

Multiple Carets, Multiple Screens and Multi-Tasking: New Behaviours with Multiple Computers

Russell Beale

Advanced Interaction Group
School of Computer Science
University of Birmingham, Edgbaston
Birmingham, B15 2TT, UK

R.Beale@cs.bham.ac.uk

William Edmondson

Advanced Interaction Group
School of Computer Science
University of Birmingham, Edgbaston
Birmingham, B15 2TT, UK

W.H.Edmondson@cs.bham.ac.uk

ABSTRACT

This study presents interview based case studies of users who work with multiple computers as well as multiple displays. Such users have not been studied before. The behaviour is discussed in terms of both technical and cognitive dimensions, and we identify the importance of having multiple carets and the complexity of multi-tasking and how it can be supported across multiple machines in a way not possible on a single system.

Categories and Subject Descriptors

H.5.2 [Information Interfaces and Presentation (e.g. HCI)]: User Interfaces; H.1.2 [Models and Principles]: User/Machine Systems.

General Terms

Human Factors.

Author Keywords

Multiple displays, multiple computers, cognition, task switching, multi-tasking, caret, design.

1. INTRODUCTION

As computing technology matures, it has become much more widely available and more cheaply priced, which has helped its penetration into many different aspects of our lives. Both monitors and computers themselves are now much cheaper and of higher specification than they were only a few years ago. Because of this, price and availability are not significant barriers for users, who are moving towards computing configurations that suit their way of working or playing, rather than being restricted to just one machine.

There have been a number of studies on the concurrent use of multiple displays in HCI, showing that they are used for a variety of purposes, from the common one of extending the visible workspace (ideal for space-hungry applications such as video editing or photo editing) through to their use to segment the tasks of the user, allowing them to maintain an overview on their concurrent tasks such as email, instant messaging, web

browsing and document preparation [2, 5, 7]. Grudin used structured interviews and found some interesting uses and usability issues with multiple monitors [4]. For example, he noted that although tasks were distributed or allocated to monitors, the applications being used did not well support 'multimon' configurations. He also found that even very small displays, such as those found on PDAs, could serve as additional displays attached to the main monitor. Ringel compared virtual desktops with multiple monitors, using a similar number of interviewed subjects, and found personal preferences to be a strong component in accounts offered by subjects [13], whilst Henderson and Card [6] discuss the Rooms metaphor as a way of reducing contention for what once was scarce window space. However, these studies have been undertaken on systems in which all the displays are connected to a single computer.

Work on multiple computers tends to come from the ubiquitous computing community, started by Weiser's vision [16], but does not really focus on the use of multiple computers at the same time within the same (usually an office) environment, though the CSCW community has contributed to this field [14, 15]. Miller and Myers [9] discuss an approach to synchronizing the clipboards of multiple computers, whilst Rekimoto [12] presents 'pick-and-drop' to allow interaction between multiple computers.

We are noticing users who routinely work with multiple displays attached to multiple computers, and because this work pattern is not well covered in reported studies we are interested to understand any differences that may exist in this setup, from cognitive, interactive and technical perspectives. We are focusing on multiple computer use in this study, where the computers are considered to be desktop or laptop machines, though we do recognize that there are often multiple electronic tools in parallel use, such as mobile phones, PDAs, and related devices, that may all interact and provide alternative patterns of usage. However, by focusing on a specific aspect, we can obtain a more detailed understanding of one area.

This paper presents the results of a preliminary study of the phenomena. Five subjects are studied in detail, referred to as Users A-E. The study was carried out primarily by observing the users in situ, discussing with them their patterns of usage and their attitudes and approaches to their multiple machines, and getting them to reflect on their activities, using open questions and asking for clarification of their answers. One subject, whose set-up was not available for the authors to visit, was interviewed on the basis of photographic evidence of his system configuration. This ethnographic approach is common in HCI [1, 3, 11], though is often used in a very descriptive manner. We agree with Dourish [10] that ethnography should not only have the goal of presenting implications for design, and so we present the detail of the study to provide an in-depth account of an

alternative usage that has not been previously reported; however, we also consider that the insights developed here do provide some useful design pointers. The paper presents the user studies in detail, draws out our explanations and understanding of the user behaviours and understanding, and then identifies some key implications for design.

2. CASE STUDIES

One of the clear outcomes of this work is that there is a spectrum of user behaviours and approaches to using multiple machines: from those that closely mirror single computer use to those that are significantly different. We have organized the users in approximate order across this spectrum from the most different and complex to the simplest. However, we have retained their labeling in terms of original interviews, and to facilitate comparison with other results, as and when reported (this does mean that they are not presented in alphabetical order).

2.1 User A

User A is a university lecturer, with a wide range of responsibilities and activities. His office environment has a large number of shelves, filled with books and journals. Despite the predominance of paper and books within the office, A considers most of his work to be done on the computer, whether it is teaching preparation, administration, or research. A's computer setup comprises four machines with six screens. From left to right in the photographic montage shown below, these are: a Mac G5 with 23" screen and 3TB of field data; a G4 laptop with screen (and a separate wired keyboard and mouse), a Samsung monitor attached to the laptop, a 23" Mac display attached to a Mac G4, with a keyboard/mouse along with an older Mac display attached to the same machine; and lastly a PC laptop and screen on the desk behind the swivel chair. Thus there are two laptops and two desktop machines (see Figure 1). All of the machines are physically networked onto the department's Ethernet network, and so are visible to each other, and to the central filestores. However, all of the machines are also wireless enabled. The Macs are MacOS X 10.4 and the PC has XP installed.

A was studied in situ on a number of different occasions over about a month; the first session was of about an hour's duration,

and was mainly observational, whilst the subsequent ones were for similar durations but involved much more discussion and questioning. Additionally, the usage and setup was discussed outside the office on numerous other occasions, and a number of short office visits also contributed to the study.

A's perception is that the tasks he undertakes are mapped onto the computers, and not so much onto the displays. In fact, it would be more correct to state that the tasks are most strongly related to the 4 keyboards that are on the desktop. Apart from the PC (with its integrated keyboard and a wired mouse) the other machines have wired keyboards and mice.

When asked to describe how his work is mapped on to the different machines, A presents a perspective on his work that partitions tasks up into specific activities: admissions, research, writing, email, and these are allocated to specific screens. However, the need for a number of computers is, at this stage, unclear – would not the same effects be visible if we have the same number of screens, all attached to one computer? We were interested to understand if there are any differences between multiple displays attached to one computer or having multiple computers.

The single computer scenario is not possible for practical reasons, since two operating systems are in use. The PC laptop is used for only one task: student admissions. The University runs a central admission system that is only accessible via a PC client, and hence the laptop supports that task. All discussions with A showed that he felt that the role of the laptop was for admissions, and felt happy with a cognitive allocation of admissions to that machine. Of interest is the fact that A went through a phase of trying to get reliable PC emulation to work on the Mac G4 so as to avoid the use of another machine. Problems with running the university's admissions software in such an environment led to A's reluctant acceptance of yet another machine into his office. Thus the proliferation of machines for A was not something desired for its own sake. The sense for A that the PC laptop is the 'Admissions' machine is thus a cognitive allocation in response to the machine's presence.

The admissions process has a strong paper-based element, and much of the actual work on admissions requires reading and understanding the forms, and recording the outcomes on the system. It has been proposed in the University that the UCAS



Figure 1. A's set-up: Six screens, four computers, one user.

forms will be distributed electronically, and when this happens the 'Admissions' PC will be needed for that work also. What is more interesting is that communications from the department's other admissions people, and from the centre, that arrive by email, are actually processed on the G4 Mac (in fact, on the older Mac monitor). This display is used primarily for processing email, whether admissions or other messages. It is the case that the emails are fundamentally meta-information: information about applications, or questions from students, but in all of our conversations regarding where the work is done for admissions, whilst A will discuss happily the emails received on the issue, when asked to identify the computers used will only present the laptop. Cognitively, he only associates the admissions task with the one machine, despite often undertaking email activities related to admissions on the other machine. When questioned on this apparent inconsistency in categorization of the PC as the 'Admissions' machine A recalls that once he was obliged to have the machine he set out to do all the emails for admissions on the PC as well. The department's email service is available via IMAP – so all the computers A uses can handle email on just one account. The email client on the PC was set up by A to have a signature referring to his role as Admissions Tutor (and A made such a signature optionally available on the G4 machine which he characterizes as the main email machine). But despite being able to access his incoming email, and to send appropriately signed outgoing email, A used his 'main' email machine instead.

A's account of the PC is simplified or 'idealized' and actually concerned with intentions rather than reality – and this is revealed in a curious way. The admissions work in his department makes extensive use of MS Access; it serves as a CRM system as well as for monitoring progress with applications paperwork and decision making. The fact that this is not available for the Mac OS X environment could be used by A to strengthen the case for calling the PC laptop the 'Admissions' machine, and thus to reinforce the cognitive partitioning of work onto machines. However, in relating the account of the usage of the machines A does not dwell on Access at all. It is clear that the cognitive partitioning is in some ways quite superficial in A's account, being linked only to one PC-specific application, though clearly it is deeper than this: there are two applications, both on the PC, that are only used for admissions and nothing else – yet A does not report this. Thus the reality of the partition of work between machines is less clear cut than A presents to himself and to others.

A's discussion of the other machines is also idealized. For example, A says of the Mac laptop and attached screen that he uses this setup for teaching work. For example, the most common usage scenario observed, and reported, is the task of lecture preparation; specifically, when creating PowerPoint slides the laptop is used to provide the overview and notes, whilst the individual slides are displayed on the Samsung monitor because of its conventional 4:3 aspect ratio. On further questioning it transpires that the portability of this machine (with its preferred operating system) is important – he takes it to lecture rooms around the University to deliver the lectures, and stores the material on this laptop. But A also takes it to meetings to make notes and to distribute those notes to others directly from the meeting room (using the wireless capability). In addition, there are aspects of teaching work (not lecture preparation) which A does on the G4, such as emailing students and colleagues, and preparing web materials for upload to the School's servers. Indeed, it transpires that despite the fact that all the machines are linked to the School's servers A considers the G4 to be the primary interface to the servers (possibly

because this setup has the most screen estate) and thus provides for effective browsing and emailing simultaneously.

A's characterization of the G5 as a research machine is the least ambiguous. The machine has 3TB of JBOD store with data gathered in the field. This machine is dedicated to the data processing – but has email facilities which are used for research focussed email, for example when sending graphical output from Matlab on the G5 to colleagues. However, A has been known to mount the G5 on the G4 as a volume, in order for graphs to be spread out over three screens for visual inspection.

A's configuration makes the functional separation of the machines very obvious. They have the potential for interlinking via Ethernet, and they all connect to the IMAP mail service so they can share email accounts. Indeed, A says that on occasion he has found it easier to move a file from one machine to another by means of attaching a document to a mail to himself, and then saving the attachment on the other machine (the alternative is ftp which simply takes longer to initiate). This practice is simply an extension of the method he uses to move files between his office and his home set-up (and vice versa).

One of the key features of the set-up is the arrangement of chairs in the room. By choosing a particular chair, A chooses a particular machine, and hence, cognitively for him, a main focus of working.

What transpires from this account is that A's characterization of the machines as functionally allocated is not adequate or complete. One way to clarify this is to discuss with A the fact that he has actually assembled in his office a set of resources which he considers to be allocated cleanly but which actually he distributes rather more fuzzily than he thinks. Importantly, he says he is able to work with this distributed resource without becoming confused, so he obviously has some success at mapping underlying activities to the resources (and he reports having had two or more machines in his office for over 10 years, so he is comfortable with multiple computer resources). But what he is actually doing is, in fact, only recoverable through introspection, observation and interview. A's account of the usage does not reveal the true complexity, and only when observation is added, and the discrepancies and different cases of use highlighted do we get the fuller picture. This is not an uncommon finding, but does support the need for direct field observation as well as accepting reported behaviours and user-based descriptions of use.

Alternatively and more objectively, perhaps, one can deepen the discussion by noting that the six screens and four keyboards permit attention switching and task suspension with minimal disruption. The issue here is that A can switch between machines and 'pick up where he left off' without having to find the caret, locate and foreground the desired client/application and so forth. This is observed behaviour. The attention switch is manifest by swiveling the chair and reaching for and using the other keyboard/mouse (of the two available when seated on the swivel chair) or getting up and moving to another chair. We will return to A's set-up, and the interpretation it prompts, in the later discussion.

2.2 User E

User E is a male University student, in a Computer Science department, who is in his final year of undergraduate study. He was recruited to the study after overhearing the researchers discussing the project, and letting them know he also used two machines. He was interviewed for approximately 90 minutes, using open questions and probing his answers for deeper understanding and self-knowledge. Being the fifth user studied,

the researchers were much more aware of the existing patterns of multiple computer use, and so used more focused questions, and were able to probe more deeply, despite the study period being shorter. However, all the material here is based on User E's reported patterns of use, dissected with questioning, but not on observation. His setup (as photographed by himself) is shown in Figure 2.

He uses two machines, both Macs: a laptop with Intel Dual core, with 80GB hard disk, and a G5 Dual Processor desktop with a 40GB disk - both running Mac OS X 10.4). Each has only one screen and they are not networked together by default, though they can be cross-mounted wirelessly for file transfer purposes. For around 6 months he tried a second screen linked to the desktop (giving a 19" and a 21" screen on that machine) - but didn't like it, citing problems in moving the mouse larger distances, in not using the second monitor but instead using tabbed browsing or overlapping windows on a single monitor, and finding it hard to work out what he wanted to put onto which monitor. E reports he does not like to move his attention from one screen to the other when using the same application, the cause of some of his difficulties in using multiple monitors, but is happy to move between multiple windows on one screen.



Figure 2. User E's working setup - laptop on left, desktop screen on right, machine under desk to right of chair.

He has used two machines for six years, upgrading the desktop and the laptop alternately, getting a new machine every twelve-eighteen months. He reports that his ideal set-up would be two desktops, but needs a laptop for University-based work. However, he doesn't see any need for more than two machines. E maintains similar file structures on the two machines - but no attempt is made to match/synch the content. One machine does not back up the other. He states that this is to reduce the cognitive load in locating files: they are in the same position in a hierarchy on either machine.

E has two main scenarios of use. The first discussed was focused around gaming activities: he plays a lot of World of Warcraft - an internet-connected, multi-user rôle playing game. This game is played in full-screen mode on the desktop machine, with the laptop used to run additional functionality needed for effective play: VOIP and text chat, to allow him to synchronize actions with others, and a web browser, which he uses to view reference material. In addition, he will run anything else he feels he needs on the laptop, in order to keep the desktop machine as uncluttered with other processes as possible. He reports that he does this in order to get the best possible performance from the game application.

The second scenario concerns his work approach. E codes on the laptop, and also runs MSN and Skype on there. It is worth

noting that he reports having in excess of 120 skype/messenger contacts, of whom 15-20 are high usage. In this scenario, he uses a web browser on the desktop, to refer to documentation and so on. He reports that he often has a lot of tabs open in the browser, but doesn't like too many windows open - and if he does have multiple windows, then he has them overlapping so he can flick between them, and does not minimize them into the dock (similar functionality as the Windows Taskbar). Note, however, that he finds it much easier to flick between overlapping open windows on one screen than in having the windows on separate screens. He will also use the desktop machine for document preparation, citing the additional screen size as the main reason for this. If this word processing requires more than one window open, he has them all onscreen and overlapping, eschewing the dock again. When word processing, he likes having the messenger application open on the other machine, since he is actively encouraging interruption, and the visibility of the application means it is easy for it to grab his attention when a contact messages him, or for him to see a new contact come online and hence instigate a conversation. This visibility and potential for interruption is enhanced by the application being on a separate monitor with more open screen estate. However, E reports that he likes being able to turn to that machine and just start typing, without having to manoeuvre the mouse and select the application: this functionality, important to User E, is only possible because of the additional caret position and foregrounded process offered by the second machine. This ability to use more than one application in the foreground crops up numerous times in E's descriptions of his usage - being able to select applications and focus by a physical switching to different keyboard appears to be much less cognitively demanding for E than managing this on one machine.

In both scenarios, he sticks rigidly to the allocation of activities to the specific machines. For each of these scenarios, however, he actually describes them as 'tasks' - so there is a 'gaming' task, and a 'programming' task, and a 'document processing' task - and for each of these tasks there are a set of applications and resources that he uses, each of which is allocated to a specific machine and which stays on that machine (only when really pressed in questioning could he recall this breaking - when programming, if he had more tabs open than could be easily displayed, he might use another web browser on the laptop as well as the default desktop). Notice, however, that the allocation is not application ⇔ machine based: for example, depending on his 'task', the web browser will be on either the laptop or the desktop.

He can and does use both machines for email, checking email only once a day, at the end. He downloads email only onto desktop, using webmail from laptop - around 10 emails per day. His rationale for this different treatment of reading email has a wider application: he wants to keep the laptop 'clean' with no superfluous material on it, and so keeps copies of emails on the desktop only, and downloads new applications onto the desktop - only if they get sufficient use will they migrate to the laptop, and then only the latest version - older versions may remain on the desktop machine, but never the laptop.

E's use of bookmarks is also interesting: the bookmarks on the browsers on his different machines are not synchronized, and match the tasks that the browser on that machine does: for programming work, bookmarks are on the desktop, but for gaming work, they are on the laptop. However, he does report that if it is a forum or something else that he likes to check regularly, it will go on both machines. He also states that the

laptop has only the most commonly used bookmarks on it (about 15) whilst the desktop has about 60. On the laptop, all the bookmarks are ones he will use often – at least once a week. The desktop browser holds the old links that may still come in useful one day – but he states that this ‘rule’ – his phrase – overrides the other, in that if it is a gaming link (which would therefore usually go on the laptop) but he thinks he’s not going to use it often, it will go on the desktop. In this case, it seems his desire to keep the laptop ‘clean’ overrides the allocation of tasks to machines.

E comments that he hardly ever does one thing at a time – a single machine feels like it forces a bottleneck in his activities – though he is aware that it is not likely to be a computational bottleneck as much as a cognitive one. He says he views the two machines as a single resource – and has no notion of ‘main machine’. This corresponds with his ‘scenario’ approach to application⇌machine allocation, in that both are used to support the same scenario, but in different ways depending on the scenario.

2.3 User D

User D responded to an emailed request for users who regularly used more than one machine at once, and was interviewed and observed on two separate occasions, for about an hour each time. Clarifying discussions and emails after the initial sessions served to clear up any confusion or lack of understanding from earlier sessions.

User D has two Apple Mac computers: a G3 and a G4 machine. The G3 runs Mac OS9.1, the G4 runs OS X. This user has the two primary monitors adjacent to each other on his desktop – the left-hand monitor for the G3 is an older CRT and the monitor for the G4 is a 23” LCD monitor. To the right of the LCD monitor is another CRT which can be called upon to provide extra screen estate when this is thought necessary (this too has a keyboard in front of it but D does not have a separate machine connected to the monitor). However, this screen has not been observed to have been used, and if it is used it is only occasionally. Beneath the two main monitors are his keyboards, with a gap of at least 12 inches from the desk edge. User D says that he likes to “have papers and things on the desk in front of the keyboard” – he has to lean forwards over these to type, is aware it’s not good for his posture, but finds it perfectly satisfactory (see Figure 3).

His allocation of functionality is primarily based on the OS used. He keeps the 9.1 machine since there is a specific application that he regularly uses which only works under 9.1, and he has some 9.1 specific visualization and graph drawing tools on that machine as well. Whenever he needs to use these tools, he is forced to use this machine, so switches to working on the left. Because of the graphical tools, which are used to produce illustrations for papers, he also has Word installed on that machine and does all his paper writing on it. The central monitor is given over to Mail and web browsing. His mail client is permanently active, though he uses filtering so that he only receives interruptive notification if something “important” comes in. For him, “important” is “non-suspected junk”. He uses Internet Explorer as his browser, with all his bookmarks set up. The central monitor also supports his X11 shells, which are used to provide a window onto the School’s research servers, on which he runs a large number of compute-intensive jobs, which he monitors via terminal windows.

A few applications exist on both machines (monitors). Word is installed on the OS X box (and hence on the middle screen) – it is set up “exactly” the same as on the left monitor – but why is

it on both? Further questioning reveals that it is used to convert Word documents to PDF, an activity easy under OS X but less simple under 9.1. Internet Explorer is also installed on the left monitor, but the bookmarks are not set up or synchronized with the ones in the middle. This browser does get used occasionally, but only when “I don’t need up-to-date bookmarks”, though because the central monitor is bigger “and better” it is usually the first choice for web browsing.



Figure 3. User D’s office layout showing three monitors.

Data are transferred between the two machines using an ftp client on the 9.1 box, chosen since it deals with the differences between line terminations across the two systems (OS X uses LF – line feed; 9.1 uses CR – carriage return), converting them seamlessly as it moves from one machine to the other. Mostly, the material transferred comprises “bits of text and files from emails and web sites”. He also uses Matlab locally on the OS X Mac, so results from that are also transferred in this way. User D also has an external Firewire drive attached to his 9.1 machine, on which all his data are stored, and occasionally uses this to move data onto the OS X machine by physically unplugging from one and plugging it in to the other. Clearly, he has a need to move large amounts of data around, though this is reported to happen about once a month only.

User D reports that he thinks about the setup as a single computer, rather than as two. It appears that tasks are identified by a key application that defines the primary operating system to be used, and then any supplementary tools for that task are on that same machine. Thus, when writing a paper, graphing tools are needed, and so the left monitor (OS 9.1) is used, with User D working in Word to create the paper. If the task is email, then the central monitor is used; ditto for web browsing. Very few of the tasks the user is involved in are split across the monitors, and this strict boundary between the physical monitors and the cognitive perceptions of task seem to make it easy for the segmentation to occur. User D says that he “never” makes mistakes in choosing the correct keyboard and mouse to use, and puts this down to having the keyboard so directly underneath the monitor, with the mouse (wired) close by. However, he is adamant that the keyboard position is dictated by his penchant for having papers on the edge of the desk close to him, not to help in identifying monitor usage. This lack of error in keyboard/mouse selection suggests to us that when focused on a task, his attention is focused on just one screen, and there is little need to shift attention to the other. He also

reports that he uses very high tracking speeds with his mice, so that they never get moved far away from their standard, base position.

Exploring this further, we can see that the left screen is used for the creation and editing of documents, and tasks related to this, and that these activities are not interrupt-driven. The central monitor contains email and windows onto results, which could, in principle, be interrupting him in his work – but they are limited to this central screen. In addition, the filtering applied to email means that interruptions occur rarely. We discussed the use of the monitoring windows for the networked server processes, and he reported that he used to periodically check these to ensure that they were running as he'd expected – but checking was an active, user-initiated action, not driven by a warning or flag from the system. Therefore, User D remains in close control of his active tasks and working context, and makes explicit switches from one to the other. He doesn't use applications such as MSN messenger, and so isn't potentially interruptible by such a program.

User D's office set-up changed whilst this paper was in preparation. His administrative duties changed and he was "forced to allow" a PC into his office, as he reports it, to deal with the overall examinations management work for the department, involving extensive use of PC based software. This set-up is illustrated in Figure 4. Being a very recent change, we have not observed his use of this system, but notice the keyboard placed directly under the monitor and away from the desk edge, and the physical separation between the existing machines and this new interloper. D reports that he will turn the Examinations PC on to deal with Examinations, then turn it off again.



Figure 4. User D's new set-up with PC on the right.

2.4 User B

The observations made above about A's pattern of computer usage contrast with the work pattern of B. B is also a University academic. B was questioned over the course of approximately six weeks, using a short series of discussions, reflection and self-reporting. B has two machines, one a Mac G4 laptop and the other a Mac G5 desktop with 23" monitor and a 19" monitor, arranged side by side. This situation is likely to be very familiar to many people. B uses the laptop for all his work when away from their base (in this case, a home office). The laptop accompanies B when at conferences, traveling on the train, visiting colleagues, and so on. For B, it represents the

subset of their activities that are current, and so contains much less data and information than the desktop machine. Whilst the desktop is regarded as the "main" machine, it is often the case that the most up-to-date documents, papers and presentations can be found on the laptop. User B has no software-based synchronization between the laptop and desktop machines, nor any informal process for undertaking such a task. Files are transferred on an ad hoc, "as needed" basis, via a wireless network, though B reports that key files are sometimes left behind (i.e., the subset conception of the laptop in relation to the desktop is not reliably implemented).

There are a number of interesting observations to be made. B makes it clear that he views the machines as quite separate, without any segmentation of tasks and activities onto the different machines – the segmentation is geographically and temporally based. B works on the laptop when out and about, and when working elsewhere in the house (whilst watching television, or in the garden). However, the additional screen estate of the desktop provides sufficient incentive to move to that machine as the need arises. Interestingly, B reports that, prior to the installation of broadband at home, his "main" machine was in the office at the University, and the laptop had completely different material on it compared to the work machine: it was used for completely different tasks. Writing papers and outline documents was done on the laptop; programming and email were done on the work machine. The introduction of ADSL broadband into his home led to a significant change in the distribution of material across machines, such that the home desktop is now the primary machine, and has copies (albeit potentially out of date ones) of the material that exists on the laptop. Email, web browsing and internet research activities are distributed randomly between the two machines, though B now uses the IMAP email protocol to maintain some synchronisation between his mailboxes and their multi-machine access patterns. The office machine is hardly used at all.

The file system on the laptop and the desktop are not the same. Whilst B does not partition work according to the machines, and regards them as both undertaking all tasks, the file structure on the laptop tends to be shallower and wider than that on the desktop. When questioned, B suggests that this is because he often has one or two specific tasks to achieve when away from his base, and so transfers just the relevant files across. Since these files may have come from a broader project, the relevant file structure may well not exist on the laptop, and they are placed nearer the top level for more immediate access. Thus it could be argued that B treats the laptop as a much more active machine with little sense of the archival rôle served by the desktop. However, the file system on B's laptop is not completely disorganised; for example, the research writings are structured around conferences rather than topic. B suggests that this is because of the nature of the task at the time: report work for the specific conference.

B does use both the desktop and laptop machines simultaneously, but not in the same way as A. In this case, the laptop is placed at about 45° to the main screen, on a side desk – sometimes to the left, sometimes to the right, depending on the piles of paper that are to either side of the monitor. Whereas A moves between computers and tasks, B tends to use just one computer in this scenario, accessing the other one to transfer data. Even when multiple tasks are being performed (e.g., email, writing papers, instant messaging, web browsing) B usually performs all of these on one machine, only occasionally using the other. B also reports that he is much happier "pulling" information off one machine, rather than "pushing" it onto the

other. This may be because he is task focussed, and pulling the information means that he is working on the machine that he wants to continue working on. In contrast, if B pushed a file from one machine to the other then he would have to switch machines in order to continue working with the file.

2.5 User C

An alternative work pattern has been discussed with another regular laptop user: they replicate exactly the file structures between their desktop machine and laptop one, and tend to keep them in synch with one another. They report liking to know that all their information is with them at all times. Here there is no separation of information across machines, and in many senses they are not using two machines at all, but are simply using the laptop as a way of being able to access their work at any time. It can be argued that this is taken to its logical conclusion for users who only have a laptop, but use additional screens and external keyboards, sometimes in a docking station, to provide them with greater facilities when at work. In particular, user C will only use one machine at any one time: their contribution to this discussion is therefore limited, but we will mention them later.

3. DISCUSSION

In the discussion that follows, we are specifically interested in what the characteristics are of multiple computers and multiple displays that make them different from a single computer with multiple monitors connected. We are interested to uncover what aspects of observed behaviour and inferred cognitive activity are influenced by these differences, whether they are at all significant, and how they can lead us to better understand how to design environments for multi-screen working. During this discussion, we assume that a single computer running multiple monitors will be sufficiently powerful to accomplish the tasks asked of it: we are not considering machine performance issues in our deliberations, though for some power users this is clearly a significant factor in them using multiple machines.

B provides an example user who is probably familiar to many readers. He uses two computers, sometimes simultaneously, but not often, and tends to use one when mobile and the other when in his home office. Whilst the tasks transfer from one machine to the other, the file structures do not, and the potential reason is that the laptop is used for fewer simultaneous tasks, and easy access to the current activities is desired over a comprehensive filing system. When in simultaneous use, B tends to use the machine with the most current information on it, though prefers the desktop to the laptop, and tends to pull work from one machine to the other. B has a simple model of task allocation, based on likely machine availability: when out and about, laptop, otherwise, desktop, and when given the opportunity to use both machines, tends to focus on one with the other acting as a data source. This pattern of usage is readily understood with few analytical challenges. User C is not actually a user of multiple computers: instead, they have multiple computers, but use only one or the other – they form one end of the spectrum of use we have observed, and are the least interesting in terms of new interactions and insights. As with User B, they contrast with the other more complex behaviours and patterns of use described above, which need an analytical framework.

There appear to be a number of possible differences between multiple computers and a single computer and multiple monitors. These can be divided into the technical, and the cognitive.

3.1 Technical Differences

The identification of the importance of multiple carets has provided us with the greatest insight into the importance of multiple computers. Compare the differences between working with multiple screens: when changing focus across screens, the single mouse is used to move the keyboard focus from one window to another. However, unless the window selection is done very carefully in the window border, the change of focus by mouse click also changes the caret insert point, which means that switching between tasks and windows requires an additional step, that of returning the caret to the previous location – and this can be cognitively complex, especially if there is a long delay between switching. Compare this to having multiple computers: there, the switch from one machine to another may require the mouse to be moved to re-activate the screen, but even then the caret point is in the same place as when the task was suspended.

For users under X11 window systems, and those that modify their default settings under Windows, it is possible to make the window under the cursor the active one (and pop to the top). This is not the usual scenario, but is a possible one. However, in this case, the order of windows is critically dependent on the route the cursor takes though the display, and so a change of task from one set of windows to another may not be easy to accomplish.

In A's case, there is a clear difference between multiple computers and a single machine with many displays, based on the ease of task switching without losing the context (mainly, the caret) of the existing task. This is also true for User E, who also takes active advantage of this.

It is also the case that multiple computers offer the opportunity for the same application to have a clearer separation of histories, allowing the user to move backwards and forwards in two threads of work without getting the histories intermingled. Web browsing is a typical example: whilst the 'back' button on the browser takes the user back through the pages that that particular window has previously accessed, the global history is merged, and this can (and does) cause confusion. Opening documents in Word changes the 'Recent Files...' list: multiple machines keep these much more separate. For User A, this is a useful feature. For User E, this is not a consideration, owing, we believe, to their strong scenario-based approach to their activities.

One of the other roles of multiple computers is that, according to A, they act as reminders of work currently in progress. For example, it has been observed that A sometimes has the left Mac laptop and associated Samsung screen off: this indicates to him that he is *not* engaged in lecture preparation, and allows him to focus more intently on other activities. Conversely, the research machine processing data showing its screensaver requires no more active attention, but provides a very different 'feel' for A, who knows it is working hard on his behalf, and knows that he can attend to it at any time. This scenario is not easy to replicate on a single machine; Hutchins and Stasko [8] have pointed out that windows can act as reminders, but this characterization does appear to take more cognitive effort than noticing a whole computer and monitor.

3.2 Physical Actions and Cognitive Implications

The next point is more complex, and concerns the need to refine our understanding of multi-tasking. The users of multi-screen systems, whether or not these are supported by multiple

computers/keyboards/mice, face a management problem also faced by users of single screen/computer configurations. But the technology of multiple screen provision teases out the difficulty that all users face. It is assumed in the following discussion that users of single and multi-screen systems will be using a multiplicity of 'windows' on the screen(s) they are using (regardless of the OS installed). Further, it is assumed that these 'windows' are deployed so as to permit some level of multi-tasking (and this can be true within an application as much as it is obvious between applications).

The difficulty is simply stated. Multi-tasking is itself a task – and it needs supporting and facilitating in ways which do not disrupt the other tasks.

Conventionally, multi-tasking is facilitated by one or another system which foregrounds one task at the expense of the others. The user has to perform one or more actions to select which task is foregrounded at any moment and he or she may be offered prompts/tricks to facilitate this – but not to avoid it. This means that for the user in a multi-tasking environment the multi-tasking becomes a task to be learned and then managed. The prompts which have to be mastered to manage window changes whilst multi-tasking can range from keyboard-based tabbing between applications, or the selection of applications or documents in the dock or taskbar, or by moving the mouse (and usually also making a selection). But they are all activities – and so all need to be managed as a task. This can interfere, we suggest, with the other task activities. Firstly, it is an additional load and thus competes for cognitive resources. Secondly, it exists at a different level cognitively, being a meta-task, and therefore requires a different form of context switching. Conventionally, then, multi-tasking is about both foregrounding and backgrounding, about management of this, and about monitoring of the relative 'positions' or 'locations' of the various tasks (the latter can be achieved via one of several tricks provided by the OS favoured – e.g., the Taskbar).

However, use of multiple machines provides a great deal more support for multiple tasking because it allows the user to leave various tasks foregrounded on different machines (specifically, the caret is not so likely to be unhelpfully relocated). The user's task switching behaviour is then reduced to the physical selection of a different keyboard and mouse, which, we suggest, may be cognitively simpler for many situations than altering the state of a single machine.

Additional issues for multi-screen systems on a single computer concern the user's behaviour with the mouse. The mouse pointer can easily get lost on two screens – having one pointer per screen makes it easier to find it. Additionally, moving the mouse large distances can be tiring and irksome: physically providing more mice with each having to travel less far resolves these issues: the space of possible locations for a 'lost' mouse pointer is reduced, and the range of movement needed is less, per mouse.

The simple fact is that use of multiple machines for different tasks is a way of providing for multi-tasking which does not require a management system more complicated than a chair. In such a set-up as constructed by A or D we see provision for task-switching provided by physical movement or relocation. In the same way, User E has recognized that, at a high level, a single task we have to undertake often requires the use of multiple applications, and managing those applications by moving between different machines is a simple one to understand and implement.

This can be taken a stage further: consider the physical separation between screens in the office setups chosen by A and D – they have made a definite effort to remove some of the screens from being adjacent to each other, something not seen in multiple monitor displays, nor in the behaviour of E. The physical distance between the systems corresponds to a cognitive separation of tasks in their minds. This helps them maintain task focus, since their visual field does not contain information from the other screen – separating other tasks out of sight also takes them, at least temporarily, out of mind. This setup offers both a physical and cognitive barrier to rapid task switching, which may actually be beneficial: it is harder to get distracted and hence attention can be more focused, but when a switch is done, it is much more immersive, and so transfers more cognitive resource to the new task. But for E, the interruption offered by an alternative application, placed near the focus of attention, provides a welcome distraction.

3.3 Further Considerations

Reflection on the differences and similarities between these users reveals several further points. Firstly, there is a spectrum of task allocation to machines. This can be done cleanly as by D, or less distinctly, as by A. Or it can be clear, but actually based on circumstance, as in E. Additionally, it may not be done at all, as in C. User B is apparently different, but actually is equivalent to C in degree of task allocation. B's concern is pragmatic, and he attempts to ensure the subset of materials on the laptop permit isolated/remote/roaming working. In fact he could do all and anything on either machine if physical location did not matter: for C, they can and do do anything on either machine.

A related point here is that some tasks naturally lend themselves to isolation on single machines – others do not. D makes this quite clear – special software requiring a special OS serves to define the reason for having a specific machine. Contrastingly, User A demonstrates that tasks are more spread across machines than initially realized. For B, tasks are often transferred from one machine to another: for C, there is only one machine (the 'current' one) on which all tasks happen.

4. IMPLICATIONS FOR DESIGN

From the studies and analysis presented, we can draw out some implications for the design of future systems. It appears to be the case that users having access to more than one machine and monitor is likely to become increasingly common, and it is clear from these studies that multiple computers and multiple displays are a very different proposition to just having one computer and multiple displays.

In this new world the existing assumptions regarding managing multiple applications do not provide the ideal solution. In particular, there is a cognitive load associated with application management, and allowing for physical distribution across machines can ease that load. It seems that the 'scenario' approach to application management, in which different applications are assigned to different machines and monitors, allowing them to be jointly in the foreground and to require physical task switching activity, is likely to become a common arrangement. Explicit support for this may be beneficial to users, since there may be many applications, in different screen locations and on different machines, depending on the scenario: providing some form of scenario management built into the system may well be beneficial. Interestingly, this often tends to be the approach adopted by users of virtual desktops, who arrange sets of applications within a single screen based on a high-level task [13].

One of the major lessons is that having multiple, active foregrounded tasks is appreciated by users, who are able to work on their current task whilst being easily able to monitor the ongoing activities of other tasks, and switch between them readily. The physical nature of the switch seems to support usability, reducing cognitive load. It also seems to be the case that providing collections of applications in a user-defined application provides a lot of innate context for the user, and gives them immediate access to the tools that they need: this scenario-based layout of programs could be better supported by application designers at present, especially across virtual desktops. We can also see that supporting scenario-based histories within both browsers and ‘recently opened...’ lists may well be beneficial for users. Also, offering support for multiple instances of applications running across multiple machines to share state, history and activity information may become useful as the number of computers increases, though clearly there is potential for adding unnecessary complexity.

Designers need to realize that their software is going to be used in conjunction with other software applications and other computers, and should therefore design it sympathetically, potentially taking advantage of its visibility onscreen, and of it being permanently in the foreground, even if not the focus of attention. For example, a screensaver could identify the scenario that the machine is being used for, or, at least, identify the applications that are active on it, in a large font legible from across the room, or in iconic form. Moreover, if any application required input or attention, this could be notified to the user by visually altering the screensaver, or suspending it, thereby identifying this fact to the user.

Some of the design implications that we can draw from this study are applicable to existing systems (such as scenario-based setups), but it is important to realize that others are directly related to the fact that multiple computers are in use (e.g. multiple foregrounded activities). Multiple computers do not behave the same as single computers with multiple monitors, and they are not used in the same way either. Such multiple use offers potential benefits in terms of task management and so we expect to see the development of applications and infrastructures that are designed specifically for more than one machine.

5. CONCLUSION

In this paper we have presented a case study of five users who use more than one computer in very different ways. User A has a clear model of task allocation to different machines, which is not as clear cut in practice as his perceptions. Users B and C use a second machine as an alternative to a primary one, transferring tasks from one to the other; C has a simple approach, whilst B’s reflects more complexity in mapping current tasks to easy accessibility on the portable machine. When using both machines, B tends to treat the laptop as not much more than a mounted disk on the main machine. D provides a clear mapping of task to machine, based on the operating system support required. User E has a set of working setups that he uses. A, E and D provide us with insights into the overheads of context switching, and we have seen there are good technical, physical and cognitive reasons that differentiate a multiple computer – multiple screen setup from a single computer – multiple screen system.

One particularly significant facet of the behaviours observed is that users have different conceptions of tasks. For some a task is constrained to be activity using a specific application. For others the operating system may be considered task specific or task defining. And again, other users may consider tasks to be

identified with machines. Of course it is also possible that users will mix their notion of tasks, rather than being cleanly differentiated in their behaviours. However, the implication for designers is clear – it is not obviously helpful to consider users as having a homogeneous approach to relating tasks to software and hardware. Task analysis and work-flow planning should not be predicated on particular configurations of software (applications) and hardware (multi-monitor versus multiple computer setups).

The study has implications both for researchers and for designers. Both groups of professionals need to understand that multi-tasking in computer use is more complex than it appears to be at face value (and thus deserving of more study). Designers need at the very least to understand that their applications are used with others, and in multiple scenarios with interleaved behaviour.

We envisage significant developments in research arising from this study through the need to understand the observed diversity. One important question to be addressed is whether the differences observed can be explained in part, if not completely, as the consequence of individual differences in experience. Is it to be expected that users will migrate to multi-monitor and multiple computer configurations? If there are other factors to be considered what are they? How will increasing divergence in workplace configurations impact on the users’ facility for migration between different configurations? How will software designers learn about, and then subsequently incorporate, differing requirements for inter-application working in varied scenarios?

The conventionally envisaged computing environment comprised one or a few applications used for tasks recognized as significant components of workplace activity; all that is dissolving into individualized scenarios and configurations. Dissatisfactions with existing applications are complemented by individual work-arounds, some of which may be instructive and generally useful to other users.

REFERENCES

- [1] Brown, B., Green, N., and Harper, R. *Wireless World: Social and Interactional Aspects of the Mobile Age*. Springer, London, 2002.
- [2] Czerwinski, M., Smith, G., Regan, T., Meyers, B., Robertson, G., and Starkweather, G. Toward characterizing the productivity benefits of very large displays. In *Proceedings of the IFIP conference on Human-computer interaction (INTERACT 2003)*. London, IOS Press, 2003, 9-16.
- [3] Forsythe, D. E. “It’s just a matter of common sense”: Ethnography as invisible work. *Computer Supported Cooperative Work*, 8 (1999), 127-145.
- [4] Grudin, J. Partitioning digital worlds: Focal and peripheral awareness in multiple monitor use. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI 2001)*. ACM Press, New York, NY, 2001, 458-465.
- [5] Guimbretière, F., Stone, M., and Winograd, T. Fluid interaction with high-resolution wall-size displays. In *Proceedings of the 14th annual ACM symposium on User interface software and technology (UIST 2001)*. ACM Press, New York, NY, 2001, 21-30.
- [6] Henderson, A., and Card, S. Rooms: The use of multiple virtual workspaces to reduce space contention in a

- window-based graphical user interface. *ACM Transactions on Graphics*, 5, 3 (1986), 211-243.
- [7] Hutchings, D., Smith, G., Meyers, B., Czerwinski, M., and Robertson, G. Display space usage and window management operation comparisons between single monitor and multiple monitor users. In *Proceedings of the working conference on Advanced visual interfaces (AVI 2004)*. ACM Press, New York, NY, 2004, 32-39.
- [8] Hutchings, D., and Stasko, J. Revisiting display space management: Understanding current practice to inform next-generation design. In *Proceedings of the 2004 conference on Graphics interface (GI 2004)*. ACM Press, New York, NY, 2004, 127-134.
- [9] Miller, R. C., and Myers, B. A. Synchronizing clipboards of multiple computers. In *Proceedings of the 12th annual ACM symposium on User interface software and technology*. ACM Press, New York, NY, 1999, 65-66.
- [10] Paul, D. Implications for design. In *Proceedings of the SIGCHI conference on Human factors in computing systems (CHI 2006)*. ACM Press, New York, NY, 2006, 541-550.
- [11] Perry, M., O'Hara, K., Sellen, A., Brown, B., and Harper, R. Dealing with mobility: Understanding access anytime, anywhere. *ACM Transactions on Computer-Human Interaction*, 8, 4 (2001), 323-347.
- [12] Rekimoto, J. Pick-and-drop: A direct manipulation technique for multiple computer environments. In *Proceedings of the ACM symposium on User interface software and technology*. ACM Press, New York, NY, 1997, 31-39.
- [13] Ringel, M. When one isn't enough: An analysis of virtual desktop usage strategies and their implications for design. In *CHI 2003: Extended abstracts of the SIGCHI conference on Human factors in computing systems*. ACM Press, New York, NY, 2003, 762-763.
- [14] Rouncefield, M., Hughes, J., Rodden, T., and Viller, S. Working with 'constant interruption': CSCW and the small office. In *Proceedings of the 1994 ACM conference on Computer supported cooperative work (CSCW 1994)*. ACM Press, New York, NY, 1994, 275-286.
- [15] Schmidt, K., and Bannon, L. Taking CSCW seriously. *Computer Supported Cooperative Work*, 1, 1 (1992), 7-40.
- [16] Weiser, M. The computer for the twenty-first century. *Scientific American*, 265, 3 (1991), 94-104.