Castaway: A Context-Aware Task Management System

Angela Kessell
Dept. of Psychology
Stanford University
akessell@stanford.edu

Christopher Chan
Dept. of Computer Science
Stanford University
cchan05@stanford.edu

CHI 2006, April 22–27, 2006, Montréal, Québec, Canada.
ACM 1-59593-298-4/06/0004.

Abstract
This paper describes the development of Castaway, a context-aware task management system. Specifically, we describe a three-week field study with thirty-five participants, the results of which illuminate the nature of people's recorded tasks. We further describe in detail iterations made to our task management interface, including a map-based view, and the insights gained that will inform future design and development.

Introduction
The increasing ability to both track people's movements and sense the environment combined with the growing ubiquity of mobile devices has lead to an exciting acceleration of research and development of context-aware computing. One potentially powerful context-aware application is the mobile management and receipt of personal tasks. Our vision of Castaway consists of three parts: 1) support for the fast and convenient input of tasks the instant they are conceived; 2) a lightweight, flexible tool to view and manage these tasks; and 3) a system for reminding users of their tasks at precisely the right place and/or time. Here we describe our progress in developing the second component. Although prior research has explored task management and the delivery of context-relevant information [1, 2, 3], the current work

distinguishes itself by exploring the potential of a novel task management interface, including an integrated map, and filling gaps in data on human practices and preferences for these domains.

Study 1: Diary Study

As this research began as fulfillment of a 10 week course requirement, time was limited. Therefore, we chose to begin developing a prototype system in parallel to collecting data characterizing personal tasks, rather than wait until data collection was completed. Our first priority was to assess the relative frequency of time and location relevant tasks, in order to determine whether and to what extent a location-based reminder system might be useful. We hypothesized that a significant proportion of tasks would be associated with particular geographic locations. To test this hypothesis, we carried out the following study.

Method

The participants in this study were 35 (11 female) undergraduate and graduate students at Stanford University. Participants were instructed to send in at least once per day all the action items they encountered, thought of, or already had recorded that they wanted to remember. Action items were defined as items such as the types of things one might put in a PDA or calendar, or on a "To Do" list or Post-It, etc.. Participants were free to submit their items via any combination of email, voicemail, text message, or paper diary. For each action item, participants were instructed to answer the following reminder basis question: *If a system existed that could actually remind you of that particular item, on what basis would you prefer to be reminded?* Participants could choose between time only, their location only, both time and/or location, or neither.

There were three conditions in this study. The one week "code during" group submitted their answer to the reminder basis question at the same time they submitted each action item. The one week "code after" group submitted their action items for one week and then after that week a list of their action items was returned to them. *Only at that time* did they learn about and have to answer the reminder basis question for each item. These two groups allowed us to test for any biases induced by the reminder basis question. The two-week "code after" group submitted their action items over a period of two weeks and were asked to answer the reminder basis question for each item at the conclusion of the two weeks. This group allowed us to track frequency of item submission over time.

Each action item submitted was coded for whether it contained an explicit reference to a time (e.g. Buy eggs *tomorrow*) or a location (e.g. Go to *Sam's house*).

Results and Discussion

Across all 35 participants, a total of 1,748 tasks were collected. The mean number of items submitted per day per participant was 7.0 (*SD* = 3.1). The mean proportion of time only based reminders was .48, of location only based reminders was .25, of both time and location based reminders was .08, and of no reminder was .19, suggesting that items requiring location-based reminders would indeed constitute a significant proportion (.33) of all items.

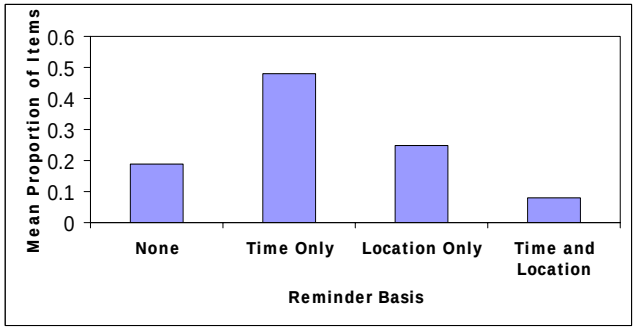

Figure 1. The mean proportion of items that required a reminder based on time only, location only, time and location, or no reminder at all.

Whether an item contained an explicit reference to a location was uncorrelated with whether the user indicated that a location-based reminder would be appropriate for that item (r = .04, $p < .10$). That is, although the correlation was marginally significant, the size of the correlation (less than .1) is insubstantial by conventional standards. This suggests it may be difficult to automatically extract the relevant location from the item, as well as to infer whether a location-based reminder is appropriate. In contrast, containing an explicit reference to a time was significantly correlated with requiring a time-based reminder (r = .20, $p < .001$), suggesting that it may be easier to automatically extract the relevant time from an item.

We expected the "code during" group to be more attentive to, and more likely to report, items requiring location-based reminders than the "code after" groups. However, the mean proportion of items explicitly referring to a location was the same for both the "code after" groups (M = .15, SD = 0.07) and the "code during" group (M = .15, SD = 0.08), suggesting that being aware of the potential for a location-based reminder did not increase their propensity to name a particular place in their item. Surprisingly, the mean proportion of items requiring location-based reminders was actually greater for the "code after" groups (M = .36, SD = 0.20) than for the "code during" group (M = .21, SD = 0.12). Statistical comparisons were not performed on these data due to very different sample sizes for the "code after" (N = 27) and "code during" (N = 8) groups.

Finally, looking at just the two week group, there was no difference between the first and second weeks in mean number of items submitted (*paired-t*(17) = 0.68, $p > .05$), suggesting that the trends found here may generalize to behavior over a longer term.

Prototype 1

The prototype we built focuses on the second component of Castaway: viewing, managing and reasoning about one's tasks.

Preliminary analysis of initial data from the diary study indicated that at least three organizational formats could be useful: time, location, and list. Therefore, we designed the interface to offer the user a choice to switch between multiple types of views, as well as the option to create custom sub-lists to be viewed in a particular format. We describe the three different types of views in more detail below.

List view

The list view displays items of all different types in a single list. This view supports real-time search, drag and drop arranging, and the creation of new tasks.

Attributes of each item, including reminder cue and method of delivery, can be easily edited.

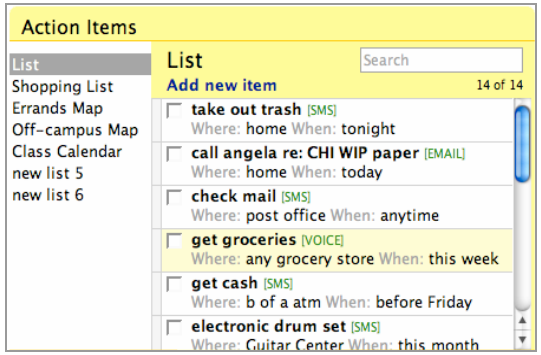

Figure 2. List view of Castaway. The green text in brackets indicates the way in which the reminder will be delivered.

Map view

The map view, implemented using the Google Maps API, offers a location-based visualization of one's tasks. Users click on interactive markers, which in turn reveal the tasks associated with a particular area. The objective of the map-based view is to facilitate spatial reasoning about one's tasks and increase efficiency in planning for and carrying out tasks.

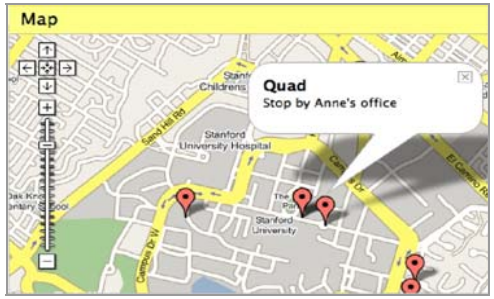

Figure 3. Map view of Castaway.

Calendar view

The calendar view is a more traditional, time-based visualization of one's tasks. Although this view was not fully implemented at the time of user testing, the feature was discussed with participants.

Custom list/map/calendar views

To maximize flexibility in task organization and management, users can drag and drop their tasks from the list view into custom lists, maps, or calendars. These custom views are automatically saved for future viewing.

Several ideas here were inspired by Apple Computer's iTunes Music Player, where tasks are somewhat analogous to songs, and custom lists, maps, and calendars to playlists.

Study 2: User Testing

In order to gauge participant reaction to the concept of Castaway and to inform the second iteration of our interface, we conducted the following user evaluation study.

Method

The participants in this study were 6 (4 female) undergraduate students at Stanford. Participants filled out a preliminary questionnaire assessing their current preferences and practices for personal information management. Next they systematically explored the following aspects of the interface: list view and manipulation, map view, calendar view, and creation of custom lists, calendars, and maps. For each component, participants were asked to assess the intuitiveness and ease of the interaction, as well as to suggest how they might change the interface. Finally,

participants filled out a follow-up questionnaire addressing what they liked and disliked about the interface and the general system concept, the likelihood of their using such a system, and what types of tasks they might manage through such a system.

Results and Discussion

Participants used a variety of tools for managing their personal tasks, ranging from Outlook, iCal, and PDAs to Post it's, and paper agendas and lists. All participants carried their cell phone with them at least 95% of the time, suggesting that a cell phone would be an effective means for delivering reminders. On average, participants rated the usefulness of the map view as 4.5 on a 7 point scale ranging from not at all useful to highly useful. Importantly, they rated their likelihood of using a system that gave location-based reminders at 5.3 on a 7 point scale ranging from highly improbable to highly likely.

Overall, the feedback received from participants was rich and informative. Most interesting to us were comments upon the map view. Some participants questioned whether a precise positioning of tasks on a map would be useful for familiar areas. Other participants commented that deciding what task locations would be "nearby" or "on the way" was dependent upon several factors, such as mode of transportation (*e.g.*, walking, biking, driving) or whether it was necessary to take the freeway. Other comments addressed cosmetic or low-level changes to the interaction, which were of less concern to us at this stage.

Prototype 2

In light of the finding that participants were satisfied with our basic interaction design (direct manipulation of tasks, customizable lists, maps, and calendars, etc.), the second prototype remains unchanged in these core aspects. We focused instead on addressing their concerns on the use of the map view.

In response to their feedback, we added an interactive slider to the map that allows for automatic clustering of tasks based on location. By default, the slider is positioned at the leftmost extreme and the map view displays tasks as before, with every marker corresponding to a single location in which a task must be completed. At the rightmost extreme, all markers for all tasks merge into a single marker. Between these two extremes, the slider position determines the tolerance level at which markers will be clustered. As the slider moves to the right, markers are progressively merged so that each marker now represents multiple locations within a certain radius, with the marker displayed in the center of these locations. When one clicks on the marker, all tasks associated with the marker are displayed. For example, if a list of location-based tasks spans multiple cities, positioning the slider midway might cluster the tasks into a single marker at every city.

This clustering feature offers several advantages over the standard, non-clustered view and addresses users' concerns in the following ways:

- Clustering items from separate lists *(e.g., Errands for upcoming trip* and *Errands for Lucy's birthday party)* allows users to quickly see which tasks may be carried out together – an association they might

not otherwise have made - without unnecessary visualization of precise geographic location.

- By controlling task clustering, users determine, on a case-by-case basis, how to define “nearby” and “on the way”.
- By adjusting clustering level to accommodate the time available to carry out tasks, users can plan task execution more effectively.
- Flexible visualization of tasks aids users in determining the optimal order in which to execute their tasks.

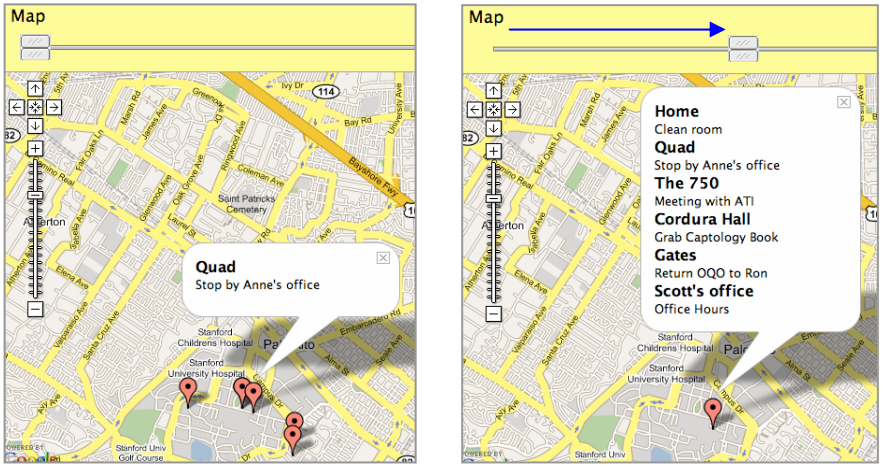

Figure 4. Clustering tasks in the map view. By default, markers correspond to single locations (above left). By adjusting the slider at the top of the screen in the direction indicated by the arrow (added for illustrative purposes), a single marker is positioned at the center of several nearby task-associated locations (above right).

Work in Progress and Future work

Much work remains to be carried out in the development of Castaway. Up to this point, we have focused primarily on task management. Presently, we are beginning to test both the usability and usefulness of our clustered map visualization. In addition, we are exploring other types of relevant contexts, such as “at a computer,” or “next time I see Bob.” In the future, we plan to focus on the first and third components of Castaway: flexible, convenient task input and spatio-temporally appropriate task notification.

Acknowledgements

We would like to thank Scott Klemmer and Robin Jeffries for valuable feedback on this project, and Doantam Phan and Ron Yeh for their substantial contributions to both project formulation and data collection.

References

[1] Anhalt, J. A., Smailagic, D. P., Siewiorek, F., Gemperle, D., Salber, S., Weber, J. Beck, and J. Jennings. Toward context-aware computing: Experiences and lessons. *IEE Intelligent System*, 3, 16 (2001), 38–46.

[2] Bellotti, V., Dalal, B., Good, N., Flynn, P., Bobrow, D. G., Ducheneaut, N. What a to-do: Studies of task management towards the design of a personal task list manager. Proc. CHI 2004, ACM Press. (2004), 735-742.

[3] Rhodes, B. J. The wearable remembrance agent: A system for augmented memory. *Personal Technologies: Special Issue on Wearable Computing.* 1, 4 (1997), 218–224.