ROLES ARE RESPONSIBILITY RELATIONSHIPS REALLY

J.E.Dobson *, I. Sommerville †

*31 Wentworth Park, Allendale, Northumberland U.K. NE47 9DR <J.E.Dobson@ncl.ac.uk>

† Department of Computing, Infolab21, University of Lancaster, Lancaster U.K. LA1 4WA <is@comp.lancs.ac.uk>

Keywords: information systems responsibility conversations ethnography

Abstract

Many approaches to role-based information systems assume that a role definition is one-sided, in that it applies to an individual person or single organisational entity. In fact, however, many roles are defined as a relationship between two entities (e.g. customer—supplier, doctor—patient), and it is the nature of the relationships in which an individual participates that defines a role. This reconceptualisation of a role as a relationship rather than as a set of properties of a single entity has a number of implications for requirements elicitation, system design, and information structures. We explore some of the implications and requirements and show how they can lead to significantly different system architectures. We also look at the use of ethnography in helping us understand responsibility relationships as they exist in practice and illustrate the complexity of responsibilities using examples drawn from ethnographic studies.

1 Introduction

This symposium is about people and computers, so we have chosen to explore one of the major differences between them: people can be given or assume responsibilities and computers can't. Things that look as if they are attempts to give computers responsibility, such as the common "It's a computer error" excuse, are really a way of saying "We know it's a human error, but we don't know whose but a computer was implicated in some way so we'll blame that". We don't really ascribe moral agency to computers, since we don't ever punish them no matter how egregious the error. We all know the computer only did as it was told, though we may not know who told it or what they told it or why. The real responsibility rests with a person.

The reason for starting from responsibility as a way of distinguishing the social from the technical in socio-technical systems is partly philosophical — that it is responsibilities that provide the basic structuring of society — and partly pragmatic because, as we hope to show, understanding responsibilities can lead to important considerations in the architecture of information systems that other current philosophical bases tend to ignore. Those who are interested in the philosophical concept of responsibility could read the book by Lucas [4], and the book by Jonas [3] for a deep discussion of its importance in society and the interpretation of responsibility in the context of the relationship between society and technology. A more general exposition of the relationship between philosophy and information systems can be found in [2].

We will have more to say about responsibilities later on, but we want here to introduce a closely related concept, namely that of role. We treat a role as a collection of responsibilities that in some sense go together. A role can be a purely social role, such as parent, a partly social and partly professional role such as doctor, or a fully professional role such as office secretary. Such roles are defined by the sorts of responsibility that they entail and the spaces or domains in which they are given (or assumed) and discharged.

Many conceptualisations of role in information systems treat it as a set of capabilities, such as information that can be accessed, tasks that can be performed. But where do these capabilities come from? Who grants them? — The answer lies in the prior notion of responsibility. Many workarounds (where capabilities have for whatever reason been inappropriately denied) can be seen as an alternative error recovery strategy to fulfil responsibilities. For example (and a colleague has reported to us something like this), a hospital ward sister might not be allowed by the computer system to reserve a vacant bed for a patient known and expected to be coming in later, but she can deliberately fail to notify the system that the bed has become vacant. In doing so she is showing her responsibility to the incoming patient by ensuring that a bed is available.

In fact this example shows a very simple example of responsibility conflict in roles: sister also has a responsibility to the hospital for efficiency in bed management which would require her to count accurately the number of vacant beds. These everyday conflicts of responsibility are probably ubiquitous in organisational life, and proper recognition, discussion and resolution of them has major implications for the process of producing a computer system architecture designed to support the work role in which the conflicts occur. This problem is well discussed in [1].

Thus a role is a set of responsibilities possibly with a number of different authorities. For example, a doctor has responsibility to the patient, to the NHS, to society. These can conflict, and there may be rules for conflict resolution: in certain circumstances, the duty of disclosure overrides the duty of patient confidentiality. In other cases, however, there may be no specific rules for conflict resolution and the

resolution process may be individual (What do I think is right in the circumstances?) or legal (What should have been done? Was the decision in fact reasonable?).

So we can define responsibility as a relationship between two roles regarding a specific state of affairs with respect to a particular mode such as bringing about, preventing, maintaining, accounting for and so on, such that the holder of the responsibility (the **responsible**) is responsible to the giver of the responsibility, the **authority**.

The important point here is that responsibilities cannot be looked at on their own but must always be considered as a relationship between two roles. The states of affairs for which responsibilities are held may be at any level of granularity of the organisation. For example the responsibilities may be at a very high level such as for the financial soundness of the organisation, for the continuity of the services provided by the organisation, for safety, for security, for the accuracy of information and suchlike, or they may be at an individual level for a very specific state such as whether a particular door is closed, or whether a particular form is correctly filled in.

In summary, then, the full characterisation of a responsibility consists of:

a) who is responsible to whom;

b) the sort of responsibility (these include accountability, blameworthiness, legal liability);

c) the mode and state of affairs for which the responsibility is held;

d) a list of obligations held by the responsibility holder (what the holder must do to fulfil the responsibility).

This reconceptualisation of a role as a set of responsibility relationships has a number of implications for role-based information system design, which we shall now explore.

2 Responsibility

It is useful to distinguish between causal responsibility, when an agent has an obligation to make something happen or prevent it from happening or to maintain a state, from consequential responsibility, when an agent is answerable when something happens or does not happen or a state is not maintained. These different responsibilities do not always rest on the same agent (the doctrine of 'ministerial responsibility'). Consequential responsibility may be held to rest with an organisation as a whole whereas causal responsibility most usually can be traced to an individual or the fact that no particular individual at the time held the responsibility. Causal responsibility may sometimes be delegated, though some responsibility remains with the delegating agent (i.e. the responsibility for having chosen to delegate), whereas consequential responsibility is not normally capable of delegation, though it may sometimes be transferred.

Causal responsibility is the responsibility for doing something. This can be something explicit (A is responsible for locking the building; B is responsible for checking that the building is secure) or implicit (A is responsible for building security). In the latter case, it is assumed that A decides on what actions are appropriate. Thus it can be reformulated as A is responsible for deciding on what to do, then doing it.

In order to complete some activity X, then a responsible A has to have:

1. A definition of activity X.
2. Sufficient time to complete activity X (before some assigned deadline).
3. Sufficient resources to complete X.
4. The required competence to carry out the activity X.
5. Necessary information on and understanding of policies and constraints that apply to activity X.

Note that this definition does not preclude A being some automated agent. In this case, notions such as 'required competence' translate into 'the right software'. If A is automated, the consequential responsibility should be, but often is not, clearly bound to a particular responsible person.

Although causal responsibility can be examined by methods such as task analysis, and asking questions about what information is needed and needs to be recorded in order to fulfil the responsibility, it is important to look at the whole set of responsibilities that define a role. It is also important to look at the nature of any shared or delegated responsibilities, and at conflicts of interest. Where a new ITC system is being procured, it is an important part of any requirements elicitation process to discover and examine all these things.

Shared causal responsibilities are a vulnerability. If there is inadequate means of communication between the parties, or an inadequate protocol, there is the possibility of a particular action being taken twice or more, or not at all. Usually a properly performed vulnerability analysis is capable of revealing these possibilities, but such an analysis is often not done as part of a requirements exercise.

Problems with responsibilities often arise when the actions associated with a consequential responsibility or the activities for which a causal responsibility exist cut across organisational boundaries. These problems may arise because of differences in interpretation of responsibilities, because of differing priorities in time and resource allocation in different organisations, because of differences in competence, because of different organisational policies, etc.

There may also be problems when there is confusion between causal and consequential responsibilities. For example, in organisation X, agent Alice may have the causal responsibility for some activity (say complete document K). This relies on contributions from agent Bob in organisation Y. However, Bob interprets the responsibility as a consequential rather than a causal responsibility (for whatever reason, perhaps organisational policy) so does not carry out the

activity himself. However, if Bob fails to inform Alice about this, Alice communicates with Bob about K. This may lead to timing problems (because Bob then has to communicate with someone else) or information problems (because Bob interprets the information provided by someone else).

In order to demonstrate that a consequential responsibility has been discharged, it is necessary to create and maintain evidence. In order to create such evidence, actions may be defined. For example, if the evidence is a test plan, then the associated action where Bob is consequentially responsible for the top-level goal is 'To ensure that the test plan is available and has been prepared according to required standards'. Notice that Bob is NOT (necessarily) responsible for preparing the evidence but is responsible for ensuring its existence.

By default, the holder of a consequential responsibility is causally responsible for the actions associated with that responsibility or for delegating that causal responsibility. Consequential responsibility does not require the use of resources to discharge the responsibility (resources, of course, are required for any associated causal responsibilities).

Thus consequential responsibility leads to requirements too, but they may be more indirect. Information recording and audit trail processes can be important in determining how a particular responsibility was discharged. Shared consequential responsibilities are particularly difficult.

In summary, roles and responsibilities are complex things, and simplistic models of them lead to inappropriate system architectures. Too often, a failure to perform a vulnerability analysis leads to a system which makes optimistic assumptions about they way that people have performed, or will perform, their duties, and about the effectiveness of their human communications. There is not enough space in this paper to describe how do a vulnerability analysis that covers all these things, but we shall look in more detail at one particular aspect, which is that of human communication and the conversations through which roles and responsibilities are negotiated.

3 Conversations

These observations, together with the fact that responsibility is a relationship, lead to the realisation of the importance of conversations. By 'conversation' here we mean an exchange between two roles during which they fulfil some of their respective responsibilities to each other. For example, a medical record is a record of a conversation — one which lasts the whole lifetime of the patient and which is distributed in time and space. Seen from the viewpoint of a particular patient, the doctor role may be bound to different individuals at different times, but the responsibilities do not change (much). The medical record is the partial state of a number of different concurrent conversations (doctor—patient, doctor—doctor). Many databases can be seen in this light.

Different conversations require different information structures. For example, when a GP refers a patient to a hospital consultant, the diagnosis as told to the patient is not necessarily the same as the preliminary diagnosis as suggested to the consultant — and this may be different again from a report to the public health or epidemiological medical authorities. Failure to appreciate these differences can be a source of system rejection. For example, our project has investigated a case in which the designers of a hospital system did not appreciate the fact that the categories used in reports to a funding authority did not match the categories used in clinical processes, with the result that the clinicians were forced to distort their accounts of their clinical work in order to fit them in to a mismatched framework for statistical purposes.

What we are going to do in this section is to look at the various properties of conversations and show how these influence the architecture of the record of the conversation, i.e. the information structures and storage which record it.

Conversations and the relationships that they define, sometimes fail. The purpose of our theory of conversations is to explain the failures associated with the intentions of the participants. It is clear that the bringing together of obligations and responsibilities can create conflicts of interest as well as synergies. It can also create overloads and imbalances which could lead to failure in operation. In addition to failures of organisational policy and design, we have operational failures due to a lack of correspondence between the expectations of the participants. In this section we identify the attributes of roles and conversations that provide a basis for analysing such situations.

Note that this analysis does not cover failures to perform the intended role by generating incorrect information, to misinterpret presentations or to proffer incorrect or inappropriate resources. These failures would be accounted for in a theory of instruments. It does not provide any explanation of failures in reading, writing or transporting data: this is the province of a theory of communication.

It is fundamental to the concept of a conversation that it is intended to have some significance for each of the participants, that they have some stake in or expectations about its outcome. The benefits generated by or exchanged in a conversation may be of different types or even belong to different value systems for each of the participants. Each makes an individual evaluation and thus a conversation has a *significance* which is particular to each of the parties. Conversations with high significance imply that the benefits at stake or the consequences of failure for one or both the parties is large. Two classes of conversations are distinguished on the basis of the intended balance of benefits.

i) Symmetrically significant conversations are intended to produce benefits which are judged as fair and more or less equal for each of the parties.

ii) Asymmetrically significant conversations occur where the main derivation of benefit is by one party. Benefits derived by the other party are the consequence of factors outside the immediate conversation such as a sense of

vocation or kinship or the acquisition of esteem from third parties.

A conversation is also defined in terms of *mutuality*. This refers to the level of responsibility each party accepts for the benefit to be derived by the other party and for protecting the other party form any harm associated with breakdown or misapplication of the conversation. This attribute has a magnitude and a distribution within a conversation; thus a relationship with high, symmetrical mutuality implies partnership and co-operation whereas high but asymmetric mutuality implies a relationship of care such as parent—child or educator—pupil. Clearly, if the significance of a conversation is asymmetrical and it also exhibits high mutuality then this will be asymmetrical. Zero mutuality is associated with the "*caveat emptor*" principle of the consumer—supplier relationship whereas negative mutuality implies competition.

A third factor which influences the acceptability and effectiveness of channels and media is concerned with the distribution of *control* between the participants. Thus, the pupil may speak when invited to do so and the teacher may demand silence. The party which is capable of initiating a conversation or causing a transition from one phase to the next derives considerable control and it is usual for imbalance in the distribution of initiative to reflect asymmetry in significance and mutuality. The issue of control is of particularly significance in multi-party conversations such as conferences and debates.

The fourth factor which characterises the parity and symmetry within conversations is the distribution of pertinent information and other relevant resources between the participants. The information referred to here is that required to conduct the conversation and derive the intended benefit from it. So a student who has not been taught to study independently is not able to benefit from a lecturer—student relationship and requires a teacher—pupil relationship with more asymmetric mutuality. In a conversation, each party must have the information and other resources appropriate for their particular role.

We refer to this as *capability,* which defines the set of resources required by each agent to properly fulfil the responsibilities of its particular role. These include the appropriate rights and capabilities in relation to the communications and information resources required to instrumentalise the conversation and also to the resources which must be deployed and possibly consumed in the discharge of the responsibilities associated with the role. Lack of an appropriate capability represents a source of failure. Market relationships, for example, imply parity of significance with low mutuality between buyers and sellers. An imbalance of information provides an opportunity for the better endowed participant to behave in an opportunistic way to achieve unfair advantage.

These properties of a conversation have important implications of the language and the medium within which the conversation is instrumentalised. The most significant of these is that in cases of high and asymmetric mutuality, the more responsible participant needs to maintain a careful check on the state of the less responsible participant throughout the course of interaction. This implies a rich communications channel where non-linguistic social and behavioural cues are accessible (e.g. primary school teaching, counselling, chairing a delicate negotiation or decision conference). Where mutuality and significance are symmetrical, the richness of the channels of communication and the medium of expression are relatively less important because the parties can rely on each other to co-operate in overcoming any channel and language limitations. One particular example would be the inappropriateness of introducing a co-operative information exchange system, such as Lotus Notes, into an organisation which deliberately fosters a competitive spirit between its employees.

One of the main uses of our conversation theory is to help analyses the composition of responsibilities into a single role. There are two basic cases of composition of responsibility relationships which can be used to illustrate the principles of conversation theory. These are illustrated in Figure1, below.

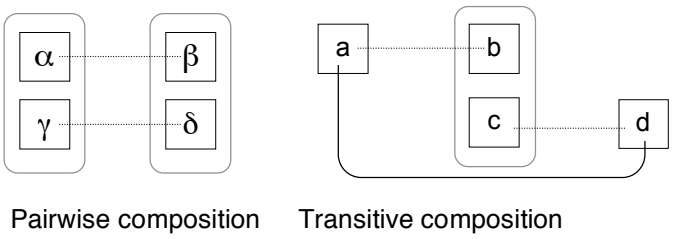

Figure 1 : Composition of responsibility relationships

In pairwise composition, the relationship α–β and the relationship γ–δ are combined. For this to be plausible and acceptable, each of the conversations need to be of comparable and compatible significance for each party. For example, the combination of a doctor—patient relationship with that of experimenter—subject, which occurs when medical research is conducted within a health care enterprise, can lead to potential conflicts of interest. Hence the special protocols which apply in such cases. In cases of asymmetric mutuality, e.g. borrowing money to create a creditor—debtor relationship, the concept of collateral is introduced to equalise the asymmetric significance. Thus, the lender's dependence on the borrower's continued commitment to repay is balanced by the borrower's dependence on the lender for continued access to pledged collateral.

In the case of transitive compositions, the key issues concern the nature of the relationship between (a) and (d). If they are independent, then the composition of (b) and (c) onto a job or mission for an individual or organisational unit is also an independent consideration. If, however, the (a-d) conversation is significant then a conflict may arise. For example, if the doctor (b) becomes the commissioned sales representative (c) for the drug company (d) and the patient (a) becomes the customer of that company because the doctor prescribes its drugs rather than drugs from some other producers, the doctor—patient mutuality has been compromised by a

conflict of interest. In contrast, if the doctor as a member (c) of a golf club (d) introduces the patient as a guest, then neither of the relationships can be said thereby to have been compromised.

In the case where the roles (a) and (d) are not independent of each other then transitive composition produces a role (bc) which is a common third party. In the case where the composite role, (bc) removes the need for direct interaction between (a) and (d), we have a intermediary or broking role. In cases where (a) and (d) continue to have a direct relationship, the third party role may either be supervisory in relation to this conversation or it may be supportive and infrastructural to it.

A different set of considerations arise when we consider capability and control in composite roles. Capabilities imply access to and use of resources, facilities, information and skills. These may interact when combined to create overloads or interferences rendering the composition of the roles inadvisable. Alternatively, combinations of roles can create efficiencies and economies through the reuse of capabilities. Finally, the distribution of control implied by roles which are to be combined must be broadly compatible: expecting the subordinate to be the teacher of the superior can be threatening and lead to tensions arising from role conflict.

We have offered a framework in which the normative definition of relationships can be analysed and compared in order to be able to reason about composite roles. The main features of this framework are as follows:

For a dyadic conversation to be self-coherent, then:

- The distribution of significance and the distribution of mutuality must be consistent. If the significance is asymmetric then the mutuality must also be compatibly asymmetric (implying that if one party has much and the other has little to gain then the second party must accept high mutuality if the conversation is to be coherent).
- The distribution of capability and control in a conversation must reflect the relative mutualities of the roles: if parties are to accept high mutuality then they must be empowered in the sense that each must have access to relevant information and be able to control the conversation as required.

For dyadic conversations to be pairwise composable then:

- The magnitude and the distributions of significance and mutuality must be broadly similar in each of the conversations.
- The combined capability for each party must be sustainable so as to avoid problems of overload and interference.
- The distribution of control in the two conversations must be similar.

For dyadic conversations to be transitively composable, then either

- The uncomposed roles are independent,

or

- The composed role is mediating between the uncomposed ones,

or

- It must be supervisory in relation to the two roles,

or

- It must be infrastructural to them.

What this is intended to show is the complexity of role, a complexity which must be represented by a role-based information system (and there are many such). A few years ago, one university, in trying to answer the question "What is a student?", found that the software it used did not have a rich enough theory of composition of roles to make it possible to combine in a coherent and conflict-free way for the purposes of a campus management information system, the various role relationships of pupil (as seen by the teaching service), tenant (as seen by the accommodation service), patron (as seen by the library), consumer (as seen by the catering service), counsellee (as seen by the student help service), patient (as seen by the health service), fee-payer (as seen by the finance office), beneficiary (as seen by the student hardship fund), employee (e.g. PhD students also engaged as teaching assistants) and so on.

Each of these roles has its own characteristic conversations defined in terms of significance, mutuality, control and capability, derived from the nature of the responsibilities involved. An information system that can properly deal with their composition (when required) and separation (when required) is not to be made by configuring a "universal" role-based package such as SAP.

4 The role of ethnography

We have discussed in earlier sections how notions such as 'role' and 'responsibility' are much more complex than they may appear at first sight. The question therefore arises – how can we understand the real responsibilities in complex socio-technical settings?

The conventional approach to this problem is to start with a formal organisation chart, identify the individuals associated with roles and interview them to discover the type and nature of their responsibilities. However, while such an approach can be a starting point, we believe that it suffers from serious flaws:

1. Organisation charts are usually formal, over-simplified views of complex organisations. The reality is inevitably more complex and messy.
2. Responsibilities are often implicit rather than explicit – people find it difficult to articulate what they really do.
3. Responsibilities are contingent and dynamic – people take on responsibilities depending on the particular tasks to be done. This is fundamental to the effective functioning of most organisations – it is only in a 'jobsworth' culture that

it is uncommon. Hence, again, articulation of these responsibilities is difficult.

An approach which we have used with some success to understand responsibilities in complex organisations is ethnography [5, 6], where an experienced social scientist spends a period of time observing the ways in which work is done, the dynamic division of labour in a particular setting and the ways in which the artefacts and the physical organisation of a setting influence the work carried out.

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.

The 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. In this case, this narrative is the means through which we can develop an understanding of the real responsibilities in an organisation..

Our motivation for exploring the use of ethnography to understand work was to better understand the real requirements of users of complex IT systems such as those used to support air traffic control, hospital patient administration, etc. Our contention was that ethnographic studies would reveal how people 'worked around' the problems with computer systems and developed coping behaviour when things went wrong. Over the last 10 years or so, we have carried out ethnographies in a range of settings from air traffic control rooms, through financial institutions to steelworks [7, 8, 9, 10].

A universal characteristic of all of the sites that we have studies is that one of the major problems that arise in the use of complex IT systems is that these systems often include an implicit model of responsibilities which:

(a) may not be configurable for each specific setting where the system is deployed;

(b) rarely if ever copes with the dynamic and contingent nature of responsibilities;

(c) does not recognise the critical distinction between causal and consequential responsibility and, hence, makes invalid assumptions about how work is actually carried out.

To illustrate the nature of responsibility as discovered through our ethnographic studies, we will draw on two examples from recent ethnographies that we have carried out in hospitals (11, 12). For each of these examples, we will highlight how commonly adopted approaches to information systems design can result in real problems for system users.

The first example we will use concerns the administration of chemotherapy treatment of patients suffering from cancer. The cocktail of drugs which is administered to each patient is prescribed by the oncology consultant who has the (consequential) responsibility for treating that patient. He or she must arrange for the appropriate drugs to be available for each chemotherapy session. The (causal) responsibility for the treatment session may, of course, be devolved to a more junior doctor. Maintaining patient records was, however, the responsibility of the nurse assisting with the chemotherapy.

In practice, hospital consultants are very busy people and they often forgot to order the required chemotherapy medication for particular patients. If this was not available, treatment session had to be cancelled and re-scheduled – a distressing experience for patients who were often very ill. To avoid the problem, the clinic staff had devised a work-around – the day before a patient was due, a nurse checked if the required drugs had been ordered. If not, with the connivance of the staff in the hospital pharmacy, he or she placed an unsigned order that was then replaced in the pharmacy with a signed prescription whenever the consultant was available.

Our studies have shown that such coping behaviour is extremely common in complex socio-technical systems and is fundamental to the dependability of these systems. In this situation, the essential problem was a problem of responsibility. There was an assumption that the consequential responsibility of the consultants for prescribing the medication was equated with the causal responsibility of physically drawing up a signed prescription.

Now consider what might happen if the process was automated. A secure system here would associated prescribing permission with a consultant and not with a nurse. It would be difficult for a nurse to repair the problem of forgotten orders. Workarounds, which are not strictly according to the rules, would be more difficult (although, in practice, human ingenuity is such that we are confident they would be discovered).

Now let is consider another example from the same domain where responsibilities are complex in practice and where this complexity only became obvious after a period of ethnographic study. We have already alluded to this example in the introduction to the paper.

When patients arrive in a hospital for in-patient treatment, they are required to be allocated to a bed. Bed management is complex as large hospitals have a constant stream of admissions and discharges as well as planned routine surgery and emergency treatments. In general, patients should not have to wait more than a few hours in a holding area before being assigned to a bed in a hospital ward.

Within the hospital, there is an administrative role of bed manager who has the (consequential) responsibility of ensuring that incoming patients are assigned to beds and who makes use of a bed database to monitor the current availability of beds. The (causal) responsibility of actually finding a bed and updating this database falls on the nursing staff in wards who change the status of beds when patients enter or leave the ward. Medical staff also feel that they have an ethical responsibility to patients to ensure that they receive

proper and timely treatment. We have not discussed ethical responsibility, in situations such as hospitals, it is important. It as the notion of 'doing the right thing' rather than working according to some job specification.

This situation works well enough most of the time. Problems arise, however, when the hospital is close to capacity. When there are several beds available, the accuracy of the information on the beds database is not critical. The database may report that there are 6 beds available when there are actually 5 or 7 but so long as a request for a bed can be satisfied, then it doesn't matter.

However, when the database reports that there are no beds available, then problems arise. Patients may have to wait longer than an acceptable time in a holding area. In order to discharge her consequential responsibility, the bed manager must, at that stage, take over the causal responsibility of finding beds for incoming patients. At this stage, she does not trust the data from the bed database and takes action to discover the situation on the ground rather than in the database.

This may involve calling round wards to discover if patients are shortly to be discharged, re-negotiating planned admissions or, in extreme cases, walking around the hospital to see if any beds are free. Once a bed is discovered, the responsibility can then be discharged by the bed manager.

It may appear that this problem is one that could be solved by technical means. If the bed database was accurate, then this problem would not arise. Our ethnographies showed that this was not in fact the case – they revealed the reasons why the inaccuracies arose (these were not generally technical faults) and, again, responsibilities had a prominent role to play.

Inaccuracies in the bed database arose from two primary causes:

1. Nurses did not update the information about bed availability in a timely way.
2. Nurses deliberately failed to update the information when beds became available.

In the first case, failure to update the database was usually a consequence of conflicting responsibilities. Nurses, primarily, have responsibilities for patient care and, generally, they see these as their prime responsibilities. In the time between a patient being discharged from a ward and the database update, patients often required care from the nurse and this distracted him or her from the database update. Updates may have been completely forgotten or significantly delayed.

This situation is predictable and understandable. The second situation, where the database was deliberately not updated generally arose because of a conflict between the responsibility to update the system and the ethical responsibility of 'doing the right thing' for patients.

A general strategy used by the hospital to make beds available was to postpone and reschedule planned surgery where this was non-critical. Therefore, an elderly patient scheduled to have a knee joint replaced might have their operation cancelled because the bed that the hospital planned to use for them had been assigned to some other patient with a shorter term (although not necessarily less urgent) demand.

Nursing staff were, of course, aware of planned surgery and often knew the patients concerned. While their surgery may have been routine, it was important to them and very distressing to have this cancelled. In some situations, the nurses used their judgement to pre-allocate beds to these patients or to keep a bed for a patient who had been discharged. This ensured that when the patients with planned operations came to the hospital, a bed would be available.

Here, the nurses were making judgements about which of their responsibilities should take priority and coming down on the side of ethical responsibility. The bed manager, of course, knew of such practices and hence used the walk around the hospital to discover beds that might be used.

Again, we have a situation where the complexity of the real responsibilities was only discoverable through ethnography – interviewing would never have revealed what really went on in the hospital. Again, we have an example of how responsibilities affect the operation of complex IT systems. Again, the IT system had a preconceived model of responsibilities which did not match the reality on the ground.

5 Summary and conclusions

Finally, we give a number of examples of how these considerations impact system design.

(1) There are a number of implications arising from the requirements of the holders of consequential responsibility.

> *Audit trail*: It is unusual for information systems to have the capability to record things that happen in the social domain, such as delegation. It is in the interest of agents who hold consequential responsibility that the audit trail is correct, complete, and easily reconstructable. They should perhaps have causal responsibility for this being the case.
>
> *Agency*: Confusion between individual and role, and between authentication and authorisation, is far too common and a frequent cause of dissatisfaction.
>
> *Answerability*: One possible organisational rule is that consequential responsibility following a technical malfunction rests with the person who chose to deploy the malfunctioning device. This answerability could be mitigated by providing evidence about the soundness of the choice process.
>
> *Organisational boundaries*: Where consequential responsibility is unclear, the social and legal processes require more than just information immediately prior to the triggering event. The nature of the contract between the parties may have implications for existing (or non-existing) systems.

(2) The importance of conversations means that what looks at first sight like a database system may turn out to be a communication system. This depends on the nature of the

conversations involved. If the conversation is equally significant or both parties and there is a high degree of mutuality, then communication structures may be more important than information structures.

(3) Looking at the way that conversations may fail is an important, and often neglected, part of system design. Recovery procedures must be considered at the outset. As we indicated earlier, failures in the communication medium and failures in the instruments of record, can often be managed through technology; but failures in the structure of the conversation based on the responsibility relationship between the parties is more subtle and complex. As a very simple example, if I wish to reconnect to a particular individual in a call centre because the advice I was given didn't work and I wish to apprise the individual of that fact and receive alternative advice, whether, and the way in which, the system facilitates this is crucial. The point for the designers to consider is how the system maintains continuity of the responsibility the centre has for my benefit from a single conversation distributed across multiple episodes, and how the system supports my responsibility for advising individuals in the centre about the incorrectness of their advice.

(4) A frequent source of IT system problems arises because IT system designers do not understand the complex model of responsibilities in large socio-technical systems. We argue that understanding this model is difficult and that our notions of causal and consequential responsibility are useful in explaining the real allocation of responsibilities. Our experience is that ethnographic field studies are an effective way of developing a 'rich picture' of the responsibilities in an organisation and help us build a model of responsibility 'as is' rather than 'as planned'.

Acknowledgements

We wish to thank Mike Martin for a number of useful discussions over the years and Karen Clarke for telling us about her observation of bed management in a hospital.

This work has been partially supported by the EPSRC through the DIRC Interdisciplinary Research Collaboration [www.dirc.org.uk].

References

[1] B. Dahlbom, L. Mathiassen (1993), *Computers in Context*, Blackwell, (1993).

[2] R. Hirschheim, H. K. Klein, K. Lyytinen, *Information Systems Development and Data Modelling*. Cambridge University Press, (1995).

[3] H. Jonas, *The Imperative of Responsibility*, University of Chicago Press, (1984).

[4] J. Lucas, *Responsibility*, Oxford University Press, (1995).

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

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

[7] Sommerville, I., Bentley, R., Rodden, T., Sawyer, P., Hughes, J., Randall, D. and Shapiro, D. 1992. 'Ethnographically-informed systems design for air traffic control'. Proc. CSCW92., ACM Press, Toronto, 123-129.

[8] Harper, R., Randall, D. and Rouncefield, M. 2000 Organisational Change in Retail Banking: An Ethnographic Perspective. Routledge, London.

[9] Martin, D. and Rouncefield, M. (2003). Making the Organisation Come Alive: Talking Through and about the technology in remote banking. *Human-Computer Interaction*, Vol. 18, No's 1 & 2, pp. 111-148.

[10] Hughes, J. Martin, D., Rouncefield, M. Sommerville, I., Hartswood, M., Procter, R., Slack, R., Voss, A. 2003. 'Dependable Red-Hot Action'. Proc. 8th European Conference on CSCW, September 2003, pp 61-80

[11] Clarke, K., Harstwood, M., Procter, R., Rouncefield, M., Slack, R. And Williams, R.(2002) Improving 'Knife to Skin Time': Process Modelling and New Technology in Medical Work. *Health Informatics Journal*, 8(1). Sheffield Academic Press, p. 41-44, 2002.

[12] Clarke, K., Hartswood, M., Procter, R., Rouncefield, M. And Slack, R. "Minus nine beds": Some Practical Problems of Integrating and Interpreting Information Technology in a Hospital Trust. In Bryant, J. (Ed.) *Proceedings of the BCS Conference on Healthcare Computing*, Harrogate, 2002. pp219-225.