Alan Dix, Julie Wilkinson and Devina Ramduny (1998).
Redefining Organisational Memory - Artefacts, and the Distribution and Coordination of Work (extended abstract). *Workshop on Understanding work and designing artefacts*, York, 21st September 1998.
https://www.alandix.com/academic/papers/artefacts98/

Redefining Organisational Memory – Artefacts, and the Distribution and Coordination of Work

Extended Abstract

Alan Dix[1], Julie Wilkinson[2], Devina Ramduny[1]

[1]School of Computing, Staffordshire University, Stafford, UK
{A.J.Dix, D.Ramduny}@soc.staffs.ac.uk

[2]School of Computing & Management Sciences,
Sheffield Hallam University Sheffield, UK
Julie.Wilkinson@shu.ac.uk

https://www.alandix.com/academic/papers/artefacts98/

Offices are full of paper: on desks, in filing cabinets, in drawers, in briefcases, perhaps even in piles on the floor. Many studies have shown *that* paper and other physical artefacts are essential for effective work. We show *why* these artefacts are important and how they form a key part of organisational memory. Using the metaphor of the organisation as a computer we see that artefacts function both as *triggers* to initiate action and also as *placeholders* within business processes. Understanding the multiple roles of artefacts is essential to avoid problems when redesigning the physical or electronic representation of documents.

Keywords: CSCW, coordination, cognition, office ecology, organisational memory, workflow

Introduction

In the 1990s the most common artefact in any office is still undoubtedly paper-based documentation. When conducting a workplace study of any kind the review of office documents is regarded as a useful data collection exercise (e.g. Orlikowski, 1992). In short, such artefacts play a key role in aiding our understanding of work – especially the flow of work. Yet what happens to those artefacts when their character is altered?

This paper explores the importance of artefacts with regard to their multiple roles in the workplace – not only do paper documents have information written on them, but they embody other functions in their very physicality and location. A letter in an in-tray says "work to be done", but the same letter in a wastebin simply says "empty me". Without this understanding of the roles of artefacts it is easy to redesign a business process or system and unintentionally lose the function of one of these roles.

In this paper we propose a framework such that interactive systems designers can appreciate the roles of artefacts and in so doing ensure that, when a paper-based system is redesigned or re-engineered, these roles are either embodied within the artefacts' new representation, or are fulfilled by other parts of the system.

Central to our analysis will be the metaphor of the computer to understand organisational information systems. This metaphor has been invaluable for understanding human thinking within cognitive science. This will help us to focus on the different kinds of memory within an organisation. In particular we will see that placeholders are essential – that is the ability to determine where you are in any process.

Whereas organisational memory research is usually focused on long-term memory of procedures and rationale, we will see that short-term organisational memory is equally important.

Background

Previous research studies of office activity have recognised the importance of the physicality of work artefacts - a specific element in the ecology of the office (Rouncefield et al., 1994; Sellen and Harper, 1997). In our own work we have sought to identify how the advantages of such artefacts can be preserved during system redesign (Dix et al., 1995, 1996 and 1998). Our particular focus has been on artefacts as triggers to action – initiating activity and driving processes to successful completion. The very existence of an artefact in a particular location acts as a reminder (Dix et al., 1995) of the next activity needed in the work process. This was achieved by analysing the flow of work at the micro level and explicitly asking what event triggered each activity to run.

Our original objective was to identify any characteristics that ‘weak’ triggers had in common and to verify that we had a complete classification of trigger types. A pleasing result of this work was the emergence of a generic pattern within work processes when decomposed – namely the ‘4 Rs’ (Request, Receive, Respond, Release). We considered this framework useful in structuring our understanding of work processes which reinforced the analytical approach of our method of investigation. In our first case study (Dix et al., 1995) a new system design was not part of the remit. However, for a second (Dix et al., 1998) major interactive systems design was intended and the framework helped to balance the tensions that fieldworkers frequently face between providing accounts of the ‘sociality of work’ and usable design recommendations (Plowman et al., 1995).

One of the theoretical underpinnings of our work has been status–event analysis, which allows human and computer phenomena to be discussed within a single analytic framework. In particular, it has highlighted common features between triggers, initiators of human action, and the propagation of events between computational agents (Dix et al., 1998; Ramduny et al., 1998). These strong parallels between human and computer systems have prompted us to look more closely.

The computer as metaphor

Since the early days of computation, the parallels between human cognition and computer systems have been a powerful source of ideas. On the one hand, most modern models in cognitive science lean heavily on computational analogies (e.g. ACT* (Anderson, 1983), ICS (Barnard 1985), SOAR (Laird et al., 1987)). On the other hand, applications of artificial neural networks have revolutionised many problems which were previously thought unsolvable.

In a similar vein, we can look at an organisation as an information processing system performing computation, or even as a locus for distributed cognition (Hutchins and Klausen, 1991; Lave, 1988), both of which suggest the computer as an appropriate metaphor.

The Socio-Organisational Church–Turing Hypothesis

In talks we referred to this parallel between organisational structures and computer architecture as the 'socio-organisational Church–Turing hypothesis'. In computer science the Church-Turing theorem proves that two early models of computation, the Church lambda calculus and the Turing machine, are equivalent. The Church–Turing thesis/conjecture is the (unproven) hypothesis that in fact any form of computation is essentially equivalent. This is all about *what* can be computed. We effectively conjecture further that there are likely to be parallels, and hence useful cross-inspiration, between the internal architectures of computers and that of both cognition and organisational structure – that is *how* computation is performed. This is effectively the assumption that has underpinned much of 'natural computing' where naturally occurring processes are used as inspiration for computer architecture and algorithms. Here we are doing the opposite looking at computer systems and using that as inspiration to explore features of organisational structures, processes and practices

The most obvious two elements in a computer are the program and data.

Take the organisation as a computer program. Most of us know what a program is - its text, its form; usually expressed in formal, explicit structures. The organisation too has its own text and form: some aspects set down explicitly in human procedures and in automated information systems; other aspects more loosely defined in people's minds. One of the roles of business process thinking is to (re-)design (or perhaps recode) this 'program' for the organisation.

The data in an organisation is also easy to see, captured in paper files and electronic databases. Indeed, the specification of this data is a central aspect of all system design methods. Although some of this information is also in people's heads, it is probably more likely to be captured than are processes. Consequently most organisational memory research is focused on externalising the process and rationale within an organisation, the 'data' part of memory is more often externalised to start with.

However, the story does not end with program and data. As well as data structures and disk files, there is a dynamic state of the computer: the program counter and temporary values stored deep within the computer's CPU. These are rarely seen by the programmer and only made visible by debuggers, or during testing if the programmer inserts 'print statements' to produce screen output that gives some feedback as each mini-task is fulfilled by the program. This dynamic state corresponds to the human short-term memory and the activation level of the long-term memory which together influence our on-going thinking and actions.

The difference between temporary variables in registers, RAM and long term data storage is mainly an implementation consideration, functionally they are simply part of a continuum of access speed vs. longevity. However, the program counter serves a fundamentally different role within the computer. It is the implicit expression of where the program is at any particular time. Similarly the dynamic state of any organisation needs to be maintained, not just the information needed to work, but the knowledge of what should be done next, where the organisation stands in each of its processes.

Imagine you are starting a new office job. No training or information is provided and no social communication with co-workers is permitted. How would you know what work needed to be done? Albeit an unlikely scenario, could you cope? Documents would prove a useful source of information: written procedures and the rubric on forms could explain the work processes. If a phonecall came asking for the price of a product you could (after some searching) find the right file in the filing cabinet and give the information.

But that is not enough. What about on-going processes such as partly processed orders? What you need is the ‘state of play’ of those processes. This may be supplied by the state of an artefact: whether the leave request form is signed or not, or may be implicit in the specific location of an artefact: whether it is in the in-tray or the out-tray. Compare this to walking into the middle of any board game – you may recognise what the game is; you may even know the rules of the game; but can you tell at a glance? Who is winning? What is the next move? What is your role in the game? To answer this you need to see the counters. Similarly an organisation must have such placeholders to record the 'state of play’ of its processes, even if this is simply each person keeping mental track of what they are doing.

It is interesting that when we first discussed these issues among ourselves, we found no obvious word for this 'program counter' for an organisation or in processes in general.

Finally, we must not forget triggers. For any active system there must be something that initiates activity. In our previous work we have seen that for

office processes this can include phone calls from other people, reminders in a diary or to-do list, regular routines and of course environmental cues – artefacts which prompt action. In cognitive science we can think of stimulus–response behaviour which is of a similar nature, although we also pursue spontaneous action, when we randomly remember. In a computing context, the equivalent is low-level interrupts and also event-driven programs as found in most graphical user interfaces.

Note the difference between placeholders and triggers. Both are essential and often may be achieved in the same artefact. For example, a letter sitting in the in-tray says "do something" (trigger), by being in the in-tray it says "I need to be read" (placeholder), and of course also contains information (data). In contrast, a lawyer might have a note in her diary telling her to work on the 'Richardson case' (trigger), but she needs to extract the file to find the current state of the case (placeholder). The essential difference is that triggers tell you *that* you need to do something whereas placeholders tell you *what* you need to do.

	computer	**cog. sci.**	**organisation**
process	program	procedural memory	processes
data	data	LTM	files
placeholder	program counter	STM/activation	location of artefacts
initiative	interrupts, event-driven programs	stimuli	triggers

Table 1 – computational parallels

event-based programming

It is interesting to note that interrupts, the equivalent of triggers in computing, are often regarded as an esoteric subject, for low-level hackers only! This is because traditional programming has had the program in charge and hence the trigger for executing a piece of code is simply that it is the next thing to do. There is one program counter and you always do it next. An organisation is more complex with many people working simultaneously and each being involved in many active processes. Hence it is natural to consider multiple placeholders with the consequent issue of which to pursue next. It is only recently, as the widespread implementation of graphical interfaces has become commonplace, that event-based programming paradigms such as Java AWT or Visual Basic have brought the issue of multiple placeholders and initiative into mainstream computing. Sadly, event-based programming is still very much an afterthought in many computing syllabi.

Physicality and distribution

The physical nature of any artefact of work matters. For instance, papers on a desk are bulky, highly visible and familiar. Change these to on-screen forms and they 'lose size', become both less visible (even invisible if the network crashes!), and still for many people, a less comfortable way of working. (How many of us can honestly say we have never printed an email?). The placeholding role of the artefact is adversely affected. Go a step further and remove that artefact altogether – the 'clean desk' syndrome. We are now in danger of losing an external representation of information in the organisation. The loss of physicality of the artefact, therefore, has consequences for the working memory of the organisation.

The very physicality of paper artefacts makes them especially suitable to fulfil several roles. A collection of documents, therefore, can explain the requirements of the job, be a placeholder for where you are and also contain useful information. However, the problem is that an outsider viewing the artefact may see that it represents the process but not necessarily the current state, or treat it solely as an information source. On the one hand, these multiple roles of such artefacts are efficient and flexible, on the other hand they can have potentially negative consequences for the smooth-running of an organisation's operations, especially when there is considerable staff movement or when artefacts are carelessly redesigned.

A further complication arises when a process is highly distributed both in time and space. The flow of work will often contain many interactions as different tasks are completed. The more complex the process, the more probable that changes will occur in the control of that process. Highly distributed flows of work can swiftly create a sense of loss – loss of control of the overall process. The many interactions involved may happen at varying pace and over long periods. We all know how easily we can be overwhelmed by the number of tasks in which we are engaged. In previous research we named this issue 'interaction-in-the-large' (Dix et al., 1998).

Paper documents play an important part in such a process, but their physicality also limits them – they must be in one place at a time! Hence the profusion of photocopies and multi-part forms in many organisations. Here a well-designed computerised artefact may have an advantage over its paper-based counterpart. The loss of physical paper documents from a desk is acute. The ability to see those same documents electronically at any time can therefore help maintain control of complex processes.

This is not to say that electronic solutions are easy. Indeed, many of the problems faced by people collaborating remotely are also encountered in distributed computer algorithms!

Distribution also highlights the issue of where control lies in a process and how it is expressed. We can imagine poles. On the one hand there may be a *global* view of a process, as each person involved in the process completes a stage the fact that it has been done is fed back to the central controller who then keeps track of it and may direct another person to continue the work. On the other hand control may be *local*, distributed amongst many people and places. For example, when James in Sales receives a two-part pink/green

order form he has to check the item costs and then pass the pink copy to Finance and the green copy to Despatch. When Sally in Despatch gets a green form she picks the items and packages them etc. No-one sees the whole process, but it works by the combination of local actions.

Both global and local control can work effectively, but they mean that organisational memory may be located in different places.

In search of organisational memory

It is 11 o'clock at night and the office is in darkness. The office workers have all gone home, the cleaners have left and only the security guard walks quietly through the dimly lit corridors.

Where is the memory of the organisation now?

From a distributed cognition or situated action viewpoint, one looks for the activity of an organisation in the ecology of interaction between people and artefacts within their working environment (Hutchins and Klausen, 1991; Lave, 1988; Suchman, 1987). During the day it is hard to disentangle these, but at night this subtle web of interaction is torn apart. We cannot simply place the organisational memory vaguely within the interaction – there is none. If anything carries over from day to day, there are only two alternatives:

- internal – within the people sleeping in their beds
- external – within the physical structure and contents of the office itself,
  either explicit: written down, the marks on paper,
  or implicit: the location of artefacts, subtle cues

Our earlier example of starting a new job in an office emphasises just how important the former is. It is typically not easy to pick up from scratch which implies that a significant part of an organisation's memory does go to bed each night with its employees. Any of the kinds of memory we have considered: process, data and placeholders may be held in people's heads. This is of course very serious for the organisation when sickness strikes or an employee gets a new job elsewhere. There are thus good reasons to ensure that this internal memory is externalised.

Some internal memory is inevitable: we do not expect to codify all the experience of a lawyer or academic. Indeed, this is the thrust of Alison Kidd's seminal paper "The marks are on the knowledge worker" (Kidd, 1994) – that knowledge is not accessed by such workers from files and records, but is embodied in the constantly changing nature of their thought processes.

As we have already noted, data memory is typically externalised well in organisations, stored in electronic documents, records and filing cabinets filled with paper. In contrast, processes may well be more diffusely held, especially how to cope with exceptions and infrequent cases and a major thrust of organisational memory research is to address the externalisation of this information.

In contrast, placeholders are easier to overlook and most likely to be held internally in people's memory or implicitly. Explicit representations of

placeholders exist, project planning do this for large scale tasks, to-do lists are a more informal form and standard office procedures are captured within workflow systems. However, the more explicit workflow systems are often criticised for being too restrictive. The reliability of greater control comes at the price of reduced flexibility. It is interesting to note that the placeholders in a workflow system closely resemble an actual computer program. In contrast the placeholders in cognitive models are often in the form of diffuse activation levels over memory, similar to the implicit placeholders found in many offices: different piles on the desk, selected papers placed at angles etc.

Imagine again it is 11 o'clock at night and you are one of the office workers. While taking your evening shower you slip and bang your head on the soap dish. Mildly concussed you are taken to hospital. Happily, the doctor declares you fit and able to go back to work straight away, except that you have partial amnesia – you can remember nothing that has happened in the past week.

- Will you encounter any difficulties?
- Will it be easier for you than for a new member of staff?

If you rely on your memory for process, data or placeholders, then the answer to the first question will be yes, but no more or less difficult than for a totally new member of staff.

If everything is external then you should be able to pick up where you left off. You peek in your out tray and don't recall writing the memo in it, but you shrug your shoulders and continue with the first item from your 'pending' pile on the left of your desk. In contrast, the new employee faced with your desk may well understand the in-tray and out-tray, but won't know that the pile on the left should be dealt with urgently and that the items on the right can wait.

The crucial difference between you and the new employee is *interpretation*. Where processes and placeholders are explicitly represented in the external environment, or where accepted standard interpretations are used (e.g. the in-tray) there is no problem, but as soon as the meaning is defined by the micro-culture or the individual (as with the 'pending' pile), the new employee falters.

In organisational memory interpretation is central, but is hardest to capture as it is implicit and assumed within the human and cultural context. But interpretation is central to understanding the meaning of all other externalised memory, from the most informal implicit cues to formal documents.

At first interpretation seems purely a human characteristic, but in fact any external signs need to be interpreted, whether explicit or implicit. If you try to run a Macintosh application on a PC you wouldn't expect it to work! Interpretation is even important within your own head. One of the effects reported by those with forms of autism mild enough to allow effective communication, is an inability to interpret the world around them. They see things, line, shapes, colours, faces, but that is all, they hold no meaning.

Interpretation is thus not just essential for organisations, but in any rational thought or computation. We must therefore add interpretation to the four elements of computation we discussed earlier.

```
2123 752f 7273 6c2f 636f 6c61 622f 6e69
702f 7265 0a6c 700a 6972 746e 2220 6548
6c6c 2c6f 7720 726f 646c 5c21 226e 0a3b
650a 6978 2874 2930 0a3b .... .... ....
```

A matter of interpretation – Macintosh or PC, or is it Java?

Summary of the design space

We now have 5 elements of computation:

- process
- data
- placeholder
- initiative
- interpretation

The first three of these form part of an organisation's 'memory', but may be held in very different ways:

- internal – in people's heads
- external explicit – in the marks on paper and electronic documents
- external implicit – in the location and attributes of artefacts

In particular, placeholders, when externalised, are often represented implicitly.

Given this we can characterise any process by its mode of memory representation for it and its placeholders. A workflow system has both an explicit (and global) representation of the process and an explicit (and global) representation of the placeholder. In contrast, an ad hoc office, where people do things as they occur to them would be at the internal/internal end. In general as one moves towards more external and explicit representation there is greater control over the process, but less flexibility for individuals. When designing a process it is important that one makes the right choice within this spectrum.

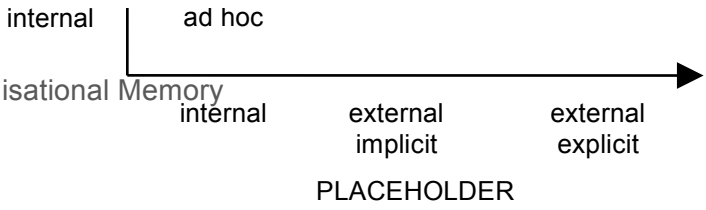

As an artefact may take on several roles, we can use a similar diagram with axes for each of its roles. For example, a file is an explicit external representation of data, but if placed on the desk may be an external implicit placeholder:

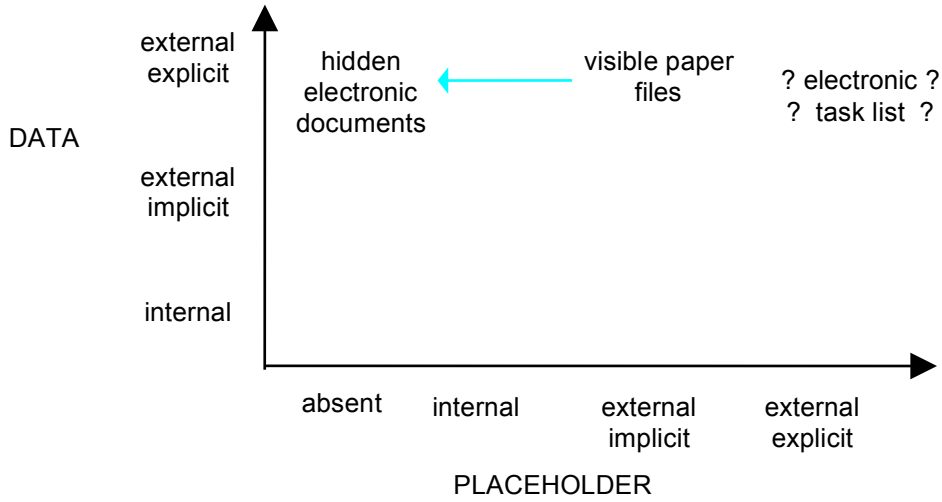

The example of making paper documents electronic can be seen here. The paper documents are an explicit external representation of data, but their presence on the desk acts as an external implicit placeholder. By making the documents electronic they become invisible and thus whilst still being an external explicit representation of the data (and perhaps more easily searched), they no longer serve a placeholder role, thus leaving a void in the process (absent placeholder). Seeing this we can start to think of solutions. One obvious solution is to represent the electronic document in such a way that the external placeholder role can still be satisfied. For example, it may be possible to lace it on the electronic desktop, similar to the physical use of location. Alternatively one might separate out the roles and add an electronic task list, completely separate from the data in the document.

Discussion

We have presented a broad view of what organisational memory is. Traditional data processing concentrated on the data aspects of memory. More recent organisational memory research includes aspects of data, but more importantly also aspects of process and interpretation. However, we have also considered placeholders as a crucial part of an organisation's short-term memory. These often rely heavily on the physicality of documents and other artefacts and thus reinforce the ecological view of information implicit in many recent studies and in the distributed cognition and situated action literature.

Organisational memory is not just an information repository, but is the rich, dynamic and distributed representation of state expressed in written

procedures, people's memory and the state of the many artefacts that comprise the organisation. Interaction and information systems design which forgets the rich multifaceted role of physical artefacts risks destroying the very processes it seeks to support.

References

Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge MA, Harvard University Press.

Barnard, P. (1985). *Interacting Cognitive Subsystems: A psycholinguistic approach to short-term memory*. In *Progress in the Psychology of Language*, vol. 2, Ed. A. Ellis. Hove, Lawrence Erlbaum Associates. pp. 197-258.

Dix, A.J., Ramduny, D. and Wilkinson, J. (1995). *Interruptions, Deadlines and Reminders: Investigations into the Flow of Cooperative Work*, RR9509, University of Huddersfield Research Report Series.

Dix, A.J., Ramduny, D. and Wilkinson, J. (1996). *Long-Term Interaction: Learning the 4 Rs*. in *CHI'96 Conference Companion*.

Dix, A.J., Ramduny, D. and Wilkinson, J. (1998). *Interaction in the Large* in *Interacting with Computers* (to appear).

Hutchins, E., and Klausen, T. (1991*). Distributed cognition in an airline cockpit*. In *Cognition in communication at work*, Eds. Y. Engestrom and D. Middleton. CUP: Cambridge.

Kidd, A. (1994). *The marks are on the knowledge worker*. In *CHI'94 – Human factors in Computing Systems*, Boston, ACM Press. pp. 186–191.

Laird , J. E., Newell, A. and Rosenbloom, P. (1987). *SOAR: an architecture for general intelligence*. *Artificial Intelligence,* **33**:1-64.

Lave, J. (1988). *Cognition in Practice: Mind, mathematics and culture in everyday life*. Cambridge University Press.

Orlikowski, W.J., (1992). *Learning from Notes: Organizational Issues in Groupware Implementation* in *Proceedings of CSCW'92*.

Plowman, L., Rogers, Y. and Ramage, M. (1995*). What are Workplace Studies For?* In *Proceedings of ECSCW'95*.

Ramduny, D., Dix, A. and Rodden, T. (1998). *Getting to Know: the design space for notification servers*. In *Proceedings of CSCW'98,* Seattle, Washington, (to appear).

Rouncefield, M., Hughes, J.A., Rodden, T. and Viller, S. (1994). *Working with "Constant Interruption" CSCW and the Small Office*. In *Proceedings of CSCW'94.*

Sellen, A. and Harper, R. (1997). *Paper as an analytic resource for the design of new technologies*. In *Proceedings of the 1997 conference on Human Factors in Computing Systems, CHI '97,* , pp. 319-326.

Suchman, L. A. (1987). *Plans and Situated Actions: The problem of human--machine communication*. Cambridge University Press.