METIS: Exploring mobile phone sensing offloading for efficiently supporting social sensing applications

Kiran K. Rachuri[†], Christos Efstratiou[†], Ilias Leontiadis[†], Cecilia Mascolo[†], Peter J. Rentfrow[‡]
[†]*Computer Laboratory,* [‡]*Department of Psychology, University of Cambridge, UK*

***Abstract*—Mobile phones play a pivotal role in supporting ubiquitous and unobtrusive sensing of human activities. However, maintaining a highly accurate record of a user's behavior throughout the day imposes significant energy demands on the phone's battery. In this paper, we present the design, implementation, and evaluation of *METIS*: an adaptive mobile sensing platform that efficiently supports social sensing applications. The platform implements a novel sensor task distribution scheme that dynamically decides whether to perform sensing on the phone or in the infrastructure, considering the energy consumption, accuracy, and mobility patterns of the user. By comparing the sensing distribution scheme with sensing performed solely on the phone or exclusively on the fixed remote sensors, we show, through benchmarks using real traces, that the opportunistic sensing distribution achieves over 60% and 40% energy savings, respectively. This is confirmed through a real world deployment in an office environment for over a month: we developed a social application over our frameworks, that is able to infer the collaborations and meetings of the users. In this setting the system preserves over 35% more battery life over pure phone sensing.**

I. Introduction

The proliferation of smartphones with sensing capabilities have created an opportunity to design systems that capture vast amounts of information about people's social behavior. By leveraging the device's sensing and communication capabilities, applications can capture an accurate depiction of the user's social context, which enables the design of novel applications that can enhance the user experience [2], improve productivity [10], or facilitates new business opportunities such as targeted advertisements [8]. It is expected that the next generation of mobile applications will use continuous sensing of social context at an extremely fine granularity.

A major challenge in achieving accurate social sensing is the significant impact that continuous sensing has on the phone's battery life. The detection of social context through mobile phone sensing, typically leverages a wide range of sensing modalities. Systems such as CenceMe [2], or Emotionsense [12], utilize a combination of accelerometer, Bluetooth, location, and microphone, in order to characterize a particular social situation. The significant energy cost of collecting such information, reduces the phone's battery life, and hinders the wider adoption of such applications. Existing efforts to minimize the energy impact rely primarily on adaptive sensing techniques, trading off energy for accuracy with the aim of reducing unnecessary sensor sampling on the device [9], [12]. Although such techniques have shown improvements in terms of energy consumption, there is still a need for more efficient solutions before continuous sensing applications can be widely accepted by everyday users.

In this work, we introduce a novel approach that can offer significantly bigger reductions in energy cost without compromising on the accuracy, by *opportunistically offloading sensing to fixed sensors embedded in the environment*. Most modern buildings are instrumented with a variety of sensors, such as RFID access control systems, Passive Infrared sensors, light sensors, etc. Intuitively, if a mobile application can take advantage of such sensing infrastructures, it could at times suspend local sensing, by leveraging remote sensors. For example, relying on a building's access control system, a mobile application can decide to suspend any localization mechanisms on the phone while the user remains in the same building or even room. The feasibility of this approach and the massive energy gains that can be achieved have significant implications both for the design of future mobile applications, and the deployment of sensing infrastructures within smart-buildings. Indeed, such approach imposes a strong argument for the need of sensing infrastructures that support open access by third party systems, enabling a wide range of mobile and pervasive applications.

In this work, we explore the feasibility of efficient sensing offloading by considering the requirements of social sensing applications operating within a smart-building environment. Through an experimental study within our research institution we demonstrate that the benefits of sensing offloading can only be achieved by considering the potential gain of offloading a particular sensing task versus the potential energy cost incurred primarily due to network communication. Estimating both metrics can depend on a number of parameters and most significantly the behavior of the user and their peers. In this paper, we present the design, implementation, and evaluation of METIS, a sensing platform that implements a novel adaptive sensing distribution scheme that automatically distributes the sensing tasks between the local phone and sensing infrastructure sensors in order to support accurate continuous sensing of social activities. The proposed scheme considers various parameters such as the mobility pattern of the user, duty cycling interval, cost of sensing to determine whether sensing offload can result in energy gain at any given situation. We show through benchmarks using real traces that the sensing distribution scheme achieves over 60% and 40% energy savings compared to static scenarios where only phone-based sensing is used, and only remote sensors are used, respectively.

Finally, we evaluate the system through a real deployment with 11 users for a month in a working environment using *WorkSense*, a social application that utilizes METIS and aims to raise awareness and improve visibility of social

interactions at the workplace, by tracking formal and informal meetings during daily routines, and inferring how social interactions may impact the user's performance. The deployment shows that the application is able to infer the effect of various interaction and social patterns on the work of the users. Furthermore, we show that the system extends battery life by more than 35% compared to when no sensing offloading to the infrastructure is used.

II. Saving Energy through Sensing Offloading

The idea of sensing offloading is built on the vision of mobile phone users living in an environment instrumented with a range of sensors that can be accessed over the internet. Within such environment certain pieces of information can potentially be sensed through either the user's mobile device or a sensor that is embedded in the environment. For example, detecting if a conversation is taking place in a room, can happen either through the mobile phone's microphone or a microphone in the room, both augmented with the necessary conversation detection software. Within this setting we attempt to explore if sensing offloading can be used to reduce the energy consumption of continuous sensing on mobile phones.

In order to understand the requirements of sensing offloading, and to help us frame our hypothesis, we conducted an *exploratory deployment* within our research institution. The primary objective of the deployment was to help us identify the conditions under which offloading could reduce the energy cost on the mobile device, and those where offloading would lead to higher energy consumption than local sensing. Furthermore, we tried to explore the feasibility of designing a system that can adapt the sensing behavior on the mobile device, selecting the most appropriate action at any given situation.

The focus of the study was to explore the support for social sensing applications. To that end we identified two key sensing modalities: location / co-location sensing, and conversation detection. The experiment was conducted in a research institution involving 10 participants, and 10 offices instrumented with sensors, and lasted one week. During the deployment we collected traces about sensor data both from the mobile devices carried by the participants, and the sensors deployed in the environment. No offloading was used in this study. Using these traces we were able to re-construct the behavior of the system when sensing is offloaded to remote sensors and compare the results with sensing taking place solely on the phone.

A. Experimental Deployment

We deployed a logging system that was able to collect sensing traces from mobile phones carried by the participants, and sensors deployed in the office environment of our research institution. Both systems were able to detect co-location, location, and conversations of the users.

B. Mobile Phone Sensing

We designed an Android application that was able to perform indoor localization, co-location detection (both using Bluetooth-based localization, utilizing Bluetooth anchors in the environment), and conversation detection (using the microphone). The application gathers data from accelerometer, Bluetooth, and microphone sensors. The conversation recognition module is based on that used in the EmotionSense system [12].

C. Sensing Infrastructure

In our deployment we aimed at exploiting the existing sensing infrastructure in the environment. As part of previous sensing experiments, within our research institution [4], two sensing infrastructures were already available: indoor localization, and room occupancy. In particular Nokia 6210 Navigator phones were tasked to act as Bluetooth anchors and assist in the localization of mobile phones. The infrastructure includes 12 such Bluetooth anchor points covering a space of 10 offices. Furthermore, each Nokia node periodically scans for Bluetooth devices in proximity using the *lightblue* module for *Python for S60* (PyS60). In order to offer additional support for conversation detection we extended the sensing module on the Nokia devices with conversation sensing functionality (using the same scheme [12]). To capture accurate room occupancy data, a network of imote2 sensors had already been deployed around the office spaces (13 nodes covering 10 offices). The sensors are attached to desks and are able to detect when a particular desk is occupied. The desk occupancy status is inferred by detecting vibration patterns using the 3-axis accelerometer sensors embedded in the node. Each of the nodes periodically sends the current state (i.e., whether the desk is occupied or not) to the root node that is connected to a server. The sensing infrastructure is interfaced with the Internet through an sMap [3] type backend.

D. Benchmarks

In order to assess the effectiveness of sensing offloading we used the collected traces to investigate the performance using the following schemes.

- **Never Offload.** This scheme resembles the typical behavior of a mobile application where sensing relies solely on the local phone sensors.
- **Always Offload.** In this scheme the mobile phone offloads sensing every time it is in an environment where appropriate sensors are available.

The benchmarks consider two sensing modalities: indoor localization using Bluetooth scanning, and conversation detection using microphones. For the always offload scheme, desk occupancy sensors are used to suspend the Bluetooth scanning on the phone when the user is at his desk, and microphone sensors in the environment are used to offload the conversation detection.

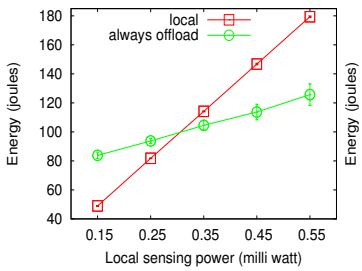

Figure 1. Energy per hour for varying sensing cost.

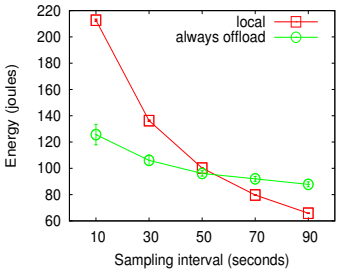

Figure 2. Energy per hour for location detection vs. sampling interval.

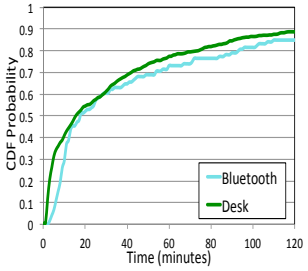

Figure 3. CDF of time spent for each location visit.

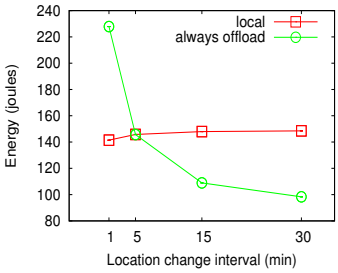

Figure 4. Energy per hour for location detection with varying mobility patterns.

In our benchmarks we estimate the energy cost of performing a sensing task locally on the phone, and the cost of communication between the phone and the infrastructure when sensing is performed on the infrastructure. The estimation of local sensing is based on the power consumption values reported by [5], [15]. The estimation of network traffic includes the baseline cost of keeping the Wi-Fi interface on and the average cost of communication per byte as it is estimated by [13]. We note that in a realistic setting, users typically enable Wi-Fi for their own purposes, which is further explored in Section IV. The communication traffic between the infrastructure and the mobile device depends on the number of events that are detected by the deployed sensors, and therefore depends on the actual behavior of the participants in the instrumented spaces. Furthermore, in order to understand how these values affect the performance of offloading, we used the same traces to simulate scenarios by modifying certain parameters such as sensing cost or sampling rate. The measurements report the average energy consumption across all the participants for the entire duration of the experiment.

E. Lessons Learnt

Which Sensing Tasks to Offload. We identified two key parameters in our scenarios that affect the energy cost of sensing: the sampling rate, and the energy cost of sensing and processing data from a particular sensor. In our deployment we had two modalities that can be considered as mid-cost and high-cost respectively in terms of energy. Based on works such as [5], [15], a low-cost sensor such as the accelerometer, for example, consumes about 30mW, the Bluetooth scanning consumes about 160mW, and an expensive sensor such as GPS consumes around 430mW. In order to generalize from our traces, we simulated the performance of the two schemes considering a range of sensor sampling costs. Figure 1, as expected, shows that for low cost sensing, offloading does not lead to energy gain, as the cost of network communication outweighs the small benefit of suspending local sensing.

The varying sampling rate has a similar effect. Higher sampling rate incurs more energy cost on the device, however, it may increase the number of events reported by the infrastructure (fewer missed events). We analyzed the impact of sampling rate by sub-sampling our original dataset. As illustrated in Figure 2 the increase in sampling interval (decrease in sampling rate) reduces the sampling cost faster than the network cost for offloading. The figure demonstrates the overall cost for Bluetooth scanning with increasing sampling interval, as shown, at low sampling rates, local sensing can in fact perform better than offloading.

The Impact of Mobility. Apart from the parameters affecting the cost of local sensing, efficient offloading depends on the varying cost of network communication that includes the cost of "hand-over" (subscribe / unsubscribe to a sensor), and the cost of receiving events that are detected by the infrastructure. Both these costs depend on the behavior of the users in the instrumented spaces. We explored the impact of the participants' mobility patterns on the cost of sensing offloading. In our traces, users spent 64% of their time inside their own office and 21% of their time in other offices with sensing capabilities. For those times, where offloading opportunities were available, we analyzed the average time that people spent when they visited a particular location. Figure 3 shows a CDF of the time each user spent in each visit. We observe a high number of short visits (50% of them last less than 20 minutes), which can be the cause of increased network traffic due to hand-overs, moreover, there is a possibility of sensing events to be either missed or wrong (reported by the wrong environment).

We investigated the impact of user mobility looking at scenarios with different average times that people spent in a given location. We used the original traces collected and modified the amount of time that each user spent in each room. As illustrated in Figure 4, the results demonstrate how high mobility can significantly increase the cost of sensing offloading. Essentially, for a given visit to a location, offloading can only deliver positive results if an offloading scheme can predict the possible time that a user will spend in that location, and ensure that the duration is enough to offer an energy gain.

The results clearly suggest that neither pure phone sensing nor exclusive remote sensing is an optimal solution for all the scenarios, and a dynamic sensing offloading scheme that considers the sensing parameters and the user's mobility is required for achieving an optimal performance.

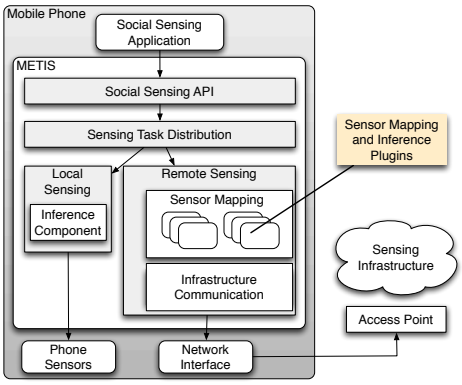

Figure 5. Architecture of the METIS system.

```
<?xml version="1.0" encoding="UTF-8"?>
<mt:METIS xmlns:mt="http://our.domain.org/metisML">
 <!-- ............ -->
 <mt:SensorNetwork>
  <mt:sMAPURI>http://our.domain.org/</mt:sMAPURI>
  <mt:SensorList>
   <mt:SensorNode>
    <mt:SensorId>BT_D11_R01</mt:SensorId>
    <mt:SensorType>BluetoothScanner</mt:SensorType>
    <mt:Location>
     <mt:Label type="office">R01</mt:Label>
    </mt:Location>
   </mt:SensorNode>
   <mt:SensorNode>
    <mt:SensorId>ACC_D1_R01</mt:SensorId>
    <mt:SensorType>DeskUseDetector</mt:SensorType>
    <!--..........-->
  </mt:SensorList>
 </mt:SensorNetwork>
</mt:METIS>
```

Figure 6. Sample sensing infrastructure manifest obtained by the METIS system from a service provider.

III. The METIS System

Motivated by the results of the exploratory deployment, we designed the METIS system. METIS is a mobile phone service that offers efficient continuous sensing for mobile social applications by leveraging both phone and infrastructure sensing. It provides an abstraction over the sensing modalities required by typical social sensing applications, such as location, co-location, and conversation detection and offers a framework for the incorporation of additional sensing modalities when needed (Figure 5). The operation of METIS includes the discovery of sensing devices that are available in the immediate environment of the mobile phone user, the identification of devices that could be used for offloading, and the decision to perform such offloading in order to maintain overall energy efficiency. If such offloading is not considered beneficial METIS falls back to local sensing utilizing the resources of the mobile device. The primary challenges and the main focus of this work are related to the utilization of remote sensors in order to improve sensing performance, which we address in the following subsections.

A. Interaction with Sensing Infrastructure

One of the key motivations for the design of METIS is the emergence of a range of Web-based architectures that allow interaction with the sensing infrastructure over IP networks. We decided to assume the presence of sensing infrastructures that follow the same architectural principles that are adopted by such architectures. Systems such as SenseWeb [7] identify two key elements in their architecture: the presence of a rendezvous point, which allows a client to query the infrastructure about available sensing resources and their capabilities, and the support for a resource communication protocol that enables clients to interact with specific sensing resources. The operation of METIS imposes the following two requirements on the sensing infrastructure: (i) the specification of the physical location of a sensing resource as reported by the rendezvous point, and (ii) the support for an asynchronous publish-subscribe interface for communication with a sensing resource. Both of these requirements are supported by most common Web-based sensing architectures. In the design of METIS we adopt the sMAP [3] communication protocol for interaction with specific sensing resources.

The information that can be retrieved through the rendezvous point plays a key role in the operation of METIS. The expectation is that METIS can identify the physical location of sensor points. In the design of METIS we target indoor environments with the aim of taking advantage of common sensing technologies that can be found in smart homes or office buildings. To that end, we define a minimal XML schema of the sensing infrastructure that incorporates information about the physical location of sensor points within a building (Figure 6). Although the format of the schema is designed to meet our needs, the same information can be easily extracted by standard-based schemata such as *SensorML*. METIS can be trivially extended with additional schema parsers to support multiple infrastructure interfaces.

Sensor Mapping. One of the primary functions of METIS is the association of social sensing modalities with possible remote sensors that can be used for offloading. METIS enables this association with the incorporation of sensor-specific drivers in the form of *plug-ins*. The *Sensor Mapping* component acts as a repository of sensing plug-ins, triggering them on demand when a particular sensor device is within the proximity of the user. Each plug-in maps a specific high-level social sensing task to a combination of subscriptions to certain sensor nodes in the environment. Plug-ins that have been implemented for the METIS system include: 1) If real-time room occupancy information is available, subscribe to receive events about the current room, and switch off the location scanning when the number of people in that room has not changed. 2) If desk occupancy information is available, subscribe to receive notifications about the current desk, and report current activity as "sitting" without using the accelerometer. 3) If noise level detection is available in the room, subscribe to receive notifications about changing noise levels, and adjust conversation detection on the phone when not needed.

B. Sensing Offloading

By analyzing the results of the exploratory study we were in a position to identify the parameters that can affect the energy trade-offs when deciding to perform sensing offloading to remote sensors. Specifically, the decision is based on the estimation of the energy cost when sensing is performed on the phone, and the prediction of the network energy cost when sensing is performed remotely. In estimating the latter, the user's mobility pattern, the time spent in a particular location, and the rate at which events are reported by the remote sensor are factors that affect the energy cost. In this subsection we describe an offloading scheme that achieves energy efficiency by considering all these parameters.

Gain Threshold based Offloading. The *Gain Threshold* scheme operates by calculating the probability that offloading a particular sensing task would result in energy gain when compared to the corresponding local sensing task. If the probability of gain is greater than 0.5, i.e., if offloading has more than 50% chance of resulting in gain then the offloading is performed. The probability of gain is calculated by estimating the possible communication costs that the phone may incur if offloading is performed. When a sensing task is offloaded, there are two types of energy costs involved, *fixed* and *variable* costs. The fixed costs are the costs involved in maintaining a network connection between the phone and the infrastructure system, subscribing to a remote service for offloading, and canceling the subscription at the end. The variable costs are the communication costs incurred by updates received as part of the infrastructure sensor events that depend on the user behavior.

A decision to offload a sensing task results in energy gain if the user stays in the location for long enough so that the sum of the fixed and variable costs is less than the cost of local phone sensing for that period. We refer to this minimum time period as *gainTimeThreshold*. We note that fixed costs can be calculated based on offline estimated values for data transfer over the network (for example, this could be measured on the Android phones using a power meter [15] and on the Nokia phones using the Nokia Energy Profiler), and variable costs (or event rate or sensor state change rate) can be calculated based on the past history of event traces as recorded by the sensing platform. The *gainTimeThreshold* varies for each of the sensors as the cost and event rate for these change.

1) Gain Time Threshold: In this subsection, we present the calculation of the *gainTimeThreshold*. Notation used:

G_s : gain threshold time of a sensor s
C_{sl} : cost per sample of local sensing
S_{sl} : sampling rate of the sensor s
C_{so} : fixed cost of offloading the sensing task
C_{sr} : cost per update of remote sensing
C_{nr} : baseline cost for maintaining network connection
U_{sr} : update rate of remote sensing
N : number of sensing tasks that can be offloaded

According to the definition of *gainTimeThreshold*, for a sensor s, the energy consumption of local phone sensing is equal to the sum of fixed and variable costs of remote sensing for *gainTimeThreshold* amount of time. In other words, if remote sensing is used for more than *gainTimeThreshold* amount of time, then it results in a positive energy gain, and if remote sensing is used for less than *gainTimeThreshold* amount of time, then it results in a negative energy gain, i.e., offloading is not beneficial in this case.

The local sensing cost for G_s amount of time for a sensor s (L_s) = The cost of local sensing per sample (C_{sl}) $\times$ Local sampling rate (S_{sl}) $\times$ G_s.

The remote sensing cost for G_s amount of time (R_s) = Fixed control traffic cost (C_{so}) + (Event update rate (U_{sr}) $\times$ Cost of network transfer per update (C_{sr}) $\times$ G_s) + Baseline network connection cost per sensor for G_s amount of time.

Per the definition, G_s is the amount of time for which, local cost = remote cost, i.e., $L_s = R_s$.

$$\implies C_{sl} \times S_{sl} \times G_s = C_{so} + C_{sr} \times U_{sr} \times G_s + \frac{C_{nr}}{N} \times G_s \quad (1)$$

$$\implies G_s = \frac{C_{so}}{C_{sl} \times S_{sl} - C_{sr} \times U_{sr} - \frac{C_{nr}}{N}} \quad (2)$$

G_s for a sensor s quantifies the minimum amount of time the sensing task should suspend local sensing and use remote sensing to achieve energy cost benefit i.e., it is the minimum amount of time the user should stay in the current location after offloading the task to achieve energy gain.

2) Probability of Gain Estimation: In this section we present the estimation of the probability that offloading a sensing task (s) will result in gain. This value is used to make a decision on whether this offloading is beneficial. Let $\{v_{j_1}, v_{j_2}, \ldots v_{j_k}\}$ be the total visits of the user to a location j, let $\{tv_{j_1}, tv_{j_2} \ldots tv_{j_k}\}$ be the total duration of each of the k visits, respectively. First, for each sensor s, we divide the total duration of l^{th} visit to a room j into two parts: *favorable time* (ft_{sj_l}) and *unfavorable time* (ut_{sj_l}). Favorable time for a sensor is the time during which the offloading of the sensing task results in a positive gain, and unfavorable time for a sensor is the time during which the offloading results in a negative gain. Therefore, for a visit to a room, the favorable time is the total visit time subtracted by the G_s value (as we need at least G_s amount of time to achieve positive gain), and unfavorable time is G_s, $i.e.$, for l^{th} visit to a room j by the user, favorable and unfavorable times for a sensor s are calculated as:

$$ft_{sj_l} = tv_{j_l} - G_s \quad (3)$$

$$ut_{sj_l} = G_s \quad (4)$$

Then, the probability (pt_{sj}) that a user stays at a location j for more than the *gainTimeThreshold* (G_s) value for a sensor s is calculated as: the total favorable time at location j for all visits divided by the total time at the location j.

$$\implies pt_{sj} = \frac{\sum_{l=1}^{k} ft_{sjl}}{\sum_{i=1}^{k}(ft_{sj_i} + ut_{sj_i})} \quad (5)$$

pt_{sj} is the probability of gain that is used to make an offloading decision for a sensing task s when the user is in room j. Finally, if this probability value is greater than 0.5, it indicates a more than 50% chance of resulting in positive gain, and in this case the sensing task is offloaded. A task that is offloaded to a remote sensor is cancelled (unsubscribed from remote sensing updates) when the user moves away from the current location, as the remote sensor may not capture the user's activities accurately, as it is not in proximity to the user. We present a detailed evaluation of this scheme through micro-benchmarks in the next section.

IV. Benchmarks

In this section we present the evaluation of the gain threshold offloading scheme through micro-benchmark tests. We compared the energy performance of the scheme with the two schemes described in Section II. The dataset collected from the initial study was used for these tests. In this evaluation, the gain threshold technique works in an online fashion, i.e., as the traces are replayed the gain threshold (Eq. 2) and the probability values (Eq. 5) are continually learned and the decision on offloading is taken accordingly.

Sensing Cost, Sampling Interval, and Mobility. We showed in the results of the initial study that the two simple sensing schemes may not be suitable for all conditions. We now compare the performance of the proposed scheme considering varying cost of sensing, sampling interval, and mobility patterns. Figures 7, 8, and 9 show the total energy consumption of the schemes with respect to these variables. We can observe that the threshold scheme tends to match the best performing scheme in all the cases as it considers these parameters in its decision to offload.

Excluding Wi-Fi Baseline Cost. In the results presented so far, we include in the network cost function the Wi-Fi baseline cost (cost of keeping Wi-Fi on). However, it is typical for many users to keep Wi-Fi on for other unrelated purposes. In order to analyze this situation, we benchmarked the behavior of the scheme excluding the Wi-Fi baseline cost from the calculation considering only the cost of data exchange. Figure 10 shows the local sensing, network exchange, and total energy consumption for location and conversation detection. We now observe that the offloading schemes result in considerable energy savings due to the fact that the Wi-Fi was already enabled. The threshold scheme resembles the remote sensing scheme in this case, achieving high efficiency, saving around 60% of energy when compared to pure phone sensing scheme.

As shown in the results so far, the threshold scheme tends to follow the most optimum scheme in different circumstances. However, the optimal strategy may not always include one of the extreme offloading schemes. To demonstrate this, we evaluate the schemes with respect to the following two scenarios.

Adaptive Sampling. The use of adaptive sampling is a typical approach in mobile applications that use continuos sensing [12]. In these cases the local phone sensor sampling rate changes in the face of changing behavior of the user. In this test, we used two adaptive sampling schemes: In *adaptive scheme 1*, we used an exponential back-off and a linear advance of sampling interval, and in *adaptive scheme 2*, we used a linear back-off and an exponential advance of sampling interval. The sampling interval is increased using the back-off function when there is no change to the user's context and advance function is used when there is a change to the user's context. The context of the user for this evaluation is defined as the co-located Bluetooth devices. Figure 11 shows this result. We observe that, in our deployment, for the purely local sensing the exponential back-off (scheme 1) saves more energy than the linear back-off (scheme 2). Also, the choice of adaptive scheme does not have a significant impact on *always offload* scheme as it uses remote sensing most of the time (except in uninstrumented areas). However, in both cases neither the *local* or *always* schemes is optimal. This is because when using adaptive sampling, each of them will be optimal for a subset of the possible situations. The threshold scheme appears to outperform the two others, by selecting the most optimal approach in every case.

Multiple Sensors. In this case, we evaluate the schemes using an inexpensive (30mW) and an expensive (1.2W) sensors. We assume that the inexpensive sensor generates large amount of data per sensor sampling (100KB per sample, similar to the audio recording for 5 seconds using the PCM format in the conversation detection module). Figure 12 shows the result of this evaluation, where we can observe that the energy consumption of the threshold scheme is much lower than the other schemes, in particular, when the Wi-Fi baseline cost is considered, the threshold schemes consumes 46%, 31% less energy than the always offload and always local schemes, respectively. When the Wi-Fi baseline cost is not considered, the threshold scheme consumes 57%, 63% less energy than the always offload and always local schemes, respectively. The always local scheme consumes high energy because of local sensing of the expensive sensor and the always offload scheme consumes high energy because of the communication cost of the inexpensive sensor. However, the energy consumption of the gain threshold based offloading scheme is much lower than the other schemes, as it selects the most optimal configuration for each sensing modality.

The results presented clearly show that the proposed scheme dynamically offloads the sensing tasks by adapting to the sensing parameters and the user's mobility patterns.

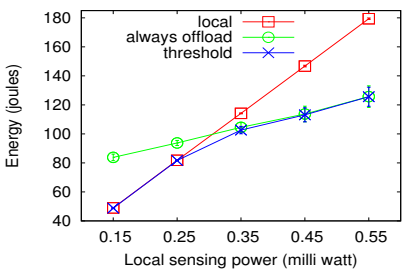

Figure 7. Total energy consumption per hour for varying sensing cost.

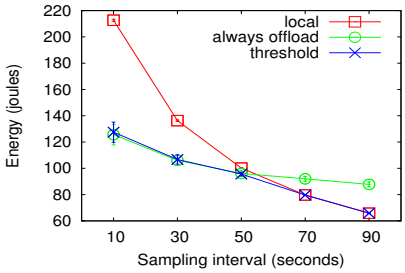

Figure 8. Total energy consumption per hour for location detection with varying sampling interval.

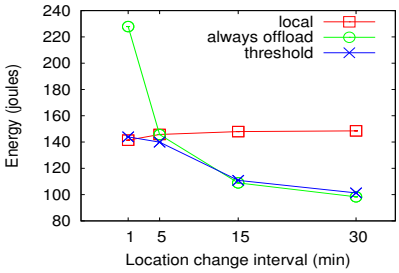

Figure 9. Total energy consumption per hour for location detection for varying mobility patterns.

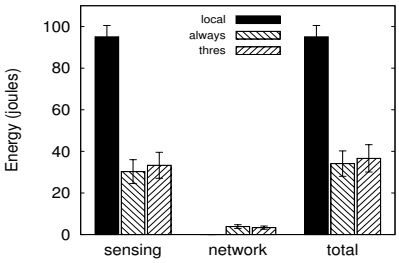

Figure 10. Energy per hour for location, conversation detection without considering Wi-Fi baseline cost.

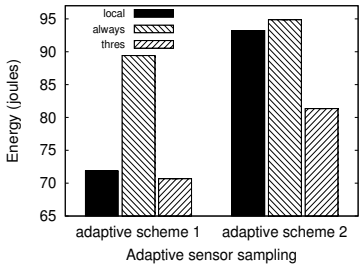

Figure 11. Energy consumption per hour for two adaptive sampling schemes.

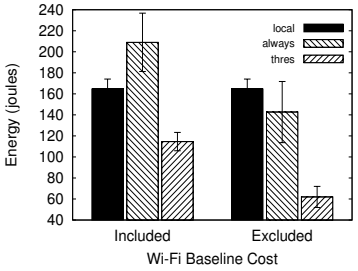

Figure 12. Energy consumption per hour for an inexpensive and an expensive sensors.

V. Case Study

The ultimate evaluation of METIS was performed through a real deployment. The main goals of the deployment were to evaluate the energy efficiency of METIS, and to show that the applications can capture the behavioral patterns of the users utilizing the services of the METIS system.

A. The WorkSense Application

The design of the WorkSense application was motivated by studies such as [10], in which, the authors showed how we can significantly increase work performance by understanding the face-to-face interactions and the formation of various social groups within the organization. We exploited the METIS framework and a number of sensing modalities that were available in our office environment to design *WorkSense*: an application that aims to infer the collaborations and meetings of users, and offers awareness on the impact social interactions may have on their work. More specifically the application includes the following functional components:

Mobile Phone Application. The application was implemented on the Android platform and captures the mobility patterns of the users during working hours (visited offices, meeting rooms etc.) and interaction patterns. This information is used to infer social context, mining working patterns, and to offer awareness to the user. The WorkSense mobile application shows to users details about their collaborations, meetings, and the effect of meetings on their work activities.

Meeting Detection. A meeting is considered a case where two or more users are co-located for more than a pre-set time threshold (currently 15 minutes) and they are having a conversation. In the definition, we did not specify a prerequisite that the location should be a meeting room, as we wanted to capture informal meetings that some times can take place in a common room or in the corridor.

Collaboration Detection. Although people in our research institution are typically formed into groups that may collaborate within the context of a particular research project, the WorkSense application attempts to capture a more objective picture of how collaborations are taking place within the research lab. Sometimes a user who is not officially assigned to a project can play an important role. Collaboration detection is achieved by applying a community detection algorithm over the co-location data. However, as there are cases where two or more colleagues may share the same office, a straight forward application of the community detection over the co-location data would result in communities that are heavily biased towards people sharing offices, even if they do not typically collaborate. In order to overcome this issue, we defined an *intended visit* as the case where a person visits another person with the intention of having a conversation ($deskEmpty(u_1) \wedge colocated(u_1, u_2) \wedge conversation(u_1)$). If sensing infrastructure is unavailable, then we do not consider the desk sensor information. The result of an intended visit is a directional edge that links the two users where u_1 is the initiator and u_2 is the target. A community detection algorithm is then applied over the directed graph to detect collaborations, which we describe in the next subsection.

From a design perspective, the operation of WorkSense *does not require* the presence of sensors embedded in the environment; however, when such infrastructure is available, the energy performance of the application can be dramatically improved and the accuracy of the gathered information can be increased as well.

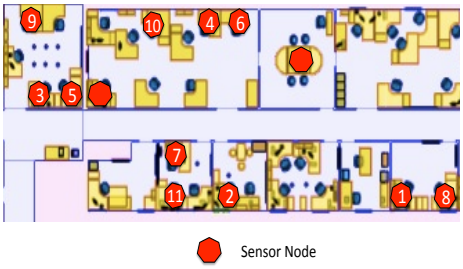

Figure 13. Sensor nodes were installed on 13 desks including two relay nodes.

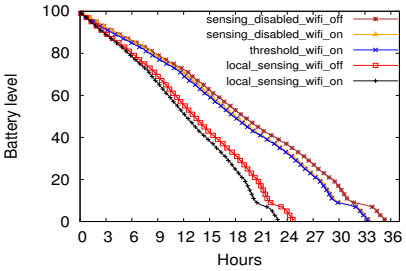

Figure 14. Battery drain of Samsung Galaxy S phone for various scenarios.

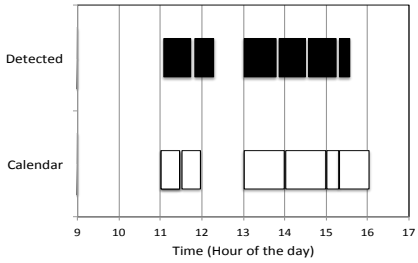

Figure 15. Timeline of detected vs. calendar meetings.

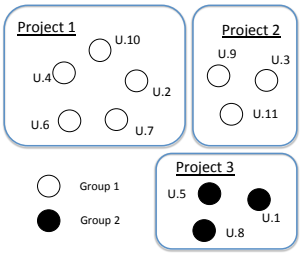

Figure 16. Communities as reported by the participants.

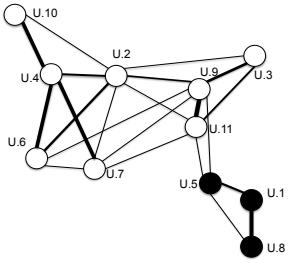

Figure 17. Automatic detection of groups (Level 2 communities).

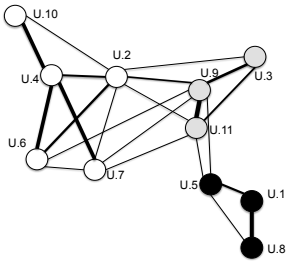

Figure 18. Automatic detection of projects (Level 1 communities).

B. Deployment

We deployed the METIS system and the WorkSense application in a working environment involving 11 users for about a month. During this period each user carried an Android phone (Samsung Galaxy S or HTC Desire). The mobile application was able to utilize existing sensing infrastructure that was deployed in our institution (Figure 13). This included desk occupancy sensors, Bluetooth sensors, and conversation detection infrastructure (see Section II).

Energy Impact Analysis. We evaluated the energy performance of METIS during the deployment. To compare the difference in battery life *with* and *without* using the sensing offloading, we used pairs of Samsung Galaxy S phones, carried by the same user. One of the phones was set to never offload data (pure phone sensing with Wi-Fi switched off) while the other followed the threshold based offloading scheme. This functionality was swapped between the two phones over consecutive measurements in order to reduce the impact that minor differences of the phone batteries may have on the results. The phones were allowed to discharge only during office hours, by switching them off overnight. During the experiment the phones were not used for any other activities. We then also measured the battery life of the phone when using: local phone sensing with Wi-Fi switched on, sensing disabled with Wi-Fi switched on, and sensing disabled with Wi-Fi switched off. In all these cases, a battery monitor was always running on the phone to measure the battery discharge. The average consumption over multiple runs is shown in Figure 14. The results show that, when using the gain threshold scheme and exploiting the sensing infrastructure, the battery lasts up to 45% longer compared to the case of local phone sensing with Wi-Fi on, and 35% longer compared to the case with the Wi-Fi off. Furthermore, the energy cost of using METIS is very close to a mobile phone with the Wi-Fi on and no sensing, and only 6% less compared to a phone that uses no sensing and no Wi-Fi. This is an indication that opportunistic sensing offloading can improve the support for the long-term deployment of sensing applications, by significantly minimizing the impact on the phone's battery life.

Meeting Detection Analysis. One of the primary motivations behind the design of WorkSense was the fact that social sensing has the potential to capture views over the interactions and activities at the workplace. The most expected result of WorkSense was the detection of formal and informal meetings. The *Meeting Detection* service was able to detect the exact times that meetings are taking place in the laboratory. Furthermore, by utilizing the calendar information such meetings were populated with appropriate meta-data. Figure 15 shows a snapshot of a timeline where the calendar schedule is contrasted with the actual meetings as they were detected by WorkSense. In this particular snapshot we see how a person that has consecutive meetings can slightly adapt the schedule based on the duration of previous meetings. In this case many meetings were moved earlier: this was not something updated in the calendar.

Community Detection Analysis. Figure 16 shows the different groups and project teams as they were reported by the users. We used WorkSense to automatically detect such communities. To do this, we first create a weighted graph of detected collaborations where a link between two users represent the amount of time they spent in meetings and discussions. Afterwards, we apply the *Louvain's* community detection algorithm [1] to detect groups and projects that users belonged to. The Louvain method is a hierarchical greedy algorithm and operates in two phases, repetitively:

first, it searches for small communities, and then it combines nodes of the same community and forms a new network. The method uncovers hierarchies of communities and allows to further identify sub-communities, sub-sub-communities etc. This algorithm generates team formations on varying levels of granularity: at the lowest level (Level 1) it identifies smaller communities and at higher levels (e.g., Level 2) it merges smaller communities to form larger clusters. The thickness of edges in the graph represent the weight/strength of the link, and the colour of a node represents its community. Figure 17 shows the Level 2 communities where WorkSense was able to identify the separation between people belonging to different departmental groups. The spatial placement of nodes is generated based on a spring model: the larger the spatial difference between two nodes the lesser is the collaboration/communication between them. More interestingly though, the output of the community detection when operating at Level 1 (Figure 18), was a more fine grained break down of groups that were not related to the information that the users offered but to the projects that users have been working on, even within the same group. Furthermore, the WorkSense application was able to identify cross team collaborations and pin point people acting as *bridges* extending collaboration links across teams.

VI. Related Work

Although, energy efficiency has been one of the key design considerations in a number of mobile sensing systems [2], [9], [12], most of them consider only the case of pure phone sensing. A number of systems attempted to address energy consumption issues by using adaptive sensor sampling techniques [9] allowing them to improve energy efficiency by trading-off accuracy. As shown in Section IV, METIS can complement adaptive sensing systems offering further improvements in energy utilization.

With respect to using sensors in the infrastructure: FollowMe [6] lets mobile applications to exploit the sensors like cameras, microphones that are available in the environment, for richer context detection. ErdOS [14] is a mobile operating system that extends the battery life of mobile handsets by managing resources proactively and by exploiting opportunistic access to resources in nearby devices using social connections among users. Comparing with FollowMe and ErdOS, our system opportunistically offloads the sensing tasks to the infrastructure to maximize the energy gain while considering the sensing parameters and the user's mobility. None of these works that exploit the sensing infrastructure provided a solution to the problem of *when to offload phone sensing to infrastructure*. In [11], the authors propose that mobile phones can serve as data mules for sensor networks due to their ubiquity and show that opportunistic mulling is suitable for office-based deployments. Even though this work involves interaction of mobile phones with sensor networks, it addresses a very different problem than METIS.

VII. Conclusions

In this paper we presented METIS, a mobile sensing platform that supports long-term deployment of social sensing applications by leveraging both phone sensors and sensors in the environment. The system implements a novel sensing distribution scheme that is able to switch between phone and remote sensors considering the various sensing parameters, and mobility patterns of the users. We showed through several benchmark experiments on real traces that the system is able to achieve significant energy savings compared to pure phone sensing and remote sensing schemes. We also reported on a real deployment of METIS through a social application. We showed that the application is able to infer the collaborations and meetings of the users while achieving a battery lifetime that is only 6% shorter in duration than a mobile phone that performs no sensing. We plan to extend the system by supporting additional sensing modalities and also enhance the application to recommend effective work behaviors to people at the workplace.

Acknowledgments

This work was supported through the Gates Cambridge Trust and EPSRC grant EP/G069557/1.

References

[1] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 2008.

[2] A. T. Campbell et. al. The rise of people-centric sensing. *IEEE Internet Computing*, 2008.

[3] S. Dawson-Haggerty, X. Jiang, G. Tolle, J. Ortiz, and D. Culler. sMAP: a simple measurement and actuation profile for physical information. In *Sensys '10*. ACM, 2010.

[4] C. Efstratiou, I. Leontiadis, M. Picone, K. K. Rachuri, C. Mascolo, and J. Crowcroft. Sense and sensibility in a pervasive world. In *Pervasive'12*, 2012.

[5] R. Friedman, A. Kogan, and Y. Krivolapov. On power and throughput tradeoffs of wifi and bluetooth in smartphones. In *INFOCOM*, 2011.

[6] B. Greenstein and B. Longstaff. Followme: Enhancing mobile applications with open infrastructure sensing. In *HotMobile'11*, 2011.

[7] A. Kansal, S. Nath, J. Liu, and F. Zhao. Senseweb: an infrastructure for shared sensing. *IEEE MultiMedia*, 2007.

[8] B. Kim, J.-Y. Ha, S. Lee, S. Kang, Y. Lee, Y. Rhee, L. Nachman, and J. Song. Adnext: a visit-pattern-aware mobile advertising system for urban commercial complexes. In *HotMobile '11*. ACM, 2011.

[9] H. Lu, J. Yang, Z. Liu, N. Lane, T. Choudhury, and A. Campbell. The Jigsaw Continuous Sensing Engine for Mobile Phone Applications. In *SenSys'10*. ACM, 2010.

[10] W. Lynn et. al. Mining Face-to-Face Interaction Networks using Sociometric Badges: Predicting Productivity in an IT Configuration Task. In *ICIS'08*, 2008.

[11] U. Park and J. Heidemann. Data Muling with Mobile Phones for Sensornets. In *Sensys'11*. ACM, 2011.

[12] K. K. Rachuri, M. Musolesi, C. Mascolo, P. J. Rentfrow, C. Longworth, and A. Aucinas. EmotionSense: A Mobile Phones based Adaptive Platform for Experimental Social Psychology Research. In *UbiComp'10*. ACM, 2010.

[13] A. Rice and S. Hay. Decomposing power measurements for mobile devices. In *Percom '10*. IEEE, 2010.

[14] N. Vallina-Rodriguez and J. Crowcroft. Erdos: achieving energy savings in mobile os. In *MobiArch'11*, 2011.

[15] L. Zhang et. al. Accurate online power estimation and automatic battery behavior based power model generation for smartphones. In *CODES/ISSS '10*. ACM, 2010.