Ramble: opportunistic content dissemination for infrastructure-deprived environments

Miguel Garcia, João Rodrigues, Eduardo R. B. Marques, and Luís M. B. Lopes

CRACS/INESC-TEC & Faculty of Science, University of Porto, Portugal

Abstract. We present Ramble, a system for geo-referenced content-sharing in environments that have limited infrastructural communications, motivated by applications such as rescue operations in disaster scenarios or citizen-science data collection in areas with low communication coverage. Ramble makes use of mobile edge-clouds, networks formed by mobile devices in close proximity, and lightweight cloudlets that serve a small geographical area. Using an Android app, users roam ("ramble") through a given area whilst generating geo-referenced content (e.g., reports, photos, or videos), and disseminate that content opportunistically to nearby peers, cloudlets, or even cloud servers, as allowed by intermittent connections of different types (e.g., respectively WiFi-Direct, WiFi, or 3G/4G). Each node in the system may thus act as a content generator, consumer, cache, and disseminator. In this extended abstract, we describe the main traits of the Ramble system and also a field experiment for evaluating the system to be conducted in the short term.

1 Introduction

Personal devices like smartphones and tablets have become ubiquitous in recent years and are increasingly powerful. Today's devices are able to perform computationally intensive tasks, such as rendering 3D graphics in games or playing high-quality video. Besides these types of tasks, smartphones have also become a source for data collection, for instance taking pictures, recording video and sound, reading fingerprints and performing geolocation. There's also been advances in the inclusion of wireless D2D communication technologies such as WiFi-Direct, WiFi TDLS and Bluetooth along with the traditional WiFi and 3G/4G. D2D communications, in particular, may greatly empower crowdsourcing applications that are now commonplace for instance in disaster scenario operations [2] or citizen science [4], given that the information is generated in a bottom-up fashion by users, but its dissemination may be impaired by infrastructural communications outages or their lack of coverage.

Infrastructure-deprived environments are one of the core scenarios considered in the scope of the Hyrax project (`http://hyrax.dcc.fc.up.pt`), that deals with the use of mobile edge-clouds, networks formed by nearby mobile devices using D2D communications, as enablers of crowd-sourced mobile applications. With this context, and inspired by web services like the Scipionus hurricane information maps [1] (among others), that aggregate and depict geo-tagged, crowd-sourced information in a map, we have developed Ramble, a system for opportunistic generation, dissemination and visualization of geo-tagged user-generated content. Ramble makes use of mobile edge clouds enabled by WiFi-Direct, lightweight proximity cloudlets that act as repositories and disseminators of user-collected data in small areas, plus, optionally, traditional cloud servers accessible via Internet. In what follows, we describe the main traits of Ramble and a field experiment we expect to conduct in the very short term to evaluate it.

2 The Ramble system

The Ramble system is illustrated in Fig. 1 in terms of (a) its architecture, and (b) the app employed by user to generate and visualize geo-referenced data in Android devices. The architecture comprises three types of network peers: mobile devices, cloudlets and cloud servers. The peers are connected through WiFi-Direct (for D2D communication), WiFi (device-to-cloudlet), mesh networking (cloudlet-to-cloudlet), and WiFi or 3G/4G Internet (cloud server communication). The main goal of the system is that user-generated content is disseminated during fortuitous connections among devices or between devices and cloudlets or cloud servers, as we envision that connections will be volatile as users roam in a reasonably large space. The Android app presents the user with a map where content markers are displayed (as shown in the app screenshot), and has the ability of generating new content in the form of text reports, photos, or videos that are geo-reference to the user's current location (a screenshot of this functionality is omitted for reasons of space).

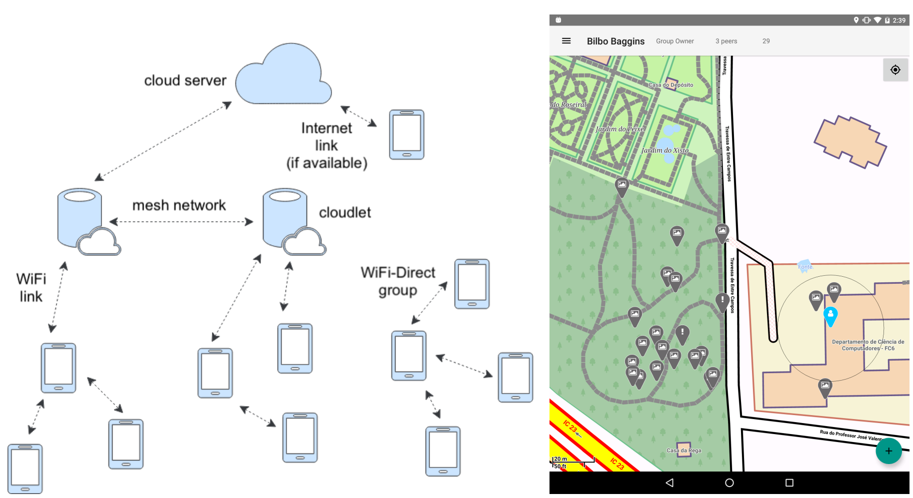

(a) Architecture. (b) Android app.

Fig. 1: Overview of Ramble.

The system is symmetric in the sense that each peer may act as a content generator, consumer, cache, and disseminator, with no pre-established hierarchy in terms of functionality. In the current version, we focus solely on mobile devices for the role of content generation, even if in principle data could be injected directly and disseminated through cloudlets and cloud servers for instance. This functional symmetry is afforded by the use of a Ramble service that deals with the core functionality shared by all nodes. The main functionalities of this service and some other important peer-specific ones are as follows:

- Local storage for (generated and received) content using an SQLite database instance;
- Location-awareness through GPS in mobile devices, used to tag generated content;
- The establishment of network links, in particular by forming WiFi-Direct groups using the Hyrax middleware [3], and mesh networking in cloudlets using the BATMAN protocol;
- The dynamic discovery of peers on top of established network links, using UDP broadcast or multicast;
- Content synchronization between two peers, such that each peer works by identifying data it does not hold stored remotely and transfers it accordingly. If desired, the synchronization may be selective through subscription filters configured per peer that account for content type, source location, and generation time. The synchronization service is implemented using gRPC (`https://grpc.io`), facilitating interoperability between heterogeneous peers (e.g., programmed in different languages).

3 Experiments

Figure 2 illustrates a field experiment we will conduct in the very short term using Ramble. The map shown depicts an area of Porto's Botanic Garden, which is just next to our department. Following the Ramble architecture, as shown in the figure as overlay, we will deploy Android devices that may form WiFi-Direct groups, Raspberry Pi cloudlets spread throughout the garden connected through a mesh network and serving as access points for mobile devices, and a cloud server hosted in our department's network that is accessible by an Internet connection.

During the experiment, in the spirit of a citizen-science data collection over a given area, volunteer users will be instructed to carry Google Nexus tablets without 3G/4G connectivity and be instructed to find and take photos of specific trees in the garden using the Ramble app. The experiment will be gamified in the sense that users should only take photos of unsampled areas/trees according to the local view of the Ramble app in the device they are carrying, As content dissemination occurs intermittently, that local view may of course be inconsistent and some areas may happen to be over-sampled in the end. We should also note that there is no public WiFi connectivity on location, except possibly in a few spots closer to our department (east/right in the figure), and that, for a

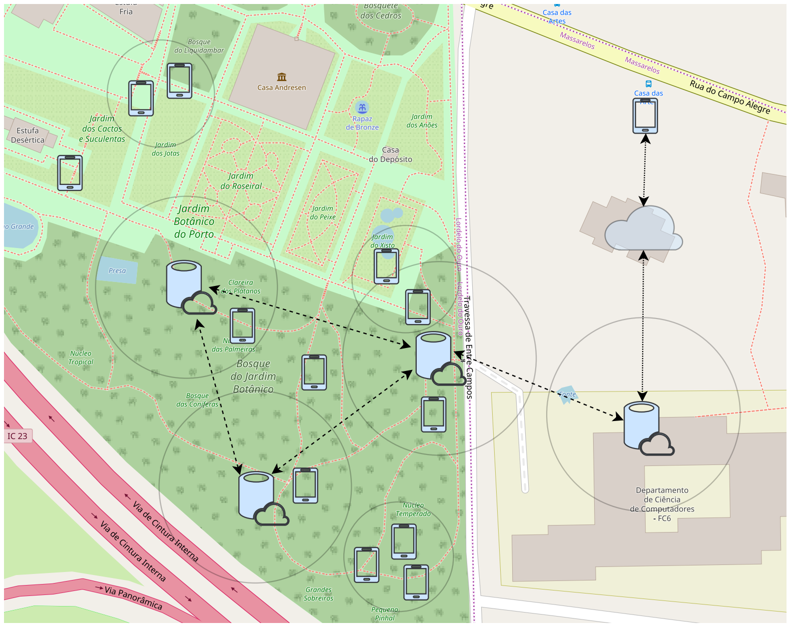

Fig. 2: Planned field experiment at Porto's Botanic Garden.

demanding scenario, the experiment will be calibrated such that some survey areas will be out of range for cloudlet connectivity.

An analysis of the results of the experiment will follow to determine metrics such as the dissemination recall and hops per generated content, and the contribution of each type of connection in that regard (WiFi-Direct, WiFi, Internet). We expect to present these results at INFORUM, if this abstract is accepted.

References

1. Scipionus - Hurricane Information Maps. `https://gregstoll.dyndns.org/scipionus/`, [Accessed: January 2018]
2. Poblet, M., García-Cuesta, E., Casanovas, P.: Crowdsourcing tools for disaster management: A review of platforms and methods. In: AI Approaches to the Complexity of Legal Systems, pp. 261–274. Springer (2014)
3. Rodrigues, J., Marques, E.R.B., Lopes, L., Silva, F.: Towards a middleware for mobile edge-cloud applications. In: Proc. 2nd Workshop on Middleware for Edge Clouds & Cloudlets (MECC'17). pp. 1:1–1:6. MECC'17, ACM (2017)
4. Simpson, R., Page, K.R., De Roure, D.: Zooniverse: Observing the world's largest citizen science platform. In: Proc. WWW. pp. 1049–1054. ACM (2014)

Acknowledgements. This work was funded by the NORTE 2020 project NORTE-01-0145-FEDER-000020, the COMPETE 2020 Programme project POCI-01-0145-FEDER-006961, and by FCT projects CMUP-ERI/FIA/0048/2013 and UID/EEA/50014/2013.