

“Constant, Constant, Multi-tasking Craziness”: Managing Multiple Working Spheres

Victor M. González and Gloria Mark
School of Information and Computer Science
University of California, Irvine
Irvine, CA 92697
{vmgyg, gmark}@ics.uci.edu

ABSTRACT

Most current designs of information technology are based on the notion of supporting distinct tasks such as document production, email usage, and voice communication. In this paper we present empirical results that suggest that people organize their work in terms of much larger and thematically connected units of work. We present results of fieldwork observation of information workers in three different roles: analysts, software developers, and managers. We discovered that all of these types of workers experience a high level of discontinuity in the execution of their activities. People average about three minutes on a task and somewhat more than two minutes using any electronic tool or paper document before switching tasks. We introduce the concept of working spheres to explain the inherent way in which individuals conceptualize and organize their basic units of work. People worked in an average of ten different working spheres. Working spheres are also fragmented; people spend about 12 minutes in a working sphere before they switch to another. We argue that design of information technology needs to support people’s continual switching between working spheres.

Categories & Subject Descriptors: H.4.1. [Information Systems Applications]: Office Automation—Time Management; H.5.2. [Information Interfaces and Presentation]: User Interfaces—Theory and Methods.

General Terms: Human Factors

Keywords: Time management, attention management, information overload, interruptions, empirical study

INTRODUCTION

More than ever, information workers are facing highly demanding workloads. In some cases, with budget cuts and layoffs, companies are struggling to do the same work with less people and consequently are increasing the number of tasks and responsibilities for each employee. In other cases,

such as academia, multiple activities are just part of the inherent nature of the work: supervising students, performing departmental duties, and participating in different research projects. In both small start-up companies and large complex organizations, information work in general is typically characterized as requiring involvement in multiple projects, initiatives and teams. This is the kind of work performed by, e.g. administrators, managers, financial analysts, consultants, and accountants. To add to the complexity, workers also use a variety of digital and physical devices to conduct their work: e.g. electronic mail, instant messaging, PDAs, cell phones and paper documents.

The purpose of this study is to examine how people manage their work on multiple projects each with different goals, deadlines, and resource constraints using a variety of technologies. Yet paradoxically, most current designs of information technology support distinct tasks such as writing and editing documents, using email, or sending text or phone messages. Technology is not organized in terms of larger themes that are associated with separate projects or duties. It is left up to people to integrate their information into cohesive task structures that make sense to them.

MANAGING MULTIPLE ACTIVITIES

In the last forty years in fields as diverse as CSCW, organizational behavior, and management science, many studies have been conducted to understand how people distribute their time at the workplace. Most studies have focused on managers, describing how they spend their time both in terms of the structure and content of their activities [4, 5, 7, 12]. For a review see [10]). These studies have consistently shown how managers engage in multiple activities. In the late 1960’s, Horne found that middle managers spend most of their time in managing various activities and very little time in reflection and solitary decision-making [4]. Mintzberg found that the activities of CEOs can be characterized by their brevity, variety and fragmentation [6, 7]. Sproull, who found that managers spend 80% of their time talking with people, proposed that managerial work be considered as multitask processing, and not in terms of tasks and interruptions [12]. A recent study by Perlow [11] of software engineers found that they conceptually separate their activities into engineering work and all else. The earlier studies have looked at time distribution before IT played a major role in the workplace.

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.

Another line of studies has examined the nature of interruptions. Gillie and Broadbent [3] found that the nature and complexity of an interruption affects how much performance will be disrupted. Surprisingly, O’Connaill and Frohlich [9] found that 41% of the time people do not resume their original task after an interruption. Interruptions have been found to be beneficial as well as disruptive [5].

Taken together, these studies show high amounts of interruptions in information work. However, we maintain it is not only interruptions that tax an individual, but also the process of frequently switching activities. Though the popular and academic press has given much recent notice to the “attention economy”, there has been a lack of empirical studies to understand how people actually switch between activities and use of paper and digital devices in the workplace. To our knowledge, no study has ever focused on the role of information technology as people switch between work activities. Our study differs from these earlier studies of time distribution in that we are studying people in a very technology-rich environment. We explore to what extent work and use of tools in a technology-rich setting is fragmented and we examine the strategies that people use to maintain continuity in their work.

RESEARCH SETTING AND METHODOLOGY

In order to gain a detailed understanding of how information workers manage multiple activities we conducted an observational study at ITS¹, an investment management company located on the west coast of the U.S. ITS acts as an outsourcer, providing information technology and accounting services for a major fund manager. Given the size and extent of operations performed by their current client, ITS services only this company. More than 250 employees work at the ITS branch where we conducted the study. We concentrated our work on the day-to-day operations of one team that we call the JEB team. JEB is in charge of developing, testing and supporting major financial software modules to be used by the client. Twenty-five information workers form the team including software developers, database administrators, financial analysts, and managers.

JEB supports the client by guaranteeing at all times the transmission, validation and booking of all the financial transactions performed by brokers in the U.S., Europe and Asia. The JEB team also acts as an intermediary between the client and a financial market information company that provides financial terminals and transaction support. Records of transactions are accessed through applications developed by JEB and used by other teams within ITS, who are in charge of the accounting system and daily consolidation with banks.

Twenty-two members of the JEB team work in cubicles in an open office environment; three have their own offices. Each individual has a cubicle with a networked computer, phone

unit, and resources such as books, binders, stationary, etc. They all have commercial email and activity management tools. Six individuals have a financial terminal to monitor the status of transactions and they can perform tests for the software modules they develop. Printer and fax machines are shared and are located at the end of an aisle.

This kind of open office setting allows team members to interact and communicate easily with other colleagues even without the need to move from their own cubicles. It is common that people chat with each other through the walls, or even walk over to join conversations in other cubicles. At the same time, the height of the cubicles is high enough to provide privacy for the occupants. The employees generally concentrate on their work within the cubicle.

Methodology

The study was based on two main ethnographic techniques: participant observation and use of long interviews. The level of detail required for our research demanded that we be able to capture the details of the informant’s behavior with respect to the structure and the content of the activities that they performed. We felt that indirect observation techniques such as asking subjects to keep diaries or to generate estimates of activities at the end of the day would be less appropriate as they would be disruptive and inaccurate. Also, videotaping was not permitted by ITS. Thus we decided to use a “shadowing” observation technique similar to the ones used in previous time management studies [11, 12]. In our case, the researcher sat with the informant at her cubicle and followed her, whenever possible, to meetings or other activities. The researcher sat just behind the informant where it was possible to fully observe what she did and to some extent to be able to read documents displayed on the computer screens, the ID caller display on the phone unit, the content of print outs, sticky notes and binders on the desk, etc.

Whenever the individual performed an action such as opening a computer application, making a phone call, writing down a note on her planner or pulling a paper note from the cubicle wall, the researcher annotated the time (to the second) and other details of the event. All interactions with others were also documented, including details about the topic of the conversation, documents used and persons involved. The researcher was very careful to capture as much detail as possible. Whenever something was not clear, the researcher noted it and asked the informant at the end of the day. Inspired by Mintzberg’s structured observation method [6], we designed an activity tracking log where we transcribed the observation notes collected during the day. In the tracking log we included the time stamps, data about the type of event (e.g. “responding email to AMX”, or “modifying the Java code for CEW module), the resources using during the action (e.g. phone, Excel, planner, sticky note, calendar), and the people participating in the event. The tracking logs were used in the analysis. A total of 477 hours was spent in observation at the field site.

¹ This name, and all other names, are pseudonyms.

Fourteen people were observed over a seven-month period. Each person was observed for a period of three and a half days. The first half day was used to get familiar with their activities and their working style. For the next three days formal data collection was done and the average time of formal observation for each individual was 26 hours. Some days after the observation we conducted a two-hour semi-structured long interview where we asked about their activity management strategies and about some of the events that happened during the observation. Those interviews were tape-recorded and transcribed. We also collected documents such as email print outs, project descriptions, group calendars and software specifications. Using a grounded theory approach [13] we analyzed our data to understand how individuals spend their time, the usage of digital and physical artifacts, the different kind of interactions and how activities switch throughout the day.

Our study started with a period of ten non-consecutive days of observation of one of the managers to become familiar with the practices, projects and people related with the JEB team. His location in the office and his position in the organization's hierarchy allowed us to gain a good overview in the early stages of the study. This observation was also useful for the other team members because they gradually became used to the researcher's presence.

The employees

Among the fourteen team members observed, six of these were business analysts who were the main point of contact with the client both for specifying long-term requirements and for reporting problems during the daily operation of the systems. Four team members were developers who were more involved in the actual design and implementation of software modules. Four individuals were managers who coordinated and planned the work of other team members, as well as interacting with the client and providers.

RESULTS: THE NATURE OF INFORMATION WORK

An overview of the data

We begin by dividing our data into five main groups of events in order to compare our results with those of earlier studies [4-6, 12] (see Table 1). Deskwork, which is defined as the time that individuals work with computers, and other physical artifacts, clearly consumes a main portion of the day. Our results reveal more time spent in deskwork compared to earlier non-intensive IT environments, and is most similar to Hudson's data [5] of managers in IT-rich environments. Unscheduled meetings (going to other cubicles, people entering cubicles, or chatting through the wall) constitutes the second largest category. This is not surprising if we consider how easy is for the JEB members to walk to another cubicle or to chat from their own. On the average, our informants spent less time in formal planned meetings which is consistent with Panko's observation about the relationship between upper hierarchical role and more time spent in meetings [10]. We studied analysts, developers and managers, while the other studies focused on managers.

The Other category includes time for lunch, personal activities (e.g. getting coffee alone), and non-identified activities outside the cubicle.

In summary, our data describes individuals who spend much of their time in cubicles and have a high level of informal interaction with nearby colleagues.

	% time	Avg. time/day (s.d.)	Horne 1965 ^a	Minzberg 1970 ^a	Sproull 1984 ^a	Hudson 2002 ^b
Desk work	36.6	3:10:40 (1:22:51)	26%	22%	19%	42%
Phone ¹	5.8	0:30:22 (0:19:14)	9	6	13	
E-mail	9.2	0:47:46 (0:21:18)				
Scheduled meetings	14.4	1:14:58 (1:17:40)	10	59	34	27
Unscheduled meetings	18.9	1:38:40 (0:40:31)	55	10	34	19
Other	15.1	1:18:39 (0:34:26)		3		
Total	100%	8:41:05 (1:03:08)	100%	100%	100%	88%²

^a Pre-email study ^b Post-email study

¹ Includes time spent on cell phones

² For this study 12% of the time subjects were "to busy to respond"

Table 1: Average percentage of time spent on activities in current study compared to previous studies (hour:min:sec).

The fragmented nature of information work

Since the first days at the site we noticed that it was a very fast-paced environment with multiple conversations, telephones ringing constantly and people walking unannounced into other cubicles or calling through the walls. The high monetary value of information that the team manages plays an important role in shaping the work rhythm, as one analyst explained: "[The client] expects 100% accuracy. They don't want 99.9%. A lot of people think: 'Oh, we can fix it tomorrow'. ITS does not work like that. In this kind of industry you have to correct it right away. Every time with one issue is like a major issue. You know, every trade is a minimum of a million bucks."

We also noticed that the nature of their work was characterized by a constant switching among physical and digital artifacts as well. Their work could be described as chains of short-term events. We coded our data into events that we defined as *any continuous use of a device or engagement in an interaction with other individuals* (e.g. phone conversation, using a spreadsheet with the PC, annotating documents, or talking "through the wall"). Following Sproull [12], we considered that in any particular event neither the structure nor the content changes. Table 2 shows the average time spent on an event per person, per

day, for all three roles combined (analyst, developer, managers). This time reflects the amount of time that people spent in *continuous uninterrupted work* on events.

Events	% entire day	Avg. Time/Day (sd)	Avg. Time/Event (sd)
Using phone ¹	5.83	0:30:22 (0:19:14)	0:02:25 (0:00:42)
Using email	9.17	0:47:46 (0:21:18)	0:02:22 (0:00:27)
Using PCs ²	29.48	2:33:36 (1:11:23)	0:02:53 (0:01:10)
Using paper documents/books	6.80	0:35:25 (0:29:48)	0:01:47 (0:00:31)
Using other tools ³	0.31	0:01:38 (0:03:08)	0:01:04 (0:00:15)
Talking through the walls	2.94	0:15:18 (0:14:12)	0:01:40 (0:00:24)
Interacting with people in their own cubicle	6.88	0:35:53 (0:29:25)	0:03:34 (0:01:57)
Formal meetings	14.39	1:14:58 (1:17:40)	0:41:47 (0:12:46)
Going to other cubicles	9.11	0:47:29 (0:27:21)	0:07:37 (0:03:24)
Other (unknown, personal)	15.09	1:18:39 (0:34:26)	0:17:27 (0:06:27)
All events except "Formal meetings" and "Other"	70.52%	0:45:56 (0:52:03)	0:03:08 (0:02:27)
All events total	100%	0:52:07 (0:55:25)	0:08:55 (0:13:23)

¹ Includes time spent on cell phones

² Includes both PCs and financial terminals – does not include email.

³ 'Other tools' include: handheld calculator, planners, and address books

Table 2: Average continuous time spent on events before switching (hour:min:sec).

What was most surprising to us was that people spend on the average slightly over three minutes on an event (leaving out formal meetings, unknown, and personal events) before another event is initiated. An ANOVA between analysts, developers, and managers shows no significant difference for all but two of the events. We found a trend that developers participate in shorter formal meetings (mean=30 min. 16 sec., sd=5 min. 54 seconds) compared with analysts and managers (combined mean=46 min. 23 sec., sd=11 min. 54 sec.), $F(2,13)=3.4$, $p<0.07$. A Tukey post-hoc test (95% confidence interval) however, shows no significant differences. Also we found that the three roles differ significantly in the continuous time spent using the computer, $F(2,13)=6.74$, $p<0.01$. A Tukey post-hoc test (95% confidence interval) showed that developers spend significantly more time using a PC (mean=4 min. 0 sec., sd=29 sec.) compared with managers (mean=1 min. 58 sec., sd=46 sec.). Four minutes on the PC before being interrupted or switching events is still quite a short period of time for developers.

People spend an average of less than two and a half minutes reading email before they switch to another event, or are interrupted. The data also show, that in fact, people spend

very little time in one continuous stretch when they engage in any kind of informal interaction with another person (combined mean=4 min. 29 sec., sd= 3 min. 23 sec.).

Device	% entire day	% of device usage only	Avg. Time/Day (sd)	Avg. Time/Event (sd)
PC ¹	37.01	72.37	3:12:52 (1:13:48)	0:02:52 (0:00:51)
Financial terminals ²	1.64	3.19	0:16:59 (0:13:13)	0:01:20 (0:00:36)
Paper documents and formats	5.01	8.92	0:26:06 (0:22:21)	0:01:33 (0:00:28)
Books, manual and other references	1.79	3.50	0:09:20 (0:12:16)	0:01:57 (0:00:55)
Hand-held calculator ³	0.05	0.10	0:01:13 (0:01:29)	0:00:48 (0:00:18)
Daily-Monthly planner (paper) ³	0.19	0.38	0:04:40 (0:04:18)	0:00:50 (0:00:15)
Address books (paper) ³	0.07	0.14	0:01:45 (0:03:04)	0:01:00 (0:00:42)
Phone unit	5.16	10.08	0:26:52 (0:18:23)	0:02:17 (0:00:43)
Cell Phone	0.67	1.31	0:04:53 (0:06:06)	0:04:13 (0:04:24)
All devices⁴	51.59%	100%	0:44:57 (1:13:27)	0:02:11 (0:01:52)

¹ Includes using email ² Only seven informants have terminals.

³ Only three informants used each of these. ⁴ Weighted average

Table 3: Average time usage per device per person (hour:min:sec).

Looking closer at the usage of the different information artifacts, we can clearly see how a short-term duration usage pattern prevails across different technologies. Table 3 shows the length of time that people spent using different electronic devices and paper documents before they were interrupted or switched to another activity. People spent an average of 2 min. 11 sec. working with any device or paper before they switched to another device or event.

An ANOVA between the three different roles shows a significant difference in phone use ($F(2,13)=7.09$, $p<0.01$). A Tukey post-hoc analysis (95% confidence interval) shows that developers (mean=10 min. 13 sec., sd=3 min. 55 sec.) speak significantly less time on the phone per day than analysts or managers (combined mean=33 min. 32 sec., sd=11 min. 27 sec.). With the exception of PC use described in Table 2, no other significant differences between roles were found.

Working Spheres: a practical unit of work

Our data so far conveys the very short-term nature of interactions with people and devices, and shows how people constantly switch between different events. In formal interviews and informal comments, the informants confirmed this behavior as typical in their day and pointed out that it does not always mean it is detrimental for their work. Constant switching at the level of events is not challenging

per se; it is the switching at higher levels of activity that we envision can be problematic.

We decided to examine how often people switch between what we consider to be higher levels of activity, such as projects, of which events are a subset. Observations gave us the opportunity to understand the different ways in which individuals define, delimit and differentiate their activities. Many times during observations, comments arose about their working on “the TRK stuff” or the “the R6 spec” or the “Clear Quest app”, or “Bill is over the Commercial Paper thing”. We analyzed our data (photocopies, email printouts, documents, pictures, observation notes, tracking logs, and transcripts). This led us to emerge with the grounded concept of *working sphere* which describes higher levels of units of work or activities that people divide their work into on a daily basis.

Type of WS	Condition	Average WS/day (sd)	Avg. Time/ WS (sd)	Avg. Time / segment (sd)
Central	Default	5.00 (1.90)	0:46:57 (1:02:40)	0:12:52 (0:11:41)
	Urgent	1.00 (1.15)	0:39:38 (0:45:45)	0:15:42 (0:11:01)
Peripheral	Default	2.90 (1.76)	0:15:10 (0:22:02)	0:08:53 (0:09:26)
	Urgent	0.90 (1.10)	0:11:40 (0:11:03)	0:07:16 (0:07:14)
Metawork	--	1 (0.0)	0:44:29 (0:29:34)	0:06:22 (0:03:59)
Personal	--	1 (0.0)	1:05:13 (0:32:48)	0:48:38 (0:30:13)
Unknown	--	1 (0.0)	1:22:24 (0:55:31)	0:12:00 (0:10:06)
Central and Peripheral WS only		9.81 (3.39)	0:33:32 (0:50:59)	0:11:28 (0:10:54)
All	All	12.81 (3.39)	0:40:41 (0:50:58)	0:13:49 (0:16:14)

Table 4. Avg. no. of working spheres (WS) per person, and avg. time spent in each WS per day (hour:min:sec).

We define a working sphere as a set of interrelated events, which share a common motive (or goal), involves the communication or interaction with a particular constellation of people, uses unique resources and has its own individual time framework. With respect to tools, each working sphere might use different documents, reference materials, software, or hardware. It is the whole web of motives, people, resources, and tools that distinguishes it from other working spheres. Examples of working spheres of the JEB team include a training effort for new UNIX programmers, an implementation of a new feature in a compliance module, a trip to a regional Park for JEB teammates, or the documentation of modules.

We characterized working spheres in two main ways: with respect to the importance or centrality that they hold for the individual, and with respect to the way that they are handled.

Working spheres can be shared with others which means that the same sphere can have different importance for different people. We identify two main levels of engagement with a working sphere: central and peripheral. A working sphere is central when it is of primary importance for an individual; otherwise it is peripheral. Working spheres are also handled in two basic ways: as urgent cases that have to be solved right away or as default cases that involve expected work with expected time schedules. The same working sphere can vary at times between being urgent (when something goes wrong) or involving its expected state of affairs. We identified and counted the number of working spheres in which each of our informants was involved in during the three days of observation and measured the length of time they spent in each. Table 4 shows the results.

We identified that each person worked on an average of ten working spheres per day, during the three days of our observation. (Note that we observed a three-day “slice of time”, so it is possible that we did not count all the working spheres that an individual might have). Table 4 shows that individuals clearly spend more time in working spheres which are central to their work. Work in peripheral working spheres generally involves answering questions or providing feedback. An individual spent about 33 min. 32 sec. per working sphere; however it does not occur as a continuous period of time. Instead the individual attends to a working sphere in small periods of time that we call *segments*. People frequently switch from one working sphere segment to another. Table 4 shows that the actual average duration of a working sphere segment is quite short (11 min, 28 sec.). An ANOVA between analysts, developers, and managers showed no significant differences in working sphere segment length.

Condition	Type of W.S.	Av. Time p/ segment (S.D.)
Central	Default	0:14:13 (0:12:55)
	Urgent	0:16:24 (0:12:25)
Peripheral	Default	0:09:03 (0:09:38)
	Urgent	0:07:41 (0:07:30)
Both	All	0:12:18 (0:11:53)

Table 5. Avg. time per WS segment without “nonsignificant” disruptions (hour:min:sec).

Metawork

Individuals spend part of their day on a set of activities that is not connected with any specific working sphere but rather related to the management of all of them. We call these activities *metawork*. People periodically conduct metawork throughout the day, which involves coordination, checking activities, organizing email, organizing their desk at the start or end of a working day, and catching up with teammates on what they have missed. People spend an average of 44 1/2 min. per day conducting metawork, and similar to working

spheres, this work is also conducted in shorter chunks of about six and a half minutes at any one time (see Table 4). Generally, we found that individuals engage in metawork whenever they conclude large activities, or when they return from a meeting.

Working spheres without “nonsignificant” disruptions

We noticed that our informants, sometimes while working in one working sphere, were interrupted by work from another working sphere, switched to the second sphere briefly, and then resumed work in the first sphere. Examples of this can be when somebody brings a document to be signed, or when they get a quick call. We realized that some disruptions are not significant and would not introduce a large overhead to resume work. We conducted a further analysis of our data where we explored the effects of disregarding short distractions on the segment length. We used a criteria of two minutes after reviewing the data and decided that this would make a feasible heuristic to begin with. We then coded the data also considering how much attention the disruption required. For example, a disruption from a central working sphere required more concentration than from a peripheral working sphere. Thus, a segment was considered as “continuous” even if people turned to another working sphere for less than two minutes. Table 5 shows that even after removing what we consider “nonsignificant” disruptions, people still spent only a short average time in a working sphere segment (somewhat more than 12 minutes).

External and internal interruptions

People can switch to a different working sphere or they can be interrupted. About two-thirds of the time people resume work in their working sphere after interruption. Most “significant” interruptions can range from two minutes up to 40 minutes. About one-third of the time people switched to another working sphere without returning to the first working sphere. In some cases, for example, when deadlines are similar, people move back and forth between working spheres. In general, interruptions could be categorized into two main groups: Internal and External, following [8]. An Internal interruption refers to self-initiated switching among working spheres. An External interruption is a condition in the environment that motivates switching. In Table 6 we can see some of the actions that resulted from internal interruptions (e.g. making a phone call or leaving the cubicle) as well as kinds of external interruptions (e.g. a person enters the cubicle).

Our data confirms previous studies indicating that people interrupt themselves as often as they are interrupted [12]. Our informants continually switched on their own volition to another working sphere and most internal interruptions were due to people leaving their cubicle to interact with other individuals. With external interruptions, individuals switch between working spheres more often due to verbal-based interruptions (such as visitors or phone calls) than to notification mechanisms from their e-mail or voice mail.

	Type	Average Interruptions per day (S.D)	% all types	Internal / External
Internal	Checking/Using Paper Docs	0.52 (0.86)	1.87	49.11%
	Checking/Using Computer	1.54 (1.47)	10.98	
	Talking t/wall	1.93 (2.15)	6.89	
	Phone call	1.14 (1.56)	4.09	
	Email use	1.04 (1.47)	7.40	
	Leaves cubicle	5.00 (2.56)	17.87%	
External	New email notif.	3.55 (3.18)	12.68%	50.89%
	Person arrives	6.00 (3.03)	21.45%	
	Status on terminals	0.36 (0.82)	1.28%	
	Phone ringing	2.62 (2.01)	9.36%	
	Voice message light	0.19 (0.45)	0.68%	
	Call through wall	1.33 (1.75)	4.77%	
	Reminder notification	0.19 (0.40)	0.68%	
Total		25.40 (8.23)	100%	100%

Table 6. Avg. number and types of interruptions per day.

Strategies for maintaining continuity in working spheres

People employ strategies to help them maintain continuity as they continually switch between their working spheres. Most of our informants commented that their preference is to work in a single working sphere until the job is completed. However, this is rarely the case because interruptions (internal or external) lead people to switch their attention into different working spheres or urgent issues arise that lead people to switch working spheres. People develop strategies to adjust to the unpredictability of their environment, such as knowing they will need to respond to urgent requests.

Because frequent interruptions are expected we noticed that some of our subjects use special artifacts that help them to prioritize and maintain their attention over their working spheres. These artifacts function like containers in that they hold information about central working spheres. The information included in the artifact plays the role of a reminder and, as pointed out by Miyata and Norman [8], it both *signals* the working sphere to be attended to and *describes* with some detail what has to be remembered. The artifact is often updated across the day with results when work within a sphere has to be postponed.

The artifact takes on different forms, each providing different properties in terms of the signaling or descriptive power. A common implementation uses a special inbox folder in the email client to contain messages related to central working spheres to be attended to. Messages are from others or from the individual herself with summaries of other messages or other documents. A second form of the artifact used by three of our informants uses printouts of email messages. People keep piles on their desks with printouts of email messages or meeting notices generated by their team’s scheduling system.

On those printouts are annotations such as clarification notes or contact information. A third form of artifact relies on more traditional activity management tools. Two of our informants used planners extensively to manage their working spheres usually by integrating conversations or email messages into the planner. For each day they listed the central working spheres to cover and if necessary they transferred pending actions. A fourth form of artifact is the post-it note. One informant places up to seven notes with references to actions to be taken in different working spheres. The notes stay on the monitor shelf for as long as the working sphere lasts and are annotated to indicate the status of the working sphere. For instance, while working on a software update for one of his colleagues, an informant wrote down comments regarding progress in that working sphere (e.g. “Test PB patch DAN’s PC –Waiting for AL”).

DISCUSSION AND CONCLUSIONS

Our study confirms what many of our colleagues and ourselves have been informally observing for some time: that information work is very fragmented. What surprised us was exactly how fragmented the work is. In a typical day, we found that people spend an average of three minutes working on any single event before switching to another event. Informal interactions average four and a half minutes each. Further, people spend on the average somewhat more than two minutes on any use of electronic tool, application, or paper document before they switch to use another tool. The longest duration of tool use is with PCs, yet this averages only slightly more than three minutes at any one time. One informant’s quote used in the title of this paper captures the extent of this fragmentation of work.

We argue that it makes more sense to understand how time is distributed among *working spheres*, activities that are thematically connected for the individual. A working sphere consists of a string of events. We found that working spheres are also highly fragmented: people spend on the average eleven and a half minutes in continuous work on a project or theme before they switch to another. After removing what we considered to be “nonsignificant” disruptions, we found that these segments of time were not much longer, averaging somewhat longer than 12 minutes.

People interrupt their work themselves (internal interruptions) about as much as they are interrupted by external influences. Most interruptions are due to face-to-face interactions, similar to what O’Conaill and Frohlich [9] found. People are just as likely to interrupt themselves as to be interrupted by external sources. An interesting area for future research is to understand the reasons that lead information workers such as these to interrupt themselves so frequently.

We examined the work of people in three different roles: analysts, developers, and managers and found (except for PC use) that no significant differences existed in the fragmentation of their work. Our work expands on the past studies of time distribution [4,5,7,11,12] that looked only at

managers’ work. We discovered that switching between multiple spheres of work is more pervasive.

Our results indicate that throughout their day, individuals are constantly moving from one topic to another and managing information streams from a myriad of sources. Our results build upon Sproull’s claim [12] that managers are multi-task processors. We provide empirical data to show that people mostly spend very little continuous time on any single activity. We found that our informants’ work can be described more specifically not as multi-task processing, but as requiring attentional resources to constantly change between different events, tools, and working spheres.

The three kinds of strategies that we found to manage different working spheres share some characteristics. First, they result from an explicit effort to help organize the work that has to be done. Information related to different working spheres is aggregated into a single type of artifact. Second, the prioritization of working spheres is not something that is explicitly indicated in the artifacts. The order in which information is listed or organized does not express an order of execution or importance. People prefer to talk of those items contained in the artifact as important items that need to be done. A third characteristic is that the artifacts always appear in a visible spot of the working space so they can be consulted constantly. For printouts and post-it notes, visibility is achieved by placing them within the normal field of view (e.g. nearby the keyboard). The inbox folders are located where they can be seen whenever email is used. For instance one of our subjects labeled his inbox “Z Immediate Attention Issues” so he would have it at the end of all his alphabetically-sorted inboxes. A final important characteristic is that individuals leave the artifacts only while the working sphere is active. When the working sphere is finished, items are transferred to other archives.

Implications for information technology design

Current information technology is designed to support individual events such as word processing or e-mail use rather than to provide mechanisms to integrate the multiple information objects required by some working spheres. The design of information technology needs to consider how information workers switch constantly among working spheres. Mechanisms should be flexible to enable people to group particular documents and applications but at the same time there should be recognition that many applications are shared among working spheres (e.g. an e-mail client or scheduling tools). In particular, when working spheres use very different information resources and applications, it will be convenient to have mechanisms that save the state of the information device particular to that working sphere, making it easier to resume work. In contrast with a similar idea proposed by Bannon [1], we argue that we need to re-examine the role of artifacts such as notepads, post-its or planners and derive hybrid designs which include computerized and non-computerized tools.

We noticed that in spite of having a commercial task management tool, our informants did not use its functionality to support their work. Perhaps the most important reason is the fact that they need to have their list of current tasks always visible, which is not possible with the tool because the window with their current activity is usually maximized. The artifacts developed by the JEB team members are useful because they are visible and available. To continually show the working sphere status requires extending an interface beyond that provided by the personal computer and points toward designs that include peripheral independent displays (or PDAs), which through wireless connections can provide portability. People also prefer more open and flexible forms of managing their activities which fit well with their forms of communication. Given the preponderance of email communication, the integration of reminder artifacts with email should be achieved very easily. Systems such as Taskmaster [2] (and the concept of *thrask*) are in line with the kind of support required for working sphere management. However, it is important to recognize that not all information related to a working sphere is delivered by email and consequently a system must include ways to integrate information from other sources.

Limitations of the study and generalizability of results

Our study has several limitations. The challenge of the “shadowing” technique is that the researcher herself cannot be a source of interruptions. To avoid this we kept our questions for the end of the day, or when walking with informants to attend meetings. In spite of our best efforts, 15.81% of the events could not be matched with any working sphere and were consequently put in an Unknown category. We believe that given the nature of the observations and the fast-pace environment we observed, the percentage of events that we clearly identified within working spheres is acceptable to support our findings.

Our observations are limited to one fieldsite. This is also true of other studies of time distribution [4,5,7,10,11,12]. It is very possible that factors unique to the organization affect the constant shifting, such as time pressures, and the high value of the product. We would need to investigate other organizations to understand more completely how organizational factors affect the management of multiple activities. We also only observed 14 people. This is a higher number observed than Hudson [5] and Sproull [12]. Perlow [11] observed 17 software developers. We therefore cannot claim that our sample represents a wide range of information workers. We believe though that the kind of fast-pace environment that we studied is widespread. Yet more research is needed to investigate the work roles and kinds of organizations where such work is typical.

ACKNOWLEDGEMENTS

We would like to thank the members of the JEB team at ITS, and Susan Knight. This research was supported by the National Science Foundation under grant no. 0093496, by the Center for Research on Information Technology and Organizations (CRITO), and UC MEXUS grant no. 32080.

REFERENCES

1. Bannon, L., Cypher, A., Greenspan, S., and Monty, M., Evaluation and analysis of users' activity organization. *Proceedings of CHI 1983*, ACM Press (1983), 54-57.
2. Bellotti, V., Ducheneaut, N., Howard, M., and Smith, I., Taking email to task: The design and evaluation of a task management centered email tool. *Proceedings of CHI 2003*, ACM Press (2003), 345-352.
3. Gillie, T. and Broadbent D., What makes interruptions disruptive? A study of length, similarity and complexity. *Psychological Research*, 50, (1989), 243-250.
4. Horne, J.H. and Lupton, T., The work activities of 'middle' managers - An exploratory study. *The Journal of Management Studies*, 2, (1965), 14-33.
5. Hudson, J.M., Christensen, J., Kellogg, W.A. and Erickson, T. "I'd be overwhelmed, but it's just one more thing to do:" Availability and interruption in research management. *Proceedings of CHI 2002*, ACM Press (2002), 97-104.
6. Mintzberg, H., Structured observation as a method to study managerial work. *The Journal of Management Studies*, 7, (1970), 87-104.
7. Mintzberg, H., *The Nature of Managerial Work*. Prentice Hall, Englewood Cliffs N.J. USA, 1973.
8. Miyata, Y. and Norman, D.A., Psychological issues in support of multiple activities, in *User Centered System Design*, D.A. Norman and S.W. Draper (Eds). Lawrence Erlbaum, Hillsdale, N.J., USA (1986), 265-284.
9. O'Conaill, B. and Frohlich, D. Timespace in the workplace: Dealing with interruptions. *Proceedings of CHI 95*, ACM Press (1995), 262-263.
10. Panko, R., Managerial communication patterns. *Journal of Organizational Computing*, 2,1, (1992), 95-122.
11. Perlow, L.A., The time famine: Toward a sociology of work time. *Administrative Science Quarterly*, 44, (1999), 57-81.
12. Sproull, L.S., The nature of managerial attention. *Advances in Information Processing in Organizations*, 1, (1984), 9-27.
13. Strauss, A. and Corbin J., *Basics of Qualitative Research*, Sage, Thousand Oaks, CA. USA, 1988.