

HOW DISTRIBUTION IN HUMAN PROBLEM SOLVING IMPERILS SYSTEMS

J S Busby (University of Bath), E Hughes (Marine and Coastguard Agency),
E Terry (Sauf Ltd), J V Sharp (Cranfield University), J E Strutt (Cranfield University),
M Lemon (Cranfield University)

The distribution of people's problem solving is an important source of failure in complex, hazardous production systems. An analysis was made of about 60 offshore accidents and incidents in an attempt to work out how people distribute their problem solving and how this fails. Inferences were made about the assumptions that were effectively being made in this distribution, and how the system could have been made less vulnerable to such assumptions in the design process. A prompting tool has been developed to help introduce this analysis into risk identification exercises - in both design and operations.

Human error, distributed problem solving, assumptions

INTRODUCTION

When accidents are categorised simply as being 'human error' the natural inference is that something is wrong with the person making the error. They had the wrong intention, a wrong belief, poor memory, a lapse of attention and so on. The environment they were working in plays an important role in inducing, suppressing or helping detect such errors, but error is still located in the mind of the individual. Some of the most influential models of error reflect this position. They associate different types of error with different levels of human performance, for example[1,2].

But this kind of model has been criticised on the basis that it does not help distinguish manifestation and cause: erroneous actions by people are not necessarily caused by erroneous effects within people[3]. Moreover, fairly recent work has emphasised the extent to which people's action and problem solving is determined by the situations they are in[4,5], and the way in which this problem solving is distributed[6]. People get parts of the solutions they need from watching other people's behaviour, custom and practice, organisational procedures and so on. The tools they use often embody other people's knowledge, the routines they follow are often the result of other people's trial and error, and the codes they observe typically have arisen from someone else's activity. This means that when people's problem solving fails (and when, for example, they have an accident) it is important not simply to look at what has gone on in someone's mind. It is important to look at how the distribution of their problem solving has failed. The distribution does not have to be planned, deliberately: it can be emergent[7]. But, whether it is planned or emergent, it can be characteristics of this distribution to which failure is most obviously traceable.

A study has been funded by the HSE to determine how the distribution of problem solving contributes to accident causation. This has involved analysing a set of about 60 accidents and incidents offshore, a substantial proportion of which were hydrocarbons leaks. In almost all cases, people's distribution of problem solving had been entrained in their activity. That is, no-one had explicitly planned this distribution: the way people habitually conducted their tasks meant that distribution was a natural part of what they did.

For this distribution to work out various kinds of implied assumption had to be satisfied - and it was failures in these assumptions that led to failure.

The purpose of the study described in this paper was to determine what these kinds of assumption were, and to find ways of making this knowledge useful to people - the people making the assumptions (operators and maintenance staff largely), people wanting to know how vulnerable a system could be to such assumptions (such as designers), and people needing to make an external assessment or audit of an installation (such as inspectors).

THE UNDERLYING STUDY

METHOD

The study used a secondary data source – a set of 59 offshore incident reports obtained from a database held by the Health and Safety Executive and a small set of 4 reports on major offshore accidents. There are obvious limitations in using such reports as data, but it is important to stress that we were not attempting to diagnose specific failures definitively. The aim of the project was to help people reason about failure in the future, and if our inferences *could* be true they ought to be useful. The reports were analysed in several stages:

- Constructing a simple causal network to represent each narrative report.
- Identifying wherever possible how problem solving had been distributed in the task implicated in the incident and how this distribution had failed.
- Identifying in each case the assumption about the world that was implied in the distribution; (for example, reusing existing solutions typically implies an assumption about the solutions' applicability in different conditions).
- Developing a taxonomy of the implied assumptions.
- Deriving a set of guidelines to express how the design of systems could reduce the potential harm arising from these assumptions.

RESULTS

The taxonomy of implied assumptions had three very general categories at the top level:

- The reasonable system assumption (the assumption that the design and disposition of the technical system was reasonable, in the operator's terms).
- The appropriate organisation assumption (the assumption that organisational systems were complete and fitting for the task in hand).
- The knowledgeable person assumption (the assumption that people involved in the task had proper knowledge of whatever was needed).

Each of these categories then consisted of a set of lower level categories. As our purpose was ultimately to influence system designers, it is sub-categories to the first category that are presented here, in Table 1.

Table 1. Sub-categories of the reasonable system assumption

Assumption	Case example
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.
No signals means all is well	Radio used for communication during crane operations. Channel had unknowingly failed and three '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.
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.
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.
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.
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.
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.
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.
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.
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.

Design guidelines were derived from these assumptions simply by asking how, in general terms, the design *might* help avoid either the assumption arising in the first place or any harm that could arise from it. There turned out to be four main categories of guideline:

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

Table 2 shows the guidelines within the last category of 'Presumption'. Most are obvious in the sense that they do not reveal new principles - but the evident failure to follow such guidelines in some cases suggests they are not systematically used.

Table 2. Sub-categories of the presumption guideline

Presumption guidelines
Determine whether use of equipment requires observers who are then vulnerable to hazards
Expect alarms to be inadvertently left inhibited
Anticipate side-effects of misdirected or excessive force in construction or use
Do not assume that the fitting of foolproofing devices is free of error
Predict the practices that operators will learn to minimise effort and maximise production
Predict that people will omit tasks when they are many and uniform
Expect operators to expect the design to be intuitive and easy to orient
Do not assume users' roles let them perform the functions you delegate to them
Expect the operator to be unpracticed in using emergency controls
Expect operators to expect designs that are reasonable in their eyes and do not expect them to test this
Determine how tasks broken down and allocated to different people could leave the system in a hazardous state
Determine how redundancy in protective actions or devices could be undermined
Anticipate that operators will believe precautionary tests to be comprehensive
Predict how an object would provide hand-holds, steps, and wedges and thereby create a hazard
Predict how the current availability of components or services or could influence operators to use the wrong ones
Do not expect people to notice objects or connections they do not expect to be there
Anticipate that a user might not test a configuration before using it in an environment that will punish incorrect configuration
Determine how interruption of dismantling sequences could be harmful

Assume that precautionary sub-tasks will be forgotten if they are not physically necessary to proceed
Anticipate that people will search for appropriate actions by trial and error

THE PROMPTING TOOL

PRINCIPLES

A tool has been developed in an attempt to make the results of the study accessible and useful to people. The basic principles are these:

- People need help testing their assumptions. The study pointed to a wide variety of assumptions that can imperil hazardous installations, and it was suggested that such assumptions tend to be implied rather than explicit. Both characteristics suggest that without a structure of some kind people will naturally find it hard to test these assumptions.
- The categories of flawed assumption that came out of the analysis should be provided as prompts to people in order to help them examine assumptions. This may be a question of helping operating or maintenance staff test their own assumptions before they are about to engage in some risky process. Or it may be a question of helping designers anticipate the assumptions that operators or maintenance staff could make, in an effort to make the design resilient to such assumptions.
- Links between the categories and accounts of the underlying accidents should be retained in the tool so that users can easily consult examples - examples both of the kind of assumption in question and of how the assumption causes the system to fail. The categories help distil the essential elements from the accident and incident reports, and provide the general type of issue that has to be examined. But they are likely to be too abstract in some cases to be applicable on their own. One of the arguments for providing this kind of tool is that it helps people think about kinds of failure of which they have not had direct experience: it helps people expand their knowledge base, as it were, many times over. But for the kinds of failure of which someone has not had direct experience it is probably necessary to demonstrate how it can happen at the level of specific, concrete events.

STRUCTURE

The structure of the tool is shown in Figure 1. There is an underlying database of accident cases, consisting of reports, causal analyses and very brief, thumbnail sketches of these reports. Overlaid on this database is the category structure of flawed assumptions. The detailed kinds of flawed assumption are then mapped to design guidelines, and there is a category structure above these.

There are two main ways of getting access to this structure. The first, which is the most obvious way for operating staff wanting to test their assumptions, is to work through the categories of flawed assumption. The screenshot in Figure 2 illustrates this. Thus the user selects one of the three coarse categories (for example the ‘reasonable system assumption’) and then selects, in turn, the lower level categories below this (for example the assumption

that ‘ambiguous things do not matter’). As each category is selected, the tool displays the thumbnail sketches of the accidents in which this assumption type is implicated.

To give the user a little more structure, a report form can then be brought up which prompts the user to write down, for the assumption type currently selected, any problems they can envisage, any actions required, the responsibility for these actions, their criticality and a deadline. This form, in common with most of the package, can be changed by the user to suit particular needs. From this point (where a particular type of flawed assumption has been selected) the user can also look in more detail at one of the underlying accident reports, or look at the design guidelines that were derived to address the assumption type in question.

An alternative way of organising these assumption types has been provided, based on their frequency. The basic rationale for this kind of organisation is that when people have insufficient time to examine their assumptions against all the 30 or so categories they need to have some rational way of rationing their attention. The tool therefore orders the assumption types according to how many cases implicate each of them. Plainly, the problem with doing this is that frequency is only a partial indication of importance. The fact that the assumption that ‘systems are left complete’ is the most common does not mean that it led to the greatest losses. And even if the tool had ordered the assumptions according to greatest loss in the past there is no guarantee that this would be reflected in the losses that could occur in the future.

The second main way of using the tool is to consult the design guidelines. Figure 3 shows a screenshot of the window for doing this. The user selects one of the four coarse kinds of design guideline, and can then work through the more detailed guidelines grouped under these. For any one of the guidelines, the user can bring up a list of implementation suggestions, and a form that helps record any thoughts the user has about changes that need to be made to a design. It is important to emphasise that at this level the user is the expert. The tool provides general kinds of failure and general, desirable constraints on a design to reflect these kinds of failure. Converting these into particular features of a particular design is something the tool cannot do. The implementation suggestions are therefore simply possibilities, derived from the particular accidents that the study happened to tackle. They do not provide a checklist for a designer. The user can also, from the design guidelines window, bring up another window explaining the flawed assumption that is linked to the selected guideline.

USES

A distinction was made in the study between uses that were essentially on-line and those that were off-line. On-line uses were to do with direct support of normal activity, while on-line uses were essentially for training - helping people with potential rather than actual problems.

The on-line possibilities are as follows.

- Hazard identification. Hazard identification exercises, however structured, rely on people’s knowledge of what can go wrong. There is plainly no guarantee that any particular group of people’s experience is broad enough to know of all relevant kinds

of failure, so supporting the process is important. A tool like the one developed here synthesises the knowledge available in quite a large number of events. The process is essentially one of inspecting all the categories of assumption and asking 'are we vulnerable to making this assumption?' (in the case of operating staff), or asking 'in what way is the design vulnerable to people making this assumption?' (in the case of design staff).

- Operator participation. Most design organisations involve operators in the design process. But operators sometimes have their own hobby-horses, and they do not necessarily know whether their particular experience and opinion is typical or unusual. A package like this could help operators participate in the design process - giving them an aide memoire and a knowledge of problems that they might personally not have encountered, but that they want to bring to the attention of designers. It seems reasonable to say that involving operators in design does not avoid the need to use a tool of the kind we have developed, and equally that using a tool of this kind does not avoid the need to involve operators in design.
- Aide memoire. Designers often need aide memoires to help them think about all possible problems - especially if these problems are associated with people misusing or misinterpreting the designed system. Each of the assumption types presented by the tool is fairly obvious, once one is told about it. The difficulty is in remembering all the items without any support or structure.

The off-line uses of the tool that we could envisage are as follows:

- Induction training. Empirical knowledge of failure is likely to be most lacking in the least experienced people, so the body of knowledge contained in the tool is likely to be most useful for an organisation's new starters. The essential structure of the tool is a set of cases and a set of generalised concepts (the categories) and one can use the links in both directions. Training could consist of looking at the cases and finding applicable concepts, or working with the concepts and finding applicable cases.
- Toolbox talks. Providing new material for short, general background briefings can sometimes be difficult. The tool could be used to support toolbox talks - for example providing one case and one kind of assumption for each session. People are thereby seeing many cases, over a period, and the message about being attentive to hazards is continually being reinforced.
- Observing others. The tool also provides checklists for observing others. The core of some safety schemes is observing others, and spotting hazardous situations and acts. Sometimes it helps to have some structure to do this - to mark someone's behaviour down as being of a particular kind.

FEEDBACK

The tool has not yet been used to any great extent so it is unclear how well it would support the kind of use that has been suggested. Favourable feedback, so far, has included the following:

- The general principle of linking risk identification strongly to accidents or incidents that have actually occurred is an important one in helping people see that implausible failures do occur. It is very easy for people to be dismissive of the possibility that failures occurring to other people could occur to themselves, and easy to dismiss misled operators as being foolish.
- The tool does seem to address a genuine gap in most risk identification processes, which sometimes lack systematic ways of dealing with the human element. Even highly structured human reliability analysis tools typically say little about the causal mechanisms by which the distribution of problem solving fails.
- The idea of helping people examine assumptions, especially implied assumptions, seems to be reasonable one. Inspecting such assumptions helps people address what have been called the 'pathogens' that contribute to failure[8]. Particular trigger events of an accident - such as an action slip - are unpredictable and hard to forestall, but the 'pathogens' that reside in the system are potentially more open to correction or containment.

Unfavourable feedback has included the following:

- The broad approach of influencing people to test their assumptions suits only cultures in which there is an assumption of empowerment. Fatalistic cultures (whether corporate or national) would probably not be influenced by this kind of work. Fatalistic operators would see the system, rather than their own thinking, as the primary determinant of hazardous outcomes, and fatalistic designers would see operators and operating organisations as the prime determinant.
- The terminology and general tenor of the language used in the package would suit only organisations in which there is a receptiveness to new concepts, particularly new concepts concerning failure and hazard.
- The package does not tell designers what to do at a sufficiently concrete level. Although design guidelines have been derived from the assumptions, and although these have been translated into specific, example implementations, the tool still cannot be applied without thinking hard about how operators or maintenance people will use the design.

ADDENDUM

There is a final part to the tool that is still being developed. A problem for managers at a certain level is less with the question as to whether a particular system design takes account of flawed assumptions, and more with the question of whether a design organisation, as a generality, tends to be good at designing systems that are robust to such assumptions. This problem has been tackled by outlining a capability maturity model. The model suggests that an organisation can sit at one of five broad levels of capability maturity, where the levels are defined qualitatively in terms of well the organisation takes account of the assumptions made by people affected by the designs. It is also possible to define transition rules that show what has to be done to move from one level to the next.

The principle is not in fact to ask which level the organisation in question sits at, since parts of the organisation are likely to be better developed than others. Instead, the user is

asked, for each of the five levels, which parts of the organisation have the associated degree of capability, and provide an argument or evidence for it. As with the other parts of the tool the user still has to do thinking: the tool simply provides a structure within which to do it.

CONCLUSIONS

WHY BOTHER?

The difficulty in advocating the use of this kind of tool is that it adds to, rather than replaces, existing processes. The tool's structure does not map directly, for example, to HAZOPs although it could support the conduct of HAZOPs by helping test whether it is likely a human action could lead to some condition and deviation. In fact one could argue that the basic starting point of one's reasoning differs between HAZOPs and the assumptions tool. So, whereas in HAZOP one is starting with deviations and working back to the possibility they could arise in human error, with the assumptions tool one is starting with possible error and working forward to consequences for the system. There is an argument that following both processes would give more assurance than either one alone, but there is obviously a cost in terms of the effort needed.

Nonetheless we felt that there were a number of reasons at least to consider using the assumptions tool. They overlap somewhat, and will not apply in all cases. But they may be compelling on occasion. And even if they are not especially compelling in isolation, it may be that the mass of these reasons together makes it worthwhile using the tool:

- Protecting oneself. The people most imperilled by the assumptions covered in this package were the people making the assumptions themselves. One of the most important reasons for using the tool is to save oneself unnecessary harm.
- Protecting others. Many of the assumptions made by one person imperilled another. So if someone wants to protect colleagues they need to make some effort to inspect their own assumptions.
- Accepting accountability. People are accountable morally and legally for hazards they contribute to, and using the tool should help people meet their responsibilities and demonstrate to others that they are meeting these responsibilities.
- The knowledge arises from harm, or the potential for harm. Some of the cases the tool draws on involved deaths and serious injuries. To ignore the knowledge, unless it is based on flawed inferences, would be to ignore the losses that generated it.
- Observing others. A critical element in the protection of hazardous installations is people's capacity to observe each other. The tool gives people a structure for observing others and testing what they are assuming.
- Common sense limits. We learn how to behave as part of our upbringing and general experience of life. Part of this involves learning what assumptions we can reasonably make about the world. But the world of a hazardous installation is different from the world of home, school and street where we learn our basic behaviour. We cannot therefore depend on our instinctive behaviour in such an installation. We have to know what it is about our assumptions about the world, learned in everyday life, that is misleading or plain wrong.

LIMITATIONS

There are limitations in this work at several levels. First, the basic principle that historical analysis helps one reason about future failures is not watertight. There seem to be some constants in human and organisational behaviour, but new technologies and new laws at the very least change the relative importance of different behaviours and beliefs. Second, the principle of sampling the past in the expectation that the sample will be representative has some obvious limitations. There is no guarantee that our list of assumption types is exhaustive, and in a qualitative study it is very hard to gauge whether one has a good enough sample. Third, the principle of using accident and incident reports as a source of data has some basic problems. The recall of the people involved in accidents can be partial, the process of investigation can be constrained, and the freedom to publish conclusions that are controversial or commercial damaging can be very limited. On the other hand, there is a strong expectation that we learn from accidents and that we can demonstrate this learning. And the complexity of failures is such that simulations and experiments have their own limitations as sources of data.

ACKNOWLEDGEMENTS

Many thanks are due to Bob Miles of the Health and Safety Executive for support and encouragement during this work. The study was funded by the HSE under contract D3916.

REFERENCES

1. Reason, J., 1990, *Human Error*, Cambridge UK: Cambridge University Press.
2. 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.
3. Hollnagel E., 1993, *Human Reliability Analysis: Context and Control*, London: Academic Press.
4. Lave, J., 1988, *Cognition in Practice; Mind, Mathematics and Culture in Everyday Life*, Cambridge, UK: Cambridge University Press.
5. Suchman, L.A., 1987, *Plans and Situated Actions*, Cambridge UK: Cambridge University Press.
6. Hutchins, E., 1995, *Cognition in the Wild*, Cambridge MA: The MIT Press.
7. Clegg, C., 1994, Psychology and information technology: the study of cognition in organizations, *British Journal of Psychology*, *85*: 449–477.
8. Reason, *op. cit.*

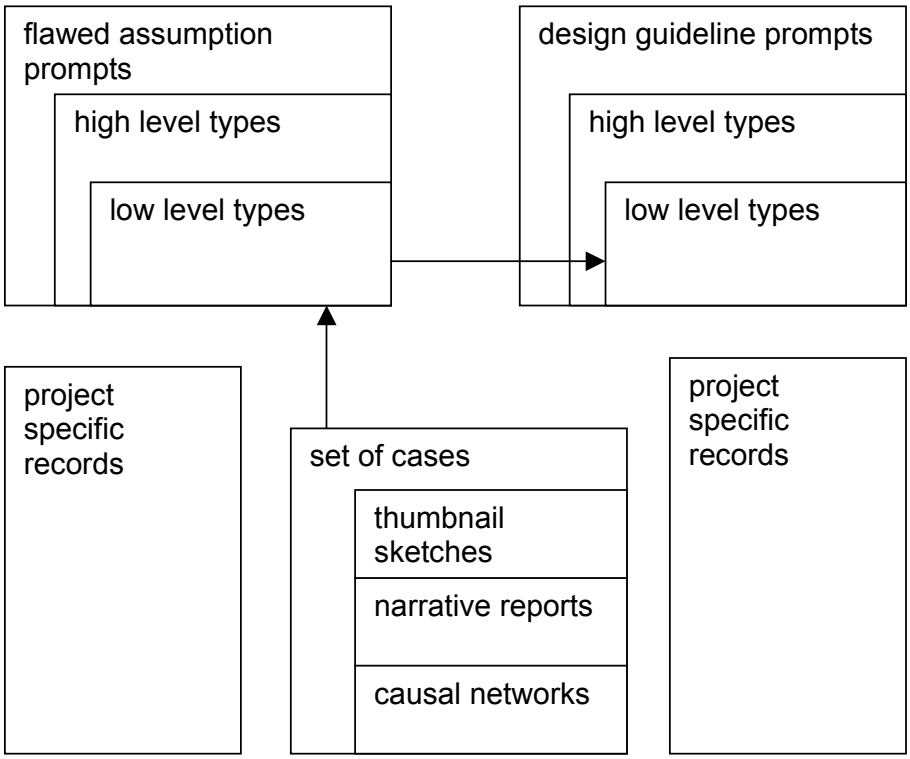

Figure 1. Basic structure of the prompting tool

Figure 2. Screenshot on categories of flawed assumption

Figure 3. Screenshot on categories of design guideline