

Characterizing User-to-User Connectivity with RIPE Atlas

Petros Gigis
RIPE NCC, Netherlands
FORTH-ICS, Greece
gkigkis@ics.forth.gr

Vasileios Kotronis
FORTH-ICS, Greece
vkotronis@ics.forth.gr

Emile Aben
RIPE NCC, Netherlands
emile.aben@ripe.net

Stephen D. Strowes
RIPE NCC, Netherlands
sds@ripe.net

Xenofontas Dimitropoulos
FORTH-ICS, Greece
fontas@ics.forth.gr

ABSTRACT

Characterizing the interconnectivity of networks at a country level is an interesting but non-trivial task. The IXP Country Jedi [8] is an existing prototype that uses RIPE Atlas probes in order to explore interconnectivity at a country level, taking into account all Autonomous Systems (AS) where RIPE Atlas probes are deployed. In this work, we build upon this basis and specifically focus on “eyeball” networks, *i.e.* the user-facing networks with the largest user populations in any given country, and explore to what extent we can provide insights on their interconnectivity. In particular, with a focused user-to-user (and/or user-to-content) version of the IXP Country Jedi we work towards meaningful statistics and comparisons between countries/economies. This is something that a general-purpose probe-to-probe version is not able to capture. We present our preliminary work on the estimation of RIPE Atlas coverage in eyeball networks, as well as an approach to measure and visualize user interconnectivity with our Eyeball Jedi tool.

CCS CONCEPTS

•Networks → Network measurement; Network properties;

KEYWORDS

Measurement, Traceroute, User Coverage, Eyeball Networks

1 INTRODUCTION

Eyeball networks, *i.e.* the networks that provide Internet access to end-users at the “last mile”, are interesting for multiple reasons. As opposed to content which can be moved around and hosted anywhere in the network [17], end-users usually access the Internet from a limited physical area, typically the area where they reside or work. Therefore, routing IP packets from end-user to end-user is really determined by how the networks serving the users are connected with each other. While the majority of traffic today is eyeball-to-content (*e.g.* users accessing content “giants” [16] such as Google, Facebook, Akamai, Cloudflare, Netflix), and these are traffic flows that are optimized for latency [9], user-to-user traffic is important for real-time communications (*e.g.* VoIP or online gaming). For reasons of efficiency—and potentially also security—local (*e.g.* country-level) user-to-user traffic should *stay local*: paths between users ought thus to traverse the fewest intermediaries possible, especially when the end-users are closely geo-located.

The RIPE Atlas measurement platform [5] has probes located inside multiple user networks; these probes can perform active Internet measurements such as traceroutes. In this work we explore

what RIPE Atlas probes can tell us about the interconnectivity of networks which serve the majority of users in a given country. The following interesting questions arise in this context:

(Q1) How many of the eyeball networks within a country contain RIPE Atlas probes, and are thus measurable?

(Q2) Do the “traffic locality” and “direct connectivity” properties actually hold for the majority of users at a country level?

Answering these questions presents a number of challenges:

(C1) What is the most suitable approach to visualize the associated statistics and measurements?

(C2) What is a suitable probe selection strategy in order to infer common connectivity between eyeball networks within a country, while also minimizing measurement costs?

(C3) How can we verify and amortize the influence of IP-to-AS mapping and IP-level geo-location artifacts?

In this short paper we answer questions (Q1) (see Section 2) and (Q2) (see Section 3), focusing on addressing challenge (C1). Approaches to address challenges (C2) and (C3) are discussed as future work in Section 4. Our intent is to understand and characterize aspects of user-to-user connectivity, even in light of limited probe coverage within their eyeball ISPs. We believe that such a characterization can help both network operators and Internet users discover interesting interconnectivity artifacts and issues [15, 18] within the countries they operate and live, and act upon them. We build a new tool to achieve this, the *Eyeball Jedi* (see Section 3).

2 ESTIMATING USER COVERAGE

As a prerequisite to understanding interconnectivity, we first need to be able to measure eyeball networks. Therefore, we need to have the necessary infrastructure deployed within the networks and understand how many users are covered by such deployments [14]. We thus measure and visualize the population coverage by RIPE Atlas probes per country *systematically in time*, taking into account the deployment of IPv6/IPv4 probes in the most populated eyeball networks. We use the following data sources: (i) RIPE Atlas probe data (including their location information), which are fetched using the RIPE Atlas API [4], (ii) Internet users per country data which are based on Internet Live Stats [3], and (iii) user population per AS data, based on APNIC estimates [1]. The latter are derived from the APNIC IPv6 measurement campaign [2]. On a daily basis we fetch the biggest eyeball networks from APNIC data for all available countries. Per country we use the inferred eyeball populations to estimate which networks (AS) are the dominant players up to a cumulative fraction of 95% of the Internet users in that country. As

a population coverage threshold *per AS*, we consider a value of 1% since this allows us to study the country-level eyeball ecosystems at a fine granularity. This method typically represents a majority of Internet users, on average covering $\sim 90.5\%$ of end-users per country, though there are outliers such as Russia with only 29.3% coverage due to a highly fragmented eyeball ecosystem. As a next step we use the readily available data from IXP Country Jedi [8] to determine if there were any probes in these networks participating in mesh-measurements during the latest runs of the tool. Eyeball networks with at least one public probe are considered as covered. Using this methodology we generate the number of Internet users being covered per AS, based on the population percentage of the AS (APNIC data) and the number of Internet users of the AS’s country (Internet Live Stats). These data are made publicly available for all countries in both tabular formats [12] and on color-coded world maps [13]. We note that the challenge of the RIPE Atlas deployment limitations and missing networks was first presented at [7].

3 THE EYEBALL JEDI

The Eyeball Jedi launches and processes traceroutes monthly, from eyeball to eyeball network, on a country level, for all available countries, after selecting the required probes per network (AS). These networks were identified in Section 2. Eyeball-to-eyeball AS-paths are extracted from IP-level traceroutes using RIPEstat [6] IP-to-AS mapping. While for this analysis we consider IPv4 datasets, IPv6 is also supported. Currently, we use the following probe selection methodology per AS per country: we select the closest and the most distant probe using the capital of a country as a point of reference, similarly to the IXP Country Jedi [8]. This strategy aims to exploit geo-related path diversity effects between eyeballs, using a minimal number of probes. The location information of the probes is provided by RIPE Atlas. Furthermore, to geo-locate the traceroute hops involved in the Eyeball Jedi measurements, we use OpenIPMap [19]. This is important to infer whether the inter-eyeball traffic remains within or exits a country. Using this approach, and assuming that the probes we select actually capture the diversity of the eyeball interconnection, and accurately represent the local market, we can estimate the percentage of user-to-user connections that stay in (or go out of) the country, as well as whether they are direct or not. It is future work to assess if this assumption holds in practice.

To visualize our results, we generate a tabular structure called the AS-to-AS matrix. Rows and columns correspond to different eyeball networks (AS), used as sources and destinations respectively. The resulting boxes are sized according to the APNIC estimations of the coverage of Internet users per AS. The colors of the boxes correspond to different types of interconnectivity information, such as out-of-country or in-country, while red borders mark indirect AS-level connections. With such a structure, we can calculate metrics related to the user population that interconnects via (direct or not) paths within or outside a country. The basic metric we use is reflected by the area of the displayed boxes, which corresponds to a product of coverage percentages. By dividing such areas with the total area, we can calculate percentages of user-to-user connections with certain characteristics. An example for Canada is depicted in Fig. 1. First, we found that 16 AS cover the 84.5% of Internet users in Canada. Second, the cumulative area of user connections, seemingly




Figure 1: Snapshot of eyeball-to-eyeball matrix for Canada (generated on 2017-04-01). Colors are mapped as follows: In-country paths are green, out-of-country orange. Eyeball networks without RIPE Atlas probe coverage are light grey; in-/out-of-country inconsistencies between probes are black. Red borders mark indirect AS-level eyeball connections; blue borders mark direct/indirect inconsistencies.

served via in-country paths, is 47.1%. Only 9% are indirect (with ≥ 1 intermediaries) on the AS-level. 3.1% leave the country. Moreover, 18.1% suffer from lack of RIPE Atlas probe coverage, and only 3.2% exhibit inconsistencies w.r.t. achieving consensus on whether traffic actually leaves the country or remains within it. 28.6% is the rest of the area not examined by the Eyeball Jedi. We note that the asymmetries displayed in the matrix may stem either from inference errors or from interesting ISP policy differentiation per traffic direction, something we plan to investigate further.

4 CONCLUSIONS & FUTURE WORK

We presented a new tool, the Eyeball Jedi, that can be used to characterize aspects of the interconnectivity between eyeball networks on a country level. The tool uses probe-to-probe traceroutes from RIPE Atlas and estimated population coverage data to yield useful user-to-user statistics and visualizations. As future work, we intend to turn our initial implementation available in [10–13] into a fully-fledged publicly available tool. In addition, we aim at gaining visibility in the eyeball-to-neighbor-country-eyeball traffic. W.r.t. challenge (C2), we plan to investigate more sophisticated probe grouping and selection strategies. Addressing challenge (C3) will require attaining proper ground truth for validation of the inferred characterization (*e.g.* via crowd-sourced verification).

5 ACKNOWLEDGEMENTS

This work has been partly funded by the European Research Council Grant Agreement no. 338402.

REFERENCES

- [1] APNIC Labs AS population data. <https://stats.labs.apnic.net/aspop/>.
- [2] APNIC Labs IPv6 measurement campaign. <https://labs.apnic.net/measureipv6/>.
- [3] Internet Users from Internet Live Stats. <http://www.internetlivestats.com/internet-users/>.
- [4] RIPE Atlas API Reference. <https://atlas.ripe.net/docs/api/v2/reference/>.
- [5] RIPE ATLAS Home. <https://atlas.ripe.net/>.
- [6] RIPEstat. <https://stat.ripe.net/>.
- [7] ABEN, E. Improving RIPE Atlas Coverage - Which Networks Are Missing? <https://labs.ripe.net/Members/emileaben/improving-ripe-atlas-coverage-what-networks-are-missing>.
- [8] ABEN, E. IXP Country Jedi. <https://github.com/emileaben/ixp-country-jedi>.
- [9] CHIU, Y.-C., SCHLINKER, B., RADHAKRISHNAN, A. B., KATZ-BASSETT, E., AND GOVINDAN, R. Are We One Hop Away from a Better Internet? In *Proceedings of the 2015 ACM Conference on Internet Measurement Conference (2015)*, ACM, pp. 523–529.
- [10] GIGIS, P. Eyeball Jedi github repository. <https://github.com/pgigis/eyeball-jedi>.
- [11] GIGIS, P., ABEN, E., AND KOTRONIS, V. Eyeball Jedi. <http://www.inspire.edu.gr/eyeball-jedi/>. Also accessible via: <http://sg-pub.ripe.net/petros/eyeball-jedi/>.
- [12] GIGIS, P., ABEN, E., AND KOTRONIS, V. Eyeball networks coverage: An estimation of Internet users being covered by RIPE Atlas probes (Online data in tabular format). <http://sg-pub.ripe.net/petros/population.coverage/table.html>.
- [13] GIGIS, P., ABEN, E., AND KOTRONIS, V. Eyeball networks coverage: An estimation of Internet users being covered by RIPE Atlas probes (Online data on world map). <http://sg-pub.ripe.net/petros/population.coverage/map.html?version=4&public.only=false>.
- [14] GIGIS, P., ABEN, E., AND KOTRONIS, V. RIPE Atlas Coverage in Eyeball Networks. <https://labs.ripe.net/Members/petros.gigis/ripe-atlas-coverage-in-eyeball-networks>.
- [15] GUPTA, A., CALDER, M., FEAMSTER, N., CHETTY, M., CALANDRO, E., AND KATZ-BASSETT, E. Peering at the Internet’s Frontier: A First Look at ISP Interconnectivity in Africa. In *International Conference on Passive and Active Network Measurement (2014)*, Springer, pp. 204–213.
- [16] LABOVITZ, C., IEKEL-JOHNSON, S., MCPHERSON, D., OBERHEIDE, J., AND JAHANIAN, F. Internet Inter-Domain Traffic. In *ACM SIGCOMM Computer Communication Review (2010)*, vol. 40, ACM, pp. 75–86.
- [17] NYGREN, E., SITARAMAN, R. K., AND SUN, J. The Akamai network: a platform for high-performance internet applications. *ACM SIGOPS Operating Systems Review* 44, 3 (2010), 2–19.
- [18] OBAR, J., AND CLEMENT, A. Internet Surveillance and Boomerang Routing: A Call for Canadian Network Sovereignty. In *TEM 2013: Proceedings of the Technology & Emerging Media Track, Annual Conference of the Canadian Communication Association (2013)*, p. 2.
- [19] RIPE ATLAS COMMUNITY. OpenIPMap: Geolocating Internet Infrastructure. <https://github.com/RIPE-Atlas-Community/openipmap>.