Characterizing IPv6 control and data plane stability

Ioana Livadariu^{*}, Ahmed Elmokashfi^{*} and Amogh Dhamdhere[†] *Simula Research Laboratory, 1364 Fornebu, Norway Email: {ioana,ahmed}@simula.no [†] CAIDA/UCSD, La Jolla, CA 92093, U.S.A Email: amogh@caida.org

Abstract-End-to-end IPv6 performance is a factor that can influence IPv6 adoption. The stability of IPv6 - both in the control and data plane - is an important determinant of endto-end performance, as it influences packet loss, network latency, and hence application performance. In this paper we compare stability and performance measurements from the control and data plane in IPv6 and IPv4. To study control plane stability, we use BGP feeds from five dual-stacked vantage points to measure routing dynamics towards IPv4 and IPv6 destinations. To study data plane stability, we probe dual-stacked webservers in 629 target ASes to determine the availability, RTT performance and RTT stability of paths toward these targets. In both control and data plane experiments IPv6 exhibited less stability than IPv4. In the control plane, most routing dynamics were generated by a small fraction of pathological unstable prefixes. In the data-plane, episodes of unavailability were longer on IPv6 than on IPv4. We found evidence of correlated performance degradation over IPv4 and IPv6 caused by shared infrastructure.

I. INTRODUCTION

The pool of available IPv4 addresses is rapidly decreasing; currently, four out of five Internet Registries are allocating from their last /8 pool of IPv4 addresses [1], [2], [3], [4]. A large scale transition to IPv6 is the long-term answer to the IPv4 address scarcity issue. However, a lack of backward compatibility with IPv4, required hardware and customer equipment upgrades, and the lack of economic incentives for deployment have delayed the widespread adoption of IPv6. Despite recent studies that showed that IPv6 is maturing [5], [6], actual uptake remains slow; only 7.73% of users access Google over IPv6, native IPv6 accounts for 7.72% of traffic [7], and only 18% of Autonomous Systems (ASes) in the BGP routing system advertise an IPv6 prefix [8]. Market needs and regulatory bodies may eventually speed up IPv6 uptake, but is IPv6 ready for prime time? We believe that for IPv6 adoption to gain traction, stability and performance over IPv6 should evolve to the point that it is comparable over IPv4.

Assessing IPv6 stability and performance involves quantifying several control plane, data plane, and application-specific metrics. We focus on the first two, as it is well known that routing instability can cause performance degradation [9], and the stability of the data plane – in terms of network availability, RTT, and the variability of RTTs – directly impacts end-toend performance. The IPv4 Internet has benefited from years of fine tuning, optimization, and measurement; the stability of the IPv4 routing system and IPv4 data plane performance are thus well-understood. On the other hand, there is relatively little work on measuring IPv6 control and data plane stability and comparing it with its IPv4 counterpart. In this paper we study IPv6 stability in depth.

To characterize control-plane stability, we measure the frequency of routing changes towards IPv4 and IPv6 prefixes, the number of prefixes that are active (experience routing changes) on a daily basis, the contribution of highly active prefixes to routing dynamics, and correlations between routing instability in IPv4 and IPv6. We measure these properties using using BGP feeds from five dual-stacked ASes (VPs) that peer with Routeviews [10]. In our measurements IPv6 was in general less stable than IPv4, with more unstable prefixes and more pathological activity. The IPv4 and IPv6 routing systems differed mainly in three aspects: the contribution of the top 1% of active prefixes to the overall routing dynamics, the contribution of prefixes of different lengths to the overall routing dynamics, and the composition of the routing event mix. We found low correlation between instability periods in IPv4 and IPv6 as seen by a VP toward the same dual-stacked AS. This indicates that the IPv4 and IPv6 routing systems do not share fate, at least in terms of control-plane dynamics.

To characterize data-plane stability, we measure network availability over IPv4 and IPv6 (reachability of target ASes from our probing vantage points), relative RTT performance of IPv4 and IPv6, and the stability of RTTs across the probing duration. We measure these properties by probing webservers in 629 target ASes (see Section III for details on the methodology) from 6 dual-stacked Ark monitors [11]. We also collect traceroutes and DNS information to infer whether IPv4 and IPv6 paths are congruent at the router level. In our measurements the overall availability of the probed target ASes was comparable over IPv4 and IPv6. The differences were in the tails of the distribution, where a small fraction of targets were unreachable for significantly longer time periods over IPv6 than over IPv4. Comparing RTT performance over IPv4 and IPv6 revealed a notable change from prior results of Dhamdhere et al. [5], who reported that IPv4 was faster than IPv6 in 78% of cases (albeit with a different set of vantage points and target ASes than we used in this study). In our measurements, performance was almost equally likely to be better over IPv4 or IPv6 for a probed target AS (54% of cases better over IPv4). We found episodes of elevated RTTs, most of which we determined were not caused by routing changes but possibly due to congested links. We found cases where IPv4 and IPv6 paths toward a target AS experienced co-ordinated level shifts due to shared infrastructure.

II. CONTROL PLANE STABILITY

A. Data set

We use BGP updates from five dual-stacked ASes that peer with RouteViews: Hurricane Electric (HE) (AS6939),



Fig. 1. CCDF of the number of events per prefix per day for HE(left) and NTT(right). The difference between IPv4 and IPv6 dynamics is most evident at the tail of the distribution, indicating that a sizable fraction of IPv6 prefixes are highly active. IPv6 activity shows more variability across years, while IPv4 activity is stable.

NTT (AS2914), Tinet (AS3257), APAN (AS7660) and IIJ (AS2497). We choose these ASes because they have IPv4 and IPv6 peering with RouteViews at the same locations according to information published by RouteViews [10]. Ensuring that IPv4 and IPv6 peering are co-located is necessary, since we plan to correlate routing activity for the same AS over IPv4 and IPv6. These five ASes represent two large transit providers (NTT and Tinet), two mid-size/small networks (APAN and IIJ) and the largest AS in the IPv6 ecosystem (HE) [5]. Our data set consists of quarterly (January, April, July, October) snapshots spanning the period from January 2009 to January 2015. The number of events varies between months and generally follows an increasing trend for both routing systems. The number of IPv4 events is between 6 and 15 times the number of IPv6 events. This ratio is less than the ratio between the number of prefixes, which is currently at 21.8 [12]. Further, the number of IPv4 events in January 2015 is between 1.6M and 3.4M depending on the monitor.

B. Comparing routing changes

When a BGP router experiences a route change to a destination prefix, it explores available alternative routes until it converges to a new route or removes the affected prefix from its routing table. Consequently, a single routing change can trigger multiple BGP updates during the path exploration phase. We compare the IPv6 and IPv4 control plane stability in terms of routing changes. To this end, we use the definition in [13] to group BGP updates for the same prefix into routing events.

Frequency of routing changes

For each of the 25 quarterly snapshots, we empirically estimate the distribution of the number of routing events per prefix per day for both IPv4 and IPv6. To capture routing changes that involve path exploration and to avoid bias due to dynamics within the monitor AS, we only consider routing events that involve more than one update. These events account for 90% of all events.

The left panel on Figure 1 shows the CCDF of the number of events per prefix per day for HE in January 2011, January 2013, and January 2015 for both IPv4 and IPv6. The right panel on the same figure shows the same CCDF for NTT in January 2010, January 2012, and January 2014. The difference between IPv6 and IPv4 is evident over the whole range of the distribution, but more clearly on the tail. For both VPs the fraction of IPv4 prefixes that experience more than 100 routing events per day is less than 0.001. The fraction of IPv6 prefixes that experience more than 100 routing changes per day varies temporally and across VPs, but can be as high as 0.05



Fig. 2. The average fraction of active prefixes per day for HE (top) and NTT(bottom). The fraction of active prefixes in IPv6 was twice that of IPv4 in 2009. In the last 2-3 years, the fractions are comparable in IPv4 and IPv6.

(HE-January 2011) and is typically around 0.01. The order of magnitude difference between IPv4 and IPv6 in the tail of the distributions indicates that a sizable fraction of IPv6 prefixes contributes a disproportionally large number of updates i.e., heavy-hitters. Looking at the head of the distributions, we observe variability in IPv6 activity across different years. For example, the fraction of prefixes that are observed to be active at least twice per day by the HE monitor is 82% and 68% in January 2013 and 2015 respectively. The variability in IPv4 activity is less pronounced across the years, which is expected since the IPv4 Internet is more mature. Other monitors and months in our data set exhibit similar characteristics. We omit those graphs due to space constraints.

To check whether the difference between IPv6 and IPv4 is diminishing over time, we calculate the Kolmogorov-Smirnov (KS) distance [14] between IPv4 and IPv6 distributions from the same month. The KS distance is the maximum vertical distance between two distributions; the larger the distance, the less similar are the two distributions. For all VPs, the KS distance has fluctuated over the years in the range between 0.1 and 0.15 without evident trend. In other words, the IPv4 and IPv6 distributions do not exhibit clear convergence.

Number of active prefixes

The previous analysis shows that the IPv6 routing system is less stable than IPv4 at the macroscopic level, but it does not reveal microscopic differences at the prefix-level, which we delve into next. For each day in our study period, we calculate the number of active prefixes (defined as the the number of prefixes that experience routing changes on that day) and then average over all days in the corresponding month. Note that we are not interested in how many times a prefix is active per day, thus we count an active prefix only once even if it experiences routing changes multiple times in the day. The plots in Figure 2 show the evolution of the fraction of active prefixes seen by the HE and NTT monitors. The fraction is calculated by dividing the number of active prefixes by the total number of observed prefixes. Across VPs, the fraction of active IPv4 prefixes was consistently between 0.03 and 0.04. The IPv6 fraction, on the other hand, shows more variability across VPs and over time. The fraction of active IPv6 prefixes was twice that of IPv4 in 2009. After 2009, and especially in the last 2-3 years, the average fraction of active IPv6 prefixes is closer to IPv4. This indicates that the fraction of active prefixes is becoming similar in both routing systems, though this fraction continues to be slightly higher in IPv6.

Which prefixes contribute most updates?

To gain further insight into the differences between IPv4 and IPv6 routing dynamics, we compute the fraction of updates contributed by prefixes of different lengths. We find that over 80% of IPv4 activity is contributed by prefixes (>/19) and 50%



Fig. 3. The average number of updates per prefix of certain length averaged over a period of one month for HE IPv4 (top) and IPv6 (bottom). In IPv4, prefixes longer than /20 do not contribute more than their share in the routing table (/20 is an exception). In IPv6, the average activity of /32 and /48 prefixes is comparable.

is contributed by /24s. Over 95% of IPv6 activity is contributed by /32 and /48 prefixes with /32s being the bigger contributor. We also note the high contribution to the IPv6 updates of /46 and /44 prefixes in 2013 and 2015, respectively. Due to the fact that we observe this contribution only in one year, we hypothesize that this is caused either by an experimental deployment of the prefixes, or by operators that were unfamiliar with the IPv6 deploying and troubleshooting We examine the relative contribution of prefixes of a certain length by dividing the average number of updates contributed by all prefixes of that length by their number in the routing table. Figure 3 shows the relative contribution from IPv4 prefixes (top panel) and IPv6 prefixes (bottom panel) of different lengths as observed by the HE VP in January 2011, January 2013, January 2015. For IPv4, we observe that prefixes longer than /17, except for /20s, do not contribute more than their share in the routing table, matching earlier results [15]. For instance, the average number of updates per /24 is comparable to that per /21, although there are many more /24s in the routing table. We also observe that the average number of updates per prefix for prefixes longer than /19 has increased markedly in 2015. We leave investigation of the recent higher activity of longer prefixes to future work. For IPv6, we observe that the number of updates per /32 and /48 prefixes are comparable.

Highly active prefixes

Previous studies showed that most BGP updates in IPv4 are generated by a few highly active prefixes [16]. This high activity is mostly due to pathological phenomena that is caused by flaky equipment or misconfiguration. Hence, the contribution of these prefixes to overall routing churn is an



Fig. 4. The contribution of the top 1% active prefixes to the overall number of BGP updates APAN (left) and IIJ (right). The top 1% active prefixes are responsible for between 40% and 60% of all IPv6 updates, which is approximately twice the contribution of the top 1% prefixes in IPv4.



Fig. 5. The contribution of the top 10% active prefixes to the number of updates APAN (left) and IIJ (right). The top 10% are responsible for between 75% and 85% of IPv6 updates and between 45% and 55% of IPv4 updates.

indicator of the health of the routing system.

We measure the contribution of the top 10% and top 1%active prefixes in IPv4 and IPv6 to overall BGP dynamics. Figures 4 and 5 show these contributions for IPv4 and IPv6 and how they have evolved over time from the perspective of the APAN and IIJ VPs. Across quarterly snapshots, the top 1% of active prefixes were responsible for between 40% and 60% of all IPv6 updates, and for about half of the IPv4 updates. Further, the top 10% were responsible for between 75% and 85% of IPv6 updates and between 45% and 55% of IPv4 updates. Interestingly, the difference between these contributions (i.e. top 10%-top 1%) was about 30% in both IPv4 and IPv6. This implies that the main difference between IPv4 and IPv6 heavy hitters was limited to the top 1%. IPv6 heavy hitters differ from their IPv4 counterparts in that they remained highly active for several days. For example, one-fifth of the top 1% active IPv6 prefixes in a given month were active for over a week in that month. However, only 5% of the top 1% active IPv4 prefixes in a month were active for over a week. We hypothesize that this difference is caused by the relative immaturity of IPv6 and the fact that relatively little user traffic is carried over it; thus, routing instability and prefix flapping may go unnoticed for days before getting fixed.

Types of routing changes

The previous results show that IPv6 is overall less stable than IPv4. We next compare the types of routing changes to see if they are significantly different. We process all routing events identified earlier and classify them into five types according to the state of the affected prefix before and after the event: 1) AW: events that withdraw an announced prefix 2) WA: events that announce a previously unreachable prefix 3) AAC: events that re-announce an existing prefix experiences a transient path disturbance then re-converges to its initial routes (i.e. the one in use before the event) 5) AAS: events where affected prefixes do not experience an AS-path change, but rather a change in one of the route attributes e.g., MED or community. Figure 6 shows the fraction of events by type from the perspective of the HE monitor for IPv4 (left panel) and IPv6 (right panel).



Fig. 6. Breaking down the routing dynamics by the type of routing event from the perspective of HE monitor V4 (left) and V6 (right). Events of type AAD dominate the IPv6 mix, while events of type AAC dominate the IPv4 mix.

The event composition in IPv4 and IPv6 exhibits one clear difference: Events of type AAD dominate the IPv6 mix, while events of type AAC dominate the IPv4 mix. Inspecting the AAD events in IPv6, we find in most cases that the affected prefix is withdrawn and then re-announced with the old path, i.e., the AAD events are essentially a combination of AW and WA events. AAD events result in a route becoming temporarily unavailable, possibly due to transient network failures, session resets, or router reboots. AAC events, on the other hand, are related to changes that last for at least more than one minute which could be related to routing policies or failures that require more time to fix. In general, we expect more AAC events in a dense topology that offers alternative routes, which is the case in the IPv4 topology. The higher occurrence of events of AW and AAD indicates the lack of path diversity in the IPv6 internetwork.

Summary

On the surface, the IPv6 routing system appears less stable than its IPv4 counterpart. Looking deeper, we find that the two routing systems differed mainly in three aspects: the role of top 1% of active prefixes (heavy hitters), the contribution of prefixes of different lengths to the overall dynamics, and the composition of the event mix. We believe that these differences can be attributed to IPv6 immaturity. The first difference will most likely disappear once users begin to depend more on IPv6, which will motivate operators to troubleshoot their prefixes. The third difference should become less prominent as the IPv6 interdomain topology becomes denser, resulting in more path diversity and reducing the likelihood of AAD events.

C. Correlating instabilities

Next, we investigate whether IPv4 and IPv6 routing changes are correlated. Strong correlations between routing events in IPv4 and IPv6 would indicate that the two routing systems share fate, possibly due to sharing the underlying infrastructure. For each VP, we processed all events that affected IPv6 prefixes originated by the same AS and group the overlapping events, giving us a list of time windows where IPv6 prefixes from a certain origin AS were active. We repeated the same process for IPv4 prefixes, and correlated the two time series as follows. We traversed all IPv6 activity windows and checked if there was an overlapping IPv4 window within δ seconds¹. We then divided origin ASes into two groups according to the AS paths seen by the monitor: The *congruent group* consisting of ASes with congruent (i.e.,



Fig. 7. Correlation of IPv4 and IPv6 routing changes HE (left) and IIJ (right) calculated for each quarterly snapshot. Origins with congruent AS level paths from the VPs show higher correlation than origins with non-congruent paths. The overall correlation is low (<0.5) both for origins with congruent and non-congruent paths.

identical) AS paths and a *non-congruent group* consisting of ASes with non-congruent AS-paths (paths differ in at least one hop). Finally, we averaged the fraction of overlapped windows across all dual stack origin per group and used it as a measure of correlation between instabilities in IPv4 and IPv6.

Figure 7 shows the evolution of the average correlation fraction within the two groups for HE (left) and IIJ (right). Origins in the *congruent* group consistently exhibit higher level of correlation than the *non-congruent* group. The correlation for the *congruent* group was higher before 2011 – about twice that of the non-congruent group. This difference has decreased after 2011. The average correlation for the *congruent* group, however, remains low at about 0.25.

Summary

As expected, IPv6 and IPv4 instabilities are more correlated for the origins in the *congruent* group for a monitor. The difference between the correlation within the *congruent* group and *non-congruent* group has decreased starting from 2011. This decrease coincides with a period with accelerated IPv6 adoption [5]. We believe that this drop and the generally low correlation between IPv4 and IPv6 instabilities are because paths that are congruent at the AS-level can be non-congruent at the router level. Overall the low level of correlation suggests that the IPv4 and IPv6 routing systems *do not share fate*, at least in terms of control-plane dynamics.

III. DATA PLANE STABILITY

Routing instability can directly affect data plane availability, e.g., causing packet loss and increased delays [17], [18]. In this section we investigate the IPv6 data plane stability and compare it to IPv4. We focus on two measures of data plane stability – network availability and performance.

A. Measurement setup

We probe dual-stacked web servers from six Ark monitors [11], situated in five different countries: Australia (per-au), Germany (bre2-de), Netherlands (ams-nl), Switzerland (zrh2ch) and US (sql-us and jfk-us). We used the Alexa top 1M list [19] and performed DNS lookups for the A and AAAA records for each domain, and then mapped the resulting IP addresses to ASes. We selected ASes that hosted both IPv4 and IPv6 webservers, and chose a maximum of 2 IPv4 addresses and 2 IPv6 addresses per dual-stacked AS to probe. This gives us a set of 629 target ASes and 1891 target IP addresses in total, which we probed from the VPs every 5 seconds for 38 days in March/April 2015. We used AS path data [10] from the same period of time to determine the AS path from

¹We varied the threshold δ to 60, 120, and 300 seconds, observing qualitatively similar results.

the VP AS to the target AS in IPv4 and IPv6. The data provides AS path information for 611 of 629 targets ASes, which we use to classify each (*monitor,target*) pair into one of two groups: *congruent* (identical IPv4 and IPv6 AS paths) or *non-congruent* (AS paths differed by at least one hop).

B. Network availability

We focus our study of network availability of the target ASes on the following metrics: the overall reachability of target ASes in IPv4 and IPv6, the length of unreachability episodes over IPv4 and IPv6, and correlation between the unreachability periods for a target AS over IPv4 and IPv6. For this analysis we divide the probing time for a target AS into periods when it was reachable from the VP (at least one of the target IP addresses in the AS was responsive) and unreachable (all target IPs in that AS were unresponsive). A caveat with using webservers as measurement targets is that unreachability towards these targets can be due to causes other than network failures, such as web service unavailability or random losses. As a consequence, we consider an unreachability period to be caused by network failures if it lasts at least 15 seconds and no longer than one hour. We also filter out targets that are unreachable for more than half of the probing period. After this filtering step, we retain 86.8% of the 629 target ASes.

Network reachability

For each target AS, we compute the *reachability fraction* over IPv4 and IPv6 as the ratio of the total reachability time and the total probing time. Over our measurement duration, we found that the reachability fraction over IPv6 was comparable with the IPv4 counterpart; 91.94% of the targets were reachable over IPv4 and IPv6 for at least 99% of the probing period. The differences between IPv4 and IPv6 manifest themselves in the tail of the distributions, where some targets were unavailable up to 40% of the time over IPv4 and IPv6. However, reachability was consistently higher over IPv4 than IPv6.

To compare the IPv4 and IPv6 reachability for the same target AS we consider only the targets for which the overlap of the probing periods over IPv4 and IPv6 was greater than 80%. This filtering results in the removal of 7.50% of the targets. For 92.47% of the remaining targets, the reachability fractions over IPv4 and IPv6 were within 10% of each other. This finding reinforces our previous observation — for most targets the overall reachability fraction was similar over IPv6 and IPv4.

Length of unreachability periods

To gain insight into typical downtime periods over IPv4 and IPv6 for the probed target ASes we plot in Figure 8 the distribution of unreachability periods. For both IPv4 and IPv6, most unreachability periods were short; only 3.03% of unreachability episodes for IPv4 and 2.77% for IPv6 were longer than 150 seconds. The unreachability period was in general longer over IPv6 than IPv4; for 30% of targets this difference was greater than 20 seconds.

Comparing network availability over time

Next we conduct a comparative analysis of network availability using two data snapshots from 2014 and 2015. We use the already described data collected during 2015, and data collected using the same methodology over a period of 37 days in August/September 2014. We exclude from this comparative



Fig. 8. CDF of the length of unreachability periods over IPv4 and IPv6. Overall most unreachability periods are short-lived.

analysis 16.37% of target ASes that were not reachable for more than half of the probing period in both in 2014 and 2015. For most targets the overall reachability was similar for the two periods of time; 84.6% and 92.5% of the targets were reachable for more than 99% of the probing period in 2014 and 2015, respectively. Moreover, for 91% of the targets the paired IPv4 and IPv6 reachability fraction differed by at most 0.1 in both snapshots. Figure 9 shows the quartiles, 10th and 90th percentile of the unreachability periods over IPv4 and IPv6 in 2014 and 2015. In both data snapshots unreachability episodes were short-lived, though they generally lasted longer over IPv6 than IPv4. We note a decrease in the median length of unreachability periods over IPv6 in 2015. This finding hints that IPv6 network availability is improving over time.



Fig. 9. Comparison of the length of the unreachability periods over IPv4 and IPv6 for two data snapshots. Unreachability periods were longer over IPv6 than IPv4. The median length of unreachability periods in IPv6 decreased between the two snapshots.

Correlating IPv4 and IPv6 downtime

We investigate the correlation between IPv4 and IPv6 downtime periods for a given (monitor, target) pair. For each monitor, target) pair we divide the entire measurement period into 5 minute bins. We define a loss event in a particular bin as the event that an unreachability episode of longer than 15 seconds occurred in that bin.² Let $P(X_6)$ be the probability that target X experiences a loss event over IPv6 in a certain bin. We compute the conditional probability $P(X_6|X_4)$ as the probability of observing a loss event over IPv6 given that we observed a loss event over IPv4 in the same bin. We compare the two probabilities by computing the conditional probability ratio $R=P(X_6|X_4)/P(X_6)$. If R is close to 1 for a target X, then the probability of observing a loss event over IPv6 given that we observed a loss event over IPv4 is the same as the probability of observing a loss event over IPv6. Thus, the

²Recall that we imposed a threshold of 15 seconds on the minimum length of an unreachability period to rule out random losses and server unresponsiveness.

downtime periods over IPv6 happen independently than those over IPv4. However, if the conditional probability is much higher and consequently the value of R is much greater than 1, then there is a correlation between the unreachability periods over IPv4 and IPv6.

We impose thresholds of 10 and 100 on the conditional probability ratio to identify significant correlations between IPv4 and IPv6 unreachability periods. We found that 72% of the *(monitor,target)* pairs in the *congruent* group and 55% in the *non-congruent* group had R > 10. Increasing the threshold to 100, 55% of the *(monitor,target)* pairs in the *congruent* group and 38% in the *non-congruent group* had R > 100. The implication is that a significant fraction of targets had a high correlation between unreachability periods in IPv4 and IPv6, even when AS paths were incongruent. Such a high correlation between the downtime periods would hint at shared infrastructure between the IPv4 and IPv6 paths. We analyze the prevalence of shared infrastructure in Section II.

Summary

The analysis showed that network availability was in general high over both IPv4 and IPv6. When unreachability periods occurred, they were longer over IPv6 than IPv4. Analysis of two data snapshots from 2014 to 2015 shows similar qualitative results, though the median unreachability period over IPv6 decreased by 10 seconds. A significant fraction of targets showed strong correlation between downtime periods over IPv4 and IPv6, and we observed this correlation both for targets with congruent and non-congruent AS paths from the VPs.

C. Performance

We focus our study of performance in IPv4 and IPv6 on the following metrics: relative difference in RTTs on IPv4 and IPv6 paths to the same target AS, and RTT stability over the measurement period.

Relative performance

Figure 10 shows the relative difference in average RTT over the whole probing period per target in IPv4 and IPv6, computed as *(slower-faster)/faster*. The region to the right of 0 constitutes cases where IPv4 was faster than IPv6, i.e, RTTs were lower over IPv4 than IPv6; this region accounts for 54.76% of the *(monitor,target)* pairs. The region to the left is when IPv6 was faster than IPv4, and accounts for is 45.24% of the *(monitor,target)* pairs. We find a notable improvement in IPv6 performance since the study by Dhamdhere et al. [5] in 2012; they reported that IPv6 was faster in only 22% of cases.

The solid lines demarcate regions where the relative difference in RTT is at least 10%, and accounts for 33.33% of the *(monitor;target)* pairs when IPv4 is faster and 30.76% of the pairs when IPv6 is faster. We confirm the results of Dhamdhere et al. [5] who showed that AS path congruency is an important factor in determining whether IPv4 and IPv6 performance are comparable. Interestingly, we now see this effect in both cases – when IPv4 performed better and when IPv6 performed better. When IPv4 had better performance, 70.33% of targets that had congruent and 51% of the targets that had non-congruent AS paths had IPv6 performed better, 75% of targets that had congruent and 64% of targets that had non-congruent AS paths had IPv4 performance within 10% of IPv6. The observed improvement in IPv6 performance confirms recent studies by Czyz et al. [6] that indicate that IPv6 is maturing.



Fig. 10. Relative RTT average values over IPv4 and IPv6. The percentage of targets for which IPv4 is faster than IPv6 and IPv6 is faster than IPv4 differs with 10%. For both categories congruency is an important factor.

RTT stability

An important determinant of end-to-end performance is the stability of RTTs, i.e., whether end-to-end paths experience episodes of potential performance deterioration due to elevated RTTs. In order to detect instability episodes we apply a CUSUM-based [20] level shift detection algorithm on the collected RTT time series data. The algorithm uses a sequential t-test analysis to detect sustained changes in RTT values, and is robust to outliers. To detect RTT changes that are potentially performance-affecting, we impose a threshold of t = 25 ms on the increase. As a consequence, 57.87% and 55.49% of the targets experience over IPv4 and over IPv6 respectively. For each target we compute the fraction of the total probing time for which the RTT is in the elevated state; we found that the level shift episodes accounted on average for a small fraction of the total reachability time - 2.82% over IPv4 and 2.07% over IPv6. However there are outliers; 4 of the 546 targets experienced level shifts in more than half of the reachable period. Of these, one target experienced AS path changes lasting several days causing the RTT level shift. Two targets saw level shifts for several hours per day, possibly due to congested links. The level shifts for one target were due to false positives of the level-shift algorithm. We divide target ASes into three classes: targets that experience RTT increases over both IPv4 and IPv6 (S_{46}) , over only IPv4 (S_4) or only IPv6 (S_6). We find that 41.94% of the total targets experienced RTT level shifts over both IPv4 and IPv6, whereas 15.93% and 13.53% experienced level shifts over only IPv4 and only IPv6, respectively.

To infer the location and cause of long-lived RTT level shifts we ran *traceroute* every two hours from each vantage point towards the target ASes. We are able to analyze cases where we had at least one traceroute during the level shift, which accounts for 28.5%, 40.74% and 38.09% of the level shifts for the targets included in S_4 , S_6 and S_{46} , respectively. We map each IP hop in the collected traceroutes to its corresponding AS and obtain its DNS name. For targets that experience level shifts over both IPv4 and IPv6, we use DNS names of hops on the IPv4 and IPv6 paths to determine if the increase occurs at the same link along the paths.

Analysis of level shifts due to forward path changes

To estimate the potential impact of routing changes on RTT instability, we analyze level shifts that coincide with a forward path change, i.e., the hops on the forward path during the shift are not present either before or after the shift occurs. We find that only a small fraction (7.25%) of the detected shifts coincided with path changes; 60% of these level shifts coincided with path changes in the Hurricane Electric (HE, AS6939) network. One of our VPs (*jfk-us*) is located in HE, and hence path changes in HE's network coincided with level shifts over both IPv4 and IPv6 for targets probed by *ifk-us*. In Figure 11 we show the fraction of level shifts that coincided with path changes as seen by each VP for IPv4 (green bar) and IPv6 (red bar). We divide each bar into cases where the level shifts coincided with path changes in HE (solid pattern) and where the level shifts coincided with path changes in other networks (hashed pattern). As expected, *jfk-us* experiences the highest fraction of level shifts that coincide with path changes in HE, over both IPv4 and IPv6. However, other monitors also see a significant fraction of level shifts in IPv6 coinciding with path changes in HE. This is due to the documented dominance of HE in the IPv6 topology, where a large fraction of end-toend paths in IPv6 cross AS6939 [5]. An implication of our observation is that routing dynamics and performance in HE have the potential to affect a large number of end-to-end paths.



Fig. 11. Fraction of level shifts due to path changes. Most of the level shifts due to path changes over IPv6 coincide with path changes in the Hurricane Electric network.

Analysis of level shifts caused by reasons other than forward path changes

Next, we focus on level shifts that do not coincide with path changes, which affect 67.72% of the targets for which we collected traceroute data during the level shifts. Our interest is in characterizing these shifts *spatially* (where do they occur?) and *temporally* (when do they occur?).

Spatial analysis

To infer the location of the RTT increase along the end-toend path, we use traceroutes collected during and before the level shift. We identify the first hop on the forward path that shows an increased RTT (as compared to the RTT at that hop prior to the level shift) as the likely location of the increase. We impose the condition that the elevated RTT must be seen on each hop following the candidate hop. We note that the location of the increase could be at the candidate hop, or at a point on the reverse path from that hop to our VP, due to the fact that traceroute does not give us any information about the reverse path. We investigate the location of the RTT increase both at the AS-level and router-level. At the AS-level, for 18% of the target ASes the RTT increase occurred in either the target or VP network; for the remaining 82% the increase occurred at an intermediate hop on the AS-path. This finding is consistent over both IPv4 and IPv6. At the router-level we find that the increase occurred on inter-domain links for 72% and 78.5% of the target ASes over IPv4 and over IPv6, respectively. Both of these results indicate that most latency increases did not occur in the *last mile*, but rather in transit networks, and predominantly at inter-domain links between networks.

Temporal analysis

We investigate whether there are temporal patterns in the level shifts, e.g., cases where level shifts occur during peak times every day, which could signify congestion [21]. We first construct the set of targets, S_r , for which level shifts did not represent isolated events. We include a target AS in S_r if it experienced level shifts on more than 10 days (approximately 25% of our measurement period). S_r contains 27.65% and 26.17% of the targets in IPv4 and over IPv6, respectively. For each target in S_r we group the level shifts based on the time of day into four time periods, each spanning 6 hours. Figure 12 shows the fraction of level shifts for each target that occurred in a given time bin, for each (monitor, target) pair. Each bar in Figure 12(a) is a target AS from S_r that experienced shifts either over IPv4 or IPv6 (but not both). Figure 12(b) shows targets that experienced shifts over both IPv4 and IPv6, and consequently we use two bars for a single target. In each subfigure we group targets by monitor and within each group by the fraction of level shifts during *peak hours*³ at the monitor's time zone.

We find that 80% of the targets experienced level shifts in three of the 4 time bins. For a small number of targets we observed a strong clustering of the level shifts in certain bins - 19 targets experienced more than 70% of the shifts within the same time interval. Of these, 9 experienced shifts during *peak hours.* We use as case studies some targets that showed strong clustering. AS 6772 (ImproWare AG) experienced 90% of level shifts during peak hours for both IPv4 and IPv6. Router level analysis of the IPv4 and IPv6 paths for this target shows that these paths shared one hop in the target network, which was the location of the RTT increase. For AS 23028, 65% of the shifts over IPv4 occurred during peak hours; the increase occurred at the hop corresponding to the IXP AMS-IX. We find that 7 targets probed by the sql-us monitor experienced level shifts that occurred during business hours (9am-5pm). For each of these targets, the paths over IPv4 and IPv6 traversed an interconnection with Cogent, which was the location of the RTT increase over both IPv4 and IPv6. Prior studies also found evidence of congestion at Cogent interconnects [23], [21].

Impact of path congruency

We take a closer look at targets that experienced repeated level shifts over both IPv4 and IPv6. Our goal is to determine whether shared infrastructure on the paths to these targets caused RTT level shifts over both IPv4 and IPv6. We are able to analyze path congruency for 12 out of 21 targets in this category using DNS data. We find that 4 targets experienced level shifts at different hops on IPv4 and IPv6; consequently the level shifts over IPv4 and IPv6 for these targets were clustered in different time bins. For 8 targets the fraction of IPv4 and IPv6 level shifts is comparable for each time bin; 6 of

³The FCC defines peak hours as between 7pm and 11pm local time [22].



(a) Targets that experience level shifts either over IPv4 or IPv6



(b) Targets that experience level shifts over both IPv4 and IPv6

Fig. 12. Fraction of latency increases that occur during 4 different time intervals per day per target. Most of the targets experienced level shifts across 3 of the 4 time intervals; For 19 of the targets most of the level shifts occur within the same time interval.

these shared a common hop on IPv4 and IPv6 where the RTT increase occurred. For 2 targets the RTT increase occurred on adjacent links. We thus find evidence that shared infrastructure on IPv4 and IPv6 paths can cause correlated increases in both IPv4 and IPv6 RTT.

We take our analysis one step further and investigate the potential role that shared infrastructure could play in influencing IPv4 and IPv6 performance. In this analysis we use the set of all targets that experienced level shifts that did not coincide with forward path changes. DNS names for 128 of the 356 targets in this set contained information we could use to identify common hops on IPv4 and IPv6 paths. We find that 89% of the targets had at least one hop in common on the IPv4 and IPv6 paths; 19% of targets had half of the hops in common; the maximum value of the fraction of common hops was 0.9. Thus, non-negligible fraction of the targets, shared infrastructure had the potential to cause cause correlated performance degradations over both IPv4 and IPv6. We believe that this fraction could increase over time, as routing and peering policies in IPv6 become similar to those in IPv4.

Summary

Our analysis revealed a notable improvement in IPv6 performance as compared to prior measurements in 2012 – in our data IPv6 was faster than IPv4 in 45.24% of cases. Most targets experienced episodes of RTT level shifts over both IPv4 and IPv6, although the total time spent in the elevated RTT state was small. Analysis of level shifts coincident with path changes revealed that a large fraction of these were due to path changes in the Hurricane Electric network, the predominant player in the IPv6 topology. Our analysis of congruency between IPv4 and IPv6 paths reveals a significant potential for shared infrastructure to cause correlated performance degradations in IPv4 and IPv6.

IV. RELATED WORK

The imminent exhaustion of IPv4 addresses has increased the attention on measuring and characterizing IPv6 adoption. In recent years a number of studies have focused on different aspects of this process, in terms of topology, traffic, routing and performance. Sarrar at al. [24] analysez traffic, application mix and tunneled traffic during IPv6 World Day. Their study reported that native IPv6 traffic almost double in terms of traffic volume. Also, the application mix in IPv6 was similar to that in IPv4. Dhamdhere et al. [5] measured whether IPv4 and IPv6 were converging in terms of topology, routing dynamics and performance. They reported that the IPv6 routing system exhibits more pathological behavior, a finding which we confirm in this work. We dig deeper into the causes of pathological routing behavior in IPv6, some of which appear to be due to the relative immaturity and topological sparseness of IPv6. Czyz et al. [6] present measurements of IPv4 adoption along several different axes: address allocation, traffic, DNS, and performance. They report that IPv6 is maturing and growing, although different metrics of assessing adoption vary widely, and geographic differences in adoption persist. These studies, however, did not look into the stability of IPv6 in terms of data plane reachability and performance.

Nikkhah et al. [25] measured IPv6 performance by focusing on web access and found that IPv4 and IPv6 performance were comparable for congruent AS-level paths, a finding that Dhamdhere et al. [5] also confirmed using the Ark measurement infrastructure. Both studies found that IPv4 was in general faster than IPv6. In the first part of the our performance analysis we also use measurements from Ark, and find a significant improvement in IPv6 performance since 2012 – IPv6 is now equally likely to be faster. We also go a step further in our analysis of performance by studying the RTT variation over IPv4 and IPv6 and highlighting episodes of elevated RTTs due path changes and congestion.

A number of studies have characterized BGP churn evolution of the IPv4 routing system [18]. Geoff Huston periodically reports on the evolution of BGP churn in IPv4 and IPv6 through his website and presentations [12]. More recent work [26] compared the BGP churn evolution in the IPv4 and IPv6 routing systems and found that churn in IPv4 and IPv6 grows at the same rate as the underlying topologies. The authors also reported that the IPv6 routing system is in general less stable than IPv4. Our analysis of the IPv4 and IPv6 dynamics confirm these findings, but also reports the main factors that cause these differences.

V. DISCUSSION AND FUTURE WORK

We presented a measurement study of IPv6 stability at the data and control plane, and compared it with its IPv4 counterpart. Our results indicate that the IPv6 routing system is less stable than IPv4. Most routing churn in IPv6 is generated by a small set of unstable prefixes. The IPv6 routing system is characterized by a large fraction of events that lead to transient unavailability of the affected prefix, as opposed to events that lead to a new route in IPv4. We believe that the differences between IPv4 and IPv6 in the control plane w.r.t. highly active prefixes and composition of the event mix are due to the relative immaturity and topological sparseness of IPv6. In the data plane we found that the overall network availability was comparable over IPv4 and IPv6, but unreachability periods were longer over IPv6 than IPv4. IPv6 performance (in terms of RTT) was comparable to that over IPv4, a notable shift from the results of Dhamdhere et al [5] who found that IPv4 was faster than IPv6 in 78% of cases.

In our analysis of data-plane stability we found that most episodes of elevated RTTs were not due to path changes but possibly due to congestion. We found evidence of congestion on both IPv4 and IPv6; moreover, correlated RTT level shifts on IPv4 and IPv6 could be attributed to shared infrastructure. It is likely that IPv4 and IPv6 paths will become more congruent as operators establish peering parity [5]. Consequently, congestion on one path would affect the other, potentially limiting the benefits of optimizations such as using multipath TCP with IPv4 and IPv6 addresses on the same interface [27]. We note that we are in the middle of a large-scale technology transition with IPv6 deployment, and several factors affecting our measurements (e.g., topological density of IPv6, peering parity between IPv4 and IPv6, shared infrastructure) will continue to evolve. We thus believe there is value in periodic reassessments of IPv6 readiness w.r.t. stability and performance.

We encountered significant challenges in determining whether an IPv4 path and an IPv6 path shared some or all hops at the router-level. We used DNS data to infer shared infrastructure, but this method has limited applicability as it depends on operators inserting relevant hints in DNS names. No method exists to infer whether IPv4 and IPv6 interfaces are on the same router (although Beverly et al. [28] present a technique applicable for servers); this is a direction we plan to pursue in future work.

Our study of data plane availability revealed episodes of network unreachability and latency increases due to congestion on IPv4 and IPv6 paths. Our ongoing work is focused on studying episodes of correlated unreachability (multiple VPs losing reachability to the same target networks) and localizing congestion to a set of candidate congested links. To this end we are devising an experiment wherein we probe a set of targets from the Alexa list from approximately 50 dual-stacked Ark VPs. The collected data will allow us to measure episodes of unreachability from all or a subset of VPs to the target ASes, and enable the use of tomographic techniques to isolate the cause of latency increases on IPv4 and IPv6 paths.

Acknowledgments

We thank Kc Claffy and the anonymous reviewers for their constructive comments. This work was supported by the U.S. NSF grants CNS-1111449 and CNS-1528148 and the Norwe-gian Research Council grants 209954/O70 and 240850/O70.

REFERENCES

- [1] APNIC, "APNIC IPv4 Address Pool Reaches Final /8," Apr. 2011, http://www.apnic.net/publications/news/2011/final-8.
- [2] RIPE NCC, "Allocations from the Last /8," Aug. 2010, http://www.ripe. net/ripe/policies/proposals/2010-02.
- [3] LACNIC, "No More IPv4 Addresses in Latin America and the Caribbean," June 2014, http://www.lacnic.net/en/web/anuncios/2014no-hay-mas-direcciones-ipv4-en-lac.
- [4] ARIN, "ARIN enters phase four of the IPv4 countdown plan," Apr. 2015, https://www.arin.net/announcements/2014/20140423.html.
- [5] A. Dhamdhere, M. Luckie, B. Huffaker, K. C. Claffy, A. Elmokashfi, and E. Aben, "Measuring the Deployment of IPv6: Topology, Routing and Performance," in ACM SIGCOMM IMC, Nov. 2012.
- [6] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey, "Measuring IPv6 Adoption," in ACM SIGCOMM, Aug. 2014.
- [7] Google, "IPv6 Adoption," July 2015, https://www.google.com/intl/en/ ipv6/statistics.html.
- [8] Geoff Huston, "IPv6: IPv6 / IPv4 Comparative Statistics," http://bgp. potaroo.net/v6/v6rpt.html.
- [9] 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 ACM SIGCOMM, Aug. 2006.
- [10] David Meyer, "University of Oregon Route Views Project," 2014, http: //www.routeviews.org/.
- [11] CAIDA, "Archipelago Measurement Infrastructure," http://www.caida. org/projects/ark/.
- [12] G. Huston, "BGP Analysis Report," http://bgp.potaroo.net/.
- [13] 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 *NSDI*, May 2005.
- [14] A. M. Law and D. M. Kelton, Simulation Modeling and Analysis. McGraw-Hill Higher Education, 1999.
- [15] L. Cittadini, W. Muhlbauer, S. Uhlig, R. Bush, P. Francois, and O. Maennel, "Evolution of Internet address space deaggregation: Myths and reality," *IEEE JSAC*, 2010.
- [16] R. Oliveira, R. Izhak-Ratzin, B. Zhang, and L. Zhang, "Measurement of highly active prefixes in BGP," in *IEEE GLOBECOM*, Nov 2005.
- [17] N. Kushman, S. Kandula, and D. Katabi, "Can You Hear Me Now?! It Must be BGP." in ACM SIGCOMM CCR, Mar. 2007.
- [18] C. Labovitz, G. R. Malan, and F. Jahanian, "Origins of internet routing instability," in *IEEE INFOCOM*, Mar. 1999.
- [19] Alexa, "Alexa Top Sites," July 2015, http://www.alexa.com/topsites.
- [20] S. Rodionov, "A sequential algorithm for testing climate regime shifts," Geophysical Research Letter, 2004.
- [21] M. Luckie, A. Dhamdhere, D. Clark, B. Huffaker, and kc claffy, "Challenges in Inferring Interdomain Congestion," in ACM SIGCOMM IMC, Nov. 2014.
- [22] FCC, "Measuring Broadband America," 2014, https://www.fcc.gov.
- [23] MLAB, "ISP Interconnection and its Impact on Consumer Internet Performance," 2014, http://www.measurementlab.net/static/observatory/ M-Lab_Interconnection_Study_US.pdf.
- [24] N. Sarrar, G. Maier, B. Ager, R. Sommer, and S. Uhlig, "Investigating ipv6 traffic: What happened at the world ipv6 day?" in *PAM*, Mar. 2012.
- [25] M. Nikkhah, R. Guérin, Y. Lee, and R. Woundy, "Assessing ipv6 through web access: a measurement study and its findings," in ACM CONEXT, Dec. 2011.
- [26] A. Elmokashfi and A. Dhamdhere, "Revisiting BGP Churn Growth," in ACM SIGCOMM CCR, Jan. 2014.
- [27] I. Livadariu, S. Ferlin, O. Alay, T. Dreibholz, A. Dhamdhere, and A. Elmokashfi, "Leveraging the IPv4/IPv6 Identity Duality by using Multi-Path Transport," in *IEEE GIS*, Apr. 2015.
- [28] R. Beverly and A. Berger, "Server Siblings: Identifying Shared IPv4/IPv6 Infrastructure via Active Fingerprinting," in PAM, Mar. 2015.