

## Interaction in the large

Alan Dix<sup>a,\*</sup>, Devina Ramduny<sup>a</sup>, Julie Wilkinson<sup>b</sup>

<sup>a</sup>*School of Computing, Staffordshire University, Stafford ST18 0AD, UK*

<sup>b</sup>*School of Computing and Management Sciences, Sheffield Hallam University, Sheffield, UK*

Received 21 August 1997; received in revised form 14 February 1998; accepted 18 May 1998

---

### Abstract

Most work in HCI focuses on interaction in the small: where tasks take a few minutes or hours and individual actions receive feedback within seconds. In contrast, many collaborative activities occur over weeks or months and the turnaround of individual messages may take hours, days or even weeks. This slow pace of interaction brings its own problems, especially when expected responses do not occur. This paper analyses these problems, focusing on the triggers which initiate activities and the way processes recover when triggers are missed or misinterpreted. Furthermore, we are able to consider processes which cross organisational boundaries. We draw on theoretical analysis, an exploratory case study of conference organisation and recent application of the techniques to a student placement office. During the studies, a pattern of recurrent activities was discovered, the 4Rs (request, receipt, response and release), which we believe to be generic to this class of collaborative process. © 1998 Elsevier Science B.V. All rights reserved

*Keywords:* Pace; Interruptions; Reminders; Events; Long-term interaction; CSCW; Cooperative work; Workflow; To-be-done-to lists; Paper documents

---

### 1. Introduction

Most modern interactive systems attempt to give instantaneous feedback to users. However, many situations both in human–computer interaction and in human–human communication involve more protracted processes. Instead of interaction over periods of seconds, we may have interactions spread over days, months or years. Interaction in the large brings its own problems distinct from those of interaction in the small. This paper brings together and builds on previously developed theoretical techniques [7,10–12].

---

\* Corresponding author. Fax: +44-1785-353431; e-mail: A.J.Dix@soc.staffs.ac.uk

These techniques are used to analyse real examples of long-term cooperative interaction, and the results of these case studies have themselves enriched our theoretical understanding.

### *1.1. Methodology*

An important aspect of this paper is the analysis-driven methodology. Many empirical studies and the ethnographic techniques popular in computer-supported cooperative work (CSCW) start from observation and subsequently analyse the results. In contrast, our start point is derived from theoretical analyses. However, this theory is validated by our initial case study which then leads to the formulation of further theoretical concepts and a second round of empirical action research.

This rich interplay of theory and observation strengthens both. We have a tighter focus in our empirical work than would be possible from an observation-centred approach and we are able to validate and investigate the practical implications of our theoretical results.

### *1.2. Outline*

Given our methodological stance, the paper starts with an extensive review of the theoretical background to this work: status–event analysis and the study of pace. Status–event analysis reminds us of the importance of status phenomena (persistent states of the world as opposed to ephemeral events) and their role in mediating interaction and themselves generating events. We consider the various factors affecting pace of interaction, especially those which serve to slow down the pace and give rise to interaction-in-the-large.

We then move on to look at specific issues of interaction-in-the-large where the cycle of interaction is longer than human short-term memory and where the slow pace of interactivity makes it difficult to recall context, to remember what we have to do and to remember what others are supposed to have done.

The issues raised allow us to propose an analytic method based on the identification of different kinds of triggers which initiate activity. We use diagrams mapping dependencies between activity, rather like forms of workflow or business process analysis; however, the key feature is the annotation with triggers. By repeatedly asking ‘what makes this happen when it does’ we are able to discover what triggers each activity and what happens if the triggers fail.

We then consider our first case study, an analysis of the processes of the HCI’95 conference. The initial intention was to use it as a test bed to check that our lists of potential triggers were complete and examine further theoretical issues surrounding long-term interaction. However, during this study we discovered a repeated pattern of activities which we call the 4Rs: request, receive, respond, release. The 4Rs pattern is described in detail in the following section and we explain why we believe it to be a fundamental pattern of long-term collaborative working.

The general analytic approach and the 4Rs pattern are then re-applied in another case study which involved a Lotus Notes implementation for the Student Placement Unit of Huddersfield University. The Unit has extensive procedures involving hundreds of

students and companies and is ‘mission critical’ for the University which has many sandwich courses that require a year’s industrial placement. The analytic approach was used to analyse the existing situation and to suggest appropriate use of Lotus Notes in support of the unit.

Finally, we compare our approach with other similar approaches in ethnography, business process re-engineering and workflow, and finally draw out some design implications of our approach.

## 2. Background

The roots of the current work lie in two principal theoretical foundations: the study of pace of interaction [7,8] and status–event analysis [1,6,13]. The primary basis is the issues surrounding pace—that is the rate at which users interact with computer systems, the physical world and with one another. Thinking about pace makes one concentrate on the timescale over which interaction occurs, both the similarities between interactions of widely different pace and also the differences. The contribution of status–event analysis is less central than pace to the work described in this paper, but does influence the way we approach the issues, in particular our understanding of triggers—events which initiate actions.

### 2.1. Status–event analysis

Status–event analysis is a collection of formal and semi-formal techniques all focused on the differences between events (things that happen) and status (things which always have a value). Applications of status–event analysis have included auditory interfaces [5], formal analysis of shared scrollbars [1] and software architectures for distributed agent-based interfaces [39].

Although we will not expand on the details in this paper, the formal analysis has been important in shaping our classifications of the kind of events that can occur and in our analysis of triggers for action. Of particular importance is the distinction between actual events (some objective thing which occurs) and from perceived events (when an agent (human or machine) notices that the event has occurred). Sometimes this is virtually instantaneous, but more often there is a lag between the two. Many formal and informal analyses of events assume a simultaneity between cause and effect. However, as we shall see, accepting that there is often a gap allows us to investigate what actually causes the secondary event to occur when it does.

Furthermore, most notations in computing focus principally on events with little if any description of status phenomena. However, we shall see that environmental cues (things that are around us) such as piles of papers or noticeboards, play an essential part in maintaining effective long-term interaction. These can only be understood when status phenomena are raised on an equal footing with events.

The analysis of status and events has also allowed us to see common features between human–human, human–computer and internal computer interactions. For example, it is common to see status mediation whereby one agent communicates an event to another by

manipulating a status that will eventually be observed by the second agent. Also polling, the periodic observation of a status phenomena to detect change, is not just a low-level computational device, but something people do as well. This rich interplay of status and event phenomena is reflected in the ecological perspective which colours the analytic stance of our current study.

## 2.2. Pace and bandwidth


The ‘pace’ of interaction with other people or with a computer system is the rate at which you send messages/commands and then receive a response. It varies from tens of milliseconds in a video game to hours or days when interacting by post.

The normal figure quoted for communication channels is bandwidth. Bandwidth measures *how much* information can be passed down the channel per second. In contrast, pace measures *how often* you can communicate using the channel. Elsewhere it is argued that, of the two, pace is the most significant for effective interaction [7]. Fig. 1 shows the typical use of a pair of communication channels: short bursts of communication in one direction await responses in the other.

## 2.3. Pace limits and mismatch

The pace of interaction is influenced by and influences three principal factors:

1. the intrinsic pace of the channels through which you communicate;
2. the pace of the collaborative tasks; and
3. the users’ own natural pace for different forms of mental activity.



Problems may occur when there is any significant mismatch between any of these and the resulting pace of interaction.




Fig. 1. Pace versus bandwidth.

The most obvious problem is when the channel is too slow. This may be because of latency of the medium, for example the time it takes for a trans-Atlantic video signal to be transmitted via a satellite or the fact that letter posts are only collected and delivered once or twice a day. Not only may channels be too slow, the intrinsic pace of the channel may also be too fast. For example, if help-desk operators have to spend a long time looking in a manual they may make small sounds as a form of ‘keep-alive’ for the channel (or play musak!). In such circumstances one might change the channel to one with a lower pace—perhaps e-mail would be more appropriate than the telephone.

Where the pace of a channel is too fast or too slow users may adopt coping strategies, patterns of behaviour to mask or overcome the problems. For example, when the channel pace is too slow, users may use multiplexing, where several conversations or activities are carried out in parallel, or eagerness, where messages make assumptions about the responses of the recipient: “Are you coming on the 5 o’clock train? If so I’ll meet you under the station clock.” [7].

The judgement of whether a channel is too fast or too slow is not absolute, but dependent on the context of the interaction, and in particular the task. Sometimes it is possible to change the task, making it slower or faster to fit a channel. For example, postal chess where the normal time rules do not apply. Typically there is less flexibility in the pace of a task, as there may be physical or computational constraints limiting its maximum or minimum pace. This is especially true of collaborative tasks: for example, it has been possible in an emergency to ‘talk down’ a plane where the pilot is incapacitated and a passenger takes the controls. However, it would be impossible to do this task by e-mail as the pace of the plane landing task cannot be slowed down! Furthermore, different subtasks each have their own natural pace; this variation can be used to accommodate poor channels. If subtasks demand a pace which is faster than that which is possible through the available channels, then we can change the roles, allocating the entire subtask to a single person or several people. Note that this pattern is typical of long-term collaborative work: substantial subtasks are completed individually within a coordinating framework of less frequent communications and collaborative actions.

Humans also have physiological and neurological limits on pace. Hand–eye coordination tasks have a pace of around 100 ms. This imposes an upper limit on the rate at which people can react and also means that computer feedback for hand–eye coordination tasks must be within this timescale. Turn-taking in conversation is itself mediated by short gaps in speech of the order of a few hundred milliseconds, which can be severely impaired by delays in inter-continental video-conferences. Short-term memory also fades over a similar timescale unless constantly refreshed by rehearsal.

Long-term interaction occurs at a much slower pace than any of these human time-scales. Over longer periods we only have the cycle of regular diurnal and weekly activities and our reactions to external irregular events. Although we may not have to worry about reaction times, the lack of short-term memory and simple sequenced interactions can make interaction at a very slow pace even more complex than at a fast pace.

### 3. Long-term interaction

The study of pace helps us to understand interactivity in a wider context. A system or collaborative process is not interactive because it is fast or it has instant feedback. Instead, interactivity is about the appropriate pace of interaction in relation to the task at hand. This is certainly the case in many collaborative situations where the pace of communication may be over days or weeks.

The reason for the prolonged nature of these interactions varies: it may be due to the communication medium (e.g. normal postal delays), or due to the nature of the task (e.g. a doctor waiting for X-ray results). One of the key points is that models of interaction which concentrate on a tight cycle between action and feedback break down [8]. This is typified by Norman's execution-evaluation cycle as shown in Fig. 2 [24,25]: a user has a goal, formulates actions which will further that goal, executes the actions, then evaluates the results of those actions against the expected outcome and the goal. This model effectively assumes that the results of the user's actions are immediately available. If the delay between executing actions and observing the results is greater than short-term memory times, then the evaluation becomes far more difficult. This problem has been called the 'broken loop of interaction'.

Another model of interaction used in more industrial settings is to treat the worker in a stimulus-response manner. Commands and alarms act as stimuli and the effective worker responds to these in the appropriate manner. However, in a pure form, this model does not allow for any long-term plans or goals on the part of the worker; the worker is treated in a mechanistic manner, a cog in the machine.

To incorporate both these perspectives we need to stretch out the interaction and consider the interplay between the user and the environment over a protracted timescale. We use the term environment to include interactions with other users, computer systems or the physical environment. Such interaction is typically of a turn-taking fashion: the user acts on the environment, the environment 'responds', the user sees the effects, and then acts again...

This process is illustrated in Fig. 3. Notice how the Norman loop concentrates on the user-environment-user part of this interaction: the user formulates goals and executes them, this affects the environment and the user evaluates the effect. In contrast, the stimulus-response model emphasises the environment-user-environment part.

Looking at this diagram we can see various ways in which long-term interaction may affect the interaction.




Fig. 2. Norman's execution-evaluation cycle and stimulus-response model.

- ① *Action–effect gap*—The user performs an action, but there is a long delay before the effects of that action occur, or become apparent to the user. For example, you send an e-mail and some days later get a reply. The problem here is loss of context; how do you recall the context when you eventually receive the feedback. When the reply comes you have to remember the reason why the original message was sent and what your expectations of the reply were. The way in which e-mail systems include the sender’s message in the reply is an attempt to address this problem. In paper communications the use of ‘my ref./your ref.’ fulfils a similar purpose.
- ② *Stimulus–response gap*—Something happens to which the user must respond, but for some reason cannot do so immediately. For example, someone asks you to do something when you meet in the corridor. The problem here is that you may forget. Hence the need for to-do lists or other forms of reminder. In the psychological literature this has been called prospective memory [28].
- ③ *Missing stimulus*—The user performs an action, but something goes wrong and there is never a response. For example, you send someone a letter, but never get a reply. For short-term interactions this is immediately obvious, you are waiting for the response and when nothing happens, you know something is wrong. However, for long-term interactions you cannot afford to do nothing for several days waiting for a reply to a letter! In this case you need a reminder that someone else needs to do something—a to-be-done-to list!

Possible design solutions to the last problem were the focus of the work by [8], as of the three this is probably the least well understood or catered for in computer systems. However, all three problems are possible causes of failure during long-term interaction.

#### 4. Analytic method: triggers and activities

The analysis in [8] had highlighted the problems due to missing stimuli and the problems of long-term interaction. Although it was clear that some of the problems could occur, some are empirical and we thus set out to validate this analysis in a real situation. Although our initial reason for starting this work was to study the problems due to missing stimuli and long-term interaction, these issues are intimately linked to issues such as interruptions as in both cases the flow of activities within a task is broken. The techniques we use are therefore designed to expose these problems as well.

Part of the data we collect is on *what* is done. In traditional workflow fashion, we catalogue the various activities performed and the dependencies between the activities. However, this is only intended as the superstructure of the analysis, not the focus. Instead,




Fig. 3. Problems for long-term interaction.

our focus is on *when* activities are performed and *whether* they happen at all. The central and distinguishing feature of our work is therefore the way we look explicitly for the triggers which initiate activities.

#### 4.1. Triggers

Triggers ensure the transition between activities. The trigger is the event which makes the activity happen when it does. The dependencies between activities tell us that one activity is a precondition for another. This is the sort of dependency which is captured in a workflow or process model [36,38]. However, there will typically be a gap between the completion of one activity and the start of the next. We therefore ask what event triggers each activity. Depending on the nature of the trigger we can determine whether it is possible or likely that an activity will be missed and, if it fails to occur, whether the failure will be noticed. For example, supposing that some individual has to remember to perform a task, we might consider that event 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 make it possible for an activity to proceed—that is the preconditions. Instead the trigger is the event which made the activity happen when it did. Consider the following scenario:

1. a letter arrives,
2. you read the letter and decide it must be answered,
3. you write the answer to the letter,
4. you post the answering letter. Each activity requires the previous activity to be completed before it can begin (precondition). However, it is likely that there will be a gap between, say, the letter arriving and the letter being read. If the letter arrives at 9am clearly the reading of the letter cannot happen before then. Let's say you eventually read the letter at 10:30am. Why did you read it then, why not at 11am or 9:30am? Some event must occur at 10:30am which causes you to read it; perhaps it is simply time for your morning cup of tea and you always read your mail then. Whatever the reason, it is that event (whether or not it is an obvious part of the process) which is the trigger.

#### 4.2. Processes and 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 Fig. 4).

There are plenty of methods for recording processes, but this is not the focus of our work, so we take a minimalist approach. We do not attempt 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 look for the corresponding trigger.

The level of analysis is also governed by this focus. In general, we put activity boundaries wherever there is the likelihood of a delay or gap. The most obvious such 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, as in the letter-reading example above. In principle such analysis could go down to the full detail found in hierarchical task analysis [34]. This would be reasonable if, for example, interruptions were possible in the middle of typing a letter. Although this would be an interesting exercise, we wish to retain a tight focus on long-term interaction and so we ignore very fine-grained tasks. 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 may also ignore details of an activity because it is uninteresting or 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 the internal processes of that firm. 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 deliberately to emphasise the gap which may occur between receipt and response (see example 1 below).

### 4.3. Types of triggers

The combination of status–event analysis and the study of pace led to an initial list of potential trigger types used in our studies. Although these were clearly informed by our common sense knowledge of the world they were obtained by analysis, not by formal observation. Instead, one of the purposes of our study was to verify that these formed a complete set and whether they failed in ways we expected. In fact, we did not add any new classes of trigger as the result of our empirical work. However, it is clear that the classification is not as ‘clean’ as we would like and we are thus currently looking again at more formal semantic models to improve the structure of the list.

The trigger classification we used in our studies was as follows.

- (a) *Completion of previous activity*—This is when one activity begins immediately after the previous activity reaches completion. But 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.
- (b) *Memory (sporadic actions)*—Frequently activities may occur simply when the




Fig. 4. Recording processes.

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.

(c) *Periodic actions*—Things which happen at regular intervals, for example if you read your mail every morning. When faced with a periodic action, how do we remember to perform the action at the relevant period? 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 (see below).

(d) *Temporal gaps*—Often it is a single significant moment or delay rather than a periodic time that is important. For example, we may need to perform a task by a deadline or expect a response by a certain date. Again we must ask what makes a person notice that the actual event has occurred.

(e) *External events*— Often periodic and temporal events are signalled by a wristwatch or automatic calendar set to give a reminder at a specific time. Also non-time based events may occur to prompt action: the completion of an automatic activity, an event in the world, even the (electronic) receipt of a message.

(f) *Receipt of a message*—This is a special kind of external event. It may be a telephone call, a face-to-face request, or the receipt of a letter or fax. Such events can only be considered to be the trigger of an action if that action occurs immediately after the request is received. If, as is more common, the request is dealt with later, then the receipt of the request and the response to the request are treated as separate activities. We also have to record the reliability of the communication media and possible communication delays.

(g) *Environmental cues*—These are things in our environment which remind us that activities ought to be done. Sometimes this is explicit (e.g. a diary entry), sometimes implicit (a half-written letter in the typewriter). Environmental cues may manifest themselves in paper forms (e.g. to-do-lists and diaries) or electronic forms (e.g. electronic mail waiting in the in-box).

Triggers of type (a) and (b) are insecure as they are liable to interruptions and poor memory, respectively. In each case we look for a secondary or back-up trigger, or, where this is absent, we look at the process as a whole and assess the consequences should the activity fail to trigger at all. Other triggers also lead to follow-on questions. For example, if a temporal event (d) is triggered because it is in a diary, what makes you look in the diary? In the case of periodic activities, how do we know when the period occurs? Environmental cues are fundamental, but even here we must ask why is it that the subject notices the cue? We 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, assess its reliability and perhaps delays associated with noticing it.

## 5. Case study—the HCI'95 conference

In order to validate our ideas on the nature of long-term work we studied the flow of work involved in the administration and organisation of the HCI'95 conference. Many activities had to be carried out prior to the actual conference and most of them required the coordination of information among several people at various sites. The central figure in many of these activities was Ann, the conference organiser. However, this was only one part of her work for the duration of the conference in addition to her normal work duties. 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. For a longer report see [10].

### 5.1. Data collection

In this study the processes we encountered were in lock-step and constituted only a small part of Ann's overall work. Direct observation was therefore impractical and instead we resorted to in-depth interviews. Interviewing 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. Our analytic focus—the structure imposed by the process flow and the specific interest in triggers—allows us to trace omissions and inconsistencies and enables us to obtain reliable results from interviews. This is important as, although we would normally expect some additional direct observation, practical design must rely principally on more directed and less intrusive techniques.

Although in our case data collection methods were severely constrained, in other cases we can use the full range of sources generally used for task analysis or requirement elicitation including documentation and direct observation. However both have special problems when trying to map out long-term, ecologically-rich, cross-organisational processes.

Documentation of long-term processes is likely to be relatively accurate, although it may omit the activities beyond organisational boundaries, and also most of the triggers. However, we can use it as an initial framework which can be filled out by observation or during subsequent interviews.

Direct observation poses special problems as the processes of interest are long-term and geographically dispersed. The necessary protracted field studies would not be acceptable as a part of normal commercial design practice. This is why we resorted to interviews in our own case study.

However, the lock-step nature of a conference is not typical of office processes and in many situations there are several instances of the same process at different stages of completion. For example, in an insurance office many claims are processed, each at a different stage. In these cases a day-in-the-life observation may be sufficient. So long as we can see each activity during the study period, we can piece them together afterwards. Even if we never see a process run from end to end we can reconstruct it from its parts. This is similar to observing a natural forest. The complete life-cycle of a tree might be hundreds of years long, but by looking at trees at different stages of growth, you can build up a full picture over a much shorter period.

Finally, the importance of environmental cues gives us another rich source of information—the work environment itself. We look at an office. There are papers and files on the desk, post-it notes, an in-tray, a wall calendar. Why is that file on the desk? What will happen to it? What would happen if it were not there? We know that environmental cues can be triggers for activities and so we take each item in the environment and look for the activity it triggers, or the coordinating role it fulfils. At the very least a piece of paper left on the desk is saying ‘file me please’.

### 5.2. Example 1—paper submissions

We will now consider part of the procedures when a paper arrived at the HCI’95 office (Fig. 5). The sub-process starts when the author sends the paper; Ann receives it through the post and then records the details of the paper in a database before filing the paper (ready for subsequent review).

For each activity we look at the triggering event.

Trigger ① is simply when the packet containing the paper arrives via a communication channel, in our case the postal mail. We could investigate the postal system in detail, but normally we would stop here, recording our expectations about its reliability and timeliness. The mode of communication therefore acts as a trigger for Ann to receive the papers. However, the failure or unreliability of the medium of interaction has serious implications for the system’s operation. A possible solution to guard against such a failure 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 postal mail; but this might result in a situation where humans, unlike software, may find the additional protocol too costly to maintain.

Ann did not immediately enter the paper’s details. Instead, when a small pile had accumulated she entered them together. Trigger ② is therefore the pile of papers on the desk. This trigger is an environmental cue which allows Ann to pick up the threads of her activities. Environmental cues are important triggers which serve as reminders. As soon as the details were recorded, Ann sent an acknowledgement and the papers were filed.

Both triggers ③ and ④ are such that in an interruption-free environment, the end of one activity is the trigger for the next. However, Ann may be interrupted for some length of time while she is in the midst of sending an acknowledgement and filing a copy of the paper. In case of an interruption, we look for a secondary trigger or a fall-back trigger. The fall-back triggers for ③ and ④ are the same as each other and the same as ②, in other




Fig. 5. Paper submissions.

words, the unfiled papers on the desk. Because the activities have the same trigger, an activity will potentially either be repeated after an interruption or omitted entirely (if Ann mistakenly thought an interruption had previously occurred). Clearly, keeping track of all the tasks in which we are currently engaged is a mental strain. If someone fails to complete or close tasks held in short-term memory, or is prevented from doing so by interference, the subject is liable to lose track of what she is doing and can consequently make errors. Happily, Ann’s memory was good enough and these problems did not arise in this case. However, interruption can have major consequences on the flow of work within a collaborative system [32]. For instance, in the next section we will see an example where failure does occur.

### 5.3. Example 2—refereeing process

Let us consider another process of the life cycle of a paper: the refereeing process. Fig. 6 below highlights the fact that the agents involved no longer reside within a single organisation. We have now crossed organisational boundaries and the whole process is entirely dependent on the referees based at several locations.

How does Ann coordinate the referees’ activities when there is a temporal gap between the dispatch of the papers and the return of the referees’ reports? Trigger ①, the deadline, enables Ann to regain control. If Ann does not receive the refereed papers by the date set for return, then she sends reminders to the referees. In our case there was only one deadline for all the papers, so that date was easy to remember. However, if each paper were allowed a different date for submission then Ann would have to keep track of deadline dates periodically (how does one remember to perform the action at the relevant time?). So we see that, in a long-term cooperative situation, especially when 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 failures can occur.

## 6. The 4Rs

Even though our initial focus was on individual triggers, we began to notice an emerging pattern as we recorded the processes during our case study. We call this pattern the 4Rs: request, receipt, response, release.




Fig. 6. Part of refereeing process.

Fig. 7 shows a simplified version of Fig. 5 which exemplifies the 4Rs. We can see a general structure emerging: *request*—someone sends a message (or implicitly passes an object) requiring your action; *receipt*—you receive it via a communication channel; *response*—you perform some necessary action; and *release*—you file or dispose of the things used during the process. At this point, if the functional goal has been achieved then the process can be considered to have reached completion.

The papers process in Fig. 7 is very similar to the process that one of the authors follows when dealing with e-mail. When the mail arrives, he reads it (or at least notes its arrival), but does not deal with it immediately—it stays in his ‘in-tray’ until he has replied or otherwise dealt with it. Only at that stage does he file it in a folder or discard it. If interrupted after replying, the original message is still in the in-tray (secondary trigger). Recently, whilst in the middle of replying to a message, the machine crashed (interruption). When some time later he again read his e-mail, he mistakenly (and unconsciously!) took the continued presence of the e-mail in the in-tray as signifying an interruption before filing (secondary trigger) and hence filed the message without replying.

Not only is the pattern of activities common between different processes, but we also see a similar pattern of triggers. ① is always simply some sort of communication mode and can be assessed for reliability and timeliness. The response activity is typically triggered by ②, the presence of a document or other object. The release activity triggered by ③, which is of the ‘immediately follows’ kind, removes that cue, but also relies on its existence as a secondary trigger. The problems with the author’s e-mail will occur elsewhere!

This pattern has various refinements: for example, when a note is made of a verbal request, adding an extra stage to receipt. Perhaps the most interesting variations are those concerning the response. In many cases there is more than one action required as part of the response. As discussed previously these may be considered to be at a finer scale than our analysis, but not always. For example, in Fig. 5 the response consists of two activities ‘enter record’ and ‘send acknowledgement’. In such cases we need to look very carefully at the triggers as it is quite likely that the two parts of the response have the same trigger. This was the case with the activity in Fig. 5 as both were triggered by the presence of the paper in the pile on the desk. On the other hand, in some situations, for example receiving information for filing, there may be no separate response as the response and release are merged. We will see further examples of complex responses in the second case study.

Fig. 5 also demonstrates a frequent aspect of the 4Rs. The response activity ‘send acknowledgement’ is itself a message to the author. It is frequently the case that the response of one 4Rs pattern forms the request activity initiating a new 4Rs pattern. A chain of such 4Rs patterns constitutes a sort of long-term conversation (see Fig. 8).




Fig. 7. The 4Rs.




Fig. 8. The 4Rs chain.

The 4Rs appears to be a pervasive, generic pattern, at a lower level than those identified in speech-act theory [37], and perhaps being the long-term interaction equivalent of adjacency pairs found in conversational analysis.

## 7. Applying the 4Rs—Student Placement Unit

The Placement Unit at the School of Computing and Mathematics at The University of Huddersfield is an extremely busy office environment. The Unit is staffed on a full- and part-time basis by administrative and academic staff, respectively, responsible for helping some 200 + sandwich course students secure 1-year placements in industry every year. Besides dealing with students seeking placements—involving processes in skilling-up, preparing students to apply and administering the recruitment cycle—the Unit also supports those already on placement—involving processes in assessment, monitoring and problem-solving. Contact with companies occurs via all media—post, fax, phone, e-mail and face-to-face—though chiefly by telephone which causes frequent interruptions. The outright winners in the interruption stakes, however, are the students for whom the Unit has an ‘open door’ policy between 10am and 4pm. For this reason alone the Unit was selected as an ideal focus for our next investigation of long-term office procedures.

In the process of our investigations a further opportunity presented itself—the Unit was to undergo major change. The MaPPiT Project<sup>1</sup> had just been launched at the School. Part of the project remit was to develop a process support system for the Placement Unit in line with a generic process model of placement activity. The decision was taken to purchase Lotus Notes to build the system. The application of our 4Rs framework was one of the methods chosen to analyse the current situation. The expectation was that we could uncover low-level issues to be addressed by Notes.

The following diagrams in Figs 9–11 exemplify the potential for the activity triggers to be seriously delayed—sometimes indefinitely—had they been left to be resolved by outside companies.

<sup>1</sup>MaPPiT—Mapping the Placement Process with Information Technology, a HEFCE-supported 2-year project under the auspices of the Fund for the Development of Teaching and Learning: Improving Student Learning on Sandwich Placements (FDTL 32/96). Details available at: <http://www.hud.ac.uk/schools/comp+maths/mappit/home.html> > .

### 7.1. Example 3—job adverts

The establishment of a new placement starts with the initial request from a company for a placement student. This is shown in Fig. 9 and we can see that it is a 4Rs pattern, but with a two-stage response. Recall that one of the dangers of multi-stage responses was that the triggers at ② and ③ are often similar, leading to problems. We were thus particularly looking for these.

In this process the request is initiated by the company. Many regular placement providers themselves diarise to send the Unit placement details and requirements for the forthcoming year. They ‘drive’ the process by setting deadlines for the Unit and the students via closing dates for applications. Another group of companies have already been triggered into sending a job description by a standard letter from the Unit. The Unit maintains a diary and companies are contacted on a fortnightly basis with the standard letter. The Unit’s work commences when job descriptions are received. So we looked at ① and asked the administrative staff how they would know if details failed to arrive, thus breaking the chain of activities. At present the only back-up is the diary so a time delay occurs between the failure occurring and the next fortnightly, and sometimes monthly, check for responses from the previous month’s companies. A follow-on question would be, how do you remember to look in the diary? At present the answer is that the paper-based diary remains highly visible on the Placement Officer’s desk. However, as the year progressed, we noticed the diary being checked less and less.

We turn now to the next activity of ‘Record Details’ which, ideally, should directly follow on from the first activity. This class of trigger is insecure as it is liable to interruptions—a common occurrence in the Unit. The staff member then has to remember what to do next. Usually there is the environmental cue of paper on the desk—jotted down from a phone call, or a fax copy or a letter—or an open e-mail message on screen. The follow-on question was asked, what if you do not record the fact of the receipt of details in the diary? Another staff member could check the diary, see the assumed non-receipt of details and annoy a company intensely by chasing for details already sent in unless that person knew to double-check the Job Adverts Log (Response b) first, or to check the company file for the ad. All this checking should nevertheless be unnecessary. The scenario is one of much paper-chasing.

With Notes in mind the project team accepted these current problems as needing resolution. In the increasingly competitive placements market the Unit can afford neither to let certain activities drift aimlessly, nor to be driven solely by companies whose priorities and timescales almost certainly do not coincide with those of the Unit. The




Fig. 9. Initial job advert.



diary could easily become electronic with built-in ‘navigators’ (agents) that automatically trigger reminders to execute activities. Gone is the need to remember to check the diary as reminders appear in individuals’ To-Do lists. Even To-Be-Done-To lists can be constructed. Receipt of a job description need only be recorded in one place so there is only one checking activity before chasing a company. Furthermore, the electronic record means that any inconsistency between the recorded job details and the Job Adverts Log can be displayed and thus act as a trigger at ③.

### 7.2. Example 4—submission of CVs

A large proportion of the Unit’s placement providers are happy to accept standard CVs from the students so it is vital that CVs are lodged with the Unit and checked by placement tutors very early in the year. Fig. 10 shows this process which is a straightforward 4Rs pattern. At ① these two activities usually happen face-to-face so there is little risk of breakdown. We noticed that some students, unfortunately, ignored the office hours of the Unit and so ‘posted’ CV disks under the door after staff had left for the day, thus risking the cleaning staff picking up the disk and it subsequently being disposed of, damaged, lost or misplaced. Similarly the Unit’s activities here are all exposed to interruptions and therefore incompleteness. Sally, the Placement Secretary, is accustomed to the sub-process of receiving and recording students’ disks. Staff changes in the Unit this year brought some new faces and when Sally took her annual leave the circumstance arose where a student made alterations to the CV, returned the disk and the new member of staff promptly lost the update having been interrupted several times to do other more complex tasks. Furthermore, disks are sometimes found to be corrupt when the CV is required, leading to another set of interactions with the student.

The planned Notes implementation for this completely bypasses the current error-prone process. Students will fill in a CV template using a web browser. The CV is then automatically submitted to a Notes database which logs the receipt and sends an e-mail to the Unit to confirm that the CV has been submitted on time. Students can update the CV at any time without bothering the administrative staff and due to Notes replication we can be sure that the latest version of the CV is being sent to potential placement employers.

### 7.3. Example 5—company decisions

After seeing the students CVs companies will decide on students to shortlist for interviewing. The final diagram (Fig. 11) demonstrates how the pace of interaction can really




Fig. 10. Students submit CVs.

slow down when pursuing students to arrange interviews or to provide feedback if rejected. At ② the pace slows considerably once the students are on vacation and hard to track down. Much time can be spent trying to contact students on a list of phone numbers where they might be located. Assuming this is successful, the next activity can be stalled by a phone call or face-to-face enquiry. We return to relying on an individual's good short-term memory and/or an environmental cue to ensure the sequence is fulfilled. Although the process appears at first to be a simple 4Rs pattern, we have put a question mark against the last activity. The release usually consumes or destroys the environmental cues which have prompted previous activity. It is not clear that this is the case for this process—what are the environmental cues? The company's decision will arrive in a letter or be recorded on paper, but the slow pace of the response means that the cue may be lost or grow 'stale', ceasing to be salient because it is around too long.

Redesigning this sequence to be supported by Notes, the project team decided it would be better to start recording receipt of the company contact when the contact happens—invariably a phone call, fax or letter—on a Notes form; then everything is on screen and, if incomplete because of an interruption, cannot be discarded without being prompted to complete the form details. Note how this has established an environmental cue within the electronic world of the Notes database. In the revised process this cue is removed when the final updates to the company details are completed, thus making the pattern a true 4Rs with robust triggers throughout. Furthermore, if the response stage becomes drawn out and relies on students responding to telephone messages or e-mails, there is the possibility of automatically signalling if the expected reply is not forthcoming, thus supplying a to-be-done-to facility.

#### 7.4. Issues

Note that, in the three examples above, different levels of automation have been suggested by the 4Rs analysis. At one extreme this has involved the complete bypassing of the human process, but in the others only parts are automated. Most important, the 4Rs analysis has ensured that the Notes implementation does not hide existing triggers, as is often the case with electronic filing, but instead is explicitly designed to enhance the triggers with automatic reminders and electronic environmental cues.

This second case study has validated the general applicability of an analysis based around triggers and the 4Rs. The 4Rs was remarkably successful in describing patterns of activity and in prompting appropriate questions to drive the Notes implementation. However, the study also brought to light some new features of the 4Rs.




Fig. 11. Company decision.

Notice that the salience of certain kinds of triggers was observed to change with time. In example 3, we saw that the diary was consulted less often as the academic year progressed. Presumably this reflects the change in the operation of the unit at different times of year, being more proactive earlier on and more responsive later. In example 5, the letter holding the company's decision would have been initially salient sitting in a desk pile, unfortunately during the time it is most salient, it must be ignored as attempts to contact the student by e-mail or letter may be outstanding. Environmental cues may therefore fail for exactly the same reasons that our memory finds to-be-done-to items difficult. In addition, in example 10, the same process had different kinds of triggers on different occasions. In general, we cannot assume that the detailed triggers are homogeneous over time, but must establish by enquiry or observation whether triggers vary in kind, or salience.

## 8. Related approaches

The nature of these studies bears some similarity to several disciplines in the general field of the 'social analysis of work', particularly workflow (including speech-act theory), ethnography and ethnomethodology. The following points, however, summarise the critical differences between our approach and the above.

Dealing firstly with workflow, the term in its precise sense<sup>2</sup> implies technological solutions 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. Similarly, they easily lend themselves as a support mechanism for the current trend in organisations to be process-focused—be it business process re-engineering [16] or any radical re-structuring of the way organisations operate.

With regard to the nature of our investigation, the principal limitation of workflow (besides the technological bias already mentioned) is that the concept of 'workflow management' hints at cultural change. The very installation of such procedural systems creates a culture of its own. To some extent or other they all have some kind of model of the user and the nature of the organisation. Stipulating procedures can be acceptable, 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 cultural change.

Previous research [12] has also highlighted the problems with workflow moving beyond the bounds of the organisation unless through some formalised collaboration and also of ignoring subtle differences between individual goals and the process goal. The process we considered, however, was as much inter- as intra-organisational. This was a major issue for us to address: as 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, initially, to seek ways of improving workflow by automating the processes of work or even facilitating them via computerisation. Still, as has become

---

<sup>2</sup>One 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> .

apparent in our case studies, our results do have design implications. To avoid confusion or disagreement over our use of the term ‘workflow’ we differentiate our approach here by referring to it as an investigation of the ‘flow of work’. The final difference for our study is its very precise focus—that is, the targeting of events triggering activity.

Despite some surface similarities, speech–act theory (SAT) contrasts quite strongly with our approach in that its basic structure comprises all possible stages in conversational interaction [37]. In some ways our approach is more abstract—for instance, the arrival of an e-mail message may be a potential trigger whereas speech–act theory would analyse the contents of the e-mail itself. In contrast, the 4Rs pattern is at a lower level of granularity than speech–act patterns such as conversation for action (CfA)—that is, each action pair in a SAT diagram expands to a complete 4R.

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 [17]. 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 by 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 [30]. In contrast, our work began with a sharper focus as previously described. On the one hand, this means that we ignore aspects of a situation that an ethnographer would record. On the other, ethnographers’ open-endedness is seen as a weakness when it is used for requirements capture [2]. By being more restricted, our approach is better suited to inform systems design.

Ethnomethodology has also been used within HCI [35] as a particular form of sociological analysis [15]. Ethnomethodologists observe, collect and analyse data and decide what is relevant about work activity as it really is, rather than as an idealised conception of work, as can be the case with process-modelling and workflow. 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. [2] 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 distinguishes 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 [3,4,18,19] for how work is executed has not escaped the notice of sociologists, least of all ethnographers/ethnomethodologists. Traditionally such studies stressed the social actors within their environment, the close teamwork, at the expense of the surroundings in which people work, but more recent studies of office work [20,32] have brought the surroundings and artefacts into the limelight, in particular the importance of paper [33]. This trend is followed in our work, but with a more specific formulation of the purpose of artefacts as triggers for activity.

More formal techniques have also been applied to the study of time and collaboration including Petri Nets [21,26], various forms of temporal and modal logic [10,31,22] and process algebras such as LOTOS [27]. Any of these could be used to capture the precedence relationship between activities, but not the significant issue for the paper which has been the nature of triggers. This is epitomised by the temporal logic eventually

operator (written as a diamond,  $\diamond$ ). This says that something will happen ‘eventually’, but not how soon. In process algebras similar issues have led to a whole debate about the semantics of ‘fairness’ [14], trying to say that things that can happen should eventually happen. There are various additions of real-time constructs to these notations which put deadlines on how long gaps can be between certain events. However, these are really about specifying what should happen, rather than looking at the rich ecological aspects of why they do happen in socio-technical systems.

Of particular interest is [23] and his use of Petri Nets to model workflow. This notation uses the word trigger: an event which “causes (an activity) *a* to be performed”. This is similar to our use of the term trigger, and does distinguish this from an enabling event, but does not go on to investigate the ecology of triggers at the level of detail found in this paper.

Finally, timeline techniques have also been found to be useful in low-level analysis and also in presenting information to end users [29]. In particular, timelines have been used in status–event analysis to look at the reasons for delays in e-mail notification and to assess feedback problems with on-screen buttons [13].

## 9. Implications for TAU research and practice

The analysis we have used was initially targeted at increasing our theoretical understanding of long-term interaction; however, in use we have realised that it has direct design implications. It can be used to determine whether a process is robust to interruptions, forgetfulness, etc. and, if not, identify why not and where the problems arise.

The reliability of the work process can be assessed by asking 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, perhaps the whole process fails to continue because something goes wrong. The combination of a process model together with a well-founded assessment 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 seize up, or will the failure eventually be noticed? Note that this is not simply an ad hoc procedure. Following our approach, one can systematically go to each trigger and ask: what happens to the entire process if the trigger fails? Furthermore, by looking at the process as a whole we can improve our assessment of the reliability of any trigger. For example, if the trigger for an activity is that a report is in someone’s in-tray, we can examine the wider context and assess the likelihood of whether the report will indeed be there when required.

We fully expected, and found in our studies, that environmental cues are one of the principal and most robust triggering mechanisms. Several ethnographic studies have noted the importance of the ecology of the workplace, including whiteboards, calendars, individual papers and piles on desks [20,32,33]. Indeed, in many cooperative processes there may be little direct communication, instead the parties coordinate by implicit communication through the artefact [9].

The distinctive nature of the work reported in this paper is that we have focused on a particular role of these environmental cues, namely their ability to remind and trigger future actions. This is especially important if there are plans to automate parts of an office procedure. Whereas many studies have concluded that paper is important, we have developed an understanding of why paper is important. This has theoretical implications for anyone investigating the ecology of the work setting and has practical design ramifications allowing us to see whether automation will break an existing work system and if so whether alternative cues can be provided.

## 10. Summary

We have seen how long-term interaction may pose problems over and above those of higher pace interaction. In particular: (i) *action–effect gap*—users may have difficulty in recalling the context of a delayed response; (ii) *stimulus–response gap*—they may forget to act themselves if they cannot react instantly to a request; and (iii) *missing stimulus*—the whole process of interaction may break down if an expected external response is not forthcoming. The second of these is the reason for to-do-lists and aide-mémoires, the last requires to-be-done-to lists or similar reminders.

The above considerations lead to an emphasis on understanding the triggers which initiate action. In particular, it is often implicitly assumed that an activity is triggered by the completion of the previous activity, whereas in practice this is rarely the case because of the competing demands and interruptions during normal office life. Triggers are important because they not only determine when an activity occurs, but whether it happens at all.

During the analysis of our first case study we noticed a recurrent pattern of activities—the 4Rs: request, receipt, response and release. We believe that this is a fundamental unit of long-term work. The existence of generic patterns makes it easier to uncover problem situations quickly and to take solutions found in one situation and adapt them to another. Our case studies show that the 4Rs is normal—the same pattern recurs with similar triggers and similar failure modes. We have also seen that it is normative—if the 4Rs pattern is nearly followed, but with some deviation, this has been seen to be an indication of possible problems.

As we have noted, problems are particularly likely when a functioning paper-based system is automated. In particular, this can often lead to the loss of important environmental triggers. Our analysis can target potential problem spots before they occur. The use of the 4Rs as part of a major Notes implementation has allowed us to design semi-automated processes where physical environmental cues are replaced or augmented with electronic cues.

Our theoretical understanding has been validated and deepened by the initial HCI'95 Conference case study and the use of our method in the Placement Unit Notes implementation. This is an area where theoretical understanding and practical application can proceed side by side. Interaction-in-the-large involves aspects of organisational modelling and CSCW, but also poses interesting system design issues. In HCI it is an under-studied, but exciting area.

## References

- [1] G. Abowd, A. Dix, Integrating status and event phenomena in formal specifications of interactive systems, in: SIGSOFT'94, ACM Press, New Orleans, LA, 1994, pp. 44–52.
- [2] R.J. Anderson, Representations and requirements: the value of ethnography in system design, *Human-Computer Interaction* 9 (1994) 151–182.
- [3] R. Bentley, J.A. Hughes, D. Randall, T. Rodden, P. Sawyer, D. Shapiro, I. Sommerville, Ethnographically-informed systems design for air traffic control, in: Proceedings of CSCW'92 Toronto, Ontario, ACM Press, 1992a, pp. 123–129.
- [4] R. Bentley, J.A. Hughes, D. Randall, S.Z. Shapiro, Technological support for decision making in a safety critical environment, Technical report, CSCW/5/92, Computing Department, Lancaster University, 1992b.
- [5] S.A. Brewster, P.C. Wright, A.D.N. Edwards, The design and evaluation of an auditory-enhanced scrollbar, in: Proceedings of CHI'94, ACM Press, Addison-Wesley, Boston, MA, 1994, pp. 173–179.
- [6] A.J. Dix, Formal methods for interactive systems, Academic Press, New York, 1991.
- [7] A.J. Dix, Pace and interaction, in: Proceedings of HCI'92: People and Computers VII, Cambridge University Press, 1993, pp. 193–207.
- [8] A.J. Dix, Que sera sera—The problem of the future perfect in open and cooperative systems, in: Proceedings of HCI'94: People and Computers IX Glasgow, Cambridge University Press, 1994a, pp. 397–408.
- [9] A.J. Dix, Computer-supported cooperative work—a framework, in: D. Rosenberg, C. Hutchison (Eds.), Design Issues in CSCW, Springer Verlag, Berlin, 1994b, pp. 9–26.
- [10] A.J. Dix, LADA—A logic for the analysis of distributed action, in: F. Paternó (Ed.), Interactive Systems: Design, Specification and Verification (1st Eurographics Workshop, Bocca di Magra, Italy, June 1994), Springer Verlag, Berlin, 1995, pp. 317–332.
- [11] A.J. Dix, D. Ramduny, J. Wilkinson, Interruptions, deadlines and reminders: investigations into the flow of cooperative work, RR9509, University of Huddersfield, 1995.
- [12] A.J. Dix, D. Ramduny, J. Wilkinson, Long-term interaction: learning the 4Rs, in: CHI'96 Conference Companion, ACM Press, Vancouver, 1996, pp. 169–170.
- [13] A.J. Dix, J. Finlay, G. Abowd, R. Beale, *Human-Computer Interaction*, 2nd edn, Prentice Hall, Englewood Cliffs, NJ, 1998 (first edition 1993).
- [14] N. Francez, *Fairness*, Springer-Verlag, Berlin, 1986.
- [15] H. Garfinkel, *Studies in Ethnomethodology*, Prentice Hall, Englewood Cliffs, NJ, 1967.
- [16] M. Hammer, J. Champy, *Reengineering the Corporation—A Manifesto for Business Revolution*, Nicholas Brealey, 1993.
- [17] M. Hammersley, P. Atkinson, *Ethnography: principles in practice*, Tavistock, 1983.
- [18] C. Heath, M. Jirokta, P. Luff, J. Hindmarsh, Unpacking collaboration: the interactional organisation of trading in a city dealing room, in: Proceedings of ECSCW'93 Milan, Italy, Kluwer Academic Publishers, Dordrecht, 1993, pp. 155–171.
- [19] C. Heath, P. Luff, Crisis management and multimedia technology in London Underground line control rooms, *Journal of CSCW* 1 (1) (1994) 69–94.
- [20] S. Herskind, Computer support for temporal aspects of coordination of cooperative work, in: ECSCW'97 Conference Supplement, Lancaster, UK, Kluwer Academic Press, Dordrecht, 1997, p. 67.
- [21] C.W. Johnson, J. McCarthy, P.C. Wright, Using Petri Nets to support natural language in accident reports, *Ergonomics* 38 (6) (1995) 1265–1283.
- [22] C.W. Johnson, The impact of time and place on the operation of mobile computing devices, in: Proceedings of HCI'97: People and Computers XII Bristol, UK, 1997, pp. 175–190.
- [23] S. Joosten, Trigger modelling for workflow analysis, in: R. Oldenbourg (Ed.), Proceedings CON'94: Workflow Management Vienna, 1994, pp. 236–247.
- [24] D.A. Norman, New views of information processing: Implications for intelligent decision support systems, in: Hollingel E. et al. (Eds.), *Intelligent Decision Support in Process Environments*, Springer-Verlag, Berlin 1986.
- [25] D.A. Norman, *The Psychology of Everyday Things*, Basic Books, New York, 1988.
- [26] P. Palanque, R. Bastide, Formal specification and verification of CSCW, in: People and Computers X — Proceedings of the HCI'95 Conference, Cambridge University Press, 1996, pp. 213–231.

- [27] F. Paternó, G. Faconti, On the use of LOTOS to describe graphical interaction, in: Proceedings of HCI'92: People and Computers VII, Cambridge University Press, 1992, pp. 155–173.
- [28] S.J. Payne, Understanding calendar use, *Human–Computer Interaction* 8 (2) (1993) 83–100.
- [29] C. Plaisant, B. Milash, A. Rose, S. Widoff, B. Shneiderman, Lifelines: visualising personal histories, in: Proceedings of CHI'96 Vancouver, ACM Press, 1996, pp. 221–227.
- [30] D. Randall, Ethnography for systems development: bounding the intersection, Tutorial Notes HCI'95, University of Huddersfield, 1995.
- [31] S. Reeves, Specifying and reasoning about CSCW, in: Design, Specification and Verification of Interactive Systems '96, Namur, Belgium, Springer-Verlag, Berlin, 1996, pp. 366–383.
- [32] M. Rouncefield, J.A. Hughes, T. Rodden, S. Viller, Working with 'Constant Interruption' CSCW and the Small Office, in: Proceedings of CSCW'94 Chapel Hill, NC, ACM Press, 1994, pp. 275–286.
- [33] A. Sellen, R. Harper, 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, 1997, pp. 319–326.
- [34] A. Shepherd, Task analysis as a framework for examining HCI tasks, in: A. Monk, N. Gilbert (Eds.), Perspectives on HCI: Diverse Approaches, Academic Press, London, 1995, pp. 145–174.
- [35] L.A. Suchman, Plans and Situated Actions: The Problem of Human–Machine Communication, Cambridge University Press, 1987.
- [36] B. Warboys, Reflections on the relationship between BPR and software process modelling, in: Proceedings of ER'94, Springer-Verlag, Berlin, 1994, pp. 1–9.
- [37] T. Winograd, F. Flores, Understanding Computers and Cognition: A New Foundation for Design, Addison-Wesley, New York, 1986.
- [38] Workflow Management Coalition, Glossary of Terms, 1994. (<http://www.aii.ed.ac.uk/WfMC>).
- [39] A. Wood, A.K. Dey, G.D. Abowd, CyberDesk: automated integration of desktop and network services, in: Proceedings of the 1997 conference on Human Factors in Computing Systems, CHI '97, 1997, pp. 552–553.