

Juggling Work Among Multiple Projects and Partner

Peter Scupelli, Susan R. Fussell, Sara Kiesler, Pablo Quinones, Gail Kusbit

Human Computer Interaction Institute

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213 USA

{pgs, sfussell, kiesler, paq, gkusbit}@andrew.cmu.edu

Abstract

Prior research suggests people have trouble juggling effort across multiple projects with multiple partners. We investigated this problem, with an experiment where groups of four participants enacted the roles of police detectives. Each detective was assigned two homicide cases, each case with a different partner. To solve each case, detectives read their case documents and discussed relevant information with their partners. Half the groups used IM to communicate and the other half used an enhanced IM tool called Project-View IM (PVIM). PVIM lists partners and joint projects and lets users know what a partner is working on. We analyzed keystroke level computer activity and the content of conversations. Generally, work unfolded as follows: coordinate across cases, start first case, read documents, coordinate within case with partner, switch to second case, and so on, but with frequent interruptions. We describe implications of our findings for theories of multitasking.

1. Introduction

Organizations increasingly use teams as a way to organize work. People often work on multiple projects and teams at the same time (e.g., [12],[13]). Police detectives work on multiple cases; consultants work on multiple accounts; university professors work on different research projects; and engineers, designers and marketing professionals work on multiple product development teams. The proliferation of popular books on organizing one's time attests to the multiplicity of responsibilities that many workers face today.

In this paper, we focus on a type of multitasking familiar to many who engage in intellectual teamwork—working on two similar projects with different partners when both have the same deadline. One familiar example of this is the preparation of paper submissions for conferences. In our own experience,

the week before a major conference deadline is chaotic, as we must coordinate paper-writing subtasks (e.g., statistical analysis and writing) with different sets of co-authors on different papers, all equally urgent. Similar examples include preparing grant proposals for submission deadlines, preparing class materials for the first day of the semester, and preparing year-end reports. Although multitasking across similar tasks is not the only type of interest (e.g., often one must work on two entirely distinct activities, such as teaching and grant preparation), this type of multitasking is common and thus a good starting point for our investigations.

When working on two similar projects with the same deadline, people face coordination problems at two levels: They must allocate time appropriately across their multiple projects and they must coordinate their efforts with their partners within each individual project. The goal of the current paper is to provide a deeper understanding of these two coordination processes. In addition, we evaluate the benefits of a prototype Instant Messaging client called Project-View IM (PVIM). PVIM is intended to help people allocate work effort more effectively both across multiple tasks and within a single task. PVIM provides users with a view of their project members' names in lists nested within their different projects. PVIM also indicates project members' project-related activity.

In the remainder of this paper we first review literature on multitasking and tools to support it. We then present a laboratory study in which participants acted as detectives working on two different cases with two different partners. We collected survey responses and analyzed keystroke logs and participants' messages to develop a finer-grained description of the multitasking process. We also evaluate the benefits of PVIM for multitasking by comparing task processes and performance for participants using regular IM and those using PVIM. We end with a discussion of the implications of our findings for a theoretical understanding of the process of multitasking.

1.1 The problem of multiple tasks

We conceive of multitasking as the interweaving of different project or task trajectories. The term “trajectory” refers to the sequence of activities through which a person, resource, or task moves [17]. For information workers, work trajectories often involve a variety of individual tasks embedded in larger team projects [5]. Trajectory management can present significant coordination challenges. Often people can concentrate on only one complex task at a time. Working on that task means not working on other tasks, thereby affecting the outcomes of the tasks that have been set aside. When working on multiple tasks of equal importance, people tend to allocate attention and effort unevenly [5], leading to stress [13].

In a detailed study of multitasking among banking professionals, Gonzales and Mark [7] found that people worked on a project an average of 11 minutes before switching (voluntarily or by being interrupted) to another project. Also, they worked continuously on individual tasks within a project, such as checking documents or using the computer, for extremely brief periods of time. For example, mean time spent working on a PC was only 2:52 minutes, and mean time talking to visitors to their office was 3:34 minutes. Thus, the employees observed in these studies multi-tasked extensively, both across and within projects.

When team members are co-located, they can coordinate work trajectories through their passive awareness of workplace activities and informal face-to-face communication [12]. For example, a worker who is currently focused on one of his or her many tasks may be reminded of other, equally important, tasks by seeing a colleague in the hallway. In addition, being co-located seems to allow collaborators to better gauge their interruptions: In Gonzales and Mark’s [7] study, people spent more continuous time on a project when they were co-located than when they were in a remote location, despite the fact that co-located workers were interrupted more than remote workers. When collaborators are remotely located, there are many fewer opportunities for workplace activity awareness and informal face-to-face communication, making trajectory management more difficult.

1.2 Tools to support multitasking

With the exception of media spaces (e.g., [12]), few of which are in widespread use, tools for remote collaboration are oriented toward supporting teamwork within a single project rather than across projects. Recent work suggests that Instant Messaging (IM) may be an exception. IM facilitates geographically

distributed work by supporting informal communication (e.g., [9],[11]). Some people even prefer IM to informal face-to-face conversation because they perceive it allows multitasking [15]. On the other hand, IM exchanges can be disruptive (e.g., [3][4]) and may not help people juggle multiple projects and teams effectively [5].

Commercial IM clients support near-synchronous dyadic text-based conversation. They also provide awareness of who is logged on and their status, (e.g., idle, away). Current IM notifications, such as “online” and “away” icons, provide social awareness of others’ availability ([11],[15]). However, because these notifications offer little information about what colleagues are actually doing, they may lead collaborators to time their interruptions poorly.

Prior work has examined the effects of using IM on individuals’ management of work effort across multiple collaborative web development projects [5]. Groups of four participants completed four web design projects. Each participant worked on two projects, each with a different partner who was either in the same room or connected via IM. Nearly all participants divided their time unequally between projects, spending about two-thirds of their time on one of their two tasks. In addition, participants tended to favor their project with a co-located partner over a remote partner. These results suggest that simply reminding people about others via an IM interface is insufficient for helping them allocate their effort equally across two tasks of the same importance.

Several applications have explored ways to improve people’s information about multiple work activities and partners. For example, Piazza shows who is doing similar tasks [10], Babble supports opportunistic interaction [2], and Rear View Mirror provides presence awareness and text communication [9]. Other tools provide peripheral awareness information (e.g., [18]). The goal of these systems, however, is typically not to provide detailed information about partners’ project-related activities. One exception is ActivityExplorer [14], which provides awareness and a log of document-related activities at the project level.

1.3 Project-View IM (PVIM)

The applications described above support individual and group awareness but most were not designed specifically to manage work effort across multiple tasks and partners. These applications provide awareness of coworkers or groups, but the information is not automatically organized by task or project. IM and other applications notify people about others’

presence and/or availability, but not what they are doing and whether it is relevant to the user.

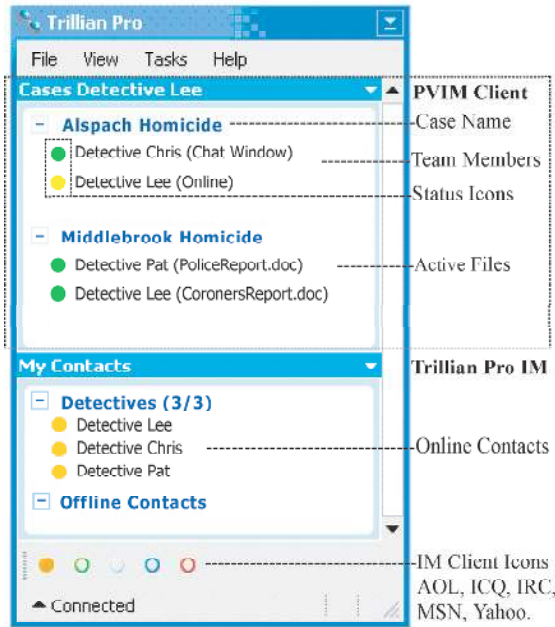


Figure 1. PVIM module in Trillian Pro.

The goal of PVIM [16] is to better support allocation of effort across multiple teams and tasks by helping people remember their total workload and maintain awareness of team members' activities on each task over time. PVIM makes team members and task responsibilities visible within an existing IM tool (see Figure 1). Different projects (here, homicide cases) and team members of each project are located in the top half of the PVIM window and a standard contact list is located in the bottom half. Status icons indicate who is working on the shared cases. This information can be used to facilitate coordination with work partners (e.g., by showing when the user and the partner are both working on the same case).

1.4 Goals of the current study

The current study had two goals. First, we aimed to enhance our theoretical understanding of the process of multitasking by examining how people allocate effort within and across projects and time with different partners. This goal is primarily descriptive and we present data showing how people's activity changes across time. Second, we wanted to evaluate the effects of PVIM on the multitasking process. We hypothesized that PVIM would help people allocate effort more easily and more effectively both across and within projects. We also hypothesized that PVIM would reduce self-reported workload by reducing the cognitive effort required to keep track of multiple

projects and partners. We pursued both goals in a laboratory experiment, described in the next section.

2. Study

Four participants at a time worked as detectives on murder cases. Each of the four detectives worked on two cases. Each participant had a different remote partner for each of his or her cases. The participant's job as detective was to examine the evidence contained in electronic documents associated with their two assigned cases and to work with their partner on each case to find the suspect who should be arrested. Half of the groups used PVIM and half used IM to discuss their cases. We studied participants' communication and work across the entire session to examine their allocation of effort across projects (cases) and across subtasks within each case (reading documents, discussing the case with their partner).

2.1 Method

2.1.1 Participants. Eighty undergraduate students (38 female, 42 male) were run in groups of four. They received \$15 for participating in the study.

2.1.2 Materials. Experimental tasks consisted of homicide cases modeled after actual detective work. Each homicide case had four suspects, one of whom was the murderer. For each case, there were eight electronic documents of 275-300 words containing prior interviews with suspects, a coroner's report, a police report, and two witness interviews. Cases were pretested to ensure that they required equivalent times to solve. Participants were given half of the documents they needed to use to solve each case. Participants had to discuss the cases with their partners to solve the homicides. We created on-line worksheets for each case for participants to check off suspects with alibis or no opportunity or motive to commit the crime. We also administered a post-task survey in which participants rated their experiences during the session and completed the NASA TLX workload scale [8].

2.1.3 Equipment and software. Participants used Microsoft Word to read case materials, Trillian Pro to communicate via IM or PVIM, and Internet Explorer to complete online case worksheets. WinWhatWhere keystroke logging software was used to create a time-stamped record of which application a participant was using and their keystrokes.

PVIM [16] was implemented as a plug-in for Trillian Pro, a commercial Windows IM client. PVIM has both a client side and a server side. The PVIM

client monitors the active window on the desktop and identifies active IM conversations. A name-matching algorithm assigns active windows and IM partners to a case and updates other team members via the PVIM server.

Figure 1 shows how the client side of PVIM would be set up for Detective Lee, who is working on the *Alspach Homicide* with Detective Chris and the *Middlebrook Homicide* with Detective Pat. The top part of the screen lists the detective's cases and the detectives assigned to each case. To the left of each detective's name is one of three status icons. Green indicates that a partner is online and working on the case; yellow indicates that a partner is online and not working on the case, and red indicates that a partner is offline. If a partner is online and working on the detective's case, the detective can see the file name of the document the partner is working on.




Figure 2. Experimental paradigm: Each detective worked on two cases, each with a different partner in a different room.

2.1.4 Procedure. Each group of four participants was randomly assigned to the IM or PVIM condition. Each participant at a desk with a PC out of view of the other participants. Each was assigned a detective role and name to use for the experiment. Participants were given a practice case to familiarize themselves with the process of solving homicide cases. They were then told that, like real detectives, they would work on two homicide cases at a time with a different partner for each case. Each detective was then randomly assigned two cases to solve, each one with a different partner, located in a different room (see Figure 2).

2.2 Measures

To understand how participants allocated their work time across their two cases and coordinated subtasks within each of their cases, we collected four sets of dependent measures: conversational content, work effort, work strategy, and self-reported workload.

2.2.1 Conversational Content. All IM messages were coded in terms of whether they pertained to coordination or task work. Coordination messages were sub-classified into one of three types: *coordinating work across cases* (e.g., discussing which cases to work on first), *coordinating work within a case* (e.g., deciding which suspects to discuss first), and *wrapping up* the case by completing paperwork. Task-related messages were sub-classified into *exchanging task documents* (copying and pasting text from case documents) and *discussing the task* (e.g., discussing the reliability of a witness' testimony). Messages that could not be coded into these five categories were coded as *other*. Examples of each category are given in Table 1. A trained coder coded all messages (N = 6608). A second trained coder coded 758 (11%) of the messages; agreement between the two coders was good (kappa = .78).

Table 1. Message coding content categories

Category	Examples
Coordinating across cases	<i>ok I have alspach ... i'm gonna work on that one second</i>
Coordinating within a case	<i>I have the interviews with Arturo, Garry, and Johnathan for this case, but I haven't finished reading them all.</i>
Exchanging documents	<i>"Jeffry and I left the party and went to the bleachers"</i>
Discussing the task	<i>there weren't any fingerprints on the gun</i>
Wrapping up	<i>did you have anything checked on the police report by hossler?</i>

2.2.2 Work effort. To measure the allocation of work effort, we computed time spent by each participant on each case from keystroke activity logs. Post-processing software was used to identify which case a participant was working on at five second intervals. Technical difficulties prevented the keystroke activity analysis for 9 PVIM logs and 8 IM logs. We ran a χ^2 test and found no significant difference between the survey

responses from the 63 participants with keystroke activity data and 17 missing such data.

Overall, participants spent more time on one of their cases in the first half hour, and spent more time on the other case in the second half hour. Thus, we defined for each participant a *first case* and a *second case*.

Next, we distinguished time spent working on documents (Word documents, online forms) from time spent chatting on the basis of the WinWhatWhere logs. Work effort spent on documents and chat was summed for each participant's first case and second case for each quarter of the hour-long session.

2.2.3 Work strategy. On a post-experimental survey, participants responded to several questions about how they allocated effort across their two assigned tasks on a scale from 1 (strongly disagree) to 7 (strongly agree). Factor Analysis with Varimax rotation indicated that three of these questions formed one dimension measuring sequential versus parallel work organization ($\alpha = .86$): "I finished all my work on one task for one case before working on a task for the other case" (reversed), "I was always interrupting one case in order to work on the other case," and "I usually worked on more than one case at the same time." Scores were averaged for the three questions prior to analysis, creating a scale in which a low score indicated working on the cases serially whereas a high score indicated working on them in parallel.

2.2.4 Workload. Participants rated their workload across both tasks on a modified version of the NASA TLX [8]. They rated their mental demand, temporal demand, effort, performance, and frustration on a scale from 1 (low) to 7 (high). The five workload questions (with performance scored in reverse) formed a reliable scale ($\alpha = .73$) and were averaged for analysis.

3. Results

We discuss the results in two sections. First, we describe the general patterns of multitasking across all participants, regardless of which IM client they used. Then, we examine how the PVIM tool altered multitasking strategies. In both sections, we discuss multitasking across cases, multitasking within cases, and the role of different types of messages in the coordination process.

3.1 Overall patterns of multitasking

In this section, we describe how people allocated their time and effort across the various activities in which they were involved.

3.1.1 Multitasking across cases. Participants worked on two cases with two different partners, and they were free to allocate their efforts across their two cases in any way that they wanted. They could work on them serially, or they could switch back and forth between them. In their self-report survey data, participants indicated the extent to which they took a serial versus parallel approach to their work on their two tasks. On the resultant scale, a high value indicates that they generally worked on both cases at once, whereas a low value indicates a serial approach. Scores on this variable ranged from 1 to 7 (the maximum), but the average of 2.59 (SD = 1.83) suggested that participants perceived they worked on one case at a time.

We next examined how people allocated their effort across their two cases as a function of time using our keystroke logging data, which allow us to identify what people are working on at any given moment. For our purposes, we equate work effort for an activity with the time spent working on that activity. We divided the work activity of each participant into four equal quarters (roughly 15 minutes each) and calculated how much time in each quarter the participant spent on each of his/her cases. In Figure 3, we show participants' percent work effort on the two cases by quarter. Not surprisingly, participants worked most of their time on one of their cases in the first two quarters and then shifted their effort to their second case in the third and fourth quarters.




Figure 3. Work effort by quarter for participants' first case (bottom) and second case (top).

Gersick has shown that teams that reconsider their work at about the midpoint of their task and then make a greater effort to reach an overall goal perform better than groups that do not reconsider their work at midpoint [6]. The participants in this study seem to have done that generally, as they started switching to their second case at the midpoint.

3.1.2 Multitasking within a case. Successful solutions to the homicide cases required both individual activity (reading and thinking about case documents) and group work (discussing the case with a partner).

Participants could divide their effort across documents and chat in any way they saw fit. Consistent with other researchers' findings (e.g. [7]) our participants spent brief amounts of time working on a document or in chat before switching to another document or chat conversation. The amount of time spent on an activity before switching or being interrupted ranged from less than one second to almost 11 minutes but on average was extremely brief (for chatting, $M = 21.43$ sec., $SD = 36.20$ sec.; for work on documents, $M = 12.79$ sec., $SD = 44.22$ sec.).

The strategy of devoting initial time to individual work on documents rather than to communication paid off in lower self-reported workload through the whole session. We found that the more participants chatted about their first case during the first quarter, the higher their self-reported workload across the whole session, as measured by the TLX survey ($r[63] = .26, p < .05$).




Figure 4. Work effort on documents and chat by quarter and case. (From bottom to top: first case documents, first case chat, second case documents, second case chat.)

Figure 4 illustrates how participants, on average, allocated their effort across documents and chat for their two cases over the four quarters of the hour-long experiment. Because work effort on documents and chat for a participants' two cases are additive, we could not analyze all of the measures simultaneously. Instead, we analyzed results for time spent on documents in each participant's first case. We ran a repeated measures ANOVA on the percent of participants' time spent on case one documents in which quarter was the repeated variable and IM condition was a between-subjects variable. Consistent with Figure 4, time spent on documents for the first case changed over quarters ($F [3, 183] = 96.69, p < .001$). Post hoc comparisons of adjacent quarters

showed that participants devoted different amounts of time on these documents in each quarter. However, the largest difference was found between the second and third quarter ($F [1,61] = 58.13, p < .001$), as participants shifted their attention from their first case to their second case after the midpoint.

3.1.3 Using IM for coordination. From Figure 4, we can see that talking with a partner played an important role in participants' work on their homicide cases. In order to understand what types of messages were exchanged and their role in the task process, participants' messages were coded into two types of task-related messages (exchanging information, discussing the case) and three types of coordination messages (coordinating across cases, coordinating within a case, and wrapping up a case).

On average, each conversation had 81.65 messages ($SD = 39.64$), each of which was 8.43 words long ($SD = 8.09$). To examine how these messages were distributed over time, we calculated the percent of messages of each type for each block of 10 messages. We used a block approach to comparing dialogues of differing lengths rather than a percentile approach because a qualitative examination of the data suggested that longer dialogues resulted from longer time spent discussing the case itself, not from longer initial coordination or final wrap-up exchanges. We examined the first 10 blocks to ensure that sufficient dyads were represented in each block (however, since conversations averaged 82 messages, later blocks represent fewer dyads than earlier blocks).

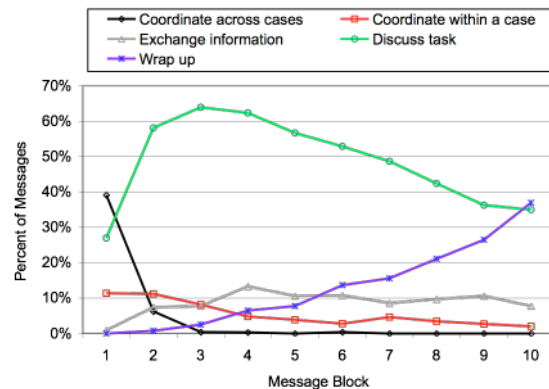


Figure 5: Distribution of messages across the first 10 blocks of 10 messages (N = 5873 messages from 80 participants).

As shown in Figure 5, the distribution of message content differs significantly by block ($\chi^2. [54] = 2709, p < .001$). At the beginning of a case, substantial messages (40%) were exchanged to coordinate across each participant's two cases, but the rate of these

messages decreased dramatically after the second block. Because participants had different partners for each case, they had to coordinate when they would start working together on their mutual task but once this plan was established, no further discussion was required. Not surprisingly, the majority of the messages in the middle of the session involved the tasks of solving the cases. As the cases came to a close, participants increased wrap up activities.

3.2 Effects of PVIM on multitasking

Half of the participants were provided with our PVIM enhanced IM tool. PVIM provided information about when a partner was working on a joint case and what document within that case he or she was currently using. We examined the effects of PVIM versus standard IM on multitasking across tasks, multitasking within tasks and use of chat for coordination.

3.2.1 Multitasking Across Cases Participants using IM and PVIM spent the first half of the experimental session mostly on one of their two cases, shifting about midway to their second case.

However, as shown in Figure 6, a closer look at the two conditions suggests that participants in the PVIM condition experienced a smoother transition from their first case to their second case. By the second quarter, participants using PVIM spent about 30% of their time working on their second case, as compared to 20% in the IM condition.

3.2.2 Multitasking within a case. As shown in Figure 7, participants in the PVIM and IM conditions differed significantly in the way they allocated effort across documents and IM chat. In the first quarter, those using PVIM spent a greater percentage of their time working on documents for both their first and second cases. Those using IM spent more of their time chatting with both partners. We think this difference may derive from problems IM groups faced in organizing their work in absence of knowledge about what partners were working on. This interpretation is supported by the analysis of their chat messages in the next section.

We ran a repeated measures ANOVA on the percent of participants' time spent on documents for their first case by each quarter. We found a significant interaction between quarter and condition ($F[3, 183] = 2.82, p < .05$). Post hoc comparisons of adjacent quarters showed that the difference between effort




Figure 6. Work effort by quarter and condition for participants' first case (bottom) and second case (top).

devoted to documents for the first case in the first and second quarters is larger for PVIM than IM ($F[1, 61] = 4.38, p < .05$), whereas the difference in effort between the second and third quarters is marginally greater for IM than for PVIM ($F[1, 61] = 3.27, p < .08$). This marginal interaction effect can be explained by those using IM who shifted from their first to their second case after the midpoint. The change in effort on documents for the first case between the third quarter and fourth quarter for those using IM and those using PVIM is not significant ($F[1, 61] < 1$).

Figure 7 shows that participants in the IM and PVIM conditions differed most in their ratio of working on documents versus chatting during the first quarter. IM groups talked more about their first case ($F[1, 61] = 5.74, p = .02$), and PVIM groups allotted more time to reading and thinking about their documents ($F[1, 61] = 2.93, p = .09$). Perhaps participants using IM spent more time deciding who would do what, whereas participants using PVIM knew who had started working on which documents and began to work earlier. Spending more time reading documents in the first quarter permitted PVIM participants to discuss their first case with their partner in the second quarter and to start their second cases.

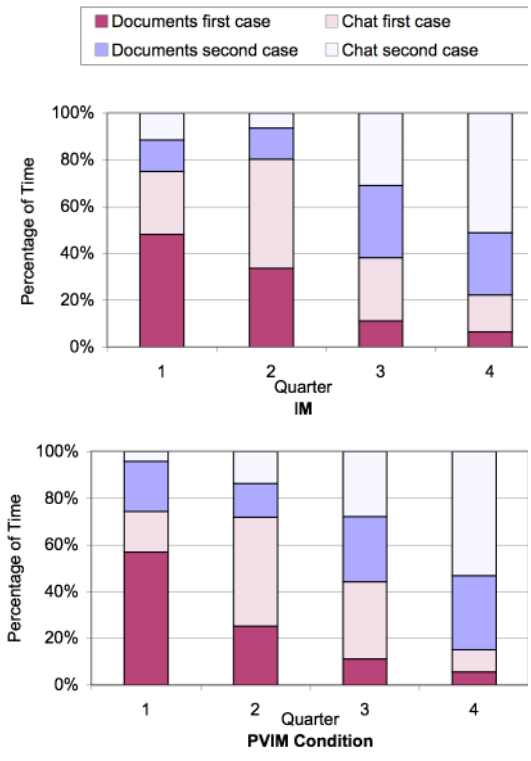


Figure 7. Work effort on documents and chat by condition, quarter and case. (From bottom to top: first case documents, first case chat, second case documents, second case chat.)

3.2.1 Using IM for coordination. Overall, the number of messages produced by participants in the PVIM and IM conditions did not differ significantly (for PVIM, $M = 78.78$, $SD = 41.54$; for IM, $M = 86.43$, $SD = 39.17$; $F < 1$). However, there was a key difference in the types of messages participants in the two conditions produced. Pairs using PVIM had a significantly lower proportion of messages pertaining to coordination across cases than did those in the IM condition ($t [18] = 3.82$, $p < .001$). Figure 8 shows how pairs with PVIM begin working quickly. Proportions of messages across the three coordination categories are shown for the first five blocks of 10 messages. People using PVIM had significantly fewer total messages coordinating across cases, and they coordinated earlier ($\chi^2 = 157.05$, $p < .001$).

4. Discussion

The results provide insights into the process of multitasking across multiple projects and collaborators. Our finer-grained analysis of how people coordinated

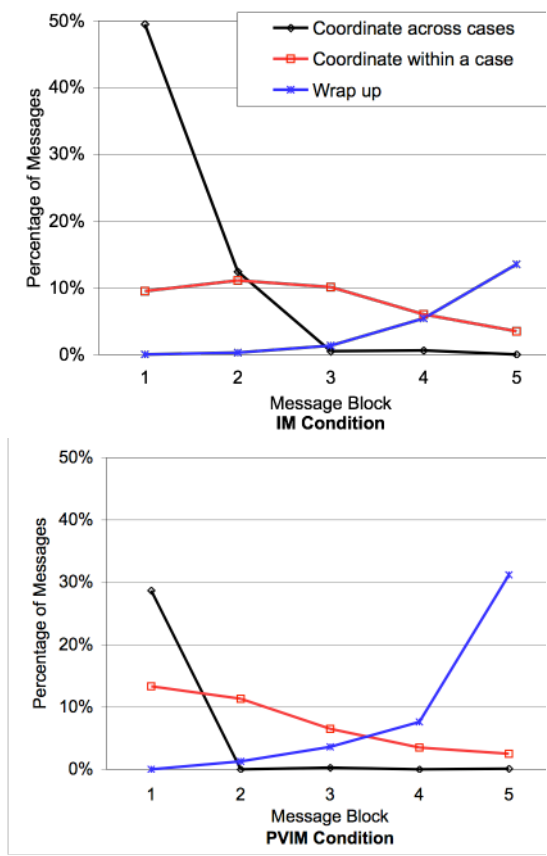


Figure 8. Percentage of three types of coordination messages over the first five blocks of 10 messages for IM (top) and PVIM (bottom). (N = 2360 messages.)

across cases, and between individual and group work within cases, suggests that people who are multitasking prefer to work on tasks more-or-less one at a time. However, the need to collaborate with different partners on each task makes it impossible for them to employ a purely serial approach. Within a case, participants spent more time on their individual work during initial phases of the case, and then shifted to discussing the case with a partner. To some extent, this phenomenon is due to the nature of our experimental task, which required that participants read the case documents to have a productive conversation about the crime. Many real-world tasks likewise require individual work before team members can collaborate.

Although patterns of effort allocation across cases were similar for participants with PVIM and IM, using PVIM with its enhanced awareness of what partners were doing shifted the process of multitasking. Those with PVIM exchanged fewer messages pertaining to the coordination of activity across cases, presumably because the display informed them about what partners

were currently doing. Overall, the results suggest that the PVIM tool could help people allocate attention across multiple tasks with multiple partners.

In the remainder of this section, we discuss some theoretical aspects of our findings, the limitations of this work, and our future directions.

4.1 The nature of multitasking

All our participants multi-tasked in that they all completed two homicide cases in a one-hour period, and we could discern a general course of multitasking over time. That is, participants typically coordinated with their partners first, to decide which case to work on first. Then they began perusing documents in this first case, and then discussed the case with their partner. At about the midpoint they switched cases, following the same general path. At the end of the session, they wrapped up by completing reports.

Within this seemingly smooth overall trajectory, we found multiple interruptions and task-switching. We also found considerable variation in participants' work strategies. Some did not complete work on the case they began first until the fourth quarter. Others started work on a second case while they were working mainly on their first case. Participants interrupted and were interrupted often, and they spent only brief amounts of time working on a document or in chat before switching to another document or chat conversation.

A key insight from our analyses is that the extent and nature of multitasking we observe will depend on the level at which the data are aggregated. If we look at what each individual is doing at each point in time, then the amount of interruption and task-switching seems almost chaotic. People appear unfocused, as if they were not following any trajectory. But if we aggregate over people and/or larger chunks of time, a much more coherent trajectory emerges, with more sequential multitasking, better efforts at preparation, and more extensive coordination with partners.

4.2 Tools to support multitasking

PVIM seemed to help participants allocate effort across tasks and partners. We had developed PVIM to provide participants with several different types of assistance—a reminder of responsibility for two different cases, information on who was currently working on a shared case, and information about what file each person was using. We do not know from our data which of these features facilitated participants' coordination. In future work, we will investigate this issue by selectively manipulating specific awareness features in different experimental conditions.

It is possible that PVIM benefits people not because of its status as a tool to support cooperative work, but because it provides a reminder of one's obligations, much as to-do lists, or because allowing a remote partner to view what one is doing creates social pressure [1]. These findings suggest that PVIM may influence multitasking through reciprocal work awareness and implicit accountability. We do not believe this explanation accounts for our results for two reasons. First, each participant was provided with a hard copy assignment sheet, which probably sufficed as a reminder of work obligations. Second, most of the benefits of PVIM occurred early in the work process, leading us to believe that it aided the coordination of work through awareness rather than helping with allocation of attention or feelings of responsibility. Nonetheless, future research should attempt to disentangle these competing hypotheses.

We had also thought PVIM might benefit multitasking by helping partners time their interruptions better. Other work suggests that displaying what a partner is doing can reduce unwanted intrusions [4]. However, our analysis of the chat logs suggests that this was not the case. Both groups chatted about their cases. The PVIM group had to chat less about coordination across cases, and got to work on their cases earlier but we have little evidence their interruption behaviors differed.

4.3 Limitations and future directions

The present study has a number of limitations that must be addressed in future research. One limitation stems from the way in which we operationalized multitasking. We looked not at the total number of projects completed in a given period of time (always two, in our study) but at the ways in which the trajectories of these projects were interwoven. An alternative conceptualization of multitasking as number of simultaneous projects might lead to other questions for investigation.

A second limitation is that we focus on one type of multitasking—scenarios in which people must work on two similar and equally important tasks with different collaborators. Future research will need to examine whether our results hold up when people are working on two entirely different projects that draw upon different skills.

A third limitation concerns our study design. In order to obtain strong experimental control, we used tasks in which each person had only one partner. In the real world, people have many more projects, with many more partners (e.g., [13]). Our assumption is that the multitasking processes observed in our experiment

will generalize to the more complex multitasking found in the real world. Although this remains to be tested, it is noteworthy that the amount of time we found participants spent on one activity before switching to another is in line with field data [7].

Other limitations to PVIM make it unsuitable in its current form for managing the complexities of real-world work. For example, people's willingness to share private information varies across partners (e.g., [9]), so future versions of PVIM will need to provide privacy controls.

5. Conclusion

Prior research suggests that people face challenges when allocating their work effort across multiple projects. To help alleviate this problem, we introduced a tool called Project View IM that provides reminders of ongoing projects and partners and awareness of partners' task-related activities within the context of an ordinary IM client. An experiment comparing how people allocate effort across two tasks with different partners using PVIM versus ordinary IM suggests that PVIM improves multitasking. Two ways in which it does so are by helping people get down to work more quickly and by helping them make smooth transitions from one task to another. The study adds to our understanding of the ways in which providing workers with detailed awareness benefits work processes.

6. Acknowledgments

This material is based upon work supported by National Science Foundation Grants #0329077 and 0325047. We thank Pablo Quinones and Mariesa Cash for running the study, Kyle Cunningham for technical assistance, Congrui Chen for programming PVIM and the log parser, and Gail Kusbit, for editorial assistance.

7. References

[1] Bradner, E. & Mark, G. (2001). Social presence with video and application sharing. *Proceedings of GROUP 2001* (pp. 154-161). NY: ACM Press.

[2] Bradner, E., Kellogg, W. E., & Erickson, T. (1999). The adoption and use of "BABBLE": A field study of chat in the workplace. *Proceedings of ECSCW 1999*.

[3] Cutrell, E.B., Czerwinski, M., & Horvitz, E. (2000). Effects of instant messaging interruptions on computing tasks. *CHI 2000 Extended Abstracts* (pp. 99-100).

[4] Dabbish, L. & Kraut, R. E. (2004). Controlling interruptions: Awareness displays and social motivation for coordination. *Proceedings of CSCW 2004* (pp. 182-191).

[5] Fussell, S. R., Kiesler, S., Setlock, S. D., Scupelli, P., & Weisband, S. (2004). Effects of instant messaging on the management of multiple projects. *Proceedings of CHI 2004* (pp. 191-198). NY: ACM Press.

[6] Gersick, C. J. G. (1988). Time and transition in work teams: Toward a new model of group development. *Academy of Management Journal*, 3, 9-40.

[7] González, V. M., & Mark, G. (2004). "Constant, constant multitasking craziness": Managing multiple working spheres. *Proceedings of CHI 2004* (pp. 113-120).

[8] Hart, S. G., & Staveland, L. E. (1988). Development of a multi-dimensional workload rating scale: Results of empirical and theoretical research. In P. A. Hancock & N. Meshkati (Ed.), *Human mental workload* (pp. 139-183). Amsterdam: Elsevier.

[9] Herbsleb, J.D, Atkins, D. L., Boyer, D. G., Handel, M., & Finholt, T. A. (2002). Introducing instant messaging and chat in the workplace. *Proceedings of CHI 2002* (pp. 171-178).

[10] Isaacs, E., Tang, J.C., & Morris, T. (1996). Piazza: A desktop environment supporting impromptu and planned interactions. *Proceedings of CSCW 1996* (pp. 315-324).

[11] Isaacs, E., Walendowski, A., Whittaker, S., Schiano, D. J., Kamm, C. (2002). The character, functions, and styles of instant messaging in the workplace. *Proceedings of CSCW 2002* (pp. 11-20). NY: ACM Press.

[12] Kraut, R. E., Fish, R.S., Root, R.W., & Chalfonte, B.L. (1990). Informal communication in organizations: Form, function, and technology. In S. Oskamp & S. Spacapan (Eds). *Human reactions to technology: The Claremont Symposium* (pp. 145-199). Beverly Hills, CA: Sage.

[13] Leroy, S., & Sproull, L. (2004). When team work means working on multiple teams. NYU Stern School of Business. Manuscript submitted for publication.

[14] Millen, D. R., Muller, M. J., Geyer, W., Wilcox, E., & Brownholtz, B. (2005). Patterns of media use in an activity-centric collaborative environment. *Proceedings of CHI 2005* (pp. 879-888). NY: ACM Press.

[15] Nardi, B.A., Whittaker, S., & Bradner E. (2000). Interaction and outeraction: Instant messaging in action. *Proceedings of CSCW 2000* (pp. 79-88). NY: ACM press.

[16] Scupelli, P., Kiesler, S., Fussell, S. R., & Chen, C. (2005). Project View IM: A tool for juggling multiple projects and teams. *Proceedings of CHI 2005* (pp. 1773 - 1776). NY: ACM press.

[17] Strauss, A., Fagerhaugh, S., Suczek, B., & Weiner, C. (1985). *Social organization of medical work*. Chicago: University of Chicago Press.

[18] Tang, J., Yankelovich, N., Begole, J.B., Kleek, M.V., Li, F., & Bhalodia, J. (2001). ConNexus to Awareness: Extending awareness to mobile users. *Proceedings of CHI 2001* (pp. 221-228). NY: ACM Press