DIRC
Dependability of Computer-Based Systems

Input to EPSRC Mid-Term Review

July 31, 2003

This document is the main written input to the "Mid-Term Review of DIRC".

Specifically, these pages coupled with the technical presentations scheduled for August 12–13 will be the basis on which DIRC's *Steering Committee* will assess the research progress made in the first three years of DIRC's funding.

Martyn Thomas (as Chairman of the Steering Committee) will inform the management review (August 19–20) of his committee's view of the research to date and plans for the future.

The other documents to made available to the review are:

1. The original "Case for Support" submitted by DIRC to EPSRC
2. A list of publications and reports by members of the DIRC team
3. Copies of slides used at the earlier Steering Committee meetings

Contents

1 Introduction and background

The first half (2000-2003) of the DIRC project has been extremely productive. An excellent interdisciplinary team has been built and there have already been over 250 publications and reports.[1] A model of what can be achieved when researchers cross disciplines is Donald MacKenzie's book [Mac01] which looks at highly technical work on theorem proving, its mechanization and its use in software development from a sociological point of view. As a direct result of the DIRC project, an interdisciplinary book on "Trust in Technology" is close to completion.

The DIRC team includes psychologists, sociologists, statisticians and computer scientists. This breadth was seen as essential if we were to make a contribution to the Dependability of the wider "Computer-Based" systems on which our society now depends. Creating the trust and cooperation between researchers from such different fields has been both exciting and difficult. In the spirit of EPSRC's request to reflect on interdisciplinarity, Section 4.1 discusses the challenges we have observed and what we have done to achieve our current state. It is extremely stimulating to see the different insights that sociologists and computer scientists bring to a discussion of a concept like "classification"; real progress in understanding how psychological type might influence skills in development teams is coming from close and long-term collaboration between psychologists and computer scientists; seeing statisticians, sociologists and computer scientists tackling problems in organising the reading of mammograms makes it obvious that it is essential that researchers embrace interdisciplinary modes of work.

Perhaps because of the success to date, we see clearly how much there is still to do. We are convinced that the six-year format of IRC funding has been essential. In particular, one needs space to persuade younger researchers to take time out of their desire to publish in the most productive vein of their own speciality or to work with a colleague whose background and values are very different. Even after having written a joint paper they can then find that there is no ideal place to publish the interdisciplinary results. Section 4 expands on these issues.

We look forward to building on the success to date over the next three years of IRC funding.

1.1 Structure of this report

The structure of this document is

- Sections 1.2–1.4 set out the context of –and background to– our project.
- Section 2 describes some of the technical challenges of Dependability.
- Section 3 reports on progress in the first half of the project including information about who is involved in DIRC.
- Section 4 attempts to identify some lessons that might be applicable to other IRCs.

[1] A list of these is available as a separate document. The references in the current report are just that: DIRC papers only appear where explicitly cited; papers from outside DIRC are freely referenced.

- Sections 5 and 6 outline how we intend to move forward: the former describes how we are restructuring our activities and the latter prioritizes the research goals.

1.2 EPSRC call for IRCs

EPSRC called for proposals for six year IT-Centric *Interdisciplinary Research Collaborations* in 1999 and received over 100 outline proposals. Twelve of these were short-listed and detailed proposals were submitted. The selection of the five successful IRCs was made after a panel in February 2000.

It is important to notice that this generation of IRCs differed from previous rounds: earlier Interdisciplinary Research *Centres* had not had such a strong emphasis on distribution and were for ten years. Section 4 describes the advantages and challenges of the twin need to cope with a mix of disciplines and geographic distribution.

DIRC was among the five IT-centric IRCs funded in 2000. The funding of five –rather than the originally foreseen four– meant that budgets were tight and, where we requested a grant of £8.4M, DIRC was actually awarded £6.8M; our way around this shortfall is discussed in Section 4.3.

A copy of the original "case for support" for *Dependability of Computer-Based Systems: A Proposal for an Interdisciplinary Research Collaboration* is attached. It was clear that the proposal went well beyond the traditional safety-critical system arena;

DIRC is interested in *computer-based* systems which includes those where people are a significant component of the overall system.

We have followed the stated plan closely over the first half of the project but are now implementing some changes which are discussed in Section 5.

1.3 DIRC organisational structure

As planned in the Project Proposal:

- each of the five research teams is led by a *Principal Investigator* (PI)[2]
- the *Industrial Advisory Board* (IAB) is chaired by the *Industrial Liaison Director* (ILD)
- the project is supported by an *Administrative Coordinator* (AC) who is funded by the IRC
- the project is managed by its *Executive Board* (EB) which is made up of the five PIs, the ILD and supported by the AC

The EB meets about quarterly. The majority of meetings have so far been physical (with a rough rotation around the five sites) but "tele-conferences" are used for urgent or brief matters. There is a wish to move further in this direction. Nominal substitutes have been identified but the attendance record of PIs has been excellent. EPSRC made clear that they expected a single point of contact and insisted that a *Project Director* (PD) was identified and they

[2] All names are given in Appendix A.

even reserved the right to reconsider the IRC if the PD were to change. Cliff Jones was named in the proposal as PD. The PD chairs the EB.

The membership of the IAB has evolved with some of our key contacts changing employers and new contacts (e.g. Philips Research) developing. The IAB is seen as an "advisory board" rather than, necessarily, as representing those outside organisations with whom DIRC is working on projects but there is an intersection of the two sets and we are in negotiation with several of the IAB members about the bigger engagements discussed in Section 6.2.[3]

The project receives, via the PD, advice from the *Steering Committee* (SC). EPSRC required the creation of an independent SC and the invitation to Martyn Thomas as chairman was by mutual agreement between EPSRC and the PD. Martyn Thomas was a member of the IRC selection panel (and is now a member of the Council of EPSRC). The membership of the SC was agreed between its chairman and the PD. The SC has had three formal meetings (2001-05-24; 2002-01-03/04; 2002-12-17/18)[4] prior to the "Mid-Term review".

1.4 Mid-term reviews

EPSRC made clear at the outset of the IRCs that there would be some form of mid-term review of progress. There was however a difference in perception of the role of these reviews between the universities and EPSRC. For DIRC, the impact of this has now passed; but it might be worth mentioning it for future IRCs. Some of the universities involved in DIRC were reluctant to have RA contract commitment run beyond the mid-term review in case the funding were withdrawn. It is –at least now– clear that EPSRC have a very different view of the mid-term review and that there is no higher probability that an IRC will be stopped prematurely than any other EPSRC-funded project.

The five IT-Centric IRCs met with Vince Osgood on February 24th in London. This was a very useful meeting at which several issues (e.g. allocation of DTAs for the whole period of the IRCs) were cleanly resolved.

DIRC welcomes the very constructive Terms of Reference (reproduced in Appendix E) for its mid-term review and has responded by being as frank as possible in its submission to the reviews. It is hoped that this openness will help EPSRC in future planning of large-project funding.

2 The challenge

Most readers of this document will be all too aware of the challenge of building large systems in a world where so much hangs on their Dependability [Nor88, Sch99]. But it is worth reviewing the challenge of improving the quality of systems because the question of how to have an impact poses a strategic question for DIRC: we continue to review our balance between research and dissemination.

To begin with a contrast, if a project promises to build a particular artefact such as an image analysis algorithm, this can have clear criteria for success. We are not for a moment suggesting that meeting such criteria is easy — nor

[3]Robin Bloomfield will report on the IAB at the SC review. He is about to release a new WWW site for IAB members.

[4]The slides presented are attached to this report.

that such projects are more development than research. (It would for example be all too easy to set precise but unattainable objectives for such software.) The point is that the criteria can be precise. For DIRC to really change the world it needs to cause other people to develop more dependable systems. It is not enough to invent methods which no one uses; nor would it be adequate to exhibit one demonstrator project which relies on all of the DIRC team focussing on that one development. Methods have to be developed and shown to be at least transferable. This is the challenge for DIRC. It is certainly one of the reasons that makes the six year funding window of IRCs essential for the area of Dependability.

Looking first at issues which affect the Dependability of purely technical systems, it is obvious that no large-scale engineering project is easy. Projects are often seen to face the tensions between Cost, Time and Quality. The design of software is peculiarly hard. Not only are large software systems among the largest engineering tasks tackled by humans, there is also something indirect about creating an artefact (the code) which is actually executed by a completely dumb servant (the computer) which, if told to obey some obviously senseless sequence of instructions, will happily do so.[5]

At its best, software development has become a proper engineering discipline. It is informed by theory and engineers know when to deploy the more difficult intellectual tools. It is our firm conviction that Dependability is an issue which has to be tackled from the inception of a system. Not only do we not believe in some magic tool which will remove errors from a badly designed system, we are also convinced that good methods can pay for themselves if (and only if) they are deployed from the earliest stages of a project. Architecture counts: a well-chosen architecture will result in a cost-effective and high-quality product; it will also stand a significantly higher chance of making it possible for the system to evolve as its environment changes. It is software that has created the IT revolution; the Moore's law improvement in hardware cannot itself create systems which change people's lives.

It must, however, also be conceded that much software development is truly awful! The rush to provide new "features" wins over Dependability in far too many cases. There is also an insidious compounding effect here which again shows the importance of architecture. A system like MicroSoft Word[6] is designed for one (in this case WYSIWYG) paradigm and progressively modified to respond to challenges from other products (such as the need for "markup" to create large documents). The resulting system becomes difficult to use in any mode. The software industry also faces the frequent criticism that its products are expensive to build and projects are difficult to estimate. (The extent to which the situation is really worse than other engineering fields is not estab-

[5]Brian Randell in [Ran00] writes eloquently about the choice he made early in his career to undertake research into software that offered some degree of Dependability in spite of imperfections and contrasts this with researchers like Dijkstra who sought ways of writing "perfect programs". Today, we see that there is no competition here: we need to use both sets of ideas. Indeed, one of the interesting challenges is to understand how to reason (in a Hoare like way) about those constructs that Randell and like-minded researchers have offered for containing and ameliorating errors. Beyond that, there are serious research challenges about how to combine evidence from the fact that certain reasoning has been checked and empirical testing data.

[6]The example is taken because it is widely known; there is no suggestion that this company or product are worse than the median in the industry.

lished.) In the weaker parts of the software industry, there is little or no process; no design documentation; no correctness statement; and no exploration of alternatives. Saying that it is a young industry with many practitioners who are not trained *engineers* cannot continue to be taken as an excuse. Things must clearly improve.

This discussion of the problems with what might be thought of as technical systems in no way diminishes DIRC's commitment to wider "computer-based" systems. Unfortunately, systems where people are significant components bring additional problems; they do not remove any of those which apply to the narrower class of system.

Today, most computer systems are deeply embedded within groups of humans. What the Moore's law reduction in cost and size has given ICT is the ubiquity and closeness to the users. Nearly all professionals, and a huge percentage of manual workers could not do their job without computers. But many of these computer systems are ill thought out and difficult to use. In many cases, the system is designed with scant attention to –or understanding of– the needs and capabilities of the user. In the highly professional contexts that DIRC has been investigating (e.g. NHS) there are crucial relationships of trust and responsibility which must be fully understood if a system is ever to be deployed successfully.

Humans possess the crucial advantage over machines that they can reason and question stupid instructions but an operator controlling –say– a chemical plant needs to be able to build a mental model of what is going on within the equipment. In an environment where human operators are under high cognitive load, the extent to which the system design ensures that the true state of the underlying physical process can be understood determines whether the operator stands a chance of playing their part in safety.[7]

To design a successful computer-based system, the wider system must be understood. Furthermore, we again see that to assess the safety of such a system we must include information about the likely behaviour of its operators. One could also add that the more humans are involved in a computer-based system, the greater will be the pressure for the system to evolve.

DIRC must therefore devise ways of understanding the human context in which a system will be placed and improve ways of developing systems. For computer-based systems, the development methods are likely to have to involve interaction with users *during* the design process. Section 6 outlines our commitment to devising a family of development methods.

There is a key distinction between the process of creating a system and the created system. Humans are the tool for the creation process. We therefore recognise that it is essential that DIRC looks at the development process itself and there is fascinating scope here for the interdisciplinary team that we have established.

It is clear that there are technical and dissemination challenges to make the area of Dependability worthy for an IRC. There is also a tension as to how much effort should be put into research as opposed to dissemination. We believe that we are balancing these pressures well so far but we welcome the fact that the IAB and SC can continue to guide us through the next three years.

[7] cf. [Vic99, Rus02].

3 Progress

This section gives an overview of what has been achieved in the first half of the initial DIRC programme. Appendices provide further details. Among the highlights are earlier than expected external workshops; publications and conference talks plus a sequence of prestigious external invitations to speak about DIRC research; and the formation of related projects.

None of this would have been possible without the team which has been formed so let's start there.

3.1 Staffing

The names of all RAs involved are given in Appendix A.1. The EPSRC and DSTL combined funding for DIRC would support just over twenty RA person years per year throughout the six years. Not surprisingly, it took time to build up to that level with City university finding it hardest to attract people because the "London weighting" is simply too small. Some RAs have moved on to "permanent" faculty positions but in each case they have stayed involved in DIRC research. Given the time that it takes to establish interdisciplinary trust, we have been fortunate that the teams have remained fairly constant. This is another regard in which the six-year funding horizon has proved indispensable.

The EPSRC funding came with four PhD studentships per year. This has been an invaluable support and we are delighted that this level of PhD starts will continue throughout the six years of the IRC. The names of the post-graduates involved (including some working within DIRC but funded from other sources) are give in Appendix A.2.[8]

In addition to the twenty or so RAs, there are more than that number of permanent "faculty staff" whose research is principally within DIRC. Obviously, the extent of such involvements varies and the list in Appendix A.3 gives an indication of the extent to which the faculty member is involved in the IRC: those with a lesser involvement are likely to be heavily involved in other externally funded projects. The senior level of some of these researchers speaks volumes for their attitude to the value of working with DIRC and makes the IRC extremely good "cost/value" for EPSRC. This is a clear case where the six year format of IRCs is much more useful than two successive three year awards.

As indicated in Section 1.3, the five PIs have the responsibility for running DIRC. It is therefore extremely fortunate that this team has remained stable. The only change is that Michael Harrison has just taken over from Alan Burns as the PI for DIRC at York. This was triggered by Alan's sabbatical but Michael had been Alan's substitute on the EB, had consequently seen all EB e-mail traffic and the transition should therefore be completely smooth.

In the Project Proposal, we had the post of Technical Director (TD) in addition to that of PD. John Dobson played this role for about two years but then asked to be relieved of the responsibility so that he could concentrate on his own research. He has subsequently taken early retirement from Newcastle but retains some involvement in DIRC via Lancaster. In particular, he is the author of one of the chapters of the "Trust in Technology" book and we hope that he will write more on ways to record Responsibility.

[8] The first theses are being prepared now; other postgraduates are also associated with DIRC [Rou02, Loe03, SO02].

EPSRC wisely foresaw the need of a mechanism to "buy out" the time of the PD of any IRC. For various tactical reasons, DIRC initially used this funding to cover (75% of) the TD's salary. This money is now used to fund replacement lecturing staff in the School of Computing. Although the School was already supportive, this arrangement has made it easier to say "no" to university distractions. We have raised with EPSRC the question of "buy out" for the time of other senior Faculty members but have not received a positive response.

DIRC's Administrative Coordinator was initially Barry Hodgson who left to join the Arjuna start-up company. He has been ably replaced by Jon Warwick who is an invaluable support to all of the project's boards and to the PD in particular.

3.2 Project Activities

DIRC's original manpower planning showed 80% of research being done within the Project Activities. PAs have focused deliverables and are the mechanism within which DIRC's involvement with outside organisations has been focused. As is discussed in Section 5, PAs have absorbed even more of the effort than foreseen in our original plan.

Overall, the PAs have been extremely successful, they have

- generated significant numbers of publications;
- generated large amounts of raw material for the Research Themes; and
- engaged with a variety of external organisations.

In this section some highlights and general points are covered. Further details about individual PAs (authored by those involved) are given in Appendix C.

Human Interaction in Real-Time Systems – PA1 The initial hope was that the case studies informing PA1 would be with the Spanish Telephone company Telefonica and with the UK Air Traffic Control System (NATS). The contact for Telefonica changed and the semi-privatisation of NATS meant that there was a period where it was difficult even to discuss collaboration. The actual case studies conducted have been in the Neonatal Unit at Leeds Hospital and with the ParcelCall project. The latter is in collaboration with Edinburgh. It was unfortunate that NATS was "in flux" at the start of DIRC since their problems present a perfect mix of human and technical challenges. It is hoped that recent negotiations spearheaded by Robin Bloomfield will result in a mutually productive collaboration in the near future.

There has been a close synergy between PA1 and RT-Timeliness. Their interest in timing has resulted in DSTL nominating this topic to be one with which they wish to interact closely and one of the DSTL RAs is located at York.

PA1 involves computer scientists, psychologists and sociologists.

Impact of Organisational Culture and Trust on Dependability – PA2 This activity has been the host to extensive ethnographic studies. PA2 (and

PA3) have exhibited very strong inter-site collaboration especially between Lancaster and Edinburgh. Within Lancaster, there has been a very productive interaction between computer scientists and ethno-methodologists. Among other things, this has resulted in the prototype tools *Scavenger* and *Strider*. The aim of *Scavenger* is to ease the process of building socio-technical system models and populating them with information derived from ethnographic studies. This is achieved through the conversion of raw, loosely formatted data produced by ethnographic activities into structured entity-relational models. As such, Scavenger is specialised to the task of data extraction from ethnographic source artefacts. *Strider* is a configuration modelling support tool which utilises the Scavenger toolset. Strider is used to build a configuration model of the static components and relationships within a system at a particular point in time. Once in place, the model created may be used to aid communication between system stakeholders, permit direct browsing and investigation by various interested parties, as well as being the basis for various automated analysis features.

An understanding of different requirements of Dependability (see also PA7) has arisen from the fieldwork. For example, a hospital bed management system will never be completely accurate but humans adapt to understood levels of inaccuracy and have workarounds when more precision is required.

A major outcome of PA2 is the "Trust in Technology" book being written jointly with PA3.

Dependable Deployment and Evolution – PA3 The Project Proposal foresaw PA3 tackling "Design for Dependability" but it became apparent that is was difficult to define what this left out of the overall DIRC objectives. At the time of the second meeting of the SC, it was clear that it was necessary to refocus PA3 even though excellent work was going on within the original description. The SC shared the view that a narrower scope would yield more coherent results.

After that meeting, it was decided that PA3 should focus on the deployment of generic systems and the related issue of coping with evolution. Both of these are major problems for Dependability. Campbell-Kelly [CK03] points out how much software is now deployed by the instantiation of generic systems. In some ways, this might be an aid to system evolution as requirements evolve but this begs the question of whether the original generic system is adequately flexible and puts a requirement on the user to make suitable provision for evolution (see PA3 ideas on "co-realization" [HPS+02]). Furthermore, such two stage creation of systems brings with it its own set of Dependability concerns.

A major outcome of PA3 is the "Trust in Technology" book being written jointly with PA2. This being one indication of the first-class collaboration between Lancaster and Edinburgh.

Decision Support for Dependability – PA4 There are many occasions where it is necessary to reason about Dependability. The best known instance is in the construction of "safety cases": is there a clear argument which justifies deploying a system? The huge state spaces of large software systems rules out exhaustive testing and their discrete nature means that arguments must always be considered carefully. Research in PA4 has looked at "dependability cases" which tackle the more subtle question of how confident one can be in an assertion

about a system.

There has been a very productive debate in PA4 about whether quantitative arguments are always possible (or indeed, most appropriate). The synergy with the RT on Diversity has been invaluable. PA4 has played the key "conscience" role of making sure that the whole of DIRC does not focus on the process of building systems. There are many points in time where it is necessary to reason about the Dependability of a system. Assessment of a completed system is one such occasion but examples like experimenting with models of bug removal in PA5 or PA8 have led to very good interactions across all sites.

Dependability Issues in Open-Source Software – PA5 This activity was unique in our original plan because it was intended to be a brief survey lasting only 12 months. Given the amount of "hype" surrounding "open source" at the time we were drafting the DIRC proposal, we considered it worth establishing whether this was a "silver bullet" that could significantly improve Dependability. It quickly became clear that the term "Open Source" was used differently by various people: some focused on the licensing model, many meant only that the source code was visible, far fewer identified the motivation and redundancy (of reviewing) aspects which looked as though they could actually change what was developed. Because of this range of meanings a paper [GLA01] was written on the "Many Meanings of 'Open Source' ". SourceForge was searched for interesting projects; ethnographic studies were conducted (cf. [MRR02]); and models of reliability growth were researched (cf. [PS01, PSL00]).

An analysis of different modes of working between projects described as "Open Source" makes it clear that there are many technical, psychological and sociological facets and –furthermore– that most of these could be deployed in "conventional" development projects. A key conclusion for the DIRC project itself was that it would not be meaningful –for example– to compare "open source" development with (conventional) "closed source" development. PA8 (see below) was set up to investigate *specific* topics of the way systems are designed and developed. Other papers include [AGL01, BLNS02].

The final report of PA5 [ABGR02] was presented by Cristina Gacek to the January 2002 meeting of the SC. Fred Schneider encouraged DIRC not to drop the topic of open source and we heeded this advice: recently we have a spate of MSc projects in the area and Budi Arief keeps a watching brief on the subject. Brian Randell and Cliff Jones are also involved as reviewers of Peter Neumann's CHATS project. A public workshop was organised in Newcastle in February 2002 and the proceedings are available as [GLA02]. It was interesting to compare our conclusions with those of a (UK) MoD study at about the same time.

Security and Privacy in Computer Based Systems – PA6 This research was nominally started at the beginning of 2002 when Peter Ryan joined DIRC. But Peter was actually funded from the EU FWP-5 project MAFTIA until the end of February 2003. So PA6 really only had significant manpower from the beginning of 2003 when Jeremy Bryans joined DIRC. Prior to that date, Denis Besnard and Budi Arief were pursuing an interdisciplinary look at "why hackers hack" — this was not a full time activity for either of them.

During this year, the research on PA6 became even more interdisciplinary with the initiation of a case study on the Chaum e-voting scheme and this will

be continued as a Targeted Activity.

Peter and Jeremy also worked together with the main proposers of a major e-science demonstrator called "Gold". This project looks at issues of "Virtual Organisations" among which Security is of course a major issue. Not only is "Gold" now funded, but DSTL are interested in using this as a joint demonstrator project.

Dependable ubiquitous computing in the home – PA7 This activity is about providing support for elderly or disabled people in their homes. The computer support envisaged is not "toys for boys" but is intended to maximise the chance that a person can continue to be safe in their own space (rather than occupy a nursing home or hospital place).

Rob Witty noted from his contacts with HSE that there is an enormous need to reduce accidents in the home. PA7 is a major collaboration between York and Lancaster; it brings together psychologists and ethno-methodologists and there is a hope to construct some software tools possibly with support from Newcastle.

Several new issues have arisen. For example, the obvious saving to society (and increased fulfilment of the individuals) of allowing someone to stay in their own home is not necessarily matched by adequate funding to that part of Social Services which would specify and provide aids in the home. Once again (cf. PA2) we find different notions of Dependability are required in the home. Our social scientists have also observed that the "confidentiality" requirements are not limited to concerns about external people. A person's dignity might make them sensitive to information becoming visible to family members.

This is a fruitful and interesting area with great social significance. PA7 will define a series of Targeted Activities.

Effective Collaboration in Design – PA8 Anecdotally, there are enormous differences between the productivity and accuracy of members of a software development team. These differences (factors of ten and above) are so large as to be interesting in their own right. But it is also clear that understanding the differences might have a major impact on the development process and the Dependability of the created products.

Weinberg in his old but seminal book [Wei71] indicated that software development was not a single activity and that different types of mind might be better equipped to design programs, to code them, to design testing strategies, to find errors and to correct them. Clearly, one could set a technical task as a predictor of technical ability. But would it not be interesting to know if some other factors correlated with ability in one or more of the software activities?

PA8 is investigating such questions. Close collaboration between psychologists and computer scientists have yielded the first results. Colleagues in Newcastle have received input from psychologists at City; experiments from one of these locations are being validated in the other.

Clearly, there will be a need to move beyond the current stage of individual activity and at this stage we hope to involve sociologists more actively. Tony Lawrie's PhD studies on "objective setting" are also in this category.

PA8 will define a series of Targeted Activities.

Dependable Service-Centric Grid Computing – PA9 PA9 is specifically funded by e-science and is described in Appendix C.9. There needs to be a careful analysis of its fuller integration into other DIRC activities because of the need to meet specific e-science deliverable commitments.

3.3 Research Themes

The intention is that each Research Theme (RT) should act as a way of gathering, analysing and recording the lasting knowledge that comes out of other DIRC research. We have said several times that, while PAs should produce a steady stream of papers, a Research Theme is more likely to produce one or more books. We have however made clear that we do not believe that "one size fits all" and where one RT could produce a single authored book, another might produce more than one collection of edited essays. The Themes differ; so can their output.

We do expect the people involved in RTs to "act as a conscience" for their topic in the sense that they should make sure their topic is not overlooked in a PA which has the opportunity to contribute to the knowledge.

It should be remembered when reading about RTs that one of the motivations in selecting the topics to be pursued in RTs was that it should be possible (and interesting) to look at them from both a technical (system) and human (user) viewpoint. Furthermore, some of our hopes to make progress on these difficult themes was that we might be able to deploy ideas from say the psychological or sociological research to the technical issues (and, of course, *vice versa*).

DIRC's original manpower plan showed a steady rate of about 20% of RA time being involved in Research Themes (RT). This has not yet been realised (the actual use of RA time on RTs in the first three years is close to one eighth). This was almost certainly optimistic for two reasons. While one fifth of the effort might be used over the full lifetime of DIRC, we should have planned a significant wedge (rising from less than 5% in the first period to something like a third in the final period). That is in the nature of synthesising knowledge from the whole project. Moreover, there is reason to suspect that a significant proportion of the work on presenting these deeper conclusions will fall to the more senior people who are involved in DIRC and "faculty" time does not show up at all in our EPSRC budget.

We have several times taken steps to ensure that the RTs are kept in people's minds. For example, they were made the main focus of the DIRC Easter Workshop in 2002. One of the major reasons for the restructuring of DIRC (cf. Section 5) is to increase the extent to which the RT agenda dictates other research in DIRC.

How then is the "gathering of wisdom" going?

It is useful to start with an example of our level of ambition. Donald MacKenzie's book [Mac01] is a perfect example of someone from one discipline making a major effort to understand research from other domains (in this case Mathematics and theoretical Computing Science).[9] This profound book is of interest to sociologists and to the disciplines studied. It has been widely and

[9]Although this was begun well before DIRC started, Donald credits EPSRC's funding of DIRC with having supported his research.

positively reviewed. Partly aided by DIRC, Donald is now looking at Financial markets and their failure modes – a talk (by Stuart Anderson) was given at the third meeting of the SC. It would be hubris to expect each RT to produce a book of this calibre. But –as a target– it is certainly a "Needham failable objective".

We are most confident about the RTs on Timeliness and Diversity producing authored books: both RTs have benefited from closely allied PAs and core researchers who have long focused on the topics.

Timeliness The Project Proposal states:

> The specific topic of Timeliness poses special concern for the safe use of systems ... there are interesting contrasts between the ways in which a system can be developed to meet guaranteed timing constraints and the expectations that one can put on the human users of a system ...

York has done research on the technical aspects of Timing over many years. Their research spans from technical aspects of notations or logics for time through to the human perceptions of time. They have done research on worst-case execution and have pioneered research on formal approaches to HCI design.

Alan Burns was already able to give a "timing roadshow" seminar at all DIRC sites[10] in 2002 and this "roadshow" format is one which will be emulated by other RTs. The report in Appendix B identifies the interdisciplinary issues including the way that humans can adjust to changes in routine. The "Parcel-Call" activity has led to strong cross-site collaboration with Edinburgh. One interesting issue that arose in Alan's seminar at Newcastle is the need for notations which make it easier to talk about cyclic behaviour. Alan Burns is about to start a long sabbatical and it is hoped that this will result in the framework for a book for the Timeliness RT.

Diversity The Project Proposal states:

> An obvious way to limit the effect of errors in systems is to introduce redundancy; although choosing how to deploy redundancy (and select between results) is non-trivial, this area is reasonably well-understood. If one is guarding against mechanical decay, it is enough to use multiple instances of the same design. If, however, one is trying to protect against design errors, there is a need to employ diversity. The informal use of diversity to achieve dependability in human affairs is ubiquitous and age-old. In IT systems significant effort is being put into understanding design diversity — how it can be measured, assessed and maximised. But it is clear that diversity can be applied in much more general forms such as the elicitation of requirements using the "diverse" viewpoints of many domain experts, and the use of "independent" argument legs in safety cases. However, there have been few formal studies of the efficacy of diversity or of ways in which it can best be deployed. By studying wider applications of diversity –between humans and computers, between

[10] A taste of this was presented at SC-3.

> different procedures and intellectual approaches, between more than one human, and so on– we will develop a better understanding of the contribution of diversity to dependability ...

Diversity has long been a focus of many researchers at City University and the RT has been greatly helped by the close association with PA4. The RT on Diversity has made a considerable effort to educate the rest of DIRC about the issues concerned. The City researchers (including their busy senior faculty members) have made a point of trawling other PAs for interesting problems: their role as conscience of their theme is perfect. They became strongly involved in PA5, in the psychological experiments within PA8 and most recently with PA6's look at the "Chaum e-voting scheme". The activities on the "Mammography Case Study" have been strongly interdisciplinary.[11]

In the next period, Bev Littlewood and Cliff Jones intend to spend more concentrated time looking how to combine evidence from testing, model checking and "verification". Lastly, Meine van der Meulen's location of well over a million programs submitted to an ACM competition site[12] offers scope for some fascinating investigations.

Structure The Project Proposal states:

> A well-chosen structure helps system designers and evaluators to understand a system by allowing them to "divide and conquer" the system's complexity, and ensures that any constraints imposed by the structure do not impose unacceptable overheads on the operation of the system. However, the issues of structure are different in computer and human contexts. It is understood how to deploy redundancy to protect against failures of physical components; although there is more research needed, there are already proposals for deploying related ideas to guard against some classes of design errors, and Newcastle is just initiating a Framework 5 project on how to guard –by system structuring– against malicious attacks. But it is difficult to apply the same techniques in a systematic way to human behaviour, and hence the issue of structuring a complex computer-based system as a whole is still a great challenge. Indeed, the IRC will probably have to tackle the issue of languages to describe human behaviour in order to reason about overall system architectures which will guard against human errors.

Although the project has experts such as Cristina Gacek on the Structure (aka Architecture) of technical systems, it is not obvious how these sorts of results can be applied to human systems. Some sociologists do categorize organisations but they do not use anything like an ADL (cf. [SG96]) to describe their findings. Somehow, we have to make progress in this area. Gillian Hardstone's use of Mary Douglas' research on *cultural composition of groups* does look interesting. The work of the Ethno-methodologists provides ample observational material of a given setting. But moving from this to something which can be used in the design process is problematic.

[11] An outline of this work was presented to SC-3.

[12] MMM = "Meine's Million Modules".

Researchers involved in DIRC were also active in the EU Framework-V projects referred to in the Proposal and have contributed to progress in the purely technical areas

of "Dependable Systems of Systems" (DSoS' "Conceptual Model" [GIJ+03]; Peter Ryan contributed to the equivalent document [PS03] for the MAFTIA project. The Laprie/Randell [Ran00] definitions of "fault", "error" and "failure" are reconsidered with a light touch of formalism in [Jon03]).

One avenue which has benefited greatly from DIRC's SVF programme is the research on "Faults as interference" – see [HJJ03].

As well as planning a (technical) book on Operational Semantics Cliff Jones is personally committed to writing a book on "Structure" but it will clearly be necessary to get more engagement from other disciplines.

Responsibility The Project Proposal states:

> ... the topic of Responsibility is currently more strongly associated with society but here again it is essential to understand how legal notions carry over to computer-based systems, and indeed how they have some analogies within complex hardware-software systems. It is therefore important to show clearly the ways in which the acceptance, recording and discharge of responsibilities are reflected in the technical systems which mediate social relationships, particularly in the presence of social and human failure. The general problem is knowing how to draw boundaries around domains of distinct responsibilities, the scope of differing sets of ethical principles, the limits of various levels of trust, and how to reflect these boundaries in an information system. Understanding these issues will help us design complex systems that can support the interaction of the participants in that system and their constant renegotiations of role. IRC research in this area will both look at ways to express responsibility structures and their use in the design of dependable systems.

It is clear that understanding notions of Responsibility in professional organisations like hospitals is fundamental to designing systems which will actually be used once deployed. It is also clear that this RT is extremely challenging. We are delighted that the combined work on PA2 and PA3 has already led to a book which is a carefully targeted collection of contributions. The book *Trust in Technology: A Socio-Technical Systems Perspective* is discussed in the Responsibility section of Appendix B. The work referred to there on recording Responsibility is a start but much remains to be done. It is, at least, reasonable to expect a further collection of essays from this RT.

Risk The Project Proposal states:

> Touching all of the issues of dependability is the central notion of Risk. Without scientific ways to understand, quantify and reason about risk there can be no real discussion of dependability. Most technical work on risk has concentrated on identifying undesirable events and assessing the severity of their consequences and their
>
> probability of occurrence. This technical view of risk is challenged by new computer-based systems that are deeply embedded

and widely distributed in our culture and by the difficulties encountered by decision takers in utilising technical risk information. We are convinced that it is necessary to take a broader view of risk in which technical developments and issues of the decision taking process are developed hand in hand. This broader view has to embrace "secondary dimensions" of risk such as whether the risk is undertaken voluntarily, whether the risk has catastrophic consequences, to what extent the consequences of the risk are reversible, etc. We believe that these are important bridges between formal, probabilistic, assessments and the decision taking context. The IRC will continue research established by City University on quantitative aspects of risk assessment and consider how these link to human questions about the degree of choice and knowledge that someone who is at risk has of the consequences of "using" a system and the degree to which the consequences can be redressed.

There is an extensive report on the Risk RT in Appendix B.5. It emphasises the need to focus with such a huge topic. The Edinburgh work is strongly informed by social science and there is excellent collaboration with the statistical expertise on Risk Assessment at City. The intention to cover individual decision making should fit well with Peter Ayton's work on psychological confidence. Donald MacKenzie's work on financial markets fits into this RT.

3.4 External visibility

In addition to standard publication activity, DIRC has created avenues for external dissemination of its research to other researchers and it has also taken opportunities to increase public awareness.

The IRC held its first public workshop in its ninth month of existence. This was much earlier than we had promised in the proposal. The workshops which have been public in the sense of calls for contributions and attendance to date are

- Dependability in Healthcare: Edinburgh 2001-03-22/23 (coordinated by Robin Procter and Mark Rouncefield)
- Workshop on Open Source Software Development: Newcastle 2002-02-25/26 (coordinated by Budi Arief, Cristina Gacek and Tony Lawrie: proceedings [GA02])
- Workshop on Ethnography, Systems and Strategy: Lancaster 2002-04-22/23 (DIRC contact John Dobson; joint with the AMASE project)

Another research interchange co-organised by Robin Bloomfield and Stuart Anderson (together with M. Heisel and B. Krömer from Germany) was the Schloß Dagstuhl event in November 2002 on "Dependability of Component-Based Systems". Cliff Jones and Massimo Felici also attended[13] and spoke at this seminar. Each of the DIRC attendees at this event became dramatically aware of the extent to which our attitude had become "interdisciplinary": it

[13] As did Meine van der Meulen who has now joined City and John Rushby (SVF).

was as though we were speaking a different language than nearly all of the other attendees.

EPSRC itself publicised them when the five "IT-centric" IRCs were initiated. Either this, or the press releases from our universities, triggered a New Scientist article about DIRC, half a dozen brief radio interviews with Cliff Jones and an article in the glossy "MIS" magazine. It would be useful if EPSRC were to run a further series of articles on the five IRCs once all of their mid-term reviews have been completed. In the first half of DIRC, we have been opportunistic about ventures into

"public understanding of science"; we need now to review how active we should be in this arena during the second half of DIRC.

Between general public awareness and specific research publication there is a key group of potential developers and users of IT systems for whom Dependability is a major issue. To raise the awareness in such groups, Denis Besnard (Newcastle) has committed some of his time to taking part in exhibitions such as SITEF (Toulouse, 2002) and CEBIT (Hanover, 2003). To date he has generated easily accessible handouts and posters and attended three such events at minimal cost to the project (the (North-East) "Regional Development Agency" has willingly subsidised Denis' travel costs).

4 Lessons we have learned

EPSRC's Terms of Reference for the Mid-Term reviews (cf. Appendix E) make clear that a prime objective is to inform their Council of benefits and pitfalls of the IRC mode of funding. This section seeks to provide input from our experience in DIRC towards this important aim.

It should already be clear that members of the DIRC team consider that the coherence provided by six year's funding has been essential in tackling Dependability in an interdisciplinary way.

4.1 About Interdisciplinary Research

We took on an ambitious challenge in the breadth of disciplines involved in DIRC. Attitudes have developed, collaborations exist and are paying off with entirely new insights. Perhaps one of the most telling comments is that it is now far less obvious from a discussion as to what was someone's original discipline. But we are aware of how much there is still to do to achieve the objectives in Section 6. A little introspection is probably good for all of us!

One of the points we made in the DIRC proposal was that most of us had had some exposure to working across discipline boundaries. To give just one example, researchers at York had pioneered research between Computer Scientists and Psychologists on HCI design. We were therefore not naively expecting interdisciplinary research to be without problems. In fact, in addition to our own experience, Cliff Jones and Brian Randell discussed the wisdom of casting the DIRC net so wide with John Goddard who had led the earlier ESRC PICT Programme and had formed the Newcastle "Centre for Urban and Regional Development Studies" (CURDS).

John warned us of two major fault lines that, interestingly, do not fall along the normal perimeters of the academic subjects. The first might be characterised

as between those people who use numbers and those who use words. Within DIRC, this could be sharpened to researchers who look for a topic that can be reduced to symbols and those who prefer to write (often long) essays about an issue. The second division that John forewarned us of was researchers who were happy to intervene and those who refuse to do anything other than observe. We have seen both of these divisions and they do not for example delimit the subjects normally known as "psychology" and "social science".

We had, of course, anticipated some difficulties in terminology between people who had studied and done their research in different disciplines. But there is something much deeper: different disciplines have different "values". Cliff Jones crystallised this by an observation on the reaction of one of his colleagues who is a psychologist: "their reaction to receiving a bunch of questionnaires was like mine when I've proved a theorem". It is as difficult for an ethno-methodologist to appreciate the pleasure that someone receives from a theorem as it is for many computer scientists to delight in a "two tailed" distribution.

How have we learned to live with each other's different values? Partly by being explicit about them; partly by natural tolerance (of most of the DIRC team). Much of this does come down to personal relationships. Some researchers (to mention some different examples: Stuart Anderson, Corin Gurr and Ian Sommerville) had previously crossed discipline boundaries. It is interesting that all of the PIs were Computer Scientists and that this has led to no significant tensions. In fact, there are occasions when the computing background appears to find it more natural to talk to both psychologists and ethno-methodologists than for people from these two disciplines to talk to each other.

Some of the DIRC team had previously worked in Operations Research (OR)[14] and it is from here that we could claim one transferable experience. The fastest way to cut across discipline boundaries is to have a mixed team look at a specific application. We have seen this several times and should perhaps have tried it even earlier than we did. The study initiated in PA4 to look at the multiple reading of mammograms brought together sociological observation of actual teams, statistical analysis of a carefully controlled experiment, analysis of the safety case, and will be enlarged to look at other aspects of a key medical procedure. Similarly, the e-voting scheme put together by David Chaum has been considered more recently in two workshops which have involved computer scientists (building formal models of the algorithms), mathematical insight on the algorithms themselves and psychologists and sociologists looking at how potential voters might react to such a scheme.

We have built an exciting interdisciplinary team in DIRC. If any of us were called independently to help with the creation of a major (Dependable) system, we would be sensitive to those aspects which were outside our own discipline and would know when to call for help. This mutual respect and understanding is a more realistic goal than the creation of universal experts who can work as effectively as someone from any discipline.[15] To design a major Computer-Based

[14] Operations Research teams in –for example– oil refining will bring together a group with different specialities; they will focus on a problem (e.g. an industrial chemist, mathematician and programmer might look at maximising crude oil processing through distillation by mixing light and heavy crude oils); any difficulties in terminology and values melt in the face of needing to come up with a concrete result.

[15] We went through a period of discussing to what extent (and how) a researcher from one discipline should strive to become completely conversant with another area. This was an

system will require people whose training is from the different disciplines.

Among the further public workshops we are considering is one on "Interdisciplinarity". This might be driven just by DIRC or possibly be combined with AKT. If this could be done in the relatively near future, there would be an opportunity to revisit this topic with even more of the IT-centric IRCs before their original six-year funding window is over.

4.2 About communication

As well as needing to find ways across discipline boundaries, the "Interdisciplinary Research *Collaborations*" have to interact over geographic space. Of course, we knew each others' strengths at the outset of the IRC and know that we can derive considerable benefit from researchers at the other locations. It is perhaps worth reviewing some of the advantages and disadvantages in such distributed projects.

The clear advantage is the wide range of expertise we can call on from the five sites. None of the universities involved could offer a similarly productive way of using such a large grant; no one of them could even approach the range of experience and talent available to DIRC. Furthermore, it is very unlikely that such a large group of RAs could have been located in one of the universities. There are people who wish to work nowhere other than one of the capital cities; there are others who would not consider doing so.

Members of DIRC have travelled extensively to take part in face-to-face meetings. Each of the PAs has held frequent meetings; every RT has met on occasions; the various boards (notably the EB) have met regularly and frequently. One influence completely outside our control was the tragic railway accidents and their aftermath. Four of the five sites are on the UK "East Coast" rail line and in mid-2000 this was a dependable service. After the three fatal accidents, travelling on this line has been reduced to a game of roulette.[16]

We have used telephone[17] and video-conferencing to a limited extent. On some occasions this has been very useful and saved both time and money. We have had two of our psychologists observe our behaviour during such meetings and we expect these results to be written up; our first experiences with the e-science funded "Access Grid" were not encouraging and we are hoping to repeat the experiments with their improved facilities before appearing to criticise a hastily constructed prototype.

We have used computer support of various kinds (and compared observations on these with the AKT IRC). DIRC's shared file store is built on *Basic Support for Cooperative Work* (BSCW). This system is far from perfect and there have been debates about moving away from BSCW. The PD's position is that nothing better has been demonstrated that meets all of the IRC's requirements. A recent study by two project RAs agreed with this position.[18] It is clear that something like Lotus Notes would offer better support but its needs (in terms of support personnel) are clearly beyond the DIRC budget.

interesting discussion but we believe that our OR-like approach is more viable.

[16] Yes, we do see the irony of this! Furthermore, the record for travel misery actually goes to Ian Sommerville who spent 10 hours on West Coast trains from Lancaster never getting south of Manchester in an attempt to attend an IAB meeting.

[17] With some expensive and fancy pod phones at each site.

[18] There was also a recent and serious "hack" of the DIRC machine on which BSCW runs. We have studied this from the point of view of Security (PA6).

DIRC (in common with other IRCs) faces a shortfall on its travel budget and would like to discuss this at the reviews. We have found that true inter-disciplinary discussion often needs face-to-face contact (and –perhaps– a trip to the pub at the end of the day). Tension is harder to dispel over a video link.

Both the project internal BSCW and the external WWW store "decay" over time and require periods of tedious activity to restructure them.

We have just gone through such a revision for the BSCW structure and that will be followed by a revision of DIRC's externally visible WWW.[19]

4.3 Difficulties

There are a few other difficulties about which it is worth warning future IRCs. The core of the team was a group of senior researchers. They have given unstintingly of their time but they all have other distractions in their roles as permanent members of departments. Bev Littlewood had to take over as acting Dean not only at short notice but also in the terrible period for City just after their fire; Cliff Jones took on the last half of the EU-funded DSoS project on Brian Randell's "retirement"; Stuart Anderson has had both the role of deputy head of a huge division and the role of Director of the Regional e-science centre. There is always a danger that one is faced with the choice between busy and successful people or ones who might be lower on both scales! There is clearly a problem here and it would be worth EPSRC looking at "buying out" the time of up to one person per partner in future IRCs.

Perhaps the most delicate balancing act for members of the DIRC EB was what to do about bidding for further funds for the IRC itself. As indicated in Section 1.2, DIRC was funded at below the level for which it was planned. David Clark of EPSRC Swindon was extremely supportive and suggested that a way forward was to apply under the joint scheme with (then) DERA. Unfortunately, DERA were just starting their long phase of (semi) privatisation and these negotiations took much longer both in elapsed time and in the effort of the PIs than could have been foreseen. The eventual £270K from DSTL is most welcome and is leading to productive collaborative work but the uncontrollable factors reduced the "return on investment".

Almost the same story *mutatis mutandis* applies to DIRC's bid to the e-science programme. The real "loss leader" has been bidding to the EU FWP-6 programme. The frequent and massive changes of advice coming from Brussels made this a much more wasteful process and took a major slice of Robin Bloomfield's time.

It is difficult to see how EPSRC could either restrict or aid such bidding. In an ideal world, they might award 125% of the requested budget; but we do not live in a world where this is likely to happen.

5 Restructuring

To date, most of the effort in DIRC has been on the Project Activities (see Section 3.2 and Appendix C). The PAs are actually large by the standard of EPSRC/ESRC-funded responsive mode projects in our fields. The first group of PAs were started as soon after July 2000 as hiring of RAs allowed.

[19] There are some interesting challenges for AKT!

The four longer term PAs (PA1–4) are either now just finishing or are due to finish within a year. As foreseen in our Proposal, we have started new PAs (PA6–9). Thus the default would be to continue as over the first three years of DIRC with PAs acting as units of work within the overall programme of our IRC. Although this would doubtless result in a further collection of valuable publications, we believe that we can achieve an even better result by changing our organisation. In particular, we want to operate in a way which makes it more likely that we achieve the long term technical objectives embodied in the Research Themes and be in a position to react more flexibly to "outside engagements".

We hope that the new *Targeted Activities* will be more focused and responsive than the PAs; we are aware of some risks which are discussed below.

The remainder of this section sets out our restructured organisation and Section 6 explains the refocused technical objectives.

5.1 Targeted Activities

We intend that *Targeted Activities* (TAs) will be shorter than the multiple year format of PAs. A TA will typically run from 3–12 months. More importantly than just the duration, a TA will normally be prompted by an identified need from a Research Theme (RT). This is not to say that there will be a top-down design of the remainder of DIRC but there will be a stronger emphasis on seeing how activities feed into the long-term goals of the themes. Proposers of TAs will also identify those aspects of dependability[20] to which they relate. Of course, in the spirit of DIRC, any TA will bring together different disciplines and will normally involve more than one site.

At its meeting of July 15th, the EB indicated that preliminary discussions on the following TAs[21] indicated that they would be likely to be among the first to get resource allocated to them:

- Complete the Trust book: the book which has grown out of PA2/3 is near completion
- Mammography case study: several papers have already been written but the work would have more lasting impact if it were carefully stored as a future resource
- Chaum Study: this activity –which started within PA6– has already generated strong interdisciplinary interaction
- Cooperation with the "Gold" project which is looking at virtual organisations and their use in the Chemical industry (this will also form the basis of a joint case study with DSTL)
- York's case study with DSTL
- Plan an engagement with Philips Medical

[20] For example: reliability, integrity, correctness (with respect to a specification), availability, safety, security and privacy.

[21] There was a strong request that TAs should not be numbered (thus denying some people of the one-upmanship of discoursing for 10 minutes and mentioning PA numbers in each sentence).

- Bev Littlewood and Cliff Jones intend to look again at building "fault freeness" and dependability arguments from different evidence
- Produce a top level "Dependability syllabus" and (potentially) define other TAs to provide detailed course material for some courses
- Continue the experiments on psychology of programming

The proposed restructuring has been discussed at three EB meetings, was presented at the 2003 (Edinburgh) Easter Workshop, and (as a result of a request there) the PD has visited each site and three PA meetings to review the idea. In these discussions, a number of risks were identified against which we must guard.

1. there is a danger that TAs are narrower than PAs involving less disciplines;
2. there is a danger that TAs are geographically focused involving only one site;
3. since TAs are short, we must avoid excessive start-up costs;
4. we must watch the management overheads in TAs;
5. longitudinal studies are extremely important and some planning of "impact reviews" as TAs might be needed.

5.2 Making Research Themes more concrete

It is clear that exhortation alone will not cause progress on the RTs. Ideas to make it easier to see progress include:

- those involved in RTs should undertake "road shows" to all DIRC sites
- inviting RTs to propose Targeted Activities
- where an authored book looks too far in the future, an edited collection of essays should be planned
- plans for books (authored or edited) should be reviewed at EB meetings
- the EB will call for progress reports to be available at each EB meeting

6 The way forward

The preceding section addresses DIRC's organisational structure; we now focus on some of our chosen goals. Our key research goal is to devise "development methods" which increase both Dependability and our (justifiable) confidence that this property has been achieved. We recognise that this is very ambitious. Although the IT-centric IRCs are very large by the standards of normal computer science research funding, their total budgets are small by industrial standards. Furthermore, there is not a distinct "software industry" (cf. [CK03]) to which researchers can communicate new ideas. Roger Needham encouraged us at the first SC meeting to be ambitious; but over-ambition could be counter-productive in that an IRC will "spread itself too thinly". It is therefore necessary to plan a strategy of both research and the attempts to influence potential users of such research.

6.1 Key objective: towards DIRC methods

DIRC's "vision" is to

> Build more dependable (computer-based) systems at lower cost

The way in which we hope to achieve this is by defining a *family of development methods*. Notice that we do not expect to come up with one single "DIRC method" which will work for all systems. Rather, we hope to identify or create various methods which might well be combined with other approaches and used in varying combinations in different applications.

It is useful to look at existing development methods and what they offer even if they apply to rather narrow technical domains. This identifies some useful characteristics. If one looks at one or other of the formal development methods for programs such as B [Abr96], one finds a way of recording what system should be built; of documenting (a range of) design decisions; and of showing that the design satisfies the specification. It is not always cost-effective to undertake formal proofs and –indeed– the major impact of "formal methods" is likely to be in helping clean up the early stages of design; but –in the extreme– the question of correctness can be settled by proof.

Before we go on to consider to what extent anything approaching this is possible with the wider notion of "computer-based" system of concern to DIRC, there is yet another desirable characteristic to be mentioned. Michael Jackson in his writings from [Jac83] to [Jac00] has sought methods which are *normative* in the sense that two people using the method on the same application would be likely to arrive at similar solutions.

What DIRC has achieved to date is a certain shared "attitude" or "culture". The most ardent formalist would not for a moment believe that usability can be achieved just by writing a formal specification. There is an "Operations Research attitude" in that all in DIRC would see the need to use people from a range of disciplines to create a usable C-BS.

At some level, all members of DIRC recognise that *abstraction* is a key tool. Obviously this is what we were doing when we sketched a formal model for the Chaum e-voting scheme but ethno-methodologists also choose to leave facts out of their records. There is obviously a wide range of opinion as to the possibility (and perhaps desirability) of being able to record ideas in formulae. The extent to which formal models can be useful in recording things like responsibility relationships is still in discussion. Perhaps the motto here should be "a little formalism goes a long way". It is probably true that graphical aids will be necessary to see such ideas widely used. It is certainly the hope that the work in [HJJ03] on enlarging the scope of system considered can be extended to tackle the sort of issue raised by [Rus99] perhaps by folding in ideas from [Rea90, Rea97]; the identification of hidden assumptions is crucial to Dependability.

It is tempting to describe the challenges of *assessment* with the purely technical systems. But it is of course necessary to look at a complete computer-based system in order to measure how well it performs. But, just like other aspects of quality, assessment cannot be considered only after a system is built. The need to be able to assess a system might need to influence its architecture.

The next three years of DIRC will focus on development methods. We recognise that there is a huge job to do in order to bring together insights from sociologists with the narrower view of "methods" which can apply to purely

technical systems. We believe that we have formed a unique team and stand a better chance of making progress than anyone else.

6.2 Impact

We intend that the major research impact of DIRC will be the publications and books it creates. But we are also keen to find avenues to demonstrate the evolving ideas.

Through its ILD, DIRC is trying to set up mutually productive projects with NATS and Philips.

It is gratifying that both Martyn Thomas and Cliff Jones were involved in the first Foresight consultation exercise on "CyberTrust". Martyn is on the Steering Committee and Brian Collins (who is coordinating the activity for DTI) has already asked for one submission from DIRC.

DIRC also needs to get inside both HMG's Procurement process and various Standards bodies.

6.3 Further external workshops

DIRC will consider workshops on at least the following topics

- Legal (and possibly economic) influences on Dependability
- Inderdisciplinary Research
- PA2/3
- the "(UKCRC) Grand Challenges" on *Dependable Systems Evolution*
- Dagstuhl seminar (4181) on "Atomicity in System Design and Execution" is scheduled for April 2004

References

[AB03] L.B. Arief and D. Besnard. Technical and human issues in computer-based systems security. Technical Report CS-TR 790, University of Newcastle Upon Tyne, March 2003 2003.

[ABGR02] L.B. Arief, D. Bosio, C. Gacek, and M. Rouncefield. Dependability issues in open source software – DIRC Project Activity 5 Final Report. Technical Report CS-TR 760, University of Newcastle Upon Tyne, February 2002.

[Abr96] J.-R. Abrial. *The B-Book: Assigning programs to meanings*. Cambridge University Press, 1996.

[AGL01] L.B. Arief, C. Gacek, and A.T. Lawrie. Software architectures and open source software; where can research leverage the most? In *Proceedings of the 1st Workshop on Open Source Software Engineering*, pages 3–5, Toronto, Canada, 2001.

[And01] Ross Anderson. *Security Engineering*. Wiley, 2001.

[APA+03] A. Alberdi, A. Povyakalo, E. Alberdi, L. Strigini, and P. Ayton. Decision support or automation bias? a study of computer aided decisionn making in briest screening. In *Proceedings of SPUDM*, 2003. in print.

[BA03] D. Besnard and L.B. Arief. Computer security impaired by legal users. Technical Report CS-TR 794, Department of Computer Science, University of Newcastle Upon Tyne, 2003.

[Bes01] D. Besnard. Attacks in IT systems: A human-factors centred approach. In *Proceedings of the International Conference on Dependable Systems and Networks (DSN-2001): Supplement*, pages B–72, Gøteborg, Sweden, 2001.

[BL03] R.E. Bloomfield and B. Littlewood. Multi-legged arguments: The impact of diversity upon confidence in dependability arguments. In *Proceedings DSN 2003*. IEEE Computer Society, 2003.

[BLNS02] D. Bosio, B. Littlewood, M.J Newby, and L. Strigini. Advantages of open source processes for reliability: clarifying the issues. In *Proceedings of the Open Source Software Development Workshop*, University of Newcastle upon Tyne, UK, 2002.

[Blo03] R.E. Bloomfield. A descriptive model of failures in complex systems. In *Proceedings of DSN 2003*, pages B76–77. IEEE Computer Society, 2003.

[BR03a] Jeremy Bryans and Peter Ryan. A dependability analysis of the Chaum voting scheme. Technical Report CS-TR-809, Newcastle niversity, 2003.

[BR03b] Jeremy Bryans and Peter Ryan. Security and trust in the Chaum voting scheme. In Theo Dimitrakos and Fabio Martinelli, editors, *Formal Aspects of Security and Trust*, 2003. submitted.

[BTN+02] G.D. Baxter, K. Tan, S. Newell, P.R.F. Dear, and A.F. Monk. Using rich pictures as a tool to facilitate the development and acceptance of an expert system in the neonatal unit. In *Proceedings of the Health Care Meets Medical Informatics and Innovation Conference (HCMMII02)*, page 36, 2002.

[BTN+03] G.D. Baxter, K. Tan, S. Newell, P.R.F. Dear, and A. Monk. Analysing requirements for decision support in neonatal intensive care. *Archives of Disease in Childhood*, 88(1):A46, 2003.

[Cha] David Chaum. Secret-Ballot Receipts and Transparent Integrity: Better and less-costly electronic voting at polling places. http://www.vreceipt.com/article.pdf.

[CK03] Martin Campbell-Kelly. *From Airline Reservations to Sonic the Hedgehog: A History of the Software Industry*. MIT Press, 2003.

[GA02] C. Gacek and B. Arief. Proceedings of the open source software development workshop. Technical Report CS-TR-812, Newcastle niversity, February 2002.

[GIJ+03] M-C. Gaudel, V. Issarny, C. Jones, H. Kopetz, E. Marsden, N. Moffat, M. Paulitsch, D. Powell, B. Randell, A. Romanovsky, R. Stroud, and F. Taiani. Final version of DSoS conceptual model. Technical Report CS-TR: 782, School of Computing Science, University of Newcastle, July 2003.

[GLA01] C. Gacek, A.T. Lawrie, and L.B. Arief. The many meanings of open source. Technical Report CS-TR 737, Department of Computer Science, University of Newcastle Upon Tyne, 2001.

[GLA02] C. Gacek, T. Lawrie, and L.B. Arief. Interdisciplinary insights on open source. In *Proceedings of the Open Source Software Development Workshop*, pages 68–82, University of Newcastle upon Tyne, UK, 2002.

[HH02a] M. Hildebrandt and M. Harrison. The temporal dimension of dynamic function allocation. In *Proceedings of the 11th European Conference on Cognitive Ergonomics*, pages 283–292, Catania, Italy, 2002.

[HH02b] M. Hildebrandt and M. Harrison. Time-related trade-offs in dynamic function scheduling. In C. Johnson, editor, *Proceedings of the 21st European Annual Conference on Human Decision Making and Control*, pages 89–95, Glasgow, UK, 2002.

[HH03] M. Hildebrandt and M.D Harrison. Putting time (back) into dynamic function allocation. In *Proceedings of the 47th Annual Meeting of the Human Factors and Ergonomics Society*, Denver, Colorado, 2003.

[HJJ03] I. Hayes, M. Jackson, and C.B. Jones. Determining the specification of a control system from that of its environment. In *FME 2003*, pages until publication, avaliable as CS–TR–808, Pisa, Italy, 2003.

[HPR+03] M. Hartswood, R. Procter, M. Rouncefield, R. Slack, and J. Soutter. The work of reading mammograms and the implications for computer-aided detection systems. In D. Barber, J. Brady, and E. Berry, editors, *Proceedings of the Seventh Medical Image Understanding and Analysis Conference*, Sheffield, UK, 2003. British Machine Vision Association.

[HPRS00] M. Harstwood, R. Procter, M. Rouncefield, and R. Slack. Order in the machine: Evaluating computer aided detection tools in mammography. In C. Johnson, editor, *Proceedings of the 21st European Annual Conference on Human Decision Making and Control*, pages 205–208, Glasgow, UK., 2000.

[HPRS01] M. Hartswood, R. Procter, M. Rouncefield, and R. Slack. Performance management in breast screening: A case study of professional vision and ecologies of practice. *Journal of Cognition, Technology and Work*, 4(2):91–100, 2001.

[HPS+02] M. Hartswood, R. Procter, R. Slack, A. Voss, M. Buscher, M. Rouncefield, and P. Rouchy. 'Co-Realisation': Towards a principled synthesis of ethnomethodology and participatory design. *Scandanavian Journal of Information Systems*, 14(2):9–30, 2002.

[Jac83] Michael Jackson. *System Design*. Prentice-Hall International, 1983.

[Jac00] Michael Jackson. *Problem Frames: Analyzing and structuring software development problems*. Addison-Wesley, 2000.

[Jon03] C.B. Jones. A formal basis for some dependability notions. In B. Aichernig and T. Maibaum, editors, *Formal Methods at the Crossroads: from Panacea to Foundational Support*, volume 7525 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.

[LAS93] Report of the inquiry into the London Ambulance Service, February 1993. published by South West Regional Health Authority.

[Loe03] K. Loer. *Model-based Automated Analysis for Dependable Interactive Systems*. PhD thesis, Department of Computer Science, University of York, UK, 2003. pending.

[LT93] N. G. Levenson and C. S. Turner. An investigation of the Therac-25 accidents. *Computer*, pages 18–41, July 1993.

[Mac01] D. MacKenzie. *Mechanizing Proof: Computing, Risk, and Trust*. MIT Press, Cambridge, Mass., 2001.

[MRR02] A. Mackenzie, P. Rouchy, and M. Rouncefield. Rebel code? the open source code of work. In *Proceedings of the Open Source Software Development Workshop*, University of Newcastle upon Tyne, UK, 2002.

[MS02] Kevin D. Mitnick and William L. Simon. *The art of deception*. Wiley, 2002.

[Nor88] Donald A Norman. *The Psychology of Everyday Things*. Basic Books, 1988.

[Per99] Charles Perrow. *Normal Accidents*. Princeton University Press, 1999.

[PS01] P. Popov and L. Strigini. The reliability of diverse systems: a contribution using modelling of the fault creation process. In *DSN 2001-The International Conference on Dependable Systems and Networks*, Gøteborg, Sweden, 2001.

[PS03] D. Powell and R. Stroud. Conceptual model and architecture for MAFTIA. Technical Report CS-TR: 787, School of Computing Science, University of Newcastle, January 2003.

[PSL00] P. Popov, L. Strigini, and B. Littlewood. Choosing between fault tolerance and increased V and V for improving reliability. In *International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA' 2000)*, Las Vegas, USA, 2000. CSREA Press.

[Ran00] B. Randell. Facing up to faults. *The Computer Journal*, 43(2):95–106, 2000.

[Rea90] James Reason. *Human Error*. Cambridge University Press, 1990.

[Rea97] James Reason. *Managing the Risks of Organisational Accidents*. Ashgate Publishing Limited, 1997.

[Rou02] M. Rouncefield. *'Business as Usual': An Ethnography of Everyday (Bank) Work*. PhD thesis, University of Lancaster, 2002. pending.

[Rus99] John Rushby. Using model checking to help discover mode confusions and other automation surprises. In *Proceedings of 3rd Workshop on Human Error*, pages 1–18. HESSD'99, 1999.

[Rus02] John Rushby. Using model checking to help discover mode confusions and other automation surprises. *Reliability Engineering and System Safety*, 75(2):167–177, February 2002.

[Sch99] Fred B. Schneider, editor. *Trust in CyberSpace*. National Academy Press, 1999.

[SG96] Mary Shaw and David Garlan. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, 1996.

[SH02a] Shamus P. Smith and Michael D. Harrison. Augmenting descriptive scenario analysis for improvements in human reliability design. In Gary B. Lamont, editor, *Applied Computing 2002: Proceedings of the 2002 ACM Symposium on Applied Computing*, pages 739–743, New York, 2002. ACM.

[SH02b] Shamus P. Smith and Michael D. Harrison. Blending descriptive and numeric analysis in human reliability design. In Peter Forbrig, Quentin Limbourg, Bodo Urban, and Jean Vanderdonckt, editors, *Interactive Systems: Design, Specification and Verification*, volume 2545 of *Lecture Notes in Computer Science*, pages 223–237, Berlin Heidelberg New York, 2002. Springer.

[SH02c] Shamus P. Smith and Michael D. Harrison. Improving hazard classification through the reuse of descriptive arguments. In Cristina Gacek, editor, *Software Reuse: Methods, Techniques, and Tools*, volume 2319 of *Lecture Notes in Computer Science*, pages 255–268, Berlin Heidelberg New York, 2002. Springer.

[SH03] Shamus P. Smith and Michael D. Harrison. Reuse in hazard analysis: Identification and support. In Stuart Anderson, Massimo Felici, and Bev Littlewood, editors, *Accepted for 22th International Conference on Computer Safety, Reliability and Security SAFECOMP 2003*, Lecture Notes in Computer Science, Edinburgh, UK, 2003. Berlin: Springer. Forthcoming.

[SO02] C. Sala-Oliveras. Suggestions for the development of software. Master's thesis, School of Computer Science, University of Newcastle upon Tyne, 2002.

[SPA03] L Strigini, A. Povyakalo, and E. Alberdi. Human machine diversity in the use of computerised advisory systems: A case study. In *DSN 2003-IEEE International Conference on Dependable Systems and Networks*, pages 249–258, San Francisco, USA, 2003.

[SPAA03] L. Strigini, A. Povyakalo, E. Alberdi, and P. Ayton. Does incorrect computer prompting affect human decision making: a case study in mammography. In H. U. Lemke, M. W. Vannier, K. Inamura, A. G Farman, K. Doi, and J.H.C. Reibe, editors, *CARS 2003: Computer Assisted Radiology and Surgery*, pages 938–943. ??, 2003.

[TBB+03] K. Tan, G. Baxter, K.G. Brownlee, S.J. Newell, P.R.F. Dear, and S. Smye. Fuzzy logic expert system for ventilation of the newborn infant. *Archives of Disease in Childhood*, 88(1):A47, 2003.

[TSE+03] K. Tan, S. Snowden, C. Evans, K.G. Baxter, and S.J Brownlee. Fuzzy logic expert system for neonatal ventilation. In *Proceedings of the First International Conference on Computational Intellifence in Medicine and Healthcare (CIMED)*, Sheffield, UK, 2003.

[Vau96] Diane Vaughan. *The Challenger Launch Decision*. Chicago Press, 1996.

[Vic99] K.J. Vicente. *Cognitive Work Analysis*. Lawrence Erlbaum Associates, 1999.

[Wei71] Gerald M. Weinberg. *The Psychology of Computer Programming*. Van Norstrand, 1971.

APPENDICES

A People in DIRC

A.1 Research Associates

A total of about 55 RA years will have been charged to DIRC's RA budget during its first three years. The names and hiring (and where appropriate leaving) dates are in the following table.

Name	Location	Started	Finished	%
E Alberdi	City	2000-11-01		100
B Arief	Newcastle	2000-12-01		100
G Baxter	York	2000-10-01		100
G Bernat	York	2000-10-01	2000-12-31	100
D Besnard	Newcastle	2000-09-11		100
D Bosio	City	2001-01-01	2002-03-31	100
J Bryans	Newcastle	2003-01-01		100
M Blythe	York	2002-10-01		100
K Clarke	Lancaster	2000-09-01		100
J Crawford	City	2002-01-01	2003-01-30	25
L D Adderio	Edinburgh	2001-01-01		100
G Dewsbury	Lancaster	2002-01-01		60
J Dobson	Newcastle	2000-07-01	2002-12-31	75
M Felici	Edinburgh	2000-10-01		50
J Küster-Filipe	Edinburgh	2000-10-01		100
C Gacek	Newcastle	2001-01-01	2002-08-31	100
C Geirl	City	2000-08-01	2000-10-31	100
D Greathead	Newcastle	2001-10-01		100
C Gurr	Edinburgh	2000-10-01		75
S Hall	Lancaster	2002-11-01		100
G Hardstone	Edinburgh	2000-10-01	2001-05-31	40
G Hardstone	Edinburgh	2001-06-01		100
M Hildebrandt	York	2001-07-01		100
C Hughes	Edinburgh	2003-01-01		100
J Mackie	Lancaster	2000-07-01	2003-03-31	100
D Martin	Lancaster	2000-10-01		50
S Lock	Lancaster	2002-07-01		100
K Loer	York	2003-01-01		100
C Sala-Oliveras	Newcastle	2001-06-01	2002-04-30	100
M Oussalah	City	2001-01-01	2003-02-28	100
A Povyakolo	City	2001-01-01		100
M Rouncefield	Lancaster	2001-03-01		70
P Ryan	Newcastle	2002-01-01	2003-02-28	0
P Ryan	Newcastle	2003-03-01		100
R Slack	Edinburgh	2002-10-01		100
S Smith	York	2000-10-01		100
M Sujan	City	2001-10-01	2001-12-31	100
M van der Meulen	City	2003-01-01		100
D Wright	City	2002-08-15		50

A.2 Postgraduate studies

Name	Location	Funding	Started	Finished
Peter Bagnall	Lancaster	DIRC DTA	2003	
Adam Betts	York	DIRC DTA	2003	
Alessandra Devito da Cunha	Newcastle	self-funded MPhil	2002	
Martin Ellis	Newcastle	DIRC DTA	2002	
Stephen Gilroy	York	DIRC DTA	2001	
Jon Gregson	Lancaster	DIRC DTA	2001	2001
John Harding	Lancaster	DIRC DTA	2003	
Tony Lawrie	Newcastle	DIRC Studentship	2000	
Russell Lock	Lancaster	DIRC DTA	2001	
Angela Miguel	York	DIRC DTA	2002	
Richard Paul	Edinburgh	DIRC DTA	2003	
Thea Peacock	Newcastle	DIRC DTA	2003	
Mark Rouncefield	Lancaster	Staff	??	2003
Carles Sala-Oliveras	Newcastle	self-funded MPhil	??	2002
Will Stephenson	Newcastle	DIRC Studentship	2000	
Jennifer Tenzer	Edinburgh	DIRC DTA	2001	
Alex Voß	Edinburgh	DIRC Studentship	2000	
Chris Wright	City	DIRC Studentship	2000	

A.3 Faculty members

This list gives the names of those "permanent" academics (and their location) who have a significant involvement in DIRC. They each have "faculty" positions and their salary is not charged to the project.

Name	Location	involvement
Stuart Anderson	Edinburgh	PI
Peter Ayton	City	medium
Guiem Bernat	York	medium
Robin Bloomfield	City	ILD
Alan Burns	York	heavy (was PI)
Alan Dix	Lancaster	medium
Cristina Gacek	Newcastle	was RA
Michael Harrison	York	PI
John Hughes	Lancaster	light
Cliff Jones	Newcastle	PD
Bev Littlewood	City	PI
Donald McKenzie	Edinburgh	heavy
Andrew Monk	York	heavy
Martin Newby	City	medium
Rob Procter	Edinburgh	heavy
Brian Randell	Newcastle	medium
Ian Sommerville	Lancaster	PI
Perdita Stevens	Edinburgh	light
Keith Stenning	Edinburgh	light
Lorenzo Strigini	City	heavy
Robert Stroud	Newcastle	light
Robin Williams	Edinburgh	heavy
Peter Wright	York	medium

A.4 SVFs

The following have visited the project already.

Name	Employer	Visits
Pierre-Jacques Courtois		one visit to date
Ian Hayes	SVRC Queensland	two visits to date
Michael Jackson	Private Consultant	nn visits to date
Tom Lincoln	Private Consultant (ex Rand)	two visits to date
Peter Neumann	SRI Menlo Park	one visit
Kristen Nygaard		one visit
John Rushby	SRI Menlo Park	one visit to date

A.5 SC members

Martyn Thomas (Chair)
Jon Warwick (secretary)
Cliff Jones (PD)
Robin Bloomfield (ILD)
Rebecca Steliaros (EPSRC)
Graham Button
John Fox
Tom McCutcheon
Colin O'Halleron
Fred Schneider
Rob Witty

Roger Needham was a member from the start of the project until his untimely death.

B Progress on Research Themes

B.1 Timeliness

Coordinating site York

Timeliness poses special concerns for the safe use of computer-based systems. There are interesting contrasts between the ways in which systems can be developed to meet guaranteed timing constraints, and the expectations that can be placed on the human users of a system. Contrasts are also evident between the ways users evolve strategies for dependable usage of systems, in a social

context, and the formal statement of operating procedures. Requirements on systems can be made too specific leading to over-engineered systems, or, more usually, temporal issues are disregarded during system development leading to inadequate performance in the deployed system.

Currently in large control systems, the temporal properties typically migrate from critical hard closed-loop deadlines at the bottom layer, to non-critical open-loop performance issues at the upper layers. The lower layers are characterized by a technical focus. The higher layers are more human-centred, with HCI, ethnographical and management issues being paramount. Future systems are likely to have requirements where the need to support and deliver timely behaviour extends up the entire hierarchy. Such systems might emerge in medical applications (remote consultations), flexible manufacture (global optimization and co-ordination) and financial management (guaranteed fair global trading).

This thematic study of Timeliness and its role in dependability is being undertaken with a focus on three phases:

1. A multidisciplinary study of distinct domains of "time" in computer-based systems.

2. An interdisciplinary identification of the key components of the thematic study.

3. An exploration of the roles "time" takes in the delivery of dependability.

There are at least four temporal domains of interest in this study: physical, computational, human and social. The various disciplines within these domains have been reviewed and a comprehensive collection of background material located. Strong concepts of time are to be found in all of the disciplines contributing to the development of computer-based systems. They provide a rich foundation to the Timeliness theme. Unfortunately, they do not always deliver a single coherent and consistent view. A multi-disciplinary study of time therefore provides useful background material for the theme. The following is a list of domains of interest identified in this study: physics and control theory, psychology of time, human factors including HCI, response time studies, sociology of time, evolutionary psychology and anthropology, economics, health and safety guidelines and work-based studies, and linguistic and conversational studies.

In a thematic study of Timeliness, it is necessary to define those components that will underpin the interdisciplinary approach and lead to improved dependability by correctly identifying the roles "time" can play in the design and deployment of computer-based systems. These components can be used to:

1. Separate concerns, where that is appropriate, between the domains of interest.

2. Consolidate notions and abstractions etc. when they are similar in different domains.

3. Identify common problems in different domains.

4. Link related issues across the different domains.

Time itself, specifically the various granularities of interest, will be used to separate concerns. Various abstractions will be used to provide consolidation. Scheduling and consensus are identified as two common problems and the necessary linkage will be provided by vocabulary, modelling and representation.

The initial period of work on this theme has derived great benefit from PA1; the next phase of the work will be concerned with a detailed assessment of a number of case studies and scenarios (already carried out within DIRC) to validate the components identified. This theme would benefit greatly from an outside engagement with NATS whose systems pose just the right sort of challenges.

B.2 Diversity

Coordinating site City

The "Diversity" theme has two components:

1. items of work carried out in various PAs, which also naturally contribute to a corpus of knowledge pertaining to the Diversity theme. For the time being, it is here that new research relating to the theme is being developed;

2. a central co-ordinating activity meant to assist the PAs in recognising the common issues and sharing insights, and to collect the pertinent outputs of the PAs for systematisation and presentation of the common Diversity-related aspects. For the time being, this part has been limited to promoting diffusion of existing results about Diversity and awareness of the pervasiveness of Diversity issues in the problems studied by DIRC.

The relevant activities within the PAs have been

- (obviously) within PA4, with the study of human-machine diversity in the context of advisory systems and of "diverse legs" in safety arguments,

- interchange with the other projects in which City and Newcastle are involved, DOTS (both universities studying Diversity as a solution for the dependability problems posed by off-the shelf procurement) and

- DISPO2 (City, studying safety systems for the nuclear industry);

- PA5, with the modelling of the effect of diversity in the population of users on the reliability growth of a software product;

- in PA6, where issues of diversity for Security have been debated at internal meetings; and

- in PA8, which has launched experimental studies about how to manage software development teams to best exploit diversity within or between them.

The co-ordinating activity has included so far, apart from the frequent presentations of Diversity-related work at PA and plenary DIRC meetings:

- a tutorial about existing mathematical models for diversity, given by City researchers at the first plenary DIRC workshop
- a "Diversity day" held in May 2002 for researchers from all sites and PAs, to offer a broader multidisciplinary picture, with DIRC and external speakers describing results about diversity in hardware and software reliability engineering, experimental studies of group decision making, management studies, etc.

B.3 Structure

Coordinating site Newcastle

Aims and Objectives Concerns about Structure are ubiquitous. The architecture of software (or hardware) governs not only its current function but also its potential to evolve as needs evolve. The social and organisational structure of the organisation into which a system will be embedded must be understood if a computer-based system is to have a chance of supporting the people involved. Furthermore, the social and organisational structure of the organisation developing a system must be taken into account in order to facilitate system development and maximise the Dependability of the resulting system.

The Structure RT seeks ways to discuss, record and understand structures of both social and technical systems. We believe that only with some ways of recording and comparing structures can one hope to review whether a system will prove fit for purpose.

The real challenge is to find something which is of use to –and usable by– both system architects and ethno-methodologists who are recording the organisations as they see them.

Although the enormous breadth of material on Diversity itself justifies keeping it as a separate theme, there will be even stronger interactions between Diversity and Structure than between other RTs.

Input from PAs All PAs have provided useful input to this RT with PA3 being a particularly rich source. The study of "evolution" of systems is fundamental to Structure. PA3 is of course also looking at how generic systems like SAP or PiMS are customized to particular "instantiations".[22] In this connection the work of Alex Voß with Deutz AG is a useful reservoir of problems faced when using SAP in a production environment.

Other discussions within PA3 which are influencing research on Structure include the purpose of "process" and the role of "classification". This latter

[22]Unfortunately, this just complicates the questions around Dependability. What are the Dependability issues of the generic system? How is the range of potential specialisations characterised?

topic has arisen in both the Deutz and PiMS engagements and was discussed in the first SVF visit by Tom Lincoln.

We often focus on the Structure of the (Computer-Based) System but there are also issues concerning the organisation which creates the system. Work in writing [Jon03] clarified the importance of one (possibly human) system "creating" another system. These discussions have arisen in PA5 and PA8 whose meetings have also benefited from Tony Lawrie's evolving PhD research on diverse objective setting in the creating organisation. Alessandra Devito da Cunha's experiments on Psychological type (MBTI) and code reading ability together with the longer-term PhD study by David Greathead on design skills also both relate to the creating organisation.

Interactions to date with PA4 have made sure that we do not exclusively view Structure as yielding a Dependable artefact but that we also remember the importance of being able to establish the degree of Dependability. One could call this "structuring for assessment".

PA6 has generated interesting issues for Structure including Budi Arief's study of "why hackers hack". Peter Ryan and Cliff Jones have made initial attempts to extend the ideas in [Jon03] to cover security issues. In this connection, it is also worth mentioning their work (respectively) on the so-called "Conceptual Models" of the MAFTIA [PS03] and DSoS [GIJ+03] projects.

Progress In early activity, we worked through several well-documented major failures such as [LT93, LAS93, Vau96] to study the fault/error/failure chains (with Tony Lawrie's "Rich Pictures" as a valuable stimulus). This informed the work that led to [Jon03].

We organised a "brainstorming" session at the 2001 Easter Workshop to identify Structure issues of concern to other members of DIRC. The list generated was subsequently written up as an internal position paper.

An important development which has benefited from SVF is [HJJ03]. This paper looks at "faults as interference" (in the technical sense of rely/guarantee research on concurrency Work with Carles Sala-Oliveras on the role of "Advisory Systems" was progressing well but his MPhil dissertation [SO02] was narrower. Carles is sorely missed but a way must be found to pick up some of the work that was not written up because of his decision to return to Catalonia.

Plans We clearly need to write more about Structure. Our early attempt to generate a "Table of Contents" for a "Structure Book" might have been premature but it would now be worth trying again. One specific hope is to extend [HJJ03] to failures of people in the way that [Rus99] tackles confusions of the operator. A harder issue is to search for notations to describe responsibility and trust. We are now at a point of trust between disciplines where we can use formal notations more openly. Equally, to be useful in DIRC, we have to find notations and ideas which will be used to record social structure work We must also relate what we are doing to UML and XML!

B.4 Responsibility

Coordinating site Lancaster

Aims and Objectives The overall goals of this theme of work are to tackle the difficult problem of recording how political issues (power, responsibility, etc.) influence system dependability and to investigate the extent to which these may be used by system designers.

Summary of progress Work in this theme has been primarily based on two areas. Firstly, as part of PA2 and PA3, we have carried out a number of field studies in different settings (NHS Trusts (in both acute and Primary care settings), manufacturing industry, local governments) to develop a deeper understanding of the relationship between responsibility and both observed failures in these organisations and successful functioning of the organisations and organisational processes. This has established a large body of fieldwork that we can now analyse from different perspectives to tease out responsibility issues and, we hope, draw some more general conclusions about the relationships between responsibility and dependability.

Secondly, in conjunction with the AMASE project at Newcastle, we have been concerned with developing a more structured approach to responsibility modelling where a defined notation is used to model responsibilities within and across agencies. This work has been slow at times but now appears to be making good progress. A paper on responsibility modelling will be included in the forthcoming book: *Trust in Technology: A Socio-Technical Systems Perspective* and work is currently underway (Dobson and Martin) to write a paper on responsibility in e-government based on field studies in a local government planning office. This book also includes other papers that address issues of responsibility.

As a background activity, we are also working on an analysis of previous accident reports from a responsibility perspective and a paper on responsibility in the Ladbroke Grove rail accident is in its final stages of preparation.

Reflections Like the other DIRC themes, it has been hard to separate work on responsibility from work on related PAs and, hence, it may appear that little progress has been made. While this is certainly the case as far as explicit outputs badged "responsibility" are concerned, we are confident that we have established a very sound basis for the development over the next few months, of a significant body of work that will focus explicitly on responsibility.

B.5 Risk

Coordinating site Edinburgh

Aim The aim of the Risk theme is to:

1. characterise Risk in computer-based systems, concentrating on risk associated with people and organisations,
2. develop techniques for assessing this type of Risk,
3. develop Risk management techniques that address this type of Risk.

Currently this type of risk is poorly dealt with by risk assessment techniques and as systems become more deeply embedded into organisations the associated

risks increase significantly. In particular, managing inter-organisational interaction while ensuring each organisation has adequate control over the distribution of risk and the control of its risk share is an increasingly pressing issue. For example, computer-based systems are increasingly integral parts of market-based systems (in the UK, energy is one instance) but the risks associated with such systems are still poorly understood.

Our Approach The literature on Risk is huge and rapidly growing. Therefore we have found it necessary to scope the theme tightly to ensure the work is manageable. Our approach is to investigate theoretical approaches to the analysis of risk drawn from the social science literature to assess their applicability to computer-based risk and to investigate whether the analysis suggests approaches to the management of the risks identified by the approach. Initially we are investigating the strengths of individual techniques but in the lifetime of the theme we intend to develop a technique that identifies the most appropriate analyses for a given situation and how the results of the analyses can be synthesised to provide an adequate risk analysis.

We are investigating the following sociological work as the basis for the analysis of risk in computer-based systems:

1. *Imitation:* recently Donald MacKenzie has investigated imitation of successful strategies amongst organisations as a potential source of "common mode" failure across organisations engaged in market-based interactions.

2. *Cultural Theory:* Mary Douglas' work on the cultural composition of groups and their approach to the identification and management of risks provides a strong framework for the analysis of risk at the level of local workplace cultures.

3. *Boundary Objects:* Leigh Starr and Geoff Bowker have developed this theory objects that are used to communicate across groups in complex "information ecologies". Agreement over the interpretation of such objects is the subject of continual negotiation and the emergence of different interpretations is a source of risk in computer-based systems.

4. *Uncertainty trough:* Confidence in technologies varies depending on "social distance" from their creation. In particular, those who are close to the technology but are not involved in its development tend to be overconfident in the technology. This phenomenon seems to play a role in risks associated with premature or inappropriate deployment decisions.

5. *Normal Accidents:* Perrow's work [Per99] suggests a framework for understanding particular classes of accident.

6. *Affordances:* Don Norman has that communities of users develop symbolic systems to extend the use of a system or ameliorate difficulties in their use. Failure to consider the use of affordances in systems is a source of risks in computer-based systems.

In the later stages of the project we want to incorporate ideas from the psychology of human decision taking to extend our approach to cover risks

associated with individual decision taking. In particular we believe that *Prospect theory* (Tversky) and *Heuristics* (Gigerenzer) could provide the basis for risk management techniques in human decision taking.

Progress to date So far we have considered individual studies involving Imitation, Cultural Theory and Boundary Objects. In all the studies it appears that additional risks are identified by the social science perspective and approaches to the management of these risks are also feasible. Our main studies include:

1. A study of the failure of Long Term Capital Management. This work has strong connections to work on Diversity in DIRC.

2. In the medical sector we have considered using Cultural Theory to account for differences in the approach to privacy and security in the development of patient information systems in a mental health setting.

3. Again in the medical domain we have considered the role of Boundary Objects (in particular classifications of activity) in communicating across social groups and have considered the potential hazards that arise from the haphazard development of classifications and their evolution and reuse as the context of use evolves.

We have also considered what this work has to say about existing standards for risk management. In particular, in association with EWICS (the European Workshop on Industrial Computer Systems), we have considered IEC 14971 the risk management standard for Medical Devices and have fed back our comments to the standardisation body via national representatives.

Future work The project now has a very rich body of empirical research with ongoing connections to a range of computer-based projects. We are in a position rapidly to consider the utility of other candidate sociological analyses of risk in computer-based systems to appropriate case studies. The following activities will be undertaken as part of the Risk theme over the next three years aiming at developing a reasonably comprehensive approach to risk associated with organisations in computer-based systems:

Coverage: we are in a position rapidly to extend the coverage of our case studies to consider all the sociological approaches listed above.

Generalisation: we believe we can generalise the approach from individual case studies, as guidance in the first instance, but it may be possible to provide more analytical procedures for some classes of risk.

Psychological aspects: we want to engage with Peter Ayton in City and Keith Stenning in Edinburgh on the articulation of the social and psychological aspects of decision taking in computer-based systems.

Risk Management: Our goal is to develop a range of techniques and guidance on their deployment that will allow adequate analysis of organisational and individual risk in the specification, design and operation of computer-based systems.

C Reports on Project Activities

C.1 PA1: Human Interaction in Real-Time Systems

Coordinator Michael Harrison

Partners involved City, York, Edinburgh

Span 2000-10-01 to 2003-08-01.

Scale of effort to date 5 RA years

Review dates (progress) 2002-05-13 (final) 2003-08-15.

Aims and Objectives The objectives of this Project Activity were fivefold.

1. The identification of Timeliness properties of interactive systems that may, in some contexts, affect the dependability of the interactive system.
2. The adoption and development of models of interactive system that may be used to analyse these dependabilities and to aid design understanding of the system.
3. The development of a technique for the structured assessment of real time interactive systems to provide arguments for the dependability of human interaction in such systems.
4. The identification of system mechanisms that shall make it possible to identify scheduling opportunities and effect process substitutions that will guarantee the meeting of timing constraints at some level of functionality.
5. The provision of a means of supporting the appropriate visualisation of timing constraints for the system operator.

The aim was that this work would be done in an interdisciplinary context with particular reference to computer science (Edinburgh and York), psychology (City and York) and sociology (Edinburgh). It was seen to be important that any work be based around realistic case studies and with this mind it was hoped that Telefonica and NATS would

be obliging in providing material. However the aim of the project was intended to be more concerned with the modelling of systems and the development of relevant concepts than, at this stage, appropriate techniques for collecting data about systems.

Outcomes:
Field study: Neonatal unit
A "cognitive task analysis" was carried out of Leeds neonatal intensive care unit where computer-based equipment is used to monitor blood gases and ventilate babies suffering from Respiratory Distress Syndrome [TBB$^+$03, BTN$^+$02]. An expert system was being introduced to assist junior doctors in making decisions about changing the ventilator settings [TSE$^+$03, BTN$^+$03]. If the expert system

was to succeed then it must be clinically useful and acceptable to the staff. The detailed cognitive task analysis raised several issues that were associated with the clinical usefulness and acceptability of the expert system that extended beyond its functionality. These issues are being used to modify the design of the expert system.

Dynamic Function Scheduling:
An alternative to current notions of *dynamic function allocation*, called *dynamic function scheduling*, has been developed [HH02b, HH02a, HH03]. An issue that arose out of exploration of the intensive care unit has been that in the face of a time critical decision a trade-off must be made between choosing to trust the results provided by the expert system or making use of other data that will cost more to extract — this cost could be in terms of time, or could be in terms of the potential effect of getting the information from an "off duty" consultant.

This problem has something in common with the problem of dynamic function allocation. An important issue in system design is to decide what aspects of the system should be automated. The York group has recently developed a method to support the decision process. Neither this method nor any other that we are aware of takes account of dynamic function allocation where levels of automation may change (for example in the face of workload or situation awareness or performance concerns).

Micro-world study:
This part of the activity is concerned with moving from conceptual considerations about Dynamic Function Scheduling (DFS) towards empirical investigations of Human Factors aspects of systems involving DFS, and to support the analysis and design of such systems. The activity is informed by the Human Factors literature on Dynamic Function Allocation (DFA). We intend to modify the methods and scenarios used in this area to investigate how temporal factors (such as the time available/time required ratio) contribute to or modify these effects. Preliminary work on micro-world simulations (heating system and paint station) has been carried out for the investigation of effects such as loss of situation awareness, complacency in automation use, and trust in automation. Of particular interest are the following questions:

- If a system provides the option to automate and/or postpone functions, do operators adequately and flexibly use these workload-balancing options to maintain both safety and production targets in the face of multi-task demands?
- What are the temporal aspects of situation awareness? How useful is the concept of temporal awareness and how can it be measured and modelled?
- Is trust in a decision support system moderated by the time pressure for making the decision such that the operator is more likely to accept the recommendation of the system under high than under low time pressure, even if the reliability of the system is sub-optimal?

Progress has been made in identifying scenarios, and in generating a micro-world simulation (the "Paintshop"). The results of a series of experiments are currently being analysed.

Exploration of formal modelling and analysis tools:
Temporal aspects of interactive systems were identified and it was explored how

these aspects can be modelled in a way that is amenable to model-checking analysis.

Two "modelling meetings" (York, February 2002 and July 2003) were attended by interdisciplinary members from various sites. During these meetings a catalogue of issues that should be considered when modelling interactive systems was created, and a number of modelling approaches were tried out.

Several model-checking tools were assessed with respect to their applicability to the problem domain and to their scalability to the required problem-size. Different kinds of modelling and analyses were applied in a case study on a heating plant and a sewage plant. The case study exhibited that model-checking and the formal modelling of devices in context can play a role in understanding the effects of tasks and the unforeseen consequences of various types of task violation. When used in an exploratory manner, as opposed to verification, the property failures that are produced by model checkers may be used to identify interesting sequences. These sequences can be particularly interesting from the timing (in the case of Uppaal and HyTech) and task sequencing (SMV) points of view.

Scenarios for DSTL case study: A set of scenarios was developed that will be used in a case study that is conducted for the DSTL customer. The cover story for the scenario is a mobile command, control and communications (C3) unit that supports the co-ordination of, and communications between a variety of ground and airborne units that are involved in a military operation.

In brief, the main goals of the C3 unit are to provide high-quality real-time information and communication without being discovered. The crew consists of operators with different skills and responsibilities. Redundancy is achieved through the overlap in the skills of crew members.

Parcel Call:
ParcelCall is a research and technology development project with several industrial and academic partners. The University of Edinburgh is one of its academic partners. The initial aim of the project was "developing and trialing an information technology system to improve business processes in transport and logistics through a 'real time', 'seamless' integrated tracking and tracing system that will operate 'end-to-end' across different carriers and transport modes at the individual parcel level." These can be found essentially at two levels: at the level of transport and logistics operations; and at the design level of the ParcelCall system itself. This case study has been used as a basis for the elicitation of timing requirements at the level of social time.

Temporal validity: In distributed real-time applications where data need to be shared among distributed components it is desirable to have overall data consistency at all times. It can be important in particular for safety-critical systems, where inconsistency can lead to catastrophic failures. Overall continuous data consistency is, however, rarely possible to achieve. For distributed systems, particularly where people are involved, a relaxed view based on temporal validity of data can be proven sufficient. If different components in a distributed computer-based system have different temporal validity assumptions or constraints for the same data, then as long as these constraints are satisfied overall system inconsistency is not harmful. A formal analysis technique is being developed by Edinburgh for determining the temporal validity of shared data in real-time system application. The approach is based on a real-time extension of

a modal logic of knowledge. Modal logics of knowledge are commonly used in formal approaches of distributed computing though they generally do not consider real-time. The logic has been kept simple to make automatic verification through model checking feasible. It makes it possible to

check that shared data in the system is consistent "enough" and cannot be a source of failure. It is believed that this approach is a valuable addition to hazard analysis and hazard-control measures for use during early stages of software development.

Future plans The following Targeted Activities will be started before the end of this year.

1. *Temporal aspects of automation use and trust in automation.*

2. *Investigation of notations and models for decision processes.*

3. *Application of formal techniques for analysis of temporal aspects of interactive systems.*

4. *Integration of timing aspects into scenario-based analysis techniques.*

C.2 PA2: Impact of Organisational Culture and Trust on Dependability

At the outset, it is important to make clear that there has been considerable overlap between the work done on PA2 and the work of PA3. Both of these activities draw on a corpus of field studies (described here with additional discussion in PA3) in a range of different settings. Furthermore, work on patterns and modelling work situations is clearly relevant to both activities.

Coordinator Ian Sommerville

Partners involved Lancaster, Edinburgh, Newcastle

Span 2000-09-01 to 2003-06-30

Scale of effort to date 6.5 RA years

Review dates (progress) 2002-08-20; (final planned) 2003-10-14

Aims and Objectives To develop conceptual and software tools that allow system designers to adapt their design to a prevailing organisational culture and to help them design systems that promote and strengthen the notion of a "dependability culture" in an organisation.

Outcomes The outcomes of the work on this activity are:

1. A corpus of field study material that can be analysed in different ways to investigate the relationships between organisational culture, trust, responsibility and dependability. These field studies have been primarily in the medical domain and in manufacturing industry. The field studies are summarised below.

2. A prototype 'method' for assessing the vulnerability of processes to human error. The aim of this work was to develop a 'method' that embodied some of the knowledge gained from the field studies and that would help system designers study situated processes with a view to identifying areas of process vulnerability where errors and omissions were likely to occur. This work was based on earlier work carried out by Adelard and Lancaster that focused on vulnerabilities in requirements engineering processes and the aim here was to generalise this to a method for generic process assessment. This work is not yet complete for the reasons discussed below but we hope to continue this as part of a Targeted Activity in the next stage of the DIRC project.

3. A method that we call *situation modelling* that is intended to provide a rich representation of a work situation. The rationale for this work was to represent these situations in a structured way that was more compatible with models that may be developed during the system development process. The work, done in conjunction with PA3, has derived a multi-perspective model of a work situation showing activities, objects and actors and allows them to be represented in a flexible way. A supporting display tool has been developed to illustrate aspects of situations with vignettes drawn from field studies and the Scavenger tool (discussed under PA3) may be used to construct these vignettes.

4. In conjunction with PA3, a series of *patterns of interaction* that embody issues of culture and trust in the workplace. These describe different patterns of interaction that have been observed in a range of settings and the standard format for a pattern description discusses the significance of the pattern for dependability. All patterns are illustrated with a series of vignettes describing pattern realisations that have been observed (many derived from the field studies discussed above) and we are in the process of annotating these patterns with comments on organisational culture and trust. The patterns web pages are at:

   http://polo.lancs.ac.uk/pointer/PatternsOfCooperativeInteraction

Changes from plan A decision was made in late 2002 to document the results of this activity as a multi-authored book that is now in late stages of preparation. This book is edited by Gillian Hardstone and Karen Clarke and the preparation involved some diversion of effort from other planned activities. We did not achieve our aim of developing software support for some of the work in this PA. This was partly due to lack

of appropriately skilled staff and partly due to delays in finishing the work on process vulnerabilities because of diversion of effort to the book and personal difficulties that seriously affected the same work.

Field studies The majority of effort has been spent collecting data on culture, trust and responsibility in a range of settings:

1. *Hospital management* This ethnographic study has involved a series of fieldwork visits spread over the last 18 months at three hospitals in the North of England. We have shadowed thirteen managers within the Trust for periods of a week at a time, observing the accomplishment of everyday managerial work and the role of ICTs within that work. The general focus of the study has been on organisational culture and trust but we have also explored specific themes in more detail as they have emerged from the fieldwork. These themes include: the everyday management and prevention of failure, trust in record-keeping systems, the affordances of existing information systems and their relevance for system development.

2. *Mammography* This study forms part of an ongoing field trial of a computer aided detection (CAD) tool in mammography. As well as extending our understanding of reading practices and considering usability issues for deployment we are also concerned to investigate some of the effects of CAD tools on reader performance, including readers' understanding of the CAD tool's behaviour. The study illustrates the general point that, in areas such as medicine, the introduction of computer-based detection and diagnosis tools highlights and problematises how the results provided by these systems are made sense of and deployed in everyday work. This raises important dependability questions for the design and use of new technologies in medical work, in particular whether some of the standardised procedural changes the technologies are intended to support are actually achievable.

3. *PiMS* This is a field study of the introduction of a new, COTS-based patient information management system (PiMS) within a large healthcare trust. The system has been introduced with the aim of integrating and standardising patient administration practices within the Trust's numerous different healthcare service units. Our interest lies in the ongoing efforts of the PiMS project team, PiMS user group, etc., to configure the system to fit the organisation (we discuss this in the description of PA3). The case study illustrates very clearly the obstacles to the creation of an organisational infrastructure or 'boundary object' capable of affording the interplay and coordination of different 'communities of practice'. It reveals that the configuration of PiMS must be supported as an evolving process as members of different communities of practice learn to work together and the implications of information integration are gradually worked out by the organisation.

4. *Manufacturing* There have been two studies in the manufacturing industry. In CORUS (formerly British Steel) and in Deutz (diesel engine manufacturers). The work in CORUS consisted of a short observational study of the Roughing Mill. The study documented everyday working practice in the Roughing Mill and identified a range of issues to do with coordination, awareness and planning. It suggests a number of problems of team-working, and computer problems related to the identification, measurement and sequencing of the slabs.

Edinburgh is engaged in a detailed participant observation study in a manufacturing plant producing mass-customised diesel engines. Through sustained presence and workgroup membership the researcher is able to explore 'co-realisation': a new form of participatory design that builds upon the everyday knowledge, practices insights, and culture of organisation members. The researcher works with control room staff in the plant, as well as company IT people, production staff and managers to analyse the detail of work practices and culture, and to develop and enhance IT systems.

Dissemination As is obvious from the appended publication list, the work in this PA has been widely disseminated in conference and journal papers. We particularly note that these publications address different disciplines thus maintaining the interdisciplinary flavour of DIRC. The work will also be summarised in a book entitled *Trust in Technology; A Socio-Technical Systems Perspective* that is in its final stages of preparation and which will be published by Kluwer in 2004.

Reflections The initial objectives of this workpackage which were concerned with understanding the difficult relationships between dependability, culture, trust and responsibility were challenging and it is not surprising that they have not been achieved in full. Nevertheless, we believe that this PA has considerably deepened our knowledge in this area and the method of process analysis forms a useful basis for

further work concerned with making social analysis accessible to systems designers and engineers. The work on patterns is, we believe, particularly significant. While this started outside the DIRC project, the focus on dependability on DIRC considerably sharpened this and the DIRC field studies were invaluable for identifying the pattern instantiations that are a critical component of pattern descriptions. It is a little disappointing that we did not manage to develop some software tool support. This would be important because it makes the work more accessible to computer scientists. This work, along with work on situation modelling will be important inputs to a new proposed TA on integrating social science and systems engineering methods.

C.3 PA3: Dependable Deployment and Evolution

Coordinator Stuart Anderson

Partners involved Edinburgh, Lancaster, Newcastle

Span 2000-07-01 to 2003-06-01

Scale of effort to date 10 RA years

Review dates (refocus approved) ??; (final (planned)): ??

Overview The study of the process of configuration and reconfiguration is central to the dependability of computer-based systems in providing the means to adapt the overall system to the evolving environment. The primary aim of PA3 is to study and characterise the relationships between dependability and evolution of the technical system, organisation and environment, focusing on (re)configuration as the means to control the effects of evolution on dependability.

The hypotheses that drive the work of this PA are based around the on failures of systems to meet the needs of users:

1. Reasonably frequent, non-catastrophic failures of sub-systems are common but the system user initiates other actions to compensate for this service loss resulting in an overall degradation in service from the enclosing system;

2. that the evolution of the system is a driver that originates new failures, and

3. that the control of these failures is often dealt with by user-controlled measures, e.g. reconfiguration both of the computer systems and the surrounding social processes.

4. these user-controlled measures often lead to significant improvements in the dependability of service provision by the socio-technical system.

All the studies described below illustrated these features and demonstrated that different domains and different work process have considerable influence on the manageability of evolution in the workplace. In some cases, the efforts of staff to work around systems takes a considerable part of their working time.

An important resource for PA3 (and other DIRC activities) is the collection of case studies which have observed situations where reconfiguration in response to evolution is being undertaken in a range of different organisational structures:

1. A study of the automation software in Deutz AG that involves modifying the locally controlled automation control host software to take account of the needs of production workers. The overall system is driven by an outsourced SAP system that provides a rigid production planning environment that results in many workarounds to enable smooth running production in the engine plant.

2. Two studies of the implementation of systems to support patient information in mental health trusts. These contrast strongly:

   (a) The Patient Information Management System (PiMS) is a configurable package that has been bought in by Lothian Primary Care Trust to provide patient information for its mental health caseload. This system has already been described in the discussion of PA2.

   (b) The other mental health system is intended to carry out many of the functions of PiMS but is being developed in the context of a smaller organisation by a development team who have worked as mental health workers in the local team. The results are strikingly different in the approach to meeting the needs of the mental health teams and responding to changes in the environment of use.

3. Our final study involves the dependable transfer of a highly automated production process from the US to a UK context inside SUN Microsystems. This involves the careful management of a process into a new operating context. The actual implementation of this policy involved detailed bottom-up work to understand the process, transfer it and measure the results of the transfer. This study provides an in-depth examination of an organisation controlling a particular evolutionary step to ensure a predictable outcome.

Progress We can assess progress in this activity by considering each of the objectives in turn:

1. To observe the long-term evolution and reconfiguration of socio-technical systems and their organisational settings and to analyse how user-led evolution and re-configuration both helps cope with and acts as a source of system failure. *We believe we have achieved this objective almost completely but we still need some work to make the detailed data from the studies more widely available across the projects.*

2. To devise and assess approaches to configuration in socio-technical systems as a means to avoid, control and manage reasonably frequent non-catastrophic failures of sub-systems. *We have made substantial progress on this objective. Our approach involves careful design of the processes used to respond to the demands of evolution and the identification of key structures requiring management if the systems is to support evolution successfully. The Deutz study has given us in-depth experience of the co-realisation approach to maintaining systems in the face of evolution. We have found the sociological notion of "boundary object" useful in identifying objects that are responsible for sharing information across organisations.*

3. To devise and assess evidence-driven approaches to support both developers and users in achieving dependability requirements across the deployment boundary. In particular, to explore lightweight methods that exploit data gathering to direct the structure of configurations and how systems should be configured to fit a particular work setting. *We have developed one prototype tool to gather configuration data (Strider, described below) to support design and use of configurational systems. This work has also shaped our work on Grid service architectures as part of Project Activity 9. Taking this work further will be part of a focused activity linking this work to the creation of dependability cases for socio-technical systems.*

Summary The work completed so far and the developments we plan over the next 6-9 months are summarised under the main work headings of the PA3 proposal.

- *Empirical Work* We have gathered a substantial body of data from each of the case studies concentrating on configuration and the evolution of configurational systems. We will maintain periodic contact with all the projects to gain more evidence over a longer baseline since seeing the consequences of earlier decisions on the trajectory of the project is an important part of the study of evolution.

- *Characterising Evolution* We have four detailed studies of evolution in configurational systems in different contexts and with different surrounding organisational structures. In Felici's forthcoming thesis he attempts to build a reasonably general framework to model evolution in the requirements of system. From other studies we have accounts of the drivers for system evolution. In health the driver is policy change, in Deutz the driver is the need to match productive activity with the production planning system and in SUN the drivers are to match environments as closely as possible.

- *Architectures and Structures for Configuration* Each of the case studies has particular processes and structures for dealing with evolution. For each case study we have some evaluation of the strengths and weaknesses of these processes and structures in dealing with change. Our work on structures has so far concentrated on structures that are negotiated as part of the development process. Work on PiMS has highlighted classifications as key structures in the sharing of knowledge across the systems. Another important structure is that of access control to data. This has arisen in the context of the health system projects that use radically different approaches to access control with consequences for the effectiveness of the system and patient privacy.

- *Evidence-Driven Support for Evolution* This is the least mature part of the activity because we had not intended much work to be done before the next 6-9 months. The aim of this activity is to link configuration and evolution to work on assessment and the dependability case. Work on evidence gathering will be facilitated by the use of the Scavenger tool (described below).

- *Supporting Decision Taking* The notion of the configuration of a socio-technical system covers all aspects of the system from the configuration data on a package through the physical arrangement of the workplace and typical information flows outside the system. Capturing this kind of data for the designer and the user is difficult. To ease the problem of managing large volumes of ethnographic and other data relating to a situation we have created the Strider tool to capture socio-technical system configurations.

- *Higher-Level Models* Work in this area has concentrated on MacKenzie's studies of catastrophic failures of market mechanisms. This explores some failure modes of market-based mechanisms and links it to a failure of diversity in the underlying model due to evolution of the market driven by imitation of the market leaders. This points to the need to maintain diversity and how, since the collapse of Long Term Capital Management, diversity has been forced internally within companies and by regulatory changes. As DIRC progresses it seems likely that market-based interactions will become more important. Recently, a new PhD student in Edinburgh has begun to explore the stability of market-based mechanisms in large-scale wireless systems.

- *Tool development* We have developed two tools to support the work on configuration. Scavenger is a tool that helps organise large volumes of field

study data and Strider is a tool that supports configuration modelling.

Scavenger has been designed to ease the process of building socio-technical system models and populating them with information derived from ethnographic study. Raw, loosely formatted data produced by ethnographic activities can be converted into structured entity-relational models. As such, scavenger acts as a sophisticated cut and paste tool, which is specialised to the task of data extraction from ethnographic source artifacts. Existing ethnographic data provides an initial source of information; scavenger then imports a set of model specific XML templates; these templates are then filled with the scavenged ethnographic information and the resultant XML document exported to the file system. This mechanism allows important entities from the ethnographic data to be identified, classified and described. In addition to this, links between defined modelling entities may also be described.

A configuration model captures the static components and relationships within a system at a particular point in time. It models the content, structure and form of all social and technical components within a system. So, for example, a configuration model would represent the software, hardware, worker and resource components of a system, as well as the organisational, structural and operational relationships between them. Strider is a configuration modelling approach and support tool which utilises the previously mentioned scavenger toolset. Strider is used to capture a very thin "skin deep" layer of system configuration whilst at the same time maximising management, support and analysis opportunities.

has been widely disseminated in conference and journal papers. We particularly note that these publications address different disciplines thus maintaining the interdisciplinary flavour of DIRC. The work will also be summarised in a book entitled *Trust in Technology; A Socio-Technical Systems Perspective* that is in its final stages of preparation and which will be published by Kluwer in 2004.

Reflections The initial objectives of this workpackage which were concerned with understanding the difficult relationships between dependability, culture, trust and responsibility were challenging and it is not surprising that they have not been achieved in full. Nevertheless, we believe that this PA has considerably deepened our knowledge in this area and the method of process analysis forms a useful basis for further work concerned with making social analysis accessible to systems designers and engineers. The work on patterns is, we believe, particularly significant. While this started outside the DIRC project, the focus on dependability on DIRC considerably sharpened this and the DIRC field studies were invaluable for identifying the pattern instantiations that are a critical component of pattern descriptions. It is a little disappointing that we did not manage to develop some software tool support. This would be important because it makes the work more accessible to computer scientists. This work, along with work on situation modelling will be important inputs to a new proposed TA on integrating social science and systems engineering methods.

C.4 PA4: Decision Support for Dependability

Coordinator Bev Littlewood

Partners involved City, York, Edinburgh, Lancaster

Span 2000-10-01 to 2004-01-01

Scale of effort to date 12 RA years

Review dates (progress) 2002-07-08; (final (planned)): 2003-11

Aims and Objectives The original project description stated "The aim of this project will be to develop methods and tools to aid decision-making in all situations where the dependability of IT systems is an issue. An important requirement will be a better understanding of means for marshalling disparate kinds of evidence to support decision-making in domains such as regulation (e.g. via the use of quantitative dependability cases)".

Most of the activity can be described in terms of two major threads, one about the structure and contents of dependability cases, the other on decisions about diversity in redundant, human-machine systems. In addition, we have run a focused study of a novel approach for the assessment of hard real time systems, and several additional exploratory activities, surveying topics at the intersection of the disciplines represented in DIRC or new techniques that appear promising.

Dependability cases (City, York)

- Fundamental concepts

  In DIRC the term "dependability case" is used (by extension from the common usage of "safety case") to denote a rigorous argument linking appropriate evidence to a claim that a particular system is sufficiently dependable to be acceptable for use. Even in industrial sectors in which the idea of a "safety case" is accepted, the cases themselves are still often lacking in rigour. Much of the work in PA4 is aimed at improving this situation, in part by formalising the way arguments are expressed, so as to clarify their meaning and make it possible for different people to communicate the arguments with precision and check each other's reasoning and evidence.

  Part of the effort in PA4 has been dedicated to "fundamental" issues of clarifying concepts used and common forms of argument. Several meetings were dedicated to discussing the advantages and the limits of quantitative methods given, for example, the uncertainty in estimating probabilities of human failures. This work is still ongoing.

  An essential issue in dependability cases is that of confidence, in the claims and the evidence and subclaims or assumptions that support them. Interestingly, the degree of confidence in the statements made is seldom addressed clearly in current practice. We have addressed this issue in formal probabilistic terms: a preliminary internal paper examines the effects of limited confidence in the claimed probabilities of failure of components

on the claims that can be made for a whole system. Another aspect of confidence concerns potential flaws in the structure of an argument. A defence used against this concern is "multi-legged" arguments, in which each "leg" is distinct and logically "independent" (as far as possible) of the others. The aim is to improve the confidence that claims made for the system would be correct even if one of the "legs" were to contain errors of reasoning, measurement etc. [BL03] illustrates the open issues in the study of such arguments, and discusses to what extent these may be amenable to the kind of analysis developed for assessing the gains brought by diversity in building fault tolerant systems.

- Argument reuse

  The observation that arguments for safety cases are often built by reuse of argument fragments has led to a study of how this reuse can be organised and supported. After a case study on the HAZOP analysis of the commercial DUST-expert software, a proptotype tool was produced and used in another in-house HAZOP exercise [SH02c, SH03].

- Other work

  A study has been conducted of the U.K. regulation of safety for ammunition in over the last 50 years, to look at the evolution of regulatory culture and formalisms and draw lessons about the feasibility of the various approaches that DIRC may propose. An internal report has been produced.

Worst case execution time (WCET) modelling (York, City) For hard real-time computer systems, designers have traditionally sought deterministic "guarantees" of schedulability of the intended workload, based on estimating the worst-case execution time of all tasks. With modern computer architecture, these worst-case calculations have become increasingly pessimistic, so that designs that enjoy such guarantees suffer from massive under-utilisation of resources. We seek to substitute these with probabilistic guarantees, i.e., proven, acceptably low probabilities of tasks missing their deadlines. We are using the mathematics of extreme value distributions to represent uncertainty in this situation, using analyses of large samples of actual execution time data to provide an empirical check on the accuracy of predictions from our models.

Assessment and design of diverse human-machine systems (City, Edinburgh, Lancaster) Diversity between redundant elements of systems and processes is recognised as an important topic in PA4. Its study (which has also gained from interaction with work on other projects) has been mostly focused on the use of diversity between humans and computers for error detection, recovery and masking, and specifically on a case study. This case study is the use of "computer aided detection" in reading X-ray films in screening programmes for breast cancer, which we hope will serve as a first example for studying the much broader category of advisory systems. The occasion for this work was a NHS-funded controlled trial (run by teams outside DIRC) of a specific "computer aided detection" machine. Edinburgh, Lancaster and City collaborated with the organisers of the trial and the machine's manufacturers, receiving access to

data and providing additional insight. The DIRC approach has combined statistical analysis, the running of small supplementary controlled experiments, "clear box" reliability models, direct observation of the medical staff at work, combining methods and staff from reliability engineering, computing, sociology, psychology and ethnography. This approach brought original insight on the results of the case study, revealing effects of the use of the machine that neither the conventional statistical analysis of the trial results, nor the manufacturers' analysis of the performance of the machine in isolation had revealed. We have published preliminary papers on the empirical aspects of the study [SPAA03, APA+03] and surrounding general issues [HPRS00, HPRS01, HPR+03], and on the modelling approach [SPA03] that we expect to be useful for advisory systems in general. The study is continuing and joint papers are in preparation both about the case study results and about the interdisciplinary methodological aspects.

Other case studies in the medical area have been considered and initial approaches made: a computerised advisory system for general practitioners, and the expert system for assisting operators in a neonatal care unit which has also been used by PA1. Preliminary meetings were held with the owners of these systems and gave DIRC researchers insight into the medical application environment, but neither had the manpower available for a sustained collaboration with DIRC.

Descriptive arguments (York, City) As part of the project activity on decision support for dependability

(PA4), York has investigated the nature of descriptive arguments and opportunities for their reuse. Descriptive arguments can be considered as informal arguments in contrast to more qualitative, numeric arguments. Such arguments are commonly produced to argue that the perceived severity of hazards are mitigated against and to document hazard barriers and defences.

Initial work at York has explored the link between descriptive arguments and numeric analysis as found in human error assessments. We demonstrated how descriptive arguments generated through scenario-based analysis can be ranked using a numeric human error assessment technique [SH02a]. Such rankings can aid the choice of re-design recommendations [SH02b].

Generating descriptive arguments is an error-prone, time consuming and repetitive process. However, within domains, it is common to find similarities between the defined arguments. We have investigated a reuse mechanism [SH02c] to exploit these similarities. This work has involved the development of a prototype support tool and several case studies [SH03] including reuse in a safety case (through an industrial partner Adelard), analysis of a computer-assisted mammography system (with PA4 colleagues at City University) and fall assessments for older adults (with external occupational therapists and PA7 "Dependable ubiquitous computing in the home" colleagues).

There are two main themes to our work and we envision that they will form Targeted Activities in the new DIRC activity structure. Argumentation for dependability involves a constant trade-off between qualitative and quantitative approaches. Our first theme is to understand this trade-off and in particular whether qualitative approaches can be used to improve the accuracy of quantitative approaches. The work will involve collaboration with DIRC colleagues at City and will focus on human reliability assessment. A second theme will

continue our work on argument reuse and will investigate how reuse plays a major role in light-weight techniques that can be actively applied in industry. In addition, bridging the qualitative and quantitative divide and the efficient reuse of arguments form the core of a research proposal under development. This proposal will investigate how confidence in a dependability analysis, and any associated reuse, can be determined and in particular verified through numeric evaluation.

Other exploratory and surveys studies

- *Adjudication and group judgement* Redundant computer systems and groups of human experts present the same problem for a designer or user to deduce a single result from their multiple outputs. We have started a survey of solutions to this problem (one of design for engineers, of 'normative' decision making for psychologists), noticing differences between treatment of these formally identical topics in the two disciplines, leading to results that appear to be "portable" between the design of systems including only machines, only computers, or both.
- *Random graphs* This category of models appears promising for describing uncertainty about the error propagation channels within a system, which in turn may be a major factor of uncertainty in a dependability case. Preliminary results were presented in [Blo03].

Supporting activities Workshop on dependability cases: PA4 work has been presented at the successive internal DIRC workshops. A special multi-site, multi-disciplinary workshop was held on 24–25 April 2002 at City, with focus on the mammography case study. The purpose was both to gain inputs for the case study from DIRC researchers, and to give the researchers an occasion to work in small groups on specific issues in dependability cases for a concrete system, thus helping to create a shared project culture on this topic. Follow-up work was scheduled for several participants, leading to a smaller follow-up meeting in York in July 2002.

C.5 PA5: Dependability Issues in Open-Source Software

Coordinator Cliff Jones

Partners involved Newcastle, City, Lancaster

Span 2000-12-01 to 2001-12-31

Scale of effort to date 3 RA years

Review dates final on 2002-01-03

Background At the time the DIRC project was being defined, "Open Source" was frequently discussed in relevant research and funding circles. It was being claimed that the approach could solve many current software and system development problems, including aspects of dependability. Consequently, DIRC's executive board decided to investigate the potential contribution of the "Open Source" approach to ameliorating Dependability problems. The decision was made to initiate a short-term Project Activity to address the potential dependability contribution of "Open Source Software". All other DIRC Project Activities were planned to last about three years; PA5 was to be a study lasting only one year, whose main task was to recommend whether there was a need for further DIRC work in this area.

Throughout the period of this study we gathered a considerable amount of data via an ethnographic study, by conducting interviews and investigating many papers and web sites. We digested the data that was gathered and discussed what could really be done in the time frame. We also considered running experiments and performing modelling work.

Understanding Open Source One of our early observations was that many organisations claimed to be using or developing "Open Source Software" (and there are even more people referring to "Open Source") without defining the term. Our investigations showed that there are as many variations among Open Source projects as there are among more "traditional" software projects and that in some instances the differences between "Open Source" and conventional regimes were smaller than within a category. The range of factors where variations occur includes: legal (differing contracts and licenses); financial (software at times being free and the possible lack of payment of developers); organisational (differing power structures controlling the project and at times differing modes of communication); and technical (the possible existence of tools supporting developers' documentation such as software architecture or test suites). The result of this work culminated in a paper [GLA01] which was unfortunately rejected on first submission but has now been submitted to a more appropriate journal.

Sociology of Open Source We also investigated the Cocoon project in detail with the aim to learn how that particular group actually functions. We studied their website, analysed their e-mail records and interviewed one of their core developers. Conclusions were derived from this work, but they are difficult to generalise given that they are based in a single project. The result of this work is described in [MRR02].

Modelling and Open Source Dependability When trying to assess the dependability claims relating to Open Source, we were confronted with the fact that current claims seem to be highly anecdotal. Some of our work has been aimed at clarifying the kinds of arguments that can be made about Open Source's dependability advantages or disadvantages and what might cause them. In order to try and validate or refute dependability claims based on quantitative measurements, we have built a speculative model of reliability growth based on bug reporting and associated fault fixing, and the diversity in users' usage

profile. Information about this model can be found on the paper in Appendix E of [ABGR02].

Conclusion and Recommendation We have found enormous variation among projects claiming to be "Open Source". Therefore, it is imperative to ask any user of the term which of the Open Source features they really mean. Actually, many features claimed of Open Source can potentially be –and are at times– used in "conventional" projects. Consequently, we recommend that DIRC should have a further project activity that leaves aside the over-used term "Open Source" and studies group problem solving in the specific domain of software design and development (see Project Activity PA8 on "Effective collaboration in design (of dependable software)").

Ongoing and Future work Although PA5 was officially finished in December 2001, research on open source is still ongoing, especially in the form of BSc/MSc projects for students at Newcastle and City universities. There has been one completed MSc dissertation at City and one completed BSc dissertation at Newcastle on the topic of open source. There are currently four MSc projects (three at Newcastle and one at City) being carried out as well as another final year BSc project planned to start in September 2003 at Newcastle.

C.6 PA6: Security and Privacy in Computer Based Systems

Coordinator Peter Ryan, Newcastle

Partners involved Newcastle, City, Edinburgh

In 2002, effort on PA6 ran at a fairly low level with most of Peter Ryan's time consumed on the MAFTIA project. From the beginning of 2003, the resourcing has been stepped up with Peter now essentially full time and a new RA, Jeremy Bryans, recruited to work specifically on PA6 (with JGS funding provided from DSTL).

Span (nominal) 2002-01-01 (significant manpower available) 2003-01-01 to 2005-08-28

Scale of effort to date 2 RA years

Review dates PA6 will be replaced by a number of Targeted Activities, starting with the Chaum voting scheme study.

Aims and Objectives The purpose of this PA is to address security and privacy aspects of dependable systems from an interdisciplinary and socio-technical perspective. It extends the scope of DIRC in two directions:

- Investigate and characterise security and privacy requirements in dependable, distributed computer-based systems.

- Produce frameworks in which socio-technical threats and countermeasures can be accurately modelled and assessed.

A guiding principal of this PA has been to consider security from a socio-technical, rather than a purely technical perspective. Virtually all prior work in information assurance has taken a purely technical approach despite the recognition that many failures of security are due to, or at least facilitated, by social factors, see for example [MS02]. A notable exception to the neglect of socio-technical aspects is the work of Anderson, [And01] and the work by Sasse's group at UCL, http://www.getrealsecurity.com/.

It is felt that this programme is best pursued by making the issues concrete by identifying suitable case studies that will engage the various disciplines in DIRC. Initially attention was directed at identifying case study in the area of electronic patient records. Although this looked quite promising at first it has proved difficult to get good access to such case studies, despite good DIRC contact with the the PIMS and HIS trails in Scotland. As a result, little progress can be reported on this front at this stage but we intend to pick this up again when the time is ripe.

Early this year, an alternative area for case studies was identified: digital voting systems. In particular, a scheme proposed by David Chaum was identified as having particularly interesting socio-technical dependability aspects. We describe progress on this in greater detail below.

Summary of Progress This project has been running for a little under a year, with two clear strands of work having been pursued.

Budi Arief and Denis Besnard at Newcastle have performed a "Hacker Study", in which they have taken a cognitive psychology approach to the question of what motivates hackers to attack a system. A number of publications have already resulted, including [AB03, Bes01]. This study has now been broadened to investigate other players in computer security, in particular the impact that legal users of a system can have on its security [BA03].

The second strand in which significant progress has been made is the socio-technical study of the requirements for digital voting schemes, focussing in particular on a scheme proposed by David Chaum, [Cha]. This has been particularly successful in engaging the interest of all the disciplines of DIRC. It has had two official meetings: a one-day introduction to the scheme and a two-day workshop, both held in Newcastle, and attended by a broad range of the DIRC disciplines. It has interested a large number of DIRC people, and sparked wide ranging discussion relating to both the technical details of the scheme and the social context in which the scheme would operate (cf. [BR03a, BR03b]).

Electronic and digital voting are currently hot topics, certainly in the US and UK. DIRC is developing a number of contacts with authorities and experts involved in assessing the prospects of digital voting such schemes. In the UK, CESG has been involved in assessing the recent UK trails and is developing a proposal on which are providing comment. In the US we are involved in some of the recent discussions on Voter-Verifiable Election Systems, http://www.verifiedvoting.org/article_text.asp?articleid=46.

We expect this study to prove highly fruitful and it will form the focus on a new Targeted Activity. We anticipate that the Chaum Scheme study will run

for approximately another 6 months. After that new security and privacy based TAs will be proposed. Possible areas under consideration include:

- Air traffic control systems
- Electronic health records
- Applications in the financial sector

C.7 PA7: Dependable ubiquitous computing in the home

Coordinator Andrew Monk

Partners involved York, Lancaster

Span 2002-01-01 to 2004-01-31

Scale of effort to date 4 RA years

Review dates (progress) 2003-06-01; (final (planned)): ??

Aims and Objectives The aim of this project activity is to investigate how existing notions of dependability developed in the workplace apply in domestic settings and to extend dependability concepts, methods and tools for application in the home. To meet these aims we use assistive technologies for elderly and disabled people as a driver application. A secondary aim is thus to improve the dependability of these assistive home technologies. The objectives are to:

- identify dependability issues in the home and in particular in the design of assistive technology for independent living;
- develop frameworks for specifying dependability concepts that can form a basis for methods for dependable system design in the home;
- develop methods for dependable system design in the home.

Introduction The networked home, involving a sophisticated infrastructure of information and communication technology, is a commonly projected future market for domestic technology. In fact, networks of sensors that connect to external services already exist as telecare systems to allow people with disabilities and the elderly to live independently in their homes when they would otherwise have to move to some sort of institutional care. A networked home where a system may have control of door locks and communication out of the home, for example, is a critical system. However, the domestic environment is different from the workplace in ways that may radically change the conception of dependability.

The work completed in the first half of this PA has been directed at the first two objectives. We are now moving to turn this work into generalisable techniques for ensuring dependability in the domestic context.

Outcomes

1. How dependability issues in the home differ from dependability issues in the workplace

   A number of field studies have been carried out to answer this question including:

   - a study of a half-way house for psychiatric patients in Carlisle;
   - a study of how someone adapted their home after having a stroke;
   - the use of cultural probes (with Age Concern Barrow, MHA and Dundee);
   - studies of Assistive Technologies in use (with Shopmobility York, York Warden Call Service, an AT installation case study);
   - interviews with professionals working with frail elderly people (occupational therapists, social services and ambulance service personnel).

   The picture that has emerged is a need for a broader conception of dependability. There is much less control over how people use technology in their homes than there is in the workplace. Indeed, if a technology does not fit in with daily domestic routines or if it is ugly or hard to work it may not be used at all. A system that is not used cannot be dependable. These field studies have yielded detailed insights in this domain of assistive technology for independent living. They have also served to develop new field study techniques for the home context. Relevant papers are listed in the appendix.

2. Frameworks for specifying dependability concepts

   The detailed results from the field studies described above have been used to develop two frameworks that will be used in the next phase of the work, i.e., to develop techniques for ensuring dependability in

   the domestic context. The first framework extends Laprie's model of dependability by adding attributes under the headings acceptability, fitness for purpose and adaptability. The second framework extends the conventional view of risk analysis by defining generic types of harm that add social and psychological harm in addition to the usual physical harms to be avoided.

Conclusions and plans for the final 18 months The research we have carried out in the last 18 months has demonstrated clearly that the domestic environment raises interesting dependability issues that have not been previously considered in dependability frameworks developed for the workplace. It has required new field study techniques that have yielded new dependability issues, and new perspectives on old dependability issues. In the next phase of the work we will turn these results into new techniques for ensuring dependability.

Of course, many of the differences between the home and the workplace are differences of degree. For example, acceptability and adaptability are critical in the home and so studying the use of technology in home has allowed us to understand these dependability attributes more clearly. While such issues have not been previously considered in the workplace they may actually be very relevant.

Dissemination In the 18 months we have been going we have had accepted for publication two book chapters, three journal articles and numerous conference papers. We are planning a workshop to disseminate our work to practitioners and researchers in the new year.

C.8 PA8: Effective Collaboration in Design of Dependable Software Systems

Coordinator Cristina Gacek

Partners involved Newcastle

Span 2002-02-01 to 2003-06-01

Scale of effort to date 3.5 RA years

Review dates (final (planned)): TBA

Aims and Objectives As a result from the Project Activity on Open Source (PA5) it was recommended that DIRC leave aside the over-used term "Open Source" and study group problem solving in the specific domain of software design and development. That recommendation inspired this specific project activity.

The aim of this Project Activity is to study the strengths and weaknesses of practices used by designers and developers working in groups in the manufacture of dependable software. The intention was to include studies of human problem-solving in relation to software design, technical issues such as the communication and justification of system design, and the sociological aspects of working in teams on complex tasks.

The objective is to provide concrete recommendations on approaches to problem solving for the creation of dependable software systems.

Summary of Progress This project activity has been running for little over one year with very low effort consumption. It has had four official meetings to date, three in Newcastle and one in Edinburgh. In these meetings there have been participants from City, Edinburgh, and Newcastle. All DIRC related disciplines have been represented in the meetings, although sociologists were only present at the meeting in Edinburgh.

The major thrust of work has been in the psychology of programming, with experiments performed using students both from Newcastle and City. Solid results have been achieved in this area and have furnished material for an MPhil dissertation by Alessandra Devito da Cunha, as well as partial material for the PhD thesis of David Greathead. The work by Alessandra Devito da Cunha has been on exploring the relationship between personality traits and the ability of successful code debugging. David Greathead's work focuses on design activity.

Since development teams often involve many people, there are many relevant social issues relating to the software design process. These include –but are not limited to– social aspects of collaborative work while performing complex tasks. To date these aspects have not been pursued in this project activity.

Reflections The progress achieved on this Project Activity has been very interesting, especially given the low effort it consumed. It is unfortunate that we were unable to really involve sociologist in this work and had relatively little interactions between different sites. Consequently, this project activity is voluntarily dissolving itself in order to fit into the new Targeted Activities structure believing that to be a better approach to actively involve other sites and disciplines. One further possibility is to get involved in the study of "MMM" (cf. footnote 12).

C.9 PA9: Dependable Service-Centric Grid Computing

Coordinator Ian Sommerville

Partners involved Lancaster, Edinburgh.

Span 2002-11-01 to 2005-10-31

Scale of effort to date 1 RA year

Review dates (progress) to be decided; (final) to be decided

Aims and Objectives The overall aim of this Project Activity is to investigate how to specify, reason about and engineer dependable grid-based systems that are constructed by composing services from a range of providers . In this context, we have identified two principal research objectives:

1. To extend existing approaches to Quality of Service (QoS) definition and specification to encompass the dependability of grid services and to provide methods for reasoning about end-to-end QoS when services are composed
2. To investigate how existing approaches to dependable software engineering may be evolved to support service-centric grid computing.

Summary of progress Three RAs have now been appointed to this project (Conrad Hughes, January 2003, Edinburgh; Stephen Hall, November 2002, Lancaster; Glen Dobson, April, 2003, Lancaster). All of these RAs have extensive practical experience and, for this reason, we have focused on the work on evolving dependable software engineering techniques rather than more theoretical work on QoS specification. The focus of the work has been to extend work on fault tolerant architectures to grid services. We are investigating how redundancy and diversity in services may be used to provide high-reliability and availability. Our aim is to demonstrate a fault tolerance controller in September 2003 and, so far, we have developed software that simulates a service grid and allows grid nodes to exhibit different failure modes. We are investigating how to wrap BPEL4WS, a service orchestration language to handle exceptions and are currently working on the specification of the fault tolerance controller. A paper entitled 'Fault Tolerant Grid Services' has been accepted for presentation at the forthcoming e-science all-hands meeting in September 2003.

Glen Dobson (who started work on PA9 earlier than originally planned) has also developed what we believe is the first implementation of the Access Grid

that utilises large plasma screens rather than data projectors. We anticipate this will be a useful tool for collaboration on this and on other DIRC projects.

Reflections We are pleased with the progress on this activity. We have concrete results and excellent collaboration has been established between the RAs working on the project. Steve spent 2 weeks working in Edinburgh and this proved to be a very successful mode of working. It is obviously too early to be confident of success here but progress so far is definitely encouraging.

D Challenging Objectives

This was a response to Roger Needham's request at the first SC meeting that DIRC should look at "failable objectives".[23] [24] The body of this report to the mid-term review focuses on a subset of these objectives but this is the first list that was presented to the December 2001 meeting of the SC.

Relating to systems themselves

- Achieve a position where critical systems can be assessed and licensed at a tenth of the current cost.
- Achieve a position where much larger systems can be assessed and licensed than currently (say, ten times).
- The preceding two points should also facilitate the construction of "everyday" systems (where the lack of dependability might be a waste of resource rather than costing lives) by using methods for improving dependability that are currently only thinkable for safety critical systems.
- Publish at least one interdisciplinary book on each Research Theme.
- Develop notations for describing structure as it affects both computer systems and the human systems into which they are embedded.
- Propose a development method that makes explicit any assumptions about failures in components (including human errors).
- Provide a framework for relating logical and stochastic arguments.
- Provide a well-founded analysis relating testing, model checking, "proof" etc.
- To improve the state-of-the-art in formally justified dependability cases/arguments.
- Propose ways of combining disparate forms of evidence (testing, "proof" etc.).
- Devise ways of reasoning about very rare events.
- The work of psychologists Reason and Rasmussen provides a starting point for an analysis of errors made by individuals; DIRC should see that the sociological dimension is tackled as well.
- Document a requirement gathering method which handles complex human systems (into which a computer-based system is to be embedded).
- Propose a notation and reasoning system for notations of trust and responsibility.

[23] "Were you to achieve them all", said the wise man, "they would have been the wrong objectives".

[24] Colleagues in DIRC took the point of raising their objectives but considered the adjective "failable" as too pessimistic; we therefore adopted the term "Challenging Objectives".

- Catalogue and document the strengths of less known code analysis techniques (e.g. symbolic execution, partial evaluation, abstract interpretation). Some of these might play a role in testing.
- Contribute to the understanding of how programmers think and work in groups.
- Provide statistical evidence on development factors like how inspections are organised.
- Recommend a notation for time which covers humans and software.
- In all areas of concern, develop or select tools to support concepts.

Wider impact on society

- Link terminology and definitions between different disciplines.
- Influence standards so that they take a rigorous multidisciplinary approach to the evaluation of (dependable) systems.
- Influence the GRID architecture so that it becomes dependable in a way which makes it capable of fulfilling the expectations of its sponsors and users.
- Come up with something like a CMM model for the development of dependable systems.
- Influence Public procurement of software to reflect Dependability (taking a broad computer-based view).
- Aid the creation of a public watchdog concerned with dependability issues of proposed and actual computer-based systems
- Increase public understanding of software difficulties, risks and expectations by pro-actively using the media.
- Influence EU programmes to recognise broad nature of computer-based systems (wrt Dependability)

Other outcomes

- Build a cadre of interdisciplinary researchers for dependable systems design.
- (From experience) make recommendations about interdisciplinary, multisite, projects.

E Terms of reference for the mid-term review of the Interdisciplinary Research Collaborations

The aim of EPSRC in forming the IT-Centric Interdisciplinary Research Collaborations (IRCs) in 1990 was to recognise the potential of IT to facilitate inter-institutional working, plus its increasing importance as a component in a wide range of interdisciplinary activities.
The remit of the mid-term review panels:
A supportive review for the benefit of the IRCs and EPSRC

- Review the IRCs as a potential funding mechanism for future large projects;
- Consider the 'added value' generated from the approach and interdisciplarity of the IRCs;
- Comment and provide recommendations for the benefit of the IRCs to the IRC and EPSRC, especially in relation to the interfaces between the IRCs, users and EPSRC;
- Identify aspects of 'good practice', rather than 'best practice', for the IRCs.

The output from the mid-term review meeting will be fed back immediately following the meeting; in the form of one or two flip chart sheets as per the evaluation criteria given to the IRCs, shown below:

- Research Excellence or Quality of science (knowledge) delivered against milestones and any mid-term aims and objectives, including recognition of the team as leading players in the field;
- New opportunities for science enabled by the IRC;
- International recognition of the IRC and its research and engagement with the wider international community;
- Level of Collaboration with users, especially industry, and participation in a meaningful way in genuine collaborative research;
- Sustainability and re-use of the infrastructure (including people) provided through the IRC;
- Dissemination and other indications of exploitation where the project has disseminated research outcomes, and IPR;
- New ways of working including research, collaboration and technology transfer.
- Good management procedures and practices, including the adoption of risk management approaches.
- Plans for the future including longer term plans beyond the life of the IRC award.