

When Traceroute Met BGP...

How to Reveal Hidden Internet AS-level Connectivity with Portolan and Isolario

Adriano Faggiani*[†], Pietro G. Giardina[†], Enrico Gregori[†], Alessandro Improta[†],
Luciano Lenzini*, Valerio Luconi*, Alessandro Pischredda[†], Lorenzo Rossi[†]
Luca Sani[†]

*Information Engineering Department, University of Pisa, Pisa, Italy
{l.lenzini, valerio.luconi}@iet.unipi.it

[†]Institute of Informatics and Telematics, Italian National Research Council, Pisa, Italy
{name.surname}@iit.cnr.it

Abstract—The incompleteness of the Internet AS-level topology is the most serious pitfall of research studies concerning the Internet inter-domain ecosystem. In this demo we show two projects aimed at filling this gap by offering services to different type of users. The first is Portolan, a smartphone-based crowdsourcing system which provides an app for Android that volunteers can install to run traceroute measurements to discover the Internet topology. Then, we introduce Isolario, which provides *real-time* services for network operators in change of their full routing table. Finally, we describe joint services provided by Portolan and Isolario and the benefit obtained.

Keywords—BGP, real-time monitoring, route collector, traceroute, crowdsourcing

I. INTRODUCTION

The incompleteness of the Internet AS-level topology is a well-known issue among researchers. Several works have been carried out on this research topic in the last decade, mostly focusing on depicting evidences of such incompleteness, or on trying to identify which parts of the Internet connectivity were actually missing [1]. Only recently some works focused on investigating the deep causes of this incompleteness [2]. In particular, in [3] was found that the main cause is related to the fact that almost only large Internet Service Providers (ISPs) are actually connected to the BGP route collectors, and use them mostly as a free of charge opportunity to advertise their network reachability. As a consequence of this view from the top of the hierarchy, a large amount of connections established among ASes placed in the lower levels of the Internet are missing, and cannot be revealed by current BGP route collectors. The only opportunity to reveal these connections would be to increase drastically the amount of ASes located in the Internet periphery that are feeding the route collecting infrastructure. To fill the gap, we developed Portolan and Isolario. Portolan exploits the contribution of volunteers to perform distributed active measurements to discover the Internet AS-level topology. The main advantage of Portolan is its crowdsourcing approach, which allows to place measurement monitors into ASes without having to obtain the explicit consensus from their administrators. Isolario instead

was designed on the fact that most of AS administrators are not interested in joining current BGP route collecting projects just to advertise their network reachability or for mere altruism. Differently Isolario aims at persuading the administrators of ASes owned by small-medium organizations to share their full routing table by offering useful services in return. In this paper we introduce some of the services already available both in Portolan and Isolario, and we show how the two infrastructures could work together.

II. PORTOLAN IN A NUTSHELL

Portolan is a crowdsourcing system aimed at discovering the Internet AS-level connectivity by exploiting smartphones as measurement monitors. Crowdsourcing is particularly suited for AS-level topology discovery, as it allows to have a potentially large number of users spread all over the Internet, especially in the lowest layers of the Internet hierarchy. In particular, crowdsourcing allows to place monitors in ASes (and thus to discover their connectivity) without having to reach an agreement with their administrators for obtaining routing information. Moreover, smartphones bring a significant advantage, due to their mobility that allows to visit multiple networks with a single device. Thus, each smartphone can act as several point of views in the Internet at different times. In addition, smartphones are easily geolocalized through the embedded GPS unit. Volunteers can participate in Portolan by installing the Portolan app for Android, freely downloadable from the Google Play Store¹. The Portolan app provides the measurement functionalities for actively probing the Internet [4], i.e. a UDP version of Paris Traceroute [5]. The measurement activities of smartphones are coordinated by a server infrastructure, which assigns measurement tasks to smartphones and collects the results [6]. Since traceroute provides a sequence of IP interfaces, the server is also responsible of performing IP-to-AS mapping in order to obtain the AS-level topology. Since a crowdsourcing system without crowd is useless, Portolan

¹<https://play.google.com/store/apps/details?id=it.unipi.iet.portolan.traceroute>

provides several useful network diagnostic services aimed at motivating users to install the Portolan app. The most relevant are a *UDP Paris traceroute with multipath detection algorithm*, a tool for measuring the maximum throughput available on the user's network (*Maximum throughput estimator*), and a tool for tracking the cellular signal coverage over a path (*Signal coverage mapping*). Currently Portolan counts over 250 active users spread all over the world (mostly Europe and U.S.).

III. ISOLARIO IN A NUTSHELL

The goal of Isolario is to improve the knowledge about the AS-level ecosystem of the Internet by increasing the amount of ASes from which BGP data is collected, hereafter *feeders*. To incentivate AS administrators to join the Isolario route collecting infrastructure, we provide them *real-time* monitoring and alerting services on the health status of their own BGP routing system. With these services, a feeder administrator would be able to detect, monitor and analyse pathological events affecting its BGP routing system, like route flapping, prefix hijacking and loss of reachability, without increasing the computational load on the router or introducing third-party software. In order to join the Isolario project, we require feeders to establish (at least) one BGP session with one of our route collectors and to announce us their routes towards all the Internet destinations. The BGP data collected from feeders is used both to provide the real-time services and to improve the quality and completeness of the Internet AS-level ecosystem. Upon agreement with the feeder, BGP data will be made publicly available in MRT format (RFC 6396) to the research community, similarly to what is done by existing route collecting projects. The Isolario system is composed by *a)* a set of Route Collectors, which establish BGP sessions with external feeders, *b)* a real-time server dedicated to analyse in real-time BGP messages coming from each feeder, *c)* an off-line server dedicated to periodically analyse collected BGP data, *d)* a set of servers dedicated to store information and results of computations, and *e)* a web server to offer our real-time services to the Isolario users and the results of our analyses concerning the Internet AS-level ecosystem to the public users (<http://www.isolario.it>).

Isolario main feature is the provisioning of real-time services to the users feeding the route collectors. The real-time feature is obtained in two steps. First, each incoming BGP flow is directly parsed, filtered and redirected towards dedicated modules which implement the services. Then, the result is provided to Isolario users through an HTML5 website which exploits the *WebSocket* protocol (RFC 6455) to update web pages only when new data related to the client become available, without additional polling traffic being generated. Some of the already available services enables Isolario users to monitor in real-time the routes that its AS is using to reach a set of Internet destinations (*Routing table viewer*), to monitor real-time occurrences of route flapping (*Route flap detector*) and to check in real-time the routes that other feeders use to reach her/his own AS networks (*My route reachability monitoring*).

IV. PORTOLAN AND ISOLARIO INTERACTION

Even if Isolario and Portolan operate on two orthogonal planes (the control plane and the data plane respectively), they cooperate in order to enhance their services and performances. On one hand, Portolan uses BGP routing tables provided on a monthly basis by Isolario, to perform focused traceroute campaigns at a regional-level. On the other hand, Isolario exploits the geographic distribution of Portolan probes to enhance some of its services. For example, in *My route reachability monitoring* the Isolario used is allowed to perform traceroutes and ping measurements to understand how the Portolan probes are reaching its own subnets. To do that, Isolario triggers traceroute and ping measurements from Portolan probing-smartphones through the Google Cloud Messaging service, and provides to the user IP-traces and latency between his AS networks and the ASes hosting at least one Portolan agent.

V. FUTURE WORKS

Portolan and Isolario are two new systems that provide services to different categories of users to receive benefits in an AS-level topology discovery perspective. Portolan has been deployed since one year, and several services are already implemented and running. In the near future we plan to implement an IP geolocation service based on RTT measurement and triangulation. Isolario is still on a very early stage of development. The listed services are only the first that have been deployed, and many more will be created in the next future. For example, we are planning to develop services to detect prefix hijack events in real-time and to develop a proper alerting systems to alarm the network administrator as soon as pathological events occur. Also the interaction between the two systems is at a very early stage. To improve it we plan to introduce a budgeting system that would allow Isolario users that agree to install the Portolan app to perform more daily measurements than a normal Isolario user.

ACKNOWLEDGEMENTS

This work has been partially supported by the European Commission within the framework of the CONGAS project FP7-ICT-2011-8-317672.

REFERENCES

- [1] Y. He, G. Siganos, M. Faloutsos, and S. V. Krishnamurthy, "Lord of the Links: A Framework for Discovering Missing Links in the Internet Topology," *IEEE/ACM ToN*, vol. 17, no. 2, pp. 391–404, 2009.
- [2] R. Oliveira, D. Pei, W. Willinger, B. Zhang, and L. Zhang, "The (In)Completeness of the Observed Internet AS-level Structure," *IEEE/ACM ToN*, vol. 18, no. 1, pp. 109–122, 2010.
- [3] E. Gregori, A. Improta, L. Lenzini, L. Rossi, and L. Sani, "On the Incompleteness of the AS-level Graph: a Novel Methodology for BGP Route Collector Placement," in *Proc. of ACM IMC*, 2012, pp. 253–264.
- [4] A. Faggiani, E. Gregori, L. Lenzini, S. Mainardi, and A. Vecchio, "On the Feasibility of Measuring the Internet through Smartphone-based Crowdsourcing," in *Proc. of IEEE WiOpt*, 2012, pp. 318–323.
- [5] B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira, "Avoiding Traceroute Anomalies with Paris Traceroute," in *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, ser. Proc. of ACM IMC, 2006, pp. 153–158.
- [6] E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio, "Sensing the Internet through crowdsourcing," in *Proc. of IEEE PerMoby*, 2013, pp. 248–254.