From Grid to Healthgrid
T. Solomonides et al. (Eds.)
IOS Press, 2005

Working IT out in e-Science: Experiences of Requirements Capture in a HealthGrid Project

Mark HARTSWOOD[a,1], Marina JIROTKA[b], Rob PROCTER[a], Roger SLACK[a], Alex VOSS[a] and Sharon LLOYD[b]
[a]*School of Informatics, University of Edinburgh*
[b]*Oxford University Computing Laboratory*

Abstract. This paper reports on our experiences of being involved in requirements capture for a HealthGrid project. Large scale, collaborative projects with multiple partners tend to experience numerous problems in the requirements capture phase (and often beyond) and HealthGrid projects are no exception. Projects with highly innovative objectives often have additional sets of problematics, however. In carving out new visions of, for example, clinical research and healthcare service delivery, HealthGrid projects have to reckon with – and work within – existing healthcare policy, legislative frameworks, professional cultures and organisational politics as well as the more common integration probkem of dealing with legacy systems. Such factors are not conducive to the achievement in healthcare of the e-Science vision of seamless integration of information and collaborative working across administrative, professional and organisational boundaries. In this paper, we document some of the challenges we encountered in investigating the requirements for eDiaMoND, a flagship pilot UK e-Science project. We discuss what we might learn from these challenges, especially approaches to requirements capture that are appropriate for projects with innovative aims and are also sensitive to representing and addressing what may be complex professional and organisational interests.

Introduction

Modern information infrastructures, capable of delivering high levels of integration, are increasingly seen as the key to the achievement of improved standards in healthcare services [1]. For example, so-called 'joined-up' healthcare envisages services being delivered to patients through flexible – and perhaps virtual – organisational structures formed around networks of healthcare professionals working within, and across, multiple service units and administrative domains. Similarly, translational medical research focuses on the integration of bench and clinical research for the benefit of patients. The aim of the latter is to reduce the turn-around time in the cycle that leads from identification of possible causes of illness (for example, particular genes or environmental factors) through their investigation, the investigation of disease mechanisms and development of treatments to clinical trials and practice.

[1]Correspondence to: Mark Hartswood, Social Informatics Cluster, School of Informatics, University of Edinburgh, 1 Buccleuch Place, Edinburgh EH8 9LW.

Grid technologies offer a potential way to address the infrastructure challenges posed by these and other areas of healthcare, and the UK e-Science research programme has invested heavily in projects aimed at developing innovative healthcare applications. eDiaMoND is one of these projects, its aim is to demonstrate the value of the Grid to the detection and treatment of breast cancer and to research into genetic factors. We present in this paper our experiences of investigating the requirements for such projects and discuss some ways in which their various challenges may be met.

An important part of the e-Science vision is collaboration. For the purposes of reviewing our experiences of the eDiaMoND project, it is useful to distinguish two separate, though interrelated, threads of collaboration that our experiences of eDiaMoND lead us to suggest that they constitutive of many HealthGrid projects. The first is concerned with the development and deployment of technologies, architectures and infrastructures for HealthGrid applications, and the second with the conduct of clinical practice or research itself.

Before we explore the eDiaMoND project in detail, it is useful to understand some of the characteristics that make the healthcare sector distinctive and challenging domain for Grid projects. There are a number of reasons for this. Healthcare services are organisationally complex, each with a bewildering array of different (and often incompatible) IT systems. Healthcare organisations themselves exhibit further complexities related to numbers of distinct roles and processes, and the richness and interrelatedness of healthcare information. Their information exchange practices and IT provision are typically rooted in local work processes as well as wider patterns of coordination and communication. The past ten years has been a period of often radical redesign of healthcare services which, in turn, has created difficulties in providing IT support that can evolve to match professional and organisational changes. Attempts to use IT as a vehicle for driving changes in practices, redefining roles, relationships, etc., may lead to resistance, if those involved have different commitments and understandings of professional cultures, organisational processes and service provision. Issues relating to different commitments, cultures and perceptions are further compounded in relation to integration across traditional service boundaries, for example, between primary and secondary healthcare sectors [2]. Finally, ambitions for deploying innovative IT solutions must be tempered by the responsibilities of having to provide a dependable service to patients.

The collective consequence of these factors is that even modest levels of system and information integration have proved difficult to achieve in practice in healthcare [3] and the sector's record in IT innovation is distinctly patchy. The UK National Health Service has, for example, launched several programmes over the last 10 years to develop an integrated care record [4].

1. The eDiaMoND Project

The UK National Health Service (NHS) runs a Breast Screening Programme (BSP) where women between the ages of 50 to 64 are invited for screening every three years at a local breast screening unit (BSU). The screening test is by mammography, where one or more X-Ray films of each breast are taken and examined for signs of abnormality by trained film readers. If suspicious features are evident, then the woman is called for further tests at an assessment clinic.

eDiaMoND is an ambitious, 2 year flagship pilot UK e-Science project with a budget of £4.1M funded through EPSRC/DTI and IBM SUR grants, with a project

team of academic and industrial collaborators over 12 sites. To deliver its aims of furthering breast cancer detection, treatment and research, the eDiaMoND project is building a grid-enabled, federated database of annotated, digitised mammograms and patient information, which is being piloted at four participating UK BSUs [5]. A number of prototype grid-enabled applications have also been developed to demonstrate how the database architecture can support screening work, radiologist training [6] and epidemiological studies. The database will also be used for research into image analysis algorithms, search mechanisms for data mining and image standardisation techniques. The eDiaMoND database embodies the e-Science vision of creating a potential for new clinical research by enabling data generated at screening to be shared between a variety of disciplines (radiology, training, epidemiology and breast imaging research) and made available to the whole breast care community, regardless of where or how it was generated.

Given the problems noted above in regard to healthcare IT projects, the requirements capture exercise undertaken for eDiaMoND was designed for robustness, relying upon a range of techniques. It drew, for example, on detailed, ethnographically-based [7–9] analyses of work practices in screening [10–13] and conducted further ethnographic studies of work practices in a number of BSUs [14]. The aim of these studies was to observe in detail everyday working practices and to explicate the numerous, situated ways in which those practices are actually achieved. This involves sometimes lengthy periods of fieldwork within the settings where the system will be deployed. Data collected included fieldworker notes of observations and discussions, and video of BSU staff engaged in routine work practices. Clinicians also participated in design meetings and discussions, intended both to elicit their views on the vision of eDiaMoND, and also to aid our understanding of the current process of breast screening. One focus for the requirements capture exercise was to understand the potential implications of transforming a mammogram into a digital artefact and what benefits this might afford for the BSP through innovations in its work practices. Once built, we also undertook quasi-naturalistic evaluations of prototype workstations with clinicians, in situ where possible [15].

We now turn to the specific issues which were raised by our experience of pursuing this requirements capture process.

2. Dealing with Uncertainty

If visionary projects are to be successful, change has to be dealt with, that is, it has to be rendered manageable, yet be given the space in which the vision can be worked out as a practical matter. Our experiences of requirements capture for eDiaMoND suggest that the challenges posed by change may be particularly problematic within HealthGrid projects. These challenges manifest themselves in two, interconnected, ways. The first of these derives from the clinical and research practices that the system is intended to support and the second stems from the technologies. Their combined effects were to cause the aims and expectations of those involved to shift a number of times as the project unfolded. This gave rise to considerable uncertainty about the nature and details of user requirements for eDiaMoND.

Taking clinical and research practices first, the concept of integrated healthcare services is new and, in many instances, are yet to be fully realised. As a consequence, ideas of how these practices might be supported by new infrastructures are themselves

still emerging. Simply put, the problem is that the potential users of the system don't really know what they want. This is not, in itself, a novel problem for requirements capture and, though it is by no means always convincingly resolved, there are a variety of requirements gathering techniques which may be deployed to address it. In particular, the need for genuine user involvement in IT systems design, coupled with iteration is now firmly entrenched as an ideal to be strived for, even though the practice may be subject to the usual project contingencies of lack of time and resources. It is, for example, now commonplace in requirements capture to begin with workshops in which users are invited to brainstorm requirements and to follow up these up with meetings where requirements engineers' ideas are presented as sketches or prototypes in order to have them evaluated, refined and validated. The difficulty is that, despite several iterations of the brainstorming-evaluate cycle and the apparent firming up of initial ideas, users' understanding of their requirements may be rather more illusory than it appears. This is seldom because of a lack of imagination or commitment on the part of users. Rather, it is because where innovations in work practices are the aim, experience suggests that users' requirements may only become concrete and detailed once the system is deployed and they get to experience first hand what the system can actually do [16]. We have frequently observed in other projects that the recognition of defects and deficiencies arises from trying to use a system in the context of doing the work [17]. When a member needs to 'get the job done' it is precisely then – when the options are foregrounded – that consideration will be given to the means of solving this problem, using these available resources. Particular artefacts and methods then become relevant to the members that were previously part of the unconsidered background of the workplace. The problem is made concrete and the contingencies associated with 'solving the problem' become recognisable. To this extent it is difficult to obtain details about requirements in the abstract in formal requirements gathering and prototyping exercises.

As we noted above, our approach to requirements capture for eDiaMoND drew on various established techniques, including user workshops and prototyping. These were supplemented by the less common (but increasing recognized as very valuable) technique of detailed studies of existing work practices. This gave us some basis (though far from perfect) for anticipating how work practices might actually change with the deployment of the eDiaMoND system in the sense that it enabled the identification of possible issues that should be addressed in the design of the eDiaMoND system. For example, it revealed potential problems underlying the vision of distributed screening relating to professional trust and judgment [18]. Nevertheless, uncertainty about user requirements was persistent in a number of areas. For example, users found it hard to specify the datasets they would need for the training tool, that is, for documenting the cases that would be used for training.

The second issue is the evolving, incomplete and sometimes unstable state of Grid technologies. The Grid is a complex set of interlocking standards and components which is still emerging and whose technical trajectory is sometimes uncertain. As yet, for example, there are few easily deployable and easily configurable components available 'off the shelf'. As a result, clinical partners may develop unrealistic expectations about what HealthGrid projects can deliver, typically expecting, for example, working systems rather than limited demonstrators. Where there is a strong project focus on delivering a working system, the choice may be made to employ conventional technologies rather than opt for the riskier strategy of implementing a Grid-based solution.

This uncertainty over requirements and technologies has had a number of consequences for the eDiaMoND project, not least of which has been the danger of that

technical and clinical research agendas become disengaged. In the unfolding of the e-DiaMoND project, we witnessed a series of re-adjustments to the technical ambition of the project, where the technical partners became increasingly focussed on solving a smaller number of technological problems, which clinically partners saw as having increasingly less relevance for themselves. We could suggest that the focus of the eDiaMoND project shifted away from addressing a series of clinical problems towards solving a few narrowly defined technical ones. While this is an understandable response to the unforeseeable contingencies associated with ambitious projects such as eDiaMoND, an important consideration is how the re-focussing can be managed so as not to exclude the interests and priorities of (particularly) clinical partners.

3. Managing User Expectations

Securing funds for HealthGrid research is done on the basis of the potential benefits to patients that will flow from the innovative clinical and research practices that lie at the heart of the HealthGrid vision. It may be the case that orienting to ambitious clinical or research impacts (for example, talk of the eDiaMoND project being 'rolled out' across the UK if successful) is one way of making a case for and securing the sorts of funding needed for HealthGrid projects. It is common and understandable in projects large and small for benefits to be 'talked up' from the outset as a way of securing user engagement.

In the case of HealthGrid projects, what is 'promised' (in one way or another) may be the delivery of a working system, but what will emerge is, perhaps (and for understandable reasons), more likely to be a 'test bed' or demonstrator. While this may be an important end in itself, there still remains the task of delivering on the vision. However, the radical transformation of practice envisaged by a project such as eDiaMoND is dependent not only on groundbreaking computer science research, but also on incrementally addressing more immediate and mundane problems of network infrastructures, usable applications, sustainability and so on. While there are interesting computer science problems at the heart of eDiaMoND, there are also more mundane difficulties such as the lack of a supportive network infrastructure (and so on) in the way of a 'for real' implementation.

4. Realising Clinical and Research Collaborations

The eDiaMoND project, as with the e-Science vision in general, aims to promote collaboration both within and across existing boundaries of healthcare practice. This calls for reaching agreement between – and possibly standardising the work practices of – the various clinical partners. This is very challenging where work practices have developed for very sound and practical reasons to fit local contingencies of the UK BSP [14]. In particular, e-Science projects in the life sciences encounter a number of common problems related to the linkage between clinical practice, audit and research. One such problem concerns the difficulties of sharing data gathered for a specific purpose, for example, as part of clinical practice, where the concerns of others wishing to make use of the data (for audit or research purposes) are not taken into account at the time of collection. Although there is a push to improve data quality throughout the health services and specifically to improve the quality of data for auditing, front-line staff will

inevitably and justifiably put the interests of care and of patients first. A lot of work is needed to repair this. Auditors routinely access various source of data, then combine and triangulate them to improve the quality of data that they extract. Similarly, researchers make use of data extraction forms designed specifically to capture the data needed for epidemiological studies and research nurses exercise considerable skill in ensuring that the data they gather is fit for the intended purposes.[2]

Further problems arise from the diverse IT infrastructures found within the health service, which are often at odds with the ambitions and the infrastructures envisaged by translational research and Health Grid projects. There are issues about access, technical compatibility, ownership and the sustainability of research-funded IT components. Different policies and cultures regarding IT management compound these problems. A number of cross-cutting concerns such as ethics and data protection, especially the assurance of valuable and sensitive study data, data linkage and quality control are also common to many studies. There are often no common frameworks in place that would allow research and clinical infrastructures to be linked, and even basic network connectivity is problematic as networks within the health services tend to be heavily policed and connections to, for example, university networks are not generally allowed.

e-Science projects are therefore in a position where they have to engage in one-off negotiations with heath service IT partners to achieve a resolution for individual projects, adding significantly to their start-up time as well as to the uncertainties referred to in the preceding section. Solutions put in place for individual studies often turn out to be highly specific and therefore fragile and idiosyncratic. A durable solution that allows e-Science to be undertaken in a permanent and sustainable basis within the healthcare sector has yet to be found. Addressing such problems requires not only a significant and sustained investment in IT infrastructure, but also that the work envisaged by e-Science becomes a core concern within the health services. This has as much to do with resources as it does with various organisational concerns and priorities, including career structures and data sharing policies.

5. Legal and Ethical Issues Surrounding the Requirements Capture Work

HealthGrids provide opportunities for collaborative working in healthcare systems both within the UK and across the world. Thus, not only has the eDiaMoND project had to understand the complexities of the volatile organisational structure of the NHS and the drive towards integrated healthcare services, we have also been required, as part of the requirements capture process, to seek out and manage the intricate web of policies, acts and organisations that govern the use of information originating from patient or case information. In as much as eDiaMoND was required to demonstrate the use of a grid-enabled digital mammography system, we had therefore to prove the HealthGrid concept by considering the use of real data in real breast screening units, hospitals and research environments.

In conforming to the ethical and legal practices surrounding patient data, during the initial phase of the project the eDiaMoND team had to ensure that their use of anonymised information from the four BSUs involved in the project satisfied the requirements under the Data Protection Act as well as gaining clearance from the Thames Val-

[2] Similar issues occur where human tissue is used for research. Standard processes in the pathology pathway may render material useless for research.

ley Ethics Committee. The situation was exacerbated by the need to use not only newly collected case data with explicit consent from the patients, but also archives of cases selected by BSUs for training purposes and for which it would have been impossible to seek retrospective consent from those patients to use the data. It has taken some eighteen months of a two year project to progress through this situation.

The ethical clearance required for the use of patient originated data for research requires individual projects to seek clearance either from a local ethics committee for research involving just a local site, or to a multi-site ethics committee for clearance to use data across many sites. Ethics clearance for the use of data in research projects is complex. At present, the process for gaining ethical clearance is extremely time-consuming. The initial eDiaMoND project has had to apply for Multi-Site Ethical Clearance (MREC) as well as Local Ethical Clearance (LREC). This enables the project to take data from the selected clinics in anonymised from and keep it for up to ten years for the purpose of the original ethics application. Any further use of this data would need further ethical clearance. It is the duty of the principal investigator to ensure the appropriate archiving of the data once the research has been completed, and to ensure that the data is kept securely at all times. For many patients, the project will have to seek explicit consent to use the data. The process of anonymisation means that we will be unable to delete records from the virtual database data once all identifying features have been removed. These issues also raise concerns over the use and management of medical research data in general. Often, valuable data is generated from innovative research, but the issue of ownership and what researchers are able to do with the data are confusing at best. The Medical Research Council (MRC) is currently embarking on a programme of research to explore this area. In addition, other projects have spawned from eDiaMoND specifically that aim to investigate issues of ownership.

Whilst these bureaucratic procedures for ethical approval were in motion, those parts of the requirements investigation that involved fieldwork (principally the ethnographic studies) were still subject to various ethical and legal processes. The NHS is not a single legal entity, but an amalgamation of over 400 legal entities. These legal entities comprise clinics, surgeries, governing bodies, trusts and regional centres which together provide the health service in the UK. The structure of these legal entities differs in each of the UK countries and within regions. However, within each of these legal entities, a Caldecott Guardian may be appointed whose role is to ensure that data is processed correctly and that the principles of the Data Protection Act are respected and acted upon. This process involves ensuring that any staff who require access to unanonymised patient data for their research or for processing to create anonymised data, are either NHS employees of that legal entity or are put onto honorary contracts to ensure that they apply the same duty of care as NHS clinicians and staff. In eDiaMoND, this has applied to researchers who are assisting with the data acquisition in the clinics as well as those who may come into contact with un-anonymised patient records through attendance at clinics for ethnographic studies. This process requires the project to have full ethical clearance before proceeding and will require staff on these contracts to supply a recent curriculum vita.

6. Methodological Challenges Revisited

We turn here to a more detailed examination of the methodological challenges with which we had to contend in the eDiaMoND project.

It is by no means unusual in IT projects to find that users often find it difficult to devote time to continuing and lengthy requirements capture activities – which often have to be fitted in and around their normal work obligations. The eDiaMoND users were no exception. Equally, requirements engineers may be under pressure to get requirements signed off so that system designers may begin their work. There may, therefore, be a temptation on the part of both users and requirements engineers to seek premature closure of the process. It would be a mistake, however, to assume that the problem with *a priori* approaches to requirements capture is simply lack of time, though that constraint was real enough in the context of the eDiaMoND project.

Part of the problem of dealing with uncertainty in requirements capture is related to the practices of those building the technologies and the sorts of technological solutions they are willing to adopt, or are familiar with. Partners responsible for the data collection component of the eDiaMoND project took the position that they should be given a completed set of requirements, a specification in effect, from which to develop the relevant tools. At the same time, it was proving difficult to achieve a consensus on a core data set for eDiaMoND. The developers' unwillingness to engage in the requirements process (other than to take receipt of its outcomes) meant that work on developing key parts of the eDiaMoND technologies was delayed. The reasons for their lack of willing might in part have been a way of managing their commitment to eDiaMoND in light of competing commercial pressures, or alternatively just representing a mode of working to which they were accustomed. Had the developers engaged with the process at an earlier stage, it would have been possible to have a looser coupling between the finalising of requirements and the production of the technology, allowing both processes to move forward in parallel to a greater degree. A closer and more active integration of the developers with the project as a whole would also have helped to have shaped the emerging requirements in ways that attended to pragmatic technological concerns, as well as perhaps engendering an awareness on the part of the developers of the sorts of problems being encountered (such as achieving consensus on data sets) that might indicate the appropriateness of certain sorts of technologies and methods.

A particular lesson here concerns the degree to which the resulting technologies are configurable. The eDiaMoND 'system' consists of a number of components (federated databases, data acquisition tools and screening applications) that have to mesh together around a representation of the eDiaMoND dataset. When, for example, the data acquisition tools were used in practice, it rapidly became apparent that additional fields or codes were required, or that the ones available were not appropriate. However, once the database schema was 'finalised', because of the dependencies between these components, it was also effectively fixed and it was not possible to make the desired changes. This was a source of discontent for clinical partners, who were concerned that the data gathered should reflect their practice and terminologies in order to be useful for what they saw as the important purposes of the exercise – to support clinical activities (for example, by producing a functional training tool). It was also a source of discontent for those responsible for providing the technologies, who in turn saw these demands as arising 'too late in the day' and attributable to a failure in the requirements process – particularly to a failure on the part of clinical partners to take advantage of the opportunities they were given to specify the data they deemed relevant. One of the ways this issue was dealt with was to orient to eDiaMoND as being concerned with the production of a prototype, where such anomalies might be expectable, but not have an impact demonstrating the *technical* feasibility of the system. The clinical impact of the

data quality could be repaired in subsequent follow-on projects that would have a greater focus on producing a working version.

The real problem, we argue, lies in the structuring of the IT systems design and development life cycle such that requirements capture and design are separated off from the deployment of the system. Emergent work practices and requirements which, by their nature, may only become evident as users attempt to apply the system to their work can never be captured and designed for in this way: user requirements that can only be identified in the context of, and through, use, are lost. The work setting is a key arena for innovation and the all-important 'domestication' of new technologies. We note, for example, how through processes such as 'learning by doing' and 'learning by interacting', users are able to experiment, share and appropriate the innovations of others, mobilising their collective resources to evolve systems, to continue 'design-in-use' [17]. In these ways, users do cope to some degree with the shortcomings of conventional IT requirements capture and design practice [19] but this nevertheless leaves outcomes very much to chance.

A potential solution is to be found in the methodology which we have called co-realisation [20], wherein the stress is on the creation of a project team without the rigid demarcations of requirements engineers, designers, developers and users. Co-realisation calls for IT professionals to shift the technical work of design and development into the users' workplace, if not completely, then at least routinely and over sustained periods of time. Through this process of committed being there, co-realisation's aim is to achieve a situation where project focus and effort can spontaneously shift between the different phases of the system lifecycle. Through this process there is a concerted and co-operative orientation to the realisation of requirements. The emphasis in co-realisation is on tightly coupled, 'lightweight' design, construction and evaluation techniques. Co-realisation seeks to bring about a context for IT design and development work where, as Buscher, Mogensen and Shapiro [20] have put it, "… effort shifts fairly smoothly between implementing or adjusting previously decided possibilities, picking up on the host of small problems that arise during work, coping with the unanticipated consequences of previous actions, talking to individuals …" (op. cit., p. 155).

7. Conclusions

HealthGrid projects have to contend with – and work within – existing healthcare policies and legislative frameworks for patient data security and confidentiality, professional cultures, organisational politics and a complex landscape of existing IT systems. Such factors are not conducive to the achievement of the e-Science vision of seamless integration of information and collaborative working across administrative, professional and organisational boundaries.

It is tempting see these as 'structural' problems to be remedied by appropriate legislation and guidance. We would argue, however, this is only part of the picture, and the difficulties faced by HealthGrid projects are actually fundamentally intertwined with the obligations, interests and concerns of participating organisations. In reality, collaboration is seldom done on a 'no strings attached' basis. Issues of data ownership, data protection, IPR, competitive advantage (both in the commercial and academic worlds), ethics, and other organisational and personal interests do not simply dissolve in the face of the collaborative ideals of e-Science, but rather they are thrown into sharp relief. We argue that the success of HealthGrid projects depends heavily on finding

practical ways by which various sorts of interests, rights and obligations incumbent on collaborating organisations and persons can be recognised and embedded within HealthGrids – that this is, in fact, necessary to make the vision a practical reality.

Where does this leave the idea of a HealthGrid? Well, on the one hand, the HealthGrid could be an important tool that might help to address some of the issues concerned, especially those about safe and secure sharing of data across boundaries. On the other, some of the issues make the implementation of HealthGrids problematic. The lack of infrastructure on the part of the health services means that there is a great imbalance between 21st century computing facilities available to researchers in Universities and in Industry and a clear lack of computing resources and research support available to clinical researchers and practitioners in the institutions such as the NHS.

HealthGrids can provide the opportunity for global, collaborative healthcare. The broad sweep of the legal and ethical issues reported here apply largely to the UK. Future investigations will need to focus on legal and ethical considerations in other countries and the complexities and challenges that arise when we attempt to share patient data worldwide. Regarding the UK specifically, there are two simple recommendations. First, it is clear that before proceeding with fieldwork, projects should examine the need for honorary contracts for team members who may come across confidential data as part of their investigations. It is crucial to conform to the data protection act and also to protect the fieldworker, the clinic and the patient data. Contracts have to be set up at specific hospitals or legal entities. Second, where projects intend to use real clinical data, the need for ethical approval must be considered and time allowed for ethical approval submissions. The fact that it has taken the eDiaMoND project eighteen months to secure MREC and LREC approval means that early considerations of these constraints is crucial. We could not begin to acquire and process data until complete ethical approval in place.

Finally, regarding requirements capture methods and the challenges of HealthGrids, our experiences of the eDiaMoND project suggest to us that there are important issues that need to be practically resolved at a number of levels. While some issues may be addressed by simply providing the required resources or channelling them in the right direction, others require the development of new ways of organising IT design and development work. Innovative technologies such as HealthGrids and visionary ways of working such as e-Science are not bounded in this way. They can not be 'inserted' or 'slotted' into a dynamic and complex socio-technical system, but are, rather, themselves dynamic and open in a way that requires their being 'grafted' into an existing (changing) socio-technical substrate, becoming a part of its dynamic – in positive, but also potentially negative ways. Co-realisation is a way of acknowledging the risks and costs of this process, it, so-to-speak, takes the 'bull by its horns'.

There is, then, no in-principle reason why the IT system design and development cycle should be organised in the 'traditional' manner and our experience suggests that there are substantial gains to be made in the adoption of the co-realisation approach. Importantly, for innovative projects like HealthGrids – and e-Science in general – co-realisation enables evolving user requirements and technologies to be accommodated, helping thereby to manage the uncertainties that are otherwise all to likely to arise. It avoids the polarisation of outcomes of technological interventions into either 'successes' or 'failures'. In contrast, we draw a picture of a process predicated on situated, practical reasoning, involving finding utility in planful assemblages of technologies where ambitions, work practices and explorations of technological limitations and af-

fordances jostle together, and are reflexively reshaped in order to accommodate one another: in other words, nothing more or less than working IT out in practice.

Acknowledgements

We wish to acknowledge the help and support of all the clinical partners in eDiaMoND, and in particular the contribution made by staff at the South East Scotland Breast Screening Service.

References

[1] Health Select Committee. (1999). Inquiry into the Relationship between Health and Social Services. London: House of Commons, January.
[2] Hartswood, M, Procter, R., Rouncefield, M. and Slack, R. (2003). Making a Case in Medical Work: Implications for the Electronic Medical Record. *Journal of Computer-Supported Cooperative Work*, 12(3), p. 241-66.
[3] Ellingsen, G. and Monterio, E. (2001). A patchwork planet: The heterogeneity of electronic patient record systems in hospitals. In Proceedings of the Information Systems Research Seminar in Scandinavia (IRIS'2000, Uddevalla, Sweden, August).
[4] Jones, M. Learning the lessons of history? Electronic records in the United Kingdom acute hospitals, 1988-2002. Health Informatics Journal, 2004 10: 253-263.
[5] Power, D., Politou, E., Slaymaker, M., Harris, S. and Simpson, A. (2004). A relational approach to the capture of DICOM files for Grid-enabled medical imaging databases. To appear in ACM Symposium on Applied Computing, Special track on Computer Applications in Healthcare.
[6] Soutter, J., Anderson, E., Campos, J.C., Hartswood, M., Khoo, L., Procter, R., Slack, R., Smart, L., Taylor, P. and Wilkinson, L. (2003). The Role of Computer Based Training in Mammography. Royal College of Radiologists Breast Group Annual Scientific Meeting, Cardiff, November.
[7] Anderson, R. (1994). Representations and requirements: The value of ethnography in systems design. Human-Computer Interaction, 9, pp. 151-182.
[8] Hughes, J., J. O'Brien, T Rodden, and M. Rouncefield (1997). Designing with Ethnography: A Presentation Framework for Design. Symposium on Designing Interactive Systems, Amsterdam.
[9] Heath, C.C. and Luff, P. (2000). Technology in Action. Cambridge University Press.
[10] Hartswood, M. and Procter, R. (2000). Designing for Breakdowns and Repairs in Collaborative Work Settings. In Fields, B. and Wright, P. (Eds.) Special issue on Understanding Work and Designing Artefacts, International Journal of Human Computer Studies, vol. 53(1), pp. 91-120.
[11] Hartswood, M., Procter, R., Rouncefield, M. and Slack, R. (2002). Performance Management in Breast Screening: A Case Study of Professional Vision and Ecologies of Practice. Journal of Cognition, Technology and Work, vol. 4(2), p. 91-1002.
[12] Hartswood, M., Procter, R., Rouncefield, M. and Slack, R. (2003). Cultures of Reading in Mammography. In Francis, D. and Hester, S. (Eds.), Orders of Ordinary Action: Respecifying Sociological Knowledge. Ashgate Publishing.
[13] Hartswood, M., Procter, R., Rouncefield, M., Slack, R. and Soutter, J. (2003). The Work of Reading Mammograms and the Implications for Computer-Aided Detection Systems. Proceedings of the Seventh Medical Image Understanding and Analysis Conference, Sheffield, July. British Machine Vision Association, pp. 89-92.
[14] Hinds, C and Coopmans, C. (2003). Screening Requirements. eDiaMoND Technical Report.
[15] Jirotka, M. Evaluation of the Screening Workstation. eDiaMoND Technical Report, 2003.
[16] Williams, R., Slack, R. and Stewart, J. Social Learning in Multimedia, Final Report to European Commission, DGXII TSER, University of Edinburgh, 2000.
[17] Hartswood, M., Procter, R., Rouchy, P., Rouncefield, M, Slack, R. and Voss, A. (2002). Co-realisation: Towards a Principled Synthesis of Ethnomethodology and Participatory Design. In Berg, M., Henriksen, D., Pors J. and Winthereik, B. (Eds.), special issue on Challenging Practice: Reflections on the Appropriateness of Fieldwork as Research Method in Information Systems Research, Scandinavian Journal of Information Systems, 14(2), p. 9-30.
[18] Jirotka, M., Procter, R., Hartswood, M., Slack, R., Simpson, A., Coopmans, C. and Hinds, C. Requirements for Collaboration and Trust in Healthcare Innovation: the eDiaMoND Case Study. Submitted to Journal of Computer Cooperative Work.

[19] Procter, R. and Williams, R. Beyond Design: Social Learning and CSCW – Some Lessons from Innovation Studies. In Shapiro D. et al. (eds.), The Design of CSCW and Groupware Systems. Elsevier Science, 1996, pp. 445-463.
[20] Buscher M., Mogensen P., and Shapiro D. Bricolage as a Software Culture. In I. Wagner (ed.) Proceedings of the COSTA4 Workshop on Software Cultures (Vienna, Dec 1996), Technical University of Vienna.