



Research article

A computer scientist's reactions to NPfIT

Brian Randell

School of Computing Science, University of Newcastle upon Tyne, Newcastle upon Tyne, UK

Correspondence:

B Randell, School of Computing Science, University of Newcastle upon Tyne, Newcastle upon Tyne, NE1 7RU, UK.

Tel: + 44 191 222 7923;

Fax: + 44 191 222 8232;

E-mail: Brian.Randell@ncl.ac.uk

Abstract

This paper contains a set of personal views relating to NHS Connecting for Health's National Programme for IT (NPfIT), and in particular its Care Records Service, written from the point of view of a computer scientist, not a medical informatics expert. The principal points made are as follows: *Centralisation*: Pulling lots of data together (for individual patients and then for large patient populations) harms safety and privacy – it is one by-product of excessive use of identification when in fact all that is usually needed is authentication. Large centralized data storage facilities can be useful for reliability, but risk exchanging lots of small failures for a lesser number of much larger failures. A much more decentralised approach to electronic patient record (EPR) data and its storage should be investigated. *Evolutionary acquisition*: Specifying, implementing, deploying and evaluating a sequence of ever more complete IT systems is the best way of ending up with well-accepted and well-trusted systems – especially when this process is controlled by the stakeholders who are most directly involved, rather than by some distant central bureaucracy. Thus authority as well as responsibility should be left with hospital and general practitioner trusts to acquire IT systems that suit their environments and priorities – subject to adherence to minimal interoperability constraints – and to use centralized services (e.g., for system support and back-up) as if and when they choose. *Socio-technical issues*: Ill-chosen imposed medical IT systems impede patient care, are resisted, result in lots of accidental faults, and lose user support and trust. All these points are attested to by rigorous studies involving expertise from the social sciences (psychology, ethnography, etc.) as well as by technical (medical and computer) experts – much more attention needs to be paid to such studies, and more such studies encouraged. *Constructive reviews*: A constructive expert review, working closely with Connecting for Health, could be very helpful, but should be evidently independent and open and thus essentially different in nature to past and current inquiries. A review of this nature could not just recommend appropriate changes of plan, and speed progress. It could also contribute to the vital task of helping to restore the trust and confidence of the public and the media in the programme and in the government officials involved.

Journal of Information Technology (2007) **22**, 222–234. doi:10.1057/palgrave.jit.2000106

Published online 24 July 2007

Keywords: electronic patient records; NHS; NPfIT; reliability; security

Introduction

This paper contains a set of personal views relating to the National Programme for IT (NPfIT), the NHS project which has claimed to be 'the world's biggest civil IT project' (Brennan, 2005). This claim is indeed believable, since the statistics are staggering – the 10-year project is intended to serve '40,000 GPs, 80,000 other doctors, 350,000 nurses, 300+ hospitals, 50m+ patients, and 1.344m healthcare workers' (Ferrar, 2006), with 'expenditure on the Programme expected to be £12.4 billion over ten years to 2013–14' (PAC, 2007).

It must be borne in mind that I am not myself a specialist in medical IT systems – I became interested in NPfIT in April 2006 when I was invited to add my signature to an open letter to the House of Commons Select Committee on Health calling for an inquiry into the programme's plans and progress. A brief investigation of a number of published articles and reports readily convinced me to sign. Since then, I have found myself spending a considerable amount of effort on tracking NPfIT, and assembling a dossier of published concerns and other relevant documents related to it.¹

Why am I doing this? The main reason is that I care deeply about the NHS. Without it I wouldn't be here today. I had an emergency triple cardiac bypass 7 years ago, at almost the exact age that my father died of a heart attack, and have nothing but praise for all the medical staff and organisations involved. (One of the same hospitals also undoubtedly saved the life of one of our sons, some years earlier.) I strongly believe they all deserve more and better IT facilities, so am very supportive of the general aims of NPfIT. But, along with my fellow signatories, I have become increasingly concerned at what I have been able to learn during this last 12 months, admittedly as an outsider, a mere computer scientist, about the directions and progress of this programme, and in particular about those aspects of NPfIT concerned with electronic patient records (EPRs), such as the National Care Records Service (NCRS) and the local NHS Care Records Services.

The background I bring to this is a long-term interest in system structuring, and in particular in system reliability and security, a subject that I took up shortly after I moved from a position at the IBM Research Center in the States to become Professor of Computing Science at the University of Newcastle upon Tyne. This was in 1969, not long after I had been involved in the first NATO Software Engineering conference (Naur and Randell, 1969). This conference led to an upsurge of research into formal development of provably correct software, but led me to wonder if it was possible to mitigate the effects of software faults that might remain in deployed systems despite such efforts. As a result, soon after joining Newcastle colleagues and I obtained a first research grant from the Science Research Council as it was then called, and so launched a still-continuing and growing programme of research concerned with computer system reliability and security. The emphasis over all these years has been on how to ensure that complex computer systems are adequately trustworthy even though they might (indeed will) suffer from faults. Such faults can be ones of design of both hardware and software, of hardware manufacture, of system operation and of careless use – including, and lately especially, what one might term 'malicious faults', arising from the activities of vandals, criminals and terrorists. Moreover, a further cause of possible failure, that is, another type of fault, is that of a project failing to keep up with changes to requirements during system development and/or deployment. (Thus my central interest is what is termed 'system fault tolerance'.)

In 2000 I helped to initiate the 6-year five-university Dependability Interdisciplinary Research Collaboration (DIRC), which was led by my Newcastle colleague Cliff Jones. This recently completed (and very successful) project involved computer scientists, psychologists, ethnographers, statisticians and others, in total some 85 researchers. Its subject was the reliability and security of large systems of *computer-based systems* (i.e., systems made up of computers *and* people) – much of the research in this project, in fact much more than originally planned, concerned healthcare systems, including EPR systems. Undoubtedly, my involvement in this project, from which I learned a great deal about the importance of socio-technical issues in system design, has coloured my attitudes to NPfIT, and hence the comments I make in this paper.

The April 2006 letter to the Health Select Committee, which in the end carried the signatures of 23 senior academics specialising in computing and systems, received far more publicity than any of us expected. Subsequently, I was one of seven signatories who accepted an invitation to meet Dr. Richard Granger and his senior staff to discuss our concerns. (Dr. Granger is the Director General of IT for the NHS and Chief Executive of Connecting for Health (CfH), which is the organization within the Department of Health that is responsible for NPfIT.) At this meeting, Dr. Granger agreed that a constructive open independent review of the type we had urged could be of benefit to NPfIT, and a press release to this effect was issued by CfH.

At this time, we also received a request from the Health Select Committee to provide them with suggested Terms of Reference for the proposed review – see Appendix. However, it was only in late 2006 that the Health Select Committee, reversing an earlier decision, decided to hold an inquiry into NPfIT. This inquiry in fact is concentrating on one crucial aspect of the NPfIT, namely its plans and provisions concerning EPRs, an inquiry that we hope will lead to a full technical review of the programme as a whole.

As a consequence of the written submission (HC, 2006; Ev164) that I made to this inquiry on behalf of the Group of 23 Signatories, I received an invitation to testify at one of the inquiry hearings. In order to prepare for this hearing I drafted an extensive set of notes, purely for my own use. However, I was subsequently urged by colleagues to produce a version of my notes for publication – hence this paper.

NPfIT and the current scepticism surrounding it

The National Programme for Information Technology in the NHS (the Programme) is a ten year programme which presents an unprecedented opportunity to use Information Technology (IT) to reform the way the NHS in England uses information, and hence to improve services and the quality of patient care. The core of the Programme will be the NHS Care Records Service, which will make relevant parts of a patient's clinical record available to whoever needs it to care for the patient... The Programme's scope, vision and complexity is wider and more extensive than any ongoing or planned healthcare IT programme in the world, and it represents the largest single IT investment in the UK to date (NAO, 2006).

The central vision of the Programme is the NHS Care Records Service, which is designed to replace local NHS computer systems with more modern integrated systems and make key elements of a patient's clinical record available electronically throughout England (e.g. NHS number, date of birth, name and address, allergies, adverse drug reactions and major treatments) so that it can be shared by all those needing to use it in the patient's care... The stakes are high. If it succeeds in its aims, the Programme could revolutionise the way the NHS in England uses information, and make significant improvements to the quality of patient care. But if it fails, it could set back IT developments in the NHS for years,

and divert money and staff time from front line patient services (PAC, 2007).

NPfIT is a 'system-of-systems' consisting of a set of interlinked systems, some provided at national level (i.e., for England as a whole), others provided by Local Service Providers (LSPs) who between them serve Primary and Secondary Care Trusts (i.e., hospitals and general practitioners) throughout five English geographical regions. The main national level systems are: the National Network for the NHS (N3), the National Data Spine (which forms the core of the National Care Records Service, and incorporates the Personal Demographics Service and the Secondary Uses Service), Choose and Book, NHSmail, the Electronic Prescription Service (EPS), the Picture Archiving and Communications System (PACS), and the Quality Management and Analysis System. LSPs provide data warehousing, applications and services for all the trusts in their respective regions, such as local NHS Care Records Services, and PACS solutions that connect to the National Data Spine.

There are many reasons why there is now much scepticism among the general public and in sections of the medical world about NPfIT:

- There have been so many occurrences of widely publicised failures of large IT systems in the UK and elsewhere, for example resulting in great cost and schedule over-runs or even complete abandonment, major reliability problems, or massive security lapses, for example, resulting in the leaking of literally millions of credit card numbers.
- There is the Information Commissioner's Report 'What Price Privacy' (ICP, 2006) indicating how easily reporters and inquiry agents can illicitly obtain information from the Police National Computer System and other allegedly secure government systems.
- Closer to home, there is the recent very embarrassing Medical Training Application Service (MTAS) blunder, in which confidential personal data from hundreds of junior doctors' job applications were made available on the Internet – one of the causes of the humiliating suspension of the service in May 2007 (e-Health, 2007a).

With regard to NPfIT itself, there have been numerous reported reliability problems. These include the 2-day outage at CSC's Maidstone data centre, and privacy failures (such as at the Melton, Rutland and Harborough PCT (e-Health, 2005), and at Leeds Teaching Hospitals NHS Trust (Collins, 2006)). There have also been well-publicised surveys of doctors expressing their concerns about the programme (e.g., Medix, 2006; e-Health, 2007b), the publicity over the Helen Wilkinson case (Evans, 2006), etc. And the NCRS, which many have regarded as the 'central vision of the Programme' through its intended provision of shared electronic patient clinical records, has very recently been found by the House of Commons Committee of Public Accounts to be 'already running two years behind schedule [since] the introduction of clinical as opposed to administrative software has scarcely begun; indeed, essential clinical software development has not been completed' (PAC, 2007).

Moreover, the way that officials and ministers have dealt with such problems has been quite counter-productive. A notable example is the initial demand by the Department of Health that doctors provide them with the names of people who wished to opt-out, and then the misleading official statement that assured the public that they could opt out of having their summary care records uploaded, without mentioning that they had no option regarding the uploading of their detailed care records (Anderson, 2006).

Achieving safety and maintaining privacy

Whether a medical IT system, such as an EPR system, is adequately safe or not depends both on how well the system requirements, both current and future, have been identified and specified, and how well the system meets these requirements, that is, how reliable and secure it is. However, identifying a medical information system's possible safety hazards (and comparing them to the hazards – which undoubtedly will also exist – of not using an IT system) is a matter for medical experts.

They will know how critical the accuracy and timeliness of particular information is likely to be to a clinician in a given situation, what would be the implications of a loss of, or a long-term unavailability of, say, an entire hospital's patient records, how dangerous could be confusions caused by badly designed computer display screens, etc. They should also be able to assess how such potential new dangers compare to existing dangers (such as the fact that paper records can be difficult to retrieve or even locate). These are matters outside my area of competence, as are issues such as whether there is greater likelihood of over-reliance on EPRs, as compared to paper records, merely because they come from a computer. However, the need for clarity and precision regarding safety and privacy requirements *is* within my competence.

Achieving IT system reliability and security in pursuit of safety, and providing guidelines as to what levels of reliability and security are achievable (i.e., what types and frequencies of failure should – unfortunately – have to be allowed and planned for) *are* issues for computer experts. (No complex IT system is completely reliable and secure – the so-called experts who claim otherwise should be shown the door.) Thus consideration of safety issues *must* involve the combined expertise of medical and IT specialists.

Like safety issues, privacy issues, such as the confidentiality of patient records, require the combined efforts of two types of expert – experts in the law, medical ethics, public policy, etc. on the one hand, and IT experts on the other, in order to determine the system requirements. The role of the IT experts will centre on explaining the possible privacy-related implications of both the chosen functional requirements placed on the IT system (e.g., whether the system has adequate provisions for protecting patient identities), and of the possible failures of the system (e.g., the likelihood of accidental or deliberate leaking of potentially compromising information about a celebrity and the possibility of sabotage by disaffected individuals). However, there is a valuable overlap in that there are some experts who have the necessary knowledge of what is likely to go wrong and what should not be able to, of data protection, etc., and also

of technical security mechanisms such as cryptography and access control.

It is vital to have a detailed specification (agreed by suppliers and users alike) not just of what a system is supposed to do, but how well it is supposed to do it. In other words, there needs to be a statement of the 'guaranteed' reliability and security levels of the various technical services to be provided to users by the NPfIT system, for example, concerning the possibility of confidential information being widely leaked from EPRs, of EPR data being corrupted or lost, or of access to EPRs being unavailable for unacceptably long periods.

System reliability and security specifications, like the functional specifications, will need agreed amendment from time to time. Such specifications are as necessary at each stage of an evolutionary system procurement process as they are in situations where the (usually misguided) aim is to produce a complete system specification *ab initio* (i.e., one that it is assumed will guide the rest of a huge and lengthy development project).

Such specifications are needed not just as part of contract negotiations, but in order that users can be informed and prepared for the incoming system, and their support for it gained. Indeed, in properly run projects well-written *agreed* specifications play a crucial role as the main channel of communication between the users and the developers, two communities that tend to speak quite different languages.

I was surprised to find, when I first started investigating NPfIT, that no such specifications were publicly available. I and my colleagues were then amazed to learn, from Dr. Richard Granger himself, that CfH themselves do not have detailed reliability and security specifications for the various major NPfIT systems (including those related to EPRs). This is because these specifications are regarded as confidential by the suppliers! (I cannot imagine the Ministry of Defence buys its aircraft without knowing beforehand how fast they are supposed to fly, or how often they will need servicing.) In the absence of such specifications, when all it has are technically vague and incomplete contracts, with 'agreements to agree', CfH is in a much weaker position vis-à-vis its suppliers than is generally admitted, and the users are not provided with the detail which would enable them to provide informed views.

However, reliability and security specifications alone are not sufficient. For any significant safety-critical system (which NPfIT most certainly is), the norm is to require that the system suppliers provide a 'safety case'. By this is meant a comprehensive and well-argued set of documents, providing credible evidence demonstrating that the agreed safety-related reliability and security specifications will be met, subject to identified (and agreed) fault assumptions. (For example, it might well be acceptable to assume that smart cards used for user authentication are in practice tamper-proof, or that cryptographic software that has been both formally verified and exposed to public assessment is indeed faultless. However, there will *always* be assumptions and it is very important that these are – as far as humanly possible – identified and agreed beforehand.)

In the absence of such safety cases and equivalent documents related to privacy protection, CfH staff have to put all their faith in the competence of the system suppliers,

and the user population their faith in CfH's ability to assess this competence. Moreover, both CfH and the users then have to await (and suffer) early and perhaps lengthy trials and several successive early deployments of the systems before they can gain any significant and justifiable confidence that the systems are likely to prove adequately reliable and secure. (This is quite apart from any questions of the systems' functionality, usability, etc.)

The nature of the challenges facing NPfIT

Like any large IT system, a large EPR system is a socio-technical system, not just a mere technical one. The extent to which it actually meets safety and privacy requirements depends on people, as well as on hardware and software. People (users, operators, etc.) can be both a cause of, and a means of preventing or minimising the impact of, safety- and privacy-compromising failures.

Reliability and security, and hence safety and privacy, are 'weakest link' properties. The larger the system, the more people involved, the easier it will be, for example, for an unscrupulous reporter or private investigator to find a weak link in the form of a legitimate user who can be fooled into committing, or bribed to commit, an act which will breach the system's privacy rules. (An experiment some years ago at the North Yorkshire Health Authority showed that about 30 phone calls were received each week attempting to trick staff into revealing confidential information (HC, 2006) Ev65. Incidentally, though such callers can be prosecuted should they be caught, lazy or corrupt insiders have little to fear from the law.)

On the other hand, doctors and nurses will, when they feel it necessary, find ways of undertaking their tasks *despite* the computer system. If security controls are too time-consuming they will be evaded, as exemplified by the recent case of smart-card sharing in response to excessively slow logging-in procedures at the South Warwickshire General Hospitals Trust (Collins, 2006). Another simple example, identified and studied by the DIRC project, was that of ward managers learning how to fool an ineffective automated bed management system into supporting a bed management policy that fitted their requirements (Clarke *et al.*, 2002).

Although intelligent and dedicated people are often driven – albeit at considerable inconvenience – to find workarounds that more or less adequately solve their immediate clinical needs, the wider consequences of those workarounds for overall system management can be adverse and unacceptable. Indeed, research studies have shown that all sorts of workarounds become the norm in situations where users have had what they regard as an unacceptable system thrust upon them from on high, resulting in an overall system that is quite dysfunctional from the viewpoint of the managers and bureaucrats, even if clinical needs are being met (Eason, 2006).

In fact there have been quite a number of careful scientific studies of healthcare IT systems, and in particular EPR systems, which have demonstrated the vital importance of socio-technical issues. One issue is the impact an EPR system can have on medical practice. To quote a

detailed scientific study of the working of Community Mental Health Teams (Hardstone *et al.*, 2004):

It would seem that integrated care records systems are, in the main, modelled along the same lines as airline reservation systems – always online, and always up to date. While this model may have its advantages in that it increases organisational control and enables strict auditing (what information was recorded in the system at a particular time and who had access to it), it fails to acknowledge and support the kinds of professional practices we have described.

Another study is of the introduction of IT into an A&E department, which involved replacing existing whiteboards with PC-based computer systems. It was found that, although the technology supported simple information requirements, complex coordination, collaboration and awareness issues were left unsupported (Broome and Adams, 2005). And a study in the context of UK psychiatric healthcare services revealed ‘important discrepancies between the assumptions of the role of the [EPR] and the ways that healthcare professionals actually use and communicate information within the particular work setting studied’ (Hartwood *et al.*, 2003).

In summary – it is sheer folly to specify and design a complex IT system such as a large EPR system with inadequate consultation and commitment from the various classes of people who will affect and be affected by it. Indeed there is an official international standard, ISO 13407, which provides guidance on achieving quality in use by incorporating user-centred design activities throughout the life cycle of interactive computer-based systems – I have seen no evidence that CfH requires its suppliers to adhere to this standard.

To quote another study (Adams *et al.*, 2005):

Traditional design and implementation approaches, isolated from communities, produce users – both clinicians and patients – who are either unaware of the technology or perceived it as complex and inappropriate for their needs. Random deployment of technology within communities, with poor design and support, is perceived by many as complex, inappropriate for their needs and a threat to current roles and practices, including the maintenance of clinician–patient relationships.

It is equally folly to assume that one can correctly determine the functional, reliability and security specifications of a complex system at the outset, then deliver the completed system in a standard form to numerous institutions and assume that people can be trained and convinced, that is, bullied, to use it ‘properly’. (A well-disciplined military unit might, just might, manage this but the NHS cannot.) It is also foolish to try to bulldoze unwilling staff into organisational change by imposing computer systems on them – recall the London Ambulance Service debacle (Finkelstein and Dowell, 1996):

This system was to supplant the existing manual system... The entire system descended into chaos (one ambulance arrived to find the patient dead and taken

away by undertakers, another ambulance answered a ‘stroke’ call after 11–5 hours after the patient had made their own way to hospital)... The Chief Executive of the LAS resigned... the implementation approach was ‘high risk’; inappropriate and unjustified assumptions were made during the specification process; there was a lack of consultation with users and clients in the development process with knock-on consequences for their ‘ownership’ of the resulting system; the poor fit of the system with the organisational structure of the ambulance service... there is a very strong message in the report about the attempt to change working practices through the specification, design and implementation of a computer system.

Let me quote two of the conclusions from the Official Report of the Inquiry into the London Ambulance Service (SWTRHA, 1993):

c) the Computer Aided Despatch system implemented in 1992 was over ambitious and was developed and implemented against an impossible timetable;

d) LAS management ignored or chose not to accept advice provided to it from many sources outside of the Service on the tightness of the timetable or the high risk of the comprehensive systems requirement.

Let me also quote two conclusions from a major multi-university study of EPR systems in practice (Proctor *et al.*, 2007):

The NHS has seriously underestimated the scale of the task involved in deploying EPRs. Constantly changing government and NHS policies has led to EPR procurement being very protracted: requirements have to be continually re-drawn and re-shaped and [this] often leads to unsatisfactory compromises. Procurement is also made problematic because these systems will be used as instruments of significant organisational change. However, the Trusts (and the NHS itself) do not have a concrete idea of what the results of those changes will lead to, consequently it is very difficult to assess system suitability.

Although ‘supporting medical practice’ and ‘patient centred’ are twin mantras of EPR design in the NHS, an over-riding design emphasis is on implementing ‘proper’ process, and on coding medical and administrative procedures ‘correctly’ so they may be standardised, counted and reported on. These ‘other’ requirements that stem from the need to provide fully technically and organisationally integrated systems can actually disrupt current medical practices. Standardisation implies that some features of local practice will be re-configured around new models that may run contrary to the way staff organise and understand their work; technical constraints can reduce flexibility. Since these ‘other’ requirements must be met, support for tried and tested local work routines may be removed with serious consequences later down the line.

As mentioned earlier, system specification, development and deployment should be gradual (‘evolutionary’)

activities, with adequate provisions for assessment, feedback and re-adjustment of plans – a well-known aphorism is ‘A complex system that works is invariably found to have evolved from a simple system that worked’ (Gall, 1975). Gradual deployment alone is not sufficient. The activities should start with relatively simple uncontroversial plans for systems that can be successfully deployed and can gain everyone’s support and trust. Indeed in general IT system specification, design and installation should be part of an overall business process re-engineering plan.

It is not surprising therefore that numerous voices are now arguing that it would have been much better if CfH had concentrated first on enabling individual trusts to acquire new or improved EPR systems that suited their particular circumstances and priorities (subject to minimal guidelines aimed at facilitating planned future interoperability). It could then have sensibly deferred attempting to construct nation-wide or LSP-wide EPR ‘systems-of-systems’ until the individual trusts’ systems were well established and accepted.

This approach is in line with the idea, now gaining ground, that the concept of *an* EPR should be abandoned in favour of that of using an ‘information broker’ to enable the accessing of information that is gathered as appropriate for each particular purpose from multiple specialised record repositories. Such a broker can be regarded as implementing what are in effect ‘virtual EPRs’. In fact, CfH already endorse one software system – Miquet (CfH, 2006) – that works this way, extracting data from different types of general medical practice computer systems, and this is the approach taken very successfully in Israel by Clalit Health services (‘the largest health organization in Israel and the second largest in the world’) (Gilon *et al.*, 2003). Moreover, the information broker approach has been shown to fit well with the use of modern system-building technologies such as web services and XML (Budgen *et al.*, 2007).

In identifying possible system requirements it is also vital to avoid placing excessive reliance on any complex critical system, especially one with demanding safety and privacy goals such as the National Care Records System. Some of the most successful complex systems owe their success not just to their high reliability and security, but also to the care that has been taken to minimise any absolute need for excessively high reliability or security – one very successful and well-documented example of this is the VISA international electronic payments system (Stearns, 2006).

Given that complex IT systems will invariably fail on occasion, it is critical in determining what services are to be provided by a system to consider how the surrounding organisation will manage to cope when the system fails (Schneier, 2000). For example, since EPRs will certainly be leaked and patient confidentiality breached on occasion, and possibly on a grand scale, it is vital to have procedures in place beforehand by means of which victims can gain prompt redress, and those responsible can be traced and penalised. (Unfortunately the Caldicott Guardian scheme (Department of Health (DoH), 2007), which might be assumed deals with such matters, is such that security breaches are not reported to patients, but only to the relevant Caldicott Guardians, and who I am told have much responsibility but little power.)

Similarly, there need to be robust plans about what to do if and when EPRs, or other critical patient records, become inaccessible for an undue period or data are lost (as has happened recently at the Nuffield Orthopaedic Centre (Bowers, 2006), and in Milton Keynes General Hospital (Gibson, 2007) for example).

The impact of centralisation and scale

There are great dangers to centralisation (whether physical or logical), and of ‘single points of failure’, as opposed to using large numbers of relatively isolated small systems, the failures of any one of which would have relatively limited effect. So a cardinal rule is to try to avoid having any ‘single points of (major) failure’. Instead, physically and logically distributed systems, employing carefully architected redundancy and diversity measures, are greatly to be preferred. (A quite different but also often important consideration is that the management of large centralized facilities tends to be much less responsive to the diverse needs of their various sets of ‘customers’ than are the managers of specialised local facilities.)

Centralisation (at LSP level, leave alone national levels) results in such massive databases, and user populations, that it inherently risks major (reliability or security) failures, from any of a variety of possible causes. Instead of having breaches of patient confidentiality, or loss of records, affecting just patients of individual GP surgeries, if these records are all pooled and hosted by an LSP a single accidental fault, or a single careless or malicious act, affecting a central server and its network and overwhelming whatever protective facilities have been provided could have a disastrous effect on all the surgeries and all their patients, in a whole region. For example, it has been reported that the recent failure at CSC’s Maidstone data centre left clinicians throughout the West Midlands and the North-West without access to their patients’ computer records for the entire 2-day outage (e-Health, 2007c). (It is worth pointing out that each LSP is intended to have responsibility for patient numbers that are comparable to the populations of various medium-sized European Union member states!)

This is not to say that all GPs and hospitals should be left to cope unaided with the responsibilities of maintaining their patients’ records reliably and securely, or that failures will not occur at the surgery or hospital level. Centralised back-up facilities, for example, could be very valuable, especially if the backed-up data are encrypted, and all the key holders can be trusted. Unfortunately, my colleagues and I have been unable to find any published discussions regarding NPfIT of such architectural issues, and of the trade-off decisions that are such a vital part of any competent large design activity.

The above points are technical ones concerning reliability and security – but they are being made against a background of there being significant controversy within the medical profession over the alleged clinical benefits of widespread access to either complete or summary EPRs. This provides an additional reason for questioning whether the evident technical risks of developing large (in effect highly centralised) ‘systems-of-systems’ tasked with

maintaining up to date widely accessible EPRs are worthwhile, even at LSP leave alone national level.

Concerns over the confidentiality of detailed patient records in NPfIT have led to the idea of centralising just what are termed 'summary care records' (though there is evident controversy concerning just what should be included in, and the clinical utility of, such summary care records), and to using a 'sealed envelope' scheme in order to allow patients more control over the use made of detailed and their summary records. However, a consultancy report – commissioned by CfH from Det Norske Veritas – determined that if a sealed envelope scheme were developed and used, it would be better employed locally, not as part of a huge centralised EPR repository (e-Health, 2006).

One further point – safety considerations indicate a need to design systems in such a way as to ensure (or at least to encourage) high data quality. The best way to do this is to arrange that EPRs be updated as an immediate by-product of clinical activities, so that these activities can directly benefit from such data capture, for example, through the immediate detection of prescription errors. (Equivalent practices are well established in other application areas.) In contrast, data that are collected afterwards and that are mainly used just for other purposes (e.g., summary care records) will never be of the same quality, or utility, because it will be of much less concern to the clinicians. (Note however that we are told that many senior staff still just will not use computers – touching keyboards is beneath them. If you insist that such a consultant uses the machine himself he will retire, emigrate or cheat.)

A critical aspect of the NPfIT system is that of 'identity management' – how information about users' and patients' identities is collected, maintained and used. In fact, in the computer world there are two very distinct approaches to 'identity management'. Large commercial and government organizations assume that it is their responsibility and right to collect, and own, and exploit identity information (whether this is to do with patients' health, customers' buying habits, or citizens' behaviour), and this is what they mean by 'identity management'. The alternative view is that individuals should be the managers of their identities, exercising control (subject to legal safeguards) as to who is allowed to see and make what use of information about them. (This point is well argued in a recent report by the Royal Academy of Engineering; RAE, 2007.) This view leads naturally to being careful to distinguish between 'identification' and 'authentication', and to use the former only when necessary, under very strict legal and technical controls.

There is in fact a growing recognition that centralised identity management, and excessive use of identification when authentication would have sufficed, is inherently dangerous from the point of view of privacy protection, avoidance of identity theft, etc. All that a concert hall manager needs to ensure is that each concert-goer has obtained a ticket – there is no reason for the manager to identify them. Yet NPfIT makes its huge demographic database play a central role in many situations where much less privacy-threatening authentication methods than those currently used could suffice. (This point is well-made in written evidence to the Health Committee from the British Computer Society (HC, 2006), Ev36).

In summary, you can collect a huge amount of valuable, readily identifiable, and hence at risk, information and then set about trying to protect it. Alternatively, you can have a number of carefully segregated sets of minimal information, each under appropriate control. In fact, the best approach is to design information services that will preserve privacy adequately even if some of their data servers have been successfully taken over by criminal gangs! Compromises are inevitable – you have to design for them.

Unfortunately, the impression that colleagues and I have gained is that little thought has been given by CfH to minimizing the need for patient (and clinician) identification (as opposed to authentication), in order to mitigate privacy concerns. Equivalently, we question what thought has been given to allowing individual patients to influence or even control (e.g., with the help of their doctor) what uses are made of their information. (However, given that there are already a number of central systems such as PACS involved in patient care, probably the only affordable immediate possibility of providing privacy to patients is simply to let them be treated under pseudonyms if they wish. Otherwise the only way for them to obtain privacy will be to go private.)

Security mechanisms and their effectiveness

No complex IT system is completely reliable and secure. Very well-argued books (Schneier, 2000; Schneier, 2003) by Bruce Schneier, now with BT, provide much evidence of this. Moreover, his and other work lead me to believe that with NPfIT security failures are more likely to be directly due to insiders than, for example, criminal hackers exploiting an insecure Internet connection, though the insider actions may well have been engineered by plausible outsiders.

The NCRS security plans rest on the use of such mechanisms as smart cards, role-based access control, and software-implemented 'sealed envelopes'. I do not claim to be particularly expert in such matters, since my security work has concentrated on system architecture issues, and how to cope with failures that could affect security.

I am prepared to assume that smart cards (as part of a well-designed overall system) can be adequately secure for use in the NHS, though less so for high-risk banking applications, given the possibility, which though low is non-zero, of them being tampered with (though it is relatively easy to clone them) and of card reading devices being compromised. The problems with smart cards in a healthcare setting are I believe more to do with misuse – such as sharing of cards, carelessness with pin codes, etc. However, though I am aware of a number of careful scientific studies of how NHS users have reacted to EPR and other IT systems, including ones that have been imposed on them, I have not found any such studies regarding smart card usage in health care, only manufacturers' claims and anecdotal reports.

The extent to which unauthorised smart card sharing can be avoided is greatly affected by the speed and convenience with which they can be used. The speed depends on the complexity of the checks that have to be carried out by the

system, and the accessibility of the data that is to be checked. Very large, distant, systems provide bigger challenges to achieving speed than small local ones.

Ideally, the speed of response of NPfIT systems should, like their reliability, have been publicly specified (based on careful usage studies), and guaranteed before implementation and deployment. In practice systems often go through numerous improvement cycles, before their performance is adequate, assuming it ever is. (Just what response times are required in what circumstances for NPfIT's various services is however something that requires careful assessment and realistic controlled trials.)

Hopefully the NPfIT response time problems are ones that can be solved. I would point out that smart cards seem to work well for bar staff who are sharing a till in a busy bar. But persuading bar staff who are working on commission to authenticate themselves using their smart cards each time they wish to register a sale is I assume much easier than persuading a busy A&E consultant to use his/her smart card properly. It is also much easier to provide the bar staff with a very fast and simple user interface.

Role-based access control, used to support the BMA security policy, has I understand proved adequate for maintaining privacy in modest-sized EPR systems (Denley and Smith, 1999) – though I do not claim to be particularly expert in such matters. I question, however, the practicality of role-based access control in very large and heterogeneous organizations such as the NHS. If there were relatively few roles, and role changes, then the technical problems of managing roles, and of verifying that role assignments are in line with security policy requirements, would perhaps be manageable – but it is very doubtful that this would be the case in the huge 'system of systems' that is NPfIT.

A further mechanism to be employed is that of sealed envelopes – this mechanism, which typically employs cryptographic techniques, ensures that information held 'within' the envelope cannot be accessed until the envelope has been 'unsealed', a carefully audited action that can only be carried out by authorised individuals. My understanding is that NPfIT's sealed envelope scheme is still being specified, so I assume that the exact scheme (and with security the devil is in the details) is both novel and untried, though the basic idea of an electronic equivalent of a sealed envelope, for example, for digital signatures, is well-known. However, I note that the Federation for Information Policy Research, whose security expertise I greatly respect, said in their written evidence (HC, 2006), Ev 63:

Sealed-envelope systems have not been built, and it is not clear that they can be.

In fact, achieving security (in particular maintaining confidentiality) in systems of the size of regional (LSP) or national level data repositories *is I believe not practicable, given the huge user population involved*, since it is more a people than a technical problem.

A fourth security technology to be employed is 'anonymisation'. Given a database with identified patient information in it, and the wish to extract information (e.g.,

for statistical or research purposes, the so-called 'secondary uses' in NHS parlance) from this database that does not need to include patient identification, then there is a need to anonymise the information adequately if patient confidentiality is to be respected. This should not just be a case of omitting patients' names and full addresses, as it can be all too easy to identify an individual by pulling together separate apparently non-sensitive information about them via a set of carefully-crafted searches.

From a US Government Report on Privacy and Health Research (Lowrance, 1997):

From a privacy-protection perspective, there is a very wide distinction between personally identifiable data and truly anonymized data. But in practice the demarcation between these extremes is not sharp. Attending assiduously to where particular data lie on the spectrum between them, and especially to data that are somewhere in the middle, is a crucial protection strategy. At present, large amounts of data lie in-between – they are not completely anonymized, but they are not readily identified, either... The power of computers to perform elaborate, powerful, rapid searches, and the pressures for access, mean that merely assigning simple pseudonyms affords little protection.

Complete anonymisation is exceedingly difficult – and it is not at all obvious that the methods employed in the Secondary Uses Service (SUS) to protect demographic information and patient records are anywhere near adequate. Indeed, I would argue the need for (independent public) studies of the effectiveness of NPfIT's planned anonymisation schemes. A big concern is that I understand some existing SUS applications rely on using postcode plus date of birth, data from which many patient identities can be easily ascertained. So either the applications have to change, which is most unlikely, or the data in SUS simply will never be anonymised to any effective extent. In this case, patients should be informed of the risk and be entitled to opt out of the SUS.

It is, incidentally, worth noting that in Iceland there was a proposal to create a national database of health and genetic information that would not just facilitate health care but be sold to drug companies for research – it was believed that the exercise could pay for itself and make a profit. The proposal was to identify records only by means of an encrypted social security number. But it became clear that patients could be identified by means of suitable queries; eventually 11% of the Icelandic population opted out, and the Icelandic Supreme Court found (on the basis of the same human rights law that is the law in Britain) that the database had to support opt-in rather than opt-out (Anderson, 1998).

However, it is well established that most security failures are not due to inadequacies in the security mechanisms employed, but to failures (such as software bugs) in the IT system in which they are employed, or through the actions of people involved with the system, and that such failures are unavoidable, and so must be coped with. (This is the burden of one of Professor Ross Anderson's most frequently referenced publications (Anderson, 1994) – one

I know well since years ago I was his Ph.D. Examiner, and the paper was part of his thesis.) All experience to date thus makes it very evident that with huge systems of the type planned patient records would frequently be divulged (or corrupted, lost or rendered inaccessible), on occasion on a grand scale, probably grand enough to destroy all trust in the entire system. (Imagine the potential impact of a blunder, such as recently happened in the MTAS system to some thousands of junior doctors' job applications, in the context of millions of patient records; e-Health, 2007a.)

To sum up, a very good summary of the fundamental security dilemma facing NPfIT is that one can (with difficulty) achieve any two of (a) high security, (b) sophisticated functionality, and (c) great scale – but achieving all three is currently (and may well remain) beyond the state of the art. (This is not my formulation, but I have no reason to doubt that it is of relevance to the NCRS.) NPfIT looks set to sacrifice security; I believe that it should instead make every effort to evade the scale problem.

What levels of system reliability are achievable?

We need to distinguish between what reliability levels can be plausibly guaranteed and relied upon beforehand (through tests, statistical and logical arguments, etc.) and what reliability levels might turn out, years later, to have been actually achieved.

The Health & Safety Executive's Nuclear Installations Inspectorate have ruled that the Sizewell B nuclear power station's computerised 'emergency shut-down' system was sufficiently complex that it was necessary to assume that it could fail to act on average once every thousand times (Littlewood *et al.*, 2001). (The suppliers originally claimed a guaranteed maximum failure rate of one in ten thousand, but HSE determined that this claim was not sustainable.) Do not worry – this is not the only means of preventing a nuclear disaster in Suffolk!

However, it is important to note that the Sizewell emergency shut-down system is trivially simple compared to NPfIT, and has had an immense amount of validation effort expended on it. In practice, however, this type of highly safety-critical system typically (but by no means always) eventually proves to be several orders of magnitude more reliable than was possible to predict.

Where one has a software system that has been delivered to and used by thousands or millions of people, who are (willingly or unwillingly) helping to complete its testing, very impressive reliability figures can *eventually* be achieved. I understand that Microsoft now claim a mean time between failures of 3000 h, that is, about 4 months, for their Windows operating system, though I doubt if many PC users believe this, since there are all sorts of other things to go wrong besides the basic operating system.

But NPfIT differs greatly from Windows – it is a massive one-off networked system, albeit constructed from a set of pre-existing (but heavily modified) sub-systems. Some failure rates I have been given for one-off networked systems (of a much smaller size than NPfIT) are

- A new military command and control system has on average been suffering one failure every 40 days, and a dangerous failure every 5 months.

- Commercial distributed process control systems using mature hardware and software, as a rule of thumb, suffer total failures about once every couple of years.

Given quite generous assumptions about the number of faults present in the system, the proportion that are dangerous, and the amount of pre-operational testing performed, recent reliability prediction theory (Bishop and Bloomfield, 1996) suggests there could be one dangerous failure every 4 days. This is orders of magnitude worse than the failure rates expected in current safety standards, where targets of one dangerous failure per year (or lower) are set for safety-related computer systems.

Achieving public trust and confidence

The general public needs to trust not just the IT systems, but also the medical staff or government officials (present and future) who control them. In particular, they need to be confident that the information that is collected about them, especially if it is gathered together into what is in effect one huge data repository, is not misused. This could for example happen through this repository being subject to 'mission creep', as various other government departments – and indeed commercial organizations – seek to exploit it. The public thus needs believable reassurances concerning what other systems (inside and outside the NHS) will be allowed to have access to the national summary care record service, and what other systems will have access to the full care records hosted by LSPs, under what legal controls. (The fact that recent official responses to concerns about patient opt-outs and the uploading of patient records have been found to have been deliberately misleading (Anderson, 2006), and the blandness of official reactions to criticisms of NPfIT and NCRS from the Public Accounts Committee (Shifrin, 2007), for example, do not augur well for any future attempts to reassure the public, and indeed the medical profession, regarding such matters.)

Trust is gained slowly and can be lost abruptly – both as a result of system failures, and of the actions (and words) of system owners. The general public's trust in the medical profession, and especially in their own GPs' respect for their privacy, is typically quite high. This provides an excellent basis on which to build, incrementally, an IT system that will also gain the public's trust – providing of course that the system gains and retains the trust of the medical profession. However, if doctors have systems imposed on them, systems that are under some distant control and ownership, then this avenue towards a well-accepted and trusted national health IT system has been closed off from the outset.

Recommendations

Let me first quote some recommendations from a Connecting for Health document – unfortunately, these recommendations are not from the British, but rather are from the *American*, Connecting for Health organization (CfH (US), 2006):

It is desirable to leave to the local systems those things best handled locally, while specifying at a national level those things required as universal in order to allow for

exchange among subordinate networks.

Avoid 'Rip and Replace': Any proposed model for health information exchange must take into account the current structure of the healthcare system...

Separate Applications from the Network: ... The network should be designed to support any and all useful types of applications...

Decentralization: Data stay where they are... [this] leaves judgments about who should and should not see patient data in the hands of the patient and the physicians and institutions that are directly involved with his or her care.

My own most confident recommendation concerns the vital urgent need for an open constructive review of NPfIT, and in particular the NCRS, by independent, including international, experts. This would need to cover, indeed be centred on, reviewing the specification by clinicians and by safety and privacy experts who know what is wanted, and simultaneously by technical people who can assess whether the proposed architecture and systems will deliver it.

Such a review is something that my colleagues and I first urged last April. At the time, in response to a request, we formulated and provided the Select Committee with suggested terms of reference (see the Appendix). In the light of developments since then we now regard a review of the type we described as even more necessary.

It is worth recalling that when we first proposed such a review Dr. Richard Granger accepted that such a review could be worthwhile, and CfH issued a press release to this effect. A constructive review of the type we envisage would in our opinion be a help rather than a hindrance, if carried out properly. Moreover, if the reviewers were evidently independent and well qualified, the review could be of great help in re-establishing trust and confidence in NPfIT among NHS staff and the general public.

The review would be best carried out by a smallish team, say of no more than six people operating full-time at least initially, who between them have expertise in all the main areas – medical informatics, computer system dependability, usability, privacy, etc. It should be undertaken in stages, since one could hope for some significant contributions quite early on, and the decision as to how long the review should continue could be taken in the light of progress. (In fact, it might be worth it continuing on a permanent but part-time basis on what I understand is the approach being used in Wales.)

The review team should not exist in a virtual ivory tower, emerging occasionally brandishing reports, but should be an integral part of the Programme, in regular contact with CfH and continually discussing issues and findings. This is essential for a constructive review, and to avoid perpetuating present adversarial attitudes, and any danger that the team will be denied key evidence (Bennatan, 2006). (Examples of the sort of review we have in mind are the US Department of Defense's scheme of 'independent and objective software and system assessments' (Baldwin 2000), and the 'independent technical assessment' schemes of the Mitre Corporation (Clapp and Funch, 2003) and the Software Engineering Institute (SEI, 2004), well-established schemes that have successfully rescued numerous large IT projects.) Incidentally, such a review might also result in

long-term recommendations that could help to avoid similar problems in future massive government IT projects.

My expectation is that such a review would provide support for my view (indeed many people's view) that specifying, implementing, deploying and evaluating a sequence of ever more complete IT systems is the best way of ending up with well-accepted and well-trusted systems – especially if this process is controlled close to the coalface, rather than by some distant central bureaucracy. (Mere evolutionary deployment, which is all that CfH is attempting, is insufficient.) Thus authority as well as responsibility should be left with trusts to acquire IT systems that suit their environments and priorities – subject to adherence to minimal interoperability constraints – and to use LSP services (e.g., for system support and back-up) as if and when they choose.

Note the fundamental inconsistency of NHS trusts being encouraged by the Government to compete in order to improve, but being subject to stultifying central control in matters of IT. One can draw an interesting contrast to the situation of local authorities – they are all required to carry out the same functions (though not in competition, except to meet central government benchmarks), but can make their own IT procurement choices.

Though such a change of direction would seem revolutionary, I feel it would be by far the best way of capitalising (to eventual splendid effect) on what has been achieved so far by Connecting for Health and its suppliers, and would help avoid a knee jerk political decision that could amount to throwing the baby out with the bathwater.

My other recommendations are more tentative – all are motivated by the concerns that I have presented above regarding the reliability and effectiveness of the overall NPfIT system, that is, the socio-technical system consisting of the computer systems, networking facilities, software and people. (None are original and all could be reconsidered in the light of the review findings.)

1. First ensure that individual trusts have well-established and well-trusted systems for their clinical records as well as for patient administrative record-keeping needs, suiting their particular circumstances and priorities (while giving due attention to issues of interoperability). Only then attempt to evolve gradually into having an adequately integrated overall system. There is thus a need for (evolutionary) central standards for interoperability of local systems, standards that will encourage (but not constrain) development of improved functionality and usability of these systems.
2. Reassess the current architectural approach to information sharing, determining the extent to which it would be feasible to move away from centralised databases of identifiable information, to sets of separate databases, each designed (by specialisation and the use of cryptography) to minimise its utility to an identity thief, bound together into one or more overall virtual databases. Thus work towards abandoning the idea of an actual EPR as such, in favour of virtual EPRs, each providing access to data appropriate to its purpose, and only these data.
3. In specifying EPR-related systems, give suitable precedence to patient and clinician needs and preferences

- over support of bureaucratic oversight and management, and where appropriate develop these systems and their specifications via a sequence of (evaluated) operational 'prototype' systems.
4. Investigate how to reduce the extent to which demographic information is used directly for authentication purposes, both locally and nationally, when other more secure methods could be used instead. In the mean time allow and facilitate the use of pseudonyms where necessary.
 5. Insist that all software developed specifically for the NHS is open (i.e., that its source code is available for general inspection) and the property of the NHS, and that all off-the-shelf software and systems purchased by the NHS have open (functionality and dependability) specifications.
 6. Identify established systems that are 'fit for purpose' and consider these as building blocks for development, not as 'brownfield sites' (e-Health, 2007d) that impede the deployment of standard solutions.
 7. Allow patients to decide (typically via their GPs) the extent to which information about them should be made widely available, and the uses that can be made of this information. (This recommendation is motivated by the practical need to gain people's trust and hence cooperation, not in order to support 'privacy fascists', to use Dr. Richard Granger's delightful term; HC, 2007: 23.)
 8. Make effective use of existing, and commission further, research into socio-technical aspects of EPR systems.

Concluding remarks

I must reiterate that I claim no expertise in medical informatics, and have had no direct involvement with NPfIT. (To avoid any misunderstanding, I should also make it clear that I have no wish to take part personally in the technical review that I and my co-signatories are advocating.) Any assessment by an outsider such as myself of NPfIT's plans and progress is very difficult, due to the Programme's size and complexity, the secrecy regarding detailed system specifications, and the atmosphere of fear that prevents many NHS staff from expressing criticisms. Such knowledge as I have gained of NPfIT has been through my work on the Dossier of Concerns² that we have assembled from a large variety of sources, ranging from newspaper articles to refereed scientific papers, but also through correspondence and conversations with a number of NHS staff who unfortunately feel a need to remain anonymous, such is the atmosphere surrounding the development of NPfIT.

Acknowledgements

I have been greatly aided in the preparation of the present paper by my co-signatories and also by colleagues from Newcastle, DIRC and elsewhere. My thanks to all, and my apologies for any accidental misrepresentations or inadequacies in this resulting account.

Notes

- 1 <http://nhs-it.info/>
- 2 <http://nhs-it.info/>

References

- Adams, A., Blandford, A. and Attfield, S. (2005). Implementing Digital Resources for Clinicians' and Patients' Varying Needs, *Medical Informatics and the Internet in Medicine* 30(2): 107–122.
- Anderson, R. (1994). Why Cryptosystems Fail, *Communications of the ACM* 37(11): 32–40.
- Anderson, R. (1998). *The DeCODE Proposal for an Icelandic Health Database*, Computer Laboratory, Cambridge University, 12pp. [<http://www.cl.cam.ac.uk/~rja14/Papers/iceland.pdf>].
- Anderson, R. (2006). Headed for the Rocks, *The Guardian* (21 December 2006) [<http://www.guardian.co.uk/comment/story/0,,1976589,00.html>].
- Baldwin, K. (2000). Help Identify and Manage Software and Program Risk, *Crosstalk: The Journal of Defense Software Engineering* 13(11): 8–11 [<http://www.stsc.hill.af.mil/CrossTalk/2000/11/baldwin.html>].
- Bennatan, E.M. (2006). *Catastrophe Disentanglement: Getting Software Projects Back on Track*, Upper Saddle River, NJ: Addison-Wesley, 288pp.
- Bishop, P.G. and Bloomfield, R.E. (1996). A Conservative Theory for Long-Term Reliability Growth Prediction, *IEEE Transactions on Reliability* 45(4): 550–560.
- Bowers, S. (2006). Fujitsu Under Spotlight for NHS failures, *The Guardian* (24 October 2006) [<http://business.guardian.co.uk/story/0,,1929770,00.html>].
- Brennan, S. (2005). *The NHS IT Project: The Biggest Computer Programme in the World... Ever!* Abingdon: Radcliffe Publishing Ltd.
- Broome, C. and Adams, A. (2005). Out with the Old in with the New: What gets missed when deploying new technologies in A&E? *Medical Informatics and the Internet in Medicine* 30(2): 34–40.
- Budgen, D., Rigby, M., Brereton, P. and Turner, M. (2007). A Data Integration Broker for Healthcare Systems, *Computer* 40(4): 34–41 [<http://doi.ieeecomputersociety.org/10.1109/MC.2007.112>].
- CfH (2006). *MIQUEST*, Connecting for Health (12 June 2006) [<http://www.connectingforhealth.nhs.uk/systemsandservices/data/miquest/>].
- CfH (US) (2006). *The Common Framework: Overview and Principles*, Connecting for Health, Markle Foundation, New York (5 December 2006), 12pp. [<http://www.connectingforhealth.org/commonframework/docs/Overview.pdf>].
- Clapp, J.A. and Funch, P.G. (2003). *A Guide to Conducting Independent Technical Assessments*, Mitre Corporation, Bedford, MA (5 March 2003), 80pp. [http://www.mitre.org/work/sepo/toolkits/assessment_guide.pdf].
- Clarke, K., Hartswood, M., Procter, R., Rouncefield, M. and Slack, R. (2002). Minus Nine Beds: Some Practical Problems of Integrating and Interpreting Information Technology in a Hospital Trust, in J. Bryant (ed.) *Proceedings of the BCS Conference on Healthcare Computing*, (18–20 March 2002), Harrogate, pp. 219–225.
- Collins, T. (2006). NHS Trust Uncovers Password Sharing Risk to Patient Data, *Computer Weekly* (11 July 2006) [<http://www.computerweekly.com/Articles/2006/07/11/216882/NHS+trust+uncovers+password+sharing+risk+to+patient.htm>].
- Denley, I. and Smith, S.W. (1999). Privacy in Clinical Information Systems in Secondary Care, *British Medical Journal* (318): 1328–1331. [<http://www.bmj.com/cgi/content/short/318/7194/1328#art>].
- DIRC. Interdisciplinary Research Collaboration in Dependability [<http://www.dirc.org.uk/>].
- Department of Health (2007). *NHS Caldicott Guardians*, Department of Health (9 February 2007) [http://www.dh.gov.uk/en/PolicyAndGuidance/InformationPolicy/PatientConfidentialityAndCaldicottGuardians/DH_4100563].
- e-Health (2005). PCT Safety Culture Needed to Prevent Errors, *e-Health Insider*, Primary Care (30 September 2005) [<http://www.ehiprimarycare.com/news/item.cfm?ID=1458>].
- e-Health (2006). Local Sealed Envelopes 'Probably Safer', *e-Health Insider* (28 Nov 2006) [<http://www.ehiprimarycare.com/news/item.cfm?ID=2302>].
- e-Health (2007a). CfH Pulls Health Informatics Community Website, *e-Health Insider* [<http://www.e-health-insider.com/news/item.cfm?ID=2645>].
- e-Health (2007b). Survey Reveals Doctors' Pessimism about NPfIT, *e-Health Insider*, Primary Care (19 February 2007) [<http://www.ehiprimarycare.com/news/item.cfm?ID=2491>].
- e-Health (2007c). 'Sealed Envelopes' on Hold as Policy Debate Continues, *e-Health Insider*, (10 May 2007) [<http://www.e-health-insider.com/news/item.cfm?ID=2680>].
- e-Health (2007d). Granger says 'Consultation' Led to Records Delays, *e-Health Insider*, (26 April 2007) [<http://www.e-health-insider.com/news/item.cfm?ID=2643>].

- Eason, K.D. (2006). *A Local Sociotechnical Design Approach to Exploiting the Potential of The National Healthcare IT Programme NPfIT*, London: The Bayswater Institute (3 Nov 2006), 9pp. [<http://www.bayswaterinst.org/downloads/Exploiting%20the%20Potential%20of%20NPfIT.pdf>].
- Evans, R. (2006). The Woman Falsely Labelled Alcoholic by the NHS, *The Guardian* (2 November 2006). [<http://society.guardian.co.uk/e-public/story/0,1937302,00.html>].
- Ferrar, M. (2006). *The NHS, Standards, Security & Identity Management*, OASIS Adoption Forum, London (28 November 2006) [<http://www.oasis-open.org/events/adoptionforum2006/slides/ferrar.ppt>].
- Finkelstein, A. and Dowell, J. (1996). A Comedy of Errors: the London Ambulance service case study, in *Proceedings of the Eighth International Workshop on Software Specification & Design IWSSD-8*, Paderborn, Germany: IEEE, pp. 2–4. [<http://www.cs.ucl.ac.uk/staff/a.finkelstein/papers/lascase.pdf>].
- Gall, J. (1975). *Systemantics: How Systems Work and Especially How They Fail*, New York: Quadrangle/New York Times Book Co, 111pp.
- Gibson, R. (2007). Hospital Patient Records System is a 'Nightmare', *Milton Keynes News* (30 March 2007) [<http://www.mk-news.co.uk/mknews/DisplayArticle.asp?ID=77775>].
- Gilon, G., Ofek, Z. and Halevy, A. (2003). dbMotion: Virtual health community, in *Integrated Care Records: Problems and Solutions Workshop*, Edinburgh [<http://www.iccs.informatics.ed.ac.uk/~mjh/chameleon/ICRworkshop/Submissions/halevy.pdf>].
- Hardstone, G., Hartswood, M., Procter, R., Slack, R., Voss, A. and Rees, G. (2004). Supporting Informality: Team working and integrated care records, in *Proceedings of the 2004 ACM Conference on Computer Supported Cooperative Work*, Chicago, pp. 142–151.
- 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): 241–266 [<http://springerlink.metapress.com/content/xw4424x3u175xx63/fulltext.pdf>].
- HC (2006). *The Electronic Patient Record – Written Evidence*, London: House of Commons Health Committee, The Stationery Office Limited (25 April 2007), 196pp. [<http://www.publications.parliament.uk/pa/cm200607/cmselect/cmhealth/422/422ii.pdf>].
- HC (2007). *The Electronic Patient Record: Uncorrected Transcript of Oral Evidence*, House of Commons, Minutes of Evidence Taken Before Health Committee (26 April 2007). [<http://www.publications.parliament.uk/pa/cm200607/cmselect/cmhealth/uc422-i/uc42201.htm>].
- ICP (2006). *What Price Privacy? The Unlawful Trade in Confidential Personal Information*, London: Information Commissioner to Parliament, The Stationery Office, (10 May 2006), 48pp. [http://www.ico.gov.uk/upload/documents/library/corporate/research_and_reports/what_price_privacy.pdf].
- Littlewood, B., Popov, P. and Strigini, L. (2001). Modelling Software Design Diversity – A review, *ACM Computing Surveys* 33(2): 177–208.
- Lovrance, W.W. (1997). *Privacy and Health Research*, Washington, DC: United States Department of Health & Human Services, (May 1997), [<http://aspe.hhs.gov/datacncl/PHR.htm>].
- Medix (2006). *Medix survey (Q1066) of doctors' views about the National Programme for IT (NPfIT)*, Medix UK plc, (21 Nov 2006), 54pp. [<http://ixdata.com/reports/106620061121.pdf>].
- NAO (2006). *The National Programme for IT in the NHS*, London: National Audit Office, The Stationery Office, (16 June 2006), 57pp. [http://www.nao.org.uk/publications/nao_reports/05-06/05061173.pdf].
- Naur, P. and Randell, B. (Eds.). (1969). *Software Engineering: Report of a Conference Sponsored by the NATO Science Committee*, Garmisch, Germany, 7th to 11th October 1968, Brussels, Scientific Affairs Division, NATO, 231pp. [<http://www.cs.ncl.ac.uk/research/pubs/proceedings/papers/16.pdf>].
- PAC (2007). *The National Programme for IT in the NHS*, London: House of Commons Committee of Public Accounts, The Stationery Office (11 April 2007), 180pp. [<http://www.publications.parliament.uk/pa/cm200607/cmselect/cmpubacc/390/390.pdf>].
- Proctor, R., Hartswood, M., Rouncefield, M., Martin, D., Mariani, J. and Taylor, A. (2007). Understanding and Improving the Design, Deployment and Use of Electronic Health Records, Chameleon Project – EPSRC GR/R86751/01, [http://homepages.inf.ed.ac.uk/mjh/chameleon/chameleon_final_report.pdf].
- RAE (2007). *Dilemmas of Privacy and Surveillance: Challenges of Technological Change*, London: Royal Academy of Engineering (March 2007), 64pp. [http://www.raeng.org.uk/policy/reports/pdf/dilemmas_of_privacy_and_surveillance_report.pdf].
- Schneier, B. (2000). *Secrets and Lies: Digital Security in a Networked World*, New York: John Wiley & Sons, 412pp.
- Schneier, B. (2003). *Beyond Fear: Thinking Sensibly About Security in an Uncertain World*, New York: Springer-Verlag Inc., 303pp.
- SEI (2004). *Software Engineering Institute annual report fiscal year 2003*, Pittsburgh: Carnegie Mellon University, 104pp. [<http://www.sei.cmu.edu/pub/documents/misc/annual-report/2003/2003.pdf>].
- Shifrin, T. (2007). *Government defends £12.4bn NHS IT scheme after MPs' damning report*, UK: Computerworld (17 April 2007) [<http://www.computerworlduk.com/management/government-law/public-sector/news/index.cfm?newsid=2627>].
- Stearns, D. (2006). In Plastic We Trust: Dependability and the visa payment system, in *DIRC Conference*, Newcastle (April 2006) [<http://www.sociology.ed.ac.uk/finance/Papers/StearnsDIRC06.pdf>].
- SWTRHA (1993). *Report of the inquiry into the London Ambulance Service*, London: South West Thames Regional Health Authority, (February 1993), 62pp. [<http://www.cs.ucl.ac.uk/staff/A.Finkelstein/las/lascase0.9.pdf>].

About the author

Brian Randell graduated in Mathematics from Imperial College, London in 1957. From 1964 to 1969 he was with IBM in the United States, mainly at the IBM T.J. Watson Research Center. He then became Professor of Computing Science at the University of Newcastle upon Tyne, where he has been Principal Investigator on a lengthy succession of research projects on system reliability and security, and is now Emeritus Professor of Computing Science, and Senior Research Investigator. He has published nearly two hundred technical papers and reports, and is the co-author or editor of seven books. He was a Member of the Conseil Scientifique of the CNRS, France (2001–2005), and Chairman of the IEEE John von Neumann Medal Committee (2003–2005), and is a Member of the ACM A.M. Turing Award Committee (2005–2009). He has received a D.Sc. from the University of London, Honorary Doctorates from the University of Rennes and the Institut National Polytechnique of Toulouse, and the IEEE Emanuel R. Piore 2002 Award.

Appendix

Proposed terms of reference

The Review should be pragmatic and constructive, and is intended to assist the NHS to achieve its overall aims. As a contribution to establishing confidence in both NPfIT and in the review itself, the review will be an open one. The final report and any interim reports will be published, and evidence given to the review will be made publicly available as far as possible. The review will be guided by an international expert advisory board. The review will undertake the following tasks.

1. Determine the detailed specifications that presently define the technical goals of the NPfIT systems, and examine the processes through which these specifications have been shown to meet the needs of all the users of the systems.
2. Consider the architectural approach that has been adopted to meet these specifications, in particular regarding the decisions made concerning centralised vs federated approaches to system construction, and the replacement or reuse of existing applications.
3. Assess the mechanisms used to control system evolution and manage change, assess the gap remaining between

- user requirements and system specification, and establish whether the rate of specification change is increasing or decreasing.
4. Assess whether the detailed technical architecture and application designs will deliver systems that match both the required functional aspects of those specifications and the required dependability aspects (safety, privacy, availability, reliability, accuracy, performance, usability, fault tolerance, and modifiability); if appropriate suggest necessary improvements.
 5. Review the programme's plans and budgets to assess whether appropriate resources are available for development, process prototyping, pilot studies, modifications, interfacing with existing systems, roll-out, training, data cleansing and maintenance.
 6. Review NPfIT risk management and consult with stakeholders to uncover major obstacles that could jeopardise the successful implementation of the new system and associated work practices; where appropriate, suggest possible ways to overcome these obstacles.

Notes

1. The Review should encompass the work of both National and LSPs.
2. In order to perform its functions, the review team should have access to all information available to the Secretary of State.
3. It shall include a formal public consultation conducted under Cabinet Office guidelines.