"It's on my other computer!": Computing with Multiple Devices

David Dearman
Department of Computer Science
University of Toronto
Toronto, Ontario M5S 3G4
dearman@dgp.toronto.edu

Jeffrey S. Pierce
IBM Research
Almaden Research Center
650 Harry Rd, San Jose, CA 95120
jspierce@us.ibm.com

ABSTRACT
The number of computing devices that people use is growing. To gain a better understanding of why and how people use multiple devices, we interviewed 27 people from academia and industry. From these interviews we distill four primary findings. First, associating a user's activities with a particular device is problematic for multiple device users because many activities span multiple devices. Second, device use varies by user and circumstance; users assign different roles to devices both by choice and by constraint. Third, users in industry want to separate work and personal activities across work and personal devices, but they have difficulty doing so in practice Finally, users employ a variety of techniques for accessing information across devices, but there is room for improvement: participants reported managing information across their devices as the most challenging aspect of using multiple devices. We suggest opportunities to improve the user experience by focusing on the user rather than the applications and devices; making devices aware of their roles; and providing lighter-weight methods for transferring information, including synchronization services that engender more trust from users.

Author Keywords
User study, multiple devices, personal computing, personal information management, cross device interaction

ACM Classification Keywords
H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.

CHI 2008, April 5−10, 2008, Florence, Italy.

INTRODUCTION
The collections of computing devices that people use to support their personal and business activities are growing. Instead of a single personal computer, users now incorporate multiple digital devices into their lives, including desktop computers, laptops, mobile phones, digital cameras and media players. Users also increasingly engage in activities that span devices, rather than just using different devices for different tasks (Figure 1). The resources for these activities, both information and applications, span multiple devices, and they may even be located elsewhere on the Internet [21].

Employing multiple devices improves access to information and computation, but it requires managing information and activities across many devices, each with different limitations and affordances. Managing personal information and files is a significant issue for a single device [3, 10, 14]; multiple devices exacerbate the issue.

Solutions to support managing information and activities across devices [3, 11, 13, 17, 22] are emerging, but research explicitly studying user practices is still sparse; [13] is an exception. Our goal is to understand how users currently employ multiple devices in order to guide future research and the development of new applications and services.

We interviewed participants from IBM Research and a research group at Stanford University about how they use

Figure 1. A participant's collection of work devices. Her workflow encompasses all three laptop computers.

computing devices at work (or school) and in their personal lives. In this paper we synthesize our interview results to identify *what* devices people in a cross-section of industry and academia use, as well as *why* and *how* they use them. We also identify broader themes in multiple device use.

Highlights of our observations include:

- Participants frequently employ multiple devices to perform and support tasks. However, current devices complicate this practice by tying information and actions to the device, rather than to the user.
- Participants assign different devices different roles within a task, but devices have no notion that their actions are part of a larger task.
- Industry participants wanted to keep work and personal activities on different devices, but had some difficulty separating them in practice.
- Participants' greatest complaint about using multiple devices is the diffusion of information across them, despite the number and variety of available tools for transferring and managing information.

Our observations suggest opportunities to improve multiple device use. Current approaches to computing and information management are primarily device, rather than user, -centric. For example, users generate web browsing or call histories through interactions that often span devices, yet these histories are typically tied to the individual devices rather than integrated across them. Computing devices should go beyond the assumption that users only employ a single, personal computer and actively be aware of and coordinate with a user's other devices. Devices should also be more aware of their roles. Finally, users need better methods for sharing information between devices.

In the next section, we provide an overview of previous research on managing information and activities. We then describe our study design and methodology, and we discuss the results of our interviews in detail. We conclude by identifying opportunities to improve multiple device interaction and suggesting future research directions.

RELATED WORK

In 1945, Vannevar Bush published his vision of the memex [7], "*a device in which an individual stores all his books, records, and communications*". Bush's physical description of the memex is grounded in his time, but the notion of a device to augment our memory is a reality. However, rather than a single memex, the current landscape of computing is closer to Weiser's vision of ubiquitous computing [26]: a variety of devices with different form factors and purposes.

Multiple personal devices are increasingly the status quo. As our collection of devices grows, so do the issues associated with effectively and efficiently managing our information across those devices. Personal information management (PIM) is a complicated task [4, 18] involving the organization of digital information that we accumulate as files [4], email [27] and bookmarks [1]. However, research on information management has primarily focused on single devices.

Barreau and Nardi [4] show that the strategy for managing information depends on how it is used. They note that a location-based filing strategy is most common because it facilitates recall. However, information used across multiple devices may have no salient home, complicating recall. Bergman *et al.* [6] show that information management across applications (e.g., email and folder hierarchies) can lead to *information fragmentation*. Multiple devices exacerbate fragmentation because projects may span devices. Whittaker and Sidner [27] describe the problem of *email overload*, where email is used to support secondary tasks, causing management problems that overwhelm the user [5]. The overload problem may transcend email into other applications as users search for ways to support communication and sharing between devices. Abrams *et al.* [1] comment on users' management strategies for web bookmarks, showing that the majority of users employ no strategy. This lack of a strategy could be problematic as resources like bookmarks that are useful across devices start to span those devices.

The type of information we manage is diverse, making a single management solution difficult. Rather than unifying PIM into a single solution, Dumais *et al.* [12] developed Stuff I've Seen (SIS) to support searching across information contexts such as email, files, contacts, etc. Although SIS is a single device application, Dumais comments that the most requested feature is *"... unified access across multiple devices"*, suggesting that cross-device support is a significant design consideration.

While how best to support activities that span multiple devices is an open question, researchers are developing tools and techniques to connect devices (*e.g.* Schilit's ensembles [24]) and transfer information (*e.g.* Miller & Myer's synchronized clipboards [19], Satya's Internet Suspend and Resume [23]) and interacting across them (*e.g.* Bandelloni & Paterno's migratory interfaces [2]). Online services are also emerging to provide limited searching (*e.g.,* Google Desktop [14]) and synchronization (*e.g.,* Google Browser Sync [13]) of information across devices, or to centralize online access (e.g., del.icio.us [11], Google Docs [15], Buzzword [8]). Each of the aforementioned projects is a step forward, but they only address a piece of a larger puzzle. To build an effective solution and complete the puzzle, we first need to understand *why* people connect their devices, and *how* they use them.

Oulasvirta and Sumari's [21] study of mobile information workers, which focused explicitly on exploring how the workers employ multiple, heterogeneous computing devices (laptops, smart phones, and mobile phones), is the closest precursor to our work. They reported that multiple devices are a significant consideration in the preparation and

Table 1. The different devices disclosed by participants. The number of devices and the percentage of the 170 disclosed devices are presented for each device type.

Device Type	Laptop	Desktop	Mobile Phone	Smart Phone	Digital Camera	Music player	PDA	Server	Other
# (%)	59 (36%)	31 (18%)	19 (11%)	10 (6%)	15 (9%)	11 (6%)	4 (2%)	9 (5%)	12 (7%)

execution of daily activities. While the workers felt that multiple devices offer benefits, managing their devices and the information on them was a constant problem, forcing them to anticipate their future needs. We extend this work by incorporating a larger pool of participants from academia and industry, and by exploring how our participants use a larger heterogeneous collection of devices to accomplish both work and personal tasks.

THE STUDY

The term "multiple devices" is broad, abstract and could conceivably include any digital device that someone uses. For the purpose of this study, we chose to focus primarily on conventional computing devices such as laptops, desktops, and mobile phones. However, we do include additional devices (e.g. digital cameras, portable media players, external hard drives, and USB flash drives) where relevant if a participant explicitly mentioned them.

Method

The study consisted of a single session semi-structured interview conducted in each participant's work environment (generally at their desk). Sessions typically took between 45 and 75 minutes. We collected notes, pictures and audio recordings during each session, and produced full transcripts of the audio recordings. We conducted the interviews in July and August of 2007.

We chose the interview questions based on 12 informal interviews with employees from several business and research divisions at IBM. We drew participants for the study from two different computer science research communities: IBM Research and Stanford University. Our focus is on researchers because they are typically early technology and device adopters. As such, they provide a window on practices and problems before they spread to the population at large. We chose researchers from industry and academia because our initial informal interviews suggested that practices differ across the two communities.

During an interview session, we worked with the participant to compile a complete list of the devices for which they are the primary user (either at work/school or home). We next asked participants about the configuration of and interactions between each device. The interview questions probed: tasks performed on each device; tasks performed across devices; the role of devices; interactions between devices; information specific to a device; information shared across devices; and application usage.

We ended each interview by asking the participant to describe the three worst and three best things about working with multiple devices. These questions served in part to inspire further dialog and elicit comments that may not have come across during the main body of the interview.

Participants

We interviewed 27 participants (19 male, 8 female); 21 from IBM Research and six from a research group at Stanford University. The 21 industry participants included 11 staff researchers, eight summer interns, two engineers and one analyst. The six academic participants included three graduate students, two faculty members and one undergraduate in his final year.

Participants varied in age; 10 were between 20-30 years; nine between 31-40 years; five between 41-50 years; two between 51-60 years; and one greater than 60 years. The level of education for the participants was high; 11 had a doctoral degree (PhD); 10 had a master's degree (MS); five are college graduates and one had some college.

We recruited participants through a combination of direct contact and word of mouth. Working with multiple devices was not a pre-condition for participating in the study. However, none of the people we contacted or who contacted us employed only a single device.

OBSERVATIONS

Before presenting our observations, we note that semi-structured interviews are an excellent tool for exploring experiences. The structured questions provide initial discussion points, and allow for unstructured follow-up questions probing interesting comments and unique experiences. However, no interview session is exactly the same, making it difficult to quantify a specific belief, experience, or behavior across all participants. Therefore when reporting our observations we may indicate, for example, how many users mentioned a specific problem. Unless otherwise noted, such a report does *not* mean other participants did not encounter a similar problem; it merely means they did not mention it during the interview.

Participants' Device Collection

The number and type of devices participants disclosed as part of their device collection varied. The average device collection consists of 5.96 devices (6.30 including servers), with a minimum of three and a maximum of 11 devices. Table 1 groups all 170 devices participants disclosed by their device type. Device types include desktop computers, laptops (including tablets), cell phones, smart phones, digital cameras, iPods, personal digital assistants (PDAs), file and network servers, voice recorders, portable media (USB disk and flash drives), and portable GPS devices.

We further categorized each device by its context of use: work/school, home, and intermediate. Work and school are analogous for the industry and academic participants, so we group them together. Intermediate devices frequently travel between work/school and home, so we consider them a separate category.

Figure 2 shows the number of devices for a mathematically average participant, categorized by device type (condensed) and context of use. We assign each device in a participant's collection to a single context of use and therefore do not report each device more than once. On average each participant has: one laptop or desktop computer (1.11) at work/school; one laptop or desktop computer (1.48) at home; one cellular device (1.07); and at least one other portable device (1.56: most frequently a digital camera or iPod). The majority of participants (20) also have a laptop computer (0.74) that they take between work/school and home on a semi-daily, if not daily basis.

The number of portable devices (4.80) such as laptop computers, mobile phones and digital cameras, outnumbers fixed desktop computers (1.15). The average number of home (1.48) devices is greater than the number of work/school (1.11) devices. The large number of intermediate (3.70) devices suggests that the majority of participant's devices are used at work/school and home.

Computers located exclusively at work/school or home generally belong to the owner of the site. For example, a desktop at work/school is typically an institutional computer, while a desktop at home is typically a personal or family computer. Participants primarily owned their small mobile devices (e.g. mobile phones and cameras) and unanimously reported using them across contexts. In contrast, participant's institutions primarily own the intermediate laptops (which 20 participants take between work/school and home).

We note that two participants currently use two cellular devices (hence the average of 1.07). Most participants only have one cellular device, but these two participants reported using multiple because they are currently transitioning from an old phone to a new phone.

Why Participants Use Multiple Devices

Interacting with and across multiple devices is a daily routine for our participants. In this section, we outline why they reported using multiple devices.

Form Factor and Device Affordances

All devices are not created equal; some are better suited for a particular activity than others. The physical design (e.g. display size and orientation) and the modes of interacting with a device (e.g. keyboard, mouse, stylus, finger, etc.) influence the suitability of a device for a specific activity. Fifteen participants highlighted this fact, commenting that using multiple devices provides access to an array of devices with different form factors and affordances.

"[The tablet PC] is easier to carry and more comfortable than the 17 inch laptop."

For this participant, he would rather use his smaller tablet PC rather than his larger 17-inch windscreen laptop when reading at home. His laptop is more powerful and has a larger display, but the tablet PC provides a small, comfortable form factor (similar to a book) and allows reading in portrait mode.

In contrast, another participant commented that at work he uses a secondary laptop (which he described as a desktop) that is permanently connected to a large display to read his email, rather than using his primary laptop where he performs the majority of his work related tasks:

"... it is the big screen that makes this a communication center during office hours."

Portability

Mobile devices allow participants to choose the setting they perform tasks. Portability (implicitly the ability to employ mobile devices as well as more capable fixed ones) is reported unanimously as the "best thing" about working with multiple devices. Woodruff's study of laptop use in the home [28], shows that laptop computers are used in a limited number of convenient settings. The previous participant who prefers to use his tablet PC to read also commented that the tablet's small size makes it easy to read in a variety of locations, such as a comfortable sofa or bed.

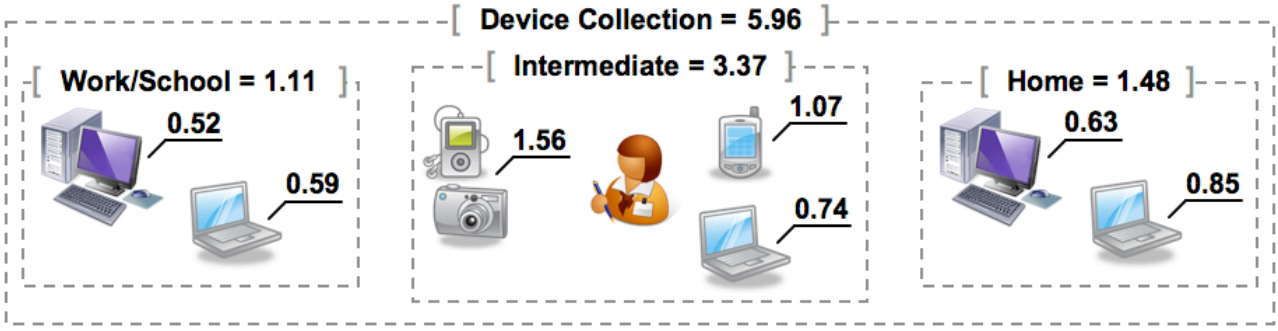

Figure 2. The device collection of a mathematically average participant, categorized by device type and context of use. The three contexts of use are work/school, home, and intermediate (devices that travel between school/work and home). The decimal values represent the average number of devices a participant uses of the respective type in the particular context.

Another participant commented that he uses his laptop to read news and blog sites in the morning, not because it is the best device for reading, but because it allows him to sit at the kitchen table with his girlfriend rather than being tied to his desk.

"I use my laptop to read my blogs in the morning not because it is really the ideal device for it. It has the smallest screen that I could reasonable read on, but it can occur in the location I want it to occur ... sitting with her at the dinner table."

Task Completion Time

Not all devices are created equal. Several participants mentioned that they chose devices for a task based on how quickly they perceive they can complete the task with the chosen devices. For example, one participant commented that when at home, given the choice between his laptop and his wife's laptop, he would choose to use his wife's laptop for quick tasks such as surfing the web.

"Mine takes forever to get started. If we wanted to pay an account, she has all the logins setup automatically, and I don't have those. If it was something I needed to do quickly I would use hers."

His choice is based partly on the speed of the hardware, but is also based on the different configurations of the devices; her computer starts faster than his and she saves passwords to allow faster logging into sites.

A Computer for Work and for Home

All of the participants have at least one computer provided to them by their institution (work or school) and at least one computer that they personally own. Participants in industry regarded separate home and work devices as particularly desirable, with 14 of them commenting that separate devices help ensure that "… w*ork is work, home is home."*

Separating work and personal activities across work and home computers provides an obvious benefit. Devices located at work are in an environment that the institution controls, increasing the security of their information and resources. Personal computers in the home provide a locus for activities and information (i.e., managing financial or medical information, storing personal media collections, etc.) that are private or inappropriate for work.

Participants in academia did not share the desire to keep a distinction between home and work devices. All of the academic participants described mechanisms they use to ensure that school-related materials are available to them no matter when or where they desire to work.

Software and Operating System Differences

Five participants commented that they use multiple devices in part because software they need or prefer to use is specific to an operating system that is different than their primary one. For example, one participant indicated that his primary computer is a MacBook Pro, but he uses a Windows PC to write papers because he prefers the Windows implementation of Microsoft Office (even though he also has Microsoft Office on his MacBook).

Special Purpose vs. General Purpose

Small portable electronic devices can be categorized as special purpose or general purpose. Special purpose devices such as a digital camera have a defined role; in this case, taking pictures. General-purpose devices such as a mobile phone, by contrast, have a variety of roles. The integration of multiple features to create a general-purpose device is not without cost. Despite increasing device convergence, many participants preferred using multiple special-purpose devices rather than a single general one.

For example, eight of the 17 participants that use a mobile phone (as opposed to a smart phone) indicated they prefer to use a digital camera rather than their phone's camera to take a picture or record a video. They indicated the picture quality is better with a digital camera and that the additional features of their phone are incredibly awkward to access.

"After I got the phone there was one situation I wanted to [take a] picture ... I wanted to get it to my laptop and apparently it came with some software for me to do that, but conveniently it didn't come with the cable to do it. There are alternatives probably on the web, but I didn't bother."

For this participant, the mentioned photo still resided only on his mobile phone at the time of the interview. Two other participants expressed a similar experience. The additional functionality could not replace the quality and easy of use of the specialized device.

Transitioning from an Old Device to a New Device

Rather than upgrade existing computers over time to better support more demanding tasks, users instead tend to purchase completely new computers. Four participants were currently transitioning from an old device to a new device at the time of the interview. However, rather than completely replacing the old device, three participants instead repurposed the old device. In effect, the old devices have entered what we call a *permanent transitional state*.

"There are three years of software and research on the XP [old] machine. 80% of which I can run on Vista [new machine] and the other 20% I can't."

For this participant, the incompatibility between Windows XP and Vista, and the sheer volume of research material on his XP computer, resulted in a division of his work environment across two computers. He described the transition from the old to the new as *"never-ending*" because many of his projects require code and libraries that will not run on Vista without re-writing them.

How Participants Employ Multiple Devices

Just as participants use multiple devices for a variety of reasons, they also use multiple devices in a variety of ways.

Blurring the Boundaries between Work and Home

We mentioned that our participants from industry use multiple devices in part to keep a separation between their work and home activities.

"I would rather keep all my work stuff on my work computer and not have to worry about [work and personal information] co-mingling."

Industry participants expressed three reasons for preserving this separation for work and home devices:

1. The perception that their company requires work information stay on a work device, *"I don't want to have to deal with the policies that are involved with having [work] information on my personal computer."*
2. A personal desire to keep a boundary between work and personal/family time, *"I have a family. I have children. If I am at my kid's games I don't want to be beeped. I am off the clock."*
3. A family computer is not a safe place for work: *"... my daughter could climb up and start hitting keys."*

However, the boundary between work and home is stronger in theory than in practice [20]. Thirteen of the 14 industry participants who desired to keep work and home devices distinct indicated that they do work tasks on home computers. One participant commented that when an idea comes to him at home he acts on it, using a personal device if he did not bring his work laptop home.

"If I want to work, [I use] whatever machine is handy."

Keeping a distinction between work and home is particularly problematic for mobile devices. When participants bring their institutionally owned laptop home, they frequently use a home computer to support their work activities. One participant commented that when writing a paper at home he prefers to use his home computer because it has a larger display. Alternately, some participants use their home computer as, in essence, a secondary or ambient display to search the web or monitor email. Some participants also reported using software on their personal computer that is not available on their institutional computer to accomplish a work task.

The participants in academia expressed a similar preference for keeping work and home distinct. However, their focus is on preserving time for personal activities, a stark contrast to the desire of the industry participants.

"Coming into work [school] is more like getting into work mode. There is a permeable and vaguely defined boundary ... more activities and context then the information."

The academic participants expressed no belief that their work should or is required to remain at school; at most they commented that school provides an environment more conducive to work. In line with this belief, academic participants reported a much higher level of connectivity between their school and home environments. Five of the academic participants make their data accessible from both home and school. In the most extreme case, one participant runs a personal Microsoft Exchange server to keep his PIM information consistent across his school desktop, home desktop, laptop, and Blackberry. In addition, the participant uses a school file server to provide universal file access across his devices. This participant justified his configuration as:

"I believe that any device I use should be throwable in the trash at any time and nothing gets lost."

Another academic participant indicated that he keeps a personal desktop computer at school (purchased with his own funds) to run file and version control servers. His department did provide similar services, but he found their servers to be slow and unreliable. None of the industry participants describing dedicating personal resources to supporting work activities.

While our academic participants are on average younger, we believe that their difference in behavior is due to a difference in attitude between academia and industry. As evidence, we note that eight of our industry participants are college interns, all of whom segregate their activities across work and personal devices (although in some cases the personal devices are actually owned by their universities!). Furthermore, even the faculty participants drew little, if any, distinction, between personal and school devices.

Devices Take on Roles within Work Flows

Participants reported frequently dividing tasks among several devices, assigning each device a role. Figure 3 shows the work environment of a participant who uses both a desktop computer with multiple monitors and a laptop to perform his work tasks. Participants mentioned three reasons why they divide a task across multiple devices: 1) a task requires a complex environment; 2) to allocate secondary tasks to a non-primary device; and 3) hardware or software constraints force the division of a task.

The task requires a complex environment. Some tasks are more complicated than others, requiring multiple devices to perform them effectively. One participant described using

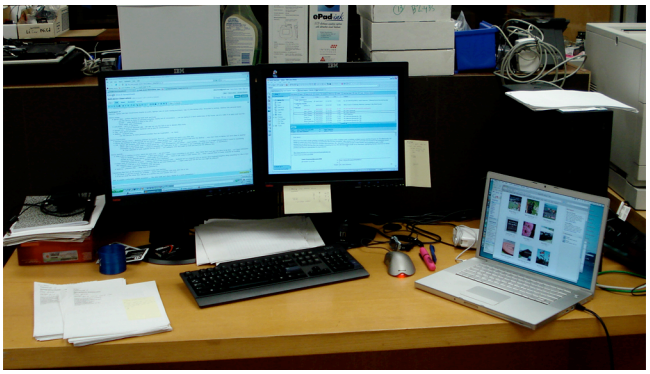

Figure 3. A participant using a laptop to support writing a paper on a multiple-monitor desktop.

two computers to develop and test her code.

"I ... code on this machine and then test on the other ..."

She develops exclusively on her primary laptop because she can write code in her chosen environment while retaining easy access to her other work information. She tests exclusively on a dedicated "testing" laptop because she can ensure its configuration remains consistent between tests.

Allocate secondary tasks to a non-primary device. Some tasks require more screen real estate for simultaneously viewing information than some individual devices can provide. When using multiple devices as multiple displays, participants often indicated designating a primary computer (typically on the basis of speed and screen space) while assigning other computers a supporting role.

Grudin [16] observed that multiple monitor users leverage a second monitor to support monitoring information and performing peripheral tasks. However, when using multiple devices, participants described carefully choosing tasks (or sub-tasks) for a secondary device in order to minimize the need to transfer information between the devices.

"I do the side by side thing, but I don't directly share anything between them"

"I do not like to context switch [between machines] ... there is too much effort involved"

Participants commonly described using a secondary device to monitor email and browse web sites. Their choices, combined with their comments, suggest that the overhead involved in transferring information between multiple devices is a significant barrier and an active influence on the strategies for partitioning tasks across devices.

Constraints require the use of multiple devices. In some cases participants did not actively choose to use multiple computers for a task; instead, software and/or hardware limitations led to it. Ten participants commented that they transfer tasks (and supporting information) across devices because part of a task requires software that is not available, supported, or installed on the primary computer.

"I have a legitimate copy of Photoshop ... at work, but I can't find the CD. So when I switched to my new computer I can't get the pass code correct, so I can't use it. So if I have to do Photoshop work, I send it to my home machine and do it on that."

Seven participants specifically mentioned software licensing as a factor when allocating tasks to a device. However, while some licenses do restrict installation to a single device, others allow installation on both a fixed and a mobile device simultaneously.

What Participants Do on Many Devices

Participants reported few activities that they could or would perform across all or most of their devices. Writing and viewing documents are two activities that participants did mention frequently performing across many of their devices, typically using a suite of applications like Microsoft Office. Web browsing is even more common, performed by all participants on all of their computers.

Web-based services enable information access that breaks the barrier between work and home devices for industrial participants. Many participants mentioned using a browser to check personal email at work. In part, this may stem from a perception that the information does not transfer to the device; one industry participant commented that when using Firefox to check her personal email at work *" ... it is like the data isn't really there. Just maybe in a cache somewhere ..."*.

The participants from academia are particularly active users of web-based services. All of the academic participants described heavy use of on-line services (e.g. Gmail, Google Documents, Remember the Milk, Meebo, Plaxo, Facebook, Buzzword, etc.) across the majority of their computing devices, making no effort to keep school and personal information and activities separate. Indeed, some academic participants went so far as to forward all of their school email to an online email account; an action that would no doubt send corporate IT departments into shock.

Managing Information across Multiple Devices

All 27 participants expressed concern and frustration over the difficulty in managing their information and activities across multiple devices. They did not, however, lack for mechanisms. Participants reported using a combination of:

- Portable media such as a USB memory stick, hard drive, CD, etc. Portable media lets users move large amounts of information, but only if it resides in files.
- Emailing a file or text to herself. Although easy, email is problematic for large (or large quantities of) files.
- Sharing directories over a network. Shared directories require configuring network access in advance and only support file transfer.
- Services on an external server such as third party offerings like Plaxo and Google Documents, or locally hosted services like Microsoft Exchange Server and Subversion. These services typically only host certain types of files or information, and third party services in particular raise the issue of privacy.

Participants reported assembling their own patchwork of solutions to manage information across their devices. No one technique could support all information types or usage patterns. The solutions varied widely, with no two participants using the exact same combination of services.

Third party services like Plaxo [22] and Google Documents are becoming more prevalent. Eleven participants commented that they would like to adopt a more centralized method of information management, either by moving data into the "cloud" or by carrying a portable drive around.

"I got an idea during class. So I put it up on Google Docs so I don't forget it on one machine. [The idea is] always there unless I delete it. There is almost no chance to forget where it is. On [my] laptop it is sometimes hard to find files. I lose them. Was it put on this [my laptops] hard drive, or one of the external drives, or laptop at home?"

Centralized solutions may in theory simplify managing activities across devices, but they are not a panacea. External servers require network access and may suffer from slow transfer speeds or unexpected outages. Portable drives, on the other hand, force users to diligently carry the drives everywhere they go; if they forget the drive the information is unavailable. User may be reluctant to use third party services, even if they can provide the desired functionality, because of privacy and how the service will manage and use their information.

"I have heard of services like Plaxo but I don't trust them ... I heard they spam you ... I have never had a good experience with online contact sharing services ... either they let you synchronize and they spam you or they don't let you synchronize."

Participants did comment on two online services that they found to be particularly successful: version control systems and email. Version control systems (e.g. CVS) are sufficiently successful that some participants applied them to manage all (or most) of their files across devices.

Participants' only complaint about email is the occasional service outage that forces them to use a different email service in the interim. In particular, participants dislike that the metadata and messages sent or read with the alternate email service are hard to access when returning to the original service.

Interaction Histories are Information Too

A common problem participants reported with information transfer mechanisms is that they focus on application data. Indeed, most mechanisms focus exclusively on files, although some online services support finer granularities: to-do list items, contact information, browser bookmarks, etc. Our participants, however, also wanted to transfer their interaction histories across devices. Eight participants commented that the interaction history used for auto-completion in web browsers and email clients would be useful across devices.

"... the history list and auto-complete ... they are completely randomly distributed [across computers] depending on who I interacted with on which computer. From my point of view I have interacted with [person], but from my computer's point of view it depends if I interacted on the desktop or laptop."

Interaction histories, such as chat logs and browser and call histories can be difficult to share between devices because the information is typically application-specific. However, when users engage in similar activities across their devices, these interaction histories can be useful across devices. Participants reported instances where their desire to have a complete interaction history and maintain its consistency is significant enough to change their behaviour:

"The reason I don't install IM here [at work] is that I always log my chats and I don't want [the logs] distributed across my multiple machines."

Similarly, another participant kept his chat interaction history consistent and accessible across his devices by utilizing VNC [25] to remotely connect to his desktop. Through VNC he could use the IM client running on his desktop while actually working at another computer.

Interaction histories also include actions to configure applications. When participants use the same application on multiple devices they must configure multiple copies. While that need is already a significant barrier for some participants, the problem is exacerbated by the trend toward pushing software updates over the Internet. Three participants commented that they regard updating the software on all their computers as a serious hassle.

"One thing that drives me up the wall is that ... something is always blinking at me. I have three machines and they all have an Adobe suite, a Microsoft suite and the Apple stuff ... when an update comes from anyone of them I have to click 'Yes' three times."

One participant added that because he does not use his computers equally, updates accumulate for less frequently used devices. This often results in a significant delay in completing his tasks while he waits for the operating system and applications to update.

Automatic File Synchronization is problematic

Automatic file synchronization is surprisingly absent from participants' information management strategy. The file synchronization mechanisms that participants did use were manual: they would store files in a central location (either online or on a portable drive) and manually synchronize the latest versions there.

The lack of automatic file synchronization adoption is not due to lack of awareness; participants expressed knowledge of a variety of automatic mechanisms. Participants' comments suggest that a lack of trust in automatic mechanisms may be the primary underlying cause. One participant commented that:

"It is scary, [you are] not sure if you can trust it ... I always do a preview."

Synchronization with one device is too limiting

In general participants commented that mobile devices such as media players and smart phones do a better job of synchronizing information than larger computers. That difference makes sense: mobile devices have limited input capabilities and are thus reliant on transferring information from more capable devices. However, mobile devices still

have room for improvement. The most common complaint from participants is that many mobile devices are only able to synchronize with a single computer.

"The iPod is something that I am disappointed with because you can only sync it with one computer, which is my desktop at home. […] I listen to podcasts all the time, it is pretty much all I use my iPod for. So it is easiest to sync it with my desktop every morning, but when I travel I take my laptop with me. Now I could download them, well, there is just no good way of doing it. I don't know a way of doing it where I can sync it with multiple devices."

Industry participants commented on the lack of partial synchronization for their personally owned mobile devices (e.g., phones) that they use for both business and personal tasks. As one participant put it, *"carrying two phones makes little sense"*. However, because devices typically only offer all-or-nothing synchronization, participants must synchronize personal contacts with a work computer, synchronize work contacts with a personal computer, or go without synchronization altogether.

OPPORTUNITIES FOR IMPROVEMENT

The results of our interviews highlight the complexity of working with multiple devices. In this section we discuss opportunities for improvement based on our observations that might facilitate a more seamless experience.

Focus on users, not devices

The predominant assumption in computing is that users interact with a single computer. The complaints of our participants reflect the consequences of this assumption. Software licenses restrict the installation of applications, forcing users to divide their activity across devices or install the software on a non-optimal device. Operating systems do not explicitly recognize a user's multiple devices, forcing the user to manage information across their devices. Applications require updates that users must accept repeatedly for each installation on each device, and make it hard to transfer the users interaction history across devices.

Researchers have proposed supporting activities that span multiple applications [10]. However, we need to go farther and give up the assumption of the single personal computer. Support for users' activities should be neither application- nor device-centric. Our devices should instead collaborate to support a notion of user-centric activities that may span multiple devices as well as multiple applications. While there are initial steps in this direction [3, 11, 13, 22], they must support a wider variety of activities and fully recognize the members of a user's device collection.

Awareness of Roles and Context of Use

A device's role in workflow is an important consideration for information management that is currently overlooked. If devices are aware of their unique and/or common roles, the devices themselves could actively support information and application management.

Consider, for example, a user writing a paper on his desktop and using his laptop to monitor email. When writing the user may switch to using the laptop to perform an unrelated task, making the laptop the primary device and the desktop a secondary device. The devices could identify the role change and move the email application from the laptop to the desktop, allowing the user to continue to monitor his email uninterrupted.

The participants' usage of multiple devices in parallel is an important design consideration for solutions that manage information across devices. Approaches like Internet Suspend and Resume [23] and SoulPad [9], which assume that users employ multiple devices serially to interact with a single virtual computer, may be problematic in practice.

Devices should explicitly support separate work and personal roles. If a user moves work information to a personal device (or vice versa) the devices could register the transfer and track the information. Each device could then help users keep track of the latest version of their information and remind them to return it to the work device when they finish with it on the personal device.

Lighter-weight Information Transfer

Numerous solutions exist for transferring information, but no solution is appropriate for all situations and all types of information. Our results suggest that even multiple solutions are not sufficient. All 27 participants described managing and accessing information across their devices as a significant problem or concern. Ideally, users need a lightweight mechanism for information transfer that works for all types of information and supports their privacy requirements. In addition, we need to consider a broader spectrum of information such as the metadata (*e.g.* browser history or call history) generated through interaction that can be useful for future interactions on other devices.

Improve Synchronization

Automatic file synchronization has the potential to simplify managing information and roles across devices, but only if users are willing to employ it. Our findings suggest that people do not trust automatic file synchronization, even though they employ automatic synchronization for other types of information: music, email messages, contact information, calendar data, and task lists.

We believe that the lack of trust in automatic file synchronization is due in part to the higher cost of failure. If a user loses an email or calendar entry, the consequences are relatively minor, whereas losing a file that contains hours of work is much more traumatic. Synchronization needs to engender trust by the user that the service will act appropriately. We believe this could be accomplished in though greater visibility: showing the user what will happen, what is happening, what has happened, and allowing them to easily undo an action.

CONCLUSION

Using multiple devices is increasingly the norm. We need to better understand how people use multiple devices so that we can design a better user experience for working across them. We interviewed 27 multiple device users from industry and academia as a step toward developing that understanding.

Our interviews yielded four primary findings. First, associating a user's activities with a particular device, rather than with the user of the device, is problematic for multiple device users. Second, users do not use all of their devices in the same ways; they assign different roles to devices both by choice and by constraint. Third, users in industry want to separate work and personal activities across work and personal devices, but they have difficulty doing so in practice. Finally, users employ a wide variety of techniques for accessing information across devices, but there is still room for improvement: participants reported that managing information across their devices as the worst part of using multiple devices.

We suggest opportunities to improve the user experience by focusing on the user rather than on applications and devices; making devices aware of their roles (both within activities and as work or personal devices); and providing lighter-weight methods for transferring information, including synchronization services that engender more trust from users. While our findings and suggestions may seem like common sense, our participants would doubtlessly assert that their devices could use a dose of common sense.

REFERENCES

1. Abrams, D., Baecker, R. and Chignell, M. Information archiving with bookmarks: personal Web space construction and organization. In *Proc. CHI 1998*, ACM Press/Addison-Wesley Publishing Co. (1998), 41-48.
2. Bandelloni, R. and Paternò, F. Migratory user interfaces able to adapt to various interaction platforms. *International Journal of Human Computer Studies*, *60*, 5-6 (2004), 621-639.
3. Bardram, J.E. Activity-based computing: support for mobility and collaboration in ubiquitous computing. *Personal and Ubiquitous Computing*, *9*, 5 (2005), 312-322.
4. Barreau, D. and Nardi, B.A. Finding and reminding: file organization from the desktop. *SIGCHI Bulletin*, *27*, 3 (1995), 39-43.
5. Bellotti, V., Ducheneaut, N., Howard, M., Smith, I. and Grinter, R.E. Quality Versus Quantity: E-Mail-Centric Task Management and Its Relation With Overload. *Human-Computer Interaction*, *20*, 1-2 (2005), 89-138.
6. Bergman, O., Beyth-Marom, R. and Nachmias, R. The project fragmentation problem in personal information management. In *Proc. CHI 2006*, ACM Press (2006), 271-274.
7. Bush, V. As we may think. *The Atlantic Monthly*, *176*, 1 (1945), 101-108.
8. Buzzword. www.buzzword.com.
9. Caceres, R., Carter, C., Narayanaswami, C. and Raghunath, M. Reincarnating PCs with portable SoulPads. In *Proc. MobiSys 2005*, ACM (2005), 65-78.
10. Cypher, A. The structure of users' activities. in Norman, D. and Draper, S. eds. *User Centered System Design, New Perspectives on Human-Computer Interaction*, Lawrence Erlbaum Associates, 1986.
11. del.icio.us. http://del.icio.us.
12. Dumais, S., Cutrell, E., Cadiz, J., Jancke, G., Sarin, R. and Robbins, D.C. Stuff I've seen: a system for personal information retrieval and re-use. In *Proc. SIGIR 2003*, ACM Press (2003), 72-79.
13. Google Browser Sync. www.google.com/tools/firefox/browsersync/.
14. Google Desktop. http://desktop.google.com/.
15. Google Docs. http://docs.google.com/.
16. Grudin, J. Partitioning digital worlds: focal and peripheral awareness in multiple monitor use. In *Proc. CHI 2001*, ACM Press (2001), 458-465.
17. Hinkley, K. Synchronous gestures for multiple persons and computers. In *Proc. UIST 2003*, ACM Press (2003), 149-158.
18. Lansdale, M. The psychology of personal information management. *Applied Ergonomics*, *19*, 1 (1988), 55-66.
19. Miller, R.C. and Myers, B.A. Synchronizing clipboards of multiple computers. In *Proc. UIST 1999*, ACM (1999), 65-66.
20. Nippert-Eng, C.E. *Home and Work: Negotiating Boundaries through Everyday Life*. University of Chicago Press, 1996.
21. Oulasvirta, A. and Sumari, L. Mobile kits and laptop trays: managing multiple devices in mobile information work. In *Proc. CHI 2007*, ACM Press (2007), 1127-1136.
22. Plaxo. www.plaxo.com.
23. Satyanarayanan, M., Kozuch, M.A., Helfrich, C.J. and O'Hallaron, D.R. Towards Seamless Mobility on Pervasive Hardware. *Pervasive and Mobile Computing*, *1*, 2 (2005), 157-189.
24. Schilit, B.N. and Sengupta, U. Device Ensembles. *Computer*, *37*, 12 (2004), 56-64.
25. RealVNC. www.realvnc.com.
26. Weiser, M. The Computer for the 21st Century. *Scientific American*, *265*, 3 (1991), 94-104.
27. Whittaker, S. and Sinder, C. Email overload: exploring personal information management of email In *Proc. CHI 1996*, ACM Press (1996), 276-283.
28. Woodruff, A., Anderson, K., Mainwaring, S.D. and Aipperspach, R. Portable, But Not Mobile: A Study of Wireless Laptops in the Home. In *Proc. Pervasive 2008*, Springer-Verlag (2007), 216-233.