Understanding Users' Personal Information Management Practices Across Multiple Devices

Manas Tungare

A Research Proposal presented to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy
in
Computer Science and Applications

Committee:

Dr. Manuel A. Pérez-Quiñones, Chair
Dr. Stephen Edwards
Dr. Edward A. Fox
Prof. Steven Harrison
Dr. Tonya Smith-Jackson

May 5, 2008
Blacksburg, Virginia.

Keywords: Personal Information Management, Multi-Platform User Interfaces, Mental Workload Measurement.

Contents

Chapter 1

Introduction

1.1 Introduction

Information is being disseminated much faster than we can assimilate it. Our tools are not adapting fast enough to keep pace with the need for ubiquitous access to information. A large sector of the economy is devoted to managing information, and information overload threatens our effectiveness. Even at home, we are inundated with information as we manage an ever-increasing library of documents, to-do lists, digital music, digital photos, and others. All of this causes stress and increases mental workload as we struggle to stay in control of our information.

One of the biggest challenges of our time is to control effectively the management of personal information. We have developed amazing capabilities to record, store, and transmit massive quantities of information with minimal effort; however this has relegated us to part-time workers and part-time librarians of our own personal information. In spite of the ease of recording, creating, receiving, storing, and accumulating digital materials, it is difficult to manage and use them sensibly [Czerwinski et al., 2006]. With time, the amount of information generated by humans can only increase, while human attentional resources have remained constant [Levy, 2005].

At the same time, advances in computer hardware have led to the miniaturization of technology that places several portable information devices at our disposal. It is common for a lot of people to carry a laptop computer or a cell phone as they go about their everyday business outside the usual contexts of an office or a home [Tungare and Pérez-Quiñones, 2008], and to expect productive work output when mobile. However, the current state-of-the-art in information management solutions sends these users into a frenzy trying to locate the most current version of their slide shows, the documents they sent around for review, and the phone number of the person they need to call right now.

In information ecosystems [Tungare et al., 2006b], a user needs to focus attention on various tasks at the same time, or in quick succession. In traditional single terminal computer systems, the majority of a user's attentional and cognitive resources would be focused on the terminal while performing a specific task. However, in an environment where multiple devices require intermittent attention and present useful information at unexpected times, the user is subjected to different mental workloads.

This proposal involves work at the intersection of three areas in HCI: personal information management (§1.1.1), multiple devices (§1.1.2), and mental workload measurement (§1.1.3).

1.1.1 Personal Information Management

Given that one's personal information exists in a continuum, and that it spans all aspects of one's life, deriving a comprehensive definition for it is challenging. [Jones and Teevan, 2007] cover the salient aspects of personal information by defining it as (1) the information a person keeps for personal use (e.g. files), (2) information about a person but kept by and under the control of others (e.g. health information), (3) information experienced by a person but outside the person's control (e.g. books or Web sites browsed), and (4) information directed to a person (e.g. email).

Various aspects of personal information management have been studied in the literature, as detailed in chapter 2. They include studies of various types of personal information, approaches and user traits (pilers versus filers [Malone, 1983], browsing versus searching [Teevan et al., 2004], etc.) and cross-project information management [Boardman et al., 2003, Bergman et al., 2006]. While some attention has been given to developing tools to assist users in personal information management, researchers have identified several problems in evaluating these tools.

The chief problem in evaluating PIM tools or systems is that personal information is, by definition, personal. Thus, it is difficult, or close to impossible to develop reference tasks that can be used by multiple users to test multiple tools and approaches. This also makes it difficult to compare the results obtained by diverse research labs while evaluating their own tools. There is a pronounced lack of measurement techniques that are known to work across tasks, across tools, and across experiments [Teevan and Jones, 2008].

1.1.2 Multiple Devices in Personal Information Management

The proliferation of portable information devices in recent years and its impact on how users manage their information has not adequately been addressed by current research — neither by PIM researchers, nor by those working in the area of multiple devices.

Though the research discipline of multi-platform user interfaces has extensively studied how applications may be written to run on many platforms [Thevenin and Coutaz, 1999, Florins and Vanderdonckt, 2004, Denis and Karsenty, 2004], not much work has focused on understanding how users access or manage their information across multiple devices. Research in this area has followed a task-oriented approach rather than an information-oriented path. The importance of following an information-oriented approach has been well highlighted [Fidel and Pejtersen, 2004].

The impact of such multiple devices on personal information management is more than that of the individual devices alone. When used together, e.g. at a desk, these devices compete for a user's attention, and require valuable mental resources to be attended to. The influence of a multi-device environment on the user's mental workload, and how it affects operator performance under these conditions has not been studied in detail.

Modern work environments consist of several information devices that operate together, in unison, to help the user perform his/her information tasks. In a way, these devices are the analogues of various organisms that constitute a biological ecosystem. Specifically, we defined a 'personal information ecosystem' as 'a system of devices and applications that are present in the information environment of a user, that interact closely and richly with one another, to help the user achieve the goal of fulfilling his/her information needs.' [Tungare et al., 2006b]

1.1.3 Mental Workload Measurement

While operator performance in a particular task situation can be measured directly by performance metrics, (e.g. the time taken to perform an experimental task, or the number and severity of errors in task performance), they cannot be used to predict performance for an unknown task [Wilson and Eggemeier, 2006]. Subjective workload assessment techniques such as NASA Task Load Index (NASA TLX) [Hart and Staveland, 1988a], the Subjective Workload Assessment Technique, [Reid et al., 1982], and Workload Profile [Tsang and Velazquez, 1996] are used to provide an estimate of mental overload. It is generally assumed that workload is related to operator performance such that low to moderate levels of workload are associated with acceptable levels of operator

performance [Wilson and Eggemeier, 2006]. Since the ratings obtained via subjective workload assessments are not task-specific, it is possible to use them to compare the workload imposed by different tasks.

Performance measurements using the NASA TLX have been conducted in airline cockpits [Ballas et al., 1992], navigation [Schryver, 1994], and in the medical field [Bertram et al., 1992]. However, there has been no study of operator workload in modern computing environments such as information ecosystems.

1.2 Problem Statement

A deeper understanding of mental workload in information ecosystems can help the design of these systems at a holistic level. In several task scenarios, workload assessments such as NASA TLX have been administered instead of direct measurement of performance metrics for several reasons: chief among them is that subjective workload assessments require less effort and instrumentation of the task, and are easier to administer. The results can be treated as valid indicators of performance metrics, because in several task scenarios, mental workload (as measured by these assessments) has consistently been shown to be negatively correlated with performance metrics [Bertram et al., 1992].

Our findings from Phase I of this research (detailed later in §3.2) indicate that, in several scenarios involving information ecosystems, users encounter difficulties in performing certain tasks. What is the mental workload incurred in these situations, identified as difficult or frustrating by users in our survey? (§3.3.1)

Users often undertake several workarounds to perform their jobs effectively. For example, because of the lack of effective synchronization between two computers, users stop using all but one of their multiple computers. Are these strategies motivated by a desire to lower their mental workload (although they may not know it at the time)? Several users also reported that they encountered errors in managing their personal information; are these signs of increased mental workload at levels that users are unable to cope with? (§3.3.2)

While the NASA TLX scale [Hart and Staveland, 1988a, Hart and Staveland, 1988b] has been demonstrated in several domains as being an accurate predictor of operator performance, none of the scenarios involved multiple computing devices that all have the capability of interrupting the user, and are used together to perform a higher-level task (i.e., personal information management).

I intend to investigate the question of whether multi-dimensional subjective workload assessment techniques such as NASA TLX are an accurate indicator of operator performance in certain instances of information ecosystems. (§3.3.3)

When the NASA TLX scale was first designed, Hart and Staveland had considered three additional dimensions, beyond the six that currently are part of the scale. They found that some of the dimensions in their experiments either did not correlate strongly, or were found to be insensitive to manipulations. Thus they were dropped or merged with other dimensions, and the scale now contains only six dimensions. As part of my research, I wish to explore whether all six dimensions of the NASA TLX scale are significant in explaining task performance in information ecosystems, or if some of them are less relevant than others to task performance correlation. Finally, there might be other dimensions (not currently a part of the NASA TLX scale) that might be better predictors of task performance.

1.3 Goals and Key Contributions

With the proposed work, I hope to answer questions about the impact of an information ecosystem on users' mental workload. Specifically, I believe that the contributions will aid three communities: users, designers, and researchers.

1.3.1 Contributions to Users

If the NASA TLX scale can be used in information ecosystems, it provides not only a standard measure for assessing workload for the specific task, but also makes it possible to compare workloads across a diverse set of tasks, across a diverse set of environments using a standard scale. Such measurements can be used to compare different configurations of information ecosystems to cause the least amount of mental workload for users, thereby maximizing productivity.

If we find that users tend to prefer workflows that reduce their mental workload, then this provides a clear method for users to evaluate multiple systems and pick the one that will work best for them. This role can also be performed by consultants and facilitators acting on behalf of the users.

1.3.2 Contributions to Designers

Better and deeper knowledge of mental workloads in information ecosystems can provide valuable formative feedback to designers, and assist them in creating systems that take these factors into account. Thus, this research can be used by designers to incorporate elements into their designs that actively aim to reduce mental workload for the operator.

E.g. systems such as Syncables, mentioned earlier [Tungare et al., 2007] were designed to reduce the workload on users by automatic support for migrating task-related data between two or among several devices in a user's computing environment. Although this was done before the development of a metric for such measurements, having a reliably established scale to measure workload and operator performance can be used to compare several such approaches to design.

1.3.3 Contributions to Researchers

Previous research in personal information management and multiple devices has not addressed the question of mental workload in performing these tasks, so this research will contribute to the field by establishing whether or not the NASA Task Load Index (TLX) can be used as an adequate metric for predicting operator performance in these environments.

This may be used to address the problems identified by the PIM research community, e.g. at the PIM 2008 workshop, those of evaluating PIM tools, and comparing results from diverse research laboratories on a single scale.

1.4 Overview of the Proposal

In this section, I introduced briefly the areas of personal information management (§1.1.1), multi-platform user interfaces (§1.1.2) and measurement of mental workload (§1.1.3). In chapter 2, I discuss the relevant literature from these areas in more detail. While this chapter introduced the problem statement in brief, chapter 3 takes a closer look at the problem definition, including work already performed (§3.2) and subsequently planned (§3.3). Chapter 4 outlines the schedule for performing the rest of the work, which consists of content analysis of data gathered during Phase I (§3.2), and experiments to be conducted as part of Phase II (§3.3).

Chapter 2

Literature Review

The work proposed here is at the intersection of personal information management, multi-device user interfaces and mental workload measurement. The chief goal of my research, as mentioned in chapter 1, is to test the validity of the various dimensions of the NASA TLX subjective workload assessment technique in the domain of personal information management performed using multiple devices. This chapter expands upon the literature and past work in the areas of personal information management (§2.2), multiple device user interfaces (§2.3), and mental workload measurement (§2.4) in preparation for the methodology to be discussed in chapter 3.

2.1 Introduction

Managing information and being able to access it whenever and wherever necessary has been a concern for mankind since long before computers arrived on the scene. The idea of digital information management can be ascribed to the vision of Dr. Vannevar Bush, as early as 1945. In his essay, "As we may think" [Bush, 1945], he led the foundation for some of the ideas that are just being realized recently. He described his vision of a Memex as a device "in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility." Information management at the office has been studied in detail at least since 1983 [Malone, 1983], and the distinction between personal information versus other kinds of information was drawn as early as 1988 [Lansdale, 1988]: "This is personal information not necessarily in the sense that it is private, but that we have it for our own use. We own it, and would feel deprived if it were taken away."

2.2 Personal Information Management

Since managing one's information is intrinsically personal, there exist a wide variety of practices dictated by internal as well as external factors such as preference, context, training, etc. Many previous studies have investigated the nature of such individual personal information practices. Malone [Malone, 1983], studied people's offices and desks before there was a computer on every desk. His study was an exploratory observation of the information organization practices of users.

With the advent of computers came the information revolution, and pretty soon, every user had turned into a part-time librarian. In the following section, we describe the various research efforts in PIM today, and the breadth of issues within the field of inquiry that are under active research.

2.2.1 Problems and Research Issues in PIM

If tasked with stating the goal of research in personal information information in brief, it would be to be able to recall the exact information that a user needs, at the exact moment and in the exact form she needs it, while retaining only that information which will be called up in the future. Of course, since such an ideal PIM system requires knowledge about the future, it is impossible to create. Researchers thus study how the problem can be optimized such that the information can be called up as fast as possible, as accurately as possible, and by needing to preserve as little extraneous information as possible.

A persistent theme of PIM research has been to examine the information in detail, but not necessarily the task of managing this information. There are several systems and studies that describe in detail how information may be captured, organized, archived and accessed in ways that make it more accessible and faster. However, existing research has not examined in detail the mental and cognitive demands placed on the information manager under these circumstances.

The next few sections discuss a few of the important themes in personal information management research.

2.2.2 Information Overload

Information overload is defined as occurring when the information processing demands on an individual's time to perform interactions and internal calculations exceed the supply or capacity of

time available for such processing [Schick et al., 1990].

As the problem of information overload has worsened over the years, human attentional resources have stayed constant [Farhoomand and Drury, 2002, Levy, 2005]. In fact, they have probably decreased because of the increasingly busy lifestyles of today. The greater the volume of information, the more we spend our resources on determining if a particular piece of information is useful. Subsequently, we spend fewer resources on actually assimilating and using that information for productive work. The situation gets worse when we take into account information fragmentation.

2.2.3 Information Fragmentation

Information fragmentation is the condition of having a user's data in different formats, distributed across multiple locations, manipulated by different applications, and residing in a generally disconnected manner [Bergman et al., 2006]. It is also referred to as 'compartmentalization of information' [Bellotti and Smith, 2000].

Bergman *et al.* [Bergman et al., 2006] describe the case of information fragmentation for a chemistry student, Jane, who has her chemistry project-related data in three different formats under three different hierarchies: documents, emails, and bookmarks. A direct consequence of such information fragmentation is that when Jane needs to work on her chemistry project, she needs to use three different applications to deal with three different sources of information, each existing in a different format, with inherently different types, and situated in different contexts.

The original definition of information fragmentation [Bergman et al., 2006] only included the fragmentation of information across different collections, e.g. files, email messages, and bookmarks all seemed to be managed within similar, yet duplicate, hierarchies [Boardman et al., 2003]. The issue of information fragmentation across multiple devices has been reported [Karger and Jones, 2006] and preliminary suggestions to solve it have been provided, but this problem has not yet been studied in depth to ascertain exactly how and where the fragments of this information live.

2.2.4 Studying Information Collections

Bellotti *et al.* [Bellotti and Smith, 2000] noted that information is managed in various collections independently, e.g. email, or documents, or bookmarks. Many of the studies that have been conducted in the area of Personal Information Management (details in [Teevan et al., 2007]) are limited

to how we manage information on a particular device (e.g. desktop), or how we manage a particular type of information (e.g. bookmarks or emails). Among the studies focused on a single information collection in isolation are a few notable examples:

Email

Whittaker *et al.* [Whittaker and Sidner, 1996] has looked at the problem of email overload in detail. Gwizdka [Gwizdka, 2000, Gwizdka, 2002, Gwizdka, 2004] contends that email is being used to perform functions that email systems were not explicitly designed to handle. Similarly, Mackay [Mackay, 1988], Ducheneaut [Ducheneaut and Bellotti, 2001], and others have looked at the many overloaded functions for which inboxes are being used. The increasingly-common use of email as a task management tool [Ducheneaut and Bellotti, 2001] has given rise to various strategies in users to manage this overload. Gwizdka [Gwizdka, 2004] examined the different management styles used for email, focusing on two types of factors affecting this: (1) internal factors, such as cognitive abilities, personality, gender, age, and email experience, and (2) external factors, such as organizational and cultural factors. He concluded that there are two groups of email users, the cleaners and the keepers, and attributed the differences between these two user groups to differences in flexibility of closure and in email experience.

Bookmarks

Abrams *et al.* [Abrams et al., 1998] studied the practice of users creating bookmarks to carve their own personal information space out as a subset of the entire World Wide Web. They probed the reasons behind creating bookmarks, how they were organized and maintained, and later retrieved. Jones *et al.* [Jones et al., 2002] studied the larger problem of how users organize web information for re-use (which involved bookmarking as well as various other techniques). Kelly *et al.* [Kelly and Teevan, 2003] performed longitudinal studies to understand users' web browsing behavior and relevance feedback.

Documents and Files

Barreau [Barreau, 1995, Barreau and Nardi, 1995] studied the contextual aspects of a person's work environment that guide the acquisition, classification, maintenance, and retrieval of documents. Her study highlights that document attributes are not the only markers that guide PIM activi-

ties; that context plays an important role in users' decisions to keep and maintain their personal information collections.

2.2.5 Studies Across Collections

Amidst these narrow studies of specific collections, a notable exception is Boardman's cross-tool study of collections, [Boardman et al., 2003], which revealed similarities in the ways we manage disparate information collections.

2.2.6 Finding Information

Teevan *et al.* [Teevan et al., 2004] studied users' strategies in locating their information, noting that users often navigated to their information in small steps instead of teleporting to it via tools such as search engines. Their study was more focused on information retrieval than organization and management, and was not restricted to personal information.

2.2.7 Tools

Various studies of personal information management have resulted in the development of tools to assist users in PIM tasks. These include various commercial offerings, as well as those built by the research community on the foundations of the findings from prior studies. Stuff I've Seen [Dumais et al., 2003], Bifrost Inbox Organizer [Bälter and Sidner, 2002], Taskmaster [Bellotti et al., 2003] are just a few examples of tools that emerged from research in Personal Information Management.

2.2.8 Difficulties in Studying Personal Information Management Practices

It is important to note some of the issues that have been reported about as difficulties in studying a phenomenon as context-dependent as personal information management. Each person works and manages personal information in his/her personalized context, thus it may prove difficult for a researcher to get accurate responses in an experimental setup that lacks the veracity of an established work context such as an office or a study desk.

Kelly [Kelly, 2006] discusses the several methodological challenges encountered in designing studies of personal information management practices. PIM tasks are often performed at unpredictable

times in response to a need for information. In her words, "PIM encompasses a range of activities and tools; understanding just one of them provides only a partial picture of what users want to accomplish and how they might accomplish it." The unique situational aspect of the working environment makes it difficult to study PIM as compared to general information storage and retrieval (ISAR) systems [Barreau, 1995].

Teevan *et al.* [Teevan et al., 2007] provide a categorization of types of studies performed in PIM: they include observational studies [Jones et al., 2001], interviews [Malone, 1983, Whittaker and Hirschberg, 2001, Teevan et al., 2004, Barreau and Nardi, 1995, Jones et al., 2001], surveys and questionnaires [Whittaker and Hirschberg, 2001, Gwizdka, 2004], log analyses [Tauscher and Greenberg, 1997, Jansen and Pooch, 2001] and laboratory studies [Kaasten et al., 2002, Capra and Pérez-Quiñones, 2004, Capra and Pérez-Quiñones, 2005].

2.3 Multiple Information Devices

One of the major causes of information fragmentation is that we no longer are restricted to a single device, or a single source of information; most of our information is scattered across multiple devices, such as desktop computers at the office, laptops at home, portable digital assistants (PDAs) on the road, and of course, cellphones. We are not aware of any study so far that has explicitly considered the presence of multiple devices in a user's personal information ecosystem, or examined how they affect a user's information management practices.[1]

2.3.1 Research Focus in Multi-Platform User Interfaces

Most of the research in multi-platform user interfaces has focused on the specifics of the interaction on each device, on moving tasks between devices seamlessly, and maintaining some form of consistency when migrating such tasks. These user interfaces have been variously called plastic user interfaces [Thevenin and Coutaz, 1999], nomadic applications [Mori et al., 2003], or multi-browsing interfaces [Johanson et al., 2001].

Various techniques for interface migration have been proposed, including model-based approaches [Mori et al., 2003, Einsenstein et al., 2001], user interface markup languages [Abrams et al., 1999], and transformation-based approaches [Richter, 2005, Florins and Vanderdonckt, 2004].

[1] As of this writing, a workshop has been scheduled at CHI 2008 to discuss Personal Information Management and the Disappearing Desktop. The author serves on the Program Committee of this workshop.

Chu *et al.* [Chu et al., 2004] take the approach of migrating an entire application to support seamless task roaming. Bandelloni and Paternò [Bandelloni and Paternò, 2004] discuss the user interaction with an application while moving from one device to another in terms of three levels of migration: total, partial and mixed. Chhatpar and Pérez-Quiñones [Chhatpar and Pérez-Quiñones, 2003] call this migration "dialogue mobility". Florins and Vanderdonckt [Florins and Vanderdonckt, 2004] describe rules and transformations for graceful degradation of user interfaces during migration. Tungare *et al.* [Tungare et al., 2006c,Tungare et al., 2007] designed a framework to enable migration of task data across multiple platforms to overcome information fragmentation.

Thus, while there is a significant body of knowledge on building multi-device interfaces, not much attention has been given to the information needs of users on these devices.

2.4 Mental Workload Measurement

Operator performance in a particular task situation can be measured directly by standard performance metrics in usability engineering, e.g. the time taken to perform an experimental task, or the number and severity of errors in task performance. Although directly measured, performance metrics are unique to a particular task, and there are several reasons why performance metrics cannot be used to predict performance for an unknown task [Wilson and Eggemeier, 2006].

Subjective workload assessment techniques such as NASA Task Load Index (NASA TLX) [Hart and Staveland, 1988a], the Subjective Workload Assessment Technique, [Reid et al., 1982], and Workload Profile [Tsang and Velazquez, 1996] are used to provide an estimate of mental overload. Workload is defined as "the difference between the cognitive demands of a particular job or task, and the operator's attention resources" [Wickens, 1992], or "that portion of an operator's limited capacity actually required to perform a particular task." [O'Donnell and Eggemeier, 1986]. It is generally assumed that workload is related to operator performance such that low to moderate levels of workload are associated with acceptable levels of operator performance [Wilson and Eggemeier, 2006].

Since the ratings obtained via subjective workload assessments are not task-specific, it is possible to use them to compare the workload imposed by different tasks. Several task performance situations involving computing devices in the environment have been examined in detail, and workload measurements have been conducted using some of the techniques listed above [Ballas et al., 1992, Schryver, 1994, Bertram et al., 1992].

Title	Endpoints	Descriptions
Mental Demand	Low/High	How much mental and perceptual activity was required (e.g. thinking, deciding, calculating, remembering, looking, searching, etc.) Was the task easy or demanding, simple or complex, exacting or forgiving?
Physical Demand	Low/High	How much physical activity was required (e.g. pushing, pulling, turning, controlling, activating, etc.)? Was the task easy or demanding, slow or brisk, slack or strenuous, restful or laborious?
Temporal Demand	Low/High	How much time pressure did you feel due to the rate or pace at which the tasks or task elements occurred? Was the pace slow and leisurely or rapid and frantic?
Performance	Good/Poor	How successful do you think you were in accomplishing the goals of the task set by the experimenter (or yourself)? How satisfied were you with your performance in accomplishing these goals?
Effort	Low/High	How hard did you have to work (mentally and physically) to accomplish your level of performance?
Frustration Level	Low/High	How insecure, discouraged, irritated, stressed and annoyed versus secure, gratified, content, relaxed and complacent did you feel during the task?

Table 2.1: NASA TLX Dimensions (from [Hart and Staveland, 1988a])

2.4.1 The NASA Task Load Index (TLX)

The NASA Task Load Index (NASA TLX) [Hart and Staveland, 1988a] was developed in 1988 as a measure of perceived work load. In the years since, it has been shown to be a highly reliable, sensitive measure of workload. It includes six bipolar dimensions, as summarized in Table 2.1.

Performance measurements using the NASA TLX have been conducted in airline cockpits [Ballas et al., 1992], navigation [Schryver, 1994], and in the medical field [Bertram et al., 1992]. However, there has been no study of operator workload in modern computing environments such as information ecosystems.

2.5 Summary

Personal Information Management is an upcoming field nearing maturity, and researchers have focused on understanding users' personal information needs, and how systems and tools can help them satisfy them. Information needs can be examined by content type (known in PIM literature as 'collections'), across collections, or in specific contexts. The field of multi-platform user interfaces has focused mainly on the specifics of interaction on various devices, but a deeper investigation of the tasks performed and users' information needs is needed. Mental workload measurement is a technique used across many task domains over several years (the NASA TLX scale was proposed in 1988), but has not yet been applied to task scenarios such as the use of multiple devices for personal information management.

Chapter 3

Methodology

3.1 Introduction

There are several factors that contribute to mental workload for a user managing personal information with multiple devices in an information ecosystem. This system is more complicated than any individual components taken individually; it is essentially more than the sum of its parts. Thus, studies that isolate the user from an integrated information environment and instead study information management in artificial single-device contexts lack certain kinds of generalizability and validity. We have explored some of these factors in past research as follows.

3.2 Phase I: Personal Information Management across Multiple Devices

As seen in chapter 2, a lot of research has been conducted in the fields of Personal Information Management and Multi-Platform User Interfaces independently. However, the intersection of the two areas — studying how users perform PIM tasks with multiple devices — has been a blind spot in current research. The PIM research community is just beginning to look at the issues raised by the proliferation of portable information management devices.[1]

[1] E.g. the theme for the 3^{rd} Invitational Workshop on Personal Information Management (PIM 2008) was 'The Disappearing Desktop', and the call for papers included a specific mention of PIM across devices.

3.2.1 Research Questions

In the first phase of this research, I was primarily concerned with exploring and understanding users' current practices in the field. Before any specific experimental design could be developed or administered, it was important to gain knowledge about the state-of-the-art in personal management practices of users. Also of interest was the spread and popularity of multiple devices among users.

The research questions of interest are:

1. Which devices are commonly used together?
2. How do users adapt their workflows to their devices?
3. What problems do they encounter and how do they avoid them?
4. What strategies evolve and how mental workload might play a role in these strategies?

3.2.2 A Survey to Study the Use of Multiple Devices for PIM

In August 2007, based on the above research questions, we conducted a survey to understand the information management practices of users who use more than one information device. Each person owns and uses a unique set of devices, and insights can be gained about the practices they develop over time.

Survey Design

The survey was targeted specifically at knowledge workers, and received responses from participants from several companies in the San Francisco Bay Area (Google, Apple, IBM, Yahoo!), several universities (Virginia Tech, Georgia Tech, Michigan State University, Bath University UK) and several companies based in Bangalore, India. (N = 220). It contained a mix of quantitative and qualitative questions, and analysis is ongoing. The questionnaire used for the survey (along with IRB approval form) is available as an appendix to this proposal (Section 5.1).

The idea behind the survey was to understand the status quo in the field today. While it is true that other experimental techniques such as ethnographic studies, personal interviews, and guided

tours would provide richer data about fewer participants, our intent behind this survey was to tap into a larger participant pool (N = 220) and gain a wider — rather than deeper — understanding of the issues. As proposed later, this will be followed up with richer studies conducted with a sub-population.

The survey which contained questions in the following four categories:

Devices and Activities

- What is the distribution of users who use multiple devices?
- Is it only a small fraction of the population, or a larger majority?
- What are a few of the most common devices?
- What are some of the common PIM tasks that users choose to perform on certain devices?
- Are there certain tasks that are bound to a particular device, such that they may only be performed on that device?
- Are there certain tasks that may never be performed on certain devices?

The Use of Multiple Devices Together

- Which devices were commonly used in groups (i.e. together with other devices, used either simultaneously, or one after the other) to perform common tasks?
- What are the methods employed to share data among these devices?
- Do users keep their grouped devices completely synchronized at all times (i.e. do they maintain a copy of the same data on both devices at all times)?
- What are some of the problems and frustrations users have faced in using multiple devices together?
- Are people completely happy with the current offerings and their own workflows, or were they frustrated by certain aspects of how they were forced to manage their information by the current crop of tools and systems?

Buying New Devices

- What are some of the factors that influence users' buying decisions for new devices?
- How important it is to them that a new device integrate well into the set of existing devices?

Device and System Failures

- How often do users encounter failures in their information management systems?
- What are some of the common types of failures?
- How do users cope with failures?
- Are there any systems in place to guard against such failures, or are there reliable means of recovery from failures, after they occur?

3.2.3 Preliminary Results from Phase I

Preliminary findings from the survey are reported in [Tungare and Pérez-Quiñones, 2008], but a complete analysis is pending, and is a part of this research proposal. A few highlights were that mobile information devices are in much wider use than their stationary counterparts: almost 96% of users reported that they used (at least one) laptop computer. Desktops were used by a modest 71% of the population studied. The numbers of devices reported by our participants are indicated in figure 3.1. Several devices are often used in groups, e.g. laptops and cell phones (reported by 52 participants). Integrated multi-function portable devices such as Palm Treos, Blackberries and Apple iPhones have begun to replace single-function devices for communication (e.g. email and IM). Figure 3.2 shows how users often use devices in groups. Users use certain features opportunistically because they happen to be carrying a multi-function device with them. There were several instances of disconnect reported when the system did not perform as the user expected, as a result of which the user's information needs remained unmet.

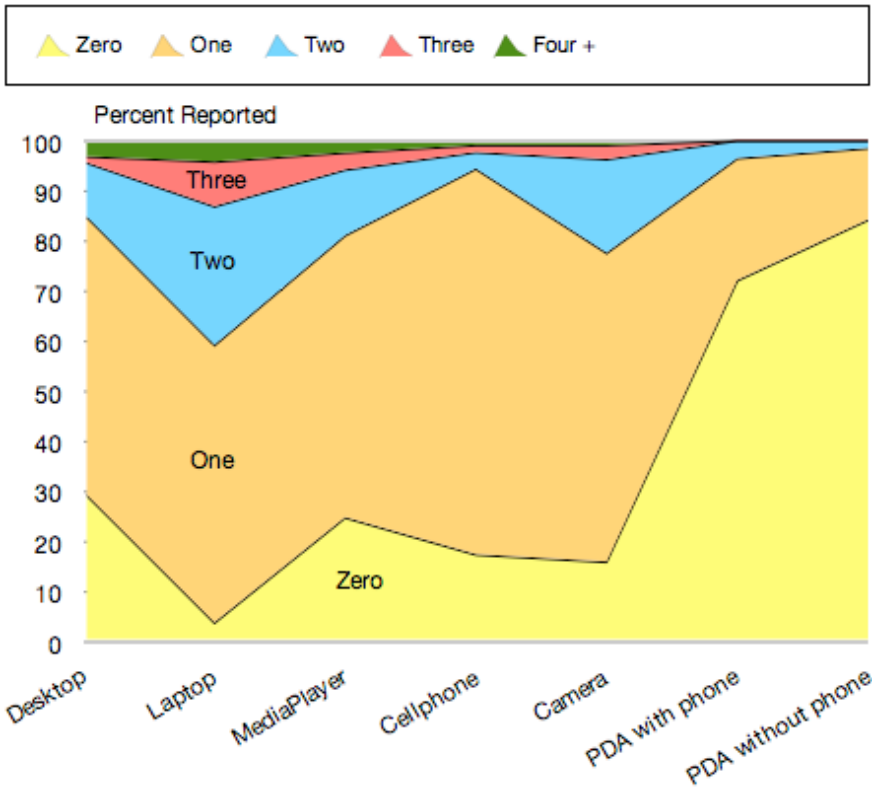

Figure 3.1: Percentages of devices used, as reported by study participants

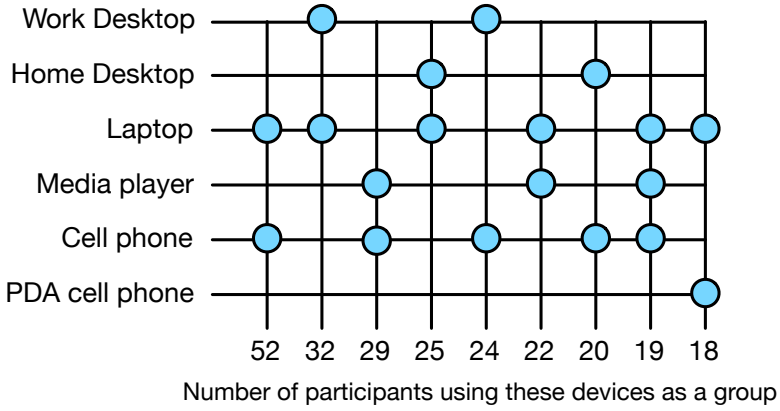

Figure 3.2: Devices used in groups, as indicated by our participants. In this visualization, groups are indicated by placing a blue dot for each device in a group, and the number of users reporting that configuration appears at the bottom of each line.

Problems Uncovered in Existing Tools, Devices and Practices

From an analysis of the survey results, and from our own observations, it is clear that there are frequent failures in information management when using multiple devices. Sometimes, fresh information is present only on a device that is not accessible in the current context (e.g. a file updated on the desktop computer is not available on a laptop during an off-site business meeting.) Worse still, stale information may be used for decision-making because of the lack of awareness of whether a document is current or not. Version conflicts are commonplace, not only in documents, but also in address book contact records, calendar events, and other personal information.

Information management tools are supposed to assist the user in her tasks, but several current tools provide only partial support for what the user actually wants to do. This potentially results in an increase in mental workload for users. The designers of such tools expect the user to stick to a premeditated plan, whereas in reality, information needs may arise anytime, anywhere. Planning ahead is impossible in most contexts. Failure to plan ahead prevents the user from making full use of the tool. Research in distributed cognition and situated action has shown that users rarely follow a rigid plan to perform actions, but instead let various external factors guide their actions until the required result is attained.

Analysis of Survey Responses

A complete content analysis and quantitative analysis of the data gathered during this survey is planned as part of this proposal. Specifically, by analyzing trends and common patterns in users' information management workflows, it is possible to design an experiment that involves a subset of these tasks. These experimental tasks are expected to be more representative of actual user behavior than others designed without such a survey to guide them.

Example of Content Analysis

Figure 3.3 shows an example comment from one of our survey participants. In figure 3.4, important elements of this comment are highlighted and tagged. Elements of interest include the devices mentioned by the user, the task(s) that the user was trying to perform, problems that the user encountered, and the outcome of the task. These elements, pooled from all participants' responses, will enable us to design representative experimental tasks with a representative experimental setup

"The last device I acquired was a cell phone from Verizon. I would have liked to synchronize data from my laptop or my PDA with it but there seems to be no reasonable way to do so. I found a program that claimed to be able to break in over bluetooth but it required a fair amount of guess work as to data rates etc and I was never able to actually get it to do anything. In the end I gave up. Fortunately I dont know that many people and I usually have my PDA with me so it isnt a big deal but frankly I dont know how Verizon continues to survive with the business set..."

Figure 3.3: Example comment from survey participant

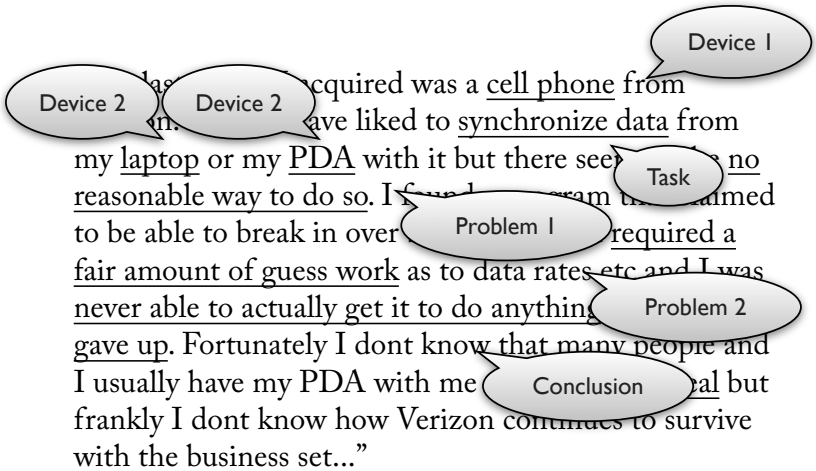

Figure 3.4: Tagging and analysis of example comment

including the specific devices used, their location, context, information stored on them, and features (or lack of features.)

I plan to recruit two additional coders for this survey, in order to make sure that my coding is unbiased. These coders will code two random subsets of the survey, and inter-rater reliability will be calculated.

3.2.4 Mobile Devices as an Information Ecosystem

We compared the many devices in a user's environment to the many organisms that form a biological ecosystem [Tungare et al., 2006b]. The devices that form a close-knit ecosystem with a user's other devices are not explicitly designed to account for their "social life". A 'personal information ecosystem' can be defined as 'a system of devices and applications that are present in the information environment of a user, that interact closely and richly with one another, to help the user achieve the goal of fulfilling his/her information needs.' Several similarities between biological ecosystems and device systems allow us to extract lessons that could help inform the design of the latter for information management.

This thought follows closely the distributed cognition approach to systems design: it cannot be assumed that one person (or one device) is the sole cognitive agent in an environment [Hutchins, 1995]. Indeed, the environment itself (in this case, consisting of other devices) constitutes a part of the cognitive system as a whole. Thus, designing for a single device at a time largely ignores the ecological context in which cognition occurs.

3.2.5 Consistency versus Holistic Design for Information Ecosystems

At CHI 2006, a workshop was organized to discuss the several forms of consistency that may or may not be essential towards the design of a multi-platform system. We argued [Tungare et al., 2006a] that the design of current mobile applications follows the design of their desktop counterparts very closely, down to the placement of buttons and other user interface elements on a smaller screen. However, this strict similarity in design and feature parity is often a hindrance to usability because users do not perform the same tasks on mobile computing devices that they do on desktop computers.

The higher goal of achieving seamless task migration — that a user should be able to suspend a task

on one device, and resume it on another seamlessly — should be the ultimate focus of interface designers. This refers to a state of extremely low mental workload for the user when relocating tasks across devices. We deemed this to be a worthy goal even if a tradeoff had to be made by forfeiting interface-level consistency. Our position was that the multiple devices used by a user for information access should be functionally complementary instead of mere replicates of each other.

The best example of this trade-off was in the design of an alarm clock application that would run on a user's desktop computer, laptop computer, cell phone, PDA and music player. Such an application, if designed with only pure consistency in mind, would result in five different sounds ringing at the chosen hour. But if it were designed with an awareness of the other devices in the user's cognitive system, the devices would negotiate amongst themselves and present only a single alarm to the user. This approach would lead to a design that reduces mental workload for the user. We termed this a 'holistic design approach' that puts users' tasks before interface-level consistency.

3.2.6 Systems that Support Task Migration

As part of previous work, I developed Syncables [Tungare et al., 2006c, Tungare et al., 2007], a software framework designed to make it easy for applications to maintain their information consistently across multiple devices without explicit actions on part of the user. The framework makes no distinction between information present on one device or another, and treats them all as part of a single system, much like the distributed cognition approach suggests [Hollan et al., 2000]. It supports all types of information, even beyond files, such as calendar events, address book contacts, and contextual information such as the location of the cursor on the screen. Thus, it removes the file-type-based separation, and instead, groups documents by project [Bergman et al., 2006].

3.3 Phase II: Measurement of Mental Workload of Users in Information Ecosystems

From the results of the survey we conducted earlier [Tungare and Pérez-Quiñones, 2008], it is clear that users are struggling with managing their personal information on multiple devices. What is not clear is a measurable indicator of mental workload for this particular task situation. NASA TLX is considered an accurate indicator of operator performance for many domains. When it was originally designed, it consisted of 9 dimensions (instead of the current six). During tests, it was observed that three dimensions (that were later dropped) did not significantly explain the variation in operator performance by more than a few tenths of a percent [Hart and Staveland, 1988a].

3.3.1 RQ: What is the mental workload incurred by users in information ecosystems?

What is the mental workload incurred by users when they are trying to use multiple devices for personal information management? Certain tasks require more attentional resources than others, and may result in increased mental workload. However, an accurate measurement of workload for such scenarios has not been conducted.

Hypothesis: Mental workload will be higher for those task contexts that users identified in Phase I as being difficult or frustrating.

While this seems obvious, it has not been experimentally evaluated. Knowledge of mental workload in various task contexts and workflows will assist users in determining the best systems for use, based on objective criteria. Measures of mental workload can be compared across tasks, thus making it possible for disparate research laboratories to compare their results on a single scale.

3.3.2 RQ: Do alternate strategies result in lower mental workload?

For those tasks that users have indicated are frustrating for them, do the alternate strategies result in lower mental workload? We know from our survey that users employ several alternate strategies when they encounter problems in their primary workflows. Are these changes in strategies motivated by a desire to move towards lower mental workload, even though they may not be aware of it at the time?

Hypothesis: Alternate strategies employed by users result in lower mental workload.

If we are able to determine that these changes are indeed motivated by a desire to move towards lower mental workload, this would provide us a way to compare strategies and pick the ones that result in lower mental workload.

3.3.3 RQ: Is mental workload a reliable indicator of operator performance?

The NASA TLX scale has been shown to be an accurate indicator of performance for several disparate task domains [Ballas et al., 1992, Schryver, 1994, Bertram et al., 1992], so it is reasonable to expect that it will be able to predict operator performance to a high degree in information ecosystems, a domain in which it has not been explicitly tested.

But this remains an unproven hypothesis. I plan to test this hypothesis as follows:

Hypothesis: Mental workload in tasks performed in certain information ecosystems can be measured with the NASA TLX scale.

The validity of the NASA TLX scale will be validated as in most other domains in the past [Ballas et al., 1992, Schryver, 1994, Bertram et al., 1992]. First, the operator's performance at a set of predetermined tasks will be measured by standard usability metrics such as time taken to perform task, number of errors encountered, total number of tasks that could not be completed (failures), etc. Immediately after each task, subjective workload assessments will be administered as surveys, either via paper, or using a computer.

The exact task-related measurements that will be taken depend on the exact tasks chosen for this experiment. The exact tasks will be chosen after a deep analysis of the 220 responses to the multiple devices survey from August 2007. An analysis of the overall score and its correlation with operator performance can help conclude whether or not NASA TLX is an accurate indicator of performance, and whether it can be substituted as an acceptable measure in place of the more rigorous operator performance tests.

How well do NASA TLX scores along each individual dimensions correlate to actual operator performance? While it is expected that they will correlate to a certain extent, the actual correlation coefficient is unknown. There may also be additional dimensions that are more accurate predictors of operator performance than the standard six NASA TLX dimensions.

Are there additional dimensions that may be able to predict operator performance better in an information ecosystem context than the existing six NASA TLX dimensions? The evaluation will be conducted exactly as above, with new dimensions introduced and measured in place of the current six dimensions of the NASA TLX scale.

3.3.4 Experiment Design for Phase II

A complete analysis of the 220 responses to the survey in Phase I will help guide the experiment design for the second phase. Specifically, the responses will help identify a set of tasks that can serve as good representatives for evaluation. Depending on the tasks, specific performance metrics will be chosen according to established usability engineering principles.

Based on necessary p-values to confirm or disprove the hypotheses, adequate sample sizes will be determined. Recruitment efforts will include posting flyers to campus notice boards and sending announcements via email lists. In addition to undergraduate and graduate students, I plan to recruit employees from companies in the nearby Virginia Tech Corporate Research Center to participate in the studies.

Since the details of the experiment design for Phase II depend upon the upcoming analysis of the survey responses from Phase I, I will meet with the committee members to discuss these after the survey data analysis is complete, and before the Phase II experiments are conducted.

3.4 Summary

A deeper understanding of users' personal information management workflows is necessary to design better tools and system. This consists of two parts: understanding current practices (§3.2), and understanding the impact on mental workload (§3.3). The question of whether mental workload measurement techniques such as NASA TLX are accurate predictors of task performance (as introduced in chapter 1) remains to be answered in Phase II of my research. The methodology outlined here is expected to be completed according to the schedule in chapter 4.

Chapter 4

Schedule

I propose to complete the work outlined in chapter 3 as per the schedule in figure 4.1.

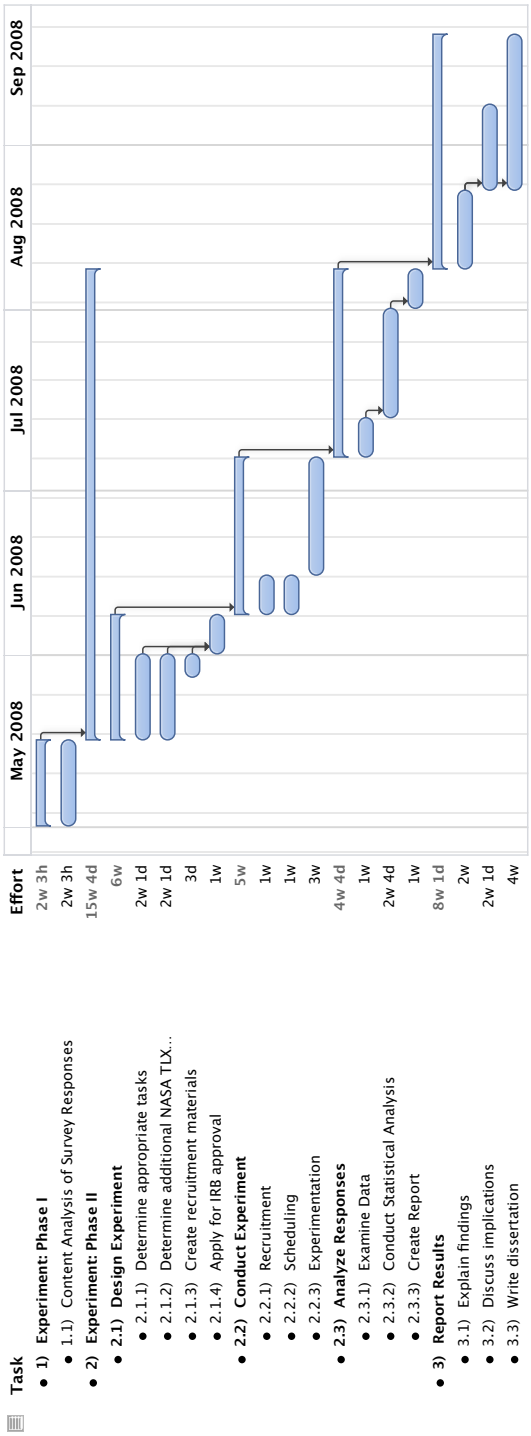

Figure 4.1: Proposed Schedule

Chapter 5

Appendices

5.1 Appendix A: Survey Questionnaire

Your Devices

1. Which of the following devices do you own and use regularly? How many of each type do you have?

	None	One	Two	Three	Four or more
Desktop computer (work)	○	○	○	○	○
Desktop computer (home)	○	○	○	○	○
Laptop computer (either work or home)	○	○	○	○	○
Portable media player (e.g. iPod)	○	○	○	○	○
Cell phone	○	○	○	○	○
Personal Digital Assistant (PDA)	○	○	○	○	○
Treo, Blackberry, iPhone, other multi-function device	○	○	○	○	○
Digital camera	○	○	○	○	○

2. What activities do you usually perform on each of your devices? If a feature exists, but you do not use it, please do not check the box.

Not all choices will apply to all devices. Please check all boxes that apply in your case. Primary home and work computers could be your desktop or laptop computer, as appropriate.

	Desktop (work)	Desktop (home)	Laptop	Cell phone	Portable media player	PDA	Treo, Blackberry, iPhone, or other	Digital camera
Browsing the Web	☐	☐	☐	☐	☐	☐	☐	☐
Read web email (work)	☐	☐	☐	☐	☐	☐	☐	☐
Read web email (personal)	☐	☐	☐	☐	☐	☐	☐	☐
Download email (work)	☐	☐	☐	☐	☐	☐	☐	☐
Download email (personal)	☐	☐	☐	☐	☐	☐	☐	☐
Instant messaging (work)	☐	☐	☐	☐	☐	☐	☐	☐
Instant messaging (personal)	☐	☐	☐	☐	☐	☐	☐	☐
Send / receive SMS	☐	☐	☐	☐	☐	☐	☐	☐
Address book / contacts (work)	☐	☐	☐	☐	☐	☐	☐	☐
Address book / contacts (personal)	☐	☐	☐	☐	☐	☐	☐	☐
Calendar (work)	☐	☐	☐	☐	☐	☐	☐	☐
Calendar (personal)	☐	☐	☐	☐	☐	☐	☐	☐
Read or edit documents (work)	☐	☐	☐	☐	☐	☐	☐	☐
Read or edit documents (personal)	☐	☐	☐	☐	☐	☐	☐	☐
To-do notes	☐	☐	☐	☐	☐	☐	☐	☐
Making phone calls	☐	☐	☐	☐	☐	☐	☐	☐
Playing music	☐	☐	☐	☐	☐	☐	☐	☐
Watching videos	☐	☐	☐	☐	☐	☐	☐	☐
Taking photos	☐	☐	☐	☐	☐	☐	☐	☐
Storing, viewing or managing photos	☐	☐	☐	☐	☐	☐	☐	☐

3. If you own a multi-function device (e.g. Treo, Blackberry, iPhone), has the presence of this device led you to abandon any other device (e.g. devices that previously performed each individual function.) If yes, please tell us more about the multi-function device as well as the others it replaced.

You can enter as many lines as you want; don't let the size of the box limit your response.

Using Multiple Devices Together

4. Which of the above devices do you frequently operate nearly at the same time? E.g. at your office, you might use your PDA and your desktop simultaneously. On the go, you might always carry your iPod and phone. In each row below, select the devices that are used together in a group.

Feel free to use as many rows as you need and leave the rest blank.

	Work desktop	Home desktop	Laptop	PDA	Cell phone	Portable media player	Treo, Blackberry, iPhone	Other multi-function device
Group 1	☐	☐	☐	☐	☐	☐	☐	☐
Group 2	☐	☐	☐	☐	☐	☐	☐	☐
Group 3	☐	☐	☐	☐	☐	☐	☐	☐
Group 4	☐	☐	☐	☐	☐	☐	☐	☐
Group 5	☐	☐	☐	☐	☐	☐	☐	☐

5. Between which pairs of devices do you usually copy or synchronize data?

Feel free to use as many rows as you need and leave the rest blank.

	First Device	Direction	Second Device	Type of data
Pair 1	Work desktop	Synchronizes both ways	Home desktop	Documents
Pair 2	-- Select --	-- Select --	-- Select --	-- Select --
Pair 3	-- Select --	-- Select --	-- Select --	-- Select --
Pair 4	-- Select --	-- Select --	-- Select --	-- Select --
Pair 5	-- Select --	-- Select --	-- Select --	-- Select --
Pair 6	-- Select --	-- Select --	-- Select --	-- Select --
Pair 7	-- Select --	-- Select --	-- Select --	-- Select --
Pair 8	-- Select --	-- Select --	-- Select --	-- Select --
Pair 9	-- Select --	-- Select --	-- Select --	-- Select --
Pair 10	-- Select --	-- Select --	-- Select --	-- Select --

6. Data synchronization horror stories: Syncing data sometimes has its own pitfalls. Have you ever been victim to a situation where synchronization failed to live up to your expectations, either due to system errors, or because of forgetting to do it, etc.? If you have a story, please share with us. When was the last time such an incident happened to you?

You can enter as many lines as you want; don't let the size of the box limit your response.

Buying a New Device

7. Please indicate your agreement with the statement below for each factor in the left column:

"Factor X is the single most important factor to me when buying a new device."

	Strongly agree	Agree	Neither agree nor disagree	Disagree	Strongly disagree
Feature richness	○	○	○	○	○
Price	○	○	○	○	○
Ease of use	○	○	○	○	○
"Hipness"	○	○	○	○	○
A good fit with existing devices	○	○	○	○	○
Manufacturer/brand, etc.	○	○	○	○	○

8. Overall, how satisfied have you been with the last device you purchased?

- ○ Very satisfied
- ○ Satisfied
- ○ Neither satisfied nor dissatisfied
- ○ Dissatisfied
- ○ Very dissatisfied

9. If you ran into any problems using your new device with your existing devices and data, please describe them. Feel free to leave blank if you were entirely satisfied with how your new device integrated into your life.

You can enter as many lines as you want; don't let the size of the box limit your response.

10. Have you ever encountered a situation where one of your devices stopped functioning, or was otherwise unusable for its normal function?

Feel free to use as many rows as you need and leave the rest blank.

	What device?	Did you lose data?	Was it a hardware or software issue?	Were you able to restore your data from a backup copy?	How soon did you get a device to replace the failed one?	How long did it take for the new device to completely replace the function of the failed device?
Failed Device	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --

1 Failed Device 2	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --
3 Failed Device 4	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --
5 Failed Device	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --	-- Select --

About you

11. Are you male or female?

◯ Male ◯ Female

12. Which of the following age groups do you belong into?

-- Select --

13. What is the highest level of education you have completed?

-- Select --

14. Do you consider yourself an information worker (or a knowledge worker)?

◯ Yes, full-time ◯ Yes, part-time ◯ No ◯ Not sure

15. Who manages your calendar appointments? Please check all boxes that apply in your case.

☐ You ☐ Your assistant ☐ Your spouse ☐ Your parent ☐ Other (please specify)

16. Which of the following does your primary work activity involve? Please check all boxes that apply in your case.

☐ Working at a desk
☐ Communicating with people
☐ Conducting research
☐ Attending classes
☐ Traveling locally (roughly within the same city, town, or metropolitan area)
☐ Traveling between local offices (but no airline travel)
☐ Airline travel
☐ Other (please specify)

17. What is your primary mode of transport for commuting to your workplace?

◯ None, I telecommute
◯ Walk
◯ Use a bicycle
◯ Drive
◯ Carpool
◯ By train
◯ By bus

18. How long is your one-way commute each day?

◯ I telecommute
◯ Less than 10 minutes
◯ 10-20 minutes
◯ 20-40 minutes
◯ 40 minutes to an hour

5.2 Appendix B: IRB Approval for Survey

VirginiaTech

Office of Research Compliance
Carmen T. Green, IRB Administrator
2000 Kraft Drive, Suite 2000 (0497)
Blacksburg, Virginia 24061
540/231-4358 Fax 540/231-0959
e-mail ctgreen@vt.edu
www.irb.vt.edu
FWA00000572(expires 1/20/2010)
IRB # is IRB00000667

DATE: August 7, 2007

MEMORANDUM

TO: Manuel A. Perez-Quinones
Manas Tungare

FROM: Carmen Green

SUBJECT: **IRB Exempt Approval:** "Understanding the Evolution of Users' Personal Information Management Practices" , IRB # 07-388

I have reviewed your request to the IRB for exemption for the above referenced project. I concur that the research falls within the exempt status. Approval is granted effective as of August 7, 2007.

As an investigator of human subjects, your responsibilities include the following:

1. Report promptly proposed changes in previously approved human subject research activities to the IRB, including changes to your study forms, procedures and investigators, regardless of how minor. The proposed changes must not be initiated without IRB review and approval, except where necessary to eliminate apparent immediate hazards to the subjects.

2. Report promptly to the IRB any injuries or other unanticipated or adverse events involving risks or harms to human research subjects or others.

cc: File

5.3 Appendix C: The NASA TLX Scale [Hart and Staveland, 1988b]

NASA Task Load Index

Hart and Staveland's NASA Task Load Index (TLX) method assesses work load on five 7-point scales. Increments of high, medium and low estimates for each point result in 21 gradations on the scales.

Name	Task	Date

Mental Demand — How mentally demanding was the task?

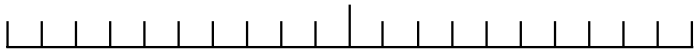

Very Low — Very High

Physical Demand — How physically demanding was the task?

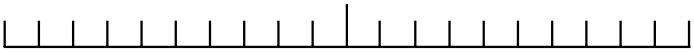

Very Low — Very High

Temporal Demand — How hurried or rushed was the pace of the task?

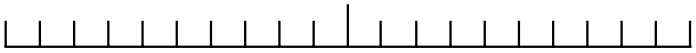

Very Low — Very High

Performance — How successful were you in accomplishing what you were asked to do?

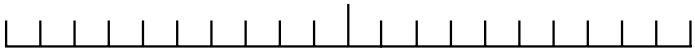

Perfect — Failure

Effort — How hard did you have to work to accomplish your level of performance?

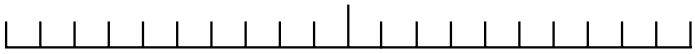

Very Low — Very High

Frustration — How insecure, discouraged, irritated, stressed, and annoyed wereyou?

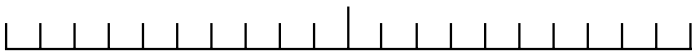

Very Low — Very High

Bibliography

[Abrams et al., 1998] Abrams, D., Baecker, R., and Chignell, M. (1998). Information archiving with bookmarks: personal web space construction and organization. In *CHI '98: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 41–48, New York, NY, USA. ACM Press/Addison-Wesley Publishing Co.

[Abrams et al., 1999] Abrams, M., Phanouriou, C., Batongbacal, A., Williams, S., and Shuster, J. (1999). UIML: An appliance-independent xml user interface language. In *Proceedings of the 8th WWW conference*.

[Ballas et al., 1992] Ballas, J. A., Heitmeyer, C. L., and nones, M. A. P.-Q. (1992). Evaluating two aspects of direct manipulation in advanced cockpits. In *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 127–134, New York, NY, USA. ACM.

[Bälter and Sidner, 2002] Bälter, O. and Sidner, C. L. (2002). Bifrost inbox organizer: giving users control over the inbox. In *NordiCHI '02: Proceedings of the second Nordic conference on Human-computer interaction*, pages 111–118, New York, NY, USA. ACM Press.

[Bandelloni and Paternò, 2004] Bandelloni, R. and Paternò, F. (2004). Flexible Interface Migration. In *IUI '04: Proceedings of the 9th international conference on Intelligent user interface*, pages 148–155, New York, NY, USA. ACM Press.

[Barreau and Nardi, 1995] Barreau, D. and Nardi, B. A. (1995). Finding and reminding: file organization from the desktop. *SIGCHI Bull.*, 27(3):39–43.

[Barreau, 1995] Barreau, D. K. (1995). Context as a factor in personal information management systems. *Journal of the American Society for Information Science*, 46(5):327–339.

[Bellotti et al., 2003] Bellotti, V., Ducheneaut, N., Howard, M., and Smith, I. (2003). Taking email to task: the design and evaluation of a task management centered email tool. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 345–352, New York, NY, USA. ACM Press.

[Bellotti and Smith, 2000] Bellotti, V. and Smith, I. (2000). Informing the design of an information management system with iterative fieldwork. In *DIS '00: Proceedings of the conference on Designing interactive systems*, pages 227–237, New York, NY, USA. ACM Press.

[Bergman et al., 2006] Bergman, O., Beyth-Marom, R., and Nachmias, R. (2006). The project fragmentation problem in personal information management. In *CHI '06: Proceedings of the SIGCHI conference on Human Factors in computing systems*, pages 271–274, New York, NY, USA. ACM Press.

[Bertram et al., 1992] Bertram, D. A., Opila, D. A., Brown, J. L., Gallagher, S. J., Schifeling, R. W., Snow, I. S., and Hershey, C. O. (1992). Measuring physician mental workload: Reliability and validity assessment of a brief instrument. *Medical Care*, 30(2):95–104.

[Boardman et al., 2003] Boardman, R., Spence, R., and Sasse, M. A. (2003). Too many hierarchies?: The daily struggle for control of the workspace. In *Proc. HCI International 2003*.

[Bush, 1945] Bush, V. (1945). As we may think. *The Atlantic Monthly*.

[Capra and Pérez-Quiñones, 2004] Capra, R. G. and Pérez-Quiñones, M. (2004). Mobile refinding of web information using a voice interface. *ACM: Computing Research Repository (CoRR)*, cs.HC/0402001.

[Capra and Pérez-Quiñones, 2005] Capra, R. G. and Pérez-Quiñones, M. A. (2005). Using web search engines to find and refind information. *IEEE Computer*, 38(10):36–42.

[Chhatpar and Pérez-Quiñones, 2003] Chhatpar, C. and Pérez-Quiñones, M. (2003). Dialogue mobility across devices. In *ACM Southeast Conference (ACMSE)*, Savannah, Georgia.

[Chu et al., 2004] Chu, H.-h., Song, H., Wong, C., Kurakake, S., and Katagiri, M. (2004). Roam, a seamless application framework. *Journal of Systems and Software*, 69(3):209–226.

[Czerwinski et al., 2006] Czerwinski, M., Gage, D. W., Gemmell, J., Marshall, C. C., Pérez-Quiñones, M. A., Skeels, M. M., and Catarci, T. (2006). Digital memories in an era of ubiquitous computing and abundant storage. *Communications of the Association for Computing Machinery (CACM)*, 49(1):44–50.

[Denis and Karsenty, 2004] Denis, C. and Karsenty, L. (2004). Inter-usability of multi-device systems - a conceptual framework. In Seffah, A. and Javahery, H., editors, *Multiple User Interfaces: Cross-Platform Applications and Context-Aware Interfaces*, pages 373–384. John Wiley and Sons.

[Ducheneaut and Bellotti, 2001] Ducheneaut, N. and Bellotti, V. (2001). E-mail as habitat: an exploration of embedded personal information management. *interactions*, 8(5):30–38.

[Dumais et al., 2003] Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R., and Robbins, D. C. (2003). Stuff i've seen: a system for personal information retrieval and re-use. In *SIGIR '03: Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 72–79, New York, NY, USA. ACM Press.

[Einsenstein et al., 2001] Einsenstein, J., Vanderdonckt, J., and Puerta, A. (2001). Applying model-based techniques to the development of uis for mobile computers. In *Proceedings IUI'01: International Conference on Intelligent User Interfaces*, pages 69–76. ACM Press.

[Farhoomand and Drury, 2002] Farhoomand, A. F. and Drury, D. H. (2002). Managerial information overload. *Commun. ACM*, 45(10):127–131.

[Fidel and Pejtersen, 2004] Fidel, R. and Pejtersen, A. (2004). From information behaviour research to the design of information systems: the Cognitive Work Analysis framework. *Information Research*, 10(1):10–1.

[Florins and Vanderdonckt, 2004] Florins, M. and Vanderdonckt, J. (2004). Graceful degradation of user interfaces as a design method for multiplatform systems. In *IUI '04: Proceedings of the 9th international conference on Intelligent user interface*, pages 140–147, New York, NY, USA. ACM Press.

[Gwizdka, 2000] Gwizdka, J. (2000). Timely reminders: a case study of temporal guidance in pim and email tools usage. In *CHI '00: CHI '00 extended abstracts on Human factors in computing systems*, pages 163–164, New York, NY, USA. ACM Press.

[Gwizdka, 2002] Gwizdka, J. (2002). Reinventing the inbox: supporting the management of pending tasks in email. In *CHI '02: CHI '02 extended abstracts on Human factors in computing systems*, pages 550–551, New York, NY, USA. ACM Press.

[Gwizdka, 2004] Gwizdka, J. (2004). Email task management styles: the cleaners and the keepers. In *CHI '04: CHI '04 extended abstracts on Human factors in computing systems*, pages 1235–1238, New York, NY, USA. ACM Press.

[Hart and Staveland, 1988a] Hart, S. G. and Staveland, L. E. (1988a). Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. *Human Mental Workload*, 1:139–183.

[Hart and Staveland, 1988b] Hart, S. G. and Staveland, L. E. (1988b). NASA TLX Scale. PDF.

[Hollan et al., 2000] Hollan, J., Hutchins, E., and Kirsh, D. (2000). Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Trans. Comput.-Hum. Interact.*, 7(2):174–196.

[Hutchins, 1995] Hutchins, E. (1995). *Cognition in the Wild*. MIT Press.

[Jansen and Pooch, 2001] Jansen, B. J. and Pooch, U. (2001). A review of web searching studies and a framework for future research. *J. Am. Soc. Inf. Sci. Technol.*, 52(3):235–246.

[Johanson et al., 2001] Johanson, B., Ponnekanti, S., Sengupta, C., and Fox, A. (2001). Multibrowsing: Moving web content across multiple displays. In *UbiComp '01: Proceedings of the 3rd international conference on Ubiquitous Computing*, pages 346–353, London, UK. Springer-Verlag.

[Jones et al., 2001] Jones, W., Bruce, H., and Dumais, S. (2001). Keeping found things found on the web. In *CIKM '01: Proceedings of the tenth international conference on Information and knowledge management*, pages 119–126, New York, NY, USA. ACM Press.

[Jones et al., 2002] Jones, W., Dumais, S., and Bruce, H. (2002). Once Found, What Then? A Study of 'Keeping' Behaviors in the Personal Use of Web Information. In *Proceedings of the American Society for Information Science and Technology*, volume 39, pages 391–402. American Society for Information Science and Technology.

[Jones and Teevan, 2007] Jones, W. and Teevan, J. (2007). *Personal Information Management*. University of Washington Press, Seattle, Washington.

[Kaasten et al., 2002] Kaasten, S., Greenberg, S., and Edwards, C. (2002). *How People Recognize Previously Seen WWW Pages from Titles, URLs and Thumbnails*, volume XVI of *Proceedings of Human Computer Interaction, BCS Conference Series*, pages 247–265. Springer Verlag.

[Karger and Jones, 2006] Karger, D. R. and Jones, W. (2006). Data unification in personal information management. *Commun. ACM*, 49(1):77–82.

[Kelly, 2006] Kelly, D. (2006). Evaluating personal information management behaviors and tools. *Commun. ACM*, 49(1):84–86.

[Kelly and Teevan, 2003] Kelly, D. and Teevan, J. (2003). Implicit feedback for inferring user preference: a bibliography. *SIGIR Forum*, 37(2):18–28.

[Lansdale, 1988] Lansdale, M. W. (1988). The psychology of personal information management. *Applied Ergonomics*, 19:55–66.

[Levy, 2005] Levy, D. M. (2005). To grow in wisdom: Vannevar Bush, Information Overload, and the Life of Leisure. In *JCDL '05: Proceedings of the 5th ACM/IEEE-CS joint conference on Digital libraries*, pages 281–286, New York, NY, USA. ACM Press.

[Mackay, 1988] Mackay, W. E. (1988). Diversity in the use of electronic mail: a preliminary inquiry. *ACM Trans. Inf. Syst.*, 6(4):380–397.

[Malone, 1983] Malone, T. W. (1983). How do people organize their desks?: Implications for the design of office information systems. *ACM Trans. Inf. Syst.*, 1(1):99–112.

[Mori et al., 2003] Mori, G., Paternò, F., and Santoro, C. (2003). Tool support for designing nomadic applications. In *IUI '03: Proceedings of the 8th international conference on Intelligent user interfaces*, pages 141–148, New York, NY, USA. ACM Press.

[O'Donnell and Eggemeier, 1986] O'Donnell, R. D. and Eggemeier, F. T. (1986). *Workload assessment methodology*, volume 2 of *Handbook of perception and human performance: Vol. 2. Cognitive processes and performance*, chapter Workload assessment methodology, pages 42/1–42/49. Wiley, New York.

[Reid et al., 1982] Reid, G. B., Eggemeier, F. T., and Shingledecker, C. A. (1982). Subjective Workload Assessment Technique. Technical report, Air Force Flight Test Center, Edwards, CA.

[Richter, 2005] Richter, K. (2005). A Transformation Strategy for Multi-device Menus and Toolbars. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1741–1744, New York, NY, USA. ACM Press.

[Schick et al., 1990] Schick, A. G., Gordon, L. A., and Haka, S. (1990). Information overload: A temporal approach. *Accounting, Organizations and Society*, 15(3):199–220.

[Schryver, 1994] Schryver, J. C. (1994). Experimental validation of navigation workload metrics. *Human Factors and Ergonomics Society Annual Meeting Proceedings*, 38:340–344(5).

[Tauscher and Greenberg, 1997] Tauscher, L. and Greenberg, S. (1997). How people revisit web pages: empirical findings and implications for the design of history systems. *Int. J. Hum.-Comput. Stud.*, 47(1):97–137.

[Teevan et al., 2004] Teevan, J., Alvarado, C., Ackerman, M. S., and Karger, D. R. (2004). The perfect search engine is not enough: a study of orienteering behavior in directed search. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 415–422, New York, NY, USA. ACM Press.

[Teevan et al., 2007] Teevan, J., Capra, R., and Pérez-Quiñones, M. (2007). *How people find information*, chapter 3, page 17. University of Washington Press, Seattle, Washington.

[Teevan and Jones, 2008] Teevan, J. and Jones, W. (2008). PIM 2008: Personal Information Management: The Disappearing Desktop, a CHI 2008 Workshop. Personal discussions with workshop participants.

[Thevenin and Coutaz, 1999] Thevenin, D. and Coutaz, J. (1999). Plasticity of user interfaces: Framework and research agenda. In *Interact*, pages 110–117, Edinburgh. IFIP.

[Tsang and Velazquez, 1996] Tsang, P. and Velazquez, V. (1996). Diagnosticity and multidimensional subjective workload ratings. *Ergonomics*, 39(3):358–381.

[Tungare and Pérez-Quiñones, 2008] Tungare, M. and Pérez-Quiñones, M. (2008). It's not what you have, but how you use it: Compromises in mobile device use. Technical report, ACM Computing Research Repository.

[Tungare et al., 2006a] Tungare, M., Pyla, P. S., and Pérez-Quiñones, M. (2006a). Multiple user interfaces: Why consistency is not everything, and seamless task migration is key. In *Proceedings of the CHI 2006 Workshop on The Many Faces of Consistency in Cross-Platform Design.*

[Tungare et al., 2006b] Tungare, M., Pyla, P. S., Pérez-Quiñones, M., and Harrison, S. (2006b). Personal information ecosystems and implications for design. Technical Report cs/0612081, ACM Computing Research Repository.

[Tungare et al., 2007] Tungare, M., Pyla, P. S., Sampat, M., and Pérez-Quiñones, M. (2007). Syncables: A framework to support seamless data migration across multiple platforms. In *IEEE International Conference on Portable Information Devices (IEEE Portable)*.

[Tungare et al., 2006c] Tungare, M., Pyla, P. S., Sampat, M., and Perez-Quinones, M. (2006c). Defragmenting information using the syncables framework. In *Proceedings of the 2nd Invitational Workshop on Personal Information Management at SIGIR 2006.*

[Whittaker and Hirschberg, 2001] Whittaker, S. and Hirschberg, J. (2001). The character, value, and management of personal paper archives. *ACM Trans. Comput.-Hum. Interact.*, 8(2):150–170.

[Whittaker and Sidner, 1996] Whittaker, S. and Sidner, C. (1996). Email overload: exploring personal information management of email. In *CHI '96: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 276–283, New York, NY, USA. ACM Press.

[Wickens, 1992] Wickens, C. (1992). *Engineering Psychology and Human Performance*. HarperCollins Publishers, New York.

[Wilson and Eggemeier, 2006] Wilson, G. F. and Eggemeier, F. T. (2006). *Mental Workload Measurement*, pages 814–817. International Encyclopedia of Ergonomics and Human Factors. CRC Press.