

Studying Interdomain Routing over Long Timescales

Giovanni Comarela
Boston University
Boston, USA
gcom@bu.edu

Gonca Gürsun
Boston University
Boston, USA
goncag@bu.edu

Mark Crovella
Boston University
Boston, USA
crovella@bu.edu

ABSTRACT

The dynamics of interdomain routing have traditionally been studied through the analysis of BGP update traffic. However, such studies tend to focus on the volume of BGP updates rather than their effects, and tend to be local rather than global in scope. Studying the global state of the Internet routing system over time requires the development of new methods, which we do in this paper. We define a new metric, MRSD, that allows us to measure the similarity between two prefixes with respect to the state of the global routing system. Applying this metric over time yields a measure of how the set of total paths to each prefix varies at a given timescale. We implement this analysis method in a MapReduce framework and apply it to a dataset of more than 1TB, collected daily over 3 distinct years and monthly over 8 years. We show that this analysis method can uncover interesting aspects of how Internet routing has changed over time. We show that on any given day, approximately 1% of the next-hop decisions made in the Internet change, and this property has been remarkably constant over time; the corresponding amount of change in one month is 10% and in two years is 50%. Digging deeper, we can decompose next-hop decision changes into two classes: churn, and structural (persistent) change. We show that structural change shows a strong 7-day periodicity and that it represents approximately 2/3 of the total amount of changes.

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network monitoring; C.2.5 [Local and Wide-Area Networks]: Internet – BGP

General Terms

Measurement

Keywords

BGP; Interdomain Routing

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
IMC'13, October 23–25, 2013, Barcelona, Spain.
Copyright 2013 ACM 978-1-4503-1953-9/13/10 ...\$15.00.
<http://dx.doi.org/10.1145/2504730.2504771>.

1. INTRODUCTION

There are many aspects of the interdomain routing system that are important to understand, including its stability, scalability, and security. However, a particularly difficult problem is understanding the overall structure of interdomain routing and how it evolves over time. The immense size, complexity, and continuous growth of the system make it challenging to gain a useful understanding of the nature of routing changes over time.

This problem is important because currently there are no metrics able to provide useful information about the rate of routing changes of the interdomain routing system. Such metrics could contribute to understand the impact of de-peering disputes, link failures, merging of autonomous systems and any other event that may affect globally the routing structure of the Internet.

To address this challenge, we focus in this paper on answering basic questions about how Internet routing has changed over time. As a first step, we are interested in characterizing the *rate of change* of the routing system and the dynamics of its change. We seek to answer these questions for the system as a whole (i.e., globally) and over long timescales. In this respect we differ from most prior work which has asked more specific questions about the evolution of Internet routing.

We believe that one reason that temporal change has not been extensively studied to date is that good methods and metrics to study it have not existed. Accordingly, the first contribution we make is to propose a metric that can be used to capture the dynamics of routing change, and show how to apply it to available BGP data. The key idea behind our approach is that the entire state of the global routing system at any time can be captured as a large set of tuples that express the next-hop decisions made by each AS with respect to each prefix. From this starting point, we define a natural measure of change for this set of next-hop decisions, and then show the utility of the measure.

This kind of approach to the study of interdomain routing moves away from studying the system in terms of an AS topology. Rather than focusing on an imperfectly understood topology, we focus on the fundamental decisions that are made by each AS, namely, its next-hop choices. Further, our approach also avoids the problems of studying routing in terms of BGP update messages, which are very difficult to interpret in a global fashion. Analyzing BGP updates is problematic, requiring special methods and heuristics, because BGP traffic represents a variety of effects (table dumps as well as route changes) and because many messages do not result in actual changes to the routing systems (e.g., updates that do not trigger new next-hop decisions in the recipient).

The new metric we develop is based on the general concept of *Routing State Distance* (RSD) [7]. However, we show that RSD as originally used is insufficient to study the dynamics of a system that

changes over time. Hence, we develop a significant improvement to RSD, called *Multiple Next-hop Routing State Distance* (MRSD). MRSD has the advantage of being applicable in a much wider set of routing systems than the original RSD, while still capturing the same concept. Further, whereas the original concept of RSD was developed to compare different prefixes, in this study we apply MRSD to compare routing to the *same* prefix at *different* times. We refer to the application of MRSD over time as *Temporal MRSD* (TRSD).

We use TRSD to address our motivating questions by applying it to a large corpus of BGP data. Our raw data comes from BGP Routing Information Bases (RIBs) made available through the Route Views and RIPE measurement projects. We convert downloaded RIBs to a cleaned format for processing. The converted data is smaller than the raw RIBs but still comprises more than 1TB (uncompressed); we store it in *Hadoop Distributed File System* (HDFS) and process it using *Hadoop MapReduce*.¹

Our study looks at routing change at two granularities: at a daily level, and at a monthly level. At the finer timescale, we study the daily evolution of the routing system during three distinct years: 2005, 2008 and 2011. At the coarser timescale, we study the monthly evolution of the routing system over 8 years, from the beginning of 2005 until the end of 2012.

Our results reveal some interesting aspects of Internet routing. We show that approximately 1% of all visible next-hop decisions change each day, and that this rate is about 10% at the timescale of one month and 50% at the timescale of 2 years. Surprisingly, we show that these values are remarkably constant over the period investigated, despite the immense change and growth of the network during our period of study. We also show a decomposition of the daily TRSD time series in two components, intended to capture the difference between sustained (policy-driven) changes in routing versus churn (temporary changes, e.g., due to equipment failures). About 2/3 of changes on a daily basis fall into the sustained category, with the remaining 1/3 classified as churn. Interestingly, we show that sustained changes show a strong weekly periodicity, with the majority of sustained changes made in the workweek. On the other hand, routing churn is better described as noise without a strong periodic component. Finally, we also study the relative rate of routing changes across different ASes. We show that this rate is very long-tailed, i.e., that a small fraction of ASes are responsible for the vast majority of changes to next-hop decisions.

The remaining of this paper is organized as follows. Section 2 briefly summarizes some related work. Section 3 presents the notation and definitions while Section 4 presents our data collection methodology. Measurement results are presented in Section 5 and finally, discussions, conclusions and future work directions are stated in Section 6.

2. RELATED WORK

The goals and methods of our study differ from or build on prior work in a number of ways.

BGP instability measurement: Considerable prior work has looked at the stability (or instability) of the interdomain routing system, including [11, 8, 12, 9, 4]. On one hand, these studies are related to ours in also performing long timescale studies of the BGP system. However, these studies do not emphasize the global evolution of routing *decisions* but focus mainly on the dynamics of routing *traffic*. They focus on characterization and analysis of BGP messages: how to understand BGP traffic dynamics over time, and how to characterize stability of BGP routing for specific des-

tinations. Our focus is on the changes to next-hop decisions made throughout the global routing system.

Longterm Internet evolution: Some other studies have looked at global properties of the AS-level Internet over multi-year timescales. In particular, [13] explores a dataset composed of AS paths collected from 2005 to 2009, in order to show that the Internet is flattening. Another study is [3], in which the authors analyzed the growth of the Internet over a period of 10 years, focusing on the economical roles played by each AS, the types of relationship between ASes, and differences across geographical regions. Our work complements those studies by looking at similarly long timescales, but focusing on more basic questions such as the rate of change in routing, the nature of routing churn, and the ASes responsible for the most routing changes.

Routing State Distance: Finally, we mention that Routing State Distance (RSD) was originally used in [7] and [6], but the definition of RSD differs in important ways from the more general MRSD metric we define here. Our definition of MRSD is a generalization of RSD; it extends RSD to additional routing configurations that RSD could not handle, while remaining equivalent to RSD for any cases in which RSD can be used.

3. NOTATION AND DEFINITIONS

In this section we define the metrics we use to analyze change in the global routing system.

3.1 Multiple next-hop RSD

Our starting point is the notion of Routing State Distance (RSD) as defined in [7]. Briefly, RSD is a metric that defines the ‘distance’ between two destinations (e.g., prefixes) as the *number of nodes* which choose different next-hops for the two destinations.

While this general concept is a good starting point for studying routing changes, it has drawbacks. The main problem with RSD as defined in [7] is that it assumes that each node in the network has a unique next-hop towards any destination. In interdomain routing, where nodes are ASes and destinations are prefixes, this property does not hold. In [7], the authors decided to proceed as [10], i.e., they partitioned Autonomous Systems in such a way that each of its parts (denoted “quasi-routers”) did not have more than one routing option to reach any prefix. Although that was a natural choice, this solution has a number of drawbacks. First, it can not be easily generalized to other routing systems with multiple next-hops, such as OSPF with Equal Cost Multi-Path (ECMP). Second, identifying quasi-routers optimally is a NP-hard problem [5]. Third, and most important for our study, there is no natural way to extend the quasi-router approach to routing systems that change over time, which introduces serious problems for longitudinal analyses.

Hence, we need to define a new metric that handles the case of multiple next-hops in a cleaner, more robust fashion: MRSD. MRSD avoids the problems of RSD by allowing next-hops decisions to be expressed as sets (instead of unique elements). MRSD then uses *Jaccard distance* to compare the next-hop sets.

Formally, let $G(V, E)$ be a directed graph where V is a set of vertices and E is the set of edges. For all source-destination pairs $(u, v) \in V \times V$ we define $N_{u,v}$ as the set of *next-hops* from u towards v . More formally, a vertex $w \in N_{u,v}$ if and only if, the edge (u, w) starts a path from u to v .

DEFINITION 1. Let $G(V, E)$ be a graph and $S (\neq \emptyset)$ and P subsets of V denoting sets of sources and destinations respectively. For each pair $(d, d') \in P \times P$ we define the Multiple next-hop

¹Our source code is available by email request to the authors.

Routing State Distance (*MRSD*) over S by

$$MRSD(d, d') = \frac{\sum_{s \in S} \delta_s(d, d')}{|S|}, \quad (1)$$

where $\delta_s(d, d') = 1 - J(N_{s,d}, N_{s,d'})$ and $J(A, B)$ denotes the Jaccard Index of any pair of sets A and B .

The intuition behind MRSD is that when $MRSD(d, d')$ is close to 0 then d and d' are very similar in terms of the next-hop choices made by all nodes in the set S . On the other hand, when $MRSD$ approaches 1 we have that those destinations are routed very differently through the network. Note that MRSD is always a value between 0 and 1, and it can be interpreted as the *fraction* of next-hop decisions that differ, across all nodes in the network. It is also important to remark that when all destinations always have unique next-hops ($|N_{s,d}| = 1$ for all s and d), this definition is equivalent to the one presented in [7], and so in that case MRSD reduces to (normalized) RSD.

3.2 Temporal MRSD

Next we show how to apply the definition of the last section in the context of interdomain routing on the Internet in order to conduct a longitudinal analysis.

Consider P and S being respectively a set of IP prefixes $\{p_1, \dots, p_{|P|}\}$ and sources (autonomous systems) $\{s_1, \dots, s_{|S|}\}$ on the Internet seen by a set of monitors at times t_1, \dots, t_T ($t_i - t_{i-1} > 0$ for $i = 2, \dots, T$). We define $N_{s,p}(t)$ as the set of next-hops (autonomous systems) that source s may use at time t in order to reach prefix p , where $p \in P$, $s \in S$ and $t = t_i$ for some $i \in \{1, \dots, T\}$.

DEFINITION 2. For a prefix p we define the Temporal multiple next-hop Routing State Distance (*TRSD*) over a set of sources S between time t_i and t_j ($1 \leq i < j \leq T$) as:

$$TRSD_p(t_i, t_j) = \frac{\sum_{s \in S} \delta_{s,p}(t_i, t_j)}{|D_p(t_i, t_j)|}, \quad (2)$$

where $\delta_{s,p}(t_i, t_j) = 1 - J(N_{s,p}(t_i), N_{s,p}(t_j))$ if both $N_{s,p}(t_i)$ and $N_{s,p}(t_j)$ are non empty sets and zero otherwise, $D_p(t_i, t_j) = \{s \in S : N_{s,p}(t_i) \neq \emptyset \text{ and } N_{s,p}(t_j) \neq \emptyset\}$ and $J(A, B)$ denotes the Jaccard Index of any pair of sets A and B .

This definition is similar to the one presented in the last section, but has some key differences. The differences are: first, TRSD is specialized to the specific case of interdomain routing. Second, TRSD compares the next-hops of the same prefix in two different times instead of two different prefixes at the same time. And third, TRSD normalizes in a slightly different way, namely, the set S of sources becomes $D_p(t_i, t_j)$ which may vary with time. This was done by design, to address the fact that networks may grow over time.

To illustrate the definition, Figure 1 presents graphs related to the routing decisions towards a prefix p (hosted on AS x_p) at times t_i and t_{i+n} ($n > 0$). Table 1 shows the intermediate computations for TRSD over the set $S = \{a, b, c, d, e, f\}$.

From Table 1 we have that $|D_p(t_i, t_{i+n})| = 5$ and as consequence $TRSD_p(t_i, t_{i+n}) = 0.283$. This can be interpreted as saying that from time t_i to t_{i+n} , the network's next-hop decisions changed 28.3% with regard to p . It is important to remark that, in this example, the last row (related to the source f) did not contribute anything to TRSD (because it was not considered in the composition of the set $D_p(t_i, t_{i+n})$ and $\delta_{f,p}(t_i, t_{i+n}) = 0$). This shows an important aspect of our definition, i.e., TRSD was designed with the intention of capturing routing *changes*, and hence does not increase simply due to the *growth* of the network.

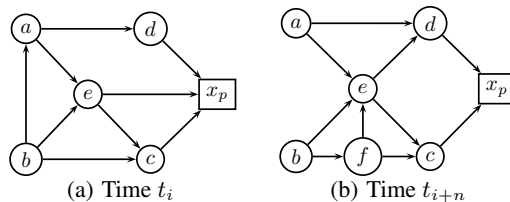


Figure 1: Routing decisions towards p at t_i and t_{i+n}

Table 1: TRSD computation for Figure 1

s	$N_{s,p}(t_i)$	$N_{s,p}(t_{i+n})$	$\delta_{s,p}(t_i, t_{i+n})$
a	$\{d, e\}$	$\{d, e\}$	0
b	$\{a, c, e\}$	$\{e, f\}$	$\frac{3}{4}$
c	$\{x_p\}$	$\{x_p\}$	0
d	$\{x_p\}$	$\{x_p\}$	0
e	$\{c, x_p\}$	$\{c, d\}$	$\frac{2}{3}$
f	$\{\}$	$\{c, e\}$	0

4. DATASET DESCRIPTION

To explore the evolution of Internet routing over time, we collected four datasets comprising Routing Information Bases (RIBs) from RIPE [1] and Route Views [2]. Datasets 1, 2 and 3 consist of all RIBs on a daily basis for the entire years of 2005, 2008 and 2011 respectively. The fourth data set consists of all RIBs for the first day of each month from 2005 to 2012. Since RIBs are made available at a coarser granularity than BGP updates (every 2 hours for Route Views and 8 hours for RIPE) we did not attempt to remove the effects of short term convergences, i.e., for each day, we kept all available distinct AS paths. It is important to remark that we collected data only for IPv4 prefixes.

Each RIB is a collection of records containing information about how to reach a prefix p from a specific autonomous system. From these records we extracted the following information: *route dumping date*, *autonomous system path* a_0, a_1, \dots, a_r ($r \geq 1$) and *destination prefix* p , hosted in the AS a_r . After that we decomposed each record into r 4-tuples of the form $[route\ dumping\ date, a_i, a_{i+1}, p]$, for $i = 0, \dots, r - 1$. The semantics of each 4-tuple is: at the time of *route dumping date*, in order to reach the prefix p , AS a_i uses AS a_{i+1} as (one of its) next-hops. In the rest of the paper, AS a_i will be referred to as a *source*, prefix p will be referred to as a *destination*, and AS a_{i+1} will be referred to as a *next-hop*.

To apply the definitions presented in Section 3 to this data, S is the set of all sources (a_i) and P is the set of all prefixes (p) found in a dataset. Table 2 presents a summary of each dataset collected after transformation into a collection of 4-tuples as described above.

Table 2: Summary of the four datasets

	Dataset 1	Dataset 2	Dataset 3	Dataset 4
Granularity	daily	daily	daily	monthly
First	01/01/05	01/01/08	01/01/11	01/01/05
Last	12/31/05	12/31/08	12/31/11	12/01/12
Size (GB)	200	500	680	160
$ S $	5,086	6,934	9,093	14,829
$ P $	316,519	517,773	616,714	1,157,670

To see the need for using Jaccard Index in Definitions 1 and 2, we note that, in our dataset on each day, approximately 7% of the sources have more than one next-hop choice towards a specific prefix.

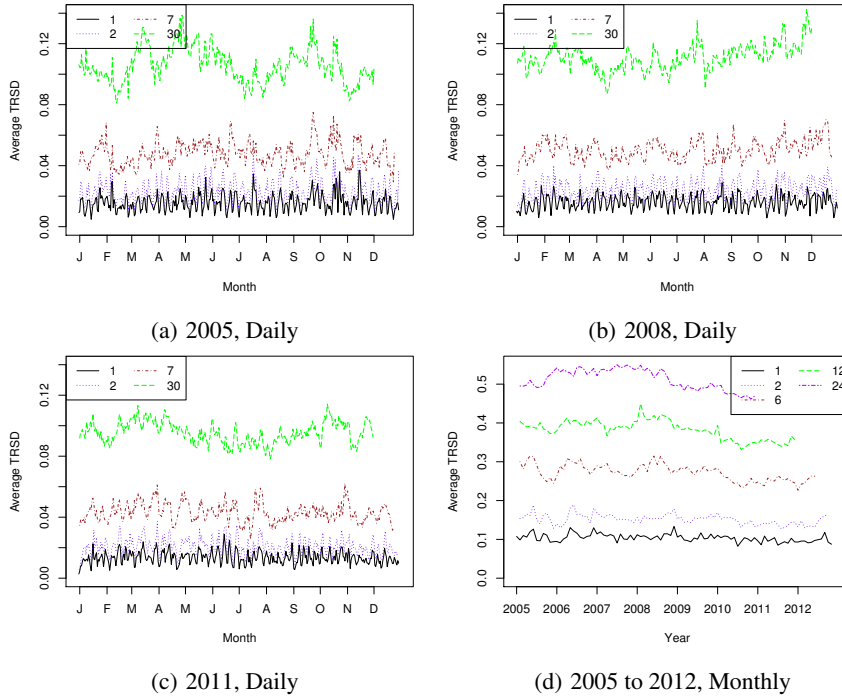


Figure 2: $r(i, n)$ versus i for different values of n

5. MEASUREMENTS

In this section we show how TRSD can be used to extract useful knowledge about interdomain routing dynamics.

5.1 Analyzing TRSD

To analyze routing dynamics in the four datasets we compute time series of TRSD, averaged over all prefixes. Specifically, let $r_p(i, n) = TRSD_p(t_i, t_{i+n})$ for $1 \leq i \leq T - n$ and some $1 \leq n \leq T - 1$. We then define $r(i, n)$ as the average of all $r_p(i, n)$ that can be computed, i.e., for all prefixes that have routing information at times t_i and t_{i+n} in our datasets. We use average of $r_p(i, n)$ because it gives a measure of the total magnitude of change. Thus, $r(i, n)$ represents average proportion of change in next-hop decisions at timescale n .

Figure 2 presents the time series of $r(i, n)$ for our 4 datasets, where n is 1, 2, 7 and 30 days for datasets 1 to 3 (daily), and n is 1, 2, 6, 12 and 24 months for dataset 4 (monthly). From now on, to simplify discussion, we assume that time indices always represent days for datasets 1, 2 and 3, and represent months for dataset 4.

The first striking aspect of these time series is that they are all approximately stationary. That is, they fluctuate around a mean value, but do not show any long-term trend. Furthermore, the mean values of TRSD do not show significant differences across years, from 2005 to 2012. This indicates that despite the considerable growth of the Internet in terms of ASes and Prefixes (which might suggest more next-hop options per AS) there is an approximately constant rate of routing decision changes over time. For example, for $n = 30$ days in the daily datasets we have that $r(i, n)$ is approximately 0.1, meaning that from month to month, on average, 10% of the Internet changes in terms of next-hop routing decisions. The same value can be seen in the monthly dataset for the curve $n = 1$ month.

The next observation is that as n grows so does $r(i, n)$. This implies that at least a portion of the routing changes that happen over time are persistent, i.e., are not undone quickly so as to return to a previous routing state. On the other hand, we cannot say that it is a system governed by system-wide changes. To see that we can refer to the curve $r(i, 24)$ of the monthly dataset. From this curve we can see that in a time window of 2 years, approximately 50% of the next-hop choices persisted (because 50% changed). This fact indicates many routes are stable over a long period of time.

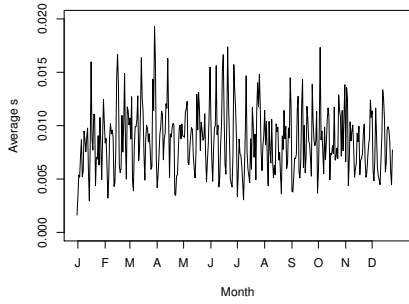
Yet another property that can be seen in the daily datasets is related to seasonality. One can note that for $n = 1$ day there is pattern of weekly variation in the data. This may be explainable as evidence of human input (network operators) in the system. We explore this conjecture in the following section.

5.2 Analyzing Routing Changes

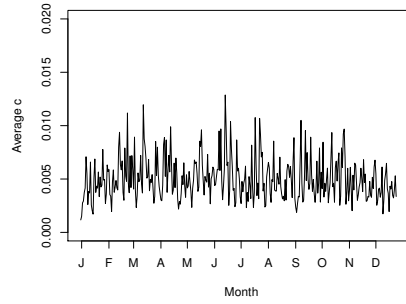
After analyzing the results of the last section one question that arises is: of the total set of next-hop changes that happen over the time, what portion is related to *sustained* (changes that persist in the system for some time) and which fraction is related to *churn* (changes that change again in the near future)?

To answer this question we decompose the TRSD time series ($r(i, 1)$) into two components: $s(i, k)$ and $c(i, k)$, standing for *sustained* TRSD and *churn* TRSD respectively. Sustained TRSD ($s(i, k)$) captures the portion of $r(i, 1)$ that represents next-hop decisions that changed from t_i to t_{i+1} but did *not* subsequently change from t_{i+1} to t_{i+k} . In the same vein we define $c(i, k)$ as the portion of $r(i, 1)$ that represents next-hop decisions that changed from t_i to t_{i+1} and then changed *again* from t_{i+1} to t_{i+k} .

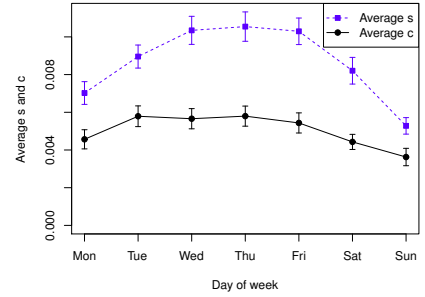
Formally, we define $s(i, k)$ as the average over all $s_p(i, k)$, where $s_p(i, k)$ is a measure for sustained changes for a specific



(a) $s(i, k)$, $k = 7$ days



(b) $c(i, k)$, $k = 7$ days



(c) weekly average of $s(i, 7)$ and $c(i, 7)$. Error bars are 95% Confidence Intervals

Figure 3: $s(i, k)$ and $c(i, k)$ for 2011, daily

prefix p given by:

$$s_p(i, k) = \frac{\sum_{s \in S} \delta'_{s,p}(i, k)}{|D'_p(i, k)|}, \quad (3)$$

where $D'_p(i, k)$ is the set of all sources for which the sets $N_{s,p}(t_i)$, $N_{s,p}(t_{i+1})$ and $N_{s,p}(t_{i+k})$ are not empty and $\delta'_{s,p}(i, k)$ is the fraction of elements, which belong to the set $N_{s,p}(t_i) \cup N_{s,p}(t_{i+1})$, that satisfy any of the two following conditions: *i*) the element is in $N_{s,p}(t_i)$ but it is not in $N_{s,p}(t_{i+1})$ nor in $N_{s,p}(t_{i+k})$; or *ii*) it is in $N_{s,p}(t_{i+1})$ and $N_{s,p}(t_{i+k})$ but it is not in $N_{s,p}(t_i)$.

Then proceeding analogously to Equation (3) we can compute $c(i, k)$. One important observation is that $s_p(i, k)$ and $c_p(i, k)$, by definition, form a partition of $r_p(i, 1)$, i.e., $s_p(i, k) + c_p(i, k) = r_p(i, 1)$, for all p and k that have available routing information at times t_i , t_{i+1} and t_{i+k} .

Figure 3 presents $s(i, 7)$ and $c(i, 7)$ for the daily dataset of 2011. (Results for 2005 and 2008 are similar). The choice of $k = 7$ means that routing changes that persist for a week are considered to be sustained. The first observation, obtained by comparing Figures 3(a) and 3(b), is that sustained routing changes ($s(i, 7)$) have more impact on TRSD than $c(i, 7)$. The average value of $s(i, 7)$ is around 0.007 (0.7%), while $c(i, 7)$ has an average value of approximately 0.004 (0.4%).

The next observation is that $s(i, k)$ inherits the weekly periodicity of $r(i, 1)$ while $c(i, k)$ is more similar to noise. This weekly periodicity suggests that there is some sort of human interaction with the system. Our conjecture is that those changes are triggered by BGP policy management and resulting changes to BGP configurations. In order to dig deeper we computed the average of $c(i, k)$ and $s(i, k)$ over days of the week. Figure 3(c) presents the results, where we can see that both curves indicate more activity in the work week and less on the weekends. However, the difference between week and weekend behavior is much more significant in the case of the sustained changes, i.e., $s(i, k)$. These results corroborate with the idea of human interaction with the system, since we can understand them as the ASes avoiding intentional routing changes on weekends (probably due to reduced number of network operators working full time).

The decomposition into sustained TRSD and churn TRSD is also informative on long timescales. Figure 4 presents $s(i, 2)$ and $c(i, 2)$ for dataset 4 (2005 to 2012, monthly). Here we adopt the assumption that routing changes that persist for one month ($n = 1$) are considered to be sustained. This Figure shows that on average, about 3% of all routing decisions churn on a monthly basis, while about 8% of all routing decisions show a sustained change. Comparing Figures 4 and 3, one can see that the sustained and churn

TRSD on a monthly basis are greater than the corresponding values on a weekly basis, which is to be expected. However, the difference between sustained and churn TRSD is much greater on the monthly timescale, showing that sustained TRSD captures the accumulation of intentional changes to the routing system over time, while churn TRSD reflects the continuous background noise in the system.

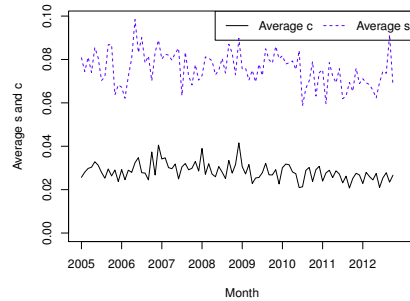


Figure 4: $s(i, k)$ and $c(i, k)$, monthly dataset with $k = 2$

The relative stability of $s(i, 2)$ in Figure 4 shows that large-scale, system-wide changes to Internet routing are rare. There are only a couple of points in time where there are noticeable peaks in the amount of sustained change in Internet routing – one peak in mid 2006 and one in late 2012. Initial investigation of these events shows evidence that system-wide routing changes took place: a large fraction of all prefixes were affected by next-hop changes during these events. A deeper investigation of these events is ongoing.

5.3 Contribution of Sources

Our final set of results aims at answering the following question: *what is the contribution of each source AS to TRSD?* in other words, are the changes represented by TRSD uniformly spread over all ASes or are they concentrated in a small set? To start this analysis we sampled randomly chosen days in 2011 and computed the contribution of each AS over the total TRSD of that day. Figure 5(a) presents the Complementary Cumulative Distribution Function (CCDF) of this contribution for the first day of 2011 (results for other days are similar). The figure shows that the contribution of ASes is long-tailed, and that most ASes, approximately 90%, make little or no contribution to TRSD.

In order to inspect whether this is common behavior over time, for the entire year of 2011 we compute the fraction of ASes that

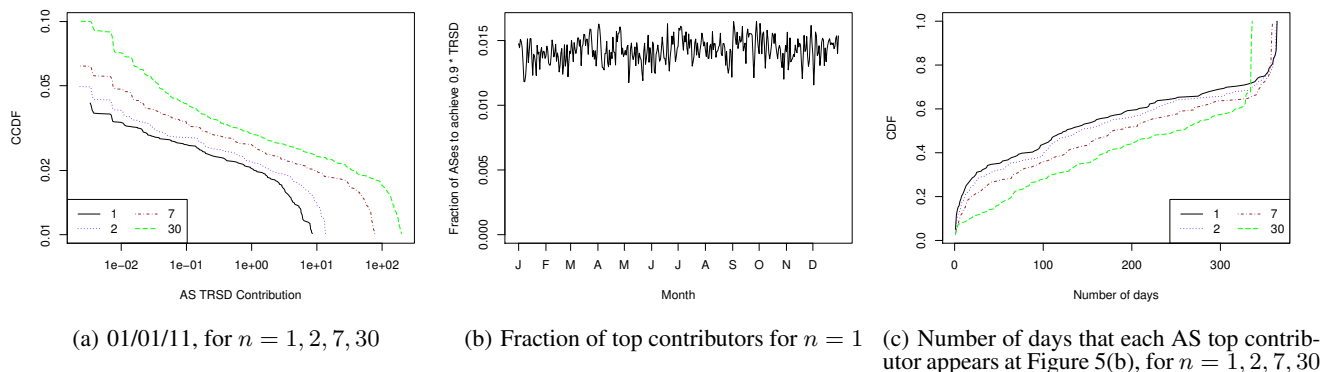


Figure 5: AS Contribution, 2011

are necessary to achieve 90% of TRSD. The result is presented in Figure 5(b), where we can see that less than 2% of all ASes are in fact necessary to capture 90% of TRSD. This result holds for the other datasets with a threshold of 2.5%.

A question that may naturally arise then is: *are these heavy-hitter ASes the same over the time?* In order to answer it we counted in how many days each AS appears during the computation of Figure 5(b), in other words, in how many days each AS was among the set of sources responsible for 90% of the daily TRSD over 2011. The Cumulative Distribution Function (CDF) of this quantity is presented in Figure 5(c). We can see that approximately 40% of ASes that appear on at least one day, appear in at least 300 days. Further, there is a group of approximately 20% of ASes that appear in all days. An initial investigation shows that many of them are in or are near to the network core.

6. CONCLUSIONS AND FUTURE WORK

In this work we proposed a new way to study the Internet at the AS level. We introduced a new measure, TRSD, allowing us to characterize rate of change in the interdomain routing system over long periods of time.

Using this measure we are able to uncover several interesting aspects of the global Internet routing system. First, we showed that the rate of change in routing decisions has been stable over time, despite the growth in the network overall. Second, we showed how to decompose TRSD into components that reflect sustained change versus churn. We showed that the rate of sustained changes shows a persistent weekly periodicity suggestive of a tendency of operators to make sustained (intentional) routing changes in the workweek. Finally, we showed that the locations of the routing changes in the Internet are generally concentrated among a small set of ASes, often those that are near or in the core. Throughout the entire study we consider various timescales and show how the magnitude of routing change (both sustained and churn) varies with timescales from days to months.

Although these are interesting results that characterize the Internet in a new way, it is important to note that this is a work in progress. In particular, there are several interesting observations that need further study in order to understand their implications. Moreover, the results presented here suggest several directions for future research: *i)* it may be valuable to group prefixes according to their TRSD similarity, as prefixes whose routing changes in similar ways may form important clusters; *ii)* TRSD may be usefully applied to individual prefixes, i.e., in a microscopic analysis as compared to the macroscopic view in this study; *iii)* it may be infor-

mative explore spatial and temporal dimensions at the same time; *iv)* TRSD may help understanding types of routing changes and how an AS's routing changes affect its neighbors; *v)* TRSD may be usefully applied in analyzing router-level topologies. Although a router-level study is a challenge due to the lack of publicly available data, an analysis at this granularity would be interesting because the same AS path may have traffic flowing through different physical paths. Finally, our long-term goal is to apply the methods of this paper to identify and understand major structural changes or reorganizations of the Internet. We hope to create a methodology that identify "interesting" events in the TRSD time series and find the responsible sources and prefixes associated with those events.

7. ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant Numbers CNS-0905565, CNS-1018266, CNS-1012910, and CNS-1117039, and supported by the Army Research Office under grant W911NF-11-1-0227. We are also grateful to the anonymous reviewers and to John Heidemann for their valuable comments.

8. REFERENCES

- [1] Ripe's routing information service raw data project. <http://www.ripe.net/data-tools/stats/ris/ris-raw-data>.
- [2] University of oregon route views project. <http://www.routeviews.org>.
- [3] Amogh Dhamdhere and Constantine Dovrolis. Ten years in the evolution of the internet ecosystem. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement, IMC '08*, pages 183–196, New York, NY, USA, 2008. ACM.
- [4] A. Elmokashfi, A. Kvalbein, and C. Dovrolis. Bgp churn evolution: A perspective from the core. *Networking, IEEE/ACM Transactions on*, 20(2):571–584, April.
- [5] Gonca Gürsun. *Inferring Hidden Features in the Internet*. PhD thesis, Boston University, 2013.
- [6] Gonca Gürsun, Natali Ruchansky, Evimaria Terzi, and Mark Crovella. Inferring visibility: who's (not) talking to whom? In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication, SIGCOMM '12*, pages 151–162, New York, NY, USA, 2012. ACM.
- [7] Gonca Gürsun, Natali Ruchansky, Evimaria Terzi, and Mark Crovella. Routing state distance: a path-based metric for

- network analysis. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, IMC '12, pages 239–252, New York, NY, USA, 2012. ACM.
- [8] Craig Labovitz, G. Robert Malan, and Farnam Jahanian. Internet routing instability. *IEEE/ACM Trans. Netw.*, 6(5):515–528, October 1998.
- [9] Jun Li, Michael Guidero, Zhen Wu, Eric Purpus, and Toby Ehrenkranz. Bgp routing dynamics revisited. *SIGCOMM Comput. Commun. Rev.*, 37(2):5–16, March 2007.
- [10] Wolfgang Mühlbauer, Steve Uhlig, Bingjie Fu, Mickael Meulle, and Olaf Maennel. In search for an appropriate granularity to model routing policies. In *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '07, pages 145–156, New York, NY, USA, 2007. ACM.
- [11] Dimitri Papadimitriou, Florin Coras, and Albert Cabellos. Path-vector routing stability analysis. *SIGMETRICS Perform. Eval. Rev.*, 39(3):22–24, December 2011.
- [12] Jennifer Rexford, Jia Wang, Zhen Xiao, and Yin Zhang. Bgp routing stability of popular destinations. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement*, IMW '02, pages 197–202, New York, NY, USA, 2002. ACM.
- [13] Yang Xiang, Xia Yin, Zhiliang Wang, and Jianping Wu. Internet flattening: Monitoring and analysis of inter-domain routing. In *ICC*, pages 1–6, 2011.