

Context-for-Wireless: Context-Sensitive Energy-Efficient Wireless Data Transfer

Ahmad Rahmati and Lin Zhong
Department of Electrical & Computer Engineering
Rice University, Houston, TX 77005
{rahmati, lzhong}@rice.edu

ABSTRACT

Ubiquitous connectivity on mobile devices will enable numerous new applications in healthcare and multimedia. We set out to check how close we are towards ubiquitous connectivity in our daily life. The findings from our recent field-collected data from an urban university population show that while network availability is decent, the energy cost of network interfaces poses a great challenge. Based on our findings, we propose to leverage the complementary strength of Wi-Fi and cellular networks by choosing wireless interfaces for data transfers based on network condition estimation. We show that an ideal selection policy can more than double the battery lifetime of a commercial mobile phone, and the improvement varies with data transfer patterns and Wi-Fi availability.

We formulate the selection of wireless interfaces as a statistical decision problem. The key to attaining the potential battery improvement is to accurately estimate Wi-Fi network conditions without powering up its network interface. We explore the use of different context information, including time, history, cellular network conditions, and device motion, for this purpose. We consequently devise algorithms that can effectively learn from context information and estimate the probability distribution of Wi-Fi network conditions. Simulations based on field-collected traces show that our algorithms can improve the average battery lifetime of a commercial mobile phone for a three-channel electrocardiogram (ECG) reporting application by 39%, very close to the theoretical upper bound of 42%. Finally, our field validation of our most simple algorithm demonstrates a 35% improvement in battery lifetime.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design

C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks

General Terms

Measurement, Performance, Design, Experimentation, Human Factors

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobiSys '07, June 11–14, 2007, San Juan, Puerto Rico, USA.
Copyright 2007 ACM 978-1-59593-614-1/07/0006...\$5.00.

Keywords

Energy-efficient wireless, multiple wireless interfaces, context-for-wireless

1. INTRODUCTION

Emerging mobile applications in healthcare and multimedia demand ubiquitously available wireless network connectivity. Despite of the wide deployment of 2.5G & 3G cellular networks and an increasing number of Wi-Fi hot-spots, it is still an open question how close we are towards achieving ubiquitous connectivity in our everyday life. In this work, we present our findings from our recent field collected data about wireless network availability and energy cost, and investigate context-based Wi-Fi estimation for energy-efficient wireless data transfer.

As a reality check and case study, 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 showed a bright picture regarding network availability: on average, 99% and 49% of our participants' everyday life was under cellular and accessible Wi-Fi network, respectively. However, the reality about the energy cost and battery lifetime is not as bright. For example, the data transfer for a three-channel ECG reporting application will reduce the battery lifetime of a commercial mobile phone to below a quarter of the original.

Our approach towards energy-efficient ubiquitous connectivity is based upon the complementary energy profiles of Wi-Fi and cellular network interfaces. Compared with Wi-Fi, cellular networks require much lower power to stay connected but incur a much higher energy per MB transfer. Our solution is to employ Wi-Fi to improve the energy efficiency for cellular data transfer. However, unlike cellular networks, Wi-Fi has limited availability and its network interface needs to remain powered-off due to its overwhelming power consumption. The key challenge is to estimate whether attempting a Wi-Fi connection is worthwhile, in terms of energy. This in turn requires estimating Wi-Fi network conditions without powering on the Wi-Fi interface.

To address this challenge, we explore the use of different context information including time, history, cellular network conditions, and device motion, to estimate Wi-Fi network conditions, which we call *Context-for-Wireless*. We devise effective algorithms to learn the probability distribution of Wi-Fi network conditions from the context information. With the probability distribution, we formulate the data transfer through multiple wireless interfaces as a statistical decision problem. We validate our solutions through both trace-based simulation and field trials. For the ECG reporting application, our most effective estimation algorithm improves the battery lifetime of a

commercial mobile phone by 39%, close to the theoretical upper bound of 42%. Our field validation for the same application and our most simple estimation algorithm showed a 35% improvement in battery lifetime.

We have made the following contributions in this work:

- To the best of our knowledge, we presented the first reality check of network availability and the energy cost of ubiquitous connectivity in people's everyday life.
- We studied the use of Wi-Fi networks to improve the energy efficiency of cellular networks. We offered 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 by opportunistically selecting between the Wi-Fi and cellular interfaces.
- We investigated various context-sensitive algorithms to estimate Wi-Fi network conditions (Context-for-Wireless). The estimated network condition is used to minimize the energy consumption of data transfers. Our measurement and field evaluation showed that our estimation algorithms can achieve energy savings close to the theoretical upper bound.

The rest of the paper is organized as follows. In Section 2, we provide the background for our study and discuss related work. In Section 3, we describe our experiments and field studies and present our findings. In Section 4, we formulate data transfer through multiple wireless interfaces as a statistical decision problem, present a theoretical analysis, and offer a suite of context-sensitive algorithms to estimate network conditions. We validate our algorithms through a real field test with a mobile healthcare application in Section 5. We discuss our study in Section 6 and address related work in Section 7. Finally, we conclude in Sections 8.

2. BACKGROUND

In addition to traditional applications, such as email and web browsing, emerging applications in mobile healthcare, multimedia, and Web 2.0 have created an insatiable appetite for ubiquitous wireless data. Emerging mobile healthcare applications [19, 22] seek to collect health information, via body-worn 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 for reporting health data and delivering messages in time. More importantly, they require a broad range of data size and allowable transfer latencies [5, 13]. The increasing information capturing and processing capabilities of mobile devices, especially mobile phones, enables users to produce and consume multimedia data in a pervasive fashion. 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 never been as urgent as it is today.

Indeed, cellular networks are becoming universal. According to the GSM Association [2], there are over 2.5 billion global mobile phone users as of October 2006, accounting for 40% of the world population. Moreover, 80% of the world population is

covered by cellular networks, double that of year 2000. As a result, the deployment of 2.5G/3G data services has the potential to enable 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 potential distance between a mobile phone and its base station limits the achievable 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 is estimated that there are more than 14 million Wi-Fi access points in US homes [8]. Several major US cities have announced plans to deploy city-wide Wi-Fi networks. Compared with cellular networks, Wi-Fi still has very limited availability. However, its relative short range enables it to have 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. REALITY CHECK

In order to check how close we are toward ubiquitous connectivity in our everyday life, over six months, we have gathered network data from a number of mobile users, and obtained the energy profiles of wireless interfaces on multiple mobile phones through measurement. We next present our findings.

3.1 Experimental Setup

We have used multiple HTC Wizard PDA phones for our data collection. The HTC 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 1250mAh at 3.7 volts.

We have developed logging software to record various network characteristics with minimal intrusion to the normal phone operation. We have converted eleven HTC Wizards into experimental mobile phones by installing our logging software.

14 volunteers from the Rice campus participated in our data gathering. They carried around our experimental 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 carrying the phone. Our participants are described in Table 1.

3.1.1 Rate Logger: Measuring Data Rates

Our first logging software, called Rate Logger, measures Wi-Fi and cellular data rates (both uplink and downlink), round-trip latencies, and signal strength, every five minutes. When a network is available, Rate Logger will attempt to send and receive a specific amount of data to and from our lab server using HTTP POST and GET. The lab server is a Windows PC on Rice campus network with a static IP address. We used transfer sizes of 20KB for cellular and 200KB for Wi-Fi. We found those sizes adequate after testing approximately one thousand cellular transfers of 5, 20, and 80 KB, and Wi-Fi transfers of 50, 200, and 800 KB. If a transfer successfully finishes, Rate Logger will record the data

Table 1. Information about data collection participants

Participant	Profession	Cellular carrier	Comments	Length (weeks)	# Cell towers observed	# Wi-Fi Access points observed		
						Preferred	Preferred + unencrypted	All
P1	Faculty	Cingular	Multiple trips out of town	4	497	113	416	1063
P2	Grad	T-Mobile	Lives close to campus	4	349	224	398	703
P3	Grad	Cingular		4	427	136	258	456
P4	Undergrad	Cingular	Lives on campus	4	445	713	982	1397
P5	Undergrad	Cingular	Lives on campus	4	701	713	1226	2048
P6	Staff	Cingular	Multiple trips out of town	3	968	113	310	846
P7	Grad	Cingular		3	706	96	210	477
P8	Grad	Cingular	Lives close to campus	4	606	626	1300	2440
P9	Grad	T-Mobile		4	287	490	913	1627
P10	Grad	Cingular	Lives close to campus	4	269	394	604	852
P11	Grad	T-Mobile	Lives close to campus	4	748	599	980	1537
P12	Grad	T-Mobile		4	963	499	1011	1898
P13	Grad	T-Mobile	Lives close to campus	4	922	628	1072	1714
P14	Undergrad	Cingular		4	248	596	972	1397

transfer time and round-trip latency. If the connection is dropped before the transfer could finish or no response is received from the server, Rate Logger records the transfer as a failed attempt. We define the *transfer success rate* as the percentage of successful transfers to all attempted transfers.

The first version of Rate Logger records the cellular and Wi-Fi signal strength reported by the operating system, Windows Mobile. Windows provides cellular signal strength in the form of a “raw signal strength” integer between 0 and 100, instead of the actual signal strength in dBm. Even worse, the Windows on our HTC Wizards only reports a limited number of discrete values between 0 and 100. It also incorrectly reports Wi-Fi signal strength. We later upgraded our software to record the driver-reported signal strength levels (in dBm), in addition to the Windows-reported ones. However, the majority of our Rate Logger measurements only had Windows-reported signal strength. Therefore, we created a lookup table based on the traces collected by the second version, and converted the Windows-reported signal strength from the traces collected by the first version to their actual levels in dBm. Because Windows reported cellular signal strength above -80 dBm as 100 (maximum), we treat all that cellular signal levels above -80 dBm as -80 dBm.

In Experiment A, we installed Rate Logger on the phones for P1 and P2, who have unlimited data plans from Cingular and T-Mobile, respectively. We used the participants as *sampling tools* to measure data rates for the different signal strength they encountered. We logged them for approximately one month.

3.1.2 Tower Logger: Measuring Network Conditions

Our second logging software, called Tower Logger, measures network availability and 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 6 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 [1], which can be placed in the phone battery compartment with a cover from a larger battery (Figure 1) and directly powered by the phone battery. The sensor board continuously samples an on-board three-axis accelerometer at 32Hz per channel. The data is buffered by the sensor board and collected by the phone every 30 seconds.




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

In Experiment B, we installed Tower Logger on the phones of all 14 participants. Participants P1, P2, and P3 were given phones equipped with the sensor board.


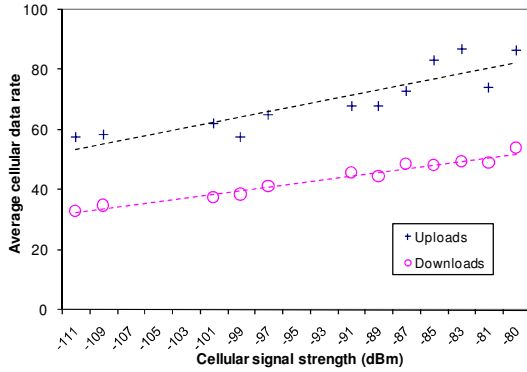


Figure 2. Signal strength impacts data rates of successful transfers for cellular (GSM/EDGE)(left) and Wi-Fi (right)

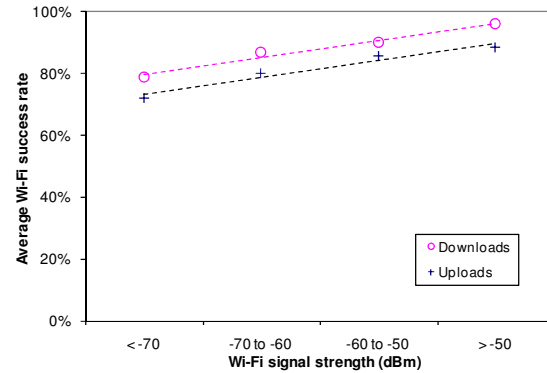
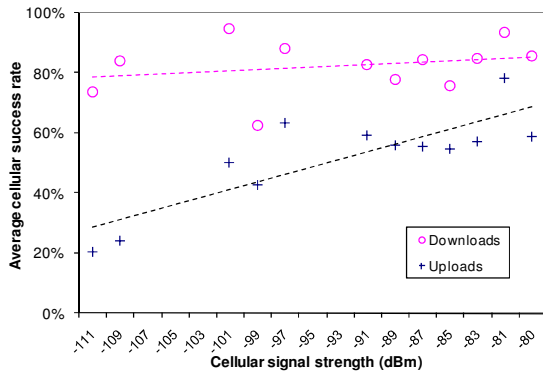


Figure 3. Signal strength impacts transfer success rates for cellular (GSM/EDGE) (left) and Wi-Fi (right)

3.1.3 Measuring Phone Power Consumption

We have developed software to measure the power consumption of various phone components under controlled conditions. Due to the smart battery interface [6], most modern phones only boot with an actual battery inside. Therefore, we had to leave the battery inside for measurements. To eliminate interference from the battery charging circuitry, we measured the power transferred between the battery and the phone, instead of between the charger and the phone. We opened the battery packaging and placed a 0.1-ohm resistor in series with the Ground pin (Figure 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 for our measurements. We found that the sampling rate of 1 KHz is adequate for the power characteristics important to us.

3.2 Real-Life Data Rates vs. Signal Strength

In Experiment A, we collected information regarding the success rates of data transfers and the data rates of successful transfers from P1 and P2. Although P1 and P2 had two different cellular carriers, our analysis showed little difference between the two carriers. Therefore, we will report their cellular data together.

Figure 2 shows data rates versus signal strength for all successful transfers. Figure 3 shows success rates versus signal strength. Despite of some diversions, the general trend is higher data and success rates with higher signal strength. However, the

diversions for Wi-Fi data rates are larger. This is not surprising: the average data rate is around 1Mbps, suggesting the bottleneck is probably the Internet connection instead of the wireless medium. We were surprised to see average cellular upload speeds were faster than downloads, although uploads had much lower success rates. We re-ran the experiment for one week for both Cingular and T-Mobile three months later, only with similar results.

3.3 Real-Life Network Availability

In Experiment B, we collected the signal strength of cellular and Wi-Fi networks from all 14 participants. Unlike war-drives or spatial coverage measurements, such as in [8, 16], we measured *personal* coverage: the signal strength seen in a person’s daily life. Figure 4 shows a sample 48-hour best signal strength trace from one of the participants.

Figure 5 summarizes the distribution of best cellular signal strength for each participant. There is no coverage when signal strength is below -112dBm. Signal strength between -112 and -94 dBm shows as one or two bars on the phone. Signal strength between -94 and -81 dBm shows as three or four bars. Signal strength higher than -81 dBm is reported as 100% by Windows.

Our first observation is that cellular network availability is extremely high (99.1%). For more than 75% time on average for all our participants, signal strength is above -94 dBm (three or four bars). This is not surprising for a major urban area like

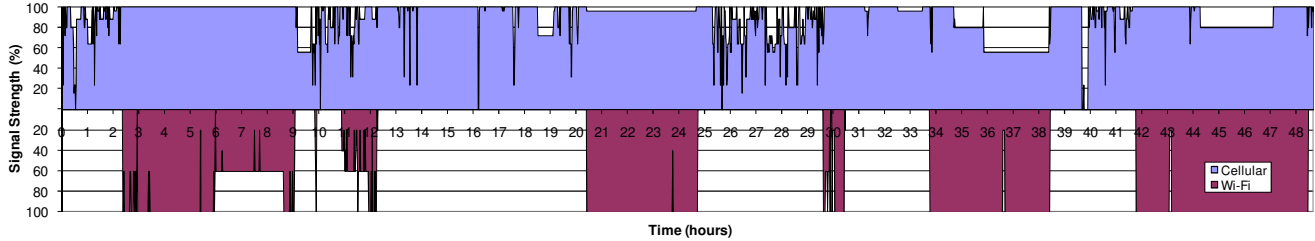


Figure 4. Best cellular and Wi-Fi signal strength in 48 hours observed by a participant

Houston. Nevertheless, while weak signal strength may be acceptable for voice communications, it is not true with data communications. Our measurement in Section 3.2 showed data transfer success rates are very low at the lower end of the signal strength range. This corroborates our real life experience that the cellular data connectivity is poor at lower signal strength.

Figure 6 shows the distributions of best Wi-Fi signal strength for each participant. Figure 6 (a) only includes participants' preferred networks, those that the user is known to have access to. We will use the preferred network traces throughout this paper. On average, our participants are covered by preferred Wi-Fi networks for 49% of their daily lives. Figure 6 (b) includes all visible unencrypted access points in addition to preferred networks. On average, our participants are covered by these networks for 77% of their daily life. This obviously overestimates network availability. Nonetheless, measurements by Nicholson *et al.* showed over 46% of all unencrypted access points in residential areas were in fact accessible [17].




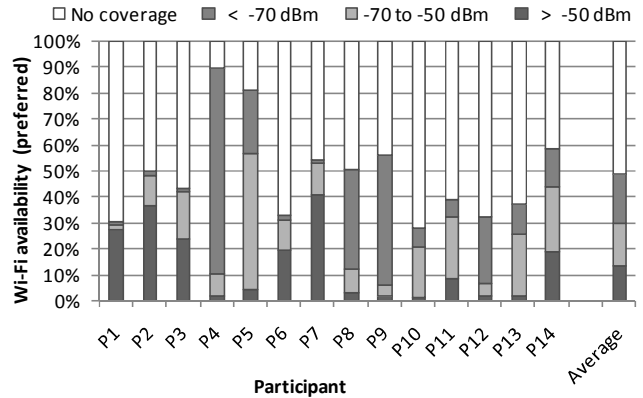
Figure 5. Distributions of best signal strength for cellular

3.4 Energy Cost & Model for Data Transfers


We next present the energy cost for various wireless activities that we obtained through direct measurements on three mobile phones, including the HTC Wizard. We also present a simple yet accurate model for the energy cost of a wireless transfer.

3.4.1 Measured Energy Costs

Table 2 provides the additional energy cost for various activities in cellular and Wi-Fi network interfaces on the HTC Wizard, 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.



(a) Preferred networks



(b) Preferred and unencrypted networks

Figure 6. Distributions of best signal strength for Wi-Fi

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_e denotes the energy cost of establishing a connection. Since we assume the phone is always connected to the cellular network, its E_e is zero. E_t denotes the additional energy cost for transferring 1MB data. We have ignored the effects of TCP and HTTP connection establishment, round trip latency (RTT), and TCP slow start, due to the much larger energy cost of data transfer. E_m denotes the energy cost for maintaining a connection for a minute, compared to that when the corresponding

Table 2. 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

network interface is powered off. For Wi-Fi, the values are shown with and without 802.11 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% reduction in E_m , compared to the default (balanced) setting.

Our measurements in Table 2 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 the cellular network interface on mobile phones is typically always on. Therefore, E_m should be regarded as zero for the cellular interface.

Since the Wi-Fi interface consumes high power even in PSM and Wi-Fi availability is low (so PSM cannot always be used), it is usually more energy-efficient to power off the Wi-Fi interface and then re-establish 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. Unfortunately, checking for Wi-Fi availability and establishing a connection consume considerable energy too. On the HTC Wizard, as Table 2 shows, $E_e=5J$, for both successful connection and a failed attempt with a five-second timeout. This large energy overhead makes Wi-Fi inefficient for small data transfers.

We also found that the built-in GPS of the HP iPAQ hw6925 Pocket PC phone consumes about 600mW (36 J per minute) additional power, which is too high to be used continually.

3.4.2 Energy Model of Wireless Data Transfers

We have built a simple energy model for wireless data transfers, assuming a constant network condition throughout a single transfer. The assumption is reasonable if the transfer time is short. We can also apply our model to a long transfer by splitting it into multiple short transfers.

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

$$E = E_e + n \cdot E_t \tag{1}$$

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

$$E = E_e + \frac{n}{S} \cdot E_t \tag{2}$$

where S is the transfer success rate.

3.5 Application Requirements

Mobile applications have diverse requirements in wireless transfer intervals and data sizes, as illustrated in Figure 7 with some representative applications.

Figure 8 highlights the challenge in battery lifetime imposed by mobile applications. We can see that cellular network, the only network with ubiquitous coverage, is inadequate for many applications because of its higher energy cost for data transfer. For most application requirements, the battery lifetime of the HTC Wizard will be reduced by more than 50% (The standby battery lifetime is eight days). Assuming Wi-Fi is always available with perfect signal strength, Figure 8 (b) offers the battery lifetime for transferring data through Wi-Fi with Wi-Fi interface powered off if not in use. The battery lifetime is more than double than that of the same data size and transfer interval with GSM/EDGE. This is essentially because Wi-Fi has much lower energy cost for data transfer. Unfortunately, Wi-Fi is far from being ubiquitous, as our reality check has shown. Even worse, with the Wi-Fi interface powered off to save energy, the mobile system does not know if Wi-Fi is available. Searching for Wi-Fi, again, costs considerable energy. Ultimately, neither cellular network nor Wi-Fi can achieve energy-efficient ubiquitous connectivity alone.

3.6 Summary of Reality Check

We summarize our findings as follows.

- We observed that for both cellular and Wi-Fi, the average data rates and transfer success rates are related to signal strength.
- While there were differences between individuals, on average, our participants spent 49% of their daily life under accessible Wi-Fi networks and more than 99% under cellular networks.
- We observed wireless energy consumption depends on network type, signal strength, connection time, and size of data transferred.
- We showed that cellular and Wi-Fi interfaces have complementary energy profiles. For cellular, there is no extra energy cost to maintain connectivity, but the energy cost per MB is more than an order of magnitude higher than that of Wi-Fi. On the other hand, the energy cost of maintaining a Wi-Fi connection is very high, even with maximal use of the power-saving mode. Moreover, establishing a Wi-Fi connection incurs considerable energy cost too.

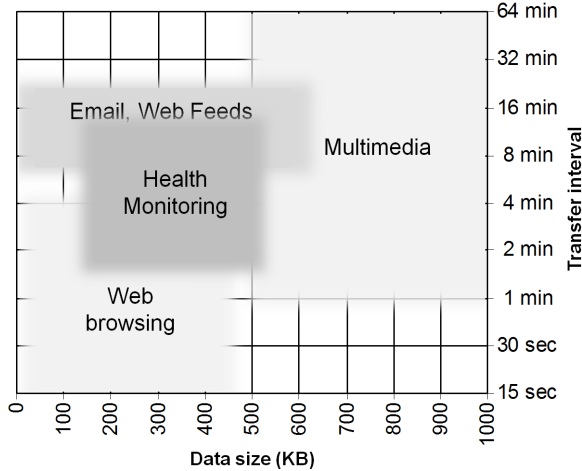
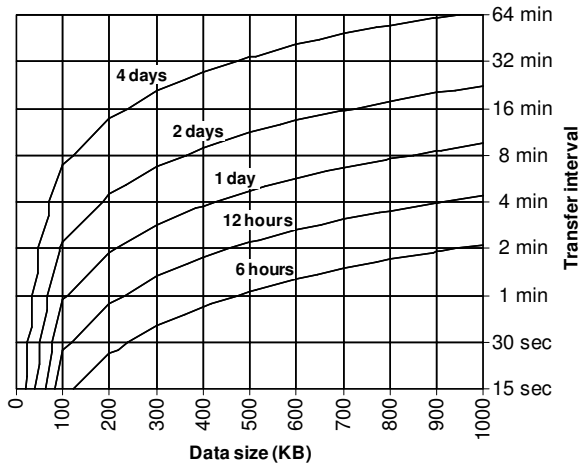
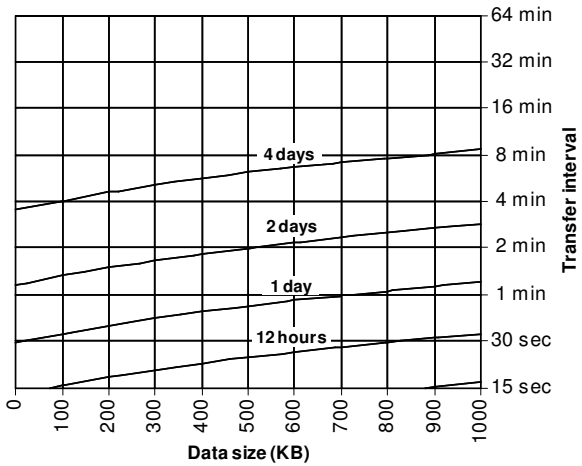


Figure 7. Data size and transfer intervals for different applications



(a) Battery lifetime with GSM/EDGE



(b) Battery lifetime with perfect Wi-Fi

Figure 8. Challenge in battery lifetime

Our reality check shows that neither cellular nor Wi-Fi alone can provide acceptable battery lifetime for future mobile applications requiring ubiquitous connectivity. This motivates our proposal to combine their complementary strength for achieving energy-efficient ubiquitous connectivity.

4. CONTEXT-SENSITIVE ENERGY-EFFICIENT WIRELESS DATA TRANSFER

We next present our problem formulation for data transfer through multiple wireless interfaces. Using the energy profiles from our measurement and our field-collected traces, we theoretically analyze the energy-saving potential of selecting between Wi-Fi and cellular. We then present our solutions to exploit the complementary strength of Wi-Fi and cellular, and compare their energy savings with the theoretical upper bound. We employ the Tower Logger traces and the energy model to evaluate our solutions through simulation. For each simulation and each participant, we use half of the trace for training, if necessary, and the rest for evaluation.

4.1 Data Transfer through Multiple Wireless Interfaces

Based on our reality check, 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 MB transfer. For mobile phones, the primary network is the cellular network (in our case, GSM/EDGE). We assume that alternative wireless networks may be available at limited locations or time. They offer higher data rates and consume lower energy per MB transfer. 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 many modern mobile handsets, one additional network is Wi-Fi.

The problem that we propose to solve is:

If the device needs to transfer n MB 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 \bar{C}_a , the expected energy cost for establishing a connection and transferring n MB data through it can be estimated with Equation (2) as

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

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, there is no energy cost for connection establishment. Therefore, the

energy cost to transfer the data through the primary interface is simple

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

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

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

Or

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

Our data transfer algorithm 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 will attempt 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 only 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. For them, we have measured S in Section 3.2, and measured Et and Ee in Section 3.3. The data size, n , can be obtained 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 Theoretical Analysis

To determine the maximum potential energy benefit of the use of multiple wireless interfaces, we first 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 network 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. Figure 9 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%. In Figure 10, we show the battery lifetime gains for hypothetical 20% and 80% Wi-Fi availabilities, assuming average Wi-Fi signal strength from the traces. Figure 10 clearly shows that the effectiveness of the use of multiple wireless interfaces, compared to a cellular-only policy, is improved with increased Wi-Fi coverage.



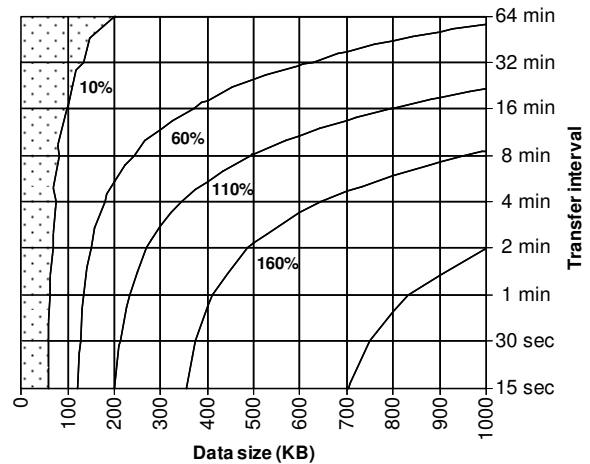


Figure 9. Potential battery lifetime gain with field-collected traces from all participants



(a) 20% Wi-Fi coverage



(b) 80% Wi-Fi coverage

Figure 10. Potential battery lifetime gain for different Wi-Fi coverage

4.3 Context-for-Wireless

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. 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.

4.3.1 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 (E_e) 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 the Simple Solution as the baseline, along with the theoretical upper bound from Section 4.1, for evaluating other algorithms in below. Figures 11 and 12 show the average data transfer energy and energy savings vs. data size using our measured traces and hypothetical 20% and 80% Wi-Fi coverage, respectively. It is clear that the Simple Solution provides substantial power saving for larger data sizes and higher Wi-Fi network coverage. However, the difference between the ideal case and the Simple Solution is substantial 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.

4.3.2 Hysteretic Estimation


People often stay at a location for a rather long time. This forms the basis for our Hysteretic Estimation. For Hysteretic Estimation, we use the previously measured network conditions for C_a and P_a until we either have a new measurement or a predetermined timeout runs out. Obviously, Hysteretic Estimation is more effective for shorter data transfer intervals, where network conditions are more likely to remain unchanged. The performance of this algorithm depends on the predetermined timeout value, and how often network conditions change. The timeout 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 timeout. It has the advantage of not requiring training.

4.3.3 History and Cell ID Estimation


People often spend days in predictable fashion, e.g., at work at home, commuting, etc. Therefore, the network conditions at the same time of different days are likely to be statistically related. This can potentially be extended to account for different days in a week. Network conditions are also highly correlated to geographical location. While a GPS can provide very accurate location information, its energy cost is too high for our purpose, as our measurement showed in Section 3.1. Furthermore, GPS systems only work outdoors. There has been considerable research on calculating location from visible cell tower IDs [15, 16, 18, 28]. However, they typically require extensive training against a known ground truth, such as a GPS. On the contrary, we



(a) Average among all participants



(b) Simulated 80% Wi-Fi availability



(c) Simulated 20% Wi-Fi availability

Figure 11. Data transfer energy vs. data size

use visible cell towers to directly train our algorithm and estimate Wi-Fi network conditions without positioning.

Our Cell ID Estimation works as follows. We store a list of all previously seen cell tower IDs. For each cell tower, i , we calculate and store three values; P_{cell_i} , the average Wi-Fi availability when that cell tower was visible, n_i , the number of samples used for calculating P_{cell_i} , and C_{cell_i} , the average Wi-Fi signal strength when Wi-Fi was available. In our simulations, for each cell tower, we calculate P_{cell_i} as the number of occasions Wi-Fi was available when that tower was visible divided by the number of occasions Wi-Fi was checked when that tower was visible. Potentially, more weight can be given to more recent measurements.

We then estimate Wi-Fi availability, P_{cell} , as the weighed mean of P_{cell_i} among all visible cell towers, denoted by V .

$$P_{cell} = \frac{\sum_{i \in V} w_i \cdot P_{cell_i}}{\sum_{i \in V} w_i}, \quad w_i = \log(n_i) \cdot (P_{cell_i} - 0.5)^4 \quad (8)$$

The weight, w_i , consists of two parts. $\log(n_i)$ gives more weight to the P_{cell_i} of towers that have been seen more often. In other words, the more samples we have of any tower, the more we trust its estimation. $(P_{cell_i} - 0.5)^4$ gives more weight to towers that have a P_{cell_i} close to 0 or 1. In other words, the more certain an 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.

For History Estimation, we divide days into 24 one-hour timeslots and compute the all time average Wi-Fi availability and signal strength, for each timeslot, and use them as the History estimated P_{hist} and S_{hist} when we are in the same timeslot. We then combine Cell ID and History Estimations using:

$$P_a = \frac{2 \cdot P_{cell} \cdot (P_{cell} - 0.5)^2 + P_{hist} \cdot (P_{hist} - 0.5)^2}{2 \cdot (P_{cell} - 0.5)^2 + (P_{hist} - 0.5)^2} \quad (9)$$

$$S_a = \frac{2 \cdot S_{cell} + S_{hist}}{3} \quad (10)$$


Again, we give more weight to estimations with higher certainty, and slightly favor Cell ID Estimation.

4.3.4 Acceleration Estimation


For P1, P2, and P3, we have recorded three-axis acceleration of their mobile phones. While there has been extensive research to extract user context information and physical activity from acceleration sensors [5], 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 potential for a change in Wi-Fi conditions, m , as:

$$m = \sum_{t=reset}^{current} [|\Delta A_x(t)| + |\Delta A_y(t)| + |\Delta A_z(t)| + c] \quad (11)$$


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 network conditions are measured. It then accumulates as time passes by. If m is below a predetermined threshold, our Acceleration Estimation assumes network conditions are unchanged and uses previous measured network conditions. 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



(a) Average among all participants



(b) Simulated 80% Wi-Fi availability



(c) Simulated 20% Wi-Fi availability

Figure 12. Data transfer energy saving vs. data size

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.

A few commercial phones today are already equipped with accelerometers, e.g., the Nokia 3220 and Sharp V603SH. 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.




Figure 13. Effectiveness of estimation algorithms for all participants

4.3.5 Combination Algorithms

To further improve the History and Cell ID Estimation, we combined it with Hysteretic Estimation and Acceleration Estimation as follows:

We use either Hysteretic Estimation 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 History and Cell ID Estimation to calculate network conditions.

4.3.6 Performance and Comparison

In this section, we present the average energy saving of our estimation algorithms. For each algorithm, we compare its energy saving to that of the Simple Solution, using that of the ideal case as 100%. For example, if the average energy cost for a data transfer is 23 J for an estimation algorithm, and 20 J and 30 J for the Simple Solution and the ideal case, respectively, the effectiveness of the estimation algorithm is $\frac{30 - 23}{30 - 20} = 70\%$.

We use the field-collected traces to calculate the energy consumption of each algorithm. We use data transfers of 270 KB every 5 minute, a data rate of a typical three channel ECG (each sampled at 200Hz with 12-bit resolution [13], without compression). Since we are normalizing with the ideal case, the result is relatively independent of data size (for larger data sizes).

Figure 13 shows the performance of our estimation algorithms, except Acceleration Estimation, for all participants in our reality check. The lowest effectiveness was for P4, who had a very high Wi-Fi coverage during out test period (90%). Our estimation algorithms had a hard time improving the already good Simple Solution for him. The highest effectiveness was for P6. We attribute the high effectiveness of our History and Cell ID Estimation to his regular work hours and locations. Our other participants, being students and faculty, typically have less regular work hours and locations. Another interesting observation was that in all cases where Hysteretic Estimation is more effective than History and Cell ID Estimation, the participants lived close to campus. We attribute this to the reduced accuracy of Cell ID Estimation for shorter displacements.

Participants P1, P2, and P3 had phones equipped with the acceleration sensor board. Figure 14 shows the performance of all estimation algorithms for them. We conjecture that Acceleration

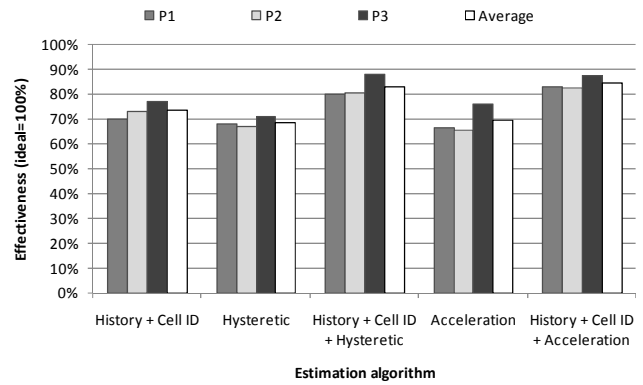


Figure 14. Effectiveness of all estimation algorithms for P1, P2, and P3

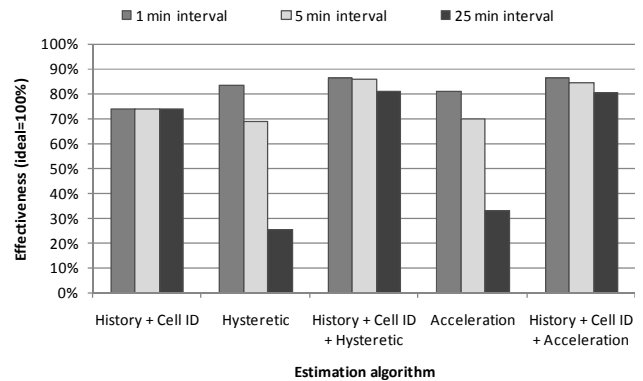


Figure 15. Effectiveness of estimation algorithms for different transfer intervals

Estimation would work well for users that regularly move between close locations with different Wi-Fi conditions, or for applications where cellular network information is not available. As previously mentioned, the data transfer interval impacts the efficiency of Hysteretic and Acceleration Estimation algorithms. To highlight the impact of transfer interval, we present their performance for 1, 5, and 25 minute intervals for P1, P2, and P3 in Figure 15. The data size for each transfer is 270 KB for all three intervals. As expected, their performance is reduced with

increased transfer intervals, when network conditions are less likely to remain unchanged.

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 strength. For example, Hysteretic and Acceleration Estimation work well for shorter transfer intervals and for participants with high Wi-Fi availability. History and Cell ID Estimation works very well for participants with regular hours and locations of Wi-Fi availability. We also observed that History and Cell ID Estimation performance is improved with increased mobility, which is opposite of Acceleration Estimation. 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 would expect our algorithms to generalize well [14], instead of only having good results for our own traces.

5. FIELD VALIDATION

To measure the effectiveness of our solution in real life conditions, we have implemented Context-for-Wireless interface selection 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 is already available on the phone. We use the same data rate and interval as in Section 4.3.6 (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-for-Wireless interface selection mode, it uses our Hysteretic algorithm, a constant 25 minute timeout, 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 4 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, who used the phones as their primary mobile phones during our experiments. We ran three experiments, each with two trials (cellular only and Context-for-Wireless), for each participant. 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-for-Wireless interface selection mode are shown in Table 3. The average battery lifetime gain was 35%.

Table 3. Phone battery lifetime during our field test

<i>Participant</i>	<i>Experiment</i>	<i>Cellular only (h)</i>	<i>Context-for-Wireless (h)</i>
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

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

6. DISCUSSION

Our reality check and studies of Context-for-Wireless data transfer are limited in 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 led to a much brighter picture for ubiquitous connectivity than what is available for the general population. Nonetheless, with the increasing availability of Wi-Fi and expanding deployment of 3G cellular networks, we believe our participants represent the trend in the development of ubiquitous wireless connectivity.

While we studied the use of Wi-Fi and cellular networks, 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 MB requirement compared to shorter range wireless technologies, such as the current and future wireless LAN technologies. On the other hand, although Wi-Fi hotspots are growing, their availability will still remain well below cellular networks, due to their short range. Therefore, we expect the availability and range vs. energy tradeoff 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 are separate. An integration of cellular and Wi-Fi networks will create additionally opportunities in estimating Wi-Fi network condition and availability. Due to E911 requirements, cellular network providers are already aware of approximate location of their subscriber phones. Therefore, they could potentially co-operate

with Context-for-Wireless and send prior known Wi-Fi availability data, based on location, to subscribers. This could be especially appealing to cellular providers that also offer Wi-Fi hotspots, such as T-Mobile. 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*. While we used context information to estimate current network conditions, it can also be employed to predict *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.

7. 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 [25] uses a low-power radio interface to transmit control information so that Wi-Fi can stay powered-off most of the time. Coolspots [20] employs Bluetooth to improve power efficiency of Wi-Fi. Nevertheless, these works and others [11, 21, 26] target at improving Wi-Fi energy efficiency and are restricted by limited Wi-Fi availability. On the contrary, 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 [3]. 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 30KB, 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 has also been widely studied [4, 7, 23, 24] through cooperation between networks on various layers. For example, Always Best Connected [12] seeks to achieve best performance and coverage. On the contrary, our solutions work at the application layer and do not require any cooperation between Wi-Fi and cellular networks.

Related to our use of Wi-Fi, Virgil [17] automatically discovers and selects access points with faster connection establishment, which will lead to energy savings too. 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.* [8] 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 [10, 16] 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. On the contrary, our reality check was targeted at the *personal* coverage of wireless networks, i.e., how networks are available throughout people's daily life.

Place Lab is also related to our endeavor on estimating Wi-Fi network condition. With positioning accuracy around 100m, Place Lab can indeed provide important clues for Wi-Fi network

condition, which is highly position-related. However, Place Lab requires a pre-mapping 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 [27], 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 [9] 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 3.4.1.

8. 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.

We showed, using findings from our recent field study, that while network availability is decent, the energy cost of ubiquitous network connectivity is overwhelming. We also showed that Wi-Fi and cellular network interfaces (802.11b and GSM/EDGE in our study) have energy profiles with complementary strength. 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 backfire.

We formulated data transfer through multiple wireless interfaces as a statistical decision problem and explored various contextual clues to estimate Wi-Fi network conditions in order to solve it. Our best algorithm without additional hardware, a combination of short and long term Wi-Fi condition measurements and cellular network information, can achieve battery lifetime improvement close to the theoretical limit. We also explored the use of a three-axis accelerometer (motion sensor) and different combinations of context information to estimate Wi-Fi network conditions. We validated our solutions using data collected from the real lives of a number of participants as well as using field trials.

As emerging mobile applications rely on ubiquitous connectivity, our work presented a challenging picture from field measurements. Our experiment showed that the data transfer required for mobile ECG reporting reduced the battery lifetime of our mobile phone to an average 15.4 hours. 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% to 20.8 hours. Yet more sophisticated algorithms, some described in this work, will produce even more opportunities for the solutions presented in this paper.

ACKNOWLEDGMENTS

This work was made possible in part through support from Enhancing Rice through Information Technology program, a seed grant funded by Sheafor/Lindsay from the Computer and

Information Technology Institute at Rice University and support from TI Leadership University Innovation Fund. We would like to thank the volunteers who participated in our data collection and field validation. We would like to thank our reviewers and our Shepherd, Professor Satyanarayanan, whose comments and shepherding have helped improve this paper greatly.

REFERENCES

- [1] Rice Orbit Platform, Rice Efficient Computing Group <http://www.ruf.rice.edu/~mobile/orbit>.
- [2] Universal Access Report, GSM Association, 2006.
- [3] Armstrong, T., Trescases, O., Amza, C. and Lara, E.d., Efficient and Transparent Dynamic Content Updates for Mobile Clients. in *Proc. 4th Int. Conf. Mobile Systems, Applications and Services (MobiSys)*, (Uppsala, Sweden, 2006), 56-68.
- [4] Bharghavan, V. Challenges and Solutions to Adaptive Computing and Seamless Mobility over Heterogeneous Wireless Networks. *Int. Journal on Wireless Personal Communications*, 4 (2). 217-256.
- [5] Bouten, C.V.C., Koekkoek, K.T.M., Verduin, M., Kodde, R. and Janssen, J.D. A Triaxial Accelerometer and Portable Data Processing Unit for the Assessment of Daily Physical Activity. *IEEE Trans. Biomedical Engineering*, 44 (3). 136-147.
- [6] Buchman, I. *Batteries in a Portable World: A Handbook on Rechargeable Batteries for Non-Engineers, Second Edition*. Cadex Electronics Inc, 2001.
- [7] Buddhikot, M.M., Chandranmenon, G., Han, S., Lee, Y.W. and Miller, S. Design and Implementation of a WLAN/CDMA2000 Interworking Architecture. *IEEE Communications*. 91.
- [8] Bychkovsky, V., Hull, B., Miu, A., Balakrishnan, H. and Madden, S., A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks. in *Proc. 12th Ann. Int. Conf. Mobile Computing and Networking (MobiCom)*, (Los Angeles, CA, 2006).
- [9] Chakraborty, S., Dong, Y., Yau, D.K.Y. and Lui, J.C.S. On the Effectiveness of Movement Prediction to Reduce Energy Consumption in Wireless Communication. *IEEE Trans. Mobile Computing*.
- [10] Chen, M.Y., Sohn, T., Chmelev, D., Haehnel, D., Hightower, J., Hughes, J., LaMarca, A., Potter, F., Smith, I. and Varshavsky, A., Practical Metropolitan-Scale Positioning for GSM Phones. in *Proc. 7th Int. Conf. Ubiquitous Computing (UbiComp)*, (2006).
- [11] Chiasserini, C.F., Rao, R.R. and di Elettronica, D. Improving Energy Saving in Wireless Systems by Using Dynamic Power Management. *IEEE Trans. Wireless Communications*, 2 (5). 1090-1100.
- [12] Gustafsson, E. and Jonsson, A. Always Best Connected. *IEEE Wireless Communications*, 10 (1). 49-55.
- [13] Hamilton, P.S. and Tompkins, W.J., Estimation of Rate-Distortion Bounds for Compression of Ambulatory ECGs. in *IEEE Ann. Int. Conf. Engineering in Medicine and Biology Society*, (1989).
- [14] Hastie, T., Tibshirani, R. and Friedman, J.H. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2001.
- [15] Hightower, J., Consolvo, S., LaMarca, A., Smith, I. and Hughes, J., Learning and Recognizing the Places We Go. in *Proc. 7th Int. Conf. Ubiquitous Computing (UbiComp)*, (2005), 159-176.
- [16] LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J. and Potter, F., Place Lab: Device Positioning Using Radio Beacons in the Wild. in *Proc. Pervasive*, (2005).
- [17] Nicholson, A.J., Chawathe, Y., Chen, M.Y., Noble, B.D. and Wetherall, D., Improved Access Point Selection. in *Proc. 4th Int. Conf. Mobile Systems, Applications and Services (MobiSys)*, (2006), 233-245.
- [18] Otsason, V., Varshavsky, A., LaMarca, A. and de Lara, E., Accurate GSM Indoor Localization. in *Proc. 7th Int. Conf. Ubiquitous Computing (UbiComp)*, (2005).
- [19] Park, S. and Jayaraman, S. Enhancing the quality of life through wearable technology. *IEEE Engineering in Medicine and Biology Mag.*, 22 (3). 41-48.
- [20] Pering, T., Agarwal, Y., Gupta, R. and Want, R., CoolSpots: Reducing the Power Consumption of Wireless Mobile Devices with Multiple Radio Interfaces. in *Proc. 4th Int. Conf. Mobile Systems, Applications and Services (MobiSys)*, (2006), 220-232.
- [21] Qadeer, W., Simunic, T., Ankcorn, J., Krishnan, V. and De Micheli, G., Heterogeneous Wireless Network Management. in *Proc. Wksp. Power Aware Computer Systems (PACS)*, (2003).
- [22] Ross, P.E. Managing Care Through the Air. *IEEE Spectrum*, 41 (12). 26-31.
- [23] Salkintzis, A.K., Fors, C. and Pazhyannur, R. WLAN-GPRS Integration for Next-Generation Mobile Data Networks. *IEEE Wireless Communications*, 9 (5). 112-124.
- [24] Sanmateu, A., Paint, F., Morand, L., Tessier, S., Fouquart, P., Sollund, A. and Bustos, E. Seamless Mobility Across IP Networks using Mobile IP. *Computer Networks*, 40 (1). 181-190.
- [25] Shih, E., Bahl, P. and Sinclair, M.J., Wake on Wireless: An Event Driven Energy Saving Strategy for Battery Operated Devices. in *Proc. 8th Int. Conf. Mobile Computing and Networking (MOBICOM)*, (2002), 160-171.
- [26] Simunic, T., Qadeer, W. and De Micheli, G., Managing Heterogeneous Wireless Environments via Hotspot Servers. in *Proc. 13th Conf. Multimedia Computing and Networking (MMCN)*, (2005), 143-154.
- [27] Sorber, J., Banerjee, N., Corner, M.D. and Rollins, S., Turducken: Hierarchical Power Management for Mobile Devices. in *Proc. 3rd Int. Conf. Mobile Systems, Applications, and Services (MobiSys)*, (2005), 261-274.
- [28] Varshavsky, A., Chen, M., de Lara, E., Froehlich, J., Haehnel, D., Hightower, J., LaMarca, A., Potter, F., Sohn, T. and Tang, K., Are GSM phones THE solution for localization? in *Proc. 7th IEEE Wksp. Mobile Computing Systems and Applications (WMCSA)*, (2006).