Distributed cognition models for human factors failures in operating and design processes

Prepared by **Cranfield University** for the
Health and Safety Executive 2004

RESEARCH REPORT 203

HSE
Health & Safety Executive

Distributed cognition models for human factors failures in operating and design processes

John Strutt, John Sharp, Ed Terry, Jerry Busby
School of Industrial and Manufacturing Science
Cranfield University
Cranfield
Bedford
MK43 0AL

A set of about 60 offshore incident reports was obtained and combined with a small set of reports of inquiries into full scale offshore accidents and then developed into cause-effect analyses. In each case the task was to identify how people's problem solving had been distributed, how this distribution had failed, and the assumptions that they had been making, by implication. It was these implied but flawed assumptions that were then carried over into the development of the workbooks. A further step was then taken to identify general design principles that would help make systems less vulnerable to each assumption type.

Computer based workbooks were developed, which allow users to consult and reason about the flawed assumptions discovered in the accident analyses. They also allow users to consult and apply the design rules that were identified in the preceding analysis.

The work also included an investigation of design error, which involved analysing a set of reports of error in the design process of firms designing both onshore and offshore installations, in combination with two group elicitation sessions. These reports were then analysed in a similar way to the accidents - by identifying how cognition had been distributed, how this distribution had failed, and the flawed assumptions that were implicated in these failures. Appropriate workbooks, similar to those developed for the accident analysis, werer also developed to help designers and managers of the design process anticipate flawed assumptions in the future by consulting assumption types known to be flawed in the past.

This report and the work it describes were funded by the Health and Safety Executive (HSE). Its contents, including any opinions and/or conclusions expressed, are those of the authors alone and do not necessarily reflect HSE policy.

HSE BOOKS

First published 2004

ISBN 0 7176 2910 4

D3916: DISTRIBUTED COGNITION MODELS FOR HUMAN FACTORS FAILURES IN OPERATING AND DESIGN PROCESSES

CONTENTS

1 INTRODUCTION

Background

The proposal originated in previous work undertaken on making inferences from accidents and failures, which were failures both in the engineering process and in engineered systems. Most of the failures that had been studied involved people's use of designed objects and could not be attributed to the person or to the design exclusively, despite the fact that most classifications did simply classify them as being one or the other. Moreover, the contributors to the failure that suggested a clear error on the part of a person typically occurred when the person took their cues, or their solutions, from some part of their environment (like another person, a procedure or a design). Contributors that suggested a clear technical failure such as an exploding vessel typically occurred when the designers had taken their solutions or analyses from some part of their environment (like standards, codes or other people). The aim of this study, therefore, was to look at 'distributed cognition' and how its failure contributed to accidents.

The basic principle is that the representation and transformation of knowledge does not just lie in the individual human mind. For example, when people solve problems, it is not just their own thinking that determines the solution. They get parts of their solution from other people's behaviour, from custom and practice, from organisational procedures, from the tools they use and so on. Sometimes they follow a routine they have watched other people perform. Sometimes they follow a code of practice developed in the past. Sometimes they read an instrument on the basis that it works in the same way as instruments they have formerly encountered. Sometimes they manipulate their surroundings to help them cope with too much information (like marking a device to remind them to operate it). In fact people would not be able to accomplish what they do accomplish without this distribution.

The principle originated in the early 1990's and is associated most strongly with Edwin Hutchins (1,2) and Don Norman (3). There have been several research papers on this subject, (4-10) and it has found its way into the analysis, mainly of real-time operating situations like marine navigation (1) and air traffic control (2). It has also influenced work on human-computer interaction (11). The Cranfield/Bath work has extended the idea into design processes (12,13) and how they have failed. It has helped to show how far people's understanding and learning is influenced by the context in which it takes place.

Distributed cognition remains controversial, however. Some members of the research community regard cognition as something that by definition is located in the human mind, and even when there are many minds at work on a problem this doesn't amount to distribution of cognition, only of the task. The idea that cognition can be distributed over inanimate artefacts makes no sense in this view. Other researchers argue, however, that cognition is concerned with the processing and representation of knowledge, and the propagation of knowledge representations across various entities - entities like people and the tools they use. This means that a natural unit of analysis when looking at cognition is not the individual human mind. One of the advantages of this latter view is that it makes cognition more observable. If knowledge is represented and processed in tools that lie outside the mind, for example, this knowledge and processing can be seen directly as what is, in a way that is impossible with human cognition.

Purpose

The purpose of this project was to tackle two main problems:

- To find out how distributed cognition failed during the operation and maintenance of complex, hazardous systems.
- To find out how distributed cognition failed during the course of the design process.

The principle was to do two main things:

- To analyse a set of past cases to determine how distributed cognition had failed, as best that could be done;
- To develop tools to help people make systems resilient to such kinds of failure in the future.

This was all to be done in the context of the offshore industry, using offshore accidents and attempting to influence operators and engineers of offshore installations.

2 PROGRAMME

Main activity

The following is a brief summary of the methodology, more details are given in annex 1:

- Analysis of operational accidents. A set of about 60 offshore incident reports was obtained from the HSE, and these were combined with a small set of reports of inquiries into full scale offshore accidents. These reports were then developed into cause-effect analyses. In each case the task was to identify how people's problem solving had been distributed, how this distribution had failed, and the assumptions that they had been making, by implication. Details of the model used are given in the articles collected in the Annex. It was these implied but flawed assumptions that were then carried over into the development of the workbooks. A further step was then taken to identify general design principles that would help make systems less vulnerable to each assumption type.
- Development of workbooks. These were computer-based workbooks, one for the industry and one for inspectors (but both took very similar forms). Essentially the workbooks allow users to consult and reason about the flawed assumptions discovered in the accident analyses. They also allow users to consult and apply the design rules that were identified in the preceding analysis.
- Investigation of design error. This involved analysing a set of about 60 reports of error in the design process of firms designing both onshore and offshore installations, in combination with group elicitation sessions at Kvaerner and AMEC. These reports were analysed in a similar way to the accidents - by identifying how cognition had been distributed, how this distribution had failed, and the flawed assumptions that were implicated in this failure.
- Development of workbooks. Again these were similar to those developed for the accident analysis, helping designers and managers of the design process anticipate flawed assumptions in the future by consulting assumption types known to be flawed in the past.

Subsidiary activity

Various subsidiary activities, mainly involving liaison with others, took place in support of the project. A brief list follows:

- Observation of HAZOP meetings at Kvaerner Process.
- Attendance of FABIG meetings by E J Hughes
- Attendance at UK Safety & Reliability Society meeting by E J Hughes
- Meeting with the University of Aberdeen on their prior work for HSE
- Meeting with DnV to discuss role of the study work in the verification process
- Contact made with Keil Centre on prior work for HSE, especially concerning handover
- Meeting with Nickleby human factors consultants on their work for the HSE
- Meeting and seminar with AMEC

Variations to the plan

All the main activities were completed on time and the outputs promised in the contract were delivered before the planned end-date. In the course of the work it became apparent that additional work would be valuable:

- An analysis was undertaken of how barriers had failed in the accidents and incidents that formed the dataset. This provided important information about how various planned and unplanned protective measures had been undermined in practice. This led to the development of a prompting tool that helps operators and designers review the hazard barriers available, and test whether they are vulnerable to being undermined by the same phenomena that were found in the analyses.
- The main outputs (the computer based workbooks) were in the form of 'Access' databases. During discussions about how these might be commercialised it was agreed that they should be re-written in part in HTML to be available as web pages.

A six month extension was made to the original plan to enable the above work to be completed.

The programme was accomplished within budget.

3 IMPLIED ASSUMPTIONS & GUIDELINES

There were about 30 distinct assumptions that were identified in the analysis as having contributed to the incidents or accidents. There were three main categories of such assumption, whose titles should be self-explanatory:

- The appropriate organisation assumption.
- The knowledgeable individual assumption.
- The reasonable system assumption.

It is the last category that is of most relevance to designers as it contains assumptions about the designed system that are implied in the way operators and maintenance staff behave. Table 1 lists the particular assumptions under this heading, together with a very brief summary of a relevant case.

Table 1: Sub-categories of the 'reasonable system' assumption

Assumption	Case examples
Ambiguous things do not matter	Fitter replaced part of a blowout preventor wrong way round. This then failed when there was a blowout. Implied assumption was that if something could be fitted in different ways then it didn't matter how it was fitted.
Boundaries are obvious	Operators had adopted a ballasting practice which allowed rapid listing. This ultimately contributed to capsize. Implied assumption was that boundaries to safe operation would be obvious.
Consequences are obvious	Damage caused to sacrificial anodes by pile driving in construction. Damage only obvious during operation. Implied assumption was that if an operation was harmful then the damage would be obvious at the time.
Function follows appearance	System started up with only a blanking plate preventing escape of gas. Possibly fitter thought it would prevent egress of vapour, not just ingress of dirt, because it was solid. Implied assumption was that the solid appearance of the plate meant gas would not escape from aperture.
Identification cannot go wrong	Drain cut to install break couplings. Second pipe with similar shape also thought to be drain so also cut, but in fact had different function. Implied assumption was that can identify the right objects to work on based on a similar appearance to other objects
Lapses will not imperil the design	Safety hatches incorporated in design for intermittent tasks. Hatches left open which contributed to capsize. Implied assumption was that the design would not be vulnerable to simple lapses and violations.
No signals means all is well	Radio used for communication during crane operations. Channel had unknowingly failed and 3 'stop' messages not heard by the operator. Implied assumption was that absence of messages meant sender had nothing to communicate - not that channel had failed.
Redundancy protects systems	Area had to be cleared for radiography. Done both by detection (sending someone to look) and self-detection (making a tannoy announcement). Both failed probably because the other was assumed to be more effective.
Sequences of actions are not interrupted	Instrument line disconnected during planned maintenance but not reconnected before startup. Implied assumption was that sequences of activity cannot be interrupted or forgotten.
Things happen in a logical order	During a crane lift slings failed when load was snagged. Operator probably not attentive, expecting that passive slings would not fail before active motors reached limit. Implied assumption was that the properties of the system would follow a natural order.
Trial and error is not hazardous	Wrong pump in a pair dismantled. Noise masked sound of running pump and poor lighting impeded visual identification. Implied assumption was that you could identify a device by trial and error and would know if you got the wrong one.
What is available is what is appropriate	Connection for low pressure line instead made to high pressure line which was the only one available. Implied assumption was that whatever was available at the time was appropriate to the task.

The next step was to formulate design guidelines that would either reduce the likelihood that vulnerable assumptions would be made, or, if they were, reduce the vulnerability of the system to them. These guidelines were then grouped. They fell fairly naturally into four categories:

- Information guidelines - guidelines that help the designer tell the operator or maintainer how or what to do.
- Salience guidelines - guidelines that help the designer show the operator what is important at a particular time.
- Restriction guidelines - guidelines that help the designer constrain what the operator does.

- Presumption guidelines - guidelines that help the designer know what to presume or predict about the operator.

Table 2 shows the guidelines within the two categories of 'Information' and 'Salience' guidelines. The guidelines in the other two categories are included on the CD-ROM.

Table 2: 'Information' and 'salience' guidelines

Information guidelines	Help users know the functions an object will not perform
	Make positive recommendations about replacement of devices that might be worn or degraded
	Help users take account of imperfect tests
	Differentiate the appearance of devices which perform the same function in different parts of a configuration
	Differentiate the appearance of devices which have different functions
	Differentiate the appearance of objects for viewing from several orientations
	Differentiate similar objects that nonetheless have to be configured or connected in opposite ways
	Provide tests of a good outcome rather than relying on perfect procedures always followed
	Provide positive indication of malfunction where this cannot be differentiated from normal absence
Salience guidelines	Make hazardous states highly visible
	Avoid giving the impression that an orientation is arbitrary
	Avoid giving the impression that choice of a material is arbitrary or will not be detectable
	Maximise the obviousness of missing components following construction or maintenance
	Maximise the visibility of the system left in a state where is should not be started up
	Avoid contradicting intuitive orderings e.g. active devices stall before passive restraints fail
	Try to match the salience of an attention seeking device with its importance

4 OUTPUTS

The following provides a summary of the outputs of the work:

- Computer based workbook for the identification of flawed assumptions during operations and associated design rules (for risk identification in operations and design)
- Computer based workbook for the identification of flawed assumptions during operations and associated design rules (for inspection, audit and investigation)
- Computer based workbook for the identification of flawed assumptions in the engineering design process (for the industry)
- Computer based workbook for the identification of flawed assumptions in the engineering design process (for inspectors)

- Web based prompting tool for the identification of flawed assumptions during operations and associated design rules
- Proof-of-concept computer based prompting package for the inspection of barrier vulnerability
- Final report

Details of the tools, their basic structures, rationales and evaluations are given in some of the articles included in the Annex.

The direction of current discussions in making the workbooks available is that the Microsoft Access-based packages could be provided at a charge, but there should be free availability of the more limited web-based tool.

5 BENEFITS

The primary benefits of the work have been:

- Exploiting and disseminating the lessons from several tens of accidents and incidents. It is valuable for this information to be extracted from within organisations, such as a regulator, and made available more widely.
- Synthesising the information available in accident and incident reports. To read them all is time consuming and takes considerable effort. The analysis and the workbooks from this study effectively do this for people.
- Introducing an important and relevant concept to the industry. The idea of distributed cognition is a little abstruse and it has academic origins. But it is universal, failures in it do lead to accidents, and it does help broaden the notion of 'human error' away from simply being a failing in people's minds.
- Helping people test their assumptions. The assumptions that people make about systems and other people are central to safe operations, and almost always implicated when there is a failure of some kind. Assumptions that are there by implication are especially vulnerable. The idea of helping people test their assumptions, and make systems less vulnerable to doubtful assumptions, therefore seems to be an important one.

6 DISSEMINATION

The main dissemination activities during the project were as follows:

- A paper to the Hazards XVI conference (14)
- A paper to the Hazards XVII conference for presentation in 2003 (15)
- A presentation to the OIMs forum in Norwich
- An article to the FABIG newsletter (16)
- A presentation at the ERA conference on Major Hazards (17)
- An article to International Journal of Risk Assessment & Management (18)

Copies of the manuscripts are included in the annex.

7 FURTHER WORK

Incremental work

The outputs that help people test their assumptions, and help designers make systems less vulnerable to poor assumptions, synthesise the knowledge available in about 70 accidents and incidents in the offshore sector. Natural extensions to this would be:

- To analyse a further body of offshore accidents (say another 100) to make the analysis more complete and the tools more comprehensive.

- To analyse accidents in other sectors, particularly those with large scale hazardous systems such as onshore process plant, marine vessels, air and rail transport. One of the advantages of doing this would be to address the problem that learning from the past may blind one to the kind of hazards arising from future technologies, particularly technologies with high levels of automation.

Fundamental work

The most promising fundamental work would be a full-scale project on barrier undermining. Some barrier-based methods of analysis are available - such as safety barrier diagrams, accident evolution and barrier models, and energy trace and barrier analysis. But none of these build in knowledge of how barriers are undermined. In the brief analysis in this study a number of quite subtle phenomena were found that undermined barriers. The most troubling were perhaps effects that undermined barriers incorporated by designers, and indicated that designers had unrealistic views of what the residual risk was in a system after they had designed protective measures of various kinds. Some barriers were so misconceived that they were introduced on the basis that operators were too unreliable to protect a system, yet induced the operators to perform actions that considerably increased a hazard. It was felt, following the very brief analysis of barrier undermining (using a set of drilling accidents) that there was considerable scope both for investigating how barriers become undermined more generally, and helping the industry in very practical ways to improve barriers.

8 ANNEXES

- 1 -Original proposal
- 2. Published articles on the work (on separate CD-ROM)

9 REFERENCES

1. Hutchins E., (1995), *Cognition in the Wild*, The MIT Press, Cambridge MA, p.155.
2. Hutchins, E. (1995). *How a Cockpit remembers its Speed*. Cognitive Science, 19, 265-288.
3. Norman D.A., (1993), *Things That Make Us Smart. Defending Human Attributes in the Age of the Machine*, Addison-Wesley, Reading MA, p.146
4. Salomon G., (1993), *Editor's introduction*. In Salomon G (ed). Distributed Cognitions: Psychological and Educational Considerations, Cambridge University Press, Cambridge UK, xi-xxi.
5. Marti P., (2000), *The choice of the unit of analysis for modelling real work settings*, Cognition, Technology and Work, 2: 62-74.
6. Lave J., (1988), Cognition in Practice, Cambridge University Press, Cambridge UK
7. Scaife M. and Rogers Y., 1996, *External cognition: how do graphical representations work?* International Journal of Human-Computer Studies, 45: 185-213.
8. Rasmussen, J. (1983). *Skills, rules, and knowledge: signals, signs, and symbols, and other distinctions in human performance models*. IEEE Transactions on Systems, Man, and Cybernetics, 13, 257-266.
9. Reason, J. (1990). *Human Error*, Cambridge UK: Cambridge University Press.
10. Marti, P., (2000) The *choice of the unit of analysis for modelling real work settings.* Cognition, Technology and Work, 2: 62-74.
11. Wright, P.C., (2000) Fields, R.E. and Harrison, M.D., *Analyzing human-computer interaction as distributed cognition: the resources model.* Human-Computer Interaction, 15: 1-41.
12. Busby J.S., (2001). Error and distributed cognition in design. *Design Studies*, **22**(3), 233-254.

13. Busby J.S., (2001), Practices in design concept selection as distributed cognition. *Cognition, Technology and Work*, **3**, 150-160.
14. Busby J.S. Hughes E.J., Sharp J.V., Strutt J.E.,Terry E., (2001).Distributed cognition and human factors failures in operating and design processes. *Hazards XVI*, Manchester, 6-8 November.
15. Busby J.S., E Terry E., Sharp J.V., Strutt J.E.,Lemon M., (2003). How distribution in human problem solving imperils systems. *Hazards XVII*, Manchester, 24-27 March.
16. *'How distribution in human problem solving imperils systems'*, FABIG Newsletter, May 2002
17. Busby, J S, Sharp J.V.,Strutt J.E.,Terry E., Hughes E.J.,Miles R.M., (2002) *Distributed Problem Solving and Offshore Accidents, ERA Conference on Major Hazards Offshore*, London, 25-26 November,2002.
18. Busby J.S. Hughes E.J, (2003). The role of distributed cognition in the causation of accidents. *International Journal of Risk Assessment and Management*, **4**, 36-51.

ANNEX 1

Original proposal

Project Title	Distributed cognition models for human factors failures in operating and design processes
Organisation Name	Cranfield University, School of Industrial & Manufacturing Science, in collaboration with the University of Bath, Sauf Ltd and Kvaerner Oil and Gas Ltd
Project Duration	**18 months**
Cost	£98,788
Summary	The distributed cognition model is an important framework for understanding the way that human problem solvers rely on the environment to accomplish their tasks. They rely on cues they receive from human co-workers, they make inferences from the appearance of the artefacts they work with, and they draw on organisational culture to work out what is expected of them. In particular, distributed cognition models have the potential to help us understand how accidents arise when people operating, installing, maintaining and repairing equipment make the wrong inferences about that equipment. They also help us understand how specialists of different disciplines sometimes fail to influence each other appropriately during the design of equipment. The purpose of this project is therefore to use the distributed cognition principle to help designers reason both about human factors failures in the operation of equipment and failures in the process of design. We also want to find out how aspects of the designers' environment, such as safety regulation and safety management systems, influence this distributed cognition. The two primary products will be 1) a workbook for offshore design organisations to help them anticipate human factors failures in both design process and designed product, and 2) a guidebook for regulatory organisations and senior industrialists on the influence of the environment on designers' thinking.

Introduction

In one of our recent studies, we analysed a case in which a maintenance engineer was killed when he used a beam that was apparently, but not actually, anchored in order to suspend a hoist. The designers could have forestalled the accident if they had been able to predict this kind of behaviour at the accident site. They could also have forestalled the accident if they had realised that there were certain cues that indicated they should consult human factors or safety engineers. And it is possible that with a more influential safety management system, a more influential professional culture and more influential regulation the designers would have actively sought such cues. The compelling lessons of accidents such as this are therefore that:

- designers need a model to help them reason both about failures in the process of designing an installation and failures in the process of operating it;
- regulators, professional institutions and managers need an understanding of how the regulatory, professional and managerial environment influences designers' reasoning about failures.

Distributed cognition provides a framework that, potentially, can provide a consistent way of tackling these needs. The great promise of having a single model that helps design organisations reason both about human factors in the operating process, and about how they use human factors knowledge in the design process, is that a much clearer connection is made between how designers think and the accidents they ultimately contribute to. Such a model also provides an anchoring principle for investigating how the environment influences designers' thinking. Most studies of distributed cognition have concentrated on the operating situation, and how the design of equipment can mislead people on flight-decks and ship's bridges. But our own recent studies have used distributed cognition to understand errors in the engineering design process, and in particular the problems that arise when people in different disciplines collaborate on design projects. This includes, for instance, the way in which the cues exchanged between disciplines are often misleading.

The model that follows shows the scope of the project. It shows distributed cognition in both design and operating processes, and shows how we would like to model both the distributed cognition and the way the environment influences it:

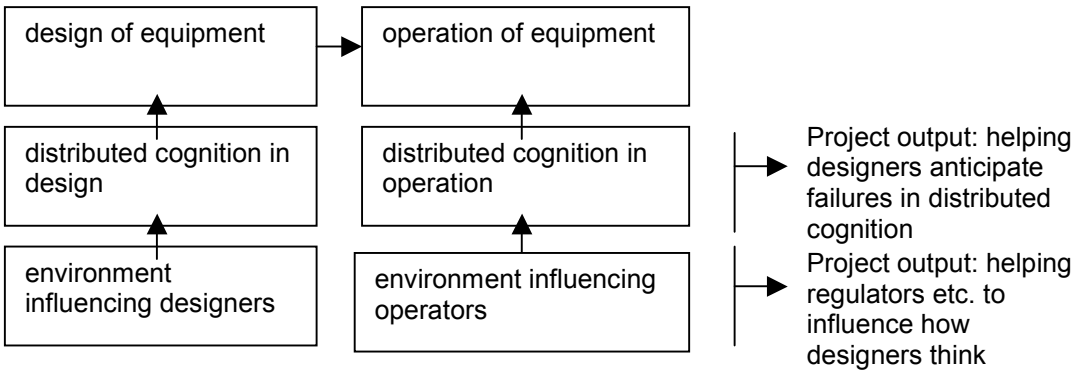

Objectives

The gist of the project is to analyse past failures, use these analyses to develop distributed cognition models of failure, and build these models into practical tools.

Detailed Objectives

- To investigate how distributed cognition failures contribute to accidents in the process of operating offshore installations.
- To investigate how distributed cognition failures contribute to shortcomings in the process of designing offshore installations.
- To investigate how the regulatory, professional and organisational environment influences distributed cognition in the processes of operating and designing offshore installations.
- To develop models and workbooks that support the design organisation in avoiding failures in both the design and operating processes.

- To develop models and guidebooks for the HSE that help staff involved in inspection, audit and policy development to assess the ability of design organisations to reason about failure and to assess the effect of regulations and guidance on designers' thinking.

Benefits

- One of the main difficulties that designers have in thinking about human factors is their lack of a general model that helps them enumerate potential problems systematically. This work ought to provide such a model.
- Similarly, one of the main obstacles in all collaborative work - but particularly collaboration involving a technical discipline and a social discipline - is the lack of a model that helps each discipline ask the right questions and send the right cues to the other discipline. Again this work ought to provide such a model.
- Because accidents are relatively infrequent, it is very hard to be confident that one organisation's historical experience is enough to help it predict and prevent future accidents. Having a model gives some assurance that there are not large gaps in an organisation's understanding of how things can go wrong.
- A model of this kind would also help reveal any limitations and gaps in current HSE guidance on incorporating human factors in the offshore design process. In particular, by understanding both the manner in which cognition is distributed, and how this is affected by such factors as regulation and knowledge of regulation, we hope that it will become clearer how to influence designers' thinking.

Deliverables

- The first main deliverable is a workbook. This will provide tutorials on distributed cognition both in the operating process and in the design process of an installation. It will provide models for each, and checklists of known failure modes associated with them. As a workbook it will be something that the users (typically lead designers) have to interact with, providing both instruction and a resource for maintaining a record of the reasoning that takes place during the design process. The workbook will be provided in computer-based form, built on a proprietary database management system.
- The second main deliverable is a guidebook to the HSE. This will provide guidance both on the distributed cognition we investigate and on what design organisations should be doing to manage this distributed cognition in a way that maximises safety. It will then provide further guidance on how the safety environment (especially regulation) influences and fails to influence this distributed cognition. Our intention is to provide this guidance in a layered form, such that the user will be able to choose from a variety of different levels of detail for consulting the material. Again this will be a computer-based package.
- A paper will be written for an academic, archival journal, and a paper for the industry will be delivered at a conference.

Methodology

Investigating distributed cognition failures in the process of operating installations

This will involve three main steps:

1. Building a case base of past accidents and incidents.
2. Analysing the distributed cognition that occurs in these cases.
3. Developing classifications of this distributed cognition and its failure modes.

The cases will be obtained from 1) public domain sources of investigative reports, particularly the MAIB digests, the HSE's UK Continental Shelf Risk Review, proceedings of public enquiries into offshore disasters, Loss Prevention Bulletin reports on process plant accidents, industry case bases such as SIREN, and case studies in case study texts (notably those of Trevor Kletz); 2) the HSE's incident and early day reports. All the investigative reports will be formalised, according to our normal practice, in order to represent the pattern of causation in a systematic way. Sub-patterns representing failures of distributed cognition will then be identified. Based on our past experience these are likely to include, for example:

- operators wrongly inferring from its external characteristics how a piece of equipment will function;
- operators searching for short-cuts in operating procedures and unknowingly violating the required sequence of operation for a piece of equipment;
- operators finding that a feature on some equipment confounds their intentions and adopting a hazardous practice;

- operators being unable to work together effectively when equipment interferes with communication or places them far apart.

The analysis will then develop a classification of such failure modes.

Investigating distributed cognition failures in the process of designing installations

The failures of the design process, such as in the collaboration between designers and human factors specialists, are naturally harder to obtain because they do not usually reach the public domain. Our intention is to draw on three sources of cases:

- An EPSRC-supported study of error in the design process that we conducted recently with a process plant design organisation. This yielded a database of 86 cases, whose causation has already been captured and formalised.
- An elicitation exercise with the offshore installation design organisation that is participating in the project. Our intention is to get designers' observations of how the design process has failed - especially in the collaboration among different disciplines.
- An elicitation exercise with the various disciplinary specialists in the research team. The purpose is to draw as systematically as possible on their experience of consulting, in particular, with firms in this and related industries on failures, breakdowns and limitations in the design process which have introduced hazards in the equipment being designed.

As with the accident analysis, the plan is to identify in these cases the nature of the distributed cognition and the modes by which it failed - and build a classification of these.

Investigating the environmental influences on distributed cognition in design

There are two elements to this part of the investigation. The first is again to elicit the observations of designers in the participating design organisation of what influences their design decisions and the associated thinking about hazard. The second is to elicit the knowledge of the disciplinary specialists in the project team, and in particular to obtain data from their different standpoints: risk management (J E Strutt), human factors (D Harris), sociology of multi-disciplinary organisations (M Lemon), regulatory regimes in the offshore sector (J V Sharp), safety engineering in the offshore sector (E Terry) and error in the engineering process (Jerry Busby). This elicitation will in turn take two forms: 1) each of the specialists will prepare a briefing document on their subject area and what is known in the subject area about the problem of influencing designers' thinking, and 2) all of the specialists will participate in joint elicitation exercises. The intention is to develop as broad a classification as possible of the influences on designers' thinking, and their thinking about other people in particular (both those they work with, and those who operate their products).

Development of distributed cognition models

The output of the three preceding stages will be three classifications: how distributed cognition fails in the operation of equipment, how it fails in the design of equipment, and how it is influenced by the safety environment. In this stage the intention is to provide a general, qualitative model that captures the failure mechanisms and their interaction with the environment. To cope with the potential complexity, this is likely to be a layered model in which detail can be exposed selectively. Our aim, however, is a model which can show how the parts are connected: how, for instance, organisational cultures influence designers' failure to anticipate operators' misuse of equipment; but also how, for instance, the nature of regulations might influence designers to make commitments before analysing hazards.

Development of the workbook

The purpose of the workbook is to make the models operational: to give, in particular, design leaders clear guidance on what they are and how to apply them - and give them a resource for recording their thinking and their consultation. Our proposal is to build the workbook on top of the knowledge base of cases so that the users can inspect cases related to the models as well as the models themselves. It is envisaged that the models will contain graphical representations of the distributed cognition involved in operating equipment and collaborating with other disciplines, attach to the elements of the model the failure modes that we know characterise them, and attach the cases to these failure modes. It is also envisaged that the workbook will have two basic parts: an offline training element, and an online support element. The first would be aimed at familiarising lead designers with the model, while the second would be aimed at providing an aide memoire for the design process. This will take users methodically through these failure modes to help them reason about the particular design and design process they are engaged in.

An assessment and refinement will be made of the workbook by 1) getting members of staff in the participating design organisation to work through it under observation, and 2) running seminars with various industry bodies (such as the British Chemical Engineering Contractors Association and the Engineering Industries Council, on whose committees Ed Terry sits).

Development of the guidebook
The purpose of the guidebook is again to make the models operational, but this time for HSE staff. As we have already described, the intention is to guide HSE staff on the distributed cognition itself, on what organisations should be doing to manage it, and on how the safety regulation and guidance environment influences it. The intention is to do this in a layered way, such that over a series of levels (probably 3 or 4) a straightforward, graphically-organised guide is given on these issues.

Management

The plan shows the main tasks described in the Methodology section, and the division of labour:

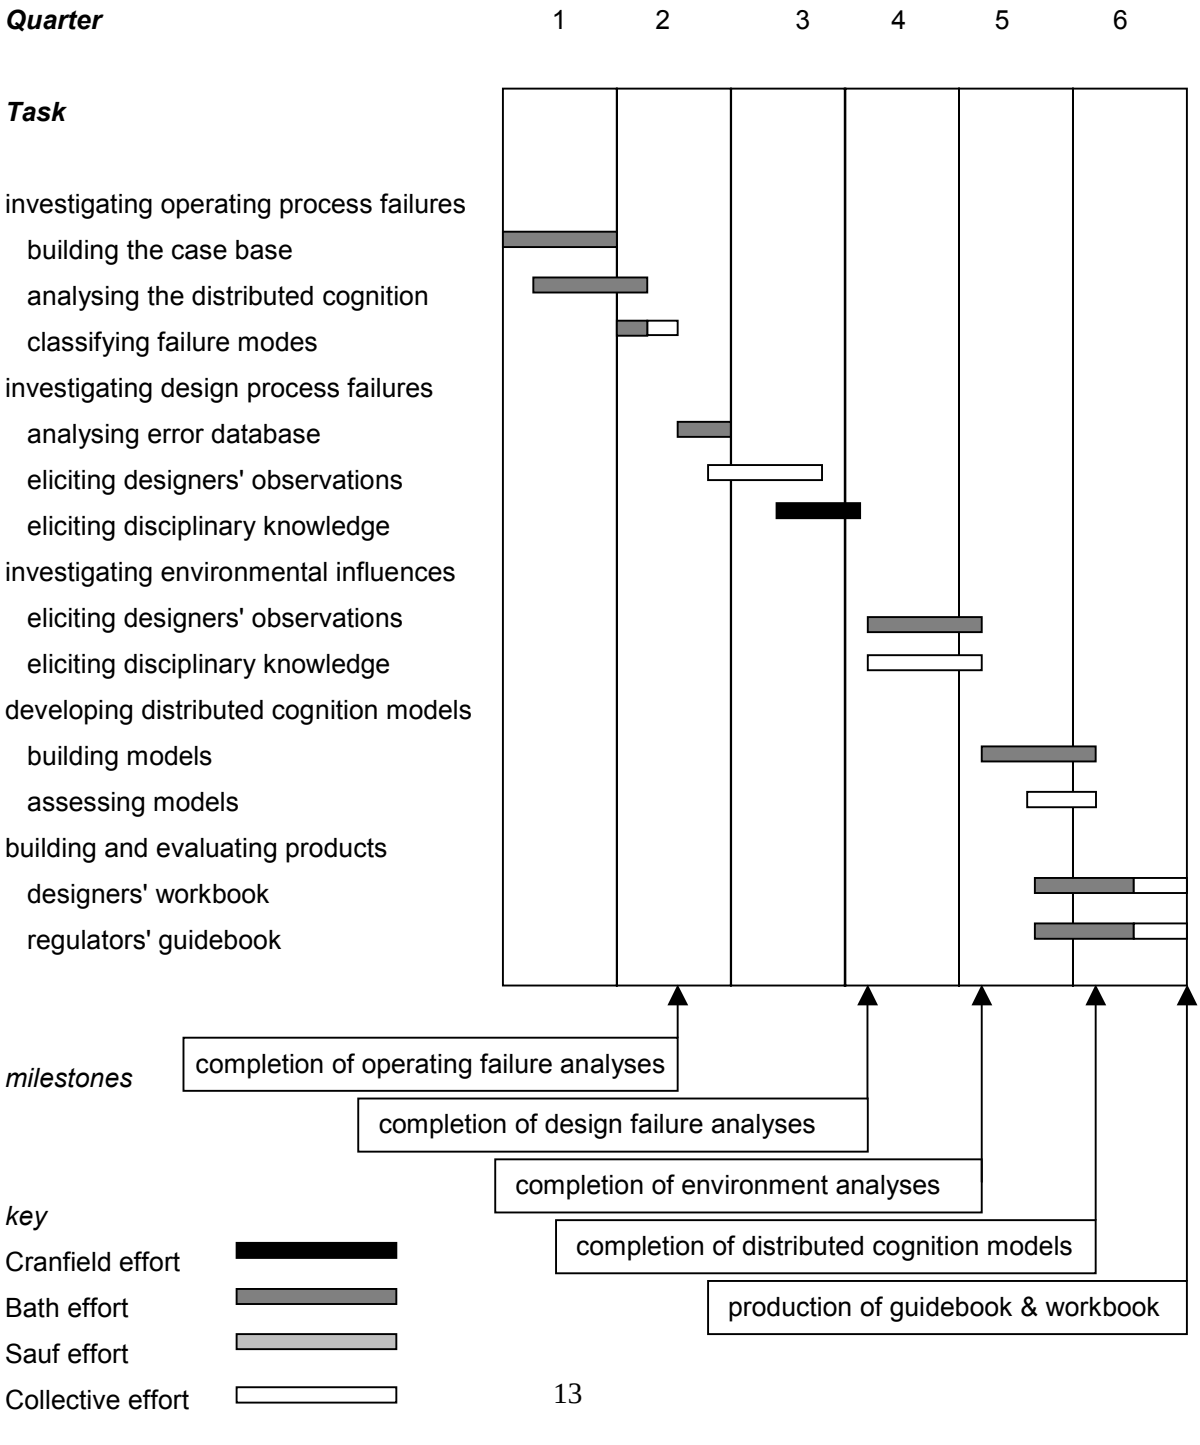

Experience

Relevant experience
The research team consists of Dr J S Busby, Prof. J E Strutt, Prof. J V Sharp, E. Terry, D Harris and Dr M Lemon. Busby and Strutt have had numerous research council grants, and research contracts from both operators and designers of installations, in the fields of engineering design, risk and safety, and experiential learning. They have worked together on a recent HSE contract, for which Cranfield was the lead organisation, on offshore design safety performance indicators (HSE8890). Busby is a lecturer in engineering design and an EPSRC advanced fellow. Strutt is professor of reliability engineering and risk management, and is chairman of the IMechE's Reliability Committee. Sharp is formerly head of HSE's offshore research programme and is a visiting professor in the School of Industrial and Manufacturing Science. He has considerable experience of participating in the drafting of national and ISO standards in the sector. Ed Terry is a safety engineering consultant in the offshore industry, formerly managed large safety engineering teams in offshore design organisations and serves on a number of government and industry committees. Harris is a senior research fellow in the Human Factors Group at Cranfield and has wide-ranging experience in the aerospace and automotive industries as well as offshore installations. He runs the biennial International Conference on Engineering Psychology and Cognitive Ergonomics. Lemon is a lecturer in sociology and, in particular, has investigated problems of multi-disciplinary working in organisations.

The team and its competences
The team draws together competences in 1) understanding organisational failures (principally Busby), 2) understanding technical failures (principally Strutt), 3) understanding the regulatory regime and role of HSE (principally Sharp), 4) understanding the role and influence of the safety engineering discipline in design organisations (principally Terry), 5) understanding human factors and the role of the human factors discipline in engineering organisations (principally Harris) and 6) understanding the problems of multi-disciplinarity in organisations (principally Lemon). This provides the rich expertise needed both to analyse the past cases and to develop the kind of models that can help integrate the design and human factors disciplines. Members of the team have particular experience in using distributed cognition to analyse failures in the engineering process, of building computer-based workbooks, of developing knowledge bases of historical accident causation, and of improving design-for-safety in the offshore industry.

ANNEX 2
(papers presented on attached CD-ROM)

J S Busby, J V Sharp, J E Strutt, E Terry, E J Hughes and R Miles, *Distributed Problem Solving and Offshore Accidents*: ERA Conference on Major Hazards, London, 2002, Copyright with ERA
How distribution in human problem solving imperils systems: FABIG Newsletter, Issue 32, May 2002, pp 24-27 and submitted to Offshore Research Focus

J S Busby, E J Hughes J V Sharp, J E Strutt and E Terry: *Distributed cognition and human factors failures in operating and design processes:* HAZARDS XIV, Manchester, 6-8 November, 2001-10-04. Copyright with IChemE

J S Busby, E J Hughes E Terry J V Sharp, J E Strutt and M Lemon: *How distribution in human problem solving imperils systems:* HAZARDS XV, Manchester, 2003.
Copyright with IChemE

Busby J.S. Hughes E.J, (2003). The role of distributed cognition in the causation of accidents. *International Journal of Risk Assessment and Management*, **4**, 36-51.

Printed and published by the Health and Safety Executive
C0.06 12/04

MAIL ORDER
HSE priced and free
publications are
available from:
HSE Books
PO Box 1999
Sudbury
Suffolk CO10 2WA
Tel: 01787 881165
Fax: 01787 313995
Website: www.hsebooks.co.uk

RETAIL
HSE priced publications
are available from booksellers

HEALTH AND SAFETY INFORMATION
HSE Infoline
Tel: 08701 545500
Fax: 02920 859260
e-mail: hseinformationservices@natbrit.com
or write to:
HSE Information Services
Caerphilly Business Park
Caerphilly CF83 3GG

HSE website: www.hse.gov.uk

RR 203

£10.00

ISBN 0-7176-2910-4

9 780717 629107