PETRA: Participatory Evaluation Through Redesign and Analysis*

CSRP No. 375, March 1995

Susi Ross **Magnus Ramage**† **Yvonne Rogers**
School of Cognitive and Computing Sciences
University of Sussex, Brighton BN1 9QH, U.K.
Phone +44 1273 678524, Fax +44 1273 671320.
susi,yvonner@cogs.susx.ac.uk, magnus@comp.lancs.ac.uk

Abstract

Compared with single user-computer interactions, evaluating CSCW is difficult. We argue for *multiplicity* - of theory, method and perspective - in CSCW evaluation. This allows us to address both theoretical concerns and practical design issues, and to incorporate the expertise and experiences of both evaluators and participants. We propose the PETRA framework, incorporating a theoretically-driven 'evaluators' perspective' (ETA) to investigate the collaborative activity, and a practical, user-focused 'participants' perspective' (ETR) to evaluate the interface of the supporting tool. Our particular instantiation of PETRA focused on collaborative writing, both in a face-to-face context, and supported by a computer-based group editor, ShrEdit. We investigated the development of shared understanding in the two different mediating settings; and used a PD-inspired rapid prototyping session to elicit participant reactions to and redesigns of the tool interface. Our ETA findings show that computer-supported shared understanding develops technologically, using social coordination as a repair mechanism; the ETR findings show that the collaborative tool must be particularly sensitive to issues of awareness, communication, focus and ownership.

Keywords: CSCW; Evaluation; Multiplicity; Collaborative Writing; Distributed Cognition; Shared Understanding; Participatory Design; Usability Criteria; User Involvement.

1 Evaluation and CSCW

Evaluating CSCW systems - be they shared drawing and writing tools, shared databases, shared to-do lists, videoconferencing or other - is difficult. One of the main problems is determining how to assess the numerous permutations of interaction that can occur between the users and the system. Compared with single-user systems there are a multitude of possibilities and problems that can develop in collaborative situations: these may prevent the supporting CSCW systems from being used optimally, or even collaboratively (Bannon, 1993; Grudin, 1988; Orlikowski, 1992). A central question, therefore, is how can we effectively evaluate systems designed to support multiple interacting users? Can we borrow and adapt existing evaluation methods or do we need to devise new techniques for analysing CSCW? In making our decision should we focus on the individual, the group or the organisation? In doing so, what level of detail should the evaluation be pitched at? Should we be striving for a full description of all the interactions at the various interfaces, or should we restrict our analyses to either general conceptual issues or

*Submitted to *Interacting with Computers*.
†Now at the Computing Department, Lancaster University, Lancaster LA1 4YR.

specific quantitative measures? Alternatively, is it possible or desirable to take the middle-road; achieving a balance between detail and generality?

The approach adopted will obviously depend on the goals of the evaluation. If the objective is to analyse the working practices and the nature of collaboration in a given organisational context, to inform system design, then the use of social methods such as ethnography would seem appropriate. Alternatively, if the aim is to analyse the usability of, say, World Wide Web browser tools for different user groups, then perhaps a cognitive task analysis would be more appropriate. In many situations however, it is not easy to determine the objectives of the evaluation because the purpose of the collaborative tool is not clear. For example, videoconferencing allows participants to see images of each other when talking at a distance. Given the ease with which people communicate via telephone, one might reasonably ask if there is any added value in being able to see distant callers in this format: much anecdotal evidence and the few empirical studies carried out suggest it is minimal and can even have adverse affects (Harper & Carter, 1994; Pagani & Mackay, 1993). This can be viewed as an example of the 'solution looking for a problem' phenomenon: potential benefits of the prototype tool are unknown, and useful contexts have yet to be discovered. Compared with evaluating single-user systems that have specific objectives (e.g. evaluating the new functionality of an update version), and for which benchmark tests can be relatively easily devised, evaluating CSCW systems is proving to be much more problematic.

Given our current lack of understanding of how best to utilise and benefit from the new and largely unfamiliar generation of CSCW systems that is emerging (e.g. video-links, shared drawing, writing and editing tools, shared repositories) - at work, home and elsewhere - we suggest there needs to be much more emphasis on determining their usefulness in different contexts. In particular, we need to develop evaluation methodologies that can show how the various tools are able (or not) to support collaboration, as well as provide HCI-style usability indicators (e.g. easy to learn, easy to use) appropriate for assessing groups of users using the CSCW tool together.

The aim of this paper is to begin addressing this situation. We start by advocating eclecticism, whereby different methods and theoretical frameworks are combined. Such an approach should allow for a more reflexive analysis, between different theoretical concerns and practical design issues, and in doing so force us to explore and make explicit many of our assumptions about the CSCW tools and the nature of the collaborative activities that they are intended to support. To this end we have developed a multi-perspective framework, called PETRA, which stands for 'Participatory Evaluation Through Redesign and Analysis'. Essentially, the framework brings together a theoretically-driven 'evaluator's perspective', based on a combination of theories that are concerned with collaboration (evaluation through analysis, or ETA), and a practically-based, user-focused 'participant's perspective', drawing from heuristic evaluation techniques and participatory design (evaluation through redesign, or ETR).[1] The primary objective in developing this framework was to obtain a detailed understanding of a collaborative activity whilst also considering how to inform the re-design of the interface for a CSCW tool. In particular, the former was intended to provide theoretical insight into the social and cognitive issues that are important for shared understanding when working together at different sites, whilst the latter was intended to provide an understanding of the system and interface attributes desirable or neccessary to successfully support the former.

In the next section we discuss the multiplicity of theory, method and perspective when evaluating CSCW; and we introduce the different parts of our PETRA framework. In Section 3 we illustrate PETRA in practice, describing how we used multiplicity in our study. Sections 4 and

[1] The participatory 'P' signifies the inclusion of both participants' and evaluators' opinions, and of multiple perspectives within the framework.

5 present findings; discussion and conclusions follow in Section 6.

2 Multiple methods, perspectives and theories

The diverse fields informing CSCW (eg. HCI, sociology, software engineering, psychology, cognitive science, management studies) bring with them a rich variety of evaluation methodologies, perspectives and theoretical and conceptual frameworks. For example, from HCI we have usability engineering, prototyping and user modelling; from sociology we have ethnomethodology, conversation analysis and ethnography; from software engineering we have requirements engineering, formal methods and soft systems methods; from psychology we have group dynamics theory, information processing and lab experiments; from cognitive science we have distributed cognition; and from management studies we have business process engineering, surveys and focus groups. There are of course many more, each having its particular strengths and limitations. In developing evaluation methodologies for CSCW, therefore, we have to decide which of the multitude of approaches are appropriate and how best to adapt and combine them.

Selecting methods

In determining which methods are appropriate, various organisational, temporal and cost constraints will dictate, to a large degree, what is feasible and practical. For example, the use of lab experiments or full-scale ethnography, where a vast amount of detailed data is generated, may be an appropriate level of analysis for a PhD study, but is unlikely to be practical in the competitive world of software and system development. To accommodate such demands on resources and time, researchers in HCI and CSCW have adapted existing lengthy and difficult methods, for example, 'quick and dirty' ethnography (Hughes et al., 1994) and 'cognitive jog-through' (Rowley & Rhodes, 1992). For our PETRA framework, we decided to tailor a form of heuristic evaluation (Nielsen, 1993) in combination with methods adapted from Interaction Analysis (Suchman & Trigg, 1991) to enable us to collect a variety of data from the different sites the tool was used in. Specifically, we decided that what was important was both to record and analyse the interactions and conversations that occured between the participants in different settings (allowing us to focus on the collaborative activity), and to obtain the users' opinions about the tool before, during and after using the tool (allowing us to focus on the appropriateness and usability of the tool for supporting the activity, and eliciting recommendations of how the interface of the tool could be improved).

Dual perspectives: evaluators' and participants'

When considering which methods to adapt and combine one needs to consider the 'source' of the data that is being collected: whose opinions or perspective does it represent? While formal methods (e.g. usability engineering, lab experiments) can elicit quantitative and so-called 'objective' data, which researchers can readily analyse and re-represent as statistical data, they do so at the expense of overlooking the users' own experience and opinions of the system. In contrast, the use of more informal methods (e.g. interviews and observations) can elicit first-hand users' impressions, but in doing so only provides qualitative and so-called 'subjective' accounts. We argue that a CSCW evaluation methodology should strive towards achieving a balance between different analytic approaches: taking into account both the users' and researchers' perspectives of the collaborative activity and how this is supported by the tool. In our PETRA framework, we call these the 'participants' perspective' and the 'evaluators' perspective'.

Theoretical frameworks

Various theoretical and conceptual frameworks are currently being developed and applied to CSCW; these include Distributed Cognition (Rogers & Ellis, 1994), Activity Theory (Kuuti & Arvonen, 1992), Organisational Theory (Sproull & Keisler, 1991), Socio-Technical Design (Mumford, 1994) and Computational Mechanisms of Interaction (Schmidt & Simone, 1993). Besides helping to construct a body of knowledge within CSCW, immersing ourselves in theory can enable us to discover aspects of the collaborative activity that we had not considered before. In addition, theory can be used as a background to frame problems, pose questions, describe and explain phenomena, inform and guide system evaluation and design (Rogers et al., 1994). Hence, having a theoretical perspective can enable the evaluator to re-describe and re-represent the 'raw' data collected in the evaluation as a richer and more abstracted account of the collaborative activity. Not only can this provide a complementary explanation to what the users themselves have interpreted, but it can also give a wider picture of the state of affairs that the participants themselves may miss.

To guide our evaluator's analysis (ETA) of the collaborative activity we chose to draw from ideas arising from distributed cognition, ethnomethodology and conversation analysis, which we ourselves were then immersed in. The conceptual notions of shared understanding, breakdowns, and the propagation of representational states across different media were used to characterise the various interactions in the collaborative activity.

Evaluation through redesign (ETR) and participatory design

The emphasis on the importance of both the participants' and evaluators' interpretations allows us to rethink the role of users in evaluation. Participatory Design (PD) has always emphasised the involvement of the end-users in every aspect of the design process, overcoming the artificial divide between *expert design* and *user testing* (Schuler & Namioka, 1993). Within PETRA, we use the methods of participatory, iterative design *as part of* evaluation, allowing participants to be *actively* involved in the evaluation process. We also sought to integrate evaluation into the design process, as "design can emerge from an understanding of current (mis)-use of systems" (Bannon, 1993, p.4). This allows for redesign proposals to be explicitly incorporated into the evaluation, further informing the emerging evaluative findings. Thus, our 'evaluation through redesign' (ETR) method uses a PD-inspired prototyping session to encourage participants to articulate directly the problems they encounter, and suggest their own solutions, by exploring ways in which the application could be redesigned.

Participants are given the opportunity to develop, discuss, and test their critiques, giving both themselves and the evaluators a better understanding of their user experiences with the system. From this perspective, the perceived problem and the proposed solution are inextricably linked: traditionally *we, the evaluators* have sought to identify problems and then to formulate solutions; in ETR *they, the participants* can say "I think I'd prefer it this way", without having to identify 'a problem' as such. Adopting this approach also enables us to avoid the problematic situation facing analysts: trying to force design implications from rich descriptive accounts of the collaborative activity or work practice (for a discussion of this see Plowman et al. (1995)). This we leave up to the participants to suggest.

In considering how to evaluate CSCW systems, we need an understanding of the collaborative group interactions and of the process of carrying out the collaborative task. That is, we need to evaluate the collaborative *activity*. We also need an appreciation of the facilities required to support the collaborative task in a computer-based, distributed setting. Thus we also need to evaluate the collaborative *tool* supporting the activity. Within PETRA, the evaluator-perspective methods are geared more towards evaluating the activity, while the participant-

perspective methods are more suited to evaluating aspects of the tool. Yet evaluating activity or tool alone results in an incomplete picture. By recognising the interdependence of collaborative activity and supporting tool we can apply both evaluator and participant methodologies within the same study. This hopefully yields *complementary* and extensive coverage of the collaboration, yet without requiring too much time or resources.

In sum, the PETRA framework integrates multiple theories, methods, and perspectives. In doing so it incorporates the expertise of both participant and evaluator; combines theory and practical experience; and integrates use, iterative design and evaluation. How we used the PETRA framework in practice is described next.

3 PETRA illustrated: our analytic framework

Here we describe one particular instantiation of PETRA, devised to suit not only the activity (collaborative writing) and the tool (ShrEdit) in question, but also our goals and resources. The 'evaluation-specific' goals were to evaluate the efficacy of ShrEdit as a tool for the support of distributed computer-based collaborative writing; and to investigate some of the processes involved in the activity of collaborative writing via various mediums. We were also interested in testing out the PETRA approach; in developing some of the theoretical concepts underlying the evaluators' perspective; and in exploring the potential for 'participatory redesign' as a part of any evaluative method. Our resources were limited: three users; two evaluators; no budget; three machines in distributed locations; one video camera; and one tape recorder. We briefly describe the activity and tool next: the remainder of the section discusses the practical activities and analytic framework we devised around these constraints.

Within CSCW, collaborative writing is concerned with the study of multi-author writing processes, and the creation of group editors which support this process. Co-authoring is an everyday occurrence in numerous paper-based tasks, and is, potentially, the most common kind of computer-supported collaborative work. Computer-based collaborative work offers the advantage of connecting multiple distributed participants (often synchronously), thus opening new opportunities for people to work together where before it was too time-consuming, expensive or cumbersome. Here there are no precedents, and even more need for effective evaluation of the resources devised to exploit such opportunities. What do we know about collaborative writing in a face-to-face context, and will the same mechanisms be relevant to a computer-based setting? From paper-based collaborative writing we need to know *who* is involved in the writing process; *what* kind of document they are creating; *why* they are co-authoring; and *how* the process is structured. Then we must consider if it is either neccessary or desirable to recreate such mechanisms in a computer-based setting (Beck, 1994), and if the essentials of interaction are not altered by the very medium being used. We hope to address some of these issues in our study, particularly through our use of the evaluator-perspective. It should be emphasised that we use collaborative writing as an exemplar in our study of evaluation in CSCW: what we describe in this context is both relevant and transferable to other kinds of CSCW.

Our collaborative tool is ShrEdit: an experimental, explorative collaborative text editing tool developed for the Apple Macintosh™ by a team of researchers at the Cognitive Science and Machine Intelligence Laboratory, University of Michigan: "Our focus in the implementation was to provide a shared workspace with group awareness functionality. We are not interested in providing a full functioning text editor" (CSMIL, 1992, p.3). It has basic functionality, supporting only the most fundamental editing and formatting facilities (e.g. cut and paste, font size and style customisation). It was designed to allow several people to edit the same document simultaneously, whilst each using their own machine (however distant the location). Each user can see all changes made by the others, and make changes themselves. A simple 'ownership'

system is used to prevent conflict, and ensure coordination: text being changed by one user is 'locked' (indicated by a padlock symbol) so that no-one else has simultaneous editing access to it. Both public and private windows are available, allowing work to be developed privately before being inserted into the shared document. Awareness of other users' actions is provided by the status panel, which allows a user to find the current editing location of another user, or to 'track' (follow) another's actions through the document. However, it is not otherwise possible to find out what another user is doing: in particular, to see who is entering a particular piece of text. We wished to evaluate the usefulness of ShrEdit as a collaborative text editor: in particular, we were interested in the attempted provision of awareness information. Were ShrEdit's mechanisms successful or sufficient to support group awareness? If not, why not, and what might be required instead?

Evaluating the activity: shared understanding through different media

We decided to focus our evaluator's analysis (ETA) of the collaborative activity on shared understanding through different communicative media (talk, paper and computer-based). By shared understanding we mean the way in which two or more people relate their common background and experiences to understand collectively what each is talking about. This concept has been characterised extensively in the literature as being central to communication, by social psychologists, cognitive scientists, sociolinguists and sociologists alike, under various guises including mutual knowledge, common ground, social organisation and intersubjectivity (Clark & Brennan, 1991; Garfinkel, 1967; Hutchins & Klausen, 1992; Krauss & Fussell, 1991). A major concern has been to explicate the mechanisms by which speakers establish and maintain shared understanding during conversations. It has been proposed that speakers achieve this by formulating their contributions in relation to their awareness of what the other persons know and do not know. Moreover, such common ground is never static, but has to be continuously updated moment by moment and from context to context.

In outlining the different mechanisms involved in shared understanding (e.g. conversational acts, rhetorical devices, repair strategies and non-verbal behaviour like gaze, nodding) some researchers have also considered how they change when different media are used. For example, Clark and Brennan (1991) summarise the various constraints that a medium can impose on communication between people. These include whether the media allow co-presence, co-visibility, co-temporality, simultaneity and revisability of messages sent. They argue that the presence or absence of these factors affect the mutual achievement of common grounding in different ways. Whereas in face-to-face situations all of the above are possible, only a restricted number of these are available with video-conferencing, telephone, email and letters. Consequently, Clark and Brennan argue that when a medium lacks one or more of these characteristics it can require the participants to use alternative grounding mechanisms, which have different costs associated with them. For example, when contextual cues are missing in communication (e.g. in email conversations) the costs can be higher; misunderstandings can arise requiring the use of repair mechanisms to maintain shared understanding.

In our study we were interested in the nature of any additional costs incurred when shared understanding is mediated through a computer-supported collaborative tool and an adjoining telephone compared with face-to-face settings utilising shared paper-based resources (i.e. notes, books). To characterise and analyse the different costs, we used notions from distributed cognition (see Rogers and Ellis (1994) for an introduction) and ideas from Heidegger. Primarily, these comprised of characterising the interactions and conversations between the participants in terms of the coordination of social (between people), cognitive (within individuals) and technological (the use of artefacts) mediations, distributed across time. In addition we examined the kind of breakdowns (cf. "the interrupted moments of our habitual, standard being-in-the-world",

Winograd and Flores (1986, p.77)) that occured in the different settings.

Previously, the distributed cognition approach has been used to analyse distributed working in a number of different environments - mainly 'control room' situations such as ship navigation (Hutchins, 1990), aircraft piloting (Hutchins & Klausen, 1992), air traffic control (Halverson, 1994) and offices such as an engineering design firm (Rogers, 1992) and hospital administration (Rogers & Ellis, 1994). Similarly, the use of breakdown analysis has been used to analyse various CSCW activities (Sharples et al., 1993; Urquijo et al., 1993). Hence, our analytic tools have already been used in various CSCW contexts. Our intention here was to use this particular combination to understand better the different mechanisms used in shared understanding when supported by different media. In particular, we wanted to find out how shared understanding develops when writing collaboratively. When a document is being constructed, we might argue that there is not just the document itself being created, but also a set of individual understandings of the document (or 'representations' in distributed cognition terms), which combine together to form a collective understanding. Specifically, we wanted to address the question of how the transition from a paper-based to a computer-based workspace, and from face-to-face dialogue to that over the telephone and through the shared work space, affected the development of shared understanding in the collaborative writing process.

Evaluating the tool: PD and the playschool

We argued above for the integration of users and evaluation, and evaluation and re-design. Our PD-derived ETR methodology allows us to incorporate both the direct experiences of the users, and the redesign process itself within the evaluation framework, giving us a strong participants' perspective. What is important is not so much the details of the methodology used but that it allow both a direct role for participants voices during the evaluation, and that it considers design strengths and weaknesses in forming an evaluation. Such a method need not be PD-based, although we found it an ideal substrate in which both participants and system, design and evaluation could co-exist. Our particular approach - the Playschool session - is described below.

The Playschool session was inspired by various brainstorming and prototyping methods developed for use during the design process: the Critique and Fantasy phases of Future Workshops (Kensing & Madsen, 1991); the design materials and video diary of PICTIVE (Muller, 1993); and the throwaway rapid-prototyping common in Scandinavian PD (Greenbaum & Kyng, 1991). Pictive, or Plastic Interface for Collaborative Technology Initiatives, combines "deliberately low-tech design components with high-tech video recording facilities" (Muller, p.211). Muller emphasises the importance of 'plasticity': the cheap plastic design materials, encouraging experimental design; the malleability of both materials and designs, to be re-shaped and re-designed; and the artificiality of the resulting prototype, emphasising that it is an ideal, fantastic solution, unconstrained (at this stage) by system limitations. The design resources are of three kinds: everyday throwaway office materials; pre-formed interface elements such as 'icons' and 'windows'; and tools, such as scissors and glue, to shape the other materials. Like Future Workshops, in which facilitators participate alongside users Kensing and Madsen (1991), Pictive uses a 'partnership model' where both developers and users shape the system. A video diary is kept of a Pictive session, recording both evolving designs, and accompanying rationale.

We wanted our ETR session to facilitate the expression of user problems through direct exploration of solutions. We called it a 'Playschool' to capture the spirit of disinhibition and playful exploration we aimed for: the name refers to an enduringly popular British childrens' television programme, which promotes learning through imaginative play. Participants were given a variety of 'cheap and cheerful' design materials: cardboard; plain, coloured and scrap paper; felt

pens; crayons; pencils; highlighters; ruler; scissors; glue; eraser; sellotape; stapler; blue tack; overhead transparencies and coloured pens; assorted coloured sticky labels; and Post-It notes. (We thought the pre-formed materials used in Pictive too restricting, and used only 'free-form' resources). These materials were available on a shared workspace table, with the participants seated around it. One evaluator acted as facilitator during the session, mainly to encourage the participants to use the materials, and to answer specific queries about system capabilities or CSCW research. The other operated the video camera, and logged interesting events. The final set-up and atmosphere was designed to enable effective exploration of usability problems through accessible and non-intimidating design materials.

There are interesting methodological issues here which merit further investigation. Should evaluators participate, or facilitate, and to what extent? How can the video diary be exploited to enhance the evaluation and redesign process? To what extent is such a Playschool dependent on the personalities of the group involved? For successful and integrated evaluation of design we believe that users should be allowed to express for themselves, both verbally and visually, *what they* found problematic; *why they* think that is; and *how they* think such problems could be solved. Undoubtedly there are many methods and techniques which can support this process, of which our Playschool session is only one example.

The study

A three-part study was conducted, using the same three participants in each part: they are referred to here as Charles, Simon and Nigel (not their real names). All three participants were students on the same Masters course, taking a class in Artificial Neural Networks.

In the first part, the participants collaboratively created a document in a face-to-face, paper-based setting, similar to Plowman's study (1992). The aim was to investigate the way in which people carry out a collaborative writing and editing task in a face-to-face environment. The task assigned to the participants was to 'create and edit a definitive/revision answer' to a sample neural nets examination question (decided jointly by the participants and evaluators in advance). An audio tape recording was taken of the task; the evaluators sat watching without participating, and logged interesting events.

In the second part, which took place a week later, the participants used a computer-based shared editing tool, ShrEdit, in three separate locations to carry out a similar collaborative writing task: the preparation of another model exam answer on a different neural nets topic. The aim here was to evaluate ShrEdit for usability and efficacy in completing synchronous shared editing tasks with distributed participants, and to see how the collaborative activity was affected by the change of medium. The participants were able to communicate by a telephone conference call, which was taped. An evaluator sat with each of two of the participants, taking notes on their actions. As there were only two evaluators, the third participant was video-taped. The audio-tape was subsequently transcribed, using standard conversation analysis notation as developed by Jefferson (1984), and a log taken from the video-tape.

In the third part, the participants took part in a rapid-prototyping 'Playschool' session, to articulate problems encountered during their use of ShrEdit by exploring possible design solutions. A set of low-tech design materials were provided. One evaluator acted as a facilitator for this session, to encourage and advise as required; the other took a log of the session. A video-tape was also taken, and subsequently logged.

4 Findings: understanding the activity

In this section we analyse the development of shared understanding in collaborative writing: in particular, how this varies when different supporting mediums (paper and computer) are used. Our analysis is based upon the transcripts and evaluators' logs from Parts 1 and 2 of the study: collaborative writing tasks in a paper-based and computer-based setting respectively.

Shared understanding in the face-to-face session

There were three main mechanisms for mediating shared understanding in the face-to-face, paper-based session: talk, writing (on shared pieces of paper) and consultation of external artefacts - books and lecture notes. Of these, the most significant was talk: the participants spent most of the time discussing their subject-matter (neural networks). This talk usually took the form of one participant raising a concern about how something worked, or explaining their understanding of something: the other participants would then agree or disagree (frequently at length) with this understanding. Through the resolution of these continual mini-disputes a new shared understanding on a particular topic was developed. The participants' experience of this, from their comments after the session, was that it was less to do with forming a perfect exam answer than with the group discussion. For example, at one point Charles says that "we've gone through a learning patch", whilst Nigel says that "the main problem was that we didn't get the writing done, we didn't get the job done"(Part 1 comments).

These feelings of not achieving their goal (given that their task was to prepare a model exam answer) do not concur with the extent to which writing formed a part of their work during this session, and shaped their shared understanding. Throughout the session, the participants wrote extensive notes on the subjects under discussion, and were constantly aware of the need to produce an answer. For example, at one point Charles says "Don't like the way this is going, there's enough to fill a page and a bit here, she [the course lecturer] wants two pages" (Part 1 transcript). Then, about two minutes later, Simon explicitly suggests: "I think we should start writing" (Part 1 transcript). By the end of the session, they had notes on most of the essay, and a plan for the rest, yet their perception was of not having completed their specified writing task.

The bulk of the writing took the form of note-taking by one participant whilst (or after) the others spoke. This 'scribe' role, identified by Plowman (1992) in a similar study of co-writing, was perceived overall as having considerably less power than the scribe in Plowman's study. However, the scribe did occasionally read back from their notes, feeding back into the discussion. Also of interest was the extensive use of paper-based shared artefacts for drawing diagrams and mathematical equations (useful when discussing technical details): this was not possible during the computer-based ShrEdit session. Overall, there was more extensive consultation of external artefacts compared with the ShrEdit session, partly because it was possible (the participants were co-located), and perhaps also due to the greater amount of 'shared remembering' going on in that session. The participants' knowledge of what others were consulting allowed them to be much more discerning in their use of sources: when there was a conflict between the lecture notes and a book over a formula, Nigel commented "I'm much more prepared to believe she's got the notes wrong" (Part 1 transcript). These artefacts were freely available to be co-consulted, which was not the case when using ShrEdit.

Shared understanding in the computer-supported session

The three basic means of mediating shared understanding in the computer-supported session (where the participants worked in separate locations) were also talk, writing, and external artefacts. However, they took different forms and characteristics. Talk took place over a telephone link; writing took place via ShrEdit; and while the same external artefacts were used as before

(books and lecture notes) the only knowledge other participants had of their consultation was by explicit spoken or written references.

Much of the shared understanding that occured took place via social coordination over the shared telephone link. Compared with the face-to-face session, where talking was continuous, there were long periods of silence (up to 90 seconds) while all the participants were focussing on using ShrEdit. An example of shared understanding occuring through telephone mediation is as follows: Simon raises a question about how the current topic-of-composition works. His understanding is confirmed by Nigel, and then strengthened by Charles, leading to a stronger understanding for the group as a whole:

```
S: I'm just - just a general thing about the back propagation itself,
   right? Em, do you - you compute the errors on the outputs, right?
N: Yeah
S: And then you compute the errors on the input units. Oh no, it's all
   right, it's all right
C: What you do is you take - what you do is you calculate the error from
   the, from the output erm for the target and then
S: Yeah, I've got that
C: And then from that, you send that back
S: The error?
C: The error from the target
S: Right
C: Back into the hidden, using a proportion of eh using a proportion that
   the hidden units actually contributed to the output unit
S: Right, I've got that, yeah
```

Such talk-based interactions are relatively rare. More frequent were those where talk-based mediation of shared understanding interlaces with references to either the document under development or to external artefacts. The following extract is an interesting example of this, where talk is initially used as the sole mediator; a conflict develops which is not resolved until Simon appeals to an external artefact.

```
S: No no, you could teach them, I'm sure
C: No! That wasn't teaching. What it was doing was, you were setting up the
   state
N: No, I think you're getting confused with Hopfield nets there, Charles
C: Simon!
N: No, I think you are. Hopfield nets you set up the er
S: Yeah I've got
N: the weights initially and it's all a bit of a fix
S: I've actually - I've got a formula in front of me em. It says perceptron
   training: delta wj equals target minus output times the input.
C: OK do it
```

The only information the other participants have about this formula is Simon's bald statement "I've got a formula in front of me". Because they are in distributed locations the amount of additional information is much reduced - the kind of 'surreptitious monitoring' observed by Heath and Luff (1991) in the London Underground control room is not possible. Thus Simon must state explicitly that he *is* quoting. Even so, the other participants do not know what weight to attach to the quote: does the formula come from the seminal book on neural nets, or from the lecturer's course notes, or from some other source? The strength of Charles's remark makes it clear that he at least has some respect for what he takes to be the source, despite his previous support for an alternative position.

We also consider the interrelatedness - what Plowman (1995) has called the interfunctionality -

of talk and writing in the mediation of shared understanding: do they work together? This raises a further question: does writing on its own (via ShrEdit) mediate shared understanding in the way that talk alone does, or is it merely an adjunct to the understanding? We suggest that co-writing via ShrEdit acts as a technological mediator of shared understanding, and furthermore, that it acts as the *primary* mediator of shared understanding. Our evidence for the existence of such technologically-mediated shared understanding is not any explicit reference to it in the transcript, but rather the existence of breakdowns in understanding that the participants notice in their ShrEdit based co-writing. We describe these as repairs to 'shared mis-understandings' that progressively happen in the writing process. The following example shows such a mis-understanding repair:

```
C: We've got a problem with um if you look the line just above um point 1
   and what you've got (.) is halfway along, threshold step function er and
   that's wrong
N: Ri::ght. Tum tum tum tum - something
C: It should be ''relies on a easily differentiable''
S: Yeah, I mean [ it's not easily taken
C:              [ threshold function
C: Yeah - a step function isn't and it can't be used =
S: = Right, sorry
C: And I explain that later on in my part
```

Hence, technologically-mediated shared understanding exists in a tacit form in ShrEdit, and is primarily manifested in breakdown/repair situations. That is, one of the roles of social coordination is as a repair mechanism for the failures engendered by technological mediation. Thus shared understanding is partially embodied within the document. As one participant writes a section, they establish their individual understanding of it; as others read that, they *add* to the shared understanding by constructing their individual understandings. Shared understanding is more than just a collection of individual understandings: yet the default assumption is that everyone has *your* understanding, until proven otherwise. Breakdown occurs when one participant's understanding differs with what another has written, necessitating group discussion and the construction of a more meaningful shared understanding.

Summary of evaluators' analysis

The main differences in the development of shared understanding between the paper-based and computer supported ShrEdit sessions were one of degree: talking, writing and artefact consultation all occured frequently during both sessions, but there was more talking and consultation in the paper-based session whilst more writing in the computer-based session. The use of external artefacts also changed substantially between the two sessions, mainly due to the fact that the participants were no longer co-present and so could not co-refer to the same artefact. In sum, the participants were able to establish and maintain shared understanding when using ShrEdit. This was partially through using technologically mediated talk (by phone) and partially through the document being co-constructed. There was, however, a greater need to utilise repair mechanisms through the occurence of more misunderstandings and fewer opportunities for tacit shared understanding.

At the same time, the participants were highly productive, being able to co-construct a joint exam paper. Unlike in the paper-based session, the partipants were not as disappointed with their work; they all felt they had achieved their goal, although still felt the result was "less coherent than if one person had written the whole thing" (Nigel, Part 2 comments). This comment, and comments in the Playschool session, seem to indicate that the participants felt their shared understanding to be less complete than in the paper session, and they found this uncomfortable (this was perhaps also influenced by the number of mini-disputes created and resolved). However more

writing did go on in the ShrEdit session, and of a much higher quality, suggesting that at least technologically-mediated collaborative writing leads to better writing, if not necessarily a higher degree of shared understanding.

5 Findings: redesigning the tool

This section describes the findings from our evaluation of ShrEdit as supporting tool in Part 3 of the study, using the Playschool ETR method. The Playschool session produces two main sources of evaluative information: the designs produced by the participants, which express both perceived problems and proposed solutions; and the video diary of the session, which captures the participants' developing rationale for problems and redesigns. Note that the role of evaluators in a Playschool session is restricted to (fairly hands-off) facilitation during the session itself, and video-camera operation. Some illustrative fragments of discussion are quoted from the video diary of the session, and from the transcript of Part 2 of the study. However, it is the drawings and other design materials which really communicate the solutions devised and problems perceived, and it is obviously difficult to capture their substance in words. Figure 1 below is an annotated schematic representation of one of the designs produced by the participants during the Playschool session, which illustrated many of the findings we describe.

HCI studies of single-user design and evaluation are often discussed in terms of the design principles and usability attributes which define the desirable properties of a usable interface. These typically include: consistency; feedback; user control; user's model; clarifying metaphors (principles); learnability; memorability; error recovery; efficiency; and subjective satisfaction (attributes). Originally developed by Nielsen (1993) for use in Heuristic Evaluation, they can be used as a convenient framework for categorising findings from any evaluative method. It is illuminating to consider which of these are still relevant when transported to the CSCW context; which require modification; and which additional issues are unique to CSCW. We discuss our results in terms of a set of 'Usability Issues' derived from those above, but extended and regrouped to cope with the differing requirements of collaborative application. Our new list comprises six categories: Feedback and Awareness; Focus; Coordination, Ownership and Control; Communication; Mental Models and Metaphors; and Consistency. These categories are convenient, being both general and inclusive: they allow us to discuss practical problems, potential solutions and theoretical ideas in an integrated manner. In each subsection below we introduce the key issues in each category, and then describe and discuss those findings relevant to it.

Feedback and awareness

Straightforward feedback has always been an important issue in HCI: a user (of any machine) must know that their actions have an effect. The classic illustration of this is a cashpoint machine, which is noticeably more difficult to use when the auditory feedback (ie. a beep) after each keypress is removed. However, in CSCW you must also be aware that there are others 'out there', and receive feedback on *their* actions. The provision of such 'awareness' information becomes crucial: giving "an understanding of the activities of others, which provides a context for your own activity" (Dourish & Bellotti, 1992, p.107).

The lack of such awareness information was the single largest problem encountered using ShrEdit. The transcript is riddled with comments such as:

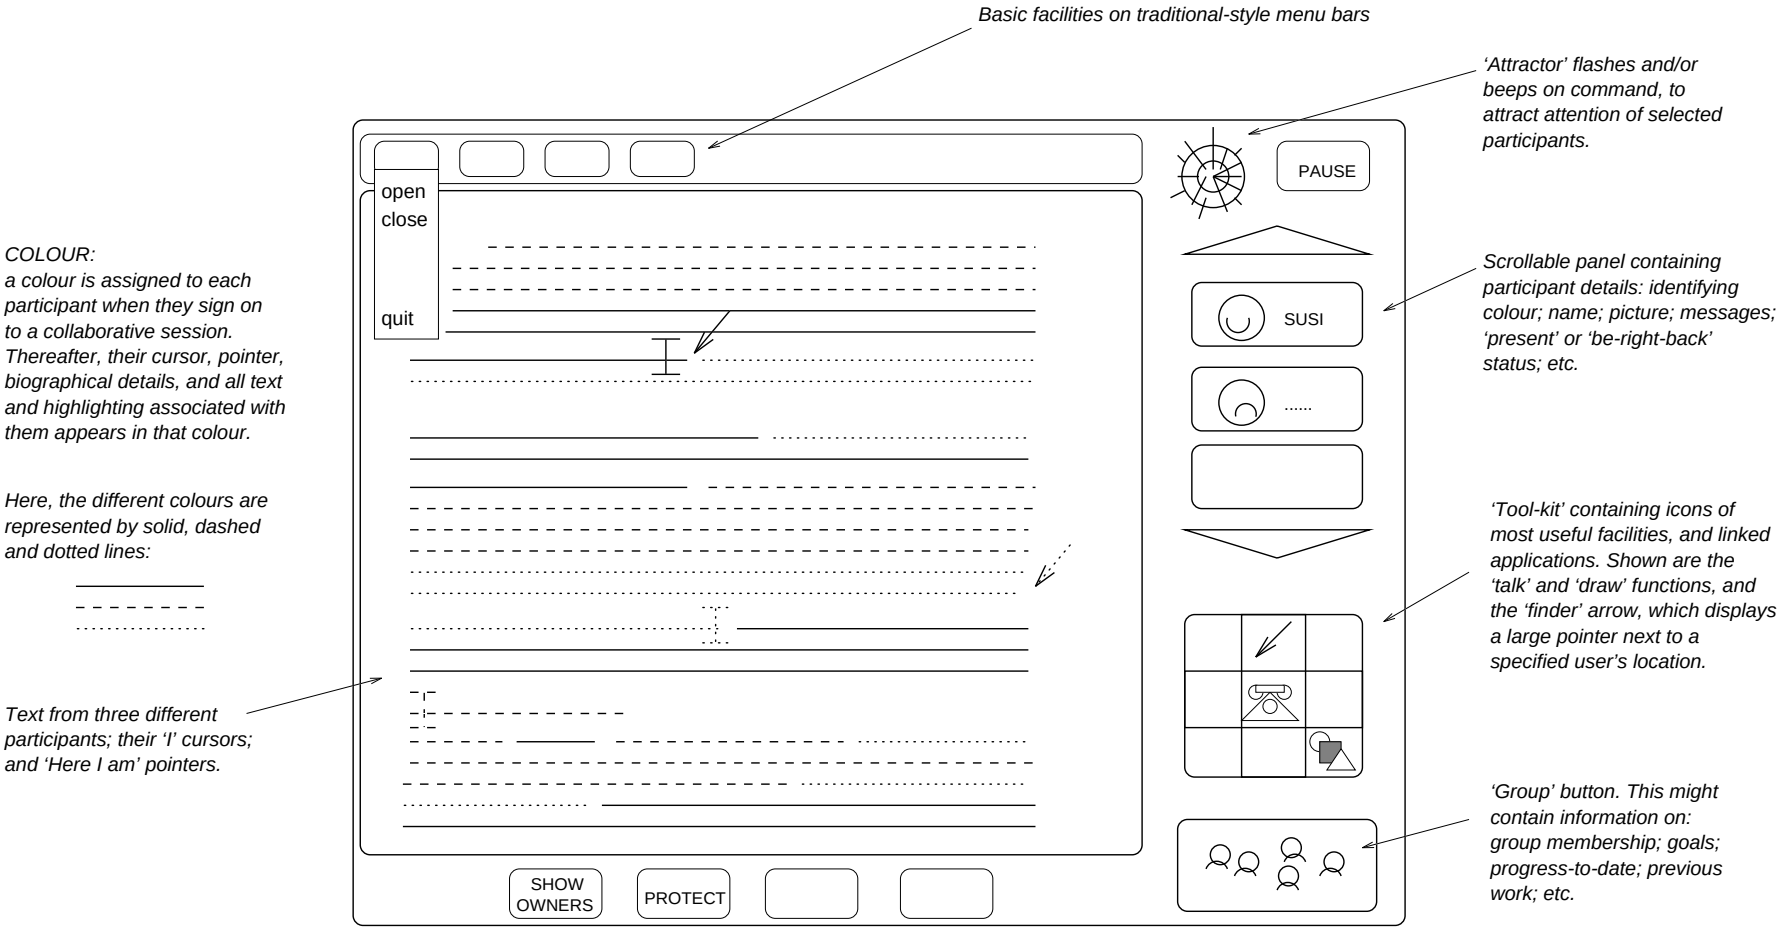

Figure 1: Annotated schematic representation of one of the ETR designs.

> 'Is someone doing something?' (Charles, 02.50);
> 'Who's typing? Charles?... Someone's typing' (Simon, 05.15);
> 'is somebody on - Nigel where are you? Are you on...' (Simon, 09.00);
> 'Frustrating! there's no clue as to where you actually are' (Charles, 09.55).

Solving this problem was the main issue tackled during the redesign session. Each editing conflict was complicated by the lack of information as to who was involved. Without a speech-based communication channel to sort out such problems, we believe it would have been impossible to use ShrEdit with any efficacy at all. Colour was the key resource proposed in dealing with the problem of providing awareness information for multiple users:

> C: COLOUR! Colour highlighting! Give everyone a colour when they sign onto a session, and then all they type will be in that colour. Cos then you can see who owns what piece of text, and you can see where people are. (Playschool log)

Highlighting and cursors, as well as text, would all vary in colour with the user. The need to see each person's cursor was considered crucial, as currently only your own cursor is visible, giving others' editing changes an uncanny disembodied feel! The participants proposed a combination of a 'normal' cursor (perhaps with face attached) and a 'big pointer', to attract attention (Stefik et al., 1987). Text currently being worked on would also be highlighted, to cut down on accidental edit conflicts. Locked text could also be designated as such in the appropriate colour.

To support awareness of who is currently using the system, the participants suggested a side-strip giving faces, names, email, and project details of each group member, which would also indicate their 'active', 'be-right-back', or 'missing' status. This strip would then be used to activate various tools: email, 'talk', or 'attract' tools would be dragged to it, and dropped onto the desired head. Finally, it is worth noting that the participants themselves discussed the need for user control of all these facilities - coloured text, flashes and sound could all be turned on or off at will.

Communication

This is not usually an issue in a single-user context. In contrast, people engaged in collaborative work (even face-to-face, paper-based work) need to communicate with each other during the process. Of interest, therefore, is what processes are involved, and how best to support them? We see two main kinds of communication: peripheral and explicit. Peripheral or background communication occurs as a side-effect of something else, yet also serves to enhance awareness of current events. Explicit communication is where one user enters discussion with any or all of the others. Both forms need supporting. The question is, should it be verbal or text-based; private or public; integrated or an add-on application?

In our study regular use was made of the telephone conference-call link, an add-on to ShrEdit's basic functionality. Much conversation time was spent discussing application problems, yet the rest of on-phone conversation was spent in discussion of the structure and content of the document, and in organising the group. The participants explored various communication-channel possibilities during the ETR session, but a built-in microphone was felt to be the best option. It should have an on-off capability, a way of attracting attention, and perhaps be mounted on a head-set. A text-based message window was thought to be too slow, and cumbersome if several other documents were already open. However an additional, private talk window was considered useful, to allow private collaboration between selected group members on an ad-hoc basis. The ability to turn both microphone and speakers on and off was felt important to ensure both privacy and user control.

Focus

In single user-computer interactions the computer simply flashes or beeps to attract the user's attention. But in a world of multiple users - all doing different things at the same time to the same document - how do you ensure that each user is aware of changes without overwhelming their eyes and ears with colour, motion and noise? How do you attract and direct the user's attention, or enable each group member to attract the notice of selected others? Moran and Anderson (1990) discuss the idea of *mundane* technology: technology that is ever-available in the background, but never impinges unduly on the user's consciousness. This becomes a much harder task when multiple active users are involved. This category is closely related to both feedback and awareness, and to communciation: in giving feedback the computer will need to draw the users attention to some outcome or other; and users will need to attract others attention before communication can take place.

The participants did have problems attracting each other's attention: ShrEdit has no in-built way of doing so, nor did they develop any social protocols to deal with the situation (such as typing a message on-screen). Simply providing an integral communication channel does not solve this problem, as a suitable means of attracting attention must also be found. The participants suggested that they should be able to direct the attention of other group members to a particular item of text, for comment or editing. Whilst one advocated flashing text, others thought this too invasive and annoying. A flashing alarm symbol with accompanying auditory feedback and message identifying the caller was finally agreed upon[2].

Coordination, ownership and control

In a single-user application, the user retains control, and is sole owner of their work. In CSCW, it is no longer possible to aim for (real or apparent) control by all users at all times: often another user will be carrying out a conflicting action at the same time. Who owns which pieces of text? Who has precedence? Who has permission to execute which action? Should these issues be enforced by the technology, or should the group itself work out social protocols to deal with them? Also, these issues are constrained as much by technological capability as by social preference.

ShrEdit uses a locking mechanism to protect a piece of text which is currently being edited. However, although the system informs users that a conflict has occurred, and displays the lock cursor, it does not say who the conflict is with (who temporarily owns that piece of text). Nor is there any clear indication as to what constitutes the locking range, for example, a message saying 'please move by a word, line, or paragraph'. When combined with the general lack of feedback and awareness information, the participants found this impossibly confusing and frustrating: however, their proposals were hampered by their lack of knowledge as to what was technologically determined and what could be redesigned. The problem of protecting finished work also interested the participants: both group work, and individual work which a user did not want changed. Should they have open access and honourable protocols, a strictly defined hierarchical editorship, or a locking mechanism which could be enabled on selected pieces of text? No final decision was reached, with debate ranging from total freedom to total control, as illustrated by the following snippets of conversation:

[2]This interface feature, unbeknown to the participants, is used in Aspects^{TM}, a commercially available shared editing tool.

```
S: Why do we have to move out of the way - isn't that assuming aggression?
C: What about negotiating changes by speech via comms channel?
S: Why not a free for all - use the colour to identify & ask permission, have
   an honour code?
C: I'd want protection, and have people ask if they wanted to change it... if
   i was happy with it i wouldn't want it changed.
N: But what about the shared editing ethos - we'd be as well doing it
   individually, then having one editor.        (Playschool evaluators' log)
```

Here, the participants showed concern for user control: all their proposed changes had to have an on-off facility, to prevent distraction and ensure privacy.

Mental models and metaphors

It is a well-known HCI design principle that we should "use the user's model" (Norman, 1986), preventing conflict between user assumptions and those the designer may unwittingly build into an application. Similarly, we should use metaphors which simplify, clarify and enable (Erickson, 1990), an unfamiliar or counter-intuitive metaphor serving only to confuse the user (cf. the infamy of the Macintosh^{TM} trashcan eject mechanism). How are these issues affected by the transfer to a collaborative context? What are our mental models of collaborative work? How does the average user make sense of self + task + computer + group + group task? Are the resulting models standard enough to be shared, or will each group member have a different understanding of the collaborative situation? Even assuming the formation of a shared group understanding (as discussed in Section 4), are there intuitive metaphors which support this?

The participants barely touched upon these issues. There is no standard accepted model of group working processes, and thus no pre-conceptions to transfer from a paper-based environment: nor did our participants have previous experience of a computer-based shared editor. We speculate however, that as more immediate (and immediately solvable) problems are sorted out, the issues surrounding users' models will become increasingly important, and the exploration of these issues will be important when considering both the experience and quality of our collaborative interactions.

Consistency

The HCI design principles recommend that consistency, of both function and metaphor, should ideally be preserved both within the application, and across applications (Grudin, 1989). Each function should perform the same way in a given context for each user of the application. When translated to a CSCW context, however, this is complicated due to the multiple synchronous users: a selected function may necessarily have different outcomes, dependent on the other users interactions with the document. Another recommendation - that consistency of 'look and feel' relative to other applications be preserved - is less relevant in CSCW, where there are (as yet) no archetypal CSCW applications to act as 'role models' Grudin. Finally, we predict that a certain duality of metaphor may be inevitable (and desirable) in a CSCW application, to support both the 'functional' and 'social' aspects of collaborative work. Whilst the standard office metaphor has proven effective in capturing the functionality of a text editor, it is unlikely to facilitate the feeling of shared purpose that a collaborative application may require. Other metaphors or visualisations, that convey the 'beingness' and 'awareness' of other users, may prove more relevant and useful.

Consistency was not seen as a major issue by the participants. Consistency of function within the application was subsumed by the problems of ownership and user control, and consistency of model and metaphor by the lack of preconceptions noted above. However, the mirroring

of typical Macintosh™ editor 'look and feel' did create expectations of habitual Macintosh™ functionality. This led to problems when the formatting action on a piece of highlighted text differed from that expected, changing the whole document instead. This lack of consistency between other editors and ShrEdit was deemed both misleading and frustrating, not to mention pointless: S: A whole document in Dingbats - ridiculous! (Playschool log). Again, we predict that this issue will only gain in importance once the basic application has matured somewhat.

Summary of participants' redesign

The participants were most concerned with issues of awareness, communication and ownership. At all times, it was pointed out by the participants that a user must know *who* is doing *what* and *where*, and how they are affected. Users must be able to communicate with others at will, selectively, and in a natural manner. They should have clear knowledge of who owns what at any particular time, and be able to retain ownership of their particular work if desired. Colour was the key resource proposed, as well as informative facilities providing information on the history and current status of each group member. All new features introduced should be user-controlled. Other issues, such as those concerning mental models, metaphors, and consistency, were considered less important: this is unsurprising given the multitude of more tangible problems.

6 Discussion and conclusions

The use of the PETRA framework enabled us to identify usability problems with the tool, and collect user recommendations as to how the interface could be improved, whilst also enabling us to analyse the different mechanisms involved in mediating shared understanding through different media.

When taken together, the findings from the evaluator's analysis (ETA) and the participants evaluation through design (ETR) can be seen as being complementary. For example, the importance of having various communicative mechanisms for shared understanding were outlined by the ETA, which were further instantiated in the participants concern with the need for feedback and awareness information, and communication channels. Specifically, the observation that the participants were required to spend considerable time talking to resolve breakdowns in their shared understanding and clarifying their written understandings when using ShrEdit (as suggested by the ETA) is directly related to the need for improved interface facilities (better feedback, awareness and control) established in the ETR session. In face-to-face situations we have already established mechanisms for coordination of the collaborative activity and shared understanding. Trying to support this kind of collaborative activity in a distributed computer-supported situation requires that the mechanisms of shared understanding be explicitly supported at the interface of the tool.

Our original concern - an investigation of the issues and difficulties involved in evaluating collaborative activities supported by CSCW systems - has led us to propose the need for a complementary framework, combining theoretical analyses and participatory design methods. We have investigated the collaborative activity of co-writing, and found that when using ShrEdit in distributed locations shared understanding develops technologically (via on-screen writing) with the use of social coordination (via telephone conversation) as a repair mechanism: this constrasts with the paper-based, face-to-face setting, where talk, not writing, has the primary role. We also discovered that the tool used to support this collaborative activity must be designed effectively to support awareness, communication, focus and ownership. In particular the tool must support awareness of not only the users' locations and actions at all times; but also an awareness of group membership and participation in the group task.

We also argued for the closer integration of use, design and evaluation. Participatory redesign, as illustrated by our Playschool, can usefully inform evaluation, and also ensure that the same users participate in both the evaluative and the iterative design stages of the development cycle. We emphasised the need for continuing to involve users as full and active participants in the evaluation process. This is a natural development of PD, maximising the richness of user opinion and experience available to the evaluators.

In conclusion we wish to emphasise the necessity of *multiplicity* in any useful, wholistic CSCW evalution: multiplicity of method, of theory, and of perspective. It needs to take theory into account, as that gives a grounding in similar or relevant experiences; but this theory should be drawn, where possible, from many disciplinary sources, to escape the theoretical blinkers that encourage us to screen out many insightful and useful ideas. Yet to use theory alone is to fall into the researcher-knows-best trap, to hold that the participants experiences with the system are less valid than outside observations of those experiences. Equal weight, therefore, must be given to the participants direct experiences, and their understandings of that experience, whether positive or negative. We propose the PETRA method: combining theoretical analysis and practical experience, evaluators' and participants' perspectives.

Acknowledgements

Our thanks to the three participants who enthusiastically and patiently donated their time; to James Goodlet for providing Macintosh™ expertise; and to Lydia Plowman and Mike Twidale for their comments on an earlier draft. Susi Ross and Magnus Ramage wish to thank Yvonne Rogers for her generosity with advice and encouragement over the course of this work. The authors were funded at various stages by the Engineering and Physical Sciences Research Council, the University of Sussex and the Digital Equipment Corporation.

References

Bannon, L. (1993). Use, design, and evaluation: Steps towards an integration. *Shaerding CSCW Workshop.*

Beck, E. (1994). Practices of collaboration in writing and their support. CSRP No. 340, School of Cognitive and Computing Sciences, University of Sussex. (DPhil thesis).

Clark, H. H., & Brennan, S. E. (1991). Grounding in communication. In Resnick, L. B., Levine, J. M., & Teasley, S. D. (Eds.), *Perspectives on Socially Shared Cognition*, pp. 127–149. American Psychological Association.

CSMIL (1992). *ShrEdit: A Shared Editor for the Macintosh.* Computing Science and Machine Intelligence Laboratory, University of Michigan.

Dourish, P., & Bellotti, V. (1992). Awareness and coordination in shared workspaces. In *Proceedings CSCW '92 Conference, Toronto.*

Erickson, T. (1990). Working with interface metaphors. In Laurel, B. (Ed.), *The Art of Human-Computer Interface Design.* Addison-Wesley: Reading, Massachusetts.

Garfinkel, H. (1967). *Studies in Ethnomethodology.* New Jersey: Prentice Hall.

Greenbaum, J., & Kyng, M. (1991). *Design at Work: Cooperative Design of Computer Systems.* Lawrence Erlbaum: New Jersey.

Grudin, J. (1988). Why CSCW applications fail: Problems in the design and evaluation of organisational interfaces.. In *Proceedings 1988 Conference on CSCW* Toronto.

Grudin, J. (1989). The case against user interface consistency. *Communications of the ACM, 32 (10)*, 1164 – 1173.

Halverson, C. (1994). Distributed cognition as a theoretical framework for HCI: Don't throw the baby out with the bath water - the importance of the cursor in air traffic control. Technical report No. 94-03, Department of Cognitive Science, University of California, San Diego.

Harper, R., & Carter, K. (1994). Keeping people apart. *Journal of Computer Supported Cooperative Work, 2 (3)*, 199–207.

Heath, C., & Luff, P. (1991). Collaborative activity and technological design: Task coordination in london underground control rooms. In *Proceedings ECSCW '91, Amsterdam.*

Hughes, J., King, V., Rodden, T., & Andersen, H. (1994). Moving out from the control room: Ethnography in system design. In *Proceedings CSCW '94, North Carolina*, pp. 429–439.

Hutchins, E. (1990). The technology of team navigation. In Galegher, J., Kraut, R., & Egido, C. (Eds.), *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work.* Lawrence Erlbaum: New Jersey.

Hutchins, E., & Klausen, T. (1992). Distributed cognition in an airline cockpit. In Middleton, D., & Engestrom, Y. (Eds.), *Communication and Cognition at Work.* Cambridge: Cambridge University Press.

Jefferson, G. (1984). Caricature versus detail: On capturing the particulars of pronunciation in transcripts of converstational data. Tilburg papers on language and literature No. 31, University of Tilburg, the Netherlands.

Kensing, F., & Madsen, K. (1991). Generating Visions: Future workshops and metaphorical design. In Greenbaum, J., & Kyng, M. (Eds.), *Design at Work: Cooperative Design of Computer Systems.* New Jersey: Lawrence Erlbaum Associates.

Krauss, R. M., & Fussell, S. R. (1991). Constructing shared communicative environments. In Resnick, L. B., Levine, J. M., & Teasley, S. D. (Eds.), *Perspectives on Socially Shared Cognition*, pp. 172–200. American Psychologial Association.

Kuuti, K., & Arvonen, T. (1992). Identifying potential CSCW applications by means of activity theory concepts: A case example. In *Proceedings CSCW '92, Toronto.*

Moran, T. P., & Anderson, R. J. (1990). The workaday world as a paradigm for CSCW design. In *Proceedings CSCW '90 Conference, Los Angeles.*

Muller, M. (1993). PICTIVE: Democratizing the dynamics of the design session. In Schuler, D., & Namioka, A. (Eds.), *Participatory Design: Principles and Practices.* New Jersey: Lawrence Erlbaum Associates.

Mumford, E. (1994). Born again consultants?. British HCI group newsletter No.25, British HCI Group.

Nielsen, J. (1993). *Usability Engineering.* London: Academic Press Inc.

Norman, D. (1986). Cognitive engineering. In Norman, D., & Draper, S. (Eds.), *User-centered Systems Design: New Perspectives on HCI.* Lawrence Erlbaum: New Jersey.

Orlikowski, W. J. (1992). Learning from notes: Organisational issues in groupware implementation. In *Proceedings CSCW '92*, pp. 362–369. ACM, NY.

Pagani, D. S., & Mackay, W. E. (1993). Bringing media spaces into the real world. In *Proceedings ECSCW93*, pp. 341–356. Kluwer, Dordrecht.

Plowman, L., Rogers, Y., & Ramage, M. (1995). What are workplace studies for?.. (submitted to ECSCW).

Plowman, L. (1992). Talking and writing in a group writing task: A sociocognitive perspective. CWRG report No. 5, School of Cognitive and Computing Sciences, University of Sussex.

Plowman, L. (1995). The interfunctionality of talk and text. CSRP No. 366, School of Cognitive and Computing Sciences, University of Sussex.

Rogers, Y., Bannon, L., & Button, G. (1994). Rethinking theoretical frameworks for HCI. *SIGCHI Bulletin, 26 (1)*, 28–30.

Rogers, Y. (1992). Ghosts in the network: Distributed troubleshooting in a shared working environment. In *Proceedings CSCW '92, Toronto*, pp. 346–355.

Rogers, Y., & Ellis, J. (1994). Distributed Cognition: An alternative framework for analysing and explaining collaborative working. *Journal of Information Technology, 9 (2)*. Forthcoming.

Rowley, D. E., & Rhodes, D. G. (1992). The cognitive jogthrough: a fast-paces user interface evaluation procedure. In *Proceedings CHI'92*, pp. 389–396. ACM, NY.

Schmidt, K., & Simone, K. (1993). Computational mechanisms of interaction for CSCW. Comic Deliverable 3.1, Esprit project 6225.

Schuler, D., & Namioka, A. (Eds.). (1993). *Participatory Design: Principles and Practices.* Lawrence Erlbaum Associates: Hillsdale.

Sharples, M., Goodlet, J., Beck, E., Wood, C., Easterbrook, S., & Plowman, L. (1993). Research issues in the study of computer supported collaborative writing. In Sharples, M. (Ed.), *Computer Supported Collaborative Writing.* London: Springer Verlag.

Sproull, L., & Keisler, S. (1991). *Connections: New Ways of Working in the Networked Organisation.* MIT Press: Cambridge, MA.

Stefik, M., Bobrow, D. G., Foster, G., Lanning, S., & Tatar, D. (1987). WYSIWYS revised: Early experiences with multiuser interfaces. *ACM Transactions on Office Information Systems, 5 (2)*, 147–167.

Suchman, L., & Trigg, R. (1991). Understanding practice: Video as a medium for reflection and design. In Greenbaum, J., & Kyng, M. (Eds.), *Design at Work: Cooperative Design of Computer Systems.* Lawrence Erlbaum: New Jersey.

Urquijo, S. P., Scrivener, S., & Palmen, H. (1993). The use of breakdown analysis in sychronous CSCW system design. In *Proceedings ECSCW '93*, pp. 281–294.

Winograd, T., & Flores, F. (1986). *Understanding Computers and Cognition.* New Jersey: Ablex.