Understanding the Reachability of IPv6 Limited Visibility Prefixes

Andra Lutu[1,2], Marcelo Bagnulo[2], Cristel Pelsser[3], and Olaf Maennel[4]

[1] Institute IMDEA Networks, Spain
[2] University Carlos III of Madrid, Spain
[3] IIJ Innovation Institute, Japan
[4] Loughborough University, UK

Abstract. The main functionality of the Internet is to provide global connectivity for every node attached to it. In light of the IPv4 address space depletion, large networks are in the process of deploying IPv6. In this paper we perform an extensive analysis of how BGP route propagation affects global reachability of the active IPv6 address space in the context of this unique transition of the Internet infrastructure. We propose and validate a methodology for testing the reachability of an IPv6 address block active in the routing system. Leveraging the global visibility status of the IPv6 prefixes evaluated with the BGP Visibility Scanner, we then use this methodology to verify if the visibility status of the prefix impacts its reachability at the interdomain level. We perform active measurements using the RIPE Atlas platform. We test destinations with different BGP visibility degrees (i.e., limited visibility - LV, high visibility - HV and dark prefixes). We show that the IPv6 LV prefixes (v6LVPs) are generally reachable, mostly due to a less-specific *HV* covering prefix (v6HVP). However, this is not the case of the dark address space, which, by not having a covering v6HVP is largely unreachable.

1 Introduction

The fundamental task envisioned for the Internet is to provide reachability for every node attached to the network. The Border Gateway Protocol (BGP) is currently responsible for the exchange of network reachability information and the selection of paths according to specified routing policies. By tweaking the BGP configurations, the network operators are able to express their interdomain routing preferences, designed to accommodate myriad economic and technical goals. However, these routing policies can at time affect the global visibility of a certain prefix, both willingly or unknowingly/accidentally [13]. Given the complex interactions between policies in the Internet, the origin AS by itself cannot ensure that only by configuring a routing policy it can also achieve the anticipated results [7]. Consequently, policies may affect the propagation of routes, making some paths unavailable at a global level, and sometimes preventing a prefix to be learned altogether. Moreover, the definition of routing policies is a complicated process, involving a number of subtle tuning operations prone to errors.

Over the last few years, much has been said about global connectivity (or the lack of it) in the IPv6 Internet due to the routing policies of a few Autonomous Systems(ASes) (e.g., [2]). In this paper, we aim to establish if IPv6 prefix visibility at the interdomain

M. Faloutsos and A. Kuzmanovic (Eds.): PAM 2014, LNCS 8362, pp. 163–172, 2014.

level has an impact on the reachability of the address space advertised in the Internet. Using the interdomain route propagation process reflected in the global routing tables as an expression of routing policy interaction, we introduce the concept of **Limited-Visibility Prefix (LVP)**. We define *LVPs* as stable long-lived Internet routes that are advertised by at least two different ASes, but visible in *less than* 95% of all the global routing tables analyzed. Though some legitimate routing policies of an AS limit the visibility of its prefixes in the Internet, the latter can also stem from human operator errors or unpredicted interplay with the external netting of otherwise correctly defined routing policies. Contrariwise, we define the **High-Visibility Prefixes (HVPs)** as the set of prefixes that are propagated in *at least* 95% of all the available global routing feeds. We also identify the **Dark Prefixes (DPs)** [9], which represent the subset of *LVPs* that are not covered by any *HV* less-specific prefix. These prefixes represent address space that, in the absence of a default route, may not be globally reachable. We use the BGP Visibility Scanner [11] to evaluate the visibility status of the IPv6 prefixes announced in the global routing system. The tool uses the routing data retrieved from the RIPE RIS and RouteViews projects to performs a differential analysis to retrieve *LVPs* on a daily basis, which are then made available on-line.

We further focus on measuring the reachability of the prefixes in all of the three above-mentioned sets of prefixes, i.e. HVP, LVP and DP. We propose a methodology for testing the reachability of an IPv6 *prefix*, which relies on the use of traceroute probes to test the destination prefix. We calibrate the proposed measurement methodology by testing a large set of so-called *anchor* prefixes, which we know a priori to contain at least one reachable address. We compile a set of approximatively 70,000 such prefixes, which we test from a major Japanese ISP using different traceroute approaches. We then apply the proposed methodology from multiple vantage points in the Internet, including 100 RIPE Atlas active probes. We thus show that the IPv6 LVPs (v6LVPs) are generally reachable, mostly due to the less-specific *HV* covering prefixes. However, this is not the case of the dark address space, which is largely unreachable.

2 The BGP Visibility Scanner for IPv6

In this section we describe the BGP Visibility Scanner - a tool we propose for identifying *LVPs* at the interdomain level. We have publicly released an initial version of the BGP Visibility Scanner[1] in November 2012, allowing any network operator to check if the AS originates LVPs. The earlier version of this tool is documented in [11]. Since it became operational, the tool has been well received by the operational community and it still attracts a large amount of attention and feedback. The methodology used for the BGP Visibility Scanner is structured in three steps: First, we retrieve the raw BGP routing data at two different times every day. Second, we clean the raw data in order to obtain the Global Routing Tables (GRTs), by applying two different cleansing filters. Third, we verify in two sub-steps the visibility of each prefix within the sample of identified GRTs using the Visibility Scanner Algorithm. We now further expand on the steps we take in order to retrieve, parse, clean and process the raw BGP routing data to distinguish the set of *LVPs* and *DPs*.

[1] The BGP Visibility Scanner is publicly available at **visibility.it.uc3m.es**

2.1 Retrieving and Refining the Raw Routing Data

We work with publicly available routing data, retrieved from the RouteViews and RIPE RIS projects. These two repositories periodically receive BGP routing table *snapshots*, i.e. one time instance of a routing table, from over 400 active BGP peers for both IPv4 and IPv6. In this first step of the methodology, we retrieve the publicly available routing data. We choose to do so at two different times during the day, i.e., at 8h00 and 16h00. We process these two different snapshots per day in order to be certain that we only work with routes that are stable expression of routing policies at the interdomain level.

In the second step of our methodology, we parse the raw data in order to identify what we define to be *global routing tables (GRTs)*. Only by comparing the GRTs from the BGP peers, we can further identify the sets of *HV* and *LV* prefixes. For the purpose of this paper, we loosely define the GRT as the entire routing table provided by a *Default Free Zone (DFZ)*[2] network to its customers requesting a full routing feed. The routing table maintained in one of the so-called DFZ routers is commonly known as the *global routing table*. Realistically speaking though, due to the current operational status of the Internet routing, such a GRT of the BGP routing is an idealized concept. However, Internet Service Providers (ISPs) do maintain their own version of the *global routing table*, which is propagated to customer networks upon request. This is not a formal definition, but it properly captures the main idea of the kind of data we require.

In order to identify the feeds which constitute a GRT, the primary characteristic of the routing feeds on which we focus is the actual size of the routing table snapshot. Based on the BGP Analysis Report [1], we consider that *a complete routing feed from a monitor should have no less than* ***10,000*** *IPv6 routing entries*. Consequently, we check over 200 routing feeds collected from the two repositories, and keep approximatively 110 BGP feeds that comply with the imposed lower-limit of prefix number.

Additionally, we perform a couple of "sanitary" checks on the data contained in the identified GRTs, in order to further discard the information that is of no interest for our study. Hence, we apply the *bogon filter* on all the GRTs. Bogon prefixes are a class of routes that should never appear in the Internet. Bogons are defined as *Martians*, representing reserved and local address space or *Fullbogons*, which include the IP space that has been allocated to a Regional Internet Registry (RIR), but has not been assigned by that RIR to an actual Internet Service Provider (ISP) or other end-user. We use the periodically updated filters from The Bogon Reference [4] in order to make sure that we eliminate any possible bogon route included in the GRTs.

2.2 The Visibility Scanner Algorithm: The Labeling Mechanism

We now apply the **Visibility Scanner Algorithm** for identifying prefixes with *stable* limited visibility in the Internet. It is important to filter out the cases of limited visibility caused by other factors unrelated to routing policies, e.g. BGP convergence or internal routes advertised only to the collector. In order to discard any internal paths leaking towards the collectors, we remove all the routes learned from only one monitor which

[2] Conceptually, the so-called *Default Free Zone (DFZ)* represents the set of BGP-speaking routers that do not need a default route to forward packets towards any destination in the Internet.

is also the route originating AS. Next, in order to further avoid that the converging prefixes emerge as false positive limited visibility prefixes in our results, we analyze two samples taken 8-hours apart of routing data. We evaluate the *visibility degree* at every sampling moment and assign *visibility labels* based on our results. We define the *visibility degree* as the number of GRTs which contain (i.e., "see") a certain prefix, and the *visibility label* as the visibility status of each prefix, i.e. *LV* for Limited Visibility and *HV* for High Visibility. We then compare the per-prefix visibility of each prefix, as observed at each sampling time and apply the prefix visibility prevalence sieve.

The Labeling Mechanism: Based on the visibility degree of the prefixes at each of the two sampling moments (i.e. 08h00 and 16h00), we assign a *visibility labels* at each sampling moment to all the prefixes discovered. *We define Limited Visibility prefixes as prefixes present in less than* 95% *of the active monitors at a sampling time.* Otherwise, the prefixes are defined as High Visibility prefixes. Ideally, a *HV* prefix should be contained in absolutely all the routing tables contained in the sample. The choice of the 95% allows for a 5% error in the sampling, including possible glitches that may appear in the data. Moreover, according to our threshold sensitivity analysis, we find that the set of *LVPs* is not particularly sensitive to the values of the prevalence sieve threshold.

Visibility Label Prevalence Sieve: When deriving the final per-day visibility label, we account for the dynamics of a prefix in time. The high visibility of a prefix in at least one monitor sample hints the fact that the route could reach all the observed ASes. Should this change during the analyzed time, it might be a cause of, for example, topology changes or failures. Therefore, we consider that *the HV label always prevails*, i.e. if a prefix is tagged as *HV* in one of the samples, it is tagged as *HV* in the final set.

Otherwise, when no *HV* label is tagged, we analyze the cases of *LV* prefixes emerging in our results. If a prefix appears only at one sampling time and it is tagged as *LVP*, this might be a sign that the prefix is in the process of being withdrawn or, contrariwise, in the process of converging after just being injected. These particular routes cannot be qualified within our study, thus we filter out any prefix with only one label in a day and that label being *LV*. The only case where a prefix has limited visibility and mark it accordingly, is when the two labels assigned at each sampling time are both *LVP*.

Identifying Dark Prefixes: Once we have identified the two main sets of prefixes, i.e. the *LVPs* and the *HVPs*, we can now identify the set of Dark Prefixes. For each of the prefix in the LVP category, we build the covering trie of less specific HV prefixes, from which we ultimately retrieve its root prefix (i.e. the smallest covering HV prefix). In the eventuality of not identifying any such globally visible less-specific prefix, we mark the LV prefix as *Dark* and continue our analysis.

3 The IPv6 Limited Visibility Prefixes

We collect more than *500* routing feeds on a daily basis, for each of the two different sampling moments, i.e., 8h00 and 16h00. After the *cleansing process*, we distinguish, in average, *110 GRTs* injected to the public repositories by unique ASes. We then compare the content of the 110 GRTs in order to identify the LVPs. In rough numbers, the daily overall total number of prefixes identified is approximatively *16,500 prefixes*. Out of these, on average *150 prefixes* are singled out as leaked internal routes and, consequently, discarded from our analysis. Furthermore, we remove the converging routes

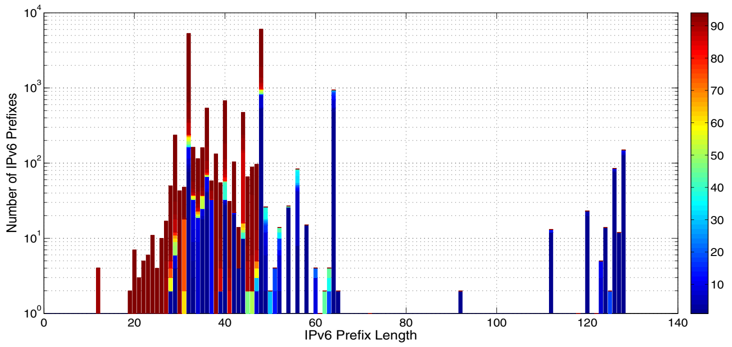

Fig. 1. Distribution of IPv6 prefixes on prefix length. The bars are color-coded to show the visibility degree of the prefixes: from dark blue for LV, going to dark red for HV.

that may otherwise emerge as limited visibility in the visibility scanner. This incurs the elimination of about *10* additional prefixes in average. For the remaining prefixes we continue our visibility analysis and assign LV/HV visibility tags.

Finally, we identify an average of *3,500 IPv6 prefixes* that are tagged *LVP* and approximatively *12,500* prefixes marked *HVP*. Therefore, 20% of all the IPv6 prefixes identified from the analyzed routing tables are LVPs. This is consistent with the result for the IPv4 LVPs, where out of all the prefixes learned, 20% have limited visibility [11]. When checking how the two sets of prefixes overlap, we find that there are more than *500* LV prefixes without a covering HVP, which we mark *DP*. This represents approximatively 14% of the whole set of v6LVPs and 3.75% of the v6HVP set. When comparing with the situation in IPv4, where in average only 3% of the LVPs (and 0.6% of the HVPs) are marked as *dark*, we conclude that we have almost 5 times more IPv6 dark address space. This is relevant because these prefixes may have limited reachability. We have observed more than 13% of all IPv6 active ASes inject LVPs, while less than 5% of all IPv6 active ASes originate DPs. In IPv4, we see that 9% of all ASes originate LVPs, while only 2% are also injecting DPs. This result further hints the early stages of development of the IPv6 architecture, previously established in [6].

For the rest of the analysis we perform in the paper, we use the LVP dataset derived on the 8th of August, 2013. The dataset consists of 12,621 v6HVPs and 3,444 v6LVPs, out of which 473 are v6DPs. Figure 1 depicts the distribution of IPv6 prefixes per prefix length, color-coded to match the visibility degree of the prefixes in question. All the prefixes with a length longer than /48 are labeled as v6LVPs by the BGP Visibility Scanner i.e. /48's do not propagate globally in the IPv6 routing system. This is consistent with the status in IPv4, where every prefix more-specific than /24 is labeled LVP.

4 Traceroute Probing for Reachability

In this section, we try to verify if the limited visibility of such prefixes have an actual impact in the reachability of the addresses in them.

We propose a methodology for determining if a prefix is reachable from a given vantage point in the Internet. The challenge for doing this with IPv6 prefixes is that it is not a simple task to find an address that is actually allocated to a host in a given prefix. The idea we put forward to probe the reachability of a prefix is to perform traceroute towards a random address within the prefix and check if the last node responding to the traceroute belongs to the origin AS of the target prefix or to one of the Internet providers of the origin AS, as observed in the BGP AS-Path. In other words, the methodology we propose for determining the reachability of a prefix is as follows. We send a traceroute probe towards a random address within the target prefix. We say that the prefix is reachable if :

1. The traceroute probe reaches the network to which the prefix has been allocated.
2. The traceroute probe traverses the second-last[3] AS along the BGP AS-Path for the target prefix.

We consider this latter hypothesis because there may be cases where, even if the probe does reach its destination, it might happen that the origin AS of the source IP for the last ICMP message received is actually the transit provider of the target AS. This happens because it is a common operational practice that ASes use addresses from their providers for their transit links. As a result, the router within the destination network that issues the last message of the traceroute process will do so using an source address from its ISP's address space. We do acknowledge that this may also be due to reachability problems in the last hop, which our methodology is unable to distinguish.

4.1 Traceroute Probing Approach

We begin by discussing the different traceroute probing methods and how we select the most suited approach. Traceroute is one of the most widely used network measurement tools, useful both to network operators and researchers. The original traceroute tool [8] sends UDP probes and it will be our *default* measurement approach. We further refer to this test as *default UDP traceroute probing*. The major weakness of the default UDP traceroute is that, in the current operational routing system, firewalls are likely to filter the probes sent to these unlikely ports, thus impacting the quality of the measurements. In order to avoid this problem, several other approaches are available. We use a modified UDP traceroute method which, instead of using high-numbered unlikely ports, sends packets on port 53. We further refer to this probing method as "UDP traceroute". A second approach we use is the so-called *ICMP traceroute*, which uses *ICMP echo request* instead of UDP probes. The last approach we use is *TCP traceroute*, which employs TCP SYN probes to port 80. The advantage of this approach is that the probes cannot be easily distinguished from normal requests to web servers, so they are less likely to be discarded along the path.

We establish which of the above-mentioned traceroute approaches is the most efficient by testing the status of a large set of control IPv6 addresses with all the listed probing methods. We use a set of 70.000 IPv6 addresses which are known to be reachable. This is made up of addresses from many sources, including DNS entries, Alexa's

[3] Usually, in the BGP AS-Path the last hop represent the origin AS of the prefix, while the first hop represents the AS whose routing table we analyze. Following this order, the second-last hop (2LH) in the AS-Path corresponds to the transit provider of the origin AS.

top sites, and several other sources. We check the reachability status of these 70,000 IPv6 addresses from a machine inside a major Japanese ISP's network. We do so by using all the above-mentioned traceroute probing approaches. Our results show that the most efficient probing method is *ICMP traceroute*, which successfully reached 99% of all the 70,000 probable IP addresses. Consequently, the traceroute probing method we further employ in our study is the ***ICMP traceroute***. This results is consistent with the observations of Luckie et al. in [10].

4.2 Validating the Measurement Methodology

We validate our methodology by testing a set of reachable IPv6 prefixes, which are known to contain at least one reachable address. The way we do this is by tacking a reverse engineering approach. For each of the previously identified 70,000 reachable IPv6 addresses, we map the covering prefix installed in the BGP routing tables. We use public routing data information to determine the most-specific prefixes covering each of these reachable addresses. The set of prefixes determined represents address space known to contain at least one address which is successful to ICMP traceroute probing. These prefixes form the target set of prefixes which we use for validation.

We start by sending ICMP traceroute probes from a machine within the major Japanese ISP towards a **randomly selected IPv6 address** within each of the prefixes determined above. According to the proposed methodology, we consider that the traceroute probe reached its destination when the traceroute probe traverses either the origin AS of the destination address, either the second-last AS appearing in the BGP AS-Path towards the target prefix. In order to identify the 2LH towards a prefix, we analyze the AS-Path information in the BGP routing table of the AS from which we are generating the traceroute messages, i.e., the major Japanese ISP.

After parsing the results of our traceroute tests, we learn that the ICMP traceroute probes successfully reached more than 96% of these *a-priori* reachable prefixes. Consequently, the methodology we propose is able to identify with 96% accuracy the reachability status of an IPv6 prefixes. For the other 4% of prefixes, our methodology is unable to determine reachability. This may be due to several reasons, including ICMP filtering or routers silently discarding packets.

5 Reachability Measurements and Results

5.1 Local Reachability Measurements

In order to establish the reachability for prefixes with the three different classes of interdomain visibility, we perform ICMP traceroute probing from a machine inside a major Japanese ISP's network. Regarding the target address space to be tested, we first re-define the set of LVPs and DPs *locally*, by analyzing only the routing table snapshot of the Japanese ISP. We are thus able to identify a total of 13,195 IPv6 prefixes present in the routing table, which we further label as High-Visibility Prefixes. These prefixes may not be globally High-Visibility, since there may be other routing tables not "seeing" some of these prefixes. We label all the rest of prefixes learned from the rest of the routing tables collected from the public repositories as Limited Visibility, which reach a total number of 2,359 prefixes. In order to check if any of the Limited Visibility prefixes are in fact Dark Prefixes from the point of view of the ISP, we check which v6LVPs have a less-specific v6HVP in the ISP's routing table to offer global reachability. We are thus able to single out a total number of 511 Dark Prefixes.

From the results of the measurements we learn that, in the case of the locally-defined v6HVPs, 92% of the target high-visibility prefixes are reachable from the ISP's network. This is consistent with the precision of our methodology, so we cannot make claims about reachability problems in the HVP set. In the case of the locally-defined v6LVPs which have a covering high-visibility IPv6 prefix (i.e., they are not dark), we observe that 94% of the prefixes are reachable from the Japanese ISP's network. Likewise, this is consistent with the precision of our tool so we cannot make any claims about reachability problems in the LVP set. We next evaluate the reachability status for the DPs and we learn that more than 95% of these prefixes traceroute ended in a network or destination unreachable error messages. Consequently, less than 5% of the dark address space is reachable from the Japanese ISP. We can then claim that within the precision of our methodology, DPs do present reachability problems.

5.2 RIPE Atlas Measurements and Results

Previously, we have seen that the non-dark LVPs defined for the Japanese ISP do not exhibit reachability issues, due to the covering HVPs. However, this was not the case for the local dark address space, which has less than 5% reachability. In this section, we use the RIPE Atlas platform [3] to run **larger-scale measurements for characterizing the reachability of the global dark address space**.

We zoom out from the previous localized analysis of reachability, and test the reachability of the DPs from 100 different probes active in the RIPE Atlas platform. We run the measurements both towards the globally defined set of IPv6 dark prefixes, i.e. the 473 v6DPs derived from analyzing 110 BGP routing tables, and also towards the set of IPv4 dark prefixes, i.e., 3,200 v4DPs derived from analyzing 154 global BGP routing tables. We send ICMP traceroute probes towards a random target address within each of the v6 and v4 DPs.. We proceed to verifying the reachability results in accordance with the methodology specified in Section 4. Point 2) of the proposed methodology requires to verify if the traceroute probe traverses the provider of the origin AS for the target prefix. As opposed to the case of the major Japanese ISP for which we have the BGP routing table to analyze, we now do not have access to the BGP routing tables corresponding to the 100 Atlas probes used. In order to overcome this issue, we build a set of *probable* second-last hops which may be traversed towards all the possible destination ASes. We do so by analyzing all the available routing tables from all the ASes active in RIPE RIS and/or Routeviews, and monitoring the ASes appearing as 2LHs towards every active destination AS. Thus, we state that the target prefix is reachable if *the traceroute probe traverses* ***any*** *of the probable second-last ASes to the origin AS of the target prefix.*

After processing all the traceroute results from each of the 100 probes towards a Dark Prefix, we conclude that the average reachability degree for a v6DP is of 46.5%, whereas for v4DPs this decreases to only 17.4%. To further understand this result, we

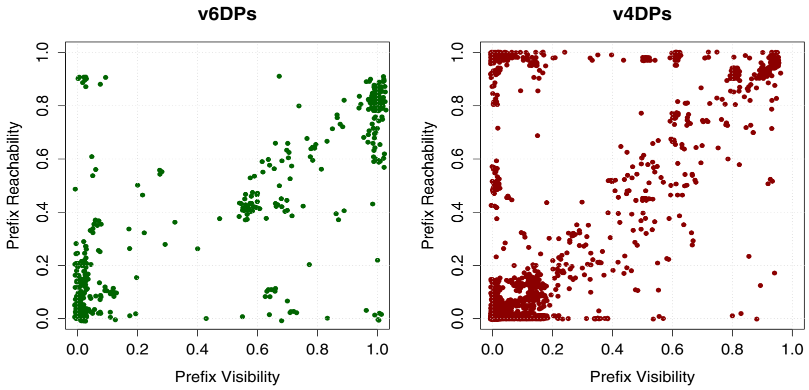

Fig. 2. Scatterplot of reachability probability against the DP's visibility, for v6DPs and for v4DPs

verify how the DP reachability correlates with the visibility degree of a DP. We show in Figure 2 the scatterplots both for IPv6 and IPv4 DPs' reachability against their visibility within the corresponding sample of ASes analyzed. We observe that for the v6DPs, depicted in the left-side plot, there is a stronger correlation between reachability and visibility than for the v4DPs. This happens because, for the v4DPs, we see a high number of prefixes with very limited visibility, but which are highly reachable from the sample of 100 probes chosen. We observe that in the v4 plot from Figure 2 there are approximatively 8% of IPv4 prefixes with visibilities smaller than 0.2 and reachability larger than 0.2. As previously noted in [5], this may be due to default routing in IPv4. In [12], the authors explain many of the real-life operational reasons for which this type if v4DPs emerge in the Internet. For example, we observe in the lower-left corner of the IPv4 plot in Figure 2 a very large number of v4DP (approximatively 72% of all the v4DPs) with a reduced visibility degree and a corresponding low reachability degree. These v4DPs may be route leaks which, as we learn from [12], often occur in the Internet. Consequently, the lack of reachability observed for v4DPs is largely explained by the fact that these prefixes are unintended to be visible in the Internet to begin with. At the same time, even if the v6DPs do not follow the known symptoms of route leaks or anomalies previously learned from the IPv4 cases, they do struggle with important lack of reachability. This further supports the hypothesis that, while in IPv4 the DPs are in majority results of mistakes or slips in the network configuration, for IPv6 we understand this as a side-effect of the early stages of development of the network.

6 Conclusions

In this paper, we perform an extensive analysis of how BGP route propagation affects global reachability of the active IPv6 address space, in the context of IPv6 penetration growing in the Internet.We proposed a methodology to measure the reachability status of the active LVP IPv6 prefixes, which represent address space that is not present in all the global routing tables of the operational networks. We find that, while the fraction of limited visibility address space is similar in the IPv4 and the IPv6 Internet (about

20% of the prefixes), the proportion of dark address space in the IPv6 Internet is significantly larger than in the IPv4 Internet (3.75% versus 0.6%). We find an important correlation between the limited visibility of a dark IPv6 prefix and its reduced reachability. Moreover, while the IPv4 dark address space can be largely explained as route leaks or mistakes, this is not valid for the v6DPs. We believe that this is a serious problem for the IPv6 Internet, as limited reachability of a non-negligible set of prefixes undermines the global connectivity of the Internet. In future work we expect to investigate the reasons behind the large amount of dark address space in the IPv6 Internet.

Acknowledgements. This work was partially supported by the European Community's Seventh Framework Programme (FP7/2007-2013) grant no. 317647 (Leone). We would like to thank Emile Aben for the discussions which helped improve this work and for his support while working with the Atlas platform.

References

1. BGP Routing Table Analysis Report, `http://bgp.potaroo.net/`
2. IPv6 internet broken - NANOG mailing list, `http://mailman.nanog.org/pipermail/nanog/2009-October/013997.html`
3. Ripe Atlas, `https://atlas.ripe.net/`
4. The Bogon Reference, `http://www.cymru.com/BGP/bogons.html`
5. Bush, R., Maennel, O., Roughan, M., Uhlig, S.: Internet optometry: Assessing the broken glasses in internet reachability. In: Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference, IMC 2009 (2009)
6. Dhamdhere, A., Luckie, M., Huffaker, B., Claffy, K., Elmokashfi, A., Aben, E.: Measuring the deployment of ipv6: topology, routing and performance. In: Proceedings of the 2012 ACM Conference on Internet Measurement Conference, IMC 2012 (2012)
7. Griffin, T., Huston, G.: BGP Wedgies, RFC 4264 (2005)
8. Jacobson, V.: Traceroute, `ftp://ftp.ee.lbl.gov/traceroute.tar.gz`
9. Labovitz, C., Ahuja, A., Bailey, M.: Shining Light on Dark Address Space. Tech. Rep. TR-2001-01, Arbor Netwoks, Ann Arbor, Michigan, USA (November 2001)
10. Luckie, M., Hyun, Y., Huffaker, B.: Traceroute probe method and forward ip path inference. In: Proceedings of the 8th ACM SIGCOMM Conference on Internet Measurement (2008)
11. Lutu, A., Bagnulo, M., Maennel, O.: The BGP Visibility Scanner. In: IEEE Global Internet Symposium, GI 2013 (April 2013)
12. Lutu, A., Bagnulo, M., Cid-Sueiro, J., Maennel, O.: Separating wheat from chaff: Winnowing unintended prefixes using machine learning. In: Proceedings of 33rd IEEE International Conference on Computer Communications, IEEE INFOCOM 2014 (to appear, 2014)
13. Zhang, K., Yen, A., Zhao, X., Massey, D., Wu, S.F., Zhang, L.: On detection of anomalous routing dynamics in BGP. In: Mitrou, N.M., Kontovasilis, K., Rouskas, G.N., Iliadis, I., Merakos, L. (eds.) NETWORKING 2004. LNCS, vol. 3042, pp. 259–270. Springer, Heidelberg (2004)