

Characterizing Address Use Structure and Stability of Origin Advertisement in Inter-domain Routing

Sophie Y. Qiu, Patrick D. McDaniel*, Fabian Monrose, and Aviel D. Rubin
Dept. of CS, The Johns Hopkins University *Dept. of CSE, The Pennsylvania State University
{yuqiu,fabian,rubin}@cs.jhu.edu *mcdaniel@cse.psu.edu

Abstract

The stability and robustness of BGP remains one of the most critical elements in sustaining today's Internet. In this paper, we study the structure and stability of origin advertisements in inter-domain routing. We visualize and quantitatively characterize the frequency, size, and effect of address assignment and origin changes by analyzing real-world BGP updates for a period of one year from multiple vantage points. Broad classes of prefix behaviors are developed. We show that a significant portion of BGP traffic is due to prefix flapping and explore the contributing factors which include a number of prefixes with abnormal short up-and-down cycles. A significant portion of prefixes have high origin stability. Most ASes are involved in few, if any, prefix movement events, while a small number of ASes are responsible for most of the origin churn. Additionally, we find that a high volume of new prefixes can be attributed to actively evolving countries, that some abnormal prefix flapping is most likely due to misconfiguration, and that some culprit ASes characterize the places where multi-origin prefixes oscillate.

1 Introduction

BGP is *the* protocol that drives contemporary inter-domain routing [1]. Its stability and robustness are often cited as key elements contributing to the success of the Internet. BGP is, however, exhibiting signs of strain resulting from the immense size, complexity, and increasing hostility of the Internet. Concerns about the resulting slow convergence time, transient performance and availability problems, and malicious behavior are growing. Of late, there have been numerous calls for new mechanisms to address these problems and for enabling new services (e.g., QoS, traffic engineering). Unfortunately, the development of such mechanisms and services has been slowed in part by an incomplete understanding of the operation of BGP.

The networking community has been accruing a vast body of knowledge about the dynamics of BGP and the oddities of its operation on the Internet. Periodic reports of the status of the global routing infrastructure and address space are widely used to ascertain the health of the Internet, and tools are used by many to construct maps of the organizational [2, 3] or topological [4, 5, 6] relationships between Autonomous Systems (ASes). BGP has also been the subject of considerable introspection within the academic community (e.g., [7, 8, 9, 10, 11, 12]).

However, one important factor that has yet to receive significant attention is the issue of assignment and ownership attestation of the IP address space. While much is known about the effects of transient network failures and topology on inter-domain routing, surprisingly little is known about how origin stability and prefix structure affect BGP performance.

In this paper, we further an understanding of address space dynamics by studying the structure and stability of origin advertisements in inter-domain routing. We begin by first exploring the macro-level structure of IP address advertisements using our `quadchart` IP address advertisement visualization tool which renders origin, prefix size, and locality information on a single two-dimensional plane. By comparing representations of routing advertisements across periods of a few years apart, one can glean valuable observations about the evolution of IP address space usage. As we show later, while these (and other) visualization techniques [13] can be a very helpful addition to quantitative analysis, a micro-level investigation is still required to better understand underlying BGP dynamics.

Our empirical analysis in Sections 3–6 attempts to do just that, and centers around an investigation of the stability of prefix origins. We characterize the frequency, size, and effect of address assignment and origin changes by studying real-world BGP traffic. The analysis presented provides empirical results that support conventional wisdom — while providing a deeper and more accurate understanding of BGP operation. That said, our findings empirically justify an important assumption which a number of BGP secu-

rity proposals rely on: a significant number of prefixes have high origin stability (see for example [9, 10, 12]). Furthermore, our analysis may be directly applicable to a number of other inter-domain routing research areas, including the recently introduced Routing Control Platform (RCP) [14] which is aimed at mitigating potentially harmful interactions between iBGP and eBGP. Our work provides a deeper understanding of factors related to the dynamics of origination, and so provides essential data for the designers of such systems.

2 Visualizing prefix advertisement

One way to illuminate hidden structure in complex datasets is through visualization [15]. Visualization tools use graphical representations of intangible physical or digital phenomena to illuminate subtle or broad characteristics. Humans are much better than general purpose algorithms at identifying subtle patterns or structure in graphical images. We exploit this ability by visualizing the advertisements of the IP address space over several years, and ascertaining the gross level features from visual inspection.

A *quadchart* is an intuitive way of visualizing hierarchical data such as the IP address space. For example, primarily used as anomaly detection, Teoh *et al.* [13] used a quadtree to relate prefix origin change events to the AS to which they pertain. Conversely, we use a quadchart herein to visualize coverage and fragmentation of address space advertisements. Illustrated in Figure 1, the quadcharts used throughout this section define advertised prefixes as the regions in a $(2^{16} \times 2^{16})$ plane. Each point in the plane represents a single IP address (i.e., /32 prefix), and prefixes are represented by square regions covering all contained addresses. The coordinates and size of the prefix in the plane are defined by its significant bits.

The odd binary digits of a 32-bit IP prefix are used to encode the y-axis coordinate, and the even digits encode the x-coordinate. (Note that only the first 6 digits are shown in Figure 1.) For example, consider the y-coordinate of the prefix 144.0.0.0/4. 144.0.0.0 is encoded as the (binary number) 10010000000000000000000000000000. Bit position 0 is the left-most digit and bit position 31 is the right-most digit. The y-coordinate is computed by recursively subdividing the plane as dictated by the odd position digits of the prefix (as 0, 1, 0, ...). Because the first digit is 0, the y-coordinate is (initially) set to $0 * (2^{16}/2)$. The second digit is 1, so it is moved down the y axis by $1 * (2^{16}/4)$. This process is repeated for all significant bits over increasingly smaller regions.

More precisely, the y-coordinate is defined as

$$\sum_{i=0}^{31} \left\{ \begin{array}{l} \text{even}(i) : 0 \\ \text{odd}(i) : \text{bit}(i) \left(\frac{2^{16}}{(\lfloor k/2 \rfloor + 1)^2} \right) \end{array} \right\}$$

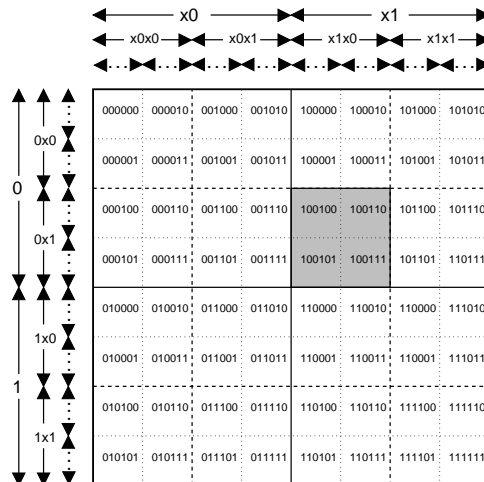


Figure 1. Quadchart visualization.

where $\text{bit}(i)$ returns the binary value at bit position i . The x-coordinate is defined similarly using the even position bits. The size of the rectangular region is determined by the mask length (say, m), i.e., the IP addresses are covered by a rectangle of size $2^{16-\lfloor m/2 \rfloor}$ by $2^{16-\lceil m/2 \rceil}$. For example, the /4 prefix is covered by a 2^{14} by 2^{14} square.

The rectangle shading is determined by the originating AS. Each AS is assigned a unique color which is shared by the rectangles of all prefixes it originates. All unadvertised *dark* address space is shaded black. The AS colors are selected to be as visually distinct as possible and the selection is also deterministic to ensure assignment coloring is consistent across graphs. Squares whose size is below the resolution depth of the graph are promoted to a size equal to a single pixel.

Figures 2 and 3 show the IPv4 space advertised via BGP on January 1st, 2001 and June 1st, 2004, respectively, as reflected in the RIB tables in the Route Views Repository [16] and ICANN reports.¹ Note that some parts of the 32-bit address space will by definition not be advertised. For example, 224.0.0.0 through 239.255.255.255 are the reserved Class D addresses (multicast) and 240.0.0.0–255.255.255.255 are the reserved Class E addresses (for experimental purpose).

While the graphs span two and a half years of rapid expansion of the Internet, there appears to be surprisingly little gross change to the structure of the advertisements. The growth in the address use is shown in the advertisements represented in the lower regions of the 64.0.0.0–69.255.255.255 (left hand side of the lower left hand quadrant) and 200.0.0.0–223.0.0.0 (lower right quadrant) in the

¹A highly detailed quadchart movie illustrating the evolution of the address space over 4 years is available at: <http://www.patrickmcdaniel.org/bgp/2001-2004.30sec.nomartians.mov>.




Figure 2. BGP advertisements on January 1st, 2001.

latter RIB. The 64/8 space is largely allocated to ARIN (the North American address registry), so it is likely that this reflects the proliferation of commercial and personal networks within the United States, Canada, and other served countries. The 200–223 address space is divided among many international registries. In particular, the growth of the use of the 200.* space is noteworthy. This address space has been assigned to Central and South America. The increased usage may indicate network growth not only in the industrialized countries in the region, but also in the extension of the Internet to historically under-served countries. The 0.0.0.0–63.255.255.255 address range (top left region) largely represents older (/8) address allocations.

Our analysis of these and many other quadcharts strongly indicates that there is significant structure in the BGP advertisements. For brevity, we omit detailed discussion but observe the following: allocation follows predictable patterns, address usage is largely static, geographically localized growth, and prefix size is highly correlated with the region in which it exists. Interested readers may obtain more detail about the visualization and our analysis of quadcharts from the extended version of this paper [17].

3 Quantitative analysis

Thus far, we have presented features of address use structure with the aid of visualization techniques. On its own, this high-level approximation of underlying events yields some interesting insights. To fully capture the micro-level dynamics that we are interested in (e.g., the nature of the origin moves between ASes), we further present our more quantitative approach to origin characterization that is driven by empirical analysis of real-world BGP traffic.




Figure 3. BGP advertisements on June 1st, 2004.

We note that while our analysis uses similar source material as in Meng *et al.*'s work [18], we characterize vastly different features of the data; that work makes no attempt to study the long-term stability of prefixes nor the reasons or effect of changing origin announcements. Additionally, our study is based on analysis of BGP updates in order to capture events that occur during the intervals between routing table snapshots. Before continuing, we first briefly present some necessary background information.

3.1 Background

Simply speaking, a BGP route is represented by a network prefix and an ordered list of ASes. The prefix denotes the destination, and the list of ASes constitutes the AS path to the destination. The last AS in the path is the origin AS, or simply the origin, of the prefix. Here we characterize BGP traffic properties and study BGP origin stability by examining archived real-world BGP data (BGP updates) gathered by Route Views [16]. During pre-processing we filter Bogon prefixes [19] and BGP Beacons [20]. Additionally, BGP updates which we believe are caused by the Route Views server, such as failures and subsequent reboots of the server etc., are ignored. Our characterization is based upon analysis of BGP updates from multiple vantage points, and spans the period of January 1st to December 31st of 2003.

We analyze BGP updates from tens of Route Views listening points. In choosing these viewpoints we attempted not only to capture some degree of diversity in terms of their connectivity, locations, and hierarchy levels, but also to include views from ASes that are affiliated with large ISPs and thus can be considered a core part of the Internet. As a result, these viewpoints provide a relatively complete

view of global activity. Since we observed similar prefix behaviors and patterns from these viewpoints, for brevity we choose to present our analyses of six of them (see Table 1). For simplicity, in the remainder of this paper we refer to these six listening points as Stockholm, AOL, Sprint, Level3, PSG, and Telstra viewpoints, respectively. In some instances we present the results from one or two particular viewpoints as representative views — when a global observation is required, we provide the results as a combined view from all these points.

Viewpoint	Organization	Location
217.75.96.60	Stockholm(P80-Net1)	Stockholm, Sweden
66.185.128.1	AOL	Virginia, U.S.
144.228.241.81	Sprint	California, U.S.
209.244.2.115	Level3	Colorado, U.S.
147.28.255.1	PSG	Washington, U.S.
203.62.252.26	Telstra	Sydney, Australia

Table 1. Viewpoints in our analysis.

We examine four types of events, namely *new prefix*, *prefix re-addition*, *removal*, and *origin change*, all of which we call *prefix movement* in general. We consider a prefix as new if this is the first announcement of the prefix (i.e, it did not appear at any prior point in the observation period). Prefix re-addition refers to the re-announcement of a prefix that was previously withdrawn and the announcing origin remains the same. Origin change, as the name implies, occurs when successive announcements of a particular prefix have different origin ASes, or if following a withdrawal, the prefix is announced by a different origin. Prefix removal simply refers to the withdrawal of a prefix.

4 Characterizing BGP traffic

For our study we analyze the BGP updates from the Route Views listening points for the entire year. The results show that origin changes do occur, though such changes are relatively few in relation to the vast amount of BGP updates. Table 2, for example, shows the new prefixes, prefix re-additions, removals, and the origin AS changes observed for each month from an example viewpoint, namely Stockholm in this case. For the most part, the results show that there is a significant number of prefix re-additions and removals. Of the announcements of prefixes that are not currently in the routing table, only 1.0%–5.3% can be considered new prefixes, while the rest represent re-additions of previously withdrawn prefixes to the routing table.

To characterize the repetitive announcements and subsequent withdrawals that are dominant in prefix movement, we examine the intervals between these events. Figure 4 shows the CCDF of prefix down time, that is, the elapsed

Period	New prefixes	Prefix re-additions	Prefix removals	Origin changes
Jan	14817	371784	386688	23214
Feb	10285	337345	347139	15766
Mar	8969	352586	362482	25009
Apr	13725	245928	258513	22699
May	11507	368131	380103	19239
Jun	9677	509752	521348	25021
Jul	6558	361901	366365	40544
Aug	6952	655482	663278	54968
Sep	7936	709223	717592	25226
Oct	8314	680055	687680	56734
Nov	6292	319019	326527	19870
Dec	10123	453389	463664	17883

Table 2. Prefix origin movement.

time from the point at which a prefix is withdrawn to when it is re-announced, and likewise, the CCDF of prefix up time. Events are clustered based on their contributing prefixes. In particular, we choose a threshold, δ , for the number of down times during the one-year period and cluster together events of those prefixes that are down more than the threshold, and label these as *frequently down* prefixes. Likewise, we consider the rest as *relatively stable* prefixes.




Figure 4. Prefix up time and down time for $\delta = 50$.

At $\delta = 50$ (see Figure 4), 8.3% of the prefixes are labeled frequently down prefixes, which manifest short *up-and-down* cycles. Moreover, 67% of the total down (up) events are contributed by these frequently down prefixes. Of these prefixes, 90% exhibit down periods of no more than 30 minutes, and 70% have up times that last no more than 30 minutes. Note that while an observed prefix outage may simply reflect connectivity problems related to it, we do assert whether these outages are observed at other viewpoints as well. This is certainly the case, and roughly 80%

of the prefixes experiencing more than 100 outages as observed from the Stockholm viewpoint, experience similarly high outages from the perspective of the other viewpoints.

5 Prefix stability

To better understand the movement patterns of different prefixes and the underlying dynamics, we further classify the prefixes based on the type of movement they undergo during the observation period. Let U be the set of prefixes that exist in the routing table at the beginning of the observation period. Our characterization is given as follows:

- *Type I Stabilized Prefix*: $\mathbb{P} \in U$, and is never withdrawn or subsequently announced with a new origin.
- *Type II New Prefix*: $\mathbb{P} \notin U$, is announced at some point, and keeps an unchanged origin.
- *Type III Removed Prefix*: $\mathbb{P} \in U$, and retains an unchanged origin until it is withdrawn.
- *Type IV Transient Prefix*: $\mathbb{P} \notin U$, is announced at some point and keeps an unchanged origin until it is withdrawn.
- *Type V Flapping Prefix*: \mathbb{P} is announced and withdrawn multiple times. However, it retains the identical origin AS between each announcement and subsequent withdrawal.
- *Type VI Origin-varying Prefix*: \mathbb{P} experiences origin change at least once during the observation period.

Table 3 presents the percentage of prefixes for each type observed at different viewpoints. The results show that new prefixes appear at a frequency slightly higher than removed prefixes, which therefore depicts the overall growth of prefix usage. Additionally, a fairly large number of transient prefixes appear in the routing system for only a short period of time. The majority of prefixes, however, experience multiple up and down times and are thus classified as flapping prefixes. On the other hand, origin-varying prefixes account for only about 9% of all observed prefixes; all other types maintain their identical origin ASes during their life period. Clearly, this indicates that a significant portion of prefixes have very high origin stability.

On closer inspection, it appears that address space stability is related to space fragmentation. For example, Figure 5 shows the percentage of each prefix type with respect to the total prefixes at a particular length. For example, from the Stockholm viewpoint we observed a total of 139,278 /24 prefixes, among which 73% are *Type V* prefixes. The results illustrate that for fragmented prefixes (e.g., /25–/32), there are relatively more cases of new prefixes (*Type II*),

prefix removals (*Type III*), and transient prefixes (*Type IV*), while for larger address spaces (e.g., /8–/24), the instability mainly manifests as prefixes being up and down (*Type V*).




Figure 5. The percentage of each prefix type with respect to the total prefixes at a particular length.

The flapping prefixes are of particular attention to us, therefore we examine the times each prefix flaps during the one-year period. It turns out that the majority of prefixes are down a few times, which is reasonable from an operational point of view; 61% are down 10 times or less, but a few outliers experience outages of more than 100 times (observed at multiple viewpoints). Clearly, this reveals the instability of these prefixes and implies potential problems related to them. We return to an examination of the ASes driving this behavior in Section 6.

For the prefixes that undergo origin changes, we further explore the degree to which these changes occur as well as their frequency. First, we examine the number of observed origin ASes for each particular prefix during its lifetime. As shown in Table 4, most prefixes have very few origin ASes. For example, about 91% have only one origin AS, 8% are observed with two origin ASes, and less than 1% have more than two. This implies that origin changes for a prefix usually occur between very few ASes.

No. of origin ASes	Stockholm		AOL	
	#	%	#	%
1	212,744	91.10%	212,170	91.0%
2	18,820	8.06%	19,965	8.52%
3	1,756	0.75%	1,910	0.82%
4	186	0.08%	212	0.09%
5	18	7.70e-05	24	1.02e-04
≥6	13	5.57e-05	16	6.83e-05

Table 4. Prefixes with multiple origins.

	Stockholm	AOL	Sprint	Level3	Telstra	PSG
Stabilized (Type I) (%)	0.2	3.4	3.0	2.8	0.8	2.7
New (Type II) (%)	2.2	2.9	2.7	2.9	2.2	2.3
Removed (Type III) (%)	1.6	1.8	1.7	1.7	1.4	1.2
Transient (Type IV) (%)	13.7	15.1	13.4	13.5	13.8	12.7
Flapping (Type V) (%)	73.4	67.3	69.9	69.9	72.4	72.1
Origin-varying (Type VI) (%)	8.9	9.5	9.3	9.2	9.4	9.0

Table 3. Prefix classification based on the movement they underwent in 2003.

Second, we examine the number of origin changes experienced by those origin-varying prefixes across the one-year period. The results of Stockholm and AOL viewpoints are given in Table 5. About 57% of these prefixes have their origin ASes changed once, indicating that these changes are indeed permanent. An additional 14% have their origin ASes changed twice, while few prefixes have relatively more changes. In the case of a large number of changes, we observe that some are due to multi-origin prefixes oscillating between their origin ASes and the oscillating cycles are even much shorter than those found in a previous study [21].

No. of changes	Stockholm		AOL	
	#	%	#	%
1	11,782	56.7%	12,702	57.5%
2	2,861	13.8%	3,010	13.5%
3	1,310	6.3%	1,413	6.4%
4	602	2.9%	783	3.5%
5	393	1.9%	431	2.0%
6	348	1.7%	396	1.8%
7	309	1.5%	344	1.6%
8	260	1.3%	227	1.0%
9	176	0.8%	155	0.7%
≥10	2,749	13.1%	2,636	12.0%

Table 5. Frequency of origin changes.

We believe it is worthwhile to characterize the major contributors with respect to each aforementioned event type, and so in what follows, we investigate this relationship by characterizing movement based on the involved origin ASes.

6 Characterizing advertisement stability from the AS perspective

To characterize origin advertisement stability from the AS perspective, we cluster the origin ASes based on the number of events the ASes are involved in during the one-year period. Figure 6 shows the results gleaned from the

Stockholm viewpoint.² As shown in Figure 6, most ASes originate few (if any) new prefixes and are involved in even fewer origin changes. About 50% did not originate any new prefixes and 68% were not involved in any prefix origin changes.

We further segment origin changes into *move-from* or *move-to* sets. For example, if the origin of prefix \mathbb{P} changes from AS A to AS B, we say that A is involved in a move-from event and B is involved in a move-to event. Figure 6 also shows how ASes are clustered based on these two events. Again, it is evident that most ASes are not involved in either event while a relatively small number of ASes are responsible for a significant number of these events.




Figure 6. Percentages of ASes that originate new prefixes and ASes that are involved in origin changes.

By contrast, substantially more ASes are involved in repetitive re-additions and subsequent removals. In particular, only 439 origin ASes were not related to any re-addition and only 209 were not related to any removal. The CDF of ASes contributing to re-additions is presented in Figure 7; about 27% were involved in 10 or less, and 75% in 100 or less. The corresponding CDF for removals is quite similar and thus omitted from our discussions. It appears that,

²The corresponding figures for other listening points are quite similar and so are omitted.




Figure 7. CDF of ASes that are involved in repetitive re-additions.

for each event type, an overwhelming number of events are contributed by a relatively very small number of ASes. That is, there are a few ASes that are actively manipulating the addresses they originate. The vast majority of others appear to alter the set of prefixes they originate infrequently.

To better examine these major contributors, we extract the top ASes and study their characteristics. For each of the six viewpoints, we retrieve the top 30 ASes that originate the most new prefixes. We then extract the common elements of the six sets. For the subsequent discussion, let \mathbb{A} denote the intersection over the 30 elements in each set. Similarly, we extract the common ASes that contribute the most to prefix re-additions, withdrawals, and prefix origin changes, and denote these sets as \mathbb{B}_1 , \mathbb{B}_2 , and \mathbb{C} , respectively. From the observed viewpoints, set \mathbb{A} has 23 elements, \mathbb{B}_1 and \mathbb{B}_2 have exactly the same 15 elements, and \mathbb{C} has 12 elements. We then examine their degrees of connectivity based on the AS topology graph that is derived from the routing tables (of the same period) of Route Views servers and calculate their degree rankings accordingly.

It is interesting that a significant number of the members of sets \mathbb{A} , \mathbb{B}_1 and \mathbb{C} are among the top ASes in terms of their degree of connectivity — 17%, 33%, and 50%, respectively, of the ASes in these sets are among the top 30 connected ASes. However, the degree of connectivity may not be enough to reflect the AS properties. Recently, Subramanian *et al.* provided a grouping of the ASes to five hierarchy levels (namely, dense core, transit core, outer core, small regional ISPs and customers) that is based on both the degree of connectivity and AS relationships, e.g., provider-customer and peer-peer relationships [3]. These hierarchy levels represent how close an AS is to the core of the Internet, with dense core being the closest and customer the farthest. In that paper, a total of 20 ASes were classified as dense core, 162 as Transit core, 675 as outer core, 950

as small regional ISPs, and 8,852 as customers. Here we adopt this classification, and present the number of ASes of different hierarchy levels that comprise the observed sets \mathbb{A} , \mathbb{B}_1 , and \mathbb{C} in Table 6.

	Dense core (20)	Transit core (162)	Outer core (675)	Small regional ISPs (950)	Customers (8852)
\mathbb{A} (23)	2	7	5	4	5
\mathbb{B}_1 (15)	4	2	1	4	4
\mathbb{C} (12)	6	0	3	2	1
\mathbb{D} (12)	6	0	3	2	1
\mathbb{E} (10)	4	0	3	2	1

Table 6. Hierarchy levels of the top ASes contributing to prefix movement.

Under this reclassification, it appears that a significant portion of prefix movement tends to happen between highly connected (i.e., core-level) ASes. Normally, the degree of connectivity is relevant to the size of the AS, so it is not surprising that highly-connected, large ASes originate more prefixes and are therefore more likely related to more events. However, it is odd that a few ASes with very low degree are among this grouping of top contributing ASes. These ASes are the primary cause of much of the instability in the origin advertisements and appear to be unnecessarily increasing BGP traffic load at a global level. In this case, the instability may indeed reflect an unreasonable policy or mis-configuration on their part.

What is particularly interesting about the top ASes that generate the most new prefixes, is that many are relatively small or edge ASes that can be attributed to organizations, regions, or countries that are actively evolving or transforming. As a case in point, 5 ASes in set \mathbb{A} are related to China which clearly reflects the fast development of the Internet and booming economy in that country. Another noteworthy example is the acquisition of Genuity by Level 3. During that transition Genuity transferred many of its prefixes to Level 3, and both of these appear in set \mathbb{C} .

Lastly, if we denote sets \mathbb{D} and \mathbb{E} as the common top 30 ASes involved in the most move-from and move-to events, respectively, then 6 ASes in \mathbb{D} (50%) and 4 in set \mathbb{E} (40%) are also among the top 30 ASes in terms of their degree of connectivity. The members of these sets are also given in Table 6. Observe that since most of the ASes in these sets are the same, this suggests that these movements are likely due to multi-homed prefixes oscillating between these ASes.

7 Other related work

For several years Huston [22] has provided reports of observed multi-origin prefixes. Zhao *et al.* subsequently showed that most multiple origin AS conflicts are short-lived (lasting a small number of days) based on analysis of daily archived routing tables [21]. Teoh *et al.* [13] introduced visual-based anomaly detection for origin AS change events. Recently, Meng *et al.* [18] measured the IPv4 address allocation and usage, the advertisement of fragmented and aggregated addresses, and their impact on the global routing table size. With respect to Internet topologies, numerous efforts [4, 5, 6] have been undertaken to measure ISP-level or router-level topologies. Subramanian *et al.* proposed a five-level classification of ASes based on the connectivity of ASes and the AS relationships [3].

8 Summary

In this paper, we study the structure and stability of origin advertisements in inter-domain routing. Via visualization, we confirm that address use exhibits consistent gross structure over long periods. In particular, this structure is manifest in, among other ways, predictable address allocation patterns, regional similarity in advertisement size, and geographically localized growth. Furthermore, we characterize the stability of origin advertisements in inter-domain routing from a variety of perspectives. We show that BGP updates can be characterized by a significant amount of repetitive prefix re-additions and subsequent withdrawals. More than 90% of the prefixes are consistently originated by the same AS for an entire year and a small number of ASes are responsible for most of the advertisement churn. Moreover, a fairly large percentage of new prefixes can be attributed to growth in actively evolving and/or transforming organizations and countries and the observed (abnormal) prefix flapping with unreasonably high frequency is likely due to misconfigurations. We argue that this understanding of the structure and stability of BGP origin advertisements is important to a number of inter-domain routing research areas.

9 Acknowledgments

This work is supported by NSF Grant 0334108.

References

- [1] Y. Rekhter and T. Li, "A Border Gateway Protocol (BGP-4), RFC 1771," 1995.
- [2] L. Gao, "On inferring autonomous system relationships in the Internet," in *IEEE Global Internet*, 2000.

- [3] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *IEEE INFOCOM*, 2002.
- [4] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet map discovery," in *INFOCOM*, 2000.
- [5] *Skitter*, <http://www.caida.org/tools/measurement/skitter/>.
- [6] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in *SIGCOMM*, 2002.
- [7] C. Labovitz, R. Malan, and F. Jahanian, "Origins of Internet routing instability," in *INFOCOM*, 1999.
- [8] A. Feldmann, O. Maennel, Z. Mao, A. Berger, and B. Maggs, "Locating Internet routing instabilities," in *SIGCOMM*, 2004.
- [9] S. Kent, C. Lynn, and K. Seo, "Secure Border Gateway Protocol (Secure-BGP)," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 4, pp. 582–592, April 2000.
- [10] J. Ng, "Extensions to BGP to support secure origin BGP (draft)," in *Network Working Group*, 2003.
- [11] G. Goodell, W. Aiello, T. Griffin, J. Ioannidis, P. McDaniel, and A. Rubin, "Working around BGP: An incremental approach to improving security and accuracy in inter-domain routing," in *IEEE Network and Distributed Systems Security Symposium*, 2003.
- [12] W. Aiello, J. Ioannidis, and P. McDaniel, "Origin authentication in inter-domain routing," in *ACM Conference on Computer and Communications Security*, 2003.
- [13] S. Teoh, K. L. Ma, S. F. Wu, D. Pei, L. Wang, L. Zhang, D. Massey, and R. Bush, "Visual-based anomaly detection for BGP origin as change (OASC) events," in *IEEE/IFIP Workshop on Distributed Systems: Operations and Management*, 2003.
- [14] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. Merwe, "Design and implementation of a Routing Control Platform," in *Network Systems Design and Implementation*, 2005.
- [15] W. Schroeder, K. Martin, and B. Lorenson, *The Visualization Toolkit*, Kitware Inc, 2004.
- [16] *University of Oregon Route View Project*, <http://routeviews.org>.
- [17] S. Qiu, P. McDaniel, F. Monrose, and A. Rubin, "Characterizing address use structure and stability of origin advertisement in Interdomain routing," Tech. Rep. NAS-TR-0018-2005, Network and Security Research Center, Dept. of CSE, Pennsylvania State University, University Park, PA, USA, July 2005.
- [18] X. Meng, Z. Xu, B. Zhang, G. Huston, S. Lu, and L. Zhang, "IPv4 address allocation and the BGP routing table evolution," in *ACM SIGCOMM Computer Communication Review, Special Issue on Internet Vital Statistics*, 2005.
- [19] *Bogon IPs*, <http://www.complethewhois.com/bogons/>.
- [20] *Historical List of RIS Routing Beacons*, <http://www.ripe.net/projects/ris/docs/ beaconlist.html>.
- [21] X. Zhao, D. Pei, L. Wang, D. Massey, A. Mankin, S.F. Wu, and L. Zhang, "An analysis of BGP multiple origin AS (MOAS) conflicts," in *Internet Measurement Workshop*, 2001.
- [22] G. Huston, *BGP Table Data*, <http://bgp.potaroo.net>.