Sustainable Computing: Informatics and Systems 8 (2015) 47–59

Contents lists available at ScienceDirect

Sustainable Computing: Informatics and Systems

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

Energy-aware collaborative sensing for multiple applications in mobile cloud computing

Radhika Loomba[a], Lei Shi[a], Brendan Jennings[a,*], Roy Friedman[b], John Kennedy[c], Joe Butler[c]

[a] TSSG, Waterford Institute of Technology, Ireland
[b] Department of Computer Science, Technion, Israel
[c] Intel Labs Europe, Ireland

ARTICLE INFO

Article history:
Received 1 October 2014
Received in revised form
17 September 2015
Accepted 21 September 2015
Available online 8 October 2015

Keywords:
Mobile cloud
Collaborative sensing
Information aggregation
Frequent pattern mining
Optimization
Human mobility

ABSTRACT

Modern mobile handsets and the myriad of wearable devices connected to them offer a wide range of sensing capabilities. The ubiquity of such sensing devices offers the potential to realise novel applications based on collaborative sensing, in which application logic makes use of sensor input from a number of handsets, typically distributed across a defined physical area. Such applications will be enabled by mobile cloud computing, with the devices transferring raw or pre-processed sensed data to application logic hosted in the cloud. This results in a trade-off between the quality of the sensed data received by applications and the energy required to transfer data from the mobile handsets. We address this trade-off by considering a scheme in which a collaborative sensing middleware mediates between multiple applications requiring sensed data and the mobile handsets located within a particular physical area. We present and evaluate an algorithm which seeks to maximise the degree to which sensed data transferred from a given mobile device can be served to more than one application. We show that this algorithm leads to better overall performance in terms of energy used than an algorithm which does not aggregate sensed information between applications.

1. Introduction

The increasing interest in the potential to offer applications making use of sensed data gathered in a collaborative manner by a number of devices, is propelled by the focus of equipment manufacturers to incorporate sensors such as cameras, accelerometers, thermometers and compasses into smartphones, tablets and wearable devices. Researchers have investigated the potential to use processed sensed data for applications in healthcare, entertainment, environmental monitoring and many other domains [1]. However, in many cases the resource limitations of mobile devices creates a significant limitation for the realisation of envisaged applications. Mobile Cloud Computing (MCC) offers a promising solution, by offloading data gathered at mobile devices to data centre hosted servers for computation and storage. As mobile applications become increasingly computation-intensive this approach offers the benefits of conserving mobile handset resources (including energy) whilst meeting application performance targets [2,3].

Several mobile sensing systems have been designed with energy efficiency in mind [4–7]. Energy efficiency has also been a major concern for wireless sensor networks, where the application of in-network aggregation and compression techniques has been widely explored [8–10]. A key concern with such systems is to balance the trade-off between the accuracy of the information received by the application logic with the volume of data offloaded (at significant energy cost) by the mobile devices. When many mobile handsets collaborate to provide data to one or more applications, one way to minimise offloaded data volume is to ensure that redundancy in this data is eliminated. Thus, an important technical challenge for collaborative sensing systems is the optimal selection of sensors from mobile devices for accurately satisfying application constraints in an energy-efficient manner. In particular, when devices provide data to more than one application it is difficult to ensure that all applications receive data relating to physical area(s) they target, and that they receive updates at the requisite frequency. This difficulty is exacerbated by the fact that handsets are typically mobile, so the sensing system needs to embody models of human mobility patterns in order to make predictions of device locations.

* Corresponding author. Tel.: +353 51 302917.
E-mail addresses: r.l.loomba@ieee.org (R. Loomba), lshi@tssg.org (L. Shi), bjennings@tssg.org (B. Jennings), roy@cs.technion.ac.il (R. Friedman), john.m.kennedy@intel.com (J. Kennedy), joe.m.butler@intel.com (J. Butler).

http://dx.doi.org/10.1016/j.suscom.2015.09.001
2210-5379/

We present a novel approach to reducing the number of activated mobile devices in a collaborative sensing system in which data served by a given device can be used by multiple applications. Our main contribution is an algorithm that seeks to identify the best combination of sensors to activate on the available devices, taking into account those devices' predicted locations, energy status and the requirements of multiple applications. The algorithm, which we refer to as 'Info-Aggregation,' applies frequent pattern mining to find a trade-off between energy efficiency and the volume of data offloaded from mobile devices. We compare its performance to an algorithm which does not recognise that there is potential to reuse the same sensed data for multiple applications. An early version of this paper appeared as [11]; the scenario addressed there made the simplifying assumption that devices were static. Here we relax that assumption and update our algorithm specification to include a prediction of the mobile device locations based on a model of human mobility patterns. Gonzalez et al. [12] have shown that human mobility can be compared to a Levy walk with heavy-tail flight distribution and Rhee et al. [13] have presented a truncated Levy walk model for the same. We use this truncated Levy walk mobility model for prediction of the mobile device locations. This mobility prediction helps in sharing resources between co-located mobile devices.

The paper is structured as follows. Section 2 describes the related work in the fields of mobile and wireless sensing networks, aggregation, frequent pattern mining and mobility models. Section 3 specifies our scenario and the mathematical formulation of the problem. Section 4 discusses the algorithm design using aggregation for sensor allocation and the mobility model used in this algorithm. Section 5 details the experimental setup and presents our results. Finally, Section 6 summarizes the paper.

2. Related work

The proliferation of mobile applications and their rapid provisioning due to cloud computing is motivating significant work regarding development of platforms for MCC. In these platforms data centre hosted servers are utilised for data storage and processing, enabling better performance for applications accessed via mobile devices. However, reliance on the cloud for computation can be costly due to data transmission overheads and performance degradation due to limited bandwidth availability. Given this, MCC platforms such as MAUI [14], CloneCloud [15] and ThinkAir [16], seek to balance the trade-off between local processing costs on mobile devices and the cost of transfer of data to/from data centre hosted servers. These platforms select parts of the mobile application to offload into the cloud server but do not consider collaboration between devices for further energy savings.

Many applications accessed via mobile devices rely on data sensed by the device itself via embedded sensors; comprehensive survey work on mobile sensing is provided by Lane et al. [1], Khan et al. [17] and Campbell et al. [18,19]. Sensing applications like CareSafe [20], BikeNet [21], CenceMe [22], NoiseTube and MobiShop [17] amongst several others use the capability of multiple sensors surrounding the human to send context feedback to mobile users. Efforts to make mobile sensing systems more energy-efficient range from low-power processors [4], sensing pipelines [6], bidirectional feedback pipelines [5] and context correlation with association rules [7]. Energy-accuracy trade-offs [23] have also been considered, with techniques such as mobile tethering with cloud gatherers and an energy-aware striper [24], multiple access links for collaborative downloading and gateways [25], and data prefetching [26,27] having been investigated. In contrast, we consider energy saving by multitasking the capability of a single mobile device to offload sensed data for more than one application. By using multiple sensors from the same mobile device instead of low-power or low-accuracy sensors, we reduce the number of mobile devices that are actively losing energy in the environment.

Collaborative sensing has also been explored by researchers. iCoMe [28] is an incentive-based cooperative resource management technique that focuses on increasing the revenue for the service provider. In contrast, our collaborative middleware focuses on satisfying multiple application requirements by using a small number of mobile devices. Serendipity [29] and Mobile Device Cloud [30] utilize connected wireless networks to use computational resources of other devices in proximity to a mobile device. Penner et al. [31] have presented 'Transient Clouds', a cloud-on-the-fly approach for collaboration between devices while AnonySense [32,33] considers privacy and threats involved in such collaborative systems. Similar to these techniques, we consider collaboration between mobile devices. However, we consider multiple applications requesting sensed data and devices interacting with a central middleware for aggregating the sensed data streams and sending data into the cloud for processing.

Our approach is based on a middleware that aggregates data from mobile devices and provides it to multiple applications. Data aggregation techniques have been extensively studied in the context of wireless sensor networks for both energy efficiency and congestion reduction. Examples include cluster-based heuristics [8], data compression [9,10], hierarchical aggregation [34,10], entropy-analysis [35], cluster formation [36], and tree selection methods [37]. We propose a revised version of the frequent pattern growth algorithm which, to the best of our knowledge, has not been used for aggregation in mobile sensing. The frequent pattern mining approach is considered as it provides an unsupervised learning technique for all kinds of data and large-datasets. The compact tree structure also reduces the complexity of searching through all possible combinations to find frequently occurring subsets. Similar algorithms have previously been applied for extraction of interesting patterns from databases, such as association rules, clusters, classifications, correlations and sequences [38,39].

Since our collaborative sensing middleware uses data from devices that are carried by humans it needs to predict the future location of a device with reasonable accuracy in order to make the decision whether or not to activate it. Random mobility models [40,41] with no restrictions in terms of memory or movement include the Random Waypoint model and its variations such as Random Walk and Random Direction. These models are unable to capture the movement of humans reliably. Although probabilistic models for user locations have been adopted by researchers [23], human mobility is most widely studied as a Levy walk model generalized from studies regarding animals like the albatross or monkeys, whose mobility pattern follows a power-law distribution. Gonzalez et al. [12] have shown that human mobility can be compared to a Levy walk with heavy-tail flight distribution. A Levy walk intuitively reflects a human mobility pattern as it consists of more short displacements than long displacements—a long tailed distribution captures this aspect, along with highlighting the probability of humans to frequent previous locations. The extension to include visiting times and frequency by Song et al. [42] provided a 93% predictability factor regarding human mobility with data history. Rhee et al. [13] have presented a truncated Levy walk model which provides a simple and realistic model for human mobility—it is this model that we have incorporated into our algorithm design.

3. Problem formulation

We seek to select the minimal number of mobile devices that can be used to gather sensor data for the benefit of a collection of