Who and what are electronic patient records for? An ethnomethdological ethnography of system deployment in the NHS

David Martin
Computing Department
Lancaster University, UK
d.b.martin@lancaster.ac.uk

Abstract
This paper reports on an ethnomethodological ethnography of a project to implement an electronic patient records (EPR) system in an NHS Trust in the North of England. The paper provides an introduction and reflection on the use of ethnographic studies for the purposes of computer system design and deployment through a discussion of the study. However, the paper also makes an more unusual turn for such work in computing by discussing the political import of the study, particularly focusing on policies related to accountability and governance and their impact on producing useful and workable EPR systems to support medical and care work.

Introduction
Electronic patient records (EPR) systems, currently being deployed within hospital Trusts across England and Wales are intended to be the infrastructural cornerstone of a modern 21st century patient-centred health service. It is envisaged that these integrated, Trust-wide systems will support medical practitioners in delivering care by ensuring better processes, supporting best practice, assisting in decision making, providing direct, timely access to patient information, and being a gateway to knowledge repositories. Such is the rhetoric that surrounds the current Connecting for Health (CfH) programme in the NHS, through which it is envisaged that all hospitals will be provided with fully functioning EPR systems, that will also integrate with GPs and a national database 'the NHS Spine' collecting various patient and medical information from across the NHS. It is envisaged that health records will be universally accessible and that GPs will be able to, for example, book hospital appointments and procedures for patients from their surgeries.

However, as with many of these technocratic visions there appear to be greater problems in actually delivering a functioning technology. Recently (early July 2006) there have been a spate of articles in the UK press [e.g. The Observer, ComputerWeekly.com, 2006] reporting on difficulties with the CfH programme. A report by the National Audit Office (NOA) had stated that the cost of the project had roughly doubled to £12.4bn and that it was currently running behind schedule. Although the NHS IT director, Richard Granger disputes the seriousness of the issue there is clearly cause for concern. Much of the media coverage has focused on GP access but it is also clear that hospital Trusts are also experiencing difficulties implementing the new EPR systems successfully. So far only around 10% of hospitals have currently deployed EPR systems, with the crucial fact that nearly all of these systems only deliver administrative support. Now, while better administrative support (if it is indeed the case) may assist medical staff and may provide benefits to the patients, it is clear that the actual functionality most heralded to benefit the patients and clinicians, such on-line medical records and prescribing, up-to-date details of best treatments and best practice support, access to knowledge repositories, epidemiological information and so forth are still some way from delivery.

Against this backdrop we want to report on a field study conducted between 2003 and 2005 at an NHS Trust in the North of England. The choice of case study is interesting because the Trust in question was an 'early adopter'. Before the current version of the NHS IT programme had been announced, the Trust had already signed a contract with an Anglo-American software firm – hereafter 'OurComp' – for them to implement and support a full blown EPR system. The NHS gave them the go-ahead, and so the study provided an opportunity to investigate what some of the issues might be when the larger scale deployment of such systems across England and Wales got underway.

We have already reported on this study in a number of articles in the domains of computing and healthcare informatics [e.g. Martin et al. 2004b, 2005, 2006b, forthcoming] where the primary aim is in some way to assist in producing 'better' computer systems. After introducing the setting and discussing the type of study undertaken – an ethnomethdological ethnography – which we will do in the next section, we want to provide a brief introduction to the use of ethnographies in computing. We believe

this may be interesting to the audience, as we believe the audience may come from different disciplines that usually employ ethnography in a different manner, aimed at understanding different phenomena and answering different questions. We believe it is also relevant for our current task in hand, however.

It became clear to us during the study that the design and the implementation of the EPR system was a perspicuous site for viewing some of the tensions and conflicts currently within the NHS. Design decisions about best supporting hospital work could only be made if they fitted with other organisational and regulatory concerns - most clearly the concerns of:

(1) Providing an accountable and standardised service, as construed by the NHS, and
(2) Providing a computerised, integrated system.

Here was a clear case where in a complex design situation, the problem of providing computer support for current medical and health practices was being made more difficult by the need to report on activities in particular ways and to standardise processes through integration. In other words, the ambitions of the EPR to provide future health benefits through the collection of standardised health data and through regulation (e.g. mass collection of data will aid the evaluation of doctors, hospitals, treatments etc., and information on patients will be available universally) produced sets of requirements that might compromise the ability of workers to deliver care successfully at the local point of care, *now*. This is not meant to suggest that there is no reason for or point to these ambitions but it is important to draw attention to the fact that there may be reasonable disputes as to whether these 'future' benefits will be as great as heralded, and it is clearly most important that current medical and health practice is not adversely affected by the implementation of new computer systems. This suggested that the research had more obvious relevance to NHS policy, particularly the relative prioritisation of different requirements, as will be discussed in relation to the findings of the ethnography.

Setting

In order to illuminate on some of the issues relating to the design and deployment of EPR systems across the NHS we report on an ethnographic study of such a project at an NHS Trust in the North of England. The setting for this research was an NHS Hospital Trust (comprising 2 hospitals) that is currently in Phase 2 of a three phase comprehensive £8.3 million EPR project, delivered as a public private partnership (PPP) between the Trust and a commercial supplier – 'OurComp'. Phase 1 went live in February 2005 (after a one year delay) and comprises the core administrative system and connected reporting system, A & E, theatres, order communications, and pathology systems, and is designed to be integrated with existing legacy systems[1]. The project entailed configuring a generic commercial-off-the-shelf (COTS) system originally designed for the US healthcare market to fit the requirements of the NHS and this particular Trust. COTS systems offer the possibility of a ready-made design solution, which can also be 'tweaked' to fit individual customer requirements. However, there is a big question as to how ready-made the solution is, how much it will have to be tweaked and how 'tweakable' it is?[2] This creates a key set of implementation problems concerning how to configure the system to fit with current endogenous work practices, existing legacy systems and new technologies, and NHS standardisation and regulatory requirements. This involves working out and deciding how the COTS system can and should be configured, and reciprocally, how existing patterns of work practice should be maintained or transformed.

Method

The research reported here employed ethnographic (observational field study) methods [see Crabtree 2003; Hughes et al. 1992, 1994; Martin & Sommerville 2004a for elaborations of how these studies can be used for computing purposes], 'shadowing' the project team as they went about their everyday work as well as observing project meetings of various kinds over a period of 18 months, and the 'go-live' of the system in February 2005. We collected a wealth of materials (field notes, tape recordings, transcriptions, and documents) and analysed them with an ethnomethodological orientation (see

[1] It is only in phases 2 and 3, however, that more obviously 'medical' features of the system will come on-line. For example, electronic medical records, drug prescribing, GP access and care pathways (processes incorporating best practice information).

[2] Due to the costs of developing and maintaining a 'generic' system (i.e. code base), COTS systems are typically not 'openly' configurable but exhibit a 'designed for' configurability that anticipates the ways in which the settings in which they might be deployed can vary.

[Button and Sharrock 1994, 1996] for comparable studies of system design). 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. Consequently, here our focus is on the everyday work, documenting how and in what ways implementation of the system was managed; how were requirements conceived, reasoned about, prioritised, worked upon and so on? We would suggest that while some of our findings are specific to this Trust, issues related to implementing COTS systems and the relative prioritisation of requirements are pertinent to the deployment of EPR systems in the NHS since most of these projects will be undertaken in similar situations and have a similar configuration of players and technologies involved.

Ethnomethodological Ethnographies and Computing

Before delivering the material and the findings we would like to contextualise the study within the tradition of using ethnomethodological studies in computing. Ethnomethodological ethnographies have recently found a home in computer science because (1) computing has increasingly focused on social interaction, (2) computer scientists and designers often wish to consider action and interaction in as much detail as possible when considering design to support those activities, and (3) ethnomethodology is seen as apolitical, not having a theoretical axe to grind (except maybe with other social scientists), instead simply describing activities as they are organised and oriented to by the participants in that setting.

The traditional role of the ethnographer in computing is that of ethnography *for* design. In this conception the ethnographer's role is generally seen as providing a service, producing rich descriptions and analyses of work or sets of activities that computer scientists wish to design to support. Examples of this type of work are numerous in the computing literature particularly in the disciplines of human-computer interaction (HCI) and computer supported cooperative work (CSCW) both of which focus on the importance of taking 'human' and 'social' factors into account in computer application and system design.

In general ethnomethodological ethnographies are seen as useful, firstly, because traditional software engineering (SE) techniques tend to focus narrowly on data and data flows while ignoring the activities of the workers in dealing with the data and making it meaningful. And, secondly, the models of human cognition that have traditionally driven interaction design are seen as based too heavily on abstracted inferences about the cognitions of single users rather than being rooted in descriptions of people interacting with each other and systems in practice. Lucy Suchman [1987] is often credited with the seminal work in this area. She used an ethnomethdological analysis of interaction with a photocopier to illustrate that the psychological models of cognition instantiated in its interface miss-matched the embodied reasoning users produced, in an on-going fashion, socially, in relation to the interface and the photocopier itself.

After Suchman, a number of computer scientists and software engineers began realising that ethnomethodological ethnographies could provide just the type of information they had been missing about the application domain; details of the people and the way they organised their work in relation to artefacts and computer systems. At least to some extent they wanted extra detail on the activities they were designing for, rather than higher level theory or politically motivated work, and this was simply a matter of pragmatic usefulness rather than a taking of sides in any social scientific debate. Following on from Suchman a famous example of this work was that undertaken in relation to air traffic control (ATC) by researchers at Lancaster University [e.g. Hughes et al. 1992, 1994]. Examples of other projects along this line are Crabtree's [2003] work in a shipping office and our work in designing an Internet banking application [Martin et al. 1998].

As might be expected, accompanying this work there have been a number of articles considering the role of ethnography for design: when should it be undertaken in the design process, what sort of requirements can and should it produce, how should it be presented, what is the role of the ethnographer, and so forth? In seeking to answer these questions ethnographers have drawn on their experiences in the 'field' but have also drawn on a second set of studies; ethnographies *of* design, which are basically studies of design projects. By understanding real life design it has been thought that ethnography might provide those involved with a form of 'reflection on practice' that might aid them in their work, it was also thought that this might provide ethnographers with further insights into how their work could best be incorporated into design projects. The work presented here is of this type.

So, to re-cap, the use of ethnomethodological ethnographies in design has predominantly followed a model whereby it serves as a rich informational resource about the activities that the system is being designed for. Or, relatedly, it provides an evaluative analysis of a technology in use (i.e. does it serve to facilitate or constrain the work of its users). In a lesser way it has also been about understanding system design itself and self-referentially about the place of ethnographic and ethnomethodological work within design. As such, with both forms of this research it tends to shy away from political issues in design. However, it is based on the position that design should try and ensure that the resultant system does not seriously disrupt important activities of the users unless this is the point of the design, or there are other crucial reasons for doing this. In the case of CfH within the NHS, the mantra of those in charge is that the new EPR systems are primarily for two groups of 'users' – health workers and patients. Consequently, successfully supporting work at the 'sharp end' should be a priority. Our ethnography, as detailed below, allowed us to consider this question.

Findings

In this section we want to focus on three long excerpts from our fieldwork. Through these examples of regular 'types' of phenomena, drawn from various design meetings and activities, I wish to demonstrate how, no matter the good intentions of those involved in the design, the desire to put the patient first, and to support the healthcare practitioner in doing their job end up being regularly de-prioritised as requirements in relation to those requirements that pertain to the bureaucracy of modern healthcare and moves towards standardisation. This is particularly obviously manifest in the design of a system in which every action is meant to be recorded in real-time so it can be counted and accounted for even when doing so according to restrictive system models has disastrous consequences on everyday work when the system goes live.

Example 1: Prioritising and Producing the 'Correct' Reports

In this example we have taken selected quotes from a longer discussion between the Trust project manager 'Helen', and 'Gail', a Trust analyst, that took place at a project developers meeting. Gail is reporting back on the findings of a project she has been given to see whether and how the new system can produce the reports they need to hand to the NHS. Gail opens by stating that the person she has been speaking to has requested some 'real-time' reports.

Gail (Trust analyst)[3] –I spoke to XXX yesterday and obviously they've produced an awful lot of reports mainly from 'focus'[4], and she's hoping that some of those will be real time reports of she's going to look through those y'know and see whether we could, turn those into real time reports

The system being configured for the Trust is a dual system. The EPR system (Patient 1) works in real-time but it is attached to a relational database 'Decision 1' (see below) that takes a sweep of data off the main system every night. Decision 1 is meant to be the system for producing reports but it is clearly not a real-time system. As Helen then states, taking reports directly off Patient One:

Helen (Trust Project Manager) - Yeah, it just seems I mean it seems so quite a drain on the system

Helen and Gail then go on to discuss just how up-to-date the reports need to be and how many of them would be required, before Helen then addresses the developers as a team to emphasise the importance of getting the reports right:

Helen – "I did meet with XXX yesterday to discuss some of the issues ... because the reports we hand into the NHS are crucial to our funding, as a Trust and obviously we have to get the reporting right and there's a huge risk to the Trust because we're going live six weeks before the end of year, and … all of our end of year reports we have to make sure are right."

[3] In each example we present the staff involved with their organisational role in parentheses in their first conversational turn

[4] 'Focus' was the old reporting system used by the Trust

Helen then continues to discuss how in her prioritisation of reporting she is getting the importance through to the supplier ('I think its hitting home to them'), and that it is crucial that they bring the testing of the reporting forward to between August and January, before an anticipated go-live in mid February.

Helen - I think its hitting home with them too, the issue of the reports we've been saying it so much how crucial this is to us, so we've tried to incorporate sooner looking at um testing for Decision One reports and making sure that they're accurate so that's going to start as soon as in August really we're going to start looking at issues around that and use that time between August and January to really make sure they are right

Gail then raises the problem that they will have a split CDS (commissioning data set - the set of data items required by the NHS)[5]. The NHS requires reporting on the 1st of April , which may raise the problem of having data on two systems. Helen responds by saying that the data should have been brought across onto Decision One:

Gail – 'Cos we'll have a split CDS as well, will have six weeks at the end of the full financial year for the HES[6] deadline as well so that's going to be a bit of an issue.
Helen – Yeah so we'll have to have a look at that but again any data that we're converting that is in Patient One and then thus Decision One we'll be able to report from Decision One on that historical data if once it's in, Decision One we should be able to report on it.

This excerpt raises a number of issues. Firstly we can see that 'correct' reporting is a clear priority in the design of the system. This is stated and shown by the fact that extra, early testing has been scheduled with the supplier being brought on-board. However, we can also see that a desire for real-time reporting needs to be tempered and de-prioritised so as not to have detrimental effects on the performance of the live system. This is also clearly crucial as the primary feature of the live system is that it should always be available and efficient in supporting on-going tasks. Another risk is the implementation of a new system just prior to reporting – the developers need to know that they can successfully manage the 'split data set' by having a successful process to transfer and match data items from the old to the new systems.

Generally in our data we found that it was clear that reporting gained a very high priority. This was because producing both timely, correct and 'favourable' reports were crucial for maintaining good levels of funding for the Trust. Bad reports and late reports would be bad for the future of the Trust. We are not suggesting that these are not important features of such systems but we might question the relative importance that they inevitably gain – an importance that leads to their potential prioritisation over other requirements and other work. That Trusts need to employ increasing numbers of administrators and bureaucrats in such situations and that more time and money is spent on these concerns is no surprise given their integral relationship to funding, but we might question whether this is a good thing.

Example 2: Usability and Security in Accident and Emergency (A & E)

Again in this example we have taken selected quotes and passages of talk from a longer discussion. The discussion arose when the laboratory systems developer (Lenny) raised a concern about how long the system gave you before a log-in timed out[7]. This is taken up as an important point by another trust developer – Pete – who raises the point that this might be a crucial issue in A & E, where the job is unpredictable and time is important. The issue is agreed by Helen who suggests a need to monitor it when demonstrating to users:

[5] The NHS commissioning data sets (CDS) are nationally stipulated sets of individual data items each with a clear label, definition and set of permissible values (codes and classifications), that are used to report on healthcare practice.

[6] Hospital Episode Statistics - the national statistical data warehouse for England of the care provided by NHS hospitals

[7] Currently there are no specifically defined standards regarding time outs.

Pete (Trust Developer) – You you're saying speed of sign on but my concern is that the only way it's going to be able to work is that you have to completely come out of Patient One y'know all the way and then log back on and re-load Patient One now that will take too long in somewhere like A & E … Cos' if it takes twenty seconds for Patient One to load up y'know and if a Doctor signs off and then they have to sign back on again and wait for Patient One to load up I mean, that might be too slow
Helen (Trust Project Manager) – Well again that's why we need to look at this and uh when they're doing the demonstrations so we can figure out what all the issues are around that it's good to flag that now

The issue is then taken up by Lenny (see below), who raises the issue of whether they could use a 'physical device' for logging on and off. This brings Neil (in charge of authentication/security) into the fray but he discounts the suggestion of physical devices straight off. This in turn brings in Bob (the A & E developer) who looks for clarification of what the procedure will be. This leads Neil to an explanation of why biometrics have not been used before Bob raises a second problem – not only will there be potential problems due to the time taken to load the application, but also one may end up with people working on the system under the previous user's log on.

Lenny (Laboratory Systems Developer) – I was just going to say that if we use some sort of device like that I think they've mentioned the button type device or a key card, that you have to physically attach or physically insert in a workstation to be able to log onto the network and then, y'know onto Patient One
Neil (Security Manager) – We're not going to do that
Lenny – I mean are you moving away from actual, physical devices inserted
Neil – Yeah
Bob (A & E analyst) – So it's a log on net ID and password
Neil – Yeah and then the registry biometrics, when it's affordable. There were issues with the eye ball the ones we looked at didn't integrate with active directories, it's a too primitive a device the ones that would be required to solve those issues was far too expensive. Originally we were led to believe that they would be about a dollar, per unit, but the more sophisticated iris devices are significantly more than just a dollar, the price we got quoted is quite expensive
Bob – That with certainly in A & E it leaves their system well not as secure as they'd like it to be
… because if they've got log out people will not log out of it they don't now
Pete – But maybe they won't have a chance because the log in time out will
Bob – Well I understand that but if it times out but if it doesn't time out before someone gets their hands on the keyboard, that next action is taking place under someone else's signature… And that's a problem

The developers then continue to discuss how the security and usability issues might be solved but at this stage we have reached a bit of an impasse. Problems and concerns have been raised and possible solutions suggested, but no easy one found. This then prompts Helen (below) to suggest a different tack, turn it into a training issue, which clearly attempts to turn the issue from a technical problem to a user's problem. Bob, however, vehemently rejects this as a possibility:

Helen –Yesterday I met with the IT trainers and we started talking about some of the issues that we need to make sure that everyone is aware of and one of them, you know is this issue now we'll add that to the list that this is one of the key ones making sure that people log out and understanding the implications because in a fact it's an electronic signature, and that's going to give a print, of where you've been on the system and if you don't log out you're allowing someone else to use that that signature
Bob – But it's not a training issue Heather
Helen – Mm
Bob – The fact is that the log out procedure will not be looked upon as important as treating a patient
Helen – Yeah so we need to look at it I think I agree it's not completely a training issue I do think it is partially a training issue
Bob – Well I understand that, yes … Confidentiality is also another issue you know another subject we talk about but that's not the priority

This example illustrates a common way in which issues are raised and discussed during design. Initially Lenny raises a question about the system 'time out'. How will the authentication and time-out security features work in A & E? Can security be achieved while practice is not disrupted in that environment? Ensuring security may undermine work practices but if the security setting is too lax, so as to not keep timing users out security may be jeopardised if the terminal is in a (semi) public place, which may allow unauthorised access. Or, if another member of staff works under the wrong signature the accountability and auditability of the system may be compromised. If the problem cannot be solved, if the right balance cannot be achieved technically, can training achieve the solution?

The excerpt shown provides a good example of the complexity of design and the types of conflicting requirements that the project team has to deal with. Clearly security, accountability and auditability are important aspects of the system, but having a system that can be worked with easily in A & E is also important. Medical staff may need to get pulled away from the system at a moment's notice but they will not want to lose whatever work they have been doing, and if it takes them a long time for the system start up again when they authenticate, in this environment it will be problematic. In computing we would suggest that the system should be designed for 'interruptibility', however, here we see that the project manager moves to turn the issue into a training issue. We noticed (as will be seen in example 3 as well) that this proved to be the 'easy' solution to difficult problems, simply because it was a way of deferring the solution, in the light of no obvious technical solution being available.

Example 3: Supporting Local Practice versus Standardisation

Again in this example we take selected quotes and conversational exchanges from a longer period of activity. This time we have taken the material from a testing session conducted by Vic (Senior OurComp Developer), Brad (OurComp Tester) and Helen (Trust Project Manager) with A & E users. The excerpt begins with Jenny (an A & E nurse) asking whether the system has some of the functionality she requires:

Jenny (A & E Nurse) – "Can we see a day's schedule… can we tell who's had x-rays.. how do we change an appointment".

Next Jenny suggests (see below) that she needs some practice with the system, as it's the *first time* she's seen a clinic screen. Helen responds by stating that another part of the reason for the design is to 'fit in with the Trust', i.e. for the purposes of integrating processes in a standard fashion across the Trust. Brian (another A & E Nurse) responds by stating that meeting the demands of this are seen as a problem when it means extra effort by local users. It is clear here that the requirement for standardised procedures and for NHS concurrence is prioritised over work practice support or usability. Helen promises placatory future efforts to 'streamline' things; however Jenny persists in describing her concerns with the new system:

Jenny – "This is the first time I've seen a clinic, before they've never been working so I'll need to go back and practice it."
Helen (Trust Project Manager) – "You need to fit in with the Trust that's why it's like this."
Brian (A & E Nurse) – "But it's a problem that fitting in with the Trust involves more work."
Helen – "Anything we can streamline we will… in the future with OurComp… and you have to realise the importance of data gathering and sharing information across the Trust."
Jenny – "I've been trying registration for months and have a problem of getting lost and not knowing where I am and I'm worried about how much training for our receptionists will be required."
Vic (Senior OurComp Developer) – "Could you drive [control the computer] and show us where you are getting lost?"

Jenny (above) notes that even though she has been practising 'registration for months' she still has difficulties, and these involve 'getting lost' on the system which suggests the proposed training for receptionists may be insufficient. This triggers a discussion regarding the user interfaces and interaction sequences required by the new and old systems. For Jenny and Brian the new system is difficult to learn, requiring a different and more complicated pattern of interaction making it easier to get lost. Finally, as shown below, Vic and Helen reiterate their comments about the need for organisational and systems integration, and that the information is required by the Trust:

Helen – "This is a Trust wide system, you get the benefits of the information gathering of other people so you need to do this….As a teaching hospital we need to do research so we need good data…since there are no A & E people on the PAS team I'll now put you on as stuff like this is a PAS requirement so it will help you to understand and keep informed of decisions."
Vic – "If a patient is sent to A & E from elsewhere you won't need to fill in these details as they will have been done elsewhere so you do get benefits."

As a 'Trust wide' system, the extra information gathered is often of benefit elsewhere, and since the hospital is a teaching hospital (required to do research) it needs 'good data'. Furthermore, users in any particular department will receive benefits from others as well as doing extra work to benefit others. However, as we can see below Jenny re-focuses on the current screen, an action which serves in some

manner as a rejection of Helen's attempts to 'bracket off' the usability issue by inviting Jenny onto the PAS team. She notes that patients are not tagged with their 'presenting complaint' and gives an example of why she would need this information – so she can readily identify patients in need of quick attention – the *'patient…with chest complaints'*. In her further comment – *'we do this now'* we can hear an appeal to the designers that the system should at least provide the resources for action that the current system does:

Jenny – "Speaking as an A & E nurse I need to know what the patient has come in with to triage…. The presenting complaint… I need to know this… it should be a mandatory field so I can look down the list and identify the patient out of 12 with chest complaints… we do this now."
[Brad suggests you can do this with the system and moves across to show them on the computer… "you can get the information from the whiteboard."]
Brian – "We don't currently work from the whiteboard, it doesn't fit with our workflow[8]…. We have a separate triage list which we can view and re-order the patients on."
Vic – "We can't change this as the screen is hardwired."
Brian – "But we currently prioritise using the triage list… it's a fundamental facility…this really worries me …There's a few things like this that worry me from having gone through the workflows."

This provides a good example of competing requirements. For the users the priority is for a usable system that supports important work activities, which it does not appear to. For the developers, (1) integration and standardisation, and (2) requirements to collect 'accurate' data of particular types for research and reporting are the key criteria to hit. It does not seem to be in dispute that it is desirable to hit all of these criteria, however, given that they must be relatively prioritised or traded off the main dispute seems to be that the developers are arguing that the system can be made usable and can be made to support work practices through training and later system development. Conversely, the users seek to emphasise the importance of the criteria they are promoting and they suggest that training may well not produce the positive outcomes that the developers suggest.

System 'Go-Live'

Our concerns for the success of the system grew during the on-going study. They were many – the system was late, there was a high staff turnover, analysts complained about difficulties in getting the information they required, there were some tensions between the US and the UK based analysts. However, with such projects problems are expected, as are delays. However, from the material presented here I want to draw out the pertinent concerns that were visible.

1. Building NHS processes, databases and reporting mechanisms to fit with NHS requirements were prioritised over other development work.
2. As other requirements impinged on how well the system fitted with work practices, the other requirements were prioritised and problems of fit with work practices were construed as training issues (and not that serious) and to be dealt with later.
3. Related to 1 and 2, the system had a very strict 'model of governance' built into it, i.e. security levels were strictly defined and implemented (who could do what, who had authorisation to change security clearances). Furthermore, the models of process offered little flexibility (processes had to begin in the right place with the right user and proceed according to the authorised path, recorded in real-time, thus offering little possibility for delay, recovery from error or exception handling).

When the system eventually went live we were lucky enough to observe the first week in the hospital. Unfortunately our concerns were proved valid – the system proved to have multiple problems. Most notably it did not fit well with a number of existing work patterns but to compound this, due it its strict model of governance, when work did not fit with the system the system broke down. On the day of 'go-live' I sat in pathology while the samples built up in reception. If a sample had not been logged on the computer by the correct person, the lab staff could not process the sample using the system, and they could not correct the system information. Levels of security authorisation were not nuanced enough, and the pathology systems manager was not authorised to change the settings. Exceptions in terms of where the samples came from or in terms of what they were (i.e. unusual in some way) could not be processed on the system. The environment became chaotic. The systems people in pathology were fire fighting, continually dealing with multiple problems at once. In the next few days as I went on various ward rounds with IT people it was clear that the situation was repeated across the Trust, and that the

[8] By 'workflow' Brian (an A & E nurse) means 'the way we currently organise our work', but he characterises this in the jargon of design.

situation was compounded by the on-going remedial work on the system that meant it was out of actions for considerable periods of time, and was in a state of flux whereby users were regularly confronted with a differently organised system when they logged on upon the start of their shifts.

The unfortunate postscript to this is that after the first week our access was denied as the drama unfolded. It was not particularly unfortunate for us as our project has gone well. However, we do feel very sorry for the members of the project team as it appears that the blame fell as some of their doors even though it was clear from our fieldwork that the system had not failed through lack of skill or effort but rather it had been bound to fail because of the massive ambition for the 'EPR' in general and the way the NHS has conceived of the requirements for such systems.

Discussion

The National Health Service (NHS) in England and Wales is currently undergoing a period of 'modernization' through computerisation. The CfH programme is meant to be the means through which 'joined up' and 'seamless' healthcare will be realised through healthcare information integration and process standardisation. Bloomfield and Verdubakis [1997] argue this 'modernisation' has been going on in different guises since the 1980s, and with CfH we see it continuing. In the 1980s and in to the 1990s the conception of what electronic heal records could and would provide gradually grew in scope and ambition from integrated hospital administration systems to the current situation, where:

[1] The systems are envisaged to support multiple and varied medical, administrative, reporting and regulatory processes, rather than just providing administrative support. Essentially, the heterogeneity of healthcare information and practice pose particular problems for integration
[2] The technologies incorporated in EPRs are more complex and sophisticated – e.g. utilizing imaging and visualization technologies, decision support and Internet-based applications, and are supported by innovative technologies such as wireless and mobile devices.
[3] Requirements as dictated by the NHS are emerging, developing and changing as successive programs are put into place and defined, meaning that system development needs to be flexible and responsive to future requirements.

An intriguing feature of the situation is that the ambitions for computerisation in the NHS keep getting more extravagant even though the delivery on previous ambitions has – to be slightly generous – only been partial. Indeed, at a recent project workshop we held in Edinburgh earlier this year it was interesting to hear that although the CfH project is in its early days and having problems 'technocratic' dreams of the next generation of systems are already running riot. As one author stated:

"In addition to this [a fully realised CfH programme], complete accurate and large volumes of data on patients, their history, their pathology and their outcomes will build a hugely powerful knowledge tool. The IHR [Integrated Health Records] will be one of the largest medical data repositories in the world. As is anticipated, if in the future this information also integrates with individual's genetic maps we should be on the verge of taking the next step in the creation of the next generation of tailored medical treatments. This will allow for the creation of drug types and treatments for the individual and not the general population as happens now. Creating medicines that will work for you but not necessarily your wife or neighbour. These new medicines could bring about significant increases in the effectiveness of treatment." [Seivwright, 2006]

Such comments cannot help but provoke a 'whoah, slow down' response from us. Without even going into the possibilities of genetically personalised treatments, it is clear that providing integrated care records solutions that work pretty well, with reasonable quality of data and that are reasonably up-to-date would be a better ambition to aim for. The big bang approach, and the desire to enforce standardisation through the incorporation of strict regulatory models in the system, and the dominance of accountability and reporting functions within the systems clearly goes against providing systems that fit well with current working practices at the sharp end of medical care. We acknowledge that it is understandable and necessary (for the purposes of integration) that some form of standardisation needs to occur and, furthermore, that systems always impact practices, practices can be improved and there is a need for transparency. However, we want to argue against an overemphasis of these priorities because of the resources they take up, but also it is arguable whether they can deliver. For instance, bad or incomplete data will always be a feature of systems, as will mistakes, and the need for recovery from error and exception handling. These features should be accounted for in design rather than attempts made to eradicate them.

To finish I would like to recount and analyse a brief discussion during one of the project meetings. Team members talked over the fact that when the 4 hour time limit for processing patients in A & E had come in they had been concerned about hitting the targets. Their way round the problem was to represent the process differently. Patients who were nearing their 4 hour limit, and had some chance they might be admitted onto a ward, were placed in a 'pre-admission' room and no longer categorised as A & E patients. This, of course, eased the figures. Now, a public reaction to this might be to condemn the hospital but this was deemed acceptable by the NHS. The clear lesson, here, is that enforcing the need to produce 'figures' does not sort everything out. Figures do not speak for themselves, they mask a whole lot of work going on in the background, the same category can stand for quite a diverse set of activities behind it [cf. Garfinkel 1967 on medical records] and notions such as *quality* are harder to locate. Furthermore, by placing such an emphasis on 'getting the figures' right, a key concern for hospitals is now to employ more administrators with the requisite skills to help them achieve this, and to implement the computer systems that support this. All of this leads to a de-prioritisation of providing systems that fit well with endogenous practices.

Conclusion

Most people would agree with the need for some transparency in healthcare. However, it is little wonder that now when sophisticated bureaucracy is crucial to hospital funding, that increasingly money is spent on this area of hospital working, most clearly here in the design of EPR systems. This study has reported findings to demonstrate the detrimental effect this has on the design of systems as they simply do not fit well with the way work is currently organised. Consequently, this suggests that if EPR systems are to better support the best care for patients their design should be re-oriented away from quantitative statistical gathering to qualitative support for healthcare practitioners.

Acknowledgements

A special thanks to all the people at 'The Trust' who assisted with this work. This work was completed with the assistance of EPSRC grant no. CSD 7680.

References

[1] Bloomfield, B. & Vurdubakis, T. (1997). Visions of Organization and Organizations of Vision. Accounting, Organizations and Society 22, 7, 639-668.

[2] Button G, Sharrock W. (1994) Occasioned practices in the work of software engineers. In Requirements Engineering: Social & Technical Issues, (Eds.) Jirotka M. and Goguen J. London: Academic Press, 1994.

[3] Button G, Sharrock W. (1996). Project work: the organisation of collaborative design and development in software engineering. Journal of Computer-Supported Cooperative Work (JCSCW) 1996; 5 (4): 369-386. Springer-Verlag.

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

[5] ComputerWeekly.com (2006). Why audit is needed more than ever. http://www.computerweekly.com/Articles/2006/06/06/216266/Why+audit+is+needed+more+than+ever.htm

[6] Dousish, P. (2006). Implications for Design. In Proceedings of CHI 2006, April 22-27, Montreal, Canada. pp. 541-550. © ACM.

[7] Hughes J, Randall D, Shapiro D (1992) Faltering from ethnography to design. Proceedings of ACM CSCW '92, Conference on Computer-Supported Cooperative Work, ACM, pp. 115–122

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

[9] Hughes J, King V, Rodden T, Andersen H. Moving Out of the Control Room: Ethnography in system design. In Proceedings of CSCW 94. 1994. © ACM press.

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

[11] Martin, D., Wastell, D. & Bowers, J. (1998) Ethnographically Informed Systems Design (EISD): The development of an Internet-based electronic banking application. Proceedings of European Conference on Information Systems (ECIS), June 1998.

[12] Martin, D. & Sommerville, I. (2004a). Patterns Of Cooperative Interaction: Linking Ethnomethodology and Design. ACM Transactions on Human-Computer Interaction. Vol. 11 Issue 1, pp. 59-89. © Association of Computing Machinery.

[13] Martin, D., Mariani, J. & Rouncefield, M. (2004b). Implementing a HIS Project: Everyday Features and Practicalities of NHS Project Work. Health Informatics Journal. Vol 10 (4): 303-313. SAGE Publications, London.
[14] Martin, D., Rouncefield, M., O'Neill, J., Hartswood, M. and Randall, D. (2005). Timing in the art of integration: 'That's How The Bastille Got Stormed'. In Proceedings of ACM Group '05, November 6-9th, Florida, USA, pp.313-322. © ACM
[15] Martin, D. & Sommerville, I. (2006a) Ethnography and the Social Structure of Work. In., Jones, C., Gacek, C. and Besnard, D (Eds.) Structuring Computer-Based Systems For Dependability. Springer-Verlag.
[16] Martin, D., Mariani, J. and Rouncefield, M. (2006b). Managing Integration Work in a NHS Electronic Patient Records (EPR) Project. Health Informatics Journal.
[17] Martin, D., Hartswood, M., Slack, R. and Voss, A. (forthcoming). Achieving Dependability in the Configuration, Integration and Testing of Healthcare Technologies. To appear in the Journal of Computer Supported Cooperative Work, Special Issue: CSCW and Dependable Healthcare Systems. Springer, Dordrecht, NL.
[18] Suchman LA (1987) Plans and situated actions: The problem of human-machine communications. Cambridge, UK: Cambridge University Press
[19] The Observer (2006). Fears over faults in NHS patient records system. http://observer.guardian.co.uk/uk_news/story/0,,1805437,00.html