Measurement of Highly Active Prefixes in BGP

Ricardo V. Oliveira, Rafit Izhak-Ratzin, Beichuan Zhang, Lixia Zhang

***Abstract*— We conduct a systematic study on the pervasiveness and persistency of one specific phenomenon in the global routing system: a small set of highly active prefixes accounts for a large number of routing updates. Our data analysis shows that this phenomenon is commonly observed from monitors in many different ISPs, and exists throughout our 3-year study period. The analysis further shows that the majority of these prefixes are highly active for only one or a few days, while a small number of them are persistently active over long period of time. Case studies demonstrate that the causes of these high routing activity include topological failures, BGP path exploration, protocol defects, and the failure of turning on protection mechanisms.**

I. Introduction

Previous measurement studies (e.g., [2][8][11]) on the Border Gateway Protocol (BGP) show that there exists a small set of prefixes which contribute a large number of routing updates. However these studies only examined data collected from one or a few ISPs with durations of two months or even shorter. It is unclear whether the small number of highly active prefixes is specific to individual ISPs and/or the limited time periods being studied, or it is a common and persistent phenomenon on the Internet. It is also unknown whether the set of highly active prefixes is stable or changing over time.

In this paper, we conduct a systematic study to answer the above questions by analyzing BGP log data collected at RouteViews [9] over a 3-year period. We define a *highly active (HA)* prefix as one whose number of updates per day exceeds a given threshold. Our findings can be summarized as follows. First, a small number of HA prefixes is observed from *all* RouteViews monitors over the 3-year period. These HA prefixes are a very small percentage (around 0.1% or less) of the routing table, but they contribute a significant portion (10% or above) of the total routing updates. Second, the set of HA prefixes changes over time; some previously stable prefixes become highly active everyday, while some existing HA prefixes become stable. Third, most prefixes become highly active at least once during our 3-year study period, however the high activity is transient for most of them, e.g., 80% of HA prefixes were only active for one day. Fourth, a small number of prefixes stayed highly active for hundreds of days continuously. Finally, through several case studies we identified a number of causes for HA prefixes, including topological failures, BGP path exploration, protocol defects, and failure of turning on protocol protection mechanisms.

The rest of the paper is organized as follows. Section II provides a quantitative definition of HA prefixes. Section III examines the persistency of HA prefixes over our 3-year study period, and Section IV measures the existence of HA prefixes across all the monitors. Section VI includes several case studies as our initial effort in investigating the causes of HA prefixes. Section VII discusses related work and Section VIII concludes the paper.

This work is partially supported by the Defense Advanced Research Projects Agency (DARPA) under Contract No. N66001-04-1-8926 and by National Science Foundation (NSF) under Contract No. ANI-0221453. Any opinions, conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the DARPA or NSF.

Ricardo V. Oliveira, Rafi t Izhak-Ratzin and Lixia Zhang are with Computer Science Department, University of California, Los Angeles. Email: {rveloso, rafi ti, lixia}@cs.ucla.edu. Beichuan Zhang is with Computer Science Department, University of Arizona. Email: bzhang@cs.arizona.edu.

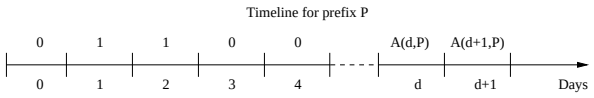

Fig. 1. Slotted time line for prefi x P

II. Methodology

A. Dataset

To assess the pervasiveness of highly active prefixes, we used BGP updates to measure HA prefixes along three dimensions:

- *Time*: how long highly active prefixes have existed;
- *Commonality*: whether HA prefixes are observed only by specific monitors, or commonly across all monitors;
- *Properties of HA prefixes*: How long an HA prefix stays active, whether the set of HA prefixes is stable or changing over time, and etc.

Due to the heavy load of data processing, we used RouteViews BGP log in the following way. To examine how long the HA phenomenon has existed, we used 3-year's of data, from October 2001 (i.e., when RouteViews started archiving BGP updates) to August 2004, but limited to 4 monitors: 129.250.0.11 (AS 2914), 144.228.241.81 (AS1239), 199.74.221.1 (AS812), and 204.42.253.253 (AS267). These 4 monitors were chosen because they were connected to RouteViews since October 2001 and they represent ISPs at different tiers in the Internet hierarchy. To assess the commonality of HA prefixes, we used data from *all* monitors, but limited to two randomly selected months, March 2002 and May 2004.

B. Classification of HA Prefixes

To identify BGP prefixes that are associated with a high level of activity, we propose a classification method based on the number of BGP updates associated with a given prefix during 1 day period. The choice of using the *day* as the interval is an engineering decision based on the assumption that most network problems occur or get resolved on a daily basis (e.g., fiber cuts, worm attacks [11]). We divided our 3-year study period into slots of 1 day as illustrated in Figure 1. Let $N_u(d, P)$ be the number of updates in day d for prefix P, the activity function $A(d, P)$ is defined as:

$$A(d,P) = \begin{cases} 0 & : \quad N_u(d,P) < T_u \\ 1 & : \quad N_u(d,P) \geq T_u \end{cases} \tag{1}$$

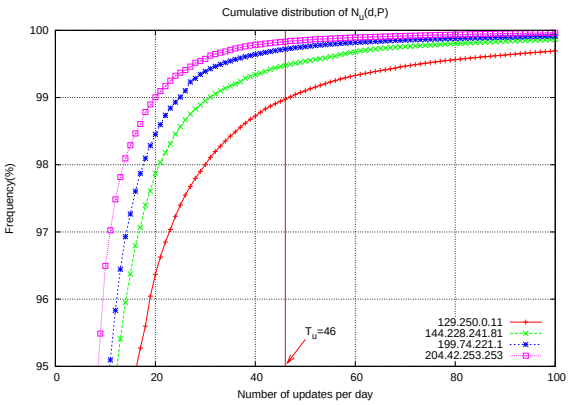

Fig. 2. Cumulative Distribution of $N_u(d, P)$ for Four Monitors

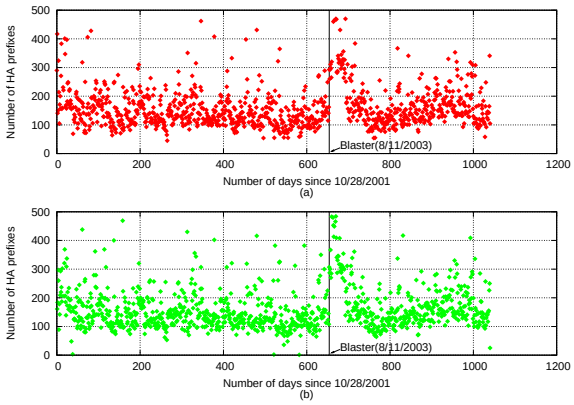

Fig. 3. (a) Number of HA prefi xes per day, from monitor 144.228.241.81 and (b) with offset of 12h in day slot.

Monitor	Table size (Oct 2001)	Table size (Aug 2004)	Growth (prefi x/day)
129.250.0.11	89406	142985	52
144.228.241.81	103510	140494	36
199.74.221.1	103789	144808	39
204.42.253.253	89395	117188	27

TABLE I
EVOLUTION OF ROUTING TABLE FROM OCTOBER 2001 TO AUGUST 2004.

A prefix P is *highly active* in day d if $A(d, P) = 1$. In section V, we will define *life time* of HA prefixes to examine how long the high activity lasts. The threshold T_u is an important parameter of $A(d, P)$ and must be chosen so that it captures the prefixes that have a "high" number of updates per day. To do so, We plotted the cumulative distribution of $N_u(d, P)$ of 4 monitors over 3 years in Figure 2. We are interested in capturing prefixes in the tail of the curves (high number of updates). Since the classification criteria of HA prefixes should be independent from any specific monitor, we need pick a single threshold T_u that can accommodate all the distribution curves. Figure 2 shows that setting $T_u = 46$ can catch the top %1 (or smaller percentage) of the curves. Actually the vertical line at 46 updates/day intersects the curves at 99.82%, 99.7%, 99.5% and 99% respectively. In next section we show that using slightly larger or smaller T_u values doesn't affect our result qualitatively. Therefore, in the rest of the paper, we will use $T_u = 46$ in determining HA prefixes.

III. PREFIX ACTIVITY OVER TIME

In this section we examine whether the existence of HA prefixes is a persistent phenomenon over time.

A. Persistent Activity

Figure 3(a) shows the number of HA prefixes over time observed from one monitor. The average is around 150 HA prefixes per day, roughly within the range of 100 to 200. During the same time period, the routing table size from this monitor increased from around 100k to 140k (Table I). This shows that (1) only a small percentage of prefixes ($\sim$0.1%) are highly active each day, and (2) the number of HA prefixes per day maintains relatively constant despite the growth of 36% in the routing table size. Data from all other three monitors show similar pattern as in Figure 3. In Figure 3 we also highlight the date of the Blaster worm attack[12] on August 11, 2003. The effects of the Blaster worm were quite visible and lasted for several days.

In determining HA prefixes, we count number of updates per *day*. Thus what time we take as the beginning of a day might have an effect on the result. For example, there could be a prefix that has more than T_u updates in a continuous 24-hour span, but these updates fall in two different days, and each day has less than T_u updates. To check whether these cases may skew our results, we shifted the beginning of the day by 12 hours, and plotted the results in Figure 3(b). We can see that there is no significant difference between Figures 3(a) and 3(b), indicating that our observations are not sensitive to the beginning time of the day.

There is a valid concern about how much our observations would differ had we chosen a slightly different threshold (T_u) in determining HA prefixes. In Figure 4 we re-plotted Figure 3(a) with $T_u = 46 \pm 10\%$. To make the curves legible, we used a weighted average $\bar{y}_n = \alpha \cdot \bar{y}_{n-1} + (1 - \alpha) \cdot y_n$ with $\alpha = 0.8$ to smooth the curves, and only plotted one day per week. Two observations are in order: (1) the shape of the curves are the same for the different thresholds and (2) the absolute values of each curve are very close to each other. This indicates that our observations are not sensitive to small changes to T_u. This was expected, since the CDF curve is already flatten around T_u in Figure 2.

B. BGP Updates Caused by HA Prefixes

Figure 5(a) shows the total number of updates per day and 5(b) shows the percentage of updates from HA prefixes in each day. Despite that the total number of prefixes in the routing table increased about 36% over the 3 years for monitor 144.228.241.81 (Table I), the average number of updates remained around 100,000 per day. The fraction of updates associated with HA prefixes have an average trend around 10% and ranges roughly between 0% and 40%. Although HA prefixes in a single day are only 0.1% of the routing table, they contribute to 10% of the updates. We did not observe a strong correlation between Figures 5(a) and 5(b), however,

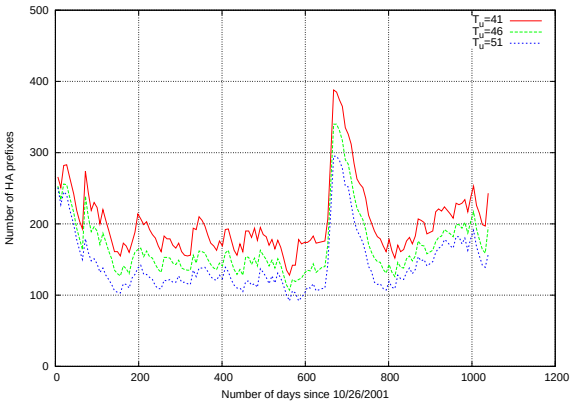

Fig. 4. Number of HA prefixes per day, from monitor 144.228.241.81 (Sprint), with different values for T_u.

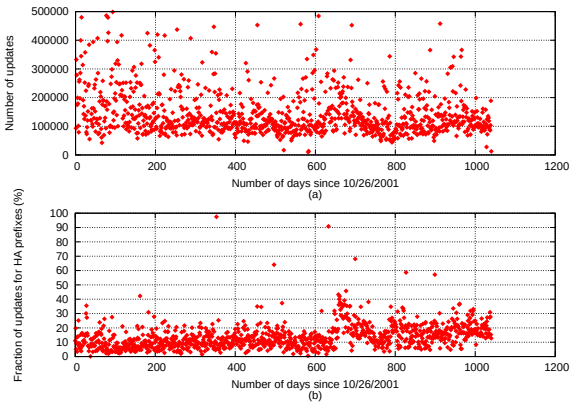

Fig. 5. (a) Number of updates per day, (b) Fraction of updates for HA prefixes from monitor 144.228.241.81.

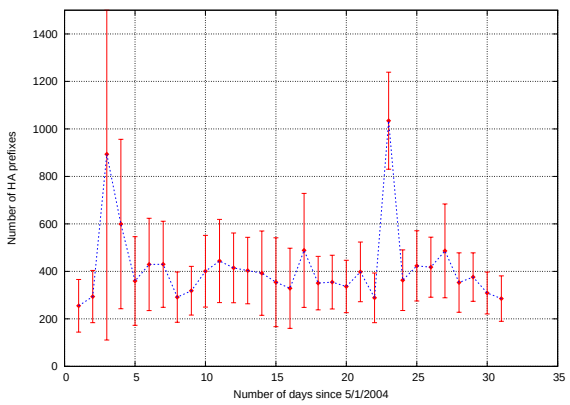

Fig. 6. Number of HA prefixes per day for 33 peers, May 2004.

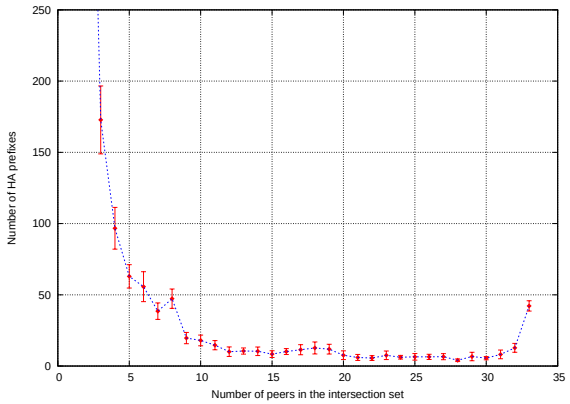

Fig. 7. HA prefix intersection over 33 different peers, May 2004

between day 600 and day 800 there is a small increment in the number of total updates, up to 40% of them were from HA prefixes.

IV. Prefix Activity Across Different Monitors

In the previous section we explored prefix activity over time by analyzing 3-year's updates from 4 monitors. To understand whether existence of HA prefixes is a common phenomenon in the Internet, in this section we look at all the monitors of RouteViews Oregon collector in 2 randomly chosen months, March 2001 and May 2004. Since the results are similar from both months, we only present the results of May 2004 here.

Figure 6 shows the average number (with confidence interval 90%) of HA prefixes observed by 33 monitors per day. Though each monitor observes relatively similar *number* of HA prefixes every day, different monitors may observe different set of HA prefixes. Figure 7 shows the intersection of HA prefixes viewed from different monitors. Each data point is the average over the month of May 2004 with 95% confidence interval. A point at (x, y) means there are y number of prefixes that are observed as highly active by x number of monitors. This figure shows that most of the HA prefixes are observed by only a small number of monitors. There is a decreasing trend on the common set of HA prefixes as the number of monitors increases. The high activity observed only by one or a small number of monitors are likely caused by network problems away from the AS that originates the prefix, and only the monitors that share the problematic path are affected. Looking closer at Figure 7, we notice that it has a "lifted tail" at 33 monitors. This means that there is a set of HA prefixes that are seen by all monitors, suggesting that the root cause of the activities is likely to be near the AS that originates the prefixes, thus it affects all monitors in the similar way. We also looked at these HA prefixes that are common to all 33 monitors and found that 38 prefixes were active for at least one month, and that most of them were originated by the same AS.

V. HA Prefix Properties

A. HA Prefix Set

We now study how the set of HA prefixes changes over time. Figure 8 plots the number of new HA prefixes appearing every day as observed from monitor 144.228.241.81. An HA prefix is "new" if it has never been highly active before that day. The average trend is around 25 new HA prefixes per day, and in most days it's within the range of 0 to 50. Combined with Figure 3, which shows that the total number of HA prefixes in each day is relatively stable, we conclude that the set of HA prefixes is not fixed, but changes over time. It is a dynamic set in the sense that every day there are new prefixes that become highly active and some previous HA prefixes stabilize. We

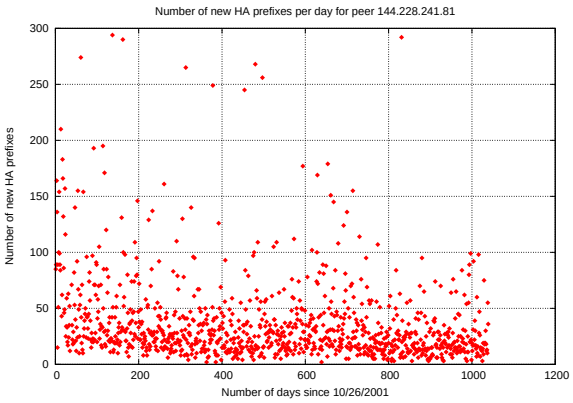

Fig. 8. New HA prefixes over time for monitor 144.228.241.81.

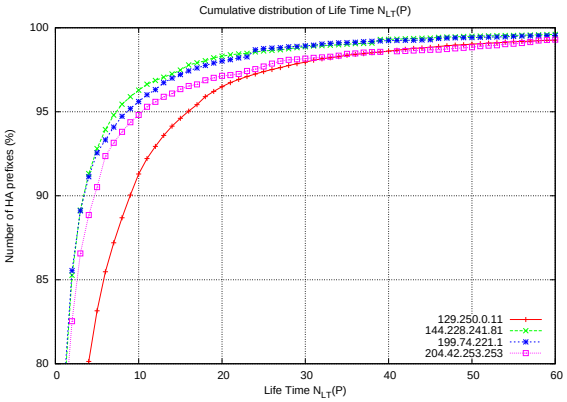

Fig. 9. Cumulative distribution of the life time $N_{LT}(P)$

believe that various topological events occurring every day in the network affect different sets of prefixes and generate new HA prefixes. Furthermore, we discovered that HA prefixes appear in groups that share the same origin AS, suggesting that routing problems in a single AS can trigger a large number of updates spanning over a range of prefixes, making them active at the same time.

B. Life Time

Given a prefix P, we define its *life time*, $N_{LT}(P)$, as the total number of days in which P is active, i.e., $N_{LT}(P) = \sum_{d=0}^{D-1} A(d, P)$, where D is the total number of days in our data set, which is 1040 days. Figure 9 shows the cumulative distribution of N_{LT} of the HA prefixes observed from the 4 monitors. More than 90% of the HA prefixes have $N_{LT}(P) < 9$, which means that most high activities are transient, lasting only a few days. In fact, more than 80% of HA prefixes we observed had a life time of only one day. This indicates that most prefixes become highly active due to localized, transient events that occur in the network. On the other hand, there is also a very small number of HA prefixes that have very long life times. Although not shown in Figure 9, we have observed cases of $N_{LT}(P) > 700$ days.

VI. Case Study

In this section we try to understand the causes of high activity through some case studies.

A. Transient HA Prefixes

We describe here a case where network outages cause a large number of HA prefixes in a single day. On April 13, 2004, one of the core routers of Internet2 [1] experienced several outages in a short time period. This router had direct connections to some of RouteViews monitors. With the help of LinkRank tool [5], we found that at the time of the first outage, one monitor switched paths for ∼1.5k prefixes, resulting BGP updates for these prefixes. Due to multiple outages on the Internet2 router, the monitor switched path back and forth several times during a 10 hour period. The BGP updates caused by these path changes made ∼1.5k prefixes appear highly active in that day. In this example, a local, transient event triggered a number of updates sufficient to make a group of prefixes highly active. We believe it reflects most of the HA cases, as more than 80% of HA prefixes have a lifetime of only one day (Figure 9).

B. Persistent HA Prefixes

From the 3 years of data, we noticed that a prefix had a life time of more than 500 days. Its high activity is observed by all the 33 monitors, with a life time between 524 and 542 days. A closer look at the BGP updates revealed that the routing path to reach this prefix alternated between two paths, one had the ATOMIC_AGGREGATE attribute set (meaning the prefix was being aggregated), and the other had no aggregation attribute. We found that the origin AS was triggering these updates in a time interval of several minutes. The rate of updates was not high enough to trigger BGP's route flap damping mechanism ([10]), but still caused lots of updates at the monitors. We believe this problem was caused by IGP misconfiguration inside the ISP that originated the prefix.

C. Path Exploration

As a path vector protocol, BGP undergoes path exploration after routing changes. During the convergence period, a router may send multiple updates before eventually settling down on the new stable paths, and this increases the number of updates propagated in the network. The best example to illustrate this is BGP beacon prefixes. BGP beacon [7] is an active measurement technique which periodically announce and withdraw its prefixes. Most beacons have 2 hours between the announcement and the withdrawal. Without path exploration, we would expect about 12 updates per day for each beacon prefix. However, we found that some beacons were included in the HA set. The numbers are shown in Table II, ordered by life time. The last column is the average number of updates per day, considering only the days during which the prefix was active.

In order to confirm that path exploration is the reason for the large number of updates, we picked beacon prefix 195.80.227.0/24, and grouped its updates between January 2004 and May 2004 into events as described in [4]. We found that during the events that were triggered by beacon announcements and withdrawals, the average number of distinct paths is 3.1 per event, and the average event duration is 76.9 seconds.

Beacon prefix	Life Time (days)	Origin	Updates per day (average)
195.80.227.0/24	60	Amsterdam, RIPE	59
195.80.225.0/24	59	London, RIPE	51
195.80.226.0/24	33	Paris, RIPE	49
195.80.229.0/24	26	Vienna, RIPE	47
195.80.236.0/24	12	Frankfurt, RIPE	47
192.135.183.0/24	7	AS5637, PSG	56
195.80.235.0/24	5	New York, RIPE	50

TABLE II
ACTIVITY OF BGP BEACON PREFIXES AS OBSERVED BY MONITOR 144.228.241.81.

And we observed individual cases that the router explored several paths before converging to the new path or withdrawing the path. This shows that path exploration contributes to the increase of routing updates.

D. MRAI Timer and Route Flap Damping

To reduce the routing instability caused by update surges, BGP has two built-in mechanisms, Minimum Route Announcement Interval (MRAI) and Route Flap Damping. A router does not send two consecutive announcements unless they are spaced out by MRAI, default at 30 seconds. With damping, a router will keep track of route instability by maintaining a penalty value. If the penalty value exceeds a threshold, the router will suppress the route, i.e., stops propagating updates for this route. It will be reused only after the route has stabilized.

MRAI and damping are implemented in all modern BGP routers, but not all vendors turn them on by default. As a result, some local instability can be propagated to remote networks unnecessarily, causing a high number of BGP updates to be injected in the network. In an extreme case, we observed that a /24 prefix was highly active for 12 consecutive days with 6011 updates/day (i.e., 11 seconds/update) on average. On the peek day November 6, 2003, one monitor observed 12k updates, one observed 8k updates, four monitors observed 49 to 59 updates, and the other monitors did not capture this prefix as HA. This example shows that MRAI and damping are not universally deployed on the Internet. Routers that are not configured with these protection mechanisms are likely to have a higher processing load (possibly overload) caused by the numerous BGP updates they originate. During stressful events such as worm attacks, the lack of MRAI and damping could reduce the system stability significantly [11][6].

VII. RELATED WORK

In [2], the authors used two months of RouteViews routing tables in 2001 to study the presence of prefixes in the global routing table. They found that half of the prefix withdrawal/re-announcement were contributed by 1.2% of all ASes. In [8], the authors used one month of routing updates from RouteViews, RIPE, and an AT&T router in 2002 and found that the majority of routing events were about a small number of prefixes. Other studies such as [3][11][6] reported that during worm attacks, BGP undergoes unusual activities that result in a large amount of routing updates being injected in the Internet. These works reveal the existence of HA prefix phenomemon at certain times. Our paper systematically studies the pervasiveness and persistency of the HA prefix phenomenon over three years, and from all RouteViews monitors.

VIII. CONCLUSION

The study of prefix activity in BGP provides an insight about the stability of routes in the Internet. In this paper we presented a simple method to classify prefixes as *highly active* (HA), based on the number of BGP updates per prefix measured in 1-day intervals. Using this method, we were able to detect hot spots of instability in the global routing system along three dimensions: time, location and prefix.

We showed that the set of HA prefixes, though only 0.1% of all global prefixes, is responsible for approximately 10% of BGP updates injected into the network every day. This behavior is persistent over time (3 years) and over location (33 monitors). The set of HA prefixes changes over time: every day there are some new prefixes become highly active and some previously active prefixes become stable. This indicates that network events are not localized in a restricted domain of ASes, but rather span all the network, affecting all IP address space over time. We find that more than 80% of HA prefixes are highly active for only one day in 3 years. This leads us to conclude that high prefix activity is, in most cases, an effect of sporadic network events. However there is a small set of prefixes that experienced long lasting activity caused by persistent network instabilities.

REFERENCES

[1] Abilene Internet2 Network. http://abilene.internet2.edu.

[2] A. Broido, E. Nemeth, and kc claffy. Internet expansion, refi nement, and churn. *European Transactions on Telecommunications*, January 2002.

[3] J. Cowie, A. Ogielski, B. J. Premore, and Y. Yuan. Global routing instabilities triggered by code red ii and nimda worm attacks. Technical report, Renesys Corporation, December 2001.

[4] Rafi t Izhak-Ratzin, Beichuan Zhang, Dan Pei, Daniel Massey, and Lixia Zhang. Quantifying Path Exploration in the Internet. under submission.

[5] M. Lad, D. Massey, and L. Zhang. A graphical tool for capturing bgp routing dynamics. In *Network Operations and Management Symposium (NOMS)*, "April" 2004.

[6] Mohit Lad, Xiaoliang Zhao, Beichuan Zhang, Dan Massey, and Lixia Zhang. An analysis of bgp update burst during slammer attack. In *Proc. of the 5th Int'l Workshop on Distributed Computing*, 2003.

[7] Z. Morley Mao, Randy Bush, Tim Griffi n, and Matt Roughan. BGP beacons. In *Internet Measurement Conference (IMC)*, 2003.

[8] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang. Bgp routing stability of popular destinations. In *Internet Measurement Workshop (IMW)*, 2002.

[9] The Route Views Project. http://www.antc.uoregon.edu/route-views/.

[10] C. Villamizar, R. Chandra, and R. Govindan. Bgp route flap dampening. RFC 2439, IETF, November 1998.

[11] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. F. Wu, and L. Zhang. Observation and analysis of BGP behavior under stress. In *Internet Measurement Workshop (IMW)*, 2002.

[12] Nicholas Weaver, Vern Paxson, Stuart Staniford, and Robert Cunningham. A taxonomy of computer worms. In *WORM'03: Proceedings of the 2003 ACM workshop on Rapid Malcode*, pages 11–18, 2003.