

Adaptive Address Family Selection for Latency-Sensitive Applications on Dual-stack Hosts

Maxime Piraux
maxime.piraux@uclouvain.be
UCLouvain
Belgium

Olivier Bonaventure
olivier.bonaventure@uclouvain.be
UCLouvain
Belgium

ABSTRACT

Latency is becoming a key factor of performance for Internet applications and has triggered a number of changes in its protocols. Our work revisits the impact on latency of address family selection in dual-stack hosts. Through RIPE Atlas measurements, we analyse the address families latency difference and establish two requirements based on our findings for a latency-focused selection mechanism. First, the address family should be chosen per destination. Second, the choice should be able to evolve over time dynamically.

We propose and implement a solution formulated as an online learning problem balancing exploration and exploitation. We validate our solution in simulations based on RIPE Atlas measurements, implement and evaluate our prototype inside a DNS resolver in four access networks using Chrome and popular web services. We demonstrate the ability of our solution to converge towards the lowest-latency address family and improve the latency of transport connections used by applications.

CCS CONCEPTS

• **Networks** → **Network algorithms**; *Naming and addressing*;

KEYWORDS

IPv4, IPv6, Latency, Dual-Stack, DNS, Happy Eyeballs, Multi-Armed Bandits, TCP, QUIC

1 INTRODUCTION

The Internet originally adopted IPv4 as its network protocol. With the important growth of the Internet over the years, the IPv4 address space became a constraint that the Internet Engineering Task Force (IETF) addressed by designing IPv6. It offers a much larger address space, which also brings many opportunities [35]. The Internet being a critical infrastructure and given the key role played by IP addresses, the network community had to plan for the coexistence of IPv4 and IPv6 as a mean of transitioning. This has led to a number of challenges in several areas of the network stack. The DNS has been extended to resolve domain names into IPv6

addresses with the AAAA [23] record type. The IPv6 addressing scheme comprises a dedicated prefix for representing IPv4 addresses [16], potentially simplifying the code of applications supporting both. The Linux socket API was also modified with *dual-stack sockets* enabling the use of IPv4 and IPv6, as well as these IPv4-mapped IPv6 addresses. Windows and macOS APIs include methods to connect a socket using domain name tackling this problem [28, 21].

Given that the adoption of IPv6 on devices, operating systems and networks is heterogeneous [31, 32, 12], very few service providers completely transitioned to IPv6 but rather became dual-stack. As a result, when an application establishes a transport connection, it needs to select one address family. This problem has seen a number of solutions over the years [42, 45, 37]. All of them made the hypothesis that IPv6 should be favoured to foster its transition and include a fallback mechanism in case of a broken IPv6 path. At the early stages of the IPv6 deployment, several transition solutions such as 6to4 [11] and Teredo [20] introduced tunnels to carry IPv6 packets over IPv4. These techniques added latency, could fail and become intermittently unavailable [47]. The first approach for applications was to start transport connections over IPv6 and use IPv4 when their establishment timer fires. Given that this time is usually set within tens of seconds, this approach greatly degrades the user experience when falling back to IPv4. A refined algorithm named Happy Eyeballs (HE) races the two connection establishments while delaying the start of the IPv4 one, and selects the one that completes first [45, 37]. Its upcoming proposed version 3 [33] leverages SVCB DNS records [38, 48] and honour the priority they set among resolved addresses. Today, as IPv6 is supported by 48% of the 1000 top websites as reported by the ISOC Pulse platform [40], a significant part of applications are facing this address family choice.

In the recent years, the greater importance in lowering the latency of Internet applications has started a series of changes in the Internet protocols. Two recent examples are the standardisation of the QUIC protocol [22], with a faster handshake and better recovery mechanisms among its key features, and the L4S architecture [8], which provides low

latency under utilisation to applications using Active Queuing Management in the network. Both help reducing the round-trip time of packets between a host and a server.

In the light of these changes, we investigate and answer two questions in this paper: (1) *Does IPv6 always provide the same latency as IPv4?* (2) *How can we improve the address family selection for latency-sensitive applications?*

Our paper addresses these questions using the following plan. First, Section 2 positions our work and research questions with respect to the state of the art. Then, in Section 3, we show, using the RIPE Atlas dataset [29], that a diversity of latency differences between IPv4 and IPv6 exists. These differences can vary across several destinations reached from a single source, with some also evolving over time. Then, in Section 4, we formulate the lowest-latency address family selection as an online learning problem balancing exploration and exploitation. Section 5 validates our design with simulations using the RIPE Atlas dataset. Section 6 presents our prototype demonstrating how DNS resolvers can aid hosts to select the lowest latency address family. Section 7 analyses results from real-world experiments evaluating our prototype using Chrome and popular web services. We will release all code from this work under an open-source license [34, 46] with the camera-ready version of this paper and Appendix B details the steps to reproduce this work.

2 RELATED WORKS

Several works have studied the end-to-end performance of IPv6 that we categorise in three phases. First, transition techniques such as 6to4 [11] and Teredo [20] that were prominent in the first phase of the worldwide IPv6 deployment had a dual-sided impact, both fostering its deployment but also often degrading user experience with added latency, sometimes poorer availability and failures [19]. Second, with more native IPv6 networks, these effects gradually decreased [13], and IPv6 and IPv4 had more often comparable performance when the inter-domain paths were similar, while differing paths significantly affected IPv6 performance [17]. These routing differences have been established as both the biggest contributor to poor IPv6 performance at scale and a key point of improvement through better IPv6 peering [32]. Third, as the support of IPv6 in applications increased with adoption from cloud companies and browsers, researchers analysed the impact of Happy Eyeballs (HE) on performance. IPv6 TCP connections times have improved over time, despite some still being higher than IPv4 [6, 5] which sometimes is due to lacking nearby IPv6 CDN caches [5]. HE was also found to select IPv6 towards a large majority of popular web services despite IPv6 being slower in most cases [7]. Halving the fallback timer of HE to 150 ms only enables marginal improvements in avoiding higher-latency IPv6 connections [7,

4]. In this work, we leverage the RIPE Atlas dataset [29] to assess the IPv4 and IPv6 performance differences today and validate our solution.

In the presence of dual-stack networks, researchers also proposed techniques beyond HE. Given several available paths, multipath transport protocols such as MPTCP [18, 36] and MPQUIC [26, 15] come as natural solutions. They can establish a connection over one path and then a subflow on the other. Then, both can be used simultaneously or selectively given application requirements, e.g. latency. An extension to MPTCP proposed a way of racing the connection establishment over several network paths and use all succeeding attempts within the MPTCP session [1]. Researchers have also used MPTCP over IPv4 and IPv6 with a coupled congestion controller and demonstrated that bandwidth can be improved by aggregating the non-congruent parts of the two paths [27]. However, MPTCP is not widely used today [2, 44] while MPQUIC is still under standardisation [26].

3 IPV4 AND IPV6 END-TO-END LATENCY

We analyse the RIPE Atlas dataset [29] to investigate on the end-to-end latency of dual-stack networks. RIPE Atlas consists of more than 10000 probes which are connected, usually using Ethernet, to networks around the world and measure connectivity with several network tests. RIPE Atlas comprises also about 1000 anchors which are dedicated servers with more capacity used to co-operate with probes in their measurements. RIPE probes regularly perform HTTP tests towards anchors. This test opens a TCP connection, sends an HTTP/1.1 GET request and receives a 4 KB HTTP response in a single initial TCP congestion window. We use the request completion time of this test as a form of smoothed latency measurement from a probe to a given anchor.

When a RIPE probe has access to a dual-stack network, the HTTP tests towards an anchor are simultaneously performed over IPv4 and IPv6 within a short window of time. When considering the HTTP tests between the 1st and 8th of June 2023, we obtain more than 156 millions of such paired tests, each happening within a two-minute time window. We filter out the completion times outside the 5th and 95th percentiles. The left part of Figure 1 reports the Cumulative Distribution Functions (CDF) of HTTP request completion times (RCT) for both IPv4 and IPv6. These are largely similar. We also compute the completion time reach ratios of IPv4 over IPv6 for each paired test.

The right part of Figure 1 reports the distribution of these ratios. We can first observe that neither IPv4 nor IPv6 has a lower completion time than the other, which confirms our first finding. However, there is an important number of cases in which one of the two has a lower completion time. When looking at the Figure, we can observe an equal




Figure 1: Cumulative Distribution Functions (CDF) of HTTP request completion times (RCT) are similar across address families but RCT ratios have significant disparities when paired temporally.

IPv4 is better	IPv6 is better	None strongly better
113092 (28.4%)	129070 (32.4%)	156212 (39.2%)

Table 1: Pairs of probe and anchor are rather homogeneously distributed between categories.

proportion of ratios have a lower completion time by 10% and more in favour of one address family as indicated by the hatched areas. We can conclude that there is an important diversity in terms of latency between the IPv4 and IPv6 paths from a source to a destination. When considering the use of Happy Eyeballs today, we observe that it can choose IPv6 when it has a higher latency than IPv4. However, further analysis is required to answer two questions: (1) *How are these differences distributed over pairs of probe and anchor?* (2) *How are these differences distributed over all anchors reached from a probe?*

(1) First, we group the HTTP tests data into about 400k pairs of probes and anchors for which at least 300 HTTP tests are available. We classify each pair into three categories by performing a *t-test* determining whether its IPv4 and IPv6 RCT distributions have different means and compute a 98% confidence interval on the means difference. When the *p-value* is lower than 0.02 and the confidence interval lower bound is higher than the supposedly-higher distribution standard deviation, we define the address family with a lower mean as better. When one of the two conditions is not met, we define the pair as having no significantly better address family. The second condition eliminates instances where the distributions strongly overlap. Table 1 reports the number of pairs in each category. We can observe that while the dominant category contains pairs with no significant difference, the others categories combined account for more than 60% of pairs. While we considered pairs with tests spanning a week, we can observe that our statistical tests establish a significant difference in a majority of cases.

Second, to better discern how stable these differences are, we deepen our analysis by performing change-point detection [43] on the time series of each pair of probe and anchor




Figure 2: An example of probe-to-anchor RCT series split into four segments using change-point detection to highlight dynamicity inside categories of Table 1 and produce Table 2.

to split them into segments of at least 30 tests. Figure 2 illustrates this process. We then perform our classification again on each segment to determine four additional groups: (A) Pairs having a significantly better address family for which at least one segment is classified in the opposite address family. (B) Pairs having a significantly better address family for which a segment is classified as having no significantly better address family. (C) Pairs having no significantly better address family for which one segment is classified in one of them. (D) Pairs having all segments matching the global classification.

Table 2 reports the number of pairs in each group. We can observe that more than 16% of the pairs of probe and anchor have time segments that diverge in terms of best address family. Pairs can temporarily either see their best address family change from one to the other (1.15% (A)), exhibit no significant difference among the two (8.15% (B)) or exhibit a significant difference for one of them (6.86% (C)). **We can conclude that differences in latency between IPv4 and IPv6 exist over a significant number of pairs of probe and anchor. These differences can be dynamic and can fluctuate over time.**

(2) Next, we investigate how the observed differences are spread among the anchors reached by a probe. We group our classification of pairs per probe and determine whether IPv4 or IPv6 performs the best in total across all destinations, i.e. having the largest number of anchors reached with the lowest latency. We then compute the ratio of anchors reached with this best address family. In the median case, only 39% of anchors can be reached with the lowest RCT using this best address family alone. When adding anchors having no strongly better address family, in the median case, 80% of anchors can be reached with the lowest RCT. **We can conclude that all dual-stack probes must use both address families to reach all anchors with the lowest completion time.**

Group	(A)		(B)		(C)		(D)
Table 1 category	IPv4	IPv6	IPv4	IPv6	None	None	Consistent
\exists segment with category	IPv6	IPv4	None	None	IPv4	IPv6	
Percentage	4569 (1.15 %)		32459 (8.15 %)		27312 (6.86 %)		334034 (83.85 %)

Table 2: While the classification of Table 1 remains consistent for 84 % of pairs, 16 % diverge in category over time.

4 ADAPTIVE FAMILY SELECTION

Based on our analysis presented in Section 3, we observe two facts. First, IPv6 can provide a higher latency than IPv4 in a significant number of cases. Second, providing the lowest latency address family to applications cannot be achieved by entirely disabling one address family for a given client. Moreover, we define two important properties for a latency-aware address family selection technique. First, the technique should be able to take different decisions for different destinations. Second, the technique must be able to adapt over time to changes in the address families latency.

This Section presents our design following these observations. Figure 3 illustrates its application on a gateway within a home or small-enterprise network. The rest of the Section further elaborates on its components.

4.1 Steering clients using the DNS resolver

There is an opportunity for hosts to share their knowledge of RTTs differences in their dual-stack network. We argue that the DNS resolver residing on the access gateway could influence the selected address family in its DNS responses. In a dual-stack network such as illustrated on the left-hand side of Figure 3, the DNS resolver can choose the address family when answering to AAAA requests by using IPv6 and IPv4-mapped addresses. The latter allows encoding IPv4 addresses as IPv6 addresses using the dedicated `::ffff:/96` prefix [16]. These IPv6 addresses are not used on the wire and represent IPv4 addresses within the IPv6 address syntax. Dual-stack applications are broadly compatible with these addresses as the Linux kernel transparently supports them when using the BSD socket API, as well as other operating systems. An application using Happy Eyeballs will directly establish a transport connection towards these addresses as they are carried in AAAA records. Using HE v3 [33] and the `SvcPriority` attribute of SVCB DNS records is a more elegant solution but the former is not widely available yet.

4.2 Selecting the best address family

When our DNS resolver receives a query in a dual-stack network, it has to decide which address family yields the lowest latency towards the requested domain name to choose which address it will answer from its cache. There are several ways of tackling this problem. A first approach is for the DNS resolver to asynchronously perform pings towards

the requested domain name using both address families to inform its next choice towards this destination. It has several drawbacks, first it considers the ping time which can differ from the latency experienced by transport protocols due to network queuing and prioritisation. Second, it narrows the available information for the DNS resolver to choose the address family to a single metric, while metrics such as loss rates could be relevant to some latency-sensitive application. Third, it adds load on the network and destinations.

Another approach would leverage passive measurements, which alleviates these three drawbacks. Researchers have proposed efficient latency monitoring solutions that could be integrated in this work [41, 39]. This approach can be taken by an on-path DNS resolver that would extract from TCP headers transport-level metrics such as connection establishment time and retransmissions. Off-path DNS resolvers could receive similar metrics using network telemetry. Based on this knowledge, the DNS resolver can discern which address family is better for each destination or prefix. For instance, a rolling history of past performance towards a given domain name could be used to select the most favourable address family. However, when steering the clients, the DNS resolver also has to explore the address families in order to rank them.

This tradeoff between exploration of the families and exploitation of the best family is classical of multi-armed bandit problems, which are a class of reinforcement learning problems. In our context, each address family is a competing alternative action, as only one can be taken at a time, i.e. a transport connection can only be established using one address family. Each established connection can be then evaluated to estimate the gain associated with using its address family towards a given destination by extracting transport metrics (e.g. RTTs, loss rate, congestion window, ...).

The EXP3 algorithm. The Exponential-weight algorithm for Exploration and Exploitation (EXP3) is a solution to this problem with several interesting properties [3, 10]. First, it makes no statistical assumptions on the process generating these gains. Second, it is resistant to adversarial models. Third, it was proven to converge towards the best action at a constant rate. Fourth, it is very computationally lightweight as its learning phase consists of a few float exponentiations and multiplications. EXP3 sets probabilities on available actions and updates them based on the gain obtained for each

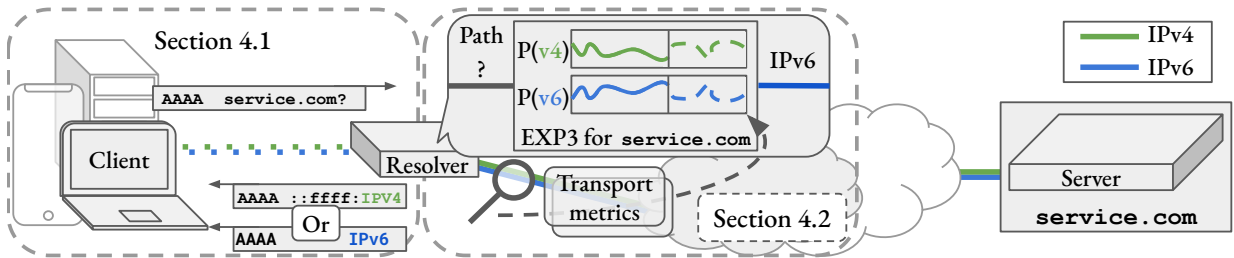


Figure 3: A client in a dual-stack network using the DNS resolver selecting IPv6 as the best performing address family towards a given domain name based on an online learning algorithm fed by transport metrics.

action taken. Initial probabilities are set equally. EXP3 defines a γ parameter within the range $[0, 1]$ which controls the balance between exploitation and exploration. When $\gamma = 1$, EXP3 only explores actions but never considers what it learns, acting as a random choice. Decreasing γ decreases both the learning rate and the exploration rate, i.e. the closer γ is to 0, the slower EXP3 learns but the more it will restrict its exploration to rather exploit the best action.

The right-hand side of Figure 3 illustrates how an instance of EXP3 estimates probabilities for each address family for a given destination ($P(v4)$ and $P(v6)$). These probabilities are updated based on the gain observed from transport metrics for each taken action. They are used by the DNS resolver to perform a probabilistic choice on the address family it selects for a given destination.

5 VALIDATION

In order to validate our solution in a controlled environment, we leverage the RIPE data presented in Section 3 to test its performance over a large amount of data. We extract more than 400k pairs of probes and anchors with at least 300 IPv4 and IPv6 HTTP tests. Given this large amount of probes and anchors, we can use this data to evaluate our solution’s ability to choose the address family with the lowest request completion time (RCT) in a large number of cases. To do so, we developed a Python simulator and an EXP3 implementation consisting of about 250 lines of code [34].

We define a simulation for each pair of probe and anchor consisting of one round for each HTTP test from the probe to the anchor performed simultaneously over IPv4 and IPv6. At the simulation start, we initialise a EXP3 instance. A round consists of the following steps. First, the EXP3 instance draws the address family by forming a probabilistic choice. Second, we record whether the RCT for the chosen family is lower compared to the other family for this round. Third, we compute the gain based on the chosen family to reward and train the EXP3 instance using the following function: we give the entire reward for the chosen address family when the obtained RCT is lower than the last RCT obtained when

EXP3 chose the other address family. Otherwise, we assign no reward to the action.

We set the γ EXP3 parameter to 0.1, which makes EXP3 reach 90 % probability for the best action within 60 rounds when it is strictly better than the other action. When performing a simulation, we split the first 60 rounds for training and the rest for evaluation, during which we compute the percentage of rounds where the family with the lowest RCT was chosen. We seed the random number generator (RNG) once and run each simulation 100 times to alleviate the RNG effect on the final score. We keep the median percentage of the runs as the final score for each simulation.

Figure 4 reports the Cumulative Distribution Functions (CDF) of three metrics computed from the results of our simulations. We compare these three metrics in three categories (Table 1) across three methods. These methods are: (plain lines) our EXP3 approach, (dashed lines) a fixed choice picking the address family with the largest amount of lower RCT *a posteriori* for this simulation, (dashed and dotted lines) a fixed choice picking the address family that performs the best for all simulations of this probe. The second method acts as an upper bound on the impact of selecting the best address family in a simulation while the third highlights the impact of selecting the best address family as a whole for a given probe rather than per destination.

On the left-hand side of Figure 4, we report the CDF of the median ratio of best choices, i.e. when the chosen address family leads to the lowest RCT, for each simulation. We can observe that when a simulation exhibits a significant difference between the completion times of IPv4 and IPv6, EXP3 is largely able to converge towards the best address family and yields a high ratio of best choices. The distance between EXP3 and our higher bound selecting the best path *a posteriori* for the simulation is constant and can be explained by the exploration factor of EXP3. Reducing this distance could be obtained by lowering the γ EXP3 parameter at the expense of a slower convergence. When our statistical test establishes no significant difference between IPv4 and IPv6, we can observe that EXP3 is adversely affecting performance




Figure 4: EXP3 is largely able to choose the address family with the lowest completion time in simulations using the RIPE HTTP tests data from probes to anchors.

in a very low number of simulations, i.e. simulations having a ratio less than 0.5. While EXP3 in most of these simulations gives a ratio lower than 0.55, it is also able to discern more subtle differences than established by *t-test* to select the best path. Indeed, the distance between EXP3 and selecting the best path a posteriori for the simulation is similar to the other categories.

When comparing EXP3 to our third method selecting the best path in average towards all destinations for a given probe, we observe an important regression for the latter with a different impact based on which family is best. When IPv4 exhibits a better performance, a higher percentage of simulations that choose the probe best path have strongly adverse performance. This can also be observed for IPv6 to a lower extent. This can be explained by several factors. First, IPv4 being better in a simulation is the rarest case, which means that for a probe, IPv4 being also the best path in average towards all destinations is rarest too. Second, when IPv4 is better, the distance from its mean completion time to IPv6 mean completion time is higher than in the opposite case. This could exacerbate the effect of the first factor.

To better understand the request completion time gain when selecting a given path, we can observe the middle and right-hand side figures which depicts the absolute expected gain per request and the gain relative to the median RCT. When compared to selecting the best path a posteriori (dashed lines), we can observe that EXP3 yields a gain close to this upper bound. When comparing IPv4 and IPv6, we can observe that IPv4 has higher expected gain in a small number of cases. When IPv4 has lower RCTs as established by our statistical test, the distance from the IPv4 to the IPv6 mean RCT is higher than in the opposite case. We argue that this could be due to networks that are still supporting IPv6 using transitions solutions introducing a greater overhead than

the non-congruency of IPv4 and IPv6 paths alone. When observing the right-hand side figure expressing this expected gain as a ratio of the median RCT, we can observe that the gain yielded by EXP3 can be higher than 20% (resp. 14%) of the median RCT for 25% of the cases when IPv4 (resp. IPv6) is best address family. **We can conclude that EXP3 is able to converge towards the best address family yielding the lowest request completion time and resulting in a significant gain for a large number of pairs of probe and anchor.**

6 PROTOTYPE

After validating the suitability of EXP3 using the RIPE Atlas dataset, we implemented our solution in a DNS proxy [46] written in Rust. Our proxy receives DNS queries, queries an actual resolver and optionally alter its answers. We added 1000 lines of code to implement a DNS record cache, to improve the library used to manipulate DNS responses, and to implement EXP3 and our path selection technique.

When our prototype receives a DNS query, it uses EXP3 to probabilistically choose which address family should be used. We took several implementation decisions in this process. First, our prototype only acts when queries can be resolved from its cache in order to avoid introducing additional delays during the resolution process. When the cache is empty, our prototype simply acts as a relay to an actual resolver and caches the received responses. Second, we group in a single EXP3 instance destinations which share a common second-level domain name. This balances the need for dedicating an EXP3 instance to each destination with the benefit of grouping destinations experiencing similar address family performance differences for accelerated learning. We make the assumption that IP addresses resolved from domain names which share a common second-level domain name are likely to experience similar performance differences. Third,

our prototype can also receive out-of-band feedback on the performance of address families towards destinations via a custom UDP message format. We use this feedback to compute EXP3 rewards using a similar function as used in Section 5: we give the full reward when the address family used gave a better performance than the last use of the other family, and none otherwise.

When receiving an AAAA query, our prototype leverages IPv4-mapped IPv6 addresses to select the address family to use towards the requested domain name. When an A query is received, our prototype respond with an empty answer to force the use of the selected address in the AAAA query. While the Happy Eyeballs algorithm as presented in Section A recommends using the address received in the AAAA record, we found Chrome to wait for the A response to arrive.

7 REAL-WORLD EXPERIMENTS

We now evaluate our prototype using real networks and towards popular Internet services. While RIPE anchors are isolated servers to which a transport connection is established directly, popular services often have distributed infrastructures in several locations and place their servers behind load-balancers that can terminate the transport connections close to the users. We argue that measuring the end-to-end latency in this context for our work is relevant as it brings a different kind of destinations into perspective.

To find popular services, we use the Tranco list [24] generated on the 5th of April 2023 [25] and choose the top 40 domain names for which an IPv4 and IPv6 address can be resolved from our campus network. The resulting list has a diversity of companies and organisations providing popular services to a significant part of the Internet users.

We use the browserless Docker image [9] to evaluate the impact of our prototype on Chrome performance in a reproducible environment. We configure Chrome to use our prototype as a DNS resolver, which leverages the DNS resolver within the network where experiments are performed. During an experiment, our prototype is started without any prior knowledge and Chrome is instructed to load websites one at a time from a list based on the top 40 domain names. To measure our prototype ability to learn and exploit the best address family towards destinations, we repeat in the list each domain name 30 times. The list is then shuffled to avoid bias. When running the experiment, we load each website from the list in three modes: IPv4 only, IPv6 only, and our prototype. This gives us a reference point for each address family to compare the quality of the choices of our prototype. These modes are also run in random order and using separate instances such that no previous run from other modes influences our prototype nor the browser. After a website is fully loaded using our prototype as a DNS resolver, we extract




Figure 5: Our prototype is largely able to choose the best address family in experiments using four access networks when loading popular websites with Chrome.

	Campus	DSL	Cable	4G FWA
Prototype	11.8 ms	15.5 ms	17.9 ms	63.6 ms
Best	7.3 ms	14.2 ms	17.2 ms	63.1 ms
Worst	14.3 ms	20.7 ms	18.2 ms	67.5 ms

Table 3: Medians reported in Figure 5.

the handshake time of all transport connections established for the website. Then, the rewards for EXP3 are computed and attributed to the corresponding domain names using the same formula as described in Section 5.

We perform this experiment from several vantage points in Belgium: (1) our university campus network, (2) a DSL line, (3) a cable line, (4) a 4G Fixed Wireless Access network. All four networks are operated by different ISPs. When measuring from a vantage point, the Linux test device is connected using Ethernet. We collect all transport connections handshake times during these experiments. Then, we process them such that the first 60 tests towards a destination are not considered as they correspond to the training phase of our prototype. Then we keep the destinations for which at least 100 tests exist, such that a significant amount of connections was steered using our trained prototype. Finally, to avoid over-representing websites that use more transport connections than others, and destinations that are more frequent among these websites, such as common CDN resources, we compute the mean handshake time for each destination. We repeat this process for the two other modes.

Figure 5 reports the distributions of the mean handshake time per destination when loading these popular websites using our prototype on these four access networks. We compare them to the best and worst address families for each destination. We determine a posteriori which address family is better or worse using the handshake times experienced in the IPv4 and IPv6 only modes. When looking first at each distribution and Table 3, we can observe that our prototype

improves the mean handshake time by selecting the best address family for each destination. In addition, the tail of the distributions is also reduced. There is no network in which our prototype degrades the handshake times. Its mean handshake time distribution is close to the ideal solution, and its small deviation can be explained by the EXP3 exploration. When looking at each network, the DSL line experiences the largest differences in terms of address family performance.

8 DISCUSSION AND FURTHER WORK

This work opens several directions for further work. First, we argue that the DNS protocol should be extended to allow indicating the address family that a host should use with a dedicated additional record. The Happy Eyeballs algorithm could be updated to take those hints into account when sorting the received addresses. Second, the problem of address family selection in dual-stack networks being a subset of the path selection problem in multihomed networks, our approach could be extended to IPv6-multihomed networks. An enterprise subscribed to several network providers could use its DNS resolver to hint the source prefix to use towards a given destination. Finally, our approach could leverage other transport metrics than the connection handshake time. For instance, some latency-sensitive applications such as video conferencing also require a low loss rate and jitter while others might require a more stable bandwidth.

ACKNOWLEDGMENTS

We thank Louis Navarre and François Michel for their comments and feedback that helped improve this work.

REFERENCES

- [1] Markus Amend, Veselin Rakocevic, Andreas Philipp Matz, and Eckard Bogenfeld. 2018. Robe: robust connection establishment for multipath tcp. In *Proceedings of the Applied Networking Research Workshop*, 76–82.
- [2] Florian Aschenbrenner, Tanya Shreedhar, Oliver Gasser, Nitinder Mohan, and Jörg Ott. 2021. From single lane to highways: analyzing the adoption of multipath tcp in the internet. In *2021 IFIP Networking Conference (IFIP Networking)*. IEEE, 1–9.
- [3] Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. 2002. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32, 1, 48–77.
- [4] Vaibhav Bajpai and Jürgen Schönwälder. 2019. A longitudinal view of dual-stacked websites—failures, latency and happy eyeballs. *IEEE/ACM Transactions on Networking*, 27, 2, 577–590.
- [5] Vaibhav Bajpai and Jürgen Schönwälder. 2015. Ipv4 versus ipv6—who connects faster? In *2015 IFIP Networking Conference (IFIP Networking)*. IEEE, 1–9.
- [6] Vaibhav Bajpai and Jürgen Schönwälder. 2013. Measuring tcp connection establishment times of dual-stacked web services. In *Proceedings of the 9th International Conference on Network and Service Management (CNSM 2013)*. IEEE, 130–133.
- [7] Vaibhav Bajpai and Jürgen Schönwälder. 2016. Measuring the effects of happy eyeballs. In *Proceedings of the 2016 Applied Networking Research Workshop*, 38–44.
- [8] Bob Briscoe, Koen De Schepper, Marcelo Bagnulo, and Greg White. 2023. Low Latency, Low Loss, and Scalable Throughput (L4S) Internet Service: Architecture. RFC 9330. (Jan. 2023). doi: 10.17487/RFC9330.
- [9] browserless.io. 2023. browserless. (2023). <https://github.com/browserless/chrome>.
- [10] Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. 2012. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5, 1, 1–122.
- [11] Brian E. Carpenter and Keith Moore. 2001. Connection of IPv6 Domains via IPv4 Clouds. RFC 3056. (Feb. 2001). doi: 10.17487/RFC3056.
- [12] Lorenzo Colitti, Steinar H Gunderson, Erik Kline, and Tiziana Refice. 2010. Evaluating ipv6 adoption in the internet. In *International Conference on Passive and Active Network Measurement*. Springer, 141–150.
- [13] Jakub Cyz, Mark Allman, Jing Zhang, Scott Iekel-Johnson, Eric Osterweil, and Michael Bailey. 2014. Measuring ipv6 adoption. In *Proceedings of the 2014 ACM conference on SIGCOMM*, 87–98.
- [14] Alun Davies. 2020. RIPE Atlas on BigQuery. (2020). <https://labs.ripe.net/tools/ripe-atlas-on-bigquery/ripe-atlas-on-bigquery/>.
- [15] Quentin De Coninck and Olivier Bonaventure. 2017. Multipath quic: design and evaluation. In *Proceedings of the 13th international conference on emerging networking experiments and technologies*, 160–166.
- [16] Dr. Steve E. Deering and Bob Hinden. 2006. IP Version 6 Addressing Architecture. RFC 4291. (Feb. 2006). doi: 10.17487/RFC4291.
- [17] Amogh Dhamdhere, Matthew Luckie, Bradley Huffaker, KC Claffy, Ahmed Elmokashfi, and Emile Aben. 2012. Measuring the deployment of ipv6: topology, routing and performance. In *Proceedings of the 2012 Internet Measurement Conference*, 537–550.
- [18] Alan Ford, Costin Raiciu, Mark J. Handley, Olivier Bonaventure, and Christoph Paasch. 2020. TCP Extensions for Multipath Operation with Multiple Addresses. RFC 8684. (Mar. 2020). doi: 10.17487/RFC8684.
- [19] Roch Guérin and Kartik Hosanagar. 2010. Fostering ipv6 migration through network quality differentials. *ACM SIGCOMM Computer Communication Review*, 40, 3, 17–25.
- [20] Christian Huitema. 2006. Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs). RFC 4380. (Feb. 2006). doi: 10.17487/RFC4380.
- [21] Apple Inc. 2023. TN3151: Choosing the right networking API. (2023). <https://developer.apple.com/documentation/technotes/tn3151-choosing-the-right-networking-api#Connect-by-name>.
- [22] Jana Iyengar and Martin Thomson. 2021. QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000. (May 2021). doi: 10.17487/RFC9000.
- [23] Vladimir Ksinant, Christian Huitema, Dr. Susan Thomson, and Mohsen Souissi. 2003. DNS Extensions to Support IP Version 6. RFC 3596. (Oct. 2003). doi: 10.17487/RFC3596.
- [24] Victor Le Pochat, Tom Van Goethem, Samaneh Tajalizadehkhoob, Maciej Korczyński, and Wouter Joosen. 2019. Tranco: a research-oriented top sites ranking hardened against manipulation. In *Proceedings of the 26th Annual Network and Distributed System Security Symposium (NDSS 2019)*. (Feb. 2019). doi: 10.14722/ndss.2019.23386.
- [25] Tranco list. [n. d.] Information on the Tranco list with ID 6JL4X. (). <https://tranco-list.eu/list/6JL4X>.
- [26] Yanmei Liu, Yunfei Ma, Quentin De Coninck, Olivier Bonaventure, Christian Huitema, and Mirja Kühlewind. 2024. Multipath Extension for QUIC. Internet-Draft draft-ietf-quic-multipath-10. Work in Progress. Internet Engineering Task Force, (July 2024). 36 pp. <https://datatracker.ietf.org/doc/draft-ietf-quic-multipath/10/>.
- [27] Ioana Livadariu, Simone Ferlin, Özgü Alay, Thomas Dreiholz, Amogh Dhamdhere, and Ahmed Elmokashfi. 2015. Leveraging the ipv4/ipv6 identity duality by using multi-path transport. In *2015 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 312–317.
- [28] Microsoft. 2023. Tcpclient.connect method. (2023). [https://learn.microsoft.com/en-us/dotnet/api/system.net.sockets.tcpclient.connect?view=net-7.0#system-net-sockets-tcpclient-connect\(system-string-system-int32\)](https://learn.microsoft.com/en-us/dotnet/api/system.net.sockets.tcpclient.connect?view=net-7.0#system-net-sockets-tcpclient-connect(system-string-system-int32)).
- [29] RIPE NCC. 2023. RIPE Atlas. (2023). <https://atlas.ripe.net>.
- [30] RIPE NCC. 2023. RIPE Atlas Service Terms and Conditions. (2023). <https://www.ripe.net/about-us/legal/ripe-atlas-service-terms-and-conditions>.
- [31] Mehdi Nikkhah and Roch Guérin. 2015. Migrating the internet to ipv6: an exploration of the when and why. *IEEE/ACM Transactions on Networking*, 24, 4, 2291–2304.
- [32] Mehdi Nikkhah, Roch Guérin, Yiu Lee, and Richard Woundy. 2011. Assessing ipv6 through web access a measurement study and its findings. In *Proceedings of the Seventh conference on emerging Networking EXperiments and Technologies*, 1–12.
- [33] Tommy Pauly, David Schinazi, Nidhi Jaju, and Kenichi Ishibashi. 2024. Happy Eyeballs Version 3: Better Connectivity Using Concurrency. Internet-Draft draft-pauly-v6ops-happy-eyeballs-v3-01. Work in Progress. Internet Engineering Task Force, (Mar. 2024). 19 pp. <https://datatracker.ietf.org/doc/draft-pauly-v6ops-happy-eyeballs-v3/01/>.
- [34] Maxime Piroux. 2023. Artefacts for this work. (2023). <https://github.com/mpiroux/cr23-artefacts/>.
- [35] Maxime Piroux, Tom Barbette, Nicolas Rybowski, Louis Navarre, Thomas Alfroy, Cristel Pelsser, François Michel, and Olivier Bonaventure. 2022. The multiple roles that ipv6 addresses can play in today’s internet. *SIGCOMM Comput. Commun. Rev.*, 52, 3, (Sept. 2022), 10–18. doi: 10.1145/3561954.3561957.

- [36] Costin Raiciu, Christoph Paasch, Sebastien Barre, Alan Ford, Michio Honda, Fabien Duchene, Olivier Bonaventure, and Mark Handley. 2012. How hard can it be? designing and implementing a deployable multipath {tcp}. In *9th USENIX symposium on networked systems design and implementation (NSDI 12)*, 399–412.
- [37] David Schinazi and Tommy Pauly. 2017. Happy Eyeballs Version 2: Better Connectivity Using Concurrency. RFC 8305. (Dec. 2017). doi: 10.17487/RFC8305.
- [38] Benjamin M. Schwartz, Mike Bishop, and Erik Nygren. 2023. Service Binding and Parameter Specification via the DNS (SVCB and HTTPS Resource Records), RFC 9460. (Nov. 2023). doi: 10.17487/RFC9460.
- [39] Satadal Sengupta, Hyojoon Kim, and Jennifer Rexford. 2022. Continuous in-network round-trip time monitoring. In *Proceedings of the ACM SIGCOMM 2022 Conference*, 473–485.
- [40] Internet Society. 2024. Pulse. (2024). <https://pulse.internetsociety.org/technologies>.
- [41] David Soldani, Petrit Nahi, Hami Bour, Saber Jafarizadeh, Mohammed Soliman, Leonardo Di Giovanna, Francesco Monaco, Giuseppe Ognibene, and Fulvio Risso. 2023. Ebpf: a new approach to cloud-native observability, networking and security for current (5g) and future mobile networks (6g and beyond). *IEEE Access*.
- [42] Dave Thaler, Richard P. Draves, Arifumi Matsumoto, and Tim Chown. 2012. Default Address Selection for Internet Protocol Version 6 (IPv6). RFC 6724. (Sept. 2012). doi: 10.17487/RFC6724.
- [43] G Dorcas Wambui, Gichuhi Anthony Waititu, and Anthony Wanjoya. 2015. The power of the pruned exact linear time (pelt) test in multiple changepoint detection. *American Journal of Theoretical and Applied Statistics*, 4, 6, 581.
- [44] Chengke Wang, Hao Wang, Feng Qian, Kai Zheng, Chenglu Wang, Fangzhu Mao, Xingmin Guo, and Chenren Xu. 2023. Experience: a three-year retrospective of large-scale multipath transport deployment for mobile applications. In *Proceedings of the 29th Annual International Conference on Mobile Computing and Networking*, 1–15.
- [45] Dan Wing and Andrew Yourtchenko. 2012. Happy Eyeballs: Success with Dual-Stack Hosts. RFC 6555. (Apr. 2012). doi: 10.17487/RFC6555.
- [46] wyhaya and Maxime Piraux. 2023. updns. (2023). <https://github.com/mpiraux/updns>.
- [47] Sebastian Zander, Lachlan LH Andrew, Grenville Armitage, Geoff Huston, and George Michaelson. 2012. Investigating the ipv6 teredo tunnelling capability and performance of internet clients. *ACM SIGCOMM Computer Communication Review*, 42, 5, 13–20.
- [48] Johannes Zirngibl, Patrick Sattler, and Georg Carle. 2023. A first look at svcb and https dns resource records in the wild. In *2023 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*. IEEE, 470–474.

A HAPPY EYEBALLS

Starting during the early deployment phase of IPv6 and its start of coexistence with IPv4, IPv6-enabled applications faced the choice of the address family when establishing a transport connection to a domain name that resolves in both families. While IPv6 could be preferred to favour the IPv6 deployment, network paths where it could be failing or be largely inferior existed. The approach taken by the IETF to solve this problem is to provide to applications an address selection mechanism favouring IPv6 that could quickly fall back in case of unexpected connection delay. The algorithm is called Happy Eyeballs and was first standardized in RFC 6555 [45], and then revised in RFC 8305 [37].

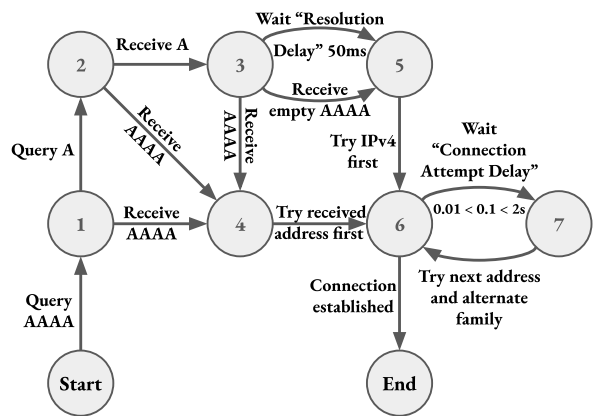


Figure 6: Happy Eyeballs Version 2 in RFC 8305.

Figure 6 illustrates the state machine of the Happy Eyeballs Version 2 algorithm as described in RFC 8305. The algorithm introduces several choices. IPv6 is favoured when choosing the first address to establish a transport connection. When the DNS resolution of IPv6 is delayed compared to IPv4 (state 3), the algorithm recommends waiting 50 ms before trying the IPv4 address first. When an empty AAAA DNS record is received, IPv4 is tried as soon as possible. When an address is tried, the algorithm specifies three time bounds before trying another address from another family. The next attempt delay must not be lower than 10 ms, with a recommended value of 100 ms and a maximal recommended value of 2 s.

B REPRODUCIBILITY

This Appendix details the required steps and material to reproduce the results presented in this work. All the code developed and data referenced in this section will be made publicly available under an open-source licence [34, 46]. An archive containing those two repositories was submitted along with this paper. The RIPE data that we downloaded and processed is subject to their legal terms [30].

B.1 IPv4 and IPv6 end-to-end latency

The RIPE data presented was first extracted using the BigQuery interface [14]. The query used is in the file `/e2e_latency/query.sql`. The resulting data should be saved in a CSV file. For this purpose, the helper Python script `/e2e_latency/download_bq_results.py` connects to BigQuery and downloads a `data.csv` file. It must be adapted with credentials and BigQuery Job ID. The expected file size is around 13.4 GB. The `/e2e_latency/figure_1.py` script produces Figure 1. We ran it on machine with 128 GB of RAM as it loads the entire CSV file.

To produce Table 1, 2 and the numbers cited in (2) of Section 3, the required data is produced by the Python simulator described in Section 5. The rest of the output of the simulator is discussed in the other parts of the Appendix. The simulator use a condensed view of the data, so the `data.csv` must be processed into the `paths.cbor` file using the `/simulations/prepare_data.py` Python script. This script was also run on machine with 128 GB of RAM. Then, running the `/simulations/ripe.py` Python script produces the required output. The rest of the Python scripts in the `/e2e_latency` folder (i.e `table_1.py`, `table_2.py`) can be run to obtain the corresponding Table.

B.2 Validation

Figure 4 requires running first the simulator as described in Appendix B.1. Then `/validation/figure_4.py` can be run to produce the Figure.

B.3 Real-world experiments

The `/experiments` folder contains the scripts necessary to run experiments and parse the results into the JSON format used to produce Figure 5. The real-world experiments require a Linux machine with Docker installed. We used the `browserless/chrome` Docker image hosted on `docker.io` with ID `e6460f3f8f60`. Our prototype is available in a separate repository [46]. Its Docker image should be built before running experiments. Then, the `/experiments/run_browserless.py` Python script starts all the experiments described Section 7 over a given network. This script needs to be adapted to specify the interface to use through indicating its IPv4 address in the `IP` global variable. The results are stored individually in a directory. This directory can be parsed by the `/experiments/parse_netlogs.py` Python script into a JSON file format summarising all results. The `/experiments/figure_5.py` Python script illustrates how Figure 5 was produced and should be modified when the number of networks tested changes.