

THE SHADOW: A PERSONAL EXPERIENCE CAPTURE SYSTEM

"Every bit of information we create is useful to someone."

James A. Landay, Mark Newman, Jason Hong

Computer Science Division
University of California at Berkeley
Berkeley, CA 94720-1776
+1 (510) 643-3043
landay@cs.berkeley.edu

ABSTRACT

We propose to build a “personal experience capture system,” which unobtrusively follows us wherever we go and knows everything that we do. This Shadow will be able to create a precise record of all of our actions and experiences. High level semantic information will also be captured. Individuals can use this data to assist in recalling information. By collectively networking the information, we can share our knowledge and experiences with one another. This paper outlines our ideas and the research problems that must be solved to achieve this vision.

1. INTRODUCTION

The system we propose is a “personal experience capture system”, which unobtrusively follows us wherever we go and knows everything that we do. This *Shadow* will make use of computing devices ubiquitously embedded in our environment to create a precise record of all of our actions and experiences. These records can include audio, video, physical location, and time data, as well as other formats, such as logs of all of our correspondence and documents we have edited. High level semantic information, such as where we have visited, whom we have met, what was said by whom when, what steps were performed in accomplishing a task, will also be captured. Individuals can use this data to assist in recalling information. Furthermore, by collectively networking the information, we can share our knowledge and experiences with one another.

A Shadow is a disembodied process that follows a particular user and is able to make use of whatever devices and network services it encounters. In the near term, location and identification of users will be accomplished by user-carried Personal Digital Assistants (PDAs) in a seamlessly connected environment. The concept is not tied to this means of identification. Image and voice recognition, when mature, could be sufficient to identify and locate a user within a smart environment.

The basic concept behind a Shadow can be traced to Vannevar Bush [1], who noted that “A record ... must be continuously extended, it must be stored, and above all it must be consulted.”

Gordon Bell has proposed a Guardian Angel that can “retrieve everything we hear, read, and see” [2]. Dan Olsen has also suggested a similar system that we would wear [3]. In effect, a Shadow would be what Don Norman calls a cognitive artifact, or a tool that aids the mind. Norman writes, “the technology of artifacts is essential for growth in human knowledge and mental capabilities” [4].

2. APPLICATIONS AND BENEFITS

2.1. Journal / Notebook / Manual

A Shadow will be extremely useful whenever we need to recall precise details of our past. While people are good at many things, recalling specific details with complete accuracy is not one of them. A Shadow can assist us by allowing us to review important details of our past experiences.

One application in this domain is the Journal. The *Journal* allows us to browse or search the data collected as a record of our personal experience. We can record and recall ideas, conversations, promises, and appointments by browsing the journal. The *Notebook* is a more active version of the journal, in which the data passively collected by the Shadow is augmented by data explicitly supplied by the user. The Notebook is specialized towards collecting and displaying information in the interest of facilitating document creation, not just information browsing and searching.

Another application in this domain is the Manual. A *Manual* assists us in remembering how we successfully accomplished something in the past, which will not only improve our performance but prevent us from wasting time. The Manual can be extremely useful in an educational environment. Although there are many things we have learned in our past, it is difficult to remember the specific steps to do something unless it is done on a regular basis. The Manual allows us to distill the essence of what we have learned, so that if it is necessary, it can be quickly recounted.

2.2. Personal Assistant

An agent could also use the data collected by our Shadow to learn our preferences and interests, and even act on our behalf. The Shadow Agent could negotiate with resources in the environment to ensure that our needs and desires are met. A simple example is a *Personal Assistant* that negotiates with a smart environment to maintain an optimal room temperature for a given user. The Agent knows the preferences of the user and tries to obtain the user's desired temperature. Meanwhile the environment attempts to balance the requests of multiple users' Agents and arrive at the best solution.

Another use of the Personal Assistant would be to suggest information relevant to the current situation. This information can be given to us directly, as in Remembrance Agents [5], or given to us by subtle cues in our environment, as in calm computing [6, 7].

2.3. Personal Manager

A Shadow will also be useful whenever we need metrics on some aspect of our behavior. Metrics allow us to predict future behavior from past behavior, aid us in planning for the future, and assist us in pinpointing the source of a problem. However, people are not very good at recording metrics because it takes enormous discipline to record metrics consistently and accurately, and because recording metrics usually distracts from the task at hand.

One useful application in this domain is the *Time Manager*. The Time Manager allows us to see how we spend our time at any granularity, whether it be hours, days, weeks, or years. This could be used to see if we are managing our time well. The Time Manager can also assist us in planning. If we can determine how long we spent on previous projects, the Time Manager can assist us in planning how long new, similar projects will take. One can imagine a whole suite of applications, such as a Money Manager, an Exercise Manager, and even a Food Manager. If combined with the Personal Assistant, the Personal Manager could analyze one's behavior and suggest improvements.

2.4. Group Shadow

A *Group Shadow* is an extension of the Shadow idea to a group of people. All of the applications aimed for a single person can also be used for a group. For example, a *Group Journal* would record all of the interesting events for a group. A *Group Manual* would enable the sharing of "how-to" information among group members by simply tracking the processes of the experts. A *Group Manager* could capture metrics on a group, assist in pinpointing problems the entire group encounters, as well as assist in planning group projects. Essentially, a Group Shadow could establish a Group Memory, so that nothing would ever be lost.

For example, suppose a well-experienced system administrator leaves a company and is replaced by an inexperienced one. By using the captured know-how of the previous system administrator, the new one should be able to do anything the

previous one could (albeit more slowly). While not a perfect replacement, the Group Shadow could help bring the new system administrator up to speed on how things are setup and how things are specifically done.

Group characteristics could also be observed from aggregated Shadow information. Information such as traffic patterns, space and resource usage, group preferences and interests could be derived from sets of Shadow-collected anonymous data.

The Bootstrap Institute, founded and directed by Doug Engelbart, uses the term "Collective IQ" to describe how quickly a group can "leverage its collective memory, perception, planning, reasoning, foresight, and experience into applicable knowledge" [8]. They note that a key factor in this Collective IQ is the quality and utility of the group's knowledge repository. A Group Shadow could greatly aid in this endeavor.

3. OBSTACLES AND APPROACHES

3.1. Infrastructure / Heterogeneous Devices

For the purposes of a Shadow, we assume seamless, constant connectivity for users in all environments. Furthermore, we assume that it is easy to locate and identify users. Work on these issues is already underway here at Berkeley in the Daedalus group [9, 10]. We do not wish to duplicate their work but to extend it in several ways.

First, even assuming seamless connectivity, we still need to develop a semantics with which to express capabilities of devices and resources so that these component resources can be composed into complex, intelligent environments. HP's JetSend protocol [11] is one step towards addressing this problem. As time progresses, however, the number of computing devices per person will increase. Adding new devices, or more abstractly, new resources to an environment thus needs to be a painless operation, i.e. there should be no special configuration necessary to incorporate a new resource into an existing environment.

Similarly, there should be no special configuration necessary to incorporate new interaction devices into an existing environment. User interfaces to Smart Spaces should be composable "on the fly," on heterogeneous client devices, from desktops to laptops to PDAs to smart phones to pagers to as yet unheard of interaction devices. Defining the description semantics that will allow arbitrary devices and resources to learn of each other's capabilities and "do the right thing" is a necessary step to make the existing Daedalus infrastructure "smart."

Lastly, our environment should be composed of both physical and virtual spaces for several reasons. First, it allows an environment to be a combination of both physical objects and virtual objects. An example of a virtual object could be software that does additional post-processing on captured information. Second, it allows us to create spaces of any granularity, from desk-sized to room-sized to building-sized spaces to city-sized spaces. Third, Smart Spaces can be composed together to form larger, virtual Smart Spaces. For example, all the divisions of a company can

have their own Smart Space, but all of these individual Smart Spaces can also be combined into a larger, virtual Smart Space. This virtual Smart Space could have all the state, behavior, and knowledge base of its components. One approach to creating a virtual space is the Jupiter System [12], a multimedia network place that supports virtual objects and virtual tools. WorldBoard [13] has a novel approach in augmenting physical space by allowing every square meter on the planet to be marked up with virtual data.

One challenge in creating a Shadow in such an environment lies in the continuous capture of interesting information. A Shadow must always be with us and must always be capturing useful information. The Factoid project [14] accomplishes this by placing all of the capture in a mobile device. However, a Shadow only requires a mobile component for when the user is outside of a Smart Space. Ideally, the Smart Space could capture all relevant data whether or not you have a mobile device.

We can create a hybrid system using both mobile devices and Smart Spaces. Since the mobile device is always with us, we can always capture certain types of information. In addition, when we enter a Smart Space, the capture capabilities of the mobile device are merged with that of the Smart Space. The mobile device can seamlessly act in conjunction with other devices to capture richer forms of information.

For example, if we are travelling and happen to be in a place that does not have any capture devices installed, then the information that will be captured is limited by the capabilities of our PDA. However, if we are in a classroom that has audio and video recording devices installed, these devices can join together with each of the students' PDAs to capture aural and visual information, any information the instructor sends to us, as well as any notes taken in class on the PDA.

3.2. Context Awareness

Another challenge in creating a Shadow is in context awareness. Certain kinds of information will be very easy for a Shadow to capture, such as position and distance. However, these kinds of information by themselves are neither interesting nor useful. Audio and video can also be captured, but besides being difficult to search through (addressed below), the computer will still need a way to capture the essence of what we are doing.

For example, while a Shadow can record the audio and video of us working on our car, we really want the Shadow to know we are working on our car, changing the oil, and adding brake fluid. If it does not know what we are doing, it will be difficult to search for when we last changed our oil. It will also be difficult to have the Shadow assist us or predict our actions.

Again, it will be relatively easy for the computer to capture certain kinds of context, such as the current time and who is nearby. Much more difficult will be for the computer to capture and infer interesting events, such as studying for a test or designing a new system. Figuring out the who, where, and when can be done with existing technology, but much more difficult (and much more useful) is figuring out the how, what, and why.

One way to start on this problem is to capture information that already has some context known to the computer, in this case, computer applications. We can easily capture three pieces of information: the name of the application, the name of the file(s) worked on, and the time spent on each file. This is a simple way of prototyping some of the basic functionality that will be needed later on.

Another way to approach this problem is to use pre-existing context, such as calendars and to do lists. For example, if a person has scheduled an event on her calendar, a computer could assume that all captured material could be indexed under that event.

A third approach, applicable to restricted domains, would involve machine learning. A person could create explicit categories, such as "recent" and "school", and explicitly place items under these categories. The computer could use pattern matching techniques to determine what similarities items placed in the same category have, and eventually be able to predict what categories an item should be placed under. This approach is used for email messages in Re:Agent [15].

3.3. Presentation / Visualization of Data

A Shadow will collect huge amounts of raw data in many different formats, including video, audio, text, and images. This data will need to be processed and presented in intelligent ways in order to be useful. The key will be to conceive of applications like those outlined above and develop presentation techniques appropriate for specific applications. These presentation and visualization tools should also be designed to exploit the smart environment, as described in Application Development below.

Using web search engines may be an initial way to start, since the web also contains massive amounts of information. One key difference is that while information on the web is extremely diverse and disparate, a person's captured information need not be. The information can be indexed by various contexts (such as people, place, or time), and can be presented to the user in familiar formats, such as calendars, photo albums, diaries, and to-do lists.

Eldridge, Lamming, and Flynn showed that a Video Diary [16] helped people remember more than a written diary did. They also discovered that people and objects were often the cues used for recall.

Lamming and Flynn have also created the Forget-me-not system, a memory prosthesis implemented on the ParcTab PDA [17]. The Forget-me-not was designed to help overcome everyday memory problems by exploiting episodic memory. People naturally organize events in memory by placing it in the current context, such as where the event happened, who was around, and so on. The Forget-me-not would capture some of these events and allow users to interactively filter the system to view past events.

Schank [18] hypothesizes that stories are an organized and compact form of information storage. Norman [4] also anecdotally notes that people often make decisions based on

stories, not hard facts. An intriguing line of research would be to see if certain mechanisms improve certain kinds of memory.

3.4. Application Development

In Smart Spaces, the “computing environment” is no longer an abstraction to describe the invisible structures inside the machine on your desk (or the network in your walls), but it is nearly synonymous with the actual environment. In such a world, how will applications be developed? User interface components, nowadays almost always presented on a single machine and a single display, can be distributed across multiple specialized devices. Data input and output functions can, and almost certainly will, be divided among several machines.

In the component-based infrastructure described at the beginning of this section, device capabilities that are known at runtime can be dynamically exploited by applications that may have been written with no knowledge of the conditions under which they would ultimately be run. What abstractions can we develop to support programming in such a world?

3.5. Agents

Creating a system that can analyze patterns and trends will be another challenge. Advances in data mining and pattern matching techniques will be needed in order to accomplish this. A related problem is predicting a user’s action, suggesting courses of action, and in some cases, executing simple actions. This problem is essentially a generalization of agent software.

SUMMARY

This paper has outlined a proposal for a “personal experience capture system.” The Shadow will be able to create a precise record of all of our actions and experiences. By collectively networking the information, we can share our knowledge and experiences with one another. In this paper we have outlined our ideas and the research problems that must be solved to achieve this vision.

REFERENCES

1. Bush, Vannevar. As we may think. <http://www.isg.sfu.ca/~duchier/misc/vbush/vbush.txt>
2. Bell, Gordon. The Body Electric. *CACM* February 1997, Vol. 40, Number 2.
3. Olsen, Dan. Interacting in Chaos. Invited talk given at Intelligent User Interfaces 1998, San Francisco.
4. Norman, Donald A. *Things That Make Us Smart*. Addison-Wesley. 1993.
5. Remembrance Agents. <http://rhodes.www.media.mit.edu/people/rhodes/RA/>
6. Weiser, Mark. “The Computer for the Twenty-First Century.” *Scientific American*, pp. 94-10, September

1991. <http://www.ubiq.com/hypertext/weiser/SciAmDraft3.html>
7. Weiser, Mark. “Does Ubiquitous Computing Need Interface Agents? No.” Invited talk given at MIT Media Lab Symposium on User Interface Agents, October 1992. <http://www.ubiq.com/hypertext/weiser/Agents.ps>
8. Bootstrap Institute. <http://www.bootstrap.org/vision.html>
9. Daedalus Project. <http://daedalus.cs.berkeley.edu/>
10. Hodes, T. D., and Katz, R. H. Composable Ad-hoc Location-based Services for Heterogeneous Mobile Clients. <http://daedalus.cs.berkeley.edu/publications/services-WINET.ps.gz>
11. JetSend Protocol. <http://www.jetsend.hp.com/>
12. Curtis, P., Dixon, M., Frederick, R., and Nichols, D. The Jupiter Audio/Video Architecture: Secure Multimedia in Network Places. *Proceedings of Multimedia 95*, ACM Press. San Francisco, CA.
13. Spohrer, Jim. WorldBoard. <http://www.almaden.ibm.com/almaden/npuc97/1997/spohrer.htm> <http://wtlinux.wisdomtools.com/wb/title.html>
14. Mayo, Bob. Factoid. <http://www.research.digital.com/wrl/projects/Factoid/index.html>
15. Boone, Gary. Re:Agent, An Intelligent Email Management Agent. <http://www.cc.gatech.edu/grads/b/Gary.N.Boone/reagent/reagent.html>
16. Eldridge, M., Lamming, M., and Flynn, M. Does a Video Diary Help Recall? <http://www.rsrc.xerox.com/publis/cam-trs/html/epc-1991-124.htm>
17. Lamming, M., and Flynn, M. “Forget-me-not” - Intimate Computing in Support of Human Memory <http://www.rsrc.xerox.com/publis/cam-trs/html/epc-1994-103.htm>
18. Schank, R. *Tell Me a Story : A New Look at Real and Artificial Memory*. Northwestern University Press. 1995.