

DESKTOP COMPUTING - DISTRIBUTED COGNITION IN A TAX OFFICE

Martin M. Nielsen
Susanne Bødker
University of Aarhus
Department of Information and Media Studies
University of Aarhus
Aarhus, Denmark
Email: mmm@imv.au.dk
Department of Computer Science
University of Aarhus
Aarhus, Denmark
Email: bodker@daimi.au.dk

ABSTRACT

Based on a detailed study of the use of representations in a tax assessment process, this paper presents an analysis of the use of the physical desktop and of paper documents, files and electronic information. This analysis challenges the ways in which the computer desktop is designed and used normally, and we present a number of challenges to user interface design. Taking these seriously, means to revisit several taken-for-granted elements of the current WIMP regime: the randomly overlapping windows on a non-structured background; the lack of traces of time and past location; and the individualised and non-activity-oriented set-up of the desktop.

Keywords

Distributed Cognition, Workplace Studies, Desktop Metaphor, User Interface, Digitalization.

INTRODUCTION

The starting point of this paper is the taken-for-grantedness of the desktop metaphor. The vast majority of computer applications today are based on a particular version of the (computer) desktop. This is the version where multiple applications run in multiple, randomly overlapping windows that are placed rather arbitrarily on the screen; where there is one cursor and only one active window at a time (van Dam 1997). This approach to the desktop is predominant in Windows applications, as well as on the Macintosh. In short, we are concerned about how the current WIMP regime rules the interaction with computer applications brought together on the computer desktop. The taken-for-grantedness of the current desktop style means that most software for e.g. administrative applications are designed with overlapping windows etc. without much deliberate concern: What if there are ways, e.g. of aligning paper documents, that get lost in the computerisation?

Historically, there was a time before this version of the desktop. Experiments were made with tiling - a variety of mechanisms to place documents next to each other on the screen, utilising the screen better, magnifying some documents while diminishing others (all of this is actually standard textbook stuff, see e.g. Shneiderman & Plaisant 2005). The Smalltalk browser used connected windows, an element that seemed to have been forgotten, until it was taken up again in Macintosh OSX. And there are many more of such experimental implementations that have been more or less forgotten in commercial settings.

In the following we will present a case study of a work process that is undergoing scrutiny by commercial software companies regarding the potential of "going electronic" while at the same time being a work process where e.g. juxtaposition of documents and use of physical desktop space matters, as we shall see in the following. The paper is based on a detailed study of the work process

based on the framework of distributed cognition. In this paper, our primary focus is on understanding how our findings challenge our understanding of desktop computing.

In HCI literature, many papers have been dedicated to accounts of studies of how people use their office/desk space, and to support this with technology. Most prominent and often still referenced is Tom Malone's paper from 1983 (Malone 1983). However, that study and many more are primarily focusing on individuals' personal information management (e.g. Belotti and Smith 2000, Hayashi et al. 1998, Kirsh 1995) – piling, filing and retrieving information that somehow “lives” in the office of the particular person studied.

Hertzum (1999), as well as e.g. Malone (ibid.) points to the many roles of documents, and indeed we find similar uses of documents here. However, our study differs from the above mentioned in that we are not as interested in documents that live in the office, as we are in detailed studies of a particular activity that occupies most of the time of the case workers for several months every year; only to be continued in the next step in the annual cycle. Furthermore we are not primarily concerned with the individual procedures as they are applied by individuals. Our focus is more on the shared procedures as they have been crystallised and are relived year after year within the community of practice of the tax office.

THE CASE STUDY

The office under investigation in the present study is assessing taxes for self-employed. In most cases the taxpayer is a regular wage earner at a firm, but runs a sideline business. The group of self-employed, however, includes traditional independent businesses such as carpentry, cleaning service, etc. as well. The self-employed are obligated to submit a balance sheet for their independent business. Usually the balance sheet accounts for expenses and earnings, profits, turnover, etc., i.e. how the business reaches the bottom line as it appears from the tax form.

The assessment of self-employed normally takes place by examining if the profits of the independent business are well accounted for; if the result is transferred to the tax form correctly; if the information from the tax form and the balance sheets are in concord with “control information” collected from various private and public sources. When the case officer at the tax office finds nothing suspicious the file is placed in the archive. When, on the other hand, something erratic is discovered, further information (vouchers and other documents) is called for.

DISTRIBUTED COGNITION

The description above (dressed in tax jargon) of the work process is a common way to think of the task at the tax office. In the following we will try to re-dress the process in the vocabulary of “distributed cognition” (Hutchins 1995a, 1995b). We do this because we find that this vocabulary provides us with a privileged access to the workings of the tax office especially with regards to the informational make-up of the tax office. This approach is, as we will argue, valuable in design of informational systems. At first, however, we need to establish the theoretical framework of distributed cognition.

Distributed cognition originates within cognitive science and the orthodox “cognition as computation” framework (Hutchins 1995a, 49). The basic presumption is that cognition and computation can be described as manipulation of *representations*. In the traditional interpretation of this view representations are mental constructs that somehow mirror salient structures in a task domain and thinking and problem solving take place through the manipulation of these mental constructs.

The cognitive ethnographer Edwin Hutchins reformulates the traditional framework in his path-breaking studies of ship navigation (1995a) and aviation (1995b). In accordance with traditional views he maintains that cognitive systems are computational and that the study of cognition is the study of representations and algorithms: “[O]ne may ask of this computation how its inputs and

outputs are represented and what algorithms are used to transform inputs to outputs” (Hutchins 1995a, 130). The clash between traditional cognitive science and distributed cognition arises on the question as to what constitutes a computational/cognitive system.

According to the traditional interpretation the cognitive system is the individual or more particular her mental apparatus. What Hutchins discovers, however, is that a whole crew of navigators on a navy vessel satisfy the same functional constraints as has traditionally been attributed to the individual. Furthermore, the navigational artefacts comprising compass, charts, alidades, hoey, etc. play an integral part in the problem solving at the navigation bridge. Hutchins takes this to show that traditional cognitivism may have mis-attributed the functional makeup of whole socio-technical systems to individuals. Instead cognitive science needs to realise that cognition is distributed across systems integrating external/socio-technical as well as internal/mental representations:

“I propose a broader notion of cognition because I want to preserve a concept of cognition as computation, and I want the sort of computation that cognition is to be as applicable to events that involve the interaction of humans with artefacts and with other humans as it is to events that are entirely internal to individual persons. (...) [T]he actual implementation of many interesting computations are achieved by other than symbolic means. For our purposes, ‘computation’ will be taken, in a broad sense, to refer to the propagation of representational state across representational media” (1995a, 118).

Representational media are thus not exclusively mental constructs according to distributed cognition. Representations may be materialised in artefacts and in work settings in general. In the navigation example, the chart (a representational media) defines a problem space where the navigators need to map their location (a representational state). As such, the computational processes of a socio-technical system can be examined by studying the “creation, manipulation, and propagation of representational states” (Hutchins 1995a, 49) of the system. A great deal of these processes are observable “on the surface” of the systems and need not be extracted from the brains.

Choosing distributed cognition as the theoretical basis is in a way an experiment in itself (see Nielsen 2004). With respect to the findings reported in this paper, we believe that it would have been possible to make rather similar analyses e.g. using the activity theoretical notion of mediation (Bødker & Bøgh Andersen, submitted, Bertelsen & Bødker 2002, Bødker 1991, Engeström 1990). We do think that distributed cognition’s explicit commitment to understanding the propagations across informational/representational media constitute a minor analytical advantage in this particular study. However, our main focus in this paper is not on the theoretical underpinnings of our analysis as much as it is on the findings and their relationship with HCI (which is discussed e.g. in Hollan et al. 2000).

STUDYING TAX ASSESSMENT

Hutchins provides a “cognitive ethnography” (2003) according to which “the mind is not in the brain. The brain is in the mind.” The brain is part of a *body* participating in socially meaningful *activity* in a particular *setting*. In order to understand a cognitive system it is crucial to successively disentangle all these dimensions. In our case study we accede to this analytical framework, that help us elucidate the work practices. We opened the study in an ethnographical manner with shorthand observations of the process, formal conversations and interviews. This helped us identify crucial processes in the tax assessment.

Following distributed cognition we focused on the propagations of representational states across representational media, i.e. the tax forms, balance sheets, and control information. Not surprisingly, the crucial propagations turned out to take place at the desktops in the case officers’ offices. As such the office in general and the desktop in particular constitutes our unit of analysis. In practice we examined how information is distributed across the various representational media at the desktops and how this information is integrated during the work processes. This distribution and integration of representational media across time and space constitutes the *metric* of the tax assessment.

Obviously, nothing would get propagated without human intervention. Equally obvious however, is the fact that no tax assessment would take place without the artefacts. We make this point in anticipation of critics of a dehumanising perspective inherent in distributed cognition (Nardi 1996, 1998). This is not an anthropological study but a study of a certain activity *distributed* across humans and artefacts (see also Berg 1997 for an elaboration of our position).

Inspired by Jordan and Henderson's (1995) "interaction analysis" we recorded the work processes on video in order to track and capture the complex interactions between case officer and artefacts. Moreover, the case officers were asked to "think-aloud" while performing tax assessment. We needed this double perspective – both the case officers' verbalisations and the material manipulations – to account for the propagations across internal and external media. An important part of the analysis is to understand the representational structures of the central artefacts at the office. To do this we studied their informational content and how this defines the problem space in which the case officers search for solutions.

By combining video recordings, verbal protocols, and tax documents collected from each case we have been able to produce detailed *interaction protocols* of the process. We collected material from 10 assessments corresponding to roughly 1 1/2 hours of video recordings and 40 pages of transcribed protocols.

THE METRIC OF TAX ASSESSMENT

We have already established the *activity*, i.e. tax assessment. Figure 1 illustrates the overall layout of the office *setting*. As the figure illustrates the office is ordered according to some kind of input-output scheme with a mediating processing/work area. The pile of still unassessed files is spatially demarcated from the finished files (dismissed or assessed). As such the physical layout of the office (comprising documents, computers, furniture, calculators, etc.) underpins the informational process, thereby making the metric of the task more explicit. As noticed the central processing area – the desktop – is where the crucial computational processes take place.

Figure 2 illustrates the desktop in *standard configuration* observed at all the examined offices. It is important to note that as soon as the assessment begins the desktop is reconfigured according to the moment-to-moment needs of the computational processes. When comparing tax form and control information for example the case officer holds the relevant documents within proximity of the relevant information on the computer screen. However, the standard configuration is reconstructed frequently during assessments.

The desktop is ordered in centre and periphery regions roughly defined by the senso-motoric reach (i.e. the *bodily* constraints) of the case officer. The centre region is occupied by the tax form and balance sheet under scrutiny. The computer screen exhibiting control information is also within the centre region. On the border of the region lies the balance sheet from the preceding year. This document displays background or reference information. Further towards the periphery is the calculator, which is employed when needed.

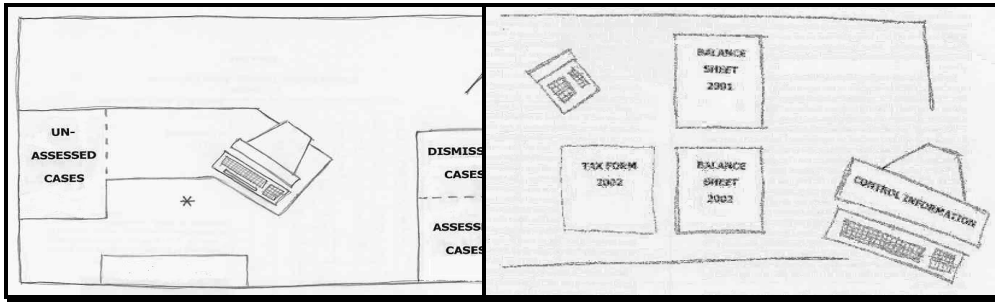


Figure 1: The overall layout of the office.

Figure 2: The desktop in standard configuration.

Having described the cognitive ecology of the tax assessment as it is constrained by activity, setting, and body, it is possible to turn to the actual computational processes facilitated by the case officer. Before moving to the finer details of these, we provide a rough overview of the tax assessment process. To acquire this overview we asked the case officers to categorise a number of the artefacts encountered at the office at their own discretion. They decided to organise the artefacts as they “fit into the work process”. Figure 3 illustrates the categorisation established by the case officers. We did make some omissions in the figure in order to obtain our analytical purposes just as we have tried to make the categorisation correspond to the 3 analytical concepts listed on the right. Below we comment on how the 3 concepts relate to the tax assessment metric.

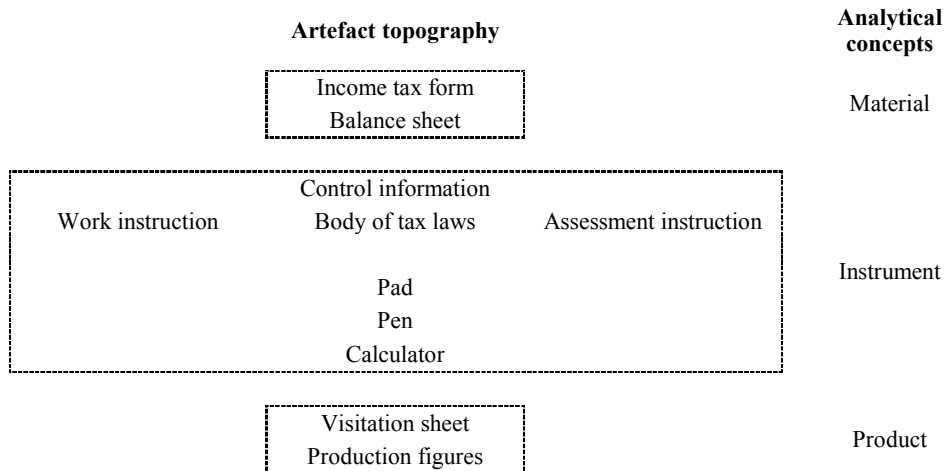


Figure 3: Artefact topography of the tax assessment

Product

The process that we are studying in detail here is part of a larger process; the annual cycle of the tax office. In this round, all tax forms are processed. The non-problematic ones are filed in the archive and not touched until next year when reaching a similar point in the cycle. The problematic or erratic forms are put aside for further processing, and classified. A tax officer later revisits them, though not necessarily the same one. When a case is found to be erratic the *product* of the assessment is a “Visitation sheet” classifying the erratic tax form according to level of urgency, expected time consumption, etc. On the sheet a short identification of the encountered problem is also annotated. The classification serves several purposes (see e.g. Bowker & Star 1999) – it is used both as a way of making visible which types of tax revisions the office carry out and how many (as established by the “production figures” in which each case officer registers the outcome of the assessment), and it is intended as a help for further processing. Hence, choosing the appropriate category is a negotiation of many concerns and it has consequences for the further treatment of the taxpayer, as well as of the actual tax forms and numbers.

Material

A number of manipulations need to take place before reaching the end-point of the process, however. When opening a new case, the officer performs a short read-through of the metadata concerning the results of the assessment of the previous year. This “unpacking” (Bannon & Bødker 1997) of information helps the case officer establish the context of the tax payer and acquire a rough overview on what kind of business it is, special issues to be aware of, etc.

Upon opening the case the officer places last year’s documents in the periphery of the work field. Subsequently, she takes the current year’s documents under consideration. The income tax form and the balance sheets are placed in the centre of the work field, and next to each other to allow easy propagation of figures. These documents constitute the *material* of the tax assessment. At first the

internal consistency of these documents is checked by comparing if the figures from the balance sheet equal the ones inserted on the tax form for instance.

Instrument

After acquiring a rough overview of the case and having placed the material within reach, the actual control of the current year's documents is initiated. Usually this implies examining whether there are any inconsistencies between the case material and the control information. We categorise the control information as an *instrument* because it is used to test the validity of the material. As such it serves a manipulative function commonly attributed to tools and instruments. Furthermore, the case officers themselves mention the control information (and their own knowledge and work instructions) among their primary tools/instruments. Interestingly, they never mention neither case files nor visitation sheets as tools.

Figures 4 to 6 reconstructs a particular assessment and exemplify a typical control process. The circles on the illustrations indicate the primary focus of the officer and what representational structures are manipulated at each step. The figures are accompanied by a short description of each step including extracts of the verbalisation provided by the case officer. At this point we withhold the theoretical concepts of distributed cognition and try to keep the description as "natural" as possible in order to let the reader receive an impression of the work process in its own right.




Figure 4: Establishing difference.

The case officer checks the control information to see if the system has discovered any inconsistencies between the information reported by the taxpayer and those registered as control information. This is possible because key information from the tax forms was transferred to the system earlier in the process. At this point the system has localised a difference in the reported and the registered interest expenditure amounting to 37.165 DKK.




Figure 5: Recalculating difference.

Next, the case officer examines how this difference was established. He consults the relevant pieces of interest information in the system and recalculates the difference between the registered and the reported figures: "Then we take 95.870, there, and up here 21.259 and 45.563 were registered, in total 162.692 [Total interest expenditure reported by tax payer]. And the system only knows about 32.655 and 92.872 [Interest expenditure registered in the system]. This means that there is a difference of 37.165. That was also what we were told."




Figure 6: Dissolving difference.

The established difference serves as a definition of the problem the case officer needs to solve. He searches through the balance sheets to see if they explain the difference. Soon the case officer localises a “loss due to depreciation” amounting to 37.313: “Here we got the explanation of the difference. It is very close to what we have got here [pointing at the display on the calculator].”

This picture fits the frame of distributed cognition well. A representational state (the difference of 37.165) is propagated across a number of representational media (control information, calculator, and balance sheet). By aligning these media in a process of re-representations a representational inconsistency is dismantled. The inconsistency is resolved by localising “a loss due to depreciation” not registered in the system. Here the officer’s knowledge of the body of tax laws enters the equation. In the wordings of one of the workers: “The system is stupid. (...) It just says that there is a difference from its perspective. We have to be able to tell if it is a genuine difference or not.”

The computations performed at the tax offices are highly unstructured, opportunistic, situated, and data-driven. The desktop reveals itself as a trade-off space where multiple – often opposing – mental, social, material, technical, and bodily constraints seek reconciliation: “The system composed of a task performer, mediating structures, and the task world settles into a solution that satisfies as many constraints as is possible. (...) This constraint satisfaction is a computation” (Hutchins 1995a, 300). We use the term “desktop computing” because the desktop serves as an *interface* between the individual case officer and the computational system at large.

This is where the analytical framework of distributed cognition leaves us. By focusing on the construction, manipulation, and propagation of representational media through the system we acquire a detailed account of the workings of the tax office. We argue that this framework grasps the *deep structures* of the work processes. It omits large parts of the activities at the office only to focus on the minimal constraints, that are crucial to the tax assessment.

We still need, however, to go beyond this descriptive stance to challenge the current exploitation of the desktop metaphor in computer applications. The analysis presented above points us towards structures and processes which are not well supported in the WIMP regime persistent today. This perspective is crucial because the tax office is under pressure for “going electronic”, i.e. to digitalise all case material.

CHARACTERISTICS OF THE TAX DESKTOP

In the following we summarise and discuss the findings from the empirical analyses. Some of these observations are evident from the example above. Others are collected from the video recordings and discussed further in (Nielsen 2004). We focus on the salient characteristics of the tax assessment as desktop computing, partly because the desktop has turned out to be a useful delimitation of the analysis of propagation and representations, and partly because we find it intriguing how different this tax desktop is from the standard WIMP desktop. We will now turn to the latter of these points while sidetracking the analytical framework of distributed cognition. The theory served its purpose, i.e. understanding the desktop computing, but now we widen the perspective again to make room for observations not well accounted for in terms of manipulation of representations.

Opening and closing of cases

The case officers rely on highly institutionalised albeit largely unarticulated rules for handling the cases. All the case officers we visited open and close their case folders in a uniform fashion. The files from each year are ordered in separate folders. Furthermore, the folders are rotated so that one year's documents are up side down to the following year's documents. It is important to the case officers to preserve the orderliness of the files, which is probably the reason for sustaining some documents in specific areas to avoid mixing documents from different years. Of course date stamps help maintaining the divisions, but by delegating parts of the responsibility to well-established spatial boundaries the distinctions are no longer in need of conscious deliberation. As such, the physical orientation of the folders supports informational distinctions written on the folders.

The browsing of the documents, where last years documents are studied first in their location at the back of the desk, follows the opening of the files. Then the tax officer turns towards the actual tax statement of this year, and its accompanying documents. When Bannon & Bødker (1997) talk about the packaging and unpacking of information this is done rather metaphorically, whereas in this case we see that this process is actually thoroughly embedded in the physical/bodily process of laying out documents and orienting them in a certain order.

The senso-motoric constitution of desktop

The physical layout of the desktop enforces certain bodily exercises. As evident from the interaction protocol presented above (figures 4-6) the case officer transports figures between computer screen and calculator. Having placed the calculator and the computer screen in opposite ends of the desktop, the case officer is gripping the table with one hand, moving closer to the monitor to read figures, drawing himself towards the other end of the table to punch in the figures on the calculator. In this case the case officer utilises his body to connect the two instruments.

In other cases, however, the case officer reconfigures the desktop to minimise his bodily engagement. When comparing the turnover figures from the balance sheet to the control information for instance, the officer places the relevant figures from the documents within proximity of the relevant figures on the computer screen. This supports direct, visual propagation of representational states across external (non-mental) representational media.

The case officers often mark relevant bits of information with fingers or pens while consulting figures in other parts of the documents. It seems to be crucial to the case officer to be able to keep track of a large number of entry points simultaneously. We have seen a tax officer using four fingers on the left hand to point at various documents, while pointing on the computer screen with his right hand.

Furthermore, when comparing the files of last year with those of the current year, the tax officer compares the material to appreciate if the "look and feel" is the same as last year. If something looks very different it raises his attention. Evidently, it is important to understand the role the body plays in the constitution of the desktop.

Physical annotations/marks

We have noticed elaborate institutionalised rules for making physical indications directly in the case material. Some signs mean that the figures match, some that they do not. It is also common to write annotations in the case files. These notes provide the case officers with salient clues, which aid the analysis and interpretation of the documents (for themselves and for others) later in the process.

With the visitation sheet presented earlier the case is categorised with respect to what steps to be taken next. This categorisation is accompanied by a short note indicating the problem. The

categorisation primarily serves to prioritise cases for the later processing whereas the notes are meant as a help to the later interpretation of the documents when they are returned to. The physical annotations serve the double purpose of helping the here and now individual information processing, as well as they serve as organisational annotations (Hertzum 1999) as they are distributed over time and tax officers (Bannon & Bødker 1997).

Functional space

The seemingly inconvenient constant manipulation and rearranging of the artefacts implies that the case officers reserve specific areas for specific functions. Apparently, one region of the desktop is for calculations; another region is for reference information; a third region is for the momentarily handled documents, etc. Rather than just segmenting the table according to reach and sight there seems to be salient functional demarcations as well. E.g. despite noticeable inconvenience, the form from last year is read at the back of the table and left there, not pulled toward the reader. The case officers are, in other words approaching the materials where they are, rather much in the fashion described by Bertelsen & Bødker (2002), which they call zooming with the feet. Obviously in this case, there are no feet involved, yet the tax officer still moves around the desktop approaching the right information in its particular location – in some instances this entails substantial bodily engagement as described above.

Case logic

Oftentimes the case officers discover inconsistencies when reading through the balance sheets. Whenever they come across figures that for one reason or another look suspicious they consult relevant control information. Instead of having a predetermined list of controls to execute every time, they adapt convenient strategies that comply with the internal logic of the case. As evident from the interaction protocol above (figure 5) the tax officer tries to reconstruct the calculation processes of the taxpayer in order to appreciate this internal logic of the case files. The calculations seem to be done to check how the difference was established, rather than to check the numbers as such.

It seems to be important to the tax officer to be able to access, manipulate, and rearrange the representational structures according to the immediate needs for propagation. As such, the desktop undergoes constant reconfiguration to support the case logic. This makes it impossible to determine a fixed algorithm that is applicable in each case. The actual computational process at the desktop is a result of negotiating the constraints imposed by the case logic, the material layout of the desktop, the individual strategies of the case officers, and the institutionalised rules for carrying out tax assessment. As such, the algorithms for the propagation of representational states across representational media can only be spelled out *post hoc*. However, as discussed by Bardram (1997), with inspiration from Suchman (1987) rules and procedures are resources for the situated action, and need to be seen as such. They are not systematically imposed on cases, but activated whenever needed.

CHALLENGING THE DESIGN OF THE COMPUTER DESKTOP

In the following we want to utilise the insights gathered from the case study to explore alternative desktop setups in future digitally enhanced settings. We want to maintain crucial structures exhibited in the current work processes at the tax office while enhancing the computational processes with some of the benefits from digitally supported work settings.

We have performed small-scale design experiments focusing on possibilities for and consequences of digitalization of the tax assessment process. Most prominently these studies suggest that

digitalisation is unlikely to provide reasonable benefit to the case officer. Paper seems to be our preferred media for reading and digesting large amounts of textual information, primarily because screen resolution limits the number of documents that can be showed side-by-side on a computer screen (even if this is projected onto a large screen display).

In this article it is not our aim, however, to discuss design possibilities for an IT-augmented taxation process *per se*. Neither will we discuss solutions applying augmented reality technologies, such as tagging and tracking of paper documents as interface to a digital solution. Instead we will take the desktop metaphor of computer applications literal and confront it with the tax desktop. In rather simple terms we see so many situations where new IT solutions continue to be implemented without questioning the current WIMP regime. In the light of this we will explore our findings. In addition, we hope that our analysis will contribute to the ongoing reformulation of alternative models of interaction, from post-WIMP to virtual and augmented reality.

We have chosen to focus on three elements:

- The spatial arrangement of the desktop
- Traces of context and processes
- Tailoring for joint practices

The spatial arrangement of the desktop

Most strikingly, our analysis shows that the physical desktop spaces are used in various ways other than just as random piles and documents on the desktop. Throughout the taxation process, each document has a location on the desktop. Other resources have locations in a similar manner, the calculator for instance. Exactly the electronic information differs from this, because it is brought up on the computer screen in windows. The physical desktop offers functional spaces that are used to separate and structure the taxation process: last year's documents at the back of the desk, etc. As we have illustrated, however, not all documents stay in one location. Some are moved around as the taxation process takes place, and spatial mobility is in this way essential.

What happens if we look at these findings as a "requirement specification" for the design of an electronic desktop? In most of what we recognise as a WIMP interface, windows are totally random, disconnected and overlapping windows. Over the years we have seen several examples of interconnected spaces and windows: The Smalltalk browser allowed for interrelated inspection of program classes and methods; Henderson & Card (1986) created ROOMS to provide an activity-centred organization of documents, and e.g. TOPOS (Bücher et al. 2002, Ørbæk & Nielsen 2001) provides means of structuring collections of documents of various types according to activity. While TOPOS does provide activity-centred collections of documents, these collections provide a bottom-up structuring of particular documents. In our case on the other hand, we are looking for a structuring of the desktop that is shared across cases, but for a particular activity. I.e. we need each case and its documents to open on the desktop, structured to depend on the step in the taxation process that is currently taking place.

In contrast to such solutions, Bertelsen & Bødker (2002) made a study of an object-oriented debugger where each break would launch a cascade of windows that were not visually or spatially connected. In the study, it was pointed out how the cascading caused a lot of confusion to the users, not least over time (as we shall return to below). We know of no applications where the screen/desktop space is pre-destined to support a particular physical arrangement, such as current case material to the right, last year's material to the left. As a matter of fact this would probably be very difficult to build using WIMP. This is both because the desktop interface is mainly seen as background, not as an active structuring mechanism, and because some sort of classification or typing of document contents is needed. However, we do see a possibility in revisiting some of the

various split screen and tiling mechanisms that were applied in the early years of the PC, since these are closer to the organization that we want.

It is furthermore important for the tax officer to be able to see and point to several documents or particular pieces of information at the same time. This is a further challenge to the ruling Windows regime, which basically supports only one cursor and one active window.

Traces of context and processes

In the study we see that packing and unpacking are processes that are integral parts of the taxation process; they are physical and take time. The relationship between documents are created and maintained in this process: The history of the case as such is maintained because the separation and ordering of year is maintained when the material is laid out on the desk, and repackaged into the file. This ordering exists alongside the ordering that is created as part of the information processing of the task at hand.

This is in contrast to the above mentioned debugger study where each new step in the debugging creates a new cascade of open documents, and where, after a short process, the connection between documents and the ordering of these in time are lost. The challenge is accordingly twofold: to provide mechanisms through which the history of placement and interconnectedness of documents can be maintained and retrieved. And to support the processes of a particular unpacking, which is the process through which the tax officer gets her initial understanding of the case. We project that this means that the user is supported in over-viewing documents as they have been left on the desktop, and opening them one at a time in the order preferred by the individual tax officer.

The tailoring for joint practices

WIMP interfaces offer far-reaching flexibility for setting up the desktop for individual needs—as many documents and applications as one may desire on the desktop (or none), launching menus to the top, bottom or to the right, etc. While these elements may be brought together into a somewhat standardised set-up, what the tax setting demands is for the desktop to fit a particular element of the entire taxation process, in a manner that reflects the common practices of the office, shared by all and by all cases. At the same time the desktop setting needs to be changeable by the individual users (in general for all cases, and as part of the processing of a particular process).

What we are after is an activity-oriented tailoring of the entire desktop. By tailoring we mean that the set-up can be created by somebody, shared by the whole group of tax officers, and adapted to their individual needs (by themselves or by somebody else) (Gantt & Nardi 1992, Trigg & Bødker 1994). This is in essence very different from the support for the individual to produce a personal desktop look and feel that lasts over time which WIMP provides today.

CONCLUSION

In this paper we have examined a tax assessment as it takes place at the (physical and virtual) desktops of the case officers. The processes observed are targets for commercial software companies regarding the potential for “going electronic” and digitalising case material. In order to recognise potential pitfalls of this process we performed a detailed field study of the tax office. We found the analytical framework of distributed cognition valuable in this study presented here, because it helps us understand the desktop computing at the tax office. Furthermore, this concept allows us to rethink interfaces between the individual and the computational system at large.

With the above discussion we hope to have illustrated how the study of a rather mundane work process such as this one challenges our conception of the desktop – the physical one, and not least

the computer desktop. Even if we stay away from augmented reality solutions and stay with computer programs that run on the computer desktop, taking the challenges seriously means that we need to revisit several taken-for-granted elements of the current WIMP regime: the randomly overlapping windows on a structureless background; the non-existing traces of time and past location; and the individualised and non-activity-oriented set-up of the desktop.

REFERENCES

- Bannon, L. and Bødker, S. (1997) "Constructing Common Information Space" in J. Hughes, T. Rodden, W. Prinz and K. Schmidt (eds.) **ECSCW'97: Proceedings of the Fifth European Conference on CSCW**, Kluwer, Dordrecht.
- Bardram, J. E. (1997): "Plans as Situated Action: An Activity Theory Approach to Workflow Systems" in **Proceedings of the 5th European Conference on Computer Supported Cooperative Work**, Lancaster, UK. Kluwer Academic Publishers.
- Bellotti, V. and Smith I. (2000) Informing the design of an information management system with iterative fieldwork, **DIS 2000**, 227-238.
- Berg, M. (1997) Of Forms, Containers, and the Electronic Medical Record: Some Tools for a Sociology of the Formal, **Science, technology, & Human Values**, 22, 403-433.
- Bertelsen, O. and Bødker, S. (2002) "Interaction through multi-artifacts" in S. Bagnara, S. Pozzi, A. Rizzo and P. Wright (eds.) **ECCE 11 - Cognition, Culture and Design Eleventh European Conference on Cognitive Ergonomics**, Rome: Istituto di scienze e tecnologie della cognizione consiglio nazionale delle ricerche.
- Bertelsen, O. and Bødker, S. (1998) Studying programming environments in use: between principles and praxis, **NWPER '98 The Eighth Nordic Workshop on Programming Environment Research**, Bergen/Norway, Sunday June 14-Tuesday June 16 1998, URL <http://www.ifi.uib.no/konf/nwper98/proceedings.html>.
- Bowker, G and Star, S.L. (1999) **Sorting things out**, MIT Press, Cambridge Massachusetts.
- Bødker, S. (1991) **Through the Interface – a Human Activity Approach to User Interface Design**, Hillsdale, NJ: Lawrence Erlbaum Associates.
- Bødker, S. and Bøgh Andersen, P. (submitted). **Complex mediation**. Paper submitted for publication.
- Büscher, M., Christensen, M., Grønbæk, K., Krogh, P., Mogensen, P., Shapiro, D and Ørbæk, P. (2000) Collaborative Augmented Reality Environments: Integrating VR, Working Materials, and Distributed Work Spaces, **Proc. Conf. on Collaborative Virtual Environments (CVE 2000)**, San Francisco, 9-12 September, ACM Press, New York, 47-56.
- Engeström, Y. (1990) **Learning by Expansion**, Orienta Konsultit, Helsinki.
- Gantt, M. and Nardi, B. (1992) Gardeners and gurus: patterns of cooperation among CAD users, **Conference on Human Factors in Computing Systems**, 107 – 117.
- Hayashi, K., Nomura, T., Hazama, T., Takeoka, M., Hashimoto, S. and Gudmundson, S. (1998) Temporally threaded workspace : a model for providing activity-based perspectives on document spaces. Conference on Hypertext and Hypermedia. **Proceedings of the ninth ACM conference on Hypertext and hypermedia**, 87–96.
- Henderson, D.A. and Card, S. (1986) Rooms: the use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface, **ACM Transactions on Graphics**, 5(3), 211-243.
- Hertzum, M. (1999) "Six roles of documents in professionals' work" in S. Bødker, M. Kyng, and K. Schmidt, **Proceedings of the sixth european conf. on computer supported cooperative work**, Kluwer, Dordrecht.

- Hollan, J., Hutchins, E. and Kirsch, D. (2000) Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research, **ACM Transactions on Computer-Human Interaction**, 7 (2), 174-196.
- Hutchins, E. (1995a) **Cognition in the Wild**, MIT Press, Cambridge Massachusetts.
- Hutchins, E. (1995b) How a Cockpit Remembers Its Speeds, **Cognitive Science**, 19, 265-288.
- Hutchins, E. (2003) Cognitive Ethnography, presentation at Cognitive Science Society Conference, Boston, 01.08.2003.
- Jordan, B. and Henderson, A. (1995) Interaction analysis: Foundations and practice, **J. Learn. Sci.**, 4, 1, 39-103.
- Kirsh, D. (1995) The Intelligent Use of Space, **Artificial Intelligence**, 73, 31-68.
- Malone, T. (1983) How do people organize their desks? Implications for the design of Office Information Systems, **ACM Transactions of Office Information Systems**, 1(1), 99-112.
- Nardi, B. (1996) "Studying Context: A Comparison of Activity Theory, Situated Action Models, and Distributed Cognition" in B. Nardi **Context and Consciousness – Activity Theory and Human-Computer Interaction**, MIT Press.
- Nardi, B. (1998) Concepts of Cognition and Consciousness: Four Voices, **Journal of Computer Documentation**, 22 (1), 31-48.
- Nielsen, M. (2004) **Desktop Computing**, unpublished master thesis, University of Aarhus.
- Ørbæk P. and Nielsen, M.B. (2001) Finding Hyper-Structure in Space: Spatial Parsing in 3D, **The New Review of Hypermedia and Multimedia**, 7, 153-183.
- Shneiderman, B. and Plaisant, C. (2005) **Designing the user interface**, Addison Wesley.
- Suchman, L. (1987) **Plans and Situated Actions – The Problem of Human-Machine Communication**, Cambridge University Press, Cambridge.
- Trigg, R. and Bødker, S (1994) "From Implementation to design: Tailoring and the emergence of systematization in CSCW" in R. Futura and C. Neuwirth **Proceedings of CSCW 94**, ACM press.
- van Dam, A. (1997) Post-WIMP user interfaces, **Communications of the ACM**, 40(2), 63 – 67.

ACKNOWLEDGEMENTS

Thanks to Michael Christensen for inviting us into the world of TOPOS and for sharing of ideas, to the tax officers and management of the Tax Office in Aarhus, and to Anders Thomsen and Ann Mølhave for useful comments.

COPYRIGHT

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