

UC San Diego

UC San Diego Electronic Theses and Dissertations

Title

Addressing the needs of mobile users

Permalink

<https://escholarship.org/uc/item/1f7645wx>

Author

Sohn, Timothy Youngjin

Publication Date

2008

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Addressing the Needs of Mobile Users

A dissertation submitted in partial satisfaction of the
requirements for the degree
Doctor of Philosophy

in

Computer Science

by

Timothy Youngjin Sohn

Committee in charge:

Professor William G. Griswold, Chair
Professor Geoffrey M. Voelker
Professor Sanjoy Dasgupta
Professor James D. Hollan
Professor Barry Brown
Dr. Anthony LaMarca
Dr. Larry Marturano

2008

Copyright
Timothy Youngjin Sohn, 2008
All rights reserved.

The dissertation of Timothy Youngjin Sohn is approved, and it is acceptable in quality and form for publication on microfilm:

Chair

University of California, San Diego

2008

DEDICATION

To my parents
Paul and Jeannie Sohn

TABLE OF CONTENTS

	Signature Page	iii
	Dedication	iv
	Table of Contents	v
	List of Figures	viii
	Acknowledgements	ix
	Vita and Publications	xii
	Abstract of the Dissertation	xiii
Chapter 1	Introduction	1
	1.1 Mobile Phones as an Enabling Technology	2
	1.2 Prior Work	4
	1.2.1 Information Seeking Behavior	4
	1.2.2 Information Deferral	5
	1.2.3 Systems to Assist Mobile Users	5
	1.3 Overview	6
	1.3.1 Contributions of the Thesis	7
	1.4 Outline	9
Chapter 2	A Diary Study of Mobile Information Needs	10
	2.1 Introduction	10
	2.2 Methods	12
	2.2.1 Participants	12
	2.2.2 Procedure	13
	2.3 Results and Discussion	15
	2.3.1 Taxonomy of Information Needs	15
	2.3.2 Deciding When to Address a Need	18
	2.3.3 How Does Mobile Internet Change Behavior	28
	2.3.4 What Prompts Information Needs?	30
	2.3.5 Where Does The Information Come From?	31
	2.4 Design Implications	32
	2.5 Conclusion	34
	2.6 Acknowledgements	35

Chapter 3	GSM-Based Localization Technology for Mobile Users	36
3.1	Introduction	36
3.2	Location Sensing	38
3.2.1	Infrastructure-Based Sensing	38
3.2.2	Beacon-Based Approach to Location	39
3.2.3	Coverage and Accuracy	39
3.3	Evaluating GSM Beacon-Based Localization	40
3.3.1	The Place Lab Model	40
3.3.2	Beacon-Based Localization at a Metropolitan Scale	42
3.4	Conclusion	43
Chapter 4	GSM-Based Mobility Detection	44
4.1	Introduction	44
4.2	Mobility Detection With GSM	46
4.2.1	Global System for Mobile Communication (GSM)	46
4.2.2	Inferring User Mobility Modes	47
4.2.3	Estimating a User's Daily Step Count	50
4.3	Experimental Evaluation	51
4.3.1	Data Trace Collection	51
4.3.2	Inferring Mobility Modes	53
4.3.3	Daily Step-Count Prediction for Data Collectors	55
4.4	Applications	56
4.4.1	Computer-Supported Coordinated Care (CSCC)	56
4.4.2	Social-Mobile Applications	57
4.5	Related Work	58
4.6	Conclusions and Future Work	59
4.7	Acknowledgements	60
Chapter 5	Place-Its: Location-Based Reminders for Mobile Users	61
5.1	Introduction	61
5.2	Related Work	63
5.3	The Place-Its Application	66
5.3.1	Mobile Phone Platform	68
5.3.2	Location-Sensing	69
5.4	User Study	71
5.4.1	Participants	72
5.4.2	Methodology	73
5.5	Observations and Initial Classification	74
5.5.1	Self-Reported Usage Data	75
5.5.2	Classifying Place-It Note Usage	77
5.6	Post-Study Responses and Discussion	79
5.6.1	Mobile phones offer available, convenient reminder creation and delivery	79

	5.6.2	Location Provides an Indirect Cue for Other Context	81
	5.6.3	Location-based reminders are useful	83
	5.7	Conclusions and Future Directions	84
	5.8	Acknowledgements	85
Chapter 6		Outlook on Mobility Research	86
	6.1	Introduction	86
	6.2	Gathering More Data About Mobile Users	87
	6.3	The Need for Just-In-Time Delivery	87
	6.4	Mobile Digital Paper	88
	6.4.1	Motivation	89
	6.4.2	Scenario	90
	6.4.3	Bridging the Paper to Digital Divide	91
	6.4.4	System Architecture	92
	6.4.5	Lessons Learned and Future Evaluation	95
	6.5	Conclusion	95
Chapter 7		Conclusion	97
	7.1	Findings	98
	7.2	Final Thoughts	99
Bibliography		101

LIST OF FIGURES

Figure 1.1: Overview of the mobile information space.	3
Figure 2.1: Taxonomy of mobile information needs	16
Figure 2.2: Pie chart of when needs were addressed	21
Figure 2.3: Pie chart of how needs were addressed	22
Figure 2.4: Pie chart of why needs were addressed later	25
Figure 2.5: Pie chart of why needs were not addressed	26
Figure 2.6: Contextual factors in addressing needs	30
Figure 3.1: The Place Lab Model	41
Figure 3.2: Wardrive map of Seattle	43
Figure 4.1: Mobility experiment figure.	49
Figure 4.2: Mobility results figure.	53
Figure 5.1: Place-Its usage screenshots	67
Figure 5.2: Place-Its location detection	71
Figure 5.3: Place-Its participants	72
Figure 5.4: Place-Its reminders usage	74
Figure 5.5: Self reported answers	76
Figure 5.6: Place-Its note classification	77
Figure 6.1: The mobile digital paper system	92

ACKNOWLEDGEMENTS

This dissertation would not have been possible without the support of many friends, family, colleagues, and professors along the way. I want to thank the many individuals who supported me along this long road, and helped me complete this project.

My advisor, Bill Griswold, was one of the biggest champions for me during my graduate studies. He pushed me to be a visionary and develop my own creative ideas. Although I questioned my research abilities at many points along the way, Bill continued to instill confidence in me to stay the course. Without Bill's guidance, coaching, and encouragement, I would have never finished this project.

Thanks to Jim Hollan who acted in many respects as a co-advisor. I'm also grateful to my other committee members who provided helpful feedback and guidance in putting this work together.

I'm also very thankful for my undergraduate advisor, Anind Dey, who coached me early on about human-computer interaction research. Even after going to graduate school, Anind helped me weather through many of the first and second year graduate school woes. Through Anind I was exposed to the Intel Research network that eventually played a large part in the direction of my research.

The many researchers and interns at Intel Research became my community, and many of my projects reflect the collaborative efforts. Thanks to Anthony LaMarca and James Scott who initially proposed an idea to pursue GSM-based localization for an intern project, which later became a foundation for the rest of my research endeavors. The following people I met through the Intel labs were great colleagues and friends that helped formulate many of my early thoughts on this thesis: Mike Chen, Dmitri Chmelev, Tanzeem Choudhury, Sunny Consolvo, Susumu Harada, Jeffrey Hightower, James Howard, Jeff Hughes, Fred Potter, Ian Smith, Karen Tang, and Alex Varshavsky.

The researchers at Motorola were helpful in brainstorming research ideas, and provided numerous resources that helped me explore ideas for mobile applications. These talented researchers include Larry Marturano, Frank Bentley, Crysta Metcalf, Gunner Harboe, and Guy Romano.

Many of the other graduate students in the Computer Science department were very supportive in helping me refine my research ideas. I'd like to thank Kevin Li, Daniel Hsu, Roshni Malani, Neil McCurdy, Patricia Shanahan, Barry Demchak, Elizabeth Bales, Lisa Cowan, and Justin Burke.

Thanks to Chunyuan Liao who provided access to the Papiercraft source code, which was the basis for the mobile digital paper work described in Chapter 6.

Thanks to the many people at Lighthouse Bible Church who quickly became my spiritual community after moving to the area. Over the years I've met great people who I am deeply thankful for their friendship and influence in my life.

I'm thankful to my parents and sister who have supported me endlessly since I moved to San Diego. Many of the seeds for my academic pursuits were planted by my parents, Paul and Jeannie Sohn. I am incredibly thankful to them for their support throughout my academic career, and encouraging me to reach my full potential. Thank you mom and dad.

Candy Lam, my fiancée, has been at my side since I started my doctoral pursuits. She supported me by trying to understand my research, walking with me through the ups and down of a PhD life, reading PhD comics to laugh with me, incessantly praying for me, and always encouraging me to keep going. I'm so thankful to have gone through this part of my life with you.

Finally, none of this would have been possible without the hand of God at work in my life. He sustained me through this whole journey, and faithfully provided for every circumstance that I found myself in. Jesus is the reason I am here today, and the reason that I can live each day.

Chapter 2, in part, is a reprint of the material as it appears in A Diary Study of Mobile Information Needs. Timothy Sohn, Kevin A. Li, William G. Griswold, and James D. Hollan. In CHI 2008: Proceedings of the Twenty-Sixth Conference on Human Factors in Computing Systems. pages 433-442. ACM press, 2008. The dissertation author was the primary investigator and author of this paper.

Chapter 4, in part, is a reprint of the material as it appears in Mobility Detection Using Everyday GSM Traces. Timothy Sohn, Alex Varshavsky, Anthony

LaMarca, Mike Y. Chen, Tanzeem Choudhury, Ian Smith, Sunny Consolvo, Jeffrey Hightower, William G. Griswold, and Eyal de Lara. In *UbiComp 2006: Proceedings of the Eighth International Conference on Ubiquitous Computing*. pages 212-224. Springer, 2006. The dissertation author was the primary investigator and author of this paper.

Chapter 5, in part, is a reprint of the material as it appears in *Place-Its: A Study of Location-Based Reminders on Mobile Phones*. Timothy Sohn, Kevin A. Li, Gunny Lee, Ian Smith, James Scott, and William G. Griswold. In *UbiComp 2005: Proceedings of the Seventh International Conference on Ubiquitous Computing*, pages 232-250. Springer, 2005. The dissertation author was the primary investigator and author of this paper.

VITA

- 2003 Bachelor of Science in Electrical Engineering and Computer Sciences
University of California, Berkeley
Berkeley, CA
- 2005 Master of Science in Computer Science
University of California, San Diego
La Jolla, CA
- 2005-2008 Research Assistant
Department of Computer Science and Engineering
La Jolla, CA
- 2008 Doctor of Philosophy in Computer Science
University of California, San Diego
La Jolla, CA

PUBLICATIONS

Timothy Sohn, Kevin A. Li, William G. Griswold, and James D. Hollan, "A Diary Study of Mobile Information Needs", In Proceedings of CHI 2008, pp. 433-442.

Timothy Sohn, Alex Varshavsky, Anthony LaMarca, Mike Y. Chen, Tanzeem Choudhury, Ian Smith, Sunny Consolvo, Jeffrey Hightower, William G. Griswold, and Eyal de Lara, "Mobility Detection Using Everyday GSM Traces", In Proceedings of Ubicomp 2006, pp. 212-224.

Timothy Sohn, William G. Griswold, James Scott, Anthony LaMarca, Yatin Chawathe, Ian Smith, and Mike Y. Chen, "Experiences with Place Lab: an Open Source Toolkit for Location-Aware Computing", In Proceedings of ICSE 2006, pp. 462-471.

Timothy Sohn, Kevin A. Li, Gunny Lee, Ian E. Smith, James Scott, William G. Griswold, "Place-Its: A Study of Location-Based Reminders on Mobile Phones", In Proceedings of Ubicomp 2005, pp. 232-250.

ABSTRACT OF THE DISSERTATION

Addressing the Needs of Mobile Users

by

Timothy Youngjin Sohn

Doctor of Philosophy in Computer Science

University of California San Diego, 2008

Professor William G. Griswold, Chair

Mobile phones hold the promise of enabling the ubiquitous computing vision, bringing computing power into our daily lives. Using sensors and actuators built into the devices we carry, applications can provide us with context-aware information delivery. Ubiquitous computing systems can use sensing capabilities and information access to help people accomplish their tasks and take advantage of unseen opportunities. The high mobile phone penetration rate across the world makes this platform ideal for building many of these advanced applications.

In order to enable more advanced mobile applications, we face several technological and human challenges. Although mobile phones today are similar to computers from a decade ago they are still limited in their computation, network, storage, and input capabilities. These areas are rapidly improving to enable further innovation in mobile computing. More important than the technical challenges are the human challenges. Unlike desktop or nomadic computing, being mobile means there is an fundamental problem of limited attention that must be dealt with. Mobile users are often engaged in several tasks at once; trying to stay active in their main task and interact with their mobile device at the same time. A lack of attentional resources naturally limits the amount of interaction one can have with his mobile device, hence the need for technology and applications that are designed with the users' attentional resources in mind. Addressing this fundamental human problem alongside the other technical challenges open the way for advancing the

ubiquitous computing vision.

This dissertation provides a structure to the rapidly expanding mobile space by contributing work that focuses on addressing the fundamental problem that mobile users face, limited attention span . The contributions of this dissertation are three-fold: 1) an understanding of the fundamental problem that mobile users face, which is often being engaged in another task while trying to gather information, 2) demonstration of using low-fidelity mobile phone sensors for context-awareness that can be used to alleviate the problems mobile users face, and 3) a description of the unique character of ubiquitous computing applications and the promise for changing the way we live in our mobile world. These core contributions demonstrate the problems that mobile users face, as well as several technological and design solutions to lessen the burden that they feel. Future mobile advancements will need to account for the lack of attentional resources in order to build systems that help improve people's lives by enhancing the ubiquitous computing vision.

Chapter 1

Introduction

People are always in the process of gathering information. Whether in a desktop setting or a mobile context, information is always needed and sought after. These information needs can be actively pursued, conscious needs, such as directions to an event, or a subtle stream of information gathering that we do not pause to think about. Regardless of the type of need, people use a variety of methods to gather information to address their needs. Phone calls, internet access, and information deferral are all ways that needs are either addressed in the moment or put off for a later point in time. The methods that people choose to address their needs are often related to the attentional resource cost of the method; the cognitive load required for a person to gather information. Attentional resources can vary depending on the situation and a person's current task.

The high desire for information becomes more apparent in a mobile context because of the set of available tools. The type of resources available when mobile are different than in a stationary setting. The mobile input and output modalities are severely limited in comparison to a twenty-inch screen and a full-size keyboard. The high desire for mobile information naturally requires a high amount of attentional resources. Unfortunately, a mobile user has limited attentional resources that vary depending on the current situation. A mobile user is typically engaged in a main task, and an information gathering task at the same time. As a result, a mobile user must resort to multitasking between these two tasks in order to manage the attentional demands of each. As each task requires more resources, balancing

the cognitive load becomes more demanding and burdensome. The burdens often begin to outweigh the benefits, which can result in an information need never being addressed at all.

The challenge of mobile information access is alleviating the burden of mobile multitasking and reducing the strain typically felt by mobile users. There are several ways to address this problem. Improvement in mobile interaction could move the burden of information gathering from the user back to the mobile device, thereby decreasing the multitasking time. Automatic information gathering would also reduce multitasking time by providing relevant information at the right moments. Tools designed around a mobile user's limited attentional resources could create opportunities for quick information deferral. Deferring an information need for a later, more appropriate time eases the burden of access a mobile user may experience in the moment. These three solutions can help reduce the amount of time a mobile user must interact with his device, and alleviate the burdens of attention that a mobile users feels. The ideal platform to explore these solutions is the mobile phone.

1.1 Mobile Phones as an Enabling Technology

In the past decade there has been a shift in technological use from stationary, desktop computing to a mobile, always-on, always-connected experience. This transition towards true mobile computing has occurred, in large part, due to the global penetration of the mobile phone. These devices are the most pervasive mobile devices available today, orders of magnitude larger than any other mobile device (e.g., PDAs, laptops). Because of their pervasiveness, mobile phones hold the promise of enabling the ubiquitous computing vision, bringing computing power, context-awareness, and information access into our daily lives. Using sensors and actuators built into the devices, applications can provide us with context-aware information delivery. The future of ubiquitous computing is filled with opportunities to use sensing capabilities and information access to help us accomplish our tasks and take advantage of unseen opportunities.

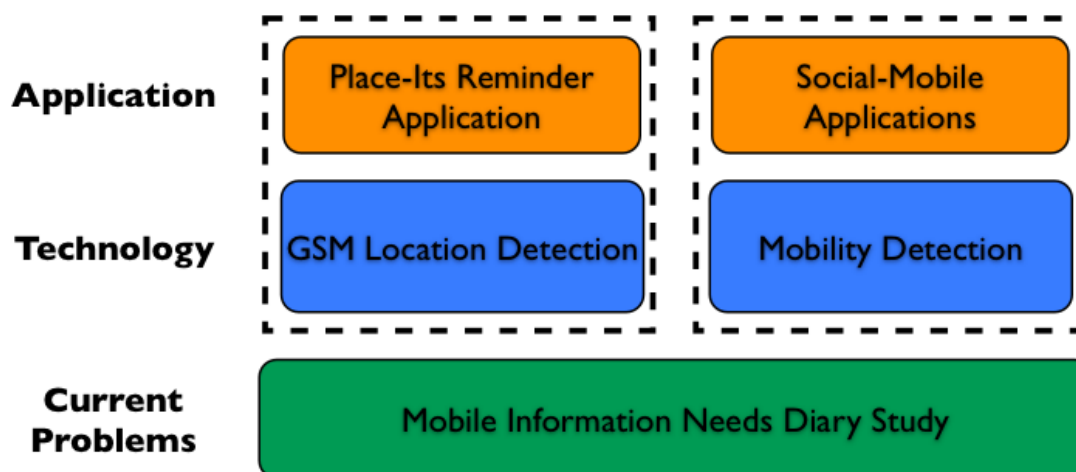


Figure 1.1: An overview of the mobile information space. There are three levels to this space: current problems, applications, and technology that supports the applications. The dotted lines represent two different applications that build on an understanding of the problems.

Mobile phone efforts today are focused on enabling information access. For example, many phones offer access to one's email, calendar, or map applications. Information access is limited though by the constraints of the mobile device. Mobile phones are constrained devices and the interactions problems extend in many dimensions. Small screens make it difficult to browse large mounds of information. Small keypads are cumbersome for input. Even with the advent of third generation cellular networks worldwide, the network is still limiting. Therefore, the burden of information access is often moved from the mobile phone to the human being. A person can use his mobile device to access the information, but must be able to interact within the constraints of the device to find the "right" information. For example, a search query provides access to information but may still require the human to browse through mounds of data for the most appropriate result. This burdensome shift magnifies the limited attentional resource problem mobile users face.

1.2 Prior Work

Progress in the mobile information space requires advances in three dimensions: understanding the problems mobile users face, building applications that alleviate these problems, and the underlying technology that supports the applications (Figure 1.1). Prior work has given insights into each of these areas, but has focused mainly on giving mobile users access to information. Access alone is not enough because of a mobile user's limited attentional resources.

1.2.1 Information Seeking Behavior

The study of information needs and information seeking behavior has typically been performed in the context of information retrieval systems. Search engines are the dominant method used for finding information on the web. Broder found that 48% of information searches are done through search queries [Bro02], while others have found the number to be upwards of 80% [JBS08]. Due to the popularity of web search engines, much of information-seeking behavior studies are grounded in the web. These studies have looked at how users navigate web sites and use search engines. Choo et al. studied how knowledge workers use the web to seek external information as part of their daily work [CDT00]. Kellar et al. conducted a field study to characterize web-based information-seeking tasks [KWS07]. Other researchers have looked at how web activities affect people's decisions to seek information. [MPC01].

Because web searches are the standard method to find information, as mobile internet access becomes increasingly popular, the mobile search domain has become the next frontier for studying information seeking. Kamvar and Baluja analyzed 1 million Google mobile search queries, noting the differences between mobile and desktop search [KB06]. Lee et. al. looked at mobile internet usage through a longitudinal study of the situations in which people use the mobile internet [KKL05]. Studies of mobile internet usage are helpful to understand what information mobile users are willing to search for. However, these studies do not help us uncover any of the unmet information needs that may never be addressed

through a web search. In order to comprehensively understand mobile users, we need ethnographic data about what mobile users actually do. Recently, Hinman et. al. used PC deprivation to look beyond formalized information needs, and try to explore the unmet needs of mobile users [HSI08]. Further studies in this space will help us grasp a more comprehensive picture about mobile users.

1.2.2 Information Deferral

One method to address the lack of attentional resources, is to create mobile tools for information deferral - a method of capturing information for later use. Some of the current methods people employ for this purpose are post-it notes, reminders, snapshot pictures, and emails. Capturing and deferring information for later use is an important way to handle information that may not be able to be addressed in the moment. An effective information deferral system is needed in order to deal with the burdens of mobility.

One mobile method for information deferral is through image capture and access [KSFS05]. Camera phones are low-cost, ubiquitous devices that allows for quick capture of information. An image can also be easily discarded after the information has been retrieved at a later point in time. Jones et. al. created a background search tool for information deferral that can also help turn visceral needs into conscious ones [JBCJ06]. The tool lets users capture search terms offline on a mobile device, and then address them when reconnected to a PC. Aridor et. al. presented another offline-online system that searches based on an intelligent agent [ACM⁺02]. Other related mobile methods for information deferral use visual codes to read data that can be accessed later [KRA01]. Unlike these methods, the applications we will describe in the following chapters utilize contextual information for a richer method of information deferral.

1.2.3 Systems to Assist Mobile Users

There have been a number of mobile context-aware systems that attempt to preempt the needs of mobile users. Automatically delivering relevant information

reduces the difficulty of mobile interaction, and moves the burden of information seeking from the person to the device. Domain-specific examples include context-aware tour guides [AAH⁺97] and local event notification systems [NTH06]. Delivering relevant content to mobile users before they request the information can reduce potential interaction time. For example, a context-aware tour guide would prepare information onscreen about relevant artifacts in the environment without a person having to search for them. Most of these mobile systems are designed around a single data source (e.g., local events). Since the needs of mobile users are various and broad, a more general system could be beneficial for anticipating a person’s everyday needs. Mining one’s personal data could provide rich data of the types of relevant information for a mobile user.

The Jimminy system was one of the first systems that fused physical context with personal data to provide a broad assortment of information that could be relevant to the user [Rho03]. The system used the author’s real personal notes, along with context information about the users location, and nearby people or objects. Although the findings showed that analyzing the current note being taken was more useful than physical context, much of the sensing ability was limited to sections of a research lab. Thus, the findings are inconclusive as to how useful the system would be if deployed in the wild and ubiquitously available.

Magitti is a mobile recommender system that suggests leisure activities within the area [BBC⁺08]. The system uses context and patterns in a user behavior to recommend relevant restaurants and activities. Despite technical challenges, their evaluation in the wild suggests that Magitti can deliver relevant content to a user. Although Magitti only predicts leisure activities, it is a first step in building systems that can predict a mobile user’s needs and reduce the burdens of mobile interaction.

1.3 Overview

Prior work gives us insights into the information seeking habits of mobile users, as well as tools and applications that can benefit them. However, there are

still areas where the prior work is lacking. The study of mobile users has mainly been focused on their internet usage, rather than an ethnographic approach to what mobile users need and the tasks they perform. Moreover, the applications that have been built are not designed with the fundamental problem of mobility in mind. Furthermore, many of these applications were evaluated in localized settings without ubiquitous coverage, which limits the systems coverage. In contrast, this thesis focuses on studying mobile user's in their natural environment, detached from other technological constraints (i.e., mobile internet). We also discuss the design and evaluation of context-aware applications that are available ubiquitously because of the worldwide penetration of mobile phone technologies.

In order to address the fundamental problem of mobility, we hypothesize that the affordances of mobile phones, in particular context-awareness, can be used to successfully meet the needs of mobile users. Despite the low fidelity of sensors and actuators on mobile phones, they can still be appropriated for detecting useful pieces of context that can lower the burden many mobile users feel. We believe that a key concept for addressing the attention deficiencies is information deferral - capturing information needs when they occur, and using contextual support to address them later.

Understanding the mobile information space is essential to designing appropriate applications that address the fundamental attentional resource problem of mobility. We need to consider this limitation as we create mobile tools and applications that can alleviate the burdens mobile users often feel. The mobile information space can be broken down into three levels (Figure 1.1). This logical distinction helps describe the contributions of this thesis.

1.3.1 Contributions of the Thesis

This dissertation provides a structure to the rapidly developing mobile space by presenting key research work in three areas. The focus of this work is on addressing the needs of mobile users, specifically looking at how to create tools that can be used in light of a person's attentional limitations. The contributions of this thesis are three fold:

1. The first contribution is a foundation for understanding the problems and challenges that mobile users face in trying to address their information need as evidenced through an *in situ* diary study. Studying what mobile users actually do provides empirical data for the many intuitions about how mobile multitasking is handled. Many believe that the mobile internet will solve the numerous information need problems mobile users face. We demonstrate that mobile internet is useful, but the lack of attentional resources is the fundamental problem mobile users face when attempting to address their information needs. Needs are often left unaddressed or deferred for a later point in time.
2. The second contribution is context detection technologies that can be used to alleviate the burdens that mobile users often feel. We show how to infer one's location and high-level mobility patterns using commodity GSM phones. Without any hardware modification, we demonstrate that GSM patterns can accurately detect context useful for a variety of applications. GSM signal patterns can be used to detect high-level mobility states such as whether a person is stationary, walking, or driving. Mobility detection can be used to schedule reminders for information deferral, or provide automatic status updates for disseminating information to a person's social network. Mobility detection creates opportunities to explore rich mobile applications.
3. The third contribution is an understanding of how ubiquitous computing applications can address the fundamental problem of limited attentional resources. We discuss the development and evaluation of two context-aware applications that enable information deferral, a key requirement for mobile users. When the burdens of multitasking are too costly, mobile users need a method for deferring their information need for a more appropriate time.

The first application is a location-based reminder system called Place-Its, that uses location to tag information at places. Location-based reminders have been proposed in previous work, but were never studied in the daily lives of mobile users. Reminders are an effective way to defer information

needs, but when fused with context they become filled with a rich mechanism for intuitive information deferral. We describe an evaluation of Place-Its in the wild using GSM location technology.

The second application is a socially-acceptable method for deferring information by using mobile digital paper. Certain situations are inappropriate for technological use, thus a socially appropriate method is needed. We describe a prototype application that uses Anoto pen technology to capture information needs and defer them for a later time. The application can be personalized to use context in order to retrieve relevant information for the user.

1.4 Outline

This thesis is divided into three parts. The first part is Chapter 2 where we will discuss a diary study about people’s mobile information needs. Our diary study shows areas where further research can address many of the problems that mobile users face today. This chapter frames the following work in the dissertation. The second part explores context-aware technologies that enable a variety of mobile applications related to the design requirements outlined previously. Chapter 3 will discuss background work on how to obtain positioning information on mobile phones using today’s commodity devices. Based on a similar principle about GSM signal characteristics, Chapter 4 presents algorithms for determining high-level mobility detection. These technologies support the applications described in the third part of the dissertation. Chapter 5 presents a location-based reminder system that is an example of a tool mobile users can appropriate for deferring their information needs. We will discuss the prototype and evaluation of a GSM location-based reminder system. Chapter 6 explores mobile digital paper, another type of information deferral application that focuses on social appropriateness. Finally, Chapter 7 concludes this thesis with a summary of contributions and descriptions of several avenues for future work.

Chapter 2

A Diary Study of Mobile Information Needs

In order to address the needs of mobile users, we need to understand the types of information people seek and the methods they use to address their needs while mobile. Although we have many intuitions about these issues, we lack empirical data to support many of our intuitions. This chapter explores what mobile users actually do when looking for information on the move as evidenced through an *in situ* diary study. Based on this diary study, in later chapters, we will describe several applications that can help address the needs of mobile users.¹

2.1 Introduction

People often need information while on the go. Sometimes the information required is essential to the task at hand, such as finding a hotel for the night. Other times, the need is associated with a question prompted by a conversation or a nearby object (e.g., a billboard). Based on the importance of the need and the amount of time available, people use a variety of strategies to obtain the desired information.

Mobile phones provide increasingly convenient ways to obtain information.

¹Note that this chapter is a reprint with minor changes of *A Diary Study of Mobile Information Needs*, a paper co-authored by Timothy Sohn, Kevin Li, William Griswold, and James Hollan.

Internet enabled phones give access to the mobile web, as well as personal email and calendars. Although mobile devices are becoming always connected, their limitations can be debilitating. Restricted input, small screens, and complex interfaces are particularly challenging when a person is mobile and rushed.

Many have proposed ways of transforming desktop browsing to the constrained display of mobile devices [BHR⁺99], [BGMPW00], [HR04], [LB05]. However, simply providing mobile users with access to the internet and desktop tools is insufficient. Mobile users need applications and services that are designed to the particular requirements of mobile context and use. As just one example, mobile users are often preoccupied with the things going on around them. As a consequence, they often need to decide if they have sufficient time and attentional resources to access potentially useful information services. Recently, some companies have attempted to bridge this gap with mobile content-driven services such as GOOG-411², Microsoft Live Mobile³, and Google Mobile⁴.

Although desktop search queries have been well studied, the mobile search space has been less adequately explored [JSBS98], [SMHM99], [SJWS02]. Kamvar and Baluja recently analyzed 1 million Google mobile search queries, finding people had different uses for mobile search vs. desktop search [KB06]. This is incredibly valuable data but only tells what people use search engines for, not what information needs they really had. This is further restricted by sampling only people who own devices that support such functionality. Lee et al. also looked at mobile internet usage through a longitudinal study of the situations in which people use the mobile internet [KKL05]. As a complement to analyzing mobile clickstream data, we need to understand what types of information people need while on the go and how they address those needs. Observing people's behaviors in such situations could point to improved mobile interface and system designs.

We report a diary study of 20 people's mobile information needs over the course of two weeks. We examined the types of information needs participants had, the strategies and methods they used to address those needs, and the contextual

²<http://www.google.com/goog411>

³<http://mobile.search.live.com/about>

⁴<http://google.com/mobile>

factors that prompted each need and influenced how it was addressed. Our study reveals that mobile contexts frequently require people to adopt novel strategies and methods. Less essential needs are often put off until later, or never addressed at all. However, needs that are essential to the task at hand are typically addressed through the perceived lowest cost method available, which is not always internet access. Based on our findings, we present several suggestions for designing future mobile technology.

2.2 Methods

There are many methods for capturing in situ data from mobile users. We considered the increasingly popular Experience Sampling Method (ESM), however we felt that the sampling frequency would need to be too high in order to capture any moment a person had an information need. In contrast, diary studies allow participants to do the capturing whenever needs arise. Diary studies suffer the drawback of potentially missing data because participants forget to record entries or are selective in reporting (e.g., thinking some events are not important enough to report). Still, we chose to use a diary study because we thought it would be the most effective technique of capturing data to reveal the nature of information needs and how they unfolded in-situ. Below we describe how we managed the burden of maintaining a diary, given the detail we desired.

2.2.1 Participants

We recruited 20 participants (10 male, 10 female) through online mailing lists and flyers. All were required to own a mobile phone and have experience sending text messages. Ages ranged from 19-58 (mean: 30.7, SD: 10.4). We selected a diverse mix of participants including undergraduate/graduate students, human resources personnel, a financial analyst, caregivers, engineers, a homemaker, project managers, a pastor, and a temporary worker. Five of our participants used mobile internet access on their phone. The other fifteen did not have mobile internet capabilities either because their phone was not capable, or because they

felt that the cost was too high.

2.2.2 Procedure

We asked participants to keep a diary for two weeks of all their information needs whenever they were mobile, defined as being away from home or work. Although information needs arise when at home or work, we wanted to focus on and capture those that occurred when people were mobile and away from the places where they spend most of their time. For participants with desk jobs, we asked them to record entries when they were away from their desk. We also asked participants to make entries about any information gathering just before they were mobile (e.g., printing directions before leaving the house). Finally, we conducted an introductory, mid-study, and final interview with each participant, spaced 7 days apart.

The notion of information need is quite broad and people experience a myriad of needs for information throughout the day. This is further complicated by how well the need can be articulated. Taylor describes four levels of information needs, varying in how well the need can be articulated: visceral, conscious, formalized, and compromised [Tay67]. Rather than focus on attempting to capture every need or rigidly defining the scope of information needs for our participants, we told them during the initial interview that one of the best ways to think about an information need is to frame it in the form of a question. We gave participants several examples and asked them to report on any desire for information regardless of whether they knew how to obtain it. Our objective was to collect needs that the participants thought to be sufficiently important to record.

Logging diary entries under mobile conditions can be especially difficult for participants. Palen proposed using voice whereby participants would dial into a voice mail system and leave a recording of their thoughts instead of writing them down [PS02]. This works well for unstructured responses, but for structured responses a voice menu is required which can quickly become more annoying than useful [YZ06]. Carter and Mankoff looked at using photo, voice and tangible objects as prompts for interviews with a researcher later [CM05]. One of their key

findings was that while these new media were more convenient, recall suffered. The txt 4 l8r system proposed by Brandt, Weiss, and Klemmer is a snippet-based system whereby participants send a small piece of information via their chosen media (paper, voice, SMS, photos) in situ, and maintain a post facto web diary where they discussed their captured data [BWK07]. The in situ capture was intended to provide personal data that would describe the situation sufficiently so the user could be reminded to describe it in detail in the associated web diary. In their user study, SMS was the most popular media. This approach provides some of the ecological validity of diary studies in a way that lowers the overhead associated with making diary entries.

We employed a similar snippet technique to capture diary entries. We used a text messaging scheme so that participants would not have to carry an extra device or paper log to record their diary entries. We asked participants to carry their mobile phones with them at all times during the study, so they could easily use it to compose diary entries. These text messages were sent to a special email address that processed the messages and posted them on a website. We asked participants to construct their text message snippets in such a way that they would be able to subsequently answer a set of diary questions on the web. To help our busy participants remember to send in diary entries, we sent them five text messages per day every three hours during the daytime to remind them to record any new information needs. At the end of the day, participants logged into a website to answer six questions about their snippets:

1. Where were you?
2. What were you doing?
3. What was your information need?
4. I addressed the need (At the time, Later, Not at all).
5. If you attempted to address the need, how did you do so? If you didnt make an attempt, why didnt you?

6. Could you have addressed your need by looking at your personal data (e.g., email, calendar, web browsing history, chat history, or other)

One concern we had was how to incentivize participation without encouraging participants to make up diary entries. We decided to reward participants for visiting the web site each night and completing the interviews, but not for constructing individual diary entries. Participants could earn a maximum of \$80 throughout the study. They were paid \$3/day for participation, with an extra \$10 bonus incentive for 12 out of the 15 days of participation. We compensated the participants \$10 for the in-person introductory and final interviews, and \$5 for the mid-study phone interview.

2.3 Results and Discussion

Our study generated 421 diary entries, with an average of 21.1 entries per person (min:7 max:45 SD:10.5). All our participants filled in the web diary portion for any text message entries they made. During interviews we asked participants for clarification of any unclear entries. In the following sections we characterize the types of information needs people had and discuss the methods and strategies they used to address those needs. We also discuss the different contextual situations that prompted information access.

2.3.1 Taxonomy of Information Needs

We sorted the diary entries into 16 broad need categories based on the participants diaries and their feedback during interviews. Figure 2.1 shows each category, representative examples taken from the diaries, the # of diary entries for each category (and percentage), as well as the number of participants that reported an entry in that need category.

We characterized the largest category of information needs as trivia (18.5%). These needs were often prompted by conversations or location-based artifacts (e.g., a billboard). They are the interesting, seemingly random thoughts that came to

Need Category	Examples From Diary Entries	% of Total Diary Entries
Trivia	"What did Bob Marley die of, and when?"	18.5%
Directions	"Directions to Sammy's Pizza"	13.3%
Point of Interest	"Where is the nearest library or bookstore?"	12.4%
Friend Info	"Where are Sam and Trevor?"	7.6%
Shopping	"How much does the Pantech phone cost on the AT&T website?"	7.1%
Business Hours	"What time does the post office close?"	6.9%
Personal Item	"What is my insurance coverage for cat scans?"	6.4%
Schedule	"Is there an open date on my family calendar?"	5.7%
Phone #	"What is the phone # for weight watchers?"	5.7%
Traffic	"How far does the traffic extend?"	4.5%
Sports/News/Stocks	"Did the Miami Heat have any free agent acquisitions?"	3.8%
Email	"Email update for work"	2.6%
Movie Times	"Are Harry Potter tickets available tonight?"	2.4%
Weather	"What will the weather be like this weekend?"	1.4%
Travel	"Flight status of my Southwest flight"	1.0%
Recipes	"Needed ingredients for hot and sour soup"	0.7%

Figure 2.1: Breakdown of information needs by categories. Examples for each category are from real diary entries. Categories are sorted based on their frequency in our study.

mind for our participants, such as *What did Bob Marley die of, and when?* or *What is the nutritional benefit of eating almonds?* Their intended use in social situations suggests, however, that they are not unimportant. When participants dismissed the need it was often because the conversation in which they were engaged had moved on to another topic, or the need was deemed not important enough within the current context.

The second highest category of need was directions (13.3%). Directions needs are ones in which participants knew where they were trying to go, but needed the address or fastest route to their destination. A related category that was also frequent was point of interest needs (12.4%). These are broader needs in

which a person needed to first find a place of interest, and then get directions to the place. Common instances involve finding a nearby restaurant, coffee shop, gas station, or bookstore. Both categories are typical mobile information needs that can be addressed through in-car GPS/Point of Interest databases. Only one of our participants owned such an in-car system.

The friend info category (7.6%) includes any information about a friend or family member that can only be obtained from that specific person. For example, *Where are Sam and Trevor?* Three people accounted for 22/32 of the reported friend info needs.

Three related categories that are close in frequency are business hours (6.9%), shopping (7.1%), and phone # (5.7%). Although determining business hours can be part of a shopping experience, there are also times when someone might want to know business hours to obtain certain services (e.g., tailor). Shopping frequently involves price check-ing with different stores, and trying to find the right store that sells a certain product. Phone #s are related because they sometimes act as an intermediary need to get business hours, or find out if a product is available. However, phone #'s can also be required to contact others.

We distinguished personal item (6.4%) and schedule (5.7%) as two separate categories because scheduling can involve both a personal and a shared schedule. Personal items are information entries that are uniquely associated with specific individuals, such as a social security number, or personal insurance policy.

Some information needs are either so predictable, or easy to access that they do not register as a need to report. For instance, traffic at specific times along a given route can be quite predictable. Weather can also be quite predictable in areas without large climate changes. Therefore, we may expect that categories such as traffic (4.5%), or weather (1.4%), as well as sports/news/stocks (3.8%), or email (2.6%) might be larger in other contexts.

Movie times (2.4%) and Recipes (0.7%) are often cited as highly desired mobile applications (e.g., Moviefone, Epicurious). Although they occurred with low frequency, they are still important needs. A person does not need movie times and recipes everyday, but when the need arises, its usually important to have the

information. Moreover, 9 people accounted for 10 of the movie time diary entries, and each of 3 different people reported a recipe needs.

We identified travel-related items (1.0%) as a separate category because this usually involves going somewhere away from ones hometown. Travel entries included getting information about flights, hotels, or activities in the given city. Any needs associated with getting from one city to another city were placed in this category. The travel category was low though because only three people went on trips during the two-week period.

2.3.2 Deciding When to Address a Need

When an information need arose, our participants could have addressed the need at the time, later, or not at all. In order to determine when to address their need, they at times appeared to be calculating a complex cost function. Our data provides evidence about the factors involved in making these decisions.

Determining the Cost/Benefit of Addressing a Need

The diary entries and interviews revealed a complex set of issues that participants weighed before deciding when to address a need. These issues can be described along four dimensions: importance (i.e., value or benefit), urgency, cost, and situational context.

Important needs are those that should be addressed, even if addressing them is not required immediately. For example, *How many units does my son have to take to be considered a student for insurance discounts?* In this entry, our participant knew that knowing how many units her son had to take to qualify for an insurance discount would save her money. The need did not have to be addressed immediately, but addressing the need at some point was important.

Urgent needs are usually related to current activities. They are time critical and demand immediate attention. For example, participant 14 got lost on his way to attend a wedding, so his diary entry read, *directions to the church where the wedding is*. Addressing this need was essential to complete his current task, getting to the wedding.

In addition to the importance and urgency of a need, our participants also considered the cost of addressing a need in terms of time and expense. When time is not critical, a person can consider multiple alternative methods to address a need. However, when a person is rushed, a method that could potentially take more time than is available isnt an effective alternative (e.g., using a cumbersome mobile browser). Participant 1, a mobile internet user, mentioned that trying to address a need with her mobile phone when shes late is not an attractive method.

You have to type stuff in to find information which can be a pain at times depending on the circumstance that youre in. Especially when you really need it and youre going some-where. You have to go there and youre late, the last thing you want to is be typing into your phone. You dont want to pull over on the side of the road to type things in.
(Participant 1)

A particular method can also have a number of monetary costs, such as calling 411 (which costs \$1.50 per/call), or accessing the mobile internet without an unlimited data plan (i.e., paying per kilobyte). One of our participants who had an internet-enabled phone rarely used it because he paid based on usage. However, in one situation he was trying to find an interesting place to visit in the nearby area with his family. After many unsuccessful attempts, he became frustrated enough to pay per kilobyte to find a point of interest.

We were actually at a different location than we were try-ing to get to. We had arrived at a random museum. It was getting more and more frustrating. I pay per kilobyte, but I figured it would be worth it to pay to get this information right now. On my way back to the car I was anxiously trying to find something [online] but nothing came up.
(Participant 18)

The situational context contributes to the cost of addressing a need. There are times when addressing a need may not be socially appropriate, such as during an important meeting.

One time I was stuck in a meeting. Since it was a high profile meeting and I couldnt open my phone and text, or laptop to email. I was wondering if my team was going to be able to do this demo. Im thinking, I need to know about this demo, but this really is an information need.
(Participant 16)

People often multitask to better manage the cost/benefit tradeoffs in satisfying an urgent information need. During multitasking activities, people divide attention between the current task and trying to address their information need, creating a challenging situational context. Driving was one of the more frequent multitasking activities cited by our participants. They all acknowledged that interacting with a phone while driving is dangerous. However, 18/20 participants admitted to doing it occasionally, employing a variety of strategies to balance safety and the urgency of their information need. For example, participants reported trying to time their phone interaction for red lights or open stretches where they felt that being distracted would be less dangerous.

Some entries I did while I was driving, and others I did at a stop sign or stop light. (Participant 1)

I do text while I am driving. Interacting with my phone is dangerous. I try to do it while Im at a red light or stopped somewhere. (Participant 10)

I will wait for stop lights, or if Im on an empty straight-away then Ill text or place a call. Or Ill write it down on a piece of paper at a stop light and then after driving Ill text it. Ill grab a scrap of paper and write it. (Participant 19)

Other participants found their needs too pressing to put off until later, and would instead use more dangerous tactics to address their information needs while driving. These participants were trying to locate their friends, obtain directions from an online map, or check their email.

I hold the phone (full keyboard) and then hold the steering wheel and text while driving. If there is heavy traffic Ill put it aside and then pick it up afterwards. I try to do my interaction in open stretches. I have to use my knees and elbows to hold the steering wheel while driving, but I try not to do that too often. (Participant 17)

More than 90% of my needs were in my commute going back and forth from work and occasionally on the weekends I try to time my texts to send them in at red lights or sitting in traffic. But I often drive with one hand and text with another. Its asking for trouble, doing too much phone interaction while driving. Its a bad habit to get into, it only takes a little lapse of concentration to get into an accident. (Participant 15)

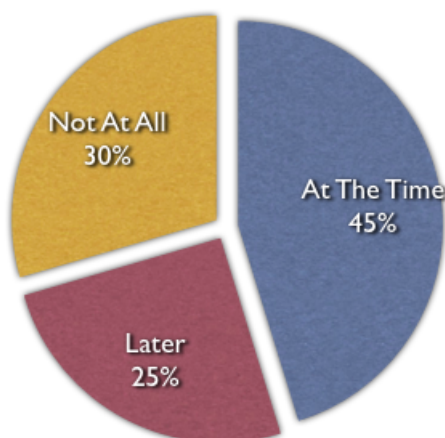


Figure 2.2: Pie chart of diary entry percentages for when participants addressed their information need.

I try not to [enter text while driving], but I found that I've done it a few times. I almost creamed somebody doing it and then I knew that this wasn't going to work. I don't think it's a good idea. I've done it, but I don't think it's good." (Participant 20)

One way we glimpsed how participants weighted various factors in addressing an information need was by having them indicate whether they addressed their need at the time they occurred, later, or not at all. If they addressed a need at the time, we asked how they addressed it. If they did not address it at the time, we asked why they did not do so. Figure 2.2 shows a breakdown of when participants addressed their information needs. 45% of the diary entries showed that people addressed their needs at the time when the need arose. However, 55% of the entries were marked as either being addressed later or not at all. Lack of internet access was not the dominant factor, as our five mobile internet users addressed their needs at the time only 58% of the time. Below, we describe the methods participants used to address needs, and reasons they did not address needs at the time they arose.

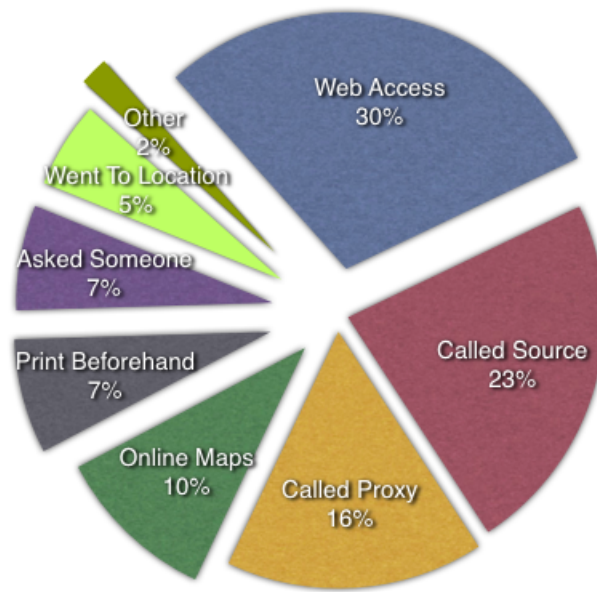


Figure 2.3: Breakdown of how needs were addressed at the time they arose across all diary entries.

How Needs Were Addressed At The Time They Arose

If needs were addressed at the time they arose, participants indicated the method they used to address the need. We categorized these methods into the following (Figure 2.3):

1. Asked Someone - Found a person and asked them face to face
2. Called a Proxy - Called a person/service that functioned as an indirect way to address their information need (e.g., call someone to access the internet, or call 411 to get a phone number).
3. Called Source - Called a service that had the desired information (e.g., call an airline company to find out flight status).
4. Web Browser - Accessed the web directly from a mobile device, or sought out web access at some place.
5. Online Maps - Used a map application such as Google Maps on a mobile phone

6. Went to Location - Went to a physical location to address the information need (e.g., driving to the post office to see if it was open)
7. Print Beforehand - Print the information before going mobile (e.g., print directions to the next location)
8. Other Means - Listening to the radio, looking at ones phone logs, or consulting the yellow pages.

Participants employed internet-based access 40% of the time when addressing information needs. We split internet-based access into two categories to differentiate between accessing a website and using a specific mobile application (30% web access, 10% online maps). When we isolate diary entries from our five mobile internet users, 73% of their needs were addressed either through web access or an online maps application. Participants without mobile internet access employed a variety of strategies to gain internet access in order to address their needs. Sometimes a person would seek internet access at a public terminal.

Since the internet has become a main source of information access today, it is not a surprise to see people rely on it, even when mobile. Still, looking up information while mobile can be difficult and cumbersome. Some of our participants printed the information they would need before they left their home or work (7% of diary entries). This was commonly directions or phone numbers associated with upcoming tasks. One participant printed 22 maps to make sure he had directions for planned destinations for an upcoming trip.

Another way that participants addressed their needs was by either calling the source of the information (23%) or calling a proxy (16%) to address the information need on their behalf. Calling the source includes calling businesses for hours of operation/level of busyness (i.e., how long is the line?), Moviefone for movie times, airlines for flight status, or stock brokers for current stock prices. Participants would often program these frequently dialed numbers as shortcuts into their phone. An interesting alternative method was to call a proxy to address their need. Instead of attempting to call a service, participants would call a general service such as 411, or a friend/family member to access the internet for them.

Some even gave out their email passwords so that the other person could retrieve information in their personal data. This proxy method was more common amongst our younger participants who usually had several people in mind who would likely be in front of a computer. Other participants felt that calling someone to look up something on the internet would be intrusive and bothersome. As a result, they would rather attempt to address the need in some other way, or not address it at the time.

As we mentioned above, circumstances can influence the method a person uses to address their need. When time is limited, a more convenient method may be chosen over other methods. Two ways that people did this were to ask someone nearby (7%), or go to a place (5%) that was the source of their information need. These two methods were chosen because at the time they were the quickest and most convenient methods available. The types of needs addressed in this way were typically direction or business hour needs. Some participants addressed their need through means (2%) other than those described above. Examples include using a yellow pages book, or listening to the radio.

Why Needs Addressed Later

In situations where participants determined that it would be difficult to address a need, they might choose to address it at a later point in time. The three primary reasons (Figure 2.4) for addressing a need later were lack of internet access, being currently involved in an activity (biking/driving, busy with task, in a meeting), or realizing that the need would naturally be addressed later. We discuss each in turn.

Since people rely on having internet access when not mobile, they naturally see it as a way to access information when they are mobile. Whether using search engines or email, having internet access gives people ways to address many information needs. The top reason participants gave for deciding to address a need later, was the lack of internet access at the time (32%, a few mobile internet users cited not having internet access as a reason only because they left their phone somewhere else). In our interviews, 13/15 (87%) of our non-mobile internet

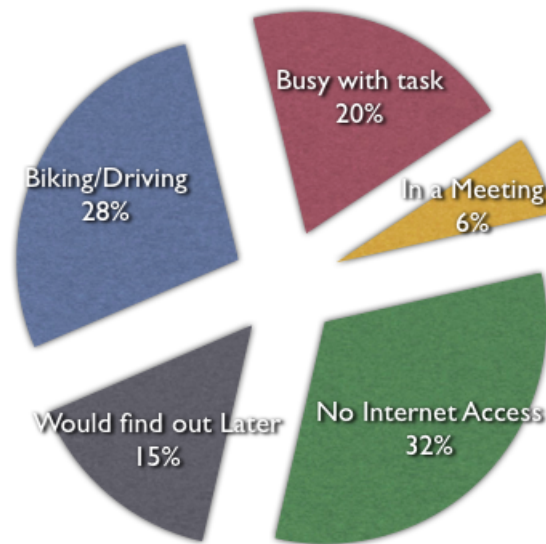


Figure 2.4: Reasons participants addressed their needs later. Percentages correspond to number of diary entries

participants felt that with mobile internet access they would be able to address a majority of their needs. However, it is important to note that the diary entries show that not having internet access only accounted for a third of the reasons a need was addressed later. This hints at there being other reasons needs are addressed later. As we describe below, many are related to the context of the current situation and the task at hand.

As depicted in Figure 2.4, more than half of the diary entries indicated that the reason participants addressed their needs later was because they were currently involved with an inhibiting activity. These activities included biking/driving (28%), busy with a task (20%), and in a meeting (6%). Both biking and driving are activities that require ones full attention. Despite this some participants reported addressing a need while driving. Others chose not to risk their safety and waited until later to address their need. The busy-with-task category reflects that important needs can arise while a person is doing something else, but the need is not urgent enough to context switch in order to address the need. Lastly, as described earlier, there are social situations where it may not be appropriate to address a need (e.g., a high profile meeting).

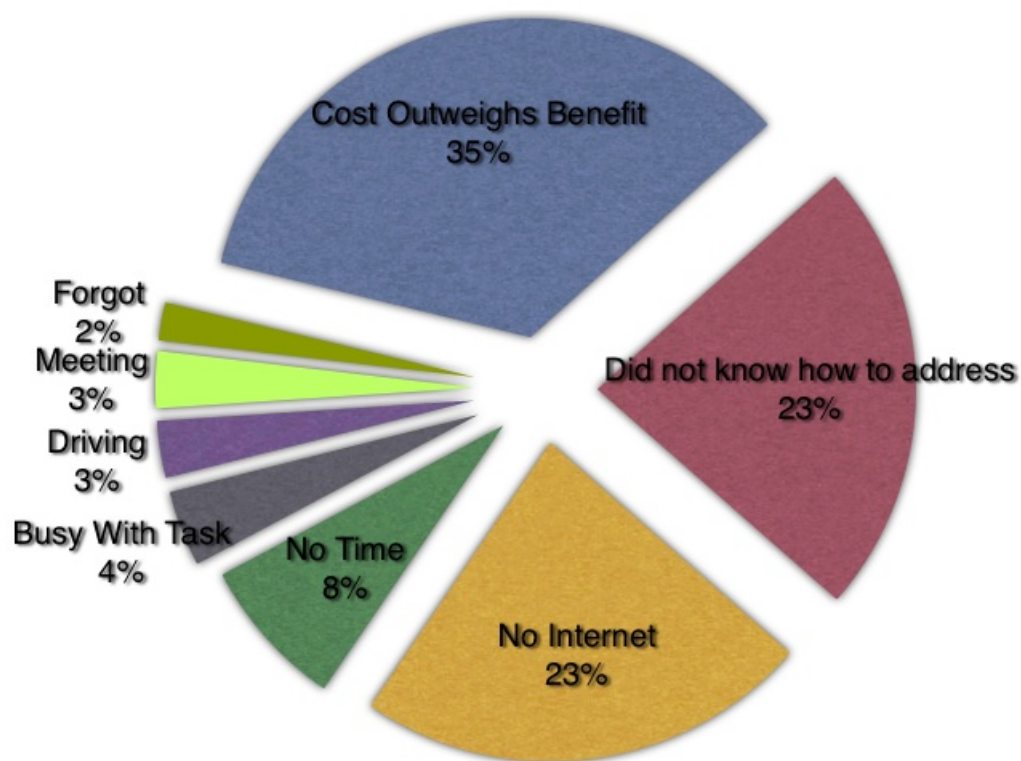


Figure 2.5: Reasons why needs were not addressed. Percentage of entries for each category.

The last category of reasons participants gave for addressing their need later was they would find out later (15%). This category involved instances in which the person knew that in the near future, another activity would address the information need. For example, one participant wanted to know the score for his favorite baseball team's game last night but was going to go see a friend soon. Instead of attempting to look up the score, he knew that his friend could tell him when they met up so he waited until then. Participants used knowledge about future situations as one factor in determining whether a need should be addressed now or later.

Why Needs Were Not Addressed At All

30% of all diary entries indicated that the participant did not address the need at all (Figure 2.2). They did not attempt to address them at the time, nor did they go back to address them later. Figure 2.5 summarizes reports of why needs were not addressed at all. The main reason that participants gave was that the need was not important (35%). The general character of their responses was that although the information would have been nice to know when a need arose, it was not worthwhile to address it at the time. Over 40% of the not important entries we categorized as trivia needs. Other not important needs included business hours, shopping, weather, and points of interest.

The second highest reason that needs were not addressed was the lack of internet access (23%). Unlike the needs that were classified as not important, participants reported that these needs would have been addressed at the time if internet access had been available. However, they appear to have become less important as time passed, evidenced by them not being addressed later.

Another reason participants cited for not addressing needs was that they did not know how to address them. Even if they had available resources, they were unsure how they would have gotten the information in an easy way. One example diary entry was, *Does the airport parking take credit card? I need to know whether to get cash or not.*

The final category of reasons needs were not addressed was that participants were engaged in other activities that pro-hibited them from addressing their needs. These categories included: in a meeting (3%), forgot (2%), driving (3%), busy with a task (4%), and no time (8%). The no time category includes instances when participants wanted to address the need, but they were late for their next activity. For example, participant 1 was on her way to see a counselor and had trouble finding the address. She reported that if more time had been available she would have addressed her need through Google Maps on her phone. Instead, she relied on her knowledge of the area to find her destination.

From a design perspective, needs that are left unaddressed provide opportunities for improving future information access facilities. Satisfying even informa-

tion needs that would be nice to know at the time but can be done without could possibly be enriching. An example of this came from Participant 7, who wrote, *I wanted to watch Harry Potter tonight with my brother so I wanted to know if tickets would be available.* Unfortunately, since she was driving at the time, the need was left unaddressed and she did not end up watching the movie. If she had found out movie times, she may have been able to get tickets ahead of time, and watch the movie with her brother.

2.3.3 How Does Mobile Internet Change Behavior

Throughout our discussion, we have not made a strong distinction between participants who had mobile internet access and those who did not. As mentioned earlier, 13 out of the 15 participants (87%) who did not have mobile internet access felt that with it, they would have been able to address a majority of their information needs. However, all 5 participants who had mobile internet access felt that it was not adequate for addressing their information needs. We noticed two characteristics of mobile internet usage that demonstrate how having unlimited mobile internet access can change behavior. The responses from the mobile internet users also capture how it was inadequate for addressing all their information needs, due often to their situational context and current activities.

First, we found that internet users appreciated being able to address their needs on their own without having to call someone for assistance. Several of our participants without mobile internet would call a friend to lookup information on the internet. However, the mobile internet participants would open their web browser, or mapping application to obtain the information directly.

I find it more convenient to search it myself. Explaining on the phone some people just dont get it when I explain something. (Participant 3)

Second, all 5 of our mobile internet users were quite savvy with their phones, for example employing shortcut tech-niques to obtain information quickly. Participant 1, who uses a Blackberry Pearl (pseudo-qwerty keypad) said:

I have shortcuts, so its almost like a qwerty keyboard. I usually type really fast. Restaurant took about 30 seconds. Its a traditional website

that you have to scroll down to see wherever you need to go. (Participant 1)

This participant had several shortcuts to common web resources that she uses to obtain information. These shortcuts help speed up her workflow, but mainly tended to work for simple information queries. In addition to having bookmark shortcuts, one participant wanted to avoid launching a web browser and would prefer a more integrated phone-based application. He said,


I wish my mobile phone had a built in Superpages, so I dont have to go to the web, launch Google, and type something in to an obscure system. I wish Superpages was integrated into the phone. Similar to how Google maps is inte-grated into my phone. (Participant 3)

Despite the mobile internet being a useful tool, all mobile internet participants alluded to it not being sufficient. Either the interaction was too cumbersome or their circumstances prohibited them from using the internet due to the extensive interaction required. Two participants described their frustration with accessing the mobile internet, one due to limited time, the other because of the cumbersome interaction.

You have to type stuff in to find information which can be a pain depending on the circumstance that youre in. Especially when you really need it and youre going somewhere. You have to go there and youre late, the last thing you want to do is be typing into your phone. (Participant 1)

The website was too much information. I just wanted to know the park hours and I got the whole water department. I said forget it. Google search turned up a few hits. That was the #1 hit. The top few were related to that same web-site. I figured that website was going to give me the information I needed, so I browsed for a while. Too much information so I gave up. I gave up after not too long. I spent maybe 3-4 minutes. (Participant 15)

The mobile internet has certainly changed the way people address their information needs. As our data indicates there are also many opportunities for improvement. However, even with better forms of interaction and increased bandwidth, users still confront issues of limited attentional resources as well as challenging contextual and situational factors.



(a) Pie chart showing percentage of (b) Percentage of different contextual factors that diary entries that were prompted by prompted information needs contextual factors

Figure 2.6: How context influenced diary entries.

2.3.4 What Prompts Information Needs?

Given that our participants were unable to address information needs 55% of the time (Figure 2.2), and frequently experienced frustration, we wanted to explore how exploiting context might alleviate some of these problems. One participant said, *In my own life, it seems so much of my needs are specific to the context.* Context-aware computing promises improved information access and more opportunistic information delivery. Mobile context-aware applications have typically been designed for one specific purpose, such as a friend finder, or shopping assistant. Since context can heavily influence when needs arise and how a person might address their needs, we examined diary entries to better understand the context that prompted the need in each situation. Our analysis is based on entries in which participants either indicated what prompted the need in their entry, or provided clarification in the interviews.

Participants indicated that 72% of their reported information needs were somehow prompted by a contextual factor (Figure 2.6(a)). The contextual prompting can be classified in four broad categories: Activity, Location, Time, and Con-

versation (Figure 2.6(b)). Activities reflect what the person was doing at the time. Location is the place where the person was at and includes any additional artifacts at that specific location. Time is the time when the need arose, and conversation is any phone or in-person conversation the participant was involved in at the time. Some diary entries were related to multiple aspects of context, such as having a conversation with someone about artifacts at the current location.

Mobile context research in the past decade has focused heavily on location acquisition [LCC⁺05]. With location being the most widely cited prompt for mobile information needs (34.6%), the possibility of applying the results of these re-search efforts is encouraging. Recent projects have begun to deploy location-based systems on mobile phones, making the technology widely available for use. There are several additional areas where context might be useful. Activity recognition can help provide a person with information related to their current task. For example, real-time traffic information, nearest gas stations, and directions could appear while driving. Time is a heavily used form of context that helps people plan their day-to-day activities. A person's calendar can help determine what types of information needs a person may have. The last form of context that prompted many information needs was conversations. We separated conversations from activities because a significant number of diary entries indicated that a participant's information need arose because of something another person mentioned.

2.3.5 Where Does The Information Come From?

Regardless of whether a need was addressed or not, our participants indicated whether they could have found the information by accessing a public, personal, or physical data source. A public data source can be accessed by anyone (e.g., the web). A personal data source is accessible only by that person (e.g., email, web history). A physical data source can only be accessed through physical methods and not electronically. Understanding which data sources are frequently helpful for addressing information needs can guide the development of further mobile technology.

Public data was the most commonly cited source (58%) for addressing infor-

mation needs. This is congruent with how our participants associated addressing their information needs with internet access. While public data provides a large sphere of helpful information, personal data was also a significant source for information access. Many of these needs were solely personal (24%, e.g., access to calendar), however a common strategy was to first pull information from a personal data source (14%, e.g., an address from email or contacts) and then use that as a pointer into or re-trieval key on a public resource (e.g., using the address to pull up a map). Although such use of personal information is vital to completing the task, it also entails a two-step retrieval process, often complicating the task.

Interestingly, 4% of the information needs indicated that they could only be addressed by visiting a place itself, at least using prevailing technologies (e.g., *Is the line long at Rubios*).

2.4 Design Implications

Many needs in our study were addressed by awkward methods or at a later time due to limited attentional resources being available or other situational factors. Based on these observations, we offer several design suggestions for ways future mobile technology might better address mobile information needs.

First, mobile technology should take into account a persons current task and enable ways to address needs at later more convenient times. More than half of our diary entries indicated that people were unable to address their need at the time it arose due to their current activity. Mobile technology that is sensitive to this could help people address their needs when situations change and resources are more freely available. Few current mobile applications are sensitive to use in multitasking situations or provide mechanisms to conveniently record needs when they arise in ways that support them being satisfied in the future.

Second, a persons context significantly influences their information needs. 72% of diary entries involved information needs that were triggered by context. As we discussed earlier, the contexts included activity, time, location, and conversations. Designing technology that considers these aspects of context could aid

in providing people the information they want at the right time and in a form appropriate to the current context.

Finally, satisfying 38% (28% solely personal, 14% indirectly to a public resource) of our participants' information needs involved access to their personal data, obtained either through web-based services or their personal devices. The personal data source often acted as a pointer to retrieve a public resource. This suggests that better access to these personal data sources as well as easy transitions between personal/public sources is an important opportunity to explore. Current mobile technology requires a person to first access a personal source in one step, and then another step to access the public resource. Seamless connections between the two could ease the burden when addressing their information needs.

These three dimensions help specify a design space for future tools to facilitate mobile information access. Our data suggest that future systems should take into account a person's context as well as their personal data stored across multiple devices to better service their mobile information needs. Such systems could be similar to recommender systems, but we envision a broader, more personalized system that would reduce the burden of mobile interaction and be sensitive to users' contexts during periods of limited attention (e.g., driving).

Designing such a system is challenging because context-based information systems have the potential to become nuisances if they misidentify the context. People are also concerned that context-based systems might enable unwanted personalized advertising and other forms of spam, and myriad privacy issues must be confronted (e.g., who has access to the system, where is data stored and how is it protected). To better support mobile information access these and many additional challenging issues need to be addressed.

We asked our participants about how they would feel about a tool that could predict their information needs and provide appropriate information at the "right" time. 17 participants responded positively to the idea.

That would be really cool. Especially you wouldn't have to type it in and look it up when you are on the road.

That would be interesting. If it relied on some sort of system where you used time and GPS so it could tell where you are. If you are at

the zoo, you could check out such and such. I dont think that much of a hindrance. It might actually be pretty cool... I also dont think itd be on the mark all the time

I can definitely see some advantages to that. If I look at the maps on my phone like google maps or traffic then I have to get to the map and wait for it to load. Or I have to dial a phone call can be a hassle. So I can definitely see an advantage to that.

I would be impressed by such a system. Definitely. I guess if there was software that knew your daily routine and gave you traffic information on your way home. Yea, I would have found such a system useful. A lot of the information needs are random, so you can't predict them. For the regular ones though yea it would have been helpful.

Our other 3 participants were more hesitant, mainly concerned with privacy and control issues. Although they use services that currently search through their private data, the thought of additional loss of control made them uncomfortable.

On the one hand I think its convenient, but on the other hand I see it as disconcerting because my technology is becoming too smart. I worry that there is a possibility for over policing of my needs. I dont want to be manipulated.

I would find it annoying. I guess I would find it violating to my privacy. I dont even like it when Gmail searches your email account and things pop up on the side.

2.5 Conclusion

In our two-week diary study we found that people on the go have a panoply of information needs, as well as challenging constraints in addressing them. We also found that people employ an ingenious variety of methods to satisfy information needs in a timely and situationally appropriate manner. Many information needs were postponed or unaddressed because of attentional costs and contextual factors. An important note is that lack of mobile internet access was not the only inhibitor by far. Even those with mobile internet access employed a wide variety of methods, and still they only addressed their needs 58% of the time, meaning that 42% of the

time they postponed or never addressed their needs. A devices sensitivity to the task at hand, situational context, and the links between personal and public data holds promise to ease mobile information access by providing the right information, at the right time, and in the right form for the current context. In addition, there is promise in enabling convenient capture of needs at the time they arise so that they might be satisfied at more appropriate times in the future.

2.6 Acknowledgements

This chapter, in part, is a reprint of the material as it appears in A Diary Study of Mobile Information Needs. Timothy Sohn, Kevin A. Li, William G. Griswold, and James D. Hollan. In CHI 2008: Proceedings of the Twenty-Sixth Conference on Human Factors in Computing Systems. pages 433-442. ACM press, 2008. The dissertation author was the primary investigator and author of this paper.

Chapter 3

GSM-Based Localization Technology for Mobile Users

The results from our mobile information need diary study showed that context can potentially play a significant role in developing appropriate applications for mobile users. Location is one of the most prominent forms of context that has been explored in the research community. The following sections will describe background work for enabling location sensing on mobile phones using commodity hardware. An effective, low-barrier to entry method for determining location is essential for enabling the applications and tools we need in order to support the needs of mobile users. We will look at GSM-based methods for localization that have helped shape the landscape of building context-based applications on mobile phones.¹

3.1 Introduction

Determining the position of a device is a key enabling factor for ubiquitous computing applications. Location can be used in a variety of ways to provide relevant information to a mobile user. Several canonical location-aware applications include tour guides [AAH⁺97], location-based annotation [GSB⁺04], and near-me

¹Much of this work is rooted in the Place Lab project [LCC⁺05] due to my involvement in building, designing, and evaluating the initial GSM localization efforts.

services [KH04a]. These applications provide services that can take advantage of digital history and reveal unseen artifacts in the environment. Location is a powerful mechanism to enhance many of the seemingly normal places that we travel through. In addition, location-aware applications can alleviate the burdens of mobile users as we will see in Chapter 5.

Dozens of location-sensing systems have been built to support applications by using various sensing technologies including infrared proximity [WHVFG92], ultrasonic time-of-flight [PCB00], optical vision [TSM⁺07], and radio signal strength [LCC⁺05]. However, location-aware computing applications continue to remain in the research lab without major deployment. Systems that use many of these technologies do not work in places where people spend most of their time; coverage in current systems is limited to outdoor environments (e.g., GPS) or certain indoor buildings with installed sensing infrastructure (e.g., the Cricket system). The value to users of dynamic location-aware applications, such as location-aware instant messaging, greatly diminishes when they only work sporadically. Even the most widely deployed location-sensing infrastructure, GPS, was found to have just 4.5% availability during everyday life activities [LCC⁺05].

In order to provide the most value to mobile users, we need a location system that is ubiquitously available, and can work without extra infrastructure. Mobile phones are a platform that meet this requirement. Mobile phones are already deployed worldwide an order of magnitude greater than any other platform. These devices are truly mobile, in that they can be used "on the move" and not just carried from place to place. Their form factor and battery life allow users to continually carry an always-on device, thus applications can be deployed that are usable twenty-four hours a day. Since users already carry these devices, no extra deployment of hardware is needed. Moreover, phones already support radio technologies that can be used for location sensing, eliminating the need for extra infrastructure deployment.

One of these supportive radio technologies is GSM, which as of today has 3 billion connections worldwide². This means existing devices can take advantage

²http://www.gsmworld.com/news/press_2005/press05_21.shtml

of powerful location-sensing capabilities to enable a new set of applications. GSM also provides nearly 100% coverage throughout our daily lives, making applications always accessible. The power of an always-on location system, paired with a pervasive device like the mobile phone are the key enablers for the ubiquitous computing vision.

3.2 Location Sensing

Location sensing on mobile phones can be accomplished in two ways, with different tradeoffs in coverage and accuracy. Coverage is defined by the amount of time that a location system is available. Accuracy is the comparison between a location estimate and a baseline position (often latitude longitude coordinates).

3.2.1 Infrastructure-Based Sensing

Whenever a mobile phone is on, its service provider is able to determine its location. This is necessary for telephony, as the infrastructure must be able to route an incoming call to the phone. Since network operators know the exact layout of their cell towers, they can use this to provide a coarse location estimate (in the kilometer-scale) based on the phone's communication. Recently, US and EU government initiatives (E911³/E112⁴) require that handset positioning be accurate to within 300 meters, specifically for emergency calls. This has forced service providers to upgrade their infrastructure using a variety of techniques such as U-TDOA (Uplink Time Difference of Arrival), which makes use of radio timing information to localize a mobile phone with greater accuracy. As a side effect of these enforced upgrades, providers are rushing to offer premium (i.e. revenue-generating) location-based services, such as allowing users to ask for nearby restaurants, gas stations, or cinemas showing a certain film.

³<http://www.fcc.gov/911>

⁴http://ec.europa.eu.int/comm/environment/civil/prote/112/112_en.html

3.2.2 Beacon-Based Approach to Location

Phone-based location sensing relies solely on the mobile device for determining location. In the absence of a built-in GPS, phones can only determine their location by sensing the radio environment they find themselves in, namely using their GSM capabilities. One disadvantage to this approach is that a mapping of radio sources (beacons) to locations must be available to allow phones to estimate their location. Based on the visible beacons in the environment, phones could use this mapping data to estimate their location. While GSM providers know their own cell tower locations, this is often regarded as commercially sensitive information and not shared freely. Therefore, mapping data is not, in general, available for GSM beacons either. However, one method for gathering map data is to enable wardriving, whereby a small community of users equipped with GPS devices gathers radio beacon locations and makes this available for other users who may lack GPS units. There are numerous public databases available today with locations for 802.11 and GSM beacons (e.g., Wigle.net).

3.2.3 Coverage and Accuracy

Good coverage is exhibited by both sensing methods due to the prevalence of GSM networks throughout the world. One area of uncertainty though is whether infrastructure based sensing support will work across provider boundaries. A mobile user may travel to areas where their home service may not be available (e.g., internationally). In these situations there is no guarantee that a provider-based approach would be able to offer location services. Phone-based sensing avoids these issues, but its coverage is limited by the need for construction of a mapping database. The beacon mapping is independent of a particular provider or country, thus providing the potential for transparent global coverage.

Looking at accuracy, infrastructure-based sensing has the advantage of enjoying cooperation from an array of cell towers to localize the mobile phone using techniques such as U-TDOA. Several mobile phone operators already provide navigation capabilities on their devices with a high degree of accuracy. Phone-based sensing using only GSM cannot rely on such infrastructure support, and will there-

fore have an accuracy related to the density of GSM cells. The size of cells range from hundreds of meters to kilometers

3.3 Evaluating GSM Beacon-Based Localization

A beacon-based approach to location opens the possibilities for research in the mobile space. Although there are coverage and accuracy limitations, being able to gain access to location data on many heterogeneous devices enables true ubiquitous deployment and evaluation of mobile applications. Several commercially available location systems use a beacon-based including Google Maps Mobile⁵ and Skyhook Wireless⁶. This section gives an overview of how the beacon-based approach works as seen through the Place Lab location system. We also highlight results for evaluating this method at a metropolitan scale.

3.3.1 The Place Lab Model

Place Lab is a beacon-based approach to localization [LCC⁺05]. As beacons are sensed in the environment, a client device checks a beacon database to see if the position of any beacons are known. The position of known beacons are used to infer the device's position. Figure 3.1 shows an overview of this approach. Every GSM tower has a unique ID that can be formed by concatenating the tower's country code (MCC), network code (MNC), area ID, and cell ID. These unique IDs are used to construct and refer to beacons in the beacon database. The beacon database is populated by using a wardriving method, or a published database of towers. The beacon database is agnostic to any client platform, making it usable by many heterogeneous devices.

This model is advantageous because it enables innovation at multiple levels.

1. Sensing Innovation. As new technologies are developed they can be easily added into the Place Lab model if they can be identified by a unique ID.

⁵<http://maps.google.com>

⁶<http://www.skyhookwireless.com>



Figure 3.1: Overview of the Place Lab model. Mobile phones sense GSM towers (beacons) with unique ids. The phone uses the id to reference a beacon database. If the beacon is known, the phone estimates its position using the position of the beacons.

Examples of other types of beacons include 802.11, Bluetooth, and Ultra-wideband technologies. As mobile phones become equipped with other types of sensors (e.g., Bluetooth, 802.11, WiMax), these technologies can be used to provide them greater accuracy and coverage based on the advantages of each technology. For instance, the range of Bluetooth is much smaller than GSM, which means more accurate positioning. Bluetooth lacks coverage though, so fusing the two technologies brings different advantages to the system.

2. Algorithmic Innovation. There are a variety of algorithms that can be used to determine a device's position. The most simple algorithm uses a centroid of all the beacons' positions' it recognizes. The Place Lab model allows for positioning algorithm research to happen without additional infrastructure. Research in positioning algorithms is important for delivering more accurate estimates. Several types of positioning algorithms include fingerprinting [OVLdL05] and particle filters [HB04].
3. Application Innovation. Using a location toolkit allows application developers to focus on the application without having to concentrate on the underlying location-sensing capabilities. The interfaces exposed by the Place Lab

toolkit are easily accessible to all types of applications. Place Lab can be statically linked into an application, accessed through an NMEA interface, or the JSR-179 location API.

These advantages make Place Lab a desirable toolkit to explore location-based computing to innovate in the mobile space. For our purposes, we want to look at how Place Lab can be used for developing mobile applications.

3.3.2 Beacon-Based Localization at a Metropolitan Scale

In 2005, several researchers analyzed the effectiveness of a beacon-based approach to location within Seattle, Washington, a large metropolitan area [CSC⁺06]. Until that point, many evaluations had been done in small campus settings, or within controlled areas. In order to study the effect of GSM beacon-based conditioning, the team collected a 208-hour, 4350 km driving trace of three major GSM networks covering the Seattle metropolitan area: AT&T, Cingular, and T-Mobile.⁷

Figure 3.2 shows a map of the wardriven area. The red street highlights show the roads driven during the data collection. The total area spanned 18.6Km x 25.0Km. Using the red highlighted as training data, Chen et. al. evaluated different positioning algorithms with a test data set in the three areas, two residential and one downtown area). Their analysis showed that GSM devices can achieve a median error of 94-196 meters using cells from a single provider. Although this may not be as accurate as an infrastructure-based approach, there are still a large class of applications for which the accuracy is sufficient. Even coarse-grained location sensing can be effective for building applications that address the needs of mobile users.

Furthermore, an entire mapping of every single street in a metropolitan is not necessary to achieve accurate results. Calibration density tests showed that only using 30% of the dataset was sufficient to provide comparable positioning accuracy - suggesting that 60 hours of driving (instead of 208 hours) could cover a similar area to the size of Seattle.

⁷Cingular recently acquired AT&T Wireless, but many AT&T-identified towers still existed at the time of this study




Figure 3.2: Red street highlights show the roads driven during a GSM trace collection spanning 18km x 25km region of the Seattle metropolitan area. Figure taken from [CSC⁺06]

3.4 Conclusion

Location is a powerful piece of context that can be used to change an application's behavior. We have provided a background and motivation for performing location-sensing on mobile phones with commodity devices available today. Typical location systems have a high-barrier to entry, but by using a beacon-based approach many researchers can use an effective location system without much cost. With a strategic beacon map of an area, GSM-based localization can achieve a relatively accurate position estimate. In Chapter 5 we will explore how a location-based application is used in practice, and find that location accuracy is not the most important component to assisting mobile user.

Chapter 4

GSM-Based Mobility Detection

A second form of technological support useful to mobile users is detecting high-level mobility patterns. Knowing a person’s mobile state can be used to support social-mobile applications, or for offering information at an appropriate time (e.g., not driving). The methods described in this chapter rely on the same principle about GSM as discussed in Chapter 3. Instead of relying on coordinates calculated from GSM traces using a beacon database, we demonstrate that mobility patterns can be detected without position information.¹

4.1 Introduction

This chapter introduces a technique for detecting a users coarse-grained mobility using commodity cell phones. Pervasive computing applications have long made use of technologies for inferring a users physical activities. Both coarse and fine-grained location systems have been used to perform location-driven activity inference [HKH05], [WJH97]. Smart spaces containing cameras, RFID tags, and the like, have been used to detect fine-grained user activities [KOA⁺99], [ILB⁺05], [PFP⁺04]. Unfortunately, the cost, complexity and maintenance overhead of such activity inference systems have hampered their main-stream adoption. Recent

¹Note that this chapter is a reprint with minor changes of *Mobility Detection Using Everyday GSM Traces*, a paper co-authored by Timothy Sohn, Alex Varshavsky, Anthony LaMarca, Mike Chen, Tanzeem Choudhury, Ian Smith, Sunny Consolvo, Jeffrey Hightower, William Griswold, and Eyal de Lara.

work has attempted to address some of these issues. An example is the belt-worn cluster of sensors developed by Lester et al. that can identify several physical activities including detecting subtle distinctions such as walking on level ground versus up stairs [LCK⁺05]. However, challenges in form factor and power usage still remain.

Fortunately, many applications do not require the detail and accuracy of the systems cited above. As an example, consider the domain of eldercare as depicted by Computer-Supported Coordinated Care (CSCC) [CRS04a]. CSCC describes the network of people who help an elder age in place and seeks to improve the quality of her care while reducing the burden of providing care on the members of her network. Many of the elders activities that are meaningful for her network members to know about involve high-level information about the elder such as whether or not she was up and about today, or if she had a sedentary day around the house.

The most widely deployed and used mobile computing device today is the mobile phone, which presents an obvious opportunity for high-level activity recognition such as that needed by CSCC applications. This paper investigates how without GPS, a commodity GSM phone could infer such high-level information without placing the types of additional burdens on the user that are typical of more heavyweight systems. Previous research has used GPS to detect the modes of transportation for an individual [PLFK03]. However, GPS positioning is available as little as 5% of a typical persons day [LCC⁺05], providing much lower coverage than we require. In contrast, cellular coverage is available throughout most, if not all, of a persons day and does not require line of sight to work [LCC⁺05]. Therefore, using a GSM sensor to detect high-level activities allows the sensing system to always be available, and allows users to continue to carry their mobile phones in their pockets, bags, etc.

The contribution of this chapter is that, with unmodified GSM mobile phones and without relying on users to modify their behavior, we can recognize several high-level activities. Using statistical classification and boosting techniques, we successfully distinguished if a person is walking, driving, or remaining at one

place with 85% accuracy. Additionally, we were able to build a GSM-based step count predictor that provides a reasonable approximation of the users daily step count compared to several commercial pedometers. Our methods were tested with real-world data from three data collectors using the two major GSM networks in the United States (T-Mobile and Cingular). The data collectors gathered GSM network trace data over a period of one month, logging a total of 249 walking events and 171 driving events. Our methods show that GSM-based sensing from commodity devices may provide enough activity information for some applications, without the overhead of requiring additional sensors.

The remainder of this chapter is organized as follows. Section 2 describes our algorithms to infer mobile activities and daily step counts. Our data collection, metrics, and evaluation results are presented in Section 3. Section 4 describes several application domains that could benefit from our mobility detection technique. Section 5 outlines related work, and we conclude in Section 6.

4.2 Mobility Detection With GSM

In this section, we offer a brief overview of the Global System for Mobile Communication (GSM) and describe algorithms that use traces of GSM signals to infer modes of mobility and to estimate daily step-count.

4.2.1 Global System for Mobile Communication (GSM)

GSM is the most widespread cellular telephony standard in the world, with deployments in more than 200 countries. As of September 2005, the GSM family of technologies has 1.5 billion subscribers and 78% of the world market². A GSM base station is typically equipped with a number of directional antennas that define sectors of coverage, or cells. Each cell is allocated a number of physical channels based on the expected traffic load and the operators requirements. Typically, the channels are allocated in a way that both increases coverage and reduces interference between cells.

²http://www.gsmworld.com/news/press_2005/press05_21.shtml

We wrote a custom application for the Audiovox SMT 5600 mobile phone to measure and record the surrounding GSM radio environment. Each reading includes signal strength values, cell IDs and channel numbers of up to seven nearby cell towers. In addition, we extract channel numbers and associated signal strength values of up to 15 additional channels. Cell IDs are uniquely identified by the combination of Mobile Country Code (MCC), Mobile Network Code (MNC), Location Area Code (LAC), and cell id. Although other cell towers may be present in the area, our application only sees those associated with the phones SIM card provider. We sampled our GSM radio environment with the mobile phone at a rate of one sample per second (1 Hz).

4.2.2 Inferring User Mobility Modes

Our method for detecting user mobility is based on the same principle as fingerprint based location systems [BP00], [OVLdL05]: namely that the radio signals observed from fixed sources are consistent in time, but variable in space. Thus, given a series of GSM observations with a stable set of towers and signal strengths, we conclude that the phone is not moving. Similarly, we interpret changes in the set of nearby towers and signal strengths as indicative of motion.

We conducted a simple controlled experiment to determine how the radio environment changes as a result of various movement activities. Figure 4.1 shows the average Euclidean distance values between consecutive GSM measurements, as the data collectors stood still, walked and drove at different speeds. Conceptually, Euclidean distance captures the similarity between GSM measurements. The smaller the Euclidean distance between two measurements, the more similar these measurements are. For example, if measurement A has 3 cells/channels with signal strengths and measurement B has the same 3 cells/channels with signal strengths, the Euclidean distance between measurements A and B will be calculated as:

Our method for detecting user mobility is based on the same principle as fingerprint based location systems [BP00], [OVLdL05]: namely that the radio signals observed from fixed sources are consistent in time, but variable in space. Thus, given a series of GSM observations with a stable set of towers and signal

strengths, we conclude that the phone is not moving. Similarly, we interpret changes in the set of nearby towers and signal strengths as indicative of motion.

We conducted a simple controlled experiment to determine how the radio environment changes as a result of various movement activities. Figure 4.1 shows the average Euclidean distance values between consecutive GSM measurements, as the data collectors stood still, walked and drove at different speeds. Conceptually, Euclidean distance captures the similarity between GSM measurements. The smaller the Euclidean distance between two measurements, the more similar these measurements are. For example, if measurement A has 3 cells/channels with signal strengths $\{R_1^A, R_2^A, R_3^A\}$ and measurement B has the same 3 cells/channels with signal strengths $\{R_1^B, R_2^B, R_3^B\}$, the Euclidean distance between measurements A and B will be calculated as:

$$\sqrt{(R_1^A - R_1^B)^2 + (R_2^A - R_2^B)^2 + (R_3^A - R_3^B)^2}$$

If a particular cell/channel is not present in one of the measurements, we substitute its signal strength with the minimal signal strength found in this measurement.

Figure 4.1 shows that the Euclidean distance between consecutive measurements is proportional with the speed of movement. During stationary periods, the distance values stay relatively small (< 5). The slow and fast walking periods show a distinct difference from the stationary period. The driving traces show the most rapid changes in the radio environment, greater than either walking or stationary. Fast walking and slow driving sometimes overlap in their range of Euclidean distance values, which may result in false recognition between the two states. For a given speed, the Euclidean distance values are not constant because changes in signal strengths are both a function of speed as well as the physical environment, such as buildings, people, or vehicles.

Based on these findings we extracted a set of seven different features to use in classifying a set of GSM measurements as either stationary, walking, or driving. Three features compare two consecutive measurements in time, while the other four features use a sliding window of measurements. We used window sizes of 10, 60, and 300 seconds. Our seven features are:




Figure 4.1: Average Euclidean distance between consecutive measurements during a stationary period, slow/fast walking periods and slow/medium/fast driving periods.

1. Euclidean distance between two consecutive measurements
2. Spearman rank correlation coefficient [Pre92] between two consecutive measurements. (This number represents how closely the signal strengths from common cell towers were ranked. A more similar ranking indicates less movement.)
3. The number of common cell towers between two consecutive measurements.
4. Mean Euclidean distance over a window of measurements where the values are calculated between consecutive measurements and then averaged together.
5. Variance in Euclidean distance values over a window of measurements where the values are calculated between consecutive measurements.
6. The variance in signal strengths for each tower seen within a given window.

(The variance values for each tower are averaged together to produce a single number representing the signal strength spread over the entire window.)

7. Euclidean distance value between the first and last measurement of a window.

We used these features to train a two-stage classification scheme. The first stage classified an instance as stationary or not. If the instance was classified as not stationary, a second classifier would determine if the instance was walking or driving. Both classifiers were trained using a boosted logistic regression technique [FHT00] using decision stumps as a single node decision tree. All algorithms were provided by the Weka machine learning toolkit [WF00]. We chose to use boosting because it has been shown to work well in a variety of classification tasks [PRGH00], [SB98]. In our own experiments we compared boosted logistics regression with naive Bayes, Support Vector Machines, AdaBoost [FS96], MultiBoost [Web00], and some heuristic-based methods; the boosted logistics classifier provided the best recognition rates. We also compared the two-stage classification approach to a single multi-class approach and found that the two-stage classifiers resulted in better accuracy. This is consistent with the findings of Viola and Jones [VJ04], which showed that cascades of classifiers can achieve better recognition rates than single multi-class approaches in face detection tasks. The other advantage of the boosted logistic regression technique using decision stumps is that after the boosting process we have a ranking of features based on how useful they are during classification. Thus the system can be used to select features as well as learn classifiers simultaneously [LCK⁺05], [VJ04]. Furthermore, using only a small subset of the most relevant features can provide computational savings, which is especially important when running inference on a mobile phone.

4.2.3 Estimating a User's Daily Step Count

A nice feature of a mobile phone being able to determine periods when a user is walking is that it can be used to approximate how much a user walks, similar to the information provided by a pedometer. Pedometers are currently popular and are used worldwide as a tool to help people track the number of steps they

take each day. The benefits of walking and the use of pedometers have been widely promoted by the healthcare community, and a popular suggestion is for people to walk at least 10,000 steps/day [TL04].

To provide a reasonable measurement of steps taken (or step count), a pedometer is clipped to the users waistband, above the thighs midline. This restriction may be problematic, as some users do not like the look of the pedometer, or may not have a place to clip it, for example, if the user is wearing a dress. The mobile phone does not have such a restriction, as it can be anywhere with the user, including in her bag. We do not expect a person to always have their phone on her, such as when she is at home. However, being able to provide pedometer-like functionality when outside the home can be useful to give a high-level report of a persons mobile activity for the day.

The GSM-based mobility recognition from the previous section allows us to add a pedometer-like capability to mobile phones. By totaling the number of walking periods and multiplying by an appropriate step rate, we can estimate the users daily step count. Although this method of calculating step count may seem crude and prone to error, we show in Section 4.3 that our GSM-based step count estimates can approximate that of several commercially available pedometers.

4.3 Experimental Evaluation

In this section, we evaluate our mobility mode detection and step-count algorithms using data collected from three people to demonstrate the feasibility of using GSM traces to recognize high-level activities. We first describe our metrics and performance for mobility detection. We then evaluate our ability to estimate a users daily step count.

4.3.1 Data Trace Collection

Three members of our research team collected GSM network traces as they went about their daily lives for one month. Each data collector carried a commodity GSM phone, the Audiovox SMT 5600, running our software for recording readings

from nearby cell towers. Two of the data collectors used Cingular, and one used T-Mobile, spanning the two major GSM network providers in the U.S.

Data collectors recorded their mobility activities using a custom diary application running on the phone that allowed them to indicate whether they were walking, driving or in one place. Each collector also carried a paper notebook where he could record any event that he forgot to indicate on the mobile phone. These paper logs were later transcribed and merged with the digital log for a complete self-reported ground truth. There were a total of 53 corrections (7% of all events) from the paper logs for all data collectors. To capture the ground truth for step counts, each data collector also wore a pedometer and manually recorded his daily step count in the paper notebook. Each collectors pedometer was calibrated with his stride length and weight to obtain the most accurate step-count estimates possible.

We chose to use members of our research team to serve as data collectors because ground-truth diary logging is a tedious, error-prone process that required significant technical expertise to trouble-shoot problems with prototype technology. Given this overhead, the lack of application value to offer data collectors, and the high reliability of data logging that we required to test our algorithms, we felt that this was a reasonable choice. Our data collectors went to common places one would expect any person to visit such as grocery stores, malls, parks, churches, and libraries.

In all, the sensor logs contained 249 walking events (avg. 9.1 min) and 171 driving events (avg. 18.5 min). Each of these mobility events provides a sequence of data points to test our algorithm because every second is one data point to test our classifier (the rate that the phone scanned the radio environment). In total we gathered 12 GB of GSM network traces, amounting to 78 days of sensor logs. Our data spans urban and suburban environments and three different metropolitan areas as the data collectors traveled during the collection period.

		Predicted Movement		
		Stationary	Walking	Car
Precision	Ground Truth	Stationary	Walking	Car
	Stationary	95.4%	12.6%	6.9%
	Walking	2.5%	70.2%	8.8%
	Car	2.1%	17.2%	84.3%

		Predicted Movement		
		Stationary	Walking	Car
Recall	Ground Truth	Stationary	Walking	Car
	Stationary	92.5%	4.5%	3.0%
	Walking	7.7%	80.0%	12.2%
	Car	4.5%	13.8%	81.7%

Figure 4.2: Precision and recall confusion matrices for all GSM network traces aggregated over all data collectors. Overall accuracy is 85%

4.3.2 Inferring Mobility Modes

Our goal was to infer one of three mobility states: stationary, walking, or driving. Periods of walking and driving were identified in the data collectors diaries. We had initially hoped to use the remaining times, which data collectors marked as being at a place, to identify periods of being stationary. Unfortunately being at a place can still involve a fair degree of mobility. In a grocery store, shoppers are in motion much of the time. Even reasonably sedentary activities such as watching TV include short periods of walking (to visit the refrigerator for example). This ambiguity prevents us from having the needed ground truth for training and testing our algorithm. To extract the most reliable ground truth from our data, we used the GSM trace data collected between 2am and 5am to represent periods of being stationary. During these times we used our data collectors logs to verify that they were at home and sleeping, thus their phones would not be moving. Although this means dropping much of our collected trace data, it provides the best possible ground truth for determining how well our classifier can differentiate properties of mobility.

Using the labeled periods of activity, we trained our classifier and evaluated it using a 5-fold cross validation method over the entire data set. This produced a single model that worked well across all three data collectors and both GSM network providers. The precision, $\frac{truepositive}{(truepositive+falsepositive)}$, and recall, $\frac{truepositive}{(truepositive+falsenegative)}$, percentages aggregated for all of our data collectors are shown in Figure 4.2. The percentages along the diagonal indicate the classifiers performance for predicting and matching the ground truth events. Precision is the

percentage of predicted events that are correct. A low precision number indicates many false positives. Recall is the percentage of ground truth events that were correctly identified. A low recall number indicates that many ground truth events were missed. Accuracy represents the percentage of predictions that are correct. Our overall accuracy, $\frac{(truepositive+truenegative)}{totalnumberofsamples}$, is 85%.

Our classification scheme performs very well for stationary periods correctly detecting most periods of no movement (recall 92.5%) and not raising many spurious stationary events (precision 95.4%). Driving also performs quite well detecting most drives (recall 81.7%) and not raising many false positives (precision 84.3%). Walking activities were also detected with high percentage (recall 80%), but exhibited the most false positives out of the three classes (precision 70.2%). Within a driving activity, there are often times when a car is moving at slow speeds such as in traffic or roads with lower speed limits. In our controlled experiment, we saw that the changes in signal strengths for slow driving speeds are similar to fast walking speeds. Thus, one would expect the classifier to predict walking movement even though a segment was marked as a driving activity. These types of misclassifications are reflected in the walking precision (17.2% driving) and driving recall (13.8% walking) numbers.

The results show that we are able to distinguish between different mobility states with high accuracy without having to instrument a person with any other additional sensors. The precision and recall numbers show that this type of scheme could be used in a persons daily life, to give an accurate diary of mobile activity. In Section 4 we will discuss several application domains where our techniques would be useful.

One question about our classification model is whether it is overfitted for our data set. As an external way to corroborate our classification model, we tested the model using the GSM traces gathered from our controlled experiment described in Section 4.2.2. These traces are independent of those used to build our model. The classifier achieved an overall accuracy of 90% on this controlled data set, with the only errors being that some portions of our slow walk were classified as stationary. Furthermore, boosting techniques have been shown to be robust to over fitting and

generalizes well to unseen data [Sch01].

4.3.3 Daily Step-Count Prediction for Data Collectors

To test the accuracy of a virtual pedometer capability, we asked our data collectors to wear an Omron Healthcare HJ-112 pedometer for a portion of the month during which they were collecting GSM data. We chose the Omron because it was rated as the overall best pedometer by Consumer Reports [Rep04]. In all, we collected 50 days worth of daily step-count totals. In contrast to inferring mobility modes, for estimating step-count we want to be able to detect any walking activity throughout the day, even if it is for short periods of walking at a place. The pedometer is always logging the steps a person takes, so our algorithm must also detect these periods of mobility. Thus, for step-count prediction we used all of the collected GSM trace data for each day.

We wanted our step-count predictor to work without any calibration for all users. This further allows us to promote ubiquitous mobility recognition with low setup costs. To predict a daily step count from our walking predictions, we used the following simple heuristic obtained by performing linear regression with a 5 fold cross validation on our data set:

$$dailystepcount = 25 * (minutesofwalking)$$

For these 50 days of pedometer data, our heuristic predicted daily step counts ranging from 1500 to 12000 steps, with an average of 5000 steps. Comparing our estimates to the Omron step counts, we saw an average difference of 1400 steps per day (std. dev. 900 steps), with a minimum difference of 1 step and a maximum difference of 3500 steps. Our step count estimation worked uniformly well for all users: the correlation between measured and predicted step counts for the three data collectors were $R=.71$, $.63$, and $.63$. The error in our step count estimation is likely due more in part to errors in mobility estimation than to the user having different step rates.

To compare how well our step count predictions compared to other pedometers, we conducted a second experiment. We purchased four additional pedometers

of varying brands, and collected seven more days of data for one data collector. For this experiment, he carried the GSM phone, while also wearing the Omron and the four other pedometers. Again, we used the Omron as ground truth in our evaluation. For these seven days, our GSM based predictions had an average difference of 1400 steps with a maximum difference of 2400. The average difference across the other pedometers varied between 500 and 900, with a maximum difference of 1500. These results show that while less accurate, our GSM-based step prediction approximates the results of off-the-shelf pedometers in predicting whether a person had a sedentary, moderately-active, or high-activity day.

4.4 Applications

Our mobility detection scheme provides a low-cost, ubiquitous method for high-level activity recognition. Since we use commodity GSM phones without any additional hardware, any owner of a GSM phone can use our mobility detection system. In this section, we describe two application domains where our mobility detection scheme would be useful.

4.4.1 Computer-Supported Coordinated Care (CSCC)

CSCC describes the network of people who help an elder age in place, i.e., avoid the transition to a care facility, and seeks to improve the quality of her care while reducing the burden on the members of her care network, such as her family and friends [CRS⁺04b]. The Digital Family Portrait [MRCJ01] and CareNet Display [CRS04a] are two applications in the CSCC domain that aim to use sensor-driven activity inference to convey care and wellness information about an elder to members of her care network. The applications report information such as: Did the elder take her medication? Did she get out of bed? Did she have any visitors? Much research has focused on inferencing these types of in-home activities, but as the CareNet Display showed, an elders care network is also concerned with activities that take place outside of the home, such as did the elder go to church on Sunday? Is she routinely late for her weekly doctors appointment?

A recent report estimated that about 50% of Americans aged 65 to 74 are wireless customers and 30% of those aged 75 to 94 have mobile phones [Bro08]. Given that so many elders already carry them, mobile phones present an interesting opportunity to provide detection of a range of activities that are meaningful to the elders care net-work and can be detected today with a device that she already carries. With just a mobile phone, an elder would be able to relay information about her daily activity level, whether or not she was up and about today, or if she had a sedentary day around the house.

4.4.2 Social-Mobile Applications

Detecting mobility patterns is useful for applications that connect people with mobile devices together in their social environment—social mobile applications [Smi05]. These applications— if one includes voice calls and SMS —are key drivers of mobile phone usage today and are likely to continue as more and more of people’s non-work lives revolve around mobile communications [IOM05]. New applications are on the horizon that will help people communicate and coordinate [SCL⁺05].

Mobility detection can provide context information to enhance these applications and provide a better experience for the user. For example, applications that prompt a user with information are competing for that persons attention and potentially inter-rupting an ongoing task. Our technique would be useful for example when driving, because the information might better serve the user if it is delayed. Mobility detection could be central to some applications such as one that computes estimated time of arrival for many people who want to rendezvous. In a scenario of this type, one user – perhaps who is holding the movie tickets– is very interested when the other 3 users will arrive. With mobility detection alone, the waiting user can discriminate that some others have parked already and are thus nearby and those who are still driving and thus distant; combining this with a location system provides an excellent tool for social coordination and obviates the need for many phone calls and SMS messages.

4.5 Related Work

The SHARP project aims to infer fine-grained activities by putting RFID tags on household objects and monitoring their usage with a wearable RFID reader [PFP⁺04]. Our approach complements the fine-grained activities SHARP can infer from instrumented objects, with high-level activities in the wider environment using low-resolution sensors.

GPS-based location sensing has been used for high-level activity recognition. Patterson et al. take a learning approach based on particle-filters to detect modes of transportation [PLFK03]. Similarly, Liao et al. extended Relational Markov Networks for learning models that, given a GPS location and the time, can differentiate among shopping, dining, visiting, at home, and at work [LFK05]. GPS sensing today still often requires purchasing and carrying additional hardware. A recent study revealed that GPS positioning is available only about 5% of a typical persons day, as it needs a wide swath of clear sky to sense enough geostationary satellites [LCC⁺05]. In contrast, mobile phones provide ubiquitous coverage, and do not require any extra hardware from what people already carry. We have shown in this paper that similar recognition performance can be achieved observing changes in cell tower signal strengths, without the need for true location. This suggests that GPS should play an assistive role in everyday inference, rather than serving as the sole environmental sensor.

Two projects have looked at using radio signals for motion detection. LOCADIO used a Hidden Markov Model to infer motion of a device using 802.11 radio signals [KH04b]. Anderson and Muller conducted a controlled, preliminary study with GSM mobile phones to detect motion of a device [AM06]. Shakra is a system that uses GSM mobility detection to encourage awareness of daily exercise awareness among friends [AMS⁺07]. Similar to these projects, our approach uses machine learning algorithms to infer motion.

A third approach to activity recognition is to use wearable sensors of a single modality [BI04] or multiple modalities [LCK⁺05]. Lester et al. use 7 different types of sensors, including light, audio, accelerometer, compass, temperature, humidity, and barometric pressure, to classify 10 activities such as sitting, standing,

walking up stairs, and walking. The GSM radio can potentially be part of the sensor ensemble to improve recognition performance. Several commercial phones are now shipping with built-in accelerometers and compass, but, unfortunately, they do not expose the sensor readings to the application developers. Finally, the Reality Mining project has used Bluetooth-capable GSM mobile phones to recognize social patterns in daily user activity, infer relationships, and model organizational rhythms [EP06]. It uses the single associated GSM cell tower, Bluetooth radio, application usage logs, and call logs to sense nearby Bluetooth phones and devices, time and duration of calls, caller ID, and so forth.

4.6 Conclusions and Future Work

We have demonstrated the feasibility of using an unmodified GSM phone, a coarse-grain but ubiquitous sensor with 1.5 billion subscribers worldwide, to recognize high-level properties of mobility that are valuable for application domains such as Computer-Supported Coordinated Care and social-mobile applications. To evaluate its effectiveness, we collected GSM traces and ground truth labels of walks and drives for a month from the everyday lives of three people, for a total of 78 days of GSM logs consisting of 249 walking events and 171 driving events. We have shown that we can recognize mobility modes among walking, driving, and stationary correctly 85% of the time, and estimate daily step counts that approximates commercial pedometers. Unlike other activity recognition systems that may require a person to wear a special device in a certain way, our approach lets users maintain their current mobile phone habits with no special requirements about where the phone is kept on their person. These results show that current mobile phones without extra sensors or devices can detect high-level activities, providing people with an estimate of their mobility patterns throughout the day.

Since our classification model was built mainly in one metropolitan area, we do not anticipate it working across different cell densities. However, building a model for our classifier with areas of different cell densities could enable our techniques to work in varying radio environments. Our future work involves exploring

how our mobility detection technique and GSM-based step predictor would work in other parts of the country.

4.7 Acknowledgements

This chapter, in part, is a reprint of the material as it appears in Mobility Detection Using Everyday GSM Traces. Timothy Sohn, Alex Varshavsky, Anthony LaMarca, Mike Y. Chen, Tanzeem Choudhury, Ian Smith, Sunny Consolvo, Jeffrey Hightower, William G. Griswold, and Eyal de Lara. In UbiComp 2006: Proceedings of the Eighth International Conference on Ubiquitous Computing. pages 212-224. Springer, 2006. The dissertation author was the primary investigator and author of this paper.

Chapter 5

Place-Its: Location-Based Reminders for Mobile Users

A key requirement for mobile users is they need a way to defer information needs for a later, more appropriate time. One way to accomplish this is by using GSM location technology to create location-based reminders. This chapter will discuss Place-Its, a location-based reminder application that can be used to defer information needs. Place-Its is supported by GSM-based location technology discussed previously. Reminders are quite versatile in their usage, but are useful as a tool to address information needs later. The following sections describe the design, implementation, and evaluation of Place-Its.¹

5.1 Introduction

Everyday we use special messages in order to help us remember future tasks. These messages, known as reminders, take many forms, such as post-it notes, emailing one-self, to-do lists, and electronic calendar alerts. For example, a student may send himself an email to remind himself to bring a book for class the next day. Reminders can be more helpful when rich contextual information is used

¹Note that this chapter is a reprint with minor changes of *Place-Its: A Study of Location-Based Reminders on Mobile Phones*, a paper co-authored by Timothy Sohn, Kevin Li, Gunny Lee, James Scott, Ian Smith, and William Griswold.

to present them at appropriate times in appropriate places. [DA00]. A grocery list reminder is more helpful while passing the supermarket en route home from work, rather than while at work or after getting home.

Several context-aware systems [DSA01], [GSB⁺04], [MS00] have prototyped reminder applications [DA00], [MS00], but the evaluation of applications built on these systems has only been conducted in limited areas. In a recent pilot study on location-based reminders, we found that the reminders that people wanted extend beyond life in the research lab into all aspects of their personal lives [LSG05]. In particular, people often set reminders because the current context, both physical and social, prohibited completing the activity at the time. Therefore, our ability to understand the role of contextual reminders in a persons natural setting depends on a ubiquitous system being available consistently in a persons life.

A compelling platform for pervasively deploying context-aware reminders is mobile phones. Mobile phones with location-sensing capabilities are becoming state of the art, and several location-aware applications are available for use. The ubiquity of mobile phone networks enables pervasive location sensing, while the always-carried and always-on nature of phones mean that reminder creation and notification are permanently available to users. These factors allow a reminder system to be omni-present in the everyday life of a user. In addition, reminder notifications on mobile phones do not require any extra hardware, and gives people a familiar device for in situ interaction.

Yet, the suitability for phone-based context-aware reminders is unclear. The sensing capabilities of phones are limited in the types of context and their accuracy of sensing. The limited input capabilities of mobile phones, combined with the tendency towards use while committed to another task, suggests a simplistic user interface that permits posting a reminder in a few key-presses. Can potentially inaccurate, one-dimensional reminders (i.e. using location) prove useful, and if so, how? Are the phones capabilities in data entry, notification, and viewing adequate?

These technical and social limitations motivate the focus on location as a context cue. Recent advances in computing and location-sensing technologies are enabling high coverage location-sensing opportunities [LRT04], [LCC⁺05] to

use in building location-aware applications. Most of these systems require some initial configuration, but can then provide pervasive location sensing throughout a persons daily life. Using ones location to trigger reminders is a potentially valuable piece of context that can improve the way people use reminders. Our aim is to find how location-based reminders are used when available throughout a persons day. Of course there are types of reminders in which location is not useful, but our focus is on those that could benefit from the additional location information. How and why does location figure into the relevance of a reminder for a person? How important is positional accuracy and timeli-ness to the usefulness of location-based reminders?

In the following sections we describe the design, implementation, and deployment of a location-based reminder application, named Place-Its, for mobile phones. This simple application, with the mobile phone as a platform, permitted the integration of location-based reminders into peoples daily practice. We then report on a 10-person user study involving Place-Its over a two week period. The study participants found location-based reminders to be useful, despite relatively low location accuracy. In particular, participants found value in having the application always naturally on-hand for posting and receiving reminders, along with pervasive location sensing. Also, the participants used location-based reminders in numerous ways, including several in which the location served as a convenient proxy for other kinds of context. Finally, we conclude with a discussion of implications for future research.

5.2 Related Work

The idea of using location information in context-aware applications is not new. Much work has been done in the past in context-aware prototypes that have all shown location to be a useful element of context.

The Forget-me-not project was one of the pioneering efforts in the area of context-aware reminders. Forget-me-not employs a small PDA-like device that associates different items of interest with icons to help the person remember various

tasks they need to attend to [LF94].

ComMotion is a more recent example of a context-aware system, supporting reminders that utilize location as contextual information [MS00]. Using GPS technology for location-sensing, people could set reminders around certain locations, with given time constraints. When the person was near that location and the timing constraints were satisfied, they would be alerted with an audio alert.

Cybreminder [DA00] took these ideas a step further, developing a reminder application based on the Context Toolkit [DSA01] that focused on using a variety of context information, including location, to determine when best to trigger reminders. This project focused on abstracting hardware technology away from the developer. Thus, it was able to create a fully featured reminders application taking into account a variety of contextual information. This toolkit relies on the existence of special sensing hardware that limits its ability to be deployed ubiquitously.

Focusing on a different aspect of reminders, Stick-e notes [Pas97] explored the post-it metaphor in the digital rather than physical world. Stick-e notes were placed at particular locations using GPS enabled PDAs, and could be made visible to others, thus emulating the affordances of physical notes in a digital environment.

Unlike the previously mentioned systems, ActiveCampus [GSB⁺04] examines mobile computing restricted to a university campus setting and provides a location-based reminders system using 802.11 radios to provide location sensing. A reminders feature is integrated into the system where people can set reminders to be triggered at predefined locations on campus, typically buildings.

Spraycan is a mobile phone digital graffiti system that has similar concepts to location-based reminders [GRCE06]. The system uses RFID tags to mark locations and 'tag' the locations with user-generated content. The application runs on a mobile phone equipped with a RFID reader.

PlaceMail is a similar application to our system that uses mobile phones and a web-based interface to deliver location-based reminders [LFR⁺06]. PlaceMail offers tools to support people's everyday tasks through a reminder and personal information management system. Although PlaceMail is designed with ubiquity in mind, it still suffers from the common coverage problems of GPS. The main

contribution of PlaceMail is an understanding of how people accomplish tasks using location.

The wearable computing research community has also had a number of projects exploring reminder notifications through extra hardware. Memoclip uses a small wearable computer that relies on location beacons distributed in the environment to trigger location-based reminders [Bei00]. The reminder bracelet involves a bracelet worn on the wrist of the user that subtly alerts the wearer of upcoming events, as entered into their PDA calendar, using temporal information only [HL00]. The Sulawesi framework discusses a spatial reminder service that uses GPS and infrared to approximate a person's location and delivers reminders accordingly [NC99]. Memory glasses is another project that attempts to augment user memory by using subliminal cues [DPC03]. The wearable remembrance agent used a heads up display to provide context-relevant information [Rho97]. The MIT Smart Helmet is an ongoing project at MIT that explores a hand-free implementation of a location-based notification system. The system delivers location cues through a bicycle helmet [JSC07]. While some of these systems use implicit reminders, with some type of context or activity inferencing, they all find that subtle cues are effective. Like the systems mentioned above, these wearable solutions also focus on exploring new services that can be provided using location as context.

In the ethnographic space, Kaasinen performed a study examining people's needs for location-based services, using existing tools and methods [Kaa03]. People were given GPS enabled devices and typical tasks to perform, with the focus being on their expectations for location-based services. The participants found the scenarios given to them to be unrealistic, as they felt the given situations did not reflect the real needs of people. They also expressed concerns about how location-aware systems might cause drastic changes in human interaction. Antifakos et al. examined the effects of imperfect memory aids when used to help recall. They found that displaying uncertainty helped to improve recall rates, especially in cases where uncertainty was high [ASS04].

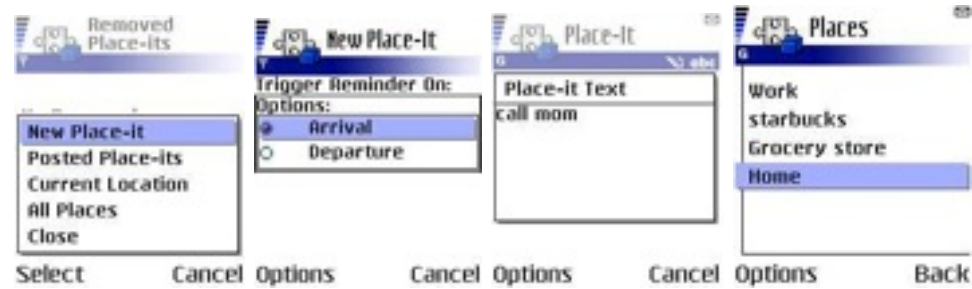
These pioneering efforts have generally not permitted evaluation in the context of people's normal lives. This is because they either used additional hardware,

such as GPS receivers, that people typically do not carry around (taking both cost and inconvenience into account), or because they are restricted to a predefined area such as the campus of a university. GPS itself provides restricted location sensing, only operating in outdoor regions where line of sight with multiple satellites is available, whereas most people spend their days largely indoors, and even while outdoors GPS may not be available in areas such as urban canyons. Studies conducted on university campuses typically use radio technology for location-sensing, requiring that beacons be placed in the environment, something that currently is not sufficiently available out-side such settings [GSB⁺04]. Other location-sensing technologies, not listed here, could be used for a reminders application but possess similar restrictions, requiring special hardware [HB01]. Wearable computing solutions also require special hardware to provide many of the developed services. As a result of these constraints, it has been difficult to examine how location-based reminders could be used when integrated into someones daily routine. The study here takes advantage of research done in client-side GSM-based location-sensing technology as well as the popularity and low cost of mobile phones to study location-based reminders, free of the aforementioned restrictions.

5.3 The Place-Its Application

Place-Its is designed around the post-it note usage metaphor, and named for its ability to place a reminder message at a physical location (i.e., a place). It usefully deviates from the metaphor in that notes can be posted to remote places. Although a person is home for the night, he can post a note at work to be retrieved the following morning upon arrival. To convey how Place-Its is used for location-based reminders, the following is a scenario based on Place-Its actual usage by one of the participants from our study (screenshots next page).

Jill is currently busy at work, and remembers she needs to call her mom. She knows she will have free time when she gets home, so she conveniently creates a Place-It note (Figure 5.1(a)) that will trigger on arrival (Figure 5.1(b)) and remind her to call her mom (Figure 5.1(c)) at home (Figure 5.1(d)). Jill glances and sees



(a) Creating a new Place-It note (b) Setting the note to be triggered upon arrival (c) Typing the text of the note (d) Posting the note to Home



(e) Showing all posted Place-It notes (f) The reminder is triggered when Jill arrives at the home and the note is removed

Figure 5.1: Screenshots of the Place-Its application.

that she has two other notes that are still posted (Figure 5.1(e)). After work Jill has dinner with her friends and arrives at home later that night. As she walks into the house, her phone vibrates and displays a message (Figure 5.1(f)), reminding her to call her mom. Immediately, she makes the call.

The three components to a Place-It reminder note are the trigger, text, and place. The trigger identifies whether the reminder should be signaled upon arrival or departure of the associated place. The text is the message associated with the note. Reminders are created with a message, and then posted to a location on the persons list of places. People could use the phones predictive text input for entering their reminder texts to ease the burden of typing on the phones keypad.

A person can view all posted reminder notes at any time (Figure 5.1(e)) and can delete or edit any of the fields associated with the note. After a reminder note is posted on a place, when the trigger (arrival or departure) occurs, the note is automatically re-moved and put in the Removed Place-Its list. Once a note is removed, it can be edited, and reposted to the same or a different location.

We had three main design principles for the Place-Its application. First, it must be an always-on service, to ensure that reminder notifications are always possible and users can have confidence that they will get their requested reminders. A reminder system that is only available a small percentage of a persons day is ineffective if the message needs to be delivered outside the operational time frame. Second, the application must be easily deployable. Requiring people to carry extra pieces of hardware can hinder their integration of the reminder tool into their daily activities. The application is best deployed through a familiar artifact that people already use or carry on a daily basis. Last, for several reasons, we purposely omitted features to set time-based reminders. This choice simplified the interface, omitting at least one step from the reminder creation process. It is also consistent with current reminder applications, which permit using just one kind of context (typically time). This was also a natural way to force all users to use location as their primary source of context when creating the reminder, even though they were not accustomed to it.

5.3.1 Mobile Phone Platform

Mobile phones offer a way to meet our first two design goals. Mobile phones are emerging as viable platforms for deploying personal ubiquitous computing applications. Multimedia, communications, and general computing capabilities are all converging to the mobile phone platform. Phones are also deployed and carried by over a billion people across the world, making them ideal for application deployment. Since people already carry mobile phones with them for communication purposes, they are unlikely to forget the device or have it hinder their normal activities. Mobile phones also present a convenience factor. Because people typically keep their phone handy for answering calls, they also can quickly create and

receive a reminder note on their phone anywhere at anytime. Finally, the public use of a phone is, in the most part, socially acceptable. Given these advantages, our aim was to deploy an application designed for the phone platform, taking care to exploit these advantages where possible.

Place-Its is targeted at the Symbian Series 60 platform and was written using Java 2 Micro Edition (J2ME), with the Connected Limited Device Configuration (CLDC 1.0) and Mobile Information Device Profile (MIDP 2.0) APIs. The application contains a small portion of C++ code to access GSM cell tower information on the device. All development and deployment was done on the Nokia 6600 phone, chosen for its large screen and good developer support.

5.3.2 Location-Sensing

Achieving pervasive location sensing is essential for Place-Its to be useful. The Global Positioning System (GPS), with all its strengths and weaknesses, is the most widely used location technology today. However, it is not widely available on mobile phones, although it is beginning to emerge on certain models. Another option would be to take advantage of the recent efforts by telephony providers to support accurate location sensing on mobile phones to meet E112²/E911³ requirements. A byproduct of this effort is the development of commercial pay-for-use location-based applications. It is possible to use a provider-based location mechanism for Place-Its, however these services are only emerging, without much support for independent developers to access location information directly. Per-use charging schemes would also cause high costs to be incurred by applications such as Place-Its, which need to continually monitor the users location.

Looking to perform all location computations on the client device to avoid any charges by the provider, a better solution is Place Lab, a location system that relies on mapped radio beacons in the environment to provide location estimates [LCC⁺05]. Place Lab for the mobile phone platform can use both GSM and Bluetooth radio technologies for location sensing. All computation is done on

²<http://ec.europa.eu.int/comm/environment/civil/prote/112/112.en.html>

³<http://www.fcc.gov/911>

the client device, preserving peoples location privacy. Previous studies have shown that Place Lab achieves approximately 100 meter accuracy and 100% coverage in urban areas. However, the initial requirement for cell tower location mappings is problematic since if a person moves outside of a mapped area, location capabilities are no longer available. Because it is costly to map every possible region where a person will go, we chose not use Place Lab. However, as the number of mapped regions grows, Place Lab would be a good solution for location-based applications like Place-Its.

We chose to use the location technique employed by Reno [SCL⁺05] for marking and detecting places, inspired by the ideas of Laasonen et al. [LRT04], but simplified some-what for our purposes. The key observation made by Laasonen et al. is that one can think of a place as a clique in a graph. This graph starts out as a set of nodes representing GSM cell towers. Edges are added to the graph between nodes as the phone observes a transition from one cell tower to another. A transition A-B means that a mobile phone was associated with cell A and now is associated with cell B. By using this strategy, the graph of nodes that a mobile phone associates with can be constructed.

When a mobile phone is stationary in one place, it does not stay always associated with the same cell tower; it "hops around" or, in the graph sense, traverses a set of edges on the graph. In a given location, the small set of nodes, a clique, that are traversed is typically quite stable for a mobile phone that is not moving. We watch for cyclic transitions such as visits to the sequence of nodes, A, B, C, A. The Reno location algorithm considers this the clique A, B, C. Reno's algorithm basically defines a place as the sequence of nodes visited in a cycle when that cycle has been repeated more than once. For example, if the sequence of nodes A, B, A C, B, A is observed, the algorithm considers the current place to be defined by the clique A, B. (C was not visited more than once.)

When trying to discriminate places, the algorithm simply takes the set of nodes seen recently (within some time window determined by hand tuning) and looks for cliques that overlap this recent list. For simplicity, the algorithm ranks the possible places (cliques) based on the amount of the clique that is "covered" by



Figure 5.2: Once Place-Its is able to determine a unique location signature, the person can name the location by typing it in, or using a pre-defined Cheat Sheet.

the recent nodes. This favors small cliques over large ones, but in practice causes few problems.

Once a place is determined, the person is able to label that place with a unique name (Figure 5.2). An advantage to this approach is that places can be identified wherever a person goes independent of any extra hardware or external data. A disadvantage is that a place must be previously visited and marked by the person before any reminder messages can be associated with that place.

5.4 User Study

In this section, we describe an exploratory user study of Place-Its performed with ten participants, over a two-week period in Winter 2005. We studied location-based reminders in the daily lives of people with different occupations, to analyze behavioral and usefulness factors for location-based reminders through a mobile phone interface. As an initial study, we emphasized naturalness over experimental controls, enabling us to observe genuine behaviors that could set directions for future research, as well as inform application design and future experiments. In this section we describe our experimental setup, and in the subsequent two sections describe our results gathered from the study.

<u>Participant</u>	<u>Age Range (Gender)</u>	<u>Occupation</u>	<u>Current Reminder Method</u>
A	18-21 (M)	Undergraduate	post-its, notepad
B	18-21 (F)	Undergraduate	paper planner
C	22-25 (M)	Graduate	email, notepad
D	22-25 (F)	Graduate	None
E	22-25 (M)	Graduate	electronic calendar, email
F	25-35 (M)	Professional	PDA
G	35-45 (M)	Professional	email, visible items
H	22-25 (M)	Professional	post-its, electronic calendar, phone alarm
I	25-35 (M)	Professional	electronic calendar, PDA
J	22-25 (F)	Professional	post-its, email, electronic calendar, phone

Figure 5.3: Demographic information about each participant and their current method of creating reminders. A-E are students and F-J are working professionals.

5.4.1 Participants

We recruited participants through mailing list postings and advertisements, seeking a group that used a variety of methods for creating reminders using other tools, had experience with mobile phones, and would manifest a variety of location-based reminder behaviors. The ten chosen participants consisted of students and working professionals, ages 18-45, three women and seven men (see Figure 5.3). Five of the ten participants were undergraduate or graduate students at different universities in the area. The other five were full-time working professionals. None of the participants had been exposed to using location technology for creating and delivering reminders before. Each of the participants had a GSM service provider, allowing them to use our application and location detection algorithm, while maintaining their communication capabilities with minimal hassle.

The dominant reminder habits of the participants fell into four broad categories. Three used personal information management tools for their reminders (e.g., Microsoft Outlook, PDA), three used mainly email, and another three wrote their reminders down in a notebook or on post-it notes. The last participant did not use any of these methods, relying purely on memory.

5.4.2 Methodology

We conducted our study in three steps, a pre-study questionnaire, a two week long deployment, and post-study interview. Our pre-study consisted of a basic question-naire regarding demographic information, mobile phone usage habits, and current methods of creating reminders. To help the participants personalize their Place-Its application, we asked each to provide, in advance, up to ten frequently visited places where they might want to set a reminder during the study. These pre-defined Cheat Sheet lists were put on the phone for the participants, enabling a participant to define a new place without having to type it in using the phones keypad (Figure 5.2). The person was still required to visit that location and mark the place before they could label it.

We provided each participant with a Nokia 6600 to use during the study. However, they transferred their personal SIM card to the new phone for the duration of the study, thereby transferring their address book and allowing them to use the same phone number and phone network account. Since the phone was unfamiliar to some participants, we also conducted a basic phone tutorial. We explained our application, using the adapted post-it note usage metaphor. Each participant was told that before a reminder could be posted at a location, they had to have visited the location and named the place. Reminder notifications would be triggered by a phone beep and vibration, and would conform to the profiles of the phone. If the phone were in silent mode, a notification would not occur. Our participants were aware that we would be logging usage data on the device for analysis after the study completed. We asked the participants to incorporate the application into their daily lives and routines, using the application to set reminders as the need arose. After one week, the participants filled out a mid-study questionnaire by email regarding their experiences with Place-Its and the types of reminders they were posting.

Near the end of the two weeks the participants were sent a post-study questionnaire by email regarding their experiences with location-based reminders and the Place-Its application. The study concluded with each participant returning the Nokia 6600 and a 30 minute personal interview discussing their questionnaire




Figure 5.4: Reminders usage for each participant. 89 total reminders were created. 67 were arrival reminders and 22 were departure reminders. Participants A-E are students, and participants F-J are working professionals.

responses.

5.5 Observations and Initial Classification

There were 89 reminders created overall, of which 67 (75%) reminders were arrival trigger reminders and 22 (25%) reminders were departure trigger reminders (See Figure 5.4). One of the main reasons we found for the smaller number of departure reminders is that many participants, after trying departure triggers for certain types of reminders (e.g., bring the book when you leave), found that the reminder came to their attention several kilometers away from where the reminder was needed. The problem is that a unique clique (see Section 5.3.2) could not be determined quickly enough as a person moved away from a location. Thus, departure reminders were generally not used and the tasks were accomplished with arrival reminders instead (e.g., pack book when you get home). This type of adaptive behavior is described further in Section 5.6.

Nineteen reminders (21%) were re-posted one or more times for a total of

63 re-postings. Reposting reminders occurred because either the reminder was triggered at the wrong location or the note was triggered at the right location, but could not be attended to at the time. Participant <I> reposted several reminders more than forty times to the same location to provide motivation to study every time he came home. Participant <E> occasionally visited a coffee shop close to the university campus on weekends. <E> posted a reminder on campus to be removed during the week when he went to campus, but it was falsely triggered during a visit to a nearby coffee shop. He therefore re-posted the reminder.

Although participants could post reminders wherever they had defined a place, 32 (36%) reminders were posted on a persons home and 39 (44%) reminders were posted on a persons workplace/campus. All of the participants reported in the post-study interviews that the majority of their desired reminders involved either the home or workplace/school. On one or two occasions, participants forgot to mark a place they had visited, preventing them from posting a reminder to that location when they wished. For example, <I> wanted to set a reminder for a place he went to once a week. During the first week of the study he forgot to mark that place, thus was unable to set a reminder for the following week. If the study was conducted for longer period of time, these initial setup issues would have a lower effect. A technical solution would be for the application to automatically log time stamped cell tower observations. A person could define a place at a later time by correlating a visited location with a specific time, and then the application could retrospectively reconstruct the clique from the cell tower logs. The time lapsed between the posting and removing of a reminder varied from minutes to days. We can discern no patterns, whether across the subjects, places, etc., suggesting that there is little time-based correlation between when a person remembers to post a reminder and when the persons location permits the activity to be performed.

5.5.1 Self-Reported Usage Data

We implemented an on-device multiple-choice questionnaire that would appear on screen after a reminder notification. The questionnaire consisted of 4 questions that could be answered quickly with a few key presses. Where are you


	Correct Place	Not Correct Place		Remembered	Not Remembered
Expected	25	0	Changed Behavior	10	6
Unexpected	5	4	Did Not Change Behavior	14	4

(a) Number of self reported responses about whether the reminder notification was expected and whether it was delivered at the correct place. (b) Self-reported responses about if the reminder was remembered before the notification, and if the reminder changed what the participant was going to do.

Figure 5.5: Self reported answers by participants

in relation to where this reminder should be delivered? Was this reminder notification expected? Did you remember this reminder before the notification? Has receiving this reminder changed what you are about to do? The responses to these questions helped us gather feedback about the timeliness of reminder notifications and behavioral changes with regards to the reminder. We did not always offer a questionnaire after each reminder notification to avoid it becoming an annoyance. If the questionnaire screen was shown (about 50% of the time), the participant had a choice to ignore the questionnaire, or to proceed with answering the questions. If the participant did not respond to the questionnaire within a two minute time interval, the form would disappear from the screen. 49 (36%) questionnaires were acknowledged, and of those, 34 (69%) were answered.

Figure 5.5 shows a table of the 34 questions and the number of responses in each category. The dominant responses in Figure 5.5(a) show that most reminders notifications were given at the correct place and the notification was anticipated. However, there were some reminders in which although the location was correct, it was delivered unexpectedly (e.g., on arrival instead of departure). Only four responses indicated that the reminder notification was not at the correct place,



Category	Communicate	Task@Place	Bring/Get	Motivate	Awareness	Go To
Description	email, call, talk to a person	a task that can only be done at a specific place	Bring/get an item from a place	motivate to do a task	being aware of an activity	go to a place
Example	"Talk to Ken"	"Start Baking!"	"Bring Book"	"Study Greek"	"Check the time"	"Go to grocery store"

Figure 5.6: A classification of Place-It note usage into 6 categories.

and hence unexpected. In Figure 5.5(b), the horizontal dimension is whether the participant remembered the reminder before getting the notification. The vertical dimension represents if after being reminded, any behavioral changes occurred because of the reminder notification. The expected value of a reminder application is in reminding people about things that they have not otherwise remembered, and this is represented by 6 out of 10 non-remembered reminders causing a change in behavior. It is interesting though, that in 10 out of 24 cases where a reminder was for something that the user had already remembered, a change in behavior was nonetheless reported. This indicates that the Place-Its application was valuable as a motivation cue, beyond its memory-aid intent.

5.5.2 Classifying Place-It Note Usage

Psychology researchers have identified three general classes of prospective memory tasks: time-based, event-based, and activity-based [EM90], [BEM95]. Our

application enables reminders in the event-based and activity-based classes. We classified the ways Place-It notes were used into 6 sub-categories, to analyze their function and use (See Figure 5.6). This is not meant to be a definitive classification of every reminder that a user might want to create, but simply provides an idea of how our application was used for the duration of the study. Some reminders could be interpreted to be in two or more categories; in these situations, participant feedback was used to disambiguate. The goto category maps directly to the activity-based class of reminders from psychology literature. The other categories are varying types of event-based reminders.

There are several noteworthy points in the frequency of the different types of reminders. The largest category is communicate reminders. These reminders involve emailing, calling, or talking to another person. This is somewhat surprising, since many forms of communication are not place-specific. An email can be sent from any location with Internet connectivity. Using mobile phones, calls can be made at any time as well. User feedback informed us that such reminders were typically created because, at the moment where the idea of communicating came to the user, they did not have enough time to actually perform this communication. This implies that location is being used as a reminder cue for other kinds of situational context (e.g., inactivity or cessation of an activity), which we expand on later.

Another surprise, given the inaccuracy of the location technology for departure reminders, was that bring/get reminders were tied for the second largest type. Participants considered this type of reminder important enough that they adapted to the limitations of the technology. Participant <F> created a reminder, *Bring metal case* to be triggered on arrival. Due to receiving the reminder, he would pack the metal case soon after he arrived home, thus not forgetting it when he left for work the next day.

Another unexpected significant type of Place-It usage was motivate reminders. These reminders did not necessarily have a high priority at the time of creation, but would sometimes increase with priority as time went on. They existed solely to motivate the person to perform a certain task such as *Study Greek*

or *Go to the gym*. Participant <I> would always re-post his motivation reminders after they were re-moved, considering them of low priority at the time of removal; however, as time for his Greek exam drew near, the motivation reminders became helpful in his time management.

5.6 Post-Study Responses and Discussion

5.6.1 Mobile phones offer available, convenient reminder creation and delivery

Mobile phones provide a means of creating and delivering reminders that makes them attractive to users. We found that 8 of our 10 participants appreciated the consistent availability of location-based reminders through their mobile phone, and lessened their use of other reminder tools in favor of Place-Its.

Quote 1 <participant J>: *Since I was out of town, I would think of things on the drive that I had to do when I got back and I'd put reminders on the phone. Even though I did remember what I had to do without the help of the reminder, it was a relief knowing I would've been reminded had I forgot.*

Participant <J> normally uses post-it notes or an electronic calendaring system to send herself reminders. However, in this situation, neither of these means were accessible to her, thus <J> found it useful to use Place-Its to create her reminder. In addition, having a reminder that would trigger after she came back to town was useful because Place-Its would interrupt her to display the reminder. With <J>s current methods, its possible had <J> forgotten the task, the reminder would have been overlooked. In our post-study interview <J> explained that she stopped using her current methods of reminders in favor of Place-Its due to its availability.

Quote 2 <participant E>: *There are certain activities that my calendaring application is not particularly good at reminding me about. Especially to do something when I'm not near a computer. So getting reminders for these types of activities was a welcome behavior [examples are] grocery shopping, and also when Im leaving work Im on my way out, done for the day, not liable to be checking email.*

The two methods that <E> uses for reminders are an electronic calendar for all reminders, and email messages for critical ones. Neither of these two methods allow <E> to trigger reminders when away from a computer, while Place-Its provided an always-available application. One of the successes was that on the way home, <E> was reminded to go grocery shopping, which would not have happened otherwise.

Quote 3 <participant F>: *I didn't use my PDA much... it's much bulkier compared to just being able to use the phone*

Although <F> normally uses his PDA for reminders, during the course of the two-week study, <F> found it more convenient to use the phone. <F> still carried his PDA with him by habit and used it for other functionality, but preferred Place-Its for reminders due to both its location capability and the fact that the phone was always within reach to use.

The convenience of mobile phones encouraged four participants, who may not have bothered to create a reminder in the past, to enter a Place-It for the sake of not forgetting it. As an interesting side note, some people found that entering a reminder helped reinforce their own memory to perform the task. One person said:

Quote 4 <participant B>: *I would pull my phone out to silence it for class and [looking at it] would remember that a reminder would be coming.*

An important concern regarding the phone platform is the text input method. Many of the participants found text entry to take too much time even with predictive input support, so would resort to one or two word phrases. Those who wanted to input a grocery list found it easier to use Place-Its for the reminder to go to the grocery store, and have a separate paper list for the actual items. One possible solution to overcome these methods would be to use the voice and picture capabilities found on many phones today. This would enable quick voice memos, or snapshot pictures that would stimulate a persons memory about a reminder.

5.6.2 Location Provides an Indirect Cue for Other Context

The kinds of reminders posted and the way that they were posted strongly suggests that the location itself is not always important, but it is just a convenient proxy for context that is not as easily sensed or readily available..

Quote 5 <participant E>: *Im busy at work, so I dont want to make the call now, but I want to remember to call my sister when I get home*

During the workday <E> was typically too busy to take the time to make a call to his sister. However, the phone call was of enough importance that he didnt want to forget about it. Since <E> knew that he would have more free time when he gets home, he set the reminder location for his home. The location itself was not important, but <E> knew that time will be more likely to be available when no longer at work. Moti-vational reminders are often similar in that the location has been chosen to catch the person in a particular frame of mind (or change a persons frame of mind). The loca-tion may afford that frame of mind, but in many reminders the relevant location was akin to no longer at work or on the road.

Certain locations imply access to tools that may be used by the person in completing a specific task. These tools may offer services not innately tied to the location, but in the persons mind, the task can be completed there.

Quote 6 <participant F>: *I was in another building at work, when I thought to myself I should create a status report e-mail to send to my boss concerning my pro-gress on a recent project. Even though I would be back in about an hour, I decided to post the Place-It on the phone. When I came back to my building, the beep went off right as I got back to my desk. Looking at the Place-It, the e-mail then became the next thing I did.*

<F> needed to write an email at a time when he did not have access to the tools he needed to compose one. Knowing that he has a computer back in his office and that he will be there shortly, he set a Place-It for his office. Sending the email is not something that is innately tied to his office; he could very well send one from any place where he has computing facilities for email access. However, he is able to take advan-tage of some knowledge about his schedule for the rest of

the day to set a reminder for a location known to provide the services he desires in a timely manner.

Similarly, location can also imply the presence of other people, but without a reasonable guarantee:

Quote 7 <participant C>: *I made a reminder for myself to ask a lab mate about a class, and I got the reminder just as he walked into the lab... I set the Place-it for the lab because I figured he would be there.*

<D> was really looking for his lab mate when he set this Place-It, but realized that his labmate would probably be in the lab at some time. Using this foresight to his advantage, he set a Place-It there, knowing that it would draw his attention at a place where he would probably be able to find the person of interest. Although many context-aware systems have supported buddy alerts, this behavior demonstrates how setting reminders on locations can be used to alert the user when someone of interest is nearby, without the explicit ability to sense when buddies are nearby. Given the nature of certain relationships between people, it is often likely to find someone of interest at a particular location within a large but acceptable time range. By using this knowledge, one can use a location-based reminder to essentially create a person proximity reminder.

Activity inference is another context-aware feature that has long been desired and seems to be supported to some degree by location-based reminders, with some help from the user. A general statement of the challenge is to determine what the user's goal is in a long-running activity. Consider inferring the activity, going to Kevin's house. When asked to explain the thought process behind setting a Place-It on departure from work to *call Kevin*, <J> responded:

Quote 8 <participant J>: *I set the place-it for departure because I knew when I would go to the guys place after I left work. Even if I didn't go there directly, I knew I would go there pretty soon.*

When she set this reminder, <J> had some notion of what her activity would be later in the day when she left work. She knew that her schedule entailed going to the guys place eventually and very likely when she left work. This allowed her to use a departure reminder as a mechanism for aiding the system in

adequately inferring her activity. In this case, leaving work meant it was probable that she was heading to Kevin’s place.

As cited earlier, departure reminders often have high accuracy requirements. Indeed, they can even require predictive power really activity inference: which of the several times that someone leaves their office during the day is the last time, so that they should be reminded to bring a book? Inaccuracy almost becomes a feature: who wants to be reminded all afternoon? Time-constrained reminders could have helped, but our users used arrival reminders as a proxy you cannot leave a place until you have arrived, after all.

5.6.3 Location-based reminders are useful

In light of the above remarks, it is understandable that participant comments regarding location-based reminders were generally positive. Two participants (F, J) requested to be future research subjects because they found Place-Its to be helpful in their daily activities. They also asked if we could build a version of Place-Its for their current mobile phone to use on a regular basis. Six participants (A, B, C, D, E, H) considered location-based reminders to be useful to them, and their use of Place-Its to be enjoyable. The remaining two participants (G, I) did not find location-based reminders to be helpful, stating that their lives revolve around a set time schedule. They only desired time-based reminders or did not need any reminders at all.

During our post-study interview we asked each participant to describe any problems they experienced with Place-Its. These responses generally fell into two categories. Four participants had problems with the application being too easy to exit, or crashing. They were sometimes unaware of these events, and hence missed reminders. More significantly, the other six participants said the location algorithm used by Place-Its was sometimes not accurate enough for their reminders. The participants would get the reminder, but not necessarily at the right location. This degree of this perception lessened over time as the participants adapted their behaviors. One participant, not surprisingly, asked for time-constrained reminders.

Due to the way location-based reminders were used and the relative inac-

curacy of location-sensing in Place-Its, we cannot claim location itself is essential context, even as we find it to be useful for triggering reminders. More than anything, its ready availability admits opportunistic use by those who can map their relevant (but unsensed) context to anticipated, coarse, location cues. Indeed, the two participants who work by a set time schedule are achieving similar results by mapping their relevant context to time cues and modifying their behavior. Providing location-triggered reminders expands the palate of context affordances that people can appropriate to guide their activities, accommodating a wider range of personal organizational styles.

5.7 Conclusions and Future Directions

The prevalence of mobile phones and the pervasiveness of their networks makes them a promising platform for personal ubiquitous computing. Our findings from a two-week deployment of Place-Its help validate that location-based reminders can be useful even with coarse location-sensing capabilities. Notably, location was widely used as a cue for other contextual information that can be hard for any system to detect. On the whole, it appears that the convenience and ubiquity of location-sensing provided by mobile phones outweighs some of their current weaknesses as a sensing platform. This bodes well for the use of mobiles phones as a personal ubiquitous computing platform. Our study revealed unexpected uses of location-aware reminders. We found that Place-It notes were often used for creating motivational reminders to perform activities that would vary in priority over time. This is similar to using post-it notes in highly visible areas for motivation. The locations for motivational reminders were often set at frequently visited places, such as home. We also found that a majority of the uses for Place-Its involved communicating with people through a variety of media (e.g. email, phone). Communication is typically not tied to specific locations, implying that location is being used as a cue for other kinds of situational context.

As a first study, the results presented here are preliminary. Our results suggest a few application modifications that are worthy of further investigation. First,

given the limited text entry mechanisms available on mobile phones, a way of associating audio messages or pictures with reminders could offer greater convenience, encouraging unique and more opportunistic use. Second, with an understanding now of how location affords certain classes of reminders, it would be interesting to investigate how adding time-constrained notifications changes user behavior. Third, research into more accurate and faster location sensing on mobile phones should reduce the need for users to adapt their reminders to the capabilities of the application.

Finally, to both account for the effects of inaccurate location sensing and naturally support the use of recurring reminders, we propose a change to the user interface. Rather than the application automatically removing a Place-It when it is detected and presenting it as an explicit reminder notification, the application would continuously display a list of nearby Place-Its as to-do items, sorted by proximity to the current location. The user would then explicitly pull down a Place-It when it is no longer relevant, rather than repost it if it is still relevant. Alerts could still be provided when location certainty is high.

5.8 Acknowledgements

This chapter, in part, is a reprint of the material as it appears in *Place-Its: A Study of Location-Based Reminders on Mobile Phones*. Timothy Sohn, Kevin A. Li, Gunny Lee, Ian Smith, James Scott, and William G. Griswold. In *UbiComp 2005: Proceedings of the Seventh International Conference on Ubiquitous Computing*, pages 232-250. Springer, 2005. The dissertation author was the primary investigator and author of this paper.

Chapter 6

Outlook on Mobility Research

6.1 Introduction

Our discussion thus far has demonstrated the power of context-aware computing for alleviating the burdens of mobile users. Chapter 5 showed that despite coarse-grained location detection on mobile phones, simple context-detection can empower mobile users with effective reminder capabilities to manage their complex mobile environment. Although reminders are a versatile tool, they have a specific application for information deferral that is made possible by context-awareness on mobile phones. The contribution of Place-Its would not have been possible without an appropriate understanding of the problems that mobile users face, and adequate advancement in context-aware research.

Novel contributions, such as those revealed by Place-Its, will emerge if we continue to innovate at each of the three levels of the mobile information space (Figure 1.1). The three main areas of exploration are: continued characterization of the problems mobile users face, exploration of context-aware technologies, and ubiquitous computing applications designed around alleviating the burdens of limited attentional resources. The interrelated dependence of each area can create appropriate solutions to a person's complex mobile milieu. The following sections describe future and ongoing research based on the insights we have discussed thus far.

6.2 Gathering More Data About Mobile Users

The study of information needs has typically focused on tasks that users perform on the web, or their intentions through search queries. Our mobile information needs diary study was one of the first to explore what people actually do when they are mobile. An *in situ* method produces real data based on a person's experience throughout the day. One of the drawbacks of this approach is that no one was physically present to observe the participants in their mobile environment. Due to the self-reporting nature of a diary study, information needs can go unreported.

There are many information needs a mobile user experiences. Sometimes these are not necessarily needs, but information gathering that is acquired through everyday experiences. Without actually observing a mobile user, it would be difficult to collect data about the latter type of information gathering. Further work in this area would involve a person following a mobile user recording their daily life and information gathering patterns. This would help complement our already rich data set, harvesting in great detail, the intentions of mobile users. Although such an approach would be burdensome, the greater body of research would benefit from having comprehensive data.

6.3 The Need for Just-In-Time Delivery

We have discussed a variety of ways to address the needs of mobile users. The applications described thus far provide tools for information deferral; a process of addressing an information need later at a more appropriate time. But what if the need could be addressed in the moment with smart, context-aware, technological assistance?

From our diary study, we observed that 72% of the diary entries were context driven, implying that context-awareness could provide ways for automatic information delivery based on contextual cues. This type of system could make mobile phones more usable by automatically inferring one's mobile information needs and making the data available for just-in-time access. An automated retrieval sys-

tem does not require extensive interaction, which is an essential requirement for mobile users. Our goal is to use a person's immediate context to display information he may want based on previous interactions with his device. Opportunistically displaying information alleviates the burden of searching mounds of data, waiting for long network times, or browsing through screenfuls of information. In addition to contextual information, we can also draw upon a person's personal data stores to gain insight into relevant types of information. Examples of this data store as seen in Chapter 2 include one's calendar, email, browser history, or chat logs.

Our personal information is often spread across multiple devices to the point that we are always trying to synchronize information amongst our phone, desktop, and laptop. A holistic approach to synchronization would allow our mobile device to use activities and data from our desktop, and vice versa. Holistic synchronization would allow for a richer understanding of a mobile user's intentions based on past digital activity. This would allow a system to opportunistically deliver information in a just-in-time manner that could drastically change the way mobile users interact with their devices.

We should mention that we do not expect mobile users to always function in a cause-effect relationship, where a system can completely predict what a mobile user may want. The decisions that people make are not always rational and formulated in a cause-effect way, but are frequently littered with preferences and other imprecise factors [Mar87]. A mobile system can be designed for the precise, static decisions that mobile users make, but trying to create a system for the imprecise, changing decisions of mobile users is difficult [Mar87]. Given these ambiguities, we still believe that a system that can present opportunistic information is still beneficial.

6.4 Mobile Digital Paper

The Place-Its reminder application described in Chapter 5 is an example of a mobile tool that helps users defer their information needs for a more appropriate time. The benefits of Place-Its are that reminders can be ubiquitously created

and delivered. Although the ability to create reminders might always be available, there may be situations and social settings where interacting with a phone may not be appropriate. For example, looking up times for the latest movie would not be appropriate during an important business meeting.

The key factor for determining social acceptance is activity and attention span. Even though many social standards have changed over the years (e.g., laptops in meetings), there are still situations where one's attention is fully required. Fully engaged with another task can often be considered rude or inappropriate.

In Chapter 2 we learned that social considerations was one reason why information needs were deferred. Even interacting with a mobile phone can sometimes be deemed inappropriate, or being divided in attention. As a result, mobile users need a different type of application that provides social acceptance, but enables them to defer and address needs when more appropriate. Our goal is to create an application that can quickly capture a person's information need without detaching them from their current situation for an extended period of time.

6.4.1 Motivation

How do you create a socially acceptable method for information deferral? In order to conform to social norms we need to use methods that are already acceptable in order to build a technological tool. Paper and pen are one of the most universally accepted methods of jotting down information for a later point in time. These types of quick personal notes that we write to ourselves are traditionally called micronotes [LLK04]. Micronotes are different from formal notetaking such as in a classroom lecture. In contrast, micronotes are focused on present information and its future use. For instance, instead of looking up movie times during a meeting, one would jot down a note to look up movie times and complete the action after the meeting. These notes are quick to write because they can be written on any surface (e.g., scrap paper, napkin). Writing is also more socially acceptable than typing on a laptop or phone, which makes it convenient to do for those quick, passing thoughts. Micronotes that are utilized for information deferral purposes often require a two-step method. The first is to write the note for a later

point in time. The second step is to act upon the note, such as searching for movie times as seen in our example. This is an area for improvement and technological assistance. A paper micronote may get lost, never be acted upon, or require further extensive interaction. Imagine leaving a business meeting to another meeting while trying to address your information need through a phone interface. This may not always be the case, but the rush of things could demand cumbersome interaction.

Bridging the physical and digital divide could address some of these problems. A digital copy could be a backup for a physical note in case it gets lost. Automatic actions based on a micronote would address issues of notes never getting acted upon or requiring additional interaction. These digital enhancements are complementary to the benefits of paper notes; the benefits do not disturb the current way people work (i.e., the physical copy and workflow are still available). The added benefit of mobile digital paper would potentially improve task completion for many mobile users.

We have designed a mobile paper application that utilizes the advantages of digital integration for automating information seeking. Instead of having the user write the note, then act upon it later, our system uses the note as a way of understanding what type of information the person wants. This helps alleviate some of the burden of trying to seek for information multiple times. Instead of a two step process, everything can be done in one step. The additional advantage of automatic retrieval is that it can be fused with additional context information that a device, such as a mobile phone, can sense. The context can help provide a richer set of results depending on the type of information desired.

6.4.2 Scenario

Here is a scenario of how one would use a mobile digital paper system.

George is in a meeting when he thinks about watching 'The Simpsons' with his brother tonight. He wants to know movie times, but doesn't want to be rude and search for them on his phone during the meeting. George writes 'The Simpsons' on his notepad and makes a search gesture. The pen recognizes the gesture and communicates with his phone. The phone begins to search for movie times with input from the pen

and other additional contextual information. After the meeting, George looks at his phone and sees movie times for the 5 closest theaters. He calls his brother and arranges a time to watch 'The Simpsons'.

This scenario demonstrates how George overcomes several barriers. First, he doesn't need to enter text on his phone because he can write with pen and paper. Second, the method maintains social acceptance because George does not need to turn his attention from the meeting for very long. Finally, context-sensing on the phone can present results that are relevant for George without much interaction. If any of these results were not acceptable for George, he could still refer to his paper note and act upon it as if the digital enhancements were not available.

6.4.3 Bridging the Paper to Digital Divide

Paper has great affordances for writing notes to defer information needs. Any scrap piece of paper can be used to create a micronote for future reference; an old receipt lying around, or a paper napkin. With the advent of better computing resources many have thought that paper would eventually become antiquated. However, as Sellen and Harper have pointed out, paper continues to thrive in the workplace, and is still a popular medium for many [SH03]. Although many of our paper items are moving towards digital versions, there are still great advantages to paper. Guimbretiere describes three fundamental affordances that make paper advantageous. Paper is easy to carry around and annotate, easy to navigate using tactile input, and provides a large inexpensive display surface only limited by our physical surfaces [Gui03]. The Paper Augmented Digital Document (PADD) system takes advantage of these affordances to enable people to annotate and use paper, while also having the archival and dynamic benefits of a digital version.

The PaperToolkit is another system that combines pen and paper with digital interaction [YKP07]. Unlike the PADD system, which can only deal with offline interactions, the PaperToolkit integrates real-time input handling to capture gestures and recognize handwriting. Similarly, PapierCraft is a command system built on PADD that can capture annotations and commands to be executed on digital documents [LGH05]. Using PapierCraft, users can use pen strokes to copy,

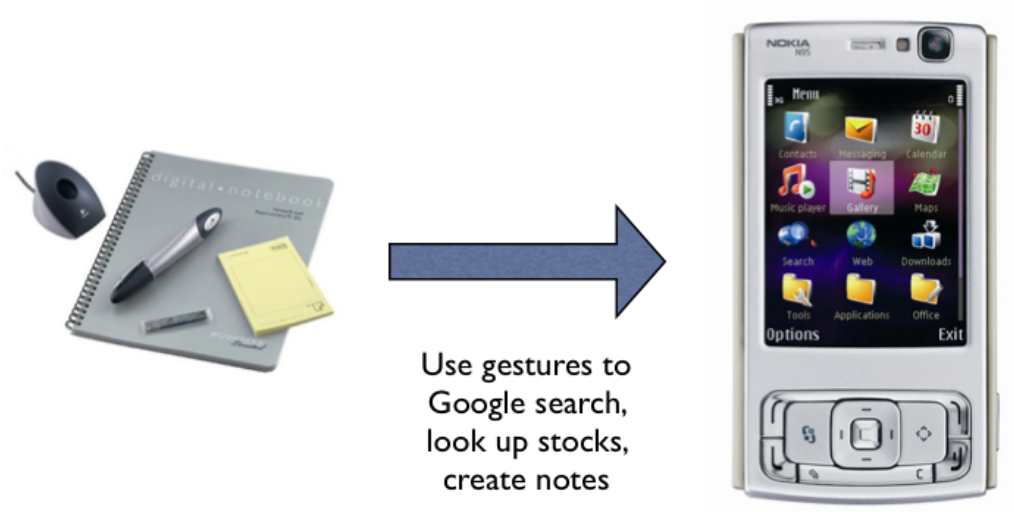


Figure 6.1: The mobile digital paper system. Anoto pen annotations on anoto-readable paper are streamed over bluetooth to a mobile phone. Example gestures include web searches or creating notes. The phone uses the data along with contextual information to fetch relevant information that a user can look at whenever time permits.

delete, paste, and hyperlink text on Anoto-enhanced paper.

Our system is implemented based on the aforementioned systems. Unlike the previous systems, our focus is on using pen and paper in a mobile setting, and exploring the types of interactions that result.

6.4.4 System Architecture

We designed a system that utilizes both digital paper and mobile phone technology to provide easy, socially permissible information deferral. The systems relies on three components: anoto pen/paper, the mobile phone, and context-sensing. These three core components enable the information deferral process.

Anoto Pen and Paper

Anoto pens are writing instruments that combine the paper and digital world. They are designed to capture input from special addressed paper through an

infrared camera built into the phone. The camera tracks coordinates on the anoto-readable paper and stores the information for future batch access, such as when docked with a computer. Several anoto pens are equipped with a bluetooth radio that can stream the data in real time to another device (e.g., phone, laptop).¹ The pen delivers information about the location, force, and time of each stroke. These low-level events can be combined to detect longer strokes that may be recognized as gestures or written words.

One of the difficulties with using Anoto technology is that it will only recognize events on anoto-readable paper. We do not assume that anoto-paper will one day be dominant throughout the world, but do believe that users of this technology can carry a notebook of anoto paper. Figure 6.1 shows an anoto pen with a digital notebook and sticky pad. These items can be used with or without anoto pens, making them advantageous for users of our system. Moreover, the cost of these alternative writing surfaces are only slightly more than non-anoto enabled paper.

Phone

The mobile phone is the second component of our system and it's purpose is to capture and interpret low-level pen data. This device could be a laptop, however we chose a phone for our system because it is the most ubiquitous mobile platform available. Phones are also highly capable devices that boast fast computation and large storage. Our implementation is built on the Nokia N95 mobile phone running Java Microedition (J2ME).

The phone first advertises a bluetooth service labeled 'ANOTOSTREAMING'. When an anoto pen is turned on, it begins to search for nearby devices with this string in the service name. Once found, the pen connects to the device and begins streaming pen data. The phone captures the pen data, which is delivered as coordinates on a piece of paper, and translates them into strokes. This clustering process involves looking at the timestamp of the movements, and how close they

¹A bluetooth-enabled anoto pen will pair with any nearby device that advertises a service labeled "ANOTOSTREAMING".

are together.²

Once these strokes are constructed they are passed to a recognizer for gesture and handwriting recognition. Our system uses the \$1 gesture recognizer proposed by Wobbrock et. al., which they showed can achieve a 97% recognition rate with only one template defined per gesture [WWL07]. The gesture recognizer also acts as a handwriting recognition system defined only for simple letters. Other handwriting recognition software packages, such as offered by the Microsoft Tablet API, could yield better results but were not included in order to enable the system to work on any java-enabled phone.³

Gestures act as commands for what actions the system should perform. We can map gestures to the common types of information needs people have as seen in Chapter 2. Examples of these gestures would be web searches, looking up sports/news/stock information, or accessing emails to name a few. The user is in control to provide the mobile device with direction on the actions it should perform. This also means that the user can define what types of data stores to look at. In our diary study we saw that 38% of the entries were related to personal data stores. Commands from the user could aid the system to find information in the right data store (e.g., email). We have implemented a simple search gesture that displays results from the Google search engine. A mobile user can write several keywords with their anoto pen, perform a search gesture, and see results on their mobile device when it is convenient. Additional gestures can easily be mapped to commands for future functionality.

Fusing With Context

Mobile phones offer the opportunity to add context information to any application. As we have seen with Place-Its, location is a powerful form of context that can enhance the opportunities for a person to defer their information needs. In our mobile digital paper application, we use context information to enhance

²More details are available in [LGH05].

³Both the PaperToolkit and Papiercraft offer solutions to this by sending pen strokes through a Windows machine that runs a recognition proxy. Although this is possible, we chose to keep all components on the mobile device.

gestures in order to deliver relevant information. This eliminates additional interaction that may be required by the user to gather the desired information. For example, instead of searching just for *The Simpsons*, we can now search *The Simpsons* looking for nearby theater listings. The additional context information creates a refined search relevant to the user.

The built-in GPS on the Nokia N95 reports coordinate information to perform location-based searches. Once a gesture is recognized by the phone, the command is fused with location information before being executed. Location is the only form of context currently supported, but activity, identity, and item could all be incorporated easily. For example, different gestures could map to different commands and specify what type of context to use.

6.4.5 Lessons Learned and Future Evaluation

The mobile paper system is a promising way to capture and defer information needs. We learned that paper is a natural input modality to a mobile device, and recognizing input can be done in realtime. However, our scenario assumes a system that correctly recognizes input each time and does not provide a feedback mechanism to the user. In a real use scenario, a mobile user would need to know if the phone correctly recognized his input text. Utilizing appropriate forms of context is also essential to identifying the information a mobile user may want. Gestures help reduce the search space with specific commands from the user, but we still need to explore which types of context are the most beneficial. In future work, we would like to explore these concerns through an evaluation of the mobile digital paper tool in a real world context. We plan to observe how people use the system, how others view its social appropriateness, and whether it serves as an effective information deferral application.

6.5 Conclusion

The possibilities of using context to alleviate the burdens felt by mobile users is promising. We have already seen how an ethnographic approach along

with context innovation can create powerful mobile tools. This chapter builds on these contributions by suggesting several avenues for future research. Each of these avenues will further our knowledge about the problems mobile users face, and how to build context-aware systems to address these problems.

Chapter 7

Conclusion

Our personal mobile experience identifies many of the problems that mobile users face today. Our high information needs while mobile place extreme demands on our lives. Although we have a high set of information needs even when we are not mobile, the constrained resources of mobility and crude tools make it difficult to deal with our information needs. The most significant, fundamental constrained resource is our attentional resources. Mobile users are often involved in a main task and trying to address their information need at the same time. This leads to a complex multitasking burden that a person must balance delicately. Since a person's attentional resources are inherently limited, change must occur on the technology side. Mobile tools need to be designed around multitasking, and provide ample methods for handling a dynamic, complex environment.

The best available device to create these tools is the mobile phone. Mobile phones provide storage and access to a variety of personal information stores (e.g., mail, calendar, web) that mobile users want. In addition, mobile phones are capable of sensing a variety contexts that can help anticipate and index relevant information. Mobile phones are limited though in several ways: small screens and keypads, slow internet connections, and hard to use devices. The problems with the device moves the burden of accessing information from the mobile phone to the human being. This leads to the fundamental problem of mobile device interaction.

This thesis has demonstrated a way to alleviate the burdens that mobile users feel. We have shown that despite the low fidelity of sensors and actuators on

mobile phones, they can still be appropriated for detecting useful pieces of context that can lower the burden many mobile users feel. Using context to assist mobile users is the effective method to change a constrained device, such as the mobile phone, into a powerful tool. We showed that with advances in our understanding of mobile users and in context-aware technologies, we can create ubiquitous computing applications for information deferral - capturing information needs when they occur, and using contextual support to address them later.

7.1 Findings

We have presented a structure to the mobile information space broken down into three areas that are essential for addressing the needs of mobile users. The three areas we discussed were the problems that mobile users face, technology that can help address some of those problems, and applications that use the technology in order to provide the right tools for mobile users. These three areas help summarize our findings.

- We had many intuitions about the current problems mobile users face, but lacked empirical data to support them. Our mobile information needs diary study helped provide the necessary evidence to support claims about what mobile users need. We uncovered a broad taxonomy of the types of needs mobile users have, as well as the methods they use to address those needs. Despite popular belief that information access through the mobile internet will solve many problems for mobile users, the majority of diary entries indicated that there were other inhibiting factors that prevented them from addressing their needs.

We found that the fundamental problem for mobile users is being engaged in another task while trying to gather information. This is the main problem for us to address through technological support. One method to address this problem is through information deferral, a concept rooted in the fact that mobile users are often busy with other tasks while mobile.

- In order to support information deferral, we need context-sensing capabilities on mobile phones. However, the mobile phone is not equipped with specific sensors and actuators for this purpose. We found that even though its sensors are not designed for sensing context, mobile phones can still be used to detect location and high-level mobility patterns. We used machine learning techniques to analyze fluctuations in GSM signal patterns and observed that changes in signal space are related to movement in physical space. Our methods were able to detect whether a person is walking, driving, or stationary with 85% accuracy. These coarse-grained detection techniques are adequate for a large class of applications including social-mobile applications, detecting appropriate times for information delivery, and creating ways for mobile users to defer their information needs.
- Our key finding was that we can use context-awareness to alleviate the burdens of mobile users. We demonstrated this through the Place-Its application that uses coarse-grained GSM location detection to create location-based reminders. Reminders are used in many situations, but Place-Its was the first system to offer a ubiquitous reminder system that can be used everywhere a person travels to. Place-Its demonstrated the power of information deferral with contextual support. Mobile users were able to defer needs from the time they were experienced, and used their knowledge about their locations to determine when a more appropriate time would be.

7.2 Final Thoughts

The mobile information space is filled with opportunities for exploration. This thesis has helped lay the groundwork in providing structure to this space, and demonstrating that contextual support is a key requirement for addressing the needs of mobile users. Chapter 6 described the outlook on the future of mobility research, and ways to complement the work presented in this dissertation. It is imperative to continue to build supportive tools and applications that are grounded in empirical user data about the problems that mobile users face. We anticipate

these problems evolving over the years, which demands a continued understanding about the landscape of the mobile world. However, the fundamental problem that is grounded in unchanging human nature is the lack of attentional resources. As we look forward to moving the burdens of mobile access from the user to the mobile device, we can begin to further alleviate this burden. The future of ubiquitous computing is promising because our mobile phones offer the opportunity to change how we interact with our world.

Bibliography

- [AAH⁺97] Gregory D. Abowd, Christopher G. Atkeson, Jason Hong, Sue Long, Rob Kooper, and Mike Pinkerton. Cyberguide: a mobile context-aware tour guide. *Wirel. Netw.*, 3(5):421–433, 1997.
- [ACM⁺02] Yariv Aridor, David Carmel, Yoelle S. Maarek, Aya Soffer, and Ronny Lempel. Knowledge encapsulation for focused search from pervasive devices. *ACM Trans. Inf. Syst.*, 20(1):25–46, 2002.
- [AM06] Ian Anderson and Henk Muller. Exploring gsm signal strength levels in pervasive environments. *aina*, 2:87–91, 2006.
- [AMS⁺07] Ian Anderson, Julie Maitland, Scott Sherwood, Louise Barkhuus, Matthew Chalmers, Malcolm Hall, Barry Brown, and Henk Muller. Shakra: tracking and sharing daily activity levels with unaugmented mobile phones. *Mobile Network and Applications*, 12(2-3):185–199, 2007.
- [ASS04] Stavros Antifakos, Adrian Schwaninger, and Bernt Schiele. Evaluating the effects of displaying uncertainty in context-aware applications. In Nigel Davies, Elizabeth D. Mynatt, and Itiro Siio, editors, *Ubicomp*, volume 3205 of *Lecture Notes in Computer Science*, pages 54–69. Springer, 2004.
- [BBC⁺08] Victoria Bellotti, Bo Begole, Ed H. Chi, Nicolas Ducheneaut, Ji Fang, Ellen Isaacs, Tracy King, Mark W. Newman, Kurt Partridge, Bob Price, Paul Rasmussen, Michael Roberts, Diane J. Schiano, and Alan Walendowski. Activity-based serendipitous recommendations with the magitti mobile leisure guide. In *CHI '08: Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, pages 1157–1166, New York, NY, USA, 2008. ACM.
- [Bei00] Michael Beigl. Memoclip: A location-based remembrance appliance. *Personal Ubiquitous Comput.*, 4(4):230–233, 2000.

- [BEM95] Maria A. Brandimonte, Gilles O. Einstein, and Mark A. McDaniel, editors. *Prospective Memory: Theory and Applications*. Lawrence Erlbaum Associates, New Jersey, 1995.
- [BGMPW00] Orkut Buyukkokten, Hector Garcia-Molina, Andreas Paepcke, and Terry Winograd. Power browser: efficient web browsing for pdas. In *CHI '00: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 430–437, New York, NY, USA, 2000. ACM.
- [BHR⁺99] Staffan Björk, Lars Erik Holmquist, Johan Redström, Ivan Bretan, Rolf Danielsson, Jussi Karlgren, and Kristofer Franzén. West: a web browser for small terminals. In *UIST '99: Proceedings of the 12th annual ACM symposium on User interface software and technology*, pages 187–196, New York, NY, USA, 1999. ACM.
- [BI04] Ling Bao and Stephen S. Intille. Activity recognition from user-annotated acceleration data. In *Pervasive*, pages 1–17, 2004.
- [BP00] P. Bahl and V.N. Padmanabhan. Radar: an in-building rf-based user location and tracking system. *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, 2:775–784 vol.2, 2000.
- [Bro02] Andrei Broder. A taxonomy of web search. *SIGIR Forum*, 36(2):3–10, 2002.
- [Bro08] Karen Brown. Newer phones, older users. *Wireless Week*, 2008.
- [BWK07] Joel Brandt, Noah Weiss, and Scott R. Klemmer. txt 4 l8r: lowering the burden for diary studies under mobile conditions. In *CHI '07: CHI '07 extended abstracts on Human factors in computing systems*, pages 2303–2308, New York, NY, USA, 2007. ACM.
- [CDT00] Chun Wei Choo, Brian Detlor, and Don Turnball. Information seeking on the web: An integrated model of browsing and searching. *First Monday*, 5(2), 2000.
- [CM05] Scott Carter and Jennifer Mankoff. When participants do the capturing: the role of media in diary studies. In *CHI '05: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 899–908, New York, NY, USA, 2005. ACM.
- [CRS04a] Sunny Consolvo, Peter Roessler, and Brett E. Shelton. The carenet display: Lessons learned from an in home evaluation of an ambient display. In *Ubicomp*, pages 1–17, 2004.

- [CRS⁺04b] Sunny Consolvo, Peter Roessler, Brett E. Shelton, Anthony LaMarca, Bill Schilit, and Sara Bly. Technology for care networks of elders. *IEEE Pervasive Computing*, 3(2):22–29, 2004.
- [CSC⁺06] Mike Y. Chen, Timothy Sohn, Dmitri Chmelev, Dirk Haehnel, Jeffrey Hightower, Jeff Hughes, Anthony LaMarca, Fred Potter, Ian Smith, and Alex Varshavsky. Practical metropolitan-scale positioning for gsm phones. In *Proceedings of the Eighth International Conference on Ubiquitous Computing (UbiComp 2006)*, Lecture Notes in Computer Science, pages 225–242. Springer-Verlag, September 2006.
- [DA00] Anind K. Dey and Gregory D. Abowd. Cybreminder: A context-aware system for supporting reminders. In *HUC '00: Proceedings of the 2nd international symposium on Handheld and Ubiquitous Computing*, pages 172–186, London, UK, 2000. Springer-Verlag.
- [DPC03] Richard W. DeVaul, Alex "Sandy" Pentland, and Vicka R. Corey. The memory glasses: Subliminal vs. overt memory support with imperfect information. In *ISWC '03: Proceedings of the 7th IEEE International Symposium on Wearable Computers*, page 146, Washington, DC, USA, 2003. IEEE Computer Society.
- [DSA01] Anind K. Dey, Daniel Salber, and Gregory D. Abowd. A conceptual framework and a toolkit for supporting the rapid prototyping of context-aware applications. *Human-Computer Interaction*, 16(2–4):97–166, 2001.
- [EM90] Giles O. Einstein and Mark A. McDaniel. Normal aging and prospective memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 16:717–726, 1990.
- [EP06] Nathan Eagle and Alex (Sandy) Pentland. Reality mining: sensing complex social systems. *Personal Ubiquitous Comput.*, 10(4):255–268, 2006.
- [FHT00] J. Friedman, T. Hastie, and R. Tibshirani. Additive Logistic Regression: a Statistical View of Boosting. *The Annals of Statistics*, 38(2), 2000.
- [FS96] Yoav Freund and Robert E. Schapire. Experiments with a new boosting algorithm. In *International Conference on Machine Learning*, pages 148–156, 1996.
- [GRCE06] Philip Garner, Omer Rashid, Paul Coulton, and Reuben Edwards. The mobile phone as a digital spraycan. In *ACE '06: Proceedings of the 2006 ACM SIGCHI international conference on Advances in*

- computer entertainment technology*, page 12, New York, NY, USA, 2006. ACM.
- [GSB⁺04] William G. Griswold, Patricia Shanahan, Steven W. Brown, Robert Boyer, Matt Ratto, R. Benjamin Shapiro, and Tan Minh Truong. Activecampus: Experiments in community-oriented ubiquitous computing. *Computer*, 37(10):73–81, 2004.
- [Gui03] François Guimbretière. Paper augmented digital documents. In *UIST '03: Proceedings of the 16th annual ACM symposium on User interface software and technology*, pages 51–60, New York, NY, USA, 2003. ACM.
- [HB01] Jeffrey Hightower and Gaetano Borriello. Location systems for ubiquitous computing. *Computer*, 34(8):57–66, 2001.
- [HB04] Jeffrey Hightower and Gaetano Borriello. Particle filters for location estimation in ubiquitous computing: A case study. In Nigel Davies, Elizabeth Mynatt, and Itiro Siiro, editors, *Proceedings of the Sixth International Conference on Ubiquitous Computing (Ubicomp 2004)*, volume 3205 of *Lecture Notes in Computer Science*, pages 88–106. Springer-Verlag, September 2004.
- [HKH05] Ramaswamy Hariharan, John Krumm, and Eric Horvitz. Web-enhanced gps. In *LoCA*, pages 95–104, 2005.
- [HL00] Rebecca Hansson and Peter Ljungstrand. The reminder bracelet: subtle notification cues for mobile devices. In *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*, pages 323–324, New York, NY, USA, 2000. ACM.
- [HR04] Vicki L. Hanson and John T. Richards. A web accessibility service: update and findings. In *Assets '04: Proceedings of the 6th international ACM SIGACCESS conference on Computers and accessibility*, pages 169–176, New York, NY, USA, 2004. ACM.
- [HSI08] Rachel Hinman, Mirjana Spasojevic, and Pekka Isomursu. They call it surfing for a reason: identifying mobile internet needs through pc internet deprivation. In *CHI '08: CHI '08 extended abstracts on Human factors in computing systems*, pages 2195–2208, New York, NY, USA, 2008. ACM.
- [ILB⁺05] Stephen S. Intille, Kent Larson, J. S. Beaudin, J. Nawyn, E. Munguia Tapia, and P. Kaushik. A living laboratory for the design and evaluation of ubiquitous computing technologies. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1941–1944, New York, NY, USA, 2005. ACM.

- [IOM05] Mizuko Ito, Daisuke Okabe, and Misa Matsuda. *Personal, Portable, Pedestrian: Mobile Phones in Japanese Life*. The MIT Press, 2005.
- [JBCJ06] Matt Jones, George Buchanan, Tzu-Chiang Cheng, and Preeti Jain. Changing the pace of search: Supporting “background” information seeking. *Journal of the American Society for Information Science and Technology*, 57(6):838–842, 2006.
- [JBS08] Bernard J. Jansen, Danielle L. Booth, and Amanda Spink. Determining the informational, navigational, and transactional intent of web queries. *Inf. Process. Manage.*, 44(3):1251–1266, 2008.
- [JSBS98] Bernard J. Jansen, Amanda Spink, Judy Bateman, and Tefko Saracevic. Real life information retrieval: a study of user queries on the web. *SIGIR Forum*, 32(1):5–17, 1998.
- [JSC07] Eric M. Jones, Ted Selker, and Hyemin Chung. What you said about where you shook your head: a hands-free implementation of a location-based notification system. In *CHI '07: CHI '07 extended abstracts on Human factors in computing systems*, pages 2477–2482, New York, NY, USA, 2007. ACM.
- [Kaa03] Eija Kaasinen. User needs for location-aware mobile services. *Personal Ubiquitous Comput.*, 7(1):70–79, 2003.
- [KB06] Maryam Kamvar and Shumeet Baluja. A large scale study of wireless search behavior: Google mobile search. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 701–709, New York, NY, USA, 2006. ACM.
- [KH04a] John Krumm and Ken Hinckley. The nearme wireless proximity server. In *UbiComp*, pages 283–300, 2004.
- [KH04b] John Krumm and Eric Horvitz. Locadio: Inferring motion and location from wi-fi signal strengths. In *MobiQuitous*, pages 4–13, 2004.
- [KKL05] Hoyoung Kim, Jinwoo Kim, and Yeonsoo Lee. An empirical study of use contexts in the mobile internet, focusing on the usability of information architecture. *Information Systems Frontiers*, 7(2):175–186, 2005.
- [KOA⁺99] Cory D. Kidd, Robert Orr, Gregory D. Abowd, Christopher G. Atkeson, Irfan A. Essa, Blair MacIntyre, Elizabeth D. Mynatt, Thad Starner, and Wendy Newstetter. The aware home: A living laboratory for ubiquitous computing research. In *CoBuild '99: Proceedings*

- of the Second International Workshop on Cooperative Buildings, Integrating Information, Organization, and Architecture*, pages 191–198, London, UK, 1999. Springer-Verlag.
- [KRA01] Naohiko Kohtake, Jun Rekimoto, and Yuichiro Anzai. Infopoint: A device that provides a uniform user interface to allow appliances to work together over a network. *Personal Ubiquitous Comput.*, 5(4):264–274, 2001.
- [KSFS05] Tim Kindberg, Mirjana Spasojevic, Rowanne Fleck, and Abigail Sellen. The ubiquitous camera: An in-depth study of camera phone use. *IEEE Pervasive Computing*, 4(2):42–50, 2005.
- [KWS07] Melanie Kellar, Carolyn Watters, and Michael Shepherd. A field study characterizing web-based information-seeking tasks. *J. Am. Soc. Inf. Sci. Technol.*, 58(7):999–1018, 2007.
- [LB05] Heidi Lam and Patrick Baudisch. Summary thumbnails: readable overviews for small screen web browsers. In *CHI*, pages 681–690, 2005.
- [LCC⁺05] Anthony LaMarca, Yatin Chawathe, Sunny Consolvo, Jeffrey Hightower, Ian E. Smith, James Scott, Timothy Sohn, James Howard, Jeff Hughes, Fred Potter, Jason Tabert, Pauline Powledge, Gaetano Borriello, and Bill N. Schilit. Place lab: Device positioning using radio beacons in the wild. In *Pervasive*, pages 116–133, 2005.
- [LCK⁺05] Jonathan Lester, Tanzeem Choudhury, Nicky Kern, Gaetano Borriello, and Blake Hannaford. A hybrid discriminative/generative approach for modeling human activities. In *IJCAI*, pages 766–772, 2005.
- [LF94] M. Lamming and M. Flynn. ‘Forget-me-not’: intimate computing in support of human memory. In *Proceedings of FRIEND21: Symposium on Next Generation Human Interfaces*, Tokyo, Japan, 1994.
- [LFK05] Lin Liao, Dieter Fox, and Henry Kautz. Location-based activity recognition using relational markov networks. In *International Conference on Artificial Intelligence*, 2005.
- [LFR⁺06] Pamela J. Ludford, Dan Frankowski, Ken Reily, Kurt Wilms, and Loren Terveen. Because i carry my cell phone anyway: functional location-based reminder applications. In *CHI ’06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 889–898, New York, NY, USA, 2006. ACM.

- [LGH05] Chunyuan Liao, François Guimbretière, and Ken Hinckley. Papercraft: a command system for interactive paper. In *UIST '05: Proceedings of the 18th annual ACM symposium on User interface software and technology*, pages 241–244, New York, NY, USA, 2005. ACM.
- [LLK04] Min Lin, Wayne G. Lutters, and Tina S. Kim. Understanding the micronote lifecycle: improving mobile support for informal note taking. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 687–694, New York, NY, USA, 2004. ACM.
- [LRT04] Kari Laasonen, Mika Raento, and Hannu Toivonen. Adaptive on-device location recognition. In *Pervasive*, pages 287–304, 2004.
- [LSG05] Kevin A. Li, Timothy Sohn, and William G. Griswold. Evaluating location-based reminders. Technical Report CS2005-0826, University of California, San Diego, 2005.
- [Mar87] James G. March. Ambiguity and accounting: The elusive link between information and decision making. *Accounting, Organizations and Society*, 12(2):153–168, March 1987.
- [MPC01] Julie B. Morrison, Peter Pirolli, and Stuart K. Card. A taxonomic analysis of what world wide web activities significantly impact people’s decisions and actions. In *CHI '01: CHI '01 extended abstracts on Human factors in computing systems*, pages 163–164, New York, NY, USA, 2001. ACM.
- [MRCJ01] Elizabeth D. Mynatt, Jim Rowan, Sarah Craighill, and Annie Jacobs. Digital family portraits: supporting peace of mind for extended family members. In *CHI '01: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 333–340, New York, NY, USA, 2001. ACM.
- [MS00] Natalia Marmasse and Chris Schmandt. Location-aware information delivery with commotion. In *HUC '00: Proceedings of the 2nd international symposium on Handheld and Ubiquitous Computing*, pages 157–171, London, UK, 2000. Springer-Verlag.
- [NC99] Neill J. Newman and Adrian F. Clark. Sulawesi: A wearable application integration framework. *iswc*, 00:170, 1999.
- [NTH06] Jennifer Ng, Jonathan Terleski, and Jason Hong. Whisper: analysis and design for a community event service. In *CHI Extended Abstracts*, pages 1151–1156, 2006.

- [OVLdL05] Veljo Otsason, Alex Varshavsky, Anthony LaMarca, and Eyal de Lara. Accurate gsm indoor localization. In *UbiComp*, pages 141–158, 2005.
- [Pas97] Jason Pascoe. The stick-e note architecture: extending the interface beyond the user. In *IUI '97: Proceedings of the 2nd international conference on Intelligent user interfaces*, pages 261–264, New York, NY, USA, 1997. ACM.
- [PCB00] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 32–43, New York, NY, USA, 2000. ACM.
- [PFP⁺04] Matthai Philipose, Kenneth P. Fishkin, Mike Perkowitz, Donald J. Patterson, Dieter Fox, Henry Kautz, and Dirk Hahnel. Inferring activities from interactions with objects. *IEEE Pervasive Computing*, 3(4):50–57, 2004.
- [PLFK03] Donald J. Patterson, Lin Liao, Dieter Fox, and Henry A. Kautz. Inferring high-level behavior from low-level sensors. In *UbiComp*, pages 73–89, 2003.
- [Pre92] William Press. *Numerical Recipes in C*. Cambridge University Press, 1992.
- [PRGH00] Vladimir Pavlovic, James M. Rehg, Ashutosh Garg, and Thomas S. Huang. Multimodal speaker detection using error feedback dynamic bayesian networks. *cvpr*, 02:2034, 2000.
- [PS02] Leysia Palen and Marilyn Salzman. Voice-mail diary studies for naturalistic data capture under mobile conditions. In *CSCW '02: Proceedings of the 2002 ACM conference on Computer supported cooperative work*, pages 87–95, New York, NY, USA, 2002. ACM.
- [Rep04] Consumer Reports. Pedometers: Walking by the numbers. *Consumer Reports*, 69:30–32, 2004.
- [Rho97] Bradley J. Rhodes. The wearable remembrance agent: a system for augmented memory. In *ISWC '97: Proceedings of the 1st IEEE International Symposium on Wearable Computers*, page 123, Washington, DC, USA, 1997. IEEE Computer Society.
- [Rho03] Bradley Rhodes. Using physical context for just-in-time information retrieval. *IEEE Trans. Comput.*, 52(8):1011–1014, 2003.

- [SB98] Holger Schwenk and Yoshua Bengio. Training methods for adaptive boosting of neural networks. In *NIPS '97: Proceedings of the 1997 conference on Advances in neural information processing systems 10*, pages 647–653, Cambridge, MA, USA, 1998. MIT Press.
- [SBC97] Ben Shneiderman, Don Byrd, and W. B Croft. Clarifying search: A user-interface framework for text searches. Technical report, University of Maryland, College Park, 1997.
- [Sch01] R. Schapire. The boosting approach to machine learning: An overview, 2001.
- [SCL⁺05] Ian E. Smith, Sunny Consolvo, Anthony LaMarca, Jeffrey Hightower, James Scott, Timothy Sohn, Jeff Hughes, Giovanni Iachello, and Gregory D. Abowd. Social disclosure of place: From location technology to communication practices. In *Pervasive*, pages 134–151, 2005.
- [SH03] Abigail J. Sellen and Richard H.R. Harper. *The Myth of the Paperless Office*. MIT Press, Cambridge, MA, USA, 2003.
- [SJWS02] Amanda Spink, Bernard J. Jansen, Dietmar Wolfram, and Tefko Saracevic. From e-sex to e-commerce: Web search changes. *Computer*, 35(3):107–109, 2002.
- [SMHM99] Craig Silverstein, Hannes Marais, Monika Henzinger, and Michael Moricz. Analysis of a very large web search engine query log. *SIGIR Forum*, 33(1):6–12, 1999.
- [Smi05] Ian Smith. Social-mobile applications. *Computer*, 38(4):84–85, 2005.
- [Tay67] Robert S. Taylor. Question-negotiation and information-seeking in libraries. *Journal of College and Research Libraries*, 29:178–194, 1967.
- [TL04] Bassett Tudor-Locke. How many steps/day are enough? preliminary pedometer indices for public health. *Journal of Sports Medicine*, 34(1):1–8, 2004.
- [TSM⁺07] Eleanor Toye, Richard Sharp, Anil Madhavapeddy, David Scott, Eben Upton, and Alan Blackwell. Interacting with mobile services: an evaluation of camera-phones and visual tags. *Personal Ubiquitous Comput.*, 11(2):97–106, 2007.
- [VJ04] Paul Viola and Michael J. Jones. Robust real-time face detection. *Int. J. Comput. Vision*, 57(2):137–154, 2004.

- [Web00] Geoffrey I. Webb. Multiboosting: A technique for combining boosting and wagging. *Machine Learning*, 40(2):159–196, 2000.
- [WF00] Ian H. Witten and Eibe Frank. *Data mining: practical machine learning tools and techniques with Java implementations*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2000.
- [WHVFG92] Roy Want, Andy Hopper, and Veronica Falc and Jonathan Gibbons. The active badge location system. *ACM Trans. Inf. Syst.*, 10(1):91–102, 1992.
- [WJH97] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office. *Personal Communications, IEEE [see also IEEE Wireless Communications]*, 4(5):42–47, Oct 1997.
- [WWL07] Jacob O. Wobbrock, Andrew D. Wilson, and Yang Li. Gestures without libraries, toolkits or training: a \$1 recognizer for user interface prototypes. In *UIST '07: Proceedings of the 20th annual ACM symposium on User interface software and technology*, pages 159–168, New York, NY, USA, 2007. ACM.
- [YKP07] Ron B. Yeh, Scott R. Klemmer, and Andreas Paepcke. Design and evaluation of an event architecture for paper uis: Developing create by copying and combining. Technical report, Stanford University Computer Science Department, 2007.
- [YZ06] Min Yin and Shumin Zhai. The benefits of augmenting telephone voice menu navigation with visual browsing and search. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 319–328, New York, NY, USA, 2006. ACM.