Energy Efficient Data Management in Smartphone Networks

Demetrios Zeinalipour-Yazti
Department of Computer Science, University of Cyprus
P.O. Box 20537, 1678 Nicosia, Cyprus
dzeina@cs.ucy.ac.cy

***Abstract*— Smartphone devices have emerged into powerful computational platforms equipped with multitude of sensors and capable of generating vast amounts of data (geo-location, audio, video, etc.) On the other hand, these devices operate on a strict energy budget, thus have a limited lifetime on a single charge. Consequently, we need to identify new energy-aware algorithms and techniques to provide innovative, feature-rich applications and services. In this white paper, we start out by providing recent trends in Smartphone technology and Smartphone networks. Our description is succeeded by an anatomy of the energy costs associated with data processing in a Smartphone Network. We conclude with prominent research directions in energy-aware data management for Smartphone networks.**

I. Introduction

The widespread deployment of Smartphone devices featuring geo-location (e.g., AGPS, Cell tower and WLAN positioning) and other sensing capabilities (e.g., proximity, ambient light, accelerometer, camera, microphone, etc.) along with Internet connectivity through WLAN, WCDMA/UMTS (3G), HSPA (3.5G) and LTE/WiMAX (4G) networks, have brought a revolution in location-oriented mobile applications and services. IMS Research and Comscore reported over 225M Smartphone sales in February 2010 (i.e., RIM, Apple, Microsoft, Google and Palm) and according to the Focal Point Group, handheld smart devices (including mobile phones and PDAs) could number 1 billion in 2010.

We define a Smartphone Network as "*a set of Smartphone devices that communicate in an unobtrusive manner, without explicit user interactions, in order to realize a collaborative or social task.*" There is already a proliferation of innovative applications founded on Smartphone networks. One example is opportunistic and participatory sensing [6], [2], [3], where applications can task mobile nodes in a given region to provide information about their vicinity using their sensing capabilities. Another example is road traffic delay estimation [9] using WiFi beams collected by Smartphone devices rather than invoking energy-demanding GPS acquisition. On the social site, Google Latitude enables users to track the places they and their social network have visited. The given service already reports millions of users, despite the controversial privacy concerns. Similarly, mobile social networking applications like Foursquare, Gowalla and Loopt enjoy enormous success in the Smartphone community.

A Smartphone is a battery-operated device, thus has a limited lifetime on a single charge. Consider a recent 4.3-inch Smartphone device, which features a rechargeable 3.7V lithium-ion battery at a capacity of 1730 mAh (i.e., 1.73 * 3600 seconds * 3.7V = 23,044 Joules). Such a device is advertised to offer 450 minutes (i.e., almost 8 hours) of talk time using WCDMA and up to 355 hours (i.e., almost 15 days!) in stand-by mode. These number assume that the users are not running any of their favorite applications nor use power hungry features such as a bright LCD, GPS, WiFi, 3G/4G and others. But even in the absence of all aforementioned features and applications, the lifetime of a Smartphone is considerably lower due to location, movement, signal strength, cell traffic and battery age. Consequently, this brings the energy lifetime of a Smartphone on a single charge, down to a day or so, depending on the Smartphone vendor and model. Finally, external battery chargers are not very practical for mobile users and fuel cells for smartphones have not reached the wide masses either (e.g., Aquafairy's AF-M3000 model requires only water to generate electricity that can charge an i-Phone in 90 minutes!)

Consequently, we need to identify new energy-aware algorithms and techniques to provide innovative, feature-rich applications and services. We start our description out with some measurements we obtained using a real Smartphone. These measurements provide an anatomy of energy costs associated with data processing on a Smartphone, enabling the following two observations: i) Local processing is an expensive operation with respect to energy consumption and must be avoided whenever possible; and ii) Data transfer over a wireless link (3G or WiFi) is again an expensive operation and must be avoided whenever possible. We conclude with prominent research directions in energy-aware data management for Smartphone networks.

II. Energy Anatomy of a Smartphone

In this subsection we present a set of real measurements we obtained from a prototype trajectory similarity search system implemented for Android smartphones [4]. Our experimental platform is an Android-based HTC Hero 2.1 Smartphone equipped with 802.11b/g and a Qualcomm MSM 7200A 528 MHz processor. We use benchmarking tools like 3gtest [1] and PowerTutor [8] to quantify the energy drain of a Smartphone device.

TABLE I

THE ENERGY ANATOMY OF A SMARTPHONE

Basic Operation on Smartphone	Power (mW = mJ/s)
CPU Idle (OS running)	175 mW
CPU Busy (Processing)	369 mW
WiFi Idle (Connected)	38 mW
WiFi Busy (Uplink 123Kbps, -58dBm)	600 mW
3G Busy	800 mW
LCD Bright. (low,hig)	300-900 mW
Function ($len(trace) = 100K$ 18B points)	**Time**
$Transmit(trace, server)$ (from Smartphone)	112 seconds
$Compare(query, trace)$ (on Smartphone)	111 seconds

In our benchmarking scenario, we are interested in transferring a GPS trajectory [11] of 100,000 18bytes data points to the query processor prior query execution. Notice that by sampling a GPS sensor every 2 seconds for one (1) year, and assuming no failures or downtimes on the Smartphone, would yield over 15M points, occupying more than 270MB of storage. We isolated the cost of uploading a single trajectory from the Smartphone to a TCP socket server over 802.11b with an uplink of 123kbps (as measured by [1]). The given operation took us 117 seconds (i.e., almost 2 minutes!), draining over 70 Joules of energy.

On the other hand, we also tried to conduct the query execution on the Smartphone device. Such a function was very processing intensive (i.e., quadratic in respect to the trajectory size), as it required the comparison of the query trajectory against the local trajectory of every Smartphone participating in the query resolution. We've isolated the time and energy cost for computing the similarity function on a single Smartphone unit. This operation took us 111 sec. (again almost 2 minutes!) and amounted to over 41 Joules of energy.

Consequently, we make the following observations: i) The asymmetric download/uplink bandwidth in these environments severely hampers the massive upload of data to a server, even under trajectory compression techniques; and ii) Local processing is an expensive operation with respect to energy consumption and must be avoided whenever possible. Notice that both aforementioned costs are accounted for a single device participating in a query, thus the costs for N smartphones participating in a query execution are much higher. Table I provides a detailed summary of our preliminary findings.

III. RESEARCH DIRECTIONS

In light of the above characteristics, we shall next identify predominant data management directions for energy efficiency in Smartphone Networks. As energy is not the sole dimension in the multi-objective optimization space of smartphone networks, additional characteristics, such as Query Processing, Efficient Data Dissemination Strategies Plans, Privacy - Security and Trust, Uncertainty, Flash Storage, Data Compression, etc. need to be taken into account in designing next generation frameworks for these environments.

A. Handle Data on the Cloud

In the Mobile Cloud Computing paradigm, Smartphone applications offload their energy-demanding functionality to powerful servers that take care of CPU-intensive tasks (e.g., voice recognition with Google Voice for iPhone or orientation processing for augmented reality apps), storage (e.g., dropbox for smartphones), network-intensive tasks (e.g., Gmail for Smartphone) and many others. This new paradigm is growing by 88% from 2009 to 2014 reaching a market of 9.5 billion USD, according to Juniper Research. This model has the following trade-off: it provides a lower duty cycle on the Smartphone device but at the same time also incorporates additional network traffic to communicate the results to the cloud. Additionally, this model suffers from low privacy levels. In particular, disclosing user data to a central entity might compromise user privacy in serious ways. That creates services that have been critized seriously in recent years [5].

B. Handle Data on the Device

In the *In-situ Computing Model* [10], [7], [4] Smartphone apps store their generated data (e.g., images, recordings, sensed parameters) on the device flash storage for latter usage. The Motorola Atrix 4G Smartphone offers 1GB of main memory and up to 48GB of secondary flash media. The local storage and indexing possibilities are certainly much more extensive with 48GB (i.e., the English version of whole Wikipedia repository is only 27GB!). Also search and retrieval (i.e., read workloads) over flash media is extremely energy efficient (as opposed to write workloads). This model has the following trade-off: it provides good privacy and energy levels but at the same time also minimizes user interactions. Consequently, to make this a viable approach would require hybrid approaches that apply concepts from distributed systems, distributed databases and peer-to-peer computing. Additionally, that would also require testbeds (such as SmartNet) for realistically measuring energy as opposed to emulations currently available.

REFERENCES

[1] 3gtest Tool, 07/2010: `http://www.eecs.umich.edu/3gtest/`.

[2] Azizyan M., Constandache I., Choudhury R.-R., "SurroundSense: mobile phone localization via ambience fingerprinting," In *MobiCom'09*.

[3] Campbell A., Eisenman S., Lane N., Miluzzo E., and Peterson R., "People-centric urban sensing," In *WICON*, 2006.

[4] Costas C., Laoudias C., Zeinalipour-Yazti D., Gunopulos D., "SmartTrace: Finding Similar Trajectories in Smartphone Networks without Disclosing the Traces," Demo in *ICDE'11*.

[5] Coursey D. "Google Apologizes for Buzz Privacy", PC World Business Center (online), Feb. 15th, 2010.

[6] Das T., Mohan P., Padmanabhan V.N., Ramjee R., Sharma A., "PRISM: platform for remote sensing using smartphones," In *MobiSys*, 2010.

[7] Konstantinidis A., Zeinalipour-Yazti D., Andreou P., Samaras G., "Multi-Objective Query Optimization in Smartphone Networks," In *MDM'11*.

[8] PowerTutor Tool, 07/2010: `http://powertutor.org/`.

[9] Thiagarajan A., et. al., "VTrack: Accurate, Energy-aware Road Traffic Delay Estimation using Mobile Phones," In *SenSys*, 2009.

[10] Zeinalipour-Yazti D., Laoudias C., Andreou M.I., Gunopulos D., "Disclosure-free GPS Trace Search in Smartphone Networks," In *MDM'11*.

[11] Zheng Y., Liu L., Wang L., Xie X., "Learning transportation mode from raw gps data for geographic applications on the web," In *WWW'08*.