

Context-Based Network Estimation for Energy-Efficient Ubiquitous Wireless Connectivity

Ahmad Rahmati, *Student Member, IEEE*, and Lin Zhong, *Member, IEEE*

Abstract—Context information brings new opportunities for efficient and effective system resource management of mobile devices. In this work, we focus on the use of context information to achieve energy-efficient, ubiquitous wireless connectivity. Our field-collected data show that the energy cost of network interfaces poses a great challenge to ubiquitous connectivity, despite decent availability of cellular networks. We propose to leverage the complementary strengths of Wi-Fi and cellular interfaces by automatically selecting the most efficient one based on context information. We formulate the selection of wireless interfaces as a statistical decision problem. The challenge is to accurately estimate Wi-Fi network conditions without powering up the network interface. We explore the use of different context information, including time, history, cellular network conditions, and device motion, to statistically estimate Wi-Fi network conditions with negligible overhead. We evaluate several context-based algorithms for the estimation and prediction of current and future network conditions. Simulations using field-collected traces show that our network estimation algorithms can improve the average battery lifetime of a commercial mobile phone for an ECG reporting application by 40 percent, very close to the estimated theoretical upper bound of 42 percent. Furthermore, our most effective algorithm can predict Wi-Fi availability for one and ten hours into the future with 95 and 90 percent accuracy, respectively.

Index Terms— Network architecture and design, local and wide-area networks.



1 INTRODUCTION

EMERGING mobile applications in healthcare and multimedia demand ubiquitously available wireless network connectivity. Despite the wide deployment of cellular networks and an increasing number of Wi-Fi hotspots, how close we are to achieving ubiquitous connectivity in our everyday life is still an open question. In this work, we show the complementary strengths of cellular and Wi-Fi networks and present context-based network estimation to leverage these strengths and provide ubiquitous energy-efficient wireless connectivity.

In order to realistically evaluate context-based network estimation for everyday life, we gathered field data about cellular and Wi-Fi networks through participants from the Rice community in Houston, Texas, a major US urban area from September 2006 to February 2007. Our data revealed that on average, our participants spent 99 and 49 percent of their everyday lives under cellular and accessible Wi-Fi networks, respectively. However, the reality of the energy costs and battery lifetime are not as bright. For example, wireless data transfer for a three-channel ECG reporting application will reduce the battery lifetime of a commercial mobile phone by over 75 percent.

Compared to Wi-Fi, cellular networks require much lower power to stay connected but consume much more energy per

bit for data transfers. Our solution is to employ Wi-Fi to improve the energy efficiency of cellular data transfer. However, unlike cellular networks, Wi-Fi has limited availability and its network interface needs to remain powered-off as much as possible due to its overwhelming idle power consumption. The key challenge is to determine if attempting a Wi-Fi connection is profitable energywise. Addressing this challenge requires estimating Wi-Fi network conditions without powering on the Wi-Fi interface.

In this work, we explore the use of context information to estimate Wi-Fi network conditions. Such context information includes time, history, cellular network conditions, and device motion. We devise effective algorithms to learn the conditional probability distribution of Wi-Fi network conditions, given context information. With the conditional probability distribution, we formulate wireless data transfer through multiple interfaces as a statistical decision problem, and validate our solutions through both trace-based simulation and field trials. Furthermore, we evaluate the performance of our algorithms for the prediction of future network conditions. On average, our most effective algorithm can predict Wi-Fi availability for one and ten hours into the future with 95 and 90 percent accuracy, respectively. For the ECG reporting application, our most effective algorithm improves the battery lifetime of a commercial mobile phone by 40 percent, close to the theoretical upper bound of 42 percent. Our field validation for the same application showed a 35 percent improvement in battery lifetime.

We have made the following contributions in this work:

- We present a reality check of network availability and the energy cost of ubiquitous connectivity in

• The authors are with the Department of Electrical and Computer Engineering, Rice University, 6100 Main Street, Houston, TX 77005.
E-mail: {rahmati, lzhong}@rice.edu.

Manuscript received 18 Oct. 2008; revised 21 Jan. 2010; accepted 25 Apr. 2010; published online 21 July 2010.

For information on obtaining reprints of this article, please send e-mail to: tmc@computer.org, and reference IEEECS Log Number TMC-2008-10-0417. Digital Object Identifier no. 10.1109/TMC.2010.139.

people's everyday life, and study the use of Wi-Fi networks to improve the energy efficiency of cellular networks. We offer a theoretical analysis on data transfer through multiple wireless interfaces. Based on our power measurements on a commercial mobile phone, we provided the theoretical upper bound of energy savings.

- We investigate several algorithms to estimate Wi-Fi network conditions using context information. Furthermore, we evaluate their performance for the prediction of future network conditions. Estimated network conditions can be utilized to minimize the energy consumption of data transfers. Our measurements and field evaluation showed that our estimation algorithms can achieve energy savings close to the theoretical upper bound.

The rest of this paper is organized as follows: In Section 2, we provide the background for our study and in Section 3, we describe our experiments, field studies, and energy model. In Section 4, we formulate data transfer through multiple wireless interfaces as a statistical decision problem and present a theoretical analysis. We present a suite of context-sensitive algorithms to estimate network conditions in Section 5, and validate them through field-collected real-world traces and a field test with a mobile healthcare application in Section 6. We present a discussion in Section 7 and address related work in Section 8. Finally, we conclude in Section 9. Early results from this work were presented in Context-for-Wireless [1].

2 BACKGROUND

In addition to traditional applications, such as e-mail and web browsing, emerging applications in mobile healthcare, multimedia, and Web 2.0 have created an insatiable appetite for ubiquitous wireless data connectivity. Emerging mobile healthcare applications seek to collect health information via bodyworn or implanted sensors and deliver health-promoting messages in situ and throughout people's everyday life. Many mobile healthcare applications depend on ubiquitous wireless connectivity from a mobile device to report health data and deliver messages on time. More importantly, they require a broad range of data sizes and allowable transfer latencies [2], [3]. The increasing information capturing and processing capabilities of mobile devices, especially mobile phones, enables users to produce and consume multimedia data pervasively. The desire to document our life and share our experience has created multimedia content for Web 2.0 applications, e.g., video blogging and YouTube. Compared with mobile healthcare applications, multimedia contents impose a much larger data size but may tolerate larger latencies. The need for energy-efficient ubiquitous connectivity that can address a broad range of data size and latency requirements has become urgent.

Indeed, cellular networks are becoming universal. According to the GSM Association [4], there are over 2.5 billion global mobile phone users as of October 2006, accounting for 40 percent of the world population. Moreover, 80 percent of the world population is covered by cellular networks, double that of year 2000. As a result, the

deployment of cellular data services has the potential to provide an unprecedented portion of the world population with increasingly ubiquitous wireless Internet access. The effective expansion of cellular network coverage roots in that it is a wireless metro-area technology and each base station covers a relatively large area. However, the potentially large distance between a mobile phone and its base station limits the energy efficiency for data transfers.

Shorter range wireless networks are also increasingly available, especially in urban, residential, and business settings. Wi-Fi, a wireless local-area network (WLAN) technology, has seen rapid expansion. It has been recently estimated that there are more than 14 million Wi-Fi access points in US homes [5]. Several major US cities have announced plans to deploy citywide Wi-Fi networks. Yet, Wi-Fi still has limited availability compared to cellular networks. On the other hand, the relatively short range of Wi-Fi enables it to achieve a much higher data rate and lower energy per MB data transfer compared to cellular. Therefore, to achieve energy-efficient ubiquitous wireless connectivity, it is important to combine the strength of both networks, as is the focus of this work.

3 EXPERIMENTAL PLATFORM

In order to realistically evaluate context-based network estimation for everyday life, we have gathered network data from a number of mobile users over six months, and obtained the energy profiles of wireless interfaces on multiple mobile phones through measurement. In this section, we present our experiment setup and findings.

3.1 Field Data Collection and Participants

We have developed and installed logging software on 11 HTC Wizard Pocket PC phones to record network characteristics with minimal intrusion to normal phone operation. At the time of the study (Fall 2006-Winter 2007), the Wizard was one of the most feature-rich commercial Pocket PC phones. The Wizard is commercially available under a variety of brands, including T-Mobile MDA and Cingular 8125. It is a Windows Mobile 5.0 GSM phone with integrated 802.11b and is capable of EDGE data connectivity. It has a battery capacity of 1250 mAh at 3.7 volts.

Our logging software, Tower Logger, measures network availability, signal levels, and context information. It records the cell tower ID, signal strength, and channel of the currently associated GSM cell and those of up to six other visible cells every 30 seconds. It also records the unique Basic Service Set Identifier (BSSID), signal strength and the security property of all visible Wi-Fi access points. With an extra sensor board, Tower Logger can also record motion information of the phone. We developed the sensor board based on the Rice Orbit sensor platform [6], which can be placed in the phone battery compartment with a cover from a larger battery (Fig. 1) and directly powered by the phone battery. The sensor board continuously samples an on-board three-axis accelerometer at 32 Hz per channel. The data are buffered by the sensor board and collected by the phone every 30 seconds.

Fourteen volunteers from the Rice campus participated in our data gathering. They carried around our experimental




Fig. 1. An HTC Wizard with the sensor board in its battery compartment (left) and an opened battery for power measurement (right).

phones for at least three weeks and could opt to use their own SIM card on the phone. We requested all participants to carry the phone as they would carry their own phone. We interviewed each participant regularly to document any significant diversion from their daily activities, for example, travels and forgetting to carry the phone. Tower Logger was installed on the phones of all 14 participants. Three participants, P1, P2, and P3, were given phones equipped with the sensor board.

Unlike war-drives or spatial coverage measurements, such as in [5], [7], we measured personal coverage: the signal strength seen in a person's daily life. More information and detailed findings regarding the participants and their traces can be found in [1].

Our first observation is that cellular network availability is extremely high (99.1 percent). On average for all our participants, cellular signal strength is above -94 dBm (three or four bars) for more than 75 percent of time. This is not surprising for a major urban area like Houston. Nevertheless, while weak signal strength maybe acceptable for voice communications, it is not true with data communications.

Our second observation is that on average, our participants spent 49 percent of their daily lives under preferred Wi-Fi networks. We define preferred networks as those from which the participants were known to have internet connectivity. We will use preferred networks for our analysis in the rest of the paper. Yet, on average, our participants are covered by preferred and/or unencrypted networks for 77 percent of their daily lives. While not all unencrypted access points are accessible, Nicholson et al. showed that over 46 percent of all unencrypted access points in residential areas were accessible [8].

3.2 Energy Cost for Data Transfers

3.2.1 Energy Model

In order to develop an energy model for wireless data transfers on the phone, we measured its power consumption under different signal conditions. The phone requires the battery in place to power up; therefore, we have to leave the battery inside for measurements. To eliminate interference from the battery charging circuitry, we measured the power transferred from the battery to the phone while the charger was disconnected, instead of measuring the power supplied by the charger. To achieve this, we placed a 0.1-ohm resistor in series with the Ground pin (Fig. 1). The phone power consumption can then be calculated by measuring the input voltage to the phone and the voltage drop on the resistor. We used the Measurement Computing USB-1608FS data acquisition device with a sampling rate of 1 KHz for our measurements.

Furthermore, we observed cellular and Wi-Fi signal strength affect data transfer speeds and success rates. To accurately account for them, we developed another logging application, Rate Logger, to measure Wi-Fi and cellular data rates and signal strength every five minutes for two of our participants, as sampling tools, for approximately one month. We utilized this information for the development of an energy model, assuming a constant network condition throughout a single transfer. This assumption is reasonable for short transfers. We can also apply our model to a long transfer by splitting it into multiple short transfers. Further details regarding the energy model are reported in [1].

We model the additional system energy cost for establishing a connection and transferring n megabytes data as

$$E = E_e + n \cdot E_t, \quad (1)$$

where E_e is the energy cost for connection establishment and E_t is the energy per bit for the transfer. To account for possible transfer failures, we assume a failed transfer will simply be retransmitted under the same network condition. Thus, when S is the transfer success rate, the energy cost is approximately

$$E = E_e + \frac{n}{S} \cdot E_t. \quad (2)$$

Table 1 provides the additional energy cost we have measured for various activities on cellular and Wi-Fi network interfaces for the HTC Wizard, in addition to the HP iPAQ hw6925, a Pocket PC phone with built-in GPS, and the HTC Tornado, a Smartphone commercially available under various brands, e.g., T-Mobile SDA.

TABLE 1
Average Measured Additional Energy Costs for Various Cellular and Wi-Fi Activities

Device	Cellular (EDGE)			Wi-Fi				
	E_m (J/min)	E_t (J/MB)		E_e (J)	E_m (J/min)		E_t (J/MB)	
		Download	Upload		PSM	No PSM	Download	Upload
HTC Wizard	1.2 – 6	40 – 50	95 – 125	5	19	61	5 – 7	7 – 11
HTC Tornado	1.2 – 2	100 – 150	170 – 300	10	6	53	4 – 6	5 – 7
HP iPAQ hw6925	1 – 2	130 – 160	220 – 330	13	4	46	5 – 14	6 – 15

It is interesting to note that we measured the additional power consumption of the built-in GPS of the HP iPAQ hw6925 at about 600 mW (36 J per minute). Such high power consumption makes it impractical to use GPS frequently, e.g. to provide context information for system management.

3.2.2 Implications

While cellular network interfaces on mobile phones are usually always on and connected, Wi-Fi interfaces are typically turned off the majority of the time due to limited availability and high power consumption, as we will discuss below. Checking for Wi-Fi availability and establishing a connection consumes considerable energy. The energy cost is closely related to the time required to either connect to an access point, or time-out assuming Wi-Fi is unavailable. We have measured the additional energy cost for the whole connection process (E_c) to be on average 5 J on the Wizard. If the phone is not associated to an access point after a specific time-out period, we assume there is no accessible Wi-Fi network available. We have found 5 seconds to be a reasonable time-out for the association process. Using this time-out, the energy cost for an unsuccessful attempt is also about 5 J for the HTC Wizard. Since we assume the phone is always connected to the cellular network, its E_c is zero.

The additional energy cost for a wireless activity refers to the extra energy consumption as compared to that if the activity is absent in an idle device. The values are averages from multiple measurements. E_t denotes the additional energy cost for transferring 1 MB data. We have ignored the effects of TCP and HTTP connection establishment, round trip latency (RTT), and TCP slow start, because the energy cost of data transfer is much larger. E_m denotes the energy cost for maintaining an idle connection for a minute, compared to the energy consumption when the corresponding network interface is powered off. For Wi-Fi, the values are shown with and without 802.11 MAC Power-Saving Mode (PSM), using the “maximum battery” setting on the phones, if available. We must note that the mobile device must stay connected to an access point to use PSM. The range for each value is based on best and worst signal strength. It is interesting that the maximum power saving setting on the HTC Wizard only provided a 3.5 percent reduction in E_m , compared to the default (balanced) setting.

Our measurements in Table 1 clearly show that cellular and Wi-Fi network interfaces have complementary energy profiles: the cellular interface can cost an order of magnitude more than the Wi-Fi interface to transfer data (E_t), but cost an order of magnitude less energy to maintain the connection (E_m). It is important to note that because the cellular network interface on mobile phones is usually left on for incoming calls, E_m should be regarded as zero for data applications using the cellular interface.

Since Wi-Fi consumes significant power even in PSM, and Wi-Fi availability is low (therefore, PSM cannot always be used), it is usually more energy-efficient to power off the Wi-Fi interface and then reestablish the connection when necessary. For example, on the HTC Wizard, it is more energy-efficient to power off the Wi-Fi interface if it has to be idle for more than 15 seconds. Checking for Wi-Fi availability and establishing a connection consume considerable energy too, E_c . This large energy overhead makes Wi-Fi inefficient for small data transfers. While newer Wi-Fi implementations

on mobile devices have been shown to be more energy-efficient [9], they still have a energy profile complementary to newer cellular interfaces because of difference in their targeted ranges and application modalities.

Indeed, neither cellular nor Wi-Fi alone can provide acceptable battery lifetime for future mobile applications requiring ubiquitous connectivity [1]. This motivates our proposal to combine their complementary strengths for achieving energy-efficient ubiquitous connectivity.

4 DATA TRANSFER ON MULTIPLE INTERFACES

We next present the problem formulation for data transfer through multiple wireless interfaces, and using our field-collected traces and energy profiles, we theoretically analyze the energy-saving potential of selecting between Wi-Fi and cellular interfaces.

4.1 Problem Description

Based on our field-collected traces, we assume that a mobile system is always connected through a low-power high-availability primary wireless network, which offers a lower data rate and consumes higher energy per bit for data transfers. For mobile phones, the primary network is the cellular network (in our case, GSM/EDGE). We assume that alternative wireless networks maybe available at limited locations or times which offer a higher data rate and consume lower energy per bit. However, they cost extra, usually significant, power to stay connected and incur significant energy and time overheads for connection establishment. For the HTC Wizard and most modern mobile handsets, this complementary network is Wi-Fi.

The problem that we propose to solve is:

If the device needs to transfer n bits of data with minimal energy consumption, should it search for an alternative network to transfer the data?

To solve this problem, we need to calculate the expected energy saving for attempting to use an alternative network, a , instead of the primary network, p , to transfer data. Assuming network a is available with condition \vec{C} , the expected energy cost for establishing a connection and transferring n bits of data through it can be estimated with (2) as

$$E_{a,available} = \frac{n}{S_a(\vec{C}_a)} \cdot Et_a(\vec{C}_a) + Ee_a. \quad (3)$$

If network a is unavailable, the energy cost of attempting an unsuccessful transfer would be the energy required to check for a connection

$$E_{a,unavailable} = Ee_a. \quad (4)$$

Since the interface for the primary network is always on and connected, there is no energy cost for connection establishment. Therefore, the energy cost to transfer the data through the primary interface is simple:

$$E_p = \frac{n}{S_p(\vec{C}_p)} \cdot Et_p(\vec{C}_p). \quad (5)$$

Let P_a denote the probability that the alternative network a is available. The expected energy saving of attempting to use network a is




Fig. 2. Potential phone battery lifetime gain achieved by optimally selecting between Cellular and Wi-Fi, for various transfer sizes and intervals, for measured and simulated Wi-Fi availability. (a) Average among all participants. (b) Simulated 80 percent Wi-Fi availability. (c) Simulated 20 percent Wi-Fi availability.

$$E_{a,p} = P_a \cdot (E_p - E_{a,available}) - (1 - P_a) \cdot E_{a,unavailable} \quad (6)$$

or

$$E_{a,p} = P_a \cdot n \cdot \left(\frac{Et_p(\vec{C}_p)}{S_p(\vec{C}_p)} - \frac{Et_a(\vec{C}_a)}{S_a(\vec{C}_a)} \right) - Ee_a. \quad (7)$$

Our algorithm for selecting the network interface for data transfer works as follows: The system calculates $E_{a,p}$ for every data transfer. If $E_{a,p}$ is negative, the system transfers the data through the primary network. Otherwise, it attempts to connect using interface a . In the case of multiple alternative interfaces, the system can choose the network with the most expected energy saving by calculating $E_{a,p}$ for all alternative networks.

In our study, the primary network is cellular and the sole alternative network is Wi-Fi. The only network condition we consider is signal strength, denoted by C_p and C_a , for cellular and Wi-Fi networks, respectively. We assume the data size, n , is provided through the software attempting the transfer. Cellular signal strength, C_p , is available without any extra energy cost. Therefore, the key to calculating $E_{a,p}$ using (7) is C_a and P_a . Before we address how to estimate C_a and P_a , we next offer the theoretical upper bound for energy savings by using multiple wireless interfaces for data transfer, assuming C_a and P_a are available with an insignificant energy cost.

4.2 Potential Energy Savings

To determine the maximum potential energy benefit of using multiple wireless interfaces, we examine the ideal case where C_a and P_a are available without an extra energy cost. P_a is equal to 1 if a Wi-Fi connection is available and 0 otherwise. Analysis of the ideal case will give the theoretical upper bound in energy savings achievable by estimating C_a and P_a .

Using the field-collected traces and the energy data for the HTC Wizard, we calculate the average battery lifetime of an otherwise idle phone using cellular-only transfer and that of one using the ideal case of data transfer through multiple wireless interfaces. Fig. 2a shows the average battery lifetime gain for different data rates and transfer intervals. We can see the use of multiple wireless interfaces has a large impact for larger data sizes and/or longer transfer intervals. Moreover, Wi-Fi network availability is a major factor too. The average Wi-Fi availability in our field-collected traces is

49 percent. In Fig. 2, we show the battery lifetime gains for hypothetical 20 and 80 percent Wi-Fi availabilities, assuming average Wi-Fi signal strength from the traces. Figs. 2b and 2c clearly show that the effectiveness of using multiple wireless interfaces, compared to a cellular-only policy, is improved with increased Wi-Fi availability.

5 CONTEXT-BASED NETWORK ESTIMATION

In Section 4.2, we calculated the theoretical upper bound of energy savings achievable by the judicious use of multiple wireless interfaces, assuming the system knows network conditions, C_a and P_a , for free. In reality, measuring Wi-Fi network conditions incurs the connection establishment energy cost, Ee_n . In this section, we present and evaluate different methods for the system to estimate C_a and P_a without powering on the Wi-Fi interface, and compare their energy savings with the theoretical upper bound. Mathematically, we can formulate the problem as estimating the conditional probability $P(WiFi|Context)$, of Wi-Fi availability given the observed *Context*. We study various sources of context information for such estimation and employ the field-collected traces to evaluate our solutions through simulation. For each simulation and each participant, we use half of the trace for training, if necessary for the estimation algorithm, and the rest for evaluation.

5.1 Naïve and Simple Solutions

A naïve solution is to have the system attempt a Wi-Fi connection for every data transfer, regardless of data size and expected network conditions. Obviously, for small data sizes, the high connection establishment energy (Ee) can easily cancel out the energy saving from the actual data transfer.

The Simple Solution uses minimal context information. It employs each user's all-time average for C_a and P_a , to decide whether to attempt a Wi-Fi connection. We use it as the baseline, along with the theoretical upper bound from Section 4.1, for evaluating other algorithms in below. Figs. 3 and 4 show the average data transfer energy and energy savings versus data size using our field-collected traces and hypothetical 20 and 80 percent Wi-Fi availabilities, respectively. It is clear that the Simple Solution provides substantial energy saving for larger data sizes and higher Wi-Fi network availability. However, the difference between the ideal case and the Simple Solution is substantial




Fig. 3. Data transfer energy versus data size. (a) Average among all participants. (b) Simulated 80 percent Wi-Fi availability. (c) Simulated 20 percent Wi-Fi availability.




Fig. 4. Data transfer energy saving versus data size. (a) Average among all participants. (b) Simulated 80 percent Wi-Fi availability. (c) Simulated 20 percent Wi-Fi availability.

for smaller data sizes and/or lower Wi-Fi coverage. In these cases, accurate network condition estimation is critical. We next propose several advanced algorithms for network condition estimation for this sake.

5.2 Hysteretic Estimation

People often stay at a certain location for rather long periods of time, therefore we expect that network conditions be related in time. This forms the basis for Hysteretic Estimation for which, the context information is previously measured network conditions (C_a and P_a). We use these values until we either have a new measurement or a predetermined time-out runs out, after which we will remeasure network conditions upon the next data transfer. Obviously, Hysteretic Estimation is more effective for shorter data transfer intervals, where network conditions are more likely to remain valid. The performance of this algorithm depends on the predetermined time-out value and how often network conditions change. The time-out can be adaptively tuned by the system based on the success rate of its previous estimations. We have tested a simple version of this algorithm with a constant 25-minute time-out. This method has the advantage of not requiring training. Fig. 12 shows, on average among all our participants, the accuracy of Hysteretic Estimation versus time for Wi-Fi availability, i.e., the probability of having the same Wi-Fi availability after a specific time.

5.3 TimeOfDay Estimation

People often spend days in a predictable fashion, e.g., at work, at home, commuting, etc. Fig. 5 shows the probability of having the same Wi-Fi availability after a specific period

of time spanning several days. It confirms that network conditions at the same time on different days are statistically related. We would expect to observe a similar correlation for the same days in different weeks had our logging period been longer.

For TimeOfDay Estimation, referred to as History Estimation in our previous work [1], we use the time of day as context information and estimate $P(WiFi|time\ of\ day)$. We divide days into 24 one-hour time slots and compute the all-time average Wi-Fi availability and signal strength for each time slot, as P_{hist} and C_{hist} , respectively, when we are in the same time slot.

5.4 Cellular Tower Estimation

Network conditions are highly correlated to geographical location. While GPS can provide very accurate location




Fig. 5. Network conditions are related at the same time of day in different days.




Fig. 6. ROC curve for Cell ID Estimation of current and future network availability.

information, its energy cost is too high for our purpose, as our measurement showed in Section 3.2.1. Furthermore, GPS systems only work outdoors. There has been considerable research on using visible cell tower IDs for localization [7], [10]. Localization requires knowing the location of existing cell towers, e.g., by means of GPS. In contrast, we directly train our algorithm and estimate Wi-Fi network conditions visible GSM cell towers and without positioning. We study the use of visible cellular towers as the context for estimating $P(WiFi|Context)$.

5.4.1 Cell ID Estimation

For Cell ID, we calculate $P(WiFi|visible\ cell\ towers)$ as a weighted sum of the probability of Wi-Fi availability for when each cell tower is visible, or $P(WiFi|i)$, where i uniquely identifies a cell tower. This can be easily recorded from previous measurements. We also store n_i , the number of previous measurements used for calculating $P(WiFi|i)$, and $C_{cell,i}$, the average Wi-Fi signal strength when Wi-Fi is available and tower i is visible

We then estimate Wi-Fi availability, $P(WiFi|visible\ cellular\ stations)$, or P_{cell} for short, as the weighed mean of $P(WiFi|i)$ among all visible cell towers.

$$P_{cell} = \frac{\sum_{i \in V} w_i \cdot P(WiFi|i)}{\sum_{i \in V} w_i}, \quad (8)$$

$$w_i = \log(n_i) \cdot (P(WiFi|i) - 0.5)^4.$$

The weights, w_i , are empirical and consist of two parts. $\log(n_i)$ gives more weight to towers that have been seen more often. In other words, the more samples we have of any tower, the more we trust its estimation. The exponent $(P_{cell,i} - 0.5)^4$ gives larger weight to towers that have $P(WiFi|i)$ close to 0 or 1. In other words, the more certain estimation is, the more we trust it.

We calculate the Cell ID estimated Wi-Fi signal, C_{cell} , as the simple mean of $C_{cell,i}$ among all visible cell towers. A receiver operating characteristic (ROC) curve shows the relationship between the fraction of true positives and false




Fig. 7. ROC curve for Fingerprinting Estimation of current and future network availability.

positives. The ROC curve for Wi-Fi availability using Cell ID Estimation is shown in Fig. 6.

5.4.2 Fingerprinting Estimation

For Fingerprinting Estimation, we use an ordered set of up to seven visible cell towers reported by the phone (i.e., the fingerprint) as the context, and calculate $P(WiFi|ordered\ set\ of\ visible\ cell\ towers)$. Similar methods have been used for localization based on Wi-Fi access points [11] and cellular base stations [12].

We calculate the Fingerprinting estimated Wi-Fi signal as the simple mean of signal strengths when Wi-Fi was available under a specific fingerprint. The ROC curve for Wi-Fi availability using Fingerprinting Estimation is shown in Fig. 7. Our results show that, on average among all our participants, fingerprinting is the most effective estimation method. However, we must note that compared to Cell ID, Fingerprinting requires considerably more memory to store training data and requires prior training at each location to be effective.

5.5 Acceleration Estimation

For P1, P2, and P3, we have recorded three-axis acceleration of their mobile phones. While there has been extensive research on extracting user-context information and physical activity from acceleration sensors [2], we use the acceleration data in a very simple fashion. We have observed the recorded acceleration values for a stationary phone are virtually constant. However, they constantly change when the phone is moving, often carried or handled by the user. In turn, wireless network conditions are expected to remain relatively constant if the phone hasn't moved much. Therefore, we compute the movement intensity, m , as:

$$m = \sum_{t=reset}^{current} [|\Delta A_x(t)| + |\Delta A_y(t)| + |\Delta A_z(t)| + c], \quad (9)$$

where $\Delta A_x(t)$, $\Delta A_y(t)$, and $\Delta A_z(t)$, are the change in acceleration along the three axes, measured at 32 Hz. c is a small positive constant to account for drift, or slow rate changes in wireless conditions. m is reset to zero every time




Fig. 8. ROC curve for the combination of Hysteretic Estimation and Cell ID Estimation.




Fig. 9. ROC curve for the combination of Acceleration Estimation and Cell ID Estimation.

network conditions are measured. It then accumulates as time passes by. Fig. 10 shows the average probability of having the same Wi-Fi availability after 30 seconds versus m for all our participants with acceleration logging. Our Acceleration Estimation assumes network conditions are unchanged and uses previously measured network conditions if the movement intensity, m , is below a predetermined threshold. Similar to Hysteretic Estimation, Acceleration Estimation is more effective for shorter data transfer intervals, where previous network conditions are more likely to remain valid. The performance of this algorithm depends on the predetermined values and how much the user actually moves around. The predetermined values can be adaptively tuned by the software based on the success rate of its previous estimation. We have tested a simple version of this algorithm with a constant threshold and c . It has the advantage of not requiring training.

At the time of our study, very few commercial phones were equipped with accelerometers, e.g., the Nokia 3220, 5500, and the Sharp V603SH. Today, accelerometers are present in an increasing number of phones. Accelerometers can be made ultra-low power. For example, the Kionix KXM52 three-axis accelerometer on our sensor board consumes less than 0.35 J/h for a 32 Hz sampling rate. As large displacements typically do not happen instantaneously, we expect the energy consumption of the accelerometer can be further reduced by reducing its measurement duty cycle, e.g. recording two seconds every ten seconds.

5.6 Combination Algorithms

One challenge we face is how to combine different estimation algorithms. The key is to combine them based on their individual strengths. The Hysteretic and Acceleration Estimation algorithms are best suited to accurately determine whether we should expect a change in network conditions. On the other hand, the Cell ID and the Fingerprinting algorithms determine network conditions irrespective of its previously measured condition. Therefore, we combine them as follows:

We first use either Hysteretic or Acceleration Estimation to determine if we should expect a change in network

conditions. If no change is expected, we will use the previously measured network conditions. If a change is expected, we will use Cell ID or Fingerprinting Estimation to calculate network conditions, as shown in Fig. 11.

Figs. 8 and 9 show the ROC curve for the combination of Cell ID Estimation and either Hysteretic or Acceleration Estimation, respectively. We can see that while each of the estimation algorithms did well alone, the key for further improved performance is combining them according to their individual strengths. The performance improvement of the combination of Fingerprinting with Hysteretic Estimation was minimal and not visible on ROC curves, as it is already extremely good (Fig. 7).

5.7 Predicting Future Network Conditions

Our above-mentioned algorithms were designed and evaluated to estimate current Wi-Fi network conditions. Yet, some applications maybe able to delay data transfer in order to utilize better network conditions. In this case, a prediction of future network conditions is necessary in addition to a cost function that accounts for the delay and power savings for each particular application. While defining and solving the cost-benefit function for particular applications is outside the scope of this work, we evaluate




Fig. 10. Probability of having the same Wi-Fi availability after 30 seconds decreases with increased movement intensity.




Fig. 11. Flowchart for combining different estimation algorithms.

the performance of several of our algorithms for the prediction of future network conditions.

Hysteretic Estimation can estimate the probability of having the same network availability as a previous measurement for a future time, based on their time difference. In contrast, the Cell ID and Fingerprinting algorithms estimate Wi-Fi network conditions irrespective of previously measured conditions. Therefore, by training these algorithms with cellular and time-of-day observations offset by Δt , we can utilize these algorithms to predict network conditions Δt in the future. The average prediction accuracies among all our participants for the Fingerprinting, Cell ID, and Hysteretic algorithms are compared in Fig. 12. For Hysteretic, we assume current conditions are known, otherwise the time difference of the last measurement should be added to the prediction time. We can see that, in particular, Fingerprinting performs very well for future network availability prediction. The ROC curve for Wi-Fi availability prediction 30, 60, and 120 minutes into the future using the Cell ID and Fingerprinting algorithms are shown in Figs. 6 and 7, respectively.

6 REAL-LIFE PERFORMANCE EVALUATION

6.1 Trace-Based Evaluation

In this section, we show how the performance of each network condition estimation algorithm translates into




Fig. 12. Prediction accuracy of different algorithms. For Hysteretic, we assume current conditions are known.

real-life energy savings. We use field-collected traces to calculate the energy consumption of each algorithm for a typical ECG reporting application with data transfers of 270 KB every 5 minutes, the data rate of a typical three-channel ECG (each sampled at 200 Hz with 12-bit resolution [3], without compression). For each algorithm, we compare its energy savings beyond that of the Simple Solution, compared to the ideal case. For example, if the average energy cost for a data transfer is 23 J for one of our estimation algorithms, 30 J for the simple case, and 20 J for the ideal case, then the effectiveness of the estimation algorithm is

$$\frac{30 - 23}{30 - 20} = 70\%.$$

The effectiveness of all estimation algorithms, except Acceleration Estimation, are shown in Fig. 13. We can see that on average, fingerprinting is the most effective estimation method. However, fingerprinting requires prior training at every specific location. In addition, Fingerprinting and Cell ID require detailed cell tower information from a GSM interface. On the other hand, Hysteretic, TimeOf-Day, and Acceleration estimation can work on devices with Wi-Fi but without GSM interfaces, e.g., the iPod Touch.

Detailed analysis of the performance of our estimation algorithms provides interesting findings. First, we can




Fig. 13. Effectiveness of estimation algorithms for minimizing energy consumption for all participants. Effectiveness is defined as the achieved energy savings beyond that of the Simple Solution, compared to the ideal case.




Fig. 14. For participants with longer commutes, the effectiveness of Cell ID increases while the effectiveness of Fingerprinting decreases.

clearly see that *combining multiple heterogeneous sources of context information provides enhanced performance*.

Second, as shown in Fig. 14, on average among all the participants, *Cell ID Estimation performs better with increased commute distance*, which is indicated by the straight-line distance between their homes and workplaces. This is the opposite of Fingerprinting Estimation, and expected; while the visible cell towers change minimally at short distances, their fingerprints can change greatly.

Third, as shown in Fig. 15, on average among all the participants, *the performance of our estimation algorithms is better with moderate or low Wi-Fi availability*. Indeed, when Wi-Fi availability is high, the Simple Solution is already very close to the ideal case.

Finally, as previously discussed, the data transfer interval impacts the efficiency of Hysteretic and Acceleration Estimation algorithms. To highlight the impact, we present the performance of those algorithms for P1, P2, and P3 in Fig. 16. The data size for each transfer is 270 KB for all three intervals. As expected, *their performance is higher for shorter transfer intervals, when network conditions are more likely to remain valid*.

In this section, we presented and evaluated several algorithms for estimating C_a and P_a with various context information. We showed that they have very good performance when Wi-Fi estimation is important, i.e., lower Wi-Fi availability. Moreover, our algorithms have complementary




Fig. 15. Our estimation algorithms performed better when Wi-Fi availability is moderate and low.




Fig. 16. Hysteretic and Acceleration Estimation perform better with shorter transfer intervals.

strengths. For example, Hysteretic Estimation works well for short transfer intervals and for participants with high Wi-Fi availability. TimeOfDay Estimation works very well for participants with regular hours and locations of Wi-Fi availability. We also observed that the performance of Cell ID Estimation is improved with increased mobility, which is opposite of Fingerprinting. Finally, combining different estimation algorithms can further improve performance.

Our goal is by no means to devise the best possible estimation algorithm. Instead, our main focus is to demonstrate how readily available context information, such as time and cell tower ID, can be used to estimate Wi-Fi network conditions. We have kept our algorithms simple and have limited the number of parameters in them. Therefore, we expect their overhead to be negligible, and that they will generalize well [13].

6.2 Field Validation

To validate the effectiveness of context-based network estimation, we have implemented it for a mobile ECG reporting application, which collects ECG data from body-worn wireless (e.g., Bluetooth) sensors and periodically reports it to an Internet server. Since the focus of this work is on data transfer between the mobile device and the Internet server, we assume ECG data are already available on the phone. We use the same data rate and interval as in Section 6.1, 270 KB every 5 minutes.

We have developed the ECG reporting software that automatically runs on the phone every 5 minutes. Our software has two modes of operation. In cellular-only mode, it transfers data only using cellular. In context-based interface selection mode, it uses our Hysteretic algorithm, with a constant 25 minute time-out, to estimate network conditions and decide whether to power up the Wi-Fi adapter and attempt a Wi-Fi transfer. We chose Hysteretic Estimation because it has relatively good performance without training or extra hardware and our simulation in Section 6.1 showed that it is effective when transfer intervals are short. If the connection fails, is dropped, or no response is received from the server, the transfer is considered unsuccessful and our software will attempt a resend. Whenever a Wi-Fi resend fails, our software will resort to cellular.

We installed our ECG reporting software on two HTC Wizard phones and gave them to participants P1 and P2,

TABLE 2
Phone Battery Lifetime during Our Field Test

Participant	Experiment	Cellular only (h)	Context-for- Wireless (h)
P1	1	12.3	16.2
	2	20.3	20.4
	3	17.5	21.5
	Average	16.7	19.4
P2	1	13.2	21.6
	2	15.0	21.4
	3	14.4	23.6
	Average	14.2	22.2
Average		15.4	20.8

who used the phones as their primary mobile phones during our experiments. We ran six experiments for each participant, half of them in cellular-only mode and the other half using context-based interface selection. We measured the battery lifetime as the operational time between when the phone is disconnected from its charger (with a fully charged battery) and when the phone automatically shuts itself off due to low battery, with no charging in between. The battery lifetimes for cellular-only mode and context-based interface selection mode are shown in Table 2. The average battery lifetime gain was 35 percent.

For comparison, the average battery lifetime gain simulated using the field-collected traces was 29 percent for the Simple Solution, 37 percent for Hysteretic Estimation, 39 percent for the combined History, Cell Tower ID, and Hysteretic Estimation, 40 percent for Fingerprinting, and 42 percent for the ideal upper bound.

7 DISCUSSION

Our reality check and studies of context-based network estimation are limited to our participants, who were affiliated with Rice University and spent a significant portion of their everyday life under the Rice campus Wi-Fi coverage. Moreover, the majority of our field-collected traces were from Houston, a major metropolitan area; cellular networks in rural or suburb areas can have different characteristics. Therefore, we expect our participants might have painted a more optimistic picture of ubiquitous connectivity than what is available for the general population. Nonetheless, with the increasing availability of Wi-Fi and expanding deployment of cellular networks, we believe our participants represent the trend in the development of ubiquitous wireless connectivity.

We studied the use of Wi-Fi and cellular networks, but the approach of selecting between multiple network interfaces to achieve energy-efficient ubiquitous connectivity is general. While newer cellular technologies such as 3G will support higher data rates, they are still metro-area networks and each base station covers a relatively large area. The long range radio communication between the mobile phone and its base station will still have high energy per bit requirement compared to shorter range wireless technologies, such as the current and future wireless LAN

technologies. Although Wi-Fi hotspots are on the rise, their availability will still remain well below cellular networks due to their short range. Therefore, we expect the availability and range versus energy trade-off behind different wireless interfaces to remain valid in the foreseeable future. Emerging technologies, such as WiMAX, can be yet another network interface to select from.

In our work, we assumed that cellular and Wi-Fi networks operated in isolation. Cooperation between cellular and Wi-Fi networks will create additional opportunities in estimating Wi-Fi network condition and availability. For example, due to E911 requirements, cellular network providers are already aware of approximate location of their subscriber phones. Therefore, they could potentially send prior known Wi-Fi availability data, based on location, to subscribers. This could be especially appealing as many cellular providers also offer Wi-Fi hotspots, including the four major US providers, AT&T, Sprint, T-Mobile, and Verizon. Nevertheless, our network estimation methods can still be useful to improve the accuracy of Wi-Fi estimation.

It is important to note that using context information to estimate wireless network conditions (Context-for-Wireless) is a mechanism. We used context information to estimate current and future network conditions. Based on the prediction, system policies can be devised to prefetch wireless data if the network is predicted to degrade, or to buffer data transfers while respecting latency requirements if the network is predicted to improve.

8 RELATED WORK

Our use of multiple wireless interfaces for data transfer resonates with a considerable body of work employing a secondary low-power wireless interface for improving Wi-Fi energy efficiency. For example, *Wake-on-Wireless* [12] uses a low-power radio interface to transmit control information so that Wi-Fi can stay powered-off most of the time. *Coolspots* [14] employs Bluetooth to improve power efficiency of Wi-Fi. However, these works and others [15], [16] target at improving Wi-Fi energy efficiency and are restricted by limited Wi-Fi availability. In contrast, our approach utilizes Wi-Fi to improve the efficiency of cellular networks and provide energy-efficient ubiquitous connectivity. Armstrong et al. also found that Wi-Fi is more energy-efficient for transfers of large data sizes and subsequently selected the wireless interface based on the data size [17]. However, they assumed that Wi-Fi is always available and sought to reduce the energy cost of data transfer without considering that of establishing or maintaining a Wi-Fi connection, which is nontrivial, as our work demonstrated. Moreover, while they only showed that Wi-Fi is more energy-efficient when the data size is larger than 30 KB, we provided more detailed energy profiles for Wi-Fi and cellular network interfaces with an analytical model. Integrating Wi-Fi and cellular networks and seamlessly switching between them on various layers has also been widely studied. For example, *Always Best Connected* [18] seeks to achieve best performance and coverage. In contrast, our solutions work at the application layer and do not require cooperation between Wi-Fi and cellular networks.

Related to our use of Wi-Fi, Virgil [8] automatically discovers and selects access points with faster connection establishment, which leads to energy savings as well. Virgil is complementary to our approach and can be readily incorporated to improve the connection establishment in multi-interface data transfer. Related to our reality check, Bychkovsky et al. [5] presented Wi-Fi access point data collected by war-drives to evaluate the possibility of a large-scale Wi-Fi network made up by volunteered home Wi-Fi hotspots. Intel Place Lab [7], [10] also collected extensive network data on GSM cellular networks for the sake of positioning. These works provide a reality check of the spatial network availability. In contrast, our work is targeted at the personal coverage of wireless networks, i.e., how networks are available throughout people's daily life.

Place Lab [7], [10] is also related to our endeavor on estimating Wi-Fi network conditions. With positioning accuracy around 100 m, Place Lab can indeed provide important clues for Wi-Fi network condition, which is highly position-related. However, Place Lab requires a premapping of GSM towers (base stations) with GPS. Instead, our estimation methods seek to learn the direct relations between Wi-Fi network conditions and context information, including cellular network conditions, thus eliminating the need for mapping. In Turducken [19], the authors used a Wi-Fi detector to reduce Wi-Fi connection attempts that fail. However, a Wi-Fi detector can only detect the existence of Wi-Fi signal; it cannot determine whether it is from an accessible network. Furthermore, the energy cost and benefit of the Wi-Fi detector were not addressed. The authors of [20] used GPS-based movement prediction to reduce wireless communication energy between two mobile nodes. They targeted at a specific application scenario with a single wireless network. More importantly, GPS only works outdoors and its energy cost is too high for our purpose, as showed in Section 3.2.1.

Our work estimates current and future Wi-Fi conditions based on context information and past Wi-Fi conditions. In contrast, the authors of [21] focus on predicting detailed Wi-Fi conditions, including available bandwidth, tens of seconds into the future using current Wi-Fi conditions, location information, and a mobility model. Yet, as our measurements indicate, measuring current Wi-Fi conditions incurs a considerable energy cost. Future network condition prediction enable the postponing of data transfers for more favorable conditions. For example, the authors of [22] employ data prefetching when more energy-efficient networks are available, in order to lower the overall energy cost. These works are complementary to Context-for-Wireless.

While we have used context information to assist in selecting the most efficient wireless interface for data transfers, context information can be used to manage a variety of other system resources, such as battery energy and power consumption [23], [24], [25]. We believe that our approach for determining network conditions from context information can be applied to other system resources as well, especially those which are location-dependent.

9 CONCLUSION

The driving vision of our work is to leverage the complementary strength of multiple available wireless networks for

energy-efficient ubiquitous connectivity. We achieve this by estimating network condition using context information.

Using findings from our recent field study, we showed that the energy cost of ubiquitous network connectivity is overwhelming. Our field measurements presented a challenging picture for emerging mobile applications that rely on ubiquitous connectivity: Our experiment showed that the data transfer required for mobile ECG reporting reduced the battery lifetime of our mobile phone to an average of 15.4 hours. We also showed that Wi-Fi and cellular network interfaces (802.11b and GSM/EDGE in our study) have energy profiles with complementary strengths. Therefore, we proposed to leverage increasingly available Wi-Fi networks to improve the data transfer energy efficiency of cellular networks. Our theoretical analysis showed that judiciously choosing between network interfaces can considerably improve battery lifetime under a broad range of application requirements, while careless use of Wi-Fi can have a severe power impact.

We formulated data transfer through multiple wireless interfaces as a statistical decision problem and explored various contextual clues to learn the conditional probability distribution of Wi-Fi network conditions in order to solve it. We showed that while individual algorithms perform significantly better than our baseline, combining different algorithms based on their individual strengths can improve performance further. Our best algorithms can achieve battery lifetime improvement close to the theoretical limit, without additional hardware on phones. Furthermore, our evaluation shows our context-based algorithms are highly effective at predicting future network conditions. In particular Fingerprinting, can, on average, predict Wi-Fi availability for one and ten hours into the future with 95 and 90 percent accuracy, respectively. Such high predictability suggests the possibility of making tradeoffs between energy efficiency and delay in wireless data transfer. We also explored the use of a three-axis accelerometer to sense motion and movement using an additional sensor board. Today, such a sensor is present in an increasing number of phones.

We validated our solutions using data collected from the everyday lives of 14 participants as well as using field trials. By judiciously choosing a wireless interface using Context-for-Wireless, with our most simple algorithm, Hysteretic Estimation, we were able to improve the average battery lifetime by 35 percent in field trials with real phone usage. Yet more sophisticated algorithms, some described in this work, will produce even more energy savings.

ACKNOWLEDGMENTS

This work was supported in part by US National Science Foundation awards CNS/CSR-EHS 0720825, CNS/NeTS-WN 0721894, and CNS/CRI 0751173, and the TI Leadership University Fund. The authors would like to thank their anonymous reviewers whose comments helped to improve this paper, and in particular, one for suggesting the use of cellular fingerprints as context information. Last but not least, the authors also thank the volunteers who participated in their data collection and field validation.

REFERENCES

- [1] A. Rahmati and L. Zhong, "Context-for-Wireless: Context-Sensitive Energy-Efficient Wireless Data Transfer," *Proc. ACM MobiSys*, pp. 165-178, 2007.
- [2] C.V.C. Bouten, K.T.M. Koekkoek, M. Verduin, R. Kodde, and J.D. Janssen, "A Triaxial Accelerometer and Portable Data Processing Unit for the Assessment of Daily Physical Activity," *IEEE Trans. Biomedical Eng.*, vol. 44, no. 3, pp. 136-147, Mar. 1997.
- [3] P.S. Hamilton and W.J. Tompkins, "Estimation of Rate-Distortion Bounds for Compression of Ambulatory ECGs," *Proc. Int'l Conf. Eng. in Medicine and Biology Soc.*, vol. 2, pp. 763-764, 1989.
- [4] Universal Access Report, GSM Assoc., http://gsmworld.com/our-work/public-policy/regulatory-affairs/universal_access.htm, 2006.
- [5] V. Bychkovsky, B. Hull, A. Miu, H. Balakrishnan, and S. Madden, "A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks," *Proc. ACM MobiCom*, pp. 50-61, 2006.
- [6] Rice Orbit Platform, Rice Efficient Computing Group, <http://www.ruf.rice.edu/~mobile/orbit>, 2010.
- [7] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hughes, and F. Potter, "Place Lab: Device Positioning Using Radio Beacons in the Wild," *Proc. Int'l Conf. Pervasive Computing (Pervasive)*, pp. 116-133, 2005.
- [8] A.J. Nicholson, Y. Chawathe, M.Y. Chen, B.D. Noble, and D. Wetherall, "Improved Access Point Selection," *Proc. ACM MobiSys*, pp. 233-245, 2006.
- [9] P. Eronen, "TCP Wake-Up: Reducing Keep-Alive Traffic in Mobile IPv4 and IPsec NAT Traversal," Nokia technical report, 2008.
- [10] M.Y. Chen, T. Sohn, D. Chmelev, D. Haehnel, J. Hightower, J. Hughes, A. LaMarca, F. Potter, I. Smith, and A. Varshavsky, "Practical Metropolitan-Scale Positioning for GSM Phones," *Proc. Int'l Conf. Ubiquitous Computing (UbiComp)*, pp. 157-169, 2006.
- [11] P. Bahl and V.N. Padmanabhan, "RADAR: An In-Building RF-Based User Location and Tracking System," *Proc. IEEE INFOCOM*, vol. 772, pp. 775-784, 2000.
- [12] E. Shih, P. Bahl, and M.J. Sinclair, "Wake on Wireless: An Event Driven Energy Saving Strategy for Battery Operated Devices," *Proc. ACM MobiCom*, pp. 160-171, 2002.
- [13] T. Hastie, R. Tibshirani, and J.H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2001.
- [14] T. Pering, Y. Agarwal, R. Gupta, and R. Want, "CoolSpots: Reducing the Power Consumption of Wireless Mobile Devices with Multiple Radio Interfaces," *Proc. ACM MobiSys*, pp. 220-232, 2006.
- [15] Y. Agarwal, R. Chandra, A. Wolman, P. Bahl, K. Chin, and R. Gupta, "Wireless Wakeups Revisited: Energy Management for VoIP over Wi-Fi Smartphones," *Proc. ACM MobiSys*, pp. 179-191, 2007.
- [16] C.F. Chiasserini, R.R. Rao, and D. di Elettronica, "Improving Energy Saving in Wireless Systems by Using Dynamic Power Management," *IEEE Trans. Wireless Comm.*, vol. 2, no. 5, pp. 1090-1100, Sept. 2003.
- [17] T. Armstrong, O. Trescases, C. Amza, and E. de Lara, "Efficient and Transparent Dynamic Content Updates for Mobile Clients," *Proc. ACM MobiSys*, pp. 56-68, 2006.
- [18] E. Gustafsson and A. Jonsson, "Always Best Connected," *IEEE Wireless Comm.*, vol. 10, no. 1, pp. 49-55, Feb. 2003.
- [19] J. Sorber, N. Banerjee, M.D. Corner, and S. Rollins, "Turducken: Hierarchical Power Management for Mobile Devices," *Proc. ACM MobiSys*, pp. 261-274, 2005.
- [20] S. Chakraborty, Y. Dong, D.K.Y. Yau, and J.C.S. Lui, "On the Effectiveness of Movement Prediction to Reduce Energy Consumption in Wireless Communication," *IEEE Trans. Mobile Computing*, vol. 5, no. 2, pp. 157-169, Feb. 2006.
- [21] A.J. Nicholson and B.D. Noble, "BreadCrumbs: Forecasting Mobile Connectivity," *Proc. ACM MobiCom*, pp. 46-57, 2008.
- [22] S. Gitzenis and N. Bambos, "Joint Transmitter Power Control and Mobile Cache Management in Wireless Computing," *IEEE Trans. Mobile Computing*, vol. 7, no. 4, pp. 498-512, Apr. 2008.
- [23] N. Banerjee, A. Rahmati, M.D. Corner, S. Rollins, and L. Zhong, "Users and Batteries: Interactions and Adaptive Energy Management in Mobile Systems," *Proc. Int'l Conf. Ubiquitous Computing (UbiComp)*, pp. 217-234, 2007.
- [24] A. Rahmati and L. Zhong, "Human-Battery Interaction on Mobile Phones," *Pervasive and Mobile Computing*, vol. 5, no. 5, pp. 465-477, 2009.

[25] N. Ravi, J. Scott, L. Han, and L. Iftode, "Context-Aware Battery Management for Mobile Phones," *Proc. Int'l Conf. Pervasive Computing and Comm. (PerCom)*, pp. 224-233, 2008.



Ahmad Rahmati received the BS degree in computer engineering from the Sharif University of Technology, Tehran, Iran, in 2004, and the MS degree in electrical and computer engineering from Rice University, Houston, Texas, in 2008, where he is currently working toward the PhD degree. He was a research intern at AT&T Labs-Research, Motorola Labs, and Deutsche Telekom R&D Lab, in Summer 2006, Summer 2008, and Spring 2010, respectively, and received the ACM MobileHCI best paper award in 2007. His research interests include the design and applications of efficient, context-aware mobile systems, system power analysis and optimization, and human-computer interaction. He is a student member of the IEEE and the IEEE Computer Society.



Lin Zhong received the BS and MS degrees from Tsinghua University in 1998 and 2000, respectively, and the PhD degree from Princeton University in September 2005. He was with NEC Labs, America, for the summer of 2003 and with Microsoft Research for the summers of 2004 and 2005. He joined the Department of Electrical and Computer Engineering, Rice University, as an assistant professor in September 2005. He received the AT&T Asian-Pacific Leadership Award in 2001 and the Harold W. Dodds Princeton University Honorary Fellowship for 2004-2005. He coauthored a paper that was selected as one of the 30 most influential papers in the first 10 years of Design, Automation & Test in European conferences. He and his students received best paper awards from ACM MobileHCI 2007 and IEEE PerCom 2009. His research interests include mobile and embedded system design, human-computer interaction, and nanoelectronics. His research has been funded by the US National Science Foundation, Motorola Labs, Nokia, Texas Instruments, and Microsoft Research. He is a member of the IEEE and the IEEE Computer Society.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.