u-map: Empowering Users and Telecom Providers by Employing the Crowd-Sourcing Paradigm

Michalis Katsarakis

Thesis submitted in partial fulfillment of the requirements for the

Masters' of Science degree in Computer Science

University of Crete
School of Sciences and Engineering
Computer Science Department
University Campus, Voutes, Heraklion, GR-70013, Greece

Thesis Advisor: Associate Professor *Maria Papadopouli*

Heraklion, March 2015

This work has been performed at the Foundation for Research and Technology–Hellas, Institute of Computer Science (FORTH–ICS), N. Plastira 100 Vassilika Vouton, GR-700 13 Heraklion, Crete, Greece. The work is partially supported by the General Secretariat for Research and Technology in Greece with a Research Excellence, Investigator-driven grant, 2012, and by a Google Faculty Research Award, 2013, PI Maria Papadopouli.

University of Crete
Computer Science Department

u-map: Empowering Users and Telecom Providers by Employing the Crowd-Sourcing Paradigm

Thesis submitted by
Michalis Katsarakis
in partial fulfillment of the requirements for the
Masters' of Science degree in Computer Science

THESIS APPROVAL

Author: ______________________
Michalis Katsarakis

Committee approvals: ______________________
Maria Papadopouli
Associate Professor, Thesis Supervisor

____
Evangelos Markatos
Professor, Committee Member

____
Xenofontas Dimitropoulos
Assistant Professor, Committee Member

Departmental approval: ______________________
Antonis Argyros
Professor, Director of Graduate Studies

Heraklion, March 2015

Abstract

Wireless access, use and traffic demand are on a fast rise. The number of mobile devices and their capabilities, accessing potentially multiple wireless network interfaces, also increase dramatically. Wireless networks experience "periods of severe impairment", causing severe degradation to the performance of the service running on wireless devices and to the respective user experience. The crowd-sourcing and participatory sensing in large-scale systems provide the capability to monitor in real-time the quality of service (QoS).

Motivated by the need to enable users to select the appropriate provider/operator, especially when visiting a new environment, we developed the u-map system. A u-map client running on a smartphone uploads in regular basis information about the user profile (e.g., service requirements, device capabilities), QoE feedback about services (as indicated explicitly by the user), traffic demand statistics, network/spectrum conditions, and position of the device. This information is stored in a spatio-temporal geo-database, attached to a u-map server. A u-map client may query this database to obtain information about service providers and their customers' satisfaction. In that way, a user can make an educated selection of the appropriate provider, when visiting a new area. The term appropriate denotes the provider that best matches the user's criteria (e.g., it is the best in term of a certain QoS metric, or in QoE, in that area, or it supports specific features). Moreover, given that a smartphone may have access to several network interfaces simultaneously, a u-map client may select the appropriate wireless or cellular operator and network interface for offloading its traffic.

This thesis presents the u-map system and analyzes its performance, in terms of power consumption, responsiveness, and scalability. To understand how users assess such systems in general, and how they perceive the performance of the u-map and its services specifically, we also performed a 2-week field study with 21 real users. The outcomes of the performance analysis and the field study are encouraging: The power consumption and response delay of the u-map are relatively low. Most of the participants indicated that the u-map can improve their experience and helps in the discovery of networks that offer improved QoS. During the study, 30% of users discovered the best available network, improving substantially their experience. Users were also satisfied with the graphical user interface (GUI) of the u-map and other specific features, as well as the overall performance of the u-map service. Based on user feedback, we also modeled the user satisfaction about the u-map as a utility function. The model can predict the overall score of the user satisfaction about the u-map accurately. Finally, we applied the u-map in other domains, such as water distribution networks and medical applications, demonstrating its generality.

Περίληψη

Η ζήτηση και η χρήση ασύρματης πρόσβασης αυξάνονται ραγδαία. Το πλήθος των φορητών συσκευών και οι ικανότητές τους, συμπεριλαμβανομένης και της δυνατότητας πρόσβασης σε πολλαπλές διεπαφές δικτύου, επίσης αυξάνονται δραματικά. Τα ασύρματα δίκτυα παρουσιάζουν «περιόδους οξείας δυσλειτουργίας», προκαλώντας σημαντική υποβάθμιση στην απόδοση της υπηρεσίας που εκτελείται σε ασύρματες συσκευές και την αντίστοιχη εμπειρία του χρήστη. Η συλλογή δεδομένων από πλήθος φορητών συσκευών σε συστήματα μεγάλης κλίμακας παρέχει τη δυνατότητα να εποπτεύουμε σε πραγματικό χρόνο την ποιότητα υπηρεσίας (QoS).

Παρακινούμενοι από την ανάγκη να επιτρέψουμε στους χρήστες να επιλέγουν τον κατάλληλο πάροχο, ειδικά όταν επισκέπτονται ένα νέο περιβάλλον, αναπτύξαμε το σύστημα u-map. Ένας πελάτης u-map (u-map client) που εκτελείται σε μια έξυπνη συσκευή κινητής τηλεφωνίας (smartphone) μεταφορτώνει τακτικά πληροφορία σχετικά με το προφίλ του χρήστη (π.χ., απαιτήσεις του χρήστη, δυνατότητες της συσκευής), εκτιμήσεις ποιότητας εμπειρίας σχετικά με υπηρεσίες (που αναφέρονται ρητά από τον χρήστη), στατιστικά ζήτησης, δικτυακές/φασματικές συνθήκες, καθώς και τη θέση της συσκευής. Αυτή η πληροφορία αποθηκεύεται σε μια χωρο-χρονική βάση δεδομένων συνδεδεμένη στον εξυπηρετητή u-map (u-map server). Ένας πελάτης u-map μπορεί να κάνει επερωτήσεις στη βάση δεδομένων για να αποκτήσει πληροφορία σχετική με παρόχους υπηρεσίας και την ικανοποίηση των πελατών τους. Με αυτόν τον τρόπο, ένας χρήστης μπορεί να κάνει μια μελετημένη επιλογή του κατάλληλου παρόχου, όταν επισκέπτεται μια νέα περιοχή. Ο όρος «κατάλληλος» υποδηλώνει τον πάροχο που ταιριάζει περισσότερο στα κριτήρια του χρήστη (π.χ., είναι ο καλύτερος με βάση κάποια συγκεκριμένη μετρική ποιότητας υπηρεσίας, ή την ποιότητα εμπειρίας, στην εκάστοτε περιοχή, ή υποστηρίζει συγκεκριμένα χαρακτηριστικά). Επιπλέον, δεδομένου ότι μια έξυπνη συσκευή κινητής τηλεφωνίας μπορεί να έχει πρόσβαση σε μερικές διεπαφές δικτύου ταυτόχρονα, ένας πελάτης u-map μπορεί να επιλέξει τον κατάλληλο ασύρματο ή κυψελωτό πάροχο και διεπαφή δικτύου για την αποσυμφόρηση της κίνησης δεδομένων της συσκευής.

Αυτή η μεταπτυχιακή εργασία παρουσιάζει το σύστημα u-map και αναλύει την απόδοσή του, σε όρους κατανάλωσης ισχύος, αποκρισιμότητας, και κλιμακωσιμότητας. Για να κατανοήσουμε πως οι χρήστες αξιολογούν τέτοια συστήματα γενικά, και πως αντιλαμβάνονται την απόδοση του u-map και των υπηρεσιών του ειδικότερα, εκτελέσαμε επίσης μια μελέτη πεδίου 2 εβδομάδων με 21 πραγματικούς χρήστες. Οι εκβάσεις της ανάλυσης απόδοσης και της μελέτης πεδίου είναι ενθαρρυντικές: Η κατανάλωση ισχύος και η καθυστέρηση απόκρισης του u-map είναι σχετικά χαμηλές. Οι περισσότεροι από τους συμμετέχοντες υπέδειξαν ότι το u-map μπορεί να βελτιώσει την εμπειρία τους και βοηθά στην ανακάλυψη δικτύων που προσφέρουν βελτιωμένη ποιότητα υπηρεσίας. Κατά τη διάρκεια της μελέτης, το 30% των χρηστών ανακάλυψαν το καλύτερο διαθέσιμο δίκτυο, βελτιώνοντας ουσιαστικά την εμπειρία τους. Οι χρήστες ήταν επίσης ικανοποιημένοι με το γραφικό περιβάλλον χρήστη (GUI) του u-map και άλλα συγκεκριμένα χαρακτηριστικά του, καθώς και με τη γενική απόδοση της υπηρεσίας του u-map. Βασισμένοι στην ανατροφοδότηση από τους χρήστες,

μοντελοποιήσαμε επίσης την ικανοποίηση των χρηστών σχετικά με το u-map σαν μια συνάρτηση ωφελιμότητας. Το μοντέλο μπορεί να προβλέπει τη γενική βαθμολογία της ικανοποίησης των χρηστών σχετικά με το u-map με ακρίβεια. Τέλος, εφαρμόσαμε το u-map σε άλλες περιοχές, όπως δίκτυα διανομής νερού και ιατρικές εφαρμογές, επιδεικνύοντας την γενικότητά του.

Acknowledgements

This thesis is a result of a three-year working experience as a graduate student at the Telecommunications and Networks Laboratory of the Institute of Computer Science (ICS) of the Foundation of Research and Technology – Hellas (FORTH).

First of all, I would like to thank my advisor, Prof. Maria Papadopouli, not only for her encouragement, but also for her guidance and her support throughout this work. Without her constructive remarks and the time she devoted to me, the completion of thesis would not be possible.

Special thanks also go to the members of my dissertation committee, Professors Evangelos Markatos and Xenofontas Dimitropoulos for their constructive comments and questions.

My warmest thanks to my classmates-colleagues-friends at the Mobile Computing Activity team: Alexandros Kostopoulos, Georgios Fortetsanakis, Maria Plakia, Paulos Charonyktakis, Charalampos Meidanis, Nikos Rapousis, Vasileios Theodosiadis, and Evripidis Tzamusis for the good time with laughter, mutual encouragement and assistance we shared together. Equally warm thanks go to the alumni of the team: Ioannis Stiakogiannakis, Ioannis Dimitriou, Dafni Zafiri, Kostas Triantafyllakis, George Vardakis, Nick Syntychakis, Menelaos Viskadouros, and Panagiotis Lionakis.

Additional thanks to Georgios Fortetsanakis for his contribution on the user profiling (Section 6.2.10), Nikos Rapousis for his contribution on the automatic network selector and traffic demand estimator (Section 4.1.1), as well as Vasileios Theodosiadis for his contribution on the power consumption measurements (Section 6.1.1). I would also like to thank all members of the TNL and SPL labs for all the helpful discussions, encouragement, and nice atmosphere.

This work would not have been completed without the valuable help and patience of the 21 volunteers who participated in the user study described in Chapter 6. Guys, thank you all.

A dedicated "thank you" to my friends Georgios Bouloukakis, Sofoklis Fasoulas, Giannis Filippakis, Manos Kapelonis, Ioanna Labraki, Apostolis Veniotis, and Despoina Xatziandreou for brightening up my days and for always providing their advice and support.

Last, but definitely not least, I would like to thank my family. They have sacrificed everything in order to help me reach my goals. Without their help I would certainly have not made it to here. Μαμά, μπαμπά, Φώτη, και Ανδρέα σας ευχαριστώ και σας αγαπώ πολύ.

To my family.
Στην οικογένειά μου.

Contents

List of Figures

List of Tables

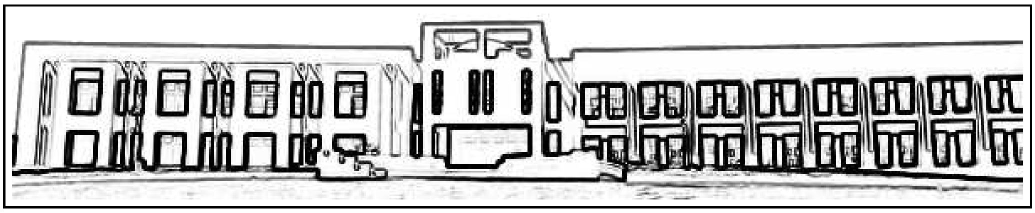

Front view of the main building of FORTH.

Chapter 1

Introduction

"This is where it all begins. Everything starts here, today."
– David Nicholls, *One Day*

Wireless access, use, and traffic demand are on fast rise. By 2019, the mobile data traffic will reach an annual run rate of 291.8 Exabytes worldwide, increasing approximately 10-fold since 2014, while video traffic will be increased 13-fold during the same period [1]. The number of mobile devices and their capabilities also increase dramatically. The wide range of applications and services running on smartphones have differentiated network requirements. New, emerging services further individualize network requirements. Examples of such services are mobile health [2], mobile machine-to-machine (M2M) [3], home automation [4], real-time monitoring, and telepresence applications. On the other hand, wireless networks experience "periods of severe impairment", causing severe degradation of the *Quality of Service (QoS)*. Traditionally, QoS indicators based on throughput, latency, jitter, and packet loss have been employed to quantify network and service performance. However, such metrics cannot capture adequately the human-perceived *Quality of Experience (QoE)*. Due to the diverse set of services, network operators, and users, the estimation of the QoE is challenging and largely underexplored. Furthermore, in various production networks, there are no automated real-time QoE diagnostic tools.

Not only the networks but also the business models of the communication services over these networks change. The technology for dynamic network selection is ready: Smartphones with dual Subscriber Identity Module (SIM) cards have already been designed [5, 6]. The flex service, which has been assumed as a typical access paradigm in wireless Local Area Networks (LANs), has also been proposed for cellular networks and studied [7, 8]. Digital marketplaces where providers, mobile network operators (MNOs), and users can trade communications services [9], or where MNOs bid for the right to transport users' requested services over their infrastructure [10] have already been proposed in the literature. Another recent trend in wireless access markets is the presence of mobile virtual network operators (MVNOs). MVNOs do not own spectrum or network infrastructure; instead they lease resources (e.g., access) in the wholesale market from MNOs and resale access services in the retail market [11].

The advances in wireless and sensor networks, cloud services, and smartphones have enabled the monitoring of large-scale dynamic environments and real-time data analysis. Crowdsourcing and participatory sensing in large-scale systems provide the capability to monitor in real-time the QoS as well as to estimate and predict the perceived QoE of services. Based on this information, personalized recommendations can be provided to customers for the appropriate services according to their profile. Moreover, smart data analytics can enable service providers to improve their services.

1.1 Motivation

We witness a paradigm shift: The traditional mobile network access paradigm is both static and fragmented: Users subscribe to or prepay for specific cellular operators/providers for long-term network access. On the contrary, the existing and emerging wireless access markets are of larger scale, with more users and more players in the game. There is an increase in the number of telecom operators and their services, as well as a dramatic growth of mobile data. Users are becoming more heterogeneous, with different demands and profiles. Customers have more options in terms of various telecom services and dataplans. These new markets are larger, more complex,and more dynamic. Network providers may form coalitions and various partnerships. Mobile Virtual Network Operators (MVNOs) provide services, often targeting specific populations. Moreover, the Body of European Regulators for Electronic Communications (BEREC) [12, 13] as well as the Federal Communications Commission (FCC) [14, 15] in recent reports envisage measures for *consumer empowerment*, boosting *consumer choice*, *network transparency* and developing mechanisms for *data collection and analysis*. By providing meaningful feedback to customers and providers, the QoE of services can be significantly improved.

These trends motivated our research in analyzing telecommunication markets.

Specifically, we have developed the u-map, a user-centric reviewing system that enables users to monitor various services (running on smartphones) and collects user opinion scores about their performance. The u-map informs users about the estimated performance of various services, and empowers them to *select the appropriate service provider/operator*, especially when visiting a new region. Providers can also access the u-map to obtain network measurements about their infrastructure and user feedback about their coverage/services performance to potentially improve/adjust their deployment and services. Thus, the u-map can act as an "early-warning" and churn-avoidance mechanism. It also allows them to gain a better knowledge about their customers and their QoE criteria and assess their strategic network planning, pricing decisions and service deployment strategies in a finer spatial granularity. The u-map can provide feedback about the "compatibility" of the device, OS, and service and their impact on the user experience, allowing the manufacturer, platform/OS developer, and service provider to early identify potential problems and constantly improve their products. Via u-map, regulators can detect whether operators comply to certain specifications. In that way, the u-map can become a powerful tool to users, providers, manufacturers, operators, and regulators, introducing a paradigm shift in wireless access markets. In this work, we have focused in two type of services, namely the Global System for Mobile Communications (GSM) and Voice over IP (VoIP) call service. The VoIP service is provided to the users through Sipdroid, an open-source Android application. Moreover, the team has developed CoRLAB [7, 8], a modeling framework and simulation platform to analyze the evolution of telecommunication markets and services under a diverse set of customer populations and network conditions. The u-map can feed CoRLAB with real-time or semi real-time data about users and QoE feedback about their service. With the u-map, we aim to assist users, providers and regulators to understand the emerging wireless markets and make more educated decisions.

1.2 Objectives

The objectives of this work are the following:

1. The design of the u-map and its implementation.
2. The performance analysis of the u-map, evaluating the following aspects:
   (a) The power consumption.
   (b) The responsiveness (e.g., the delay that a user experiences when querying the server).
   (c) The scalability of the u-map server, i.e., the delay as the traffic demand of the server increases.
   (d) The overall performance of the u-map, its Graphical User Interface (GUI), and its main functionalities, e.g., provider selection, improvement of QoE.

   (e) The dominant parameters that affect the user experience and future adoption of the u-map.
   (f) The main weaknesses of the u-map.

3. Knowledge about the user requirements and profile.

4. A generalization of the u-map, by extending its functionality in other domains.

1.3 Methodology

We first specified the system requirements and the desired properties of the system: A proper data representation scheme should be determined to ensure an efficient and scalable data management and access control. Furthermore, the set of supported services and queries should be designed in such a way, that the various entities (e.g., users, providers) can access the information they need in a reliable and efficient manner. The provision of data integrity and fault-tolerance, detection of erroneous data (e.g., injected by mis-configured or malicious users), support of non-repudiation, protection of user privacy, and appropriate access control are also critical for the adoption of the u-map system. Having specified the system requirements, we developed the u-map prototype.

The second step of our methodology regards the evaluation process. The evaluation was two-fold: a) we analyzed its systems performance by measuring the power consumption, response delay, and scalability, and b) we performed a subjective study to understand the perceived user satisfaction about the u-map with 21 participants. The testbeds, experiment methodology, data sources, data preprocessing and analysis methods of this step are thoroughly presented in Chapter 6. Finally, we applied the u-map in other domains, such as water distribution networks and medical applications, demonstrating its generality.

1.4 Challenges

The limited resources and capabilities, as well as the diversity of smartphones, the trade-off between thorough monitoring and user privacy pose various challenges in the development of an efficient implementation of the u-map client. The limited memory of the Android smartphones requires the development of low space complexity algorithms everywhere in the u-map client. For instance, the data uploading process uses a threaded implementation to stream data from the client's local database to the u-map server, allocating only a circular buffer in the smartphone memory. Limitations on WiFi and CPU performance confine the sampling rate for link-level WiFi measurements, as well as the frequency of location updates by indoor fingerprint-based positioning algorithms. Namely, the hardware of Google Nexus One takes 1200 ms to report the results of a scan, while other devices (e.g., Sony Ericsson Xperia X10 mini) take even longer. However, capturing the location where network measurements and user feedback are collected is crucial for the u-map service. To deal with this problem, only a simplified version of the percentiles [16] algorithm was used for indoor positioning.

The battery consumption is another challenge. Power-greedy applications can be a source of great inconvenience to users. The u-map requires access to some power consuming modules (e.g., GPS sensor, WiFi interface for scans). We implemented a *monitor rate selector* module, that controls the power consumption of the u-map.

Access to system-wide packet-level measurements (e.g., with tcpdump or Wireshark) would provide a detailed view of QoS and allow assessing the QoE about large number of applications. However, running such tools requires root privileges, which are by default restricted on Android. However, installing the "su" binary on an Android (i.e., rooting it) imposes various security risks and may void the manufacturer warranty. Conducting a user study would be a lot harder with an application requiring root privileges. Instead, we decided to follow a more secure approach, compatible with not-rooted Android devices: for every supported application, a receiver module is implemented in the u-map client, that receives transport- and application-level network measurements, as well as session information. The network measurements are obtained by a module that we develop an add on each supported application.

In general, Android devices vary with respect to their capabilities (e.g., installed sensors, network interfaces, dual-SIM card support), characteristics (e.g., screen resolution, CPU cores and frequency, RAM size), and Android Operating System (OS) version. The collection of measurements from different devices can be challenging. The debugging and testing requires running the application on different devices. Also, designing a GUI for a wide range of display sizes requires a lot of effort.

For the implementation of a testbed in a dynamic, loosely-controlled environment (in contrast with emulating a system with controlled in-lab experiments), the developer has to face various non-trivial aspects. The innate complexity of networks and services, the dynamic nature of the Internet, as well as the wide range of psychological, socio-economical, and contextual factors that influence the user satisfaction and the perception of QoE pose make the development of accurate models challenging. The diversity in the user profiles requires a large number of users to reach statistically significant conclusions. The protection of user privacy, and the design of an appropriate access control mechanism are also critical for the adoption of the u-map system.

1.5 Contributions

This work makes the following contributions:

1. The u-map system and its implementation. The u-map is a user-centric recommendation system that enables users to upload measurements that their

devices collect about network conditions, interference, and coverage as well as their feedback about their profile and QoE for certain types of services. The collected data are uploaded to the u-map server and stored in a geo-database. This database can be accessed by users to make an educated selection of the provider/service as well as by providers to obtain information about the customer experience and profiles, in order to improve/adjust their deployment, services, and prices.

2. The performance of the u-map. The outcome of the performance analysis is encouraging:

    (a) The power consumption of the u-map is relatively low. The u-map consumes less power than other popular applications, such as Skype, YouTube, and Sipdroid.

    (b) The responsiveness of the u-map is relatively low, too. For example, the median total delay a user experiences when querying the server is 187 ms.

    (c) To highlight the scalability of the u-map, we measured the delay as a function of the number of concurrent requests (i.e., queries, data uploads). The querying delay is small. Using the Hadoop Distributed Filesystem (HDFS) and the HBase could improve the scalability of the u-map.

    (d) To understand how users assess the u-map and its services specifically, we performed a 2-week user study with 21 real users. The mean overall score of the u-map is 3.58 out of 5, while the mean overall score for the GUI is 3.90 out of 5. Most of the participants indicated that the u-map can improve their experience, helps in the discovery of networks that offer improved QoS. During the study, 30% of users discovered the best available network, improving substantially their experience.

    (e) Another contribution of the user study is the determination of the factors that affect the use of the u-map. The three most dominant factors are the “Privacy”, “Improved QoE of VoIP”, and “Battery consumption”.

    (f) The network selection has been simplified to a one-click operation. The automatic network selection has also been implemented. New networks are automatically discovered and presented to the user. The implementation of a robust access control mechanism is an ongoing effort.

3. Knowledge about the user requirements and profile:

    (a) The user satisfaction is modeled as a utility function based on a set of parameters and their significance level. The proposed model accurately predicts the overall score.

    (b) The services of the u-map have been assessed with both objective measurements (i.e., power consumption, responsiveness, scalability), and subjective evaluations provided by a user study. The user study includes two questionnaires that capture the user profile and its expectation about the services, while the second evaluates various functionalities of the u-map and provides an overall score.

(c) Clustering algorithms can find user groups with similar profiles. Then, the correlation of profile characteristics, traffic demand, and evaluation of the u-map can provide an insight of user requirements.

4. Demonstration of the generality of the u-map, by extending its functionality in other domains, such as water distribution networks and medical applications.

1.6 Roadmap

In Chapter 2, we overview the related work. Chapter 3 introduces the tools that the team has developed for the analysis of large-scale wireless markets and the empowerment of users. Then, Chapter 4 focuses on the u-map, which is the main topic of this M.Sc. Chapter 5 describes the prototype implementation the u-map. Chapter 6 presents the evaluation of the prototype, while Chapter 7 describes the improvements that followed the evaluation, and the application of the u-map on other domains. Finally, Chapter 8 summarizes our conclusions and future work plan.

1.7 Related publications

The system, its implementation, and its evaluation have been presented in the following publications:

1. G. Fortetsanakis, M. Katsarakis, M. Plakia, N. Syntychakis, and M. Papadopouli, “Supporting wireless access markets with a user-centric QoE-based geo-database,” in *ACM MobiArch*, Istanbul, Turkey, 2012

2. M. Katsarakis, V. Theodosiadis, M. Dramitinos, and M. Papadopouli, “u-map: a user-centric QoE-based recommendation tool for wireless access markets,” in *ACM S3*, Miami, FL, USA, 2013

3. M. Katsarakis, G. Fortetsanakis, P. Charonyktakis, A. Kostopoulos, and M. Papadopouli, “On user-centric tools for QoE-based recommendation and real-time analysis of large-scale markets,” *IEEE Communications Magazine*, vol. 52, no. 9, pp. 37–43, 2014

4. M. Katsarakis, V. Theodosiadis, and M. Papadopouli, “Evaluation of a User-centric QoE-based Recommendation Tool for Wireless Access,” in *ACM SIGCOMM Workshop on Crowdsourcing and crowdsharing of Big (Internet) Data (C2B(I)D)*, London, UK, 2015

Additionally, the u-map has earned distinctions in contests organized by international conferences:

1. M. Katsarakis, G. Fortetsanakis, M. Viskadouros, N. Syntychakis, and P. Lionakis. "Who is afraid of active users? The u-map as a Catalyst of Wireless Access Markets", 2nd position at the *Student Video Contest, DySpan 2012*, Bellevue, Washington, USA, October 19, 2012.
   `http://dyspan2012.ieee-dyspan.org/video.html`

2. M. Katsarakis, V. Theodosiadis, M. Dramitinos, and M. Papadopouli. "u-map", finalist at *The MobiCom First Mobile App Competition, MobiCom 2013*, Miami, Florida, USA, October 2, 2013.
   `http://www.sigmobile.org/mobicom/2013/app_finalists.html`

The application of the u-map paradigm for the monitoring of *water distribution networks (WDNs)* and assessment of water quality is an ongoing work of the team. The following publication describes the system architecture and a preliminary analysis:

1. N. Rapousis, M. Katsarakis, and M. Papadopouli, "QoWater–A crowd-sourcing approach for assessing the water quality," in *Cyber-Physical Systems for Smart Water Networks (CySWater), 2015 ACM 1st International Workshop on.* ACM, 2015

Chapter 2

Related Work

> *"We are like dwarfs sitting on the shoulders of giants. We see more, and things that are more distant, than they did, not because our sight is superior or because we are taller than they, but because they raise us up, and by their great stature add to ours."*
>
> – John of Salisbury, *Metalogicon*

2.1 Databases for wireless network measurements

Databases have been proposed to maintain spectrum or physical-layer based information (e.g., [22, 23, 24, 25, 26, 27]), focusing on spectrum availability/usage, "whitespace", and interference. For example, [26] *et al.* presented a cooperative sensing algorithm that uses information about the network topology and signal propagation characteristics to detect attacks. Gurney *et al.* [27] focused on a geo-database for TV-band incumbents that allows dynamic updating techniques to correct interference problems. Karlsson *et al.* [28] proposed an open spectrum approach in which providers consult a database with user feedback to improve their network infrastructure and reduce interference.

2.2 Collecting wireless network measurements

2.2.1 Drive tests

Providing network coverage and quality of service (QoS) is an vital task for a network provider. To improve their networks, providers perform drive tests to collect radio measurements and discover potential problems (e.g., areas of bad coverage, and interference problems). The field measurement equipment usually consists of a laptop PC connected with a Global Positioning System (GPS) receiver, a spectrum analyzer, a User Equipment (UE) device (e.g., mobile phone), and possibly external antennas. These drive tests are typically of high cost and provide data only for a limited geographical area and time span. The approach of field measurements provides accurate data, but has several drawbacks: The high cost of drive tests usually limits the extent of measurements to specific locations and time periods. Also, the drive test teams typically cannot collect measurements from indoor areas. As the approach of drive tests is not the focus of this thesis, the interested reader is refered to [29, 30, 31] for further reading.

2.2.2 Crowdsourcing

Crowdsourcing is a technical paradigm and an emerging bussiness model in the Internet, according to which, a task (e.g., computation, problem-solving, environment sensing) is outsourced to a crowd of anonymous users. Employing the mobility of users and the capabilities of their sensor-enhanced smartphones to collect local data (e.g., sensor readings) and extract information is often referred to as participatory sensing [32] or crowdsensing [33] in the literature. The collected data is shared among the participating smartphones or uploaded to a central location (e.g., server, cloud infrastructure).

Some early crowdsourcing applications [34, 35] for mobile phones record a wide range of parameters and upload them on their server. Micro-Blog [34] is an application for mobile phones that allows users to record multimedia blogs. Blog entries can be enriched with sensor inputs (e.g., accelerometer data, health sensors, and WiFi SSIDs). The phone application associates a blog entry with the time and location of the device. The Micro-Blog entries are uploaded to a web-accessible database and presented on a spatial platform (e.g., Google Maps). Then, web services can access and mine the Micro-Blog data. MyExperience [35] passively logs device usage (e.g., communication and application usage, and media capture), user context (e.g., calendar appointments), and environmental sensing (e.g., Bluetooth and GPS). In addition, user experience sampling can be targeted to moments of interest by triggering off sensor readings. MyExperience can be used effectively to understand how people use and experience mobile technology.

Other systems [22, 36, 37] focus on collecting link-layer information from cellular and WiFi networks by crowdsourcing the "drive test" task to participating users and their smartphones. OpenSignal maps [22] collects Received Signal Strength Indication (RSSI) measurements from Android and iPhone smartphones to build cellular and WiFi coverage maps. It also provides a compass-like GUI that directs the user to nearby locations with higher signal strength. The RF Signal Tracker [36] is an engineering application for doing drive-tests with an Android smartphone. The user can monitor the cellular and WiFi RSSI as well as the serving cell locations and WiFi Access Points (APs), describe a cell site's zone of coverage, identify changes in technology and handover points, and save and playback that data. Farshad *et al.* [37] have used Android smartphones with IEEE 802.11a/b/g/n radios and the RF Signal Tracker [36] application, to perform an urban WiFi characterization study. During the study they recorded Service Set Identifier (SSID), Basic Service Set Identifier (BSSID), channel, and RSSI information, as well as the security scheme, by listening to AP beacons.

Some applications [38, 39, 40, 41, 42] also collect network- or transport-layer measurements. Sensorly [38] is another application for Android and iPhone smartphones that builds coverage maps from RSSI measurements. In addition, a Sensorly user can trigger a speed test that measures download and upload throughput, as well as Round-Trip Time (RTT) delay through active probing a server. Portolan [39] uses the smartphones of participating volunteers as mobile monitors to collect measures about the network that surrounds them: The Portolan server receives measurement campaign specifications from human users, and translates them into smaller jobs which will be executed by mobile devices. The tasks are executions of traceroute or bandwidth measurements to a specific target (i.e., destination IP address), or signal strength measurements. WiScape [40] is a framework for measuring and understanding the behavior of wide-area wireless networks (e.g., city-wide or nation-wide cellular data networks) using active participation from clients. In their approach, a centralized controller instructs clients to collect measurement samples over time and space in an opportunistic manner. The measured parameters of WiScape are the TCP and UDP throughput, UDP packet loss rate, jitter, and RTT. The implementation and testbed described in the paper [40] uses three client devices, namely three laptops with cellular data cards and GPS sensors. The MCNet [41] measures the WiFi performance using the crowdsourcing paradigm on Android devices. It passively measures RSSI and data rate from AP beacons, and conducts active throughput and latency measurements. The MCNet also uses contextual data (i.e., WiFi connection events, power state, accelerometer readings) to intelligently schedule measurements. The application supports three operating modes (i.e., inactive, active, and aggresive) with different sampling rates. WeFi [42] monitors RSSI throughput, bandwidth and signal strength of both cellular and WiFi networks. The WeFi client allows the automatic connection to the best available WiFi network in terms of throughput. Also, a WeFi user can tag its favorite WiFi networks and indicate their category (e.g., home, work, onboard or public Hotspot).

Unlike these approaches, the u-map integrates a richer set of data, encompassing cross-layer measurements (e.g., RSSI values to QoE scores) and user preferences (e.g., traffic demand, preferences about data rate and price).

2.3 Metrics

QoS metrics include the achievable data rates, throughput, end-to-end delay, packet loss, number of resource units (e.g., time slots in TDMA, frequency-time in OFDMA), blocking probability for calls, RSSI, SINR, and bit error rate. However, such metrics may not be capable of capturing adequately the human perception of the quality of various services. The inherent complexity of services and applications, the dynamic and transient nature of the Internet, the limited insight in the psychological and sociological factors pose many research challenges in modeling the perceived QoE. One well-known metric for the VoIP quality is the E-model [43]. Such user-centric QoE-driven spatio-temporal data repositories, as the proposed one, could assist in modeling the QoE of various services.

Chapter 3

Tools for Large-Scale Market Analysis

This chapter provides an overview of the tools that the team has developed for the empowerment of users with QoE recommendations and smart real-time analytics. It presents the u-map, a user-centric recommendation system that enables users to collect network measurements and subjective opinion scores about the performance of various services. It also presents the CoRLAB, a modular multi-layer framework for modeling and assessing various markets, services, and their evolution under a diverse set of customer populations and conditions. The u-map feeds CoRLAB with user measurements and feedback in (semi) real-time. The chapter discusses how the u-map and CoRLAB can been used to analyze telecommunication markets and services. It highlights the main research results, challenges, and potential research directions.

New paradigms in both the wholesale and retail service markets are being formed and accelerated by the technological advances (e.g., in networking, virtualization), booming of content delivery, and regulatory changes on access and competition rules. These paradigms can enrich the roles of service providers, differentiate traditional pricing schemes, and enable new business models. For example, cost reductions and higher-efficiency can be achieved through increased multiplexing due to the pooling of existing infrastructure, resource sharing, crowdsourcing, and partnerships among providers.

The advances in wireless and sensor networks, cloud services, and smartphones have enabled the monitoring of large-scale dynamic environments and real-time data analysis. Crowdsourcing and participatory sensing in large-scale systems provide the capability to monitor in real-time the QoS as well as to estimate and predict the perceived QoE of services. Based on this information, personalized recommendations can be provided to customers for the appropriate services according to their profile. Moreover, smart data analytics can enable service providers to improve their services. This chapter focuses on telecommunication markets, although these issues are also relevant in a plethora of other application domains and markets.

These trends motivated our research in analyzing telecommunication markets. Specifically, we have developed the u-map [17], a user-centric reviewing system that enables users to monitor various services (running on smartphones) and collects user opinion scores about their performance (Fig. 3.1). These data are uploaded on a geo-database in a crowdsourcing manner. Data analytics can be then applied for customer profiling, new service planning and testing, pricing, and market refinements. the u-map informs users about the estimated performance of various services. Providers can also access the u-map to obtain network measurements about their infrastructure and user feedback about their coverage/services performance to potentially improve/adjust their deployment and services. Thus, the u-map can act as an "early-warning" and churn-avoidance mechanism. It also allows them to gain a better knowledge about their customers and their QoE criteria and assess their strategic network planning, pricing decisions and service deployment strategies in a finer spatial granularity. the u-map can provide feedback about the "compatibility" of the device, OS, and service and their impact on the user experience, allowing the manufacturer, platform/OS developer, and service provider to early identify potential problems and constantly improve their products. Via u-map, regulators can detect whether operators comply to certain specifications. In that way, the u-map can become a powerful tool to users, providers, manufacturers, operators, and regulators, introducing a paradigm shift in wireless access markets. Moreover, the Mobile Computing Activity team has developed the CoRLAB [7, 8], a modeling framework and simulation platform to analyze the evolution of telecommunication markets and services under a diverse set of customer populations and network conditions. In general, the analysis of such markets either adopts agent-based simulations (e.g., [44]) or employs a macroscopic approach that models an average behavior of various entities (e.g., [45]). Moreover, such approaches rarely incorporate user-centric data or models in real time. Unlike them, the CoRLAB provides a modular *multi-layer* framework, including models at multiple levels of detail: microscopic models that describe each distinct entity (e.g., user) as well as mesoscopic or macroscopic models based on aggregations of entities (e.g., homogeneous user populations). Furthermore, the u-map feeds CoRLAB with real-time or semi real-time data about users and QoE feedback about their service.

3.1 u-map

The u-map [17] is a recommendation system running on smartphones that follows the client-to-server architecture. In a crowdsourcing manner, it collects various objective and subjective measurements that indicate the performance of a certain service. Regularly, these measurements are uploaded on a spatio-temporal geo-database and are processed by the u-map server. Based on this information, the u-map server can then provide user-centric QoE feedback and recommendations about the availability and performance of services in a region. In addition, CoRLAB

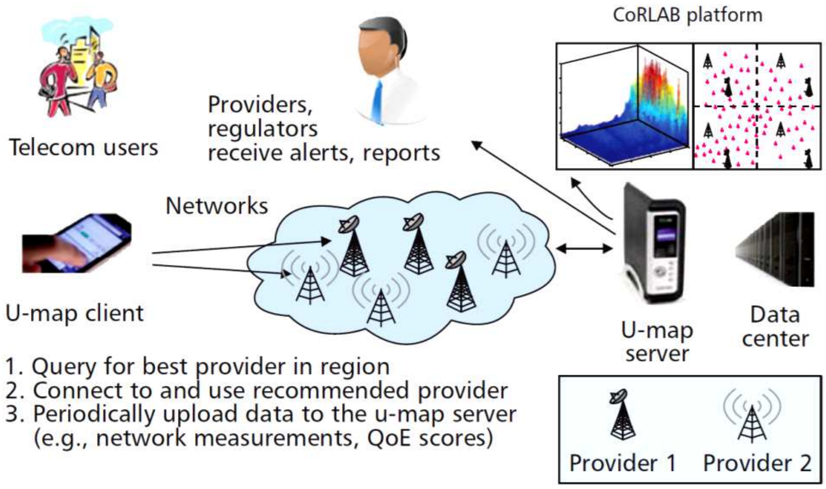

Figure 3.1: The u-map use case in telecommunication markets: users select the best provider/service, providers receive reports about potential network failures.

obtains these measurements from the u-map server (Fig. 3.3) and provides various geo-statistics about the evolution and performance of these services. To cope with large user populations, the u-map can be developed as a cloud service. The key modules of the u-map architecture includes the monitoring, data analytics, QoE modeling, and privacy and access control. The following paragraphs discuss them in more detail.

3.1.1 Monitoring

Monitors have been developed for user devices or network infrastructures. For example, monitoring tools may collect various network measurements (e.g., Portolan [39]) or physical-layer information to enable the detection of spectrum availability/usage, whitespace, and interference [27]. The reliability of the measurements and security are significant aspects that need to be addressed. The monitoring tools should not compromise the user privacy. Monitoring on mobile devices needs to also address the energy cost and dependencies on the hardware, while monitoring at the infrastructure, the deployment cost. In the context of the u-map, monitoring can be performed at the user mobile devices as well as at various gateways. The u-map data collection departs from state-of-the-art techniques in the following manner: unlike existing tools which focus on specific network measurements [39, 27], The u-map obtains and correlates a rich set of multi-sourced data with QoE feedback (opinion scores) and geographical information. The collected data can be active or passive cross-layer QoS measurements, user demand statistics, and user preferences. The measurements are uploaded to the u-map server data repository for analysis (e.g., user profiling, QoE inference).

Extensive monitoring and collection of data can improve the accuracy of the performance estimates but also increase the energy consumption and detection delay, as the network interfaces need to monitor the channel over longer time periods and then send this information to the u-map server. The u-map can employ advanced signal processing and data mining techniques to determine the appropriate spatio-temporal granularity in the sampling process and address the accuracy, privacy, and energy constraints tradeoffs. The identification of the appropriate parameters that need to be analyzed in order to characterize the specific condition of interest is an important first step in the monitoring process. This will determine at which layers and network points and at which spatio-temporal granularities the monitoring needs to be performed. Section 3.1.2 provides input on these aspects.

3.1.2 Data Analytics

The user profiling, clustering of user population, QoE modeling, geo-statistics, and warning generation are the primary objectives of the data analytics. Prior to data analysis, a sanitization treatment needs to be performed, since the collected measurements can be noisy, of high dimensionality, erroneous, and sparse. The data sanitization module detects erroneous entries, misconfigured/malicious data sources, and missing values. The statistical analysis method for treating these issues depends on the specific application characteristics, objectives, and requirements.

The profile of a user "integrates" a number of parameters associated with the user preferences, requirements, demand, constraints, and capabilities, with respect to services and context. For example, it may include information about the user's willingness-to-pay, data rate, traffic demand, QoE requirements, feedback consistency, and device characteristics. User profiling can be performed at different levels of detail, which may then impact the computational complexity of the follow-up market analysis. For example, an aggregate-level approach ignores specific individual user aspects and develops general, often less detailed, macroscopic models. On the other hand, a user-centric approach takes into consideration fine-level user information, aiming to form detailed profiles for individual users. A third approach employs clustering algorithms to determine homogeneous user populations and find representative user profiles for each population. Related to user profiling is the QoE modeling, which will be discussed in more detail in the next Section.

The spatio-temporal analysis focuses on the geographical and temporal distribution of specific features or metrics, identifying the locality and evolution/spreading of various phenomena. Examples include the analysis of user workload (e.g., amount of traffic, flow arrival process, interactivity model, application, and usage pattern). The u-map provides an early-warning mechanism that appropriately notifies users, providers and regulatory authorities about various failures or deficiencies, reducing maintenance costs (e.g., when certain conditions on QoE and traffic load hold). Understanding how user behavior and expectations change depending on the context (e.g., network topology, network conditions, device/technology characteristics, mobility, location, and environment) is challenging and the performance of empirical-based modeling studies is required.

3.1.3 QoE Modeling

Network benchmarks, such as jitter, latency, and packet loss, have been extensively used to quantify network performance. The evaluation of their impact on the user experience has received a lot of attention from the research community. Especially in the case of high-dimensional measurements, it becomes important to understand which network metrics have dominant impact on the performance of certain applications, distinguish the conditions that degrade substantially the performance of a given application, and investigate the predictability of these conditions.

The QoE is influenced by a diverse set of technical (e.g., QoS, and device features), socio-economic (e.g., social network, advertisements, and brand name), human-related (e.g., sentiment, and age), business (e.g., pricing, and willingness-to-pay), and contextual (e.g., time, and location) factors [46], some of which are difficult to capture and model. For example, important parameters for VoIP are the call setup time (time from the call initiation request to the beginning of the call), Mel-Frequency Cepstral Coefficients (MFCCs) for audio quality measurement, round-trip time (RTT) statistics, packet loss ratio and burstiness, jitter, and number of retransmissions. For video streaming, important parameters are the playback quality (the bit rate being delivered), startup delay (time between the user clicks on the play button to the time the video starts playing), buffering ratio (the percentage of the session duration spent in buffering state, and rate adaptation/temporal dynamics (e.g., change of the rate during the session), RTT statistics, packet loss ratio, jitter, and number of retransmissions. Parameters to be measured for web browsing include page load time (the difference between the time the URL is requested from the browser and the time when the objects are fetched) and lower-layer metrics such as, DNS lookup time, RTT, and TCP retransmission rates. To make the study amenable to theoretical analysis, QoE is usually expressed through simplified utility functions with nice mathematical properties that consider only a subset of these parameters. For instance, E-model [43] and PESQ [47] have been used for modeling the QoE in VoIP. Based on network measurements and subjective feedback collected via u-map, we aim to develop user-centric QoE models that can accurately reflect the user perspective for various services considering different techno-economical factors.

Classification and regression methods based on machine learning, data mining, and statistical modeling algorithms have been also employed for the prediction of QoE [48, 49]. The u-map applies a number of state-of-the-art machine learning algorithms, develops and trains models based on the collected network measurements and user feedback, and *dynamically* selects the best one, for predicting the QoE in a *user-centric* manner. A longer-term objective is the inference of QoE *without* necessarily the user intervention or feedback.

3.1.4 Privacy and Access Control

Anytime data is shared with a third party, there is a potential for abuse. Therefore, a great deal of effort has been made to design privacy protection techniques for publishing anonymized records. Even though there are clear benefits from sharing data, users have privacy concerns and value the benefits (provided by the service provider) differently. Some may prioritize anonymity, while others may be willing to share data unconditionally, with most users falling somewhere in between.

To protect the user privacy, the u-map requires authorization for granting access to the database. The client-to-database connection relies on end-to-end security that protects the integrity and confidentiality of the submitted data, by leveraging standard technologies (e.g., public-private key pairs, TLS). For further protection of sensitive information, like the user location, access is allowed only to aggregate statistics. Obfuscation approaches (e.g., spatial/temporal cloaking) could also protect user location privacy, in the cost of degrading the user experience. For example, if we assume a high level for user privacy, the responses of a location-based service would be inaccurate or untimely. Last but not least, the u-map provides a user-centric access control module that allows users to control the information revealed to third parties, through a fine-grained discretionary approach. More precisely, access control rules define *who* has access to *what* data. *Who* can be a user, a role (e.g., operator, application), and *what* is a query over the data. *Query rewriting* is then applied, "injecting" the access control rules in the request, so that the rewritten request filters out the inaccessible data.

3.1.5 Development

The main functionalities of the u-map have been developed and a pilot testbed has been deployed. Using this testbed, we performed a field study with real users and evaluated the impact of the u-map on the user experience in the context of a VoIP service [50]. Based on the collected data, we have also modeled the QoE for VoIP. We plan to extend this study to consider a larger user population.

3.2 CoRLAB

Unlike the traditional markets, emerging ones are larger (in the number of users and providers), more heterogeneous (in terms of services), and more complex and dynamic (e.g., in the interactions of providers and clients and their decision making). Modeling such markets is challenging due to a plethora of business, network, and service related phenomena that manifest at different spatio-temporal scales. Furthermore, the computational and scalability issues when analyzing such markets for long time periods or at a nation-wide level are prominent. The modeling approaches can be classified into two general categories, the microscopic- and macroscopic- level ones. Microscopic-level approaches model each entity, and its interactions with other participating entities, at a fine level of detail. However,

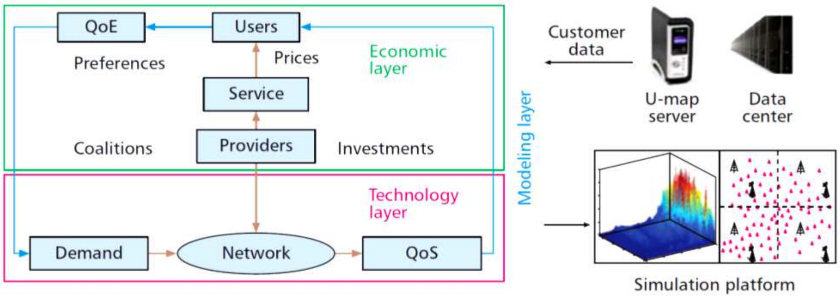

Figure 3.2: CoRLAB is integrated with the u-map. It can model and simulate markets using data provided by the u-map.

due to the high computational complexity, they typically assume a limited number of such entities [44, 51]. On the other hand, macroscopic-level approaches model the "average" behavior of certain types of entities (e.g., user population, service infrastructure) to make the analysis more tractable [45]. However, in many cases they result in inaccuracies and suboptimal performance.

In response to these challenges, CoRLAB, a modeling framework for large-scale diverse and dynamic markets has been developed. In contrast to the previous approaches that are either purely microscopic or macroscopic, CoRLAB is a complete multi-layer framework that allows the instantiation of a market at multiple levels of detail. At the microscopic level, the various entities are modeled in fine temporal and spatial detail, while at the macroscopic level, entities are described as a homogeneous population. Between these levels, various mesoscopic levels are defined (Fig. 3.2) in which entities are grouped in clusters. In a "coarse-graining" procedure that results to a loss of information in a controlled and hierarchical fashion, the individual entities of the microscopic level (e.g., users) are replaced by clusters with certain attributes, computed based on data mining algorithms. Instead of modeling the decision making of each distinct user, the mesoscopic levels consider a number of user clusters reducing significantly the computational complexity. Then, based on the requirements of a specific study, the appropriate mesoscopic level can be selected that achieves the desired tradeoff between accuracy and complexity (Fig. 3.2c).

The mathematical models in CoRLAB are selected in such a way that they can be studied analytically. However, there are some markets of interest that due to their inherent features, they are analytically intractable. In these cases, CoRLAB employs *empirical game theory*, a recent research direction that solves complex games (i.e., estimates Nash equilibriums) via simulations. The efficiency of this methodology strongly depends on the computational complexity and accuracy of the simulator and the selection of the appropriate mesoscopic level.

CoRLAB incorporates an economic and a technology layer (Fig. 3.3): at the economic layer, the decision making of providers and users is modeled using game

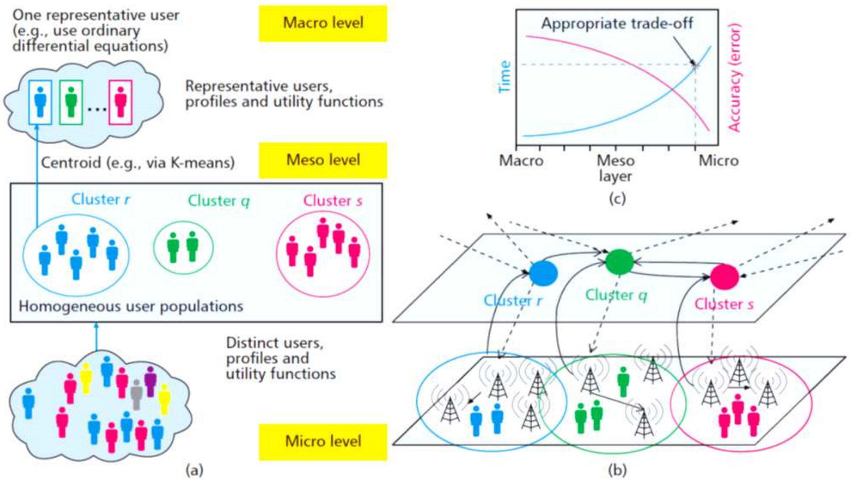

Figure 3.3: Multi-layer modeling framework developed in CoRLAB for the analysis of a wireless duopoly. (a) The clusters can be derived based on hierarchical clustering algorithms. (b) A decomposition method enables one to analyze a cluster microscopically and the other clusters macroscopically. (c) An example of the accuracy and scalability tradeoff. It does not correspond to a specific use case.

theory, while at the technology layer, the QoS is estimated using appropriate queuing theoretical models. The QoS metrics and the offered prices of providers are translated through appropriate utility functions to an expression of the user QoE. For the selection of the input parameters of the queuing theoretical models and the parameters of the user utility functions, the contribution of the u-map is important, since it reflects what happens in actual markets.

CoRLAB has been used to evaluate two scenarios, namely, a telecommunication market that offers the flex service and subscriptions, and a telecommunication market that supports mobile network virtualization. The next sections discuss the outcome of this analysis.

3.2.1 Flex Service

Traditionally, customers subscribe to specific providers and are served by accessing base stations (BSs) of the network of their provider. Inevitably, subscribers with relatively high usage pattern and data rate requirements are subsidized by the ones with lower usage and data rates. As the wireless network technology advances, a more diverse set of services is made available. To this end, we introduced the paradigm of a "flex user" who is not "locked" to a specific provider but can dynamically select base stations of different providers based on various criteria, such as network conditions, and offered prices [7, 8]. Flex users are flexible to select the appropriate provider even on a per-session basis. This flex service paradigm, which has been assumed as a typical access paradigm in wireless LANs, could be a new type of service offered in cellular markets. A similar concept is the "soft" (or virtual) SIM cards. The flex service could be also viewed in the context of a user with multiple SIM cards available. CoRLAB considers a diverse customer population with respect to their demand, their preference on data rate over price, their tolerance on the blocking probabilities of their sessions, and their willingness-to-pay for certain services. Users can dynamically decide to buy a long-term subscription or become flex users. We analyzed the evolution of a duopoly, focusing on whether the flex service can improve the QoS, social welfare, flexibility and further enhance the competition among providers. The analysis assumes two infrastructures of TDMA-based BSs and considers the perspective of clients, providers, and regulators.

This work demonstrates the benefits of markets with the flex service paradigm, compared to the ones with only subscription contracts. The merits are prominent also for specific user populations. A user can select the most suitable product that matches its profile, thus increasing participation in the market. In cases in which the user population is sensitive to blocking probability and data rate, the flex service can improve the revenue of providers. Furthermore, even under different utility functions and price-setting algorithms, several trends persist: for large user populations, the provider with most resources (i.e., TDMA slots) outperforms in terms of revenue and market share. The reverse trend holds for small user populations. This difference between providers is reduced in cost-driven markets compared to QoS-driven ones.

The presence of multiple services may also enhance the competition. Depending on user preference, the flex service may suppress or increase the prices (revenue). In some cases, a myopically greedy pricing results in suboptimal performance. In general, the flex service allows a finer market segmentation, offering more options to providers and users. Providers can also employ sophisticated pricing schemes that take into consideration the user demand to increase their revenue.

3.2.2 Pricing for Mobile Virtual Network Operators

Wireless communications change in a fast pace. However, not only the networks but also the business models of the communication services over these networks change. For example, the network virtualization allows network operators to lease parts of their infrastructure: Mobile Virtual Network Operators (MVNOs) do not own a network; instead they lease resources (e.g., access) in the wholesale market from a Mobile Network Operator (MNO) and resale access services in the retail market. Typically, MVNOs target specific user populations, services, and regions, such as low-willingness-to-pay users (e.g., with low charge-per-month offers), youngsters (e.g., with unlimited SMS, voice minutes over the same network), web users (e.g., with only data offers), ethnic (e.g., with special offers for calls to specific destinations) and roaming users (e.g., with calls, data). Instead of directly competing with MNOs, MVNOs often aim to widen and deepen the market through brand appeal, targeting of niche markets, and alternative distribution channels. The consumer-welfare impact of MVNOs is in offering extended and innovative services as opposed to lower prices. CoRLAB analyzed the pricing model of an MVNO in the retail market and the charging scheme of an MNO in the wholesale market. Furthermore, it assessed the impact of the knowledge about users on the profit of MVNOs and MNOs. Typically, operators collect such information from market surveys, while CoRLAB employed the u-map to obtain this information. The analysis has shown that, in the absence of capacity constraints, the optimal pricing strategies for both operators depend only on the distribution of users' willingness-to-pay [52]. The impact of the availability of information on the design of tariff plans was then analyzed via simulations.

3.3 Summary

This chapter provided an overview of the u-map and CoRLAB: The u-map is a system running on mobile devices that collects various objective and subjective measurements about the performance of certain services and provides recommendations to users. the u-map is integrated with CoRLAB, a multi-layer modeling framework and simulation platform for the analysis of large-scale markets.

CoRLAB has been employed to analyze the flex service in telecommunication markets [7, 8]. Moreover, we have studied pricing strategies in the context of MNOs and MVNOs and the impact of the u-map. Furthermore, we have performed a pilot study and simulation experiments to show the benefits of the u-map to wireless users. Users may obtain aggregate statistics about the QoE of various services within a specific region and can also query for the best provider according to their profile. We plan to extend the deployment in several cities to evaluate the benefits of the u-map.

The u-map can be beneficial to network/service operators by providing: (a) a better "picture" about their customers, profiles, and user satisfaction in a cost-effective manner (e.g., without questionnaires), (b) real-time geo-statistics and early-warning notifications about possible customer problems, (c) data to assess their coalitions and roaming agreements with other providers, (d) opportunities to extend/improve their services, or roll-out new ones in a cost-effective manner, and (e) the opportunity to act as "role model" and satisfy the transparency requirements. Regulators can also audit the providers' conformance to certain regulatory agreements. Via u-map, the market transparency is also promoted.

Crowdsourcing systems become useful only after achieving a critical mass of participating members. Sensing on smartphones consumes battery, which might discourage user participation in some scenarios. Similarly, the privacy requirements could also prevent users from participating. Part of our current research is the development of energy-efficient monitoring, effective incentive mechanisms, and business models that encourage participation. Moreover, the QoE estimation, with only limited user intervention and training is of interest.

Chapter 4

The u-map system

The u-map is a review and recommendation system that follows the client-server architecture (Fig. 4.1). A u-map client, running on mobile devices, stores locally cross-layer network measurements collected at the background during a session, such as a GSM or VoIP call, web browsing, game, or streaming of a video. Additionally, the u-map enables a user to indicate his/her QoE feedback at the end of the session. The u-map dynamically uploads these objective network measurements and subjective user feedback to the u-map server. The u-map server collects, processes, and stores these data in its spatio-temporal geo-database. In this work, we have focused in two type of services, namely the GSM and the VoIP call service. The VoIP service is provided to the users through *Sipdroid*, an open-source Android application.

Design objectives: The large size and heterogeneity of the collected measurements from different devices impose various challenges. A proper data representation scheme should be determined to ensure an efficient and scalable data management and access control. Furthermore, the set of supported services and queries should be designed in such a way, that the various entities (e.g., users, providers) can access the information they need in a reliable and efficient manner. The provision of data integrity and fault-tolerance, detection of erroneous data (e.g., injected by mis-configured or malicious users), support of non-repudiation, protection of user privacy, and appropriate access control are also critical for the adoption of the u-map system.

Communication, query, retrieval, searching: The u-map clients (or simply clients) may connect and exchange information with the u-map server. The u-map clients also record the status with which their sessions were terminated, namely whether these sessions were terminated successfully or abruptly or were blocked, along with the QoE scores which were provided by the user (to assess the services of those sessions). Monitors that run at the u-map client record network-related information, such as traffic demand, network data (e.g., interference, RSSI, packet loss) with timestamps. In addition, the position information of the device is recorded. All this information is stored locally in the client device. During a connection, a client may upload this information to the u-map server.

Access control: A user-centric access control will allow users to control the information revealed to third parties. When a user reports its experience, it provides possibly sensitive data. The u-map will define rules that determine "who" has access to "what" data. "Who" can be a user or a role (e.g., operator, application), and "what" is a query over the data.

Security and user privacy: The database should be protected from compromise through standardized widely-accepted techniques that prevent unauthorized access and intrusions. The protection of the collected data and customers can rely on strong authentication of the querying customers. The data stored at the u-map can be anonymized to not reveal private sensitive information.

In general, the u-map system can be developed and provided as a service by different entities, employing various business models. For example, a user community can offer this system as a review mechanism in an altruistic and "grass-root" manner to enable users to make a more "educated" selection of their service provider. A different approach is that of a third party that provides this service to registered users or network operators/service providers. Alternatively, a network operator/service provider may develop and support such a system and set of services. Each of these paradigms and underlying business models affect the design of the system and its architecture.

4.1 Architecture

4.1.1 u-map client

The u-map client (henceforth referred to as *client*) includes an *indoor localization* module, *GSM receiver*, *Sipdroid receiver*, *monitor*, *performance estimator*, *traffic demand estimator*, *auto network selector*, *database* (i.e., DB in Fig. 4.1), *back-end interface*, and its *GUI*. The following paragraphs describe these modules in more detail.

Indoor localization: This module provides functionality for accurate indoor localization, utilizing RSSI fingerprints collected from IEEE 802.11 APs. The indoor localization module employs statistical-based fingerprint methods. These methods divide the localization in two phases, namely, *training*, and *runtime*. An interesting area of the map (e.g., the interior of a building) is divided in cells that form a grid. During the training phase, measurements are collected at every cell of this area and fingerprints are generated that associate the corresponding cell of the physical space with statistical measurements based on signal-strength values acquired at those cells. The fingerprint of a cell is a vector of training signatures with size equal to the number of APs deployed in the area. Each entry of the vector corresponds to one AP. Then, during the runtime phase, a fingerprint generated on-the-fly at an unknown position is compared with all the training fingerprints. The fingerprint of the unknown position is the corresponding vector of the runtime signatures, generated using the same statistical method. The cell with a training fingerprint that has the smallest distance from the runtime fingerprint is reported as the estimated position. The indoor localization module provides a service similar to the *Location manager*[1] service that is provided by the Android application framework for accessing the GPS sensor. Both location providing modules are used by the monitor in order to achieve ubiquitous location awareness in both outdoor and indoor environments.

GSM receiver: This component listens for changes in the telephony state of the device. More specifically, the GSM receiver identifies when a GSM call is started and ended, and GSM link-layer information (e.g., RSSI, Cell ID, network operator) for GSM and mobile packet-switching technologies (e.g., GPRS, HSPA, and LTE). When a change is identified, appropriate events are sent to the monitor, traffic demand estimator, and performance estimator modules.

Sipdroid receiver: This component listens for messages from the Sipdroid application that also runs on the Android smartphone. The Sipdroid receiver identifies when a Sipdroid call is started and ended. Furthermore, it receives transport-level measurements (e.g., packet loss, delay, and jitter) collected directly from the RTP streams used for the VoIP call. The Sipdroid receiver sends messages to the monitor, the traffic demand estimator, and the performance estimator, to forward Sipdroid call events and RTP measurements.

Monitor: The monitor runs in the background and records cross-layer network-related information, such as network type, base station, and provider, as well as RSSI, interference, packet loss, delay, and jitter, together with timestamps and position information. Session data is also recorded. For each session, the start and end time, along with the termination status (i.e., terminated successfully or abruptly or was blocked) are recorded. Each session is associated with the corresponding network measurements collected during that session. The design of Fig. 4.1 focuses on two types of sessions: GSM and Sipdroid calls. By extending the architecture with additional receivers, other services (e.g., video streaming, and web browsing) can be supported, too. The location information is provided by the Location manager and the indoor localization modules, while the network- and session-related information are provided by the WiFi manager, GSM receiver, and Sipdroid receiver modules.

The monitor features a *Monitor rate selector* sub-module that utilizes information from the *Battery manager* to control the power consumption of the u-map. More specifically, the monitor rate selector confines the network measurement collection frequency using temporal and spatial constraints (e.g., next sampling of monitored parameters when 1 minute has elapsed and the device has moved 10 meters or more). These constraints are dependent on the battery status (i.e., battery level, USB or AC charging), session information (e.g., "in call" status), and a crude device speed estimation (i.e., stationary, moving with pedestrian or car speed). It also suspends the background processes of the u-map, when the battery level falls below a specific threshold. The monitor rate selector allows to gradually

[1] http://developer.android.com/reference/android/location/LocationManager.html

reduce the power consumption of the u-map, as the battery level decreases.

Performance estimator: The performance estimator receives events about sessions from the GSM receiver and the Sipdroid receiver. It assesses the services of those sessions. At the end of a session, the performance estimator launches a GUI form, prompting for user feedback. The user feedback consists of a QoE score, followed in some cases by an explanation (e.g., as in the Skype application). The user feedback is associated with the specific session and the network measurements collected during that session by the monitor module.

Traffic demand estimator: This module estimates the user's demand for various services. More specifically, the traffic demand estimator calculates: a) the number and duration of GSM calls, b) the number and duration of Sipdroid calls, c) the number of exchanged Short Message Service (SMS) messages, d) the number of Bytes transferred by the WiFi interface of the device, and e) number of Bytes transferred by the cellular interface of the device. GSM calls and SMS are distinguished to outgoing and incoming ones. Outgoing GSM calls, and outgoing SMS are further categorized regarding the provider of the called party/SMS recipient. For this, the traffic demand estimator resolves the provider of the called party/SMS recipient from its telephone number, by requesting *number-provider* pairs from the u-map server. Sipdroid calls, as well as WiFi and cellular Bytes are also distinguished in incoming and outgoing ones. The traffic demand history is illustrated through charts in the GUI. Also, statistics of the traffic demand history are used for defining the user profile. Also, estimating the traffic demand for the current time period (e.g., month) can be used for cost-aware service/network recommendations.

Database: The u-map client features an internal SQLite relational database. To facilitate database operations, the *OrmLite* Object-Relational Mapping (ORM) library is used. All data collected by the monitor, performance estimator, and traffic demand modules are stored temporarily at the database, before being uploaded to the u-map server.

Back-end interface The back-end interface handles the communication with the u-map server. It includes a secure HTTP client and the functionality that connects the client to the HTTP service using the JavaScript Object Notation (JSON) data-interchange format for communication. During a connection, a client may upload the recorded information to the u-map server.

Auto network selector: This module enables the u-map client to act as a software agent that makes network-selection decisions of behalf of its user. The auto network selector listens for application launch events, such as Sipdroid, or “Phone” (i.e., the default application of Android for GSM calls). When an event is received, the auto network selector automatically sends a query to the u-map server and connects the device to the recommended network. The query type can be specified by the user from the “Settings” menu of the GUI.

GUI: The GUI of the u-map client enables the user to form queries for network/service recommendations. The GUI also includes forms for requesting user feedback at the end of sessions, as well as charts that illustrate the traffic demand history.

4.1.2 u-map server

The u-map server (referred to simply as *server* herethereafter) includes the set of services and mechanisms to receive data from clients, process, store, and analyze it, as well as allow controlled access to it. The u-map server consists of the PHP application and the analytics module that the team has designed, as well as several dependent services (i.e., the HTTP and memcached services, the geo-database, the job scheduler, and the Certificate Authority (CA)). The PHP application is a content management system that consists of distinct components and together with the HTTP service enables the u-map server and client communication. The following paragraphs provide detailed descriptions for all components of the u-map server.

HTTP service: The u-map clients communicate with the u-map server using the request-response communication model and the JSON data-interchange format. Standard technologies (e.g., public-private key pairs, TLS) are employed to ensure identification, authentication, confidentiality, and integrity for the client-server communication. The HTTP service directs the clients' requests to the *PHP application*.

Client controller: The first component of the PHP application to be executed when a request arrives is the client controller. This component manages registration of new clients, authentication of existing clients, and maintaining client-related information. The registration process includes the signing of a client certificate by the *CA*, and the creation of a VoIP (e.g., Sipdroid) account by the VoIP server. If the client is authenticated, its request is directed to the proper component of the PHP application.

Query handler: This component answers queries from the u-map clients. A query has a spatial parameter which is either a polygon, or a place name (e.g., city name). Place names have to be resolved to the respective polygons. This is a two-stage procedure: a) The *Google geo-coding API* returns a geo-point that shows where the place is located, taking as input the place name. b) The *Wikimapia API* takes as input both the place name and the geo-point, and returns the polygon. Resolved place names and their polygons are stored in the geo-database. When the polygon is determined, it is passed in the SQL query and evaluated on the database. The evaluation is a table containing the networks/services that are available in this polygon area, together with a score. The score is dependent on the query type (e.g., based on QoE or on a specific QoS metric). Query evaluations are cached using the *Memcached service* and used to faster respond to future queries regarding the same area.

Data receiver: The data reveiver receives JSON-formatted datasets from the u-map client and performs the first pre-processing steps before storing them in the geo-database. These steps include a) verifying that the uploaded data have the proper data types and values within accepted value ranges, b) verifying that database schema constraints (e.g., not null, unique values) are fulfilled, and c) avoiding duplicate entries in the geo-database (e.g., provider names) and updating references appropriately.

Traffic demand: The traffic demand module responds to requests about the network provider that serves a specific telephone number during a specific time period. This information is retrieved from number portability databases (e.g., National Reference Database for Portability (NRDBP)) and cached in the geo-database.

Analytics: The analytics module implements Machine Learning (ML) and statistical analysis algorithms. The clustering of user population, and the QoE modeling are the primary objectives of this module. The MLQoE sub-module builds models for the prediction of QoE of a certain service given some network benchmarks. A special model is built for each client. The model can predict the QoE using network measurements stored in the database or submitted by the specific client. The query handler can respond to queries using these predictions.

Access Control: All database operations are passed through the access control module. Users can define access control rules. These rules define *who* has access to *what* data. *Who* can be a user, a role (e.g., operator, application), and *what* is a query over the data. *Query rewriting* is then applied, "injecting" the access control rules in the request, so that the rewritten request filters out the inaccessible data. To further protect sensitive information, access is allowed only to aggregate statistics.

Geo-database: The geo-database is managed by a Relational Database Management System (RDBMS) with geo-spatial extensions that provide the capability to form and efficiently evaluate spatial queries. All data of the u-map server (e.g., network measurements and feedback from the u-map clients, place names and their respective polygons, <telephone number, serving provider> pairs, and generated statistics) are stored in the geo-database.

Job scheduler: The job scheduler allows the PHP application and the Analytics modules to schedule jobs (e.g., query pre-caching by the query handler, or QoE model building by the MLQoE) for periodic execution (e.g., every night or every week). A job is registered to the job scheduler in the form of a shell command, together with an execution plan. Both the PHP application and the analytics module provide Command Line Interfaces (CLIs) that define commands and parameters for triggering certain jobs.

4.1.3 Sipdroid

We chose the Sipdroid application for the role of a VoIP client for Android devices. Two significant factors for our choice are that Sipdroid is free, and open-source. We can extend Sipdroid to communicate with the u-map client through Android Inter-Process Communication (IPC) and freely distribute our own Sipdroid edition. The users of the u-map and user study participants can download our Sipdroid edition without having to pay for it. A detailed description of the Sipdroid architecture is out of the scope of this thesis. Instead, only the components of interest, in the context of conducting network measurements and reporting them together with session events to the u-map, are described.

RTP implementation: Both the Session Initiation Protocol (SIP) and the Real-time Transport Protocol (RTP) protocols are entirely implemented in Java and included in the Sipdroid source code. SIP messages are exchanged between each Sipdroid instance and the VoIP server. On the contrary, the Sipdroid and the VoIP configuration is such that Sipdroid instances transport their RTP streams directly to each other, without the VoIP server being intervened. The *Receiver* and *Sender* sub-modules of the RTP implementation are instantiated during a Sipdroid call to receive and send RTP packets, respectively. We have developed the *Performance estimator* sub-module that records arrival and departure timestamps and sequence numbers of the transferred RTP packets. Sequence numbers are read from the RTP packet headers. The recorded data are used to calculate packet inter-arrivals, delays, jitter, and packet loss.

GUI: We have inserted code at specific locations in the source code where the GUI interacts with the SIP protocol, in order to detect session events (i.e., when a Sipdroid call is started and ended).

u-map sender: This module communicates with the u-map client through Android IPC. It sends RTP measurements and session information to the Sipdroid receiver of the u-map client. Additional events (e.g., when Sipdroid is launched) are sent, to allow for timely automatic network selection by the u-map.

4.1.4 VoIP server

We have designed a VoIP server architecture for switching VoIP calls among Sipdroid instances and managing their VoIP accounts. The VoIP server consists of the HTTP and Asetrisk services, as well as the Account manager and its database.

Asterisk service: The VoIP server uses the Asterisk service for switching VoIP calls among Sipdroid instances. The Asterisk service requires authentication and features a voicemail service for when a Sipdroid instance is offline. Each Sipdroid instance is assigned an ID (i.e., telephone number), together with authentication information. These data are stored in the database of the VoIP server.

HTTP service: The VoIP server includes an HTTP service for communicating with the u-map server. The HTTP service responds to requests for account management (e.g., creation, deletion) from the Client controller of the u-map server.

Account manager: The account manager is a PHP application offered as as a web service to the u-map server through the HTTP service. It connects to the database and implements all account management actions.

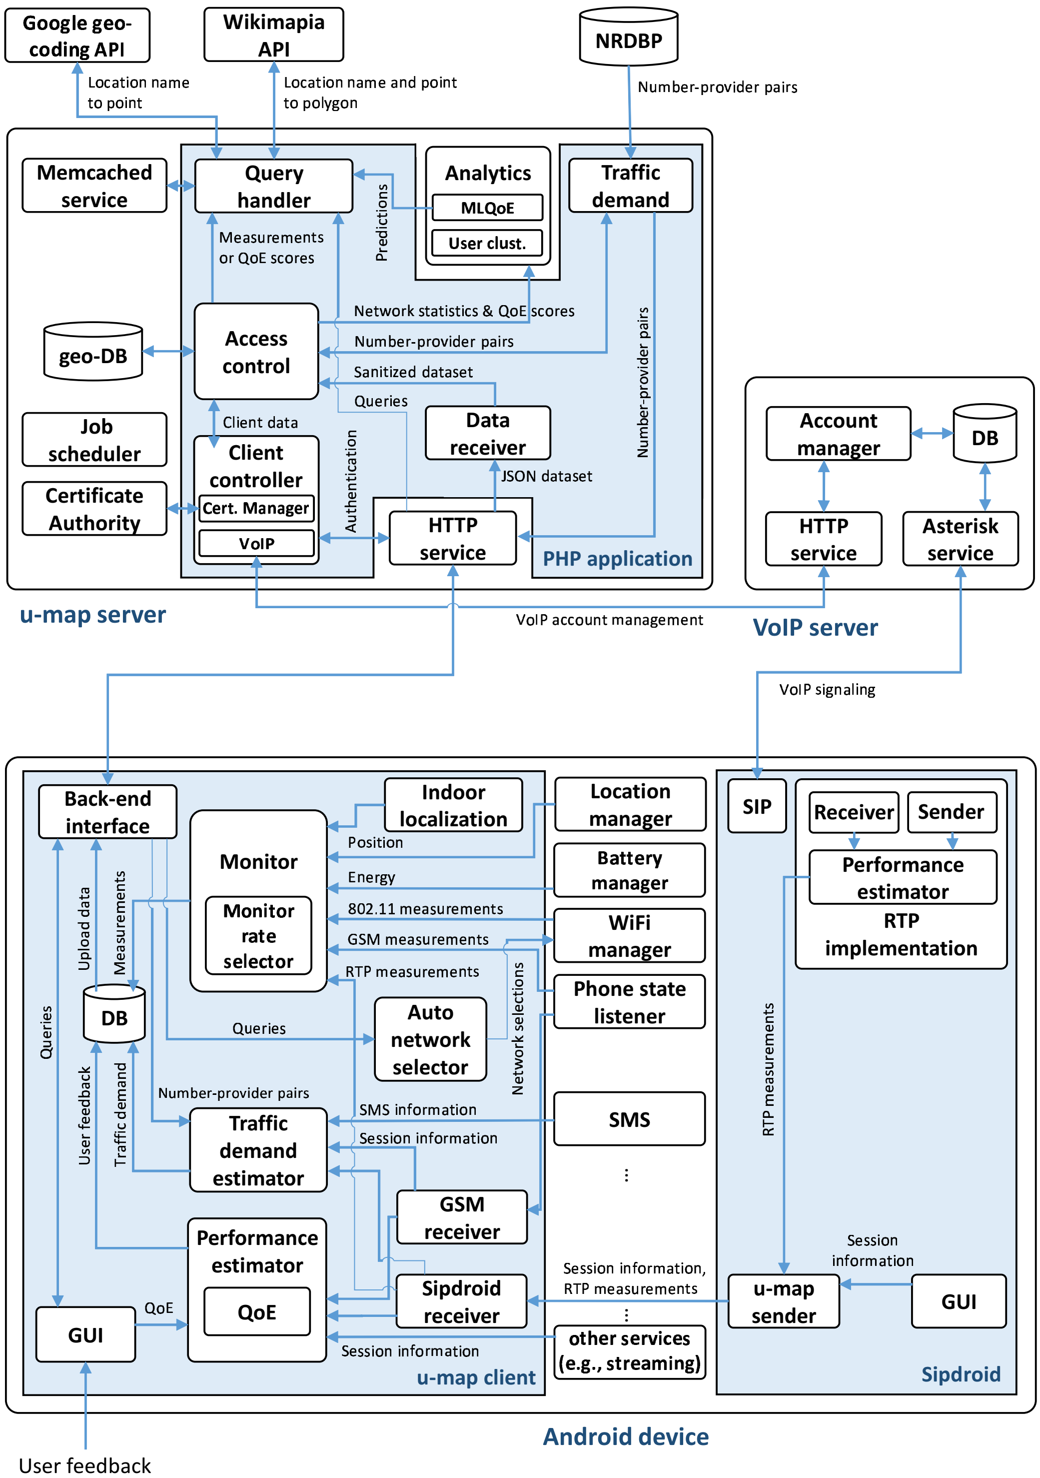

Figure 4.1: The architecture of the u-map system: The u-map client communicates with the u-map server to send queries and upload data. The Sipdroid communicates with the VoIP server using the SIP protocol. The u-map server communicates with the VoIP server for the management of the VoIP accounts, as well as with two APIs and a database for the retrieval of geographical and telephone number-provider information.

Chapter 5

Proof of concept

> *"Such a simple concept, yet so true: that which we manifest is before us; we are the creators of our own destiny. Be it through intention or ignorance, our successes and our failures have been brought on by none other than ourselves."*
>
> – Garth Stein, *The Art of Racing in the Rain*

Chapter 4 described the architecture of the u-map system. This Chapter presents the proof-of-concept version of the implementation. This version includes only a subset of described architecture. It was used in the evaluation of Chapter 6. After the evaluation, the prototype was refined with improvements and extensions. The improved implementation is described later, on Chapter 7.

5.1 Main modules of the architecture

The proof-of-concept implementation does not implement every component of the architecture presented in Fig. 4.1. More specifically, the u-map client does not implement the *traffic demand estimator*, the *auto network selector*, and the *monitor rate selector*. The u-map server does not implement the *access control*, the *analytics*, the *traffic demand* modules, nor the *VoIP* sub-module of the client controller. Also, the *query handler* does not support pre-cached queries, stored in the memcached service. Regarding the *performance estimator* module in Sipdroid, it measures packet loss, and jitter, but not the arrival and departure timestamps and sequence numbers of every transferred RTP packet. Also, the VoIP server consists solely from the Asterisk service. VoIP accounts are are defined manually in the Asterisk configuration files.

5.2 Prototype development

The u-map server consists of the *communication protocol*, *PHP application, database*, and *security and privacy* components. The database manages the content and serves queries, while the PHP application is a content management system that enables the u-map server and client communication. The registration and login processes are also implemented in the PHP application using the certificate authority to sign and validate client certificates. The u-map client includes a *monitor*, *database*, *performance estimator*, its *back-end interface*, and a *GUI*. The back-end interface includes a secure HTTP client and the functionality that connects the client to the HTTP server using the JSON data-interchange format for communication. The u-map server runs on a Linux machine and uses Apache HTTP server. The u-map client runs on the Android OS. Fig. 5.1 shows screenshots of the GUI of a u-map client.

To store a large size of data in a structured and easily manageable manner, we used relational database schemes at both the u-map client and server. The server employs the PostgreSQL RDBMS with the PostGIS extension, which allows the formation of spatial queries. SQLite is the only *database management system* supported by the Android OS. The client internal database stores *temporarily* the u-map data using the OrmLite, an object-relational mapping package.

Fig. 5.2 shows a representation of the ER model of the u-map server database. This ER model contains various entities, namely the clients, sessions, traces, and network providers. Sessions can be distinguished in two categories, namely the GSM and Sipdroid calls. For each session, the start and end time, the status, and the QoE score that is provided by the user are recorded. Each trace contains network data (e.g., RSSI values) that are measured during a session as well as the position and time at which these data have been collected. Finally, each trace indicates the network provider which the client uses to connect. The ER representation in Fig. 5.2 is not complete. For the sake of simplicity, only the most important attributes are presented. Several entities, such as, the service type and charging scheme, have been omitted.

Figure 5.1: GUI screenshots from a client: procedure for (a, b, c) querying the u-map server about the best network in a user-defined area, (d) connecting to the suggested network, (e, f) making a VoIP call using Sipdroid, and (g, h) evaluating the QoE of the call, when it is ended.

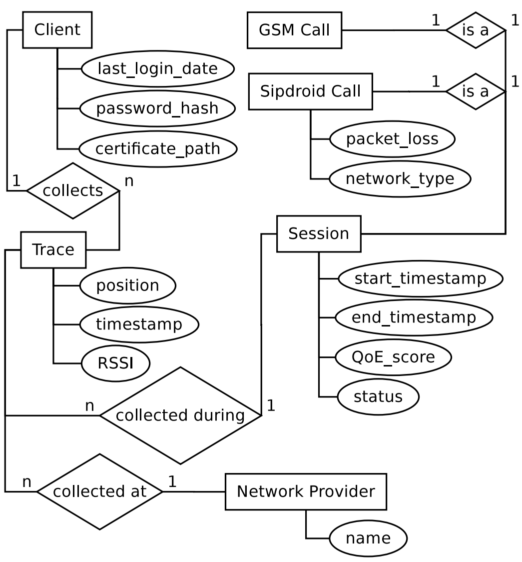

Figure 5.2: The ER model of the u-map server database.

Chapter 6

Evaluation

> *"For this purpose I determined to keep an account of the voyage, and to write down punctually every thing we performed or saw from day to day, as will hereafter appear."*
>
> – Columbus, Christopher, and Bartolomé de las Casas, *Personal Narrative of the First Voyage of Columbus to America: From a Manuscript Recently Discovered in Spain*

This section evaluates the u-map, its systems aspects and usability, and its impact on improving the QoE of VoIP calls.

The evaluation involves measurements of the power consumption, responsiveness, and scalability of the u-map. For the assessment of the usability, functionality and services of the u-map, a *field study* with real users was conducted. We use the terms subjects, participants, or volunteers, to refer to these users interchangeably throughout the text. During the field study, *objective measurements* and *subjective scores* were collected. Specifically, via u-map clients, the *subjects evaluated the QoE of their VoIP calls*, and these QoE scores, in conjunction with network measurements, were uploaded to the server. Furthermore, at the start and end of the study, the participants provided feedback by answering *two questionnaires* (available online at [53]). The network measurements, QoE scores, and questionnaire responses were analyzed using various statistical analysis and machine learning algorithms.

6.1 Systems performance

6.1.1 Power consumption

The power consumption of the u-map, and its main operations, was measured and compared with the power consumption of some popular applications. The estimation of the power consumption on smartphones has received considerable attention in the literature. Some approaches, such as [54, 55], use external monitoring tools (e.g., multimeters) to measure the power consumption. However, these approaches are highly intrusive, since they often require to dismantle the Android device in order to connect the probes of the monitoring tool at the hardware components of interest. Other approaches, such as [56, 57, 58, 59], perform the power measurements internally, using existing functionalities of the hardware of the Android device. For example, Dong *et al.* [56] estimate the overall power consumption of the Android, by automatically modeling the power consumption through system statistics. PowerTutor [57, 60] accesses usage statistics from *procfs* and *BatteryStat* to provide application-specific energy information, but this approach has several limitations affecting the accuracy. PowerProf [58] focuses on power modeling rather than application energy metering. In this work, we employed AppScope [59], a non-intrusive, online power analysis tool that estimates the energy consumption on a per-process and per-hardware component basis.

For the power consumption measurements, we set up a Desktop computer and an Android smartphone, connected through a USB cable. The smartphone runs the applications that will be evaluated, and a dynamic Android kernel module, which estimates online the power consumption. For each application, the power consumption of the CPU, OLED display, GPS receiver, and WiFi interface of the Android is measured. The kernel module sends the measurements via the USB cable to the Desktop computer, where the analysis is performed.

Using the above testbed, we performed a series of experiments. We defined several scenarios in which we measured the power consumption of various operations of the u-map, namely its data uploading, querying, and VoIP call. Then, the power consumption of popular applications, such as YouTube, Skype, and Sipdroid was measured and compared with that of the u-map. Each scenario was executed 100 times. To achieve consistency and synchronization, we automated the scenario execution and measurement process, using the testing program "UI/Application Exerciser Monkey" [61]. This testing program emulates the "tapping" of a user on the screen given specific coordinates. During an execution, the power consumption of each hardware component, namely CPU, display, GPS, and WiFi of the tested application was measured. The energy consumption was sampled with a rate of 1 sample/second.

Table 6.1 reports the median power consumption during the 100 repetitions of each scenario (in mW). The largest power consumption of the u-map is observed during a VoIP call (893 mW). The display is the main source for this consumption. Although the u-map runs in the background, the display is active due to the Sipdroid application through which the VoIP call is performed. Therefore, the consumption due to the u-map in this scenario is approximately 333 mW. Similarly, when uploading, the power consumption of the display is the main source of the overall consumption (567 out of 675 mW). However, when a u-map client uploads data, it is not necessary for the display to be active, thus, the power consumption due to the u-map can be reduced (to approximately 108 mW).

Even during querying, which is the most expensive operation in terms of power consumption, the u-map consumes less power than other popular applications, such as Skype (804 mW), YouTube (816 mW), and Sipdroid (780 mW). GPS is another power greedy component. When querying and during a VoIP call, the consumption due to GPS is relatively high (177 and 313 mW, respectively). Thus, a less frequent collection of position information can further improve the energy savings. Interestingly, the consumed power due to WiFi is relatively low (11.6 % of the total consumed power in the uploading scenario).

6.1.2 Responsiveness

We measured the delay a user experiences when uploading traces or querying. For that, we set up a testbed (Fig. 6.1) that consists of a client, running on an Android and communicating with the u-map service via a wireless AP. An external monitor that runs on a laptop captures all data exchanged between the client and AP. Using this testbed, we performed two experiments: the u-map client a) performs 50 queries, and b) uploads 50 datasets of 500 measurements. We defined a sequence of events (Table 6.2) and the following types of delay: a) *Server delay*: the total time elapsed between the reception of a client request by the server and the transmission of the response (i.e., $T_4 - T_3$). The distance estimation in the geospatial queries and information retrieval are the dominant causes of this delay. b) *Network delay*: it consists of the time required for the client request to reach the server and the corresponding response from the server to reach the client (i.e., $T_3 - T_2 + T_5 - T_4$). Note that the timestamps T_2, and T_5 are recorded at the external monitor, while the T_3, and T_4 are recorded at the server. Possible asynchronicities between the clocks of the two devices do not affect the measured network delay. c) *Android delay*: it consists of the time for the client request generation and display of the server response on the screen (i.e., $T_2 - T_1 + T_6 - T_5$). d) *Total delay*: it is defined as the sum of the Android, network and server delays (i.e., $T_6 - T_1$) (Figs. 6.2 (a) & (b)). The sequence of events and definition of the various delays are shown in Table 6.2 and Fig. 6.1.

In the case of sending queries, the median Android, network, server, and total delays are 43 ms, 8 ms, 135 ms, and 187 ms, while on uploading traces they are 49

Table 6.1: Median power consumption (mW).

Scenario	**CPU**	**Display**	**GPS**	**WiFi**	**Total**
u-map: uploading	30	567	0	78	675
u-map: querying	33	563	177	6	779
u-map: VoIP call	20	553	313	7	893
YouTube	23	787	0	6	816
Skype	111	639	0	56	806
Sipdroid	33	657	0	90	780

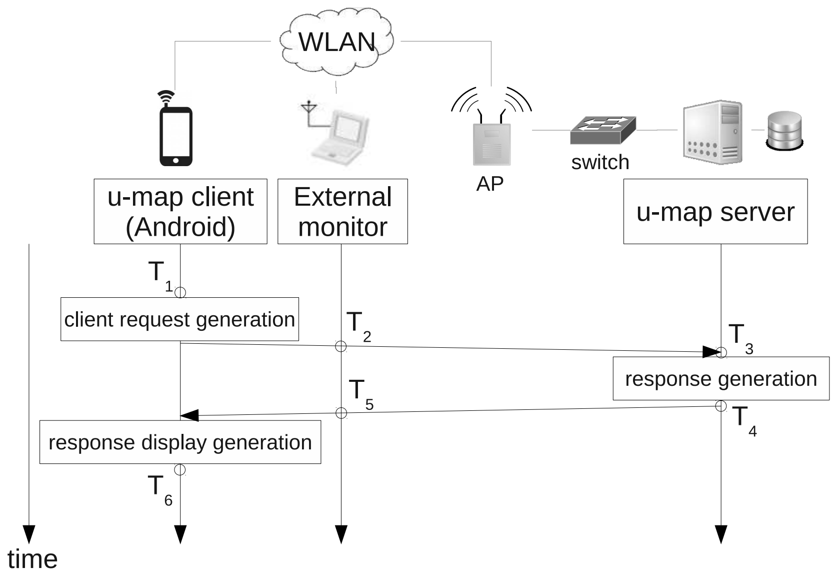

Figure 6.1: The u-map testbed for the delay measurements.

ms, 397 ms, 159 ms, and 604 ms, respectively.

6.1.3 Scalability

To highlight the scalability of the u-map, we measured the server and total delay as a function of the number of concurrent requests. For the experiment, a desktop PC emulated a number of the u-map clients by sending requests in parallel, in a range of 5 to 50 (step of 5). Each experiment was repeated 20 times. The timestamps T_1 and T_6 were recorded at the desktop PC. Fig. 6.3 shows the server and total delay. Each point corresponds to the mean value across all emulated clients and runs of the experiment. The delay increases when the number of concurrent requests becomes 40 or greater, due to the substantial increase in the memory required for processing. The querying delay is small. A prominent increase of the delay occurs, when uploading traces, especially at a large number of concurrent users (Fig. 6.3) due to the large amount of traffic exchanged between a client and the

Table 6.2: Events for the delay measurements.

Time	Event	Recorded @	By application
T_1	Client triggered to query/upload traces	client	System.currentTimeMillis
T_2	Client sends query/traces	monitor	Tcpdump
T_3	Server receives the query/traces	Server	PHP microtime
T_4	Server responds to the client	Server	PHP microtime
T_5	Client receives the response	monitor	Tcpdump
T_6	Client displays the result	client	System.currentTimeMillis

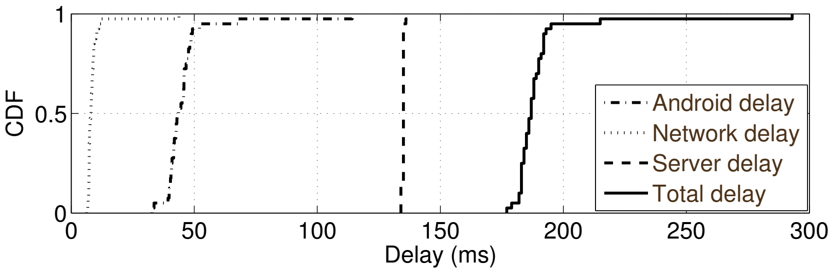

(a)

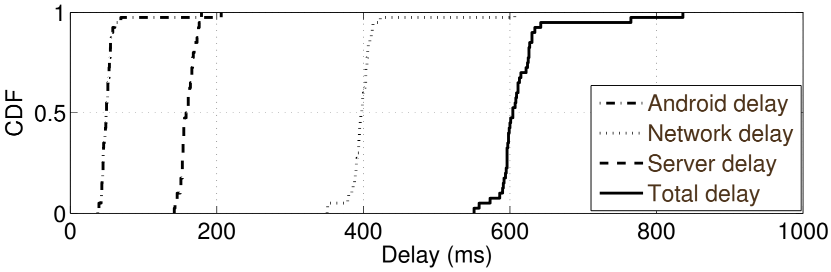

(b)

Figure 6.2: Delay for sending queries (a) and uploading traces (b).

server. However, the uploading runs in the background and does not impact the perceived user experience.

To improve the performance and achieve a higher aggregate I/O throughput, we are currently exploring the use of the Hadoop Distributed Filesystem (HDFS) and the HBase. The impact of delay, as perceived by the users, was also assessed in the context of the subjective field study, described in the following subsection.

6.2 User study

6.2.1 Methodology

To evaluate the user experience when running the u-map, a field study was conducted at the premises of FORTH, for a two-week period, with the participation of 21 volunteers. Common practices and guidelines for performing subjective studies [62, 63] were followed. The subject population was composed by graduate students and junior researchers from FORTH, excluding the u-map development team.

The study encompassed the following four phases: (i) A 30-minute tutorial was given on the first day of the study by one of the team developers. The tutorial introduced the u-map, its functionality, and the client GUI. The users were instructed to perform *at least 3 VoIP calls per day*, while at FORTH. For that, they had to use the Sipdroid VoIP client with the u-map and assign their QoE feedback (score) via the u-map GUI. The score is an integer ranging from 1 to 5 (1 corresponds to a poor performance, while 5 the excellent one). (ii) At the end of the

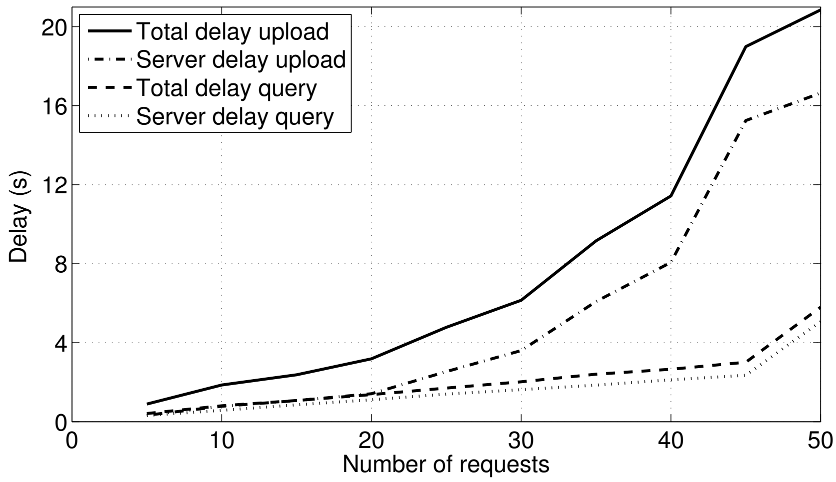

Figure 6.3: Scalability of the server.

first phase, the subjects also answered the *first questionnaire*. This questionnaire focused on the user background and profile (e.g., gender, familiarity with various technologies, Android usage), preferences (e.g., QoE-vs-cost trade-off), traffic demand (e.g., GSM and VoIP call frequency), and expectations about the u-map. (iii) For a *2-week period*, the participants used the u-map. (iv) At the end of this two-week period, the subjects answered a second questionnaire to assess the u-map. They evaluated various functionalities (e.g., GUI, ability to discover new networks and improve the overal QoE) and also provided an overall score. Moreover, participants indicated the significant factors that affect the performance and future adoption of the u-map. Feedback for improving the u-map was also collected. During the phases (iii) & (iv), there was no interaction between the subjects and the u-map developers.

The study was anonymous. Each participant was identified by a unique id generated by the u-map client application. The submitted questionnaires included this user id.

6.2.2 Testbed of the user study

The testbed includes a server, 21 clients, five IEEE 802.11 networks, and a monitor. The server runs on a Linux machine (CPU: 2.66 GHz Core 2 Duo, RAM: 2048 MB) located at FORTH. Eleven users installed the u-map client on their own Androids. The remaining ten users were provided with HTC Nexus One devices (part of our laboratory equipment).

IEEE 802.11 networks: Four WiFi networks have been deployed at FORTH, namely the *forth public access*, *forth authorized access*, *eduroam*, and *netlab-QoS-test* (experimental single-AP network). One objective is the evaluation of the functionality of the u-map to discover a new network that potentially can improve the QoE of VoIP services. For this, a *new* WiFi network was deployed with SSID *PublicWiFi!*. This network was configured to offer an improved VoIP quality compared to the other networks. It used a MAC address filter to allow only the devices of the participants to connect to its APs. The *PublicWiFi!* had the same number of APs and covered the same area as the *forth public access*. To avoid introducing any bias, we did not indicate/mention the name or the presence of any network to the participants. The monitor was running on a Desktop computer and captured the network traffic of all APs. The collected data, along with the answers to the questionnaires, were analyzed and cross-validated.

6.2.3 Data sources

The field study produced a rich dataset that includes (a) network measurements collected by the monitor and the u-map, (b) users' QoE feedback about their calls (collected by the u-map), and (c) subjects' questionnaire responses. A *variable* corresponded to a question (in the two questionnaires) and was *numerical* or *categorical* (with no numerical ordering). Most of the numerical variables measured the perception of subjects (about the quality of a certain aspect of the u-map) using the Likert scale from 1 to 5, ranging from "No" to "Yes". They were encoded as discrete numerical variables. Few numerical variables corresponded to questions about the frequency of performing an activity. Other multiple-choice questions were encoded as categorical variables. Finally, three questions required text: two questions required the names of the networks the subjects used, while the other one the characterization of the user familiarity with the WiFi configuration.

6.2.4 Data pre-processing

The questionnaires were submitted in paper form. Prior to the analysis, we sanitized/pre-processed (when necessary) and encoded the questionnaire answers based on their type (e.g., numerical, categorical, binary, string) and domain (e.g., range). An example of pre-processing is for the case of a question about the significance of various factors for the use of the u-map, where subjects are asked to assign a weight to each factor, with the weights summing up to 100. We normalized these weights when the sum was a different number. In another question where users had to provide the WiFi network names in a "free"-text form, only obvious spelling mistakes were corrected. Regarding the data collected by the u-map, call and network statistics were retrieved from the geo-database with the proper SQL queries, and saved in files for later analysis. The collected data were also visually inspected.

Missing values: Often data collection via questionnaires suffers from *missing values*. In general, the methodology for treating the missing values depends on the specific objective and may impact the results by introducing bias, leading to erroneous results. For example, in some cases ignoring the missing values or substituting them with a mean of the corresponding variable could be a valid approach.

The *k-nearest neighbors imputation* considers the statistical dependencies between variables to predict the missing values and is used in modeling the user satisfaction (Section 6.2.11). However, the k-nearest neighbors imputation method is not suitable when categorical variables exist in the data. So, in the case of user profiling using clustering algorithms (where categorical variables are present, Section 6.2.10), the missing values were substituted with the median of the non-missing values of the respective variable. The treatment of the "I don't know" or "I didn't understand the question" answers is a related issue; for the general statistics of our subject population (Section 6.2.6) these answers are considered valid, while missing values are ignored.

6.2.5 Statistical tests about the independence of various factors

To assess the statistical (in)dependence between various parameters, we employed three statistical tests, namely the χ^2 test of independence, the Spearman's rank correlation coefficient (for ordinal variables), and a G-test, based on a standard Monte-Carlo permutation procedure [64]. An advantage of this G-test is its robustness for small sample sizes.

6.2.6 General statistics about the user population

To understand our user population, we computed some general statistics based on their feedback from the first questionnaire (Fig. 6.4). Most of the subjects have a substantial level of familiarity with Android or are experts (71%). Similarly, a large percentage of users can configure their wireless interface without facing problems (48 %). In the first questionnaire, the subjects indicated their most frequent activities using smartphones. Most of them send SMS, access local data, and read e-mails (Fig. 6.5). A large percentage of participants browse the Web, watch videos, or play on-line games, while a relatively small percentage performs commercial product/service surveys or online booking and shopping.

6.2.7 Significant factors for the use of the u-map

Subjects indicated the important factors that affect the use of the u-map (or a related system) with their answers in the two questionnaires. The questions were phrased in a slightly different manner. Specifically, in the first questionnaire, subjects indicated the factors, specifying also their significance level (weight). The top-3 factors were: "Privacy", "Improved QoE of VoIP", and "Easy & friendly GUI" (Fig. 6.6 (a)). In the second questionnaire, the subjects were asked to classify these factors in a scale from 1 to 5, where 1 means "not significant" and 5 means "very significant". There, the top-3 factors were: "Privacy", "Improved QoE of VoIP", and "Battery consumption" (Fig. 6.6 (b)). Interestingly, the subjects had changed their opinion about the significance level of these factors. Specifically, the "Battery consumption" became an even more significant factor, reaching the third

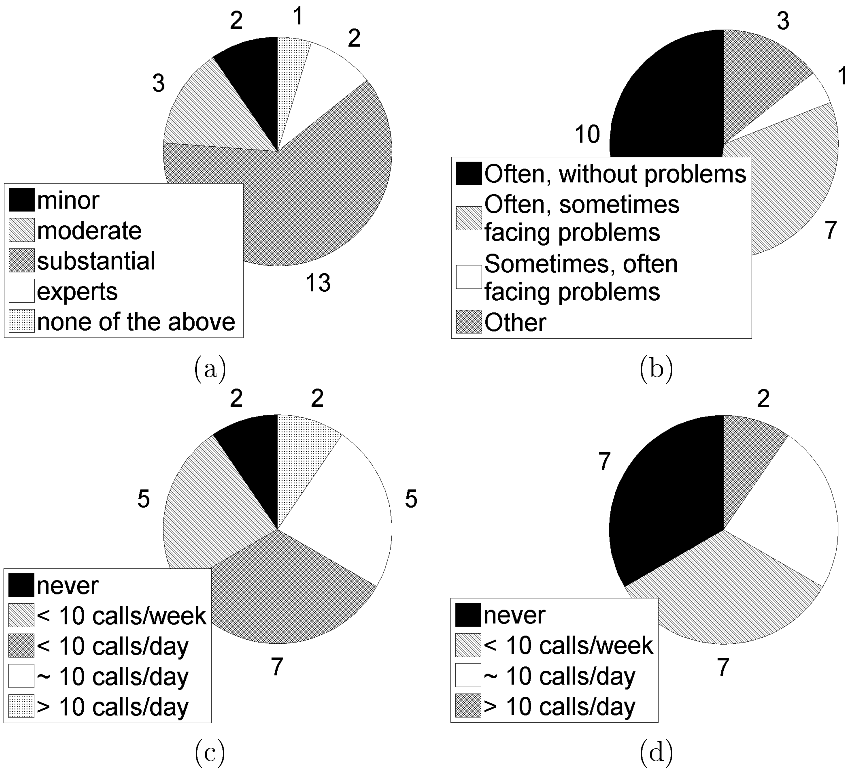

Figure 6.4: User familiarity with (a) Android & (b) WiFi configuration and frequency of making (c) GSM & (d) VoIP calls from mobile devices.

position, while the "Easy & friendly GUI" was downgraded to the fifth position. The "Privacy" steadily remained the most significant factor.

The "Low search time" was not one of the important factors (7th & 11th position, in Figs. 6.6 (a) & (b), respectively). During the field study, the load of the server was low, resulting in small response delays, which did not affect the user experience. This is also consistent with the fact that users evaluated the responsiveness of the u-map positively, regardless of their *responsiveness tolerance threshold*. Subjects were asked to indicate their delay/responsiveness tolerance threshold when using any Internet retrieving-based service/application at the first questionnaire. We assessed the statistical (in)dependence between the responsiveness of the u-map and the overall responsiveness threshold. The χ^2, the Spearman's correlation, and the permutation-based test report p-values of 0.23, 0.11, and 0.25, respectively, indicating a statistical independence.

Interestingly while many participants indicated that they did notice an increased battery consumption (question Q2.18 in Appendix B), it seems that this did not impact negatively their overall evaluation. To statistically assess this correlation, we performed the above tests. With p-values of 0.55 (χ^2 test), 0.33 (Spearman's correlation), or 0.51 (permutation-based test), we cannot reject the null

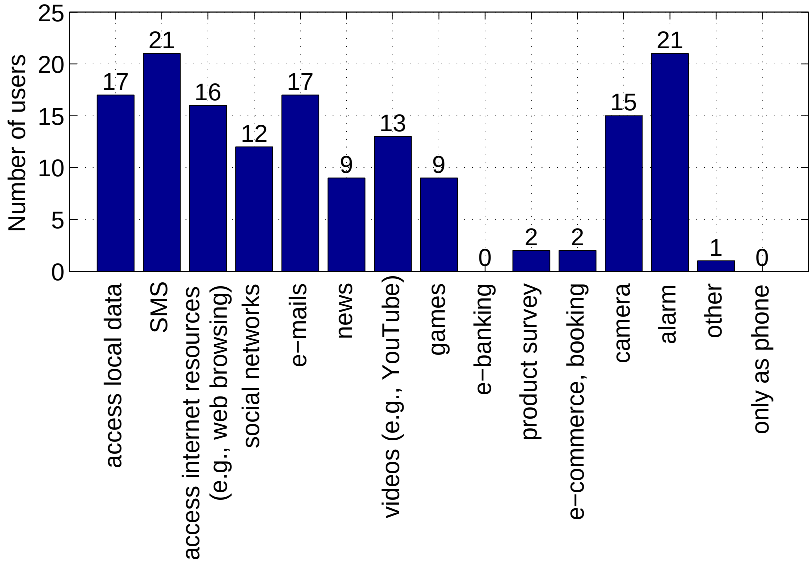

Figure 6.5: Frequent activities using smartphones.

hypothesis about the independence of the observed power consumption and overall score of the u-map. More specifically, there is evidence that the two parameters are independent. We assume that due to the small size, we cannot conclusively verify it. We speculate that most of the participants were able to alleviate the negative effect of the increased power consumption by performing some sort of "power management" (e.g., deactivating the wireless network interface). Unfortunately we do not have any measurements to analyze such actions. Note that a typical user may not be able to perform such a task, an thus, an automated dynamic power management becomes important.

6.2.8 On discovering new networks & improving QoE

The study showed that the u-map can assist in the discovery of new networks! When a new network becomes available in a region, a u-map user can eventually discover it and connect to it. The discovery of this new network was manual in our study but in a more general setting, it could take place through announcements or friends. The measurements and QoE feedback collected from this (first) user were uploaded to the server, and the name of this new network was then included in the u-map recommendations for that region. From that point on, the entire community of the u-map users was informed about the performance of this new network. In this study, the first Sipdroid call using the *PublicWiFi!* network took place during the fourth day of the study. Until the end of the study, 6 out of 21 users performed VoIP calls using this network. Figs. 6.7 (left) present the QoE score, packet loss

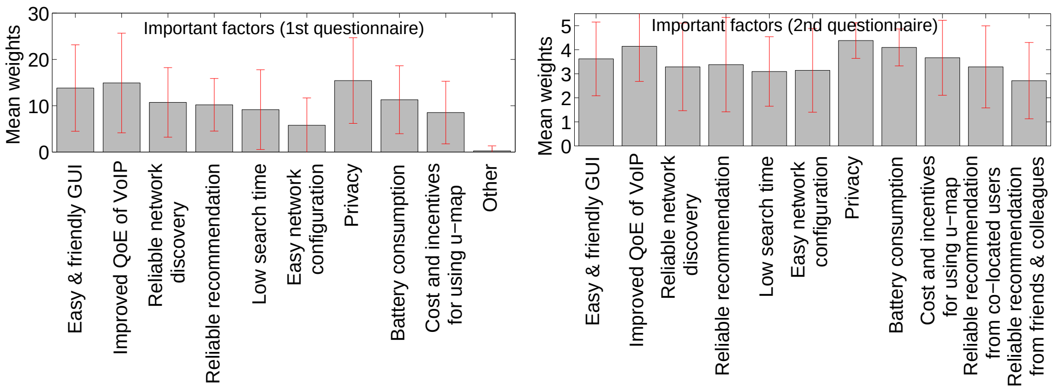

(a) (b)

Figure 6.6: Important factors that affect the use of the u-map based on data collected from, (a) the 1st questionnaire (before using the u-map), and (b) the 2nd questionnaire (at the end of the field study).

ratio, and burstiness of packet loss of each VoIP call that was made by one of those users (in a temporal order). Low QoE scores are correlated with high packet loss ratio and burstiness. It is evident that by discovering *PublicWiFi!*, this user improved its QoE. *PublicWiFi!* performed better than the other networks, as it can be verified by the collected network measurements and QoE statistics. Fig. 6.8 shows the complementary CDFs (CCDFs) of the performance of all WiFi networks from which Sipdroid calls were made. Due to its light traffic, the *PublicWiFi!* network could be beneficial to all the participants. However, only some users took advantage of its presence. We asked the users to report (in a related question of the second questionnaire) whether they noticed a difference in the QoE of their VoIP calls over the different networks. Then, we compared these responses with the QoE scores that the users assigned to their VoIP calls (that may have taken place in different networks). Indeed, the users who indicated that they perceived a better quality of VoIP when using a specific network had also provided higher scores for the QoE of the VoIP calls made from *PublicWiFi!*, compared to the calls made from *forth public access* (Fig. 6.9, red solid line vs. blue solid one). Similarly, the stochastic difference between the QoE scores assigned to the calls performed through different networks by the users who did not perceive any difference on the QoE of their calls is less prominent.

Although several participants agreed that the improvement of QoE is important, they evaluated the “reliable network discovery” and “easy network configuration” as less significant factors (Fig. 6.6). Users were interested in increasing the quality of their experience through the u-map without having to perform the network discovery and configuration themselves. This phenomenon was also consistent with what was observed via the u-map traces: there were users that despite their low QoE, did not change network (as the one in Figs. 6.7 (right)). This could be due to their unwillingness to “navigate” through the network discovery and configuration GUI or their privacy concerns about a new and unknown network (e.g., *PublicWiFi!*). This is an example of the trade-off between privacy protection and the automation of the network discovery and configuration process (i.e., reducing the user intervention by making the calls an almost "1-click" operation).

6.2.9 GUI evaluation

The evaluation of the GUI focused on the *functionality, visual clarity, consistency, compatibility, informative feedback, explicitness*, and *flexibility & control* of the u-map (Table 6.3). The users evaluated these aspects in the second questionnaire using the Likert scale. The mean overall score for the GUI is 3.90 out of 5, while the mean score for each aspect is shown in Fig. 6.10. Regarding the GUI that reports errors (question Q2.33 in Appendix B), there was a noticeable dichotomy: there were participants that assigned high score (4 or 5), while others provided a low score (2). During the field study, we discovered some malfunctions of the u-map uploading related to the concurrency of the reading and writing at the local database. This was corrected only after the completion of the field study.

6.2.10 User profiling

Could we distinguish different groups of users, with respect to their background (e.g., gender, degree of familiarity with respect to various technologies), demand, expectations about the u-map and its services, evaluation of the u-map, QoE requirements (e.g., responsiveness, QoE vs. cost)? To answer this question, we applied a clustering algorithm, namely an extended version of the K-means algorithm [65], on a data set that consists of traffic statistics and responses from the questionnaires. This algorithm handles mixed categorical and numerical datasets. Table 6.4 describes all used variables, which are ordered based on their significance in the determination of user profiles ("weight" column). The impact of the overall evaluation of the u-map and the traffic demand are of significant importance. Also, the preference among QoE and price was not found to be significant, since all used WiFi networks were free of charge. Under different conditions, this factor should have higher importance.

Table 6.3: Metrics for the evaluation of GUI.

Name	Description
Functionality	Tools for the u-map services (query, call, configure)
Visual clarity	Information in straightforward, organized, clear way
Consistency	Information always depicted in same format/layout
Compatibility	Compatibility with user conventions and expectations
Informative feedback	About actions that result in errors
Explicitness	Information in a complete, clear, & structured way
Flexibility & control	Control and recovery from erroneous actions

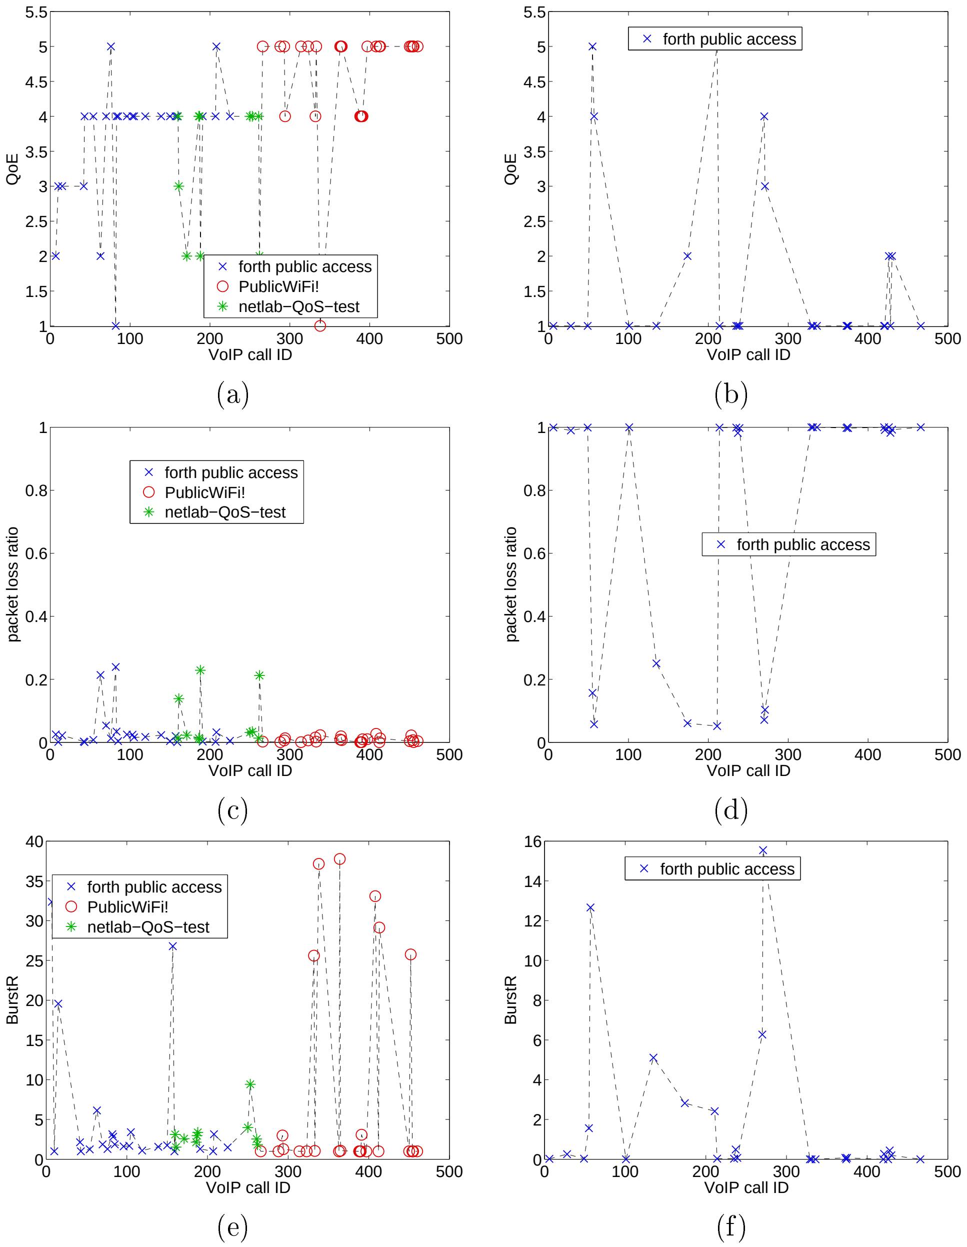

Figure 6.7: QoE score (a & b), packet loss ratio (c & d), and burstiness of packet loss (e & f) of the calls of two participants, in temporal order. Figures at the left correspond to a participant that discovered the new network (*PublicWiFi!*), while Figures at the right belong to one who decided to remain in the well-known public wireless network of FORTH, despite the poor network conditions during his/her calls, which affected their quality.

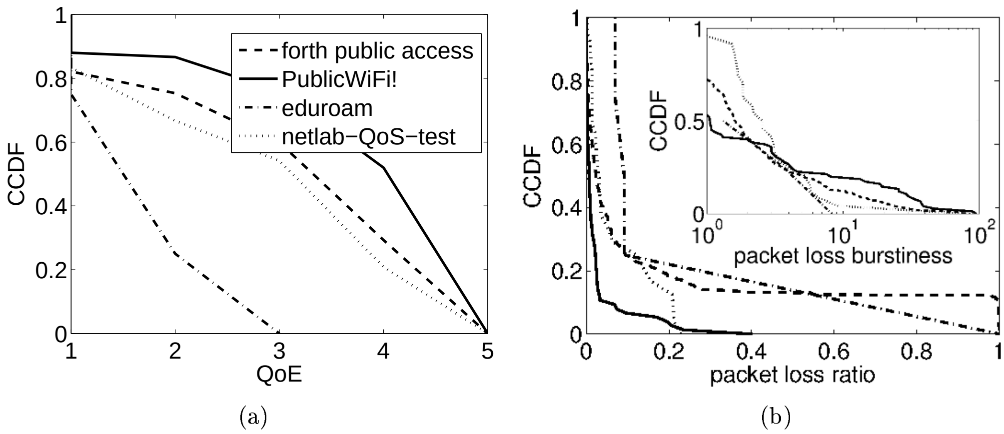

Figure 6.8: Performance of WiFi networks, in terms of QoE, packet loss ratio & burstiness of packet loss ((b) has the same legend as (a)).

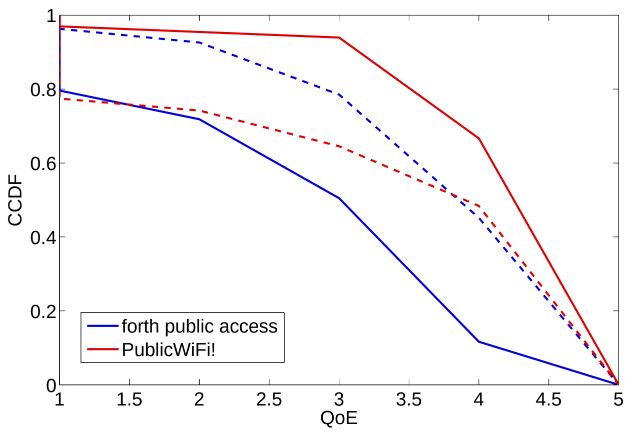

Figure 6.9: The QoE scores of VoIP calls of all users who indicated that they perceived a better quality of VoIP when using a specific network (solid line) vs. of the ones who did not observe any difference (dashed line), in the "Improved QoE of VoIP" related question, according to the network where these calls took place (*forth public access* & *PublicWiFi!*).

The clustering algorithm reports three clusters of 4, 6, and 11, users, respectively. Table 6.5 shows the centroids of the numerical attributes in each cluster, while Fig. 6.11 shows how users of each cluster answer various types of questions. With respect to the evaluation of the u-map, users of cluster 1 give the highest score (Table 6.5) and are more willing to adopt the u-map compared to users of other clusters (Fig 6.11 (a)). They also give the highest score to the provider selection capability (Table 6.5). This is in accordance with the high score for the u-map, as provider selection is one of the key features of the u-map. Furthermore,

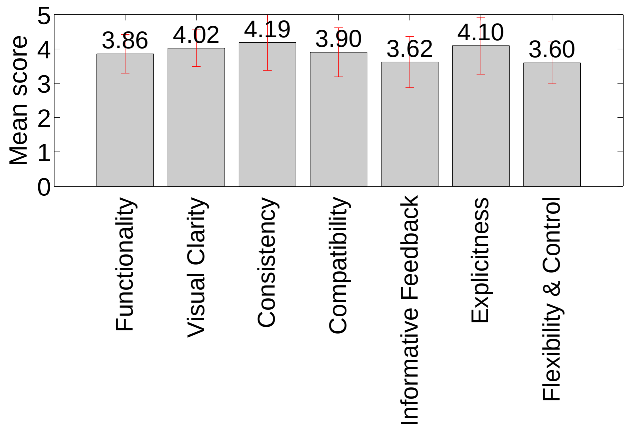

Figure 6.10: GUI evaluation.

users of clusters 1 and 3 expect that the u-map will have a positive impact on the quality of VoIP calls, while users of cluster 2 are more skeptical (Fig. 6.11 (b)). Users of cluster 3 evaluate poorly the quality of VoIP calls on Android (Table 6.5). This belief strengthens their motivation to adopt the u-map. On the other hand, users of cluster 2 give a relatively high score to the quality of VoIP calls on Android. This belief in conjunction with their expectation that the u-map will not significantly improve the quality of VoIP calls explains why they provide the lowest score for the u-map compared to other clusters. Another interesting result for users of cluster 2 is that their majority has a higher level of familiarity with Android compared to users of other clusters (Fig. 6.11 (d)). It seems that users with low familiarity with Android provide a higher overall score compared to users with high familiarity. We speculate that the u-map can improve the experience of users that are not very familiar with new technologies. Finally, users of clusters 1 and 3 have the highest and lowest levels of traffic demand, respectively, while users of cluster 2 are characterized by a medium level of demand.

6.2.11 Modeling the user satisfaction about the u-map

In general, the user satisfaction about a system can be parameterized by the perception of the user about the quality of its services and functionalities, given his/her requirements. Requirements exist with respect to the battery consumption, responsiveness, GUI, security, fault-tolerance, and scalability. In this work, we focus on the battery consumption, responsiveness and GUI. In the second questionnaire, the participants were asked to assess the u-map services, namely the *improvement in the QoE of VoIP*, *reliability of network discovery*, *reliability in network recommendation*, and *easiness in network configuration*. The objective of this work is to model the user satisfaction as a utility function based on the above parameters and their significance level. The design of these utility functions is still largely

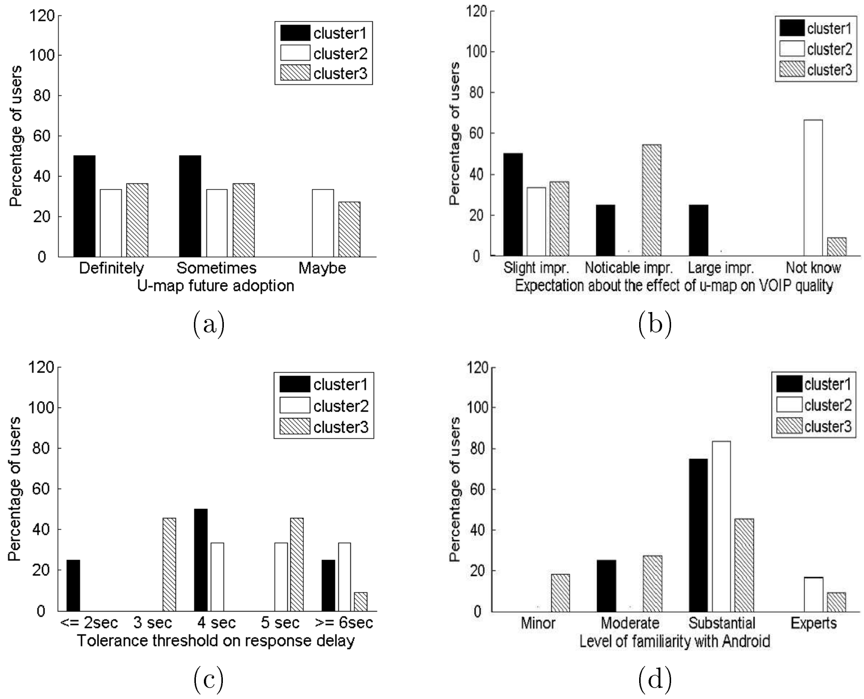

Figure 6.11: Population distribution per cluster according to the categorical variables.

under-explored. To the best of our knowledge, there are no such research works that perform real-world user studies to model the user satisfaction.

In the second questionnaire, at the end of the field study, the subjects assessed their satisfaction about the u-map, providing an overall score and an indication

Table 6.4: User profile parameters and their weight.

Weight	Parameters of user profile
0.6234	Overall score for u-map
0.5738	Expectation about the effect of u-map on VOIP quality
0.5698	Significance of provider selection capability
0.5502	Frequency of VOIP calls from mobile device
0.5212	Total call duration
0.5132	Tolerance threshold on response delay
0.5131	Frequency of GSM calls
0.4910	Usage of recommendation systems
0.4907	Evaluation of VOIP quality on Android
0.4686	Median call duration
0.4682	Level of familiarity with Android
0.4340	Quality of service vs. cost (preference)
0.3846	gender
0.3327	u-map future adoption

about their intention for future adoption of the u-map (Fig. 6.12).

Following the approach of Grigoroudis *et al.* [66], we model the user satisfaction of the u-map as a weighted sum of n criteria, based on user feedback. These criteria, with their respective score and significance weight, are presented in Fig. 6.13 and Fig. 6.6 (b), respectively. Fig. 6.13 shows the mean opinion score of users, per criterion. Specifically, the mean user satisfaction about the u-map is defined as a utility estimator $U(u)$,

$$U(u) = \sum_{i=1}^{n} w_i(u)s_i(u) \tag{6.1}$$

where $\sum_{i=1}^{n} w_i(u) = 1$, and $w_i(u)$ and $s_i(u)$ indicate the weight of the i-th criterion and its score, normalized in the interval $[1, 5]$, respectively. Fig. 6.14 illustrates the user utility function. 33 questions are involved in the generation of the utility: 20 of them evaluate the GUI, 6 of them the remaining criteria, and 7 others provide the criteria weights. For the evaluation of the GUI and the assessment of the user satisfaction, the missing values needed to be filled, and the *k-nearest neighbors imputation* method was used. The predicted user satisfaction of the u-map (according to Eq. 6.1) is consistent with the overall score provided by users. The mean user satisfaction is 3.65 (out of 5) compared to the mean overall score of 3.58. The proposed model accurately predicts the overall score: for 10 out of 21 users, the model estimates the exact score, while for 7 other users the absolute error is 1 (in a range from 1 to 5). There were 4 users that did not provide any overall score.

Table 6.5: The mean of each numerical parameter per cluster.

Description	**Cluster 1**	**Cluster 2**	**Cluster 3**
Overall score for the u-map	3.75	3.50	3.73
Significance of provider selection capability	4.25	3.83	3.73
Total call duration (sec)	2868.2	1632.5	560.2
Evaluation of VoIP quality on Android	3.25	3.50	1.45
Median call duration (sec)	59.63	54.83	31.55
Significance of QoS	47.50	50	48.18
Significance of price	52.50	50	51.82

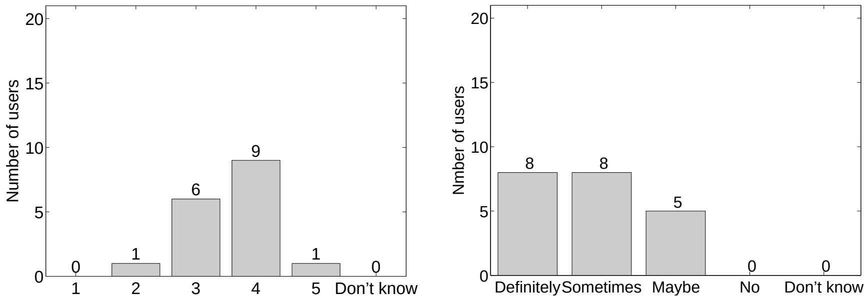

(a) (b)

Figure 6.12: Overall evaluation of the u-map (a), and future adoption of the u-map (b).

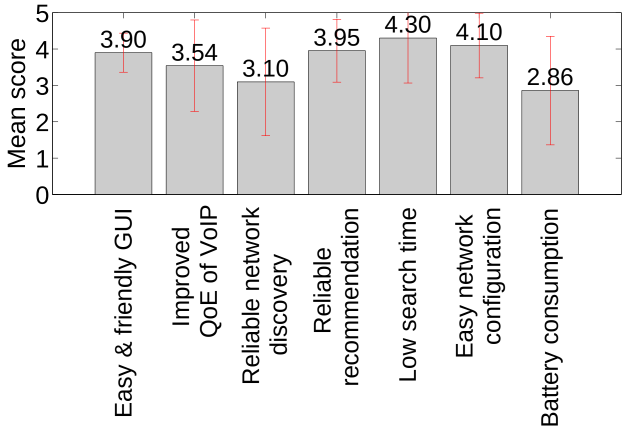

Figure 6.13: Scores of the user satisfaction criteria.

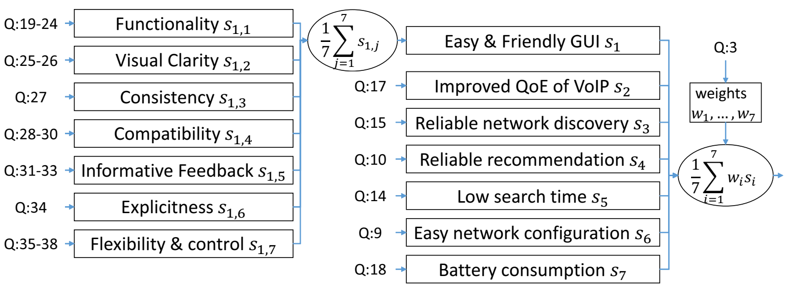

Figure 6.14: Model about the user satisfaction about the u-map. The "Q:X-Y" indicates the corresponding questions in the 2nd questionnaire (Appendix B).

Chapter 7

The evolution of the u-map

"All men can see the tactics whereby I conquer, but what none can see is the strategy out of which victory is evolved."

– Sun Tzu, *The Art of War*

This Chapter presents the improvements and extensions of the u-map system, that followed the evaluation of Chapter 6. Also, it presents how the u-map has been applied on different domains, namely on water distribution networks and medical applications, demonstrating its generality.

7.1 Prototype refinement

Following the directions for improvement highlighted by the performance analysis and the user study, we have refined the implementation of the u-map system.

One-click operation: In the u-map version of Chapter 5, the user had to specify a polygon, before submitting a query. Then, the presentation of the query response showed a list of networks and the respective scores, together with four icons that launched the cellular, and WiFi network settings, as well as to the Phone, and Sipdroid applications. The effort required by the user for the execution of a query and the connection to a recommended network has been minimized. The user can submit a query without specifying a polygon. In this case, the area of the map that is currently shown on screen is considered for the polygon parameter of the query (Fig. 7.1 (b)). The GUI that presents the query results has been improved, too. Each network included in the query result is represented by a list line, that includes its name, its score (using a graphical element with 5 stars), and an icon (Fig. 7.1 (b)). The icon indicates the network type and status. The network type can be WiFi, GSM, or 3G. The color of the icon represents the status: green icons

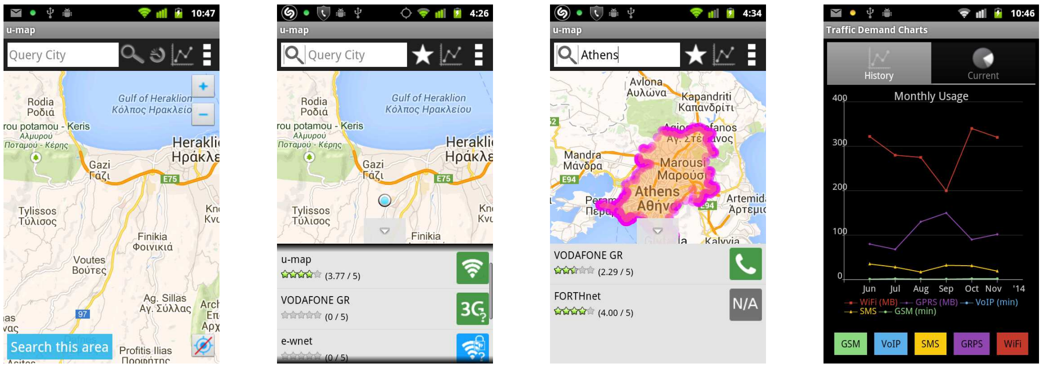

(a) (b) (c) (d)

Figure 7.1: GUI screenshots from the refined the u-map client: (a) main screen, presentation of a query result considering (b) the currently shown area and (c) a city name, as well as (d) traffic demand history chart.

depict networks currently connected at, yellow icons depict networks for which the u-map server has QoS or QoE evaluations in the specified area, blue icons depict networks for which the server has not any data in the specified area (newly discovered networks), and gray icons depict networks that are included in the query response, yet are not available currently (e.g., the network has been shutdown, or the device is out of range). Clicking on a line on the network list connects the device on the specific network and launches either the Phone or the Sipdroid application, depending on the network type. This new implementation of the querying GUI allows the user to query the server with a single click, and connect to one of the recommended networks with one more click.

Automatic network connection: To completely relieve the user from the burden of querying the server and selecting a network, we have fully automated the network selection process. If this feature is enabled by the u-map settings, when the Sipdroid application is launched, a query is automatically sent to the u-map server, and the device is connected to the best available WiFi or cellular network, based on the server-provided recommendations.

Querying by city names: Another extension of the querying mechanism is the ability to define areas by their name. Geo-coding web services (i.e., the Google geo-coding API and the Wikimapia API) are employed to resolve area names (e.g., city names) into polygons that define the querying area precisely (Fig. 7.1 (c)). Resolved area names are cached, to answer faster subsequent queries regarding these areas.

Pre-cached queries: Evaluating complex queries with spatial parameters on a large database can take a while. Ensuring an efficient database schema, optimizing the SQL queries, as well as defining proper indexes on the database can improve reduce the query evaluation delay. To further improve the responsiveness of the u-map server, regarding the responses to queries from clients, a pre-caching approach has been implemented. Queries are pre-evaluated on specific areas and the results are stored in memory, using the memcached service. Then, when a client queries the server, the result can be retrieved from the cache, instead of being computed at the database.

Traffic demand estimation: The traffic demand estimator is a module implemented after the user study. It monitors the traffic demand of the user, using different counters for outgoing (possibly charged) traffic to every distinct provider. The u-map client illustrates the traffic demand history using GUI charts (Fig. 7.1 (d)). Traffic demand statistics are used for the user profiling. Also, monitoring the traffic demand allows to precisely estimate the cost of telecommunication services, provided that the pricing information (e.g., tariff plans) is available. Cost-based queries can then be implemented, providing recommendations for the most cost-efficient networks/services.

Integration with the VoIP service: Registering for a VoIP service and configuring the Sipdroid application requires some effort. To facilitate this procedure, we are hosting a VoIP server that is cooperating with the u-map: When the u-map client is installed, it checks whether the Sipdroid application is installed. If not, the u-map client requests permission from the user to download and install our Sipdroid edition. When the Sipdroid is installed and launched, it automatically retrieves the VoIP account configuration from the u-map client, through Android IPC. The VoIP account is created when the u-map client is registered to the u-map server.

RTP measurements: The network delay significantly impacts the perceived quality of a VoIP call [43]. We have improved the performance estimator module that we have implemented inside Sipdroid, so as to measure the RTP-packet delay.

Access control: The access control module has been designed. Providing a robust implementation with a low footprint on query evaluation delay is an ongoing effort.

Analytics: Our approach for analyzing the collected data has been described in Chapter 3. The team has implemented the analytics module in Matlab and the dataset collected during the field study has been analyzed [67]. The integration of the analytics module with the PHP application of the u-map server is another ongoing task of the team.

7.2 Connection with other applications

Functionality or components of the u-map system have been used on other systems. The successful reuse of the u-map on different application domains provides an insight about the modularity of the system and the generality of the u-map paradigm. The following paragraphs describe how the u-map has been applied at each of these domains.

7.2.1 u-plan

U-plan is another system that the team has developed. The u-plan system empowers users with tariff plan recommendations, based on the user profile and tariff plan cost. Specifically, the u-plan server models the tariff plans with which the providers charge their customers for the offered telecommunication services. The u-plan server includes a database that contains the tariff plans available in the wireless market (e.g., tariff plans offered by providers). The u-plan client shares the traffic demand estimator module of the u-map. By applying various tariff plans on the traffic demand history of the user, u-plan calculates the cost that the user would pay by choosing it.

7.2.2 Water distribution networks

Water management is another application domain in which the mechanisms of the u-map and CoRLAB could be applied. Fresh water resources are threatened by overexploitation, poor management, ecological degradation, and pollution. Water distribution networks should maintain microbial quality and prevent the growth of pathogenic organisms. In urban environments, such water distribution networks require prevention and monitoring, planning, and management systems to maintain drinking water within acceptable standards.

Sensors measuring the water quality in terms of mineral concentration, chlorination level, and bio-agent (e.g., protozoa, E. coli, and salmonella) and noxious-substance (e.g., pharmaceuticals) contamination could be installed on various parts of the network (e.g., tanks, large transmission pipe water networks, distribution water networks, and home water networks). Apart from placing sensors and monitoring devices at various points along the path from the main water distribution network to the end-users, the u-map can empower end-users by enabling them to provide real-time feedback about a wide range of factors, such as subjective opinion scores about the quality of water (e.g., unpleasant odor or taste, presence of impurities). This is analogous to our approach on the telecommunication markets, since it enables users to collect network measurements and provide feedback/opinion scores. These measurements would be uploaded on a central server, building a geo-database with combined objective and subjective quality indicators. The aim is to inform users about the water quality in specific regions in order to make optimized decisions about water usage and capacity planning in their infrastructure. Such information can be either automatically transmitted to user smartphones or upon request.

The interested parties (e.g., health centers, hospitals, water distribution operators, and governmental/consumer protection bodies) may also build a profile based on their demand and requirements in order to receive personalized notifications. For example, depending on the type of environment (e.g., a residential area, hospital, industrial environment, or farm) and user activity, the quality, demand, and availability requirements may change considerably. Furthermore, different access policies can be applied according to each profile.

The proposed system provides the capability to aggregate measurements about water quality and demand in large spatio-temporal databases, as well as to utilize and correlate a set of multi-sourced data by employing data mining techniques. Based on these data analytics, the u-map could provide customized recommendations to users, predictions or early-warnings to various interested parties (e.g., health centers, hospitals, water distribution operators, and governmental/consumer protection bodies) and trigger various adaptation mechanisms and events (e.g., change the sampling rate, replace filters).

The team has extended the functionality of the u-map for water distribution networks and conducted a preliminary analysis [21]. The proof-of-concept implementation has demonstrated the generality of the u-map.

We foresee that adopting the u-map for water management purposes may offer great opportunities for early problem identification and alleviation, improving the performance and possibly preventing unnecessary maintenance.

7.2.3 Medical applications

The team (not including the author) has developed mMamee [68], a platform that monitors and assesses the exposure of pregnant women on environmental risks. The platform integrates sensing data with the descriptive input on maternal daily habits. The mMammee system adopts the design principles of the u-map.

7.2.4 Forthnet

The Forthnet Group is the largest privately owned group that provides broadband and pay TV services in Greece. At the time of writing this thesis, a contract has been signed between FORTH and Forthnet Group, for the application of the u-map paradigm on a video streaming service that the Forthnet Group offers. This collaboration provides the context for extending the u-map for more services, and an opportunity to apply the research produced by this thesis at the industry.

Chapter 8

Conclusions and Future Work

"I tell my students, 'When you get these jobs that you have been so brilliantly trained for, just remember that your real job is that if you are free, you need to free somebody else. If you have some power, then your job is to empower somebody else. This is not just a grab-bag candy game.'"

– Toni Morrison

8.1 Conclusions

This thesis presents the u-map and an extensive evaluation of its performance. The u-map is a review and recommendation system that follows the client-server architecture. A u-map client running on a smartphone uploads in regular basis information about the user profile (e.g., service requirements, device capabilities), QoE feedback about services (as indicated explicitly by the user), traffic demand statistics, network/spectrum conditions, and position of the device. This information is stored in a spatio-temporal geo-database, attached to a u-map server. A u-map client may query this database to obtain information about service providers and their customers' satisfaction. In that way, a user can make an educated selection of the appropriate provider, when visiting a new area. The term appropriate denotes the provider that best matches the user's criteria (e.g., it is the best in term of a certain QoS metric, or in QoE, in that area, or it supports specific features). Moreover, given that a smartphone may have access to several network interfaces simultaneously, a u-map client may select the appropriate wireless or cellular operator and network interface for offloading its traffic.

The analysis demonstrates that the u-map has low response delay and power consumption. The display is the main drain of power consumption. Energy savings can be achieved by less frequent position estimations and more efficient GUI navigation. The outcome of the field study is also encouraging. Most of the participants indicated that the u-map can improve their experience, helps in the discovery of networks that offer improved QoS. During the study, 30% of users discovered the best available network, improving substantially their experience. Users were also satisfied with the GUI of the u-map. The presence of users that despite the low QoE of their calls did not change network motivates us to jointly consider the problem of improving the automation of the network discovery and configuration process, addressing the privacy concern in joining a new network. Finally, the users evaluated positively the overall performance as well as specific features of the u-map. Based on user feedback, we modeled the user satisfaction about the u-map as a utility function. The model can predict the overall score of the user satisfaction about the u-map accurately. Finally, we applied the u-map in other domains, such as water distribution networks and medical applications, demonstrating its generality.

8.2 Future work

From a signal-processing perspective, the monitoring mechanism of the u-map system includes a number of the u-map clients that sample the QoS and QoE at various locations and feed the server with multi-dimensional data. Experimenting with signal processing algorithms for real-time analysis of multi-dimensional time-series is part of our future work plan. Related to this is the estimation of data uncertainty (e.g., confidence level), possibly by considering the cross-correlation between the sampled parameters. Also, anomaly detection algorithms could distinct faults (or bad performance) at the infrastructure from faults at monitoring nodes (e.g., smartphones, sensor nodes) and inaccurate/mistaken user feedback. The detection of malicious users is also an interesting problem.

Deploying the u-map as a cloud service is another direction for future work. Clusters of commodity machines as units of scaling allow a service to scale up while maintaining a high performance/cost ratio. This allows to upgrade the service infrastructure without having to throw out the current machine and replace it with a larger one. Instead, new commodity building blocks are added to the cluster. Also, clusters have natural redundancy due to the independence of nodes, ensuring stability and robustness. In this context, we plan to design the u-map with a layered architecture that includes load balancing machines, front-end instances that serve incoming requests from the u-map clients, and back-end workers that provide the storage and computational resources for running analysis algorithms. Worker instances should host a software stack providing them with a distributed file system, a fault tolerant, distributed database capable of storing large quantities of spatial, sparse data, and a way of processing large data sets with parallel, distributed algorithms. The Apache Hadoop software stack, with the Hadoop Distributed Filesystem (HDFS) and the HBase database could be integrated with the current implementation of the u-map.

Crowdsourcing systems become useful only after achieving a critical mass of participating members. Sensing on smartphones consumes battery, which might discourage user participation in some scenarios. Similarly, the privacy requirements could also prevent users from participating. Part of our future research plan is the development of energy-efficient monitoring, effective incentive mechanisms that can attract more user participation. User reputation or altruism, as well as the improved QoE in services could be some incentives for using the u-map. Micropayments (e.g., free SMS, calls) could further motivate users to participate.

Developing an efficient access control mechanism that allows users to define various access control rules without significant reduction on the responsiveness of the system is an ongoing effort.

A strong point of the u-map is its generality: The concept of cross-correlating objective QoS measurements with subjective QoE feedback from the users can be applied in a plethora of domains. The u-map functionality has successfully been extended for water distribution networks and medical applications. We plan to experiment with applying the u-map paradigm on various domains.

Appendix A

List of Abbreviations

AC Alternating Current

AP Access Point

API Application Program Interface

BEREC Body of European Regulators for Electronic Communications

BS Base Station

BSSID Basic Service Set Identifier

CA Certificate Authority

CPU Central Processing Unit

CDF Cumulative Distribution Function

CCDF Complementary Cumulative Distribution Function

CLI Command Line Interface

DNS Domain Name System

FCC Federal Communications Commission

GPRS General packet radio service

GPS Global Positioning System

GSM Global System for Mobile Communications

GUI Graphical User Interface

HDFS Hadoop Distributed Filesystem

HSPA High Speed Packet Access

HTTP Hypertext Transfer Protocol

IP Internet Protocol

IPC Inter-Process Communication

JSON JavaScript Object Notation

LAN Local Area Network

LTE Long-Term Evolution

ML Machine Learning

MNO Mobile Network Operator

MVNO Mobile Virtual Network Operator

NRDBP National Reference Database for Portability

OLED Organic Light-Emitting Diode

ORM Object-Relational Mapping

OS Operating System

PC Personal Computer

QoE Quality of Experience

QoS Quality of Service

RAM Random-Access Memory

RDBMS Relational Database Management System

RSSI Received Signal Strength Indication

RTP Real-time Transport Protocol

RTT Round-Trip Time

SIM Subscriber Identity Module

SIP Session Initiation Protocol

SMS Short Message Service

SQL Structured Query Language

SSID Service Set Identifier

TCP Transmission Control Protocol

TDMA Time Division Multiple Access

TLS Transport Layer Security

UDP User Datagram Protocol

UE User Equipment

USB Universal Serial Bus

VoIP Voice over IP

WiFi Wireless Fidelity

Appendix B

Questionnaires

B.1 Questionnaire 1

User ID: ....
Gender: ....
Age: ....
Have you studied at CSD-UoC or a relevant Department of Computer Science? ....

B.1.1 Experience on Android devices

Q1.1 How often do you use *Android devices* in your everyday life?

- Never before this study.
- Rarely until this study.
- In daily basis, for the last .... months.

Q1.2 Evaluate your experience and familiarity with the use of *Android smartphones.*

- Minimum familiarity: the interface seems strange and I feel uncomfortable. I use rarely Android smartphones.
- Moderate familiarity: I use the Android smartphone mainly for phone calls.
- High familiarity: I use the Android smartphone for many applications.
- Expert: I use the Android smartphone for many applications, and I am developing Android applications.
- None of the above. (Specify ....)

B.1.2 Use of Android device and relevant applications

Q1.3 How often do you use your *mobile phone for phone calls (receive/perform) via GSM* in your everyday life?

- Never.
- I make less than 10 calls per week.
- I make less than 10 calls every day (on average).
- I make about 10 calls every day (on average).
- I make more than 10 calls every day (on average).

Q1.4 How often do you use *VoIP, (e.g., Skype or other VoIP application) from mobile phone* in your everyday life?

- Never.
- I make less than 10 calls per week.
- I make less than 10 calls every day (on average).
- I make more than 10 calls every day (on average).

Q1.5 When do you prefer the *use of the following devices/applications* for your phone calls, given that you have all these three choices available?

- Skype or relevant VoIP application from your mobile phone? ....
- GSM from your mobile phone? ....
- Landline phone? ....
- I do not understand the question.

Q1.6 Choose your most *frequent activities with smartphones* (Android or other platforms). You must have performed an activity more than 3 times in order to choose it. You may select multiple options.

- Access to information stored locally in the device
- Sending messages
- Retrieving information from the Internet (e.g., web browsing)
- Social networks (e.g. Facebook, Twitter)
- Email
- Read newspapers/news sites
- Watching videos (e.g., access to YouTube)
- Games
- Banking
- Searching for products and services
- Buying products, reservation
- Camera
- Alarm clock

- Other
- Just as mobile phone

Q1.7 When sending a query to *retrieve/find information from the Internet (e.g. a Web page while browsing)*, how much time are you willing to wait for the response? Choose one of the following:

- 2 seconds or less
- 3 seconds
- 4 seconds
- 5 seconds
- 6 seconds or more

Q1.8 Which *wireless Internet connection* technologies do you use from your mobile phone? (You must have used a technology more that 3 times in order to choose it. You can choose multiple answers.)

- 3G
- 4G
- WiFi
- I do not know
- I do not have wireless connection

Q1.9 Evaluate your experience and familiarity in *configuring the WiFi* of your smartphone, when visiting a new place (e.g., cafe, shop) and wanting to connect to the Internet. Choose one of the following:

- I do it often. I face problems with the WiFi settings rarely, even in places that I visit for first time.
- I do it often, but sometimes I face problems with the WiFi settings in places that I visit for first time.
- I have done it a few times, I often face problems with the WiFi settings, and I ask for help.
- I have not tried it yet.
- I have not tried it yet and I am not willing to try.
- I do not know/understand the question.
- Other: ....

Q1.10 *Comment on your use of online recommendation systems* (e.g. booking.com, imdb.com) from mobile phones.

- I haven't used any recommendation systems, neither from my mobile phone nor my computer.

- I have used them only from my computer.
- I have used some of them from my mobile phone, but not systematically.
- I systematically use them from my mobile phone.
- I do not know what online recommendation systems are.

Q1.11 Which is your overall *experience on the quality of VoIP calls performed using Android* smartphones until now (before this user study)?

- (Unacceptable) 1 2 3 4 5 (Excellent)
- I do not remember
- I do not have used VoIP applications on Android

Q1.12 How much do you expect *to change your overall experience on the call quality via VoIP applications* at ICS-FORTH, because of the use of u-map?

- Worse No improvement Least Perceivable Great
- I do not know

Q1.13 Which *parameters* do you consider as the most *important* for the use of a system like u-map? Classify the following parameters in an order from the most to the least significant one, and assign them a weight in the range from 0 to 100, where the sum of all weights equals to 100.

A. Easy and friendly GUI environment
B. Improved quality of VoIP calls
C. Reliable information about the presence of new networks and providers
D. Reliable information about the performance of networks and providers in a region
E. Small delay on the search of the appropriate wireless network
F. Easy configuration for the appropriate wireless network, for the conduction of a call
G. Privacy
H. Battery consumption
I. Cost of the service and incentives (i.e., free calls or extra services) for using the system
J. Other

B.2 Questionnaire 2

User ID: ....

Q2.1 Evaluate the quality of the u-map service in general.

- (NOT USEFUL) 1 2 3 4 5 (VERY USEFUL)

B.2.1 Service evaluation of u-map

Future adoption of u-map

Q2.2 Suppose that you are visiting a new area and you are interested in finding the most appropriate WiFi provider for your phone calls. Would you use u-map for finding the appropriate WiFi provider for your phone calls?

- Yes, surely.
- Yes, sometimes.
- Maybe.
- No.
- I do not know.

Q2.3 Classify the following factors in the range of *NOT IMPORTANT (1)* to *VERY IMPORTANT (5)*.

A. Easy and friendly GUI environment
B. Improved quality of VoIP calls
C. Reliable information about the presence of new networks and providers
D. Reliable information about the performance of networks and providers in a region
E. Small delay on the search of the appropriate wireless network
F. Easy configuration for the appropriate wireless network, for the conduction of a call
G. Privacy
H. Battery consumption
I. Cost of the service and incentives (i.e., free calls or extra services) for using the system
J. Reliable information about provider evaluations from colocated users
K. Reliable information about provider evaluations from friends and colleagues

1. (NOT IMPORTANT): ....
2. ....
3. ....
4. ....
5. (VERY IMPORTANT): ....

Q2.4 How important is for you the capability of being able to *choose, before a call, the provider you will use*?

- (NOT AT ALL) 1 2 3 4 5 (VERY IMPORTANT)

Q2.5 Assign a significance *weight* to each of the following two call parameters *(should sum up to 100).*

- Call quality ....
- Call cost ....

Q2.6 Did you encounter problems in performing calls during the study? Fill in for each problem a number from 1 to 5, where "1" corresponds to "no problem" and "5" to "many problems".

- The u-map application was not working properly on my device ....
- The sipdroid application was not working properly on my device ....
- My friends were usually offline and i could not call them ....
- Other, specify ....

Q2.7 Evaluate Sipdroid.

- (NOT USEFUL) 1 2 3 4 5 (VERY USEFUL)

Use of u-map during the current study

Q2.8 How did *the use of Sipdroid* affect your experience regarding the use of u-map?

- (VERY NEGATIVE) 1 2 3 4 5 (VERY POSITIVE)

Q2.9 How easy is *the WiFi network selection for VoIP calls* using u-map?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.10 Do you find the results of your *u-map queries* reliable (i.e., consistent with your view of these networks)?

- (NO) 1 2 3 4 5 (YES)
- I did not send a query.
- I do not know/ I do not understand the question.

Q2.11 On average, how good was the quality of the calls that you made?

- (BAD) 1 2 3 4 5 (GOOD)
- I do not remember.
- I do not understand the question.

Q2.12 How many times *have you changed the WiFi network* for your phone calls during the study?

- Never.
- Once.
- Between 2 to 5 times.
- More than 5 times.
- I do not remember.
- I do not know/ I do not understand the question.

Q2.13 For which reasons did you *choose a WiFi network for your phone calls*?

- It was free.
- The call quality was better using WiFi than GSM.
- It was recommended by u-map.
- It was already configured in the settings of the phone.
- I had to do so because of study.
- I chose it randomly.
- I do not remember doing something relevant!
- Oops... I did not realize that i had to choose a WiFi network.

Q2.14 Evaluate *the system's responsiveness* during the provider selection process.

- The response delay was low.
- The response delay was tolerable most times.
- The response delay was long most times.
- I stopped using the system because of the long response delay.
- I did not realize something relevant to the response delay.
- I do not know/ I do not understand the question.

Q2.15 Did you notice the presence of *different WiFi networks through u-map*?

- Yes, I usually used ....
- Yes, through u-map I was informed about ....
- No, I did not notice that there were different WiFi networks during the study.
- I do not know what you mean by different WiFi networks.

Q2.16 Did you notice differences in coverage/quality of *WiFi networks* that were available at FORTH?

- Yes. The .... has/have larger coverage than the .... in the region of ....
- No, I did not notice anything.
- I do not know/ I do not understand the question.

Q2.17 Did you notice any change in the call quality when *using different WiFi networks*? Choose one of the following sentences that better describes your experience.

- There was a WiFi that has considerably better quality performance during the call, systematically for at most calls that I made.
- The most of the times I used the same WiFi network, because .... (choose only one of the following)
  - I already knew it.
  - I was bored to change it.
  - It was complicated to change it.
- I did not notice a WiFi network that had better performance systematically, during the most of the calls I made.
- I did not notice any difference between WiFi networks.
- I do not know/ I do not understand the question.

Q2.18 Did you notice any significant *reduction in battery lifetime* while using u-map?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.

B.2.2 Evaluation of u-map GUI

Functionality

Q2.19 Does u-map satisfy your *needs and requirements*, while performing tasks?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.20 Does u-map simplify the *discovery of a WiFi network*?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I did not use this action.

- I do not know/ I do not understand the question.

Q2.21 Is the *configuration/connection on a WiFi network* easy?

- (NO) 1 2 3 4 5 (YES)
- I did not use the Android settings for a WiFi network.
- I do not remember.
- I do not know/ I do not understand the question.

Q2.22 Does the system give the appropriate *feedback*, in order to facilitate the *conduction of a call*?

- (NO) 1 2 3 4 5 (YES)
- I did not conduct any call.
- I do not remember.
- I do not know/ I do not understand the question.

Q2.23 Does the system give the appropriate *feedback* for the *general application settings* of u-map?

- (NO) 1 2 3 4 5 (YES)
- I did not use the general application settings of u-map.
- I do not remember.
- I do not know/ I do not understand the question.

Q2.24 Does the *"Help" section* represent the actual use of u-map? Is it helpful?

- (NO) 1 2 3 4 5 (YES)
- I did not use the "Help" section.
- I do not remember.
- I do not know/ I do not understand the question.

Visual Clarity

Q2.25 Is the displayed information straightforward, organized, and clear?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.26 Is the *important information* presented emphatically?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Consistency

Q2.27 Is the *same type of information* presented (i.e., instructions, menu, messages, titles) *everywhere at the same location* using the same layout?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Compatibility

Q2.28 Does u-map follow your *conventions and preferences*?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.29 Does the u-map graphical representation follow the *ordinary conventions*?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.30 Are you *familiar with the terminology* used by u-map in the GUI?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Informative Feedback

Q2.31 Do the instructions indicate clearly what you can do?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.32 Does u-map clearly inform you when an action is complete?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.33 Do the error messages clearly explain the error type and what caused them?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Explicitness

Q2.34 Is it clear what you have to do to complete an action?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Flexibility & control

Q2.35 Is the interface *fairly well structured*, in order to give you the feeling that you control the system?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.36 Do you feel that you can easily “undo” an action?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.37 Are the *irreversible actions*, that can be caused *by mistake*, doubleconfirmed?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Q2.38 Is *navigating back to the main application menu easy* from anywhere in the system?

- (NO) 1 2 3 4 5 (YES)
- I do not remember.
- I do not know/ I do not understand the question.

Bibliography

[1] CISCO, “Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2014–2019.” February 2015.

[2] P. Klasnja and W. Pratt, “Healthcare in the pocket: mapping the space of mobile-phone health interventions.” *Journal of biomedical informatics*, vol. 45, no. 1, pp. 184–198, 2012.

[3] I. Stojmenovic, “Fog computing: A cloud to the ground support for smart things and machine-to-machine networks.” in *Telecommunication Networks and Applications Conference (ATNAC), 2014 Australasian*. IEEE, 2014, pp. 117–122.

[4] S. G. Sakamoto, L. C. de Miranda, and H. Hornung, “Home control via mobile devices: State of the art and hci challenges under the perspective of diversity,” in *Universal Access in Human-Computer Interaction. Aging and Assistive Environments*. Springer, 2014, pp. 501–512.

[5] L. Wong, “Method of Operating a Dual SIM Wireless Communications Device.” Jun. 26 2007, uS Patent App. 12/303,772.

[6] B. Jiang and J.-F. Mao, “Design of an PIFA-IFA-monopole in dual-SIM mobile phone for GSM/DCS/Bluetooth operations.” in *Microwave and Millimeter Wave Technology, 2008. ICMMT 2008. International Conference on*, vol. 3. IEEE, 2008, pp. 1050–1053.

[7] G. Fortetsanakis and M. Papadopouli, “On multi-layer modeling and analysis of wireless access markets,” *IEEE Trans. Mobile Computing*, (to appear).

[8] G. Fortetsanakis, M. Papadopouli, G. Karlsson, M. Dramitinos, and E. A. Yavuz, “To subscribe, or not to subscribe: Modeling and analysis of service paradigms in cellular markets,” in *IEEE DySPAN*, 2012.

[9] G. Le Bodic, D. Girma, J. Irvine, and J. Dunlop, “Dynamic 3G network selection for increasing the competition in the mobile communications market.” in *Vehicular Technology Conference, 2000. IEEE-VTS Fall VTC 2000. 52nd*, vol. 3. IEEE, 2000, pp. 1064–1071.

[10] J. Konka, I. Andonovic, C. Michie, and R. Atkinson, "Auction-based network selection in a market-based framework for trading wireless communication services." *Vehicular Technology, IEEE Transactions on*, vol. 63, no. 3, pp. 1365–1377, 2014.

[11] C. Camaran and D. Miguel, "Mobile Virtual Network Operator (MVNO) Basics: What is behind this mobile business trend." *Valoris Telecommunication Practice, Tech. Rep*, 2008.

[12] BEREC, "Report on the Public Consultation on the Draft BEREC Medium Term Strategy Outlook." February 2012.

[13] BEREC, "Monitoring quality of Internet access services in the context of net neutrality. Draft BEREC report." March 2014.

[14] FCC, "FCC unveils new, free speed test app to empower consumers with U.S. mobile broadband performance information." November 2013.

[15] FCC, "FCC Speed Test (Android App)." https://play.google.com/store/apps/details?id=com.samknows.fcc&hl=en, September 2014.

[16] D. Milioris, G. Tzagkarakis, A. Papakonstantinou, M. Papadopouli, and P. Tsakalides, "Low-dimensional signal-strength fingerprint-based positioning in wireless lans," *Ad hoc networks*, 2012.

[17] G. Fortetsanakis, M. Katsarakis, M. Plakia, N. Syntychakis, and M. Papadopouli, "Supporting wireless access markets with a user-centric QoE-based geo-database," in *ACM MobiArch*, Istanbul, Turkey, 2012.

[18] M. Katsarakis, V. Theodosiadis, M. Dramitinos, and M. Papadopouli, "u-map: a user-centric QoE-based recommendation tool for wireless access markets," in *ACM S3*, Miami, FL, USA, 2013.

[19] M. Katsarakis, G. Fortetsanakis, P. Charonyktakis, A. Kostopoulos, and M. Papadopouli, "On user-centric tools for QoE-based recommendation and real-time analysis of large-scale markets," *IEEE Communications Magazine*, vol. 52, no. 9, pp. 37–43, 2014.

[20] M. Katsarakis, V. Theodosiadis, and M. Papadopouli, "Evaluation of a User-centric QoE-based Recommendation Tool for Wireless Access," in *ACM SIGCOMM Workshop on Crowdsourcing and crowdsharing of Big (Internet) Data (C2B(I)D)*, London, UK, 2015.

[21] N. Rapousis, M. Katsarakis, and M. Papadopouli, "QoWater–A crowdsourcing approach for assessing the water quality," in *Cyber-Physical Systems for Smart Water Networks (CySWater), 2015 ACM 1st International Workshop on*. ACM, 2015.

[22] "OpenSignal: 3G & 4GLTE Cell Map," http://opensignal.com/.

[23] R. Murty, R. Chandra, T. Moscibroda, and P. Bahl, "Senseless: A database-driven white spaces network," in *IEEE DySPAN*, 2011.

[24] M. Mishra and A. Sahai, "How much white space is there?" EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-3.

[25] J. van de Beek, J. Riihijarvi, A. Achtzehn, and P. Mahonen, "UHF white space in Europe - a quantitative study into the potential of the 470 - 790 MHz band," in *IEEE DySPAN*, 2011.

[26] A. Min, K.-H. Kim, and K. Shin, "Robust cooperative sensing via state estimation in cognitive radio networks," in *IEEE DySPAN*, 2011.

[27] D. Gurney, G. Buchwald, L. Ecklund, S. Kuffner, and J. Grosspietsch, "Geo-location database techniques for incumbent protection in the TV white space," in *IEEE DySPAN*, 2008, pp. 1–9.

[28] G. Karlsson, E. A. Yavuz, and G. C. Polyzos, "A future wireless Internet beyond generations," in *ACM CFI*, 2010.

[29] J. Lempiäinen and M. Manninen, "Umts radio network planning, optimization and qos management," *Dodrecht Kluwer Academic Publishers*, 2003.

[30] I. Samuel, K. Arora, and B. Narasimhan, "Location-based performance-measuring techniques in umts," *Bell Labs technical journal*, vol. 8, no. 2, pp. 15–32, 2003.

[31] J. Matamales, D. Martin-Sacristan, J. F. Monserrat, and N. Cardona, "Performance assessment of hsdpa networks from outdoor drive-test measurements," in *Vehicular Technology Conference, 2009. VTC Spring 2009. IEEE 69th*. IEEE, 2009, pp. 1–5.

[32] J. A. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," *Center for Embedded Network Sensing*, 2006.

[33] R. K. Ganti, F. Ye, and H. Lei, "Mobile crowdsensing: current state and future challenges," *Communications Magazine, IEEE*, vol. 49, no. 11, pp. 32–39, 2011.

[34] S. Gaonkar, J. Li, R. R. Choudhury, L. Cox, and A. Schmidt, "Micro-blog: sharing and querying content through mobile phones and social participation," in *ACM MobiSys*, 2008.

[35] J. Froehlich, M. Y. Chen, S. Consolvo, B. Harrison, and J. A. Landay, "Myexperience: a system for in situ tracing and capturing of user feedback on mobile phones," in *ACM MobiSys*, 2007.

[36] "RF Signal Tracker," https://sites.google.com/site/androiddevelopmentproject/home/rf-signal-tracker.

[37] A. Farshad, M. K. Marina, and F. Garcia, "Urban wifi characterization via mobile crowdsensing," in *IEEE NOMS*. IEEE, 2014.

[38] "Sensorly: Coverage Maps," http://www.sensorly.com/.

[39] E. Gregori, L. Lenzini, V. Luconi, and A. Vecchio, "Sensing the Internet through crowdsourcing," in *IEEE PerMoby*, 2013, pp. 248–254.

[40] S. Sen, J. Yoon, J. Hare, J. Ormont, and S. Banerjee, "Can they hear me now?: a case for a client-assisted approach to monitoring wide-area wireless networks," in *ACM IMC*, 2011.

[41] S. Rosen, S.-J. Lee, J. Lee, P. Congdon, Z. Morley Mao, and K. Burden, "MCNet: Crowdsourcing wireless performance measurements through the eyes of mobile devices," *IEEE Commun. Mag.*, 2014.

[42] "WeFi," http://www.wefi.com/.

[43] J. Bergstra and C. Middelburg, "ITU-T recommendation G. 107: The E-Model, a computational model for use in transmission planning," 2003.

[44] L. Duan, J. Huang, and B. Shou, "Duopoly competition in dynamic spectrum leasing and pricing," *IEEE Trans. Mobile Computing*, vol. 11, no. 11, pp. 1706–1719, 2012.

[45] D. Niyato, E. Hossain, and Z. Han, "Dynamics of multiple-seller and multiple-buyer spectrum trading in cognitive radio networks: A game-theoretic modeling approach," *IEEE Trans. Mobile Computing*, vol. 8, no. 8, pp. 1009–1022, 2009.

[46] K. U. R. Laghari and K. Connelly, "Toward total quality of experience: A QoE model in a communication ecosystem," *IEEE Commun. Mag.*, vol. 50, no. 4, pp. 58–65, 2012.

[47] ITU, "ITU-T recommendation P.862: Perceptual evaluation of speech quality PESQ: An objective method for end-to-end speech quality assessment of narrow-band telephone networks and speech codec," 2001.

[48] K. Mitra, A. Zaslavsky, and C. Ahlund, "Context-aware QoE modelling, measurement and prediction in mobile computing systems," *IEEE Trans. Mobile Computing*, 2013.

[49] A. Bhattacharya, W. Wu, and Z. Yang, "Quality of experience evaluation of voice communication: an affect-based approach," *Human-centric Computing and Information Sciences*, vol. 2, no. 1, pp. 1–18, 2012.

[50] M. Katsarakis, V. Theodosiadis, and M. Papadopouli, "On the evaluation of a user-centric qoe-based recommendation tool for wireless access," ICS-FORTH, Heraklion, Crete, Greece, Tech. Rep. 445, May 2014.

[51] L. Gao, X. Wang, Y. Xu, and Q. Zhang, "Spectrum trading in cognitive radio networks: A contract-theoretic modeling approach," *IEEE J. Selected Areas in Comm.*, vol. 29, no. 4, pp. 843–855, 2011.

[52] C. Meidanis, I. Stiakogiannakis, and M. Papadopouli, "Pricing for mobile virtual network operators: The contribution of u-map," in *IEEE DySPAN*, 2014.

[53] "u-map," http://www.ics.forth.gr/mobile/umap/.

[54] A. Pathak, Y. C. Hu, M. Zhang, P. Bahl, and Y.-M. Wang, "Fine-grained power modeling for smartphones using system call tracing," in *ACM EuroSys*, 2011.

[55] A. Carroll and G. Heiser, "An analysis of power consumption in a smartphone," in *USENIX ATC*, 2010.

[56] M. Dong and L. Zhong, "Self-constructive high-rate system energy modeling for battery-powered mobile systems," in *ACM MobiSys*, 2011.

[57] L. Zhang, B. Tiwana, Z. Qian, Z. Wang, R. P. Dick, Z. M. Mao, and L. Yang, "Accurate online power estimation and automatic battery behavior based power model generation for smartphones," in *CODES+ISSS*, 2010.

[58] M. B. Kjærgaard and H. Blunck, "Unsupervised power profiling for mobile devices," in *MobiQuitous*, 2012.

[59] C. Yoon, D. Kim, W. Jung, C. Kang, and H. Cha, "Appscope: Application energy metering framework for android smartphone using kernel activity monitoring," in *USENIX ATC*, 2012.

[60] "Powertutor," http://powertutor.org.

[61] "UI/Application Exerciser Monkey," http://developer.android.com/tools/help/monkey.html.

[62] B. A. Kitchenham, S. L. Pfleeger, L. M. Pickard, P. W. Jones, D. C. Hoaglin, K. El Emam, and J. Rosenberg, "Preliminary guidelines for empirical research in software engineering," *IEEE TSE*, 2002.

[63] L. Wilkinson, "Statistical methods in psychology journals: guidelines and explanations." *American psychologist*, 1999.

[64] I. Tsamardinos and G. Borboudakis, "Permutation testing improves bayesian network learning," in *ECML PKDD*, 2010.

[65] A. Ahmad and L. Dey, "A k-mean clustering algorithm for mixed numeric and categorical data," *Data & Knowledge Engrg*, 2007.

[66] E. Grigoroudis and Y. Siskos, "Musa: A decision support system for evaluating and analysing customer satisfaction," in *PCI*, 2003.

[67] P. Charonyktakis, M. Plakia, I. Tsamardinos, and M. Papadopouli, "On user-centric modular QoE prediction for VoIP based on machine-learning algorithms," *IEEE Transactions on Mobile Computing*, 2015.

[68] K. Karagiannaki, S. Honianakis, A. Panousopoulou, and M. Papadopouli, "mMamee: A mHealth Platform for Monitoring and Assessing Maternal Environmental Exposure," in *Computer-Based Medical Systems (CBMS), IEEE International Symposium on*, 2015.