

Interruptions, Deadlines and Reminders: Investigations into the Flow of Cooperative Work

Technical Report RR9509
School of Computing and Mathematics
Huddersfield University

Alan Dix^{1,2}, Devina Ramduny², Julie Wilkinson³

<http://www.soc.staffs.ac.uk/~cmtajd/topics/time/>

current affiliations:

1. aQtive limited, <http://www.aqtive.com/>
2. Staffordshire University, UK
3. Sheffield Hallam University, UK

Interruptions, Deadlines and Reminders: Investigations into the Flow of Cooperative Work

Alan Dix, Devina Ramduny, Julie Wilkinson

*HCI Research Centre, School of Computing & Mathematics,
University of Huddersfield, Queensgate,
Huddersfield HD1 3DH, UK.*

Email: {alan@zeus,scomdr@zeus,scomjw@pegasus}.hud.ac.uk

Abstract

This is a broad theoretical investigation into how events trigger action. In the course of studying cooperative workflow activity in a structured fashion it has emerged that a significant proportion of such activity involves long-term interactions. Human cognitive abilities are such that we usually cannot retain mental lists of all the tasks we are (meant to be) engaged in. As a result of this, the flow of work activities can be adversely affected. The potential problems that can ensue are:

- recalling out-of-context events
- reacting to long-term deadlines
- resuming activity (post-interruption)
- remembering the non-occurrence of anticipated events

In the application of a systematic analysis of workflow activities we have gained some insight into predicting which of the above problems are likely to occur and how we might offset any disruptions to the flow of cooperative work. Unlike most workflow studies - in the strict sense of the term - the work reported here is concerned far more with inter-organisational interactions and processes. This in turn produces another problem: that of maintaining a flow of interaction when there may be no central control over the whole activity.

1 Introduction

The issues of long-term interaction have been discussed by Dix (1994a, 1992). Because the work was mainly of a theoretical nature, we felt the need to extend our understanding of the problems of long-term interaction by carrying out a more focused empirical study. This would challenge and augment our existing work in two principal ways.

First, it enables us to verify and extend our theoretical understanding of the problems. Second, it allows us to apply the existing analysis in a more systematic fashion by aiming at a methodology which can be used by others to look for potential problems in long-term interaction. Moreover,

it can guide the design process to solve and ultimately avoid those problems. There is also the danger that theoretical work can become too general and so this was an opportunity to ground the work in a more specific context. In particular, we decided to look at cooperative tasks in an office-based environment.

One of the major problems identified by Dix (1994) was the non-occurrence of expected events. For example, you send someone a letter, but never receive a reply. For short-term interactions this is immediately obvious, you wait for the response and when nothing happens, you know something is wrong. However, for long-term interactions you cannot afford to twiddle your thumbs for several days waiting for a reply to a letter! You need a reminder that someone else needs to do something – a to-be-done-to list! The above problem is also closely related to issues such as interruptions. So the techniques we used were designed to expose these problems as well.

The key to our approach is therefore not so much to look at what happens, or even when it happens, but *whether* it happens at all. To do this we look at long-term cooperative processes. We divide these into activities performed by individuals or groups and record the interdependencies between these activities. However, the distinguishing aspect is that we look explicitly for the triggers which initiate activities.

The next section compares and contrasts the nature of our work with other related disciplines within the area of work analysis. Section three introduces our case study and describes the approach and the method we used to represent the flow of work. We then give a detailed account of our study. Finally we analyse the emerging issues.

2 Related approaches

The nature of this study bears some similarities to several disciplines in the generic field of the 'social analysis of work', namely workflow, ethnography and, more specifically, ethnomethodology. Nonetheless the following points highlight the radical differences between our approach and those above.

Dealing firstly with workflow, the term in its precise sense¹ implies technological solutions and/or opportunities to develop new technologies to improve the current nature of work. This is hardly surprising as most workflow systems originate from technological need, with office automation systems as a close and earlier cousin. They frequently go arm-in-arm with the growing trend for organisations to be 'process-focused'. In

¹One of the main centres within the workflow community - the Workflow Management Coalition 1994 - has defined all the terms relating to workflow in organisations - see their web site at <<http://www.aiai.ed.ac.uk/WfMC>>

accordance with this new thinking - be it business process re-engineering (Hammer & Champy, 1993), business process re-design, re-structuring or any other moves to alter radically the way organisations operate - workflow easily lends itself as a support mechanism in these contexts.

Workflow does have its limitations with regard to the type of investigation we were undertaking. Apart from the technological bias already mentioned, the concept of 'workflow management' hints at cultural change, as defined below:

the automation of procedures or workflows where documents, information or tasks are passed from one participant to another in a way that is governed by rules or procedures.

(Workflow Management Coalition 1994, Glossary of Terms)

Some degree of dictating the rules is suggested here, stipulating how and even when the different activities that constitute a process should be done. The very installation of these systems creates a culture of its own. To varying degrees they all have some kind of model of the user and the nature of the organisation. This 'laying down' of procedures is perhaps acceptable, possibly even desirable, within the confines of the organisation. Doubtless some may question the ethics of a system that imposes its own model of work activity on the users. It was not our aim to dictate any such cultural change.

The background theory to and origins of workflow also stem in part from the study of the flow of information around organisations and the attempts to make that more efficient. There appears to be little discussion of moving beyond the bounds of the organisation unless through some formalised collaboration of a long-term, often contractual nature - for example, a manufacturer 'locking in' a supplier (Carey, 1993). Even when internal to the organisation workflow is criticised for ignoring individual goals within a process (Warboys, 1994).

However, our process under scrutiny was as much extra- as intra-organisational. The overall control of the process may have resided ultimately in an office within one organisation, but significant control elements of it existed elsewhere, at other sites within different organisations. This was a major issue for us to address - since we were operating in a less predictable environment, how could we ensure that the links of communication and activity remained intact?

Our purpose was not to seek ways of improving workflow by automating the processes of work or even facilitating them by some means of computerisation. Nevertheless, our results do have some design implications as discussed in section 5. To avoid confusion or disagreement over our use of the term 'workflow' we have opted to differentiate our approach here by referring to it as a study of the 'flow of work'. The investigation at hand still differs further from workflow in its very precise

focus - that being on one process of the many that constituted the employee's daily work and also the targeting of events triggering activities.

Turning to ethnography, again some similarities may be drawn with the chosen approach. Ethnography is committed to inquiring into patterns of interaction and collaboration, based on the assumption that human activities are socially organised (Hammersley & Atkinson, 1983). We too were inquiring about a particular pattern - but with a difference.

Ethnography has an open-ended approach to what it may find through the social analysis of work. Indeed this approach is founded on the belief of ethnographers that one cannot know in advance of inquiry which elements of organisational life will prove to be of interest, value and importance for work (Randall, 1995). In contrast our work began with a sharper focus which will be described later. However, we omit aspects of a situation that an ethnographer would record. But ethnographers' open-endedness is often seen as a weakness when it is used for requirements capture (Anderson 1994). By being more restricted our approach is better suited to drive systems design.

Ethnomethodology has also been used within HCI (Suchman, 1988). Ethnomethodologists observe, collect and analyse data and decide what is relevant about work activity as it really is, not as an idealised conception of work - as can happen with process modelling and workflow. As a particular form of sociological analysis (Garfinkel, 1967) ethnomethodology is not interested in the fact that the social world is ordered but in how it becomes ordered in and through the processes of interaction. The main contrast between ethnomethodology and other modes of sociology is that it seeks to describe *from within* how people actually order their work activities through mutual attentiveness to what has to be done. Anderson (1994) calls it 'society's *lived-work*'.

We too were seeking to describe people's work activities but again the a priori focus on specific aspects of work distinguish our approach. Armed with the knowledge of what work had to be done we were interested in establishing 'breakdowns' which could affect the completion of that work process.

The importance of the environment (Bentley et al., 1992; Heath & Luff, 1992; Heath et al., 1993) for how work is executed has not escaped the notice of sociologists, least of all ethnographers/ethnomethodologists. These studies have stressed the other social actors within that environment - the close teamwork - rather than the surroundings in which people work. Recent studies of office work (Rouncefield et al., 1994) have brought the environment into the limelight. This trend is further followed in our work, but with a more specific formulation of the purpose of artefacts within that environment.

3 The case study

The case study is about the flow of work involved in the administration and organisation of HCI'95 that was held at the University of Huddersfield. The activities involved are typical of an academic conference. The call for participation was first issued by the conference committee whereby submissions for papers, industrial reports, panels, posters, short papers, doctoral consortium, tutorials, demonstrations and lab overviews were invited. A set of deadlines was attached to each submission category. Booking forms were made available to delegates who wanted to participate in the conference at a later stage.

There were many activities which had to be carried out prior to the actual conference and most of them required the coordination of information among several people. The parties concerned were either located within the same establishment or spread across various sites. Ann Jones assumed the role of conference organiser in addition to her normal work duties of administrative assistant to the research office. She acted as the first point of contact in any enquiry. We looked at an extensive range of activities which Ann had to coordinate but the flow of work during the life cycle of a paper was examined in the greatest detail.

3.1 Data collection method

Because of the similarity of our study to traditional task analysis, we could use many of the same sources for data collection: documentation, observation, interviews, etc.

Documentation of long-term process is likely to be relatively accurate, although it may omit the activities beyond organisational boundaries, and also most of the triggers. Direct observation, on the other hand, poses special problems as the processes of interest are long-term and often include geographically dispersed people. Moreover, due to the fact that the conference activities were a relatively small subset of Ann's overall work, this technique is not appropriate and therefore not employed.

Instead we conducted informal interviews to ascertain all the activities within the life cycle of a paper. We then refined the knowledge we gained and armed with specific questions, we approached Ann on numerous other occasions. The interviewing approach is often regarded as problematic since the accounts people give of their actions are frequently at odds with what they actually do. However, we are in a strong position as we approach such interviews - the structure imposed by the process flow and the specific interest in triggers allows us to trace omissions and inconsistencies.

Finally, the importance of environmental cues cannot be discarded. It gives us another rich source of information - the work itself. We know that environmental cues can be triggers for activities and so we can take

each item in the environment and look for the activity it triggers. However, environmental cues did not make a major contribution for this particular study for the very fact that Ann had other tasks to perform besides being the coordinator of the conference and hence she had to handle several types of paper work.

3.2 Processes and activities

The procedures for processing papers are described in section 4. The narrative account is mapped onto a model which depicts the flow of the various activities. We record the processes as a series of circles or bubbles, one for each activity. Each bubble names the activity and the person or persons who perform it. Lines between the bubbles record dependencies and arrows at the beginning of each bubble record the trigger for the activity (see Figure 1).

There are several methods for recording processes, but this is not the focus of our work, so we take a minimalist approach. We are not attempting to record all the complexities of real processes in a single diagram. Instead we use many separate diagrams, often concentrating on specific scenarios. The crucial thing is that for each activity we record we look for the corresponding trigger.

The level of analysis is also governed by this focus. In general, we place activity boundaries wherever there is the likelihood of a delay or gap. The most obvious case of such a break occurs when subsequent activities in a process are performed by people at different sites. However, there are often distinct activities performed sequentially by an individual.

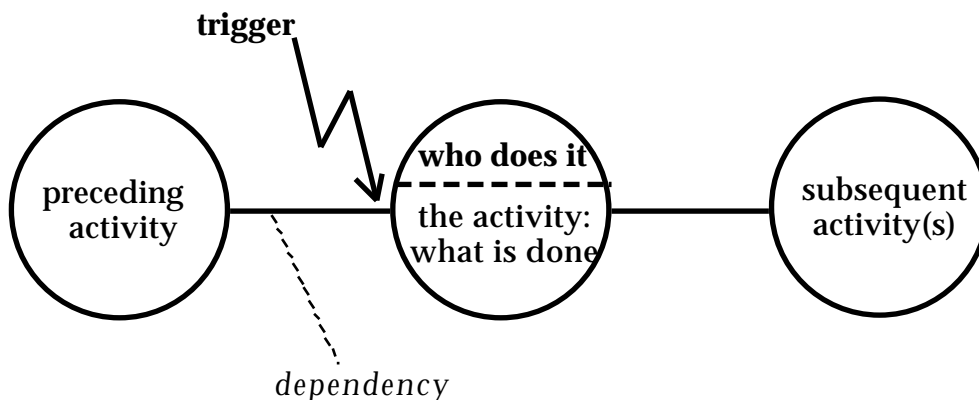


Figure 1: Recording processes

We also wish to retain a tight focus on long-term interaction and so we ignore very fine-grained tasks such as interruptions in the middle of an activity - for example, whilst typing a letter. We deliberately use the term activity rather than action to emphasise that the lowest level of our analysis is far from atomic. Activities may be shared between individuals, for example, having a meeting or dictating a letter would be regarded as a

single activity involving several people. Again, one could dissect such an interaction, but this would be the remit of conversational analysis.

We also ignore details of an activity when it is irrelevant or when we do not have sufficient knowledge about it. For example, if we issue an order to an external organisation and then wait for the goods to arrive, we may not be interested in their internal processes. Finally, we include some activities which would normally be omitted in a traditional process model. In particular, we often include the receipt of a message as a distinct activity. This is done deliberately to emphasise the gap which may occur between receipt and response.

3.3 Triggers

Triggers ensure the transition between activities. The dependencies between activities tell us that one activity is a pre-condition for another. This is the sort of dependency which is captured in a workflow or process model. However, there will typically be a gap between the completion of one activity and the start of the next. We therefore ask precisely which event triggers an activity.

Depending on the nature of the trigger we can determine the likelihood of an activity being missed and if the activity fails to occur, whether the failure will be noticed. For example, if the event is that some individual has to remember to perform a task, we might consider it a fragile part of the process, especially if it is performed in a complex and busy environment. Note that the triggers we are looking for are not the events which enable an activity to proceed – those are the preconditions. Instead the trigger is the event which makes the activity happen when it does.

4 Life Cycle of a Paper

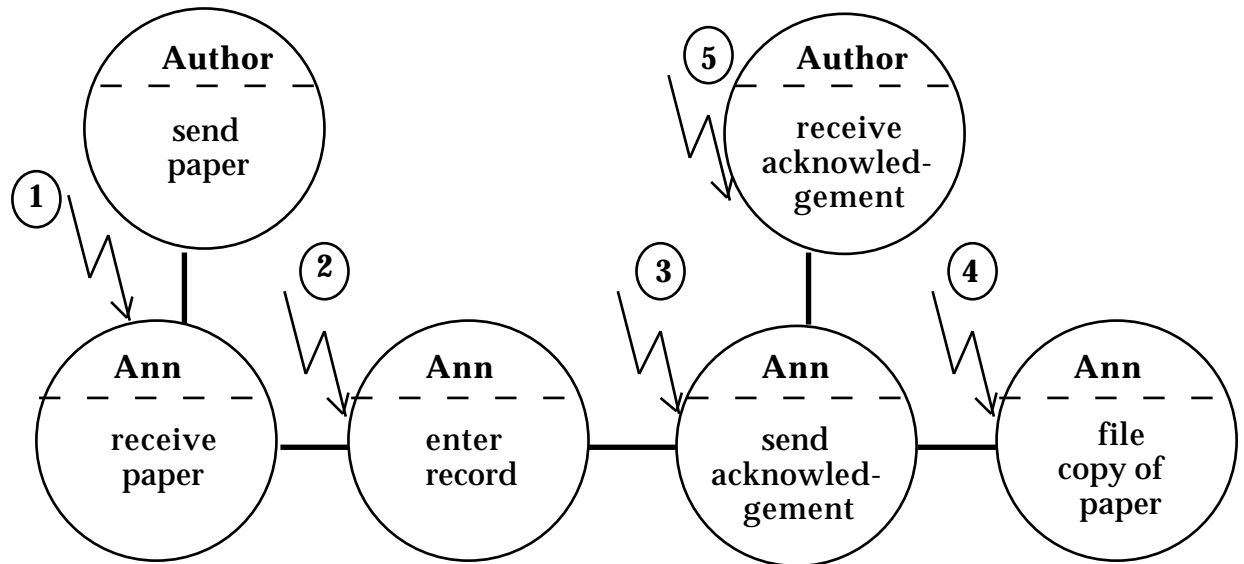
We have considered the underlying methods used in our study. We now describe the procedures followed during the life cycle of a paper.

Stage 1

Papers were dispatched by post to the conference organiser of HCI'95. Those papers reached Ann's mail pigeon-hole. The sender was either the author of the paper in question or the contact person in the case of multiple authors. Ann therefore checked her postal mail on a daily basis and left the papers lying on her desk until she had a sufficient number to process. She first entered a record on the electronic spreadsheet. Then she sent an acknowledgement to the contact person via post or email. Finally she filed a copy of the paper. The process flow is shown in Figure 2.

The first activity represents authors sending papers to Ann. They normally do so via a communication channel; in our case they used the postal mail. This mode of communication therefore acted as a trigger for

Ann to receive the papers. We can ask ourselves the question: what happens if the communication medium fails? The very fact of sending and receiving papers, assuming the recipient's address was correctly written, initiates the whole process. However, if the medium fails, the papers would never reach Ann and authors would not be aware that Ann has not received their papers. Likewise if the medium is unreliable, there is a probability that Ann would eventually receive the papers but after a considerable length of time.



Triggering events :

- ① & ⑤ direct communication
- ② paper on desk
- ③ & ④ directly follows preceding activity

Figure 2

Hence the failure or unreliability of the medium of interaction has serious implications for the system's operation. A possible solution to guard against the failure of a communication medium is to build a more reliable protocol on top of it. For instance, in our case the electronic mail could be used in parallel with the post. This might result in a situation yet where humans unlike software, may find the additional protocol too costly to maintain. Moreover, the reliability of the electronic mail could be questioned as well.

Triggers 3 and 4 are both such that, in an interruption-free environment, the end of one is the trigger for the next. However, in case of interruption, whenever we see such a trigger we always look for a fall-back trigger. For instance, Ann may be interrupted for some length of time while she is in the midst of sending an acknowledgement or filing a copy of the paper. The fall-back triggers for 3 and 4 are the same as each other and the same as trigger 2 - the paper lying on the desk. Because the activities have the

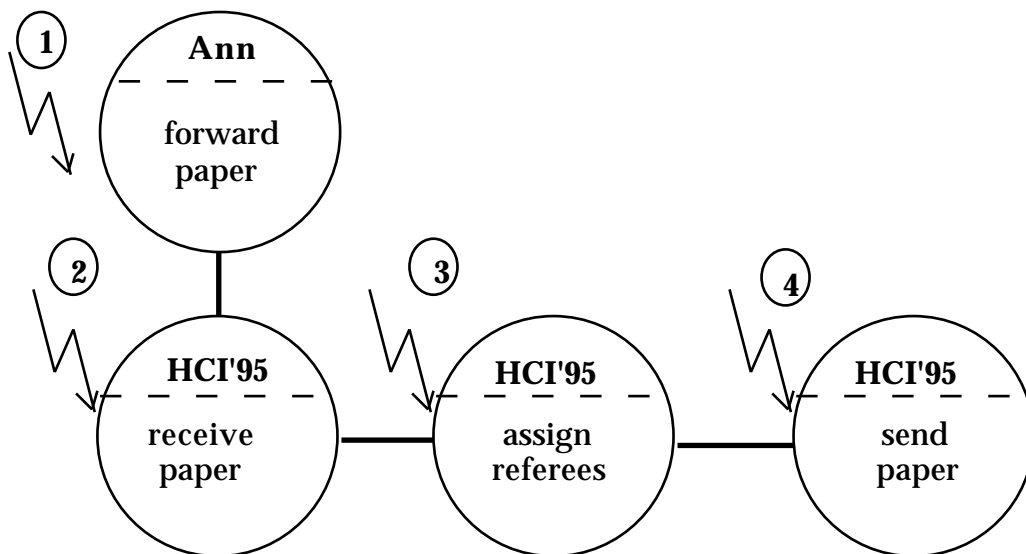
same triggers, Ann must remember the stage she reached before being interrupted.

It is the nature of human cognitive abilities that we cannot usually keep mental lists of all the tasks we are engaged in. Besides, if someone fails to complete or close tasks held in short-term memory or is prevented from so doing by interference, the subject is liable to lose track of what she is doing and can consequently make errors (Dix et al.,1993). So if Ann forgot to send an acknowledgement to the author and proceeded to file a copy of the paper, the author would have no immediate knowledge of the fact that his paper had been received until he receives some feedback at a later stage.

Interruption therefore does have major consequences on the flow of work within a system (Rouncefield et al., 1994). It can, in the least case, disrupt the flow of work or, in the worst case, cause a system failure.

Stage 2

Once the deadline for receiving papers had passed, Ann forwarded them to the HCI'95 committee. The task of forwarding the papers to the committee was not too difficult for Ann because it was governed by the deadline. There was only one deadline for all the papers, so that date was easy to remember.



Triggering events :


- ① after deadline
- ② direct communication
- ③ directly follows preceding activity
- ④ after assignment exercise

Figure 3

However, if each paper was allowed a different date for submission (most unlikely in this case) then it would give rise to a new scenario. Ann would have to keep track of deadline dates periodically. When faced with a periodic action one always asks, how does the person remember to perform the action at the relevant period?

Stage 3

When the referee assignment procedure was completed by the committee, Ann updated the records on the electronic spreadsheet. Afterwards she sent the papers to the relevant referees.



Triggering events :

- ① direct communication
- ② paper on desk
- ③ directly follows preceding activity

Figure 4

The process flow shown in Figure 4 poses the question: how does Ann ensure that she updates the records before sending the paper to the referees? Trigger 2 acts as a very crucial environmental cue which enables Ann to pick up the threads of her activities. The very fact of having the papers lying on the desk reminds her that she has to update a particular record.

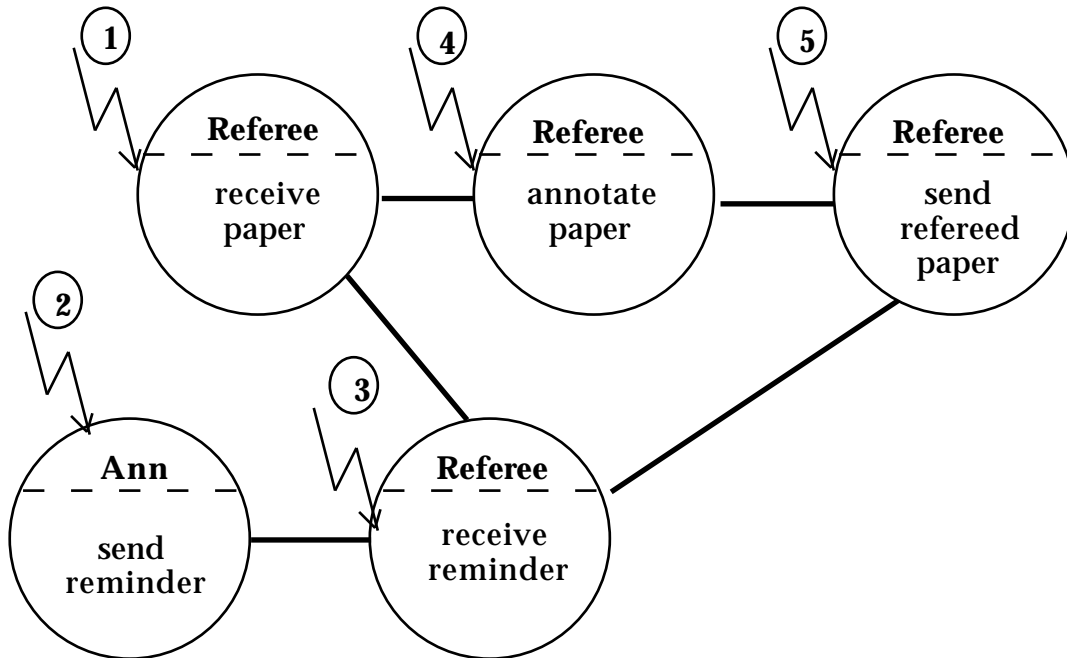
Environmental cues are therefore important triggers which serve as aide-mémoires. Reminders can manifest themselves in paper form: to-do-lists, diaries; or electronic form: emails; or even as knotted handkerchiefs!

Stage 4

After receiving the papers, referees annotated them with comments and returned the papers together with acceptance or rejection instructions to the HCI'95 committee. Occasionally, Ann sent reminders to the referees if she had not received the refereed papers by the date set for submission.

An interesting point which should be highlighted from Figure 5 is that the agents who are performing the activities are no longer within a single organisation. We have now crossed organisational boundaries and the

paper-processing activities are dependent on the referees spread across several locations. Thus Ann had the difficult and tedious task of coordinating the referees' activities at the different sites. Once she has sent the papers to be refereed, she loses control over the referees' actions. There is a temporal gap between Ann dispatching the papers and receiving them back duly refereed.



Triggering events :

- ① & ③ direct communication
- ② after deadline
- ④ directly follows preceding activity
- ⑤ after refereeing process

Figure 5

However, trigger 2 in Figure 5 enables Ann to regain control and subsequently allows her to coordinate the activities. The deadline prompts Ann to send reminders to non-responding referees. For instance, some referees annotated the papers but forgot to send them while others disregarded the refereeing exercise due to more pressing matter.

Therefore in a long-term cooperative situation where the control resides among different agents and when there is a gap between an event and its action, it is vital to prevent activities getting out of synchronisation otherwise a range of things can go wrong.

Stage 5

The HCI'95 committee reviewed the referees' comments and examined the remarks thoroughly. Ann was simultaneously notified so that she could update the spreadsheet records. Based on the outcome of the analysis, the committee decided whether a paper should be accepted or rejected. Subsequently, letters of acceptance or rejection were issued to the authors.

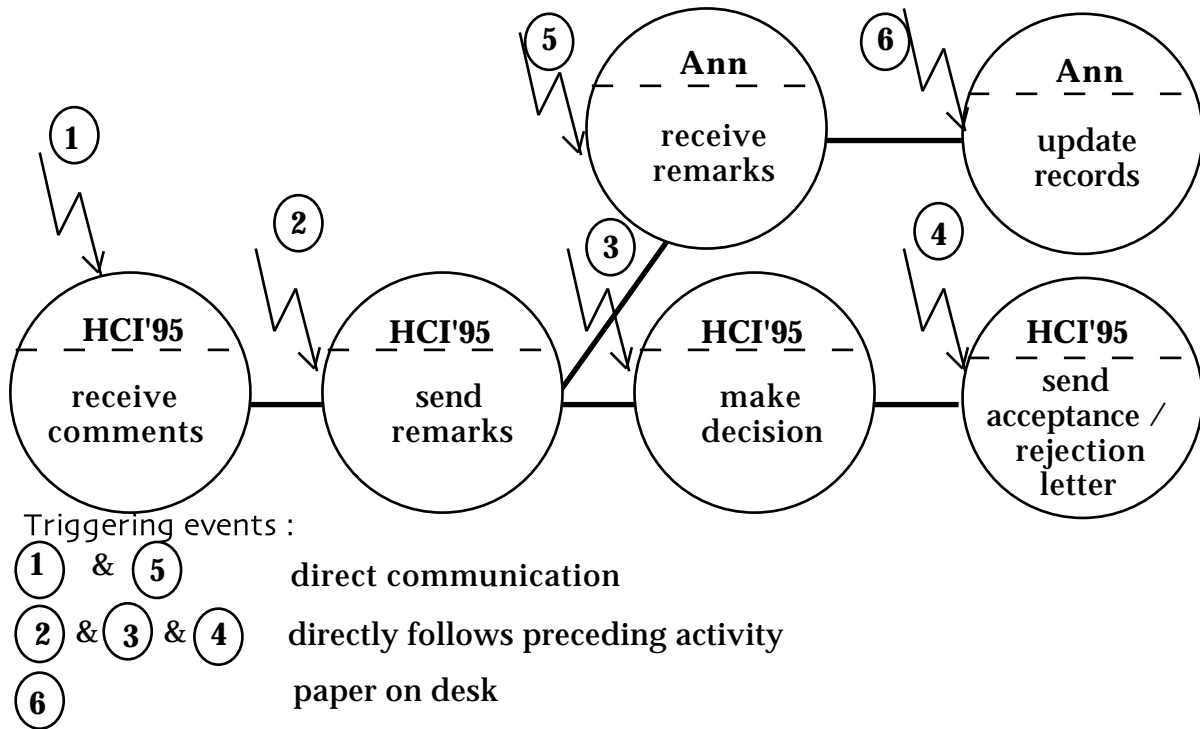


Figure 6

Stage 6 (a)

If the paper was rejected, the author(s) normally had two alternatives: either the paper was modified and sent back as a short paper or no further action was taken.

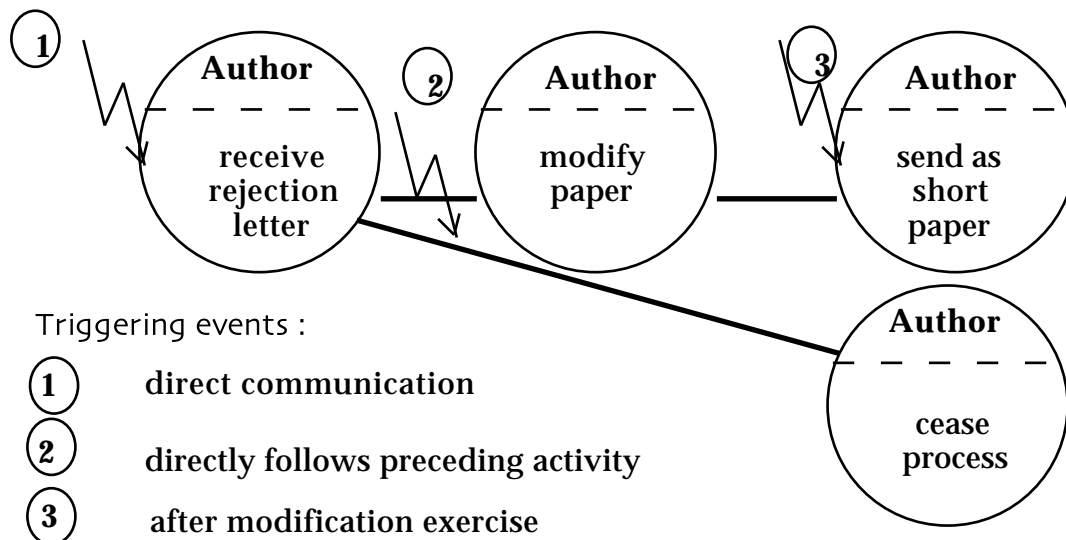
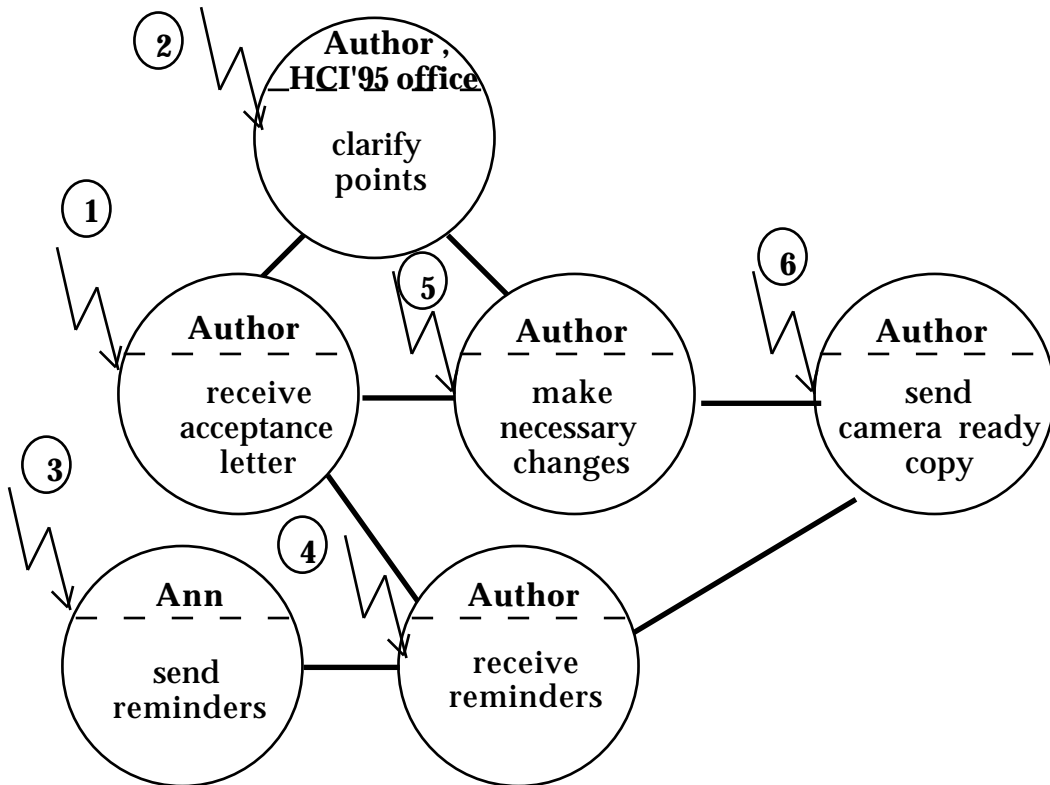


Figure 7

It should be noted that the branching in Figure 6 implies concurrent activities taking place while in the above illustration it represents an alternative set of actions. We have made no attempt to record this difference diagrammatically.

Stage 6 (b)

Instead if the paper was accepted, the author(s) made the necessary changes as stipulated by the referees' comments and sent a camera ready copy. However, there were some instances when Ann had to issue reminders to authors who had not submitted the camera ready copy by the deadline.



Triggering events :

- ① & ④ direct communication
- ② enquiries
- ③ after deadline
- ⑤ follows preceding activity
- ⑥ after modification exercise

Figure 8

Trigger 2 in above figure is an external event which prompts authors to clarify some details with the HCI'95 office before making the relevant changes to their papers. Because this exercise is shared between two agents and is of an ad hoc nature it is regarded as a single activity. Once all the

camera ready copies were received, they were dispatched for publication in the conference proceedings.

5 General Issues

The outcome of this study has provided us with a deep understanding of how to model activities during a work process. The major issues which have emerged are discussed below.

5.1 Types of triggers

In section 4 we encountered a number of recurrent triggers in the various scenarios. We are now in a position to identify some general classes of triggers.

- **Communication mode** – This might be a telephone call, face-to-face request or the receipt of a letter or a fax. The areas of concern are the reliability of these media and the consequences of their failure.
- **Environmental cues** – Things in our environment which remind us that things ought to be done. Sometimes this is explicit (e.g. a diary entry) sometimes implicit (a half written letter in the typewriter).
- **Temporal gaps** – The expectation of receiving a response by a certain date or the generic task of reminding people based on some temporal interval.
- **Completion of previous activity** – This is when one activity begins immediately after the previous activity reaches completion. However, we may treat this with suspicion. Does the second activity always proceed immediately? If there is any chance of a gap or interruption we must look for secondary triggers.
- **External events** – This might be from a wristwatch or automatic calendar set to give a reminder at a specific time. Alternatively, it might be due to a specific event occurring; the completion of an automatic activity, an event in the world, even the (electronic) receipt of a message.

Some additional types of triggers which can be identified implicitly are :

- **memory (sporadic actions)** – Frequently activities may occur simply when the responsible individual remembers that they must be done. It is often the case that when a request is made verbally, the recipient has to remember that the request is outstanding until either it can be performed or some record is made of the commitment. Note also that in the latter case, the recording of the commitment is itself an important activity.
- **Periodic actions** – Things which happen at regular intervals - for example, reading your mail every morning. When faced with a periodic action, how do we remember to perform the action at the

relevant moment? If it is something like consulting a diary every morning, we can believe it is part of a routine. However, an hourly activity should prompt further questions – how does the person know when it is the hour? Perhaps the clock strikes an external signal.

When we described the triggers in section 4 we had subsidiary questions in many cases. In the case of temporal gaps – how do we recall the non-occurrence of anticipated events? In the case of having the same trigger for activities in immediate succession – what happens if there is an interruption? The environmental cues seem fairly fundamental, but even there we must ask why is it the subject notices the cue? In the case of a diary entry, perhaps the subject consults the diary each morning – a periodic activity. One could continue asking such follow up questions indefinitely, but at some point we must stop and either believe that a trigger does always occur as specified, or if not, to assess its reliability and perhaps delays associated with noticing it.

5.2 Robustness

The reliability of individual parts of a work process can be assessed by asking penetrating questions about the triggers for activities. However, nothing is ever 100% correct and it is inevitable that triggers will fail for some reason, activities may be missed and perhaps the whole process may fail to continue because something goes wrong.

The combination of a process model, together with a well-founded evaluation of the reliability of each activity, can allow us to assess the robustness of the whole process. If someone fails to complete some activity, and hence quite probably the next activity is never triggered, what happens? Does the whole process grind to a halt, or will the failure eventually be noticed? This is not an ad hoc procedure, but one can systematically go to each trigger and ask – what happens to the process as a whole if the trigger fails? Furthermore, by looking at the whole process we can improve our assessment of the reliability of any trigger.

5.3 The 4 Rs

The flow of activities in Section 4 shows an emerging pattern. We have decided to call it the 4Rs : request, receipt, response, release.




Figure 9: The 4Rs

Figure 9 shows a simplistic version of Figure 2 which exemplifies the 4Rs. A *request* is first made by an agent. This is closely followed by the *receipt* of that request via a communication channel. Then there is a *response* which manifests itself as some sort of action. The final stage is the *release* of the thread of activities when the request has been dealt with. At this point, if the functional goal has been achieved, then the process can be considered to have reached completion. Nevertheless, the 4Rs process is rarely isolated since a response may itself be the request of another 4R sub-process.

Not only is the pattern of activities common between different processes, but we also see a relevant pattern of triggers. Trigger 1 is always simply some sort of communication mode and can be assessed for reliability and timeliness. Trigger 2 is often an environmental one (in this case the paper on the desk). The release activity removes this cue - for example, filing a paper, throwing away a post-it-note. However, trigger 3 is typical of the "immediately follows" kind and its fall-back trigger is identical to trigger 2 leading to the danger of repeating or omitting the response activity.

5.4 Design Implications

Our study highlighted the importance of reminders as an aide-mémoire for us to respond to a certain request. If we have responded to that request, we might still encounter further problems. Should there be a long delay before the effects of our action occur or become apparent we risk losing the context of that event. One way email systems tackle this problem is to include the senders' message in the reply. Similarly, if something goes wrong and there is never a response, especially in the case of long-term interactions, then we need another reminder that someone else needs to do something.

We have also considered the effect of interruptions when the user is engaged in some action. Some interruptions can be on a very short-term basis and they may not have a serious impact upon the current activity. Others, however, may be more prolonged. The essential difficulty rests in resuming the activity. This is twofold: remembering where you were before being interrupted, or even remembering to resume the activity at all. Many researchers have noted the importance of paper on the desktop and other external reminders (Rouncefield et al. 1994). Indeed, in many cooperative processes there may be little direct communication. Instead the parties coordinate by implicit communication *through the artefact* (Dix 1994b).

Our adopted approach to investigate the flow of work therefore enables us to identify some potential breakdown points. By being aware of the importance of triggers we can avoid losing them when systems are automated. Environmental cues are of prime significance and they should be given due consideration when proposing automated solutions or when implementing any changes to the flow of work.

References

Anderson, R.J., (1994) Representations and Requirements : The Value of Ethnography in System Design. In *Human-Computer Interaction*, Volume 9, Lawrence Erlbaum, pp 151-182.

Bentley, R., Hughes, J.A., Randall, D., Rodden, T., Sawyer, P., Shapiro, D. & Sommerville, I. (1992) Ethnographically-informed systems design for air traffic control. In *Proceedings of CSCW'92* (Nov., Toronto, Ontario), ACM Press, 1992, pp. 123-129.

Bentley, R., Hughes, J.A., Randall, D. & Shapiro, S.Z. (1992) *Technological support for decision making in a safety critical environment*. Technical report (CSCW/5/92), Computing Department, Lancaster University. Available via anonymous ftp from comp.lancs.ac.uk.

Carey, D. (1993) *Reengineering success*, I.T. Magazine, Nov. 1993, v25, n11, pp 12-18.

Dix, A, Finlay, J., Abowd, G., & Beale, R. (1993) *Human-Computer Interaction*, Prentice Hall.

Dix, A. (1994a) Que sera sera — The problem of the future perfect in open and cooperative systems. In *Proceedings of HCI'94: People and Computers IX* (Aug., Glasgow), Cambridge University Press, 1994, pp 397-408.

Dix, A.J., (1994b) Computer-supported cooperative work — a framework. Design. In *Issues in CSCW*, Eds. D. Rosenburg and C. Hutchison, Springer Verlag, pp 9-26.

Dix, A.J (1992). Pace and interaction. *Proceedings of HCI'92: People and Computers VII*, (Sept. York) Cambridge University Press, 1992, pp 193-208.

Garfinkel, H. (1967) *Studies in ethnomethodology*, Prentice Hall.

Hammer, M. & Champy, J. (1993) *Reengineering the Corporation - A manifesto for business revolution*, Nicholas Brealey.

Hammersley, M & Atkinson, P (1983) *Ethnography: Principles in Practice*, Tavistock.

Heath, C. & Luff, P. (1992) Crisis management and multimedia technology in London Underground line control rooms. In *Journal of CSCW*, Vol. 1 (1-2), 1994, pp 69-94.

Heath, C., Jirokta, M. Luff, P. & Hindmarsh, J. (1993) Unpacking Collaboration: The Interactional Organisation of Trading in a City Dealing Room. In *Proceedings of ECSCW'93*, (Sept. Milan, Italy), Kluwer Academic Publishers, 1993, pp 155-171.

Randall, D. (1995) *Ethnography for Systems Development: Bounding the Intersection*. Tutorial Notes HCI'95, University of Huddersfield.

Rouncefield, M., Hughes, J.A., Rodden, T. & Viller, S. (1994). Working with “Constant Interruption” CSCW and the Small Office. In *Proceedings of CSCW'94*, (Oct. Chapel Hill, North Carolina), ACM Press., 1994, pp 275-287.

Suchman, L. (1988) *Plans and situated actions*, Cambridge University Press.

Warboys, Brian (1994) Reflections on the Relationship Between BPR and Software Process Modelling. In: *Proceedings of ER '94*, (Loucopoulos ed.), Springer-Verlag, 1994, pp 1-9.