# Farther Than You May Think: An Empirical Investigation of the Proximity of Users to Their Mobile Phones

Shwetak N. Patel, Julie A. Kientz, Gillian R. Hayes, Sooraj Bhat, and Gregory D. Abowd

College of Computing & GVU Center
Georgia Institute of Technology
801 Atlantic Drive, Atlanta GA 30332-0280 USA
{shwetak, julie, gillian, sooraj, abowd}@cc.gatech.edu

**Abstract.** Implicit in much research and application development for mobile phones is the assumption that the mobile phone is a suitable proxy for its owner's location. We report an in-depth empirical investigation of this assumption in which we measured proximity of the phone to its owner over several weeks of continual observation. Our findings, summarizing results over 16 different subjects of a variety of ages and occupations, establish baseline statistics for the proximity relationship in a typical US metropolitan market. Supplemental interviews help us to establish reasons why the phone and owner are separated, leading to guidelines for developing mobile phone applications that can be smart with respect to the proximity assumption. We show it is possible to predict the proximity relationship with 86% confidence using simple parameters of the phone, such as current cell ID, current date and time, signal status, charger status and ring/vibrate mode.

## 1 Introduction and Motivation

Mobile computing systems have been one of the fastest evolving and growing technologies of the last decade. Simple mobile phones and pagers have given way to simultaneously complex and small computing artifacts that provide a myriad of services and can often even interact with each other. The increasing power and ubiquity of these mobile technologies make it possible to realize many of the early visions of ubiquitous computing. Many argue that the mobile phone, with its expanded capabilities, can be the platform of choice for applications that once required customized mobile hardware [23]. Examples of research focused on these expanded uses include memory aids [7], augmented cognition [8, 25], location-based services [12], medical data collection [27], authentication mechanisms [3, 20], and personal information stores [28].

The topic of location has been a common discussion point in ubiquitous computing with researchers making the mobile phone the platform of choice for location-aware computing. The PlaceLab effort at Intel Research and other location systems (see

P. Dourish and A. Friday (Eds.): Ubicomp 2006, LNCS 4206, pp. 123 – 140, 2006.

Hightower & Borriello [9] for a survey) have demonstrated that ubiquitous location-awareness can be delivered on commodity hardware, most interestingly mobile phones [12, 16]. This advance creates many opportunities for developing knowledge on the mobile phone of where a person has been and what they have been doing. However, this approach assumes that the mobile phone is an accurate proxy for the location of its owner. Intuitively and anecdotally, we know that people do in fact carry their mobile phones with them *much* of the time, but these same phones are not physically on their bodies nor within arm's reach at *all* times.

Many researchers and application designers make the implicit assumption that people are likely to have their mobile phones with them and available most of the time. However, little empirical evidence on the actual proximity relationship between a mobile phone and its owner exists. The results presented in this paper provide in-depth empirical results uncovering the habits of a small set of representative users in a major metropolitan US city (Atlanta, Georgia). This work not only tests the hypothesis that a user's phone is available to her most of the time but also provides an exploration of the situations in which the proximity assumption is broken and attempts to select the factors that best predict the proximity relationship. Through this evidence, we create concrete design advice for mobile phone applications that require knowledge about the proximity of the user to her phone. We also define for the first time a way to predict proximity using standard features of the phone, such as current cell tower, time, charger status, or ring profile. With an accurate predictor, an important piece of context, namely, how near the owner is, can inform context-aware behaviors on the mobile phone.

This paper provides four contributions. First, we present the design and creation of a proximity-sensing technique and the design of an empirical proximity study that can be replicated by others. Second, we present empirical evidence directly testing the strength of the assumption that the mobile phone is a good proxy for its owner's location. Third, we present a classification of situations that break the proximity assumption, information that can be interpreted as design advice for mobile phone applications. Fourth, we present a decision tree method for predicting proximity to mobile phones based on readily accessible features on the phone itself.

## 2 Related Work

The mobile phone, initially a device simply for strategic communications, has gone through a long evolutionary process. Originally designed primarily for durability, they were not particularly usable and certainly not stylish. In the 1980's, this trend shifted, and they moved into the consumer product space, complete with the power and status of a high-end watch or automobile. Now, they come in a wide variety of form factors with numerous possible combinations of services. During this evolution, people have been studying mobile phone usage patterns. Marketing firms and mobile phone manufacturers study a variety of user needs, from the calling plans that are most appealing to certain demographics to the usability of the handset itself.

Much of this research has focused on the design of new handsets and/or new services. For example, in 1998, Vaananen-Vainio-Mattila and Ruuska presented an ethnographic study of mobile phone users conducted at Nokia [26]. In this study, the authors used contextual inquiry to uncover both the sociological and cultural considerations affecting mobile phone usage and the design challenges and some potential basic solutions for the handset itself. Palen *et al.* took a slightly different approach, focusing on the use of the mobile phone *system*, including everything from the sales people to the phone itself to the service contacts [19]. Schlosser investigated the ways in which mobile phones are appropriated into organizations and daily activities [24]. She used interviews to uncover both these details and, in turn, how those individuals, organizations, and activities change based on this use.

This related research shows the power and the limitations of these types of studies with real mobile phone users. Palen and Salzman [18] note that although direct naturalistic observation can help investigators to understand interactions as they "really happen… tracking particular participants requires getting access to the many places participants spend their time while also involving a large time commitment for all parties." Thus, they chose to supplement interviews not with observation but with voice mail diary entries, in which mobile phone users called a voice mailbox daily to report their interactions and troubles. McGuigan explores how social science methods can and should be used to study mobile phone usage, describing in depth the strengths and weaknesses of four different sociological methods: social demography; political economy; conversation, discourse, and text analysis; and ethnography [14].

Automatic logging, in which software automatically records the user's actions for later analysis, provides many of the benefits of observation methods without some of the problems. Researchers can gather data across all times, locations and activities without being excessively intrusive to the participant. For example, Demumieux and Losquin developed a tool that collects logs of applications used on mobile devices, including both mobile phones and PDAs [5]. MIT's Reality Mining Project has used Bluetooth proximity and phone context to predict things such as daily routine, but has not yet explored proximity to mobile devices [6]. The ContextPhone system [22] also logs a variety of information on phone use both for research purposes and for context-aware applications. Similarly, the Mobile Media Metadata system leverages such context information to assist users in annotating images on their camera-phones (digital camera equipped mobile phones) [4], whereas other systems use cellular identification to predict user routes and location [11] and develop context-aware contact lists [17]. We are interested in inferring user proximity from contextual information already available on the mobile phone with minimal additional sensors.

## 3 Experimental Design

This paper presents a mixed-method approach, in which we use both interviews and automatic logging to develop a full picture of user practices. We collected data about the phone and about its owner to help understand and potentially predict the proximity relationship between owner and phone. These data, initially collected automatically, are verified using self-report. In this section, we describe the study design and participants and elaborate on the technology for automatic collection of data.

### 3.1 Study Method and Participants

A primary goal of this work includes gathering information about users and their mobile phones *all* day *every* day for some extended period introducing minimal burden and without relying on self-report. Due to the increased capabilities of mobile phones, we were able to gather much of the data using software custom-developed for this purpose and running on the phone itself. Recording the user's physical relationship to the phone, however, requires a reliable proxy for the user. We used small, plastic beacons that users wore on lanyards around their necks nearly all the time (shown in Figure 1). We were thus able to measure the phone's distance from the tag and assume this roughly equated to the phone's distance from the individual.

The tags measure 40 mm by 25 mm by 5 mm, approximately the size of an automobile keyless entry remote control. Each weighs approximately 20 grams, including the battery, and can be worn around the neck as a pendant. The tag is also weatherproof, shockproof, and hypoallergenic. The tag emits a Bluetooth signal detected by a custom-built application on the user's Bluetooth-enabled phone. The application on the phone pings the tag every 60 seconds and approximates the distance based on the strength of the signal received. This method allows for the determination of three levels of proximity: within arm's reach (strong signal); in the same room (signal is weak or varied); or unavailable (signal could not be detected). A separate application on the user's phone records contextual information, including signal strength, battery level, charge status, current running application, cell tower ID, area ID, ring volume, ring type, and vibration status. A third application inherent to most mobile phones logs incoming, outgoing, and missed calls and data usage.

Sixteen individuals participated for at least three weeks each. All of the subjects lived in the greater metropolitan area of Atlanta, Georgia, USA and were recruited via word of mouth and Internet classified advertisements. We compensated participants with $200 for completing the entire three-week study and returning the equipment. Participants ranged in age from 21 to 66 and included 9 males and 7 females. Self-reported phone plans ranged from a 5000 minute per month contract to a prepaid, "emergencies only" service plan. Participants also had a wide variety of professions and income levels (see Table 1 for details on each participant). Each participant completed a background interview to provide basic demographic information and data on perceptions of individual phone usage patterns. These questions included those about current phone-charging patterns, applications used on the phone, service plan information, and the perceived phone proximity throughout the day.

After the initial interview, we replaced the participant's phone with one of several form factors all capable of running the logging software, accomplished by a simple swap of SIM cards. We copied all contact list information to the new phone by using the SIM card's memory or, in rare situations, manually entering the information. We provided phones in a similar form factor and with similar software and menu structures to the phones already in use by the participants. Thus, we believe that the phones had minimal impact on the practices of the participants. Participants received a beacon pendant and instructions about charging the phone and the tag. We instructed them to use their phones as normal and to wear the tag at all times. Notable exceptions included while showering or swimming. Most subjects wore the tag while sleeping, but others preferred to place it next them while sleeping. If they removed the tag, we asked them to note the time and duration and keep the tag as near as possible.

During the three weeks of participation, the individuals met with us once per week when we downloaded the logging data from the phone and interviewed them about their usage patterns for the week. At the beginning of each interview, the participant completed a detailed diary of the previous 24-hour period, as suggested by the Day Reconstruction Method [10], breaking the day into episodes described by activities, locations, and the phone's location during these times. During these interviews, participants self-reported reasons they did or did not have their phones for various episodes reported on the diary. Together with the participant, we then compared this diary to the data recorded by the logging application (showing them visualizations of phone proximity similar to that in Figure 1) and asked clarifying questions for any inconsistencies. The diary and visualization could disagree for three reasons:

(1) The participant may have an error in recollection.
(2) The logging application and/or the hardware itself could produce an error.
(3) The tag was not an appropriate proxy (*e.g.* the user was not wearing it).

The interview closed with general questions about the remaining days from the preceding week, such as whether it was a normal workday or a day off, but did not include any specific details about the days that were further in the past. On the third and final interview for each participant, equipment was returned and replaced with the participant's original phone.

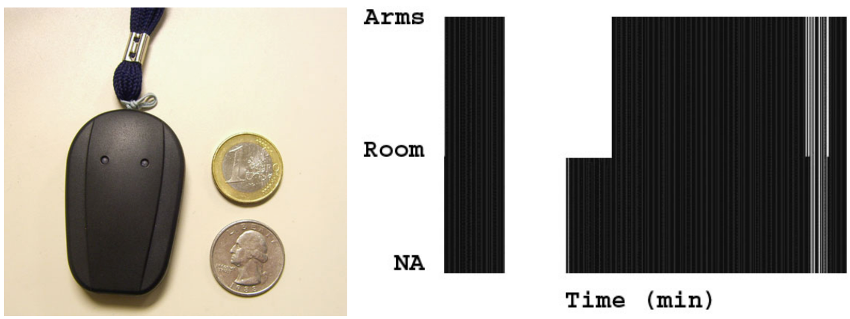


**Fig. 1.** The left image is a picture the Bluetooth tag. The right shows sample minute by minute proximity data from laboratory tests of the Bluetooth tag and mobile phone application. The full solid lines indicate the tag is within arms reach, the white indicates that it is not available and halfway between or oscillating indicates the tag is room level.

### 3.2 Phone Proximity Hardware Implementation Details

Although many other contextual indicators of phone usage could be gathered automatically using simply the mobile phone, detection of the user's proximity was more complicated and required development of a proxy for the user. The chosen embodiment of a proxy for the user was a Bluetooth tag from Bluelon Inc. using modified radio parameters, partially inspired by the SPECs project at HP Labs [13]. An application running on the mobile phone continuously records signal values from associated tags. The tag includes a low-power CSR BlueCore-02 Class 2 Bluetooth

RF module with an integrated antenna and a 3.7 V 345 mAh lithium ion battery (see Figure 1). The tag can signal every minute for approximately five days with a single two-hour charge. A buzzer and LED on the tag indicate when the battery is low.

The tag uses a Class 2 Bluetooth module with a 10 meter range, which is sufficient for registering the levels of proximity of interest to the study and uses much less power than the longer range Class 1 modules. The Bluetooth stack implements the Serial Port Profile (SPP) running over L2CAP and RFCOMM for firmware programming. The Bluetooth radio in the user's beacon was reduced to -22 dB to extend battery life and limit the maximum range at which the mobile phone can detect the tag to around 5 to 6 meters. The design of the radio output and subsequent distance analysis assumes a tag placed around the neck of an average adult. Thus, we assume a 5 dB signal loss from the human body due to absorption, with one tag reconfigured to account for a participant of larger size.

Rather than use a Received Signal Strength Indicator (RSSI), which is implemented inconsistently across mobile phones if at all, we implemented our own simpler signal strength indicator for proximity detection. In this solution, the round trip time of the Service Discovery Protocol (SDP) packets are used to estimate the distance between the tag and the mobile phone. As the distance increases between the mobile phone and the tag, the link quality should degrade. The lower link quality then increases the bit error rate and thus the number of packet retransmissions. The retransmits in turn increase the service discovery time. Despite the simplicity of this approach, it was more than sufficient for the level of granularity desired for this study.

By reducing the radio output of the tag, we can specify a rough range at which the bit error rates increase by a set amount. After some experimentation in lab settings with humans of average size, we determined that the appropriate range for human body signal absorption is around ten feet. A phone within arm's reach typically shows a service discovery time of about 2000-4000 ms, room level distance of about 4000-7000, and no returned service discovery information is interpreted as the phone being out of range or further than room level (5-6 meters). In practice, physical room level distance can result in fluctuating values between 4000 ms and no discovery. This fluctuation is likely due to a bit error rate that is so high the Bluetooth module times out and does not report a successful service discovery. One serious issue with this phenomenon is the difficulty that results in determining whether the phone is transitioning from "room level" and truly out of range or whether the phone is consistently at room level with the erroneous fluctuation described. Thus, if we observed high rates of fluctuation (*e.g.* alternating with every reading) over extended periods of time (more than five minutes), we classified the reading as room level.

We developed two applications for gathering the empirical evidence of users' relationships to their mobile phones. The Bluetooth distance logging application was written in Java 2, Micro Edition (J2ME) using the MIDP 2.0 and JSR-82 Bluetooth specification. This application, designed for the Nokia Series 60 platform, also works with other mobile phones that feature the Symbian 7.0 or higher operating system. The application used to log other contextual information, such as cell tower ID and battery status, was written in Symbian using C. Both applications log pertinent information to the removable memory card once per minute and run constantly in the background with minimal interference on other phone functionality. A watchdog application automatically restarts both logging applications in the event of a restart.

## 4 Results

All participants successfully completed the study for at least three weeks. In every case, at least one of these weeks represented what they considered "typical" patterns, and in many cases, all three weeks were "typical." Participants reported the tag was comfortable to wear and did not interfere with their day-to-day lives. In some instances, participants reported forgetting to put on the tag first thing in the morning after leaving it off for sleeping. We adjusted to account for these errors. In this section, we present the results from both the automatically-collected proximity and phone context data and the self-reported results from the interviews.

### 4.1 Proximity Levels

As described in Section 3.2, three levels of proximity between the user and the phone can be determined using the Bluetooth tag and application scheme. These are:

- Within arm's reach (within 1-2 meters of the tag)
- Within the same room (within 5-6 meters of the tag)
- Unavailable (beyond 6 meters from the tag)

From the minute-by-minute readings taken each day during the three-week study, we obtained between 6190 and 35791 proximity measurements, with an average of 1175 readings per day per person. When a phone is turned off, no proximity ratings can be logged, but the very nature of the phone being off indicates it is unavailable. Given the large quantities of data, we were able to analyze different scenarios which may or may not affect proximity. In this section, we report those scenarios that showed the most significant trends:

- In and out of the home (determined by cell ID)
- Waking vs. sleeping hours (determined by hours reported during interviews)
- Weekend vs. weekday (weekend being 12 AM Saturday to 12 AM Monday)

Overall, participants varied in their proximity levels, ranging from 17% of the time within arm's reach to 85% of the time, with an average of 58% of the time within arm's reach (see Figure 2). All but two users kept the phones on more than 85% of the time. Participant 7 had a prepaid plan and only had her phone on 21% of the time to conserve minutes, and Participant 14 turned his phone off almost every night while sleeping reportedly to avoid being disturbed.

Interestingly, participants showed a significant increase in the average percentage of time the phone was within arm's reach during times they were away from home ($p < 0.0001$; see Figure 2). Users were more likely to keep the phone at room level or even further while at home. In fact, two participants with the lowest overall proximity data (2 and 13) were "stay at home" mothers who spent a significant time at home.

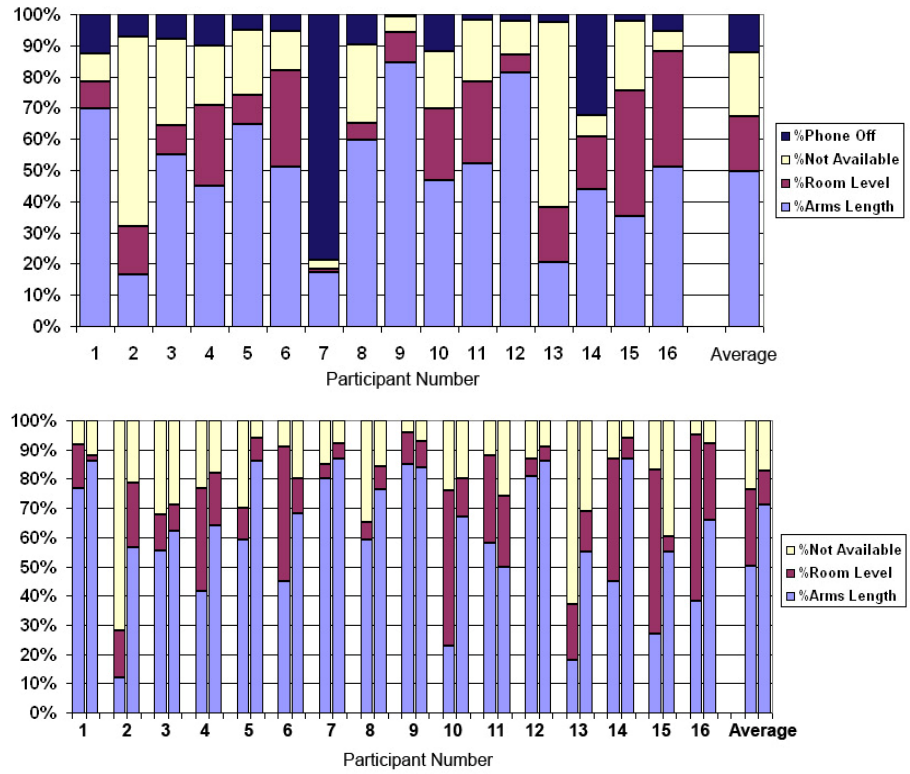


**Fig. 2.** Individuals varied in proximity levels, but on average people kept their phone within arm's reach half the time (Top). Most users carried the phone close to them at all times when away from home if the phones were turned on (Bottom: Left bar is at home, Right is away).

We compared proximity trends for individual users for times they were sleeping versus times they were awake. Participants had their phones within arm's reach more often while awake than they did while they were asleep (61% while awake, 52% while sleeping). Also, participants tended to keep their phones within arm's reach slightly more often on weekdays (59%) as opposed to weekends (53%). Although these values showed interesting trends, they did not demonstrate statistical significance.

The phone was within arm's reach and turned on, thus highly available, 50% of the time ($\sigma = 20.4$). Thus, we categorized users further than one standard deviation below the average (29.6%) as "Low" availability; users within one standard deviation as "Medium"; and users above one standard deviation (70.4%) as "High." Table 1 lists the proximity category. Contrary to our initial hypothesis, the number of minutes used per month on the phone does not correlate to the proximity relationship throughout the day. For example, the person with the highest number of minutes used per month, Participant 3, fits the Medium availability category. On the other hand, Participant 12 was highly available to the phone, but used it infrequently to make or to receive calls.

**Table 1.** Demographic information, basic data logged during study, and proximity levels

| Participant | Gender | Age | Profession | Minutes Per Month | # Cell towers logged | % Phone off | %Arm's Reach | Proximity Category |
|---|---|---|---|---|---|---|---|---|
| 1 | M | 24 | Graduate Student | 645 | 168 | 13 | 70 | High |
| 2 | F | 36 | Homemaker | 253 | 130 | 7 | 17 | Low |
| 3 | M | 46 | Sales Rep. | 4402 | 696 | 8 | 54 | Med. |
| 4 | F | 50 | Graduate Student | 344 | 253 | 10 | 45 | Med. |
| 5 | M | 41 | Software Sales | 1068 | 258 | 5 | 65 | Med. |
| 6 | M | 40 | Mail Carrier | 2905 | 269 | 6 | 51 | Med. |
| 7 | F | 47 | Dry Cleaner | 25 | 52 | 79 | 17 | Low |
| 8 | F | 23 | Admin. Asst. | 468 | 204 | 10 | 60 | Med. |
| 9 | M | 25 | Consultant | 559 | 139 | 1 | 84 | High |
| 10 | M | 61 | Lecturer | 384 | 414 | 12 | 47 | Med. |
| 11 | F | 21 | Childcare Provider | 1394 | 227 | 2 | 52 | Med. |
| 12 | M | 33 | Project Manager | 189 | 198 | 2 | 81 | High |
| 13 | F | 35 | Homemaker | 148 | 133 | 2 | 20 | Low |
| 14 | M | 32 | Sales/Marketing | 1769 | 900 | 32 | 44 | Med. |
| 15 | M | 66 | Retired | 984 | 227 | 2 | 35 | Med. |
| 16 | F | 24 | Financial Associate | 2075 | 254 | 5 | 51 | Med. |

### 4.2 Factors Affecting the Proximity of Mobile Phones to Users

The weekly participant interviews served two purposes. First, during these interviews, participants described their own recollections of proximity to the phone. At times, their recollections and the automatically gathered data, visualized as in Figure 1, appeared to disagree. The interviews provided a time to discuss these discrepancies. Occasionally, participants recalled more accurate information about daily activities after being prompted by the automatically collected data. The interviews also provided an opportunity to note times that the tag was not an adequate proxy, typically because the participant could not wear the tag for some reason. For cases in which no resolution could be reached through the interview discussions, the discrepancy was attributed to technical error. These error cases always registered for only short time periods (a single data point) and occurred on average 3 times in a given day for a participant out of over 1200 data points each day.

Second, the interviews provided participants an opportunity to reflect on and explain activities grounded in the data, both automatically-collected and self-reported. These discussions resulted in better understanding of the factors that contributed to an individual's phone being near or not than would have been possible using solely the automatically-logged contextual information. These factors were determined in two ways: by examining the variables that most directly affected the learned model of the user, as described in Section 5, and by using affinity clustering [21] to group the self-reported reasons for the phone's proximity from the interview data.

Specifically, the affinity clustering results were produced by three researchers. The first recorded each participant's stated reasons for the phone's proximity and availability during the interviews. Two researchers then independently categorized these statements using affinity clustering to determine themes, producing 15 unique themes, 13 of which were shared. The two coders agreed on categorization of most of the cases (105 of 120), and after discussion, agreement was made to include all 15 unique themes. One of the original coders then re-categorized all of the statements while the third researcher categorized them independently using the 15 themes (see Table 2 for inter-coder reliability). The 15 emergent themes follow:

1. *Routine*: The phone's proximity is related to anything that is part of a common routine for the individual, particularly those things that might help them to remember the phone's location or be within range of its use. Example: User always leaves phone on kitchen counter while at home.
2. *Environment*: The phone's proximity is related directly to the distance at which the user believes the phone should be due to the physical constraints of the space. Example: In a car, the phone is rarely out of arm's reach.
3. *Physicality of person/Activity*: The phone's proximity to the user is related directly to something physical about the person or the activity in which he/she is engaged. Example: Phone is awkward to carry while working out.
4. *Disruption to others*: User makes a choice about the phone's proximity and/or on/off status based on how that choice affects other people or the environment. Example: User turns off phone car during a client meeting.
5. *Disruption to self*: User makes a choice about the phone's proximity and/or on/off status based on that choice's effects on self. Example: User turns off phone at home after a long day of calls at work.
6. *Regulations*: Legal or other specific regulations prevent use, carrying, and/or powering of phone. Example: User has to turn off phone while in a hospital.
7. *Use of phone by self*: The phone's proximity is affected by the owner using it or anticipating use. Example: Phone is nearby while user is on a phone call.
8. *Need for use of phone by others*: The phone's proximity is affected by expectation that others may need to reach owner or otherwise make use of owner's phone. Example: Phone is nearby when the user is expecting a call.
9. *Need for use of phone both by self and by others*: The expectation of needing features for self as well as availability of self to others through the phone's features. Example: The user keeps the phone close while trying to coordinate a group of people at a social event.
10. *Use of handset by others*: Someone else is physically using the handset. Example: User loaned phone to spouse while she was out running errands.
11. *No need for use of phone*: Phone's availability directly affected by the belief that no use is imminent. Example: While at home, others can use a landline to reach the user.
12. *Technical resource issues*: The phone's availability and proximity are directly affected by technical considerations inherent to the phone or the network. Example: User moves close to a window to obtain a better signal or moves it to where the charger is located when the battery is low.

13. *Quick trips*: The timing (or expected timing) of an activity affects the user's choice about whether to explicitly consider/act on phone's proximity or not. Example: Phone is on the desk at work while taking a coffee break.
14. *Memory and forgetfulness*: Phone lost (at least temporarily) or unintentionally left behind due primarily to user's forgetfulness or memory error. Example: Phone is left behind while leaving the house.
15. *Protection of phone from others*: The user's choice about phone placement is directly related to protecting the physical handset or the resources that can be accessed through the phone from tampering or use by other people. Example: Phone on a high shelf out of the reach of children.

**Table 2.** Inter-coder reliability for each thematic cluster was determined using two measures: (1) Observed Agreement represents a measure of simple agreement between two coders for each theme and is measured by agreements divided by total number of statements coded; (2) Cohen's Kappa measures how much better than chance the agreement between the two coders is [1]. Both measure between 0 and 1 with 1 indicating perfect agreement between coders.

| **Cluster** | **1** | **2** | **3** | **4** | **5** | **6** | **7** | **8** |
|---|---|---|---|---|---|---|---|---|
| **Observed Agreement** | .93 | 1 | .99 | .99 | .98 | 1 | .98 | 1 |
| **Cohen's Kappa ($\kappa$)** | .59 | 1 | .83 | .83 | .81 | 1 | .85 | 1 |

| **Cluster** | **9** | **10** | **11** | **12** | **13** | **14** | **15** |
|---|---|---|---|---|---|---|---|
| **Observed Agreement** | .95 | .99 | .93 | .99 | .98 | .99 | .96 |
| **Cohen's Kappa ($\kappa$)** | .58 | .91 | .74 | .75 | .89 | .95 | .67 |

## 5 Predicting User's Proximity

The preceding section presents empirical findings on the proximity relationship as well as reasons for why a phone is not near its owner. An open question was whether this proximity relationship could be determined on the phone itself, without the aid of a proxy tag. We collected approximately 30,000 proximity readings per participant. In addition to proximity data, we also collected a variety of contextual data from the mobile phone. This information includes hour of day, day of week, cellular tower ID, cellular area ID, signal strength, battery level, charging status, ring volume, ring type or mode, vibration status, foreground application, idle status of the phone, missed calls, time and duration of incoming and outgoing calls, SMS messages, and GPRS data usage. We wanted to investigate whether the real proximity value could be deduced from some subset of those available data already on the phone and whether such a correlation was independent of the individual. If so, mobile phone application developers could create context-aware behaviors triggered by proximity.

The descriptive statistics and interviews summarized throughout this paper suggest that some features do have predictive power. For example, for participants with very structured work schedules, day and hour were effective features for predicting proximity to the phone. Cell tower IDs and charging status were two other contextual features that also showed promise. Many of our participants tended to have their phones on their bodies every time they were away from the cell towers near their homes. Some people only charged their phones while in the car, which made using the charger status (*i.e.,* whether or not connected to charger) one way of inferring that those users would be arms-length from their phones at those times. Figure 3 shows evidence of patterns existing between features and the user's proximity to their phone.

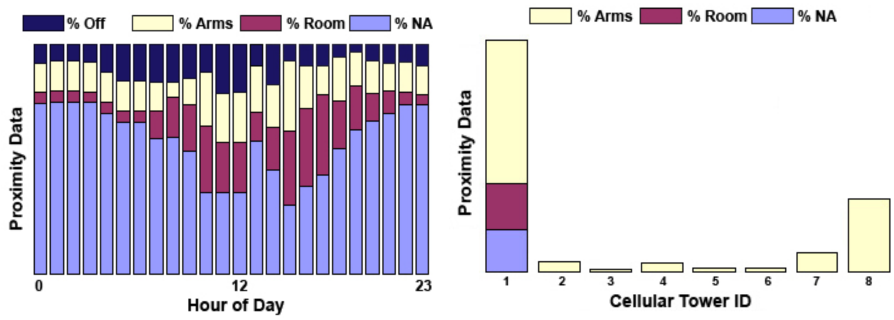


**Fig. 3.** Left: Proximity percentages for each hour of the day for Participant 2 (a homemaker). Right: Proximity percentage for each cellular tower ID, again for Participant 2. Cell ID #1 is the participant's home and is the only one that has variability on proximity level.

### 5.1 Proximity Classifier

From a machine learning perspective, the real question is whether the features that predict proximity were general across individuals. We created a model that could classify and predict the proximity of an owner to the mobile phone based on the logged contextual features. We employed a decision tree classifier using the ID3 algorithm [15]. Decision trees have several important advantages for our aim as compared with classifiers such as neural networks, support vector machines or boosting methods. First, they are lightweight yet effective predictors that can function on a mobile phone. Second, the internal representation of a decision tree is highly human-interpretable and thus can inform application design decisions. Furthermore, a decision tree built with the ID3 algorithm doubles as a feature selection mechanism. ID3 works by greedily applying an information gain criterion for selecting which features to use for prediction. Thus, the features near the root of the tree have high predictive power and can be thought of as the most important features. Therefore, the initial question of whether there are common features of significance across all individuals can be addressed by determining whether the root of the decision tree varies across users and, if so, how.

This challenge can be formulated as a supervised learning problem, in which the class labels are the three levels of proximity, and each instance is a feature vector encoding the logged contextual information. We first tested the performance of the decision trees by using all three weeks of data for each user. We used 10-fold cross-validation to ensure effective use of the entire data set without biasing the test phase. To ascertain how many weeks of training were actually necessary for high accuracies, we conducted tests using restricted training sets. In one test, we used the first two weeks of data for the training set and the third week for the test set. In the second, we used the first week of data for training and the other two weeks as the testing set.

**Table 3.** Classification accuracies in percentages. The test using 3 weeks of data was conducted using 10-fold cross-validation over the entire data set.

| Participant | Majority Classifier | Tree: 3 wks training | Tree: 2 wks training | Tree: 1 wk training | Tree: cell id, hour, & day | Tree: the top 4 features |
|---|---|---|---|---|---|---|
| 1 | 77.0 | 90.1 | 89.2 | 87.7 | 85.0 | 88.1 |
| 2 | 53.6 | 88.9 | 87.2 | 85.8 | 80.2 | 86.5 |
| 3 | 59.4 | 93.1 | 88.0 | 85.8 | 78.4 | 85.4 |
| 4 | 49.8 | 86.1 | 85.1 | 84.0 | 85.8 | 86.0 |
| 5 | 68.3 | 85.0 | 83.7 | 82.9 | 73.5 | 82.5 |
| 6 | 53.8 | 86.3 | 85.9 | 85.1 | 79.2 | 84.0 |
| 7 | 81.5 | 90.1 | 90.0 | 88.2 | 90.0 | 90.0 |
| 8 | 65.9 | 88.7 | 87.9 | 86.5 | 82.1 | 85.3 |
| 9 | 84.6 | 91.0 | 90.8 | 89.4 | 88.1 | 88.9 |
| 10 | 52.4 | 90.1 | 89.2 | 88.3 | 86.6 | 88.5 |
| 11 | 52.8 | 84.1 | 82.6 | 80.9 | 82.4 | 83.9 |
| 12 | 83.1 | 89.6 | 86.8 | 84.4 | 84.8 | 87.4 |
| 13 | 60.5 | 85.6 | 84.3 | 84.2 | 75.2 | 82.8 |
| 14 | 64.7 | 87.4 | 85.2 | 84.8 | 84.0 | 85.6 |
| 15 | 40.7 | 87.3 | 86.9 | 85.3 | 83.8 | 86.1 |
| 16 | 53.6 | 90.1 | 89.4 | 89.0 | 84.4 | 87.5 |
| Averages | 62.61 | 88.34 | 87.01 | 85.77 | 82.72 | 86.16 |

As a baseline, we compared decision trees against a majority classifier to demonstrate how much additional predictive power a decision tree actually provides (see Table 3). On average, the decision tree classifier ranged in accuracy from 85-90%. The subjects with high majority classification accuracies tended to be arms-length from their phones for significant periods; however, the prediction accuracy still improved when using decision trees. For subjects with an even distribution among the three proximity levels, the majority classifier performed poorly, as expected, and the decision trees dramatically increased the accuracy. Reducing the training set down to one week still provided classification accuracies between 84-88%, suggesting that one week of training is sufficient to provide near-optimal prediction accuracies.

### 5.2 Analyzing the Decision Trees

We analyzed the decision tree for each of the 16 participants and determined the most important features for classification (see Table 4). For every subject, either cell tower ID, hour of the day, or day of the week was the root node (*i.e.,* the most predictive feature). These three features also appeared in the top four contributing features for each subject. Depending on the user, the remaining best feature was one of signal strength, charger status, or ring/vibrate status. To test the power of these top features, we calculated the decision tree accuracies for each subject using only their top four features (see Table 3). Restricting the feature set to only the top four features does not result in a large decrease in accuracy. We also computed the classification accuracies using only the cell tower ID, hour, and day features for each subject and observed an average accuracy of 83% (an average accuracy loss of 5%). Thus, comparable classification accuracies can be obtained with a common set of minimal features, meaning we do not even need to have a training phase in practice.

**Table 4.** This table shows that three groups of users emerged based on their top four features. Note that we present the features in no particular order of predictive power.

| % of Participants | Feature 1 | Feature 2 | Feature 3 | Feature 4 |
|---|---|---|---|---|
| 19 | Cell ID | Hour | Day | Ring/Vibrate |
| 50 | Cell ID | Hour | Day | Charger |
| 31 | Cell ID | Hour | Day | Signal |

Time and location are major factors for predicting user proximity to a mobile phone. Participants with hour or day as their top feature typically had structured workdays in which they interacted with the phones in a consistent pattern. For some users, the ring or vibrate status was a good indicator for proximity to the phone. The phone being within arm's reach often correlated to the acts of disabling the ring volume and activating the vibrator. On the other hand, a high ring volume often correlated to the phone being distant from the user. Many users typically carried the phone very close to them when they were away from home, as determined by cell tower IDs. Charger status and signal strength may have provided some subtle location information as well. For example, participants that only charged the phone in the car were very likely to be within arm's reach of the phone during charging. Users that only charged at home tended to be further away from the phones during charging. Often, the signal strength branched from cell tower IDs in the decision trees, indicating that the signal strength was playing a disambiguating role in those cases.

## 6 Discussion

### 6.1 Potential Alternative Data Gathering Methods

We considered several different methods of data collection when designing this experiment. In addition to the constant logging method, we considered conducting an experience sampling method (ESM), such as that used by Consolvo and Walker [2], or using solely self-reported data via a diary study, surveys, or interviews. We did not believe that self-report would give us the fine-grained, accurate information we needed, and when conducting the actual study, we observed that individuals often could not even accurately report where the phone was in the past day, even though they could remember the episodes of the day clearly.

Diary studies would likely have required too much work from the participants to get a broad range of samples. ESM, on the other hand, might have been appropriate, but we were concerned that random sampling would not uncover the subtle details inherent to user habits with mobile phones. To test this hypothesis, we randomly selected 16 data points per day (1 per hour) from each participant (one probe per waking hour) and calculated the average proximity level for each participant. We calculated this average for 100 random samplings and an overall average was calculated (see Table 5). The simulation assumed participants would be willing and

**Table 5.** Comparison of empirical proximity data to percentages from a simulated ESM study

| Level | Overall Empirical | Overall ESM | Weekend Empirical | Weekend ESM | Weekday Empirical | Weekday ESM |
|---|---|---|---|---|---|---|
| Arms | 58 | 61 | 53 | 55 | 59 | 64 |
| Room | 20 | 14 | 18 | 13 | 20 | 15 |
| NA | 23 | 24 | 28 | 32 | 21 | 21 |

able to respond to 16 queries per day for a three-week period. The percentages provided by the simulations were close to the actual data for most participants, despite only having 16 samples per day, compared to 1440 per day for our empirical study.

Although the overall percentages were similar, the ESM data would not include some of the more fine-grained details we were able to harvest from the high resolution, automatically-collected data, including times when the user was away from their phone for a short period. For example, we calculated the number of times per day each person was away from his or her phone for a short amount of time (2-20 minutes), *i.e.*, a "quick trip" as defined above. The participants each reported from 1 to 20 of these quick trips away from the phone per day. This information could be crucial to applications on a phone that assume the user is nearby all the time, such as reminder systems or constant health monitoring. Furthermore, with a sampling method, overall trends in some of the other features, such as phone call usage and number of cell towers detected by the phones would likely be missed. Lastly, and possibly most importantly, the empirical study did not require users to be conscious of their phone's location at all times, thus allowing us to capture a more realistic data set.

### 6.2 Design Considerations for Mobile Applications

The empirical results presented in this paper begin to uncover some interesting insights for mobile phone application designers to consider.

- Mobile phones may not be as good of a location proxy as many people believe. The participant with the closest overall proximity level was within arm's reach of his phone 85% of the time, despite his strong intuition that he carried the phone nearly 100% of the time.
- Certain features, such as the number of minutes of "talk time" are not as good predictors as intuition might have us believe.
- When considering a particular group of users, designers can leverage simple information available on the phone itself (time and location) to infer a user's proximity to the mobile phone.
- When away from home, the phone is more likely to be with the individual. Thus, designing applications for use in home would need to make different assumptions about a mobile phone's proximity than those for outside the home.
- The effect of physical activity on participant choices about phone proximity is an indicator that those potential applications that focus on monitoring of

physical fitness activities should consider the physical reasons users might avoid carrying a phone during these periods in designing their form factors.

- As expected, users keep their phones near in the presence of perceived needs. Thus, one solution for applications that require the user and phone being close most of the time is to build in functionality that the users may need regularly.
- Control over disruptions is important to users. Thus, any applications relying on interruptions must consider social, regulatory, and personal reasons for minimizing disruptions.

### 6.3 The Value of Proximity Modeling

A small number of features can predict the likelihood of proximity with fairly high confidence. Some of these features are the same across all participants, such as cell tower ID, hour and day, and together yield 83% predictive power. The ease of sensing these features on a mobile phone and the availability of lightweight machine learning techniques suggest that it is possible to build a context-aware mobile phone that can predict relatively easily the user's proximity to the phone. Such a system is valuable for applications that rely on the mobile phone as a proxy for a user, because it would allow for appropriate adaptation to situations in which proximity is a concern.

Central to the development of such a system is the model of the individual user. A simple tagging scheme, such as the one used for this study, can result in an accurate model with one week of "typical" use. However, this tagging method may not be practical for everyone. Thus, creating predefined or easy to construct models *a priori* for particular types of mobile phone users is an important consideration for effective adoption of this kind of system. Having identified common features across users and categories of user with respect to these features, we believe it possible to devise a survey mechanism in which users answer high-level questions. The answers to these questions could then translate into low-level modeling information to form the basis of a proximity-aware mobile phone.

## 7 Conclusions and Future Work

We presented an empirical study of sixteen users and the proximity relationships to their phones over a three-week period. We demonstrated that people often have higher expectations of their own proximity and availability to their mobile phones than is accurate in reality. Furthermore, we catalogued those factors that may influence them to carry (or not) their phones with them. Based on this empirical evidence, we present information for application designers attempting to meet the needs of users with particular patterns, including providing an understanding of the likelihood of availability to the phone in general as well as ways to model the proximity relationship automatically to provide individualized services and preferences.

While this study provided insight into the proximity relationship between a diverse group of users and their phones and how this might be predicted, there is still more that needs to be done to obtain generalizable information. Participants were particular to one metropolitan area, and phone use is likely to differ substantially by geographical and cultural regions. Additionally, due to the narrow but deep nature of our investigation, generalization of results to the larger population, particularly across varied demographics is difficult. The promise of similar proximity results via our simulated ESM study, and the need for only one week's worth of data for our classifier suggest that a more large-scale investigation can be done across varied populations with a less involved experimental design. Furthermore, this same information may be useful in determining an appropriate set of simple questions to allow end users to "train" the phone without collecting a large number of samples.

The nature of this type of study allowed us to obtain realistic proximity data for users that may not have otherwise been obtained with more low fidelity studies. Although a sampling of data points obtained through ESM can come up with similar proximity relationships, it runs the risk of altering the user's proximity relationship to the phone by continually reminding users about their phones' whereabouts. We believe the type of study reported in this paper useful to obtain ground truth data about a user's proximity relationship to the phone. Perhaps more significantly, however, it can also result in baseline data to compare against similar proximity evidence that would result from the effects of new mobile phone applications, such as location-based services, continual health monitoring systems, or context-aware applications, will have on that proximity relationship. Finally, this same technique may be used to evaluate proximity relationships between collections of mobile phones and their owners as well as the proximity relationships between people and other technologies, mobile or stationary.

## Acknowledgements

This work was sponsored by Intel Corporation and the National Science Foundation and is covered under IRB protocol H05205 at Georgia Tech. We also thank Nokia Research, and in particular Dana Pavel, for the donation of the mobile phones used in this study and the members of the Ubicomp Research Group at Georgia Tech.

## References

1. Cohen, J., A Coefficient of Agreement for Nominal Scales. *Educational and Psychological Measurement*, 1960. **20**: p. 37-48.
2. Consolvo, S. and M. Walker, Using the experience sampling method to evalute ubicomp applications, in *IEEE Pervasive Computing*. 2003. p. 24-31.
3. Corner, M.D. and B.D. Noble. Zero-Interaction Authentication. in *Mobicom '02*. 2002. Atlanta, GA USA: ACM.
4. Davis, M., N. Good, and R. Sarvas. From Context to Content: Leveraging Context for Mobile Media Metadata. in *International Multimedia Conference: 12th annual ACM international conference on Multimedia*. 2004. New York, NY, USA.
5. Demumieux, R. and P. Losquin. Gather customer's real usage on mobile phones. in *Conference on Human Computer Interaction with Mobile Devices & Services*. 2005.
6. Eagle, N. and A. Pentland, Reality Mining: Sensing Complex Social Systems. *Personal and Ubiquitous Computing*, 2005. **10**(4): p. 255-268.
7. Hayes, G.R., S.N. Patel, K.N. Truong, G. Iachello, J.A. Kientz, R. Farmer, and G.D. Abowd. The Personal Audio Loop: Designing a Ubiquitous Audio-Based Memory Aid. in

*Mobile HCI 2004: The 6th International Conference on Human Computer Interaction with Mobile Devices and Services*. 2004. Glasgow, Scotland.
8. Helal, S., C. Giraldo, Y. Kaddoura, C. Lee, H.e. Zabadani, and W. Mann. Smart phone based cognitive assistant. in *UbiHealth 2003: The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications*. 2003. Seattle, Washington.
9. Hightower, J. and G. Borriello, A Survey and Taxonomy of Location Systems for Ubiquitous Computing. 2001, University of Washington.
10. Kahneman, D., A.B. Krueger, D.A. Schkade, N. Schwarz, and A.A. Stone, A Survey Method for Characterizing Daily Life Experience: The Day Reconstruction Method, in *Science*. 2004. p. 1776-1780.
11. Laasonen, K. Clustering and Prediction of Mobile User Routes from Cellular Data. in *PKDD 2005*. 2005: Springer Verlag.
12. LaMarca, A., et al. Place Lab: Device Positioning Using Radio Beacons in the Wild. in *Pervasive 2005*. 2005. Munich, Germany.
13. Lamming, M. SPECs: Another Approach to Human Context and Activity Sensing Research, Using Tiny Peer-to-Peer Wireless Computers. in *Ubicomp 2003*. 2003.
14. McGuigan, J., Towards a Sociology of the Mobile Phone. *Human Technology*, 2005.
15. Mitchell, T., *Machine Learning*. 1996, New York, US: McGraw Hill.
16. Otsason, V., A. Varshavsky, L. A., and E. de Lara. Accurate GSM Indoor Localization. in *The Seventh International Conference on Ubiquitous Computing (Ubi-Comp 2005)*. 2005.
17. Oulasvirta, A., M. Raento, and S. Tiitta. ContextContacts: Re-Designing SmartPhone's Contact Book to Support Mobile Awareness and Collaboration. in *MobileHCI 2005*. 2005.
18. Palen, L. and M. Salzman, Beyond the handset: designing for wireless communications usability. *ACM Transactions on Computer-Human Interaction*, 2002. **9**(2): p. 125-151.
19. Palen, L., M. Salzman, and E. Youngs. Going Wireless: Behavior & Practice of New Mobile Phone Users. in *CSCW '00: Computer Supported Cooperative Work*. 2000.
20. Patel, S.N., J.S. Pierce, and G.D. Abowd. A gesture-based authentication scheme for untrusted public terminals in *17th annual ACM symposium on User interface software and technology* 2004. Santa Fe, NM, USA: ACM Press.
21. Preece, J., Y. Rogers, and H. Sharp, *Interaction design: beyond human-computer interaction*. 2002: John Wiley & Sons, Inc.
22. Raento, M., A. Oulasvirta, R. Petit, and H. Toivonen, ContextPhone - A prototyping platform for context-aware mobile applications, in *IEEE Pervasive Computing*. 2005.
23. Satyanarayanan, M., Swiss Army Knife or Wallet?, in *IEEE Pervasive Computing*. 2005.
24. 24. Schlosser, F.K., So, how do people really use their handheld devices? An interactive study of wireless technology use. *Journal of Organizational Behavior*, 2002. **23**(4): p. 401-423.
25. Sullivan, J. and G. Fischer. Mobile Architectures and Prototypes to Assist Persons with Cognitive Disabilities using Public Transportation. in *26th International Conference on Technology and Rehabilitation*. 2003. Atlanta GA, USA.
26. Vaananen-Vainio-Mattila, K. and S. Ruuska, User Needs for Mobile Communication Devices: Requirements Gathering and Analysis through Contextual Inquiry, in *First Workshop on HCI and Mobile Devices*. 1998.
27. Wang, X.H., R.S.H. Istepanian, and Y.H. Song, Mobile e-Health: The Unwired Evolution of Telemedicine. *Telemedicine Journal and e-Health*, 2003.
28. Want, R., T. Pering, G. Danneels, M. Kumar, M. Sundar, and J. Light. The Personal Server: Changing the Way We Think about Ubiquitous Computing. in *Ubicomp 2002*. 2002. Sweden: Springer-Verlag.