# 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
Department of Computer Science *Department of Computer Science and Engineering
The Johns Hopkins University 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. Using our `q-chart` IP address advertisement visualization tool, we explore the gross structure of IP address advertisements and show that it exhibits considerably consistent structure. We further quantitatively characterize the stability of origin advertisements by analyzing real-world BGP updates for a period of one year from multiple vantage points. We show that while repetitive prefix re-additions and subsequent withdrawals constitute a major volume of BGP updates — due in part to a number of frequently flapping prefixes with short up-and-down cycles — a significant portion of prefixes have high origin stability. In particular, origin changes account for less than 2% of the BGP update traffic, with more than 90% of the prefixes being consistently originated by the same AS for an entire year. For those prefixes involved in origin changes, approximately 57% have only one change across the year, implying that these changes are indeed permanent. We also show that most ASes are involved in few prefix movement events, if any, while a small number of ASes are responsible for most of the advertisement 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 most of the origin changes are a result of multi-homed prefixes oscillating between their origins. This work not only contributes to a better understanding of BGP dynamics, but also provides insights for other research areas such as BGP security that rely on key assumptions pertaining to origin stability.**

## I. 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 not only new mechanisms to address these problems but for enabling new services (e.g., QoS, traffic engineering) as well. 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 [2] and address space [3] are widely used to ascertain the health of the Internet, and tools are used by many to construct maps of the organizational [4], [5], [6] or topological [7], [8], [9], [10], [11] relationships between Autonomous Systems (ASes). BGP has also been the subject of considerable introspection within the academic community. Key among these are studies of the dynamic behavior of BGP, particularly as it relates to instability [12], [13], [14], convergence [15], [16], [17], [18], [19], and security [20], [21], [22], [23], [24], [25].

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. Moreover, while periodic reports and recent research present broad views about what parts of the address space are being advertised, they provide little insight on how origin advertisements evolve. We argue in this paper that such an understanding is essential to building a robust, responsive, and secure Internet.

In this paper, we develop a deep understanding of address space dynamics by studying the structure and stability of origin advertisements in inter-domain routing. We begin by first exploring the gross structure of IP address advertisements using our `q-chart` 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. For instance, it becomes immediately apparent that the address advertisements exhibit considerably consistent structure which is manifest in, among other ways, predictable address allocation patterns, (address range) regional similarity in advertisement size, and geographically localized growth. Our visualization methodology is simple, straightforward and vivid. As we show later, while these (and other) visualization techniques 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 III–VI 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. Broad classes of prefix behaviors are developed from the analysis of BGP updates. 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. We further show that origin changes account for less than 2% of the BGP update traffic, with more than 90% of the prefixes consistently originated by the same AS for an entire year. For those prefixes involved in origin changes, approximately 57% have only one change across the year, implying that these changes are permanent. Likewise, in the case of frequent changes, most are due to multi-origin prefixes [26] oscillating between their origin ASes.

Furthermore, we show that most ASes are involved in few, if any, movement events, while a small number of ASes are responsible for most of the origin churn. We extract the ASes that introduce the largest number of prefix movement (e.g., prefix re-additions, removals, origin changes). We then examine their degrees of connectivity and classify them by their positions in BGP topology. A fairly large percentage of new prefixes can be attributed to growth in actively evolving and/or transforming organizations and countries (e.g., China). We also show that the observed (abnormal) prefix flapping with unreasonably high frequency is likely due to misconfigurations, and that some culprit ASes characterize the places where multi-origin prefixes oscillate.

The analysis of BGP traffic and study of prefix behaviors presented in this paper provide 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 security proposals rely on: a significant number of prefixes have high origin stability (see for example [21], [22], [24]). Furthermore, our analysis may be directly applicable to a number of other inter-domain routing research areas. For example, Rexford *et al.* recently introduced the Routing Control Platform (RCP) [27], [28] which is aimed at mitigating potentially harmful interactions between iBGP and eBGP, as well as simplifying the administration of an AS. The protocol and associated infrastructure, however, require significant changes in the way BGP operates, and hence should take into consideration the dynamics of origination. Similarly, Aiello *et al.* noted that the performance of any security solution for origin authentication is directly related to the structure and stability of IP address use [24]. We believe our work provides a deeper understanding of such factors and thus represents essential data for the designers of such systems. The paper also reveals some potential problems that need further investigation and improvement; for example, repetitive prefix re-additions and subsequent withdrawals still constitute a significant part of BGP traffic, routes originated from some ASes manifest instable characteristics, and a small number of culprit ASes are responsible for a significant number of movement events.

PAPER OUTLINE: We begin in the following section by exploring the characteristics of address use on the Internet via visualization. Sections III–VI describe our empirical efforts at origin characterization using statistical analysis of the real-world BGP update streams. We conclude with a brief discussion of related work and the implications of our findings to routing protocols, security, and network stability.

## II. VISUALIZING PREFIX ADVERTISEMENT

One way to illuminate hidden structure in complex datasets is through visualization [29]. 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. Visualization using A Quad Chart

*Quad chart* is an intuitive and simple way of visualizing hierarchical data (such as the IP address space). Illustrated in Figure 1, the quad charts used throughout this section define advertised prefixes as the regions in a ($2^{16}$x$2^{16}$) 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.

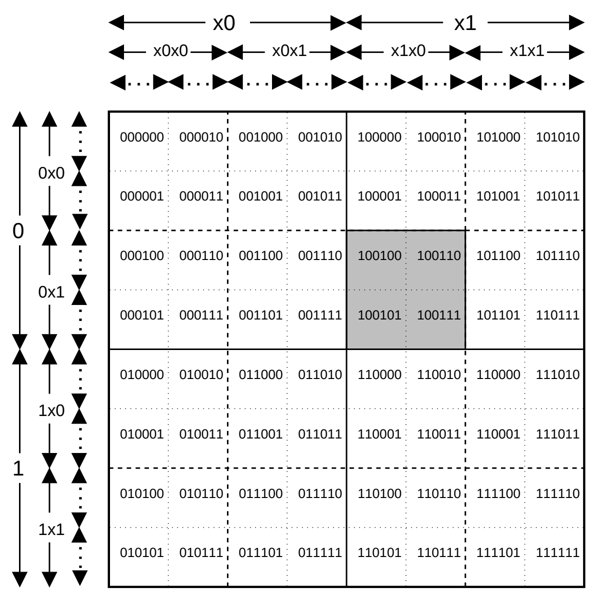


Fig. 1. Quad chart - prefixes are encoded to cover rectangular regions of the plane, where x and y axes are defined by the most significant bits and the size of the square is defined by the prefix length. For example, the shaded region represents the prefix 144.0.0.0/4.

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

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

where $bit(i)$ 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 (e.g., farthest distance in hue/brightness encoding, and extremely dark colors are avoided to ensure differentiation from unused space). Color selection is also deterministic to ensure assignment coloring is consistent across graphs (e.g., the same AS is assigned the same color in all graphs). Squares whose size is below the resolution depth of the graph (e.g., JPG) are promoted to a size equal to a single pixel.

### B. Visual Usage Features

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 [30].[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 addresses) and 240.0.0.0–255.255.255.255 are the reserved Class E addresses (for experimental purpose). These ranges consume one eighth of the total IP address space, and are shown as the large black region on the right half of the lower right quadrant of both graphs. Much of the following discussion is aided by the vast body of useful information provided by ICANN about the current use of the IP address space [31].

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 latter RIB. The 64/8 space is largely allocated to ARIN (the North American address registry[2]), so it is likely that this reflects the proliferation of commercial and personal networks within the United States, Canada, and other served countries.

Interestingly, 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 address allocation. Organizations which would receive Class B or smaller prefixes today received Class A prefixes in the 1980s and early 1990s. For example, the Ford Motor company received a /8 (one 256th of all the IPv4 addresses space) in May of 1995. It appears that this and similar address assignments have led to under-utilized address space. Only a small fragment of these large address allocations are actually being advertised.[3] In the case of the Ford allocation, the /8 was not advertised in 2001 or 2004, and remains unused as of the time of this writing. The fact that the number of allocations in the 0–63/8 range has not increased demonstrably also supports this thesis. Note that several /8s that were advertised in 2001 were not in 2004. This probably indicates that the owning organization either no longer uses it (e.g., preferring to use address space assigned to it by an ISP), or has returned it to ICANN as Stanford University did in April of 1993.

The graphs also expose a key feature of the allocation of address space, where each block appears to be allocated from the top left to the bottom right. This makes sense because the prefixes numerically increase in this direction. Hence, at apparently all levels, addresses are generally allocated from the numeric bottom of the larger assigned block toward the top, i.e., the `*.1`s are allocated before `*.2`. This assertion is strengthened when looking at the time-lapse graphs. Early half-used (and presumably partially allocated) address spaces are subsequently fully advertised in the later graph.

Finally, the graphs illustrate an interesting fact about the interaction between delegation, advertisements, and prefix size. Class Bs (/16s) are largely assigned out of the 128.0.0.0–191.255.255.255 block (top right quadrant), and Class Cs are largely assigned out of 64.0.0.0–126.255.255.255[4] (lower left) and 196.0.0.0–255.255.255.255 (lower right). This viewpoint allows for some interesting insights about the use of this address space: large (/16) delegations of address space occur out of the 128–192 address regions, while smaller ones (/24) are allocated (either by registries or by ISPs) out of the others. This is largely a reflection of the behavior of registries and

[1] A highly detailed q-chart movie illustrating the evolution of the address space over 4 years is available at: `http://www.patrickmcdaniel.org/bgp/ 2001-2004_30sec_nomartians.mov`

[2] ARIN currently serves areas other than North America, notably including parts of Africa.

[3] The lack of advertisement could be explained by the addresses not being connected to the Internet. However, given the sheer number of unused address and the reliance of modern computing on global resources, this seems unlikely to be a major factor.

[4] 127.0.0.1/8 is reserved, as are 0.0.0.0/8 and 14.0.0.0/8.

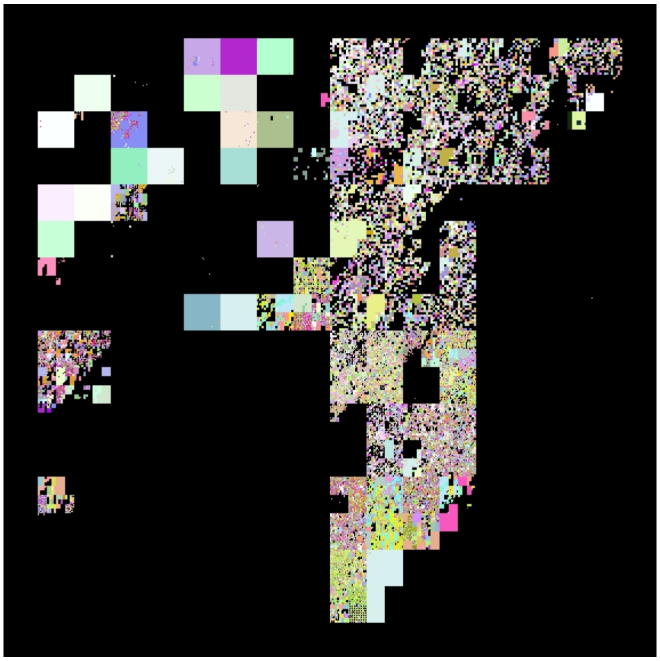

Fig. 2. BGP advertisements on January 1st, 2001.

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

history of the underlying allocation. The term Class B (Class C) was originally defined to not only refer to prefix length 16 (24) but also require that the first bits of the address be 10 (110). This restricted the Class B networks to addresses in the 128.*–191.* range and Class C networks to the 192.*–223.* range. Hence, while the current allocation procedures for Class Bs largely respect the original definition, assignments for Class Cs do not.

This brief and informal visualization strongly indicates that there is structure in the BGP advertisements. Allocation follows predictable patterns, usage in certain regions is very unlikely to change over time, and prefix size is highly correlated with the region in which it exists. This study also suggests that much of the IP address space is under-utilized. Vast regions of the 0–64/8 address space remain unused and unadvertised, and the vast majority of the 64–195* address space is being held as reserved. Hence, while possibly controversial, nearly half and possibly more of the address space could be recovered and used as is needed.

## III. Quantitative analysis

Thus far, we have presented features of address use structure on the Internet with the aide 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. Our study is based on analysis of BGP updates so that we can capture events that occur during the intervals between routing table snapshots. We note that while our analysis uses similar source material as in Meng *et al.* 's work [32], 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. Before continuing further, we first briefly present some necessary background information.

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 [30]. During pre-processing we filter Bogon prefixes [33] and BGP Beacons [34], [35]. 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 $1^{st}$ to December $31^{st}$ of 2003.

Table I lists the six Route Views listening points used in our study. In choosing these viewpoints we attempted not only to capture some degree of diversity, but also include views from ASes that are affiliated with large ISPs and thus can be considered a core part of the Internet. As a result, the chosen viewpoints provide a relatively complete view of global activity. 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. Furthermore, because we observed similar prefix behaviors and patterns from these viewpoints, in some instances we only provide the results from one or two particular viewpoints.

| 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 I

THE LIST OF ROUTE VIEWS LISTENING POINTS USED AS VIEWPOINTS IN OUR ANALYSIS

| Viewpoint | Prefix | Origin AS | % of path changes | % of origin changes |
|---|---|---|---|---|
| Stockholm | 233537 | 18152 | 46% | 0.51% |
| AOL | 234297 | 18214 | 35% | 1.64% |
| Sprint | 220494 | 18007 | 31% | 1.38% |
| Level3 | 219727 | 18052 | 35% | 1.14% |
| PSG | 230537 | 18157 | 45% | 0.63% |
| Telstra | 226959 | 18183 | 42% | 1.02% |

TABLE II

BASIC BGP TRAFFIC STATISTICS OF 2003

In what follows, 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.

## IV. CHARACTERIZING BGP TRAFFIC

For our study we analyze the BGP updates from the Route View listening points for the entire year. Table II lists the number of distinct prefixes, origin ASes, the percentage of BGP announcements that are related to changes in advertised AS path, and the percentage of AS path changes that are involved in origin changes. Since different viewpoints exhibit different connectivity and operational characteristics one can expect slight differences in the absolute values. The results show that origin changes do occur, though such changes are relatively few in relation to the vast amount of BGP announcements, and BGP updates. Table III, 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[5]. For the most part, the results show that there is a significant amount 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 (more than 94%) 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 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

[5] We picked the Stockholm viewpoint here simply because it was the first viewpoint used in our study and data from Stockholm was representative of all the viewpoints.

| 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 III

PREFIX ORIGIN MOVEMENT IN 2003, FROM STOCKHOLM VIEWPOINT

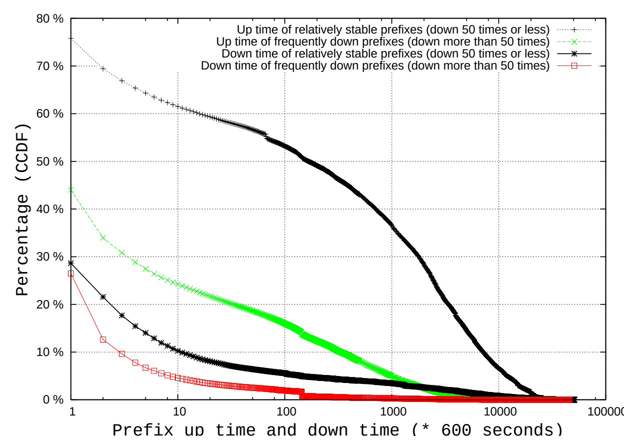


Fig. 4. Prefix up time and down time observed in 2003, from Stockholm viewpoint, with $\delta = 50$. X axis is in logarithmic scale and each unit represents 600 seconds.

are clustered based on their contributing prefixes. In particular, we choose a threshold, $\delta$, 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.

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; 69% have up times of no more than 30 minutes, while 80% are up for no more than 5 hours. At $\delta = 100$ (not shown), 3.4% of the prefixes are considered frequently down, with 54% of the overall down (up) events attributable to them. In this case, 92% of these prefixes have down times of no more than 30 minutes, and 80% have up times that last no more than 1 hour. 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.

## V. 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. Such a classification allows us to gain a microscopic view on the relationship present within different classes of movement. 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 different origin.
- *Type II New Prefix:* $\mathbb{P} \notin U$, is announced at some point, and remains with its origin intact throughout the observation period.
- *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 up to the point of withdrawal.
- *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-nonstatic Prefix:* $\mathbb{P}$ experiences origin change at least once during the observation period.

Table IV 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 amount 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-nonstatic 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 there are 101,989 *Type V* prefixes, i.e., 73%. Since very few prefixes with length below 8 are observed, they are omitted in Figure 5. 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 being up and down (*Type V*).

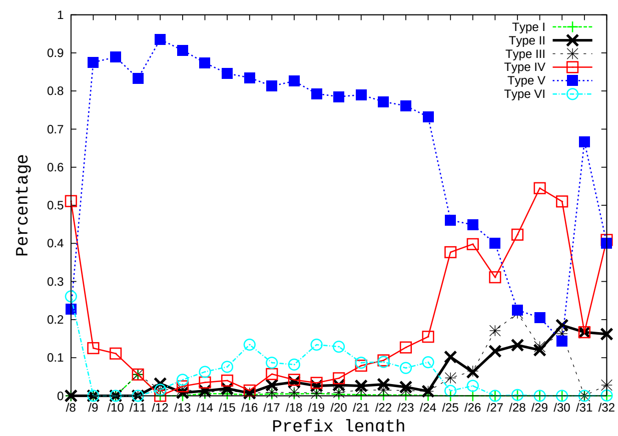


Fig. 5. The percentage of each prefix type with respect to the total prefixes at a particular length.

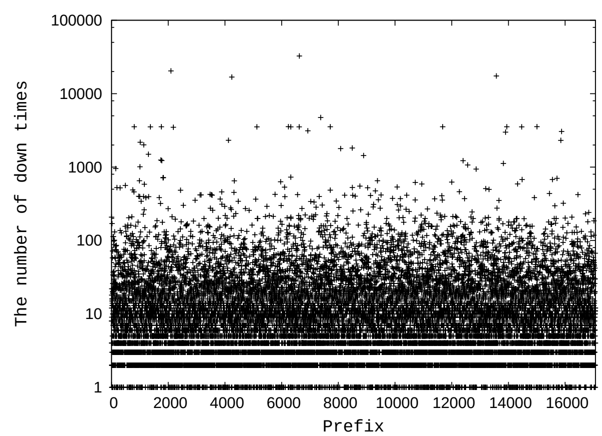


Fig. 6. The down times for prefixes announced and withdrawn repetitively in 2003. To aid visualization, we randomly dropped 90% of the prefixes (i.e., points) in this figure.

The flapping prefixes, which are announced and subsequently withdrawn repetitively, are of particularly attention to us, so we examine the times each prefix flaps during a one-year period. Figure 6 illustrates the down times of these flapping prefixes. It turns out that the majority of prefixes are down a few times over the year, which is reasonable

| | 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-nonstatic (Type VI) (%) | 8.9 | 9.5 | 9.3 | 9.2 | 9.4 | 9.0 |

TABLE IV

PREFIX CLASSIFICATION BASED ON THE MOVEMENT THEY UNDERWENT IN 2003

from an operational point of view; 41% of prefixes are down 5 times or less, 61% 10 times or less, 92% 50 times or less, and 97% 100 times or less. As shown in the figure, most down times are less than 100 while a few outliers are up and down far more than the others. Nearly 80% of the outliers experiencing outages more than 100 times from this viewpoint are also observed at other viewpoints with similar frequency. Clearly, this reveals the instability of these prefixes and implies potential problems related to them. For example, the most frequently down prefixes observed from the Stockholm viewpoint are down more than 60,000 times during the one-year period, which is equivalent to 7 times per hour on average. A natural question arises as to the nature of these outliers and the ASes they belong to. Undoubtedly, the distinct pattern within the outliers implies some underlying cause-effect. We return to an examination of those ASes driving this behavior in Section VI.

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 V, most prefixes have very few origin ASes. For example, about 91% have only one origin AS all along, 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 V

THE NUMBER ("#") AND PERCENTAGE ("%") OF PREFIXES THAT HAVE A PARTICULAR NUMBER OF ORIGIN ASES

Second, we examine the number of origin changes experienced by those origin-nonstatic prefixes across the one-year period. The results of Stockholm and AOL viewpoints are given in Table VI. 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 do check and verify that some are due to multi-origin prefixes [26] oscillating between their origin ASes. Other possible causes include traffic engineering, and faulty or malicious configurations [36].

| 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 VI

THE NUMBER ("#") AND PERCENTAGE ("%") OF PREFIXES THAT UNDERGO A PARTICULAR NUMBER OF ORIGIN CHANGES IN 2003

We believe it is worthwhile to characterize the major contributors with respect to each of the aforementioned event types and thus in the following section, we investigate this relationship, by characterizing movement based on the involved origin ASes.

## VI. CHARACTERIZING ADVERTISEMENT STABILITY FROM THE AS PERSPECTIVE

To characterize origin advertisement stability from the AS perspective, we examine the BGP updates of 2003, and investigate which ASes originate new prefixes, how they contribute to prefix re-additions and removals, and which ASes are involved in origin changes.

We cluster the origin ASes based on the number of events the ASes are involved in during the one-year period. Figures 7 and 8 show the results gleaned from the Stockholm viewpoint. The corresponding figures for other listening points are quite similar and so are omitted. As shown in Figure 7, most ASes originate few (if any) new prefixes and are involved in even fewer origin changes. More specifically, about 50% did not originate any new prefixes, 20% originated one new prefix, and 92% originated 10 or less; 68% were not involved in any prefix origin changes, 8% involved in one origin change, and

88% involved in 10 or less. Furthermore, we 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 7 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.

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. Out of the ASes that were involved in at least one re-addition, 27% were involved in 10 or less, 75% 100 or less, and 90% 340 or less. Figure 8 presents the CDF of ASes contributing to re-additions. The corresponding CDF for removals is quite similar and thus omitted from our discussions.

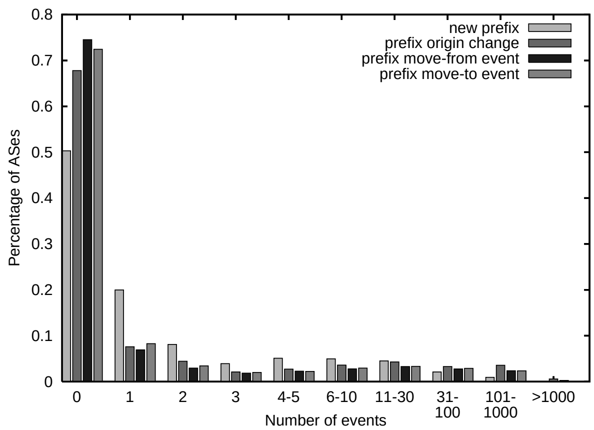


Fig. 7. Percentages of ASes that originate new prefixes and ASes that are involved in origin changes.

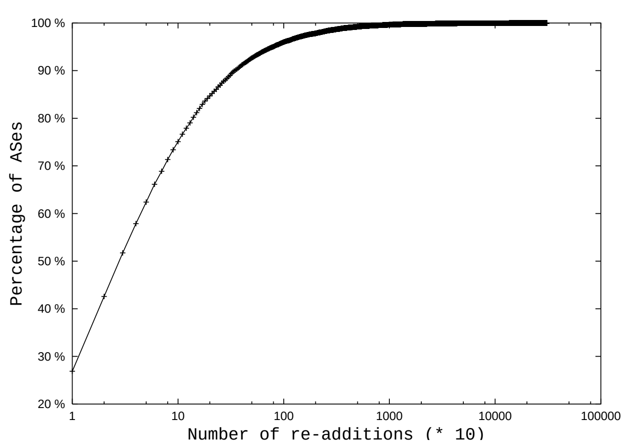


Fig. 8. CDF of ASes that are involved in repetitive re-additions. Each unit of X axis represents 10 re-additions.

It appears that, 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. Table VII shows the percentages of events contributed by the top ASes that introduced the most movement, observed from the Stockholm and AOL viewpoints, respectively. Note that in the table, the top ASes for different event types may denote different sets of ASes.

| | | Top 5 | Top 10 | Top 20 |
|---|---|---|---|---|
| Stock-holm | new prefix | 13.0% | 17.5% | 24.0% |
| | prefix re-addition | 16.8% | 21.1% | 27.9% |
| | prefix removal | 16.4% | 21.7% | 27.6% |
| | move-from event | 23.2% | 30.0% | 39.9% |
| | move-to event | 23.0% | 29.8% | 39.7% |
| AOL | new prefix | 16.3% | 22.3% | 29.5% |
| | prefix re-addition | 13.5% | 18.0% | 24.1% |
| | prefix removal | 13.1% | 17.8% | 23.9% |
| | move-from event | 25.7% | 34.2% | 43.1% |
| | move-to event | 25.5% | 34.2% | 43.4% |

TABLE VII
THE PERCENTAGES OF EVENTS CONTRIBUTED BY TOP ASES

To further 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. Tables XI, VIII, and XII (Tables XI and XII are in the Appendix) show the ASes in sets $\mathbb{A}$, $\mathbb{B}_1$, and $\mathbb{C}$ as well as their degrees of connectivity and degree rankings.[6] The elements are listed based on their orders observed from Stockholm viewpoint. The degree of connectivity of the ASes are calculated from the AS topology graph that is derived from the routing tables (of the same period) of Route Views servers. The degree rankings are calculated accordingly. Note that some ASes, especially those sparsely connected ASes, have the same degrees and thus the same rankings.

It is interesting that a significant number of the members of sets $\mathbb{A}$ (Table XI), $\mathbb{B}_1$ (Table VIII) and $\mathbb{C}$ (Table XII) 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 [5]. These hierarchy levels represent how close an AS is to the core of

[6]Set $\mathbb{B}_2$ is omitted as it has exactly the same elements as $\mathbb{B}_1$.

| AS | Organization | Degree | Degree rank |
|---|---|---|---|
| 14846 | NBCI NBC Internet | 1 | 13883 |
| 1295 | GE-AEROSPACE-AS General Electric Company | 1 | 13883 |
| 3921 | GEIS-OOCL General Electric Co. | 1 | 13883 |
| 702 | ALTERNET-AS UUNET Technologies, Inc. | 768 | 9 |
| 705 | ALTERNET-AS UUNET Technologies, Inc. | 6 | 1664 |
| 4755 | VSNL-AS Videsh Sanchar Nigam Ltd. Autonomous System | 72 | 119 |
| 209 | ASN-QWEST Qwest | 1296 | 4 |
| 701 | ALTERNET-AS UUNET Technologies, Inc. | 3213 | 1 |
| 7713 | TELKOMNET-AS-AP PT TELEKOMUNIKASI INDONESIA | 8 | 1243 |
| 1239 | SPRINTLINK Sprint | 2279 | 2 |
| 4621 | UNSPECIFIED UNINET-TH | 22 | 428 |
| 2048 | LANET-1 State of Louisiana | 2 | 7044 |
| 724 | DLA-AS724-ASN DISA PACIFIC | 39 | 214 |
| 7018 | ATT-INTERNET4 AT&T WorldNet Services | 1995 | 3 |
| 17557 | PKTELECOM-AS-AP Pakistan Tele. | 16 | 610 |

TABLE VIII

THE COMMON ASES AMONG THE TOP 30 ASES, OBSERVED FROM ALL SIX LISTENING POINTS, WHICH ARE RELATED TO THE MOST REPETITIVE RE-ADDITIONS IN 2003

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 IX.

| | 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 |

TABLE IX

HIERARCHY LEVELS OF THE TOP ASES CONTRIBUTING TO PREFIX MOVEMENT

Under this reclassification, it appears that there is a strong correlation between the origin advertisement stability and the hierarchy levels of origin ASes — 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 (e.g., General Electric, NBC, and State of Louisiana in Table VIII have degree $\leq 2$ ) 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 or mis-configuration on their part.

What is particularly interesting in the case of 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 and/or transforming. As a case in point, note that for example, 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 transfered 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 given in Table X. 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.

| | Dense core (20) | Transit core (162) | Outer core (675) | Small regional ISPs (950) | Customers (8852) |
|---|---|---|---|---|---|
| $\mathbb{D}$ (12) | 6 | 0 | 3 | 2 | 1 |
| $\mathbb{E}$ (10) | 4 | 0 | 3 | 2 | 1 |

TABLE X

HIERARCHY LEVELS OF THE TOP ASES RELATED TO MOST MOVE-FROM AND MOVE-TO EVENTS

## VII. BACKGROUND AND RELATED WORK

To date, numerous efforts have sought to measure BGP characteristics. For example, Huston [2] maintained BGP table statistics that cover BGP table entries, advertised address space, unique ASes, unique AS paths, prefix length distribution, and AS connectivity. The CIDR Report [3] recorded added and withdrawn routes of the past week based on BGP data snapshots and listed the top originating ASes in terms of added and withdrawn routes, respectively. Recently, Meng *et al.* [32] measured the IPv4 address allocation and usage, the advertisement of fragmented and aggregated addresses, and their impact on the global routing table size. Additionally, Bu *et al.* [37] attributed BGP table growth to multi-homing, failure to aggregate, load-balancing and fragmentation. To reduce routing table size and mitigate its growth, RFC 2519 [38] provided insights into source based aggregation, both within routing domains and towards upstream providers. With respect to Internet topologies, numerous efforts [7], [8], [9], [10], [11] have been undertaken to measure ISP-level or router-level topologies. Moreover, Gao introduced a heuristic algorithm for inferring AS relationships from BGP routing tables [4]. Similarly Subramanian *et al.* proposed a five-level classification of ASes based on the connectivity of AS and the relationship of neighboring ASes [5].

Dynamic behaviors of BGP have always been of tremendous interest to the community. Labovitz *et al.* extensively studied BGP routing instability, that is, the rapid fluctuation of network reachability information [12], [13]. Their findings showed that pathological, or redundant updates are major contributing factors for routing instability. Their work has led to improvement in router architectures and protocol implementations. Similarly, the convergence time of BGP was well studied [15], [16], [17], [18], [19]. For the most part, these studies have shown that there can be considerable delay in convergence after faults or routing changes, resulting in intermittent connectivity failure and packet loss.

RFC 1930 [39] advocated that a prefix should belong to only one AS (since at each point there should be exactly one routing policy for network traffic to each destination prefix). However, people have found cases that do not obey this advocation. Huston began recording the number of observed multi-origin prefixes [40]. Zhao *et al.* subsequently studied BGP routes that originate from multiple ASes, called MOAS conflicts, and showed that they are not uncommon [26]. Their results showed that most MOAS conflicts are short-lived (with a typical duration of a few days), and arise due to a myriad of reasons including exchange points addresses, multi-homing without BGP, multi-homing with private AS numbers, and faulty or malicious configurations.

Improving the reliability and security of BGP is increasingly being recognized as one of the most important tasks in securing the Internet today. Mahajan *et al.* [36] studied two types of BGP configuration errors: the accidental insertion of routes into the global routing tables (origin misconfiguration) and the accidental propagation of routes that should be filtered out (export misconfiguration). To their surprise, configuration errors are pervasive, with 200–1200 prefixes experiencing errors every day. Besides misconfigurations, intentional attacks are another concern and may cause even more serious damages considering their malicious nature. It was suggested that security risks may arise from several major BGP vulnerabilities [41]. In particular, messages are not guaranteed secrecy, integrity and freshness, nor does BGP verify an AS's authority to advertise a prefix or validate the announced AS path. These last two issues are typically referred to as the origin authentication and the path verification, respectively.

Addressing security risks associated with some of the aforementioned issues has been of much interests of late. Authentication of the origin ASes of prefixes is one of the primary goals of most of the security designs proposed thus far [21], [22], [23], [24]. In Secure BGP (S-BGP) [21], an address allocation public key infrastructure (PKI) was proposed to support origin authentication. The PKI reflects the ownerships of address space by binding address block(s) to a public key belonging to the organization which owns the address block(s). Similarly, Secure Origin BGP (SoBGP) [22] introduced Authentication Certificates to verify the origin AS and the prefix, but mandated less authority. Recently, Aiello *et al.* [24] formalized the semantics of address advertisement and proofs of origin authentication, and proposed a model of IP address delegation graph in addressing specifically the problem of creation and validation of proofs for origin authentication. To investigate the stability of the delegation hierarchy, the authors compared snapshots of BGP routing tables over different time periods and observed some address space moves, though the nature of these moves were not thoroughly explored.

## VIII. Summary

In this paper, we studied the structure and stability of origin advertisements in inter-domain routing. Using our `q-chart` IP address advertisement visualization tool, we explored the gross structure of IP address advertisements and found that it exhibits considerably consistent structure. In particular, this structure is manifest in, among other ways, predictable address allocation patterns, (address range) regional similarity in advertisement size, and geographically localized growth.

Furthermore, we quantitatively characterized the stability of origin advertisements in inter-domain routing from a variety of perspectives. Our study was based on analysis of BGP traffic from multiple vantage points across a one-year period. We showed that BGP updates can be characterized by a significant amount of repetitive prefix re-additions and subsequent withdrawals. Origin changes, on the other hand, are relatively few and account for less than 2% of the BGP update traffic. More than 90% of the prefixes are consistently originated by the same AS for an entire year. Moreover, most ASes are involved in few, if any, movement events, while a small number of ASes are responsible for most of the advertisement churn. We also showed that a fairly large percentage of new prefixes can be attributed to growth in actively evolving and/or transforming organizations and countries (e.g., China), and that the observed (abnormal) prefix flapping with unreasonably high frequency is likely due to misconfigurations. Lastly, our event characterization showed that a significant numbers of the observed origin changes can be attributed to multi-homed prefixes oscillating between their origins.

We argue that understanding of the structure and stability of BGP origin advertisements is important to a number of inter-domain routing research areas, including but not limited to, BGP dynamics, security, and convergence. We believe this work provides operational insight into the problem constraints and solution principles of existing and future network systems. For example, this empirical study shows that a significant portion of prefixes have high origin stability, which provides supporting evidence for key assumptions in security solutions on origin authentication. Moreover, this first-hand analysis of prefix behaviors and traffic patterns can provide a foundation to better analyze and evaluate the performance, cost, and operational impact of a number of proposed designs and systems [21], [22], [23], [24]. We expect to further explore these issues in future work.

## REFERENCES


[1] Y. Rekhter and T. Li, "A border gateway protocol (BGP-4), RFC 1771," 1995.

[2] G. Huston, *BGP Table Data*, http://bgp.potaroo.net.

[3] T. Bates, P. Smith, and G. Huston, *CIDR report*, http://www.cidr-report.org.

[4] L. Gao, "On inferring autonomous system relationships in the Internet," in *IEEE Global Internet*, 2000.

[5] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, "Characterizing the Internet hierarchy from multiple vantage points," in *IEEE INFOCOM*, 2002.

[6] F. Wang and L. Gao, "Inferring and characterizing Internet routing policies," in *Internet Measurement Conference*, 2003.

[7] H. Burch and B. Cheswick, "Mapping the Internet," *IEEE Computer*, 1999.

[8] R. Govindan and H. Tangmunarunkit, "Heuristics for Internet map discovery," in *INFOCOM*, 2000.

[9] *Skitter*, http://www.caida.org/tools/measurement/skitter/.

[10] P. Barford, A. Bestavros, J. Byers, , and M. Crovella, "On the marginal utility of network topology measurements," in *Internet Measurement Workshop*, 2001.

[11] N. Spring, R. Mahajan, and D. Wetherall, "Measuring ISP topologies with Rocketfuel," in *SIGCOMM*, 2002.

[12] C. Labovitz, R. Malan, and F. Jahanian, "Internet routing instability," *IEEE/ACM Transactions on Networking*, 1998.

[13] C. Labovitz, R. Malan, and F. Jahanian, "Origins of Internet routing instability," in *INFOCOM*, 1999.

[14] A. Feldmann, O. Maennel, Z. Mao, A. Berger, and B. Maggs, "Locating Internet routing instabilities," in *SIGCOMM*, 2004.

[15] A. Shaikh, L. Kalampoukas, R. Dube, and A. Varma, "An experimental study of BGP convergence," in *SIGCOMM*, 2000.

[16] T. G. Griffin and B. J. Premore, "An experimental analysis of BGP convergence time," in *IEEE International Conference on Network Protocols*, 2001.

[17] C. Labovitz, A. Ahuja, R. Wattenhofer, and S. Venkatachary, "The impact of Internet policy and topology on delayed routing convergence," in *INFOCOM*, 2001.

[18] D. Pei, X. Zhao, L. Wang, D. Massey, A. Mankin, S. Wu, and L. Zhang, "BGP convergence through consistency assertions," in *INFOCOM*, 2002.

[19] D. Obradovic, "Model and convergence time of BGP," in *INFOCOM*, 2002.

[20] B. R. Smith and J. J. Garcia-Luna-Aceves, "Securing the border gateway routing protocol," in *Proc. of Global Internet*, 1996.

[21] S. Kent, C. Lynn, and K. Seo, "Secure border gateway protocol (secure-BGP)," *IEEE Journal on Selected Areas in Communications*, 2000.

[22] J. Ng, "Extensions to BGP to support secure origin BGP (draft)," in *Network Working Group*, 2003.

[23] G. Goodell, W. Aiello, T. Griffin, J. OIoannidis, 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.

[24] W. Aiello, J. Ioannidis, and P. McDaniel, "Origin authentication in inter-domain routing," in *ACM Conference on Computer and Communications Security*, 2003.

[25] Y. Hu, A. Perrig, and M. Sirbu, "SPV: Secure path vector routing for securing BGP," in *SIGCOMM*, 2004.

[26] 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.

[27] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. van der Merwe, "The case for separating routing from routers," in *ACM SIGCOMM Workshop on Future Directions in Network Architecture*, 2004.

[28] R. Teixeira and J. Rexford, "A measurement framework for pin-pointing routing changes," in *ACM SIGCOMM Network Troubleshooting Workshop*, 2004.

[29] W. Schroeder, K. Martin, and B. Lorensen, *The Visualization Toolkit*, Kitware Inc, 2004.

[30] *University of Oregon Route View Project*, http://routeviews.org.

[31] ICANN, "The Internet Corporation for Assigned Names and Numbers," 2004, `http://www.icann.org/`.

[32] 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.

[33] *Bogon IPs*, http://www.completewhois.com/bogons/.

[34] *BGP Beacon Info*, http://www.psg.com/ zmao/BGPBeacon.htm.

[35] *Historical List of RIS Routing Beacons*, http://www.ripe.net/projects/ris/docs/beaconlist.html.

[36] R. Mahajan, D. Wetherall, and T. Anderson, "Understanding BGP misconfiguration," in *SIGCOMM*, 2002.

[37] T. Bu, L. Gao, and D. Towsley, "On characterizing BGP routing table growth," in *Computer Networks*, 2004.

[38] E. Chen and J. Stewart, "A framework for inter-domain route aggregation, RFC 2519," 1999.

[39] J. Hawkinson and T. Bates, "Guidelines for creation, selection, and registration of an autonomous system (AS), RFC 1930," 1996.

[40] G. Huston, http://bgp.potaroo.net/as6447/bgp-multi-orgas.txt.

[41] S. Murphy, "BGP security vulnerabilities analysis," in *Internet Research Task Force*, 2002.


## APPENDIX

Table XI shows the intersection of the sets of the top 30 ASes that originate the most new prefixes, observed from six viewpoints, as well as the degrees of connectivity of the ASes and their corresponding degree rankings. Similarly, Table XII shows those that are involved in the most origin changes.

| AS | Organization | Degree | Degree rank |
|---|---|---|---|
| 705 | ALTERNET-AS UUNET Tech., Inc. | 6 | 1664 |
| 209 | ASN-QWEST Qwest | 1296 | 4 |
| 702 | ALTERNET-AS UUNET Tech., Inc. | 768 | 9 |
| 11305 | INTERLAND-NET1 Interland Inc. | 8 | 1243 |
| 6198 | BATI-MIA BellSouth Network Solutions, Inc. | 31 | 280 |
| 4151 | USDA-1 USDA | 1 | 13883 |
| 3908 | SUPERNETASBLK SuperNet, Inc. | 240 | 37 |
| 6197 | BATI-ATL BellSouth Network Solutions, Inc. | 22 | 428 |
| 2150 | CSUNET-SW California State Univ. | 17 | 575 |
| 17964 | DXTNET Beijing Dian-Xin-Tong Net Tech. Co., Ltd. | 10 | 1010 |
| 2686 | AT&T Global Network Services - EMEA | 121 | 70 |
| 7843 | ADELPHIA-AS Adelphia Corp. | 29 | 312 |
| 19029 | NEWEDGENETS New Edge Networks | 205 | 45 |
| 4323 | TW-COMM Time Warner Communications, Inc. | 452 | 15 |
| 9394 | CRNET CHINA RAILWAY Internet(CRNET) | 10 | 1010 |
| 4134 | CHINANET-BACKBONE No.31,Jin-rong Street | 64 | 129 |
| 6147 | PE-TPSA-LACNIC Telefonica del Peru S.A.A. | 1 | 13883 |
| 1239 | SPRINTLINK Sprint | 2279 | 2 |
| 1237 | KREONET-AS-KR Korea Institute of Sci. and Tech. Information | 35 | 244 |
| 9121 | TTNet TTnet Autonomous Sys. | 72 | 119 |
| 4814 | CHINANET-BEIJING-AP China Telecom (Group) | 1 | 13883 |
| 577 | BACOM Bell Advanced Communications Inc. | 112 | 75 |
| 852 | ASN852 Telus Advanced Communications | 138 | 164 |

TABLE XI

THE COMMON ASES AMONG THE TOP 30 ASES, OBSERVED FROM ALL SIX LISTENING POINTS, WHICH ORIGINATE THE MOST NEW PREFIXES IN 2003

| AS | Organization | Degree | Degree rank |
|---|---|---|---|
| 701 | ALTERNET-AS UUNET Technologies, Inc. | 3213 | 1 |
| 1 | GNTY-1 Genuity | 695 | 11 |
| 16422 | NEWSKIES-NETWORKS Newskies Networks, Inc. | 12 | 832 |
| 1239 | SPRINTLINK Sprint | 2279 | 2 |
| 6197 | BATI-ATL BellSouth Network Solutions, Inc. | 22 | 428 |
| 6198 | BATI-MIA BellSouth Network Solutions, Inc. | 31 | 280 |
| 2907 | ERX-SINET-AS National Center for Science Information Systems | 51 | 161 |
| 7018 | ATT-INTERNET4 AT&T WorldNet Services | 1995 | 3 |
| 209 | ASN-QWEST Qwest | 1296 | 4 |
| 7713 | TELKOMNET-AS-AP PT TELEKOMUNIKASI INDONESIA | 8 | 1243 |
| 4795 | INDOSAT2-ID-AP INDOSATNET-ASN | 5 | 2047 |
| 3356 | LEVEL3 Level 3 Communications, LLC | 1164 | 5 |

TABLE XII

THE COMMON ASES AMONG THE TOP 30 ASES, OBSERVED FROM ALL SIX LISTENING POINTS, WHICH ARE INVOLVED IN THE MOST ORIGIN CHANGES IN 2003