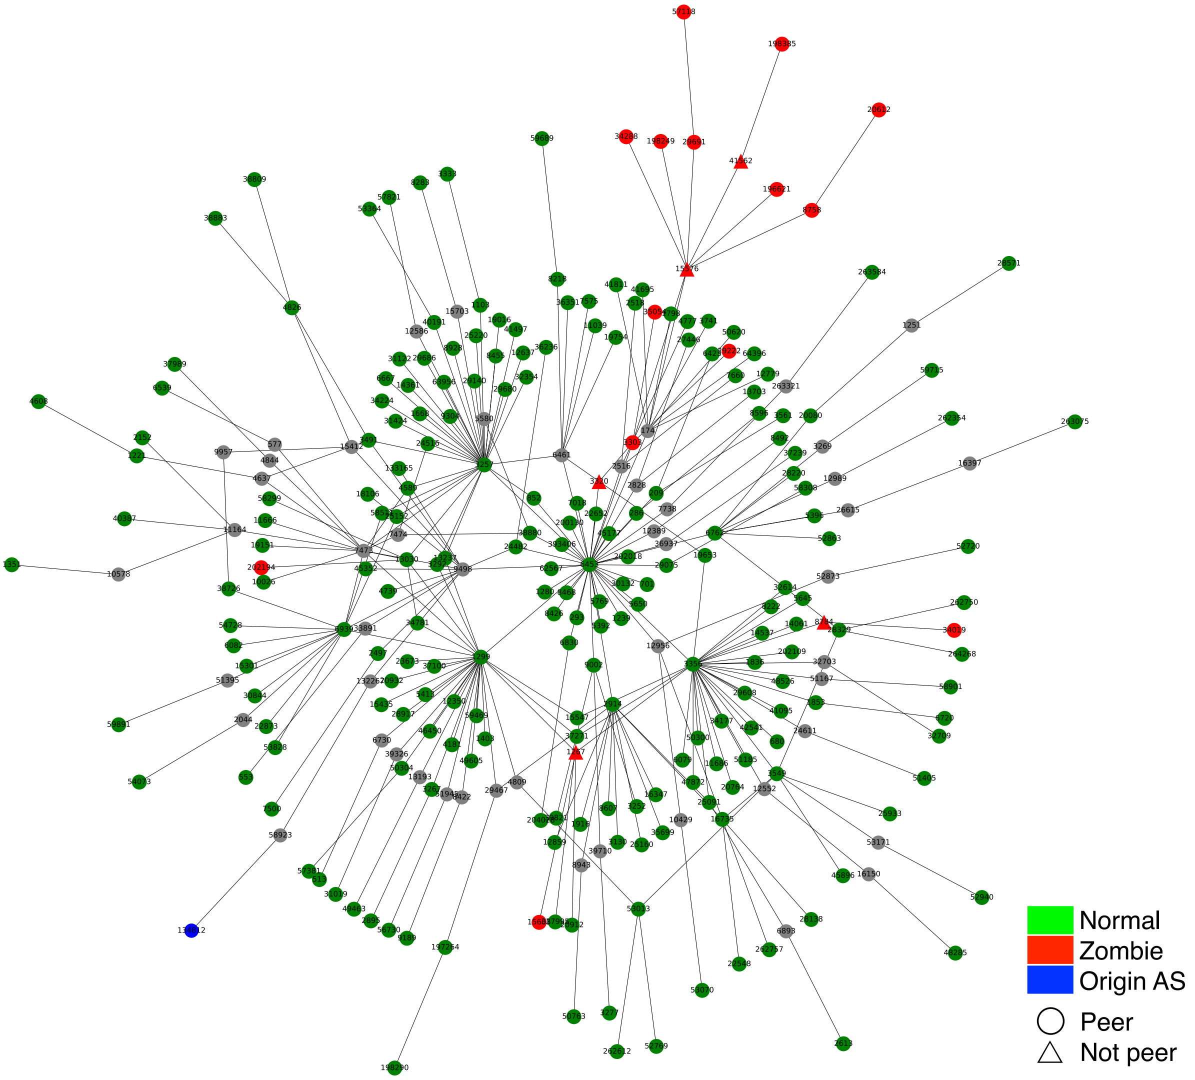

Fig. 4.3. AS paths from the zombie outbreak in 2017-04-12 for 103.16.24.0/23. Each node is an AS.

ASes but they appear in collected AS paths. Traiangle nodes represents zombie ASes that are not peer ASes detected by AS paths and show the detection method in next section.

4.3.1 Zombie path and normal path

To infer zombie ASes that are not peering with collectors, we examine whether or not not peering ASes are included in the zombie paths and the normal paths. Figure 4.4 shows Model of AS paths. Each node is an AS. Blue node is origin AS. Red node is zombie AS. We can find 2 zombie peer ASes and 4 normal ASes. In other words we can find 2 zombie paths and 4 normal paths. We cannot monitor RIBs of not peer ASes, ASa, ASb, ASc and ASd. However, we can

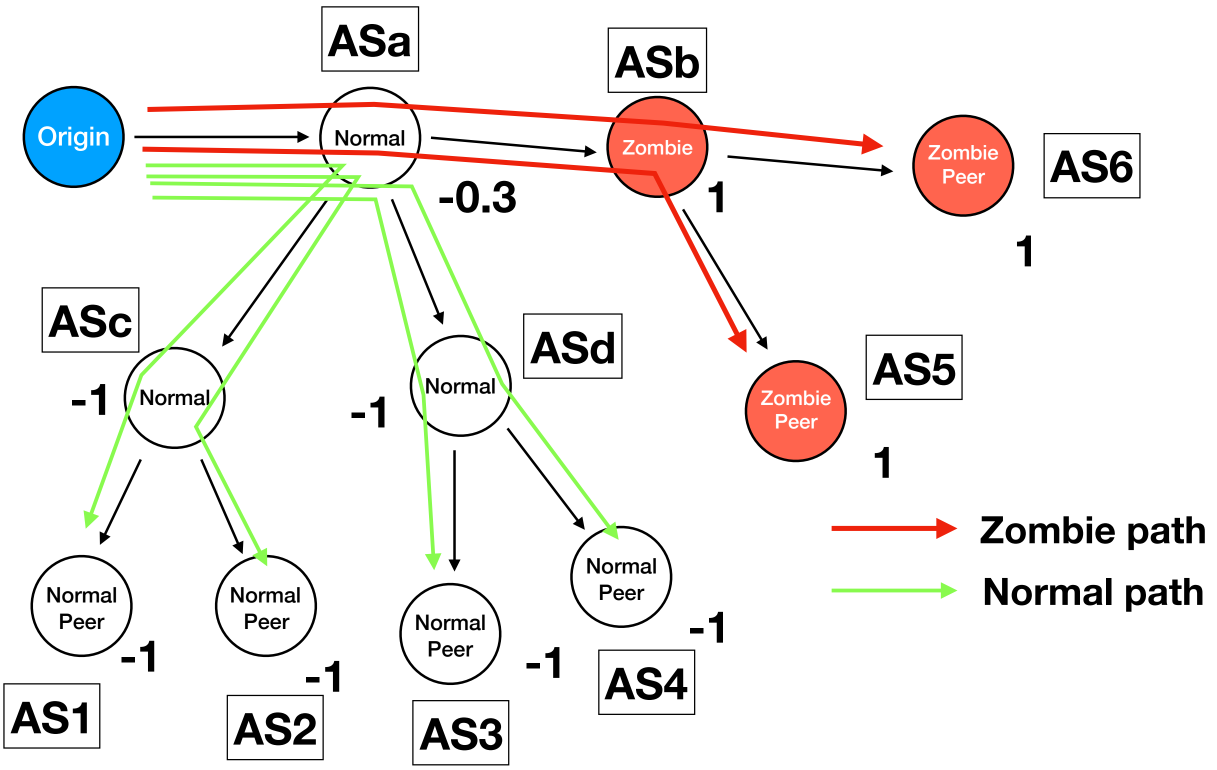

Fig. 4.4. Model of AS paths and zombie scores

peer AS	path	status
peer AS1	AS1 ASc ASa O	normal
peer AS2	AS2 ASc ASa O	normal
peer AS3	AS3 ASd ASa O	normal
peer AS4	AS4 ASd ASa O	normal
peer AS5	AS5 ASb ASa O	zombie
peer AS6	AS6 ASb ASa O	zombie

Table 4.2. Normal paths and zombie paths from peer ASes to Origin AS in figure 4.4. "O" means Origin AS, blue node.

expect whether these not peer ASes are zombie ASes or not with the state of paths including them. ASc and ASd are contained in normal paths and we can expect these ASes are not zombie ASes. If ASc or ASd are zombies, neighbor ASes of ASc or ASd may be reported as zombie ASes. On the other hand, all downstream ASes of ASb are reported as zombie ASes and we can expect ASb is a zombie AS. It is a little difficult to judge whether ASa is a zombie AS or not. If ASa is a zombie AS, ASa never send withdraw messages to tis neighbor ASes. It is in conflict with the existence of normal peer ASes. So we can expect ASa is a normal AS. If the AS correctly send withdraw messages to its peering ASes, this AS is likely to be included in normal paths. However, the AS that fail to send withdrawals is probably included in zombie paths. For each ASes, we examine

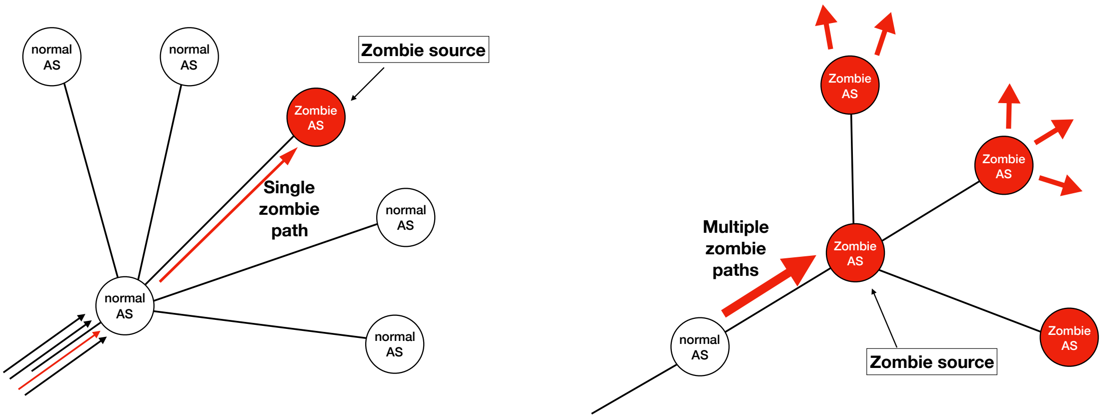

(a) Non-transit zombie source AS.

(b) Transit zombie source AS..

Fig. 4.5. Zombie source AS.

how many normal paths or zombie paths contain the AS as follows.

$$\frac{\#Z - \#N}{\#Z + \#N} \tag{4.1}$$

#W means the number of normal paths. #Z means the number of zombie paths We call the output of 4.1 **zombie score** of the AS. Table 4.2 shows paths peers receive and whether or not peers withdraw the path in figure 4.4 For example, let me calculate ASd in figure 4.4. According to table **??**, ASd is included in two normal paths and we can find the zombie score of ASd is $\frac{0-2}{0+2} = -1$. The zombie score of ASb = $\frac{2-0}{/}1 + 0 = 1$ and the zombie score of ASa = $\frac{2-4}{2+4} = -0.3$.

The zombie score ranges from -1 to 1. The closer the zombie score of the AS gets to 1, the AS is likely to be a zombie AS: the AS that has the route to the prefix when the prefix is withdrawn.

4.3.2 The source of zombie

We can observe zombie ASes regardless of peering with RIS or Route-Views collectors by couting the number of times each AS is included in the zombie path or the normal path. In this section, we show how to find the source of the zombie outbreak. The source of the zombie outbreak means the AS that fails to send withdraw messages to its peer ASes when one prefix is withdrawn and the AS that causes the zombie outbreak is included in zombie paths. So we track zombie paths from the origin AS to RIS or Route-Views peer ASes and find BGP sessions in which the previous AS (close to the origin AS) is the normal AS and the next AS (close to RIS or Route-Views peers)is the zombie AS. We call the next AS of this BGP session the zombie source AS. We also count how many zombie paths path across the session. It is not necessarily the case that the zombie outbreak caused by one AS. The zombie source AS that has many peering ASes downstream can cause a large zombie outbreak illustrated in figure 4.5b. On the other hand, the zombie source AS with a few peering ASes downstream can cause a small zombie outbreak illustrated in figure 4.5b. Large zombie outbreaks and small zombie outbreaks can exist together when a prefix is withdrawn.

Figure 4.6 illustlates the coexistence of differents kind of zombie outbreaks. We can find many zombie peer ASes at the lower left of the figure. These zombie ASes are downstream of the

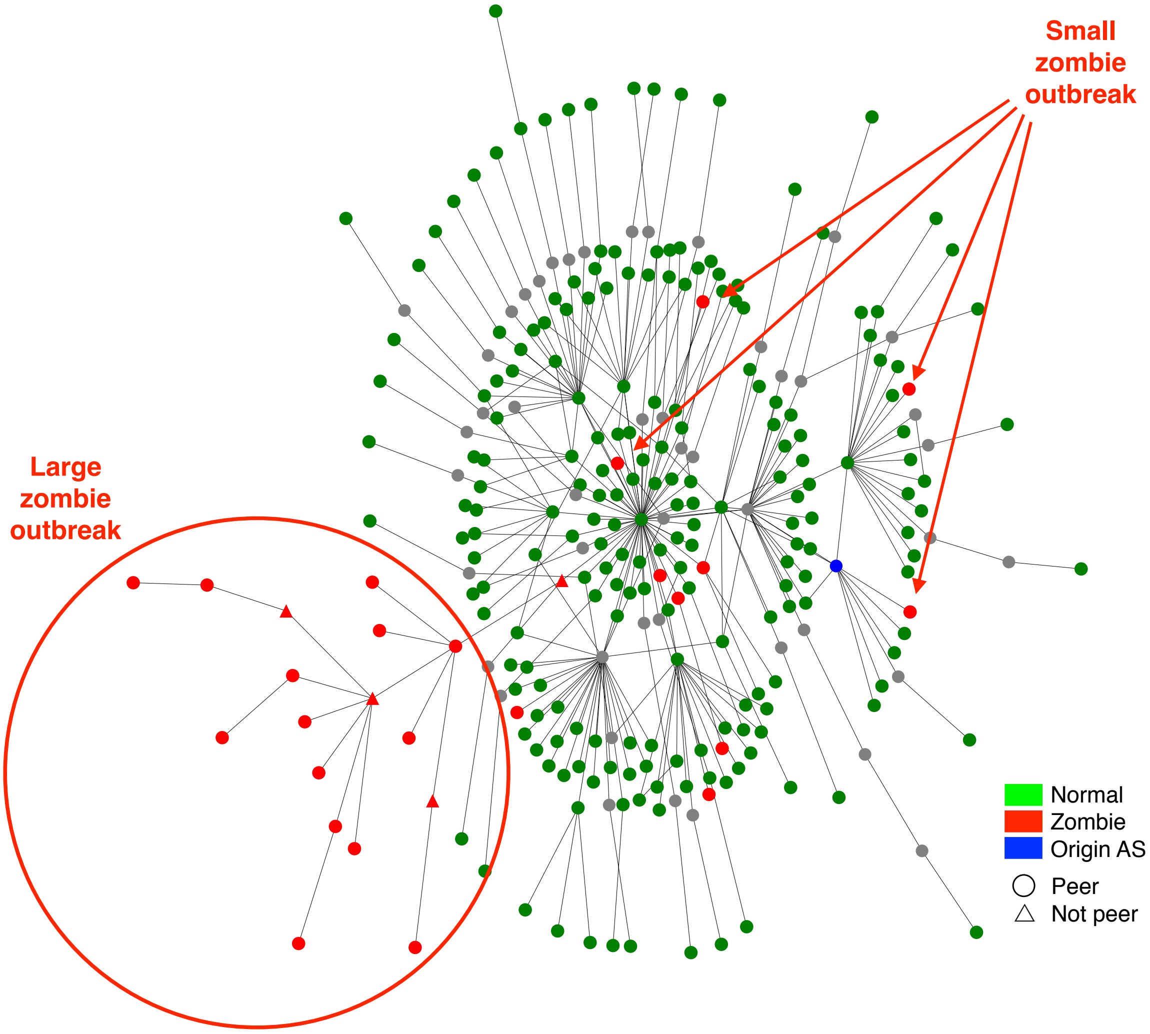

Fig. 4.6. Withdraw prefix 138.117.120.0/22 from 2017/4/13 23:30 to 2017/4/14 05:30.

AS13030, Init7. It is possible that the previous AS of AS13030 or AS1299 cause the large zombie outbreak. So we can determine that this large zombie outbreak is caused by AS13030 from this figure. On the other hand, we can find several zombie peer ASes that has no downstream peering ASes causing small zombie outbreaks.

Chapter 5

Zombie characteristics

We now investigate temporal and topological characteristics of zombies directly observed at RIS or Route-Views peers and those inferred using zombie score method. Our purpose here is to quantify the frequency of zombies, uncover the locality, and estimate the scale of zombie outbreaks. In this section, we show characteristic of zombie outbreaks or zombie ASes mainly detected during April, 2017.

The zombie outbreak is not a rare phenomenon. Figure 5.1a shows the number of zombie outbreaks normalized by the number of all prefix withdraws in our dataset, April 2017, October 2017 and April 2018. We can find more than 90% prefix withdraws that include zombie peer ASes. Intuitively, it seems that the Internet is breaking or not working but it is not true. Zombie outbreaks are local incidents in the Internet. We compute the zombie peer AS's emergence rate, that is the number of times zombies are reported for each peer AS normalized by the number of times prefixs have been withdrawn during our measurement study. Figure 5.2a shows the distribution of the values during April, 2017. We observe 25.3% of peer ASes with no zombie for IPv4 prefixes (40.8% for IPv6 prefixes). 75% peer ASes become zombie ASes for less than 1.7% zombie withdraws (1.1% for IPv6). For example, AS34288, Educational Network in the Canton of Zug, is reported as a zombie AS in 61.6% of all zombie outbraks for IPv4 prefixes. Though

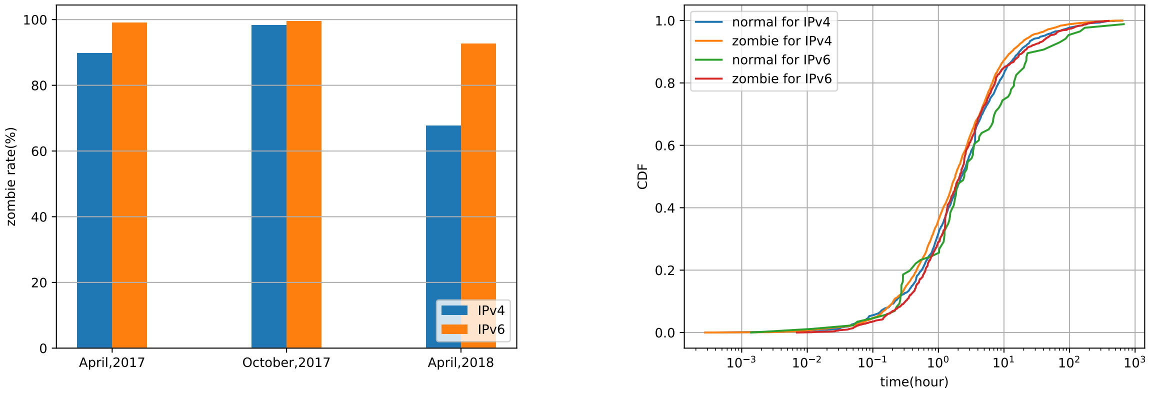

(a) Emergence rate of zombie outbreaks.

(b) Distribution of the periods of prefix withdraws.

Fig. 5.1. zombie withdraws.

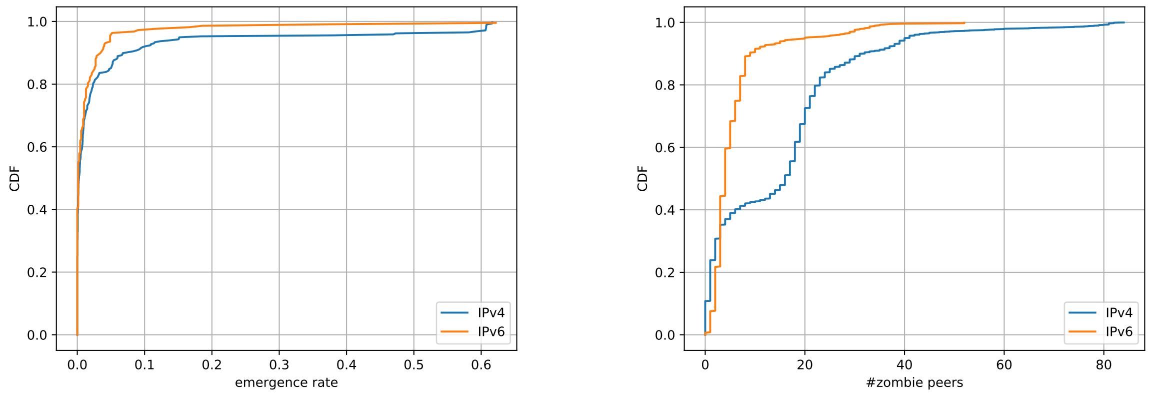

(a) Frequency of zombie appearance for each RIS or Route-Views peer AS.

(b) Distribution of the number of zombie peer ASes.

Fig. 5.2. Zombie peer ASes.

most peer ASes become zombie ASes with a small number of withdraws, some peer ASes become zombie ASes for more than half IPv4 prefix withdraws. In October, 2017 most prefix withdraws become zombie withdraws. However, we observe 19.1% peer ASes with no zombie for IPv4 (32.0% for IPv6). 75% peer ASes become zombie ASes for less than 2.3% zombie withdraws (4.0% for IPv6). In April, 2018 the probability of zombie outbreaks is less than before for IPv4. We observe 25.4% peer ASes with no zombie IPv4. This probability is not different from the one in April, 2017. These obervations imply us that a few RIS or Route-Views peer ASes are more prone to zombies.

Then, we examine the scale of zombies. Figure 5.2b illustrates the distribution of the number of zombie peer ASes in April, 2017. For IPv4 about 80% of zombie outbreaks have less than 20 zombie peer ASes. An average of the number of peer ASes for each prefix is 300 and about 90% of prefix withdraws have zombie outbreaks. From these observations, when the prefix withdraw happend, zombie peer ASes often exist somewhere, however, the number of the zombie peer ASes is small.

Figure 5.1b illustrates the distribution of periods of prefix withdraws from 2018-04-01 rto 2018-04-30. We examine how long zombie outbreaks last and compare the period of zombie outbreaks with the one of normal prefix withdraws. We can find there is no difference between zombie withdraws and normal withdraws. About 40% prefix withdraws last for less than 1 hour and about 80% of prefix withdraws last for less than 10 hours.

We compared the zombie AS paths to the paths that are advertised before the prefix withdraw. For IPv4, 25% of the zombie paths are different from the paths that are used before the withdraw (8% for IPv6). Figure 5.3 illustrates the distribution of path length for zombie paths, paths that were previously advertised by zombie ASes (Normal path(zombie peer)), and paths that were advertised by peers that correctly withdrawn the prefix (Normal path(normal peer)). The distribution of zombie paths is cleary shifted to the right hence zombie paths are usually longer. These observations imply that zombie paths are mostly different from the paths that are selected during BGP path convergence, and numerous zombies appear during path hunting.

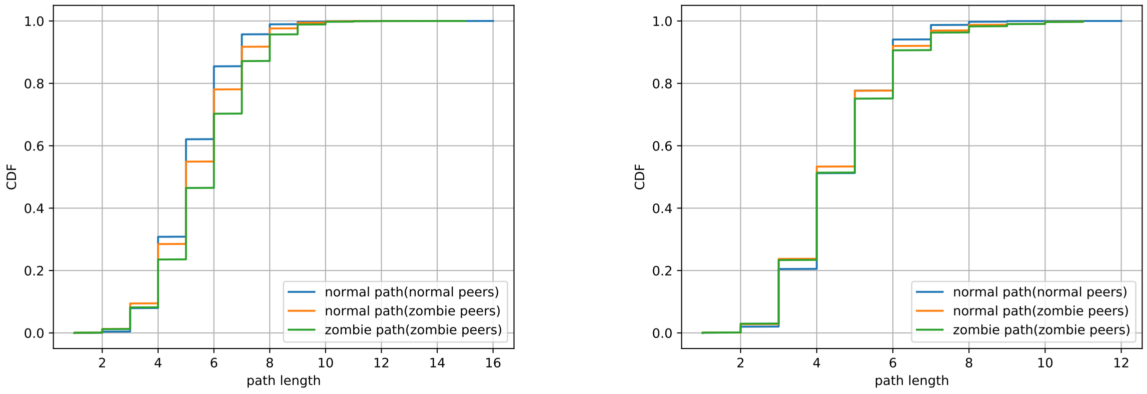

(a) For IPv4 prefix.

(b) For IPv6 prefix.

Fig. 5.3. AS path length from 2017/4/1 to 2017/4/30

Chapter 6

The source of zombie AS

We can observe several zombie peer AS during one prefix withdraw and we can detect some zombie sources from zombie paths. If different paths become zombie paths at the same point, we can determine the point (AS or BGP session) as the prime zombie root source. A transit AS is likely to become the prime zombie root source.

The number of wild preifxes is too large and it takes too long to examnie zombie outbreaks for all prefixes. We select one prefix for each AS when we read prefixes from RIBs of RIS or Route-Views collectors and we read update messages of selected prefixes during each period.

6.1 Transit zombie source ASes (in April, 2017)

We read BGP messages from 2017/04/01 to 2017/04/30 and examine zombie outbreaks and zombie sources. The number of prefixes is 53,385. The number of zombie outbreaks is 4271 for IPv4 and 698 for IPv6. Figure 6.1 illustrates the change of the number of zombie outbreaks over time normalized by the number of all prefix withdraws for IPv4 prefix. We classify the zombie outbreak according to the number of peers that have routes to withdrawn prefixes. We can observe that zombie outbreaks happened anytime and large zombie outbreaks with more than 40 zombie peers happened end of the month. Table 6.1 shows top 5 zombie source ASes that appeared a lot in zombie outbreaks. For IPv4 we found about half of all zombie outbreaks are affected by Deutsche Telekom and Swisscom. Figure 6.3a shows AS paths for 103.207.11.0/24 on 2017-04-11. We can find that all peers that are downstream of Swisscom (AS3303) are zombie ASes and AS3303's provider AS is Deutsche Telekom (AS3320). We can guess that these downstream peer ASes become zombie ASes because of Deutsche Telekom (AS3320). and Swisscom (AS3303) become a zombie AS because of zombie Deutsche Telekom (AS3320). NTS workspace (AS15576) is a downstream AS of Swisscom (AS3303). Though Swisscom and NTS workspace are not zombie

IPv4			IPv6		
3320	Deutsche Telekom	33.0%	16347	ADISTA SAS	58.5%
3303	Swisscom	12.4%	263674	JSneT	36.2%
15576	NTS workspace	12.1%	7575	AARNet	17.7%
34288	Kantonsschule Zug	8.1%	48166	FORTEX	13.6%
8220	TAL.DE Klaus	5.0%	32614	High Desert	9.6%

Table 6.1. Zombie source ASes.

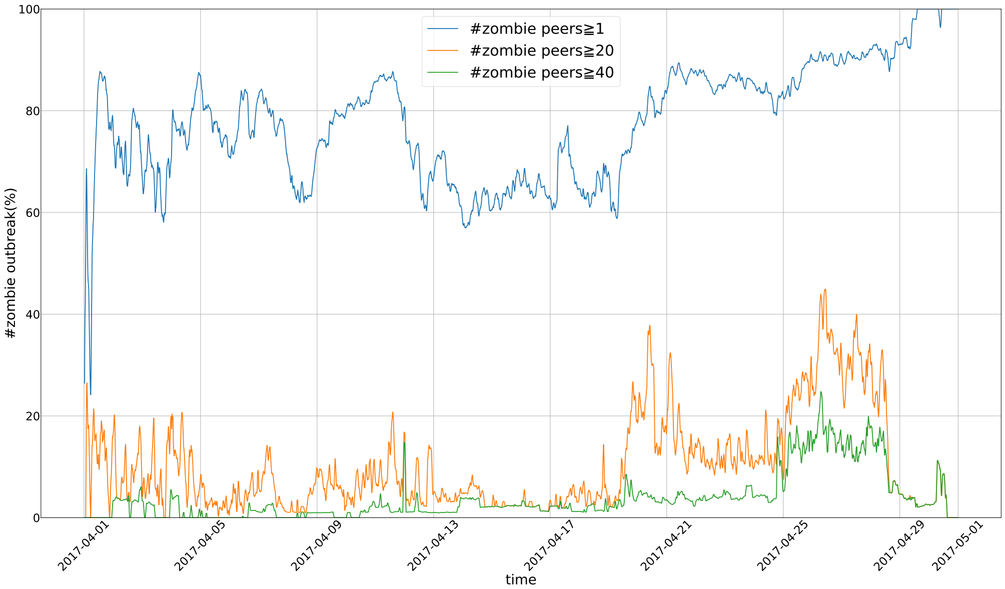

Fig. 6.1. Zombie outbreaks for IPv4 in April, 2017.

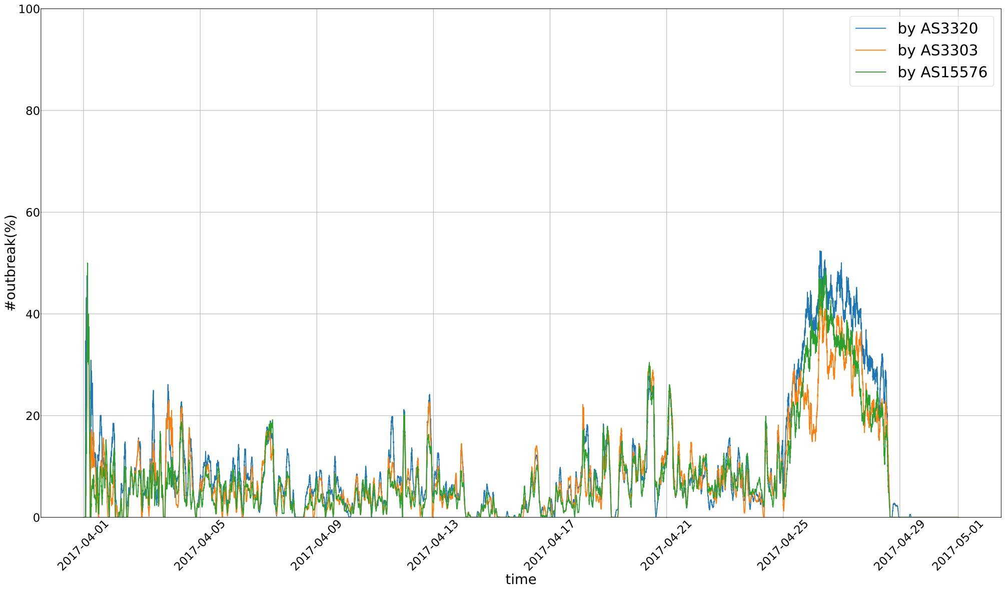

Fig. 6.2. Zombie sourcess for IPv4 in April, 2017.

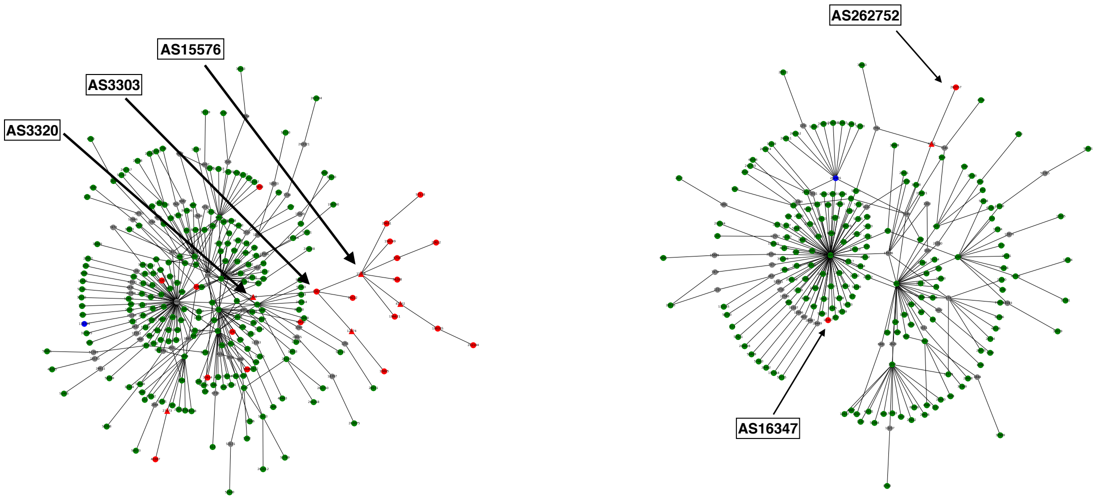

(a) Zombie caused by AS3303 and AS3320 for preifx 103.207.11.0/24

(b) Zombie caused by the session between AS6939 and AS16347.

Fig. 6.3. AS paths.

IPv4			IPv6		
34288	Kantonsschule Zug	61.6%	16347	ADIDAS SAS	62.1%
202194	EDSI-Tech Sarl	61.6%	263674	JSnetT	36.9%
20612	SwissIX	60.8%	7575	AARNet	18.6%
57118	CommunityRack.org Verein	60.7%	48166	Fortex	16.6%
196621	Matthias Cramer	60.7%	32614	High Desert	16.6%

Table 6.2. Top 5 zombie peer ASs in April, 2017.

sources strictly speaking, these ASes are in table 6.1. The reason seems that Swisscom and NTS workspace have many provider ASes other than Deutsche and the propgation of routes that pass accross the provider ASs stuck at the Swisscom and NTS workspace.

Table 6.2 shows top 5 zombie peers in April, 2017. Kantonsschule Zug (AS34288) and Matthias Cramer (AS196621) receive advertisements from NTS workspace. EDSI-Tech Sarl (AS202194), SwissIX (AS20612) are CommunityRack.org Verein (AS57118) are downstream ASes of NTS workspace. We can determine that zombie outbreaks because of Swisscom (AS3320) happened many times and this lead to high emergence rates of its downstream peer ASes. Figure 6.2 illustrates changes in the number of zombie outbreaks caused by AS3320, AS3303 or AS15576 over time. We can observe that zombie outbreaks caused by AS3320 happened mainly from 2017-04-25 to 2017-04-29. We can guess that zombie outbreaks end of April, 2017 had a great effect on the whole results of the month. We also ovserve the relationship between AS3320 and "AS3303 and AS15576" that are downstream ASs of AS3320. When AS3320 become a zombie AS, AS3303 and AS15576 also become zombie ASs.

For IPv6 we can observe that most zombie outbreaks are caused by ADISTA SAS(AS16347)

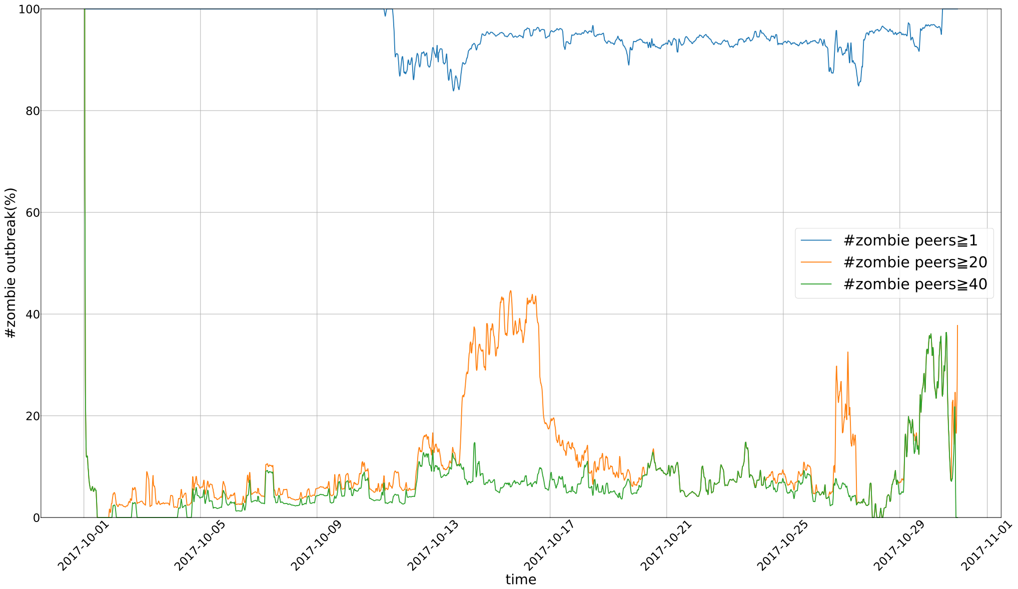

Fig. 6.4. Zombie outbreaks for IPv4 in October, 2017.

in table 6.1. AS16347 is peer AS of RIS or Route-Views collectors and become zombie peer with highest emergence rate in table 6.2. We can find other peer zombie ASs in table 6.2 as sources of zombie outbreaks. We can guess that the size of zombie outbreaks for IPv6 prefixes is small in this measurement period and few transit AS became a zombie AS. One peer AS became a zombie AS many times and the zombie AS did not spread this zombie information to its peers. Figure 6.3b shows a zombie outbreaks for 2804:3d4::/32 from 2017-04-08 20:00 to 2017-04-18 22:10. There were only two zombie peer ASes, ADISTA SAS (AS16347) and Insidesign Tecnologia (AS262757) alone. This solely peer zombie AS can be found many times for IPv6 prefix zombie outbreaks.

6.2 Peer zombie ASes and transit zombie source ASes

In previous section, we can find zombie source ASes simply. But if a lot of zombie peer ASes that has few downstream ASes are reported, it is difficult to distinguish a transit zombie source AS from a zombie peer AS that cause a small zombie outbreak. In this section we explain zombie sources during October, 2017 and compare the number of zombie paths that pass across the zombie ASes. The number of monitored prefixes is 44,208. We can observe 5415 prefix withdraws and 98% of withdraws include at least one zombie peer AS. Figure 6.4 illustrates changes in the number of zombie outbreaks for each number of zombie peers normalized by a total number of prefix withdraws for IPv4 prefix. We can observe that at least one zombie peer AS is reported in most prefix withdraws. As described in the chapter 5 this high probability of zombie outbreaks does not indicate problems in the global network. Our examination of zombie peers emergence rate shows about 20% of RIS or Route-Views peers with no zombie for IPv4 (about 30% for IPv6). 75% of peer ASes is reported as zombie peers for less than 2.3% zombie outbreaks for IPv4 (3.9%

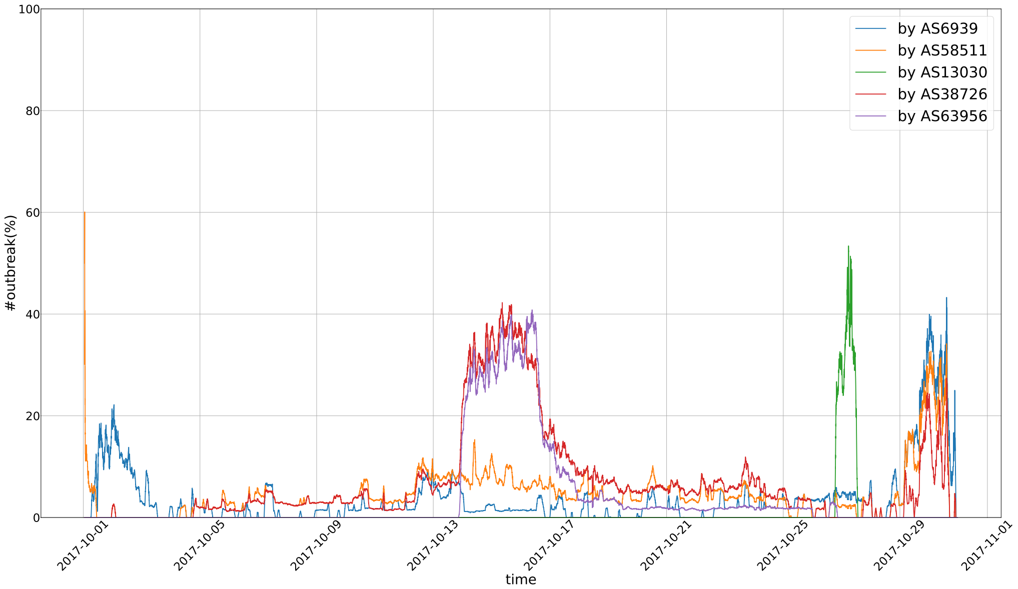

Fig. 6.5. Zombie sourcess for IPv4 in October, 2017.

IPv4			IPv6		
36236	NetActuate	48.7%	29691	Nine Internet Solututions	88.4%
37648	UEMOA	44.3%	16347	ADISTA SAS	68.7%
31122	Viatel	35.0%	37468	Angola Cables	40.5%
6082	IPiFony Systems	17.9%	264268	TELETURBO	16.5%
38726	VTC DIGICOM	14.5%	264911	POWERNET TELECOM	14.5%

Table 6.3. Top 5 zombie peer ASes in October, 2017.

for IPv6). These observations shows that a few zombie peers is likely to be reported as zombie peers many times. The number of peer ASes with more than 10% emergence rate is less than 10 for IPv4. Table 6.3 shows top 5 zombie peer ASes with high emergence rate in October, 2017. We can find NetActuate (AS36236) and UEMOA (AS37468) became zombie peers for about half of all zombie outbreaks for IPv4. When we focus on a large zombie outbreaks, different large zombie outbreaks happened together from 2017/10/13 to 2017/10/17, from 2017-10-26 to 2017-10-28 and from 2017-10-29 to 2017-10-30.

Table 6.2 shows top 5 zombie sources for each number of paths that pass across the zombie source AS. We have shown the classification of zombie outbreaks when a prefix is withdrawn in chapter 4. A transit zombie AS can cause a large zombie outbreaks. So we set the threshold that is referred to the number of paths passing across a source of a zombie outbreak in order to extract transit zombie source ASes. Left column in table 6.2 shows top 5 zombie sources with at least one zombie path, that is, we examined all zombie source ASes. We can find that 38.5% of zombie outbreaks are caused by Digiweb (AS31122). However we can also find that Digiweb (AS31122) has an small influence on other ASes because it does not appear in the ranking of zombie source ASes that cause more than 5 zombie ASes. Digiweb (AS31122) is a RIS or

	$\#path \geq 1$		$\#path \geq 5$		$\#path \geq 10$	
	AS	rate	AS	rate	AS	rate
IPv4	31122	38.5%	6939	47.7%	6939	64.8%
	36236	37.4%	58511	24.8%	13030	16.4%
	37468	16.0%	13030	22.6%	11058	1.5%
	6082	14.2%	24482	6.7%	24490	0.9%
	38726	9.6%	3491	1.5%	28338	0.9%
IPv6	29691	82.5%	13030	54.9%	13030	74%
	16347	64.1%	1299	13.7%	1299	14.2%
	37468	38.8%	24482	7.8%	8167	5.7%
	264268	16.3%	7473	5.8%	3549	2.8%
	264911	12.1%	15576	5.8%	3356	2.8%

Table 6.4. Top 5 zombie source ASes for each number of pathes that pass across zombie causes.

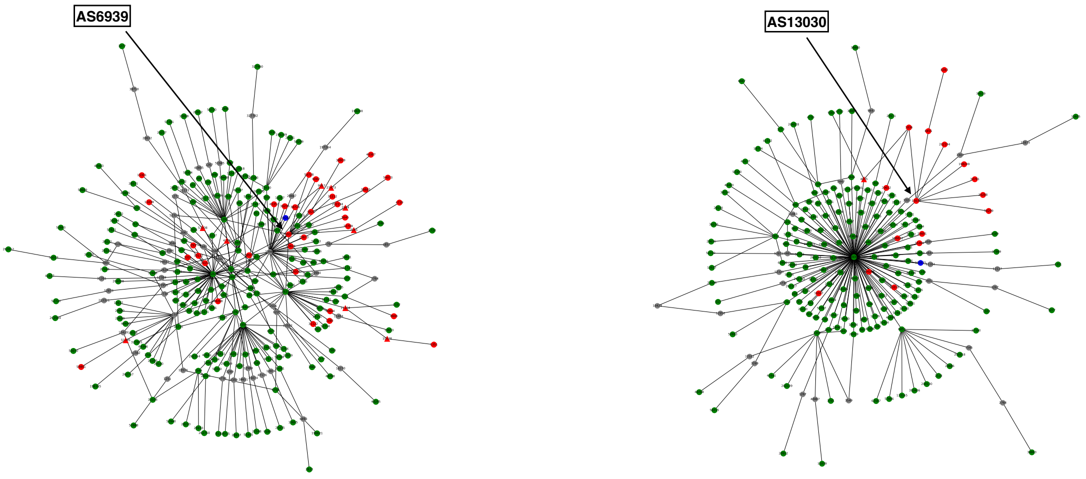

(a) Caused by Hurricane Electric (AS6939).

(b) Caused by INIT7 (AS13030).

Fig. 6.6. AS Paths in October, 2017

Route-View peer AS and 35% of zombie outbreaks contain Digiweb as a zombie AS according to table 6.3. Similiar characteristics can be seen for other zombie peers with high emergence rate in table 6.3. On the other hand, a transit zombie source AS has a big influence on other downstream ASes. We show transit zombie source ASes in center and right columns of table 6.2. For IPv4 we can observe AS6939 Hurricane Electric caused large zombie outbreaks. Hurricane Electric operates a large IPv4 and IPv6 transit networks globally. Hurricane Electric is connected to over 200 major exchange points and exchanges traffic directly with more than 7,300 different networks [20]. So the possibility of zombie outbreaks caused by Hurricane Electric gets higher. Figure 6.6a illustrates a large zombie outbreak caused by Hurricane Electric (AS6939). Many zombie

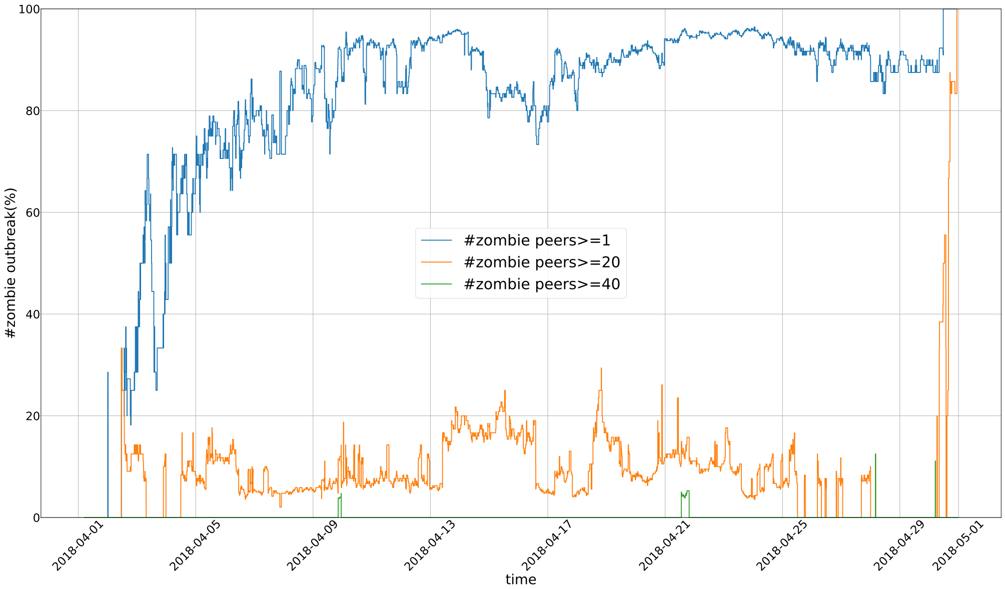

Fig. 6.7. Zombie outbreaks for IPv6 in October, 2018.

peers are reported when 103.228.244.0/24 is withdrawn from 2017-10-29 16:17:56 to 2017-10-30 00:44:51 and Hurricane Electric's downstream zombie peers are outstanding. All downstream ASes become zombie ASes in figure 6.6a. We can observe not all but many ASes that have BGP sessions with Hurricane Electric became zombie ASes.

For IPv6 Init7 (AS13030) is the prime source of large zombie outbreaks in October, 2017. Figure 6.6b illustrates a IPv6 prefix 2400:ec80:1101::/48 was withdrawn from 2017-10-27 07:19 to 2017-10-27 08:27 and a large zombie outbreak caused by Init7.

Figure 6.5 illustrates changes in the number of zombie outbreaks of which source ASes are AS6939 (Hurricane Electric), AS58511 (Anycast Global Backbone), AS13030 (Init7), AS38726(VTC DIGICOM) and AS63956 (Colocation Australia). AS6939, AS58511 and AS13030 are high frequency zombie source ASes. AS13030 became a source of zombie AS around 2017-1-27 in figure 6.5 and many zombie outbreaks with more than 20 zombie peers are reported during the same period in figure 6.4 Similarly when zombie outbreaks with more than 40 zombie peers happened many times from 2017-10-29 to 2017-10-30, AS6939 become a source AS of zombie outbreaks outstandingly. AS58511 is downstream of AS6939 and follow AS6939. On the other hand, from 2017-10-13 to 2017-10-17 zombie outbreaks with more than 20 zombie peers happened many times. However, we can observe AS6939 and AS13030 have small influence on these zombie outbreaks. So we examine outstanding zombie source ASes in this period and detect AS38726 and AS63956. AS38726 and AS63956 are not outstanding in the examination of zombie outbreaks in a whole month.

6.3 Long-lasting zombie source for IPv6 (April, 2018)

In this section we explain the examination of zombie outbreaks in April, 2018. We shows

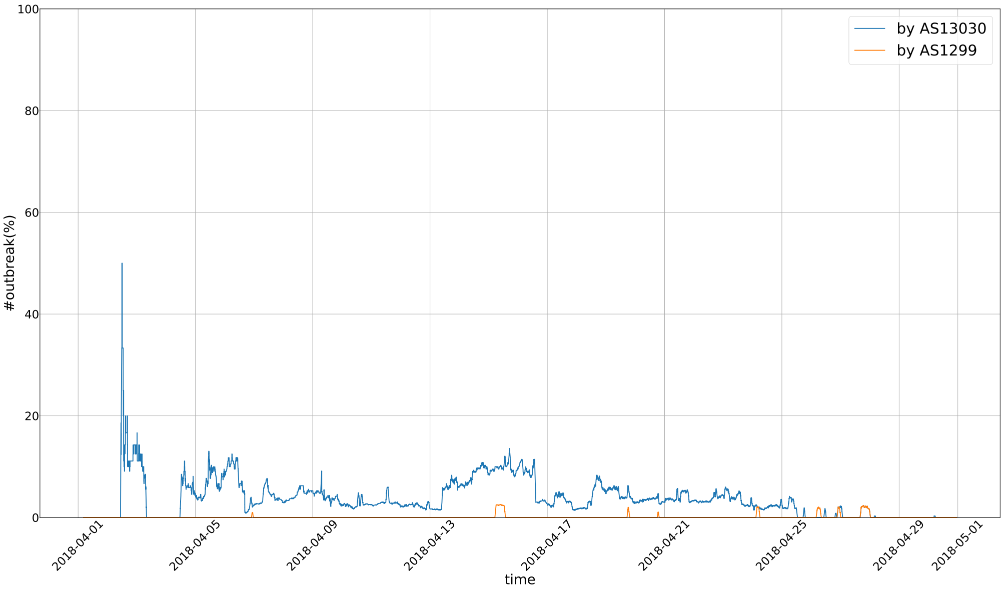

Fig. 6.8. Zombie sourcess for IPv6 in October, 2018.

IPv4			IPv6		
34224	Neterra	66.1%	13030	Init7	72.3%
3741	Internet Solutions	64.4%	1299	Telia Company	12.3%
1239	Sprint	55.6%	24482	SG.GS	4.6%
53364	ZEROFAIL	49.4%	3549	Level 3 Communications	4.6%
262612	Piotr Piwowar	48.8%	3356	Level 3 Parent, LLC	3.0%

Table 6.5. Top 5 zombie peer ASes in April, 2018.

the result of zombie outbreaks for IPv6. Figure 6.7 shows changes in the number of zombie outbreaks for each number of zombie peers normalized by a total number of prefix withdraws for IPv6 prefix. We can observe about 90% prefix withdraws with at least one zombie peer AS and about 15% prefix withdraws with more than 20 zombie peer ASes. Table 6.5 shows top 5 zombie source ASes for each number of paths that pass accross the zombie source AS for IPv6. Top 5 ASes that caused at least 1 zombie path are peer ASes of RIS or Route-Views collectors. Neterra (AS34224) became zombie peer AS in about half of zombie outbreaks for IPv4 in this period. Internet Solutions (AS3741) and Sprint (AS1239) became zombie peer ASes in about 45% of zombie outbreaks for IPv4. These peer ASes do not propagate update messages to its downstream ASes. On the other hand, Init7 (AS13030) is mainly the source of large zombie outbreaks in this period for IPv6. Figure 6.8 illustrates changes in the number of zombie outbreaks caused by AS13030, AS1299 over time. The change of the number of zombie outbreaks with more than 20 zombie peers is similar to one of the number of zombie outbreaks caused by AS13030. So we can determine most large zombie outbreaks in this period were caused by Init7 (AS13030). For IPv6, AS13030 is also the prime source of zombie outbreaks from 2017. So this problem last for a long time and it seems difficult to solve this problem.

We can contact with network operators at Init7. They recognized these issues with IPv6 routers, likely due to misbehaving vendor software, and expressed the need for zombie reporting systems, as it creates customer complains every few months. Mitigation of the BGP zombies usually required the cleaning of some Route Reflector iBGP sessions within Init7's network. which seems to be uncommon. Upgradeing to later firmware version did not resolve the problem.

Chapter 7

Disscussion

In this paper we show the characteristics and examine sources of BGP zombie outbreaks. While BGP zombies are not unusual, limited ASes become zombie ASes many times in the Internet. We can detect zombie peer ASes by monitoring routing information and advertisements collected by RIS and Route-Views projects straightforward and we can observe multiple zombie peer ASes are reported at the same time when a transit AS fails to send withdraw messages to its downstream ASes.

We examine prefixes in all ASes collected by RIS and Route-Views projects so we can inspect zombie routes in wide portion of the Internet. On the other hand we exclude prefix withdraws that last for less than an hour because we take into account of update messages propagation time and delay. So we can not examine short prefix withdraws. Furthermore in this paper we collect and examine past BGP routing information and we do not conduct realtime detection.

We can detect zombie ASes that have route information of withdrawn prefix in the global network. It is a challenge to identify a routing configuration change intended to limit the visibility of a prefix. We distinguish zombie ASes that are downstream of a transit zombie AS from a single zombie AS and refer to the transit AS as the source of zombie. Network operators in downstream ASes of the transit zombie AS can suffer from BGP zombies but it may be difficult to identify the source of the problem in their AS. It is also challenging for network providers with large transit networks to examine the status of all routers. We can provide one solution to check the health of routers or dubug routers by pinpointing the root source of BGP zombies.

In this paper, we focus on a large zombie outbreaks caused by one transit zombie AS. We discussed the possibility of router failures that cause zombie outbreaks but we do not have enough data to certify it. According to the observations in our study and the investigations conducted with network operators, we believe BGP zombies are primarily caused by software bugs in routers, BGP optimizers, and route reflectors. We can identify the source zombie AS but we do not address the router's behavior in the AS. Finer investigation of sources of zombies can help us and network operators to provide good internet service.

Chapter 8

Conclusion

In this paper, we examined BGP zombie outbreaks in the Internet. We monitored routing changes for one month and observed BGP zombies every day. However these observations do not mean the vulnerability of the Internet. We observed that zombie ASes are located locally and the same ASes are seen many times as zombie ASes. Many zombie paths are revealed during path hunting and a large zombie outbreak is caused by a transit AS. Identifying the source of zombie outbreaks can help network operators to debug and improve their routing systems.

References

[1] RIPE NCC, RIS Raw Data. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/ris-raw-data.

[2] The RouteViews project. http://www.routeviews.org/.

[3] ARIN. https://www.arin.net/.

[4] Geoff Huston, Mattia Rossi, and Grenville Armitage. A technique for reducing bgp update announcements through path exploration damping. *IEEE Journal on Selected Areas in Communications*, 28(8):1271–1286, 2010.

[5] Curtis Villamizar, Ravi Chandra, and Ramesh Govindan. Bgp route flap damping. Technical report, 1998.

[6] Geoff Huston. Analyzing the internet’ s bgp routing table. *The Internet Protocol Journal*, 4(1):2–15, 2001.

[7] Curtis Villamizar, Ravi Chandra, and Ramesh Govindan. Bgp route flap damping. Technical report, 1998.

[8] Timothy G Griffin and Brian J Premore. An experimental analysis of bgp convergence time. In *Network Protocols, 2001. Ninth International Conference on*, pages 53–61. IEEE, 2001.

[9] Philip Smith and Christian Panigl. Ripe routing working group recommendations on route-flap damping. *ripe-378, May*, 2006.

[10] Cisco Field Note. Endless bgp convergence problem in cisco ios software releases. oct. 2001.

[11] Tony Bates, Ravi Chandra, and Enke Chen. Bgp route reflection-an alternative to full mesh ibgp. Technical report, 2000.

[12] Paul Traina, Danny McPherson, and John Scudder. Autonomous system confederations for bgp. Technical report, 2007.

[13] Anindya Basu, Chih-Hao Luke Ong, April Rasala, F Bruce Shepherd, and Gordon Wilfong. Route oscillations in i-bgp with route reflection. In *ACM SIGCOMM Computer Communication Review*, volume 32, pages 235–247. ACM, 2002.

[14] Ratul Mahajan, David Wetherall, and Tom Anderson. Understanding bgp misconfiguration. In *ACM SIGCOMM Computer Communication Review*, volume 32, pages 3–16. ACM, 2002.

[15] Craig Labovitz, Abha Ahuja, and Farnam Jahanian. Experimental study of internet stability and backbone failures. In *Fault-Tolerant Computing, 1999. Digest of Papers. Twenty-Ninth Annual International Symposium on*, pages 278–285. IEEE, 1999.

[16] RRIPEstat. https://stat.ripe.net/index/about-ripestat.

[17] BGPstream. https://bgpstream.caida.org/.

[18] NLNOG RING looking glass. http://lg.ring.nlnog.net/.

[19] RIPE NCC, Current RIS Routing Beacons. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/current-ris-routing-beacons.

[20] About Hurricane Electric. http://www.he.net/about_us.html.

[21] Jun Li, Michael Guidero, Zhen Wu, Eric Purpus, and Toby Ehrenkranz. Bgp routing dynamics revisited. *ACM SIGCOMM Computer Communication Review*, 37(2):5–16, 2007.

[22] RIPE NCC, Atlas. `https://atlas.ripe.net`.

[23] Alexander Marder and Jonathan M. Smith. Map-it: Multipass accurate passive inferences from traceroute. In *Proceedings of the 2016 Internet Measurement Conference*, IMC '16, pages 397–411, New York, NY, USA, 2016. ACM.

[24] Matthew Luckie, Amogh Dhamdhere, Bradley Huffaker, David Clark, and kc claffy. Bdrmap: Inference of borders between ip networks. In *Proceedings of the 2016 Internet Measurement Conference*, IMC '16, pages 381–396, New York, NY, USA, 2016. ACM.

[25] About Hurricane Electric. `http://www.dataswitchworks.com/MLXe-4.asp`.

[26] Abraham Yaar, Adrian Perrig, and Dawn Song. Siff: A stateless internet flow filter to mitigate ddos flooding attacks. In *Security and Privacy, 2004. Proceedings. 2004 IEEE Symposium on*, pages 130–143. IEEE, 2004.

[27] Lan Wang, Xiaoliang Zhao, Dan Pei, Randy Bush, Daniel Massey, Allison Mankin, S Felix Wu, and Lixia Zhang. Observation and analysis of bgp behavior under stress. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, pages 183–195. ACM, 2002.

[28] Jong Han Park, Dan Jen, Mohit Lad, Shane Amante, Danny McPherson, and Lixia Zhang. Investigating occurrence of duplicate updates in bgp announcements. In *International Conference on Passive and Active Network Measurement*, pages 11–20. Springer, 2010.

Acknowledgements

First and foremost, I would like to express my sincere gratitude to Professor Hiroshi Esaki for giving me a great opportunity to be a part of his laboratory, which made my one year of bachelor's course and 2 years of master's course priceless and unforgettable. I would like to express my gratitude to Associate professor Hideya Ochiai, Dr. Manabu Tsukada. They taught me what is a research, how to organize a thesis and how to construct a presentation. Without these advices, I could not have finished my thesis. I also deeply grateful to Romain Fontugne at Internet Initiative Japan Innovation Institute and Emile Aben at Reseaux IP Europeens Network Coordination Centre for directing my research and giving me many valuable advices. I owe my deepest gratitude to Ryo Nakamura, Hiroyuki Ikegami, Satoru Kobayashi to have a discussion about my research and giving me technical advices. I would like to thank all of our lab members, especially Kazuki Otomo, Toki Suga, Masahiro Kitazawa, Genta Imai and Daichi Kouda for being good friends of mine, and making my lab life joyful. I owe a very important debt to secretaries, Aiko Iwai and Fumi Takahashi for assisting my administrative tasks and making our life in the laboratory comfortable. Lastly, my heartfelt appreciation goes to my family for letting me study in this University and support my student life. Without their support, I must have not experienced such a great time like this.