Pervasive and Mobile Computing 10 (2014) 66–82

Contents lists available at ScienceDirect

Pervasive and Mobile Computing

journal homepage: www.elsevier.com/locate/pmc

Service-oriented middleware for large-scale mobile participatory sensing

Sara Hachem *, Animesh Pathak, Valerie Issarny

Inria Paris-Rocquencourt, France

ARTICLE INFO

Article history:
Available online 26 October 2013

Keywords:
Participator sensing
Scalability
Probabilistic registration
Truncated Lévy walk

ABSTRACT

In this paper, we introduce MobIoT, a service-oriented middleware that enables large-scale mobile participatory sensing. Scalability is achieved by limiting the participation of redundant sensing devices. Precisely, MobIoT allows a new device to register its services only if it increases the sensing coverage of a physical attribute, along its expected path, for the set of registered devices. We present the design and implementation of MobIoT, which mobile devices use to determine their registration decision and become accessible for their services. Through experiments performed on real datasets, we show that our solution scales, while meeting sensing coverage requirements.

1. Introduction

In its traditional definition, *Participatory Sensing*, also known as mobile crowdsensing [1] assigns mobile devices to form participatory sensor networks that enable public and professional users to gather, analyze and share local knowledge [2].

In existing participatory sensing solutions, it is common to adopt a decoupled approach where the sensing of and the querying for information operations are *asynchronous* (as seen in solutions presented in [3]), i.e., the sensor is tasked to sample the real world periodically, continuously or based on some event, etc., but is not triggered by a user's request. In many cases, the latter leads to more appropriate results. Indeed, sensing the world independently of the measurement request time and intended use of the data is not necessarily the best option; such approach may lead to large out-dated streams of data that can be of low benefit while more useful up-to-date measurements are missing. The issue of data freshness versus response time arises in this case. Our position is that many scenarios can handle a slower response as long as data is more up to date. Further, most existing solutions in the participatory sensing field are domain-specific, designed to monitor either the environment, traffic, health, or social interactions, etc. [3]. However, with the increasing popularity of participatory sensing, providing common support to different sensing applications belonging to different categories becomes a must, as it saves development efforts and time.

To that end, our work focuses on domain-agnostic support for large-scale mobile participatory sensing. When performing large-scale sensing – where sensing devices are not necessarily in the geographical vicinity of the requestors – adopting a distributed approach does not perform well since routing and communication costs will grow unmanageable when millions of devices become involved. Therefore, the information about devices willing to provide their sensing measurements at the time of the query must be made available somewhere, e.g., in a registry. Hence, we augment the definition of participatory sensing with the identification of two main components:

1. *Mobile devices*: Computers, smartphones, tablets, etc., hosting different sensors and actuators. Those devices are numerous and have limited energy resources.

* Corresponding author. Tel.: +33 615724357.
E-mail addresses: sara.hachem@inria.fr (S. Hachem), animesh.pathak@inria.fr (A. Pathak), valerie.issarny@inria.fr (V. Issarny).

1574-1192/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.pmcj.2013.10.010

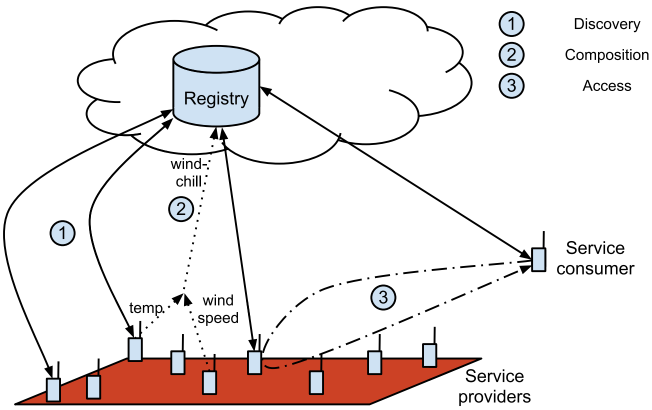

Fig. 1. The operations of a service-oriented mobile participatory sensing system.

2. *Registry*: Web-based component that holds metadata of sensing services. The component has performance capacity directly proportional to monetary cost.

The two components interact with each other in two phases (shown in Fig. 1):

1. *Registration*: Each mobile device is expected to periodically advertise its services to the Registry which holds their metadata.
2. *Lookup and access*: When data about the physical environment of a remote region is required, the Registry is queried for devices with the relevant sensing services to then access them and acquire their measurements, which can be processed to obtain the desired results.

For example, consider a query of the form "what is the average crowd level at the *Jardin de Tuileries* in Paris right now?". The answer can be generated by sensing the sound level surrounding each mobile device at the location of interest and using it to compute the crowd level through some mathematical expression. Although the solution above seems straightforward, there are several important challenges to address (for a detailed discussion, we refer readers to [4]). The first issue to overcome is the *heterogeneity* of devices as participatory sensing networks are bound to contain devices from an assortment of vendors, with varying sensing characteristics such as error distributions, sampling rates, spatial resolution, and so on. Second, the network is characterized by an *unknown and dynamic structure* resulting from the mobility of devices [5]. Last but not least, any approach must be *scalable* as mobile participatory sensing involves millions of mobile devices, a number that has important repercussions on communication, storage, memory and time consumption costs.

To address the above challenges, we adopt the service-oriented paradigm [6,7] to decouple the high level sensing system logic from the low level network heterogeneity and mobility issues. We provide a service-oriented middleware that supports the following, traditional functionalities: *Discovery* of, *Composition* of, and *Access* to *sensing services* (Fig. 1). Sensing services are, in this context, software logic that abstracts sensors and have functional (e.g., the attribute to measure) and non functional (e.g., measurement accuracy) properties. *Discovery* is used by devices wishing to provide measurements of the real world, to publish (register) their sensing services in registries that hold service metadata and to look for sensing services that can satisfy a sensing request. *Composition* of services is used when discovered sensing services are unable to fulfil the request. In such case, other registered sensing services are combined to provide a new, more convenient functionality. The composed services can further be used for more complex compositions, all of which should be defined, then executed in a composition engine provided by the middleware. Finally, *Access* enables interaction with discovered services.

Traditional approaches in mobile participatory sensing require input from as many sensors as possible to have the most accurate information about the real world [8]. We postulate that this should not be the governing rule for two main reasons. First, given today's abundant sensor availability through smart mobile devices, involving them all to partake in a sensing task provides redundant measurements that do not necessarily benefit the sensing quality noticeably, yet increase the communication and energy consumption costs. Second, in many cases, having the most accurate answer is not always the priority. In some situations, users might be satisfied with less accurate answers that have a "good enough" geographical sensing coverage of the real world. Consequently, this allows us to address the data freshness versus response time issue, mentioned earlier, that on-demand sensing may face. That is because sensing tasks triggered by users' requests can now provide up-to-date data with lower measurement latency as less devices are involved, and therefore less communication and computation efforts must be paid. Our contribution builds on that logic as we design our middleware to limit the number of participating sensing devices, based on the devices' mobility characteristics, while satisfying required sensing coverage through a revisited version of service discovery, namely in registration. Sensing coverage refers to the geographical area covered, i.e., sensed by a sensing device.

We first provide an approach that determines, with high accuracy, if the path of a new device will intersect with those of registered devices by acquiring information on their mobility. It is then possible to determine, from the computed result,