Less is More: Energy-Efficient Mobile Sensing with SenseLess

Fehmi Ben Abdesslem
School of Computer Science
University of St Andrews
St Andrews, Fife, UK
fehmi@cs.st-andrews.ac.uk

Andrew Phillips
School of Computer Science
University of St Andrews
St Andrews, Fife, UK
ap267@st-andrews.ac.uk

Tristan Henderson
School of Computer Science
University of St Andrews
St Andrews, Fife, UK
tristan@cs.st-andrews.ac.uk

ABSTRACT

We present SenseLess, a system that leverages the different energy consumption characteristics of sensors to maximise battery life in mobile-sensing applications. We use the less expensive sensors more often, thereby enabling us to use the more expensive sensors less frequently. In the context of location-aware services, experimental results indicate that for a typical indoor and outdoor walk, compared to a simple GPS-based system, our SenseLess system can reduce energy consumption by more than 58% when determining a user's location, while maintaining the fidelity of the sensed data. This extends the battery life of a typical handheld device from 9 hours to 22 hours.

1. INTRODUCTION

Mobile phones now increasingly contain a number of sensors, which can be leveraged to create interesting new applications. For instance, accelerometers can be used to determine movement or orientation, which aids functionality of games or photography. Global Positioning Service (GPS) receivers can be used to enable location-aware services.

The use of these sensors, however, can have a severe impact on battery life. To demonstrate these effects, we conducted an experiment using a typical smartphone, the Nokia N95 8GB. We used firmware version v30.0.018, with all default settings and brand new[1] 1200 mAh batteries (Nokia BL-6F Li-Ion). Simple Python scripts were run to trigger a particular sensor continuously; for instance, the 802.11 radio was set to scan continuously, while the microphone was set to record audio continuously. Using the Nokia Energy Profiler v1.1 software, we then measured how long the battery took to deplete when the sensor was continously triggered (all other sensors and radios, including the GSM/UMTS cellular radio were switched off). Table 1 indicates that applications which use sensors have the potential to run down a mobile device's battery very quickly.

From this very simple experiment, we propose the following hypothesis: by choosing to use more energy-efficient sensors where appropriate, it is possible to decrease the energy consumption of mobile sensing applications. In other words, we can sense *more* (in that the device batteries will last longer) by doing *less* (using the less power-hungry sensors). We thus propose SenseLess, a system for handheld devices that uses sensors in this manner to save energy. Our intent is for this to be a general system, but for now we focus on one particular application: obtaining a user's location in an energy-efficient manner.

[1]A brand new battery was used for each experiment (first cycle after full charge) to ensure that battery age was not a factor.

MobiHeld'09, August 17, 2009, Barcelona, Spain.
ACM 978-1-60558-444-7/09/08.

Table 1: Energy consumption of different sensors. Each sensor was run continuously on a Nokia N95 8GB smartphone until the battery was depleted.

Sensor	*Approximate battery life (hrs)*	*Average power consumption (mW)*
Video camera	3.5	1258
IEEE 802.11	6.7	661
GPS (outdoors)	7.1	623
GPS (indoors)	11.6	383
Microphone	13.6	329
Bluetooth	21.0	211
Accelerometer	45.9	96
All sensors turned off	170.6	26

2. THE SENSELESS SYSTEM

Table 1 shows that the accelerometer is the cheapest sensor in terms of energy consumption. For this reason, SenseLess makes extensive use of the accelerometer to sense movements. We argue that sensing more movements can allow us to use high-energy sensors less frequently, thus reducing the overall energy consumption. For example, in a location-aware application, using sensors for localisation is unnecessary when the user is not moving. The accelerometer can thus be used to detect when the user is not moving and then stop sensing to save energy.

We conducted a simple experiment where 3 different users were asked to keep a Nokia N95 8GB in their pocket, and to sit down normally for 20 minutes while continuing their current task (namely watching TV, attending a talk or participating in a meeting). We also asked 3 volunteers to walk normally in the streets for at least 20 minutes. During the experiments, the data returned by the accelerometer were logged every 10 seconds into a file, along with timestamps. The data collected during our experiment show that the Euclidean distance computed by SenseLess, between the current coordinates returned by the accelerometer and the last one, is clearly greater when the subjects are walking than when they are sitting down. To differentiate the 2 groups of curves, we have experimentally chosen a threshold Euclidean distance value of 50. A user, however, might also briefly move while sitting down, and briefly stop while walking (e.g., before crossing a road). Hence, the threshold might be reached, incurring false positives when detecting movements or stops. False positives when detecting movements are rare enough to be ignored and allow the sensors to activate, even if the user is not moving. The reason behind this design choice is two-fold. First, occasional triggering of the sensors may not consume a significant amount of energy. Second, SenseLess has been designed to be responsive and avoid missing user's movements, which would affect location accuracy. On the contrary, false positives when detecting stops must be avoided, as they can potentially alter the location accuracy. Hence, SenseLess does not switch off sensors when only one value is under the threshold. Instead, switching off the sensors requires 3 consecutive values under the threshold, which corresponds to 30 seconds of perceived inactivity.

SenseLess uses both GPS and 802.11 localisation methods, switching between the two methods to use the most energy-efficient sensor. By default, the GPS is used. When no GPS signal is available, SenseLess starts scanning for 802.11 access points. This is mostly to overcome the lack of GPS signals indoors. At the same time, the GPS receiver keeps searching for signals, in case the user goes outside or comes into GPS coverage. Keeping the GPS switched on indoors consumes less energy than outdoors, as shown in Table 1, because when no signal is received, the position is not computed and hence the CPU load is reduced. Once a GPS signal is available, SenseLess stops scanning for 802.11 access points and uses the available GPS position instead. When available, preference is always given to the GPS for two reasons. Firstly, the GPS sensor uses slightly less energy than scanning for 802.11 access points, and secondly because it is more accurate in general, especially in areas with a low density of access points.

3. EXPERIMENTAL EVALUATION

We implemented SenseLess on Nokia N95 8GB smartphones, using the Python PyS60 API. To test the system, a user placed one device in each of his trouser front pockets. One device ran the SenseLess system, whereas the other was set to poll the GPS receiver every 10 seconds (hereafter referred to as "GPS-only"). Over the course of one hour, the user conducted several activities: a 150m outdoor walk, followed by a 20 minute period of sitting down, followed by a 350m walk. The user then entered a building, walked 100m inside the building, exited and entered another building, walking for another 100m and sat down for 25 minutes. This experiment corresponds to a typical walk, both indoors and outdoors, and including stationary periods.

Table 2 shows that during this experiment, SenseLess consumed less than half the energy consumed by the GPS-only system (58.8% less). Whereas SenseLess uses more CPU because of frequent accelerometer usage, it consumes less power. If a user keeps the same behavior in average (walking approximately 700m per hour, staying half of the time outdoors), the estimated battery life would be around 22 hours, against only 9 hours for a GPS-only system, that does not give positions when indoors.

Table 2: Resource consumption of SenseLess compared to GPS-only sensing.

	GPS-only	SenseLess
Average power consumption (mW)	454	118
Average CPU load (%)	1.5	2.6
Electric charge (mAh)	131	54
Estimated battery life (h)	9.2	22.2

To calculate the fidelity of the sensed data, we calculate the average distance between the positions sensed by both systems at each 10 second interval. When the user was not moving, SenseLess reported the last sensed position, hence this position was used to compute the distance with the GPS-only positions when the user was sitting down. The period of time the user was indoors was not included to compute the difference between positions. We find that the overall average difference between the two systems is 8m, with values ranging from 0.4m to 41m, and a median of 3m. This average distance between the two tracks is within the range of the GPS accuracy of our N95 devices. Indeed, the GPS accuracy varies according to various factors. The best GPS accuracy is around 10m, but since the GPS receiver was in trouser pockets in our experiment, we expect the accuracy to be worse. Two GPS coordinates computed from the same position with different devices can both be at least 10m away from the real position, and hence up to 20m away from each other. Thus, it appears that the fidelity of our sensed location data is maintained.

4. RELATED WORK

The notion of switching between radio interfaces to save energy has been examined before. Ionut et al.[2] propose EnLoc, a system that uses habitual activity of individuals to switch between 802.11 and GPS localisation methods. The Context-for-Wireless system [4] aims to save energy by switching between 802.11 and EDGE cellular interfaces depending on network availability and conditions. Ananthanarayanan and Stoica [1] also look at selecting different network interfaces to save energy — in this case 802.11 and Bluetooth.

5. CONCLUSIONS AND ONGOING WORK

We have described SenseLess, a system for saving energy consumption in sensing applications for mobile phones. We presented preliminary results from an early version of the system which demonstrate that we are able to save energy in a localisation application. In particular, our localisation runs show an increase in battery life from 9.2 hours using GPS to 22.2 hours using a combination of accelerometer, GPS and 802.11, with little change in the fidelity of the location data.

We are now interested in extending SenseLess to other services than localisation and, as part of the PVNets project[2], including our system in an experimental testbed for conducting context-aware experience sampling method [3] studies in mobile social networks.

6. REFERENCES

[1] G. Ananthanarayanan and I. Stoica. Blue-Fi: Enhancing Wi-Fi performance using Bluetooth signals. In *Proceedings of ACM MobiSys 2009*, Krakow, Poland, June 2009.

[2] I. Constandache, S. Gaonkar, M. Sayler, R. R. Choudhury, and L. Cox. EnLoc: Energy-efficient localization for mobile phones. In *Proceedings of INFOCOM Mini Conference*, Apr. 2009.

[3] S. S. Intille, J. Rondoni, C. Kukla, I. Ancona, and L. Bao. A context-aware experience sampling tool. In *CHI '03 extended abstracts on Human factors in computing systems*, pp 972–973, Ft. Lauderdale, FL, USA, Apr. 2003.

[4] A. Rahmati and L. Zhong. Context-for-Wireless: context-sensitive energy-efficient wireless data transfer. In *Proceedings of ACM MobiSys 2007*, pp 165–178, San Juan, Puerto Rico, June 2007.

[2] http://www.pvnets.org/