1st International Workshop On Services For Mobile Data Collection

Referenceable mobile crowdsensing architecture: A healthcare use-case

Muntazir Mehdi[a], Guido Mühlmeier[b], Kushal Agrawal[c], Rüdiger Pryss[c],
Manfred Reichert[c], Winfried Schlee[d], Franz J. Hauck[a]

[a]Institute of Distributed Systems, Ulm University, Germany
[b]Department for ENT/HNS, Federal Armed Hospital of Ulm, Germany
[c]Institute for Databases and Information Systems, Ulm University, Germany
[d]Department of Psychiatry and Psychotherapy, University of Regensburg, Germany

Abstract

Smartphones have become an integral part in life of users, mainly because over the course of recent years, they have become extremely mainstream, cheap, flexible, and they pack high-end hardware that offers high computational capabilities. Many, if not all of today's smartphones are equipped with sophisticated sensors which enable smart mobile sensing. The programmable nature of these sensors in the smartphones enable a wide array of possibilities to achieve user-centric or environmental sensing. Even though there have been different approaches proposed to develop a smartphone app, platform, design frameworks, APIs, and even application-specific architectures, there is a lack of generalized referenceable architecture in the literature. In this paper, we propose a generic reference architecture, which can be derived to create more concrete mobile sensing or mobile app architectures. Furthermore, we realize the proposed reference architecture in a healthcare use-case, specifically in the context of applying smart mobile sensing to support tinnitus research.

© 2018 The Authors. Published by Elsevier B.V.
Peer-review under responsibility of the Conference Program Chairs.

Keywords: mobile crowdsensing ; mhealth ; tinnitus ; reference architectures

1. Introduction

Smart mobile phones have changed the way information is perceived. Specifically from the perspective of a smart phone user, smart phones have become personalised, mainstream, flexible, and cheap. The current generation phones have drastically evolved from merely being phones to ubiquitous computing and communication devices[1]. Today, most of the mainstream smart phones are embedded with sensors, which enable sensing motion (using accelerometers, gravity sensors, gyroscopes, and rotational vectors), environment (using barometer, thermometer, and photometer), and position (using GPS, orientation sensors, and magnetometers). Furthermore, in addition to packing state-of-the-art computing hardware, these smart phones are equipped with sophisticated communication technologies such

as Bluetooth, Cellular, Wifi, NFC that enable them to communicate with other smart devices (for instance external sensors, laptops, tablets).

The combination of diverse set of sensors, sophisticated hardware, and ease of communication make smart mobile phones an important tool in both research and industry. People coming from different aspects of research and industry are readily interested in leveraging the potentials of mobile sensing or mobile crowdsensing[2] in their domains. Specifically, mobile crowdsensing is applied in two different modes, 1) Opportunistic[3], and 2) Participatory[4]. Herein, opportunistic sensing aims at keeping the user involvement to the minimum, while participatory sensing thrives on the continuous input from the user. Regardless of the user involvement, both opportunistic and participatory sensing modes can be applied on a level of an individual, a community, or on a larger scale (urban-scale)[5]. Furthermore, the applications of mobile crowdsensing are mainly categorised as People-centric[6] or Environment-centric[7].

Irrespective of the mode, level, or category of application of mobile crowdsensing, it is more advantageous than traditional sensor networks. Primarily, because of high and cheap availability of smart mobile phones, secondarily, because of almost negligible deployment costs[8]. In addition to this, from perspective of a research community, mobile crowdsensing offers affordable means for conducting surveys as well as performing large-scale studies to better understand the user behaviour or preferences[9]. Be it industry or research community, the interest in developing mobile crowdsensing applications is high, mainly because of extensive and effortless access to the mobile application development tools, frameworks, and application distribution mechanisms. This is why we see a myriad of existing mobile crowdsensing applications, and high interest in developing new. Even though, the applications of mobile crowdsensing resonate in different sectors of our economy, for instance from healthcare, e-governance, education, environmental monitoring, to transportation, the challenges faced by mobile crowdsensing are not limited[10].

Addressing these challenges, a number of frameworks, architectures, and design strategies exist in the literature. However, in order to achieve an optimal architecture to address the challenges posed to mobile crowdsensing, we believe that a Reference Architecture (RA) will sustain the process of deriving more concrete architectures. In this paper, we propose a criterion to formulate an RA, and propose a Reference Architecture strictly within the bounds of mobile crowdsensing. Furthermore, we support the idea of our proposed RA using a use-case of applying mobile crowdsensing in healthcare domain, specifically, to support patients suffering from Tinnitus. The general concept behind the RA, its importance and the criterion is further detailed in Section 3. The motivational scenario driving the proposed work and importance of mobile crowdsensing in the healthcare domain is detailed in Section 2. And before we conclude our paper in 5, the proposed reference architecture in explained in Section 4.

2. Motivational Scenario

The proposed work presented in this paper is motivated and driven by the needs of tinnitus research within the context of European School for Interdisciplinary Tinnitus Research (ESIT) Project. Among many objectives of the ESIT project, one core goal is to develop a generic, robust and flexible middleware for mobile crowdsensing to monitor real-time measurements of tinnitus-related parameters as well as electroencephalographic and physical activities.

Tinnitus is a common disorder, which is associated with the perception of a ringing sound or noise in the ears. The causative factors of tinnitus are unknown, however, it usually is intertwined with an underlying condition in the ear. In addition to general health complications, tinnitus might be also responsible for provoking other psychological disorders (stress, anxiety, depression, or obsessive-compulsive disorder) and may affect the common as well as social lifestyles. Furthermore, tinnitus has also been described to may have a direct proportionality with migraine and vertigo. The perception of tinnitus is driven by the filtering work of subconscious areas in the brain-stem influenced by the limbic system, and increases in episodes of interior or exterior stress. Changes in the atmospheric surrounding may have direct or indirect effects on the annoyance caused by tinnitus.

In this context, a plethora of scientific literature has reported an alleviation in the tinnitus condition, which can be related to the surrounding atmospheric and environmental conditions of the patient[11]. Among others, factors like a decrease in the atmospheric pressure, a weather change (specifically rainy and cold weather), intensity of light, the current sound environment, or (sudden) change in altitude are some of the most pertinent. For instance, weather is no causative factor for tinnitus, but may take influence on the perception of the sounds. Very often patients address that they hear intermittent secondary tones to their common tinnitus sound or describe their common tone changing from compensated to annoying expressed by terms like hammering, beating, fizzling. Similarly, air pressure

is considered to be responsible for episodes of migraine in many patients[12]. In rare cases patients describe a change in tinnitus perception corresponding to their migraine in major air pressure changes. On the other hand, the orientation, movement speed, light intensity, and direction of movement of the patient are some of the less common, yet still significant set of factors that may also induce a spike in tinnitus symptoms.

Mobile crowdsensing can be applied to monitor the aforementioned atmospheric factors (most, if not all) that provoke tinnitus. Moreover, the raw data gathered with these sensors can be processed to identify (drastic) changes in the circumstances (motion, environment, and position) of a tinnitus patient. In conclusion, we believe that a sophisticated mobile crowdsensing app monitoring patient surroundings and profiling patients for personalised therapy may foster controlling and mitigating of tinnitus symptoms, as well as promote community or participatory sensing.

3. Reference Architecture

There are various ways to define an RA as there is no straightforward definition of it[13]. In conclusion, an RA is a diagram, or a pattern, or specification or set of diagrams, patterns, or specifications that; 1) depict the administration of system functions among components in the infrastructure and 2) provide a map for how those functions inter-relate. Given the understanding that RAs are becoming an integral part of software design and planning, few questions arise, where does the RA incept from and how do RAs evolve to become a necessity? The increase in complexity of applications that cater with the current business needs in organizations and enterprises can answer the confusions related to evolution and inception of RA. The increased complexity in development and implementation of distributed systems to achieve high level of interoperability, by creating robust components within the systems is the basis of high involvement of RAs in the software design process. Hence these RAs can be used as a mechanism for: 1) the development of concrete architectures, 2) standardizing tool that guarantees the interoperability between systems, 3) standardizing tool that guarantees the interoperability between system components by validating the original purposes and 4) making sure that the basic requirements, specified during the problem definition phase, were addressed[14].

A Reference Architecture is a combination of *Business Architecture*, *Technical Architecture*, and *Customer Context*[14]. Therefore, we can infer that an RA ensures that end-user and participants have the confidence to deploy the technology. With this in mind and the aforementioned definition of RA, we can safely conclude the high reliance of software applications on RAs. However, the guidelines or a criterion on which a RA is selected or defined for a particular derivation of concrete architectures is not very well established. Herein, during the review of existing literature on RAs, we found a minimalist selection criterion for a good RA. Some of the highlights of the criterion presented in[14] dictate that, the RA should be understandable for all stakeholders (customers, product managers, project managers, engineers etc.). An RA should add value to the business by providing consistent models, capabilities, and equipment. It should be of satisfactory quality, up-to-date, and maintainable. Furthermore, an RA must address the key issues of the specific domain.

Fig. 1. TinnituSense Reference Architecture
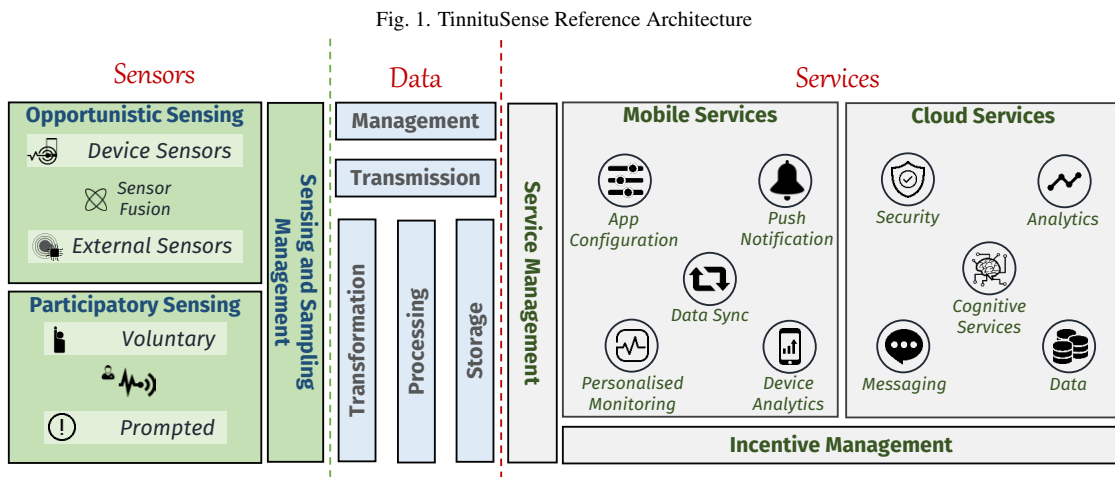
In addition to this, from our understanding of RA and its importance, we also define or extend the discussed criterion for RA. As we already know that an RA provides a template for architecture of a particular domain, that is, it provides a set of functions with their respective interfaces and interfaces for other domains to communicate with it. Therefore, the *level of abstraction* at which an RA is selected for a particular domain plays a vital role. While considering the *level of abstraction* as a pivotal point for RA, the generalization of a system with respect to 1) itself, 2) the subsystems and 3) other domains should also be considered [13]. *Context* is another consideration that has proved to be vital for an RA [15]. While dealing with *context* for RA, the aspects of *design and application context* that affects the business goals and design of RA should be investigated. The investigation is a result of answers obtained by answering some basic questions like 1) Where will it be used? 2) Who defines it? and 3) When is it defined? Since *context* might affect the main goals of RA, therefore *Goal* should also be selected as a consideration for RA criterion. The *goal* consideration is investigated by addressing the main intentions of use of a particular RA, that is, why a particular RA has been defined? And does it address its major purpose? Furthermore, the *Design* consideration is the most important consideration for an RA in order to encompass the major responsibilities and purposes of its usage. In *design* consideration a specification is formed which contains information regarding the main RA itself, level of concreteness and the way it is represented.

4. TinnituSense Reference Architecture

Based on the general understanding, and conceptual analysis of the RA, as well as the proposed criterion for defining RAs, we propose a layered RA targeted specifically for mobile crowdsensing, called *TinnituSense*. The layered RA is given in Figure 1, and is composed of three main layers, namely, 1) Sensors, 2) Data, and 3) Services. Details on the individual constituents of the layered RA are further discussed in the following subsections.

4.1. Sensors

The entire principle of mobile crowdsensing revolves around the sophisticated, and diverse sensors embedded in most of the mainstream smart phones. As previously mentioned, these sensors are capable of sensing motion, environment, and position. Within the sensors layer of our RA, we propose two sublayers of Oppotunistic Sensing, and Participatory Sensing. This is mainly done in order to define a clear distinction between the two modes of sensing. Furthermore, an additional layer of 'Sensing and Sampling Management', which encompasses both sensing modes is also shown in Figure 4. This sublayer is mainly responsible for implementing services which are responsible for controlling the sensors, specifically to activate and deactivate the sensor listeners. This is mainly done to preserve the power consumption that results from continuous activation of sensors. In addition to this, the rate at which the sensors sample the data also plays a vital role to manage the overall battery consumption of the smart device. Therefore, we propose that the services responsible for implementing sampling scheme (for example, adaptive, fixed, timed, or on demand sampling) should be consolidated with services responsible for controlling or managing sensors.

Both opportunistic and participatory sensing modes should be implemented to acquire data from the user. The main target of opportunistic sensing should be to completely remove the user involvement and acquire sensor data in an automated fashion. Herein, since the sensors embedded in smart phones, even though programmable, and state-of-the-art are still general-purpose, therefore, there is a chance of accuracy-related problems. This is therefore, we propose that the output of these sensors should be fused together with more specialised sensors (for instance, in case of temperature sensor, the output can be fused with sensors from the nearest met-office). In case of participatory sensing, we propose that the services should be implemented with the sole purpose of involving user in the data acquisition process. Herein, the user can provide data in two different ways, 1) Voluntarily, whenever an interesting event happens, the user inputs the data using the app (for example, considering the use-case of tinnitus, a user filling out an explicit questionnaire⋆ when there is a sudden rise in tinnitus symptoms). 2) Prompted, on continuous monitoring the surroundings of the user, the device detects an interesting event and prompts user for the input (for example, in scenario of tinnitus use-case, a sudden fall in the atmospheric pressure detected by the barometer pressure prompts the user to fill out the questionnaire⋆.

⋆ The questionnaire or survey will be designed based on the Tinnitus Handicap Inventory [16]

4.2. Data

The backbone of our RA is the data management, which is done on both mobile device and the cloud (a light term which specifically refers to the back-end of the entire system). By data management, we refer mainly to acquiring, transforming, and processing the data from its raw source of sensors. Data acquisition, as explained earlier in this Section, is mainly done on the mobile device. The raw data, coming directly out of the sensors may be pre-processed and transformed on the mobile device in order to focus on relevant data and interesting events. Some sensors may read a huge amount of data which cannot be immediately sent to cloud services. In an ideal scenario, this large amount of data cannot be processed entirely on the mobile device, mainly because in doing so will exhaust the mobile device resources. Therefore, there has to be a fine balance between data processing on mobile device and cloud needs to be identified. Such pre-processing (which is done on the mobile device) may consist of aggregation, filtering, compression, and even first analytical functions. One example for the latter is the detection of interesting phases by identifying trigger events in the data stream. Before the data (processed or raw) is forwarded to cloud services, its storage has to be managed on the mobile device. Since storage, communication bandwidth, and processing capacities on the mobile device may be scarce due to limited memory resources, fluctuating, and cost-relevant network connectivity, and bounded available energy, the data management layer has to decide on processing, storage, transformation and transmission in a controlled and user-aware way.

Data transmission connects the mobile device to back-end services. Transmission has to preserve data formats, especially timing information, in order to correctly identify the data items and their meaning at receiver side. In addition to this, the data transmission services have to be aware of the connectivity to the cloud. The transmission layer is thus responsible for sustaining the balance between processing the data on device and cloud, as well as finding a proper balance between amount of data that is transmitted to the cloud. On the cloud, data has to be managed, too. Again, data storage, processing, and transformation are the basic building blocks of data management in our RA. Even though, in the cloud, storage capacity is no longer scarce but, the complexity here is that there is a large amount of data that has to be stored in a way that data can be retrieved by cloud services in a scalable way. Processing and transformation is mainly done by analytical services within the cloud.

4.3. Services

The services of the proposed RA are divided into mobile services and cloud services. Mobile services are mainly executed on the mobile device (and practically reside on the mobile device), whereas cloud services are mainly hosted on one or multiple cloud systems, and may even be provided by different associated providers. Further, mobile services are typically personalised and work for their particular user, whereas cloud services typically have many users at the same time and are more generic.

Mobile Services: App Configuration service allows the user to configure the sensing app and may be associated with system configurations within the mobile device. *Push Notifications* allow cloud services to raise attention of the mobile user, for instance, to notify user of an interesting event, or to fill in a questionnaire of an app. *Data Synchronisation* services allow to backup mobile data, e.g. configurations, preferences, files, in the cloud. *Personalised Monitoring* are services on the mobile that provide personalised feedback about the user activity. There could be services monitoring the usage of the mobile, the battery consumption, the mobility of its user, etc. It also can be a domain-specific service for e-health that in principle is configured by the user and/or by therapists from within backend services to monitor certain events of the user in order help on medical conditions, as for example in case of tinnitus. Finally, *Device Analytics* gives the user feedback about the behaviour of his device. Examples are battery and network usage, but also domain-specific analytics like frequency and number of prompted user interactions.

Cloud Services: Security services employ authenticated provision and authorised access to data. Furthermore, these services should be capable enough to address the data policies, secure transmission, as well as storage of data. The *Analytics* services comprise basically all kind of analytical evaluations of the user's data with respect to the domain specific setting, for instance in our use-case scenario of tinnitus, the user gets the information about the critical environmental factors that has impacted the rise in symptoms in past. This can be both analytics on a per user basis that can be pushed to the user as personal feedback, and analytics in comparison of the user's data to the data sets of different users. *Cognitive Services* are a mesh of tools combining the use of machine-learning capabilities and artificial intelligence to enable intelligent features. For instance, in tinnitus use-case, detection of the possible context

and surrounding (e.g. shopping, driving, riding) to analyse and associate the noise levels. A *Messaging* service enables the user to have a direct interaction with other users of the system, for instance, a tinnitus patient can directly message his therapist about his concerns regarding personalized therapies. It can be used to establish an easy-to-use communication channel among users and between users on one side and support staff on the other side. Messaging can even be used in combination with cognitive services, e.g., automatic and AI-based help services. *Data Services* store and provide access to the sensed raw data as well as processed, aggregated and analysed data in databases.

Incentive Management is a sublayer primarily responsible for incentivizing both users and developers of the app. From the perspective of the user, this layer should dictate that the user receives all necessary incentives to persuade them to keep using the app and provide quality data. Furthermore, based on the users preferences of the app usage, the sensor data, the developers should improve the quality of the provided services, identify new services, and formulate frameworks to offer additional incentives. And from the perspective of developer, it should enable the developers to balance between the development time of the app and the quality of mobile crowdsensing data.

5. Conclusion

We presented an RA for mobile crowdsensing applications, which we refer to as TinnituSense Reference Architecture. Our main motivation is an application to help tinnitus patients and therapists to gain new insights to the disease and its user-related characteristics. It is also worthwhile to mention that within our motivational scenario, we highlight the importance of mobile crowdsensing in healthcare domain, specifically, the importance of mobile crowdsensing in tinnitus research and how it can play a pivotal point in mitigating the symptoms of tinnitus. In addition to the TinnituSense RA, we also propose a generic criterion that could lead to the design and selection of any particular type of RA in domain related implementations. The presented TinnituSense RA has multiple layers, sensing, data management, and services that we further dissected and explained. Herein, it is pertinent to mention that we also introduce the term 'Prompted Participatory Sensing' within the layer of Sensors. Even though, the proposed RA is targeted for mobile crowdsensing within healthcare, we argue that it can be easily tailored for other application domains.

References

1. N. D. Lane, E. Miluzzo, H. Lu, D. Peebles, T. Choudhury, A. T. Campbell, A survey of mobile phone sensing, IEEE Communications magazine 48 (9).
2. R. K. Ganti, F. Ye, H. Lei, Mobile crowdsensing: current state and future challenges, IEEE Communications Magazine 49 (11).
3. L. D. L. V. Duc, J. Scholten, P. J. Havinga, Towards opportunistic data dissemination in mobile phone sensor networks, in: The Eleventh International Conference on Networks, ICN 2012, International Academy, Research, and Industry Association (IARIA), 2012.
4. J. A. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, M. B. Srivastava, Participatory sensing.
5. A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, People-centric urban sensing, in: Proceedings of the 2Nd Annual International Workshop on Wireless Internet, WICON '06, ACM, New York, NY, USA, 2006. doi:10.1145/1234161.1234179.
   URL http://doi.acm.org/10.1145/1234161.1234179
6. A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, G.-S. Ahn, The rise of people-centric sensing, IEEE Internet Computing 12 (4).
7. S. S. Kanhere, Participatory sensing: Crowdsourcing data from mobile smartphones in urban spaces, in: Mobile Data Management (MDM), 2011 12th IEEE International Conference on, Vol. 2, IEEE, 2011, pp. 3–6.
8. D. Christin, A. Reinhardt, S. S. Kanhere, M. Hollick, A survey on privacy in mobile participatory sensing applications, Journal of systems and software 84 (11) (2011) 1928–1946.
9. T. S. Behrend, D. J. Sharek, A. W. Meade, E. N. Wiebe, The viability of crowdsourcing for survey research, Behavior research methods 43 (3) (2011) 800.
10. J. He, K. Kunze, C. Lofi, S. K. Madria, S. Sigg, Towards mobile sensor-aware crowdsourcing: Architecture, opportunities and challenges, in: International Conference on Database Systems for Advanced Applications, Springer, 2014, pp. 403–412.
11. W. Schmidt, C. Sarran, N. Ronan, G. Barrett, D. J. Whinney, L. E. Fleming, N. J. Osborne, J. Tyrrell, The weather and menieres disease: a longitudinal analysis in the uk, Otology & neurotology: official publication of the American Otological Society, American Neurotology Society [and] European Academy of Otology and Neurotology 38 (2) (2017) 225.
12. K. Kimoto, S. Aiba, R. Takashima, K. Suzuki, H. Takekawa, Y. Watanabe, M. Tatsumoto, K. Hirata, Influence of barometric pressure in patients with migraine headache, Internal Medicine 50 (18) (2011) 1923–1928.
13. R. Cloutier, G. Muller, D. Verma, R. Nilchiani, E. Hole, M. Bone, The concept of reference architectures, Sys. Eng. 13 (1) (2010) 14–27.
14. G. Muller, A reference architecture primer, Eindhoven Univ. of Techn., Eindhoven, White paper.
15. A. Bürkle, W. Müller, U. Pfirrmann, Towards a reference architecture for context-aware services, Adv. in Human-Comp. Interaction 31.
16. C. W. Newman, G. P. Jacobson, J. B. Spitzer, Development of the tinnitus handicap inventory, Archives of Otolaryngology–Head & Neck Surgery 122 (2) (1996) 143–148.