

On the Impact of Filters on Analyzing Prefix Reachability in the Internet

Ravish Khosla
School of Electrical
and Computer Engineering
Purdue University
West Lafayette, IN 47907
Email: rkhosla@purdue.edu

Sonia Fahmy
Department of Computer Science
Purdue University
West Lafayette, IN 47907
Email: fahmy@cs.purdue.edu

Y. Charlie Hu
School of Electrical
and Computer Engineering
Purdue University
West Lafayette, IN 47907
Email: ychu@purdue.edu

Abstract—The reachability of IP address prefixes exhibits significant fluctuations due to changes in both physical connectivity and ISP routing policies. In the late 1990s, Labovitz *et al.* performed an extensive study of inter-domain path stability by analyzing BGP routing data. To reduce the noise in the BGP data, e.g., transient updates during route convergence, they applied several filters to preprocess the raw BGP data.

In this work, we investigate prefix reachability as advertised by BGP, while revisiting the preprocessing filter design problem. We show that the reachability analysis results are highly sensitive to the specific filters applied and the parameters that control the strength of the filters. In particular, we compute the Mean Time to Failure and Recovery (MTTF and MTTR) as well as the up-to-downtime ratios of prefixes, and find that these can fluctuate by a factor of 10 by varying the filter parameters. We analyze the impact of recent fiber cuts in the Mediterranean sea and the Middle East, and study prefix reachability during a nine-month period in 2007 to evaluate the general health of the Internet¹.

Index Terms—BGP, prefix reachability, prefix availability

I. INTRODUCTION

While universal Internet reachability is taken for granted, failures and attacks may lead to a significant loss in connectivity across continents. Recent incidents illustrating this problem include the Taiwan Earthquake of 2006 [1], and the undersea cable cuts in 2008 [2], [3] which caused significant disruptions and increase in web latencies to much of the Middle East, Asia, and North Africa [4], [5]. Understanding how inter-domain routing via the Border Gateway Protocol (BGP) handles failures is challenging because of the immense scale of the Internet, and BGP flexibility, for instance in implementing policies based on business agreements.

In the late 1990s, Labovitz *et al.* [6], [7] performed an extensive study of Internet routing instability by analyzing BGP routing data. To reduce the effect of noise, they applied several filters to preprocess the raw BGP data, where the term “filter” denotes data processing operations that remove biased data points. They quantified *reachability*, defined as the existence of any path to a prefix, and examined how prefix reachability changes over time, computing the Mean Time to Failure (MTTF), i.e., mean time to loss of reachability of a

prefix, and the Mean Time to Recovery (MTTR), defined as the mean time until a prefix is reachable.

In this work, we revisit the problem of prefix reachability using BGP data collected from RouteViews [8]. Labovitz *et al.* [6] had used filters with parameter values chosen either without in-depth investigation, or based on measurements like convergence time which can change over time based on the BGP implementations deployed. Our contributions include: (1) systematically studying the impact of the filter parameters used in data preprocessing on the MTTF, MTTR, and prefix availability, (2) comparing Internet “health” results obtained in [6] to recent data, and (3) investigating prefix reachability during the 2008 cable cuts in the Mediterranean sea and the Middle East. We find that while the filters are useful in studying the general health of the Internet, one has to judiciously use them based on the goal of the measurement study. We also find that today’s Internet is healthier than that of the late 1990s, confirming the conclusions in [9]. Finally, our results for the 2008 cable cuts demonstrate that the impact of transient stress events can only be seen with minimal filtering.

The remainder of this paper is organized as follows. Section II summarizes work related to our study, and Section III describes our datasets. Section IV explains how we compute MTTF and MTTR. In Section V, we discuss the design of the filters. Section VI presents our results and observations. Section VII concludes the paper.

II. RELATED WORK

In this section, we briefly review previous work closely related to prefix stability and reachability.

Labovitz *et al.* studied Internet instability in the late 1990s [6], [7]. In [6], the authors compute the MTTF and MTTR for both inter-domain and intra-domain paths. They find the Internet to be remarkably robust with inter-domain path MTTF close to 10 days and MTTR around 10 minutes, and with intra-domain paths performing even better. Several *filters* are applied to the dataset before computing these values. These include choosing prefixes that were present in the routing table for more than 60% of a 9-month period, counting multiple failures within a 15-minute window as a single failure, and excluding the loss of a peering session as a

¹This research is sponsored in part by NSF grants 0238294 and 0831353.

source of failure. We explore how the filter parameters affect the results, while repeating the reachability study for today's Internet. Recently, Li *et al.* [9] showed that BGP dynamics are busier but healthier, compared to Labovitz's study. They classified BGP updates while studying temporal characteristics and a peer's contribution to the dynamics. Their work does not follow-up on the availability of prefixes, which we address in this paper. Closely related to our work, in [10], the authors studied the reachability of a particular set of prefixes (DoD prefixes). They found global unreachability (when none of the peers can reach a prefix) durations to be fairly high with 17% of unreachability durations being longer than one hour. Among data-plane-based monitoring systems, Hubble [11] continuously monitors reachability of prefixes in the Internet using many probing end hosts.

Several recent studies have focused on the impact of BGP instability. The authors of [12] designed an online system that not only identifies significant routing disruptions (by correlating updates), but also estimates their impact on the flow of traffic. Wang *et al.* [13] injected routing changes on the PlanetLab testbed and studied changes in the end-to-end performance using the metrics of packet loss, delay, and reordering. A related problem (addressed in [14]) is how instabilities in the control plane affect the data plane. Their measurement study monitors traffic towards a sink behind a BGP beacon with the beacon going up and down at regular intervals. Finally, Wu *et al.* [15] studied the impact of large-scale failures, e.g., the Taiwan earthquake, under a general umbrella of analyzing routing resilience, and Renesys researchers studied the effects of the Taiwan earthquake and the Middle East and Mediterranean cable cuts, analyzing the number of prefixes that were down or unstable [16].

Several studies aim at finding the root cause of BGP dynamics [17], [18]. Several papers cluster updates into events [19], [20], [21] by grouping updates within a certain time threshold. Our work is orthogonal to these studies as we do not study the root cause of routing events but rather their effect on prefix reachability.

III. DATA SETS

We use RouteViews [8] as our source for BGP routing tables and updates. Specifically, our data was obtained from routeviews2.oregon-ix.net which contains the Routing Information Bases (RIBs) and Updates in MRT format. We used this source as it was the largest data source in the MRT format, with data dating back to October 2001. We used the RIBs and update files for a period of 9 months from March to November 2007. These months were chosen to represent typical Internet health since, to our knowledge, no major stress events were reported during this period. The period length of 9 months was selected to match Labovitz's study [6], which enables a direct comparison of the results.

We also used selected periods when some known events occurred. We used January and February 2008, as they were associated with Middle East and Mediterranean cable cuts [2] and a YouTube prefix hijack [16]. The cable cuts occurred

between January 30 and February 2, 2008 and most of the repairs were completed within 14 days [16]. The YouTube prefix hijack happened on February 24, 2008 for a period of approximately 2 hours.

A. Preprocessing RouteViews Data

The RIBs and updates available are in .bz2 format with typical sizes of 0.8 GB per day of RIB files (sampled every 2 hours) and about 25 MB per day of update files (written every 15 minutes), which total about 25 GB per month of data. We preprocess the data, using libbgpdump version 1.4.99.7 [22] to convert the files from the MRT format to text. We reduce the storage space required by removing unused fields. We only keep the timestamp, peer IP, prefix, and the type of update (Announcement or Withdrawal). After preprocessing and filtering of table transfers (described in Section III-B), we have about 13-15 GB of gzipped RIB and update files per month of data.

B. Filtering Table Transfers

One problem with using the raw updates from RouteViews is that they also include routing table transfers which are caused by session resets between a monitor and a peer [23]. These spurious updates are an artifact of the update collection methodology and should not influence the state of a prefix. Zhang *et al.* [23] developed the Minimum Collection Time (MCT) algorithm to identify BGP routing table transfers.

We executed scripts (with default parameters as in [23]) kindly contributed by the authors from the point of view of every peer in our dataset. We define a peer as any peer which is present in any routing table entry and at least one update. This definition yields 45-47 peers in our dataset. The scripts report tuples of the form (PeerIP, Starting Time, Duration) which identify a table transfer as observed by the peer specified by its IP, starting time, and duration. We developed a script that uses this information to remove the table transfer updates from the update files obtained from RouteViews. We use these filtered updates for all later processing.

IV. MTTF AND MTTR COMPUTATION

We define a *combination* as a (peer, prefix) tuple which implies that the prefix was observed by the peer. The MTTF and MTTR are computed for each combination that we are interested in using the following methodology. We maintain the state of each combination at each point in time, and at the time of each state change, we record a downtime or an uptime. If the state of a combination changes from Announced (A) to Withdrawn (W), an uptime is recorded whereas a change from W to A leads to the recording of a downtime.

With this methodology, we encounter the problem of boundary effects. Since the updates only show incremental changes, we do not know the state of a combination till we observe the first update for that combination. For stable combinations, this may be a significant time period. To address this problem, we treat every entry of the first routing table of our data period as an announcement. This enables us to initialize the

state of each combination. As the state changes, we continue recording uptimes and downtimes, while updating the state of the combination and the time of the state change. After processing all the update files, we add an extra up or downtime depending upon the last state of the combination. For example, if the last state change was to W and was reported at time t_1 , and if the data period ended at time t_2 , we add an extra downtime with value $t_2 - t_1$. Then, we compute the mean, median, standard deviation, and coefficient of variation of the recorded uptimes and downtimes. The mean uptime is the MTTF and the mean downtime is the MTTR. For each combination, we also compute the ratio of cumulative uptime and cumulative downtime and report the statistics of this up-to-downtime ratio. The ratio can be infinite if the downtime is zero.

Observe that our approach of measuring the uptime and downtime of a prefix is similar in principle to Oliveira *et al.*'s work [24], which studies the evolution of Internet topology by computing births and deaths of ASes and inter-AS links.

V. FILTER DESIGN

In this section, we devise one new filter, and revisit the design of filters in [6]. We use both routing tables and BGP update messages to infer the existence of paths to a prefix.

A. Observed Combinations Filter

We design a filter, the *observed combinations* filter, to remove bias introduced by combinations for which we have insufficient data, because we observed them towards the end of our dataset. This eliminates *boundary effects*. The filter outputs combinations which either exist in the first routing table of the period under consideration, or are present in the first few update files of the period. The percentage of update files at the beginning of the period that we check is a parameter α of this filter. The choice of α is explored in Section VI, along with the choice of parameters used in other filters. Small values of α bias the results towards relatively stable combinations, since most of the output combinations will exist in the first routing table. Since, at this point, we have not employed filtering other than the *observed combinations* filter, we refer to results on this data as the *initial filtering* results.

B. Stable Prefixes Filter

We implement a filter which we refer to as *stable prefixes* filter, based on the work by Labovitz *et al.* [6]. This filter outputs the prefixes that are available in the routing table for more than an aggregate percentage of time. We refer to this percentage parameter as β . Labovitz *et al.* use 60% as the value of β for a 9-month period. The goal of this filter is to eliminate transient prefixes which can go up and down frequently, confounding the results of stable prefixes.

C. Route Convergence Filter

We implement a filter that we refer to as the *route convergence* filter, which considers multiple failures within a time window as a single failure. The motivation behind this filter

is to remove the updates generated while the routes are still converging. During such convergence periods, it is usual to have a route withdrawn to a prefix only to be followed by an announcement of a different path to that prefix. Such a change does not truly reflect a failure. The filter operates by deleting all pairs of withdrawals followed by announcements (for the same (peer, prefix) combination) if the time between the withdrawal and announcement is less than a time window.

We call this time window parameter γ , whose value was chosen to be 15 minutes in [6]. The effect of increasing γ is to increase the value of both MTTF and MTTR, because as withdrawals are removed, uptimes are increased, thereby causing an increase in MTTF. An increase in γ also increases the value of MTTR because shorter downtimes are eliminated as no recorded downtime can be less than γ seconds.

D. Summary

Table I summarizes the filters. If we define a “healthy” Internet as one with high up-to-downtime ratios, the filters impact the computed “health of the Internet” in different ways. The *stable prefixes* filter removes unstable prefixes causing the Internet to appear healthier than if these prefixes with low uptimes and high downtimes had not been removed. The *route convergence* filter increases the uptimes of certain combinations but at the same time, it increases the downtimes. The impact on the perceived health of the Internet is unlikely to be as significant as the *stable prefixes* filter. The *observed combinations* filter removes the combinations first seen in the latter part of our data period, so the Internet appears healthier by applying this filter.

VI. RESULTS AND DISCUSSION

In this section, we quantify the impact of the filter parameters on the MTTF and MTTR values.

A. Impact of Filter Parameters

We first study the dependence of MTTF and MTTR on filter parameter choices. We use January 2005 and January 2007 as two representative months of data for investigating the impact of filter parameters, and use the control variable technique to select parameter values to use in our later experiments.

1) *Observed Combinations Filter Parameter*: We implement the *observed combinations* filter with $\alpha = 10\%$, 25% and 50% and compare the MTTF and MTTR obtained with the one obtained when $\alpha=100\%$, i.e., when all the combinations are used. The MTTF and MTTR are computed for twenty typical cases, which include the *initial filtering* case, the *stable prefixes* filter with $\beta=0\%$, 30% , 60% and 90% , the *route convergence* filter with $\gamma = 120$ s, 300 s and 900 s, and with the two filters put together using the twelve permutations of parameters. We investigate the value of α which showed, on the average, the largest increase in uptime and the largest decrease in downtime w.r.t. the 100% case. This is because this filter aims to eliminate boundary effects and the new combinations advertised close to the end of the observation

TABLE I
SUMMARY OF FILTERS AND CORRELATION OF THEIR PARAMETERS WITH THE MTTF/MTTR MEASURES

Filter	Description	Impact on MTTF	Impact on MTTR	Value in [6]	“Healthy” Values
Observed Combinations	Choose the combinations that exist in the initial $\alpha\%$ of update files in our dataset	$\alpha \uparrow \Rightarrow MTTF \downarrow$	$\alpha \uparrow \Rightarrow MTTR \uparrow$	Not Used	10%
Stable Prefixes	Choose the prefixes that are present for more than $\beta\%$ of time in routing tables	$\beta \uparrow \Rightarrow MTTF \uparrow$	$\beta \uparrow \Rightarrow MTTR \downarrow$	60%	99.5%
Route Convergence	Count multiple failures that happen within γ seconds as a single failure	$\gamma \uparrow \Rightarrow MTTF \uparrow$	$\gamma \uparrow \Rightarrow MTTR \uparrow$	900 s	200 s

period are typically less stable, thereby having smaller uptimes and larger downtimes.

The results for Jan. 2007 are shown in Table II. As expected, for lower values of α , the MTTF increase is higher and the MTTR decrease is larger so that overall, the difference between the MTTF increase and MTTR decrease is the highest. We do not reduce the value of α below 10%, since we believe that ignoring more than 90% of the update files (more than 27 days in a month’s period) is eliminating too much data. For Jan. 2005, the results were similar, except that the MTTF–MTTR increase for $\alpha=10\%$ is 21.1%, vs. 14.37% for $\alpha=25\%$. With 10% as the choice of the parameter, we have about 9 million combinations or about 91% of the total combinations.

TABLE II
OBSERVED COMBINATIONS FILTER PARAMETER α FOR JAN. 07

α (%)	MTTF Change (%)	MTTR Change (%)	MTTF–MTTR Change (%)
10	5.06	-26.30	31.36
25	3.98	-26.42	30.40
50	1.12	-2.52	3.64

2) *Stable Prefixes Filter Parameter*: If we aim to study healthy prefixes, we should select a value of β that maximizes the difference between the MTTF and MTTR while ensuring some minimum percentage of prefixes are available. We use an *observed combinations* filter parameter $\alpha=10\%$ as the control variable and vary β from 0% to 98% for Jan. 2007. We do not choose 100% because, for longer dataset durations (a month or more), we found that there are no prefixes that exist in all the routing tables. Since routing tables are recorded every two hours, it is likely that even a relatively stable prefix may be transiently withdrawn and may not appear in the table at the time when it is recorded. It is also interesting to note that we discovered that the results for $\beta=0\%$ were not the same as the ones for *initial filtering*. We traced this to the fact that there are some prefixes which appear in update messages but never manifest themselves in routing tables, excluding them from results with the *stable prefixes* filter, but not from the *initial filtering* case. We believe that these prefixes (henceforth referred to as “hidden prefixes”) are transient (perhaps used for malicious purposes), and are announced and withdrawn very frequently. This discovery gives further credibility to this filter.

The actual numbers of stable prefixes for various values of

β for Jan. 2005 and 2007 are listed in Table III. It is worth noting that we have excluded the hidden prefixes from the second and fourth columns of these tables; hence, the result for $\beta=0\%$ is the baseline for 100% of the prefixes. Further, the number of prefixes visible in our data is about 25% lower for Jan. 05 compared to Jan 07. This is consistent with the reported growth of the Internet and the addition of eight extra peers (about 20%) to RouteViews during this two year period.

TABLE III
STABLE PREFIXES FILTER PARAMETER β FOR JAN. 07 AND JAN. 05

β (%)	Jan. 07		Jan. 05	
	Output Prefixes	% of Total Prefixes	Output Prefixes	% of Total Prefixes
0	233537	100	180229	100
30	225685	96.64	171289	95.04
60	221620	94.89	166447	92.35
90	217607	93.18	162619	90.23
98	212883	91.16	145633	80.8

As we increase the value of β , we compute the changes in MTTF and MTTR compared to the *initial filtering* case. The results, listed in Table IV, show that as the value of β increases, the MTTF increases and the MTTR decreases, which is a characteristic of stable prefixes. Similar results are also observed for Jan. 05, where the MTTF–MTTR increase varies from 5% to about 63% for the same values of β as in Table IV.

In subsequent results, we will use $\beta=0\%$ (to eliminate the hidden prefixes), 60% (to compare with Labovitz *et al.*’s [6] results) and a high value like 99.5% or 100% depending upon whether $\beta=100\%$ yields a significant number of prefixes or not. The work in [6] does not explain the reasoning behind the choice of $\beta=60\%$; we experiment with the same value, although in most of our results, the number of prefixes filtered is less than the 20% reported in [6]. This is expected since their study was conducted in the late 1990s, and the Internet health has improved [9], leading to a higher percentage of stable prefixes.

3) *Route Convergence Filter Parameter*: The results of the effect of the filter parameter γ on MTTF and MTTR w.r.t. the *initial filtering* case for Jan. 2007 are given in Table V. The table columns are the same as previous tables for ease of comparison. As can be seen, the MTTF–MTTR decreases with the increase in γ . Hence, increasing the parameter is reducing the perceived “health” of the Internet. This filter is useful as it reduces the impact of transient updates on

TABLE IV
STABLE PREFIXES FILTER PARAMETER β FOR JAN. 07

β (%)	MTTF Change (%)	MTTR Change (%)	MTTF-MTTR Change (%)
0	0.48	-7.61	8.09
30	2.44	-29.88	32.32
60	5.12	-41.07	46.19
90	8.55	-46.30	54.85
98	20.03	-41.5	61.53

perceived failures. Thus, a tradeoff exists between reducing this impact and adversely impacting perceived Internet health. We believe that the minimum value of γ that is indicative of route convergence should be used. Labovitz *et al.* [6] specified γ to be 15 minutes, so we use 900 seconds as one of the values of γ in our runs; however, this is longer than the time required for convergence in today's Internet: Burkle [25] points out that most BGP convergence happens within 2 minutes. In fact, Labovitz in one of his talks [26] reports that the convergence time for announcements is about 60 seconds and for withdrawals, the median convergence time is about 100-180 seconds, and equals 15 minutes only in the worst case. Hence, we use 200 s and 300 s as additional values for γ .

TABLE V
ROUTE CONVERGENCE FILTER PARAMETER γ FOR JAN. 07

γ (s)	MTTF Change (%)	MTTR Change (%)	MTTF-MTTR Change (%)
60	59.56	102.47	-42.91
120	81.76	156.27	-74.51
200	97.8	203.24	-105.44
300	111.32	249.66	-138.34
900	149.14	429.61	-280.46
1500	167.59	560.27	-392.68

B. Results for a Nine-Month Period

In this section, we investigate the effect of the filters on 9 months of data from March through Nov. 2007. We applied the *observed combinations* filter with $\alpha=10\%$ and found that the total number of combinations exceeds those for Jan. 2007, which is expected since we have a longer period of data. However, the percentage of observed combinations that were output by this filter was far less (61.3% of total 16,654,558) than the percentage output for $\alpha=10\%$ for Jan. 2007 (91.5% of total 9,862,508), which is also explained by the longer remaining (90%) observation period.

Table VII shows the changes in median and mean uptimes and downtimes as parameter β of the *stable prefixes* filter is changed. There are no prefixes that are available in every routing table ($\beta=100\%$) so we use 99.5% as the highest value of β (Table VI). Table VII indicates that as the *stable prefixes* parameter β increases, the MTTF increases by as much as 51% and the MTTR falls by up to 73%. The trend is easily explained by the fact that the stable prefixes have higher average uptimes and lower average downtimes. The advantage of using the median is that it is not affected by

outliers like the mean. The median uptime increases by a much larger percentage (up to 406%) than the mean uptime when β increases. However, the median downtime also increases slightly with β . This increase is not unusual since the mean is much higher than the median in this case. Hence, a fall in the mean does not necessarily mean a fall in the median.

TABLE VI
STABLE PREFIXES FILTER PARAMETER β FOR MAR.-NOV. 07

β (%)	Output Prefixes	% of Total Prefixes
0	341122	100
60	227291	66.63
99.5	180489	52.91

TABLE VII
STABLE PREFIXES FILTER FOR MAR.-NOV. 07. INITIAL FILTERING MTTF = 827989 s, MTTR = 125479 s, MEDIAN UPTIME = 962 s, MEDIAN DOWNTIME = 64 s. ALL % CHANGES ARE RELATIVE TO INITIAL FILTERING

β (%)	MTTF Change (%)	MTTR Change (%)	Median Uptime Change (%)	Median Downtime Change (%)
0	0.25	-11.75	0.93	-1.56
60	10.86	-71.33	71.62	1.56
99.5	51.52	-72.61	406.76	17.19

Combining the results of Table VI and Table VII gives us valuable insights into the stability of prefixes, which is indicative of the general health of the Internet. An average prefix in the Internet has an MTTF of about 9.5 days and an MTTR of about 1.5 days. However, these observations are heavily affected by outliers since the median uptime is about 16 minutes and the median downtime is about a minute. Hence, while failures happen fairly frequently for a typical prefix, the recovery is quick as well. There is also a significant number of prefixes that are down for days at a time (attributable to physical failure) and which are up for a long periods of time (very stable prefixes, e.g., of a large ISP). The stable prefixes filter enables us to infer that about 53% of the more stable prefixes have MTTF of about two weeks and MTTR of about 10 hours, with median uptime of about 1.4 hours and median downtime of about 53 seconds. The filter also enables us to rank prefixes in order of their stability. There is significant churn in today's Internet and, while uptimes and downtimes do point to a healthy Internet, one needs to study other aspects like route convergence and availability for a holistic view. We investigate these aspects below.

We now apply the *route convergence* filter with parameter $\gamma = 200$ s, 300 s, and 900 s to the *initial filtering* data. Its effect on the MTTF and MTTR data is given in Table VIII. The results indicate that, as the value of γ increases, both uptimes and downtimes increase significantly. The median uptime increases by a much higher percentage than the median downtime, suggesting that the filter is improving the perceived "health" of the combinations.

TABLE VIII

ROUTE CONVERGENCE FILTER FOR MAR.-NOV. 07. INITIAL FILTERING
 MTTF = 827989 s, MTTR = 125479 s, MEDIAN UPTIME = 962 s,
 MEDIAN DOWNTIME = 64 s. ALL % CHANGES ARE RELATIVE TO INITIAL
 FILTERING

γ (s)	MTTF Change (%)	MTTR Change (%)	Median Uptime Change (%)	Median Downtime Change (%)
200	167.27	182.95	4755.93	1817.19
300	204.74	226.89	6972.56	2459.38
900	342.48	398.91	23633.89	6482.81

TABLE IX

STABLE PREFIXES AND ROUTE CONVERGENCE FILTERS MAR.-NOV. 07.
 INITIAL FILTERING MTTF = 827989 s, MTTR = 125479 s, MEDIAN
 UPTIME = 962 s, MEDIAN DOWNTIME = 64 s. ALL % CHANGES ARE
 RELATIVE TO INITIAL FILTERING

β (%)	γ (s)	MTTF Change (%)	MTTR Change (%)	Median Uptime Change (%)	Median Downtime Change (%)
0	200	168.46	150.29	4998.86	1807.81
0	300	206.2	189.30	7169.13	2446.87
0	900	344.74	341.82	24726.4	6375
60	200	196.71	-18.15	8902.39	1631.25
60	300	240.52	-4.57	14634.09	2234.37
60	900	402.79	49.64	49557.69	5346.87
99.5	200	279.80	-25.37	17362.99	1529.68
99.5	300	334.29	-12.87	26439.19	2128.12
99.5	900	539.55	39.32	93618.61	5284.37

Next, we put the *stable prefixes* and *route convergence* filters together as listed in Table IX. The results show that the trends of median and mean uptimes and downtimes with each of the filters are coalesced. For example, with *stable prefixes* filter parameter $\beta = 0\%$, the uptimes get a small boost over the *route convergence* filter case and the downtimes are slightly reduced in their gains from Table VIII. However, with $\beta = 60\%$, the MTTF increases by a larger amount over the *initial filtering* case and the MTTR is actually reduced. The effect of lowering the MTTR is reversed as γ crosses 300 seconds, which shows that route recovery is being (mis)classified as route convergence with higher values of γ . The median exhibits similar trends except that the increase in downtime due to the *route convergence* filter is high enough that it cannot be offset by any value of β .

We also compare our results of MTTF and MTTR obtained above to the results of Labovitz *et al.* [6]. Specifically, we compare the values to those depicted in Fig. 6 in [6] where the authors study the MTTF and MTTR of inter-domain routes from three ISPs and find the MTTF to be about 12 days (read from the graph) and the MTTR to be about 15 minutes. The comparison is shown in Table X.

One immediate difference is in the values of our MTTR which are very high compared to their counterparts in [6]. This is because the MTTR values in [6] are those for paths connecting three ISPs, which are backbone paths and expected to be more stable than a general Internet prefix. Since our dataset spans 9 months, it is not uncommon to have a significant number of combinations whose MTTR is of the order

TABLE X

COMPARISON OF MTTF AND MTTR VALUES OF MAR.-NOV. 07 TO THOSE IN [6] WHERE MTTF = 1036800 s, MTTR = 900 s

β (%)	γ (s)	MTTF (s)	MTTR (s)	MTTF Change (%)	MTTR Change (%)
None	None	827989	125479	-20.1	13842.11
60	900	4163080	187775	301.53	20763.89
99.5	200	3144780	93639	203.33	10304.33
99.5	900	5295400	174815	410.74	19323.89

of days (physical failure). Our median downtime ranges from around 17 minutes to up to 1 hour, indicating the reliability of the median.

However, our MTTF values are significantly higher than the ones in [6] (except for the *initial filtering* case). This confirms that the Internet has become healthier than in the late 1990s. These results also point to the importance of filter parameters. If we want to report the “best” possible results, (i.e., focusing on the healthy Internet), it seems logical that we should choose the most stable prefixes (hence $\beta = 99.5\%$), and a time sufficient for route convergence in most cases ($\gamma = 200$ s). This data point shows a reasonably high MTTF and reasonably low MTTR. This also illustrates the tradeoff between reporting a high MTTR and a high MTTF; in many cases, a high MTTF implies a high MTTR as well. This leads us to investigate the average uptime to downtime ratio, which cannot be computed by MTTF and MTTR alone. Hence, we add an instrumentation hook to record for each (peer, prefix) combination, the uptime and downtime, so that the up-to-downtime ratio of that combination, or equivalently the *availability* of the combination = $(1/(1/\text{up-to-downtime ratio})+1)$, can be computed. We use this hook to output the Cumulative Distribution Function (CDF) and histogram of the availability ratio of all the combinations. Many up-to-downtime ratios are infinite, since the combinations have a zero downtime at the end of the run (no failure was reported or the failure was removed by one of the filters). These ratios have an availability of 1.

The CDF of availability of various combinations for the *initial filtering* case is given in Fig. 1. From this plot, we can see that about 10% of the combinations are available for less than 60% of the time. About 80% of the combinations are available more than 90% of the time. That such a result is observed even without applying any filter points to the health of the Internet. However, the long tail of combinations having low availability cannot be ignored, and it can be concluded that the Internet is not as healthy as some would like it to be, e.g., compared to telephone networks.

We also plot the CDF of the availability when the *stable prefixes* and *route convergence* filters are applied with the same parameters as [6], i.e., $\beta = 60\%$ and $\gamma = 900$ s. This plot is shown in Fig. 2. Comparing this CDF plot with Fig. 4(a) of [6], we find some combinations with availability around 60% whose number appears to be higher than the plot in [6]. However, the 90th and the 95th percentile values reasonably

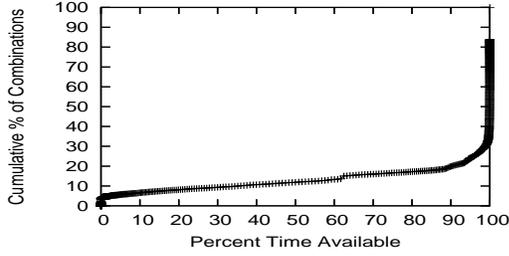


Fig. 1. CDF of combinations for the initial filtering case for Mar.-Nov. 07

match those in their paper. With $\beta = 99.5\%$ and $\gamma = 200$ s, we obtain a CDF with less weight at lower percentages and a higher weight at cumulative percentage close to 100%.

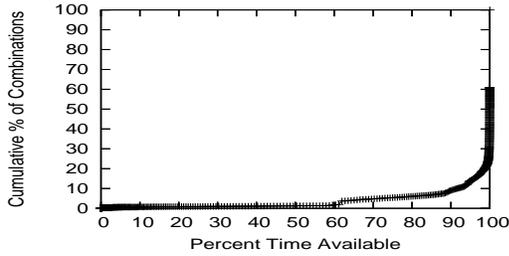


Fig. 2. CDF of combinations for Mar.-Nov. 07 with all filters, $\beta = 60\%$ and $\gamma = 900$ s

To better understand the impact of filters, we also plot the histograms of the availability. For the *initial filtering* case, the histogram is shown in Fig. 3. The plot shows that almost a constant number of combinations (around 0.25%) is available from close to 0% to about 90% of the time, except for 1.5% of the combinations which are available around 60% and a higher percentage is available for more than 90% of the time. The histogram corresponding to Fig. 2 is shown in Fig. 4. This plot shows that most of the combinations that have availability less than 60% are filtered, which is expected as $\beta = 60\%$. The combinations causing the spike around 60% were found to be prefixes originating from a single peer 206.24.210.99, which belongs to AS 3561 (SAVVIS Corporation). Almost all prefixes advertised by that peer have availability around 60%. This indicates why β should be large: since the routing tables are sampled fairly infrequently, it is easy for a relatively low availability combination to be present in a high $\beta\%$ of the routing tables.

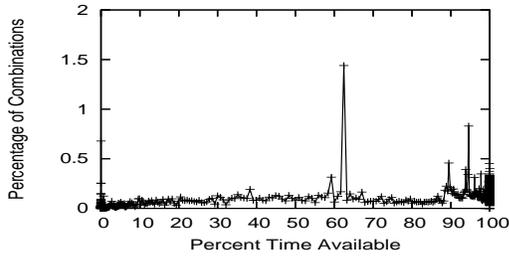


Fig. 3. Histogram of combinations for Mar.-Nov. 07 for the initial filtering case

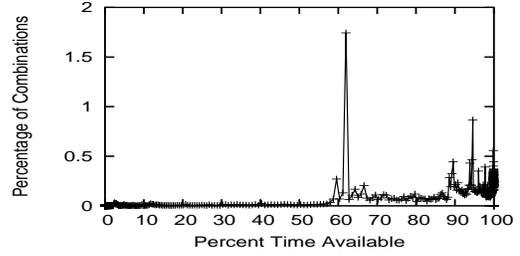


Fig. 4. Histogram of combinations for Mar.-Nov. 07 with all filters, $\beta = 60\%$ and $\gamma = 900$ s

C. Impact of Cable Cuts

In this section, we investigate the cable cuts in the Mediterranean sea and in the Middle East in early 2008. Since their effects are transient in nature (on the order of a few days), we use a one week window to study the uptimes and downtimes. We slide the window one day at a time. This allows sufficient time to investigate up and downtimes, while mitigating the impact of boundary effects. We consider Jan. and Feb. 2008. Hence, the first window is from Jan. 1st - Jan. 7th, 2008, the second is from Jan. 2nd to Jan. 8th, and so on. Since there are 60 days in these two months, we have 54 windows of observation, with indices 1–54.

For each window, we use the *observed combinations* filter with $\alpha = 10\%$. We evaluate the *initial filtering* case, cases with $\beta = 0\%$, 60%, and 100%, cases with $\gamma = 200, 300$ and 900 s, and finally the filters put back-to-back with the nine permutations of β and γ for the above stated values. Thus, there are 16 runs computing uptimes and downtimes for each window. Note that the maximum uptime or downtime that we can see will be limited by the observation period of one week, compared to the previous one or 9 month periods. However, the trends of uptimes and downtimes with various filters with time remain the same.

The change of uptimes for the *initial filtering* case against window index is depicted in Fig. 5. The figure shows that the median uptime holds steady at about 600,000 s for most of the windows, except for window indices between 26 and 34. This is intuitive since the cable cut occurred around the middle of window 26 (Jan. 26-Feb. 1). The lowest uptime (434833 seconds) is seen during the 28th window (Jan. 28-Feb. 3) and is about 28% lower than the value for most of the windows.

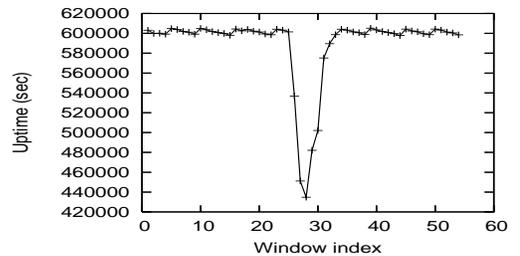


Fig. 5. Median uptime vs. window index with initial filtering for Jan.-Feb. 08

The lowest observed uptime starts increasing as we apply

the *stable prefixes* filter and the *route convergence* filter. This is because these filters prefer stable prefixes and filter away dynamics so the transient nature of this stress event is filtered out. The plot of median uptime for $\beta = 60\%$ and $\gamma = 900$ s (the parameters used in [6]) is shown in Fig. 6. The sharp drop from 600000 s to about 435000 s has completely disappeared. While this demonstrates the effectiveness of filtering, it reiterates the notion that one cannot use filter parameters without thorough investigation. In fact, we argue that *if we want to observe stress events*, we must use the *initial filtering* case, whereas if want to make general conclusions about the health of the Internet, we should use filters with carefully chosen parameters.

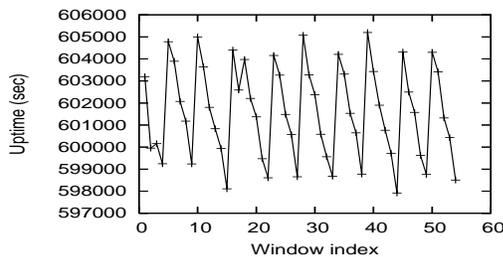


Fig. 6. Median uptime vs. window index for Jan.-Feb. 08 with all filters, $\beta = 60\%$ and $\gamma = 900$ s

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we explore the methodology for quantifying reachability and availability of Internet prefixes. We find that although preprocessing filters on BGP data offer significant power in eliminating pathological updates and unstable prefixes, they should be used judiciously. In fact, one can modify the filter parameters to make the Internet look healthy or unhealthy. We provide recommendations for the values of the parameters based on our observations, and suggest that they should be chosen according to the goal of a study. For example, if the goal is to study the health of the Internet, we should use both the *stable prefixes* and the *route convergence* filters with high values of β and the minimum value of γ indicative of route convergence, whereas if the aim is to identify stress events, it is advisable not to filter the data.

We compute the MTTF and MTTR for a 9-month period in 2007 with and without filtering. We show the (peer, prefix) combinations filtered by using a histogram of availability of the combinations under different filter parameters. We find that 53% of the most “healthy” prefixes have MTTF of about two weeks and MTTR of about 10 hours, with median uptime of about 1.4 hours and median downtime of about 53 seconds. We also study prefix reachability during cable cut events and show that the median uptime of prefixes fell by as much as 30%.

As future work, we aim to study prefix reachability over a wider range of periods, and hone the parameter selection of the filters by studying the distributions of prefix lifetimes and routing update arrivals. We also plan to investigate the impact of prefix aggregation. Tying in the prefix reachability to the resilience of the Internet is our continuing goal.

REFERENCES

- [1] International Herald Tribune, “Telephone and Internet services cut off after Taiwan earthquake,” http://www.iht.com/articles/ap/2006/12/27/asia/AS_GEN_Taiwan_Quake.php, 2006.
- [2] Cable News Network (CNN), “Internet failure hits two continents,” <http://edition.cnn.com/2008/WORLD/meast/02/01/internet.outage/index.html>, Jan. 2008.
- [3] Fox News, “Severed Cables Cut Egypt’s Internet Access Once Again,” <http://www.foxnews.com/story/0,2933,470224,00.html>, Dec. 2008.
- [4] Network World, “Mideast cable cuts tripled Web latency on some routes, Akamai reports,” <http://www.networkworld.com/news/2008/021408-cable-cuts-tripled-web-latency.html>, Feb. 2008.
- [5] Akamai, “Mideast Outage,” <http://www.akamai.com/mideast-outage>, Jan. 2008.
- [6] C. Labovitz, A. Ahuja, and F. Jahanian, “Experimental Study of Internet Stability and Backbone Failures,” in *Proc. of FTCS*, 1999.
- [7] C. Labovitz, G. R. Malan, and F. Jahanian, “Internet Routing Instability,” *IEEE/ACM Transactions on Networking*, vol. 6, no. 5, pp. 515–528, 1998.
- [8] University of Oregon, “University of Oregon Route Views Project,” <http://www.routeviews.org/>.
- [9] J. Li, M. Guidero, Z. Wu, E. Purpus, and T. Ehrenkrantz, “BGP Routing Dynamics Revisited,” *SIGCOMM Computer Communication Review*, vol. 37, no. 2, pp. 5–16, 2007.
- [10] X. Zhao, D. Massey, S. F. Wu, M. Lad, D. Pei, L. Wang, and L. Zhang, “Understanding BGP Behavior through a Study of DoD Prefixes,” in *Proc. of DARPA Information Survivability Conference and Exposition*, April, 2003.
- [11] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson, “Studying Blackholes in the Internet with Hubble,” in *Proc. of NSDI*, 2008.
- [12] J. Wu, Z. M. Mao, J. Rexford, and J. Wang, “Finding a Needle in a Haystack: Pinpointing Significant BGP Routing Changes in an IP Network,” in *Proc. of NSDI*, 2005.
- [13] F. Wang, Z. M. Mao, J. Wang, L. Gao, and R. Bush, “A Measurement Study on the Impact of Routing Events on End-to-End Internet Path Performance,” in *Proc. of ACM SIGCOMM*, 2006.
- [14] J. Li, R. Bush, Z. Mao, T. Griffin, M. Roughan, D. Stutzbach, and E. Purpus, “Watching Data Streams Toward a Multi-Homed Sink Under Routing Changes Introduced by a BGP Beacon,” in *Proc. of PAM*, 2006.
- [15] J. Wu, Y. Zhang, Z. M. Mao, and K. G. Shin, “Internet routing resilience to failures: analysis and implications,” in *Proc. of CoNEXT*, 2007, pp. 1–12.
- [16] E. Zmijewski, “Threats to Internet Routing and Global Connectivity,” in *Proc. of 20th Annual FIRST Conference*, 2008.
- [17] M. Caesar, L. Subramanian, and R. Katz, “The Case for an Internet Health Monitoring System,” in *Hot Topics in System Dependability (HotDep)*, 2005.
- [18] A. Feldmann, O. Maennel, Z. M. Mao, A. Berger, and B. Maggs, “Locating Internet Routing Instabilities,” in *Proc. of ACM SIGCOMM*, 2004.
- [19] R. Oliveira, B. Zhang, D. Pei, R. Izhak-Ratzin, and L. Zhang, “Quantifying path exploration in the internet,” in *IMC ‘06: Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*. New York, NY, USA: ACM, 2006, pp. 269–282.
- [20] D.-F. Chang, R. Govindan, and J. Heidemann, “The temporal and topological characteristics of BGP path changes,” in *ICNP ‘03: Proceedings of the 11th IEEE International Conference on Network Protocols*. Washington, DC, USA: IEEE Computer Society, 2003, p. 190.
- [21] J. Rexford, J. Wang, Z. Xiao, and Y. Zhang, “BGP Routing Stability of Popular Destinations,” in *Proc. of ACM IMW*, 2002.
- [22] RIPE, “RIPE Network Coordination Centre,” <http://www.ris.ripe.net/source/>.
- [23] B. Zhang, V. Kambhampati, M. Lad, D. Massey, and L. Zhang, “Identifying BGP routing table transfers,” in *Proc. of ACM MineNet workshop*, 2005.
- [24] R. V. Oliveira, B. Zhang, and L. Zhang, “Observing the evolution of internet as topology,” *SIGCOMM Computer Communication Review*, vol. 37, no. 4, pp. 313–324, 2007.
- [25] S. Burkle, “BGP convergence analysis,” Ph.D. dissertation, 2003.
- [26] C. Labovitz, A. Ahuja, A. Abose, and F. Jahanian, “The Problem with BGP,” <http://www.nanog.org/mtg-0002/converge2.html>, 2000.