

What a To-Do: Studies of Task Management Towards the Design of a Personal Task List Manager

Victoria Bellotti, Brinda Dalal, Nathaniel Good[†], Peter Flynn[‡], Daniel G. Bobrow
and Nicolas Ducheneaut

Palo Alto Research Center
3333 Coyote Hill Road,
Palo Alto, CA 94110
bellotti@parc.com
650 812 4000

[†]School of Information
Management & Systems
University of California, Berkeley,
CA 94720
ngood@sims.berkeley.edu

[‡]Cornell University
Ithaca, NY 14853
psf2@cornell.edu

ABSTRACT

This paper reports on the results of studies of task management to support the design of a task list manager. We examined the media used to record and organize to-dos and tracked how tasks are completed over time. Our work shows that, contrary to popular wisdom, people are not poor at prioritizing. Rather, they have well-honed strategies for tackling particular task management challenges. By illustrating what factors influence task completion and how representations function to support task management, we hope to provide a strong foundation for the design of a personal to-do list manager. We also present some preliminary efforts in this direction.

Categories & Subject Descriptors: H.4.1 [Information Systems Applications]: Office Automation – *time management*; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems – *evaluation/methodology*; **Keywords:** ethnography.

General Terms: Design, Human Factors.

Keywords: Task management, to-dos, ethnography.

INTRODUCTION

A major initiative being considered at DARPA is directed towards developing cognitive systems that can support busy professionals in government or military roles in managing and even performing office and military tasks. These ‘cognitive assistants’ will be capable of reasoning and learning and will be aware of and able to explain their own behavior as well as accepting direction from users [8]. One of the possible embodiments of this type of system is a task list manager system (TLM) that could help users manage and execute their to-dos. Such a system would:

- Capture the person’s daily tasks.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2004, April 24–29, 2004, Vienna, Austria.

Copyright 2004 ACM 1-58113-702-8/04/0004...\$5.00.

- Plan and execute simple actions.
- Prioritize, manage, and reason about tasks.
- Learn to improve by being told, observing the user, asking questions, and reflection.
- Record notes, action items and ideas.
- Answer questions and offer advice and assist in planning and problem solving.

As part of this initiative, an initial effort at our laboratory was undertaken to understand natural practices of task management and types and quantities of tasks taken on by people similar to prospective users of DARPA technology, namely busy professionals and managers. In particular we sought to discover what kinds of task management demands might be supported by a TLM.

Prior Work

There are a number of best sellers and tools available on how to organize one’s time and prioritize work [e.g., 1, 7]. The market for these resources seems to thrive on the fact that many people worry about whether they are prioritizing and meeting their many obligations effectively [4].

Personal information management (PIM; organization, note taking, reminding and calendaring) has also been examined in the HCI literature, but mainly focusing on the problem of organizing documents, files and notes for the purpose of reminding and efficient retrieval, rather than task management [2, 6, 9, 12, 13, 14]. There are a number of studies on how people use calendars in practice (for far more than just event scheduling), [5, 16, 17]. But this literature focuses only on a single resource that mainly serves time management needs. And the many readings available on cognition, planning and task execution in the classical psychological literature [10, 15] have little to say about task management and planning in work practice.

Distributed cognition analyses [11, 18] comes closest to the kind of analysis we seek here showing that external resources are critical in performing complex tasks, but the literature does not look at how external representations function to help their creators assess the current state, extent and priority of many tasks to be completed. The focus has been on resources supporting task execution rather than articulation work (the

work required to plan and organize work) [19]. And task management from an articulation standpoint goes beyond simply organizing physical and virtual collections and putting events in a calendar.

Recent work has begun to tackle the challenge of general task management at it plays out in email [3, 20, 23]. Given that many PC users are overwhelmed by the number of tasks handled through email, this trend is not surprising. This work suggests that any successful task management resource must be closely integrated with email functions. But since email is unlikely to be the entire story, we felt we needed to look at task management practices more broadly.

In the remainder of this paper we briefly discuss results of a “snapshot” pilot study and then focus on a longer-term study of task management representations and practices. We also discuss our embodiment of some of our fieldwork-driven ideas in a preliminary design. In our reporting, all people, project and organization names have been changed or obscured to protect the privacy of our participants.

A SNAPSHOT STUDY OF TO-DO'S

Prior to beginning an in-depth study of task management, we conducted a pilot study in which 3 administrative staff, 4 researchers and one manager were asked to show all the resources where they kept to-dos and count all of the active to-dos they were currently tracking in each one. Table 1 shows average numbers of to-dos and resources used to represent them, based on 595 counted by the 8 participants. We also captured explanations of each resource and 79 detailed explanations of examples of to-dos, taking one or two from each resource each person showed us.

Per Person	Ave	Median	%
No. of resources used to represent to-dos	11.25	11.5	
No. of active to-dos	74.4	65	
In email	26.6	20	35.8
In online calendar	8.6	9	11.6
In paper list or paper notepad	7.2	0	9.7
In online folders	3.5	1.5	4.7
In online special purpose to-do list	3	0	4.0
In PDA calendar and list combined	2	0	2.7
In daytimer/bound notebook/planner	1.6	0	2.2

Table 1. Number of to-do resources and active to-dos per person. Online resources are shaded.

Some key properties of effective to-dos quickly emerged, a selection of which are discussed in the following:

To-dos are *made expending minimal effort*, so most of them do not describe the task, they typically are only elaborated enough to provide a salient cue. For example one to-do was some text on a pad of paper; ‘Joe the attorney.’ The explanation was, “A reminder to send him mail. I think I was supposed to ask him about this [...] lawsuit. I can't remember.” Interestingly, *to-do text is often not grammatical* as in, “Send Mother's Day” or even “Beth blah blah”. The cue is so minimal that it is only effective for a limited period




Figure 1. To-dos showing various levels detail from telling someone (Name) to take an SGBS course (bottom right) to preparing a presentation (Subject) presentation (top left).

of time while the task stays in memory. In some cases an item such as a book or a pile serves as enough of a cue to recall the task, without creating a note.

Only a minority of to-do reminders appear in lists. We found only 14% of the to-dos we counted were in a list.

To-dos are used in multiple ways. Sometimes they are part of a list that provides a sense of the amount of work to do. Sometimes they are *resources* supporting consultation, linking to work objects, or are work objects themselves, displaying state as well as to-do-ness.

Many to-dos are prompts placed *in-the-way in anticipation of a routine practice* that will occur at the right moment for the to-do to be discovered. For example, “When I go to grab my bag to go home, I'll go, ‘I must take that [object next to bag] home.’” Emails left in the inbox in particular serve this prompting function during email-centered work.

To-dos may be *represented at any level of abstraction or detail* (see Figure 1). We saw one to-do (not shown) that was detailed enough to support preparation of some slides, but other “give presentation” to-dos only referred to the subject of the presentation as shown in Figure 1 (top right).

To-dos don't all get done. People procrastinate about some, and deem some of lower importance. Two participants mentioned deliberately keeping low priority to-dos in an electronic form that would be lost if an application or machine crashed. An effective way of reducing the task list!

Implications for Task List Manager Design

There were some clear design implications from the above findings. First, of our participants’ 70 or so to-dos, about half are already online, even without the incentive of smart to-do management support. The main resources are email and the electronic calendar, but these have weaknesses [3, 20, 5, 17]. Thus, 50% of what is going on may be tractable to improvement with a TLM if it can capture this activity, and possibly more if users are motivated to move more to-dos online because of the benefits of system support. Many challenges are apparent however. A TLM must offer:

- *Diverse ways to view and manipulate* to-dos to emulate advantages of existing resources, going beyond lists.
- *The in-the-way property*, e.g., by becoming the habitual place for routine activities where reminders might be encountered.
- *Instantly on*, to support quick and easy input and clear

visualization. PDA's are often abandoned due to slow laborious input and attenuated output [6].

- *No formal task description*, categorization or decomposition required from users, and any level of abstraction must be allowed for atomic task entries.
- *A mechanism for handling stale to-dos* of low importance that are diminishingly likely to ever get done but have not been explicitly deleted.

LONGER-TERM STUDY OF TASK MANAGEMENT

The snapshot study, while useful, left many important questions for a TLM unanswered, for example:

- What help do people need with prioritization?
- What factors drive people to prioritize and complete tasks such that a TLM can propose appropriate action?
- What is the lifespan of a useful to-do?
- What kinds of task management resources are appropriate for different challenges?

In order to address these questions, a longer-term study was undertaken to capture and track a large number of to-dos. Due to the heavy time-commitments required, only 7 participants were engaged. Participants were specifically selected for having highly multitasking and diverse work regimes but none were using any task-management techniques such as the Franklin Covey system:

- M1: Manager of between 5 and 9 research staff in our organization. Reviews intellectual property (IP), plans and conducts research and obtains external funding.
- M2: Manages 15 to 20 researchers in our organization. Tracks IP and does business development.
- Prof: Professor and co-director of a research laboratory in a university and manager of 5 to 20 people. Conducts and manages research, and obtains funding.
- SPM: Senior product manager signing up and managing 200 tour operators who sell their products through his online company. Works in a small office in the USA while his head office is in London, UK.
- DDM: Director of Development and Marketing for a charity. Writes grant proposals, liaises with donors and supervises 3 to 4 staff members.
- DM: District Manager of ten stores in a chain of retail food and beverage stores. Visits stores, supervises store managers (20 people), tracks and develops business.
- SAM: Sales account manager in a large office-supplies retail and wholesale company. Has approximately 300 ongoing and prospective accounts.

The data collection method (executed at all of the participants' places of work) was as follows:

- Preliminary 2-hour background information interview.
- Four 1-hour, weekly task-tracking interviews (referred to as TT1, TT2, TT3 and TT4 respectively) in which (usually) 10 to-dos that might be done by next week were elicited, discussed and rated for importance and urgency on a scales of 1 to 5. All to-dos from each week were followed up at the next interview. Participants were

prompted not to focus on a single to-do resource or only important and urgent to-dos.

- One day of shadowing in which the participant was observed doing normal daily tasks.
- Final 1-hour interview to answer remaining questions and follow up on the final disposition of all to-dos.
- All sessions were video taped and transcribed.

All 287 to-dos were coded with a value of 'yes' or 'no' for 37 codes such as 'done,' 'on-hold,' 'common,' 'discretionary.' Cross-coder reliability was obtained for a subset of 50 tasks (50x37=1850 codings) with agreement of 92%. Correlations between coded factors were obtained with both Pearson's R and Spearman Correlation tests.

Task Tracking Findings

A minority, 104, of our tracked 287 tasks represented 45 tasks that were carried over one or more weeks. Thus we have 228 unique tasks.

We were unable to observe the tasks we were tracking but we got estimates for completion times ranging from 30 seconds to 5 days with 63% being between 10 minutes and 4 hours. Recollections were the same as the predictions in only 25% of cases. In 40% of cases, recalled times were shorter, averaging 77% of the duration of predicted times. In 35% of cases they were longer, averaging 130% of the duration of the predicted time. This could be because tasks took more or less time than planned, and/or because people are poor at estimating and recalling time required to do a task. In either case (and this was confirmed informally in discussion with our participant) predictions of task time are often inaccurate (sometimes grossly so). Thus TLM users should not be relied upon for precise time planning.

Medium of reminder	No.	% of R	I	U	%Done in 1 wk	%Dealt by end
Email	88	34.8	3.2	2.9	68	82
Formal paper to-do	61	24.1	4.3	3.2	52	77
Print out(s)	28	11.1	3.1	2.7	64	79
E-Calendar Entrv	28	11.1	3.3	3.2	61	82
Index card	17	6.7	3.6	3.4	94	94
Paper sheet/bad note	13	5.1	2.9	2.8	69	77
Object: book/bag/etc.	9	3.6	2.7	2.7	56	56
Note in notebook	4	1.6	3.3	2.3	100	100
Vertical folders	3	1.2	3.2	3.7	100	100
Voicemail	1	0.4	4.5	4.5	100	100
Open window on PC	1	0.4	4.0	3.0	100	100
In head (no remindr)	34	0	3.8	3.3	88	85
Totals or averages	287	100	3.5	3.0	68	81

Table 2. Media used to record to-dos, their prevalence as a percentage of all to-dos with reminders (R: 287-34=253), average importance and urgency of to-dos per medium and percentage of each type done in a week or by end of study.

Table 2 summarizes how the 287 to-dos were represented (in fact 34 had no reminder representation) and shows that, on average, participants did not offer only important (I) and

urgent (U) tasks to track. Urgency was not found to be significantly correlated with any reminder medium.

The use of to-do lists was higher in this longer study (24%) than in Table 1 (9% + 4.7%). Being on a list was positively correlated ($p < 0.001$) with importance and negatively correlated with getting done in a week ($p < 0.001$). Our data thus confirms the idea that more important to-dos get onto lists as they risk being overlooked while being delayed.

To-dos as email, print-outs, notes on paper or objects were less important on average ($p < 0.05$). Index cards were positively correlated with getting done in a week, (but these were used by only one person and so the 'done' factor is confounded with 'participant'). And objects were negatively correlated with getting done in a week ($P < 0.05$) suggesting these object reminders often have low priority.

Having no reminder at all was significantly *positively* correlated with completion in a week ($p < 0.01$). However, it is not correlated with importance, so it is not criticality that is making these tasks easier to remember. We did get some verbal reports that routine tasks tended not to be recorded as to-dos because they are easily remembered habits.

By the end of the study, though, we found that there was no statistical relationship between reminder medium and task completion, except for object to-dos, which were significantly less likely to be dealt with in the end ($p < 0.05$).

Table 3 shows further significant ($p < 0.05$; less than 5% chance of this data occurring randomly) correlates of task completion within a week (see Table 3).

Factor	Done if high/yes	Done if low/no	Sig p level
Urgency (high/low)	93%	44%	0.001
Meeting (yes/no)	87%	66%	0.05
Importance (high/low)	81%	42%	0.00
Someone expecting (yes/no)	77%	53%	0.00
Other(s) involved; not mtg (yes/no)	82%	62%	0.00
Non-discretionary (yes/no)	79%	63%	0.02

Table 3. Factors related to task completion within a week.

Highly important and urgent tasks (rated 5 on a 5 point scale) were twice as likely to get done as those rated low importance or urgency (rated 1; note that these factors were scalar, not binary like the rest). However, having a deadline was not significant suggesting that people do *not* rate tasks as important or urgent just because they have a time limit.

Tasks that were meetings or simply involved other people were both more likely to get done. Someone else expecting the task to be done was also highly significant. Whether or not a task was self assigned or assigned by someone else was not a factor, but discretionary tasks were less likely to be completed than non-discretionary.

One of our most surprising findings is that our participants seem to be prioritizing very competently. Only 3% of tasks

were dropped with no good reason. Since this finding flies in the face of popular literature on task management (which argues that people often fail to do important tasks due to poor prioritization) we will examine it in more detail.

Of all 287 tasks tracked across our participants:

- 68% of all tasks were done in a week and 81% were dealt with (as far as possible) by the final interview.
- By the final interview, 79% of TT1, 81% of TT2, 83% of TT3 and 80% of TT4 tasks were completed. Note that the final interview was usually two weeks after TT4 explaining the higher completion rate than 68%.

The fact that there is no difference between older and newer tasks being completed suggests that tasks planned for the weeks ahead that are not completed after two weeks are unlikely ever to get done. Qualitative data showed that to-dos were sometimes held up by unplanned overload. For example DDM admitted that she did not complete most of her to-dos one week due to an unexpected and onerous task with a tight deadline. However, all but one of these tasks were completed by the following week.

A further 16% of tasks were not done by the final interview, with this being a satisfactory outcome. In these cases, participants gave us a very good reason for not completing these tasks, for example (paraphrasing for brevity):

- I don't *have* to do this for several months (M2)
- We don't qualify for this funding cycle (DDM)
- The customer wasn't interested in the product (SAM)

The 3% of tasks dropped *without* good reason were always either possible to do at some later point or were minor slip-ups. We have no reason to suspect our participants were dishonest about failure, as they showed no hesitation when admitting to it (usually smiling or laughing).

So a surprising conclusion is that *the problem with task management is not failure to prioritize well*. We would argue rather that it is the effort that must go into *making sure that the important tasks get done*, even if the unexpected occurs, that is the real challenge. In previous research we found people express great concern about this and frequently switch between PIM resources to try to improve efficiency [4]. So our next section discusses how professionals manage to be so competent at prioritization. By learning from their success a TLM can emulate and improve on the resources and mechanisms they use.

We conclude this section, with the caveat that an 81% to-do completion rate might be higher than the true completion rate. Inevitably, participants chose the tasks that we tracked and these may have been ones that were less likely to be overlooked. However, as mentioned, to minimize this effect, we encouraged participants to include many less obvious tasks (hidden in piles, the email in-box etc).

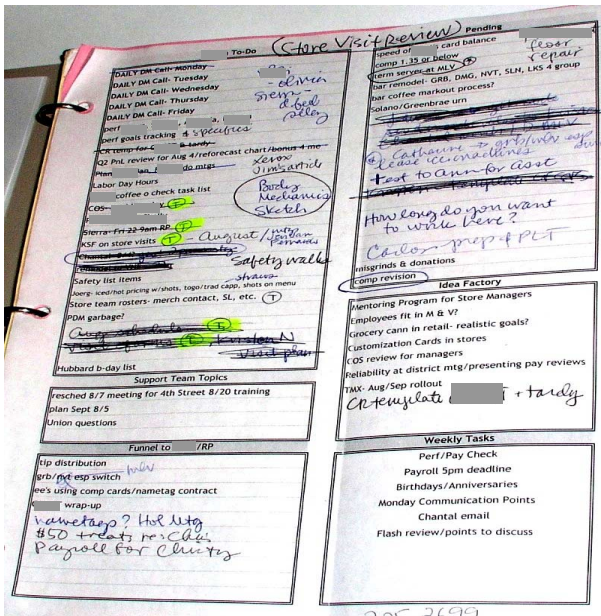


Figure 2. DM’s to-do list: A systematic task vista.

Task Management Strategies

A number of task management strategies emerged that seemed, not only to reflect personal preferences, but often to be closely tied with the demands of the participants’ particular job pressures. We will discuss some of these in the following sections (with the participants who used these strategies shown in brackets):

Task Vistas (DDM, DM, M1)

Task vistas are fairly comprehensive lists used for planning, ensuring that nothing (that could be forgotten if it were not listed here) falls through the cracks. They reflect a desire to be able to see all to-dos together on one page (all 3 listers in this study used paper), hence our term ‘task vista.’ As DM put it “I’ll shrink the font and change the margins to keep everything all on one page. [...] It just makes it more digestible for me I think and I can see the connections....”

DDM and DM produced weekly updated task vistas with systematic categories of tasks to assist in planning. Many items reappeared week after week. DM’s list is particularly interesting as she created many categories (see Figure 2). For example she keeps things to talk to her manager about in one place “Funnel to [blacked] /RP”, (bottom left) and things for her support team in another (lower left). These persistent categories reflect and support her routines and also break the many disparate tasks that are a feature of her job into more manageable chunks that she can plan around.

M1 showed us a different kind of ad hoc, occasional task vista that he created when he felt overwhelmed, and organized around projects, unlike the systematic lists. M1’s lists were not updated, but simply replaced with a new list with different structure when needed. He was managing and/or participating in about 10 complex projects and proposals during our study.

Informal Priority Lists (Prof, SAM)

Informal priority lists are selective task views that ensure near-term execution of priority actions. These were not organized and were jotted on small scraps of paper. Both Prof and SAM created them at home for the coming work day, daily in SAM’s case, and most days in Prof’s case. The small size of the paper seems to be important for Prof as she carries the list in a pocket to work with her.

State Tracking Resources (SPM SAM)

State tracking resources are needed when many similar (and confusable) threads are on the go at once. They show what actions were performed when and support looking back in order to infer what must be done instead of explicitly outlining what must be done. SPM and SAM both need to keep track of a great many simultaneous ongoing threads of activity. SPM has 200 or so prospective tour operators and SAM has 300 or so leads and accounts. Since SPM is not highly mobile, he uses a whiteboard with all of the accounts listed with the last action and date usually recorded. SAM needs a mobile solution as he is often out meeting clients and prospective clients face-to-face, so he staples a business card to an index card and write dates and notes on status and other useful information on that.

Time Management (M1, M2, Prof, DDM, DM)

Time management resources are needed when time becomes a scarce or inflexible commodity. The five participants who seemed to pay most attention to timing were, unsurprisingly the ones who were most subject to deadline pressures (accounting for 93% of the deadline task codes). We saw three strategies for time management:

- The first is to make sure that things get noticed later, when no time is available now. This is done by placing a to-do ‘event’ in the calendar at some future date.
- The second is to carefully preserve time for critical tasks in a busy schedule where appointments could fill up the days. Thus time would be blocked out in the calendar for things that were not events. As Prof put it “...it’s a time slot that I wanted to protect.”
- The third, used by DM, is to categorize some tasks on her list into either under 15 minutes or 30 minutes to 1 hour or more. She had around a dozen in the former and half that number in the latter category, which she stated made her task list seem far less daunting.

Value Extension (M1, Prof, SAM)

Value extension is the practice of making one task serve multiple goals and further, to choose tasks to maximize this effect. As M1 and Prof put it:

M1: “I spent some time talking to the people who are going to commercialize it [software]. [...] Thinking about how we could actually use this for what we want to do. What can I do to make it easy for those [interns] to try and do things with that [software]. [...] what are the small collaterals to make it work in these different contexts.”

Prof: "... we're actually leveraging this work across a number of projects."

When shadowing SAM, we also noticed him stopping by offices opportunistically on his way to clients' offices to see if he could drum up new business; having two clients close together obviously saves on travel time.

Key Phenomena

Among our findings a number of phenomena, that are not task management resources, are nonetheless key in determining the outcomes of critical tasks:

Social Relationships: The Network Effect

People are highly sensitive to who is important to them in their network when assigned tasks or getting requests:

SPM: "There are some people [...] there [in London] that work very, very hard for me, and if I have a choice, if I have time to answer only one email, I'm going to answer that person's email first because I know that they are going to act on [it...], right away. They'll be appreciative."

Another phenomenon is that when one person fails to complete a task that others are waiting on, the network compensates (a reminder or a demand materializes). And when a person consistently fails, their colleagues develop compensation strategies:

Prof: "[...] one thing I try really hard to do is get things off my plate, so I ask for confirmation from unreliable resources, 'I transferred this to-do' to you. Can you confirm that you have it?' And then, if I don't trust them, I ask them to get back to me on date x and date y."

Thus, the implications of failure to complete a task are nuanced since they depend on the nature of relationships between people. And our participants often mentioned that tasks that were important to other people were not necessarily important to them. Whether they completed the task or not depended on the relationship between them.

Working Away from the Desk

Among our participants, only SPM spends nearly all day in his office. With the others, we saw various instances where task management away from the desk was crucial. Many tasks were created away from the desk through extensive interaction in meetings and customer visits. For example, during our shadowing, DDM and her marketing manager created many actions for each of them in one meeting.

Rhythms and Routines

As with Tolmie et al. in their study of domestic life [21], we noticed tasks had temporal and procedural regularities. One kind of regularity is that a certain kind of task is likely to take place at a certain time—of day, day of the week, time of year, etc. For example, everyone begins the day by dealing with email and/or voicemail. And DM generally visits particular stores on the same day each week. Another kind of regularity is that some tasks follow a pattern, which has temporal and

action-based consistency. For example: SAM makes a cold call; if the client is interested, he follows up with a visit and product information; the client (hopefully) opens an account; SAM waits for some period and then goes online to inspect activity in the account to see if a follow up is needed to uncover any problems.

Implications for Task List Manager Design

Our findings have a number of implications for a TLM:

A TLM should support the viewing of entire task vistas, but also allow different perspectives for different kinds of planning. By this we mean a TLM must emulate the properties of the different kinds of to-do lists we have observed. Such views must be able to be sort and filter activity to show a day or, week or project scope. It should allow *top priority items to be made apparent* or items with similar properties to be viewed together, e.g., items that can be executed quickly, or those to be completed with a particular person. As noted, DM has special categories for both these types of items on her to-do list.

Task histories and state should be captured for the purpose of being able to determine information such as, when the last action occurred, what the current state is, who is responsible for the next action and so on. This capacity would support users who, like SPM and SAM, have many tasks ongoing at once.

Time constraints should be captured such that decisions about workload and scheduling can be made with a clear idea of how the overall picture will be affected. Combining these with *knowledge about rhythms and routines* could allow for visualizations of truly free time, as opposed to unscheduled time and support time and interval sensitive re-prioritizations of a user's tasks. However, a TLM should *not rely on users making accurate predictions about the time required to complete their tasks*. Indeed it should *not require tasks to be completed at all*. However, if a TLM is to succeed at representing time commitments, then users must be encouraged to enter many tasks that they might not normally represent as to-dos, simply to capture their time commitments. This suggests that a TLM must *offer more benefits than just reminding for entering tasks*, such as using list entries as organizational or time management resources or to launch the related tasks themselves.

The properties of tasks must be modeled in such a way that it is possible to *practice value extension more explicitly*. For example, making the location of a task explicit (as in the Llamagraphics task management tool Life Balance™) allows multiple tasks to be planned retrieved and executed at the appropriate location, requiring only one excursion.

Social relations should be captured and modeled, perhaps using information such as email responsiveness images (time taken to respond to specific people [22]), which seem predictive of the importance of social relationships. Doing so can help prioritize tasks such that those requested by known valued colleagues are pushed to the top of the stack.

Finally, a TLM must *support the capture of notes and task lists away from the desk* in order to be effective.

TOWARDS DESIGN FOR A TLM

During the period when we were intensively engaged with fieldwork and developing our ideas about task management, we felt it was also important to keep ourselves grounded in solving design problems relevant to support for task management. So we designed and prototyped an early component of a TLM system that addresses some (but not all) of our findings. This is described in the following section. To assist the reader in linking its design features to our ethnographic findings, we have italicized the properties that our research indicates are likely to be valuable.

TaskVista

TaskVista (see Figure 3) is a *lightweight resource* for collecting and listing to-dos and conveniently *launching tasks* from them. Unlike existing tools such as the Microsoft Outlook™ Task List, it is compact and sits on the desktop (or PDA) and is *in-the-way but not obtrusive when users reach transition points between tasks* (e.g., when they switch applications). It is a *comprehensive to-do list* that easily handles a *realistic number of active to-dos*. *Old to-dos do not need to be deleted* but are filtered out of sight when they become defunct or are done, to avoid clutter.

Users can *easily create a new to-do* by typing one in, or, dragging an item (e.g. a file or email) into the list. The (editable) title defaults to the subject or title of a dragged-in item to *reduce the user's need to type*. Additional items such as notes, documents, etc., can be dragged in to a to-do, so *the to-do becomes resource for saving content and launching activity* on the task, like a pile or folder. But unlike a pile or folder, *a to-do has computational properties that support task management*. For example, each to-do has the property of importance that determines its *priority* or position in the list when the list is sorted for importance. *Users can change importance easily* by dragging a to-do up and down the list. A to-do can also have *time constraints* (e.g. a duration as in DM's list or Prof's calendar time-slots, or a deadline). Green 'warning bars' turning red (also used in [3]) are a *salient*

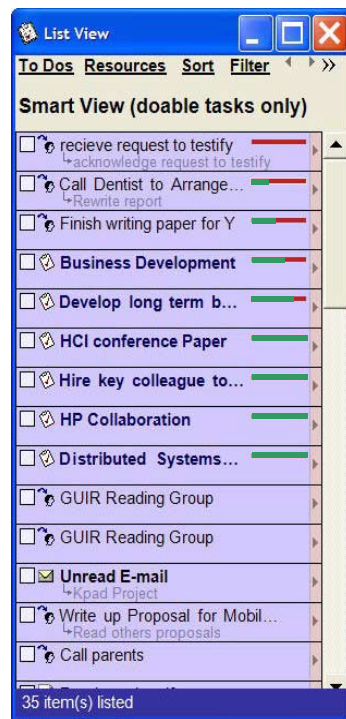


Figure 3. TaskVista: a personal task list manager component.

visualization to cue users of the urgency of approaching deadlines.

TaskVista also has a pop-out visualization (it slides out to the right) that shows tasks in a temporal view laid out as bars in a Gantt chart across a horizontal timeline to *support time management*. Users can see which tasks coincide and can directly manipulate date and time information.

Since we found that *relationships are an important factor* in task management, to-dos also have the *property of participants*. These are names extracted from sender and recipient information in email messages or from document content or properties and can be matched with details in a contact list (which the user can mark up to record who is important; this information could also be inferred from email response profiles [22]). TaskVista provides a contact widget (also used in [3]) for each task, so users can easily open a message to all participants with two mouse-clicks or select a subset one by one and then open a message to them.

To-dos can have other properties that, for example, show location, task or participant dependencies, and whether they are a project. Properties do not need to be specified up-front, and can change over time. The more information TaskVista has, the better job it can do at what we call 'smart' prioritization. The user can elect to have the system prioritize tasks, for example, because they are non-discretionary, because the participants are important, or because a deadline is approaching. It does so simply by pushing them up the list in a 'smart sort' view. The list can be sorted or filtered based on single or multiple properties to *support different kinds of task management activities*.

Finally, users do not have to decide if a to-do is a simple task or an entire project. They can *specify to-dos at any level of abstraction* and turn tasks into projects (encompassing other tasks) at any time depending on needs.

TaskVista was developed quickly in C# to allow us to conduct a very early informal evaluation with 9 volunteers. Users were guided through 10 task management exercises based on a real experience of one of the authors who was asked to insert an important and urgent presentation into an already crowded schedule. This required creating a new task from an email invitation to give a presentation, rating it as important, mailing people involved in existing tasks to notify them that one is now too busy to work on them, etc. This effort, even with a buggy prototype, provided a great deal of valuable feedback about how to refine our ideas for a TLM towards something we feel would be truly effective.

Overall, even with an early prototype, users were positive about many of its features, (rating them on average 3.94 in a scale from 1 = hate-it to 5 = love-it; standard deviation (stdev) was 0.84). The most popular features are the drag and drop creation of to-dos (4.67 on our scale, stdev 0.71) viewing the list on the desktop (4.56, stdev 0.73), the ability to sort by importance and deadlines (4.2, stdev 0.83) and the fact that it's easy to get at task participants (4.4, stdev 0.73).

The temporal view has so far been the least well received feature (3.0, stdev 0.71); it seems somewhat unnatural for those not used to looking at Gantt charts, even though we know that this is an accepted type of view for planning complex projects.

A critical challenge and valuable feedback for DARPA is that most users disliked the idea of explicitly adding a lot of metadata such as participants, planned start and end times, or tasks dependencies. Thus, the design of a TLM must encompass lightweight or automatic mechanisms to capture or infer such information. Fortunately one of the main thrusts of the DARPA program is to create just such mechanisms, which will incorporate reasoning and learning to lighten the load of users in prioritizing and planning.

CONCLUSIONS

We propose that the principal problem of task management is not poor prioritization, but the effort it requires and have outlined resources and methods people use that help ensure they are effective at this. We designed TaskVista as a tool to reduce this effort. It is faithful to field-derived insights into what factors and resources relate to task management. Designing and evaluating it early, even as we conducted our fieldwork, helped us maintain our focus as ethnographers in generating design recommendations. Our early evaluation of a TLM prototype confirmed our ethnographically derived requirements for an interactionally lightweight tool with intuitive visualizations and the capacity to work with underspecified and arbitrarily abstract content as it naturally occurs in task management practice.

ACKNOWLEDGMENTS

This work was supported in part by DARPA/IPTO under SPAWAR Contract N66001-03-C-8010. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied of the Defense Advanced Research Projects Agency, or the U.S. Government.

REFERENCES

- Allen, D. *Getting Things Done*. Viking (2001).
- Barreau, D. & Nardi, B. Finding and Reminding: File Organization from the Desktop. *SIGCHI Bulletin*, 27, 3 (1995), 39-43.
- Bellotti, V., Ducheneaut, N., Howard, M., & Smith, I. Taking Email to Task: The Design and Evaluation of a Task Management Centered Email Tool. In *Proc. CHI 2003*, CHI Letters 5 (1). 297-304.
- Bellotti, V. & Smith, I. Informing the Design of an Information Management System with Iterative Fieldwork. *Proc. DIS'00, Designing Interactive Systems*. ACM Press, (2000) 227-237.
- Blandford, A. & Green, T. Group and Individual Time Management Tools: What You Get is Not What You Need. *Personal and Ubiquitous Computing* 5, 4 (2001) 213-230.
- Campbell, C. & Maglio, P. Supporting Notable Information in Office Work. *Extended Abstracts CHI'03 Conference on Human Factors in Computing Systems*, ACM Press (2003), 902-903.
- Covey, S. *The Seven Habits of Highly Effective People*, Simon & Schuster (1989).
- DARPA IPTO BAA 02-21: http://www.darpa.mil/ipto/Solicitations/open/02-21_PIP.htm
- Hayes, G., Pierce, J. & Abowd, G. Practices for Capturing Short Important Thoughts. *Extended Abstracts CHI'03 Conference on Human Factors in Computing Systems*, ACM Press (2003), 904-905.
- Hoc, J. *Cognitive Psychology of Planning*. Academic Press (1988).
- Hutchins, E. *Cognition in the Wild*. Cambridge, Mass: MIT Press (1995).
- Jones, S. & Thomas, P. Empirical Assessment of Individuals' 'Personal Information Management Systems'. *Behaviour and Information Technology*, 16.3, (1997) 158-160.
- Lansdale M. The Psychology of Information Management. *Applied Ergonomics*, 19, (1988), 55-66.
- Malone, T. How Do People Organize Their Desktops? Implications for the Design of Office Information Systems, *ACM Transactions on Office Information Systems*, 1, (1983) 99-112.
- Miller, G., Galanter, E. & Pribram, K. *Plans and the Structure of Behavior*. Holt, Reinhart & Winston Inc. (1960).
- Palen, L. Social, Individual & Technological Issues for Groupware Calendar Systems. *Proc. CHI'99 Conference on Human Factors in Computing Systems*, ACM Press (1999) 17-24.
- Payne, S. Understanding Calendar Use. *Human-Computer Interaction*, 8 (1993) 83-100.
- Scaife, M. & Rogers, Y. External Cognition: How Do Graphical Representations Work? *International Journal of Human-Computer Studies*, 45 (1996) 185-213.
- Strauss, A. Work and the Division of Labor. *The Sociological Quarterly*, 26,1, (1985), 1-19.
- Takkinen, J. & Shahmehri, N. Task-Oriented Restructuring of an Application Domain: A Multi-Agent Architecture for Doing Things in Internet E-Mail," *Proc. 32nd Hawaii International Conference on System Sciences* (1999).
- Tolmie, P., Pycoc, J. Diggins, T., MacLean, A. & Karsenty, A. Ubiquity. Unremarkable Computing. *Proc CHI 2002*, CHI Letters 4(1).
- Tyler J. & Tang J. Email Responsiveness and Rhythms: Study and Design Implications. *Proc. ECSCW'03 European Conference on Computer Supported Cooperative Work*. Kluwer (2003) 238-258.
- Whittaker, S., & Sidner, C. Email Overload: Exploring Personal Information Management of Email. In *Proc. CHI'96, Conference on Human Factors in Computing Systems*, ACM Press (1996) 276-283.