Social analysis in the requirements engineering process: from ethnography to method

Stephen Viller & Ian Sommerville
Computing Department, Lancaster University, UK
{viller, is}@comp.lancs.ac.uk

Abstract

Over a number of years, we have been involved in investigations into using workplace observation to inform requirements for complex systems. This paper discusses how our work has evolved from ethnography with prototyping, through presentation of ethnographic fieldwork, to developing a method for social analysis that has been derived from our experience of applying ethnographic techniques. We discuss the strengths and weaknesses of each of these approaches with a particular focus on our most recent work in developing the Coherence method. This method is based on a fusion of viewpoint-oriented and ethnographic approaches to requirements engineering and uses an industry-standard notation (UML) to represent knowledge of work. We use a common example of an air traffic control system to illustrate each approach.

1. Introduction

This decade has seen ethnography become increasingly popular in the requirements engineering (RE) community. Studies have been performed in a variety of domains, including: underground control rooms [8], air traffic control [1], and banking [3]. By spending time alongside workers, observing what they do, ethnographers develop a deep understanding of the work. Ethnographers can therefore provide designers with detailed insights into the work *as it is actually performed*, presented in the language and terminology of the users. Ethnographic studies can uncover subtle features of the social nature of work that are vital to successful operations, yet at the same time appear to be so trivial that other techniques can miss them. This is in contrast to some other human-centred approaches, which tend towards studying simulations of work in artificial laboratory-based settings, and imposing their own vocabulary to describe the work.

Ethnography has a great deal to offer as a technique for RE, but a number of issues limit its use in practice:

- *Time*. Ethnography can be a very lengthy process, lasting months or even years in the context of social research. Much shorter time-scales are required for the RE process.
- *Results*. Ethnography tends to produce a great deal of detailed, textual description as a result of performing a study. Effectively communicating these findings to requirements engineers is not straightforward.
- *Culture*. There are significant differences in language and culture between sociologists and software developers. This can lead to problems of communication between the two groups.
- *Abstraction*. It is difficult to draw design principles and other abstract lessons from a technique that is concerned with the detail of a particular situation.
- *Skill*. The lack of a systematic approach to conducting ethnography makes the technique dependent on the individual ethnographer's skill.

We have been addressing these issues in our research for several years. We are not alone in wanting to improve how RE deals with social and organisational factors. Our experience with ethnography is analogous to many others investigating the boundaries between the social and the technical. These include other ethnographically informed approaches [2, 4, 15], Participatory Design [17], and sociotechnical methods [5, 7].

Rather than engaging in a broader theoretical debate, this paper wishes to convey our practical experiences. The following sections describe our experience with three different approaches to ethnographically informed design. They present not only a chronological view of how the work has moved on, but also a progression from the locus of ethnography being in the ethnographers' heads, towards being embodied in the method itself [4].

2. Ethnographers working with designers

Early work at Lancaster on ethnographically informed design took place on a project concerned with developing displays for air traffic controllers [1]. This was the first project at Lancaster to involve sociologists and software engineers collaborating together in the design process. Reports from both sociologists [12], and software engineers [18] reflected their respective concerns about the nature of this collaboration.

The model of work involved periods of fieldwork by the ethnographers, in parallel with prototype development by the software engineers, followed by de-briefing meetings where their respective findings were reported back to the rest of the team. This model was subsequently referred to as *concurrent ethnography* [9] (Figure 1).

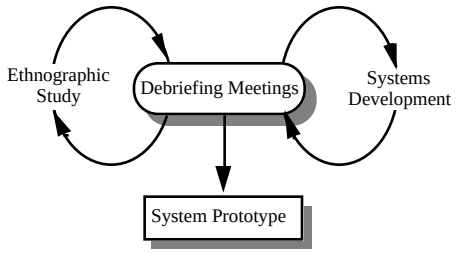

Figure 1: Concurrent ethnography

During the periods of fieldwork, the ethnographers observed controllers at work, making audio recordings of the conversations, supplemented with notes. The tapes were transcribed, and the transcriptions augmented with further details from the notes. These notes were used in the debriefing meetings to illustrate points that the ethnographer wished to make and in support of comments about the latest version of the prototype or design suggestions proposed by the software engineers. Figure 2 illustrates this pattern of working.

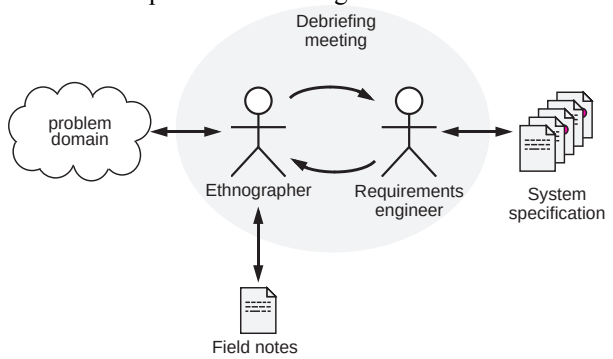

Figure 2: Ethnographers collaborating with software engineers

Fieldnotes take the form of a diary or journal, with short descriptions of the interaction taking place between controllers and with pilots. These are supplemented with comments on the nature of the work being observed, reports from other activities (e.g., a report on a session using an ATC simulator), transcriptions of interviews with controllers, and so on.

Early on, the notes contain detailed descriptions of the various activities engaged in by the controllers. As the observations continue, however, the descriptions gradually adopt the terminology and abbreviations used by the controllers, without further explanation. For example, the following excerpt from the first visit describes the *coordination* of a flight, which takes place whenever a flight is not going to pass between sectors according to a 'standing agreement'.

> I watched as the controller began to write '↓260L' in red on a strip, whist at the same time instructing a plane "descend flight level 260 to be level at the Isle of Man …" This strip was towards the bottom of one of three vertical positions, each separated by yellow plastic strips with beacon information printed on them eg IOM, along with some heading and frequency information…

Later on in the study, the description in the notes becomes much closer to a transcript of the talk alone, as in this excerpt.

> ...
> Controller: 'Speedbird 799L … delays running at 13-15 minutes at the moment … I'll try to keep you advised.'
> Pilot: 'Speedbird 799L … thank you.'
> Controller: 'Speedbird 799L … descend flight level 120'
> Pilot: 'Speedbird 799L … roger, descend flight level 120'
> Writes it on strip
> (takes place over about a three minute period)

In this second excerpt, there is more taking place, but the ethnographer's familiarity with the domain has removed the need for detailed description.

In the debriefing meetings, however, it was not the fieldnotes that communicated information to the software engineers. Rather, it was the ethnographer who acted as a 'proxy' for the field site. Their familiarity with the setting enabled them to respond to questions and suggestions on the controllers' behalf.

2.1 Strengths

Involving ethnographers in a prototyping cycle has a number of benefits. The ethnographer can bring insights about the domain to the debriefing meetings, that other approaches would struggle to obtain. The strength of these insights is that they are what *actually happens*. Observation can also uncover workplace features that are not obvious to the participants being observed, because they are so deeply embedded in what they do on a day to day basis.

Ethnographers' understanding of the workplace can be used when evaluating prototypes, where they can act as a 'proxy' for the actors in the domain. Disruption can therefore be minimised for 'real users'.

2.2 Weaknesses

This early work highlighted the issues identified in section 1. A good ethnographer can provide useful, and sometimes counterintuitive (for software engineers, at least) insights into the effects of introducing technology into the workplace. However, the wealth of information that is contained within an ethnographer's notes remains largely untapped, because of their detailed, unstructured, and highly personal nature.

2.3 Outcomes

A number of specific findings from the ATC project changed our approach to the development of interactive systems. For example, 'common sense' approaches to automation in software engineering could turn out to be detrimental for what appear to be routine manual tasks, which actually incorporate error checking procedures or that allow system users to build and maintain a mental domain model. There is also a tension between allowing users to tailor their interfaces while other users need to be able to understand the displayed information 'at a glance'.

3. Modifying ethnography

Following on from the ATC project, a number of collaborative projects coordinated at Lancaster extended our experience with using ethnography in design. Most of the work was concerned with how ethnographic techniques could be modified to orient the results towards the needs of RE. The modifications proposed were both in terms of the *process* by which ethnography is integrated into RE, and also the form in which the results of ethnographic studies are *communicated* to designers.

3.1 Moving out from the control room

In a paper that recognised, and attempted to address, ethnography's limited contribution to RE in general [9], a number of different processes for integrating ethnography into design were examined. These included:

- **Quick and dirty ethnography**. This directly addressed the criticisms of the time scale involved with ethnography. Several short, focused studies take place, interspersed with debriefing meetings, similar to those in concurrent ethnography. The shorter time scale for this approach is justified because diminishing returns set in quite early with ethnography in RE. Quick and dirty ethnography allows the effort of the ethnographer to be redirected towards new points of interest as they arise in the ongoing fieldwork.
- **Evaluative ethnography**. This is a focused version of quick and dirty ethnography. Where it differs is that it is intended for evaluating an existing design, or in other words performing a 'sanity check' on an already formulated design.

The above are concerned with the process of ethnographers and designers working together, which address some of the issues presented in section 1. For example, quick and dirty ethnography greatly reduces the lengthy duration of study. Evaluative ethnography opens up new possibilities for using ethnography in RE. However, these alternative processes still follow the basic model of ethnographers understanding and describing the problem domain, and designers making use of the information provided by them.

3.2 Presenting ethnography in RE

Of all the issues identified in section 1, the most difficult to address are those concerning how the results of ethnography are communicated to designers. Field notes are detailed, discursive documents, with little structure, and are oriented towards understanding and describing the social nature of a workplace. Further, the backgrounds of software engineers and sociologists make communication between the two disciplines open to misunderstanding. The approach taken at Lancaster was to examine how the results of ethnography can be modified to be more appropriate for RE. In particular, this took the form of *presentation viewpoints* [10] and, later, a *presentation framework* [11].

Presentation viewpoints. The first attempt to modify how ethnographic analyses are presented used a number of viewpoints, each one addressing a particular aspect of the social organisation of the work [10]. The viewpoints focus how ethnographers present their understanding of the domain, as recorded in their field notes. This model of working differs from the previous situation (Figure 2) by the fact that communication between ethnographers and designers is mediated by the description that is structured and presented as viewpoints (Figure 3).

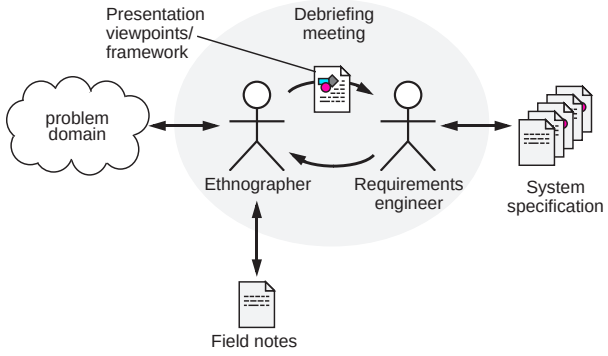

Figure 3: Modifying ethnographic records

Three viewpoints were identified in this work:

The ecology of work is concerned with representing the spatial distribution of the workplace. The participants, where they work, and the resources that they use are relevant to this viewpoint.

In the ATC study, the spatial layout of each suite, and its situation relative to other suites bears a direct relationship to the physical organisation of the airspace. ATC suites controlling neighbouring sectors of airspace are placed next to each other. This enables controllers to coordinate flights by talking with their neighbours, and allows them to get a picture of their future workload thanks to the 'at a glance' availability of paper flight strips in the racks of neighbouring suites.

Views of work is a collection of possible viewpoints from which the detailed ethnographic data can be presented. Conversations with controllers about the nature of their work appear here, alongside the ethnographer's observations and remarks.

The flow of work can also be considered as a collection of potential viewpoints focused on workflow, information flows, etc. In the ATC setting, this viewpoint includes the tracking of an aircraft through controlled airspace, and the related process of working with the paper flight strips.

The presentation viewpoints were supported in an augmented version of a tool previously developed at Lancaster [20]. This provided a means for cross-referencing the ethnographic material with abstract models of the system as they were developed.

Presentation framework. Following on from the work on presentation viewpoints, the Presentation Framework [11] structures ethnographic data in terms of three

dimensions of work; 'Distributed Coordination', 'Plans and Procedures', and 'Awareness of Work'.

Distributed coordination is concerned with how tasks are performed within the broader context of the organisation, as steps in continuing processes, and as part of a division of labour. Workers rapidly build up an understanding of what constitutes their work, and what is "somebody else's". Looking at work from this perspective is useful for understanding the roles played by different individuals as they collaborate together.

Returning to the ATC study, this dimension is concerned with how the paper flight strips, for example, which represent the progress of a flight through a sector, are used to coordinate the work of various actors. Features of the flight strip determine actions to be taken by the various actors, and colour coded annotations on the strip record who did what with it.

Plans and procedures focuses on how the organisation supports distributed coordination through job descriptions, workflow diagrams, instruction manuals and procedures, etc. Of particular interest here is the way that work in practice can differ from documented procedures.

An example of this from the ATC study was the practice of controllers deliberately placing flights on conflicting paths to solve an immediate problem before returning the flights to safe routes.

Awareness of work refers to how individuals perform their tasks so that what they are doing is made 'visible' or 'available' to others. Two people working alongside each other will have a good impression of what each other is doing, without being explicitly informed by their neighbour.

An example of this that has emerged from a number of control room based studies, including ATC, is gaze awareness. A controller will often know which flight another is talking about before they explicitly identify it, simply because one can see where the other is looking.

Cutting across these dimensions are a number of other aspects of work. These are used to further structure the presentation of features in the ethnographic record that are pertinent to the design in question.

3.3 Strengths

Of the work focusing upon the process of ethnography, the 'quick and dirty' approach directly addressed the time taken to perform ethnographic studies. Useful results can be obtained within a shorter time-scale, and debriefing meetings can be used to redirect the ethnographer's effort to other aspects of interest in the domain as they emerge.

The main benefit of the work on the presentation viewpoints and framework was the structuring of ethnographic data, and orienting it towards RE. Coupled with 'quick and dirty' ethnography, this makes the approach accessible to designers and procurers [3]. The tool support made the fieldwork notes more accessible to designers and enabled cross-referencing between design models and the fieldwork.

3.4 Weaknesses

Whilst cited as a strength, the tool used was non-standard, and therefore difficult to integrate with other tool support. None of the modified approaches to ethnography addressed the problems of communication between sociologists and software engineers.

3.5 Outcomes

The work that followed on from the ATC project has shown that a number of different approaches to using ethnography in design are possible. The 'quick and dirty' approach has been used extensively, and these studies have reinforced the appropriateness and utility of the categories used in the presentation framework. The presentation viewpoints served to demonstrate that viewpoints are a useful concept for structuring ethnographic fieldnotes and making them available to the design process.

4. Ethnographically informed method

The most recent work at Lancaster in this area has been conducted in the Coherence project. This represents a fusion of the research on the presentation framework with other work conducted at Lancaster on viewpoints for RE [19]. The difference between this and the work on presentation viewpoints is that Coherence can fit into an established framework for eliciting and reasoning about requirements from a number of perspectives.

In the Coherence project, we have been concerned with addressing the issues in section 1 from a different approach to previous work. Rather than modifying the process of ethnographic study or its outputs, we have developed an ethnographically informed approach to RE. Figure 4 presents how this approach differs from the previous models of work.

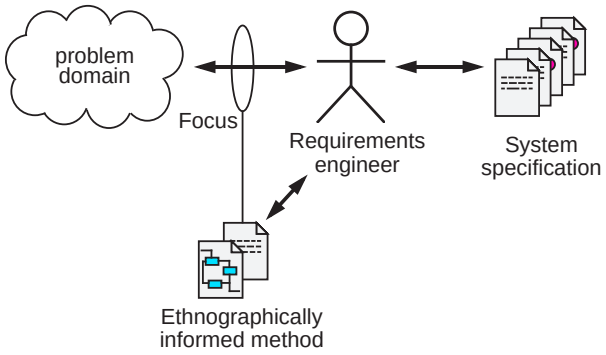

Figure 4: The Coherence method

Rather than requirements engineers relying on an ethnographer for analysis of the social nature of the workplace, they now perform the analysis themselves, supported by ethnographically-informed guidance. The Coherence method delivers this guidance in a systematic manner by using viewpoints to structure the requirements elicitation and analysis.

4.1 Viewpoint-oriented requirements

Coherence brings together two well-established threads of research at Lancaster University: the presentation framework for design described in section 3.2, and viewpoint oriented RE.

Viewpoints are encapsulations of information about a system or process, each one being a partial analysis of the workplace as seen from a particular perspective. For example, a pilot or controller would both be viewpoints in an analysis of an ATC system, as would be the *Manual of Air Traffic Control* which defines standard ATC procedures. Reconciling and integrating the separate viewpoints leads to a complete analysis of the system.

Viewpoints have been conceptualised and implemented in a variety of ways in RE [6, 14]. The particular viewpoint oriented approach we use is called PREview [19]. What distinguishes PREview is the use of *concerns* to drive the analysis. Concerns shift the perspective of PREview from what the system should do, to how it can best serve the organisation. They explicitly link organisational goals and objectives with system requirements. Concerns such as safety, compatibility, etc. are elaborated into *questions*, which must be asked of every viewpoint source to collect information about the system, and/or *external requirements*, which apply across all viewpoints to ensure compliance with the organisational concerns.

The structure of viewpoints, each of which focuses on a particular aspect of the workplace, and concerns, which cut across all of the viewpoints, is analogous to the structure of the dimensions and aspects of work adopted in the presentation framework. This is exploited in Coherence, where the framework is represented as social viewpoints and concerns.

4.2 Social viewpoints and concerns

Coherence structures the three dimensions of work in the presentation framework as viewpoints. To assist their elaboration in a given context, Coherence provides a number of questions to consider when elaborating each viewpoint (Table 1). These questions are not checklist items that *must* be responded to but are proposed as guidance for analysts, to sensitise them to the social features of the workplace. Elaboration of the social concerns is similarly supported by a number of questions for each concern, and these are presented in Table 2.

In determining the relevance of the concerns in a given context, each of the questions in Table 2 are re-cast to make questions of the form "*Does X exist?*" rather than "*How is X manifest in the workplace?*" For example, "Do forms and other artefacts on paper or screen act as embodiments of the process?" or "Are there aspects of the work to be supported that are time dependent?" If the answer is mainly no, then the concern can be eliminated from the analysis, thus reducing the amount of work required to reconcile each concern with every viewpoint.

Table 1: Social viewpoint questions

Distributed Coordination
• How is the division of labour manifest through the work of individuals and its coordination with others?
• How clear are the boundaries between one person's responsibilities and another's?
• What appreciation do people have of the work/tasks/roles of others?
• How is the work of individuals oriented towards others?
Plans and Procedures
• How do plans and procedures function in the workplace?
• Do they always work?
• How do they fail?
• What happens when they fail?
• How, and in what situations, are they circumvented?
Awareness of Work
• How does the spatial organisation of the workplace facilitate interaction between workers and with the objects they use?
• How do workers organise the space around them? Which artefacts that are kept 'to hand' are likely to be important to the achievement of everyday work?
• What are the notes and lists that the workers regularly refer to?
• What are the location(s) of objects, who uses them, how often?

Table 2: Social concern questions

Paperwork and computer work
• How do forms and other artefacts on paper or screen act as embodiments of the process?
• To what extent do the paper and computer work make it clear to others what stage people are at in their work?
• How flexible is the support for the work process—is a particular process enforced, or are alternatives permitted?
Skill and the use of local knowledge
• What are the everyday skills employed by individuals and teams in order to get the work done?
• How is local knowledge used and made available, e.g. through the use of personalised checklists, asking experts, etc.?
• To what extent have standard procedures been adapted to take local factors into account?
Spatial and temporal organisation
• How does the spatial organisation of the workplace reflect how the work is performed?
• Which aspects of the work to be supported are time-dependent?
• Does any data have a 'use-by-date'?
• How do workers make sure that they make use of the most up-to-date information?
Organisational memory
• How do people learn and remember how to perform their work?
• How well do formal records match the reality of how work is done?

Once irrelevant concerns have been eliminated, those remaining must be elaborated, with the help of the

questions in Table 2. Each concern record consists of the concern name, a description, and any external requirements and/or questions arising from the concern. Table 3 presents an example from the ATC study to illustrate a social concern that has been elaborated. In the table, the first cell contains the title of the concern, and a description of workplace features relevant to the concern, inspired by the pertinent questions in Table 2. The second cell includes a number of external requirements (ERs) arising from the concern, which must be considered by each viewpoint in the analysis. Questions arising from the concern, which would be asked of each viewpoint source, appear in the third cell (none in this instance).

Table 3: Example social concern for ATC

Paperwork and computer work The main feature of a control suite that this concern is interested in is the flight strip itself. As a consequence, the following concern questions responses all focus on how the flight strip is used by controllers in the course of their work. Flight strips embody the process of an aircraft's progress through the sector of airspace controlled by a suite. As an aircraft approaches the sector, its strip is moved progressively to the bottom of the rack until it becomes the current strip for the controller to deal with. The work of the controller can therefore be viewed in terms of dealing with the flow of strips as aircraft enter, traverse, and leave the controller's sector. The collection of strips in various racks in a suite provide an 'at a glance' means of determining the current and future workload of a particular controller. The practise of 'cocking out' strips—raising them slightly in the racks—informs the controller that there is something non-standard about the flight concerned. Glancing at the strips provides a controller with an indication of their current and future workload, in the same way as it allows other controllers to see the relative loading on other sectors. Flight strips provide incredibly flexible support for the work of controllers. Different practices exist regarding whether strips are placed into the racks in a top to bottom sequence or vice versa. All instructions given by controllers to pilots, and the pilots' acknowledgements, are recorded onto the relevant flight strip. These annotations are made using a standard set of symbols, and different coloured pens according to the annotator's role within the controlling team. In this way, flight strips record a flight's progress through a sector.
ER1. The system shall support controllers 'getting the picture' by providing the ability to determine current and future load for a sector 'at a glance' ER2. The system shall provide a facility to mark exceptional or non-standard flights requiring special attention ER3. Annotations to flight records shall be recorded and presented in such a way that they identify the person who made them.
No questions for this concern.

Returning to the ATC example, Table 4 illustrates the Awareness of work viewpoint. Here, it can be seen that the Organisational memory concern has been eliminated from the analysis as not relevant. There are also two organisational concerns, namely *Safety* and *Capacity*. These reflect two goals that the organisation has for the introduction of any system in that safety should not be compromised, and the capacity (i.e. the number of aircraft that can be handled in a given time scale) should be maximised. Requirements are generated by considering the pertinent questions for the viewpoint in Table 1, as well as through attention to the workplace from the viewpoint's perspective. They are noted in the viewpoint record, and elaborated elsewhere. In this case, the three requirements listed relate to:

- providing support for controllers to make their work available for scrutiny by others (AW1);
- providing information on controllers work so that it may be scrutinised (AW2); and
- attending to how the physical layout of the control suites maps onto the layout of the airspace (AW3).

Table 4: Example social viewpoint for ATC

Name:	Awareness of work
Focus:	How the physical organisation of the control suites affects how controllers can make sense of each other's activities. How controllers monitor the work of other controllers, and how controllers orient their work to facilitate others monitoring it.
Concerns:	Paperwork and computer work Skill & the use of local knowledge Spatial and temporal organisation ~~Organisational memory~~ Safety Capacity
Sources:	Controllers, and observation of controllers at work
Requirements	AW1 (Making work available) AW2 (Availability of awareness information) AW3 (Relationship of suite layout to controlled airspace)

4.3 Linking with system models

One of the major concerns we had when developing Coherence was that it should be of use to requirements engineers in practice. Therefore, we decided to look for an existing notation which could be used to document the social analysis. The notation we decided on is the Unified Modelling Language (UML) [16]. Our first task was to establish that UML is capable of expressing information about the social nature of workplaces [21]. We have subsequently been concerned with the process of undertaking social analysis in a systematic way, and providing links from our work to standard approaches to requirements analysis and systems development.

Coherence is particularly suited to use case driven requirements analysis [13], as there are a number of correspondences between the two approaches. For example, transcripts that result from observation of work such as that presented in section 2 can be modelled as sequence diagrams, which describe the scenarios which make up use cases (Figure 5)

When Coherence is used alongside PREview analysis, a number of additional domain viewpoints are identified and elaborated. An important class of viewpoint is known as *interactor* viewpoints, which correspond directly to *actors* in use case models. Use cases themselves are generated from the *plans and procedures* viewpoint, which focuses analysis on workflow. Being informed by ethnography, Coherence is concerned with how work is

actually performed, and as a consequence how it may differ from documented procedures. It can therefore be used to generate initial use case models, based upon observed interaction, such as that presented in Figure 6.

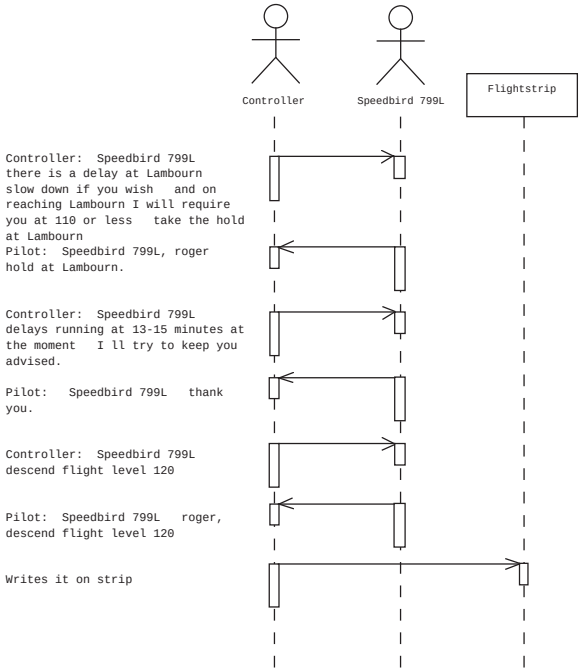

Figure 5: Observed interaction modelled in a sequence diagram

Thanks to the lessons learned from experience with 'quick and dirty' ethnography, the time scales involved in Coherence are similarly reduced in comparison with the initial work on the ATC project. Using a systematic approach to the analysis also means that requirements engineers who do not necessarily have any training in sociology can apply the lessons embodied in the presentation framework. Conversely, this work also enables ethnographers to use viewpoints and a standard notation to present the results of their studies to designers.

4.4 Strengths

The approach to ethnographically informed design in Coherence has a number of strengths. First and foremost, we are using an industry standard notation to describe the social features of a domain alongside the rest of the requirements. Our approach is flexible in that it can complement a viewpoint-oriented approach to RE, or be used to provide social analysis as a 'front end' to any other preferred approach. We have examined use case driven design as one approach in particular, and found that Coherence analysis can assist in the identification and description of use cases, scenarios, and actors and objects in the domain. A further benefit of adopting and integrating with standard approaches to analysis and design is that it will be easier to draw out general and abstract lessons about the social nature of work, and how to reuse them in future designs.

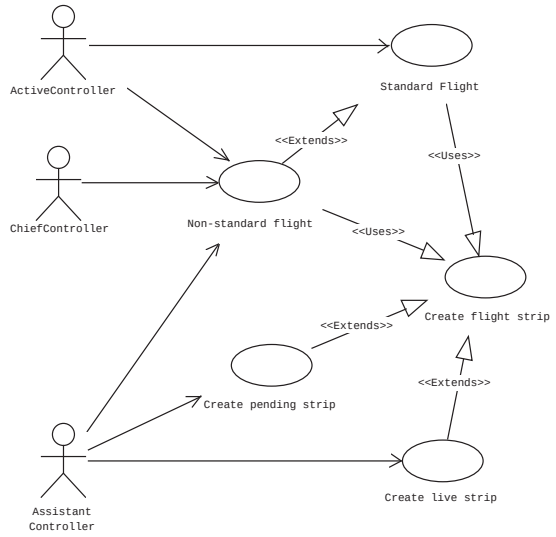

Figure 6: an initial use case model for the ATC study

4.5 Weaknesses

Whilst it could be claimed that Coherence has removed the communication problem between sociologists and software engineers, the effect has been to shift the problem from one of understanding people in conversation, to understanding people via the method's documentation. This places an onus on Coherence to communicate the intricacies of understanding social interaction to an audience with little or no background in the social sciences.

4.6 Outcomes

In the Coherence project, we have tried to demonstrate that ethnographically informed RE can make the transition into a systematic approach for social analysis. We have demonstrated that an industry standard notation can be used to represent the type of information that ethnographic studies reveal about the social nature of work [21]. We have also shown how our approach can be used in concert with other established approaches to RE to make an ethnographically informed understanding of a domain available to the rest of the RE process.

5. Conclusions and further work

This paper has presented a review of the work at Lancaster University on integrating ethnography into the requirements process, which has culminated in the ethnographically informed Coherence method of social analysis. Previous work in this field has made progress by modifying the process of ethnographic study, and how its results are presented. Our approach in Coherence has been from the other direction, aiming to produce a method that is informed by ethnography, rather than modify ethnography to suit the needs of design. The resulting

method has a number of attributes which facilitate the incorporation of social issues into system requirements.

- The method itself is informed by ethnography, providing an encapsulation of our experience of using ethnography in RE.
- We have used an existing viewpoint-oriented approach to structure analysis in Coherence, and to enable requirements arising from the social analysis to be reconciled against those from other sources.
- In contrast with ethnography, Coherence provides a systematic approach to social analysis which does not require a background in sociology to apply.
- Coherence uses a standard notation (UML) to communicate its results to the rest of the RE process, rather than burden requirements engineers with yet another notation to learn.
- Finally, we have provided links to use case and object models to enable Coherence to act as a 'front end' to existing approaches which do not currently address social issues in the workplace.

We believe that Coherence is a significant step forwards in terms of making ethnographically informed analysis acceptable in mainstream RE. Yet there are still areas requiring further attention. Users other than the authors have not yet used Coherence in a 'real' development context. Feedback from usage such as this will be invaluable for making the method more usable, and understanding where extra guidance is required for the social analysis. Tool support does not exist for Coherence, but is in development for the PREview approach which underlies it.

A working document is available from the Coherence project's world wide web pages at http://www.comp.lancs.ac.uk/computing/research/cseg/projects/coherence/ describing the method in more detail.

6. Acknowledgements

We wish to acknowledge the contribution to the ongoing work on ethnographically informed RE by various members of the sociology department, in particular John Hughes, Jon O'Brien, Dave Randall, and Mark Rouncefield. Thanks also to Tom Rodden and the anonymous reviewers for comments on this paper. The Coherence project is funded by the UK's Engineering and Physical Science Research Council.

7. References

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

2 Beyer, H. and Holtzblatt, K., *Contextual Design: Defining Customer-Centered Systems*. Morgan Kaufmann, San Francisco, CA, 1998.

3 Blythin, S., Rouncefield, M. and Hughes, J.A., Never mind the ethno stuff—what does all this mean and what do we do now?: Ethnography in the commercial world. *Interactions 4,* 3 (1997) 38-47.

4 Button, G. and Dourish, P., Technomethodology: paradoxes and possibilities. In *Proc. CHI'96* (Vancouver, 1996) ACM Press, 19-26.

5 Eason, K., *Information Technology and Organisational Change.* Taylor & Francis, London, 1988.

6 Finkelstein, A., Kramer, J., Nuseibeh, B., Finkelstein, L. and Goedicke, M., Viewpoints: A framework for integrating multiple perspectives in system development. *International Journal of Software Engineering and Knowledge Engineering 2,* 1 (1992) 31-58.

7 Harker, S.D.P., Eason, K.D. and Dobson, J.E., The change and evolution of requirements as a challenge to the practice of software engineering. In *Proc. RE'93* (San Diego, CA, 1993) IEEE Computer Society Press, 266-272.

8 Heath, C. and Luff, P., Collaboration and control: crisis management and multimedia technology in London Underground control rooms. *Computer Supported Cooperative Work 1,* 1 (1992) 69-94.

9 Hughes, J., King, V., Rodden, T. and Andersen, H., Moving out from the control room: ethnography in system design. In *Proc. CSCW'94* (Chapel Hill, 1994) ACM Press, 429-439.

10 Hughes, J., O'Brien, J., Rodden, T., Rouncefield, M. and Sommerville, I., Presenting ethnography in the requirements process. In *Proc. RE'95* (York, 1995) IEEE Computer Soc. Press, 27-34.

11 Hughes, J.A., O'Brien, J., Rodden, T. and Rouncefield, M., Designing with Ethnography: A Presentation Framework for Design. In *Proc. DIS'97* (Amsterdam, 1997) ACM Press, 147-159.

12 Hughes, J.A., Randall, D. and Shapiro, D., Faltering from Ethnography to Design. In *Proc. CSCW'92* (Toronto, 1992) ACM Press, 115-122.

13 Jacobson, I., Christerson, M., Jonsson, P. and Övergaard, G., *Object-Oriented Software Engineering: A Use Case Driven Approach.* Addison-Wesley, Reading, MA, 1992.

14 Leite, J.C.S.P., Viewpoint analysis: a case study. *ACM Software Engineering Notes 14,* 3 (1989) 111-119.

15 Potts, C. and Hsi, I., Abstraction and context in requirements engineering: toward a synthesis. *Annals of Software Engineering 9,* (1997) 1-39.

16 Rumbaugh, J., Jacobson, I. and Booch, G., *The Unified Modelling Language Reference Manual.* Addison-Wesley, Reading, MA, 1999.

17 Schuler, D. and Namioka, A., Eds., *Participatory Design: Principles and Practices*. Lawrence Erlbaum Associates, Hillsdale, NJ, 1993.

18 Sommerville, I., Rodden, T., Sawyer, P. and Bentley, R., Sociologists can be surprisingly useful in interactive systems design. In *Proc. HCI'92* (York, 1992) Cambridge University Press, 341-353.

19 Sommerville, I., Sawyer, P. and Viller, S., Viewpoints for requirements elicitation: a practical approach. In *Proc. ICRE'98* (Colorado Springs, 1998) IEEE Computer Soc. Press, 74-81.

20 Twidale, M., Rodden, T. and Sommerville, I., The Designers' Notepad: supporting and understanding cooperative design. In *Proc. ECSCW'93* (Milan, 1993) Kluwer, 93-108.

21 Viller, S. and Sommerville, I., Coherence: an approach to representing ethnographic analyses in systems design. *Human–Computer Interaction 14,* Special issue on representations in interactive systems development (1999)