



# The (Thin) Bridges of AS Connectivity: Measuring Dependency Using AS Hegemony

Romain Fontugne<sup>1(✉)</sup>, Anant Shah<sup>2</sup>, and Emile Aben<sup>3</sup>

<sup>1</sup> IIJ Research Lab, Tokyo, Japan  
romain@iij.ad.jp

<sup>2</sup> Colorado State University, Fort Collins, USA  
akshah@cs.colostate.edu

<sup>3</sup> RIPE NCC, Amsterdam, Netherlands  
emile.aben@ripe.net

**Abstract.** Inter-domain routing is a crucial part of the Internet designed for arbitrary policies, economical models, and topologies. This versatility translates into a substantially complex system that is hard to comprehend. Monitoring the inter-domain routing infrastructure is however essential for understanding the current state of the Internet and improving it. In this paper we design a methodology to answer two simple questions: Which are the common transit networks used to reach a certain AS? How much does this AS depend on these transit networks? To answer these questions we digest AS paths advertised with the Border Gateway Protocol (BGP) into AS graphs and measure node centrality (i.e. the likelihood of an AS to lie on paths between two other ASes). Our proposal relies solely on the AS hegemony metric, a new way to quantify node centrality while taking into account the bias towards the partial view offered by BGP. Our analysis using 14 years of BGP data refines our knowledge on Internet flattening but also exhibits the consolidated position of tier-1 networks in today's IPv4 and IPv6 Internet. We also study the connectivity to two content providers (Google and Akamai) and investigate the AS dependency of networks hosting DNS root servers. These case studies emphasize the benefits of our method to assist ISPs in planning and assessing infrastructure deployment.

## 1 Introduction

Networks connected to the Internet are inherently relying on other Autonomous Systems (ASes) to transmit data. To determine the path of ASes to go from one place to another, the Internet relies solely on the Border Gateway Protocol (BGP). Computed AS paths are the result of an involved process that considers various peering policies set by each connected AS. BGP exposes only paths that are favored by ASes hence concealing peering policies and the exact routing process. However, as the connectivity of a network depends greatly on the connectivity of other ASes, operators need to clearly understand ASes that are

crucial to their networks. Identifying these AS interdependencies facilitates decisions for deployments, routing decisions, and connectivity troubleshooting [17].


In this paper we aim at estimating the AS interdependencies from BGP data. We devise a methodology that models AS interconnections as a graph and measure AS centrality, that is the likelihood of an AS to lie on paths between two other ASes. We identify in Sect. 2 shortcomings of a classical centrality metric, Betweenness Centrality (BC), when used with BGP data. From these observations we employ a robust metric to estimate AS centrality, called AS hegemony (Sect. 3). We demonstrate the value of the proposed method with 14 years of BGP data (Sect. 4). Overall we found that AS interdependencies in IPv4 are decreasing over time which corroborate with previous observations of the Internet flattening [3]. But we also found that the important role played by tier-1 ISP is reinforced in today’s Internet. The Internet flattening for IPv6 is happening at a faster rate, but we found that Hurricane Electric network is utterly central for the last 9 years. We also investigated the AS dependency of two popular networks, Akamai and Google, showing that their dependency to other networks is minimal although their peering policies are completely different. Finally, we look at two networks hosting DNS root servers and show how recent structural changes to these root servers have affected their AS dependencies.

We make our tools and updated results publicly available [1] hence network operators can quickly understand their networks’ AS dependency.

## 2 Background

**Related Work:** The essence of this work is the estimation of AS centrality in AS graphs. In the literature AS centrality is commonly measured using Betweenness Centrality (BC). This is one of the basic metric used to characterize the topology of the Internet [12, 18]. It was also applied for similar motivation as ours. Karlin et al. [9] consider Internet routing at the country-level to investigate the interdependencies of countries and identify countries relying on other countries enforcing censorship or wiretapping. BC is also used to identify critical ASes for industrial and public sectors in Germany [17]. Similarly, Schuchard et al. [15] select targets for control plane attacks using a ranking based on BC. Finally, researchers have also applied BC to detect changes in the AS-topology. For example, Liu et al. [11] employ BC to monitor rerouting events caused by important disruptive events such as sea cable faults. Following past research, we initially conducted our experiments using BC but faced fundamental shortcomings due to the incomplete view provided by BGP data. To introduce these challenges let’s first review BC.

**Betweenness Centrality:** BC is a fundamental metric that represents the fraction of paths that goes through a node. Intuitively one expects high BC scores for transit ASes as they occur on numerous AS paths, and low BC scores for stub ASes. Formally, for a graph  $G = (V, E)$  composed of a set of nodes  $V$  and edges  $E$ , the betweenness centrality is defined as:  $BC(v) = \frac{1}{S} \sum_{u,w \in V} \sigma_{uw}(v)$  where  $\sigma_{uw}(v)$  is the number of paths from  $u$  to  $w$  passing through  $v$ , and  $S$  is



(a) Simple graph with three viewpoints (illustrated by looking glasses). The sampled BC and AS hegemony are computed only with best paths from the three viewpoints, the expected BC is computed with all best paths. (b) Sampling error for BC and AS hegemony in function of the number of viewpoints.

**Fig. 1.** Comparison of Betweenness Centrality (BC) and AS hegemony with a toy example and BGP data.

the total number of paths. BC ranges in  $[0, 1]$ , but the relative magnitudes of the scores are usually more significant than the absolute values.

**Challenges:** In theory, to compute BC one needs the set of all paths in the graph. With BGP data, however, we are restricted to paths bounded to a small number of viewpoints. We found that this singular type of path sampling greatly impairs BC results. To illustrate this, Fig. 1a presents a simple example with 13 ASes and three viewpoints. If we had viewpoints in all ASes, thus access to all paths in the graph, we would obtain the highest BC score for the transit ISP (.62) and lowest scores for the stub ASes (.15). But, using only paths bound to the three viewpoints, the computed BC scores are substantially different (Sampled BC in Fig. 1a). Because a third of the paths converge to each viewpoint, BC values for ASes close to the viewpoints are undesirably high. This bias is so pronounced that the BC for stub ASes accommodating viewpoints (.38) is twice higher than the BC of one of the regional ISP (.16). Theoretical studies have already reported the shortcomings of BC with sampled data [10], but this issue has been rarely acknowledged in the networking literature. Mahadevan et al. [12] reported that BC is not a measure of centrality when computed with network data, but we stress that this issue comes from the non-random and opportunistic sampling method used to collect BGP data rather than the metric itself.

In our experiments we construct a global AS graph using all data from the Route Views, RIS, and BGPmon project on June 1<sup>st</sup> 2016. This corresponds to an AS graph of more than 50k nodes with 326 viewpoints (we consider only full-feed BGP peers), and only 0.6% of all the AS paths on the Internet (16M paths out of the 2.5B). As collected paths all converge to the 326 viewpoints, ASes accommodating viewpoints and their neighboring ASes are seemingly more central than other ASes. To measure the bias obtained with real BGP data we conduct the following experiment. First, we compute the BC for all ASes with data from all 326 viewpoints, then we compare this distribution of BC values to BC values obtained with a smaller set of randomly selected viewpoints. The distance between two distributions is measured with the Kullback-Leibler

divergence. Figure 1b shows that changing the number of viewpoints invariably reshapes the BC distribution, meaning that the obtained BC values are conditioned by the number of viewpoints. From these results, we hypothesize that having more viewpoints would yield different BC values, thus the BC values obtained with the 326 viewpoints might not be representative of AS centrality.

### 3 Methodology

To address the above BC shortcomings, we devise a monitoring method based on a robust centrality metric called AS hegemony. The proposed method consists of two basic steps. First we generate graphs from AS paths advertised via BGP. Then, using AS hegemony, we estimate the centrality of each AS in the graphs. We consider two types of graphs, global and local graphs.

**Global graph:** A global graph is made from all AS paths reported by the BGP viewpoints regardless of the origin AS and announced prefix. Consequently, these graphs represent the global Internet and central nodes stand for transit networks that are commonly crossed to reach arbitrary IP addresses. In 2017, IPv4 global graphs typically contains about 58 k nodes and 188 k edges (14k nodes and 43 k edges for IPv6). The structure of these graphs is complex, yet they are valuable to monitor the Internet altogether and reveal Internet-wide routing changes.

**Local graph:** A local graph is made only from AS paths with the same origin AS. Thereby, we compute a local graph for each AS announcing IP space globally. Each local graph represents the different ways to reach its origin AS and dominant nodes highlight the main transit networks towards only this AS. These graphs are particularly useful to monitor the dependence of an AS to other networks. In addition, structural changes in local graphs can expose important routing changes that may be detrimental to the origin AS reachability.

**AS Hegemony:** The core of the proposed method is to quantify the centrality of ASes in the generated graphs. To circumvent BC sampling problems we extend the recently proposed AS hegemony metric [5]. This metric measures the fraction of paths passing through a node while correcting for sampling bias.

Computing the hegemony of AS  $v$  from AS paths collected from several viewpoints consists of the two following steps. First, AS paths from viewpoints that are biased towards or against AS  $v$  are discarded. A viewpoint bias towards AS  $v$  means that the viewpoint is located within AS  $v$ , or topologically very close to it, and reports numerous AS paths passing through AS  $v$ . In contrast, a viewpoint bias against AS  $v$  is topologically far from  $v$  and is reporting an usually low number of AS paths containing  $v$ . Therefore, viewpoints with an abnormally high, or low, number of paths passing through  $v$  are discarded and only other viewpoints are selected to compute the hegemony score. Second, the centrality of  $v$  is computed independently for each selected viewpoint and these scores are aggregated to give the final AS hegemony value. That is, for each selected viewpoint  $j$  the BC of  $v$  (hereafter referred as  $BC_{(j)}(v)$ ) is computed

only from AS paths reported by  $j$ . And the average BC value across all selected viewpoints is the AS hegemony score of  $v$ .

These steps can be formally summarized into one equation. Let  $n$  be the total number of viewpoints,  $\lfloor \cdot \rfloor$  be the floor function and  $0 \leq 2\alpha < 1$  be the ratio of disregarded viewpoints. That is, the top  $\lfloor \alpha n \rfloor$  viewpoints with the highest number of paths passing through the AS and do the same for viewpoints with the lowest number of paths. Then the AS hegemony is defined as:

$$\mathcal{H}(v, \alpha) = \frac{1}{n - (2\lfloor \alpha n \rfloor)} \sum_{j=\lfloor \alpha n \rfloor + 1}^{n - \lfloor \alpha n \rfloor} BC_{(j)}(v),$$

where  $BC_{(j)}$  is the BC value computed with paths from only one viewpoint  $j$  (i.e.  $BC_{(j)}(v) = 1/S \sum_{w \in V} \sigma_{jw}(v)$ ) and these values are arranged in ascending order such that  $BC_{(1)}(v) \leq BC_{(2)}(v) \leq \dots \leq BC_{(n)}(v)$ .

Figure 1a depicts the AS hegemony obtained for the simple graph with three viewpoints ( $\alpha = .34$ ). Unlike the sampled BC, the AS hegemony is consistent for each type of node: transit ( $\mathcal{H} = 0.58$ ), regional ISP ( $\mathcal{H} = 0.25$ ) and stub AS ( $\mathcal{H} = 0.08$ ). AS hegemony scores are intuitively interpreted as the average fraction of paths crossing a node. For example, on average a viewpoint has one fourth of its paths crossing a regional ISP ( $\mathcal{H} = 0.25$ ).

In order to compare the robustness of AS hegemony and BC with real data, we reproduce the experiment of Sect. 2 with AS hegemony. Figure 1b shows that the hegemony values with 20 or more viewpoints are very similar to the ones obtained with the 326 peers. Meaning that path sampling has significantly less impact on AS hegemony than on BC. Note that we randomly select peers across different projects (e.g. Route Views, RIS) to obtain a diverse set of viewpoints. Selecting viewpoints from the same collector may yield poor results [5].


**Path Weights:** We extend AS hegemony to account for path disparities. In a nutshell, we weight paths according to the amount of IP space they are bound to. For example, a path to a /24 IP prefix represents a route to a smaller network than a path to a /16 IP prefix, thus we want to emphasize the path to the /16. The network prefix length alone is however not sufficient to resolve the IP space bound to a path. IP space deaggregation [2, 6] should also be taken into account. For example, a viewpoint reports the path ‘X Y Z’ for the prefix  $a.b.c.0/17$  and the path ‘X W Z’ for the prefix  $a.b.0.0/16$ . Meaning that BGP favors path ‘X Y Z’ for half of the advertised /16. Here there is no need to give more emphasis to the path bound to the /16 as each path represents a route to  $2^{15}$  IP addresses.

Consequently, we modify our definition of BC to account for the size of the IP space reachable through a path. Formally,  $\sigma_{uw}(v)$  is the number of IP addresses bound to the paths from  $u$  to  $w$  and passing through  $v$ . That is the number of IP addresses corresponding to the advertised IP prefixes minus the number of IP addresses from covered prefixes (i.e. deaggregated and delegated prefixes [2]) that are not passing through  $v$ . In the rest of the paper this weighted version of BC is applied for the calculation of AS hegemony in IPv4, but as the relation between number of addresses and prefix size in IPv6 is more ambiguous we keep the classical BC definition for the calculation of AS hegemony in IPv6.

## 4 Results

Our Python implementation of the above method fetches BGP data using the BGPStream framework from CAIDA [13] and computes AS hegemony of all ASes in the global graph as well as AS hegemony for ASes in all local graphs. Our tool and updated results are made publicly available [1].

**Parameter tuning:** Setting the value of  $\alpha$  is a trade-off between sampling robustness ( $\alpha \approx 0.5$ ) and sensitivity to local routing changes ( $\alpha \approx 0$ ). For example, setting  $\alpha \approx 0.5$  achieves the most robust results to path sampling but conceals routing events affecting less than half of the monitored BGP peers. To monitor routing changes we seek for a small value of  $\alpha$  that is still robust to path sampling. In Fig. 2 we compare robust AS hegemony scores ( $\alpha = 0.49$ ) to scores obtained with different values of  $\alpha$ . For the following experiments we set the parameter  $\alpha = 0.1$ , as it provides results similar to those obtained with  $\alpha = 0.49$  but is more sensitive to local changes.

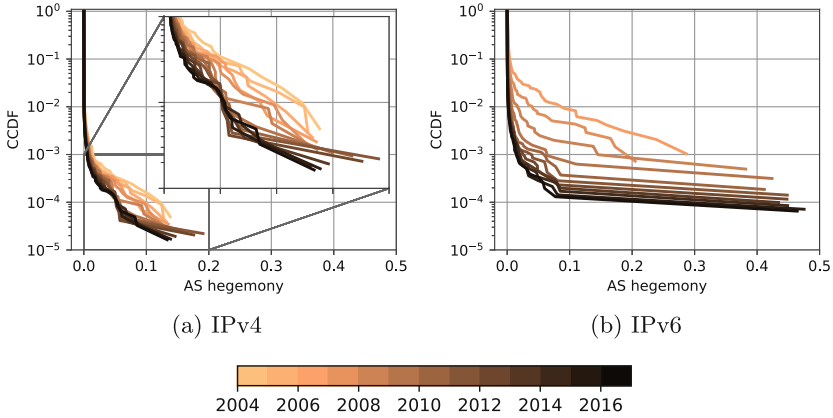


**Fig. 2.** KL divergence between AS hegemony scores obtained with  $\alpha = 0.49$  and different values of  $\alpha$  (global graph on 2017/12/15 with rv2, LINX, rrc00, and rrc10 collectors).

**Dataset:** The following results are all obtained using BGP data from four BGP collectors, two from the RouteViews project (route-views2 and LINX) and two from the RIS project (rrc00 and rrc10). These four collectors are selected from the collectors sensitivity results presented in [5]. For IPv4 they represent from 51 to 95 BGP peers respectively in 2004 and 2017. For IPv6, however, as the number of BGP peers is rather small before 2007 (i.e. less than 10 peers) and AS hegemony values might be irrelevant with such low number of peers (see Fig. 1b), we report only results obtained from 2007 onward using from 11 to 44 peers. The results presented below are obtained with RIB data of all peers for the 15<sup>th</sup> of each month from January 2004 to September 2017.


### 4.1 IPv4 and IPv6 Global Graphs

As the starting point of our analysis, we investigate the AS interdependency for the entire IP space. We monitor the evolution of AS hegemony scores in the global AS graph from 2004 to 2017. Here large AS hegemony scores represent transit networks that are commonly crossed to reach arbitrary IP addresses. Figure 3 depicts the distribution of the yearly average AS hegemony for all ASes in the IPv4 and IPv6 global AS graphs. In these figures each point represent an AS, and those on the right hand side of the figures stand for nodes with the highest hegemony values.



**Fig. 3.** Distribution of AS hegemony for all ASes in the global graph.

As the distribution of AS hegemony values for IPv4 is overall shifting to the left over time (Fig. 3a), we observe a global and steady decrease of AS hegemony values. This is another evidence of Internet’s flattening [3], as networks are peering with more networks we observe less dominant ASes. Nonetheless, Fig. 3a suggests that the AS hegemony for the most dominant networks (i.e. points on the right hand side) is quite stable.



**Fig. 4.** AS hegemony for Tier-1 ISPs from 2004 to 2017 (global graph, IPv4).

We further investigate this by selecting the eight most dominant ASes found in our dataset and monitor their yearly AS hegemony (Fig. 4). The AS hegemony for these networks is indeed either steady, or increasing, which is contradictory with the global Internet flattening observed earlier. These two observations provide evidences of dense connectivity at the edge of the Internet but the role of large transit ISP is still very central to connect remote places in the Internet. This can be explained by the growth of public peering facilities (IXP) that allows regional networks to keep traffic locally and peer directly with content providers. Yet transiting to remote locations requires the international networks

of tier-1 ISPs. In recent years this distinction between tier-1 ISP and other networks is even more visible, as we observe in Fig. 3a a clear gap between most networks ( $\mathcal{H} < 0.03$ ) and tier-1 ISPs ( $\mathcal{H} > 0.05$ ).

Figure 4 also depicts the dominance of Level(3) through the entire study period. After Level(3) acquisition of Global Crossing (AS3549) in 2011, it reached in 2012 the highest AS hegemony score monitored for the IPv4 global graph ( $\mathcal{H} = 0.19$ ). We also found that from 2008 to 2010 Global Crossing was the most dominant AS in Level(3) local graph, meaning that it was the most common transit network to reach Level(3). These results thus assert that Global Crossing acquisition was the most effective way for Level(3) to attain new customers. It also illustrates the benefits of our tools for deployment and business decisions.

For IPv6 (Fig. 3b) we observe a faster Internet flattening than for IPv4. We hypothesize that this is mainly because the Internet topology for IPv6 in 2007 was quite archaic. But IPv6 has drastically gained in maturity, the AS hegemony distribution for IPv6 in 2017 is then very close to the one for IPv4 in 2009. The most striking difference with IPv4 is the central role played by Hurricane Electric (HE) in the IPv6 topology. After doubling its number of peers in 2009 [8], HE has been clearly dominating the IPv6 space from 2009 onward. It reaches an impressive AS hegemony  $\mathcal{H} = 0.46$  in 2017, largely above the second and third highest scores (0.07 and 0.05), respectively, for Level(3) and Telia. Consequently, our tools confirm the dominant position of HE observed previously [4] and permit to systematically quantify the overall IPv6 dependency to HE.

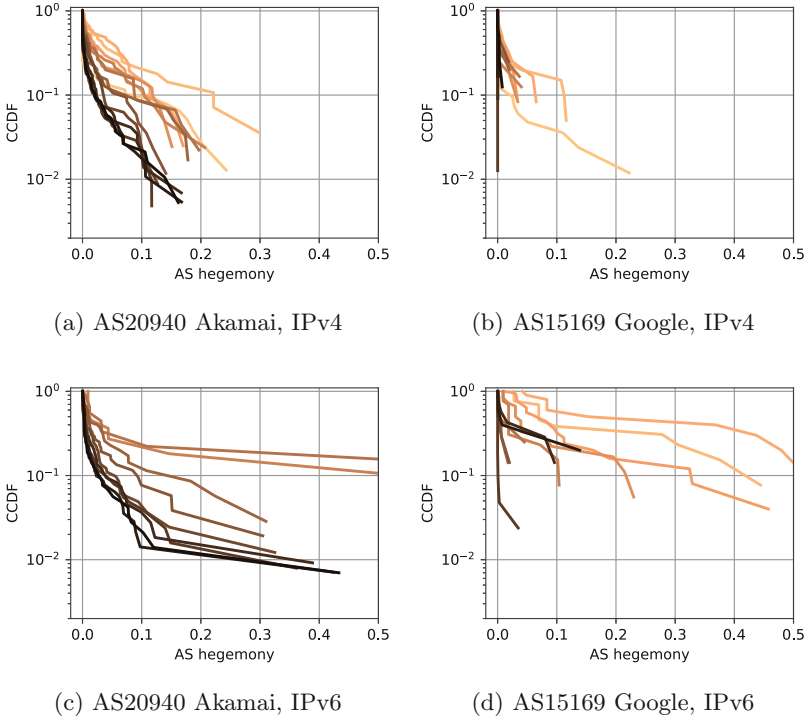
## 4.2 Case Studies: Local Graphs

Our analysis now focuses on results obtained with local graphs. Unlike the global ones, local graphs shed light to AS dependency only for a specific origin AS. We found that the structure of local graphs is very different depending on the size and peering policies of the origin AS. On average in 2017, an IPv4 local graph contains 98 nodes but 93% of these nodes have an hegemony null ( $\mathcal{H} = 0$ ). Typically ASes hosting BGP peers have an hegemony null and AS hegemony scores increases as the paths converge towards the origin AS. Thereby, the upstream provider of a single-homed origin AS gets the maximum hegemony score,  $\mathcal{H} = 1$ . By definition the origin AS of each local graph also features  $\mathcal{H} = 1$ , therefore, we are not reporting the AS hegemony of the origin AS in the following results.

In 2017, local graphs have on average 5 ASes with  $\mathcal{H} > 0.01$ , which usually corresponds to a set of upstream providers and tier-1 ASes. We also noticed interesting graphs containing no dominant AS, and other graphs containing numerous nodes with non-negligible AS hegemony scores. To illustrate this we pick a local graph from both end of the spectrum, namely, AS20940 from Akamai and AS15169 from Google.

**Akamai and Google:** The IPv4 graph for Akamai’s main network, AS20940, is the local graph with the largest number of nodes in our results. In 2017, it contains on average 30 nodes with an AS hegemony greater than 0.01 (see Fig. 5a). Meaning that accessing Akamai IP space relies on a large set of transit networks. This is true for our entire analysis period as shown in Fig. 5a. Our manual inspection of Akamai BGP announcements reveals that Akamai is heavily fragmenting





**Fig. 5.** Distribution of AS hegemony for Google and Akamai local graphs. Same color scale as Fig. 3.

its IP space and advertising small prefixes at various Points of Presence (PoPs). Consequently, each prefix is accessible only through a very limited number of upstream providers and all BGP peers report AS paths going through these providers. In summary, Akamai local graph contains a lot of nodes with weak but non-negligible AS hegemony scores implying that Akamai has numerous weak AS-dependencies.

On the other hand, the IPv4 graph for Google (AS15169) in 2017 contains no node with an hegemony greater than 0.01 (see Fig. 5b). Our manual inspection of Google BGP advertisements reveals that, unlike Akamai, Google announces all their prefixes at each PoP. Because Google is peering at numerous places, all BGP peers report very short and different AS paths with almost no AS in common hence no relevant hegemony score. Nonetheless, Google's local graphs before 2012 feature a different AS hegemony distribution with a few high AS hegemony scores (Fig. 5b). Level(3) is the most dominant AS observed until 2012. But then Google has clearly succeeded to bypass Level(3) and alleviate its dependency to this AS (usually  $\mathcal{H} < 0.00005$  from 2014). Now Level(3) is rarely seen in paths towards Google. In summary, we observe that Google used to depend on a few ASes but it is now mostly independent from all ASes. This

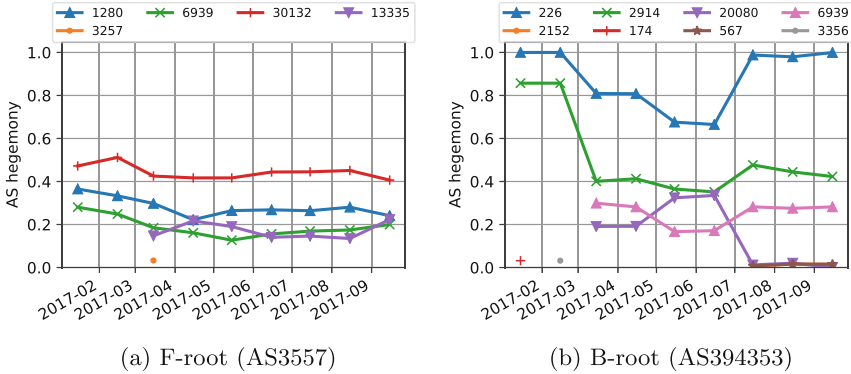
is not an isolated case, we measured no AS dependency for a few other ASes, notably, Microsoft (AS8075), Level(3) (AS3356), HE (AS6939), and Verisign (AS7342).

For IPv6, the situations for Akamai and Google are a bit different. The local graph for Akamai contains a lot of nodes with a high AS hegemony (Fig. 5c). But HE is quite outstanding and features an AS hegemony ( $\mathcal{H} = 0.43$ ) very close to the one observed for HE in the IPv6 global graph (Fig. 3b). HE is also the dominant node in Google's IPv6 local graph (Fig. 5d) but at a much lower magnitude ( $\mathcal{H} = 0.12$ ). Thereby, our results show that Google's aggressive peering policy has partially succeeded to bypass HE ubiquitous IPv6 network.

**DNS root servers:** Monitoring an AS with our tools provides valuable insights into its AS dependency. This is particularly useful for networks hosting critical infrastructure, as operators of these ASes try to minimize their dependencies to third-party networks. To illustrate the benefits of our tools, we present results for the local graphs of ASes hosting DNS root servers. Notice that understanding AS dependency of root servers is usually a complicated task as most root servers are using anycast and more than 500 instances are deployed worldwide. Due to space constraints, we detail only IPv4 results for networks hosting the F-root (AS3557) and B-root (AS394353) servers as they had significant structural changes in 2017.

In early 2017, we observe three dominant transit ASes for the network hosting the F-root server (Fig. 6a). AS30132 and AS1280 are direct upstream networks managed by ISC, the administrator of the F-root server. AS6939 is HE, the main provider for AS1280, and is found in about a third of the AS paths toward the F-root server. From March, Cloudflare (AS13335) starts providing connectivity to new F-root instances [7]. This new infrastructure is clearly visible in our results. Starting from March 2017, Cloudflare hegemony is fluctuating around 0.2 and seems to divert traffic from other instances as the three other transit networks have their hegemony proportionally decreased. From these results we deduce that the addition of Cloudflare has successfully reduced F-root dependencies on other ASes.

For the B-root server (Fig. 6b), we observe two dominant ASes in January and February 2017, Los Nettos (AS226) and NTT America (AS2914). Los Nettos reaches  $\mathcal{H} = 1$  because at that time the B-root server was unicasted and Los Nettos was the sole provider. NTT also has a very high AS hegemony score, in fact more than 80% of analyzed AS paths also cross NTT's network. From March 2017, we observe two other transit nodes AMPATH (AS20080) and HE (AS6939). Our manual inspection of the advertised paths reveals that a single /24 prefix is advertised with AMPATH as the first hop and usually HE as the second hop. This prefix is one of the two /24 prefixes advertised by the network hosting the B-root server (AS394353) but is not the one containing the server IP address. We believe that B-root operators were testing anycast in preparation for the deployment of the second instance of B-root at Miami that happened in May [14]. In May we acknowledge the deployment of the second instance hosted at AMPATH as the hegemony of that AS is raising again and the one for Los Nettos had significantly decreased. From July onward, however, we observe a sudden decrease of AMPATH hegemony while hegemony for Los Nettos is getting back



**Fig. 6.** AS hegemony for nodes in F-root (AS3557) and B-root (AS394353) local graphs from 15<sup>th</sup> January to 15<sup>th</sup> September 2017.

close to 1. A manual comparison of AS paths in June and July reveals that Los Nettos is trying to fix this by prepending its ASN to paths through HE. Despite these efforts, most of the paths that were transiting through HE and AMPATH in June are replaced by paths going through HE and Los Nettos in July. The addition of the second instance in Miami had uncertain benefits, first, it considerably mitigated the dependence on NTT and Los Nettos networks in May and June, but then, from July Los Nettos is once again totally dominating the B-root connectivity. Results for IPv6 are quite different, after the deployment in Miami we observe higher hegemony values for AMPATH. Both the IPv4 and IPv6 observations have been confirmed by the B-root operators.

**Future directions:** The structural changes observed for the F and B root servers illustrate the value of AS hegemony to monitor significant routing events. We are now designing an automated detection process to identify significant changes in AS hegemony scores. This detector reports sudden routing changes such as the recent BGP leak from Google [16]. During this event Google became a transit provider for NTT OCN, which exhibits a sudden and significant increase in Google’s AS hegemony for NTT’s local graph. Thanks to AS hegemony detecting this type of event is fairly easy, while state of the art tools employed by network operators (e.g. BGPmon provided by OpenDNS) have usually missed this significant event. As the details and evaluation of this detector go beyond the scope of this paper we leave this for future work.

In the future we are also planning to investigate different weighting schemes. For example by assigning paths’ weight based on traffic volume an ISP can emphasize destinations that are favored by its customers.

## 5 Conclusions

We presented a methodology to quantify the AS interdependency in the Internet. It deals with the various AS paths reported via BGP and produce AS hegemony scores, that are robust estimates of the ASes centrality. Using 14 years of BGP

data we proved that this method permits to monitor structural changes in the Internet and identify most important ASes to reach a certain part of the IP space. We also demonstrated with case studies the benefits of our tools to help ISPs to plan and assess infrastructure deployment. To assist network operators in these tasks we make our tools and results publicly available [1].

## References

1. AS Hegemony Results (2017). <http://ihr.iijlab.net/ihr/hegemony/>
2. Cittadini, L., Muhlbauer, W., Uhlig, S., Bush, R., François, P., Maennel, O.: Evolution of internet address space deaggregation: myths and reality. *IEEE JSAC* **8**(28), 1238–1249 (2010)
3. Comarela, G., Terzi, E., Crovella, M.: Detecting unusually-routed ases: methods and applications. In: *IMC*, pp. 445–459. ACM (2016)
4. Dhamdhere, A., Luckie, M., Huffaker, B., Elmokashfi, A., Aben, E., et al.: Measuring the deployment of ipv6: topology, routing and performance. In: *IMC*, pp. 537–550. ACM (2012)
5. Fontugne, R., Shah, A., Aben, E.: As hegemony: a robust metric for as centrality. In: *Proceedings of the SIGCOMM Posters and Demos*, pp. 48–50. ACM (2017)
6. Gamba, J., Fontugne, R., Pelsser, C., Bush, R., Aben, E.: BGP table fragmentation: what & who? In: *CoRes* (2017)
7. Grant, D.: Delivering Dot (2017). <https://blog.cloudflare.com/f-root/>
8. Pan, E.H.: Gigabit/ATM Monthly Newsletter. Information Gatekeepers Inc., November 2009
9. Karlin, J., Forrest, S., Rexford, J.: Nation-state routing: censorship, wiretapping, and BGP. *CoRR*, abs/0903.3218 2009
10. Lee, S.H., Kim, P.-J., Jeong, H.: Statistical properties of sampled networks. *Phys. Rev. E* **73**, 016102 (2006)
11. Liu, Y., Luo, X., Chang, R.K., Su, J.: Characterizing inter-domain rerouting by betweenness centrality after disruptive events. *IEEE JSAC* **31**(6), 1147–1157 (2013)
12. Mahadevan, P., Krioukov, D., Fomenkov, M., Dimitropoulos, X., Claffy, K.C., Vahdat, A.: The internet as-level topology: three data sources and one definitive metric. In: *SIGCOMM CCR*, vol. 36(1), pp. 17–26 (2006)
13. Orsini, C., King, A., Giordano, D., Giotsas, V., Dainotti, A.: BGPstream: a software framework for live and historical BGP data analysis. In: *IMC*, pp. 429–444. ACM (2016)
14. Root Operators. B-Root Begins Anycast in May (2017). <http://root-servers.org/news/b-root-begins-anycast-in-may.txt>
15. Schuchard, M., Mohaisen, A., Foo Kune, D., Hopper, N., Kim, Y., Vasserman, E.Y.: Losing control of the internet: using the data plane to attack the control plane. In: *CCS*, pp. 726–728. ACM (2010)
16. Toonk, A.: BGP leak causing Internet outages in Japan and beyond, August 2017. <https://bgpmon.net/bgp-leak-causing-internet-outages-in-japan-and-beyond/>
17. Wählisch, M., Schmidt, T.C., de Brün, M., Häberlen, T.: Exposing a nation-centric view on the german internet – a change in perspective on AS-Level. In: Taft, N., Ricciato, F. (eds.) *PAM 2012*. LNCS, vol. 7192, pp. 200–210. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-28537-0\\_20](https://doi.org/10.1007/978-3-642-28537-0_20)
18. Zhou, S., Mondragón, R.J.: Accurately modeling the internet topology. *Phys. Rev. E* **70**(6), 066108 (2004)