Poster Abstract: Network Sensing through Smartphone-based Crowdsourcing

Adriano Faggiani*, Enrico Gregori*, Luciano Lenzini†, Valerio Luconi†, Alessio Vecchio†
*IIT-CNR, Pisa, Italy
firstname.lastname@iit.cnr.it
†Dip. Ingegneria dell'informazione, University of Pisa, Pisa, Italy
firstname.lastname@iet.unipi.it

ABSTRACT

Portolan is a crowdsourcing system, aimed at monitoring and measuring large-scale networks, that uses smartphones as mobile observation elements. Currently, Portolan is able to collect information about both wired and wireless networks, in particular it is used to obtain the graph of the Internet with unprecedented resolution and to associate performance indexes (received signal strength, maximum throughput) of cellular networks to geographic locations.

Categories and Subject Descriptors

C.2.3 [**Network Operations**]: Network monitoring

Keywords

Crowdsourcing, network sensing, smartphone

1. INTRODUCTION

Information about the performance and the structure of large networks, such as the Internet, is extremely valuable both from a research and a commercial point of view. However, collecting data about such complex and large-scale systems is a challenging task, given the massive requirements in terms of computing and economical resources.

Crowdsourcing represents a possible answer: *i*) the required amount of work can be split over a large set of participating users; *ii*) the evaluation of performance can be carried out at the periphery of the network, where the vast majority of end users is located; *iii*) the geographic distribution of participants enables the collection of fine-grained measurements scattered over wide regions.

At the same time, the adoption of smartphones as a network sensing platform brings several advantages. First, smartphones are ever increasing in number (the sales of smartphones have surpassed the sales of PCs). Thus, the limited resources of single devices, in terms of computational and networking power, are balanced by their aggregate power. Second, smartphones are highly mobile and always on. As a consequence, a single device may provide different points of observation of the network (at different times), and automated actions can be scheduled quite easily. Finally, they embed a GPS unit for geolocalization. This adds a new class of network monitoring scenarios where the geographic dimension plays a relevant role.

These ideas led us to build Portolan, a smartphone-based crowdsourcing system for measuring large networks. Currently Portolan is used in two network measurement scenarios: *i*) Internet mapping, and *ii*) cellular signal coverage.

SenSys '13 , Nov 11-15 2013, Roma, Italy
ACM 978-1-4503-2027-6/13/11.
http://dx.doi.org/10.1145/2517351.2517397

2. THE PORTOLAN ARCHITECTURE

Several challenges have to be faced when building a crowdsourcing system based on smartphones. Smartphones have low computational power, limited communication capabilities, especially if connected via cellular network, and are battery operated. These aspects limit the amount of work that a single smartphone can carry out without consuming too many resources and bothering the user. Moreover, since smartphones are under the user's control, they can come and go with unpredictable patterns. The system has also been designed in a modular way, in order to ease the implementation and integration of additional network sensing subsystems.

The Portolan system is organized as a distributed client-server architecture. Smartphones play the role of clients and run an app that implements measurement functionalities. The Portolan app is currently available for Android at the Google Play Store[1]. Each participating smartphone performs a small amount of local measures. The contributions of all participants are collected and aggregated on the Portolan server. The server is also responsible of coordinating the smartphone activities and of assigning the measurement tasks to the smartphones. Tasks are specified by the system administrator and submitted to the server, which, after performing some consistency controls, splits the task in smaller parts (microtasks) that can be executed by mobile devices. Tasks include the specification of properties that smartphones need to satisfy for performing them, e.g. the geographic position or the network address. A detailed description of Portolan implementation is provided in [1, 2].

For motivating people to participate in the Portolan activities, the app includes also a suite of network diagnostic tools (e.g. traceroute, maximum throughput estimator), which can be manually launched by the interested users.

[1] `https://play.google.com/store/apps/details?id=it.unipi.iet.portolan.traceroute`

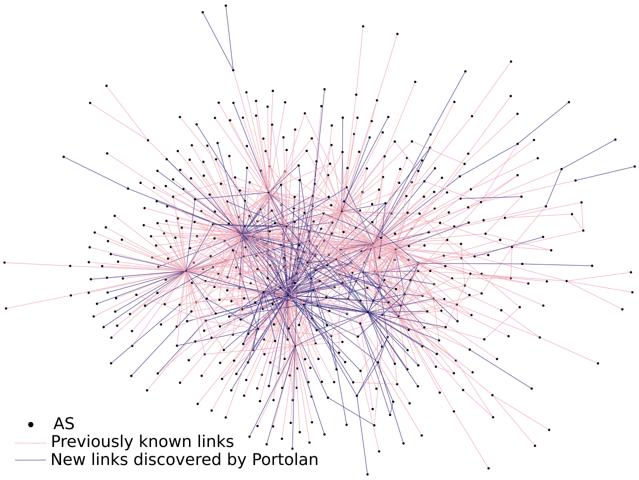

Figure 1: Italian Internet structure discovered by Portolan

3. MAPPING THE INTERNET

The nowadays Internet is an enormous interconnection of millions of networks, such as transit providers, academic networks, content delivery networks, etc. The vast majority of Internet actors are private companies, and the connections that they establish with each other are driven by business interests. Thus, they are unwilling to share their network structure and interconnections. Nevertheless, such information would be useful for several reasons, like planning business strategies between Internet Service Providers (ISPs) or analyzing topological properties of the Internet itself. These reasons drove several organizations (private or academic) to build measurement infrastructures for discovering the Internet topology. However, despite such effort the complete knowledge of the Internet structure is still to be achieved.

We used Portolan for discovering the topology of the Internet at the Autonomous System (AS) level of abstraction (for our purposes an AS is equivalent to an ISP). The approach of Portolan is based on active measurements, performed with a UDP version of Paris Traceroute, implemented in the Portolan app. Data aggregation is performed on the Portolan server, by running an IP-to-AS mapping process, which produces the AS-level topology.

To evaluate the soundness of Portolan's approach, we submitted a measurement campaign for discovering the Italian Internet structure as seen by three major Italian ISPs. The traceroute target was the entire Italian address space. Portolan discovered 1117 links that we compared with a dataset provided by CAIDA[2], another Internet mapping project (*IPv4 Routed /24 AS Links* dataset, January 2013). 244 links out of 1117 are not present in CAIDA's dataset (21.8%). An intuitive representation of results is shown in Figure 1.

4. SIGNAL COVERAGE MAPPING

The campaign dedicated to mapping the cellular signal gathers Received Signal Strength (RSS) samples and associates them to geographic positions provided by the GPS unit. Such information is useful both for cellular service providers, for improving their coverage, and for end-users, to select the most suitable operator for their needs.

Since the GPS unit is extremely battery-consuming, we decided not to remotely start RSS measurements. Instead samples are gathered in two different ways. The first is user-driven. Users can manually start/stop measurement for mapping the RSS along a specific path. The second uses the Android passive location provider. If another app starts the GPS, the collected positions are also forwarded to the Portolan app which uses them for geolocalizing RSS samples. Data is then sent to the Portolan server.

Actually over 1,000,000 RSS samples have been collected and are shown on a map on the Portolan website[3]. Figure 2 shows the signal coverage in the surroundings of Pisa. An analysis of the map shows interesting aspects. For example the quality of the signal can experience significant variations even in a few meters. Moreover, the RSS in a given location can vary significantly with different operators.

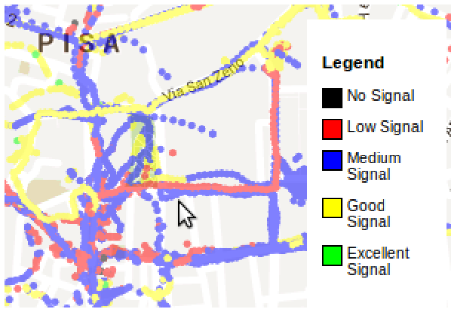

Figure 2: Signal coverage map (Map data ©2013 Google)

5. CONCLUSION AND FUTURE WORK

We briefly presented Portolan, a crowdsourcing system for measuring large-scale networks which uses smartphones as measuring agents. The philosophy of Portolan is to perform short range measurements for evaluating network characteristics where the end-users are. Currently Portolan is able to map the Internet at the AS-level of abstraction and to build signal coverage maps. The results obtained so far confirm the soundness of the Portolan approach.

As future work, we plan to develop and integrate in Portolan two new measurement subsystems. The first is aimed at geolocalizing Internet hosts, by using smartphones as mobile landmarks; the purpose of the second is detecting if ISPs differentiate traffic without informing the user.

6. REFERENCES

[1] 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 WiOpt'12*, pages 318–323, May 2012.

[2] E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio. Sensing the Internet through crowdsourcing. In *Proc. of PerMoby'13*, May 2013.

[2] http://www.caida.org/

[3] http://portolan.iet.unipi.it/