Invaluable Measurement Networks And Their Critical Infrastructure: Uptime Discrepancy Localization And Network Improvements

G. Vredeveld
University of Twente
P.O. Box 217, 7500AE Enschede
The Netherlands
g.w.r.vredeveld@student.utwente.nl

ABSTRACT

RIPE NCC is one of five Regional Internet Registries who provide essential resources to keep the core logistics of the Internet operational by assigning resources such as IPv4, IPv6 and Autonomous System Numbers (ASN). Another task RIPE NCC has taken to heart regards the global traffic, connectivity and reliability throughout the Internet. To do so properly, together with sponsors, associates and other interested parties, RIPE NCC has set up a project called Atlas which contains an enormous network of devices, called probes, to take diverse measurements to assess the state of the Internet. However, this network of probes is not perfect. Regularly probes are down which can negatively affect the consistency, connectivity and reliability of the measurement network. This might in turn lead to poorer results and statistics which will subsequently be used and studied in research. To summarize, it is important to ensure that the probes used for essential measurements concerning the state of the Internet are in fact reliably connected. This project aims to take a critical look at the uptimes of the probes in relation to several characteristics to identify, assess and possibly resolve irregularities due to bad probes throughout the RIPE Atlas network. As a result, the state of an essential measurement tool used by scholars, researchers, institutions and companies can be further improved. As there are indeed several characteristics where significant differences regarding the uptimes of probes have been discovered.

Keywords

connectivity, reliability, Regional Internet Registry, Measurement Network, RIPE NCC

1. INTRODUCTION

Since the inception of the Internet, the allocation of resources such as IPv4, IPv6 and autonomous systems numbers have been delegated to so-called regional Internet registries. These registries are important instances to keep the logistical setup of the Internet operational. One of these registries is RIPE NCC, which covers the areas of Europe, the Middle East and parts of Central Asia. Figure 1 illustrates whom cover other designated areas. Furthermore, they also initiated the RIPE Atlas project to gain better insights on the global development, connectivity and reliability of the Internet on an unprecedented scale. Initially, only specific hardware devices could be used as probes, but now software packages are also available to improve extension capabilities of the network.

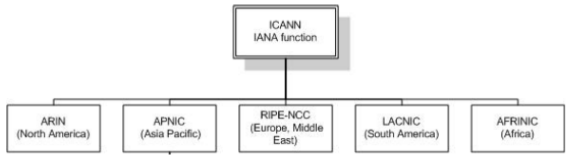

Figure 1. delegation RIRs

The resulting network has extensive network measurements capabilities. It allows for Internet measurements such as Ping, traceroute, DNS, SSL, NTP and even custom measurements. The infrastructure also helps to observe the flow, density and activity of traffic over time. For example, research on the density of traffic assesses the influence of coronavirus and the resulting effects of massive homeworking and schooling.[7] It also helps discover network interference[1]. All of these insights can help to further strengthen and harden the Internet, which closely relates to another core task they fulfill, namely Internet governance. Additional focuses of their operations are described in their 2018 activity and budget plan: "With each passing year, the Internet becomes more integral to all aspects of society and it is reasonable to expect that outside attention on the RIR system will increase. The RIPE NCC will work to proactively strengthen the RIR system by increasing its accountability, transparency and resilience."

Also, the measurements generated by Ripe Atlas are used in longitudinal studies such as [5] and [6], which would make unreliable probes detrimental to their results. Because unreliable probes might have executed less measurements, delayed or other fallible effects. Of course, this will negatively affect the validity of their research, because both scrapping the results of these probes might skew the results as well as keeping them. Thus, the best way to address this iffy issue would be to identify and possibly fix these probes.

Furthermore, research into their probes can help to improve the resilience of their network and possibly expose irregularities or other (localized) discrepancies. Research such as [8] and [2] show issues with both measurements networks and RIPE atlas in specific, but also their invaluable usage for outlining the Internet and the need to improve them for the betterment of society. For this reason,

33^{rd} *Twente Student Conference on IT* July 3^{rd}, 2020, Enschede, The Netherlands.

this paper will take a closer look at probes, their characteristics and possible discrepancies in reliability. This will substantiate previous research and provide valuable information to future endeavors regarding measurement networks.

2. PROBLEM STATEMENT

As described in the , the enormous network erected by RIPE NCC and its associates provides invaluable resources and services to assess, observe, and perform measurements on the Internet. While the main focus of the network of probes is to obtain a better understanding of the functioning of the Internet through measurements, the characteristics of the probes can potentially provide details concerning the reliability of the underlying (physical) infrastructure. RIPE NCC provides broader statistics concerning the probes, such as the owner, which can be an ISP or other organization, but also the location, times and duration of disconnection and type of probe (e.g. physical, software-based).

Furthermore, the underlying network is not flawless and information such as mentioned above, can be used to discover, assess and possibly improve the resilience of the network in case significant discrepancies have been discovered. For example, such findings might highlight profound problems concerning certain actors or the stability of the underlying infrastructure. In this research the assessment of reliability and connectivity of probes will be deduced by the uptime (connected and working) of a probe.

Therefore, the research questions will be used to assess the resilience in terms of reliability and connectivity of the probes in regards to several characteristics. The first research question is the key the main focus and its answer can help further improve the reliability of the RIPE Atlas network.

1. Can potential reliability issues of probes within the RIPE Atlas project be characterized according to certain attributes?
    - geographical location
    - ISP-based
2. Do the addition of software-based probes prove to be as reliable and connected as their hardware-based counterparts?
3. If there are significant differences according to certain attributes, could it affect the metrics taken with the Atlas network in a meaningful way?

To elaborate on the second question, software-based probes allow for slight modification of its configuration which can, in turn, negatively reflect on its uptime. Although, the availability effects might be small, they may in fact be less reliable which can have broader implications on measurements taken through the network. These effects might not be as straightforward and is therefore delegated to the next research question.

Any relevant findings on the network can be relayed back to RIPE NCC to possibly improve the reliability of the RIPE Atlas network as a whole. As a result, the projects and researches which build on top of this extensive network of probes can reap benefits in terms of a better quality measurement network as well.

3. BACKGROUND & RELATED WORK

Over the past 30 years the Internet has become such an essential and powerful tool and part of our lives.[3] Currently, the traffic alone flowing over the Internet was already 100 exabytes per month in 2017 and appears to grow exponentially. This amount is of unfathomable scale and yet excludes all the data that is generated by these large-scale measurement networks.

[3] also unfortunately states that the founders of the Internet were initially more preoccupied with improving speed, capacity and coverage. As time progressed bad actors and malicious intent has drastically grown and to counter these and scaling issues are great problems with dynamic targets that shall remain unresolved in the near future.

The RIPE Atlas project was initiated to gain better insights on the global development, connectivity and reliability of the Internet on an unprecedented scale. Argued is that the large scale continuous network measurements provided by projects such as RIPE are essential in discovering and potentially resolving fundamental flaws in the Internet, as stated by [13]. Similar projects include BISmark and the SamKnows platforms.[2]

To do this, they crowdsource devices, referred to as probes, to sponsors, ambassadors and partners to effectively scan the underlying network of the Internet. These measurements that this network allows to be taken assist in the ability to comprehensively observe the effective flow and routes information takes over the Internet. This is also used by researchers to assess problems with the network, also deliberate ones, such as traffic interference.

Additionally, an important observation made in [8] discovered that the paralleled execution of measurements by networks such as RIPE atlas can have unexpected consequences. They discovered that the vast amounts of concurrent requests can cause the measurements to be completely out of sync, even up to one full hour. This further support the presence of significant effects according to varying sets of probes.

RIPE themselves, scholars, researchers and companies alike make great use of the network. As previously stated, it is used to observe the Internet and provides real-time insights, including changes caused by events such as the corona-virus.

Furthermore, research indicated that the reliability of probes can differ according to their hardware version. For example, version 1 of the devices had problems with memory management, which was subsequently fixed by relaying this to the supplier. They also expect the required deployment of 100,000 probes to gain a complete overview of all ASNs, which includes sufficient redundancy.[14] Currently they are at around 10,000, which means they still have a long way to go.

The RIPE Atlas measurement network provided invaluable resources to geolocate IP infrastructure. The research took a different approach to do so, utilizing active measurements which turn out be feasible as a means to precisely and accurately map the entire system in a scalable manner.[4] Moreover, the RIPE network can even assist in discovering routing anomalies such as bogon filtering, Internet censorship and BGP prefix high-jacking.[17]

So the use cases of the network RIPE has been operating are incredibly diverse and provides new angles to tackle issues.

4. METHODOLOGY

The details of how each of the research questions can be answered will be elaborated here. This includes the tools to obtain the data sets as well as the cleaning process and other resources used to complement information.

4.1 data sets

This section describes the data sets which were used for the research.

The data will primarily be sourced from the RIPE NCC website and their resources. To be more specific, RIPE NCC provides an API showing information concerning the state of probes, which can be accessed through python. This tool is named Cousteau and probe details include the type of probe, geographical location, Autonomous System Number (ASN) and a range of other valuable information.

RIPE Atlas provides access to data through their rest APIs. This data includes probe attributes as referenced in table 1. This data is only scraped one-time as many attributes, such as country_code, first_connected, etc, are static or change rarely.

The total amount of probes was 31,501 with a combined size of around 25 MB. These are included in the git repository.

address_v4	address_v6
asn_v4	asn_v6
country_code	description
first_connected	geometry
is_anchor	is_public
last_connected	prefix_v4
prefix_v6	status
status_since	tags
total_uptime	type
comb_asn	name_holder
total_up	

Table 1. probe attributes

Furthermore, it was possible to obtain data concerning the disconnection of probes. The data was requested through an API call on the probe's page. This data was limited to a time frame of 152 days (and 6 hours) starting at March 12th 2019 and ending at May 3rd 2020. The specific attributes are shown in table 2.

The total amount of probes this could be obtained for was 6374 and accumulated to a size of around 35 MB which is included on the repository. See https://github.com/Gibson223/ripe-atlas/ for the used data.

Both sets were obtained as jsons and distributed likewise so others can use it and reproduce the results to also validate them.

probe	from
to	ip
controller	latest
on	

Table 2. Connection attributes

4.2 data cleaning & modifications

After obtaining the aforementioned data sets they required some cleaning and modifications, as well as additions.

First of all, the analysis regarding the data contained many explorative avenues of which the most important for the research will be elaborated on.

After importing the data, it was validated and converted to their correct types to allow proper handling. Attributes like dates regularly contained null data. For 161 rows there was no "to" attribute specified for connections as it was still online, so this was replaced by the end time. The API also returned times before and after (when connection was lost but not running at the moment). These were consequently replaced by the start and end times, respectively. As the given "total_uptime" was unreliable, the attribute "total_up" was added which aggregated the uptimes given in the "connections" data set.

To answer RQ1, geographic properties were used from the probe attributes. Attributes like "country_code" were used to group the data. Another geographic property of probes that was not directly available, concerns the organization. Information regarding the organization, such as the ISP, was not present. To add this attribute to probes their ipv 4 and IPv6 AS number were combined, as many only contained one. With help of a file containing mappings between AS numbers and name holders the correct organization was attributed. Some, however, were not present and obtained through requests to online services or referred to as "UNKNOWN". The resulting "name_holder" attribute can also be seen in table 1. This attribute was used to group and discover national and country-dedicated ISPs. These were then compared amongst their respective groups and the groups between each other. The resulting data was displayed through easier to understand grasps, which can be seen and are elaborated in the results section. Also a statistical analysis with one way ANOVA tests were conducted to validate assumed differences and potentially discover this group.

RQ2 focused on abstract, non-localized attributes to discover potential discrepancies between the uptime of probes. Therefore, a look was taken at probe tags. Only system tags were used and these were assigned by RIPE. So, no user-defined tags were used nor information in the "description" attribute. The "total_up" attribute for the probes is the culmination of time a probe was online according to the connections data set.

For all of the resulting figures and statistics no sampling was conducted and inconclusive information was disregarded. The graphs were normalized and primarily not cumulative. Also, the reliability of these probes according to the chosen parameters, will be assessed on their uptimes.

5. RESULTS

In this section some resulting figures and statistics will be described.

To start off, figure 2 shows the distribution of uptime per country. In this figure, one can see that the medians of all countries are relatively equal, all around 140. It is important to note that there a quite some outliers, which might skew the means.

Table 3 showcases the mean and standard deviation per country in regards to their uptimes. The statistics indicate that France (μ=116 days) and the Netherlands (μ=116 days) have a lower reliability than the other countries with the most probes.

A one way test on the variances (ANOVA) for the countries gives an F-value of 2.5561 and p-value of 0.0257.

F-value=2.5667, p-value=0.0364 (without France) F-value=2.2750, p-value=0.0779 (without France and Netherlands)

Figures 3 and 4 show the distribution of uptimes per coun-

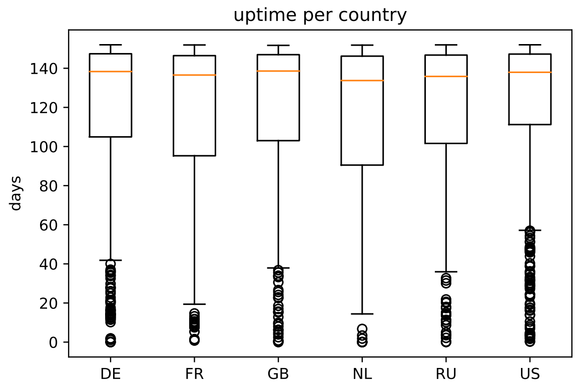

Figure 2. distribution of uptimes for countries with the most probes

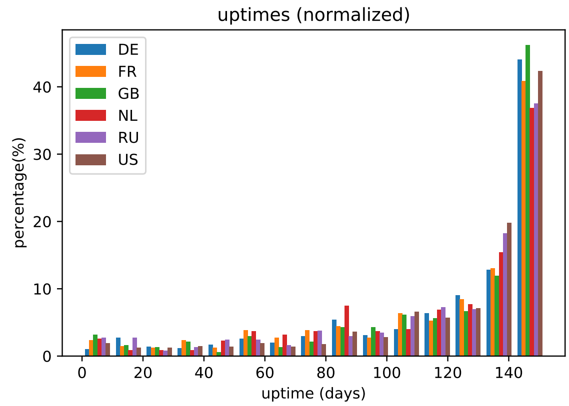

Figure 3. distribution of uptimes of several countries normalized

	mean	std
country_code		
DE	120 days	37 days
FR	116 days	39 days
GB	119 days	39 days
NL	116 days	38 days
RU	117 days	40 days
US	123 days	35 days

Table 3. mean and standard deviation countries

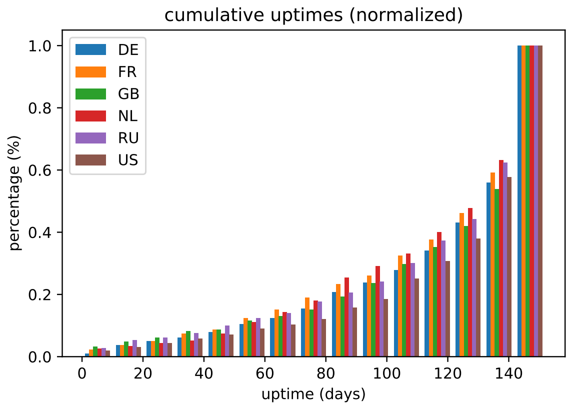

Figure 4. distribution of uptimes of several countries accumulated and normalized

try normalized and/or accumulated. Both France and the Netherlands have around 5% lower in the last bin. Besides that, they also have slightly higher peaks in the lower spectrum. The figures show the other countries, like Great Britain and Russia, seem to be consistently a bit higher.

Another attribute looked into regarding reliability of probes, concerns the ISPs. The ISPs are aggregated through AS numbers. Some ISPs are operating nationally only and other ISPs operate on a global scale.

Figure 5 showcases the distribution of uptimes for key country ISPs. It shows that operator att is peaking more in the lower region and seemingly has fewer probes with 100% reliability.France Telecom performs slightly higher in the middle segment. Statistical analysis confirms that there is a significant difference between means ($p=0.028 < 0.05$), with further research indicating that only att significantly deviate from the others.

F-value=2.7364, p-value=0.0283 (country ISPs) F-value=1.4599, p-value=0.2250 (country ISPs without att)

Figure 6 showcases the distrubution of the uptimes of several, mainly globally-operating ISPs. Comcast is a provider mainly servicing in America. Following ANOVA results determine that Comcast is the only provider significantly deviating in regards to uptimes. F-value=8.2742, p-value=0.0002 (all) F-value=2.4397, p-value=0.1192 (all, excluding comcast)

Figure 7 showcases the distribution of uptimes for the largest and country ISPs. The figure shows that there is no clear distinction between national and regional ISP's regarding uptime. an ANOVA test gives F-value=0.686686 and p-value=0.4074865, thus the comparison does not yield any conclusive information about differing uptimes.

Figure 8 showcases the distribution of uptimes compared by system version. It shows that although newer versions achieve less in the highest echelon, statistical tests indicate that all vary significantly in mean uptimes. Looking at the figure makes clear that the most reliable version concerns version1 then version2 and lastly the newest version.

The question concerning potential differences between hardware and software probes will be based on the figure 9.

In figure 9 one can see probes grouped on type. For the software probes the distribution is clearly skewed to the lower end compared to the hardware probes. There is an obvious distinction between software and hardware. The hardware probes perform significantly better in uptimes overall. Furthermore, a p-value of under 1% substantiates this assumption.

6. DISCUSSION

In this section the results will be discussed in regards to the uptimes of probes of the RIPE Atlas measurement network.

In regards to answering RQ1, figures 2, 3 and 4 have been analysed. These figures display the distribution of several key countries, answering the geographical locations aspect. Furthermore, the ISP-based differences can be seen with figures 5, 6 and 7. Additionally, some descriptive statistics were conducted to substantiate inferences over differences for countries, see table 3. A difference due to version is also looked into (see figure 8).

Firstly, the results of the ANOVA test and figures concerning several highly concentrated countries (see figures

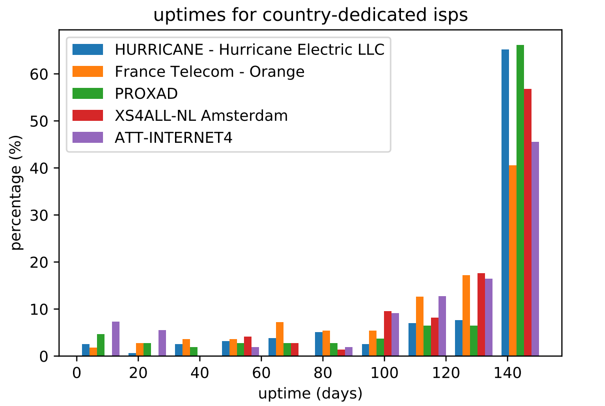

Figure 5. distribution of uptimes for key country ISPs

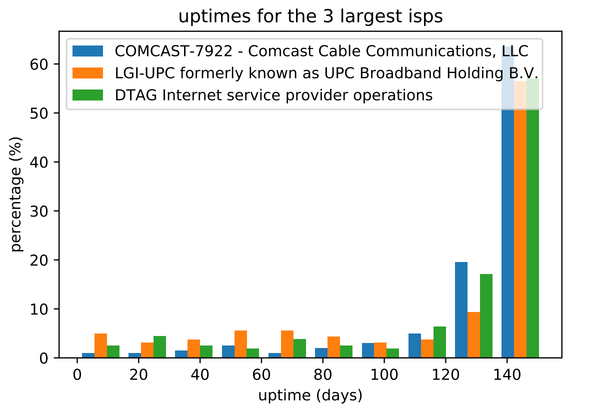

Figure 6. distribution of uptimes compared based on ISP

2, 4 and 3) illustrate significant differences amongst them. The Netherlands and France turn out to perform relatively worse regarding uptime.

These results are surprising as the Netherlands is among the best regarding broadband speed. As speed is clearly of importance there, it is logical to assume that this would be paired with a high reliability. However, the aforementioned results do not properly line up with previous results, see [11], and reliability seems to be relatively lacking in France and the Netherlands in regards to several highly developed countries.

On the basis of results of figures 5, 6 and 7 several observations can be made. For example, the operator att appears to be worse in regards to other country-dedicated ISPs. This is an american operator and that is where their probes are situated. A plausible reason for this might be the consistent neglect of american ISPs to upgrade their infrastructure, even when money is provided.[9]

Likewise, such a comparison was made in regards to the largest global ISPs and they were compared amongst each other. Surprisingly, comcast, like att, is a predominantly american operator however they perform better to a significant degree on uptimes than att. This appears to contradict what is mentioned above. A reasonable explanation involves their service areas, comcast appears to be more scattered with smaller contiguous plots. Whereas att has vast lands, likely including many rural areas. Research

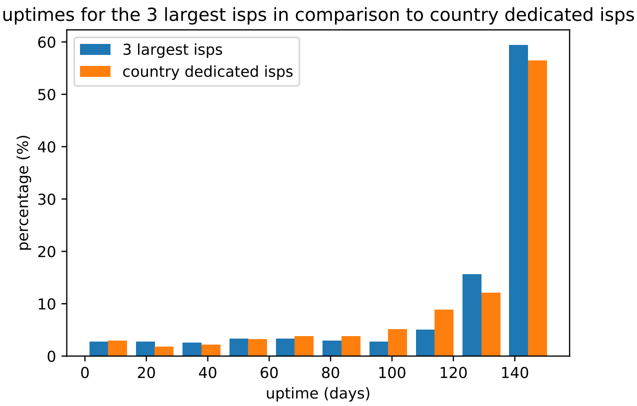

Figure 7. distribution of uptimes for the largest and country ISPs

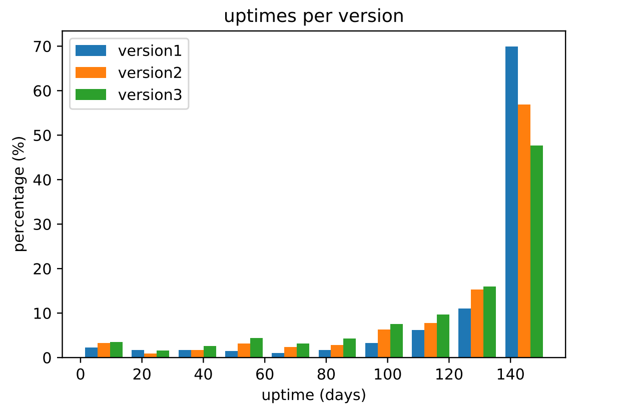

Figure 8. distribution of uptimes compared by system version

from the us government, see [15], substantiates the lack of proper Internet access for rural areas, explaining the seemingly contradictory findings.

The results regarding version numbers are overwhelming in stating a significant difference in uptimes. One widely supported reason in general that each version proves to be less reliable concerns the phenomenon of software aging.[12] This implies that every subsequent update boosts the size of the code base, more bugs and less flexibility to anticipate worthwhile changes, but also worsened documentation. Thus, each newer version will have aged causing a degrading product. Furthermore, the latest version of the Atlas probes, version3, functions as a wireless router. This is the first and currently only version to do so, also substantiating why they perform significantly worse than its predecessors. These devices simply disconnect more often and are thus less reliable.[11]

To answer RQ2, the results of figure 9 are used. The results from analyzing the type of probe shows that software probes are significantly less reliable in regards to their physical counter-parts. Reasons for this includes RIPE's stance on access to the configuration of software probes but also their policy on extending their measurement network. To further elaborate on this, probe hosts have the ability to modify the configuration of software probes which might cause throttling or other limiting factors resulting in more packet loss or power issues, which in turn make them a less reliable and useful addition to the Atlas

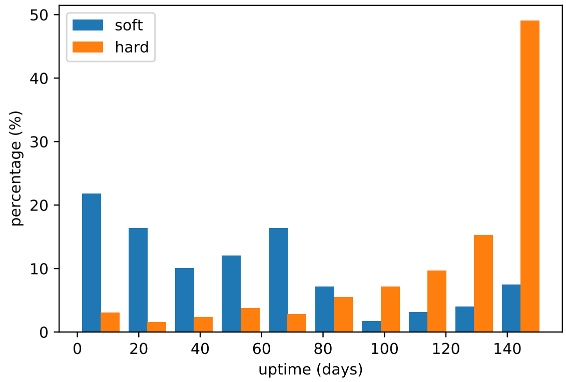

Figure 9. distribution of uptimes compared on type of probe (hardware/software)

network than their hardware-based counterparts. Additionally, one of RIPE's motivation for providing software-based probes concerned the difficulty to obtain coverage for certain areas. Thus, they thought it would be more beneficial to provide a software alternative so locals could choose an appropriate physical location where hopefully connectivity and reliability of the probe can be guaranteed. However, a caveat may be in the aforementioned assumption as well since such physical locations might simply not provide and sustain the necessary underlying network for a probe. Also, these software probes will be run on virtual machines which makes it impossible to discover how they are connected. Thus, they might be connected wirelessly, definitely impacting reliability, see [10] and [11]. This worse performance for probes of version 3 is also mentioned [2].

The impacts of the modifiability to these software probes can potentially be reflected in the metrics generated by the network as well.

Regarding RQ3, it is highly likely that the differences in attributes for probes reflect similarly on measurement metrics. However, due to time limitations and the vast amounts of metrics these probes generate and would have to be delved through, it was infeasible to adequately address the question. However, studies such as [2] acknowledges differences between first-hop latencies in regards to probe versions, with v1/v2 performing significantly worse. In [8] also shows how measurements from and to probes, leading to significant and varying latencies.

To summarize, as with all types of networks, there are bound to be irregularities and unexpected or unseen issues that might negatively effect it. An analysis on the uptime of probes through a range of characteristics such as ASN, location, type of probe, etc, has yielded correlations and shown differences in uptimes. These can be used to improve the understanding of a possible concentration and localization of bad probes in the network.

6.1 limitations

During this study, various limitations have been identified that have affected this research's results.

First of all, a preferred metric for the reliability would be as described in [16] where it would be calculated over a given time span. However, this was not possible due to some lacking and inadequate data concerning the first and last_connected attributes. Therefore a comparative approach on sorted groups of the same sample was more valuable.

As mentioned in the discussion as well and research like [2], there appears to be an inherent bias in who and where these probes are deployed. Such biases are unfortunately intrinsic to any research but should be mentioned nevertheless. The effects, however, might be minimal as the research also tried to discover and highlight such biases.

To elaborate, the first_connected as well as last_connected attributes showed to be inconsistent as many did not even contain it and thus it was difficult to presume the exact amount of time a probe would be online in a time span. Furthermore, both connections where the probe was online or offline the entire time frame resulted in an empty connection log. Therefore, statistical analyses concerned relative uptimes in groups without comparing in between.

Another attribute, "is_anchor", was not used due to the high likelihood of these probes performing excellent. These probes are situated at universities or IXP's and of great importance for the overall network. Therefore, preliminary research ruled out their significance, as expected due to their overall importance for RIPE NCC and their measurement network.

Furthermore, due to the privacy-centric approach of RIPE NCC, many probes were set to private by their owners. As a result, it was not possible to obtain metrics of around 8,000 probes. Furthermore, of the 30,000 probes about half were already abandoned. Nevertheless, 7,000 probes could be utilized for this research which showcases an extensive sample size, rendering the results as described valid.

Lastly, another tag that seemed promising was "system-wifi". This tag indicated whether a device was connected over WiFi, however, the amount of such probes was only around 100. As this sample size was small and appears to not be applied to all matching probes, strongly undermines any conclusions drawn from it.

7. CONCLUSION

To conclude, this research has provided insights in understanding, assessing and possibly improving the RIPE Atlas network in regards to several characteristics, ranging from geographical to more abstract attributes. For example, ISPs, location but also the version and type of probes.

Furthermore, analysing reliability in regards to ISPs can help RIPE further strengthen their network. There seem to be some networks with lower uptimes and reliability. Thus efforts to go for their redundancy plan regarding probes might be less beneficial than addressing core issues with probes in their current measurement network. It might be worthwhile to consider the benefits of taking a closer look at potentially strengthening the measurement network at weaker ISPs. Similar differences have been observed in regards to version which substantiates observations in previous works.

Lastly, after a formal preliminary statistical analysis of reliability and connectivity of the RIPE Atlas infrastructure in regards to certain characteristics, there are results signifying discrepancies in uptimes and thus reliability. These observations are valuable to relay back to RIPE NCC for them to possibly improve the resilience, connectivity and reliability of their measurement network. This would in turn improve reliability of consequent results based on measurements and data from the network which would be great for the scientific community overall.

8. FUTURE WORKS

Depending on the localization of the worse-off probes one could try and deduce educated assumptions on why the probes are worse off. For example, in case a set of probes are part of the same ASN it might be due to faulty BGP routing or other network related issues. This can be checked for with help of spitting through relevant measurement results if necessary. When the set only spans a small geographical location one could argue that the probes in that area are provided a worse underlying network, such as its power grid.

An alternative source for connection data from RIPE was also found. This way involved a live feed which would not provide enough longitudinal data given the time frame. Through this information source, one could verify or compare to the stated findings. Furthermore, this data might include smaller downtimes for simple missed calls. This might in fact make it more valuable as a source to answer RQ3 and finding probes with reliability issues.

When significant irregularities in connectivity and reliability are discovered, it might be worthwhile to analyze for possible effects on the metrics conducted by the relevant probes. For example, there might be clear delays in measurements such as DNS resolution times. To analyze such data another python library developed by RIPE NCC can be utilized, namely Sagan. This tool allows for the extraction of conducted measurements given their measurement id. It is possible to obtain all the id's for a specific (public) probe as well through the web interface.

9. REFERENCES

[1] Anderson, C., Winter, P., et al. Global network interference detection over the {RIPE} atlas network. In *4th {USENIX} Workshop on Free and Open Communications on the Internet ({FOCI} 14)* (2014).

[2] Bajpai, V., Eravuchira, S. J., and Schönwälder, J. Lessons learned from using the ripe atlas platform for measurement research. *SIGCOMM Comput. Commun. Rev. 45*, 3 (July 2015), 35–42.

[3] Brownlee, N., and Claffy, K. Internet measurement. *IEEE Internet Computing 8*, 5 (2004), 30–33.

[4] Candela, M., Gregori, E., Luconi, V., and Vecchio, A. Using ripe atlas for geolocating ip infrastructure. *IEEE Access 7* (2019), 48816–48829.

[5] Cicalese, D., and Rossi, D. A longitudinal study of ip anycast. *ACM SIGCOMM Computer Communication Review 48*, 1 (2018), 10–18.

[6] Cicalese, D., and Rossi, D. A longitudinal study of ip anycast. *ACM SIGCOMM Computer Communication Review 48*, 1 (2018), 10–18.

[7] Fontugne, R. Network delays in times of corona, March 2020. Accessed May 7th 2020.

[8] Holterbach, T., Pelsser, C., Bush, R., and Vanbever, L. Quantifying interference between measurements on the ripe atlas platform. In *Proceedings of the 2015 Internet Measurement Conference* (2015), pp. 437–443.

[9] Jon Brodkin Aug 16, . . p. U., JerkfaceMcGee Smack-Fu Master, i. t. j. t. p., and jump to post, T. A. S. P. Isps say they can't expand broadband unless gov't gives them more money, Aug 2018.

[10] Lehr, W. H., and Chapin, J. M. On the convergence of wired and wireless access network architectures. *Information Economics and Policy 22*, 1 (2010), 33–41.

[11] Leighton, T. Improving performance on the internet. *Communications of the ACM 52*, 2 (2009), 44–51.

[12] Parnas, D. L. Software aging. In *Proceedings of 16th International Conference on Software Engineering* (1994), IEEE, pp. 279–287.

[13] Schönwälder, J., and Bajpai, V. Large-scale network measurements what? why? how? findings? impact?, 2013.

[14] Staff, R. N. Ripe atlas: A global internet measurement network. *Internet Protocol Journal 18*, 3 (2015).

[15] Strover, S. Rural internet connectivity. *Telecommunications policy 25*, 5 (2001), 331–347.

[16] Van Slyke, R., and Frank, H. Network reliability analysis: Part i. *Networks 1*, 3 (1971), 279–290.

[17] Yakimov, T., van der Ham, J., and van Kampen, B. Detecting routing anomalies with ripe atlas, 2014.