

BGP Routing: A study at Large Time Scale

Georgos Siganos

U.C. Riverside

Dept. of Comp. Science

siganos@cs.ucr.edu

Michalis Faloutsos

U.C. Riverside

Dept. of Comp. Science

michalis@cs.ucr.edu

Abstract—In this paper, we conduct a longterm analysis of BGP routing properties such as path stability. Our work complements previous studies that examine the BGP routing behavior at smaller time scales such as its convergence to a routing update. We focus on properties that would reflect the BGP evolution due to growth, policy, and business reasons. We use daily snapshots of a number of BGP tables and we study the evolution of the paths for every IP-prefix advertised between two Autonomous Systems (source AS, destination AS, IP-prefix).

In a high level, we observe that BGP routing is characterized by a) fairly robust routing with usually few dominating paths per prefix, b) routing richness in the advertised paths and c) significant number of short lived IP-prefix advertisements.

I. INTRODUCTION

“How often does a BGP path change?”

“How many *legitimate* paths exist between two ASes?”

Our initial goal was to develop a simulation model for BGP research, but we stumbled on many questions like the two above. We found that it was impossible to conduct realistic BGP studies without having a general framework for fundamental properties such as routing stability, path multiplicity, and lifetime of advertised IP-prefixes. As a result, this paper attempts to provide guidelines for BGP model. In contrast to previous work, we focus on the longterm properties: we study daily snapshots of routing tables over the duration of three years for a total of more than 7 million advertised paths.

Most previous studies examine the stability and convergence properties of routing at small time scales [7], [4], [5], [6], [3], [8], [10]. Most of them focus at the level of routing updates, and examine the time needed to converge after a routing change. There is no recent analysis of the long term properties of routing. The most relevant study we found [10] examines routing properties during 1994 and 1995. Although that work is focused on routing at the router level, it does provide some statistics on the ASes that traceroutes traverse.

In a nutshell, we find that the routing is fairly robust, but there is also a lot of “routing noise”. More specifically, our analysis leads to several observations. First, one out of three prefixes appears for less than ten days. Second, approximately half of the prefixes are dominated by one path, which suggests good routing stability. Third, we observe that 70% of the IP-prefixes use at least three distinct routing paths in their lifetime, but practically all of the IP-prefixes have less than 10 distinct paths. Fourth, we observe that 65% of the IP-prefixes have one

path that follow continuously at least for 50% of the duration of their life.

The rest of this paper is structured as follows. In section II, we present some previous work and definitions. In section III we study the stability in terms of number of prefixes advertised and the richness of paths. In section IV, we analyze the prevalence of the routing paths. In section V, we study the persistence. In section VI, we conclude our work.

II. MODEL AND BACKGROUND

We use real BGP routing tables in our analysis. We provide a simplified explanation of the data of a BGP table for completeness. A BGP table consists of advertisements from the BGP routers of different ASes. Each BGP router belongs to an AS and follows the policy of that AS. Each AS advertises through each routers what are the IP-prefixes (chunks of IP address space) that it is capable of reaching. These IP-prefixes are either internal to the AS, or they can be reached through a neighboring AS. The type of IP-prefixes that an AS advertises is subject to its policy and business relationships with the neighboring ASes. This way, each BGP router maintains a table of IP-prefixes and the corresponding AS path that can use to reach the destination AS that owns the IP-prefix. The path in the routing table of a router is the path, that a packet whose address matches the IP-prefix will follow from that router to the destination AS.

Let us clarify the routing elements that we measure in our analysis. First, an AS may advertise many different IP-prefixes to another AS. Each of these IP-prefixes may be associated with a different path between the two ASes. For this, we use the term **triplet** to refer to: a source AS, a destination AS, and the IP-prefix that the destination AS advertises to the source AS. Intuitively, a triplet represents a path for a given prefix between two ASes. This path can change in time, and it is exactly the stability of this path that we want to examine. We will often use the term IP-prefix to refer to a triplet, when it is clear that we refer to the IP-prefix for a specific pair of ASes.

We need to develop metrics to quantify routing stability. We use the concepts of prevalence and persistence of a path to quantify the stability of the path of a triplet [10]. **Prevalence** of a path between two nodes is the unconditional probability of observing the given path or, simply, the percentage of time that two nodes communicate over the given path. For example, a path that appears two thirds of the observed time has a prevalence of 66%. We define the **dominant** path to be the path that we observe more often among all paths of a triplet. Another aspect of stability is whether the paths appears in consecutive

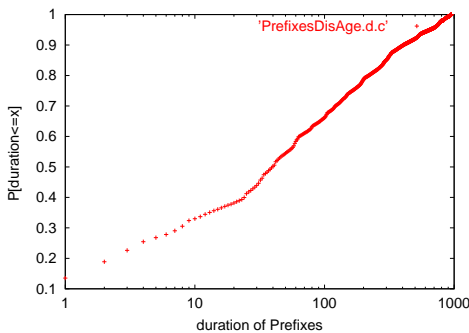


Fig. 1. CDF of number of days we observe a Prefix.

intervals. In the previous example, we can have two extreme cases: the path can appear for the first two thirds of the time continuously or it can appear for two days at a time followed by a third day with a different path. More rigorously, given that we observe path r at time t , how long before that path is likely to have changed? We refer to this notion as **persistence**. We use the term **longest continuous or persistent path** of a triplet for the path with the longest continuous presence among all paths of the triplet.

In order to describe the distribution of a variable we will use both the cumulative distribution function or **CDF** and the complementary cumulative distribution function or **CCDF**. CDF of a variable is the percentage of instances of the variable, that have a value less or equal to a given number. The CCDF of a variable is the percentage of instances that have a value greater to a given number.

In the most related study, Paxson [10] analyzes the stability of routing paths at the router, city and AS level, with the focus on the router level, using traceroute analysis. He studied 40,000 paths between 1994 and 1995. Paxson reports that more than 60% of the measured triplets follow only one path and approximately 90% of the triplets have prevalence 70%.

The Data of Our Analysis. We use data that is collected by a route server at Oregon Route Views Project [9], [2]. The data contains a number of real BGP routing tables of real BGP routers geographically distributed. We examine these tables from the end of 1997 till the start of 2001. It has been reported [1] that this data sees only a subset of the total AS topology. However, in our study, we examine “what a router sees” and our results do not require a complete topology. Consistency is more of an issue, and we take this into account by measuring the time both in absolute term (days) and in relative term (lifetime that we actually observe an AS).

Our data set is over **107 Gbytes** and contains over **one billion** paths which correspond to the evolution of over **400 thousands** unique source, destination pairs of AS, and over **7 million** triplets, over a period of more than three years.

III. ADVERTISEMENTS BETWEEN BGP ROUTERS

In this section, we examine what does a BGP router hear from other BGP routers. We want to obtain a better understanding of what our data shows before we analyze it for stability. First, we examine the lifetime of a triplet: how long does

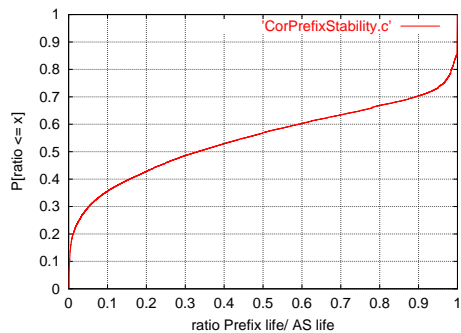


Fig. 2. CDF of the ratio of the number of days i observe a prefix, and the number of days i observe the AS that advertises the prefix.

a router receive the same IP-prefix advertisement from a destination AS¹. Second, we examine how many distinct routing paths² does a routing table have between two ASes at a given time independently of which IP-prefix they are related to.

A. Lifetime of an IP-Prefix

We find that a large portion of the IP-prefixes advertised between two ASes are short lived and disappear after a small interval of time.

14% of the IP-prefixes are visible for only one day. In figure 1 we plot the CDF of the total number of days that we observe a triplet. Some other results from this graph are that 14% of the prefixes are visible only for one day, 33% of the prefixes are visible for less than 10 days, and 44% of the prefixes are visible for less than 30 days.

The previous surprising result made us investigate this issue further. The question was: could it be that an IP-prefix disappears very quickly because I stop seeing the destination AS altogether? The answer is no, although this could be a contributing factor. We find that IP-prefixes are much more short-lived than their ASes.

Approximately one third of the IP-prefixes is short-lived while another one third is long-lived. We focus on the relative lifespan of an IP-prefix that is defined as the ratio of the days we observe the IP-prefix over the number of days we observe the AS that advertises it. In figure 2, we plot the cumulative distribution function of the relative duration of an IP-prefix between two ASes. From the figure, we have that 30% (y-axis) of the IP-prefixes are short-lived and are present for less than 5% (x-axis) of the time. On the other hand, we have that approximately one third 30% of the prefixes are long lived and are present for more than 90% of the life of an AS. The rest of the prefixes have a relative lifespan between 5-90%.

Such short-lived advertisements are most likely due to errors and misconfigurations as it is discussed in [7] [8]. Address aggregation and traffic engineering likely contribute to that effect too.

We were concerned that short-lived IP-prefixes could bias the results. For example, a one-day IP-prefix has a unique path that

¹The destination AS is the AS that sends the advertisement but it is the destination of the related routing path.

²Clearly, we can only measure the advertised paths. This is a lower bound of the possible paths. There may be paths that are not advertised unless the current advertised path becomes unavailable.




Fig. 3. Percentage of source-destination pairs versus the number of distinct paths, that the pair has. Each line corresponds to a different time.

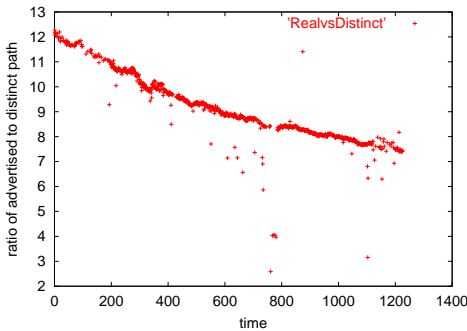


Fig. 4. Ratio of the number of advertised per distinct paths in time.

prevails for 100% of the time, which gives an artificial indication of stability. For this reason, we filter out IP-prefixes that appear for less than 30 days. This filtering is also aligned with our intention to study the long term behavior of the routing.

B. Number of Distinct Paths Between ASes

We observe that IP-prefixes between two ASes follow different paths. This suggests two things: a) policy and traffic engineering have a significant presence, b) there is routing richness, so even in case of a path failure there will be other paths that can continue the communication.

20% of the source-destinations pairs use more than one distinct path simultaneously. In figure 3 we plot the percentage of source destinations pairs versus the number of distinct paths for three different time instants, the first day of observance, the 600th day and the 1200 day. From the figure we observe that for all three plots the percentage of source destination pairs that have one distinct path is around 80%, for 2 distinct paths the percentage goes to 17% and for 3 distinct paths we have a percentage of 3%. This means that for 20% of the source, destination pairs, the source uses more than one path to reach the prefixes that the destination advertises.

The ratio of the number of advertised per distinct paths decreases steadily with time. In figure 4, we plot the ratio of the number of advertised per distinct paths in time. As we can see from the plot this number is decreasing steadily with time.

IV. PREVALENCE OF PATHS

In this section, we study the properties of the dominant path of triplets that appears for more than 30 days. First, we observe

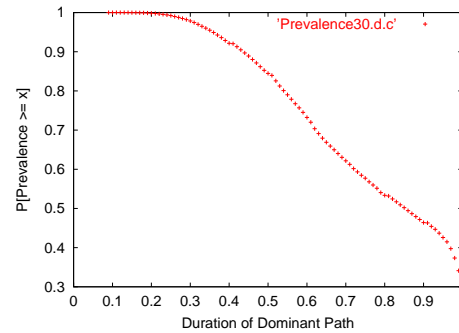


Fig. 5. Complementary CDF of prevalence of the dominant path

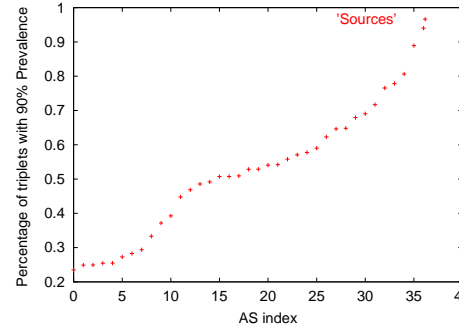


Fig. 6. Percentage of triplets that have more than 90% prevalence for each source AS, versus the AS index in increasing order.

that in terms of prevalence, routing seems to be stable. Then, we study which parameters affect the prevalence. We find that prevalence is affected negatively by the length of the dominant path and the lifespan of the triplet. We also find that the address space of the IP-prefix does not affect its prevalence in a significant way.

Every triplet has a dominant path that appears for some percentage of X% the lifetime of the triplet. We will call X% the **duration of the dominant path or the prevalence of the triplet**.

The majority of triplets have a dominant path for 84% of their time. In figure 5, we plot the CCDF of the duration of the dominant path. We find that 30% (y-axis) of triplets have one path, in other words they have a dominant path with 100% duration (x-axis). 50% of the triplets have a dominant path with more than 84% duration. On the negative side, we observe that 10% of the triplets have a prevalence less than 43%.

Routing prevalence is affected by the source or the destination of the AS triplet. We find that the prevalence properties are not the same for different source or destination ASes. Some ASes are more likely to have prevalent triplets. To measure this, we group the prevalence of the triplet (source, destination, IP-prefix) by the AS source and study the prevalence of each set of triplets. We do the following steps. First, for each set of triplets, we calculate the distribution of the duration of the dominant path for each source AS. (Each such distribution would look like our figure 5 but for only one source). For each of such distribution, we find the percentage of triplets that have a dominant path with duration 90% or more. Intuitively, we can think of this as the percentage of “stable” triplets that are served

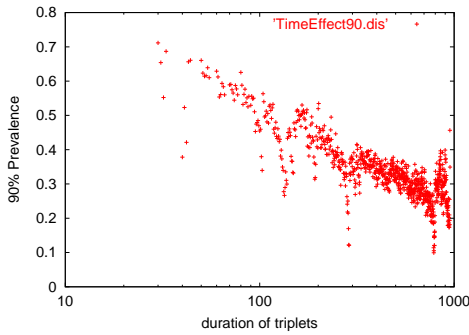


Fig. 7. Prevalence decreases with the observed life of the triplet.




Fig. 8. Complementary CDF of prevalence versus the length of the prevalent path.

by mainly one path. We study what is the percentage of “stable” triplets for each AS. In figure 6, we plot this percentage of “stable” triplets for each AS source. We observe that there is a high variation between different ASes. We arrive also at variable distribution when we group triplets by AS destination.

From this analysis, we conclude that routing prevalence is a function of the AS source or destination. As a result, if we create a BGP model, we should not assume uniform properties of routing prevalence and more generally stability.

Prevalence decreases with the life-span of the triplet. We group triplets by the duration of their observed life and repeat the analysis above for each group of triplets. In figure 7, we plot the percentage of “stable” triplets versus the lifespan of those triplets. From the figure, we observe that prevalence decreases with the observed duration of a triplet. This is something that agrees with our intuition that the longer the observation the more likely it is that the routing path will change in a substantially way.

Prevalence decreases with the length of the path. In figure 8 we plot the CCDF of prevalence for various lengths of the dominant path. Clearly the shorter the path the higher the prevalence we have. In order to quantify this relation, we take from the previous plot the percentage of triplets that have a specific prevalent value, for example greater than 90%, and we plot this value versus the length of each path. In figure 9 we can see the result for various percentages, like 90%, 70% and so on. We find that prevalence is depended on the length of the path, and that prevalence **decreases roughly exponentially** with the length of the path. We can approximate each plot using linear regression with correlation coefficient always higher than 98%.




Fig. 9. Percentage of Paths with a given Prevalence versus Length, for various length of Paths




Fig. 10. Complementary CDF of prevalence for various IP Prefixes sizes

$P_x = e^{a_x * L}$, where L is the length of the dominant path, a_x a constant, and P_x is the percentage of triplets that have prevalence more than x , i.e. a dominant path with duration more than x .

Prevalence does not depend on the size of the IP-prefix advertised. We group triplets according to the size of their IP-prefix. In figure 10, we plot the prevalence distribution as seen for three different IP-prefixes sizes: 8, 16, and 24. We observe that the three distributions look practically the same. From the figure, we see that all three sizes of prefixes have similar distributions. We actually see the 8 and 24 sizes to be almost identical, with 16 showing some small divergence. This is a counter-intuitive observation, since we expected large prefixes to be more stable.

V. PERSISTENCE OF PATHS

In this section, we study the persistent or longest continuous path of a triplet. Our goal is to quantify what is the likelihood that once a path appears it will stay for a long time. Even for a triplet with a dominating path of 66% duration it is possible to have an alternate path every third day. Again, we use the triplets that appear for more than 30 days. First, we want to find how many different paths a triplet uses in its duration.

70% the triplets use at most three distinct paths. In figure 11, we plot the CDF of the distinct number of paths of a triplet. 31% of the triplets have one path for the duration of their existence, and 70% use no more than three distinct paths in their lifetime.

Next, we want to study for how long we observe the longest continuous path. We want to express this duration relatively to




Fig. 11. CDF of the Number of Distinct paths per IP Prefix.




Fig. 12. CDF of the percentage of duration of the longest continuous path

the lifespan of the triplet.

BGP routing paths are persistent. In figure 12, we plot the CDF of the duration of the longest continuous path divided by the lifespan of the triplet. We observe that 65% (y-axis) of the triplets have a persistent path that last more than 50% (x-axis) of their duration, and 40% of the triplets have a persistent path for more than 80%. It is interesting to note that more than 99.7% of the triplets have a persistent path for at least 10% of the lifespan of the triplet.

We have seen that the most persistent path of each triplet is quite persistent in time relative to the lifespan of the triplet. We examine also the duration of the persistence in absolute time (days).

90% of the most persistent paths last more than 27 days.

In figure 13 the line 'longest' represents the CDF of the duration of the persistent path of a triplet. We observe that 50% of the most persistent path of a triplet last more than 80 days, and 20%

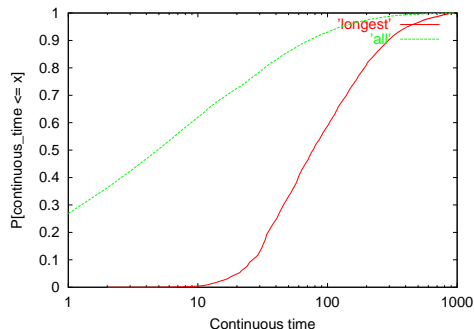


Fig. 13. CDF of the continuous time we see a path.

of those paths last 193 days.

So far our analysis considered only the most persistent path from each triplet. For completeness, we also plot the continuous time for all the paths and all their occurrences.

62% of all paths last for less than 10 days. In figure 13 we plot the CDF of the continuous time we have the same path. 62% of the paths appear for less than 10 days continuously and 78% appear for less than 30 consecutive days. Note though that a large part of these paths, 26%, are present for only one day.

VI. CONCLUSIONS

We presented a longterm study of the stability of routing at the Autonomous Systems level. Our analysis leads to several interesting observations:

- 1) We found an unexpected large number of short lived prefixes. These prefixes correspond to 1/3 of the total, and last for less than 10% of the duration of the AS that advertises them.
- 2) 20% of the source destination pairs use more than one distinct path, this seems to be a manifestation of traffic engineering.
- 3) fairly robust routing with few dominating paths per prefix.

These observations provide some initial guidelines to develop a simulation model for BGP research. First, we can assume that path changes, due to growth and policy, are for the most cases infrequent. Note though that we can not assume uniform properties of routing stability, since as we saw prevalence and subsequently persistence depends also on the AS. Second, we should include in the BGP model that many ASes advertise more than one prefix, and sometimes different prefixes follow different paths. Assuming that each AS advertises only one prefix, which is observed for the duration of the life of an AS, simply fails to characterize the much more dynamic nature of the actual exchange of advertisements.

This work is complementary to other people's work that focuses on short term convergence behavior of BGP. The final goal is to develop sufficient accurate model for BGP simulations that should consider both the longterm and the short term convergence properties.

REFERENCES

- [1] Qian Chen, Hyunseok Chang, Ramesh Govindan, Sugih Jamin, Scott J. Shenker, and Walter Willinger. The origin of power laws in internet topologies revisited. *infocom*, 2002.
- [2] National Laboratory for Applied Network Research. Online data and reports. Supported by NSF, <http://www.nlanr.net>, 1998.
- [3] Lixin Gao, Timothy Griffin, and Jennifer Rexford. Inherently safe backup routing with bgp. *INFOCOM*, 2001.
- [4] Timothy Griffin and Brian Premore. An experimental analysis of bgp convergence time. *ICNP*, 2001.
- [5] Timothy Griffin and Gordon Wilfong. An analysis of bgp convergence properties. *ACM SIGCOMM*, 1997.
- [6] Craig Labovitz, Abha Ahuja, Abhijit Abose, and Farnam Jahanian. An experimental study of bgp convergence. *ACM SIGCOMM*, 2000.
- [7] Craig Labovitz, G. Robert Malan, and Farnam Jahanian. Internet routing instability. *SIGCOMM*, 1997.
- [8] Ratul Mahajan, David Wetherall, and Tom Anderson. A study of bgp misconfiguration. *ISMA*, 2001.
- [9] University of Oregon Route Views Project. Online data and reports. <http://www.routeviews.org/>.
- [10] V. Paxson. End-to-end routing behavior in the Internet. *ACM SIGCOMM*, August 1996.