Public Review for

BGP Routing Dynamics Revisited

Jun Li, Michael Guidero, Zhen Wu, Eric Purpus,
and Toby Ehrenkranz

BGP is indeed a critical protocol of the Internet. Understanding its behaviour, robustness and dynamics is therefore of utmost importance. This paper proposes an update of our knowledge in this area with respect to the comprehensive work by Labowitz et al (Sigcomm'97). The authors use recent BGP data from RIPE and pay particular attention to the changes in Internet routing dynamics since Labovitz work. Deriving observations from measurements is an important work if properly done.

The paper is concerned by producing an analysis as in Labovitz, using a similar methodology and taxonomy in order to be able to benchmark the current BGP with respect to the observations carried out 10 years ago. This is a general concern in any area to be able to derive observations of a system over time, develop its analysis and build an understanding of the causes that produce the observations.

The main founding of this paper is that, a decade later, BGP dynamics are now “busier” but “healthier.” Others informative results are derived and provide an interesting update on the subject though many other studies of BGP were published since 97.

Nevertheless, comparing dynamic systems over a ten year period is not obvious as the original classification needs to be deepened and extended: to extract a better knowledge about the system dynamics, and to observe new behaviours.

This paper has therefore been considered as valuable for the networking community. It presents a snapshot of the current BGP dynamics, discusses its evolution with respect to the work by Labovitz and, indirectly, questions the issue of the methodology to observe the same dynamic system over time and understand the causes behind the dynamics.

Public review written by

Serge Fdida
University Pierre & Marie Curie, LIP6-CNRS Lab

BGP Routing Dynamics Revisited*

Jun Li, Michael Guidero, Zhen Wu, Eric Purpus, and Toby Ehrenkranz
Dept. of Computer and Information Science, University of Oregon
Eugene, OR, USA
lijun@cs.uoregon.edu, guiderom@cs.uoregon.edu, zwu@cs.uoregon.edu
epurpus@cs.uoregon.edu, tehrenkr@cs.uoregon.edu

ABSTRACT

Understanding BGP routing dynamics is critical to the solid growth and maintenance of the Internet routing infrastructure. However, while the most extensive study on BGP dynamics is nearly a decade old, many factors that could affect BGP dynamics have changed considerably. We revisit this important topic in this paper, focusing on not only comparing with the previous results, but also issues not well explored before. We have found that, compared to almost a decade ago, although certain characteristics remain unchanged (such as some temporal properties), BGP dynamics are now "busier," and more importantly, now have much less pathological behavior and are "healthier"; for example, forwarding dynamics are now not only dominant, but also more consistent across different days. Contributions to BGP dynamics by different BGP peers—which are not proportional to the size of a peer's AS—are also more stable, and dynamics due to policy changes or duplicate announcements are usually from specific peers.

Categories and Subject Descriptors

C.2.2 [**Computer-Communication Networks**]: Network Protocols—*Routing Protocols*

General Terms

Measurement, Analysis

Keywords

BGP, routing dynamics, routing instability, network measurement

1. INTRODUCTION

Internet routing is dynamic in nature. Caused by the regular or irregular exchange of routing updates between routers, routing dynamics has always been a major concern of the Internet engineering community. Irregular dynamics can not only cause high bandwidth and processing overhead on routers, but may also lead to packet forwarding failures, including packet delay, jitter, drop, reordering, duplication, or other difficulties in reaching destinations. With the Internet being indispensable to modern communications and the economy, it is critical to understand the characteristics of routing dynamics.

The most comprehensive study [1] of Internet routing dynamics is from nearly a decade ago, a substantial period for the fast-evolving Internet. Since then many factors that could affect routing dynamics have changed, and it is important to thoroughly re-examine this topic. These factors include routing protocol implementation by vendors, network engineering practices, and Internet topologies.

The Border Gateway Protocol (BGP) is the *de facto* standard inter-domain routing protocol on the Internet. Not only has the BGP specification been evolving [2] since it became an IETF standard in 1995 [3], but the specification also provides a great deal of flexibility. Individual vendors can add new features such as communities [4] and route flap damping [5] in response to the problems encountered by Internet operators.

Engineering practices also constantly brought changes. Multi-homing, load balancing, and address fragmentation, have led to substantial increases in BGP routing table sizes [6]. Route flap damping has been continuing to change over time; for example, operators have put more lenient thresholds than those recommended in [5] after noticing that the recommended thresholds heavily penalized flapping paths.

The Internet itself has been growing at a staggering rate. The number of ASes has been increasing at a linear rate, and both the number of prefixes and the size of BGP routing tables have been growing almost exponentially [7, 8], resulting in more complicated traffic engineering and routing policies. Moreover, while BGP updates are carried over the same link as the ordinary traffic, the usage patterns of the Internet change over time and new applications and traffic patterns could have an impact on BGP as well.

Essentially, the constant growth and changing features of the Internet create the necessity of revisiting the topic of BGP routing dynamics in order to capture new statistics and trends. In this paper, we investigate characteristics of the BGP routing dynamics on the Internet using recent BGP data from August 2005 to January 2006. Our primary focus in this work is to observe and understand how BGP dynamics look nowadays and how it has changed since the previous work in [1]. We will also provide our insights when appropriate on what may have caused various changes.

Table 1 shows the comparison between the major results from [1] and ours. Moreover, our research results also do not agree with the early expectation on the trend of routing dynamics. It was expected that with the rapid addition of multi-homed domains and new peering relationships (thus an increasing number of globally visible address prefixes as well as globally visible routes), we would experience

---

*This material is based upon work supported by the National Science Foundation under Grant No. 0520326.

Table 1: Summarized BGP dynamics comparison between now and almost a decade ago

Major results from [1]	Our major results
"The number of BGP updates exchanged per day in the Internet core is one or more orders of magnitude larger than expected."	The number of updates exchanged is now basically as expected (Section 4).
"Routing information is dominated by pathological, or redundant updates, which may not reflect changes in routing policy or topology."	BGP is now "healthier" as routing information is dominated by forwarding dynamics and some policy changes; pathological updates are now about 16%, with most being redundant announcements (Section 4).
"Instability and redundant updates exhibit a specific periodicity of 30 and 60 seconds."	Certain predominant inter-arrival times still exist, but different dynamics types demonstrate different patterns of inter-arrival times (Section 7).
"Instability and redundant updates show a surprising correlation to network usage and exhibit corresponding daily and weekly cyclic trends."	These trends remain (Section 7).
"Instability is not dominated by a small set of ASes or routes."	This remains to be true (Sections 5 and 6), while BGP dynamics are more stable in terms of contributions by different BGP peers and more consistent from day to day and across different dynamics types.
"Instability and redundant updates exhibit both strong high and low frequency components. Much of the high frequency instability is pathological."	There are now much fewer pathological updates (Section 4). Predominant frequency still exists while low frequency component is stronger (Section 7).
"Discounting policy fluctuation and pathological behavior, there remains a significant level of Internet forwarding instability."	Even without discounting policy fluctuation and pathological duplicates, the amount of forwarding instability is more significant (Section 4).

more routing dynamics [1]. However, although it is true that the Internet has become more complex, we found that BGP routing dynamics are also becoming "healthier," and there are now a reasonable amount of dynamics while BGP accomplishes its job.

Authors in [1] also pointed out that one of the dynamics types, AADup, needs a closer analysis as a future topic in order to understand the effect of policy fluctuation on routing dynamics. Toward this end, our research provides a refined taxonomy of dynamics types; as we will describe later, this new taxonomy is key to capturing policy fluctuation as well as identifying the new dominant pathological duplicates.

The rest of this paper is organized as follows. Section 2 provides an overview of work related to BGP dynamics. In Section 3, we outline our methodology in observing and studying BGP dynamics, comparing that from almost a decade ago. Our findings and analysis are described in Sections 4 to 7. We conclude our paper in Section 8.

2. RELATED WORK

The related work for this study includes a number of BGP and Internet routing measurement studies, including the original study by Labovitz *et al.* [1]. Many of them were conducted several years or even more than a decade ago. Given the changing nature of the Internet even over a relatively short period, these studies provide both a foundation and a motivation for our work.

Using archived BGP updates from route servers at major US network exchange points, Labovitz *et al.* found several unexpected trends in routing dynamics [1]. In particular, the number of BGP updates per day was several orders of magnitude greater than expected, with pathological duplicates accounting for the majority of unstable dynamics. BGP updates also exhibited daily and weekly patterns that corresponded to network usage patterns, indicating a correlation between routing instability and network usage. Moreover, they carefully measured BGP dynamics at a finer grain to study the BGP dynamics contributions from different ASes or BGP dynamics on a per-route basis.

In their follow-up work [9], Labovitz *et al.* further investigated the origins of BGP dynamics. They confirmed that the high volume of duplicate withdrawals was largely due to the stateless BGP implementation from a particular vendor. Whereas the exact determination is difficult, they also posited possible causes for duplicate or different BGP announcements for the same prefix, as well as the reasons for the periodicity of BGP dynamics. Finally, whereas the most extensive study on BGP dynamics is still from [1], certain aspects of BGP dynamics were also reported in [9].

Other representative studies of routing dynamics, mostly old, are as follows. Chinoy [10] found that most routing fluctuations involve only a small number of networks and that they exhibit periodic patterns on the order of 10 minutes. This study focused on NSFNet in 1993, however. Govindan *et al.* [11] studied the routing stability based on BGP updates between 1994 and 1995, and found that the stability had degraded as the Internet grew. Their work was in the context of prefix availability and prefix steadiness. Research in [12] studied the BGP routing stability of popular destinations, and also focused on prefix reachability instead of aggregated routing dynamics. Through extensive traceroute measurements, Paxon [13] studied routing pathologies and stability by focusing on the end-to-end behavior instead of routing updates. He found that the observed routing pathologies (loops, erroneous routing, transient failures, etc.) doubled between 1994 and 1995.

Instead of measuring and characterizing routing dynamics, some works have focused on identifying the origin of unstable dynamics. In addition to the work by Labovitz *et al.* in [9], a representative study in this direction is research in [14], where a BGP health inferencing system is designed to find what the cause of a routing change is and where a routing change originates. Research in [15] introduced “BGP Beacons” to observe BGP behavior under artificially scheduled path announcements and withdrawals. Feldmann *et al.* [16] correlated the instability from three dimensions—time, prefix, and point of observation—in order to pinpoint which AS or which BGP session is responsible for specific unstable dynamics. However, none of these studies attempted to measure and characterize the routing dynamics or study their prevalence, which is the goal of this paper.

In a more general context, a number of studies have attempted to characterize various specific aspects of BGP dynamics. Research from [17, 18, 19, 20, 21] found that, due to various reasons such as routing policies, topologies, protocol timers, and route flap damping mechanisms, the BGP convergence process is much slower than previously thought. Many recent works ([22, 23, 24, 25]) also investigated the effects of Internet worms, electricity outage, and other significant events on BGP dynamics. These works discovered that the Internet experienced a much higher level of dynamics under severe conditions.

3. METHODOLOGY

3.1 Overview

We assess BGP dynamics by collecting and observing recent BGP updates and further classifying them into different dynamics types. In particular, we identify dynamics that result from BGP doing its job as a routing protocol, and those that result from problems with BGP. The principle of our methodology is to (1) use basically the same procedure as [1] so that we can compare the new results with those from nearly a decade ago in [1], but (2) further introduce new procedures, such as refining the original classification, in order to more closely study new phenomena.

Since our results need to be comparable with the seminal work a decade ago, introducing a completely new methodology may not be wise. We believe that the work by Labovitz *et al.* should be repeated at least in several major aspects as the Internet grows so that we can evaluate BGP over time.

Note that as we study the BGP routing system, we prefer the term *dynamics* to the term *instability.* While *instability* was suitable in the earlier work when most BGP updates were caused by negative events, as we will report in Section 4, BGP is now “healthier” and sees much fewer pathological updates. Using the neutral term *dynamics* can avoid the negative connotation of *instability.*

3.2 Dynamics Types and Classification

We classify routing dynamics basically in the same manner as the *instability* types defined in [1]. A major difference in our classification from that in [1] is that, we refine the original AADup type into two sub-types: AADupType1 and AADupType2. Labovitz *et al.* also modified their classification in their follow-up work [9], but because we specifically target the very original study (which had the most extensive and systematic study on BGP dynamics) and the follow-up work has a different focus (which mainly aims to identify the origins of various BGP dynamics), we inherit our classification from [1] instead of [9].

AADiff: After a route to a destination is announced as available, an alternative route is announced. This is an implicit withdrawal. This occurs when a previously preferred route is not available or a newly preferred route is.

WADiff: After a route to a destination is explicitly withdrawn, a different route (with different ASPATH or NEXTHOP attributes) to the same destination is announced as available.

WADup: After a route to a destination is explicitly withdrawn, the same route is re-announced as available. This could be caused by transient topological failure, or a pathological oscillation.

WWDup: After a route is explicitly withdrawn, it is explicitly withdrawn again. This is a pathological behavior.

AADupType1: After a route to a destination is announced as available, the same route is announced again, with *all* the attributes of the route unchanged. This can indicate a pathological behavior since BGP is an *incremental* protocol and repeated announcements are not expected.

AADupType2: Similar to AADupType1, except that while the ASPATH and NEXTHOP are duplicated, one or more of the other attributes (such as MED or community) is different. This can reflect a routing policy change. Policies sometimes are dynamically changing such as based on some intra-AS routing protocol characteristics.

These dynamics types are grouped into three classes: **forwarding dynamics** that reflect topological changes and affect the forwarding paths, **policy fluctuations** that reflect changes in routing policy attributes but may not affect forwarding paths, and **pathological duplicates** which are redundant BGP updates that reflect neither topology nor policy changes. Both AADiff and WADiff reflect possible exogenous network events, such as router failures or link disconnectivity, and they are forwarding dynamics. WWDup along with AADupType1 are pathological duplicates. AADupType2 reflects policy fluctuations. WADup may be a forwarding dynamics type or pathological duplicate.

We view the BGP routing system as being **healthy** if the BGP routing is filled with mostly forwarding dynamics and policy fluctuations and no, or at most a small amount of, pathological duplicates. For example, with mostly duplicate withdrawals, BGP was not healthy almost a decade ago. The healthiness can also be viewed from finer-grained characteristics, for example, if BGP dynamics are stable and consistent in terms of contributions to dynamics by different peers over different days (Sections 5 and 6). Finally, in our study, we also define BGP to be **busy** if there are many instances of forwarding dynamics (as we will describe in Section 6).

The separation of AADupType1 and AADupType2 allows us to distinguish the dynamics likely caused by policy fluctuation and that caused by pathological behavior, rather than a combined AADup dynamics type that mixes the two. In fact, research in [1] counted AADup as a pathological instability but discovered that AADup at that time was actually dominated by policy changes, and stated that a closer look into AADup is necessary. We address that deficiency in our study. Furthermore, as we will report, this separation leads to two important discoveries: the pathological duplicates nowadays are not dominated by WWDup as in [1], but instead AADupType1; and AADup nowadays are not dominated by policy changes, but equally divided into AADupType1 and AADupType2.

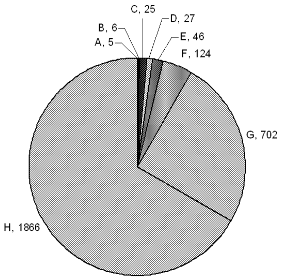

Figure 1: The AS degrees of a RIPE monitor's eight peers

3.3 Data Source

In choosing data sources for studying BGP dynamics, both our study and that in [1] adopt the same philosophy that the BGP data used should provide a good sampling of BGP dynamics. The data source in [1] are route servers at major network exchange points, with representative data from the largest exchange (Mae-East exchange point in USA). Our data sources come from RIPE monitors that many BGP studies have used; in particular, we choose the *RRC00 monitor* [26], which archives BGP updates received from 12 peers. Each BGP update is time-stamped with one-second precision as the monitor receives it. During the six-month period from August 2005 to January 2006 that we conduct our measurement, there are 253,179,846 updates, and for the month of January 2006 that we conducted more analysis, there are 44,384,297 updates in total.

In the course of the investigation, we discovered that only eight of the 12 peers of the RIPE monitor transmitted announcements and withdrawals during the month of January, 2006. Thus, these are the peers used for the detailed portion of this investigation. These eight peers cover a wide range of sizes. Figure 1 shows the relative size of each of the eight peers in terms of *AS degree*, which is the number of adjacent ASes in the network topology. We infer the network topology from our observation period and from table dumps archived by the RRC00 monitor in January, 2006.

Our selection of RRC00 as the data collection point is analogous to the choice of Labovitz *et al.* in concentrating on data from a single exchange point [1, 9]. The eight active peers of the RRC00 monitor cover providers from all over the world: two from the Asia-Pacific region, two from the U.S., one from Africa, and three from Europe. Nonetheless, one might still be concerned whether or not this data source has a bias towards a particular region (such as Europe). In our opinion, given the recommendations of the IETF, router vendors, and years of engineering practice, there is probably not a significant regional or cultural difference between the operational procedures of a sampling of providers in one region versus a sampling of providers in another.

In a word, RIPE provides a good data source for our study. As we described above, the size of the dataset is large and the coverage is reasonable. And we also found the RIPE data easy to retrieve and process. There are also other BGP data archives, such as RouteViews [27], that have similar properties, but for this study the RIPE data proved to be easier and more reliable in filtering out noise that was caused by BGP session resets between a monitor and its peers. We borrowed the algorithm described in [28] to filter out table dumps resulting from BGP session resets.

3.4 Result Presentation In Light Of the Previous Research

We do not present our results exactly the same way as the prior work in [1], for the following reasons:

- While whenever appropriate we do draw direct comparisons between our results and those of the prior work, our emphasis is to learn how every fundamental aspect of BGP dynamics has changed after a decade, not simply how every graph from almost a decade ago may look differently nowadays. In other words, we draw comparisons primarily at a macro level instead of a micro level. Plus, while we do try to repeat many things, we cannot repeat everything exactly the same anyway (it is not the same data source).
- We refined the methodology to explore more interesting information, thus we also have more results. For example, as the dominant pathological BGP dynamics now shifted to a different type (Section 4), which was not even defined previously, we need to put more focus on this new type.
- We discovered that certain results can be represented in better forms. *E.g.,* originally the breakdown of different dynamics types is presented in a figure using raw data (Fig. 2 in [1]), but now we replace it with a figure (Fig. 2(b)) with much cleaner, more statistically meaningful results.

We mainly compare the following fundamental aspects of BGP dynamics with the previous work in [1]:

- General observations of BGP dynamics;
- Fine-grained dynamics statistics, including peer contributions to BGP dynamics and BGP dynamics on a per-route basis; and
- BGP dynamics over time, including possible daily and weekly patterns and inter-arrival time of each BGP dynamics type.

For each aspect, in addition to comparing to corresponding results from [1], if there is any new phenomena, we also analyze them closely.

In [9] Labovitz *et al.* further explored the *origins* of BGP dynamics. Whereas the most extensive BGP dynamics study was still in [9], in this followup work they also analyzed BGP data up to June 1998 to report the evolution of BGP dynamics in certain aspects. So we also compare our results with those in [9] when appropriate.

4. GENERAL OBSERVATIONS OF BGP DYNAMICS

We now describe our general observations of BGP dynamics from August 2005 to January 2006. Gross observation from [1] exposed that pathological instabilities dominated BGP dynamics. In our study, we mainly focus on the following questions: Do BGP dynamics still include as many pathological duplicates as a decade ago, or mainly "healthy" components? What pathological BGP updates, if any, are still around? Are policy changes commonly reflected in BGP dynamics? Can we learn from BGP dynamics whether or not the Internet topology is steady?

Is BGP nowadays performing well? One of our major discoveries is that BGP now experiences more forwarding dynamics and fewer pathological duplicates. As shown

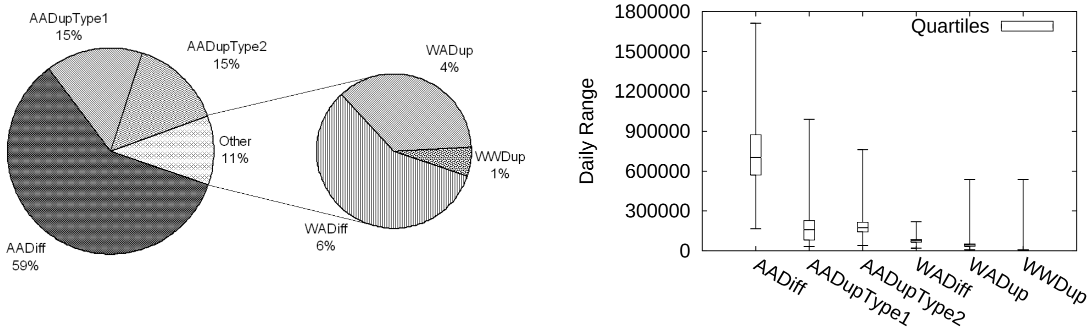

(a) Proportion of each dynamics type

(b) Daily dynamics ranges

Figure 2: Statistics of BGP dynamics over six months

in Figure 2(a), AADiff and WADiff—the two forwarding dynamics types—are now 65% of the total dynamics over the six-month window. If we add the 15% from policy changes (AADupType2), and a certain portion of WADup which could also be forwarding dynamics, we can obtain an even higher percentage, $80 - 84\%$. Figure 2(b) also shows the distribution of each dynamics type (aggregated daily). Although the daily range of every type of dynamics sometimes can be large, it is apparent that their relative size is consistent, and forwarding dynamics now dominate. (Interested readers can compare Figure 2 with Figure 2 in [1].)

Results above are different from those about a decade ago, and are better. The breakdown of routing updates from April through September 1996 [1] shows that the two forwarding dynamics types—AADiff and WADiff, were not only much less than the then dominant pathological WWDup, but also consistently less significant than AADup (the sum of AADupType1 and AADupType2) and WADup. Actually, even after the pathological WWDup dominance went away (because of specific router vendor software changes), the follow-up work of Labovitz *et al.* [9] still shows that the healthy AADiff is much less than today's 59%.

With the current proportions of forwarding dynamics, policy changes, and pathological duplicates, we can view today's BGP healthy, or at least healthier than almost a decade ago. We believe the healthiness of today's BGP is largely due to the continuous efforts by the research community, industry, and BGP operators. For example, research in [1] identified a stateless BGP implementation as the origin for the originally dominant WWDup, and research in [9] helped vendors further reduce duplicate announcements.

How many pathological duplicates are there now? The pathological duplicates as of this study is disparate from those in [1]. Back then, the majority of BGP dynamics consisted of pathological duplicates, and the majority of pathological duplicates were WWDup. In contrast, our measurement shows that pathological dynamics, represented by WWDup and AADupType1, are only 16% of the total dynamics. Furthermore, duplicate withdrawals (WWDup) now make up a mere 1% of observed dynamics, as opposed to the 99% of BGP updates about a decade ago, and duplicate announcements (AADupType1) are now much more dominant than duplicate withdrawals.

In the previous study, some of the WWDups were attributed to the stateless BGP implementation by a specific major router vendor. After the vendor corrected this implementation strategy, the number of duplicate withdrawals dropped by an order of magnitude as observed in 1998 [9]. Our findings as of 2006 indicate that with the effort of the Internet community, this pathological behavior continues to be more insignificant.

The significant reduction of WWDup leads to an increased percentage of another pathological dynamics type: duplicate announcements. In the follow-up work of Labovitz *et al.* in [9], duplicate announcements (shown as AADup in [9]) then made up approximately one third of all dynamics—more significant than today. Duplicate announcements, classified as AADupType1 in our work, made up 15% of all the dynamics from all peers. The hourly rates were typically less than 5,000, with occasional spikes reaching above 100,000 and even 200,000, as shown in Figure 3.

Research in [9] gave two plausible explanations for duplicate announcements: non-transitive attribute filtering, or stateless BGP implementation coupled with BGP minimum advertisement timers. Since the router vendor responsible for these two causes has developed software updates to fix the problem, if these two causes have indeed gone away, there should still be other reasons for today's duplicate announcements. One such reason is related to BGP session reset. Recall BGP is running on top of TCP, and BGP control traffic (*i.e.,* BGP updates) utilize the same links as data traffic. If a link is congested and the congestion reaches a certain level, a BGP session over the link will break, and then restart once the link becomes available again. Every time the session is restarted, the BGP routers of the session will exchange with each other their entire BGP routing tables by sending each other BGP announcements. Because of the ever-increasing demand on bandwidth by data traffic, links used for BGP sessions on the Internet can often be congested, thus causing repeated announcements. (Note that we have filtered out the duplicate announcements caused by session resets between the RRC00 monitor and its peers; the duplicate announcements we collected are those that really happened at those peers and those propagated from other BGP routers.)

Are policy changes common? Dynamics that reflect policy fluctuations (AADupType2) made up 15% of the dynamics from all peers. Hourly rates were typically less than 5,000, seldom above 50,000, and there was only one instance of a spike above 100,000, occurring on November 10, 2005. This spike is clearly visible in Figure 4.

Making a policy change is essentially a matter of a particular AS, and it is not related to topology changes but rather administrative reasons, a human factor. Our study on BGP

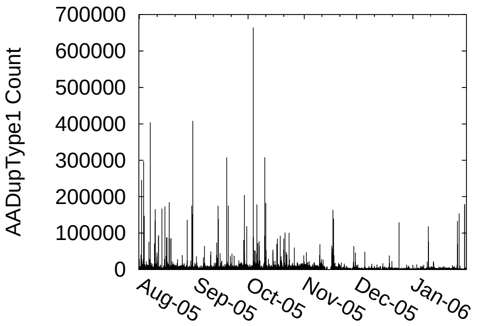

Figure 3: AADupType1 occurrences over six months, aggregated hourly

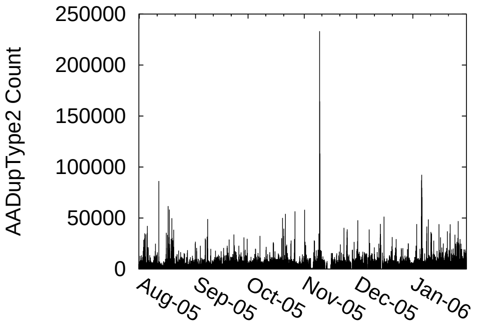

Figure 4: AADupType2 occurrences over six months, aggregated hourly

dynamics contributions by different peers (Section 5) will show that two particular peers (*E* and *G*) were actually responsible for most AADupType2 instances.

Is the Internet topology steady? While only 4%, WADup is very useful in estimating the steadiness of the Internet's topology. Although WADup could simply be a pathological behavior, a very likely reason for WADup is transient topological failure. For example, if the link between a router and its next hop toward a destination is congested, or the link is broken, or the next-hop router is not steady, the router may find its next hop not reachable (no KEEP-ALIVE messages). If it has no alternative route to that destination, it will withdraw its route to that destination. But if the next hop is only temporarily unavailable, the router may find it can use the same route again, thus announcing the same route that was just withdrawn. So, when WADup occurs, BGP is probably doing its job, but the topology is oscillating. Fortunately, with 4% WADup, we can see that at most 4% of BGP updates are caused by topology unsteadiness.

5. PEER CONTRIBUTIONS TO BGP DYNAMICS

A key question for BGP dynamics is: Does every BGP router play basically the same role in contributing to BGP dynamics? If a BGP router is from a bigger AS or more richly connected (as measured by its AS degree), does it contribute more to the dynamics? In this section, we inspect and analyze the BGP updates from the eight peers in our study. We first compare their amounts of announcements and withdrawals, and then compare their contributions to each dynamics type. We conclude that the dynamics level of every BGP router is probably different, but it is unlikely determined by the size of the router's AS.

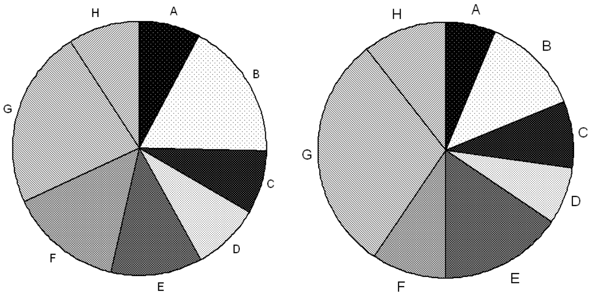

(a) Announcements (b) Withdrawals

Figure 5: Announcements and withdrawals from eight BGP peers in January 2006

We first compare the amount of raw output from each peer. Noticing each peer is from an AS with a different AS degree, as shown in Figure 1, one may expect that a peer with a higher AS degree—thus more BGP peers—will have a larger amount of announcements and withdrawals. Surprisingly, we have found that, even if a BGP peer is associated with a high AS degree, it may only produce as many announcements or withdrawals as a BGP peer with a low AS degree, or even less. For example, peer *H* is from an AS with $1,866$ adjacent ASes (Figure 1), and peer *B* is from an AS with only 6, but as shown in Figure 5, peer *B* generates both more announcements and more withdrawals than *H*. Our understanding of this phenomenon is as follows. If a router is from a bigger AS with a larger AS degree, it will have more BGP neighbors and can probably hear more announcements or withdrawals, thus giving it more opportunities to replace the router's currently preferred paths. However, the probability that the router further propagates every newly received announcement or withdrawal is also lower: the more neighbors the router has, the less likely it is for a new announcement to advertise a more preferred path, and also less likely for a new withdrawal to affect the router's current selection of the preferred path. Also note that a larger AS tends to have multiple connections to the same neighbor. If there is a failure in one of the connections, the AS will not change its AS-level path for reaching a prefix through that neighbor, and hence will not generate a BGP update.

Every peer may also have a different level of contribution to each BGP dynamics type. To study this, we first measured the forwarding dynamics contribution by each of the eight peers during the month of January 2006, as shown in Figure 6. In comparison to those observed in 1997, daily peer contributions are relatively stable. Wide day-to-day swings in peer contribution to each dynamics type do occur, but much less often. (Interested readers can compare Figure 6 with Figure 6 in [1].)

We also measured each peer's contribution to pathological duplicates (Figure 7(a)) and policy changes (Figure 7(b)). WWDups were not graphed due to their negligible presence. What we found interesting is that, compared to peer contri-

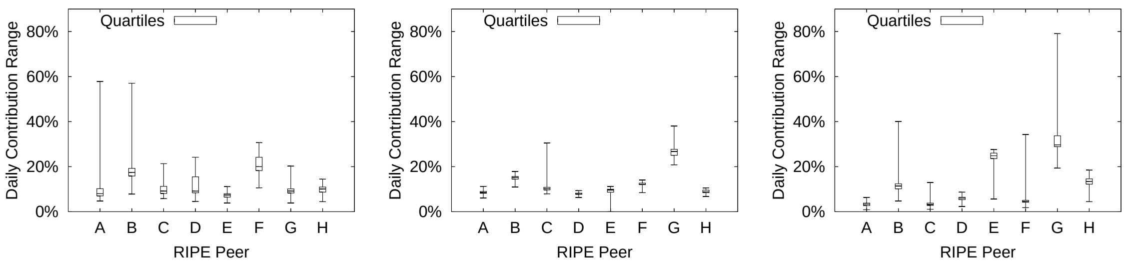

(a) AADiff (b) WADiff (c) WADup

Figure 6: Contributions to forwarding dynamics by individual BGP peers in January 2006

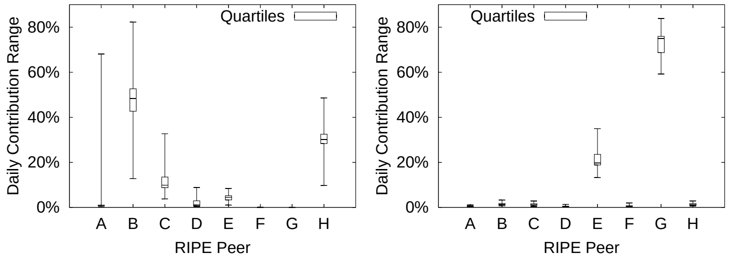

(a) AADupType1 (b) AADupType2

Figure 7: Contributions to policy changes or pathological duplicates by individual BGP peers in January 2006

butions to forwarding dynamics, peer contributions to policy changes or pathological duplicates are much more specific to individual BGP routers. Whereas the contributions for forwarding dynamics are roughly at the same magnitude between different peers, BGP dynamics caused by either policy changes or pathological duplicates are much more disparate. For example, peers B and H are the main sources of pathological duplicates, while the rest of the peers are much less troublesome. Also, about 70% of the observed policy changes are caused by peer G, plus approximately 20% from peer E, and the remaining six peers only contribute an insignificant amount. Peer contributions to these dynamics were not studied in [1], and here our separation of AADup into AADupType1 and AADupType2 again proves to be useful.

6. BGP DYNAMICS BY PREFIX+AS PAIRS

We now examine BGP dynamics on a *per-route* basis. We follow Labovitz *et al.* and aggregate updates into Prefix+AS pairs, each pair about reaching a given prefix via a specific AS. This study can produce more finer-grained results than a study at simply a per-prefix level; for example, it can help answer whether BGP dynamics are dominated by the dynamics of a small collection of routes.

Figure 8 represents daily contributions of Prefix+AS pairs to each dynamics type. To make our results directly comparable to those in about a decade ago, graphs to the left are in the same presentation format as those in [1], and by presenting graphs to the right we further show the general tendencies of daily dynamics contributions by prefix+AS pairs.

We found that the daily dynamics pattern by Prefix+AS pairs as of January 2006 has three major characteristics:

- *It is more intense than that from August 1996, showing BGP is "busier."* In August 1996 there were very few days in which a Prefix+AS pair has more than 200 AADiff instances or 100 WADiff instances. But our results for January 2006, as depicted in Figure 8, show that all three forwarding dynamics types climb to 95% or 100% much more slowly. There are clearly many days in which many prefix+AS pairs have more than 200 AADiff updates or more than 100 WADiff updates. In August 1996, the Prefix+AS pairs that have 10 or less AADiffs can contribute up to 85% of the total AADiff dynamics; but in January 2006, the contribution from the same pairs would at most be only 50%. This phenomenon is understandable—the Internet is becoming bigger, and nowadays every BGP router sees many more address prefixes as well as more updates than before. In other words, we now see more BGP updates for each Prefix+AS pair.
- *It is more consistent than that from August 1996, showing again BGP is "healthier."* For every forwarding dynamics type, the thirty-one curves for January 2006 are much more similar to each other than those for August 1996. The comparison would be easier if the previous study also reported the standard deviation over the thirty-one days in August 1996. Nonetheless, the change over almost a decade is clear. For instance, in August 1996 the prefix+AS pairs with exactly one AADiff instance contributed approximately $3-18\%$ AADiff dynamics, a range of 15%, and the range in January 2006 became $6.5\%(1.3-7.8\%)$. This improved consistency indicates that BGP routing is more stable in January 2006 than August 1996. Random factors, such as sudden link failures, that could cause one day to be significantly more dynamic than the other, are probably much less.
- *AADupType1 and AADupType2 are less consistent than other dynamics types.* For daily BGP dynamics per Prefix+AS pair, we see more variations for dynamics of pathological duplicates and policy changes than forwarding dynamics. Figure 8(d) shows that the pathological duplicates of announcements can vary from day to day a lot. In one day, the Prefix+AS pairs with one or two AADupType1 instances can make up 90% of the total AADupType1 dynamics for the day; in another day, it may be all Prefix+AS pairs with up to 500 AADupType1 instances that make up 90% of the total. This phenomenon agrees with the idea that pathological behavior is typically unpredictable. The day-to-day variation of the policy change dynamics (represented by AADupType2) is less than that of AADupType1, but it is still moderately large due to the nature of policy changes. An important observation here,

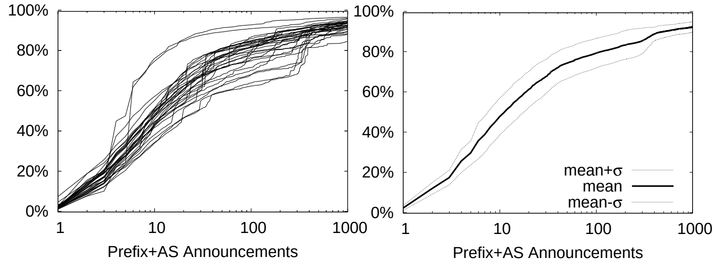

(a) AADiff

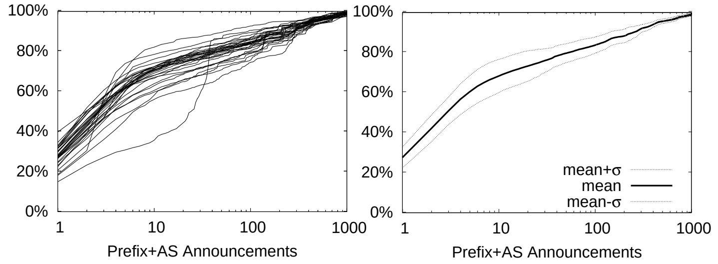

(b) WADiff

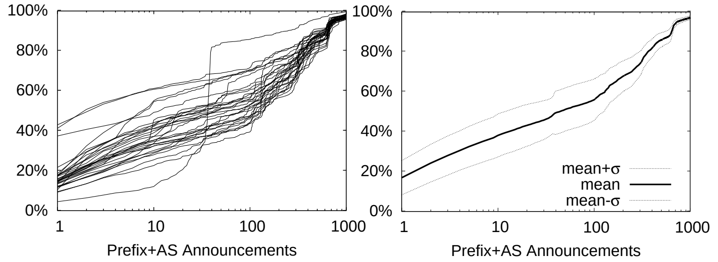

(c) WADup

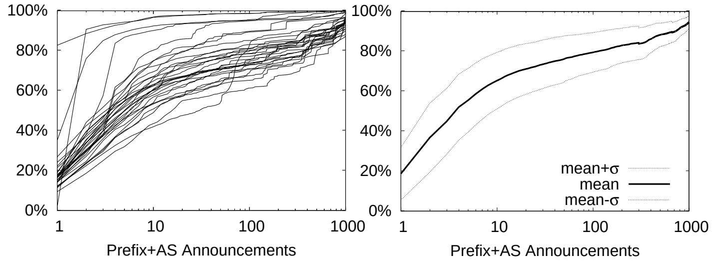

(d) AADupType1

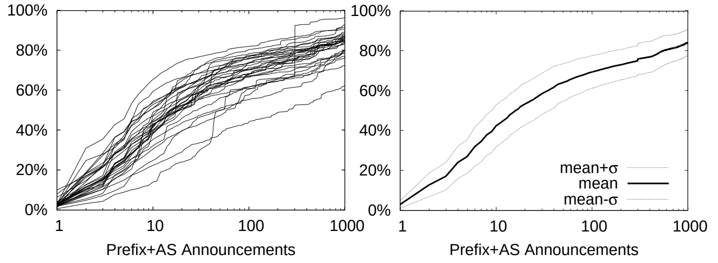

(e) AADupType2

Figure 8: Daily (left) and mean ± standard deviation (right) Prefix+AS contributions to BGP dynamics. Each line on a graph to the left represents the cumulative distribution function for a dynamics type of a single day from January 2006. Any point (x, y) on the line can be read as: "On a given day, all the Prefix+AS pairs that contributed x or less instances of a given dynamics type contributed y percent of the total dynamics of that type."

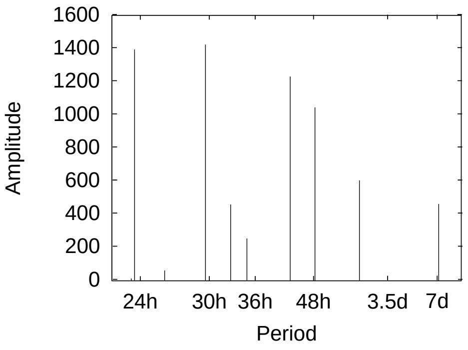

Figure 9: Results from harmonic inversion of January 2006 data (aggregated every 10 minutes)

probably key to improving BGP, is that there is no single route that consistently dominates the dynamics of either AADupType1 or AADupType2. Note that the AADupType1 and AADupType2 dynamics per Prefix+AS pair were not studied previously.

7. BGP DYNAMICS OVER TIME

We also examine BGP dynamics with respect to time. Specifically, we identify whether there are daily or weekly patterns associated with network usage patterns, and examine the inter-arrival times of different dynamics types.

Daily and Weekly Patterns. By performing a spectrum analysis on their data, the previous study in [1] identified daily and weekly patterns in BGP dynamics. To determine if there is any daily or weekly patterns in our data, we aggregated January 2006 AADiff, AADupType1, and AADupType2 events at 10-minute intervals and applied the harmonic inversion technique [29] against the resulting data using the *Harminv* program [30]. Our result as shown in Figure 9 agrees with the previous study. Clearly, there is a strong component with a period very close to 24 hours as well as a component with a seven day period. (We also observe components with periods that are not easily attributed to high times of usage. These additional components suggest that there are other factors that shape the cyclical features of routing dynamics. Identifying such factors is a topic for future research.)

Inter-Arrival Time of Different BGP Dynamics Types. For each dynamics type, the inter-arrival time between routing updates reflects how frequently updates for that dynamics are sent. Since routing updates, especially those from forwarding dynamics, reflect exogenous events such as random node or link failures, one would expect that the inter-arrival time follows an exponential distribution. The previous studies in [1] and its follow-up work [9] discovered, however, all dynamics types have a predominant inter-arrival time of $30s$ and 1 minute. The authors stated that such predominance was mostly due to the minimum BGP advertisement timer at BGP routers; and thanks to adding a random jitter to the originally fixed timer of $30s$ at BGP routers, the follow-up study showed a less significant predominance of these two intervals. We now report our results based on data from January 2006. We studied not only the inter-arrival times for AADiff, WADiff, WADup—the

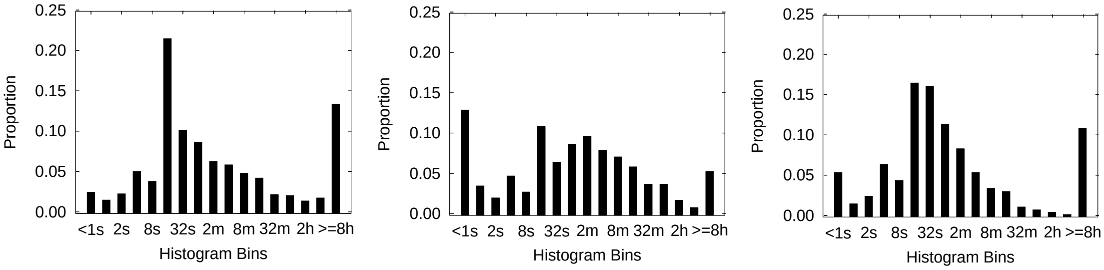

(a) AADiff (b) WADiff (c) WADup

Figure 10: Inter-arrival distribution for forwarding dynamics during January 2006

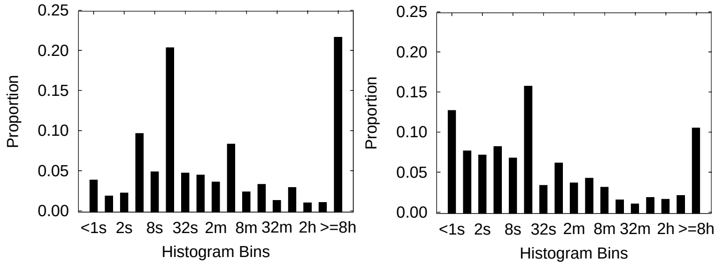

(a) AADupType1 (b) AADupType2

Figure 11: Inter-arrival distribution for pathological updates and policy changes during January 2006

three forwarding dynamics types, but also those for AADupType1 and AADupType2. Figure 10 shows our results of inter-arrival time for forwarding dynamics, and Figure 11 for pathological updates (AADupType1) and policy changes (AADupType2). (In the figures we represent aggregated inter-arrival times as histograms, and divide the histogram into bins based on a modified exponential scale, where every x-axis label marks the starting time of a bin.)

As in previous studies, we also see predominant inter-arrival times. Large portions of dynamics of all types have inter-arrival times around 30*s* (thus showing high proportions for 16*s* and 32*s* bins). We believe this predominance at high frequency can be explained with the same reason as almost a decade ago (*i.e.,* the effect of the minimum BGP advertisement timer). Furthermore, compared to the previous study, all dynamics types have a much larger proportion of inter-arrival times that exceed eight hours, suggesting that many prefixes are stabler than before. (Interested readers can refer to Figure 8 in [1] for direct comparisons.) In short, we now still see predominant high frequency components whereas the low frequency components are stronger.

We further compare the three types of forwarding dynamics with one another. Because a BGP router must wait for a minimum amount of time—typically 30*s*—before announcing a route to the same prefix, there are not much AADiff dynamics with short inter-arrival times. However, compared to AADiff, WADiff instances favor either much shorter inter-arrival times (less than 1*s*), or longer ones. In a WADiff instance, in which a router replaces an originally preferred route with a new route, the router either almost immediately advertises the new route (if available), or waits for the new route to be available—and waits longer on average than AADiff intervals. Finally, as WADups reflect topological instability, we can see the inter-arrival times of WADups are more evenly distributed than AADiffs, especially in the range from 16*s* to 4 minutes. Compared to WADiff, the inter-arrival times of WADup also has a much smaller proportion for the $< 1s$ bin, *i.e.,* after withdrawing a path it takes longer to announce the same path than a different path—we suspect this is related to route flap damping [5].

We also studied the inter-arrival times of AADupType1 and AADupType2, which show significantly different histograms (Figure 11): AADupType1 dynamics have an inter-arrival distribution leaning toward longer times, while AADupType2 dynamics significantly lean toward shorter times. The difference in inter-arrival distribution reinforces our separation of AADup dynamics into two classes.

8. CONCLUSIONS

Understanding the dynamics of BGP routing is critical to the solid growth and maintenance of the Internet routing infrastructure. As almost a decade has passed since the most comprehensive study on this topic, by following and refining the methodology in that study, we now revisit BGP dynamics using six months of recent BGP updates. While certain characteristics remain the same, significant changes have occurred.

We have found that, unlike almost a decade ago when the majority of updates are pathological withdrawals, BGP dynamics are now dominated by forwarding dynamics, showing that most of the time BGP is really doing its job—discovering routes for packets. Also at most 4% of the dynamics are caused by transient topological failures. We believe this “healthier” behavior of BGP is largely a consequence of the maturity of vendor implementations of BGP and the experience of BGP operators and network engineers over the last decade.

By enhancing the previous dynamics classification, we are also able to discover that 15% of updates nowadays are probably related to policy changes. Furthermore, the Internet is still plagued with duplicated updates, except that most duplicates are now announcements (15%) instead of withdrawals (only 1%).

A closer look at BGP dynamics reveals that, just as it was almost a decade ago, BGP dynamics are not provided by only a small number of BGP peers; all of them contribute. A surprising discovery here is that a peer’s contribution may not be proportional to its AS’s size. Also a new findings is that, for forwarding dynamics each peer’s contributions are more stable than almost a decade ago, and the dynamics due to policy changes or duplicate announcements are often from specific peers.

While, as before, BGP dynamics are not dominated by a small collection of routes, we now also see forwarding dynamics more intense—thus busier, and more importantly, more consistent from day to day—thus healthier. Moreover, BGP dynamics still display certain temporal properties, and we can still observe daily and weekly patterns as well as some predominant inter-arrival times; but we now see more very long inter-arrival times (as many routes are more stable), and the inter-arrival times clearly differ between different dynamics types.

Many open issues still exist. BGP routing is very complicated and certain phenomena we observe still warrant more inspection. More BGP data over a longer period with more peers, probably with a mechanism for continuously monitoring BGP dynamics, would greatly advance this work. And it is also critical to see what insights vendors and network operators can draw from today's BGP dynamics.

9. REFERENCES

[1] C. Labovitz, G. Malan, and F. Jahanian, "Internet routing instability," in *Proc. ACM SIGCOMM*, September 1997.

[2] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, 2006.

[3] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771, 1995.

[4] R. Chandra, P. Traina, and T. Li, "BGP Communities Attribute," RFC 1997, 1996.

[5] C. Villamizar, R. Chandra, and R. Govindan, "BGP Route Flap Damping," RFC 2439, 1998.

[6] T. Bu, L. Gao, and D. Towsley, "On characterizing BGP routing table growth," in *Proc. IEEE Global Internet*, November 2002.

[7] G. Huston, "Analyzing the Internet's BGP routing table," *The Internet Protocol Journal*, vol. 4, no. 1, March 2001.

[8] "BGP reports," `http://bgp.potaroo.net/`.

[9] C. Labovitz, G. Malan, and F. Jahanian, "Origins of Internet routing instability," in *Proc. IEEE INFOCOM*, March 1999.

[10] B. Chinoy, "Dynamics of Internet routing information," in *Proc. ACM SIGCOMM*, September 1993.

[11] R. Govindan and A. Reddy, "An analysis of Internet inter-domain topology and route stability," in *Proc. IEEE INFOCOM*, April 1997.

[12] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, "BGP routing stability of popular destinations," in *Proc. ACM Internet Measurement Workshop*, November 2002.

[13] V. Paxson, "End-to-end routing behavior in the Internet," in *Proc. ACM SIGCOMM*, 1996.

[14] M. Caesar, L. Subramanian, and R. Katz, "Towards localizing root causes of BGP dynamics," Tech. Rep. CSD-04-1302, UC Berkeley, 2004.

[15] Z. Mao, R. Bush, T. Griffin, and M. Roughan, "BGP beacons," in *Proc. ACM Internet Measurement Conference*, October 2003.

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

[17] T. Griffin and G. Wilfong, "An analysis of BGP convergence properties," in *Proc. ACM SIGCOMM*, August 1999.

[18] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, "Delayed Internet routing convergence," in *Proc. ACM SIGCOMM*, August 2000.

[19] C. Labovitz, A. Ahuja, R. Wattenhofer, and V. Srinivasan, "The impact of Internet policy and topology on delayed routing convergence," in *Proc. IEEE INFOCOM*, April 2001.

[20] T. Griffin and G. Wilfong, "Analysis of the MED oscillation problem in BGP," in *Proc. ICNP*, November 2002.

[21] Z. Mao, R. Govindan, G. Varghese, and R. Katz, "Route flap damping exacerbates Internet routing convergence," in *Proc. ACM SIGCOMM*, August 2002.

[22] J. Cowie, A. Ogielski, B. Premore, and Y. Yuan, "Internet worms and global routing instabilities," in *Proc. of SPIE International symposium on Convergence of IT and Communication*, July 2002.

[23] J. Cowie, A. Ogielski, B. Premore, E. Smith, and T. Underwood, "Impact of the 2003 blackouts on Internet communications," `http://www.renesys.com/news/2003-11-21/Renesys_BlackoutReport.pdf`, November 2003.

[24] Mohit Lad, Xiaoliang Zhao, Beichuan Zhang, Dan Massey, and Lixia Zhang, "An analysis of BGP update burst during Slammer worm attack," in *Proceedings of 5th International Workshop on Distributed Computing*, December 2003.

[25] Jun Li, Dejing Dou, Zhen Wu, Shiwoong Kim, and Vikash Agarwal, "An Internet routing forensics framework for discovering rules of abnormal BGP events," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 55–66, October 2005.

[26] RIPE NCC, "RIPE routing information service raw data," `http://data.ris.ripe.net/`.

[27] "University of Oregon Route Views Project," `http://antc.uoregon.edu/route-views/`.

[28] L. Wang, X. Zhao, D. Pei, R. Bush, D. Massey, A. Mankin, S. Wu, and L. Zhang, "Observation and analysis of BGP behavior under stress," in *Proc. ACM Internet Measurement Workshop*, November 2002.

[29] V. A. Mandelshtam and H. S. Taylor, "Harmonic inversion of time signals," *Journal of Chemical Physics*, vol. 107, no. 17, 1997.

[30] Steven G. Johnson, "harminv," `http://ab-initio.mit.edu/wiki/index.php/Harminv`.