9: Ethnography and the social structure of work

David Martin & Ian Sommerville

Lancaster University

1 Introduction

Achieving dependable systems design and implementation is now considered to be a process where attention needs to be paid not only to the technical system but also to the social and work environment into which the system will be placed. Dependability is seen as a property of the whole socio-technical system. Socio-technical systems comprise, holistically, computer based systems and the social systems of work of the people that work with, through and around those computer based systems. It is acknowledged that particular consideration is required to understand how well the technical system will fit with the activities of the users in the proposed setting (the application domain). For instance, highly dependable technical systems may be part of an undependable socio-technical system because they are inappropriate to the setting and users. This chapter discusses the relationship between the social structure of work and the technical structure of a system and the importance of understanding that relationship when designing dependable socio-technical systems.

Ever since the 'turn to the social' in systems design [12]., areas of computer science and systems design – most notably computer supported cooperative work (CSCW) and human-computer interaction (HCI) – have been increasingly interested in and have widely acknowledged that design may be enhanced by a better understanding of the social 'systems' of work into which computer systems will be imbedded. This is because researchers and professionals now understand that the functions and processes of any technical system need to mesh well with the work practices of

Joanne 27/6/05 11:57

Comment: could you please add a sentence to establish the relationship between the terms socio-technical and computer-based systems as you are using them.

personnel (or that the people or the system will be able to adapt such that they will mesh well) or problems will occur with the use of that technical system. This may range from staff producing workarounds to fit the system to their work – meaning the system is being used non-optimally – to acts of sabotage or the rejection of the system. Concurrently (but not just coincidently), this period has also seen the rise to prominence of ethnography (or observational field studies) as a key method for studying social systems of work – the interaction of personnel with each other, with computers and other artifacts, in their 'home' environment – the shop floor, the office, the control room, the home and so forth ([2]: [21]; [21]).

A major strand of ethnographic work within the systems design area draws on the program of ethnomethodology ([9] see [18]; [7] for studies in computing). Ethnomethodology eschews theorising and instead takes an approach to field studies whereby 'work' is analysed and explicated in the terms in which it is organised as a recognisable social accomplishment by the participants in that setting, rather than describing it in relation to extrinsically generated theoretical constructs. Ethnomethodology is interested in explicating the *social structure of action*, as it is produced in a setting. It takes the position that activities are structured from *within* in response to developing situational contingencies and as such this means that every given occasion of e.g. a telephone banking call will have its own unique structure. Different calls will be structurally similar, and differences in calls will be accountable in terms of how they emerge in the particularities of an unfolding situation.

The job of the researcher following an ethnomethodological program is to explicate what these methods and practices are and how they are deployed in unfolding action, to note regularities (and exceptions) in action and to delineate the circumstances that provoke them. For example, they will describe the methods by which control room workers coordinate their work around various computer systems and paper artifacts to achieve an ordered airspace in air traffic control ([14]); or, explicate the practices by which telephone operators in a bank call centre achieve smooth interaction with their computer system and customers ([3]). Such studies have a definite sociological interest – how is this work actually done? However, they have also been *"surprisingly useful"* ([30]) for systems design as many involved in that field have realised that these studies can fill in the 'just what actually goes on', 'just how is this actually done' - the *haecceities*, or 'just thisness', (cf. [19]; [10]) – of action and interaction in a situation that are missing from purely technical, data centric representations of work.

Data and object centric representations and abstractions are crucial for design but these omit social details of work that are equally important. Understanding work more holistically, as involving groupings of people interacting with, through and around both paper and computer-based technologies, provides a deeper, broader and more subtle understanding of work organisation that can, for example, enable better

Joanne 27/6/05 15:03
Comment: There are no 1995 references in your list

Ian Sommerville 29/7/05 15:32
Deleted: [13]

Joanne 28/6/05 9:18
Comment: Please check - as Heath is not the first name listed in the reference list I have had to make an assumption here

Ian Sommerville 29/7/05 15:32
Deleted: [7]

Ian Sommerville 29/7/05 15:32
Deleted: [18]

Joanne 27/6/05 11:59
Comment: I've made this plural to agree but you might want to change to singular (an unfolding situation). Let me know if so.

Joanne 27/6/05 16:44
Comment: There is no Bowers & Martin publication from 2000 in your references list.

Ian Sommerville 29/7/05 15:32
Deleted: [29]

Ian Sommerville 29/7/05 15:32
Deleted: [9]

Ian Sommerville 29/7/05 15:32
Deleted: [19]

modelling of processes, and more accurate and enhanced models of the artifacts and data used in that work [30]; [31].

In this chapter, we consider ethnography and structure from a number of perspectives. We start by discussing how ethnomethodological ethnography can be used in practice to reveal the social structure of work. We go on to argue that the collected ethnographic data can itself be usefully structured and we suggest a structure that helps organise this data in such a way that it can support systems design. This structure has been carefully designed to help reveal the social structure of work. We then examine this notion and its relationship with the structure of supporting computer-based systems, and discuss how organisational structure imposed by rules, plans and procedures influences and affects the social structure of work. Finally, we reflect on social structure and system dependability, making the point that the key requirement for socio-technical systems dependability is that there should not be a major mismatch between the structures assumed by that system and the practical social structure of work.

Ian Sommerville 29/7/05 15:32
Deleted: [33]

2 Ethnography

Ethnography is a method of data capture that works through the immersion of the researcher within the environment being studied, collecting detailed material (notes, documentation, recordings) on the 'real-time real-world' activities of those involved. Periods of immersion can range from intensive periods of a few days to weeks and months (more common in systems design studies), and even years. A primary product of most ethnographies is the development of a 'rich' description – a detailed narrative – of the work or activity in question, which may then be further *analysed* or *modelled* for various means, taking various approaches. The means may be for the purposes of answering sociological, psychological or systems design research questions, with the different approaches for analysis arising from various theoretical and methodological perspectives within these areas.

Ethnographers are interested in studying the work going on in settings rather than just computer systems in a narrow sense – they are interested in studying computer systems in operation, being used by people, within an organisational context and therefore shaped by organisational norms, rules, procedures, ethos, culture etc. In this conception we can think of, for example, a tax office as a complete organisational system – it comprises various technologies (e.g. computer and paper-based), organisational rules, processes (and methods for implementing them) and so forth, and staff who draw on their everyday and specialised social and vocational skills, abilities and knowledge to operate the technologies and work according to organisational requirements.

For ethnomethodological ethnographies in the computing literature, the control room (e.g. for air traffic, underground trains, ambulances) has been described as the multi-media field site *par excellence* with many studies focusing on these (e.g. [14]; [14]; [20];,[24] respectively). Control rooms involve small groups of workers, co-located within a setting, working in a coordinated manner on a set of tasks. There are two main reasons for their popularity. Firstly, they are readily amenable to ethnomethdological study, as they are self-contained and small scale. When taking an approach that focuses on the details of social action, scaling up, or investigating coordinated activity in distributed sites creates issues to do with general resource concerns and needing to be in more than one place at once to understand distributed coordination ([15]). Secondly, much ethnomethodological work has been oriented to concerns relevant to computer-support for cooperative work. Therefore, studies of settings which involve varieties of different technologies being used collaboratively in a number of ways have been seen as particularly important to provide knowledge about collaborative practices 'in the wild' and to inform the design of systems to promote and support collaboration.

When studies of control rooms are conducted they usually have the form of either a 'concurrent' or an 'evaluative' ethnography (cf. [15]) or as a mixture of the two. Concurrent ethnographies are conducted during a design process to inform the design of a new system in some way. They focus on describing current socio-technical system operation – how participants interact with, through and around current technical systems – and are aimed at highlighting important features of the work to support in any subsequent new computer system design. Findings are usually communicated in project meetings with designers who may then direct the researcher to focus on particular practices they are interested in understanding when they are making design decisions or dealing with design problems concerning current practice.

Evaluative ethnographies were first discussed as a means of deciding whether a proposed system fitted well with work patterns – the ethnography of current practices would be used to assess the potential fit of a new system with them. However, any ethnography of current practices may well reveal situations where workers have to work around technical systems, or recover their faulty operation, or make them more dependable through their social practices. Findings of this nature may simply provide an evaluation of a current system or may be translated into a resource in a redesign process.

Control room studies often 'draw a line around' the system at the interface between the control room and the outside world. Their interest is generally in how, given the inputs – radio calls, and various forms of information, displayed and visualised in various ways – do the workers manipulate and transform this information with the support of the technical system into the requisite outputs (e.g. 'correct' instructions to ambulance drivers). Later studies have extended this work to consider how the work of the control room relates to that going on outside for example by

Ian Sommerville 29/7/05 15:32
Deleted: [13]

Joanne 27/6/05 16:51
Comment: There are two items in the References list that would fit this description

considering how the work of ambulance crews relates to that going on in the room ([25]).

While control room settings offer the opportunity for a relatively comprehensive ethnomethodological study, it should be noted that studies like Pettersson and Rouchy's ([25]) point to the fact that boundaries always tend to get drawn somewhat *arbitrarily* around the subject matter. When we look again at the other control room studies we can see that they tend to focus on certain jobs, practices and technologies over others. For example, in ambulance studies, some focus on call taking, some on selecting and managing ambulances for dispatch to incidents, others on the coordination between the two groups of workers. Nevertheless, control rooms represent fairly straightforward settings with established topics and ways to carry out research.

Following on from the control room studies, ethnographies have generally been employed for design by looking at small scale settings and activities, generally by adopting similar approaches. However, in larger, more complex and distributed settings (like, for example, the headquarters of a bank, or a hospital) scaling up the study is a problem, and therefore the relation to design may be more partial, more abstract, and more complex to understand. Given that these field sites have many more participants working in them, carrying out many more activities, simultaneously in different places sometimes as part of a coordinated process, sometimes as different processes realised in parallel, the following problems often ensue for the fieldworker. Firstly, given that a comprehensive ethnography is often outside the scope of the project what should they focus on? Secondly, if they look for a subsection of people, activities, technologies and processes it may be harder to draw a boundary around them in the way that a control room forms a nice unit.[1] Thirdly, how easy is it to collect and present data in an accurate form about people collaborating synchronously in different places?[2]

Originally, it was suggested that 'quick and dirty' ethnographies (cf. [15]) would be useful in this situation. Quick and dirty ethnographies were conceived as being relatively short in duration (up to a few months) but their title actually refers more to the amount of understanding that could be gained from them in relation to the scope of the setting and the project. Quick and dirty ethnographies were used to study, in a relatively brief manner, a variety of areas of operation in a complex organisation. From this some general issues for design were extracted. Hughes and colleagues

---

[1] This problem is not insurmountable. Boundaries are always drawn, but it is important to indicate how, where and why they have been drawn, what has been included and what the limits of the research are.

[2] There are various ways in which distributed coordinated activity can be 'observed' and analysed. Often interactions on the telephone, or via computers (emails, instant messaging etc.) can be captured and analysed, and this material can be supplemented by direct observation of participants in these interactions.

([15]) also suggest that discussions over the findings of quick and dirty ethnographies amongst stakeholders may be used to select certain areas and activities for more detailed ethnographic inquiry.

To expand on this idea we can usefully think in terms of how to *time* and *target* ethnographic interventions. Ethnographic studies can be used to explicate the details of current practice during build and configuration in situations where design concerns are raised over the fit of a new system with current practice, and the potential ramifications of a new system being disruptive to current work practices. Through discussions with the design team, the ethnographer can be directed to study certain areas of work that utilise certain technical systems (or certain modules, functions etc. of the technical system) and concentrate on these. Or the focus may be on looking at various interactions between various systems.

Ethnomethodological studies for design may serve as a *resource* for evaluating current socio-technical system operation and as a *resource* for considering which aspects of the current system are important to support in any re-design. They can help identify areas where socio-technical systems are not very dependable – e.g. where people are making mistakes, where the system 'design' or outputs are causing problems for those using them, or where human ingenuity is making up for problematic technical systems. However, they will not tell you in themselves how to build a system, transform work or make something more dependable. It is in this way that the studies serve an informational input (rather than an automatic solution generator) to help think about possibilities in these situations and weigh them up against an in-depth understanding of the current situation.

3 Structuring ethnographic data

Ethnographic records are collected opportunistically and, consequently, cannot be planned, organised and structured during the ethnography itself. However, when the raw data is analysed, we believe that it is useful to organise and structure this data in such a way that it is more accessible to system designers. We do this by providing a series of topics that can be used to guide observations and organise (or *structure*) fieldwork data. These topics have been developed in the ethnomethodological literature, particularly through the studies that have been presented for computing audiences (see, for example, [1]; [4]; [5]; [16]; [17]). The topics provide a comprehensive framework for considering features of social systems of work and how social systems interact with technical systems, rules, plans and procedures and the spatial arrangement (ecology) of the workplace.

We suggest eight different headings that may be used to structure and organise ethnographic data. However, we are not suggesting that these are the only ways to impose structure on this data or that the headings proposed are necessarily relevant to

every study and setting. Rather, from extensive experience, we have found that these structural devices allow a mass of data to be organised so that it becomes more accessible to system designers who can relate the ethnographic structure to the structure of the requirements and the design of the computer-based system.

3.1 Temporality and sequentiality

Since ethnomethodological studies are primarily interested in the production of order in social action it is easy to see the relevance of this topic. This has been a primary concern of such studies since their inception ([9]). Here the focus is on the actual, embodied achievement of a sequence of action (or interaction) from within. The meaning of actions for participants in a setting is at least partially determined by the context in terms of where something has occurred as part of a sequence (even *history*) of actions. That activities are part of a sequence, that things get done one after the other, that activities happen closely in sequence, further apart and have a precise placing is important to the meaning they have and the sense they make to those involved.

Clearly linked to this is the importance of the temporal dimension to how action and interaction unfolds. Within the flow of action or interaction the notion of how actions relate to previous actions and preface future ones is essential to understanding. Structurally, ethnomethodology is interested in the emergent order (temporal and sequential structure) of activities over time and the practices employed to achieve this order. In work and technology studies the interest is, for example, with how this structure relates to (or works with or against) the temporal and sequential structure of procedures that may be instantiated in computer systems.

Joanne 27/6/05 12:07
Comment: There's either something missing or something worng with this sentence - it doesn't hang together in any case

3.2 The working division of labour

Many workplaces are characterised by an organisationally explicit, formal division of labour. People are given 'positions', 'job titles', or 'roles', to which 'duties' and 'responsibilities' are attached and may well be documented as 'job descriptions'. The ethnomethodological take on formal descriptions of divisions of labour is to offer a re-specification by including 'working' to focus on the fact that a division of labour must be achieved in practice, in situ, by the personnel. Where formal descriptions or representations of the division of labour and its operation exist there is often an interest in the relationship between these and the manner in which the division of labour works in practice.

The 'egological' and 'alteriological' principles refer respectively to how individuals within a working division of labour, in an on-going fashion, firstly, delineate their work from the work of others. Secondly, they also orient their activities such that they fit with the work of others (or make other's work easier). These concepts as

a means for understanding the actual operation of a division of labour were first suggested by Anderson and colleagues ([1]) and are a useful means for considering how the formal delineation of different 'jobs' or 'positions' is made to work by the social system of work.

The structural separation of ambulance control workers into call takers, dispatchers, supervisors and so on in a formal scheme is made to work on the control room floor by workers in the different positions doing various forms of 'boundary' work. Call takers carry out work that is oriented to dispatch in ways that go above the letter of their position, but they also do work that delimits their job as separate from dispatch. Supervisors supervise but also take calls and dispatch ambulances and so forth. The formal scheme is made to work by a social system that cannot be described so cleanly – it instead operates according to the manner in which requirements develop, for jobs to be done or actions undertaken ([20]).

3.3 Rules, plans and procedures

Rules, plans and procedures are often written down in various documents, (e.g. lists, charts, reports, process maps) or are embedded in artifacts (computer systems, checklists etc.). In the case of procedures they may encapsulate a temporal and sequential structure that is formally specified outside their situations of use. In all cases, they have to be 'written in' to sets of actions – they need to be translated to 'how to do this, here' – and they serve as a means for interpreting actions – demonstrating and deciding whether actions are in line with them. These features of their use and relation to action, in turn become the means for understanding e.g. 'what the rules are' or 'how you carry out procedure *x* correctly in this case' or 'the different ways in which plan *b* can be realised'.

A classic unsophisticated take on this is to state that rules, plans and procedures do not capture the full details of work or activity as it is played out but the more crucial point is to examine this *mutually constitutive and elaborative* relationship between rules, plans and procedures and the actual work undertaken. Where do they (and in what way), guide, constrain or drive action and interaction? How is action and interaction conducted as to orient to rules, plans and procedures and so forth? Clearly, the relationship is variable – sometimes people are strongly constrained by process and action has a more 'set' quality. Other times the relationship between the two is far looser. Here we have the nub of our argument on structure – formal descriptions of procedure, structured, for example, as a process model or as a workflow on a technical system have particular types of relationship to social systems of work. Explicating this relationship is of special interest when considering socio-technical systems.

3.4 Routines, rhythms, patterns

Human activities have an order and an orderliness that follows routines, rhythms and patterns. This is the case irrespective of whether the activity is formally planned or not, or whether officially documented rules and procedures are being followed or adhered to. However, the 'rules' for 'getting served in this bar' or the 'process for getting the kids to school' are rarely explicitly documented, but these activities of course exhibit routine qualities. Importantly, one should realise that this orderliness is something that is achieved in the doing rather than as something that can be specified extant to these situations. Often, such mundane (everyday) routines are not marked out (i.e. remarked upon), they are just carried out as such, with no explicit or formal representation. Indeed, their routine (and ordered) nature can be revealed by the fact that noticeable deviations are marked out, commented on, shown to be non-routine, clearly repaired and so forth. Researchers (e.g. [6]; [32]) have discussed non-work related activity in the home in such terms.

'Patterns' is another term used to capture routine aspects of activity. 'Rhythms' (e.g. see [26]) too is similar, but nicely brings to mind the importance of the temporal dimension to activity. Therefore, although social activity is structured from within there is also a sense in which it exhibits regularities or patterns. Many work situations differ only in the sense that routines, patterns and rhythms are seen to have specific relationships with formally specified rules, plans and procedures.

3.5 (Distributed) coordination

Ethnomethodological studies commonly describe the means by which people coordinate their activity, whether this is people working in a division of labour or collaborating in some activity. They can focus on coordination in fine grain detail or on a more general level. Coordination may be achieved face-to-face, as in the workings of a team in a control room, or may be remote and distributed and achieved through technology, e.g. CSCW or CMC (computer mediated communication). Coordination between participants, achieved through talk and action, may be a routine or regular feature of work or may be more ad hoc, happening occasionally. But from an ethnomethodological perspective, coordination is seen as something that is *always occasioned*, that is motivated by something specific and is directed to achieving something specific, whether this happens often and regularly or only now and then.

It is not just the activities or means of cooperation that are of interest but what gives rise to it and what it is directed at achieving. As with awareness (a means by which coordination is achieved) below, this topic focuses on the reasons initiating group collaboration and methods through which group work is achieved. Formal descriptions (e.g. procedures) or inscriptions (system workflow) embody a distinction of tasks and often roles. These are made to work through workers making each

other aware of aspects of their work and through coordinating their work through e.g. talk, hand offs and so on.

3.6 Awareness of work

This topic concerns the means by which co-participants in a working division of labour or in a concerted activity become aware and make others aware of important aspects of their activities for getting the job done. For instance, this can involve looking at the methods by which participants make their activity available for others to pick up on, or through observing the ways in which participants seek out information on the activity of others. In face-to-face situations, being there, in a shared situation may provide a ready context within which awareness 'needs' may be worked out. In distributed situations such 'awareness work' may be computer supported or more explicitly achieved. Understanding how and why this works (or fails), has been an important topic in these studies of work and technology use.

3.7 Ecology and affordances

The spatial structure of a setting – the arrangement of people and the configuration of artifacts (pre-designed and designed through use) are related to the ways in which activity gets done, what participants can see, do, how they may interact with others and through which means. For example, co-location in part of an office may allow participants to oversee and overhear one another, providing for on-going supervision of work, ready assistance and the ability to coordinate activities tightly. Distributed settings may create greater separation of activity or may require more work to coordinate activities or may require different types of support. A related notion is that of *affordances*, originally derived from the ecological theory of visual perception ([11]). Slightly different conceptions of affordances exist, but all are related to the way in which aspects of the environment and objects in it provide resources for the purposes of action and interaction. For example, a cup might be said to afford *picking up* and *drinking from*.

Joanne 27/6/05 12:16
Comment: I've reordered here to avoid the split infinitive

The ethnomethdological perspective on affordances stresses their inherently social, as well as learned, nature ([28]). It is through being regular participants in a setting that people can readily infer details on the status of work and what other people are doing through looking and listening. The competent participant can look at another worker looking at a screen and know that they are working on the dispatch of an ambulance or can tell that a pile of paper in that person's in-tray means that there is a backlog of invoices to be signed off. When considering ecology, affordances and structure we can see that the social system of work in a setting develops partially in response to the spatial structure of the environment (the arrangement of the workplace, what the artifacts are and where they are placed, where people sit etc)

Joanne 27/6/05 12:18
Comment: There's something wrong here - it is not a complete clause. It may just be that'where' can be delted from the start.

but that people also deliberately shape their environment. For example, objects and artefacts are placed and arranged to better suit the work practices of the personnel.

Joanne 27/6/05 14:01
Comment: It's not clear what other part of the sentence this relates to. Could you perhaps simplify the whole thing (maybe even split it over two sentences)

3.8 Skills, knowledge and reasoning in action

A final topic of interest relates to skills, knowledge and reasoning *in action*. This topic is related to psychological approaches that focus on *conceptual*, *cognitive* or *mental models* of users. These are usually inferred from user actions, reports and accounts but are taken to reside 'in the head' of users and are then exhibited in some way in their actions. The ethnomethdological re-specification of these topics is to acknowledge that while people do 'think to themselves' it is not our interest to try and extract this by whatever technique. Instead, the studies focus on explicating skills, knowledge and reasoning as it is exhibited in everyday action and interaction.

Joanne 27/6/05 14:05
Comment: it's not clear what is exhibited. I wonder whether you mean ***and*** *are then exhibited* rather than ***which*** *are then exhibited*but I'm not sure.

The way people reason about their activities is shown in what they say about them and how they carry them out, just as their knowledge and skills are exhibited in their talk and actions. When system analysts work through whether a new system fits with current work practice they show how they reason about the nature of 'fit', and other aspects of design – e.g. as a process of *transformation* or *standardisation* ([23]). When customers are engaged in calls to telephone banking the ways in which they talk through sequences of actions on their accounts (with reasons, justifications, explanations etc.) exhibit how they reason about their financial matters ([3]). In both these situations knowledge is brought into play by participants and skills are deployed in getting the work done. System analysts demonstrate their knowledge of the social patterns of work, of the developing system, of previous experience etc. in making design decisions, one of their key skills being the ability to cooperatively sift through disparate material to come to agreed design decisions.

In telephone banking, operators are seen to use their knowledge of everyday financial reasoning and customer histories to suggest courses of action for customers, while skillfully guiding the customer through bank processes carried out on a computer system. It is in these types of ways that skills, knowledge and reasoning are readily accessible in people's talk and action as observed. When considering structure, of particular interest here, is how these *practical epistemologies* (particularly as related to the kinds of knowledge and usage that stand invisibly (until revealed by ethnographic work) behind the world of representation, object, function and process) relate to aspects of the structure of technical systems.

4 The Social Structure of Work

In the previous section, we introduced a set of headings under which we believe it is useful to organise the ethnographic record and, in some cases, they may be effective

in focusing ethnographic studies. These headings are not, of course, arbitrary, but reflect perspectives through which we believe it is possible to discern *the social structure of work*.

The social structure of work can be thought of as the way in which work is organised as a social process – how organisations perceive how work should be done by their employees and how this is reflected in actual practice by the people doing the work. Unlike a system architecture, say, it is a more subjective, dynamic concept and cannot reliably be expressed as a set of static models.

Broadly speaking we suggest that there are three relevant forms of structure which are central to the social structure of work:

(1) **Temporal and sequential structure:** how processes and practices unfold – the relationships between entities, actions, utterances etc. over time in sequence.
(2) **Spatial structure:** related to the spatial relationships between objects, persons, actions and so forth.
(3) **Conceptual structure:** (sometimes also termed *ontological*, in a particular usage in computing) what a set of objects, entities, people, actions are, how they can be individuated and how they relate to one another conceptually[3].

Of course, these notions are also applicable to some extent to the structure of technical systems. The temporal and sequential structure reflects the assumptions of systems designers as to the sequences of operations that the system will support and the dependencies between the members of these sequences. The conceptual structure is, in essence, the system and data architecture and the abstractions used in the system design. The spatial structure is, perhaps, less significant because of the intangibility of software but may be reflected in some systems where the physical positioning of hardware is significant or in the layout and organisation of the system's user interface.

Ethnomethodological studies of work are often interested in the temporal and sequential structure of processes in the technical system (structured as a series of definite steps – *'workflow without'*, cf. [2]) and how well these processes mesh with the ways in which the social practices are structured temporally and sequentially from within. Commonly, the temporal structures of the technical system are much more rigid than the fluid, reactive structures of the social system and this leads to a mismatch where users are frustrated by the restrictions imposed by the technical system.

To understand the fit between the temporal structures of the social and technical systems, the ethnographer has to ask questions such as:

---

[3] Dourish states that questions of ontological or conceptual structure "…address(es) the question of how we can **individuate** the world, or distinguish between one entity and another; how we can understand the relationships between different entities or classes of entity; and so forth." ([8], p.129)

- Does the technical system assume that actions are carried out in a particular sequence and does it attempt to enforce this sequence? Do users have any control over the assumptions made and can they modify the temporal assumptions embedded in the technical system?
- Can sequences of actions in the technical system be interrupted and resumed without extensive rework? Can privacy be maintained during an interruption? Are there reminder mechanisms to show users where they are in a sequence of actions? Can previous sequences be consulted to illustrate what to do next?
- How are exceptions handled? Are there mechanisms for managing exceptions built into a sequence or must the user leave that sequence to handle the problem? Is it clear how and where to resume a sequence after an exception? Can actions be 'undone'?

When considering the relationship between social practice and the technical system the conceptual structure is particularly significant for three reasons:

- The conceptual structure of the technical system may not accurately match the *practical epistemology* of participants – the structure of objects, entities *in action* as currently configured and understood by participants in the workplace? [4]
- The conceptual structure of a technical system may constrain or facilitate possibilities for realising processes in a flexible manner[5]
- The degree of understanding that the personnel in the social system have of the conceptual structure of technical systems, as a faulty or incomplete understanding may cause problems with their use of the system.

Usually, the users' access to the conceptual structure of systems is through their interaction (individual and collaborative) with the system. Their experiences are of

---

[4] We want to suggest that ethnographic results should not be seen as producing an ontology but what we will prefer to call a *practical epistemology*. This refers to the kinds of knowledge and usage that stand invisibly (until revealed by ethnographic work) behind the world of representation, object, function and process. Referring to it this way avoids some of the confusions that Dourish ([8]) mentions in pointing to the way the term, 'ontology', is used both to describe, the 'internal representational structure of a software system' and 'the elements of a user's conceptual model; the model of either of the user's own work or their model of system operation' (ibid: p130).

[5] For example, the demands for integrating previously non-integrated processes, during computerisation, by producing a coherent, all encompassing conceptual model may then place restrictions on the variety of ways in which individual processes may be realised for different local user groups.

some of the temporal and sequential manifestations of the conceptual model as workflow, and of the spatial (or quasi spatial) manifestation of aspects of the conceptual structure in the menu and icon arrangements on the graphical user interface (GUI). Needless to say, these manifestations provide a partial, abstracted and sometimes misleading view of the conceptual structure of the system. The extent to which users should be made aware of the conceptual structure of a technical system is a moot point. However, as Whalen and Vinkhuzen ([34]) note in their study of a photocopier help desk, where the conceptual structure of an expert problem diagnosis system was concealed from the users:

> *...users, despite [the company's] intentions, are regularly and necessarily engaged in various kinds of analyses but are denied full access to knowledge that would make such analysis more effective, accurate and reliable. (p.16)*

4.1 Rules, plans and procedures

We defined the social structure of work in the previous section as a reflection of both how work is perceived by an organisation and how that work is actually carried out by people. The organisational view is normally defined in sets of rules, plans and procedures. Rules define conditions that must be maintained (e.g. credits and debits must balance), plans (or processes) define workflows (e.g. what steps are followed to close an account) and procedures define the particular ways in which activities are carried out (e.g. how to validate a customer's identity).

Practical experience, as well as a wide range of ethnographic studies, tells us that the way in which work is actually done and the way in which it is set out in the rules, plans and procedures is often markedly different. Different people interpret the organisational rules, plans and procedures in different ways depending on their competence, knowledge, status, experience and the contingencies of each particular situation. Drawing on Suchman ([31]), Schmidt ([27]) and Wittgenstein ([35]) we can understand the relationship between rules, plans, procedures and social action as one where:

- Social action and practices do not follow rules, plans and procedures to the letter as these can never exhaustively specify how they should be put into practice for *these* circumstances in *this* given situation.
- Social action and practices have a variable relationship with rules, plans and procedures which sometimes have a strongly constraining influence on what actions may be taken in practice, other times they offer great flexibility – it depends on the rules and the social practices surrounding their use.
- Social action and practices, on the one hand, and rules, plans and procedures, on the other, are mutually constitutive and elaborative: social agreement that a set of actions and practices falls within the specifics of a rule in any given

case elaborates, in an on-going and incremental sense, shared understandings of just what a rule covers; and also (re)constitutes the 'set' of activities that are agreed upon as rule following.

Given this understanding we can look at the relationships between the social and technical systems and rules, plans and procedures. Firstly, it is important to draw distinctions in these relationships. The participants in a setting (the social system) *interact* with technical systems and the rules, plans and procedures of the organisation in similar ways – and they, in turn, shape the work of the personnel. Technical systems and rules, plans and procedures tend to be related. Organisational rules, plans and procedures, as formally specified, often become (or are the key resource for designing) the procedures instantiated in technical system. Conversely when a technical design project is used to produce an integration of procedures or manage organisational change, the production of a technical system may lead to a re-description of organisational rules, plans and procedures.

Therefore, to re-iterate, rules, plans and procedures are often manifested in technical systems in terms of permissions and rules, workflows and access rights. It is the personnel engaged in their activities and interacting with the systems who make this work in practice. Furthermore, (from an ethnomethodological perspective) the social system is not amenable to the same formal, extrinsic description as technical systems and rules, plans, and procedures. Technical systems, and rules, plans and procedures structurally have a static quality unless they are being re-designed. However, as discussed above, the people working with these artifacts are necessarily engaged in interpretive work, a process of deciding what they should be doing now given the rule, or the system output. And consequently this leads to an understanding of what constitutes the process, the rule, the work; and what is allowable within the rules, and what actions stand outside them.

A key design consideration that flows from this analysis is that of the desired flexibility that is built into rules, plans and procedures and technical systems. The main point to note is that when rules, plans, and procedures are instantiated strictly, to the letter, in technical systems they may allow less flexibility than when they were only documented on paper. Taken out of the contexts in which they were followed, or adhered to, their structure may suggest a rigidity, a specific set of interpretations, a finite set of actions, that was not borne out, as such, in practice. If the technical instantiation of rules and procedures is based on the documentation alone rather than its elaboration in practice this can lead to inflexible systems that can cause serious usability problems

In some cases, often safety critical situations such as power plant operation, there may be a good case for making procedures particularly strict, constraining the space for interpretation, for options, for different possibilities for achieving the same goal. In these situations, embedding the rules and procedures strictly in technical systems may well be desired. However, in many other settings, the system may end up incor-

porating an unforgiving rationalisation of work that interferes with work practice. Bowers' and colleagues' ([2]) study of the introduction of workflow technology in the print industry is a prime example of this. The workflow technology presumed that print jobs would be owned by a user, begun and then followed through to completion, with the technology providing an accurate audit trail of this process. It was only the introduction of the technology that revealed that smooth workflow of this type was just one of the crucial processes that workers oriented to. They also oriented to maximising the use of machines, dealing with breakdowns, prioritising jobs and so forth. This meant that in reality print jobs were stopped, re-ordered, re-allocated and so on. The workflow software was so unforgiving in allowing this other work that personnel ended up retrospectively creating a smooth workflow record of all the print jobs at the end of the day instead of in real-time.

For the reasons discussed above, it is crucial to consider how rules, plans and procedures are elaborated in practice when using them as a resource for design. Whether the desire is to maintain or alter their flexibility, understanding their current use provides a better resource for making these decisions and working out what the possible consequences might be. It should be noted that the same argument applies to why it is important to study technical systems *in use* when considering re-design, rather than just considering the technical system *in isolation*.

5 Social structure and system dependability

How do we now approach the dependable design of socio-technical systems given this understanding of 'social structure' in socio-technical systems? The key issue here is to understand the way in which the structure of technical systems and the structure of rules, plans and procedures, fit with that of the social system. The relationship cannot be adequately described in formal structural terms, i.e. it is not possible to produce an accurate *model* of a socio-technical system because social practices are structured from within while technical systems are structured from without (they have a structure that can be specified separately to the technology they are implemented in).

Technical systems can be, and are, modelled, social practices are emergent, dynamic and are always responsive to the contingencies of *this situation, this time*. Models of social practices abstract, gloss and rationalise these features of them, giving them a rigid, formulaic structure not found 'in the wild'. Therefore, while models of social practices can be made *commensurate* with those of technical systems, i.e. by encapsulating a user model in the structure of the GUI, caution needs to be applied when considering how usable the system will be (how well it will fit in with social practices). The abstractions, glosses and rationalisations of practice used to construct the *idealised* user model may have problematic consequences when im-

plemented in a real, dynamic and contingent situation. Social practices will have to adapt in a way that enables users to carry out what they need to do, in each case, in response to the idealised user model encapsulated in the system. The idealised user model will not match what they already do, and it may well clash quite badly with certain crucial aspects of everyday practice.

Technical systems, however, need to be built using user models and models of work. Does this necessarily set up a serious problem? Fortunately the answer to this is no, for two reasons. Firstly, humans and the social systems they form are necessarily adaptive. They respond to the contingencies of *this* situation, *this* time, and they can also adapt their practices over time to work successfully with a computer system that initially fitted badly with their work practices. Secondly, user models can be created through observation 'in the wild' rather than theoretically conceptualised. A user model or model of work based on a faulty or incomplete understanding or work, or created through imagining what users do, rather than discovering what they do runs serious risks of misunderstanding the users or misrepresenting their work.

A key feature of system dependability concerns efficient and effective socio-technical system operation such that personnel will be able to achieve work with technical systems successfully. This includes the extent to which technical systems will not have to be worked around, and will not inhibit important social practices, or getting the job done. Achieving dependability also includes an assessment of how reliable, safe, secure, resistant to failure these processes and practices are. A design process therefore involves an assessment of current working, and is often characterised by a desire to transform things to make them better or more dependable. The desired design is envisaged to 'preserve' certain adaptive, or desirable, patterns of work, while transforming inefficient, maladaptive or inconsequential practices for organisational gains. Better decision making in this process should be facilitated by a detailed understanding of current process and practice.

Sommerville et al. ([29]) point out that dependable design involves successfully sorting out how the following four non-trivial 'problems' relate to sets of requirements[6]:

[1] What characteristics of the existing manual system are unimportant and need not be supported in an automated system?

[2] What are important manual activities which need not be supported in an automated system because the activities are a consequence of the fact that no automated support is available?

---

[6] A fifth problem might be added to Sommerville's list, namely: 'What activities not present in the manual system become necessary following the introduction of the automated system? How can these be integrated with residual manual activities and how might they need to be supported through training?'

[3] What characteristics of the manual system must be replicated without change in an automated system?

[4] What activities from the manual system may be supported in a way which is different from that used in the manual system?

Reviewing this list (with the proposed fifth question added) we can make the following comments. Firstly, the questions can be seen to apply equally to cases where one technical system is being replaced by another (rather than a manual system by a technical). Secondly, these are *perennial* questions for designers and developers, and can be characterised as being to do with deciding what current work activities need to be preserved, what can be transformed and in what way. Thirdly, issues of 'fit' between the new system and current practice are central to considering these questions. For example, if a new system does not fit well with the current social structure of work, the question then is posed as to whether this matters, and if it does, how will this be solved. Will the technical system need to be reconfigured, or will users need to be trained to carry out their work in a different manner, to acquire new skills and so forth. Finally, it is important to acknowledge that other matters also impinge on this problem solving or decision making activity. Building a new system to be more reliable, efficient, to transform work is likely to produce requirements that go against the new system fitting well with existing patterns of work. The question then becomes one of trying to ensure these benefits can be achieved as envisaged, since if the new system proves problematic to work with, many sorts of unforeseen consequences may well follow.

Joanne 27/6/05 14:15
Comment: should this be *manner* rather than *matter*?

These sets of questions are clearly tricky, as it is not possible to know accurately in advance *just what effects* a new system will have on existing work practices. However, it is clear that the type of material produced by ethnomethodological ethnographies can serve as a useful *resource* to aid designers and analysts in trying to sort through these problems.

Joanne 27/6/05 14:17
Comment: I have added a comma to this sentence as to work with many was difficult - I think it's easier to read now but let me know if you think I have altered the meaning.

In a study we have been conducting of the design process to configure and deploy a customisable-off-the-shelf (COTS – after looking this up I understand your point but still found that customisable-off-the-shelf systems are also called COTS systems) system for a UK hospital Trust ([22]; [23]), we have been struck by how often problematic areas in the design are identified as situations where there is a perceived bad fit between the developing system and the current structure of work. We discovered that the project team in this complex setting (1) had problems finding out exactly how work was carried out; (2) perceived problems where the system did not fit with the rhythms and routines of current work practice; but (3) had difficulties working out how serious these problems would be (especially in cases where they realised that since no easy technical solution could be found the issue would have to be dealt with by training); (4) needed to treat perceived problems as more serious when a negative clinical impact was suggested; and (5) needed to treat other requirements,

Joanne 28/6/05 10:34
Comment: My understanding is that the C usually stands for common (or commercial)

like various integration demands, as more important than achieving a good fit with current work structures.

In this case, given that a key question in the analyst's discussions of potential problems of fit was always about gaining an accurate, independent description of the actual working practices, this suggests that ethnographies could be successfully *targeted* to explicate the practices in question. Often discussions centred around questions like 'Are tasks in A & E often interrupted, what does this mean for logging off and security?' or 'When the users tell us that the system means more work or that they cannot do a current activity, will it really mean more work, or just a couple more button presses, and is the activity actually crucial to their work?'. While ethnography would not tell you how to design the system it can answer questions like those above, and in doing so become a useful resource when making design decisions during build and configuration.

All of this leads to an understanding of how dependable design may pragmatically proceed. Attention must be to the detail of the actual operation of sociotechnical systems; the details of how the social setting is organised, uncovered by studying the work and activity going on in that actual setting. Understanding the work better should help minimise chance being involved in constructing models of users and their work. We have argued that ethnomethodological studies are particularly well suited to be used as a resource in designing systems (particularly their models of users and their work) that fit better with current practices.

When considering how well a system under construction (or as an evaluation, a computer system in use) fits with a social system of work we can decompose structure as having three elements: (1) temporal and sequential structure, (2) spatial structure, and (3) conceptual structure. We can then perform matching exercises by looking at how current work practices relate to system models of work. How well does the temporal and sequential structure of a technical systems fit with that of the social system? Does it facilitate work, or does it have to be worked around? If there are failures, where do they lie, do we need to adjust the technical process or seek to retrain users to achieve better practice? These are the questions designers need to sort out.

With conceptual structure, we are interested in the restrictions or flexibility the technical system places on how workflow may be realised for particular groups of users. Does the conceptual structure mean processes must be realised in a particular restrictive manner? Generally, we would argue that attention needs to be paid to try and support local practice as it is and even when systems are meant to be transformative of local practice, flexibility is often required as the processes and practices need to be evolved together during a period of domestication.

We have also discussed how another topic of study can be to find out what users' understandings of the conceptual models of technical systems are, and then to consider whether these understandings are useful in achieving dependable operation. Do

they/would they lead to errors? If so, attention needs to be paid to changing the conceptual models of the technical systems, or, so to speak, of the users. When dealing with the users one might want to consider how a different temporal and sequential structure to process, or a different spatial arrangement to a GUI might help them to understand the conceptual structure if not in a more truly accurate fashion, in a functionally dependable one.

Finally, ethnomethodological studies have a particular interest in the spatial structure (the ecology) of the workplace and how the arrangement facilitates or constrains work achievement. As control room studies have shown us, system dependability can be facilitated by the particular arrangement of people and technologies in a setting. That certain technologies are public and shared, that staff can oversee and overhear one another and so on has been shown as crucial for dependability. Doubtless some of the practices we see have developed in response to arrangements that were not deliberately designed for those settings. However, it is important and useful to understand these relationships when thinking about re-design or design for new settings.

6 Conclusion

Approaching the issue of dependable design in organisationally embedded socio-technical systems from the perspective of the social structure of work enables a new way of thinking about the relationships and interactions between social systems and technical systems. Although we have argued that the structures of these systems cannot accurately be united in a single form in a model (or series of models) we believe that a structural approach allows us to make deeper, more fundamental, connections between them and may aid in mediating between the rich descriptions of field studies and the abstract modelling of technical design.

This can be achieved in actual projects through a *practical* (rather than theoretical) consideration of fit between the structure of the social systems and that of the technical systems. This can proceed in a number of ways as outlined above. Understanding the social systems of work in the first place, prior to design, is a good place to start. This should facilitate the design of more appropriate workflows, conceptual models of users and so on. However, all design is transformative – social practices will necessarily change in some ways - and the consequences cannot all be correctly imagined in advance. Therefore design is likely to proceed from there adjusting the technical structure and adapting the work practices through a period of evolution until the desired level of dependability is achieved. Ethnomethodological studies can be a surprisingly useful resource to better enable decisions regarding the ‘design’ of work and the design of technology during this process.

References

[1] Anderson R, Hughes J, Sharrock W (1989) Working for profit; The Social Organisation of Calculation in an Entrepreneurial Firm. Aldershot: Avebury.

[2] Bowers, J., Button, G., & Sharrock, W. (1995). Workflow From Within and Without: Technology and Cooperative Work on the Print Industry Shopfloor. In Proceeedings of ECSCW '95, Kluwer, Stockholm.

[3] Bowers, J. & Martin, D. (2000). Machinery in the New Factories: Talk and technology in a bank's call centre. In Proceedings of CSCW 2000, Philadelphia, © ACM Press.

[4] Button G, Dourish P (1996) Technomethodology: Paradoxes and Possibilities. In Proceedings of ACM CHI 96 Conference on Human Factors in Computing Systems 1996, v.1, pp.19-26 © Copyright 1996 ACM.

[5] Crabtree A, Nichols DM, O'Brien J, Rouncefield M, Twidale MB (2000) Ethnomethodologically-Informed Ethnography and Information System Design. In Journal of the American Society for Information Science, 51(7), pp. 666-682.

[6] Crabtree A, Hemmings T, Rodden T (2002) Pattern-based support for interactive design in domestic settings. Proceedings of the 2002 Symposium on Designing Interactive Systems. London: ACM Press.

[7] Crabtree A (2003) Designing Collaborative Systems: A practical guide to ethnography. London: Springer-Verlag.

[8] Dourish, P. (2001). Where the Action Is: The Foundations of Embodied Interaction. Cambridge: MIT Press.

[9] Garfinkel H, (1967) Studies in ethnomethodology. Englewood Cliffs, N.J.: Prentice-Hall.

[10] Garfinkel H, (2002) Ethnomethodology's Program: Working out Durkheim's Aphorism. Rowman & Littlefield, Oxford.

[11] Gibson JJ (1979) The ecological approach to visual perception. Boston, Houghton Mifflin.

[12] Grudin J (1990) The Computer Reaches Out: The Historical Continuity of Interface Design. In proceedings of ACM Conference on Human Factors in Computing Systems. CHI'90: Seattle, Wv.1, pp.19-26 © Copyright 1996 ACM.

[13] Heath, C. & Luff, P. (1992). Collaboration and control: Crisis management and multimedia technology in London Underground line control rooms. Computer Supported Cooperative Work, 1, 69-94.

[14] Hughes J, Randall D, Shapiro D (1992) Faltering from ethnography to design. Proceedings of ACM CSCW '92, Conference on Computer-Supported Cooperative Work, pp115-122. © Copyright 1992 ACM

[15] Hughes J, King V, Rodden T, Andersen H (1994) Moving Out from the Control Room: Ethnography in System Design. Proceedings of ACM CSCW '94, Conference on Computer-Supported Cooperative Work, pp429-439. © Copyright 1994 ACM.

[16] Hughes J, O'Brien J, Rodden T, Rouncefield M, Blythin S (1997) Designing with Ethnography: A Presentation Framework for Design. Proceedings of DIS'97: Designing Interactive Systems: Processes, Practices, Methods, & Techniques 1997, pp.147-158 © Copyright 1997 ACM.

Joanne 27/6/05 14:35

Comment: Is this really J Rodden or should it be T?

[17] Hughes J, O'Brien J, Rodden T, Rouncefield M (1997) Ethnography, Communication and Support for Design. CSEG Technical Report Ref: CSEG/24/1997. http://www.comp.lancs.ac.uk/computing/research/cseg/97_rep.html

[18] Luff P, Hindmarsh J, Heath CC (eds.) (2000) Workplace Studies: Recovering work practice and informing system design. Cambridge: Cambridge University Press.

[19] Lynch M (1993) Scientific Practice and Ordinary Action: Ethnomethodology and social studies of science. Cambridge: CUP.

[20] Martin D, Bowers J, Wastell D (1997) The Interactional Affordances of Technology: An Ethnography of Human-Computer Interaction in an Ambulance Control Centre. Proceedings of the HCI'97 Conference on People and Computers XII 1997, pp 263-281.

[21] Martin D, Rouncefield M (2003) Making The Organisation Come Alive: Talking through and about the technology in remote banking. In Human-Computer Interaction. Vol 17, No's 1 & 2, pp. 111-148.

[22] Martin D, Mariani J, Rouncefield M (2004) Implementing an HIS project; Everyday features and practicalities of NHS project work. Health Informatics Journal. Vol. 10 (4), pp 303-313, Sage, London.

[23] Martin D, Rouncefield M, O'Neill J, Hartswood M, Randall D (2005) Timing in the Art of Integration: 'That's how the Bastille Got Stormed'. Paper Submitted to Group 05.

[24] Pettersson M, Randall D, Hegelson B (2002) Ambiguities, Awareness and Economy: A Study of Emergency Services Work. In Proceedings of ACM CSCW 2002.

[25] Pettersson M, Rouchy P (2002) 'We don't need the Ambulance then' – Technological Handling of the Unexpected. Presented at the XV World Congress of Sociology, 7-13 July 2002, RC.25 Language, Technology and Work, Brisbane, Australia.

[26] Reddy, M. and Dourish, P. (2002). A Finger on the Pulse: Temporal Rhythms and Information Seeking in Medical Work. Proceedings of the ACM Conference on Computer-Supported Cooperative Work CSCW 2002 (New Orleans, LO), 344-353. New York: ACM.

[27] Schmidt K (1997) Of Maps and Scripts: The status of formal constructs in cooperative work. In Proceedings of Group 97. © ACM Press.

[28] Sharrock W, Anderson R (1992) Can organisations afford knowledge? Computer Supported Cooperative Work, 1, pp. 143-162.

[29] Sommerville I, Rodden T, Sawyer P, Bentley R (1992) Sociologists can be surprisingly useful in interactive systems design. Proceedings of HCI 92, 341-353. Cambridge, UK: Cambridge University Press.

[30] Sommerville I, Rodden T, Sawyer P, Twidale M, Bentley R (1993) Incorporating Ethnographic Data into the Systems Design Process. In Proceedings of RE 93: International Symposium on Requirements Enginnering, January 4-6, San Diego, IEEE Press, pp. 165-174.

[31] Suchman, L. A. (1987). Plans and situated actions: The problem of human-machine communications. Cambridge, UK: Cambridge University Press.

Joanne 27/6/05 14:46
Comment: Really I Rodden or should it be T?

[32] Tolmie, P., Pycock, J., Diggins, T., Maclean, A. and Karsenty, A. (2002) "Unremarkable computing", Proc. of CHI 2002, pp. 399-406, Minneapolis: ACM Press.
[33] Viller S, Sommerville I (1999) Coherence: an Approach to Representing Ethnographic Analyses in Systems Design. Human-Computer Interaction 14: pp. 9-41
[34] Whalen, J. and Vinkhuyzen, E. Expert Systems in (Inter)Action: Diagnosing Document Machine Problems over the Telephone. in Luff, P., Hindmarsh, J. and Heath, C. eds. Workplace Studies: Recovering Work Practice and Information System Design, Cambridge University Press, Cambridge, UK, 2001, 92-140.
[35] Wittgenstein L (1958) Philosophical Investigations, Oxford, Blackwell.