Thinking Outside the (Beige) Box: Personal Information Management Beyond the Desktop

Manas Tungare
Center for Human-Computer Interaction
Virginia Tech
Blacksburg VA USA.
manas@tungare.name

Manuel Pérez-Quiñones
Center for Human-Computer Interaction
Virginia Tech
Blacksburg VA USA.
perez@vt.edu

ABSTRACT
As personal information ventures out with users from the desktop into the wild, researchers in personal information management are confronted with many new issues, and an amplification of the effects of existing issues. In this paper, we take a detailed look at the mobile context and how it differs from the stationary context, specifically in relation to PIM. Other changes in PIM, as information moves beyond the personal and into the social are also interesting to the PIM researcher. Based in part on comments made by participants of a study we conducted, and from our own experience in designing PIM for multiple devices and other scenarios, we highlight a few specific instances where existing PIM research falls short of addressing the issues, and pose a few open questions for the PIM community to consider as we move beyond the disappearing desktop.

Author Keywords
Personal Information Management, Multi-Device User Interfaces

ACM Classification Keywords
H.5.2 Information Interfaces and Presentation

INTRODUCTION

Humankind has been managing personal information of various types since before computers arrived on the horizon; e.g. office files, documents, charts, etc. The advent of computers not only brought forth more information, but also newer types of information (new collections). Word processor documents, spreadsheets, calendar events, etc. were among them. The rising popularity of computer networks, and specifically, the Internet, unleashed a torrent of information, making the user drink from the proverbial firehose. Email, Web bookmarks, digital music libraries, and other information collections burgeoned as more and more users connected to the Internet.

Today, we are witnessing a similar boom in personal information as we embrace the next wave of technological innovation: the arrival of mobile information devices. Technology is fading into the background and slowly disappearing [22]. As in the previous two technological revolutions, these devices bring with them the need to include newer information collections into the boundary of our personal information store. In addition, many users have begun to use multiple devices to manage their already-existing information, a trend that raises novel issues that were never encountered in a single device paradigm.

In this paper, we review some of the unique aspects of the mobile context that make it fundamentally different when it comes to managing personal information. We also look at some of the existing problems in PIM that are further complicated when multiple portable information devices are in use. We present some of the issues we are likely to face as we move forward into an age of handhelds and ubiquitous information capture devices such as the SenseCam. Backed by data from a survey we conducted in July/August 2007 and by some of our on-going work, we discuss some of these problems and hope to encourage discussion about the open questions posed.

INFORMATION BEYOND THE DESKTOP

To understand the far-reaching changes that mobile devices are bringing upon us, we present a few examples of ubiquitous information typically encountered in a mobile setting. These examples explicitly highlight problems that have not been adequately addressed by existing PIM research.

Electronic Billboards, Posters and Flyers

While walking along a street, driving on a highway, or strolling in a mall, we are surrounded by advertising billboards, posters announcing local events and flyers publicizing special deals at every nook and corner. When a person interested in these offerings notices the billboard, she might be interested in capturing some of the information for later use: e.g. date and time of the concert, phone number of the contact person, website of the manufacturer, etc. The use of small devices is particularly appropriate in this context. A quick photo of the billboard using a cell phone camera supports quick capture. However, this presents yet another challenge for the user to organize the information. The picture captured is not for the 'Pictures' collection but instead for a 'to do list' or a reminder for some other activity. The lightweight capture of information afforded by the small device in this example complicates information management at a later stage.

Music on the Car Radio

The car has already taken over the home as the primary location for radio listening among young adults [12]. In the context of driving a car, listening to the radio is a passive activity. However, if the user wants to "bookmark" a song, or

record a phone number from an advertisement, there is currently no easy way to do that. On this, and many other occasions, the mobile context does not leave enough attentional resources for users to be able to perform any complex tasks with their information [16]. Traditional ways of capture and keeping, such as typing a note, or writing on a Post-It, are extremely bad choices for a car driver.

Keeping and Organizing in a Mobile Context

The previous two examples highlight why users prefer to (or are forced to) postpone 'the keeping decision' [13] until they are no longer mobile and/or until after they are able to resume their access to a full computing environment. Since this information is organized at a different location than where it was encountered, several contextual cues that were available and handy at the moment of the encounter are no longer available when performing the keeping and organizing tasks. We believe that this will be a prevalent situation as more and more information is managed in a mobile context.

The Need for Lightweight Information Capture

In our previous examples, noting down and filing away details while walking on the road or listening to a car radio were not viable options. Thus, mobile users need a way to "tag", "pin", or bookmark items of information for deferred keeping. Portable information devices should support such lightweight interaction that does not require the user's full attention, but still enables future access to that information. Such interaction would support post-facto browsing and retrieval on a different device and make available the past context to avoid the location post-value recall problem.

Even within the desktop context, several tools have acknowledged and tried to address this problem. QuickSilver[1] and OmniFocus[2] both have a *quick entry* or *quick capture* mode, in which users can create notes quickly for later processing. A system of lightweight capture is encouraged by proponents of the Getting Things Done [1] approach as a way to reduce the cost of keeping without cluttering the permanent information store with poorly organized information. In all these approaches, an Inbox is used to collect notes and supports lightweight capture.

Location and Contextual Data

Some forms of information are captured at will in a mobile context, but organized (or edited) in another context later on. Examples include photographs, video footage, business cards received at conferences, voice memos, quick handwritten notes captured during interpersonal interaction, etc. This illustrates that PIM systems and principles that guide development of desktop tools can no longer function optimally in a mobile context.

Certain kinds of information — contextual data — are only captured in mobile settings: the best example is perhaps location tracking logs, also known as GPS (Global Positioning System) tracks. Many products on the market enable users to keep complete and detailed logs of their movement annotated by location information (latitude, longitude and altitude). Life-logging projects such as the idea of a Memex proposed by Vannevar Bush [6], Microsoft's MyLifeBits [8] and numerous projects involving the SenseCam (several examples are mentioned in [18]) generate huge amounts of contextual data that PIM researchers are interested in studying [7]. Moving forward, we will need efficient ways to organize, integrate, peruse through and search this information.

Another example of contextual information capture is found in photography. Photographs are often taken in mobile settings and some photographers automatically annotate their digital photos with location data from GPS track logs. By effectively combining these two disparate data collections, we can generate visualizations that allow photos to be superimposed on a map, along with other users' photos taken at nearby locations[3]. Integrating these sources of information within a personal information store still remains a research topic.

Other Examples

Many other types of information are intimately connected with mobile contexts (though they often also exist in stationary contexts.) Among these are phone calls and meta-data about them, voice mail [23], contact information exchanged via non-electronic means [24], and others. Ways and means of taking personal information from the desktop into the outside world, e.g. by writing notes for users based on their location, have also been developed [9, 19].

MOTIVATIONS FOR STUDYING THE MOBILE CONTEXT

Doing personal information management is slightly different when you bring into consideration the mobile context. There are several reasons for this. First, users might encounter information while mobile, that is, while being away from their main device. Second, users now have more devices with which to perform information management. Third, more and more information management is being done for reasons other than work. The following sections discuss these observations in some detail.

Increasing Variety and Adoption Rates of Devices

In July/August 2007, we conducted a survey study among knowledge workers (technology workers, university professors, graduate and undergraduate students) about the devices they use, activities they perform, groups of devices that are often used together, and how the failure of a device or the introduction of a new device influenced their practices. 220 participants took the survey; some expressed dissatisfaction about not being able to manage their information effectively with their many devices; some told us about the unique strategies and techniques they had devised to work around the problems they encountered. Although the results of that survey have been only partially analyzed at the time of this writing, our initial findings indicate that the variety of devices and their adoption rates are on the rise.

[1] http://www.blacktree.com

[2] http://www.omnigroup.com/applications/omnifocus/

[3] http://panoramio.com/

Many more knowledge workers reported using laptops (>96%) and cell phones (83%) than desktops (71%). 26.4% of all users reported using both laptops and cell phones (plus other devices as well.) Other recent studies conducted over a broader population also hint at the increasing diversity in the use of information devices [11]. The largest product launch of 2007, by some accounts [15], was the Apple iPhone, a mobile personal information management tool. Thus, not only is there an increasing variety of devices on the market, but users are also adopting several of these devices voraciously.

Many of these portable information devices are somewhat limited in processing power, but their limited input and output capabilities place them in a separate category. Although Moore's law has permitted a large amount of raw processing power to be compressed into smaller and smaller devices, the physical limitations encountered by input-output devices are much more difficult to overcome. These limitations make PIM on these devices an even more challenging task.

Not all Information Seekers are Knowledge Workers

After the arrival of computers, the image of the information seeker is that of a person working at a desk using a computer. With the presence of mobile devices, the information seeker of today is not at a desk, or at home, but out there, mingling among others.

Even the moniker *"knowledge worker"* is indicative of the fact that any interaction with knowledge is considered related to work. Today, we must also study knowledge consumers — information seekers — who perform PIM for reasons other than work, outside of the office, away from a desk. Some areas of study include: how people organize and share their music collections, photo libraries, grocery shopping lists, DVR recordings, etc. Of particular interest is how the *work* and *home* environments relate to each other. Do users keep completely separate identities when it comes to their information management practices, clearly delineated by their work/home lives?

Challenges Raised by the Mobile Context

Context has been recognized as an important influencing factor in the decisions of keeping and organizing information [2]. Mobility brings a whole new dimension to the concept of context. For example, Kristoffersen *et al* [14] identified four major factors that makes the mobile context more challenging than the stationary workplace or home, especially as it relates to the use of computing devices:

- Computing is often a **secondary task**; tasks external to operating the mobile computer are the most important.
- Users' **hands** are often used to manipulate physical objects, as opposed to users in the traditional office setting, whose hands are safely and ergonomically placed on the keyboard.
- Users' tasks may demand a high level of **visual attention** as opposed to the traditional office setting where a large degree of visual attention is usually directed at the computer.
- Users may be **highly mobile** during the task, as opposed to in the office, where doing and typing are often separated.

Furthermore, users are usually unable to devote sufficient attentional resources to keeping and organizing information encountered in a mobile context [16].

RESEARCH OPPORTUNITIES IN MOBILE PIM

In the next few sections, we present three general areas of research that arise because of the availability of ever more advanced mobile devices, the availability of ubiquitous information, and the users' need for and ability to manage their personal information in mobile settings.

No Device is an Island

We argue that a new step in the evolution of design practices is needed to make possible the synergistic use of multiple devices to fulfill a user's information needs. Designers now have the responsibility of thinking about multiple devices, information collections, and users' tasks when designing new devices. Thus, it is necessary to consider all platforms together and distribute or replicate functionality according to the affordances and contexts of use of each device [17]. This may require forfeiting interface-level consistency between two or more platforms in favor of presenting a holistic interface to the user. In our opinion, this represents a design approach unlike building multiple applications for the same desktop computer, where consistency is essential.

Portable devices today have been designed as disparate islands of information. Some devices and software were designed as functional replicates of their desktop counterparts (e.g. Microsoft Pocket Word and Pocket Excel), while others are designed without consideration for how they fit into a user's overall environment (e.g. many cell phones keep an independent copy of a user's contact list and are unable to synchronize with other devices that also keep similar lists). When many such devices are used together, functionality often overlaps instead of being mutually complementary [20]. To use these devices effectively, they must artificially be bridged together using synchronization software, network file systems, USB thumbdrives, or by sending files back to oneself using email.

From our survey (*N=220*), we learned that fully 78% of users use more than one device. However, a lot of these users had horror stories to narrate about how their individual devices failed to communicate amongst themselves and live happily together. Instead of providing added value, they created more problems. Users also reported halting the use of multiple devices in favor of a single one because of the inevitable synchronization problems rampant in a multi-device environment. E.g. the use of laptops (96%) has far surpassed the desktop (71%) for precisely this reason, as reported by participants.

Information Sub-Collections

Another area where mobility affects PIM is in how we think of our information collections. To make devices play better

with each other, it is not necessary to keep all information consistent over all devices. Strict synchronization of all information on each device with all the other devices owned by a user is generally not a good idea; e.g. a user may not wish to synchronize music from his desktop to his laptop, because it already exists on his music player, and is thus accessible on-the-go. To address this concern, we introduce the term *'sub-collection'* to describe a subset of an information collection that may be kept synchronized across devices.

In the survey, we encountered many different ways of how users think of their sub-collections. In one example, a user found that navigating the address book on his/her phone was too cumbersome, so they copied only a few phone numbers to the phone. Here is the direct quote from our study:

> *"I find my cell phone's contact navigation to be a real pain [...] — although there are quite a few people on there, I usually don't call them. Thus I find it tedious/ somewhat pointless to put more people on there — after all it will just cause me more pain when I am navigating to people I really want to call."*

In another case, the different components of an information collection are maintained in different devices. For example, the quote below shows how this user prefers to have name and phone number in the phone address book but name and email in the computer address book:

> *"Usually my contacts on the phone are just with numbers while my contacts on the computer are just with email addresses (makes sense since I'm using the former to make calls and the later to send emails). [...] The name of the contact is usually different for emails (e.g. full name instead of only first name or last name first or use of title in front of name.)"*

A third example shows how users might have different sub-collections based on the social role these play and on the device where they are used. Here, a user has a music sub-collection for exercising and another one for long trips.

> *"I have two MP3 players: A small one for the gym and large one for long travel, etc. and I do not have the same music on both of them. It is generally difficult to make the synchronization software for each player understand that I do not want it to grab my entire music library, only the portion that I want to send to that particular player."*

Other examples abound: while it is useful to keep complete past and future calendar data on a user's desktop computer, keeping a copy of all that information on a cell phone is wasteful of resources and negatively affects the time and attention it requires for common tasks such as looking up today's schedule. Synchronizing browser bookmarks between computers and mobile devices[4] is often sub-optimal, since the sites visited while mobile are completely different from the sites visited in a desktop context.

Overall, users think of their information collections in different ways when multiple devices are in use. From all the examples above, we see the need to split information up based on the device it will be accessed on and/or the purpose of accessing it. This is, in some ways, the opposite of the information fragmentation problem [3]. On one hand, personal information that is currently locked in individual silos based on their origin and type of data (i.e. information collections [4]) needs to be integrated. On the other hand, the above examples illustrate the need to meaningfully split collections into sub-collections according to some complex, idiosyncratic, and as yet unknown set of rules.

Information Fragmentation in Syncables

We are exploring a solution to the problem of information sub-collections in our design of the Syncables framework [21] by adding explicit support for sub-collections through configurable filters and transcoders.

The Syncables framework [21] provides applications a way to migrate their data and context across any number of devices, utilizing only an HTTP-based protocol for maximum device compatibility. Syncables addresses both problems, information fragmentation and sub-collections, as described below.

Syncables treats all data collections as parts of a single data tree. Each item has a unique 'address', and Syncables resolves these addresses across collections. Information is grouped into projects, such that information from multiple collections can be browsed together. A user can open calendar events such as *'sync://cluster/pim2008/calendar/2007/10/15/submission-deadline'* in exactly the same manner as opening a document *'sync://cluster/pim2008/files/latex/beyond-the-desktop.tex'*, or a bookmark *'sync://cluster/pim2008/bookmark/http/chi2008-.org/workshops'*. These address links can be embedded in other documents, so all three of these could be inside a to-do note. Clicking a link sends a request to the Syncables framework to open the corresponding information in the program best suited to display it.

Syncables allows filtering and transcoding content for suitability on the target platform by the insertion of pluggable modules into the information migration path. For example, a calendar filter would allow only events from the current month to be synchronized to a cell phone, while a document transcoder would transcode HTML documents to plain text when sending these to a reader device. Address book entries could be filtered to include only the phone numbers on a cell phone, with the entire entry on a desktop computer.

New Information Collections

Sub-collections are not the only challenge for mobile PIM. As the world ventures out with new information capture devices in tow, newer types of information will undoubtedly enter our already-burgeoning information stores. Though we already have collection managers for existing types of information, e.g. email clients, music browsers, photo man-

[4]E.g. Apple iPhone automatically synchronizes all bookmarks with the Safari browser on a Mac.

agement software, there is a need to develop specialized collection managers for the new collections.

The use of technologies such as GPS generates track logs containing information about us in four dimensions: latitude, longitude, altitude and time. As more and more people start tracking their movement, software such as Google Earth (that can let you view these tracks superimposed on digital photographs of the earth) may become part of their arsenal of personal information tools.

The SenseCam project from Microsoft[5] involves a camera that a user wears around her neck to capture all aspects of her environment to create a life log. Among the types of information captured by this contraption are light intensity, light color, body temperature, ambient temperature, audio level, GPS location, and data from a multiple-axis accelerometer. The current version of the SenseCam viewer lets the user view photos chronologically and the measured values as a graph. Future research in this direction will lead to visualizations that let the user slice and dice this information in many meaningful ways and trace life events any way she wants.

Beyond the Personal — Into the Social

Information is moving off the desktop in another significant way: it is moving from one user's desktop onto another user's desktop. A user is never alone; she is in the midst of other people, socializing, collaborating, learning, teaching, commuting, or simply hanging out at a bar. Some of these social interactions inevitably lead to information exchange, and this presents an interesting challenge for PIM researchers. Mobility also contributes to this: mobility engenders sociability.

Influence of the Social on the Personal

Most personal information enters a user's personal space from the external public space. Gwizdka suggests that characteristics of that public space influence the user's decision to keep and organize the received information [10]. But this public space is defined as "all information resources and channels out there", and is not necessarily social in nature; it is publicly accessible, but it need not involve the participation of any other users.

We believe that social information spaces (different from public spaces, as illustrated above) play an important role in a user's keeping, organizing and finding decisions. The permeability of the boundary between personal and social information spaces raises new issues that PIM research must look at, while tools that take advantage of this relationship can offer newer and better ways for users to manage their personal information.

Information Osmosis

Personal information is no longer restricted to personal spaces. It may be created, stored, managed, archived, and transmitted entirely outside the personal space, and brought into the personal space at will. Most information is eventually exchanged with other users. Even the etymological root of information, *'inform-'*, refers to the transmission of a message from one entity to another.

With the advent of ubiquitous electronic social spaces and our increasing dependence upon machines for information processing, it is but natural that information will flow from personal spaces to social spaces. We term this natural process of the movement of information across the permeable boundary between the personal and social spaces as *'information osmosis'*. The choice of term suggests that the two information spaces do not collapse into one, neither is there a strict wall between them. Instead, the boundary is permeable and allows the selective passage of information across it.

Some current examples of this flow of information include:

- People upload photos to Facebook and tag them with their friends' names. This causes the photo to appear also on the tagged user's profile, although it was not uploaded by that person. This reflects an ability of friends to "push" information into others' personal spaces, with or without the consent, or even knowledge, of the target.
- A similar pattern is seen in the use of corporate calendars. Although each user's individual calendar is his/her personal information, one's colleagues often schedule meetings by adding an event to her public calendar.
- Blogs consist of personal information published to the Web (social space). Regular readers track their favorite blogs by subscribing to feeds in their blog reader software, thus incorporating this information into their *own* personal spaces.
- Users prefer to upload videos (undeniably, part of their personal information) to public sites like YouTube, and then embed the video into their own blogs. This is a transition of the information from personal (the hard disk) to the social (YouTube) to personal (writing the blog post) to the social (publishing the blog post) to personal again (entering the reader's personal space).

This free flow of information between the personal and social spaces brings with it unique issues and amplifies the impact of existing ones.

Privacy

Privacy is undoubtedly an issue inherent in any kind of personal information management. But due to the closed nature of our individual information spaces, this issue has received relatively lesser attention from the research community.

There have been several incidences reported in the popular press about potential employers scanning social spaces for personal information about candidates they are looking to hire. Quite often, some bits of unsavory personal information sneak out into the wider world unintentionally, and may have serious implications for the person involved. A single casual email that may have been considered acceptable for

[5] https://research.microsoft.com/sendev/projects/sensecam/

personal use may wreak havoc upon the career of its unsuspecting sender if made available for public scrutiny.

What can PIM research do to understand this movement of information from personal to social, to channel it through acceptable means, without removing from the user the flexibility to use her data as she pleases?

Information Osmosis as a Method of Recommendation
Can we consider information that is purposely exported from the personal to the social as a recommendation, or a vote, from the user for that information? Many people maintain short lists of their links on their home page, usually to public web sites and their friends' personal pages. The purpose of such links is neither to manage them, nor to help others find them (because those web sites are already much more popular than the site that links to them, e.g. http://bbc.co.uk/) Is the act of providing such links, then, an indication of a recommendation from the user?

The Web, authored by a billion different people, has an intrinsic recommendation system based exactly on this property. Internet search engines such as Google rely on treating a hyperlink from one site as a personal recommendation for another site [5]. However, the creator of the first site did not intend the link as a vote; she simply considered it part of her personal Web site.

Motivation for Making Personal Information Social
What might be some of the motivations of users seeking to make their information available to their social group? Anecdotally, some users reported that their primary motivation was for sharing with others, but that was not the only reason. Other compelling arguments included the fact that information placed into a shared social space had greater longevity: by being placed in a social space, there was less likelihood of it getting corrupted or lost, and thus the perceived need for local backups was lower. Being able to access their information from a different computer than their own was also a crucial need (though it is obvious from this comment that users imagine a strong conceptual connection between two orthogonal ideas — storing information off the desktop computer, and making it socially available.) These ideas need to be studied better in order to harvest this phenomenon profitably in favor of users.

Augmenting Personal Information with Social Information
These trends are indicative of a larger phenomenon: that of the thinning of boundaries between the personal and social (or at least, the evolution of a more permeable boundary between the two.) How, then, can we take advantage of this to assist the user in performing information management better?

PIM tools can use the permeability of information to advantage by incorporating what we call *socially-assisted personal information management*. There already exist tools based on this principle, but there is also a vast untapped gold mine of interactions that researchers can exploit to advantage.

There exists a tool, Facebook Address Book Synchronizer[6] which compares email addresses, telephone numbers and photos in the local address book with information provided by the respective people on the social networking site, Facebook. If any outdated information is found locally, it is automatically updated with more recent information from Facebook user profiles.

Now imagine if when a user changes her personal information on her own address book, this change is automatically propagated to Facebook. The above software, in turn, propagates the change down to the user's friends. Thus, changes in personal information propagate to the social space to enhance and correct locally stored personal information of another user. As a result, personal information flow will move through the internet and personal information collections in much the way that Web 2.0 technologies quickly propagate other types of information today.

When importing music from a newly-purchased CD into a computer in MP3 format, track names are often not available from the disc alone. Software such as iTunes, WinAmp and several others consult the Gracenote CD database (CDDB), a collective social database of information about track names in CD albums. This information used to be the personal information of someone, until that person decided to share it with the world by uploading it to Gracenote CDDB. When it is downloaded by a second user to tag recently-imported music, it now becomes that person's personal information. Thus, information moves between a user's personal information space to the shared social space and may eventually enter another person's personal space.

What if future email clients could provide us similar functionality? Along with a message from Person X to Person Y, Y could also receive hints on filing it: "X has applied the following labels to this email: *research*, *pim* and *notes*. Would you like to apply the same labels?" Not all such labels would apply, but the thinking is that our tools can and should adapt to leverage some of the organizational strategies used by our peers, and these can be piggy-backed on top of the information they already send us.

Social Data Processing
Using individual people's personal information in aggregate form can yield benefits to the entire community in managing their own niche of personal information. Consider, for example, that certain kinds of spam filters are able to identify patterns in large email data sets by examining them collectively. When one recipient of a suspected spam message flags it as spam, these systems can automatically flag it as spam for the hundreds of others recipients of that (or similar) email.

Internet startup company Wesabe.com helps users manage their money: users have to provide their bank account details to Wesabe, and Wesabe will keep track of their finances. The web site looks at aggregate spending characteristics of its users and suggests them ways and means to save money.

[6] http://fsbsoftware.com/

These tips are based on other users' personal information, but suitably anonymized and aggregated to allay legitimate concerns about privacy.

CONCLUSION

Via the examples presented, we focused on some of the problems and issues that need to be tackled to take personal information management beyond the desktop. Several reasons highlight why the mobile context is different from the desktop context in relation to PIM: using a mobile device is not the primary task, and hands and visual attention are in use for the primary task. This presents opportunities for PIM researchers to study and analyze these issues, and possibly develop solutions for them. The three major areas of research that we believe are crucial to mobile PIM are: the complementarity of devices and how designers should design holistically for all the devices instead of each one alone; the need for information sub-collections and new collections as data gets split across multiple devices in meaningful ways, and understanding and harnessing the transition of information from personal to social and vice-versa.

ACKNOWLEDGMENTS

The authors would like to thank the members of the Dialogue Research Group at Virginia Tech, among whom a few of the topics presented in this paper were discussed. Some of the issues raised were brought up in conversations with Dr. Rob Capra and Dr. Pardha Pyla.

REFERENCES

1. D. Allen. *Getting Things Done: The Art of Stress-Free Productivity*. Penguin, 2002.

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

3. O. Bergman, R. Beyth-Marom, and R. Nachmias. 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, 2006. ACM Press.

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

5. S. Brin and L. Page. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 30(1-7):107–117, 1998.

6. V. Bush. As we may think. *The Atlantic Monthly*, 1945.

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

8. J. Gemmell, A. Aris, and R. Lueder. Telling Stories with MyLifeBits. *ICME 2005. IEEE International Conference on Multimedia and Expo*, 06(06):1536–1539, July 2005.

9. W. G. Griswold, P. Shanahan, S. W. Brown, R. Boyer, M. Ratto, R. B. Shapiro, and T. M. Truong. Activecampus: experiments in community-oriented ubiquitous computing. *Computer*, 37(10):73–81, 2004.

10. J. Gwizdka. Finding to keep and organize: Personal information collections as context. In *Proceedings of the 2nd Invitational Workshop on Personal Information Management at SIGIR 2006.*, 2006.

11. J. Horrigan. A typology of information and communication technology users. Technical report, Pew Internet and American Life Project, 2007.

12. F. Jacobs. The Bedroom Project. http://www.thebedroomstudy.com/, Last Accessed: October 2007.

13. W. P. Jones. Finders, keepers? the present and future perfect in support of personal information management. *First Monday*, 9(3), 2004.

14. S. Kristoffersen and F. Ljungberg. Making place to make IT work: Empirical explorations of HCI for Mobile CSCW. In *GROUP '99: Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work*, pages 276–285, New York, NY, USA, 1999. ACM Press.

15. A. Lipsman. Searches for Apple iPhone spike in weeks prior to launch. http://www.comscore.com/press/release.asp?press=1518, June 2007.

16. A. Oulasvirta, S. Tamminen, V. Roto, and J. Kuorelahti. Interaction in 4-second bursts: The fragmented nature of attentional resources in mobile hci. In *CHI '05: Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, pages 919–928, New York, NY, USA, 2005. ACM Press.

17. P. S. Pyla, M. Tungare, and M. Pérez-Quiñones. 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.*, 2006.

18. A. J. Sellen, A. Fogg, M. Aitken, S. Hodges, C. Rother, and K. Wood. Do life-logging technologies support memory for the past?: An experimental study using SenseCam. In *CHI '07: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 81–90, New York, NY, USA, 2007. ACM Press.

19. M. Tungare, I. Burbey, M. A. Pérez-Quiñones, and A. Raghavan. Evaluation of a location-linked notes system. In *Proceedings of the 44th ACM SouthEast Conference (ACM SE 2006)*, 2006.

20. M. Tungare, P. S. Pyla, M. Pérez-Quiñones, and S. Harrison. Personal information ecosystems and implications for design. Technical Report cs/0612081, ACM Computing Research Repository, 2006.

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

22. M. Weiser. The computer for the 21st century. *Scientific American*, 265(3):66–75, 1991.

23. S. Whittaker, R. Davis, J. Hirschberg, and U. Muller. Jotmail: a voicemail interface that enables you to see what was said. In *CHI '00: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 89–96, New York, NY, USA, 2000. ACM Press.

24. S. Whittaker, Q. Jones, and L. Terveen. Managing long term communications: conversation and contact management. *Proceedings of the 35th Annual Hawaii International Conference on System Sciences, 2002.*, pages 1070–1079, 2002.