Information Re-finding Through Physical-Social Contexts

Blake Sawyer
Virginia Tech
2202 Kraft Drive
Blacksburg, VA 24060 USA
basawyer@vt.edu

Francis Quek
Virginia Tech
2202 Kraft Drive
Blacksburg, VA 24060 USA
quek@cs.vt.edu

Wai Choong Wong
National University of Singapore
4 Engineering Drive 3
Singapore 117576
elewwcl@nus.edu.sg

Mehul Motani
National University of Singapore
4 Engineering Drive 3
Singapore 117576
elemm@nus.edu.sg

Sharon Lynn Chu Yew Yee
Virginia Tech
2202 Kraft Drive
Blacksburg, VA 24060 USA
sharilyn@vt.edu

Manuel Pérez-Quiñones
Virginia Tech
2202 Kraft Drive
Blacksburg, VA 24060 USA
perez@cs.vt.edu

ABSTRACT

Modern operating systems allow users to organize and re-find information using many contextual keys such as timestamps, content, custom tags, origin (e.g., a website) and people (i.e., through email). As humans naturally engage in activities with people, we want to investigate how a person's physical-social interactions, in the form of collocation, can be employed to support information re-finding. We designed a system that associates digital information used in social situations with co-present individuals. The system makes use of automatic or manual tagging and enables users to re-find information based on people and/or groups that were present at the time the information was used. This paper presents findings from three, two-month long case studies of individuals using our system. All three participants organized and accessed information using the context of a person or group. However, real-time use of digital information during social situations differed among participants and affected our system's overall use. Our results show that our initial idea is promising. In the future we would like to expand our studies to many users who collaborate with each other. We discuss improvements to our system and additional research questions we want to investigate.

Keywords

personal information re-finding, social interactions, collocation, context, tagging

Categories and Subject Descriptors

H5.2.m [**Information interfaces and presentation: User Interfaces.**]: Evaluation/Methodology and Graphical User Interfaces.

General Terms

Human Factors and Design

1. INTRODUCTION

Humans organize their activities around different parts of their social world. However, with the exception of information in email clients or a conscious effort by the user, a person's information space shows little resemblance to their social world. Our research explores how a person's physical-social interactions (i.e., collocated interactions) may be connected to their ever growing digital-information space. More specifically, we ask the following research question: *How can our physical-social activity be employed to support our digital information management needs?* We posit that associating a user's social world with his/her digital information may provide a natural way for organizing and re-finding files. To explore this research question, we designed a system that connects physical-social interaction with one's use of information and evaluated this in extended-use case studies with three individuals.

2. RE-FINDING USING SOCIAL CONTEXT

Research in supporting information filing and re-finding [19, 2, 26, 15] have already impacted the way we think about information access. Modern operating systems allow users to organize and re-find information using many contextual keys such as timestamps, content, custom tags [4], origin (e.g., a website) and people (i.e., through email). With the growth of mobile computing, tags can include the location of where information is encountered or encoded (i.e., geo-tagging, -caching, -blogging, etc.) These contextual keys enable users to use a step-by-step approach to re-find a piece of information, a process called *orienteering* by Teevan et al. [23]. This orienteering behavior was also observed in the use of the "*Stuff I've Seen*" application [10] and in re-finding information on the web [8]. Our research adds a new contextual key based on the user's *physical-social interactions*.

Re-finding using the social context of information is limited either through a user's conscious filing decision or through information shared in virtual communication tools (e.g., email). Both the ContactMap [25] and SNARF [13] applications,

for example, support information tasks by allowing users to think in terms of the people and groups they interact with in email. This technique allows the user to perform tasks for managing workflows, projects and goals more easily. Elsweiler et al.'s study on human memory in email shows that sender information is one dominant way by which users remember their emails [11]. These research only address information shared virtually and downplays the importance of real-time physical-social interactions. While our focus is on using the context of physical-social interactions for information re-finding, our system does include sender/receiver information for e-mail and attachment re-finding. This allows us to also support virtual interactions as well.

3. SYSTEM DESIGN AND IMPLEMENTATION

Our research question elicited four core requirements for our system:

1. The physical-social interactions of the user need to be tracked;
2. The user's information-related activities and the use of 'information objects' must be tracked;
3. User information objects should be tagged based on the overlap of the user's social and information interactions; and,
4. The physical context of the social interactions should be used for re-finding information.

We accomplish these four requirements by using a list of objects called *social orbits*. The user configures these objects and then uses them to tag and re-find their information.

3.1 Configuring Social Orbits

Social orbits are similar to Facebook's groups and Google+'s circles. This list of social orbits is the user's ego-centric network [1] where each social orbit is made up of one or more people that the user interacts with in some way. People in a user's network can belong to more than one social orbit. By default each person added to the list is also a single-person social orbit. The names of each social orbit are used for tagging the user's information objects. During a re-finding task, the user can use the person or group name they created to *orienteer* and browse for a piece of information on their computer.

Our system requires that the user explicitly create the social orbits for his/her use. The automatic discovery of one's social network has been the subject of much research in information mining [14, 18]. However, we are concerned with the *use* of social interactions as a kind of *handle* to re-find information. Automatically mined social networks may be too expansive and unwieldy for such use. Furthermore, research has shown that at any one time, a person's active set of social interactions typically number in the range of 7-12 groups [1]. This "live" [21] network of people is continuously changing on a need-to-need basis and is a much smaller subset of all the people one knows. Hence, making conscious and deliberate configuration of one's social orbits renders them more practicable and effective as handles for information access.

3.2 Tagging with Social Orbits

We implemented our prototype system on Mac OS X 10.5+, which allows us to use DTrace [9] to monitor information activity, and Spotlight [22] to tag different information objects. DTrace supports the deployment of system probes that inform us of such information activity as files being opened, written, and closed, e-mails being received along with their attachments, and web pages being accessed. These constitute 'meaningful' information activity associated with specific information objects. As discussed below, our system supports dynamic tracking of the physical-social presence of individuals in one's social orbits. If at least one person in a social orbit is present, we deem the social orbit to be present with a 'strength of presence' measured by the proportion of the members of that orbit who are presently co-located. When an information object is accessed, it is tagged with all social orbits that are wholly or partially present.

We chose to tag documents in Mac OS X Spotlight because, like [24], we believe it's critical to not disrupt a user's normal information management practices (i.e., filing files, managing folders, etc.). Our system does not care where a file is stored and our Spotlight tags can be used in queries the same way a user searches by filename or content.

The application determines which social orbits are present using two methods:

1. **Automatic**: A light-weight custom awareness device was seeded to participants and people in their social orbits. Our prototype device employs the low-power XBee radios [27] that can report the presence of other registered similar XBee-based devices. This device automatically informs the user's application of who is present so that information can be tagged with the appropriate social orbit participants. The system also automatically tags emails and attachments using sender/receiver information, independent of who is present.
2. **Manual**: The user is able to set people, or whole social orbits, as being co-present through an interface. Manual control overrides the automatic awareness device.

3.3 Re-finding with Social Orbits

Figure 1 shows the user interface of the social orbits application. It takes the form of a slide-out drawer that activates with a mouse-over action on the black bar on the left of the user's screen. This makes the interface less intrusive. What is seen in this figure is the list of social orbits, first arranged by 'presence', shaded in green, and then alphabetically. This list can be scrolled to show non-present social orbits. Double-clicking a social orbit can either activate or deactivate it as being present, overriding any triggering from the automatic awareness device. Right-clicking a social orbit brings up a Spotlight SmartFolder that lists all information objects that are tagged with that social orbit. These SmartFolders are simply saved search folders; users can interact with them in the same way as with other standard folders, e.g., ordering by modified date.

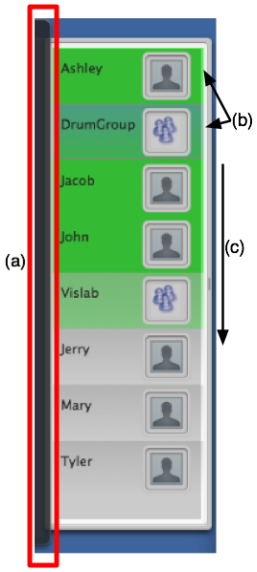

Figure 1: The interface consists of a list view that slides out when moused over; the black bar (a) is all that is seen when the application is not used. Each social orbit has an icon and name (b) and they are ordered by number of people from the social orbit who is present (c).

4. CASE STUDY

4.1 Methods, Procedures and Data Collection

Typical evaluations of information systems, including applications for re-finding, are problematic. Our hypothesis is grounded on the idea that social connections bring meaning to our information objects, and this meaning participates in the process of re-finding. This suggests that a staged, 'one-shot' study is not appropriate to explore how one's sense of social meaning helps in information re-finding. Evaluation of our application requires real use, whereby real social connections are available and meaningful information use is meaningful. Therefore, we chose to take an in vivo approach [3] to evaluate our application. We seeded the application with three professors in our Center. While the "professor/academic environment" may not be fully representative of all real-world use scenarios, our three participants exhibit a wide range of information management behaviors that may allow us to gain greater insight. They run three of the largest laboratories in our Center, and they have to juggle multiple projects and attend many meetings throughout their day. In brief, they are representative of users who have a need to re-find documents many times during the study period, providing the most opportunities to evaluate our application and hypothesis.

Each participant was asked to run our application on their laptops for two months. They were given a tutorial on how to create and use multiple social orbits for tagging and re-finding information. The participant and a number of people in their respective social orbits were given the automatic sensing device. Due to the small number of devices built (10), some social orbits relied only on manual activation by the user. The three case studies were run sequentially, beginning with one participant then moving to the other two participants a few months later.

The application collected timestamped logs of all tracked social orbits, names of information objects accessed and tagged, and social orbit SmartFolders opened. We also presented the participants with a use-diary to log their experience with the system. The participants provided comments on system-use and the usability of the system. Our first participant volunteered to do an hourly log to give us a fine-grained understanding of the workings of our system, and to report on usability issues. Our other two participants were polled for diary entries on a daily basis. We also conducted pre-interviews to get a sense of the participant's information practices prior to the study, including how they organize and re-find digital information and how they use it in social situations such as meetings. Lastly, exit interviews were conducted to probe deeper into the system's overall adoption, use and effects.

4.2 Prior Filing/Re-finding Usage

The three participants were distinctively different in their information filing strategies, and in their use of information during meetings:

Participant 1 (P1) is a loose filer/spring-cleaner. He employs many folder structures, organized by people, projects, papers, grants, etc. However, he noted that he does have trouble keeping up with his structure, especially when dealing with information that cut across multiple domains – e.g., a paper related to a project given to him by a student could be filed in any of three (papers, projects, and students) folders. He also mentioned that he often leaves information "inside" his email client because it makes them "easier to search for". His meetings often take place in a secondary office equipped with a large 56" display. He prefers to use documents such as outlines, progress reports, data analysis reports, etc. on this display to "drive" the meetings. Prior to the study, he used a computer that was hooked up to the large display in his secondary office for the collaborations. He would synchronize collaboration folders with his laptop that he uses as his primary machine. For the study, he stopped using the collaboration machine, and simply hooked up his laptop that was running the social orbits program for meetings. This was so that the social orbit tags and SmartFolders could be used to support both his personal work and meetings.

Participant 2 (P2) is a strict filer who has clear rules for where different information should go. Similar to P1, he used many folder structures such as projects, people, papers, etc. However, different from P1, he filed new items immediately and kept a clean desktop and workspace. His email client contained "two-hundred [folders] at least, with several levels deep in hierarchy". He has "a pretty good mental map of where all those folders are," and is "pretty confident that (he) can get to any topic very quickly." The only information he admitted to not organizing was his downloads folder saying, "I don't feel the need really to organize that stuff for the most part because I know that I can get to it easily with search." During meetings he would seldom use any form of digital information. When this did happen, he used a note-taking application that he goes "straight to." Each note represented a meeting – "I have 79 notes...that represent a lot more meetings because I get a lot of my regular

meetings in just one note, where I just keep appending." He said that re-finding these notes was never a problem for him. Other information used at his meetings was displayed on the computer screen of another person – "...they have their laptop, and they open up on their laptop and we looked at it together. So I'm not necessarily doing anything with documents on my laptop."

Participant 3 (P3) has a somewhat unconventional information management behavior, in the sense that he was "not good at being organized." His strategies consisted of piling everything on the desktop until it was full, then dumping all files into a "Desktops" folder. Interestingly, his method of re-finding information was through his students – "So often times...whenever I need a piece of information I don't go looking for it here [points to computer], I go ask one of them for it." His information use during meetings is similar to P2 – "When I meet with people, I tend not to work that way, I tend not to work with stuff on here [points to computer]." Like P2, he also noted that sometimes his students or collaborators would discuss information on their own laptop screens.

P1 used the system for two months with eight collaborators carrying the automatic awareness device. P2 and P3 used our system for two months concurrently. P2 and P3 each had four people with devices, but some of these collaborators worked with both P2 and P3. We note that the social orbits used by the participants can and did exceed the number of individuals who carried the automatic awareness devices, through the use of the manual 'click on/off' method of marking collaborators' presence.

4.3 Adoption and Use Experience

Table 1 summarizes results from the system logs for each participant. The social orbit configurations for each participant can be seen in Figure 2. The number of people and groups each participant added was relatively small and each participant also accessed about the same number of total information objects. We now discuss each participant's experience with the application that we gathered from the diaries and exit interviews.

	P1	P2	P3
Days Installed	56	52	50
Days Used	48	10	5
Social Orbits	15(5)	9(2)	6(1)
Accessed Objects	1945	2148	2360
Tagged Objects	747(410)	36(24)	49(35)
SmartFolders Opened	48	5	2

Table 1: Summary of system logs for each participant. The social orbits row's notation is total number of people with number of groups in parenthesis, adding the number of groups and people together will be the total number of social orbits. The tagged objects row notation represents the total number of objects tagged where the number in parenthesis is the number of objects tagged automatically through the device and/or email.

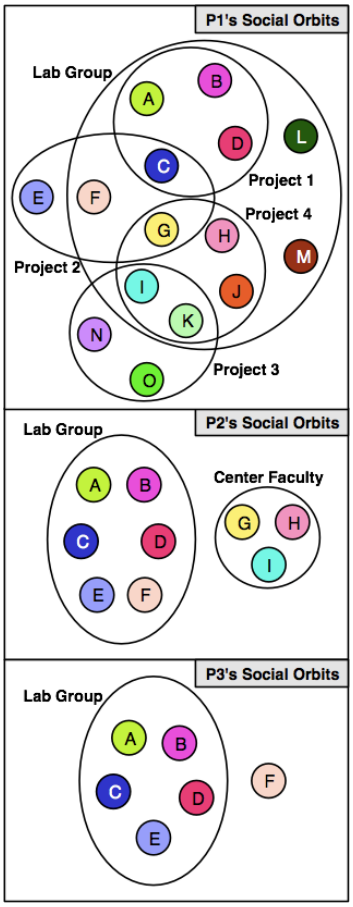

Figure 2: Illustration of the three participant's social orbit configuration.

P1: From the logs and exit interviews, P1 clearly used the system more than P2 and P3. In his logs, it can be seen that he had more information objects tagged either manually or automatically. He also opened a social orbit SmartFolder about once a day. In his exit interview, he noted that he often left these SmartFolders open and used them much more than once a day. He created most of his social orbits during the first two weeks. As his work changed, he added more or changed older ones. P1 liked how the application allowed him to "decouple" his tasks of information storage and information re-finding. He did not have to remember where he filed a piece of information, a process he highlighted as bothersome in his pre-interview. With the system, he relied on the context of how he used it and not on a "filing decision" he made days or weeks ago. Another interesting point is that on several occasions, P1 set social orbits to present even while alone. He used this feature to "remind" himself about sharing a webpage or paper with someone the next day when he popped up their SmartFolder.

P2: P2's use of the system was somewhat limited as compared to P1. After a couple of weeks, his system use declined due to conference travels and vacations. During the early parts of the study the participant set up his social orbits and preferred to use the manual method to set them to present. Even though he accessed about the same num-

ber of information objects as P1, only a few were actually used in the course of meetings (and were thus not tagged). It was not "too much of a burden having to remember to double-click those people" for him, but due to his practice of not using much digital information during meetings, setting people to present was not very useful. For the few files that were tagged, he did not use the application to re-find them. For example, if it was a note, he would go straight to the note-taking application or if it was an email, he would go straight to the email client, reducing his "motivation to keep double-clicking on people and tagging things." The few times that he used a SmartFolder was more out of curiosity and not for re-finding a piece of information. He found that he rarely had an occasion when his "normal [re-finding] mechanisms" would not work.

P3: P3 used the application for the first week and, after conference travel, did not go back to using it. During the first week of use, he created his social orbits and used both the automatic awareness device and manual method to tag information in a few meetings. After this first week, however, his use of the system stopped. He said he never remembered to use the device or set people to present. He also never used the application to re-find documents. When asked why, he said, "in the heat of the moment when I'm trying to find something, that thought wouldn't occur to me...I'm just not aware of that possibility." The majority of information objects tagged were done automatically through email activity. During the exit interview, we reviewed some social orbit SmartFolders. P3 was surprised "how we knew about this stuff" because the information all made sense in terms of relating to the correct person.

4.4 Participant 1 Re-visited

With regard to each participant's prior information practices, all three participants used both people and groups to organize a large portion of their information. It could be argued that a large part of P1 and P2's folder structures can be broken down or re-organized into either a person and/or group (e.g., papers, grants, classes, etc.) P3, while not having any folder structures, relied heavily on students to re-find information. His first instinct was to "go ask one of them" for it. With this type of behavior used by all three participants, our system had some potential to be useful.

However, a significant challenge to our hypothesis, that real-time use of digital information in physical-social interaction provides a means for information organization, arises from the non-use of digital information during meetings by P2 and P3. Without co-temporal physical-social activity and information use, there is just nothing to tag even when social configurations are tracked. This seems counter-intuitive in our increasingly digital world. A corollary question is how digital information use may become more integrated into our physical-social interactions. A clue to answer this question may be seen in the difference between P1 and P2, both of whom organized their information along the lines of people they work with, and activities (such as projects, paper writing, and proposal preparation).

One critical difference between the physical-social interactions of P1 and P2 is in the way P1 mediates his meeting activity with the use of a large display. We had the serendipitous opportunity to explore this difference, while still running our studies with P2 and P3.

As we described, P1 holds meetings in a secondary office with a large screen display. This display provides a point of shared focus among all present at a meeting. This practice also aligned with our initial assumption that people tend to use digital information during meetings, which explains how our application was more useful to P1. After P1's study period was over and toward the end of P2 and P3's study, P1 happened to lose the use of his large display for a period of about four weeks. This allowed us to explore how this changed his behavior in terms of using digital information in physical-social interactions.

P1's use of digital information dropped suddenly with the loss of his large display. His practices began to resemble that of P2 and P3 during meetings. He showed us the different notes he kept for each meeting and they had not been updated during the period when the large display was absent. The only time some form of digital information was used in a meeting was when the student used their laptop. This behavior, demonstrated by all three participants, led us to consider an important question: how does a shared focus display affect the use of our system and overall information use during physical-social interactions?

We posit that the large display creates a locus of digital interaction for the meeting, in a sense allowing the digital world a foothold in the physical-social configuration. By materializing the information in a 'physical way' the data become social, therefore incentivizing their use. This insight is analogous to the way social media facilitated the explosion of digital photography. An article in the New York Times [17] describes how photographs on Facebook increased while Flickr's use diminished. The former makes photographs social while the latter focuses on the technical aspects of picture tagging and finding. Users being able to tag and thus browse photos by friends (and friends of friends) creates a natural way for people to organize, share and even discover (i.e., navigating to friends of friends' photos). When people could share their pictures, it made more sense to take them in the first place – nostalgia brings to mind the Polaroid camera.

Returning to P2's post interview, we find that he did take notes on his laptop when he led faculty discussions and was responsible for coordinating activities, but not when "[he] is just a spectator at the meeting" or when "[he is] not in charge." Digital notes became useful when he had the responsibility to share them with other meeting members. In other words, it appears that P2 uses digital information when it serves some *social purpose*. This extends Erickson and Kellogg's assertion that we need new software tools to support social processes [12], by adding the need to support use of digital information in real-time physical-social interactions. Our observations raise the possibility that technologies that afford information use in a physical and social way are still absent from widespread use. This provides tremendous opportunity for further research and innovation.

Returning to our initial hypothesis, information use in physical-social interaction may well support re-finding and reuse.

This may, however, be predicated upon the introduction of real-time tools for physical collaboration. In the future, we plan to investigate this idea further. We also plan to scale our system and studies to a larger group of users who all interact with one another (e.g., lab group, industry team, etc.) In the next section, we discuss the logistics for running a study such as this and additional research questions we would like to explore in these studies.

5. FUTURE WORK

Our overall goal is to investigate what role a user's social interactions play in organizing and accessing information, with an emphasis on the user's physical-social interactions. The system and three case studies provide the initial ground work for investigating this research question, yet there is one main limitation. In each case study, we observed one individual working with their social orbits (i.e., a 1xN study). We want to scale this to an NxN study, where we observe many individuals working together, each using our system on their computing devices. This type of study requires extensions to our overall system design and implementation and also brings to light more concrete research questions related to our main question.

5.1 System Design

To scale up our studies to an NxN configuration, we plan to improve our user interface in three areas and redesign our physical sensing devices that will better support many users interacting with each other.

5.1.1 User Interface

Our studies are limited to users who run Mac OS X, so we would like to begin porting our system to Windows 7. Windows 7 now supports tagging of files similar to Mac OS X. While DTrace does not run on Windows 7, debugging software such as Windows Event Tracing [20] can be employed for information activity tracking.

In addition to porting our system to Windows 7, we would like to extend our user experience in three ways:

1. **Temporal Re-finding**: Our system only adds tags to files and users re-find by browsing a large folder for a particular social orbit. In addition to re-finding using social context, we would like to capture and use more temporal information. For example, presenting a calendar view of social interactions with any information accessed. We also would like to support some type of decay or strength of tagging data (e.g., partial social orbit tagging, file loses tagging information over time, etc.)

2. **Social Orbit Management**: We need an interface where the user creates and maintains their social orbits. This also includes assisting the user in their creation and maintenance. We mention we do not automatically create social orbits for the user. However, we would like to investigate how our system could assist a user during this process (e.g., suggest a new person or new group to add, detect when you have not met with someone in a while, etc.), building off research that mines social interactions.

3. **Social Orbit Sharing**: All social orbits, device IDs, files and tags are stored on the user's computer. For our studies, we would like to move the social orbits and device IDs databases to an online server. Users will be able to reveal their device IDs to a person they know. We also envision users being able to share their social orbit configurations with other related users. This could even be built on top of Facebook groups or Google+ circles.

5.1.2 Sensing Device

Our case study devices were built for studying one individual, carrying the master device, who is interacting with multiple peripheral people that carry a slave device. For our NxN study, we need one small device that serves both purposes. We designed a new device that obtains its power and connectivity to the user's computer by USB. In the parlance of our pilot design, it functions in 'master mode', informing the user's computer of the dynamic social configuration via XBee ID resolution while it charges the device's battery. When the device is disconnected from USB, it will function in 'slave mode' as a tracked device. In slave mode, the device will also keep a log of the user's dynamic social configuration. To ensure its use, the device has to be made small enough to carry in one's pocket and easily plugged in via USB.

We stress that this device is to enable our studies in advance of widespread adoption of a low-powered awareness technology. We anticipate that our everyday devices will soon support this type of proximity awareness. For example, new Bluetooth software will support several proximity profiles for sensing when other devices are nearby [5]. Also, the Bluetooth Smart chips will be more low-powered where a single button-cell battery can keep a device running for over a year [6].

5.2 Research Questions

Using our findings from our case studies we now have some more corollary questions related to our main research topic. We want to investigate these further through future NxN studies with our proposed system discussed above.

Main Research Questions: Is social presence an attribute that can be effectively used for information organization? Are users aware of the social context of when information was used? Can social proximity be one more handle to access information, much like an email sender is today?

- How does the use of real-time collaboration resources (i.e., shared large displays) affect the usefulness and impact of our system?
- What are the effects of our system in an NxN use environment? How does our system change users' information practices?
- What can users learn from others' social orbits?
- What other information can we share among users besides their social orbits? Can users share their social interactions? If so, can we share tags and information? How can we use social orbit tagging with information in the cloud?

- If we log users' social and information activities, how can we use this data to improve the users' information management? Can we use [16] and predict [7] locations, social interactions and information needs from patterns of user activity (i.e. pre-finding)?

- Finally, How will issues of trust and privacy impact our overall system design and usage?

6. CONCLUSIONS

Previous research in socially-based information re-finding have concentrated on the social interactions that occur in the cloud, i.e., email. We are interested in employing the context of physical-social interactions as a way to index a person's information. Our prototype application enables users to tag and re-find their information objects based on people and groups they physically interact with. We presented our findings from the two-month long studies of three participants using our application. All three participants had different ways of managing their information and yet all demonstrated some form of information access using people. However, the amount of digital information used during physical-social interactions affected the overall impact of our application. We want to explore this notion of using a large display as the focus point for interacting and sharing digital information. Much of the digital information we create and organize is meant to be shared and we may not have technologies that properly support this type of real-time information use and organization. In the future we would also like to expand our system and studies to many users who all collaborate with each other (i.e., NxN study) which will help us investigate our more concrete research questions.

7. ACKNOWLEDGEMENTS

We would like to thank all of our participants, including those who carried a device, for their time and useful feedback. We also thank the additional authors below for providing a fresh look at our research and giving guidance while formulating our future work.

This work is partially support by project grant NRF2007IDM-IDM002-069 on "Live Spaces" from the IDM Project Office, Media Development Authority of Singapore and through the project grant NSF CNS- 0551610 and IIS-0954048.

8. ADDITIONAL AUTHORS

Tom Martin (ECE Department at Virginia Tech, email: `tl-martin@vt.edu`), Ingrid Burbey (Hume Center for National Security and Technology, email: `iburbey@vt.edu`) and Lisa McNair (Department of Engineering Education at Virginia Tech, email: `lmcnair@vt.edu`).

9. REFERENCES

[1] P. Adams. Communication mapping: understanding anyone's social network in 60 minutes. In *Proceedings of the 2007 conference on Designing for User eXperiences*, DUX '07, pages 7:1–7:8, New York, NY, USA, 2007. ACM.

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

[3] V. Bellotti and I. Smith. Informing the design of an information management system with iterative fieldwork. In *Proceedings of the 3rd conference on Designing interactive systems: processes, practices, methods, and techniques*, DIS '00, pages 227–237, New York, NY, USA, 2000. ACM.

[4] O. Bergman, R. Beyth-Marom, and R. Nachmias. The user-subjective approach to personal information management systems. *J. Am. Soc. Inf. Sci. Technol.*, 54:872–878, June 2003.

[5] S. Bluetooth. Bluetooth sig finalizes proximity profiles for bluetooth version 4.0. http://www.bluetooth.com/pages/press-releases-detail.aspx?itemid=134, 2011.

[6] S. Bluetooth. Bluetooth smart: Low power. http://www.bluetooth.com/pages/low-energy.aspx, 2011.

[7] I. Burbey and T. L. Martin. When will you be at the office? predicting future locations and times. In *Second International ICST Conference on Mobile Computing, Applications, and Services*, MobiCASE 2010, New York, NY, USA, 2010. ACM.

[8] I. Capra, R.G. and M. Perez-Quinones. Using web search engines to find and refind information. *Computer*, 38(10):36 – 42, 2005.

[9] DTrace. Sun microsystems inc. http://opensolaris.org/os/community/dtrace/, 1995-2011.

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

[11] D. Elsweiler, M. Harvey, and M. Hacker. Understanding re-finding behavior in naturalistic email interaction logs. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information*, SIGIR '11, pages 35–44, New York, NY, USA, 2011. ACM.

[12] T. Erickson and W. A. Kellogg. Social translucence: an approach to designing systems that support social processes. *ACM Trans. Comput.-Hum. Interact.*, 7(1):59–83, 2000.

[13] D. Fisher, A. J. Brush, B. Hogan, M. Smith, and A. Jacobs. Using social metadata in email triage: lessons from the field. In *Proceedings of the 2007 conference on Human interface: Part II*, pages 13–22, Berlin, Heidelberg, 2007. Springer-Verlag.

[14] D. Fisher and P. Dourish. Social and temporal structures in everyday collaboration. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 551–558, New York, NY, USA, 2004. ACM.

[15] W. Jones. *Keeping Found Things Found: The Study and Practice of Personal Information Management (Interactive Technologies) (Interactive Technologies)*. Morgan Kaufmann, Nov. 2007.

[16] L. Kelly, D. Byrne, and G. J. F. Jones. The role of places and spaces in lifelog retrieval. *PIM Workshop*

2009, Vancouver, 2009.

[17] V. Kopytoff. At flickr, fending off rumors and facebook. *New York Times*, page B3, January 30, 2011.

[18] D. MacLean, S. Hangal, S. K. Teh, M. S. Lam, and J. Heer. Groups without tears: mining social topologies from email. In *Proceedings of the 16th international conference on Intelligent user interfaces*, IUI '11, pages 83–92, New York, NY, USA, 2011. ACM.

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

[20] C. Microsoft. Windows event tracing. http://msdn.microsoft.com/en-us/library/windows/desktop/bb968803(v=vs.85).aspx, 2011.

[21] B. Nardi, S. Whittaker, and H. Schwarz. It's Not What You Know, It's Who You Know: Work in the Information Age. *First Monday*, 5(5), May 2000.

[22] Spotlight. Apple inc. http://www.apple.com/macosx/what-is-macosx/spotlight.html.

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

[24] K. Voit, K. Andrews, and W. Slany. Why personal information management (pim) technologies are not widespread. *PIM Workshop 2009, Vancouver*, 2009.

[25] S. Whittaker, Q. Jones, B. Nardi, M. Creech, L. Terveen, E. Isaacs, and J. Hainsworth. Contactmap: Organizing communication in a social desktop. *ACM Trans. Comput.-Hum. Interact.*, 11:445–471, December 2004.

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

[27] XBee. http://www.digi.com, 2011.