

Measuring Web Similarity from Dual-stacked Hosts

Steffie Jacob Eravuchira*, Vaibhav Bajpai†, Jürgen Schönwälder†, Sam Crawford*

*SamKnows Limited, London, UK

(steffie | sam)@samknows.com

†Computer Science, Jacobs University Bremen, Germany

(v.bajpai | j.schoenwaelder)@jacobs-university.de

Abstract—We compare the similarity of webpages delivered over IPv4 and IPv6. Using the SamKnows web performance (`webget`) test, we implemented an extension (`simweb`) that allows us to measure the similarity of webpages. The `simweb` test measures against ALEXA top 100 dual-stacked websites from 80 SamKnows probes connected to dual-stacked networks representing 58 different ASes. Using a two months-long dataset we show that 14% of these dual-stacked websites exhibit a dissimilarity in the *number* of fetched webpage elements, with 94% of them exhibiting a dissimilarity in their *size*. We show that 6% of these websites announce AAAA entries in the DNS but no content is delivered over IPv6 when an HTTP request is made. We also noticed several cases where *not all* webpage elements (such as images, javascript and CSS) of a dual-stacked website are available over IPv6. We show that 27% of the dual-stacked websites have some fraction of webpage elements that fail over IPv6, with 9% of the websites having more than 50% webpage elements that fail over IPv6. We perform a causality analysis and also identify sources for these failing elements. We show that 12% of these websites have more than 50% webpage elements that belong to the same origin source and fail over IPv6. Failure rates are largely affected by DNS resolution error on images, javascript and CSS content delivered from both same-origin and cross-origin sources. These failures tend to cripple experience for users behind an IPv6-only network and a quantification of failure cases may help improve IPv6 adoption on the Internet.

I. INTRODUCTION

With the World IPv6 Launch Day in 2012 [1], several notable content and service providers joined efforts to expedite IPv6 adoption [2], [3], [4]. Within a span of 4 years since the initiative, global adoption of IPv6 has increased to ~13.25% (as of Aug 2016) according to Google IPv6 adoption statistics [5] with Belgium (~45.0%), US (~28.8%), Switzerland (~26.0%), and Germany (~24.4%) leading IPv6 adoption rates. This has largely been possible due to spearheaded IPv6 deployment by ISPs both in the fixed-line (such as Telenet in Belgium, Comcast in US, Swisscom in Switzerland and Deutsche Telekom in Germany) and cellular (such as AT&T, Verizon Wireless and T-mobile USA) space.

Jakub Czyz *et al.* in [3] provide a nice overview on the state of IPv6 adoption on the Internet. They show that 3.5% (350) of ALEXA top 10K websites announce AAAA in DNS, with 3.2% (320) of these being reachable over IPv6. Recent studies [6], [7], [8], [9], [10] have compared performance of these dual-stacked websites over IPv4 and IPv6. However, there has been no study on comparing the similarity of webpage content delivered over IPv4 and IPv6. This is important since applications running on top of TCP will prefer fetching

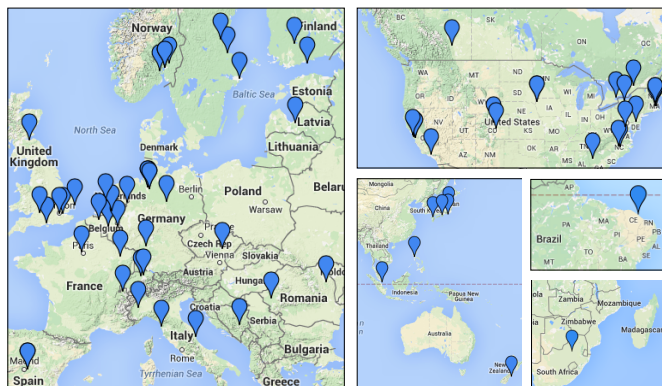


Fig. 1. Geographical distribution of our measurement trial comprising of 80 dual-stacked probes as of Aug 2016. Each vantage point is a SamKnows probe which is part of a larger SamKnows measurement platform.

webpages over IPv6 due to the default address selection policy [11] which prefers IPv6. We have recently shown [6] that this policy in conjunction with the Happy Eyeballs (HE) algorithm [12] leaves a dual-stacked host with less than 1% chance to prefer a connection over IPv4. As such, content providers need to ensure that the content delivered over IPv4 and IPv6 is identical. We want to know –

- How similar are the webpages accessed over IPv6 to their IPv4 counterparts? – In situations where the content is dissimilar over IPv4 and IPv6, what factors contribute to the dissimilarity?

We develop an active test (`simweb`) that uses well-known content and service complexity metrics [13], [14] to quantify the level of webpage dissimilarity. We deploy this test on 80 geographically distributed SamKnows [15] probes connected to dual-stacked networks (see Fig. 1) representing 58 different ASes to provide diversity of network origins. The test measures against ALEXA top 100 dual-stacked websites. Using a two-months long dataset, we quantify the dissimilarity of dual-stacked webpages. In situations where there is a dissimilarity we also perform a causal analysis and identify sources responsible for the difference. Our **contributions** –

- `simweb`: A tool for measuring webpage similarity (see Section III) over IPv4 and IPv6. The tool is written in C and open-sourced [16] for the measurement community.
- 14% of the ALEXA top 100 dual-stacked websites exhibit dissimilarity in the *number* of fetched webpage elements

with 6% showing more than 50% difference. 94% of dual-stacked websites exhibit dissimilarity in *size* with 8% showing at least 50% difference. This dissimilarity (see Section IV) in number and size of elements negatively impacts webpages fetched over IPv6.

- 27% of dual-stacked websites have some fraction of webpage elements that fail (see Section V) over IPv6 with 9% of the websites having more than 50% webpage elements that fail over IPv6. Worse, 6% announce AAAA entries in the DNS but no content is delivered over IPv6 when an HTTP request is made.
- 12% of dual-stacked websites have more than 50% webpage elements that belong to same-origin source and fail over IPv6. Failure rates are affected (see Section VI) by DNS resolution error on images, javascript and CSS delivered from both same-origin and cross-origin sources.

To the best of our knowledge this is the first study to measure webpage content similarity over IPv4 and IPv6. This is also the first study to investigate IPv6 adoption that goes beyond the root page (see II) of a dual-stacked website.

II. RELATED WORK

Mehdi Nikkiah *et al.* in [9] measure webpage performance over IPv4 and IPv6. They measure object size of the downloaded root page (without downloading embedded objects) and filter out websites where these sizes are not within 6% (over IPv4 and IPv6) of each other. Going forward, they measure average download speeds of the rest of the identical pages. Amogh Dhamdhare *et al.* in [10] take this forward to measure webpage performance by downloading the smallest webpage element which is at least 10K bytes in size. They also filter out websites where these sizes are not within 1% (over IPv4 and IPv6) of each other. No analysis on content dissimilarity is performed in either work. In this study, we close this gap by quantifying the amount of dissimilar dual-stacked websites with a causal analysis to identify potential areas for improvements.

III. METHODOLOGY

We use the SamKnows web performance test (*webget*) as a baseline to implement an extension (*simweb*) that we use to measure similarity of webpages delivered over IPv4 and IPv6. *webget* [17] (also called Mirage within the BISmark [18] platform) is written in C and is part of the SamKnows measurement test suite [19]. For a given website, *webget* downloads the root webpage and all its referenced webpage elements one-level deep over IPv4. In the process it calculates the DNS lookup time, time to first byte, HTTP request time, total size and download speed to fetch all webpage elements of a given website. While *webget* provides an aggregated statistics report across all webpage elements of a website, we are interested in the individual statistical report of each webpage element. As such, *simweb* builds upon this test to also report the content type, content size, resource URL, and IP endpoint used to fetch each webpage element. These properties are reported both over IPv4 and IPv6. Given a

hostname can point to multiple IP endpoints in DNS, *simweb* picks up the first IP endpoint returned by `getaddrinfo()` to establish TCP connections both over IPv4 and IPv6. In addition, HTTP status codes and `curl` [20] response codes are also used to identify the network level status of each request. The similarity is calculated in the data analysis phase using well-known metrics [13], [14] that measure the content and service complexity of a webpage. We use the number and content size of fetched webpage elements to quantify the *content complexity* of a website. The *service complexity* of a website is quantified by classifying webpage elements to belong to same and cross-origin sources (see Section VI for details) using hostnames derived from resource URLs.

We cross-compiled *simweb* for the OpenWrt platform [21] and deployed it on SamKnows probes. These probes, in addition to the *simweb* test, also perform standard SamKnows measurements [19]. The *simweb* test runs twice, once for IPv4 and The *simweb* test runs twice, once for IPv4 and subsequently for IPv6 and repeats every hour. This is to ensure that the first HTTP request to fetch the root webpage and all subsequent HTTP requests to fetch the webpage elements are made over one specific address family only. The test measures against ALEXA top 100 dual-stacked websites [8], which covers around top 350 ALEXA websites as of Aug 2016. Since the test repeats every hour, we limit probing to 100 websites to not overwhelm the volunteer’s broadband connection with our measurement traffic. It uses the user-agent string `Mozilla/4.0` when establishing HTTP session with the servers. Due to the inherent storage limitation of the probes, the locally collected measurement results are pushed every hour to our data collector.

Our measurement trial consists of 80 SamKnows probes (see Fig. 1) connected to dual-stacked networks that represent 58 different origin ASes. The dataset that we use in our analysis consists of *simweb* measurements collected for 65 days between April 2015 and June 2015. This includes around 207M data points captured from 80 SamKnows probes.

IV. COMPARING CONTENT SIMILARITY

We begin by defining our terminology. Let u denote a website identified by a URL. We call the HTML page returned by fetching the URL u as the root page of u denoted by $r(u)$. The root HTML page contains a set of embedded objects (such as images, CSS or javascripts) which we denote by $O(u)$. Furthermore, the root HTML page usually has a set of embedded links denoted as $L(u)$. Since we study the impact of accessing websites using different network protocols, we denote the root page of u accessed over IP version v as $r_v(u)$. Similarly, we refer to the set of embedded objects in $r_v(u)$ as $O_v(u)$ and the set of links in $r_v(u)$ as $L_v(u)$. Given $r_v(u)$ for IP version v , the set of objects and links that we are able to retrieve successfully using IP version v is given by $O'_v(u)$ and $L'_v(u)$ respectively. This terminology will be used in the rest of the data analysis.

In order to estimate the difference in the number of objects fetched for a website u over IPv4 and IPv6 we used Eq. 1. For




Fig. 2. Distribution of fractional difference in total number and size of successfully fetched objects over IPv4 and IPv6 for ALEXA top 100 dual-stacked websites. 14% of websites exhibit dissimilarity in the number of fetched webpage elements and 94% exhibit dissimilarity in size.




Fig. 3. Distribution of success rates towards ALEXA top 100 dual-stacked websites. 27% of dual-stacked websites have some fraction of webpage elements that fail over IPv6 with 9% having more than 50% webpage elements that fail over IPv6. 6% websites exhibit complete failure over IPv6.

a dual-stacked website u , the equation calculates the fraction of difference between number of fetched objects using IPv4 and IPv6, over the total number of objects fetched using IPv4 where $\hat{n}_v(u)$ represents the median of the sample of $n'_v(u)$ values across all measurements from all probes.

$$\Delta n(u) = \frac{\hat{n}_4(u) - \hat{n}_6(u)}{\hat{n}_4(u)} \times 100\% \quad (1)$$

$$n'_v(u) = |O'_v(u)| + |L'_v(u)|$$

Fig. 2 shows the distribution of $\Delta n(u)$ across ALEXA top 100 dual-stacked websites. It can be seen that 14% of websites exhibit dissimilarity in the number of fetched webpage elements with 6% showing more than 50% difference. This dissimilarity in number of elements negatively impacts webpages fetched over IPv6.

Similarly, in order to estimate the difference in the size of objects fetched for a website u we used Eq. 2. For a dual-stacked website u , the equation calculates the fraction of difference between size of objects fetched using IPv4 and IPv6, over total size of objects fetched using IPv4 where $\hat{s}_v(u)$ represents the median of the sample of $s'_v(u)$ values across all measurements from all probes.

$$\Delta s(u) = \frac{\hat{s}_4(u) - \hat{s}_6(u)}{\hat{s}_4(u)} \times 100\% \quad (2)$$

$$s'_v(u) = s(O'_v(u)) + s(L'_v(u))$$

TABLE I

TOP 27 DUAL-STACKED WEBSITES ORDERED BY THEIR SUCCESS RATES OVER IPV6. POPULAR WEBSITES ARE HIGHLIGHTED IN RED. WEBSITES WITH CHECKMARKS PARTICIPATED ON THE WORLD IPV6 LAUNCH DAY.

#	Webpage	Success Rate (%)		W6LD
		IPv6(↓)	IPv4	
01	www.bing.com	0	100	✓
02	www.detik.com	0	100	✓
03	www.engadget.com	0	100	✓
04	www.nifty.com	0	100	
05	www.qq.com	0	100	
06	www.sakura.ne.jp	0	100	
07	www.flipkart.com	09	99	✓
08	www.folha.uol.com.br	13	100	
09	www.aol.com	48	100	✓
10	www.comcast.net	52	100	✓
11	www.yahoo.com	72	100	✓
12	www.mozilla.org	84	100	✓
13	www.orange.fr	86	100	✓
14	www.seznam.cz	89	100	✓
15	www.mobile.de	90	100	✓
16	www.wikimedia.org	90	100	
17	www.t-online.de	93	100	✓
18	www.free.fr	95	100	
19	www.usps.com	95	100	
20	www.vk.com	95	100	✓
21	www.wikipedia.org	95	100	✓
22	www.wiktionary.org	95	100	
23	www.elmundo.es	96	100	✓
24	www.uol.com.br	96	100	✓
25	www.marca.com	97	100	✓
26	www.terra.com.br	98	100	✓
27	www.youm7.com	99	100	

The reported content size is the size of the payload (excluding the header). In situations where the response is HTTP chunked encoded [22], the payload is the sum of the size of all chunks (excluding the chunked metadata). In situations where the response is compressed, the content size reports the payload size before the receiver decompresses the data.

Fig. 2 shows the distribution of $\Delta s(u)$ across ALEXA top 100 dual-stacked websites. It can be seen that 94% of dual-stacked websites exhibit dissimilarity in size with 8% showing at least 50% difference. This size dissimilarity also negatively impacts webpages fetched over IPv6. This dissimilarity is due to lower success rates when fetching content over IPv6.

V. COMPARING SUCCESS RATES

For a website u , we define *success rate* as the fraction of number of successfully fetched webpage objects over the total number of objects and links embedded in the root page. We use Eq. 3 to estimate whether all webpage elements can be successfully fetched over IPv6.

$$p_v(u) = \frac{n'_v(u)}{n_v(u)} \times 100\% \quad (3)$$

We consider $\hat{p}_v(u)$, the median of the sample of success rate $p_v(u)$ values across all measurements from all probes to

Website failing over IPv6




Fig. 4. The causal analysis of websites failing over IPv6 at the network (left), content (middle) and service level (right). The percentage next to each website is its failure rate. Failure rates are affected by DNS resolution error on images, javascript and CSS content delivered from both same- and cross-origin sources.

a website u over IP version v as the representative success rate value for that website over IP version v . A website u with a $x\%$ value of $\hat{p}_v(u)$ means that only $x\%$ of the total objects embedded in the root page can be fetched over IP version v . Fig. 3 shows the distribution of $\hat{p}_v(u)$ across ALEXA top 100 dual-stacked websites. It can be seen that over IPv4, all webpages except www.flipkart.com have a median success rate value of 100% over the entire measurement duration from all probes. However, we see that 27% of websites show some rate of failure over IPv6, with 9% of websites exhibiting more than 50% failures over IPv6. Worse 6% of websites shows complete failure (0% success) over IPv6. 18% of these websites participated in the World IPv6 Launch Day [1] in 2012 with a promise to provide production-ready IPv6 access towards their services. Table I shows the success rate of these top 27 websites arranged in ascending order of their success rate over IPv6. The special case of www.bing.com, which has globally stopped providing IPv6 services in 2013 was recently identified by us in [8]. Apart from www.bing.com we further identified 4 (www.detik.com, www.engadget.com, www.nifty.com and www.sakura.ne.jp) websites that exhibit complete failure over IPv6. These websites were accessible over IPv6 in the past, but have stopped providing AAAA entries and therefore exhibit 0% success rate over IPv6. Furthermore, it's also surprising to see popular websites such as www.comcast.net (52% success rate),

www.yahoo.com (72%), www.mozilla.org (84%), www.wikimedia.org (90%) and www.wikipedia.org (95%) exhibiting partial failure over IPv6.

VI. CAUSALITY ANALYSIS

We further performed a causal analysis on the failure of 27% of ALEXA top 100 dual-stacked websites (see Table I) to investigate the network, content and service level source of the issue. We investigated the spectrum of libcurl error codes reported by `simweb` for each object of a failing website u . Fig. 4 (left) shows the percentage contribution of error codes to each failing website u . The numbers next to each failing website are the failure rates, $100\% - \hat{p}_6(u)$ flipped over from Table I. The error code `CURLE_OK` contributes to the success rate, while rest of the error codes contribute to the failure rate of each website u over IPv6. It can be seen that `CURLE_COULDNT_RESOLVE_HOST` is the major contributor to failure rates. This goes to show that most of the webpage elements of an IPv6-capable website fail due to a DNS resolution error. This is caused due to missing AAAA entries for these webpage elements in the DNS.

We also investigated MIME types reported by `simweb` for each object of a failing website u . Note that, `simweb` would not be able to return MIME types of objects that fail to be fetched. We therefore ran a post-processing function to fetch the response headers (and consequently the MIME type) of these objects over IPv4. Fig. 4 (middle) shows the percentage




Fig. 5. Contribution (in percentage) of cross origin sources towards the failure of webpage elements over IPv6. The percentage next to each website is its failure rate. These cross-origin sources can be largely classified as third-party advertisements (*.doubleclick.net), analytics (*.scorecardresearch.com, *.quantserve.com et al.), user-centric (*.facebook.com, *.ajax.googleapis.com et al.) and static content (*.wikimedia.org, *.creativecommons.org et al.) that tends to fail over IPv6.

contribution of MIME types to each failing website u . It can be seen that for websites which have AAAA entries (websites with less than 100% failure rate) – images, javascripts, and CSS content contribute to the majority of the failure over IPv6.

We further investigated the URLs reported by `simweb` for each object of a website u failing over IPv6. We used the URLs to identify the hostnames of these failing objects. We used these hostnames to classify objects into *same origin* and *cross origin* sources. We classify objects of a website u , to belong to a cross origin source whenever their hostnames do not match the hostname of the website u . The rest of the objects (including subdomains) belong to the same origin source. Fig. 4 (right) shows the contribution of same and cross origin sources to webpage elements that failed over IPv6. It can be seen that all failing (27% of dual-stacked) websites have some fraction of webpage elements that belong to the same origin source and fail over IPv6. Worse, Table II shows that 12% of dual-stacked websites have more than 50% webpage elements that belong to the same origin source and fail over IPv6. For instance, 48% of the webpage elements in `www.comcast.net` fail over IPv6. 85% of these failing webpage elements belong to the same origin (*.comcast.net) source. Similarly, popular websites such as `www.yahoo.com` (83% failing webpage elements

with a same-origin source), `www.wikimedia.org` (65%), `www.wiktionary.org` (22%), `www.wikipedia.org` (22%) and `www.mozilla.org` (7%) experience a high percentage of webpage elements that belong to the same origin source and that fail over IPv6. This is because the infrastructure that serves the content of a website does not have IPv6 turned on by default for all same-origin webpage elements. Furthermore, Fig. 5 shows the contribution of webpage elements that belong to cross origin sources. These cross-origin sources can be largely classified [23] as third-party advertisements (*.doubleclick.net), analytics (*.scorecardresearch.com, *.quantserve.com et al.), user-centric content (*.facebook.com, *.ajax.googleapis.com et al.) and static content (*.wikimedia.org, *.creativecommons.org et al.) that tends to fail over IPv6. In this way, it can be seen that for websites which have AAAA entries (websites with less than 100% failure rate) – both same and cross origin sources contribute to the failure of webpage elements over IPv6.

Given some of the cross-origin sources contribute to the failure of multiple websites, we also identify cross-origin sources that would help benefit more websites if their content was available over IPv6. Fig. 6 shows the distribution of cross-origin sources that contribute to the failure of more than




Fig. 6. Boxplot (left) of cross-origin sources that contribute to the failure of more than 1 dual-stacked website. The right y-axis shows the number of websites spanned by each of the cross-origin source. A table (right) shows median contribution of each cross-origin source to all spanned websites. *doubleclick.net* has the highest span across 5 websites, with *creativecommons.org* having highest median contribution to the failure rate of 3 websites.

TABLE II
CONTRIBUTION OF SAME ORIGIN SOURCES TOWARDS IPV6 FAILURE OF
WEBPAGE ELEMENTS WITH POPULAR WEBSITES HIGHLIGHTED IN RED.

#	Webpage	Same Origin (↓)
01	www.bing.com	100%
02	www.detik.com	100%
03	www.engadget.com	100%
04	www.nifty.com	100%
05	www.usps.com	100%
06	www.qq.com	100%
07	www.sakura.ne.jp	100%
08	www.comcast.net	85%
09	www.yahoo.com	83%
10	www.terra.com.br	74%
11	www.marca.com	70%
12	www.wikimedia.org	65%
13	www.elmundo.es	37%
14	www.vk.com	31%
15	www.t-online.de	30%
16	www.youm7.com	24%
17	www.wiktionary.org	22%
18	www.wikipedia.org	22%
19	www.free.fr	13%
20	www.folha.uol.com.br	12%
21	www.mozilla.org	7%
22	www.uol.com.br	7%
23	www.mobile.de	7%
24	www.aol.com	5%
25	www.orange.fr	5%
26	www.seznam.cz	4%
27	www.flipkart.com	1%

1 dual-stacked website. It can be seen that the cross-origin source *doubleclick.net* has the highest span across 5 websites, with a 0.54% median contribution to failure rates. The cross-origin source *creativecommons.org* on the other hand has 76% median contribution to the failure rate of 3 websites. The number of spanned websites show how many can be benefited, with the median contribution exhibiting how much the failure rate can be reduced by enabling IPv6 content delivery from these cross-origin sources.

VII. CONCLUSION

We presented *simweb*, a tool that can be used to measure similarity of webpages delivered over IPv4 and IPv6. Using a *simweb* dataset collected from 80 dual-stacked SamKnows probes, we showed that there is dissimilarity in the number and size of webpage elements that negatively impacts the user experience over IPv6. We showed that websites which used to be IPv6 capable once did not remain as such forever. We witnessed cases where a dual-stacked website (such as *www.bing.com*, *www.detik.com*, *www.engadget.com*, *www.nifty.com* and *www.sakura.ne.jp*) stopped announcing AAAA entries in DNS over time. Metrics that measure IPv6 adoption should account for such changes. We also showed that metrics that limit only to the root webpage of a dual-stacked website can lead to an overestimation of IPv6 adoption numbers on the Internet. We witnessed several cases where images, javascript and CSS content of a dual-stacked website did not have AAAA entries in the DNS. In fact, 27% of ALEXA top 100 dual-stacked websites have some fraction of webpage elements that fail over IPv6 with 9% websites having more than 50% failure rates. Furthermore, 12% of these websites have more than 50% webpage elements that belong to the same origin source and fail over IPv6. This includes popular websites such as *www.comcast.net*, *www.yahoo.com*, *www.mozilla.org*, *www.wikimedia.org* and *www.wikipedia.org*. It remains unclear whether such websites can be deemed IPv6 ready. Finally, we identified IPv4-only cross-origin sources (such as **.creativecommons.org*) that would help improve IPv6 experience of a number of dual-stacked websites once they enable content delivery over IPv6.

ACKNOWLEDGEMENTS

We thank all probe hosts and Jamie Mason (SamKnows) for providing us technical support on the SamKnows infrastructure. This work was partly funded by Flamingo, a Network of Excellence project (ICT-318488) supported by the European Commission under its Seventh Framework Programme.

REFERENCES

- [1] "Internet Society - World IPv6 Launch," <http://www.worldipv6launch.org>, [Online; accessed 11-January-2016].
- [2] M. Nikkha, and R. Guérin, "Migrating the Internet to IPv6: An Exploration of the When and Why," *IEEE/ACM Trans. Netw.*, vol. 24, no. 4, pp. 2291–2304, 2016. [Online]. Available: <http://dx.doi.org/10.1109/TNET.2015.2453338>
- [3] J. Czyz, M. Allman, J. Zhang, S. Iekel-Johnson, E. Osterweil, and M. Bailey, "Measuring IPv6 Adoption," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: ACM, 2014, pp. 87–98. [Online]. Available: <http://doi.acm.org/10.1145/2619239.2626295>
- [4] L. Colitti, S. H. Gunderson, E. Kline, and T. Refice, "Evaluating IPv6 Adoption in the Internet," in *Passive and Active Measurement, 11th International Conference, PAM 2010, Zurich, Switzerland, April 7-9, 2010. Proceedings*, 2010, pp. 141–150. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-12334-4_15
- [5] "Google IPv6 Adoption Statistics," <http://www.google.com/intl/en/ipv6/statistics.html>, [Online; accessed 11-January-2016].
- [6] V. Bajpai and J. Schönwälder, "Measuring the Effects of Happy Eyeballs," in *Proceedings of the 2016 Applied Networking Research Workshop*, ser. ANRW '16. New York, NY, USA: ACM, 2016, pp. 38–44. [Online]. Available: <http://doi.acm.org/10.1145/2959424.2959429>
- [7] I. Livadariu, A. Elmokashfi, and A. Dhamdhere, "Characterizing IPv6 control and data plane stability," in *35th Annual IEEE International Conference on Computer Communications, INFOCOM 2016, San Francisco, CA, USA, April 10-14, 2016*, 2016, pp. 1–9. [Online]. Available: <http://dx.doi.org/10.1109/INFOCOM.2016.7524465>
- [8] V. Bajpai and J. Schönwälder, "IPv4 versus IPv6 - who connects faster?" in *Proceedings of the 14th IFIP Networking Conference, Networking 2015, Toulouse, France, 20-22 May, 2015*, 2015, pp. 1–9. [Online]. Available: <http://dx.doi.org/10.1109/IFIPNetworking.2015.7145323>
- [9] M. Nikkha, R. Guérin, Y. Lee, and R. Woundy, "Assessing IPv6 through web access a measurement study and its findings," in *Proceedings of the 2011 Conference on Emerging Networking Experiments and Technologies, Co-NEXT '11, Tokyo, Japan, December 6-9, 2011*, 2011, p. 26. [Online]. Available: <http://doi.acm.org/10.1145/2079296.2079322>
- [10] A. Dhamdhere, M. J. Luckie, B. Huffaker, kc claffy, A. Elmokashfi, and E. Aben, "Measuring the deployment of IPv6: topology, routing and performance," in *Proceedings of the 12th ACM SIGCOMM Internet Measurement Conference, IMC '12, Boston, MA, USA, November 14-16, 2012*, 2012, pp. 537–550. [Online]. Available: <http://doi.acm.org/10.1145/2398776.2398832>
- [11] D. Thaler, R. Draves, A. Matsumoto, and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)," RFC 6724 (Proposed Standard), Internet Engineering Task Force, Sep. 2012. [Online]. Available: <http://www.ietf.org/rfc/rfc6724.txt>
- [12] D. Wing and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts," RFC 6555 (Proposed Standard), Internet Engineering Task Force, Apr. 2012. [Online]. Available: <http://www.ietf.org/rfc/rfc6555.txt>
- [13] M. Butkiewicz, H. V. Madhyastha, and V. Sekar, "Characterizing Web Page Complexity and Its Impact," *IEEE/ACM Trans. Netw.*, vol. 22, no. 3, pp. 943–956, 2014. [Online]. Available: <http://dx.doi.org/10.1109/TNET.2013.2269999>
- [14] —, "Understanding website complexity: measurements, metrics, and implications," in *Proceedings of the 11th ACM SIGCOMM Internet Measurement Conference, IMC '11, Berlin, Germany, November 2-, 2011*, 2011, pp. 313–328. [Online]. Available: <http://doi.acm.org/10.1145/2068816.2068846>
- [15] V. Bajpai and J. Schönwälder, "A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts," *IEEE Communications Surveys and Tutorials*, vol. 17, no. 3, pp. 1313–1341, 2015. [Online]. Available: <http://dx.doi.org/10.1109/COMST.2015.2418435>
- [16] "simweb - A tool to measure similarity of webpages delivered over IPv4 and IPv6," <https://github.com/steffiejacob/simweb>, [Online; accessed 30-Aug-2016].
- [17] S. Sundaresan, N. Feamster, R. Teixeira, and N. Magharei, "Measuring and mitigating web performance bottlenecks in broadband access networks," in *Proceedings of the 2013 Internet Measurement Conference, IMC 2013, Barcelona, Spain, October 23-25, 2013*, 2013, pp. 213–226. [Online]. Available: <http://doi.acm.org/10.1145/2504730.2504741>
- [18] S. Sundaresan, S. Burnett, N. Feamster, and W. de Donato, "BISmark: A Testbed for Deploying Measurements and Applications in Broadband Access Networks," in *2014 USENIX Annual Technical Conference, USENIX ATC '14, Philadelphia, PA, USA, June 19-20, 2014*, 2014, pp. 383–394. [Online]. Available: <https://www.usenix.org/conference/atc14/technical-sessions/presentation/sundaresan>
- [19] "SamKnows - Test Methodology," <https://goo.gl/LHRZks>, [Online; accessed 11-January-2016].
- [20] "curl and libcurl," <https://curl.haxx.se>, [Online; accessed 11-January-2016].
- [21] M. Petullo, "Building Custom Firmware with OpenWrt," *Linux Journal*, vol. 2010, no. 196, Aug. 2010. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1883498.1883501>
- [22] R. Fielding and J. Reschke, "Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing," RFC 7230 (Proposed Standard), Internet Engineering Task Force, Jun. 2014. [Online]. Available: <http://www.ietf.org/rfc/rfc7230.txt>
- [23] D. Gugelmann, B. Ager, and V. Lenders, "Towards classifying third-party web services at scale," in *Proceedings of the 2014 CoNEXT on Student Workshop, CoNEXT Student Workshop '14, Sydney, Australia, December 2, 2014*, 2014, pp. 34–36. [Online]. Available: <http://doi.acm.org/10.1145/2680821.2680838>