

The Complex Structure of Office Work: Tasks, Activities & Modes

Wally Smith

Department of Information Systems
The University of Melbourne
email: wsmith@unimelb.edu.au

Abstract

An observational study of the busy week of a solicitor is reported in which the deep structure of work is revealed in terms of tasks, activities and modes. The need to trade-off two sorts of contextual continuity (task vs mode) is presented as a fundamental cause of work fragmentation. This analysis extends recent research that has tended to analyse surface activity when investigating the role of informational artifacts to aid task management. The findings here show how meta-work is heavily integrated into task work and point to the limitations of task management tools that separate things-to-do planners and schedules from the content of work.

Keywords

multitasking; interruptions; task switching; knowledge work

INTRODUCTION

Most forms of knowledge work demand individuals to divide their time across a number of ongoing tasks. People frequently switch task in response to changing priorities and external interruptions. Computers and other informational artifacts help people to maintain the continuity of their various streams of work (Kirsch, 2001) through planning aids, reminders, or by more subtle contextual support for re-entering previously suspended tasks. The ability to multi-task over long periods is distinct from our limited ability to carry out multiple practiced activities simultaneously; a phenomenon studied extensively by psychologists. Studies of the broader type of multi-tasking and the role of informational artifacts has a long tradition in ergonomics (e.g., Beishon, 1969) and was an early focus of interest in Human-Computer Interaction during the 1980s as the graphical user interface and the desktop metaphor took hold (Miyata & Norman, 1988; Card and Henderson, 1987). As Jeff Raskin points out (Raskin, 2001), just as the physical document left on a desk supports the resumption of an associated task, so too the file icon left on a virtual desktop can serve as an indirect reminder to act. This early practical interest was cemented by growing theoretical insights into the intimate relationship between intentional action and informational environments (Suchman, 1987).

More recently the fragmented nature of work, and the role of informational artifacts as supporting or disrupting multi-tasking, has re-emerged as a focus of interest (e.g., Kirsch, 2001; McFarlane, 2002; McFarlane & Latorella, 2002; Czerwinski, Horvitz & Wilhite, 2004; Gonzalez & Mark, 2004). A website now exists that points to a large body of associated literature: <http://interruptions.net/literature.htm>. This re-emergence of interest is associated with the broader problem of how *context* is maintained or disrupted by the design of mobile technology, ubiquitous computing and virtual work teams. In a special issue of the journal Human-Computer Interaction, Kirsch (2001) makes the point that if we are to design truly ubiquitous computing, then we must understand the 'deep structure' of work, so that the context(s) of situation A can be recreated sufficiently in situation B. Kirsch also raises a number of questions that remain poorly understood. What is the deep structure of context in a typical office? How are the various contexts of a person's work manifest in the assortment of informational artifacts: desk, papers, diaries, scribbled notes, etc? How is the cognitive capacity of remembering tasks distributed across an interactive system of a person plus artifacts? (See Malone, 1983, for an earlier expression of these issues.)

The motivation for this paper is to further our understanding of the deep contextual structure of work, and the role of informational artifacts in aiding continuity across fragmented multiple tasks. The importance of *context* for understanding multi-tasking is realized when we consider why discontinuity in surface behaviour may or may not constitute a problem. A user may have executed a contextually seamless sequence when transferring a column of figures from a spreadsheet to a word-processor, answering a telephone call from a colleague about progress with that same document, and telling them it is about to be sent it by email. In another situation, a movement from spreadsheet to word processor to telephone to email might constitute a disruptive series of interruptions in which a new context applies at each stage.

To develop this point further and to sharpen the present research questions, it is instructive to examine two recent studies. Both observed the nature and effect of interruptions in modern knowledge work. Gonzalez and Mark (2004) carried out 447 hours of direct observation of 14 employees of an IT outsourcing company providing support for share-trading. They manually recorded every event, defined as a 'continuous use of a device or engagement in an interaction with other individuals (e.g. phone conversation, using a spreadsheet with the PC ...)'. They reported that changes from one event to another occurred roughly every 3 minutes - a finding expressed as a change of 'task' every 3 minutes. Through analysis of these event patterns, the researchers noted that each event might be related to different longer running work objectives. The term 'working spheres' was coined for these longer units of work. The employees typically spent 12 minutes in one working sphere before moving on. Gonzalez and Mark have provided valuable data on working patterns, but a problem with the study is that they observed the surface structure of work and only later made the retrospective discovery of 'working spheres' - a notion close to the idea of a 'task' as conventionally defined. The interesting questions about knowledge work in this situation, but not transparent in Gonzalez and Mark's analysis, are What was the deep structure of goal-directed tasks, and How did it determine the surface structure of observed behaviour?

A second study by Czerwinski *et al* (2004) used a diary method of observation to investigate multi-tasking in the work of various different knowledge workers, including a web designer, a stock broker, an academic and others. Participants in the study completed a special spreadsheet over a period of one week, recording when they switched task and what prompted the switch. The average task duration was reported as 53 minutes with a standard deviation of 90.9 minutes. Confirming Gonzalez and Mark's findings, most task switches were *internal* (driven by the participant themselves) rather than *external* (caused by some event in the environment). Internal task switches were classified as self-initiated breaks (40%) or simply moving on to the next task (19%). External switches were identified as: telephone interruptions (14%), resuming suspended tasks (7%), and appointments (9%). Again these data are extremely valuable but they face interpretative problems as acknowledged by the researchers. Individuals were idiosyncratic in what they considered to be a task and a task switch. Some reported examples in the paper, 'Checking email, synch Pocket PC', suggest that typically it was discontinuity in surface activity that was recorded by participants. Again, there is not a clear notion of task in this study that relates to the deep structure of goals behind the participants actions. By pooling data across knowledge workers from various domains, the notion of task becomes further blurred because is likely to be idiosyncratic.

The two studies by Gonzalez and Mark (2004) and Czerwinski *et al* (2004) are valuable because they provide rich empirical pictures of the fragmented nature of modern knowledge work. However, they suffer two problems. First is the conceptual and operational ambiguity in basic terminology like 'task'. Second, the structure of work is likely to vary according to work domain; share trading is probably very different from wedding planning. Therefore, generalised findings of time spent per task are highly problematic. Smith, Hill, Long and Whitefield (1997) and Hill, Long, Smith & Whitefield (1995) attempted to solve the terminological ambiguity by articulating a conceptual framework of the 'planning and control of multiple task work', based on a more general framework of human-computer interaction (Dowell & Long, 1989). Some of the basic ideas of the framework are needed here to further our understanding of the deep contextual structure of work.

Following General Systems Theory (e.g., Boulder, 1956), we need to distinguish between a *designed system*, comprising organizations, people and technology, and an *external environment* in which work has its consequence. So, for example, in the study of a lawyer reported here, we must distinguish the designed system of a lawyer's office and its various informational artifacts, from the environment of clients, social workers, police, courts, laws and other lawyers. Herein also lies a useful distinction between the terms *activity* and *task*. When we ask what is the lawyer doing at a particular moment, we can answer from a system-perspective by describing activities (e.g., answering the telephone, reading email, organizing files) or from an environment-perspective by describing tasks (e.g., providing legal assistance for client X). The current research set out to maintain and explore this distinction to bring clarity to our understanding of multi-tasking.

To understand the deep structure of work and of multi-tasking behaviour is a large project that must be conducted on many fronts by many researchers. Investigations into different work domains are needed, and the design of systems that work in those areas. The research reported here is intended as a contribution to that larger project. It focuses on the work of a lawyer. The specific aims are to continue the questions raised by the studies discussed above for one kind of work domain. How are multiple concurrent tasks managed? What causes switches between tasks? Are internal interruptions prevalent, and if so, why are they? What is the role of informational artifacts in maintaining the continuity of fragmented work? A broader research aim is more exploratory, and asks How should we conceptualize the deep structure of discontinuous work? What are the multiple contexts that are brought to bear when performing work?

The study reported here is a naturalistic study of the work of a solicitor. In the UK, a solicitor is a lawyer who works directly with a client, and is distinguished from a barrister who specialises in making higher level court appearances on behalf of the client. This domain was chosen because there is a clear operational definition of task. Each task is *the provision of legal advice and support to a specific client* - often referred to as a case. The study presented the possibility of identifying and tracing individual tasks during observation and analysis. This

improves on the other studies described above in which analysis focused on surface discontinuities in activity, and task threads remained opaque. A second reason for choosing this domain is that it is intensely multi-tasking. At any one time, a solicitor typically represents several clients, running to hundreds in some situations depending on circumstances such as the area of law.

This study differs from the previous studies in that it observed only one participant over a period of one week. However, through observation and subsequent analysis, the study was able to trace both patterns of activity and patterns of tasks, allowing a more complete picture of the deep structure of work. Although the focus on one individual is clearly a limitation, the study is able to identify some generic characteristics of the work and to advance our approach to conceptualising its structure. The study offers general insights into the use of informational artifacts intended to maintain disparate contexts in an extreme multi-tasking situation.

METHOD

The study focused on the work of a single solicitor (S) working in the area of Family Law in the UK. The study was carried out in two phases: (i) a series of three structured interviews, and (ii) a week of direct observation by continuous shadowing.

The structured interviews were designed to give the researcher sufficient understanding to interpret events during the direct observation. Interviews established the physical context and processes of work, including the layout of offices, the role of secretarial support, the filing of cases, the handling of in-coming post and telephone messages, and the scheduling of client meetings, case reviews and court appearances.

In the observation phase, the researcher stood or sat by S for a continuous week and recorded events of the work to the nearest minute on paper. Observations recorded both tasks and activities. The operational definition of task used - of providing legal advice and support for an individual client - is shown in Table 1 along with the logically discrete sub-tasks that this entails. Table 1 also shows the types of various parties involved in each task (this aspect is discussed later). The operational definition of activity is shown in Table 2, with the list of those activities identified in the study. In subsequent analysis, a distinction was made between task-work, which progressed legal advice or assistance directly, and 'meta-work' (Kirsh, 2001) which concerned organizing when task-oriented work was to be carried out.

Therefore, noted down in the original protocol recording was a continuous sequence of the following: the name of the client for the case being worked on (e.g., JONES), the activities carried out (e.g., commence dictation of letter to Dr X, hand-correct errors in affidavit; schedule an appointment), interruptions (e.g., an incoming telephone call, or a visit by a colleague asking for advice) and other non-interrupting events (e.g. a delivery of files to S's office). Although the work was fast-paced and often involved considerable concurrency of activity and the rapid interleaving of tasks, the work was generally transparent to the researcher both at the surface level of activity, and also at the deeper level concerning task progress. Where the speed of switching was too fast to observe and note, S was asked to carry out a 'running commentary' (but not a think aloud protocol) in which client case names and the activity were called out. The few missed details and anomalies were clarified with S during breaks and in later analysis.

The only parts of S's work not observed, for reasons of confidentiality, were client interviews and some court appearances. This did not affect the findings because S did not switch tasks during interviews or court appearances and the content of individual cases was not the focus of investigation.

FINDINGS

Work processes and task management

During the structured interviews, the informational artifacts of S's work were analysed at the physical level by identifying the input and output channels, and the physical means of specifying actions to be carried out. Figure 1 shows the output channels of S's work were: producing various documents (e.g., letters, court documents, applications for legal aid, etc), communicating with other parties through face-to-face meetings (e.g., client interviews), communicating with other parties by telephone (see lower panel of Table 1), and by making court appearances on behalf of the client. Documents were produced by dictating instructions to a secretary who prepared them and then returned them to S for checking and signing before posting them out. More than one dictation tape was used to separate work of different priority. On the physical input side, S received post and telephone messages taken by the secretary, and also verbal information through face-to-face meetings. Input on these channels could be controlled by S, while two further channels of input could not be controlled and were interruptive: telephone calls and visits by colleagues or the secretary.

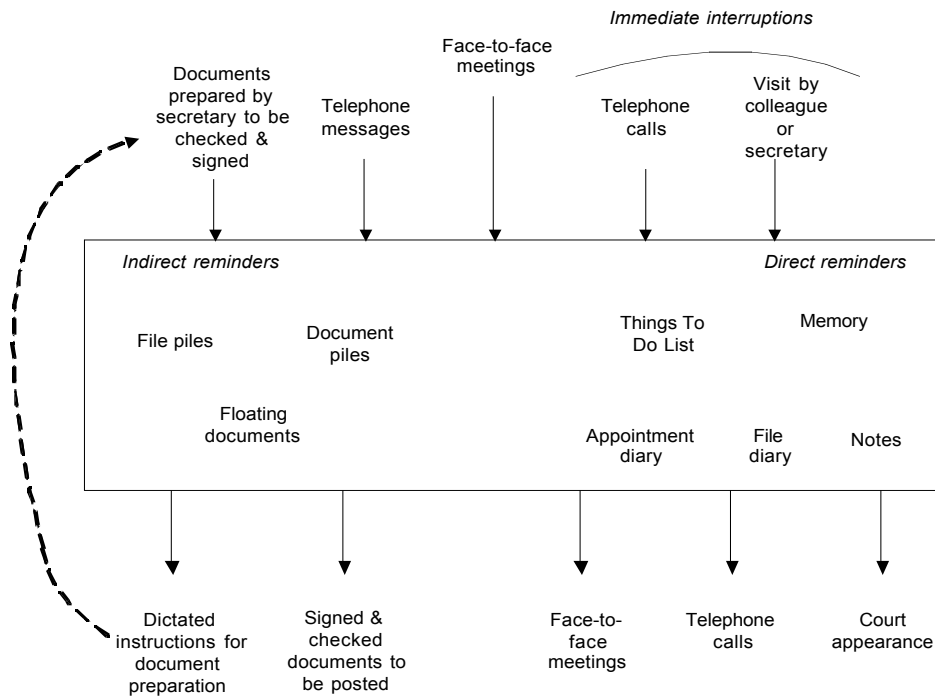


Figure 1. Physical channel of input and output and the physical representations of work to be carried out for the observed solicitor (S)

Within S's office, a variety of informational mechanisms were used to specify the work to be carried out; these also serving as reminders to preserve continuity. Direct reminders were those mechanisms that explicitly represented actions, tasks, activities or modes (explained later) to be carried out or adopted. These were: a Things To Do List (TTDL), an appointment diary (for meetings, client interviews and court appearances), a file diary (for scheduled reviews of cases), handwritten notes about urgent work, and human memory. Indirect reminders were those mechanisms that implied work to be done by the arrangement of the objects of that work. These were: the arrangement of files into piles (each file contained the documents relating to one task) in piles; the arrangement of documents in piles (these documents being temporarily out of their files); the arrangement of floating paperwork. Files and documents were organized into piles according to priority or the nature of work needed. Piles of files and documents were positioned carefully on the desk, nearby shelves and often on the floor as the work grew in volume.

The multi-tasking nature of the work

From the original protocol of observed activities and tasks, a second condensed protocol description was constructed which identified the basis on which each task switch was made (Table 7, presented later, shows a small section). A rule was used in the analysis that a switch to a task was only recorded if it was considered necessary for S to mentally recover the context of that task to carry out the behaviour needed. Thus diarising the file for a specific case to be reviewed in two weeks time was considered a task even if it took just a few seconds. However, assembly and delivery of a pile of dictated work did not count as task work if the contexts of the tasks concerned did not need to be individuated.

Over the observed week, S worked approximately 53 hours on 162 different tasks and made a total of 481 switches between tasks. Table 3 shows a breakdown of the 53 hours into the time spent in states called hard modes (explained below). The average time spent on a task before switching away, in office mode, was 3:45 minutes, ranging from less than one minute to a rare 40 mins (caused by a sequence of task-related telephone calls). More typically, 15 minutes was a long time to spend continuously on a single task.

The total amount of effort spent on each task across the week varied considerably. For some tasks it was as little as checking one letter. At the other extreme, the busiest task, as measured by total time spent, involved a client interview, several court appearances and substantial document preparation. By the nature of the work, all 162 cases were important, but they varied in their urgency. There were 12 cases that were judged as 'hot' during the week on the criterion that switches to them was often based on memory rather than other types of prompting. A

further 34 cases had 5 or more switches to them during the week which gives some indication of the level of activity.

Table 1. Tasks

Definition of task: Providing legal advice and support for an individual client		
Logical sub-tasks		
<ul style="list-style-type: none"> take instructions from client determine legal course of action prepare and distribute documents to pursue legal action represent client in court to pursue legal action 	<ul style="list-style-type: none"> obtain reports from other parties keep other parties informed (including client) coordinate the activities of parties involved keep record of actions taken 	
Typical parties involved in tasks (not always all relevant)		
<ul style="list-style-type: none"> client other-side client other-side solicitor barrister 	<ul style="list-style-type: none"> colleague client's guardian doctor client's social worker 	<ul style="list-style-type: none"> police officer court official other-side guardian other-side social worker

Table 2. Activities

Definition of activity: A type of behaviour carried out by S	
task-work	meta-work
<ul style="list-style-type: none"> listen to / speak to client face-to-face read post (official documents, letters, notes) read case papers (official documents, letters, notes) dictate documents: official documents / letters / notes write documents: official documents / letters / notes sign documents make & receive telephone calls look up/record contact details handle documents (e.g. papers in/out of file) 	<ul style="list-style-type: none"> write / read TTDL assign file to pile assign case paper in pile organize layout of piles scan piles and TTDL scan case papers diarise meeting check appointment diary dictate instructions: priority, preparation (e.g copy, form) select tape urgency level telephone other party write attendance note

Table 3. Overview of the week of observation, showing hours spent in each Hard Mode (See Table 5 for definitions of Hard Modes)

HARD MODES	Office & Circulating	Client Interviews (i)	Court (c) & Visits (v)	Breaks	TOTAL
Monday	7:18	1:36 (5 interviews)	1:49 (3 court)	0:51	11:34
Tuesday	6:09	2:25 (4 interviews)	3:58 (1 court, 1 visit)	2:04	14:36
Wednesday	-	-	8:10 (1 court)	0:50	9:00
Thursday	8:23	1:21 (3 interviews)	-	0:18	10:02
Friday	7:11	-	-	1:00	8:11
TOTAL	29:01	5 :22	13 :57	5:03	53 :23

The causes of task switching

In studies of task switching, the distinction between external and internal switches (Miyata & Norman, 1986) is often used. When interpreting the current data, it was apparent that the distinction marks the ends of a continuum. An incoming telephone call, that appears like an external interruption on the surface, may be delivering a recently requested piece of information for the current task. Similarly, what appears like a self-driven switch to a different task, may be forced by an imminent scheduled interview. Each interruption is caused by a mixture of external and internal forces. From this perspective, the following types of task switch were identified and are shown in Table 4 with their frequency of occurrence in the week of observation.

An *external interruption* was classified as an incoming telephone call or visit by a colleague that had not been recently solicited. After this, there were five types of mixed internal/external switches. An *expected interruption* was one that had been specifically and recently solicited. A *scheduled* switch was an interview or court appearance that was known about by S but nevertheless forced a discontinuity. An *opportunity* was where S was in communication with a particular party (see Table 1) and either S or the other party took the opportunity to discuss another case in which they were both involved. Also in the mixed category were switches driven by S's *memory for an impending event*, and specifically the need to prepare for it. Similarly task switches based on the *resumption* of a suddenly suspended task were considered mixed because they relied on memory, but typically occurred in conjunction with an earlier external interruption.

In the category of internally-initiated interruptions, are those that were driven by S's *memory*. S's memory was, of course, continually refreshed by a rich array of direct and indirect reminders (see Figure 1). Furthermore, task switches made through memory often involved confirmation by looking back at the associated reminder. However, it was possible to distinguish switches that were caused by this augmented human memory, from those that were directly prompted by artifacts as for example where S worked through a pile of documents, taking appropriate actions on each. In this later situation, task switches were made through direct *prompt* either by *note*, *message*, *file* or *document*. A file consisted of all the papers for one case, and a document was one of those papers. A note was a reminder hand-written by S. A message consisted of incoming post and telephone messages. A letter arriving by post was considered a message when first consulted, but then became a document once filed away as known information.

As shown in Table 4, and consistent with findings in other studies, only a small proportion of interruptions were truly external (12%). Even extending this to include the mixed external/internal category (expected interruptions, scheduled events, opportunity, memory for impending events, and resumptions) brings the total to 28.7%. The remaining 71.3% of task switches were strongly internal (by S's memory, by direct prompting by artifact).

Table 4. Reasons for making task switches, and the frequency of each type of switch

source	type	description	number observed	%
external	external interruption	a telephone call from another party, or visit by a colleague or secretary	62	12.9%
mixed external / internal	expected external interruption	an interruption that has been recently directly elicited by S	8	1.7%
	scheduled event	a scheduled client interview or court appearance	3	0.6%
	opportunity (during communication)	asking a party who is already engaged in communication about another task	25	5.2%
	S's memory for an impending event	memory for impending interview/court appearance	4	0.8%
	resumption	picked up directly following an interruption (context judged to be still active in mind)	36	7.5%
internal	S's memory (continually maintained by artifacts)	from memory, continually maintained by artifacts (diary, TTDL, document array, file array) but not directly prompted	81	16.8%
	prompt - note	prompted by a note written to self/ arranged as reminder	25	5.2%
	prompt - message	prompted by a written message usually from Secretary or colleague	5	1.0%
	prompt - file	prompted by file when going through piles	111	23.1%
	prompt - document	prompted by document when going through document pile	121	25.2%
		TOTAL =	481	

The multi-mode nature of the work

During the analysis of tasks, it was noted that S's time was organized into larger patterns, such as working in S's office, circulating around other offices, attending court, making client visits, and interviewing clients in S's office. These patterns are neither tasks nor activities as defined in Tables 1 and 2, and are described here as modes.

A *mode* is defined as a state of S which constrains the types of activity that can be carried out. Table 7 shows a small section of the second condensed protocol which illustrates the relationship between modes and tasks.

A further distinction was made between hard modes and soft modes. In a *hard mode*, such as attending court, the state is physically defined so that switches in and out of the mode are heavily constrained. These hard modes are defined in Table 5 which also shows the number of times that S switched into each hard mode. *Soft modes*, in contrast, are defined by the mental orientation of S in conjunction with a readily changeable physical set-up. For example, the soft mode of reading messages means that S takes up a pile of messages and intentionally scans through them to get an overall sense of which tasks have new information to consider, and whether any new information requires urgent attention. Table 6 shows the soft modes identified in the study and the frequency with which they occurred during the observation period. The various hard modes are mutually exclusive, as are the various soft modes, but soft modes co-exist with a particular hard mode. Nearly all of the soft mode switches occurred in the hard mode of office.

Modes have complex relationships with activities and tasks. For example, the mode of action to out-post pile involves checking and signing letters and documents prepared by the secretary. This might involve various activities: reading documents; writing; maintaining the orderliness of a document pile. Also, it can relate to various different logically defined sub-tasks, for example, taking instructions from a client; advising a client; coordinating the activities of parties involved.

It is clear from the analysis so far that there are three sorts of discontinuity relating to switches in activity, task and mode. An activity switch occurs, for example, when S moves from talking on the telephone to writing an attendance note. A task switch occurs when moving from case A to case B. And a mode switch occurs, for example, when moving from the normal office mode into a client interview or into circulating the offices.

The focus in this paper is on task switches and mode switches. A shift of one type may be aligned with a shift of another type, but often it is not. So for example, finishing a client interview was often followed by normal office work on the same case (mode shift, but no task switch). Alternatively, S sometimes returned to a more urgent case immediately following the interview (both a mode switch and a task switch).

The nature and causes of the fragmentation of work

We now return to the question of what causes the highly fragmented structure of multi-tasking work? A number of causes were identified. The most important point here was that the fragmentation of S's work was unavoidable because of the need to preserve two sorts of contextual continuity which worked against each other: *mode continuity* and *task continuity*. This created a *task-mode continuity trade-off*. For example, when S expended effort to attend court it was effective to make many court appearances relating to different cases. Similarly, when circulating in the offices, it was effective to make many visits to colleagues and secretaries about different cases rather than just one. This logic extends to soft modes. When S assumed the mental orientation of reading messages or responding to a pile of out-post, then it was effective to work through a number of different cases. In all these situations, the physical and mental costs of switching from the context of one mode to another were reduced, while the contextual cost of switching between tasks was increased. Conversely, if greater effort was made to work more continuously on individual tasks, then greater discontinuity in modes were encountered.

Having identified that some mixture of mode discontinuity and task discontinuity is unavoidable, there were further reasons evident for the high degree of internally-driven task switching. Most obviously perhaps, was *external dependency* of a task on information or confirmation by other parties, and where S had to wait until it was received. Similarly, S had to *maximize external windows of opportunity*, by sending requests for information as early as possible, or by regularly resuming some tasks where it was difficult to 'get hold of' another party. Both of these necessary techniques served to fragment tasks.

Beyond this, there was another fundamental reason for task switching which relates to the reduction of uncertainty. The point here is that the plan of action for each case was at least partly constructed by working directly on it. Hopping from task to task was not merely getting the work done, it was part of a process of *continual discovery*: finding out what work was yet to be done. Through task work, new priorities and deadlines were regularly established. Therefore it was not safe to leave some tasks unworked on, for fear that they hid some unexpected difficulty. Once priorities across the board were appreciated to a satisfactory degree, S worked by *levelling urgency*, going through the most urgent tasks and taking sufficient action to reduce their urgency to a level commensurate with other tasks. Again this led to fragmentation.

Finally, S's work structure can be interpreted as reflecting subtle context preserving techniques. Frequent task switches, as well as being needed for discovery and levelling urgency, also served to keep S mentally refreshed on the most urgent cases. This primed S to exploit opportunities for progressing tasks when communicating with relevant parties, and also to more quickly realize the significance of incoming information. Other context preserving techniques were observed. One was to prepare for upcoming interviews by switching into the relevant

Table 5. Hard Modes (showing number of switches into the each hard mode)

Definition of hard mode: A state of S that restricts the range of activities that can be carried out. In a hard mode, the state is physically defined so that switches in and out of the mode are heavily constrained.		
Hard Modes		number observed
Office	S working in own office	42
Circulating	S is moving through the other offices and workspaces, to collect post, visit secretaries, visit colleagues, etc	21
Client interview	S in own office conducting an interview with client	13
Court	S travelling to or attending court	7
Visit	S travelling to or visiting a client	1
Break	S taking a lunch, coffee or dinner break	18
		TOTAL = 98

Table 6. Soft Modes

Definition of soft mode: A state of S that restricts the range of activities that can be carried out. In a soft mode, the state is defined as a mental orientation of S in combination with a readily changeable physical configuration of artifacts.		
Soft Modes		number observed
action general	working through tasks according to priority, current context, opportunity, etc	42
action to floating documents	responding to loose documents to a tidy reminder system while also progressing work	3
action to file pile	working through a designated pile of files (one file per case) and taking appropriate action for each	12
action to out-post pile	working through a pile of documents prepared by the secretary, checking, correcting and signing them	15
read messages	reading and scanning all new messages (post and telephone)	4
evaluate & plan	evaluating the status of all tasks and setting up a plan of tasks and modes; including sorting through post, scanning all notes, rearranging files piles and document piles, revising TTDL	24
		TOTAL = 100

task in the middle of others a short time before the interview, but not immediately before. Another was to negotiate the exact point of interruptions (see McFarlane, 2002), by making visiting colleagues wait and by delaying the point of commencing a client interview.

Implications for the design of informational artifacts to support task management

We now turn briefly to the implications of this study for the design of information artifacts to support task management. Care must be taking in generalizing S's behaviour, but some points are pertinent to broader design issues. S spent less than 5% of time in the mode *evaluate & plan* which was devoted to assessing the overall picture of work. Of this time, most was spent arranging documents and files, and rarely did S consult the things-to-do-list which typically served as a back-up in case other reminders had failed. Most meta-work, including decisions about task priorities, was integrated into the work on individual tasks. Related to this point, most reminding was through indirect artifacts; that is, by the arrangement of files, documents and piles, on which action was to be taken, rather than by explicit specifications to do X, Y or Z. An exception to this was the vital use of hand-written notes about specific individual actions (e.g., 'phone Dr X'). These direct single action notes became objects that could be arranged among the indirect reminders of documents and files. S manipulated the intrusiveness of these object-based reminders (see Kirsh, 2001); that is, the extent to which they intruded into awareness and could not be overlooked. Piles of documents and files were continually arranged and re-arranged with different salience on the desk, shelves and floor. Urgent notes floated prominently on the desk, and were stuck on the telephone in the most extreme situations.

What this suggests is tools to support metawork should be integrated into the objects of work, so that they allow users to create systems of indirect reminders. As widely noted, the power of the file management interface based on windows, icons and direct manipulation, is that it supports an integration between the use of information and

Table 7. A small section of the condensed protocol showing the relationship between mode and task (labelled as cases). Hard modes are shown in UPPER CASE and soft modes in lower case.

onset	duration	mode / task (case)	switch cause
11:05	0:11	mode COURT	
11:16	0:24	case LANDON	scheduled, court appearance
11:40	0:09	mode OFFICE read messages	
11:49	0:00:30	case PETERS	prompt, message
11:49	0:04	case LANDON	prompt, message
11:53	0:00	mode OFFICE evaluate & plan	
11:53	0:05	metawork	
11:58	0:00	mode OFFICE action general	
11:58	0:01	case HILLIER	ext.interruption, solicitor-X
11:59	0:00:30	case FLAYER	opportunity, solicitor-X
11:59	0:03	case ARCHER	opportunity, solicitor-X
12:02	0:04	case LANDON	memory
12:06	0:01	case JOHN	ext.interruption, secretary-B
12:07	0:02	case LANDON	resumption
12:09	0:00	case JOHN	ext.interruption, secretary-B
12:09	0:01	case CONSTABLE	ext.interruption, colleague-S
12:10	0:05	case LANDON	resumption
12:15	0:04	case RODIN	memory
12:19	0:01	case GIBBON	memory
12:20	0:00	mode OFFICE despatch	
12:20	0:02	arrange collection	
12:22	0:00	mode OFFICE action general	
12:22	0:03	case ORANGE	prompt note
12:25	0:03	case LANDON	prompt note
12:28	0:01	case PETERS	prompt note
12:29	0:04	case FLAYER	prompt note
12:33	0:13	case HILLIER	opportunity, guardian

its organization. However, the findings here point to the limitations of typical contemporary user interfaces that separate time-based reminders (diary), actions to be taken (an electronic task list), and the objects of work (documents, email messages, etc). This problem of poor metawork-work integration, has been noted by other researchers who are currently building prototypes to propose solutions. Czerwinski *et al's* (2004) 'groupbar' attempts to provide 'project-specific' integration, and Bellotti *et al's* (2003) 'Taskmaster' provides 'task-centric' organization for emails to maintain continuity across communication-based tasks. Beyond these developments, the current study demonstrates further subtle complexity of work structure in the domain of the solicitor.

SUMMARY & CONCLUSIONS

This paper has used a distinction between task, activity and mode to understand the 'deep structure' (Kirsch, 2001) of the work of a solicitor; one type of intensely multi-tasking knowledge work. It is demonstrated that such a minimal conceptual framework is needed to account for different types of discontinuity in surface behaviour. The study showed that the widely-used distinction between external and internal interruptions (Miyata & Norman, 1986) are ends of a continuum; most interruptions are some mixture between the external forces of the task environment and the internal choice of the designed system. Consistent with the findings of other studies, only a small proportion of interruptions were *truly external* (12%), and even extending this to include mixed external/internal interruptions (*expected interruptions, scheduled events, opportunity, memory for impending events, and resumptions*) brings the total to 28.7%. The remaining 71.3% of task switches were strongly internal (by *S's memory*, by direct *prompting* by notes, documents, files or messages).

The solicitor's work pattern over the week of observation showed that the internally-driven fragmentation of work had a number of causes, but most significant was the need to preserve two types of contextual continuity: task continuity and mode continuity. Preserving the context of a mode, such as responding to an out-post pile, necessitated the fragmentation of task context because it involved working through many different legal cases; and vice versa for efforts to preserve task context. Other causes were external dependency on relevant information, the

need to maximize external windows of opportunity, the continual discovery of new priorities through carry out tasks, and the need to level the urgency of tasks.

The subject of the study used a range of informational artifacts to support task management, but made more frequent use of indirect reminders (e.g., the arrangement of objects, file, documents, and messages) than direct reminders that specified actions explicitly (e.g., things to do lists, diaries), though the latter were essential. This suggests that informational tools to support task management should allow the user to integrate metawork into normal work.

REFERENCES

- Beishon, R.J. (1969) An analysis and simulation of an operator's behaviour in controlling continuous baking ovens, in F.Bresson and M. Montmollin (eds) *The Simulation of Human Behaviour* (Duvied, Paris), 79-90.
- Bellotti, V., Ducheneaut, N., Howard, M. & Smith, I. (2003) Taking email to task: the design and evaluation of a task management centred email tool, *Proceedings of CHI 2003, Ft Lauderdale*, 345-352
- Boulding, F.H. (1956) General Systems Theory - the skeleton of science. *Management Science*, 2(3).
- Card, S.K. and Henderson, A. (1987) A multiple, virtual-workspace interface to support user task switching, *Proceedings of the ACM Human Factors in Computing Systems conference, CHI'87*, 53-59.
- Czerwinski, M., Horvitz, E. & Wilhite, S. (2004). A diary of task switching and interruptions. *Proceedings of CHI 2004, Vienna*, 175-182.
- Dowell, J. and Long, J.B. (1989) Towards a conception for an engineering discipline of human factors, *Ergonomics*, 32, 1513-1536.
- Gonzalez, V.M. & Mark, G. (2004) Constant, constant, multi-tasking craziness: managing multiple working spaces. *Proceedings of CHI 2004, Vienna*.
- Hill, B., Long, J.B., Smith, W. & Whitefield, A.D. (1995) A model of medical reception - the planning and control of multiple task work. *Applied Cognitive Psychology*, 9, 81-114.
- Kirsh, D. (2001) The context of work. *Human-Computer Interaction*, 16, 305-322.
- Malone, T. (1983) How do people organize their desks? Implications for the design of office automation systems. *ACM Transactions on Office Information Systems*, 1, 99-112.
- McFarlane, D.C. & Latorella, K.A. (2002) The scope and importance of human interruption in human-computer interaction. *Human-Computer Interaction*, 17, 1-61.
- McFarlane, D.C. (2002) Comparison of four primary methods for coordinating the interruption of people in human-computer interaction. *Human-Computer Interaction*, 17, 63-139
- Miyata, Y. and Norman, D.A. (1986) Psychological issues in support of multiple activities, In D.Norman and S.Draper (eds) *User-Centered Systems Design*, Lawrence Erlbaum Associates, Hillsdale, NJ, 266-284
- Raskin, J. (2000) *The Humane Interface: New Directions for Designing Interactive Systems*. ACM Press, Upper Saddle River NJ.
- Smith, W., Hill, B., Long, J.B. & Whitefield, A.D. (1997) A design-oriented framework for modelling the planning and control of multiple task work in secretarial office administration. *Behaviour and Information Technology*, 16(3), 161-183.
- Suchman, L.A. (1987) *Plans and Situated Actions: the Problem of Human-Machine Communication*, Cambridge University Press, Cambridge.

COPYRIGHT

Wally Smith © 2004. The authors assign to OZCHI and educational and non-profit institutions a non-exclusive licence to use this document for personal use and in courses of instruction provided that the article is used in full and this copyright statement is reproduced. The authors also grant a non-exclusive licence to OZCHI to publish this document in full in the Conference Papers and Proceedings. Those documents may be published on the World Wide Web, CD-ROM, in printed form, and on mirror sites on the World Wide Web. Any other usage is prohibited without the express permission of the authors.