Rusty Clusters? Dusting an IPv6 Research Foundation

Johannes Zirngibl
Technical University of Munich
Germany
zirngibl@net.in.tum.de

Lion Steger
Technical University of Munich
Germany
stegerl@net.in.tum.de

Patrick Sattler
Technical University of Munich
Germany
sattler@net.in.tum.de

Oliver Gasser
Max Planck Institute for Informatics
Germany
oliver.gasser@mpi-inf.mpg.de

Georg Carle
Technical University of Munich
Germany
carle@net.in.tum.de

ABSTRACT

The long-running *IPv6 Hitlist* service is an important foundation for IPv6 measurement studies. It helps to overcome infeasible, complete address space scans by collecting valuable, unbiased IPv6 address candidates and regularly testing their responsiveness. However, the Internet itself is a quickly changing ecosystem that can affect long-running services, potentially inducing biases and obscurities into ongoing data collection means. Frequent analyses but also updates are necessary to enable a valuable service to the community.

In this paper, we show that the existing hitlist is highly impacted by the Great Firewall of China, and we offer a cleaned view on the development of responsive addresses. While the accumulated input shows an increasing bias towards some networks, the cleaned set of responsive addresses is well distributed and shows a steady increase.

Although it is a best practice to remove aliased prefixes from IPv6 hitlists, we show that this also removes major content delivery networks. More than 98 % of all IPv6 addresses announced by Fastly were labeled as aliased and Cloudflare prefixes hosting more than 10 M domains were excluded. Depending on the hitlist usage, *e.g.*, higher layer protocol scans, inclusion of addresses from these providers can be valuable.

Lastly, we evaluate different new address candidate sources, including target generation algorithms to improve the coverage of the current *IPv6 Hitlist*. We show that a combination of different methodologies is able to identify 5.6 M new, responsive addresses. This accounts for an increase by 174 % and combined with the current *IPv6 Hitlist*, we identify 8.8 M responsive addresses.

CCS CONCEPTS

• **Networks → Network structure**; **Naming and addressing**; **Network measurement**.

KEYWORDS

IPv6, Hitlist, Aliased Prefixes, Target Generation

1 INTRODUCTION

The usage and importance of IPv6 are steadily increasing [6, 24, 37]. With the IPv4 address depletion of all but one Regional Internet Registry (RIR), the necessity to deploy IPv6 is prevalent for more and more operators and content providers. While this development is generally positive, it imposes fundamental difficulties to research and network analysis. The immense size of the address space, combined with the sparse distribution of used addresses renders active IPv6 measurements difficult. While tools like ZMap effectively scan the complete IPv4 address space, complete scans for IPv6 are impossible.

With the publication of the *IPv6 Hitlist* service in 2018, Gasser *et al.* [22] established an ongoing service that collects IPv6 address candidates, identifies aliased prefixes, and tests the responsiveness in respect to different protocols, namely ICMP, TCP on port 80 (HTTP) and 443 (HTTPS), and UDP on port 53 (DNS) and 443 (QUIC) (see Figure 1). The service is well maintained, up to date and has been used as de-facto standard for IPv6 analysis and scans, *e.g.*, [2, 7, 16, 38, 42, 51]. However, it has not seen significant updates or analyses since its initial publication. Changes in the usage of IPv6 and input sources might have influenced the quality of the service since 2018. Furthermore, a variety of methodologies to *generate* likely responsive addresses emerged [13–15, 29, 49] but have not been established as ongoing service publishing data.

Our main contributions in this work are:

(i) We evaluate the development of the *IPv6 Hitlist* over the last four years and new biases introduced by the accumulation of new addresses. Our findings allow us to filter targets incorrectly tested as responsive. We identify 134 M addresses falsely reported as responsive to UDP/53 by the *IPv6 Hitlist* since 2018 due to the Great Firewall of China's DNS injection.

(ii) We analyze aliased prefixes in more detail and investigate whether the initial definition of a single host responsive to a complete prefix remains correct or whether a set of addresses needs to be treated differently. We show that aliased prefixes host at least 15 M domains including ranked domains from different top lists [1, 12, 31]. In combination with additional findings, we suggest
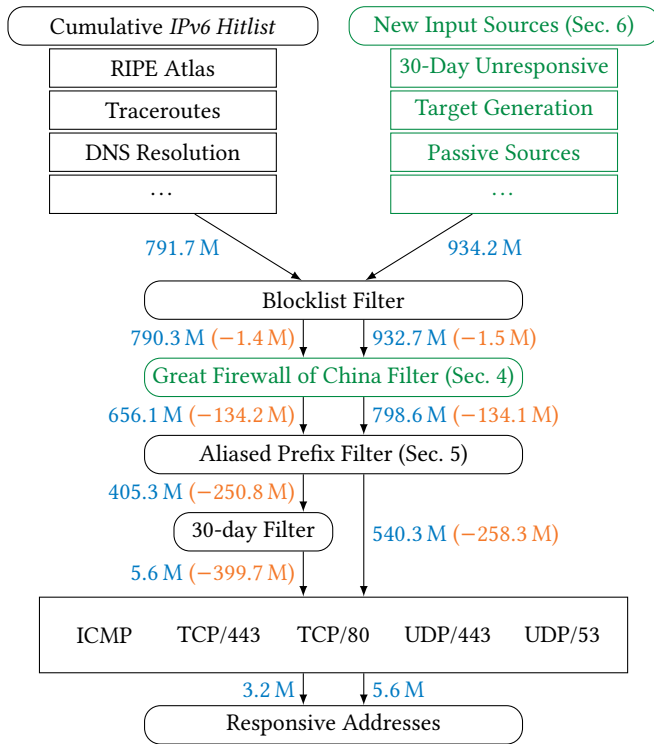
Figure 1: *IPv6 Hitlist* pipeline. This work analyzes the existing steps and adds a Great Firewall of China filter (see Sec. 4). Furthermore, the filtered, aliased prefixes are analyzed in Sec. 5 and new input sources are evaluated in Sec. 6. Newly implemented service components are indicated by green borders.

users of the hitlist to include subsets of these prefixes in future research.

(iii) We evaluate additional input sources and address generation mechanisms, and whether we can identify new, responsive addresses to improve the existing hitlist and support future research. We identify 5.6 M new, responsive addresses, potentially improving coverage of the *IPv6 Hitlist* by 174 %.

(iv) We update the *IPv6 Hitlist* service [22] to allow future research to use our findings within the established service.

<div align="center">https://ipv6hitlist.github.io/</div>

We present related work in Section 2 and introduce used data sources in Section 3. In Section 4, we analyze the development of IPv6 and the *IPv6 Hitlist* service [22] in detail. We analyze aliased prefixes in Section 5, followed by an evaluation of potential new input sources in Section 6. Finally, we discuss our findings in Section 7 and conclude in Section 8.

2 RELATED WORK

Initial attempts to establish IPv6 hitlists have been conducted, *e.g.*, by Gasser et al. [23] in 2016 and Fiebig et al. [19] in 2017. Gasser et al. [23] collected addresses from passive traces and active sources such as traceroutes and Domain Name System (DNS) resolutions, while

Fiebig et al. [19] relied on reverse DNS to identify used addresses. In 2016, Foremski et al. [20] modeled IPv6 addresses based on their entropy and derived *Entropy/IP*, relying on structural similarities between addresses to generate target candidates based on address seeds. Based on these findings, Murdock et al. [35] developed 6Gen in 2017 and derived an IPv6 hitlist containing 55 M active addresses. However, they identify 98 % of these as aliased. In 2018, Gasser et al. [22] combined the earlier approaches and analyzed the composition of the resulting hitlist and its value. They showed biases within different data sources and the importance of detecting aliased prefixes to remove additionally induced biases. Based on this work, they established an ongoing service accumulating addresses from a variety of services and regularly testing their responsiveness.

Throughout the following years, different IPv6 address generation algorithms were published [13–15, 25, 29, 47, 49]. All approaches rely on the assumption that IPv6 addresses contain patterns due to assignment strategies, that allow research to guess new, responsive addresses. They differ based on selected address representations and machine learning approaches. Cui et al. [14] use General Adversary Networks while Liu et al. [29] represent addresses as a space tree. The according studies mostly rely on the *IPv6 Hitlist* [22] and evaluate their generated lists with active scans. They were able to generate hitlists containing more than one billion candidates and reach hit rates of up to 50 % [47]. However, as of May 2022, we only found a snapshot of generated addresses from DET [47] and other sources could not be verified or reused for further studies. We use 6Tree [29], 6Graph [49], 6GAN [14], and 6VecLM [15] to generate addresses as new input sources in this work and show the potential of some algorithms. While we are not able to reproduce their published hit rates, we show that some algorithms can be used as valuable new sources in addition to the existing *IPv6 Hitlist*.

Besides the collection or generation of new address candidates, the detection of aliased prefixes is important to understand and remove induced biases to IPv6 hitlists. Different alias detection methodologies have been proposed to detect different addresses of the same router [10, 30, 32, 39, 48] and to detect IPv4 and IPv6 siblings [44]. However, most of these exploit side channels not available on every target. Therefore, Murdock et al. [35] proposed an alias detection for IPv6 hitlists based on the responsiveness of random addresses within prefixes of size /96. Gasser et al. [22] extended this idea to a multi-level aliased prefix detection on different prefix lengths. Furthermore, they combine TCP/80 and ICMP probes together with results from previous days to account for probe timeouts. They verified the effectiveness of their approach to identify aliases with TCP fingerprints.

Song et al. [47] suggested using the Too Big Trick (TBT) introduced by Beverly et al. [10] to evaluate real aliases. They show that some prefixes identified as aliased by the *IPv6 Hitlist* might contain multiple hosts, although being fully responsive. In this work, we combine detected aliases based on the multi-level detection method, with TCP fingerprinting, the TBT, and information about specific Autonomous Systems (ASes) and hosted domains to shed further light on identified aliased prefixes.

While the Great Firewall of China does not necessarily seem related to IPv6 hitlists, we show its relation and impact in Section 4,

mainly its DNS injection behavior. It has been analyzed in a different context [4, 5, 18] and has been recently used as a side channel to actively analyze DNS root server performance in China by Zhang et al. [50]. Most importantly, these works find that the GFW is used to inject DNS responses for censored domains at the border of Chinese networks. Anonymous et al. [5] analyzed the behavior in 2020 observing multiple responses that can be mapped to different injectors. They further find that erroneous responses normally contain generally routed, valid IP addresses. However, these addresses can be mapped to operators unrelated to the requested domain. We identify the GFW to be responsible for a majority of addresses responsive to DNS probes. Results show similar behavior as reported by related work, highly impacting the quality of the *IPv6 Hitlist*. However, we currently observe different addresses in all responses.

3 DATA SOURCES AND SCANS

Our work relies on different public data sources but also individual scans. We introduce all sources and used tools in the following section. For our own scans, we strictly follow ethical considerations as explained in Section 3.3.

3.1 IPv6 Hitlist Service

Our primary source is the *IPv6 Hitlist* published by Gasser *et al.* [22]. The service collects address candidates from multiple sources, including DNS AAAA resolutions, conducted traceroutes, and public sources, *e.g.*, from RIPE Atlas and CT logs. The service frequently *(i)* updates addresses, *(ii)* uses *all* collected addresses as input, *(iii)* applies multiple filters, and *(iv)* tests the responsiveness of addresses in respect to different protocols. Figure 1 shows the service pipeline.

The first filter removes the addresses of operators who requested exclusion of regular scans following ethical considerations described in Section 3.3. The most important existing filter is the aliased prefix detection. The initial definition of aliased prefixes describes a single host responsive for all addresses in a prefix. Each individual aliased prefix may be infeasible to scan, offers limited value to following scans, and introduces a bias. As shown in Section 5, aliased prefixes in our data had different lengths between /28 and /120. While earlier work from Murdock et al. [35] tests for aliased prefixes with a fixed length of /96, the aliased prefix detection of the *IPv6 Hitlist* tests prefixes of different lengths, including:

- IPv6 prefixes announced in BGP;
- all /64 prefixes with at least one address contained in the *IPv6 Hitlist* service input;
- and prefixes longer than /64 (in steps of four bit) with at least 100 addresses.

The implemented detection [22] relies on the assumption that it is highly unlikely that multiple randomly selected addresses within an IPv6 prefix are responsive. Therefore, the detection selects one random address within each of the 16 more specific prefixes (0-f) and uses ZMapv6 [21] to test responsiveness. To test whether prefix 2001:db8::/32 is aliased, a single, random address is tested within all subprefixes 2001:db8:[0-f]000::/36. This address generation distributes the pseudo-random targets evenly across the complete prefix. If all 16 addresses are responsive, the prefix is labeled as aliased. ICMP and TCP/80 are tested, and results are merged across

protocols and with the previous three scans. This reduces misclassification of prefixes, *e.g.*, due to random network events or packet loss during individual scans.

The final filter in Figure 1 removes all addresses that are unresponsive for at least 30 days. This filter reduces the required scan load drastically. However, these addresses are never tested for responsiveness again after exclusion. We re-scanned these addresses to test whether addresses are responsive after 30 days again as discussed in Section 6.

After all filter steps, the service executes traceroutes using Yarrp [9] to all targets to potentially identify new targets. Furthermore, ZMapv6 scans ICMP, TCP port 80 (HTTP) and 443 (HTTPS), and UDP port 53 (DNS) and 443 (QUIC). While scans were executed on a daily basis initially, the growth of the input set increased the overall runtime to several days. The data covers more than 750 scans between July 2018 and April 2022.

3.2 Additional Data Sources

Besides the *IPv6 Hitlist* service, we use additional data throughout this work to analyze the development of the hitlist and improve its current state for future research.

CAIDA Ark. While the *IPv6 Hitlist* contains its own traceroutes and data from RIPE Atlas, traceroutes conducted on a regular basis from the CAIDA Archipelago (Ark) Infrastructure are not yet included [11]. We use a snapshot of the data from March 2022 to analyze the value additional vantage points can offer to reveal new routers or whether the existing data set covers most contained targets.

DET. Song et al. [47] collected IPv6 addresses from different services similar to Gasser et al. [22] and additionally used target generation algorithms. They advertise an ongoing service and data publication. However, we are only able to download a single snapshot of responsive addresses and did not receive a reply requesting additional data. Thus, we only use this snapshot as a new source of addresses and evaluate its value.

IPv6 Scans. Besides the publicly available data from the *IPv6 Hitlist* service, we use the existing pipeline to conduct additional scans testing new IPv6 address sources and target generation methodologies analyzed in Section 6 (see Figure 1). We use ZMapv6 with the same probe modules, configurations and payloads as the *IPv6 Hitlist* service from the same vantage point. We further deploy the multi-level aliased prefix detection from Gasser et al. [22] and filter our scans with already known aliased prefixes and blocklists collected by the existing service.

DNS Scans. We received access to DNS scans regularly conducted by our institution. While these scans have not been explicitly set up for this study, ethical considerations described in Section 3.3 are still applied. These scans are already used as input for the existing *IPv6 Hitlist* service. The scans resolve domains from various input sources covering more than 300 M domains from the Centralized Zone Data Service (CZDS), including .com, .net, and .org, Certificate Transparency logs, and cc-TLDs. Additionally, the Alexa Top 1M [1], Majestic Top 1M [31], and Umbrella Top 1M [12] are resolved.

These scans resolve domains to IPv6 addresses (AAAA records) but also to name server (NS) and mail exchanger (MX) records. These resource records are resolved to their respective IPv6 addresses. While the direct AAAA resolution of domains is already used as input of the *IPv6 Hitlist* service, the name server and mail exchanger domains were not explicitly included. We use a single snapshot of these scans from April 7th, 2022 to analyze aliased prefixes in Section 5 and evaluate the quality of these domains as new input in Section 6 and include them into the *IPv6 Hitlist* service.

3.3 Ethics

We primarily rely on available data from the *IPv6 Hitlist* service. However, some analyses are based on new scans, *e.g.*, the study of aliased prefixes and new IPv6 address sources. During all scans, we follow ethical measures, *i.e.*, informed consent [17] and community best practices [40]. We scan with a limited rate and use a request-based blocklist. Furthermore, we use the current blocklist of the existing *IPv6 Hitlist* service to seed our blocklist. This prevents interferences with networks that already requested to opt-out of existing IPv6 scans. Our measurement vantage point is clearly identified. This includes reverse DNS, WHOIS information and a hosted website informing about our measurements. We did not receive any inquiries related to our scans during this work.

Additional DNS scans we rely on are conducted independently of this work but follow the same ethical principles, including a blocklist and public information about the scans and vantage point.

4 IPV6 HITLIST DEVELOPMENT

As explained in Section 3, the *IPv6 Hitlist* [22] is implemented as an ongoing service, regularly updating its input, identifying aliased prefixes (analyzed in more detail in Section 5), and testing the responsiveness of addresses. We analyze the overall hitlist and responsiveness of addresses in the following.

4.1 Input Development

Starting with 90 M addresses in July 2018, the service has accumulated more than 790 M addresses until April 2022. They cover 22 074 ASes compared to 10 866 in 2018 [22]. The 22 k ASes cover 76 % of ASes announcing at least one IPv6 prefix in BGP at the time of writing. The number of announced prefixes and ASes is based on a routing information base from RIPE RIS collector rrc00 [36].[1] Furthermore, the current input list covers four times more (97 k) announced BGP prefixes compared to 2018, 62 % of all announced prefixes. The visible growth is similar to the general growth of IPv6 deployments and usage on the Internet and shows that the hitlist is able to adapt to it.

Figure 2 shows the cumulative distribution of IPv6 addresses in the hitlist input across ASes. Without any filter, the most prominent AS is Amazon (AS16509), covering 32 % of all addresses. However, 99.6 % of these are filtered due to the aliased prefix detection. Section 4.2 explains the additional plots covering the GFW impact and responsive addresses.

Nevertheless, after filtering aliased prefixes, 80 % of the current input is still covered by only 10 ASes, mostly from Internet Service
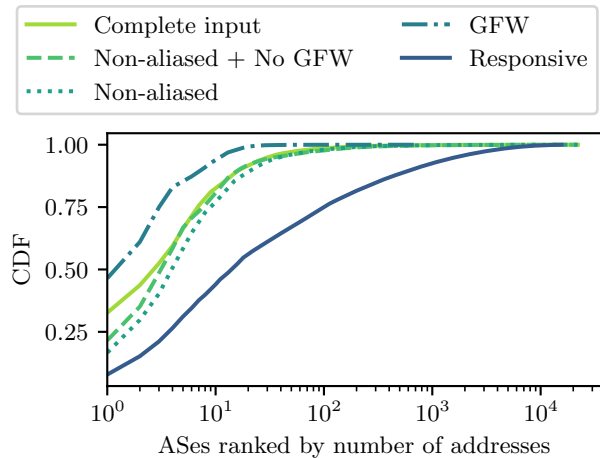

Figure 2: Distribution of addresses in the input list across ASes. Additionally, the effect of the alias and GFW filter are displayed. Note the log x-Axis. While the complete input is biased towards some ASes, established and new filters result in a well distributed set of responsive addresses.

Providers (ISPs) such as ANTEL (AS6057, 16 %) or DTAG (AS3320, 10 %). Analyzing the source of these addresses reveals that they are mostly accumulated due to regular traceroutes, especially from RIPE Atlas, changing prefix assignments and rotating addresses. 282 M addresses from the input contain a EUI-64 IPv6 Interface ID (IID) including ff:fe and are based on a MAC address. Extracting the IID reveals that these addresses are only derived from 22.7 M distinct MAC addresses. Grouping addresses based on the EUI-64 values shows that 9 M occur only within one IPv6 address each, while the remaining are seen in multiple addresses. The most frequent EUI-64 value can be seen in 240 k distinct IPv6 addresses. The Organizationally Unique Identifier (OUI) of the MAC address is mapped to a vendor (ZTE) and all addresses are part of the same /32 prefix, but within different subnets. Similarly, Rye et al. [43] found a variety of Customer-premises equipment (CPE) used EUI-64 IIDs and whose ISPs regularly rotated prefixes. They were able to track individual devices within ISPs.

These rotating addresses result in a visible bias of the complete input list towards a small set of ASes due to the ongoing accumulation of addresses. While the filter of unresponsive addresses after 30 days removes most of these addresses from regular scans and the responsive results afterward, the overall input list has to be treated carefully by other research relying on this data.

4.2 Address Responsiveness

Figure 3 shows the responsiveness of addresses throughout the *IPv6 Hitlist* service lifetime. It depicts the responsiveness for each tested protocol and a total count of addresses responsive to at least one protocol within each scan. The left half represents the published state, with clearly visible spikes which will be explained in the following, and on the right, a *cleaned* version is presented.
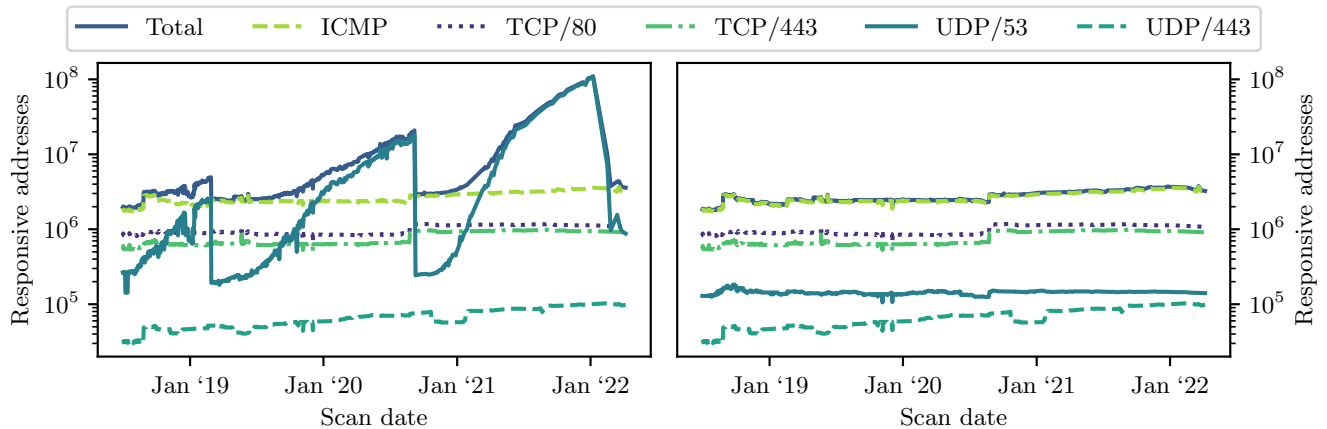
---

[1]https://data.ris.ripe.net/rrc00/2022.04/bview.20220407.0800.gz

Figure 3: Comparison of the *IPv6 Hitlist* as published (left) and cleaned from GFW injection (right). Without injected responses by the GFW, a steady development of the *IPv6 Hitlist* and each protocol is visible. Note the log y-Axis.

The Prevalence of DNS. The published hitlist contains significant spikes in responsive addresses as seen in Figure 3. During these events, the number of addresses responsive to DNS probes is larger compared to all other protocols. After each peak, the number of addresses drops to a similar level as before the event. The last event started in early 2021 with the most drastic growth of addresses, peaking at more than 100 M. In comparison, only 3.5 M and 1.4 M addresses were responsive to ICMP and TCP/80 respectively, and no increase was visible during the same periods. The peak drops in February 2022, after we implemented a filter based on the following findings and updated the service.

To understand the origin of these peaks, we describe the scan configuration and analyze results in detail. The service sends a DNS query requesting a AAAA record for `www.google.com`. All peak events share similarities but differ slightly. During the first two events, a significant fraction of addresses responded with A records only containing an IPv4 address. During the third and most recent event, responses carried valid AAAA records but contained Teredo addresses. Note that Teredo is a deprecated standard embedding an IPv4 into an IPv6 address [26]. Furthermore, ZMap accumulated two or three responses for each scanned address, with up to 440 responses in the worst case.

Analyzing the erroneous IPv4 addresses contained in A records during earlier events and embedded in the returned Teredo addresses in the latter event reveals that none can be associated to the Google ASes but other companies like Facebook, Microsoft or Dropbox. Collecting all IPv6 addresses that responded with a clearly erroneous record (IPv4 or Teredo address) throughout the four years accumulates to more than 134 M addresses (17 % of the cumulative *IPv6 Hitlist* input on April 7th, 2022).

Querying a different domain shows that these targets are not responsive themselves, but responses are injected and falsely interpreted as success by ZMap. Most addresses with this response behavior are announced by ASes of Chinese networks, *e.g.*, China Telecom Backbone (AS4134) and China Telecom (AS4812) originating 46.44 % and 14.59 % of impacted addresses respectively. Figure 2

compares the AS distribution of IPv6 addresses responding with these incorrectly responded addresses to the complete input. These addresses cover only 695 ASes. 93 % of addresses are located in only 10 Chinese ASes. See Table 5 in Appendix A for an overview of the Top 10. We used MaxMind GeoLite2 [33] as an additional indicator of network location. While we are aware of potential inaccuracies especially on a city level [41, 45], it mapped a majority of impacted IPv6 addresses to China. Given these indicators a strong relation of these addresses to Chinese networks can be seen.

The overall behavior has been described similarly in related work [4, 5, 18]. We see multiple responses to a single query indicating multiple injectors, responses are mostly in relation to addresses from China and `www.google.com` is a blocked domain. Querying different blocked domains from these addresses shows similar behavior. In contrast, a domain owned by ourselves, most likely not blocked, results in no response at all, not even a DNS error. A difference to existing related work is that the currently injected responses carry a Teredo address not explicitly reported in previous findings [4, 5, 18]. However, contained IPv4 addresses show similar behavior (cf.[5]) and can be mapped to previously identified incorrect networks.

The source of these addresses is the regularly conducted traceroutes by the *IPv6 Hitlist* service using Yarrp. Traceroute captures regularly changing addresses mostly with randomized IIDs and visible as the last responsive hop. The targeted address is not responsive itself. Scanning these addresses in the following with ZMap triggers a DNS injection by the GFW, but for a majority no other protocol is responsive. In some cases, these targets are actually responsive for other probed protocols. Thus, invalid DNS responses should be filtered but individual addresses should remain in the *IPv6 Hitlist* if responsive to other protocols.

Reducing the Great Firewall of China Impact. The *IPv6 Hitlist* service sends DNS queries for a blocked domain, namely `www.google.com`. These queries result in DNS injections from the GFW which were labeled as successful responses by ZMap. While changing the queried domain could prevent this, *e.g.*, to

Table 1: Development of responsive IPv6 addresses and covered ASes over four years. Results are based on cleaned data, removing GFW injected responses. For each year, a snapshot representing a single scan is used. The cumulative result covers all scans since the start of the *IPv6 Hitlist*.

Year	ICMP		TCP/443		TCP/80		UDP/443		UDP/53		Total	
	Addr.	ASes	Addr.	ASes	Addr.	ASes	Addr.	ASes	Addr.	ASes	Addr.	ASes
2018-07-01	1.7 M	10.1 k	550.6 k	5.8 k	832.1 k	6.2 k	31.0 k	0.9 k	129.1 k	5.1 k	1.8 M	10.3 k
2019-04-01	2.4 M	11.0 k	645.8 k	6.2 k	919.2 k	6.6 k	50.4 k	1.0 k	145.4 k	5.2 k	2.5 M	11.2 k
2020-04-01	2.3 M	11.7 k	632.8 k	6.6 k	836.2 k	6.9 k	67.7 k	1.3 k	148.4 k	5.1 k	2.4 M	11.9 k
2021-04-02	3.0 M	13.7 k	954.8 k	7.4 k	1.1 M	7.7 k	83.0 k	1.3 k	148.0 k	6.0 k	3.1 M	13.9 k
2022-04-07	3.1 M	15.4 k	910.8 k	7.9 k	1.0 M	8.2 k	98.1 k	2.0 k	140.7 k	6.0 k	3.2 M	15.7 k
Cumulative	45.3 M		6.7 M		8.6 M		2.5 M		200 k		46.8 M	

example.com, we decided to operate the *IPv6 Hitlist* service with its current configuration. This consistency in the service behavior increases the comparability of results over time. With a consistent scan, different results indicate changes in the target behavior and are not induced by the scan itself.

This decision to focus on consistency requires an ongoing data cleaning step. We initially implemented a filter removing 134 M IPv6 addresses to reduce the GFW impact. The *IPv6 Hitlist* service saw at least one DNS injection for these addresses during July 2018 and April 2022 but no response to any other protocol. This filter immediately reduced scan duration and impact on the Internet.

To reduce the future impact on the ongoing *IPv6 Hitlist* service, we filter the DNS/53 results directly after the scan. Therefore, the results correctly reflect the responsiveness of newly scanned input addresses. Thus, if addresses are responsive to any other protocol, they remain in the scan input. Otherwise, they are filtered by the 30-day filter (see Figure 1).

Evaluation of Remaining IPv6 Addresses Supporting DNS.
After removing GFW injected DNS responses, 140 k addresses responsive to DNS probes remain. We evaluate the quality of the remaining IPv6 addresses with an experiment including a domain under our control. To analyze the behavior of each scanned target, we query for a subdomain including a unique hash instead of a static domain. The name server returns a valid AAAA record for the requested domain. This approach allows us to map outgoing probes towards individual addresses to incoming requests at our name server. Within this filtered set of targets, 131.8 k (93.8 %) of probes result in valid DNS responses with status codes indicating errors, because these targets are either name servers or resolvers unwilling to resolve the requested domain recursively. 6.5 k (4.6 %) return regular responses with the correct AAAA record and according requests from the same IPv6 addresses are visible at our name server. 593 targets respond with a referral to name servers of the root or our domain's parent zone. Only 15 IPv6 addresses return a correct record, but the source addresses of incoming requests on our name server do not match the probed targets, *e.g.*, due to proxies or the usage of another interface.

The remaining 1.1 % of targets respond incorrectly but an analysis of results reveals nothing similar to the GFW injection. Responses contain for example incorrect status codes or referrals to *localhost*.

***Cleaned* IPv6 Hitlist**. We cleaned the historical data from GFW injected responses resulting in the timeline shown in the right part of Figure 3. This results in a relatively stable number of responsive addresses for all protocols respectively, each with a slight increase throughout the four-year period. This development is in line with statistics reported by RIPE NCC regarding the general growth of IPv6 assignments [37]. Most addresses responsive to at least one protocol are responsive to ICMP, followed by TCP/80 and TCP/443. Figure 10 in Appendix B shows the exact overlap between protocols. For ICMP, the responsiveness increased from 1.78 M addresses covering 10.1 k ASes to 3.15 M addresses in 15.5 k ASes. Results for the remaining protocols can be seen in Table 1.

While UDP/443 (QUIC) increased the most by a factor of three, it still shows the worst response rate below UDP/53, even though QUIC was finally standardized in May 2021 [27]. Interestingly, Zirngibl et al. [51] reported 210 k IPv6 addresses supporting QUIC in 2021, more than twice as many addresses as reported by the *IPv6 Hitlist*. However, they include addresses from DNS resolutions many of which are located in aliased prefixes (*e.g.*, from Cloudflare, Fastly, and Amazon) and thus filtered by the hitlist service. We analyze this effect in more detail in Section 5.

While an overall growth is visible throughout the four years, the number of responsive IPv6 addresses to ICMP, TCP/80 and TCP/443 slightly decreases in between the analyzed scans from 2019 and 2020 (see Table 1). This is primarily due to input sources only added once, *e.g.*, rDNS data. A fraction of these addresses are not responsive anymore after several scans. However, without updates of the input, newly responsive addresses are not discovered.

Figure 2 shows the distribution of responsive addresses across ASes. Compared to the overall input, the distribution is flatter. The top AS, Linode (AS63949), covers only 7.9 %, followed by China Telecom (AS4812). 50 % of responsive addresses are covered by 14 ASes. A distribution of addresses responsive to individual protocols can be seen in Figure 9 in Appendix B.

Whereas the set of responsive addresses changes over time, the cardinality of the set remains relatively stable. Since July 2018, 46.8 M non-aliased addresses have been responsive at least once and to one protocol. This is primarily due to ICMP with 45.3 M (96.8 %) addresses followed by TCP/80 with 8.6 M (18.4 %). This shows that frequent change is visible, and an up-to-date service is necessary to provide a high-quality service for further research.
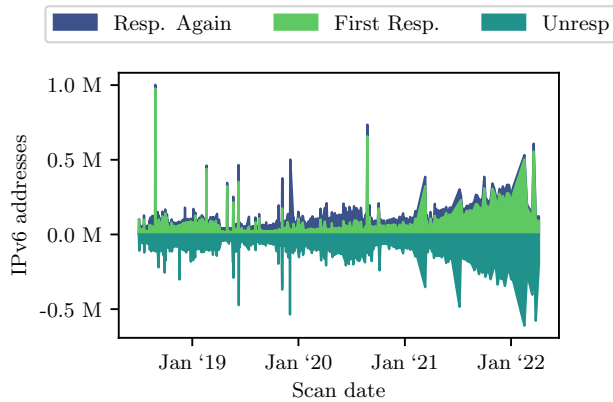
Figure 4: Development of responsive addresses over time. Later scans have a larger runtime covering up to 7 days, increasing the seen churn in between scans.

Figure 4 indicates the stability of the responsive address set over time. It shows how many addresses are newly responsive or unresponsive for each scan compared to the previous scan. In case of new IPv6 addresses, the figure differentiates between addresses that are completely new or were already responsive in a previous scan but not the last. Generally, a frequent churn of 200 k to 500 k IPv6 addresses can be seen between two consecutive scans within one to five days. However, unresponsive addresses are frequently recurring afterward. Nevertheless, completely new IPv6 addresses are regularly visible. As explained in Section 3 the *IPv6 Hitlist* service is not executed daily anymore but takes multiple days resulting in the increased churn in between scans towards the end of the analyzed period as seen in Figure 4.

More significant changes are either due to issues with the aliased prefix detection, if a drop is directly visible after a large increase or due to missing sources for a scan if the drop is before the rise. Increases without a close drop are primarily due to added sources, *e.g.*, an addition of IPv6 addresses from rDNS scans as used by Fiebig et al. [19].

4.3 Takeaways and Suggestions

The analysis of the current hitlist state shows that a regularly updated, ongoing service is required to provide an up-to-date view and adapt to regular changes within the Internet. While 176.6 k addresses were responsive throughout the entire period (5.4 % out of 3.2 M on April 7th, 2022), regular churn is visible across all protocols. On the other side, the ongoing accumulation of addresses leads to an overall input list containing unresponsive addresses. Especially regular traceroutes identify a large set of addresses derived from a fraction of EUI-64 IIDs but located within regularly changing prefixes (see Section 4.1).

Furthermore, ongoing services require regular monitoring to understand the impact of network changes on the scanning methodology, *e.g.*, the impact of the GFW that slowly ramped up over time and changed behavior frequently. The GFW injections additionally led to increased problems in the comparability of results. Chinese

vantage points are most likely affected by the GFW injection as well but on the complete opposite set of addresses, namely targets outside Chinese networks.

We suggest frequent monitoring of the hitlist in the future by the operators but also the community. Furthermore, we plan to frequently clean the overall input of specific addresses, such as outdated EUI-64 based addresses to better support scans of other protocols.

5 ALIASED PREFIX ANALYSIS

Besides the difficulty of IPv6 scans due to the inherent problem of the large address space itself, challenges occur due to the fact that individual addresses do not necessarily identify individual targets. With IPv6, servers can be reached using multiple IP addresses or even complete prefixes. They often appear as fully used prefixes with each address responsive, *e.g.*, to ICMP or TCP handshakes. The most commonly assumed reason is aliasing, where a *single* target is reachable using a complete prefix. The *IPv6 Hitlist* tries to identify these prefixes as described in Section 3. Regarding scans, this is mainly a problem with IPv6. However, similar occurrences have also been mentioned in combination with IPv4 address scans, *e.g.*, by Izhikevich et al. [28] or Alt et al. [3].

The implemented mechanism by the *IPv6 Hitlist* tests for responsiveness. However, identified prefixes do not have to be used by a single endpoint, *i.e.*, as an alias. Other reasons can be for example Content Delivery Networks (CDNs), using complete IP prefixes for multiple servers or middle-boxes and proxies preemptively terminating connection attempts. Especially in the case of CDNs, (a subset of) these fully responsive prefixes might be a valuable input for advanced scans, *e.g.*, to analyze protocol deployments such as QUIC or TLS1.3. In the following, we investigate these address regions in more detail to allow a better understanding of the *IPv6 Hitlist* and a better usage of its outcome.

In 2018, the *IPv6 Hitlist* service identified 12 k aliased prefixes of different sizes. The number of aliased prefixes increased steadily throughout the year, reaching 42.8 k in January 2022 and a sudden increase to 111.5 k prefixes afterward. The latter growth by 66.4 k prefixes is due to a single AS, namely Trafficforce (AS212144), a Lithuanian network starting to announce a larger number of prefixes in February 2022, solely limited to IPv6. All aliased prefixes were /64, responding to ICMP but not TCP/80. This sudden increase by a single operator stands out. The classification was based on successful ICMP probes and is reproducible. We contacted the technical support of the AS regarding this steep increase but did not receive further information.

We analyze whether the increasing number of addresses in the overall hitlist (cf. Section 4) results in the remaining increase from 12 k to 42.8 k labeled prefixes. As explained in Section 3, the service only checks for aliased prefixes of size /68 and longer, if more than 100 addresses from this prefix are part of the *IPv6 Hitlist* service input. Thus, aliased prefixes might remain unrecognized if too few addresses were in the input. Therefore, we check the size of aliased prefixes shown in Figure 5. It shows the distribution of aliased prefixes across prefix sizes for a single snapshot each year. It has been similar throughout the years with a small percentage of prefixes within /28 and /60. The shortest aliased prefixes are several /28s announced by EpicUp (AS397165) a US-based cloud provider.
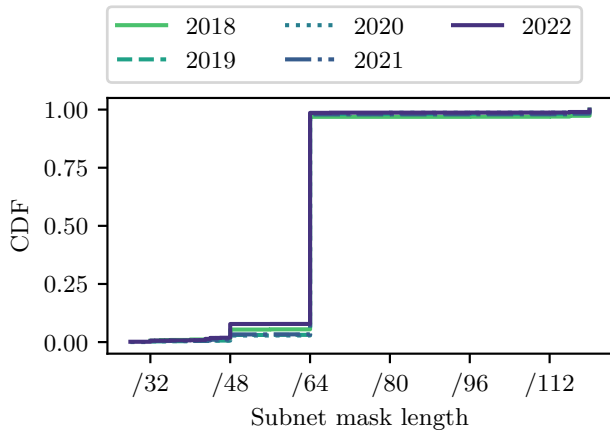
Figure 5: Distribution of aliased prefix sizes over time. The plot for 2022 excludes Trafficforce (AS212144) accounting for 60 k (61.6 %) aliased prefixes mostly /64. For each year, a snapshot representing a single scan is used. More than 90 % of aliased prefixes had a length of /64.

However, most aliased prefixes constantly have a length of /64 where the hitlist does not require a threshold. For this prefix length, only a single address is required to trigger the detection mechanism. Newly identified prefixes throughout the years were either not aliased in 2018 or did not contain a single address. Therefore, testing /64 prefixes if a single address is known to the *IPv6 Hitlist* proves to be effective to detect a majority of aliased prefixes.

5.1 Fingerprinting Aliased Prefixes

We use fingerprinting approaches to analyze aliased prefixes identified by the *IPv6 Hitlist* service in more detail. These methodologies allowed us to evaluate whether aliased prefixes were a single responsive host, or whether some identified aliased prefixes contained multiple hosts. We combine two methods to fingerprint hosts in identified aliased prefixes, namely fingerprinting based on TCP response features as conducted by Gasser et al. [22] and the Too Big Trick (TBT) as presented by Beverly et al. [10].

TCP Fingerprints. We use TCP fingerprints extracted from the scans for the aliased prefix detection [22]. It relies on different features from the TCP handshake, namely the *Optionstext*, an order preserving string representation of TCP options, the TCP window size and window scale option, the Maximum Segment Size (MSS) and the iTTL. The iTTL rounds the TTL to the next power of 2 [8, 34] and thus represents the likely selected initial TTL. Therefore, it reduces inconsistencies based on routing changes or middleboxes. We omit the timestamp analysis as Linux machines using kernel 4.10 or newer randomize the value and are expected to show increased deployment since 2018. In contrast to the TBT, the same values between two IPv6 addresses do not necessarily indicate the same host, but similarly, varying values suggest different hosts.

TCP fingerprints can be derived for 33.5 k aliased prefixes. The remaining prefixes were detected based on ICMP scans. For 33.3 k (99.5 %) of these prefixes, all values match while for 160, differences

are visible. The iTTL, MSS, window scale option and *Optionstext* only differ for addresses within up to 13 prefixes. In contrast, the TCP Window Size is different for addresses in 154 aliased prefixes. However, the Window size can change on a single host within different connections and different values are not necessarily a discriminating factor. In comparison to the evaluation by Gasser et al. [22] we see similar results. Most fingerprintable prefixes respond with a uniform behavior, while a small fraction shows variable behavior. This supports the initial assumption that detected fully responsive prefixes are often aliases for the same host. However, some fingerprints indicate that some detected address blocks are potentially used differently and by multiple hosts.

Too Big Trick. We used the TBT as an additional indication of prefix usage. It relies on the characteristic of IPv6 that only end hosts are allowed to fragment packets and relies on a shared Path Maximum Transmission Unit (PMTU) between aliased addresses for the same host. In general, if a router receives an IPv6 packet that is too big, it has to notify the sender using an ICMPv6 Packet Too Big Messages. Afterward, the sender should update its PMTU cache and fragment the respective packet. While this process was initially utilized to identify alias addresses of routers, Song et al. [47] proposed it to analyze aliased prefixes. We shortly describe the required steps in the following:

(i) The Too Big Trick verifies that a set of addresses (8) under test within a prefix replies to ICMP Echo Requests of size 1300 B, slightly larger than the minimum required MTU for IPv6 of 1280 B, without fragmentation.

(ii) It sends ICMPv6 Packet Too Big Messages itself to *one* of the addresses and verifies that the next round of ICMP Echo Request messages is in fact fragmented.

(iii) It sends ICMP Echo Requests to the remaining addresses under test *without* the preceding error message. In case all addresses are aliases for the same device and interface, they share the same PMTU cache and should now fragment the response.

The methodology only provides insights if targets respond to the initial ICMP Echo Requests without fragmentation. We used this methodology on the 111 k prefixes identified by the *IPv6 Hitlist* service on April 7, 2022 and received successful results for 29.4 k. Out of these, for 27.6 k (93.75 %) all eight responses were fragmented after the initial error message to a single address indicating a shared PMTU cache. Only for 249 (0.85 %) prefixes, no request resulted in a fragmented response but each individual address required a new error message.

Interestingly within the remaining prefixes (1592, 5.4 %), between two and seven addresses share a PMTU cache but not all. This effect is mostly seen with Akamai and Cloudflare with 1 k and 268 prefixes respectively. This supports our assumption that a fraction of identified *aliased* prefixes is not completely in line with the initial definition of a single host with the complete prefix as an alias. Nevertheless, addresses from these fully responsive prefixes are still not assigned to single hosts used within a load balancing setup, *e.g.*, of CDNs.

5.2 Characteristics of Aliased Prefixes

Besides these technical fingerprints, we analyze additional characteristics of aliased prefixes in addition to the originating ASes. We
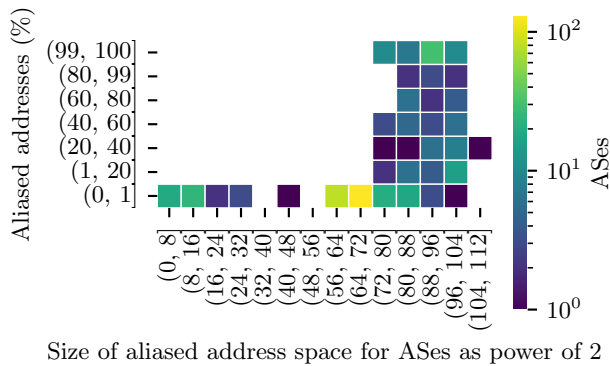
Figure 6: Number of aliased addresses (power of two) within ASes in relation to their overall number of announced addresses. If multiple prefixes within an AS are aliased, the number of addresses are summed. For some ASes, more than 99 % of the announced IPv6 addresses were covered by aliased prefixes. Results are based on the aliased prefix detection of April 7th, 2022.

argue that these characteristics allow for a more informed evaluation and usage of the *IPv6 Hitlist* results in the future.

Are Individual Prefixes Fully Responsive or Complete ASes? To analyze whether aliased prefixes are more likely due to individual network entities or set up on an AS-level, we analyze the fraction of aliased addresses within each AS in respect to the total number of announced IPv6 addresses by an AS. Figure 6 shows this relation for all ASes with at least one aliased prefix. The total number of addresses from aliased prefixes is given as power of two on the x-axis ranging from 2^8 to 2^{112} addresses. This value does not necessarily represent single prefixes but sums all aliased prefixes within each AS. The highest number of IPv6 addresses from aliased prefixes is again due to EpicUp (AS397165) announcing 61 fully responsive /28 prefixes. The y-axis depicts the fraction in respect to all announced IPv6 addresses by the respective AS. The axes are binned for better visibility.

While the fraction for many ASes is less than 1‰, for 80 ASes more than 50 % of announced addresses are located in aliased prefixes, and for 61 ASes even more than 90 % are reached. The most prominent candidates in the latter category are Fastly (AS54113) with 95.3 %, but also AS33905 owned by Akamai and AS209242 owned by Cloudflare both aliased to 100 %. We argue that in these cases even without exact fingerprinting, it is highly unlikely that all addresses are an alias of a single host because these CDNs serve numerous websites and clients. The complete exclusion of all addresses might almost exclude complete ASes.

Are Domains Hosted in Aliased Prefixes? To further evaluate the effect of an exclusion of all aliased prefixes, we analyze whether domains are hosted within these networks. In the case of higher layer protocol evaluation, including TLS, QUIC or HTTP but also the analysis of Internet consolidation, these targets can be highly relevant and should not be excluded completely. We use a single

Table 2: Responsiveness of aliased prefixes. For each prefix one random address is tested to reduce impact.

Protocol	# Prefixes	# ASes
ICMP	39.0 k	270
TCP/443	31.9 k	155
TCP/80	32.3 k	179
UDP/443	28.8 k	41
UDP/53	172	32

snapshot of our DNS scans introduced in Section 3 from April 7th, 2022 to analyze whether identified, aliased prefixes host domains and to which extent.

Based on our data 15.0 M domains resolved to 5.2 k aliased prefixes in total. 133 different ASes announced these prefixes. Note that this is a lower bound, given that we resolve a subset of the DNS namespace and that load balancing might impact our scans. Nevertheless, it shows that a fraction of hosting infrastructure can be missed by research based on the *IPv6 Hitlist*. The most prominent AS with aliased prefixes hosting domains was Cloudflare (AS13335) with 115 prefixes each hosting a mean of 167.0 k domains. The highest number of domains even reached 3.94 M domains within a single fully responsive /48. CDNs such as Fastly (AS54113), Amazon (AS16509), and Google (AS15169) were affected.

We investigated how many of these domains were on top lists. As explained in Section 3, three top lists were resolved, all containing 1 M domains respectively on April 7th, 2022. Domains from all top lists resolved to IPv6 addresses within aliased prefixes:

- Alexa Top 1M: 177.0 k
- Majestic Top 1M: 170.2 k
- Umbrella Top 1M: 118.0 k

Considering the Alexa Top 1M, 129 and 22.6 k domains were within the Top 1k and Top 100 k respectively, including `facebook.com` and `spotify.com`. Affected domains from the Majestic Top 1M were of similar ranks. In contrast, only 53 affected domains from the Umbrella top list were within the Top 1k.

Including addresses out of aliased prefixes contained in AAAA records would include 2.7 M distinct addresses, a small fraction of the respective aliased prefixes. We argue that all of these or at least a subset of IPv6 addresses should be considered by researchers in the analysis of protocols on top of IPv6 even though they are identified as aliased prefixes.

Are other Protocols Responsive? Based on the previous findings, we tested aliased prefixes for their responsiveness to all protocols. We excluded 66.4 k aliased prefixes announced by Trafficforce, to reduce the impact as they were only responsive to ICMP probes during the initial multi-level aliased detection. We only selected a single, random address from each prefix to reduce traffic on real aliases and due to our assumption that all addresses behave the same. Scan results for 42.8 k aliased prefixes can be seen in Table 2. Most probed addresses were responsive to ICMP or TCP/80, and already tested during the multi-level aliased prefix detection. Additionally, TCP/443 and especially UDP/443 were supported by a majority of tested addresses as well. As shown in Table 1, UDP/443 accounted

Table 3: New input sources for IPv6 address candidates evaluated in this work. The covered ASes are set in relation (%) to the total number of ASes (29 k) announcing IPv6 prefixes based on RIPE RIS [36].

			ASes	
Source	Information	Addresses	Total	%
Passive sources	Extracted addresses from e.g., MX/NS records, CAIDA Ark [11], DET [47]	356.7 k	3.6 k	12.5
Unresponsive addresses	All addresses excluded from scans due to the 30 day unresponsive filter	638.6 M	18.5 k	64.9
6Graph [49]	Applied on responsive addresses in December 2021	125.8 M	18.9 k	65.2
6Tree [29]	Applied on responsive addresses in December 2021	37.6 M	15.0 k	51.7
6GAN [14]	Applied on responsive addresses in December 2021	3.3 M	249	0.8
6VecLM [15]	Applied on responsive addresses in December 2021	70.3 k	278	0.9
Distance clustering	Extending clustered address regions	5.3 M	7.2 k	25.0

for the lowest number of responsive addresses with 98.1 k in the *IPv6 Hitlist*. Therefore, using a single address from each aliased prefix increases the amount of responsive addresses by 29.4 %. This is in line with our observations that aliased prefixes are frequently seen in combination with CDNs and the findings by Zirngibl et al. [51] that large providers mainly drive the deployment of QUIC. In contrast, only 172 addresses were responsive to UDP/53 probes, *e.g.*, from Cloudflare or Misaka (AS50069), an anycast DNS service.

In this scan, aliased prefix were responsive to at most four protocols. Only Cloudflare originates at least one prefix responsive to each probe respectively. In no prefix was UDP/443 and UDP/53 seen in combination.

5.3 Takeaways and Suggestions

The *aliased prefix* detection of the *IPv6 Hitlist* service is an important feature, necessary to allow an ongoing, feasible service and prevent biases in the set of responsive addresses, *e.g.*, towards Amazon with more than 200 M addresses. However, the initial definition as an alias for a single host does not necessarily hold and the number of aliased prefixes is increasing over time. Their complete removal can result in the exclusion of complete ASes, *e.g.*, Fastly, or targets hosting multiple millions of domains, *e.g.*, within Cloudflare. An informed assessment of these fully responsive prefixes depending on the usage of hitlist results is essential and can drastically change results and improve insights.

We suggest extending information regarding these address regions, e.g., statistics about hosted domains, to allow an informed selection of address candidates out of these. Furthermore, at least a single address out of each aliased prefix can be added to the hitlist. Even if the complete prefix is an alias for a single host, testing its responsiveness with one of these addresses can cover its behavior and offer a valuable foundation for future research. As shown in Table 2, random addresses are often responsive to different protocol scans. However, known addresses contained in the input from DNS scans or passive sources should be used if possible because these addresses are either actively announced by operators in DNS or known to be used by network devices (e.g., responsive to RIPE Atlas traceroutes).

6 DISCOVERING NEW ADDRESSES

Besides the analysis of the historical state of the *IPv6 Hitlist* and identified aliased prefixes, we analyze new potential address sources

to further improve the service. As shown in Figure 1, we use the *IPv6 Hitlist* service pipeline to filter and scan these addresses.

6.1 New Sources

The *IPv6 Hitlist* was initially created with a variety of sources [22]. However, other sources are available via different protocols (*e.g.*, MX and NS records) or scanning from different vantage points (*e.g.*, CAIDA Ark [11]). Furthermore, Gasser et al. [22] showed the potential of target generation algorithms as part of the *IPv6 Hitlist* input collection. The reachability of IPv6 addresses from available vantage points impacts the identification of candidates, *e.g.*, due to location specific load balancing. Additionally, active sources such as DNS scans and traceroutes are biased towards available scan targets. They can only detect IPv6 addresses mapped to a domain or responsive to traceroutes. Target generation algorithms try to mitigate these disadvantages from other sources.

We identified the following list of sources as potential candidates to improve the *IPv6 Hitlist* and consolidate its value to the community. An overview of all new sources can be seen in Table 3.

Passive Sources. Due to changes in the availability of data sources (*e.g.*, the discontinuation of the free availability of Rapid7 data), an ongoing update of the initial data collection efforts is necessary. We collected a set of new passive sources, trying to identify additional responsive addresses that can be used as an extension of the existing service. This includes results from our DNS scans as explained in Section 3. While our institution actively conducts these scans, they were not specifically set up for this work and data was available for use. Furthermore, we extract IPv6 addresses from the CAIDA Ark traceroute efforts and the published list of responsive addresses by Song et al. [47]. In total, these sources result in 3.5 M candidates.

However, 90 % of these addresses were already contained in the service, *e.g.*, from its DNS resolutions or traceroutes, and 7.5 % were additionally located in aliased prefixes. 369.1 k (71 %) of the IPv6 addresses related to NS and MX records were located within Amazon, a highly aliased prefix. Therefore, these sources result in only 356.7 k new addresses in total out of which, 84.9 k were not aliased.

Unresponsive Addresses. As described in Section 3, the *IPv6 Hitlist* service stops scanning addresses unresponsive for more than 30 days and never re-evaluates their responsiveness. This

Table 4: Responsive addresses for new sources divided by protocol. The top ASes for each source based on the number of responsive addresses indicates potential biases in each data set. ASes are abbreviated as symbols explained in the footer of the table. *IPv6 Hitlist* results are from April 7, 2022.

Source	Responsive Addresses						ASes					
	ICMP	TCP/443	TCP/80	UDP/443	UDP/53	Total	Top 1		Top 2		Total	
6Graph	3.8 M	428.4 k	491.1 k	121.9 k	78.6 k	3.8 M	52.1 %	◆	5.1 %	▼	10.7 k	
6Tree	2.2 M	374.2 k	425.5 k	116.6 k	62.8 k	2.2 M	41.0 %	◆	8.0 %	▼	11.5 k	
Unresponsive addresses	1.2 M	274.8 k	282.3 k	18.6 k	51.6 k	1.3 M	34.4 %	▲	6.2 %	▼	9.0 k	
Distance clustering	637.1 k	167.7 k	193.4 k	85.1 k	32.4 k	651.0 k	14.9 %	●	10.9 %	★	5.5 k	
Passive sources	21.0 k	1.5 k	1.9 k	358	3012	21.6 k	6.7 %	●	3.2 %	★	2.9 k	
6GAN	4.3 k	27	28	2	2	4.3 k	82.8 %	◆	12.3 %	▲	39	
6VecLM	990	103	116	38	22	1.0 k	17.1 %	◆	14.9 %	◆	105	
New Sources	5.4 M	764.9 k	843.4 k	164.0 k	144.3 k	5.6 M	26.8 %	◆	5.8 %	▲	14.6 k	
IPv6 Hitlist	3.2 M	910.8 k	1.1 M	98.1 k	140.7 k	3.2 M	7.9 %	▲	7.4 %	▼	15.7 k	
Total	8.6 M	1.7 M	1.9 M	266.2 k	287.4 k	8.8 M	25.5 %	◆	5.5 %	▲	17.3 k	

◆ Free SAS (AS12322), ▲ VNPT (AS45899), ▼ DigitalOcean (AS14061), ● China Mobile (AS9808), ★ Racktech (AS208861),
◆ CERN (AS513) ▲, ARNES (AS2107), ▼ home.pl (AS12824), ● Deutsche Glasfaser (AS60294), ★ Akamai (AS20940),
◆ Level3 (AS3356), ▲ Linode (AS63949), ▼ China Telecom (AS4812)

list accumulates 787.7 M addresses. We applied the blocklist filter and removed all candidates that showed GFW injection, resulting in 638.6 M remaining candidates. The explicit inclusion of these addresses into the scan requires deactivating the final filter of the *IPv6 Hitlist* service (see Figure 1). We re-scanned these addresses once to get insights into whether addresses are responsive again and find that these should be re-evaluated regularly as shown in the next section.

Target Generation. We applied different target generation algorithms to the set of responsive addresses of December 2021, including 6Graph [49], 6Tree [29], 6GAN [14] and 6VecLM [15]. We did not try to optimize or tune the algorithms but follow the respective explanations and standard parameters. Furthermore, we generated addresses with a simple approach named distance clustering (DC), extending more densely clustered address regions that show high entropy in the last nibble(s) of the address. Note, these regions were not fully responsive but only densely populated. Therefore, we collected clusters of addresses with at least 10 addresses and a distance of at most 64 between two addresses. Given the vast address space of IPv6, even a few addresses (10) within this comparably small distance are highly likely not assigned randomly but based on active assignment policies. We generated missing addresses within these clusters. We tested for potential new aliased prefixes after extending these clusters before our scans.

Table 3 lists how many addresses were generated by each applied method. While the respective publications often limit the number of generated addresses, *e.g.*, to 50 k, we invest more computation time to increase the number of generated candidates. However, we did not test different parameters or subsets of our input. We used published code and the contained default configuration.

In theory, 6Tree actively scans candidates during target generation and uses results to improve detection. To reduce scan impact, it contains functionality to detect aliases. However, it did not detect aliased prefixes effectively in our initial tests, quickly inducing bias

towards fully responsive regions. 6Tree generated a set of more than 8.3 M IPv6 addresses part of a single /48 prefix originated by Akamai (AS20940). Most of these addresses were incrementally assigned and responsive but not labeled as aliased by the given 6Tree implementation. However, the aliased prefix detection of the *IPv6 Hitlist* identified these addresses as aliased. Therefore, we prevented active scans, limited 6Tree to target generation only, and used the detection proposed by the *IPv6 Hitlist* service during our scans.

6.2 New Responsive Addresses

We used ZMapv6 to scan all previously introduced, non-aliased addresses for ICMP, TCP/443, TCP/80, UDP/443, and UDP/53. We scanned multiple times across four weeks to account for packet loss or network events and aggregated the results afterward. The only source we did not completely scan multiple times is the set of unresponsive addresses due to its size and thus ethical reasons. Here, we only included responsive addresses during the first test into following scans.

Even though we filtered responses injected by the GFW before generating addresses, some algorithms generated many addresses located within Chinese ASes and thus were affected by the GFW. Considering 6Graph, 18.5 M (14.5 %) out of the generated 125.8 M addresses were affected. We filtered these responses for the following analysis. Table 4 shows the cleaned number of responsive addresses for each source.

For the new passive sources, 25 % of non-aliased addresses were responsive (21 k out of 84.9 k). For 638.6 M previously excluded addresses, unresponsive for at least 30 days, more than 1.2 M addresses were responsive again. However, the responsiveness of these addresses decreased after the initial scan to nearly half the number of addresses. We plan to regularly include a subset of these addresses into the existing ongoing service. This allows testing addresses regularly but spreads the scan load over a longer period.

Regarding target generation algorithms, our naive approach to extend densely assigned address regions achieved better results than more sophisticated approaches, namely 6GAN and 6VecLM. The latter methods only resulted in 1 k and 4.4 k responsive addresses respectively. The hit rate of both algorithms was low even before filtering GFW injected responses. Our naive distance clustering approach resulted in 651 k addresses with a hit rate of nearly 12 %.

6Tree and 6Graph resulted in the highest number of responsive addresses with 2.2 M and 3.8 M, respectively. However, due to their large number of generated addresses, their discovery rate was 6 % and 3 %, respectively.

In total, we were able to identify and generate 5.4 M new or previously removed addresses responsive to at least one protocol. This results in nearly twice as many responsive addresses for UDP/53 in total and three times more addresses for UDP/443. Interestingly, slightly fewer addresses were responsive to TCP/80 and TCP/443 compared to the current *IPv6 Hitlist*. However, new sources identify 168 % more responsive addresses to ICMP probes compared to the *IPv6 Hitlist*. This is mostly due to the bias of target generation algorithms towards specific providers which we discuss next.

Overlap. Individual sources do not necessarily contribute unique responsive addresses to the *IPv6 Hitlist* but an overlap between new sources is visible. Figure 7 shows the overlap between all new sources in relation to the total number of responsive addresses for each row in percent. Thus, 89.34 % of all responsive addresses generated by 6Tree were also generated by 6Graph. All target generation algorithms identified responsive addresses which were also part of passive sources and thus visible in traceroutes or DNS data. Furthermore, all sources provided unique responsive addresses and thus show potential to improve the hitlist in the future. Even though new passive sources analyzed in this work only provide 21.6 k addresses (0.7 % of to the total of 3.2 M on April 7th, 2022), they offer new responsive candidates that are also not covered by target generation algorithms. Especially with the end of freely available data sources, *e.g.*, Rapid7 forward DNS data, a frequent evaluation of new sources will be valuable.

Distribution Across ASes. To analyze whether new addresses impose a new bias to the *IPv6 Hitlist*, we analyze their AS distribution. Figure 8 shows the distribution of responsive addresses across ASes for each new input while the total number and most common ASes are listed in Table 4. 6Graph and 6Tree, contributing most new addresses, both show a visible bias towards Free SAS/ProXad covering up to 52 % of the respective candidates. The second most prominent AS, DigitalOcean, is only at around 5 % to 8 %. We verified the correct classification of these addresses as non-aliased and came to the same conclusion as the automatic detection. While many addresses were responsive, we can identify unresponsive addresses. The existing *IPv6 Hitlist* already contains 149.8 k responsive addresses from Free SAS on April 7th, 2022.

Other sources provide different AS distributions and top hitters, *e.g.*, many previously unresponsive addresses were from VNPT. The distance clustering approach and new passive sources show the most even distribution. The latter even covers 2.9 k ASes with only 21 k responsive addresses.
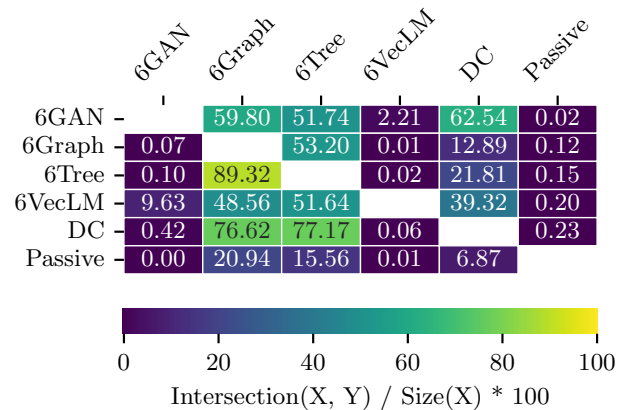

Figure 7: Overlap between responsive addresses from new sources in % in respect to the total number of responsive addresses for each row.
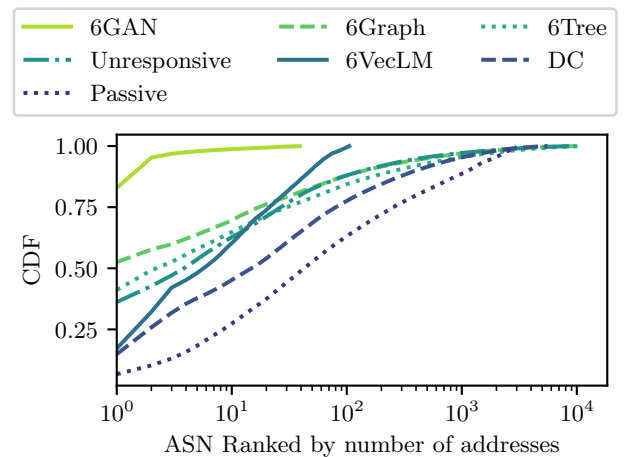

Figure 8: AS distribution of responsive addresses from new inputs.

6.3 Takeaways and Suggestions

90 % of passively extracted addresses from existing data sets were already contained in the hitlist or unique addresses which were however aliased. In contrast, address generation approaches were able to identify new, previously unknown but responsive addresses. While their individual hit rate was below 10 %, combining all responsive addresses from passive sources and address generation can more than double the number of responsive addresses currently reported by the *IPv6 Hitlist*.

Often, a major difficulty with these techniques is their incomplete documentation. Hence, we were not able to reproduce results of 6GAN [14] but it only generated 4 k responsive addresses.

Nevertheless, we suggest to regularly include addresses from some of these approaches and to include new candidates in the future.

7 DISCUSSION

In this section, we summarize key-findings of our work and discuss results.

Fully Responsive Prefixes. The *IPv6 Hitlist* tries to identify *aliased* prefixes, where a complete prefix is used by a single host. However, we show that identified prefixes are not necessarily used by a single host, but these are sometimes a result of CDN or middlebox deployments. As shown in Section 5, multiple CDNs assign large parts of their owned address space to their fleet of servers, *e.g.*, Fastly or Cloudflare. Amazon alone would introduce 200 M addresses from fully responsive prefixes. These addresses still do not represent individual hosts each, introduce significant biases, and result in a massively increased scan load. However, these prefixes are used in combination with multiple hosts for load balancing as seen with the Too Big Trick and host multiple millions of domains including highly ranked domains according to top lists but also name servers or mail exchangers. We suggest naming identified prefixes *fully responsive prefixes* in the future and to analyze their characteristics in more detail to allow for better differentiation. Fully responsive prefixes are a superset of aliased prefixes, identifiable by the implemented multi-level aliased prefix detection. However, not all identified prefixes are actual aliases on a single machine.

While we will exclude most of these targets from the ongoing scans of the *IPv6 Hitlist*, we argue that a subset of these addresses can be included and should be used by follow-up research. Especially in the case of higher layer protocol scans and analyses, such as TLS or QUIC scans, these targets should not be excluded but considered, *e.g.*, based on up-to-date DNS resolutions.

We suggest that for each fully responsive prefix, at least a single address can be tested for responsiveness by the service. Even if the prefix is an actual alias for a single host, it is an actual host, is part of the Internet ecosystem, and should thus be represented in the *IPv6 Hitlist* and considered in the future. Regarding fully responsive prefixes related to CDNs, we suggest a use case specific selection of addresses. Single addresses can be enough to identify server setups, *e.g.*, in respect to TLS where a centrally administered configuration within a CDN can be expected. In contrast, multiple addresses might be required to analyze hosted domains, certificates or websites to spread the load, *e.g.*, for Cloudflare with multiple million domains within a single fully responsive prefix.

New Input Sources. As shown in Section 6, updating the *IPv6 Hitlist* with new input sources can be valuable to improve its quality and increase the amount of identified responsive addresses. We collected and successfully applied a set of different target generation methods. While we were not able to reproduce most hit rates published with target generation mechanisms, we can identify new responsive targets, more than doubling the number of responsive addresses. However, a comprehensive and public evaluation is only possible if generated candidates and results are shared. This would further improve the impact of these methodologies in general and on the *IPv6 Hitlist*.

While target generation mechanisms sometimes bias towards certain providers, such as Free SaS, generated targets still cover a variety of ASes and show a similar distribution across protocols as the current *IPv6 Hitlist*. Nevertheless, solely relying on those generated candidates shifts the focus towards different providers and deployments and changes the view on the IPv6 ecosystem. In our opinion all sources (existing passive sources, frequent traceroutes, generated candidates) provide individual value, and their contribution to the *IPv6 Hitlist* offers a valuable foundation.

We incorporated new sources into the ongoing *IPv6 Hitlist* service, and plan to update the service with new sources on a regular basis. A new target generation algorithm announced by Song et al. [46] additionally proposes a methodology to identify potential address candidates in previously uncovered ASes. However, their results and tool were not yet accessible as of September 6th, 2022. This research direction shows the potential to increase the coverage of announced prefixes by the *IPv6 Hitlist* from currently 62 % of announced IPv6 prefixes (see Section 4).

Service Maintenance. An intrinsic motivation of this research was to improve the foundation for IPv6 research and thus improve a valuable existing building block, namely the *IPv6 Hitlist* service [22]. Ongoing measurement studies that freely publish results and data are valuable research resources. The *IPv6 Hitlist* has been operated for more than four years and has been used by a multitude of research [2, 7, 13–16, 29, 38, 42, 49, 51]. However, its operation and maintenance requires continual effort which can be difficult especially in an academic environment. While different approaches have evaluated the hitlist and proposed new address generation methodologies, they did not establish a new continuously running service. We updated the service and included newly identified addresses into the service to improve the hitlist for future use.

To better enable future collaborations to maintain the *IPv6 Hitlist* service and allow reproducible network measurement studies, we additionally share all data used for this work and used analysis scripts [52]. This includes our adaptation to 6Tree, the distance clustering implementation and a tool to filter ZMap output as published by the *IPv6 Hitlist* service from the GFW injection.

8 CONCLUSION

In this work, we took steps to dust the *IPv6 Hitlist* service after four years of ongoing operation and prepare it for the following years. Our steps include the analysis of the current state, cleaning the list of addresses caused by the Great Firewall of China, and evaluating the development of the *IPv6 Hitlist* over time. We further analyzed identified aliased prefixes and highlight that their strict omission excludes large CDNs such as Fastly or Cloudflare and thus multiple million domains from research based on the *IPv6 Hitlist*. Due to their load balancing mechanisms, announced prefixes appear to be aliased even though the backing infrastructure can be expected to be more than a single host and is valuable to analyze. Finally, we evaluate different address candidate sources, their potentially induced biases and responsiveness, which can more than double the set of responsive addresses of the *IPv6 Hitlist*. The impact of our findings regarding the GFW and new address sources is already visible in the current state of the hitlist and further steps are planned.

ACKNOWLEDGMENTS

REFERENCES

[1] Alexa. 2022. *Top 1M sites.* https://www.alexa.com/topsites
[2] Rafael Almeida, ítalo Cunha, Renata Teixeira, Darryl Veitch, and Christophe Diot. 2020. Classification of Load Balancing in the Internet. In *Proc. IEEE Int. Conference on Computer Communications (INFOCOM).*
[3] Lance Alt, Robert Beverly, and Alberto Dainotti. 2014. Uncovering Network Tarpits with Degreaser. In *Proceedings of the 30th Annual Computer Security Applications Conference* (New Orleans, Louisiana, USA).
[4] Anonymous. 2014. Towards a Comprehensive Picture of the Great Firewall's DNS Censorship. In *4th USENIX Workshop on Free and Open Communications on the Internet (FOCI 14).* USENIX Association, San Diego, CA. https://www.usenix.org/conference/foci14/workshop-program/presentation/anonymous
[5] Anonymous, Arian Akhavan Niaki, Nguyen Phong Hoang, Phillipa Gill, and Amir Houmansadr. 2020. Triplet Censors: Demystifying Great Firewall's DNS Censorship Behavior. In *10th USENIX Workshop on Free and Open Communications on the Internet (FOCI 20).* USENIX Association. https://www.usenix.org/conference/foci20/presentation/anonymous
[6] APNIC. 2022. *IPv6 Capable Rate by country (%).* Retrieved 2022-05-07 from https://stats.labs.apnic.net/ipv6
[7] 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).*
[8] Michael Backes, Thorsten Holz, Christian Rossow, Teemu Rytilahti, Milivoj Simeonovski, and Ben Stock. 2016. On the Feasibility of TTL-Based Filtering for DRDoS Mitigation. In *Research in Attacks, Intrusions, and Defenses.*
[9] Robert Beverly. 2016. Yarrp'ing the Internet: Randomized High-Speed Active Topology Discovery. In *Proc. ACM Int. Measurement Conference (IMC)* (Santa Monica, California, USA).
[10] Robert Beverly, William Brinkmeyer, Matthew Luckie, and Justin P. Rohrer. 2013. IPv6 Alias Resolution via Induced Fragmentation. In *Proc. Passive and Active Measurement (PAM).*
[11] CAIDA. 2022. *Ark IPv6 Topology Dataset.* Retrieved 2022-05-11 from https://catalog.caida.org/details/dataset/ipv6_allpref_topology
[12] Cisco. 2022. *Umbrella Top 1M List.* https://umbrella.cisco.com/blog/cisco-umbrella-1-million
[13] Tianyu Cui, Gaopeng Gou, and Gang Xiong. 2020. 6GCVAE: Gated Convolutional Variational Autoencoder for IPv6 Target Generation. In *Advances in Knowledge Discovery and Data Mining.* Springer International Publishing.
[14] Tianyu Cui, Gaopeng Gou, Gang Xiong, Chang Liu, Peipei Fu, and Zhen Li. 2021. 6GAN: IPv6 Multi-Pattern Target Generation via Generative Adversarial Nets with Reinforcement Learning. In *Proc. IEEE Int. Conference on Computer Communications (INFOCOM).*
[15] Tianyu Cui, Gang Xiong, Gaopeng Gou, Junzheng Shi, and Wei Xia. 2021. 6VecLM: Language Modeling in Vector Space for IPv6 Target Generation. In *Machine Learning and Knowledge Discovery in Databases: Applied Data Science Track.*
[16] Casey Deccio, Alden Hilton, Michael Briggs, Trevin Avery, and Robert Richardson. 2020. Behind Closed Doors: A Network Tale of Spoofing, Intrusion, and False DNS Security. In *Proc. ACM Int. Measurement Conference (IMC)* (Virtual Event, USA).
[17] David Dittrich, Erin Kenneally, et al. 2012. The Menlo Report: Ethical principles guiding information and communication technology research. *US Department of Homeland Security* (2012).
[18] Oliver Farnan, Alexander Darer, and Joss Wright. 2016. Poisoning the Well: Exploring the Great Firewall's Poisoned DNS Responses. In *Proceedings of the 2016 ACM on Workshop on Privacy in the Electronic Society* (Vienna, Austria).
[19] Tobias Fiebig, Kevin Borgolte, Shuang Hao, Christopher Kruegel, and Giovanni Vigna. 2017. Something from Nothing (There): Collecting Global IPv6 Datasets from DNS. In *Proc. Passive and Active Measurement (PAM).*
[20] Pawel Foremski, David Plonka, and Arthur Berger. 2016. Entropy/IP: Uncovering Structure in IPv6 Addresses. In *Proc. ACM Int. Measurement Conference (IMC)* (Santa Monica, California, USA).
[21] Oliver Gasser. 2022. *ZMapv6: Internet Scanner with IPv6 capabilities.* https://github.com/tumi8/zmap
[22] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczynski, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Proc. ACM Int. Measurement Conference (IMC)* (Boston, MA, USA).
[23] Oliver Gasser, Quirin Scheitle, Sebastian Gebhard, and Georg Carle. 2016. Scanning the IPv6 Internet: Towards a Comprehensive Hitlist. In *Proc. 8th Int. Workshop on Traffic Monitoring and Analysis.* Louvain-la-Neuve, Belgium.
[24] Google. 2022. *Statistics: IPv6 Adoption.* Retrieved 2022-05-07 from https://www.google.com/intl/en/ipv6/statistics.html
[25] Bingnan Hou, Zhiping Cai, Kui Wu, Jinshu Su, and Yinqiao Xiong. 2021. 6Hit: A Reinforcement Learning-based Approach to Target Generation for Internet-wide IPv6 Scanning. In *Proc. IEEE Int. Conference on Computer Communications (INFOCOM).*
[26] Christian Huitema. 2006. Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs). RFC 4380.
[27] Jana Iyengar and Martin Thomson. 2021. QUIC: A UDP-Based Multiplexed and Secure Transport. RFC 9000. https://doi.org/10.17487/RFC9000
[28] Liz Izhikevich, Renata Teixeira, and Zakir Durumeric. 2021. LZR: Identifying Unexpected Internet Services. In *Proc. USENIX Security Symposium.* https://www.usenix.org/conference/usenixsecurity21/presentation/izhikevich
[29] Zhizhu Liu, Yinqiao Xiong, Xin Liu, Wei Xie, and Peidong Zhu. 2019. 6Tree: Efficient dynamic discovery of active addresses in the IPv6 address space. *Computer Networks* (2019). https://www.sciencedirect.com/science/article/pii/S1389128618312003
[30] Matthew Luckie, Robert Beverly, William Brinkmeyer, and kc claffy. 2013. Speed-trap: Internet-Scale IPv6 Alias Resolution. In *Proc. ACM Int. Measurement Conference (IMC)* (Barcelona, Spain).
[31] Majestic. 2022. *The Majestic Million.* https://majestic.com/reports/majestic-million/
[32] Alexander Marder. 2020. APPLE: Alias Pruning by Path Length Estimation. In *Proc. Passive and Active Measurement (PAM).*
[33] Maxmind. 2022. *GeoLite2 Free Geolocation Data.* https://dev.maxmind.com/geoip/geolite2-free-geolocation-data?lang=en
[34] Ayman Mukaddam, Imad Elhajj, Ayman Kayssi, and Ali Chehab. 2014. IP Spoofing Detection Using Modified Hop Count. In *2014 IEEE 28th International Conference on Advanced Information Networking and Applications.*
[35] Austin Murdock, Frank Li, Paul Bramsen, Zakir Durumeric, and Vern Paxson. 2017. Target Generation for Internet-Wide IPv6 Scanning. In *Proc. ACM Int. Measurement Conference (IMC)* (London, United Kingdom).
[36] RIPE NCC. 2022. *Routing Information Service (RIS).* Retrieved 2022-08-29 from https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris
[37] RIPE NCC. 2022. *Total IPv6 Allocations and Assignments.* Retrieved 2022-05-07 from https://www.ripe.net/analyse/statistics/total-ipv6-allocations_and_assignments
[38] Yevheniya Nosyk, Maciej Korczyński, and Andrzej Duda. 2022. Routing Loops as Mega Amplifiers for DNS-Based DDoS Attacks. In *Proc. Passive and Active Measurement (PAM).*
[39] Ramakrishna Padmanabhan, Zhihao Li, Dave Levin, and Neil Spring. 2015. UAv6: Alias Resolution in IPv6 Using Unused Addresses. In *Proc. Passive and Active Measurement (PAM).*
[40] Craig Partridge and Mark Allman. 2016. Addressing Ethical Considerations in Network Measurement Papers. *Commun. ACM* 59, 10 (Oct. 2016).
[41] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. 2011. IP Geolocation Databases: Unreliable? *ACM SIGCOMM Computer Communication Review* (2011).
[42] Nils Rodday, Lukas Kaltenbach, Italo Cunha, Randy Bush, Ethan Katz-Bassett, Gabi Dreo Rodosek, Thomas C. Schmidt, and Matthias Wählisch. 2021. On the Deployment of Default Routes in Inter-Domain Routing *(TAURIN'21).*
[43] Erik Rye, Robert Beverly, and K C Claffy. 2021. Follow the Scent: Defeating IPv6 Prefix Rotation Privacy. In *Proc. ACM Int. Measurement Conference (IMC)* (Virtual Event).
[44] Quirin Scheitle, Oliver Gasser, Minoo Rouhi, and Georg Carle. 2017. Large-scale classification of IPv6-IPv4 siblings with variable clock skew. In *Proc. Network Traffic Measurement and Analysis Conference (TMA).*
[45] Quirin Scheitle, Oliver Gasser, Patrick Sattler, and Georg Carle. 2017. HLOC: Hints-Based Geolocation Leveraging Multiple Measurement Frameworks. In *Proc. Network Traffic Measurement and Analysis Conference (TMA).*
[46] Guanglei Song, Jiahai Yang, Lin He, Zhiliang Wang, Guo Li, Chenxin Duan, Yaozhong Liu, and Zhongxiang Sun. 2022. AddrMiner: A Comprehensive Global Active IPv6 Address Discovery System. In *2022 USENIX Annual Technical Conference (USENIX ATC 22).* Carlsbad, CA. https://www.usenix.org/conference/atc22/presentation/song
[47] Guanglei Song, Jiahai Yang, Zhiliang Wang, Lin He, Jinlei Lin, Long Pan, Chenxin Duan, and Xiaowen Quan. 2022. DET: Enabling Efficient Probing of IPv6 Active Addresses. *IEEE/ACM Transactions on Networking* (2022).
[48] Kevin Vermeulen, Burim Ljuma, Vamsi Addanki, Matthieu Gouel, Olivier Fourmaux, Timur Friedman, and Reza Rejaie. 2020. Alias Resolution Based on ICMP Rate Limiting. In *Proc. Passive and Active Measurement (PAM).*
[49] Tao Yang, Bingnan Hou, Zhiping Cai, Kui Wu, Tongqing Zhou, and Chengyu Wang. 2022. 6Graph: A graph-theoretic approach to address pattern mining for Internet-wide IPv6 scanning. *Computer Networks* (2022). https://www.sciencedirect.com/science/article/pii/S1389128621005430

[50] Fenglu Zhang, Chaoyi Lu, Baojun Liu, Haixin Duan, and Ying Liu. 2022. Measuring the Practical Effect of DNS Root Server Instances: A China-Wide Case Study. In *Proc. Passive and Active Measurement (PAM)*.

[51] Johannes Zirngibl, Philippe Buschmann, Patrick Sattler, Benedikt Jaeger, Juliane Aulbach, and Georg Carle. 2021. It's over 9000: Analyzing early QUIC Deployments with the Standardization on the Horizon. In *Proc. ACM Int. Measurement Conference (IMC)* (Virtual Event, USA).

[52] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. 2022. *Data and Analysis at TUM University Library*. https://mediatum.ub.tum.de/1686542 doi:10.14459/2022mp1686542.

A AUTONOMOUS SYSTEMS IMPACTED BY THE GFW

Table 5 shows the top 10 ASes accounting for addresses in the *IPv6 Hitlist* impacted by the Great Firewall of China. The effect and impact on the *IPv6 Hitlist* is explained in detail in Section 4. The *IPv6 Hitlist* vantage point is located in Germany and thus, probes sent into Chinese ASes, crossing the GFW are impacted. The general distribution of impacted addresses across ASes can be seen in Figure 2.

Table 5: Top 10 ASes of addresses impacted by the GFW. The total accumulated number of impacted addresses is 134 M.

ASes	# Addresses	%	CDF
4134	62.3 M	46.44	46.44
4812	19.5 M	14.59	61.03
134774	18.6 M	13.88	74.92
134773	10.7 M	8.04	82.96
140329	3.1 M	2.37	85.34
134772	2.5 M	1.93	87.28
4837	2.5 M	1.87	89.17
136200	2.3 M	1.76	90.94
140330	2.3 M	1.72	92.66
140316	1.6 M	1.24	93.91

B ADDITIONAL INFORMATION ABOUT RESPONSIVE ADDRESSES PER PROTOCOL

Figure 9 and 10 extend the results from Section 4 with further details regarding probed protocols. Figure 9 shows the AS distribution of responsive addresses from the *IPv6 Hitlist* for each protocol as extension to Figure 2. Addresses responsive to UDP/53 show the most even distribution while UDP/443 is limited to the smallest number of different ASes. Figure 10 shows the overlap between

addresses responsive to each protocol. Addresses responsive to TCP and UDP on each port are mostly also responsive to ICMP. Additionally, large overlaps can be seen between TCP/80 (HTTP), TCP/443 (HTTPS) and UDP/443 (QUIC, HTTP3).
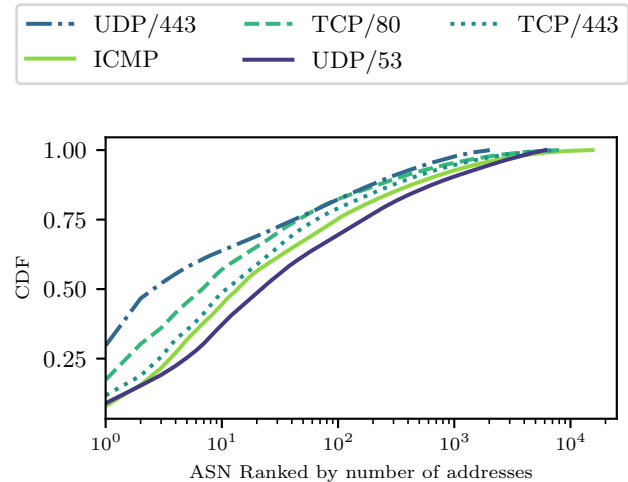

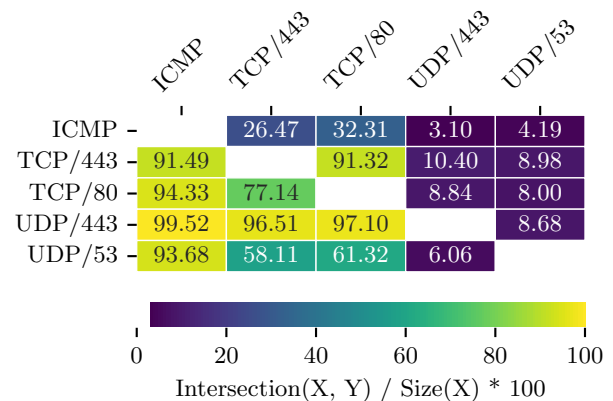**Figure 9: AS distribution of addresses responsive to each protocol on April 7th, 2022**

Figure 10: Overlap of addresses responsive to each protocol in the *IPv6 Hitlist* on April 7th, 2022