Error and distributed cognition in design

J.S. Busby, School of Industrial and Manufacturing Science, Cranfield University, Cranfield, Bedford, MK43 0AL, UK

In this study of the errors made in the design of complex process plant, 75 cases were analysed as failures of distributed cognition. Error arose in the interactions of different designers, of designers and design tools, designers and the formal organisation, and designers and the environment beyond the immediate organisation. The analysis attempted to describe the failures of each type of interaction, but there were some patterns of error that emerged in all types. For example, many errors arose from absent cues or inattention to cues. Many also arose from problems with norms (in the form of codes, standards and procedures)—including, ironically, norms that had been implemented in the light of earlier errors. One of the main practical implications of these findings is that designers could benefit from thinking about their tasks in terms of distributed cognition since this suggests several heuristics which, among practitioners, seem to be frequently neglected.

Keywords: case studies, collaborative design, design cognition, design management, distributed design

The purpose of this paper is to describe the nature of errors in organised, full scale design processes—and in particular to explain these errors in terms of distributed cognition.

Error has a basic importance in most human tasks. It is a necessary element in learning a task and adapting it to changing needs, but is also one of the main influences that limits performance in the task. This importance increases at the organisational level, where error can be especially costly yet especially informative. In design tasks error can be still more important because of the possibility that the consequences can be very widespread. Error is also revealing: it often helps us understand the nature of a task that has become habitual, automated or just taken for granted. When an organism is well adapted to its environment satisfactory performance says

ELSEVIER
www.elsevier.com/locate/destud
0142-694X/01 $ - see front matter *Design Studies* **22** (2001) 233–254
PII: S0142-694X(00)00028-4

more about the environment than its internal nature[1]. It is when performance fails that this internal nature becomes evident.

The study reported here took place in a firm designing complex industrial equipment, primarily certain types of chemical process plant and power generating plant. These are complicated products, and the design process for each distinct project takes a long time (typically one to two years) and involves a large number of people (typically a hundred design engineers in the firm itself, with perhaps the same number again in its immediate sub-contractors). The process is vulnerable to error not just because there is a lot of it but also because it is highly interconnected, like the designed product. Error is also very costly: it typically affects many people, can incur costs of hundreds of millions of pounds, and can in extreme cases lead to fatalities.

1 Background

The background to the work comes from a number of areas. Of most direct relevance is prior research on error in engineering design, which has tackled a number of issues:

- *A priori* categorisations of possible errors derived from some systematic framework[2]. Categorisations of this kind, developed in advance of empirical work, help guide the search for actual error modes during the design process and at the same time help illuminate the nature of the design process.
- Case histories of engineering design failures—typically large, catastrophic ones in small numbers[3,4]. These are interesting in their own right, but in very small numbers there is obviously a question about how far they can be generalised. The intention behind this study was to analyse at least several tens of errors as a single dataset.
- Empirical observations of error-making in the course of studies on the design process, design profession, design cognition or design organisations[5]. Because these findings of error tend to be by-products of studies of something else, they might not represent systematic accounts of error-making itself.
- Simulation models that might help predict the occurrence of error in design tasks[6]. These are, probably as a matter of necessity, limited in their scope and make some quite demanding assumptions about how an individual error-maker interacts with the world at large.
- Analytical approaches to failure in engineering design[7]. These concentrate on the physical mechanisms of failure rather than failures on the part of the designer, but are instructive about the limitations of designers' dealings with physical failure—for example the assumptions that accompany safety margins.

1 **Simon, H A** *The sciences of the artificial* 2nd edition MIT Press, Cambridge MA (1981)
2 **McMahon, C A, Cooke, J A and Coleman, P A** 'Classification of errors in design' *Proc. International Conference on Engineering Design ICED97*, Tampere, 19–21 August (1997) pp 119–124
3 **Hales, C** 'Five fatal designs' *Proc. Int. Conference on. Engineering Design ICED95* Prague, 22–24 August (1995) pp 662–667
4 **Petroski, H** 'Paconius and the pedestal for Apollo: a case study of error in conceptual design' *Research in Engineering Design* Vol 3 (1991) pp 123–128
5 **Dwarakanath, S and Wallace, K M** 'Decision-making in engineering design: observations from design experiments' *Journal of Engineering Design* Vol 6 (1995) pp 191–206
6 **Stewart, M G** 'Simulation of human error in reinforced concrete design' *Research in Engineering Design* Vol 4 (1992) pp 51–60
7 **Carter, A D S** *Mechanical reliability and design* Macmillan, Basingstoke (1997)

Perhaps the most obvious need in order to advance our understanding of error in engineering design is for systematic surveys of relatively large numbers of errors in design organisations. Without this it is hard to know how much variety there is in the origins of error in design, how comprehensive the proposed taxonomies in fact are, and the fidelity of models that attempt to represent and predict error.

A second field in which design error has been studied is that of construction. There are several themes, mostly similar to those in mechanical engineering design:

- Case studies of failure involving error[8,9].
- Empirically-derived categorisations of errors[10,11].
- Analytical methods for representing and predicting error[12].
- Computer based tools for managing error[13].
- General insights into the nature of error and its role in the development of engineering knowledge[14–16].

Civil engineering design is associated more with single constructions than mass-produced artefacts, which, in combination with the scale of most civil projects, arguably makes design error more problematic. There are fewer opportunities for early detection of error, and possibly more opportunities for inconsistent interfaces. This is also true of process plant design, in which error making has been of particular concern because of the safety hazards it presents[17].

In systems engineering, human error has been the subject of considerable attention. Again, the subject has been tackled in a number of ways with no obviously dominant approach. There are, for example:

- general models that capture modes of human reasoning and ways they go wrong[18,19];
- analyses of error making in specific domains such as fault-finding[20,21];
- studies of the errors people make as a result of their difficulties in dealing with their environment[22];
- observations on the strategies people should adopt to limit error making[23];
- quantitative models for the prediction of human error[24].

It is probably true to say that the systems community's concern has been almost exclusively with the errors that take place at the sites of accidents and failures: with people in control rooms, for example. Where it has tackled design it has tackled design as a contributor to error making on the part of users of artefacts—rather than investigating error in the design

8 Dias, W P S and Blockley, D I 'Galileo's confirmation of a false hypothesis: a paradigm of logical error in design (Discussion)' *Civil Engineering Systems* Vol 11 (1994) pp 75–77

9 Petroski, H *Design paradigms. Case histories of error and judgment in engineering* Cambridge University Press, Cambridge (1994)

10 Blockley, D I 'Analysis of structural failures' *Proceedings Institution of Civil Engineers Part 1* Vol 62 (1977) pp 51–74

11 Brown, C B and Yin, X 'Errors in structural engineering' *Journal of Structural Engineering* Vol 114 (1988) pp 2575–2593

12 Lind, N C 'Optimal allocation to control gross errors in structures' in **Grigoriu, M** (ed.) *Risk, structural engineering and human error* University of Waterloo Press, Waterloo, Canada (1984) pp 192–204

13 Blockley, D I 'An AI tool in the control of structural safety' in **A S Nowak,** (ed.) *Modelling human error in structural design and construction* American Society of Civil Engineers, New York (1986) pp 99–215

14 Blockley, D I and Henderson, J R 'Structural failures and the growth of engineering knowledge' *Proceedings Institution of Civil Engineers Part 1* Vol 68 (1980) pp 719–728

15 Brown, C B 'Incomplete design paradigms' in **A S Nowak,** (ed.) *Modelling human error in structural design and construction* American Society of Civil Engineers, New York (1986) pp 8–12

16 Allen, D E 'Structural failures due to human error: what research to do?' in **M Grigoriu,** (ed.) *Risk, structural engineering and human error* University of Waterloo Press, Waterloo, Canada (1984) pp 127–138

17 Kletz, T *An engineer's view of human error* Institution of Chemical Engineers, Rugby (1991)

18 Rasmussen, J 'Skills, rules, and knowledge: signals, signs, and symbols, and other distinctions in human performance models' *IEEE Transactions on Systems, Man, and Cybernetics* Vol 13 (1983) pp 257–266

19 Reason, J *Human error* Cambridge University Press, Cambridge (1990)

20 Johnson, W B and Rouse, W B 'Analysis and classification

process itself. Similar comments apply to work on error in ergonomics, where again past research has addressed a number of themes:

- methods for analysing artefacts that are tolerant of human error[25];
- understanding what predisposes people to error, or 'shaping factors'[26];
- theoretical insights into the role of error in learning and organising[27,28];
- empirical studies of error making in specific tasks, like driving[29].

Much of this work does not explicitly exclude design tasks from its remit, but equally it does not say much about design error in particular, and what it is about design that makes it more or less significant as a process that can fail.

The upshot is that our knowledge of error in design organisations remains unsystematic. At the same time it is hard to see how to generalise research on error in real-time tasks to design, and it is hard to see how to generalise research on error on the part of individuals to error on the part of organisations. Without this systematic knowledge of error in organised design processes, specifically, it is hard to know how much of the failure that occurs in the design process is an acceptable part of planned experimentation, and how much is avoidable, hazardous and damaging to organisational survival.

2. Method

2.1 Research design

A pilot study was conducted in four design organisations producing industrial equipment of various kinds. This study involved semi-structured questionnaires, asking designers to describe particular characteristics of specific cases of error, and the observation of post-project review meetings. The results are not given here but can be found elsewhere[30]. Methods that rely on self-reports, despite their problems, seem to be common in research on human error, where it is usually difficult to undertake direct observation—not least because errors are typically infrequent, and people tend not to make the most serious errors when under observation. Following this the main study was conducted in a single design organisation, and involved unstructured interviews in which interviewees were asked to describe specific cases of error at length.

2.2 Data collection

Some 22 interviewees were each given an interview of between 50 minutes and one and a quarter hours. The sample of interviewees was drawn up by managers in the organisation and is shown in Table 1. In the interviews, all interviewees were briefed on the purpose of the study and asked to list

of human errors in troubleshooting live aircraft power plants' *IEEE Transactions on Systems, Man, and Cybernetics* Vol 12 (1982) pp 389–392

21 Duncan, K D 'Reflections on fault diagnostic expertise' in **J Rasmussen, K Duncan, and J Leplat,** (eds) *New technology and human error* J Wiley, Chichester (1987) pp 261–269

22 Dorner, D 'On the difficulties people have in dealing with complexity' in **J Rasmussen, K Duncan, and J Leplat,** (eds) *New technology and human error* J Wiley, Chichester (1987) pp 15–22

23 Moray, N 'The role of attention in the detection of errors and the diagnosis of failures in man–machine systems' in **J Rasmussen and W B Rouse,** (eds) *Human detection and diagnosis of system failures* Plenum Press, New York (1980) pp 185–198

24 Gibson, H and Megaw, E 'Quantifying human error for predictive safety assessments' *Loss Prevention Bulletin* Vol 146 (1999) pp 13–18

25 Baber, C and Stanton, N A 'Task analysis for error identification—a methodology for designing error-tolerant consumer products' *Ergonomics* Vol 37 (1994) pp 1923–1941

26 Park, K S and Jung, K T 'Considering performance shaping factors in situation-specific human error probabilities' *International Journal of Industrial Ergonomics* Vol 18 (1996) pp 325–331

27 Brehmer, B 'Variable errors set a limit to adaptation' *Ergonomics* Vol 33 (1990) pp 1231–1239

28 Rasmussen, J 'The role of error in organizing behaviour' *Ergonomics* Vol 33 (1990) pp 1185–1200

29 Fuller, R 'Learning to make errors: evidence from a driving task simulation' *Ergonomics* Vol 33 (1990) pp 1241–1250

30 Busby, J S 'Human error in the engineering design process' *Proceedings Engineering Design Conference EDC2000* Uxbridge, 27–29 June 2000 (forthcoming)

Table 1 The interviewees

No. (in chronological order)	*Job title*
1	Senior specialist electrical engineer
2	Chief piping engineer
3	Mechanical/project engineer
4	Chief civils and structural engineer
5	Principal piping design engineer
6	Chief mechanical engineer
7	Process engineer
8	Heat transfer engineer
9	Chief electrical engineer
10	Project engineer
11	Project manager
12	Business development manager, formerly environmental engineer
13	Chief instrument engineer
14	IT manager
15	Health and safety manager
16	Senior specialist structural engineer
17	Construction engineer, formerly piping engineer
18	Senior project manager
19	Instrument specialist
20	Project engineering manager
21	Engineering systems administrator
22	Process engineering manager

the design projects they currently and had recently worked on. They were asked to think about the stages each project went through, and to identify errors in their direct experience. Errors were defined as occurrences which were unexpected and which could not be attributed entirely to chance or circumstances. This is similar to Brown and Yin's[11] principle that errors should involve surprise, although this has to be a specific rather than generalised surprise since someone might not be surprised that some failure happens in the course of a design, yet still be surprised at the particular failure. Error was not restricted either to wrongdoing or carelessness, and informants were told that errors having slight consequences were as of just as much interest as those that happened to have had severe consequences.

The interviewees were then asked to elaborate on these errors by describing the events both at the point of error and in the process leading up to it. They were also asked to describe the consequences of the error, but there was no attempt made to quantify outcomes—partly because of the difficulty of doing so reliably, and partly because there is often no meaningful connection between the nature of the error and the outcome. Errors are like gates that allow unfavourable consequences to flow but of themselves do not determine the consequences, so it is a mistake to think that someone should learn primarily from the 'worst' errors of the past.

2.3 Analysis

The analysis essentially involved two steps. In the first, every case was first transcribed verbatim and then represented by a causal network—in an attempt both to differentiate specific causes and to determine their interrelationship. These networks (examples of which are given later) simply provide a partial ordering of factors which appear to be causally related in a particular direction. In the second step, every case was tested against a model of distributed cognition. Distributed cognition is essentially concerned with solving problems by collaboration, where none of the collaborators individually can have a full appreciation of the problem. The 'collaborators' can be tools or artefacts of some kind, as well as human information processors, and activity is dynamically referred to parts of the system in both planned and emergent ways[31]. Typically the participants inter-communicate without being fully aware of the extent to which they need to in order to maintain smooth operation of the system[32]. Hutchins[33], for instance, conducted a frequently-cited study of distributed cognition in marine navigation. There is an obvious connection between the work in distributed cognition and that in the related themes of distributed problem solving[34] and distributed decision making[35].

In this study, the particular approach that was taken was to assume that error arose from failures to meet the needs of distributed cognition—and to enquire what these needs were that had not been met. These unmet needs were then differentiated according to whether:

(1) they involved interactions among several *participants* in the design process;
(2) they involved interactions between participants and *designs*;
(3) they involved interactions between participants and *tools* used during the design process;
(4) they involved interactions between participants and the *organisation*;
(5) they involved interactions between participants and the *environment*.

Participants were taken to be designers in past design projects whose solutions or part solutions contributed to the current design process, as well as designers actively involved in the current process. The cases collected under each of these headings were then further differentiated according to the nature of the interaction that went wrong. Figure 1 shows the rudimentary model underlying this categorisation.

Of the 86 usable cases of error, 75 were categorised according to this

31 **Clegg, C** 'Psychology and information technology: the study of cognition in organizations' *British Journal of Psychology* Vol 85 (1984) pp 449–477

32 **Norman, D A** *Things that make us smart* Addison-Wesley, Reading MA (1993) p 142

33 **Hutchins, E** *Cognition in the wild* MIT Press, Cambridge MA (1995)

34 **Chandrasekaran, B** 'Natural and social system metaphors for distributed problem solving: introduction to the issue' *IEEE Transactions on Systems, Man, and Cybernetics* Vol 11 (1981) pp 1–4

35 **Brehmer, B** 'Distributed decision making: some notes on the literature' in **J Rasmussen, B Brehmer, and J Leplat,** (eds) *Distributed decision making* J Wiley, Chichester (1991) pp 3–14

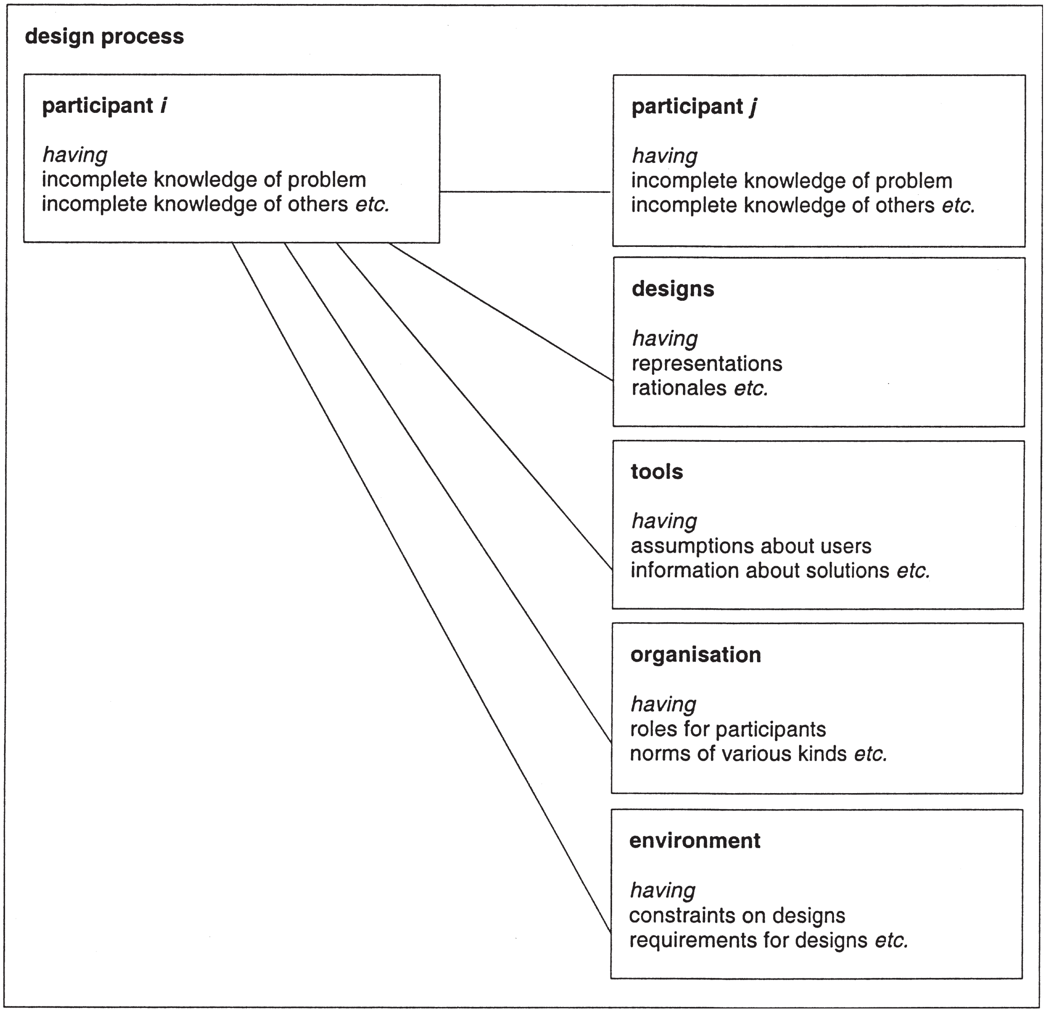

Figure 1 The basic model used for differentiating the cases

scheme. The omitted cases were those that apparently involved no failure in distributed cognition but an error committed purely by an individual designer.

3 *Findings*

Table 2 shows the number of cases brought together under each of the main headings. It is important to reiterate that, because the loss associated with each case was not quantified, these numbers cannot be linked to the seriousness of the failures they are associated with. This section describes the nature of the errors under each heading in turn. Specific cases are described by causal trees, in which the convention is that the left-most

Table 2 Numbers of cases involving the principal types of element

Element	*Cases*
Participants	27
Designs	19
Tools	5
Organisation	19
Environment	5

factor is the outcome, and factors indented once to the right are its immediate causes. If these causes have more fundamental causes, these are shown indented once more to the right and immediately below, and so on. Information that might identify the company has been removed, but where possible the particular artefacts being designed are shown.

3.1 Failures involving interactions among several participants

First come the failures in the distributed cognition among human participants in the design process. Of the analysed cases 27 were of this kind, and Table 3 shows the types of interaction that were involved in the failure, together with the nature of the failures. *Co-option* essentially involved the process of getting the right people to participate in the process; *cueing* involved one participant signalling the need for another to take a particular action; *elicitation* involved a participant actively seeking information about another par-

Table 3 Failures in the interaction of participants

Type of interaction	*Nature of failures*
Co-option of others	Omitting to involve others in design decisions
Cueing of others' actions	Omitting to tell others the assumptions they can make, omitting to state exceptional circumstances, omitting to state normal needs, not knowing there are mutual waiting loops
Elicitation of others' goals	Omitting to elicit other's needs, omitting to elicit other's schedules
Prediction of others' qualities	Not knowing the effect of one's action on another, not knowing the effect of a change on another, not knowing another's circumstances, not knowing another's current state, not knowing another's capabilities, not knowing another's specific experience, omitting to test robustness of one's solution to uncertainties about another's action, not knowing another's strategy
Reconstruction of others' thinking	Not understanding the history of another's problem solving in a replicated design
Scoping of tasks with others	Omitting to define the scope of tasks offered or allocated to others, omitting to find out the scope of imposed tasks
Verification of task by others	Omitting to verify others' decisions, attributing discrepancies when checking other's work to the wrong cause

No. cases = 27

ticipant; *prediction* involved one participant making a prediction about some quality of another participant; *reconstruction* involved one participant working out what another one had done in the past; *scoping* involved participants determining a consistent distribution of tasks between them; and *verification* involved one participant checking the work of another. The types are fairly well-defined but some cases could be allocated to more than one so the taxonomy implied by the table is an approximate one.

Errors in making predictions about qualities of other people in the design process were the most common failings. In the example that follows a mechanical designer changed a type of pump because he had strong commercial incentives to do so. He expected a piping designer, who needed to know various things about the pump, to consult him to find out these details. He did not realise that the piping designer, knowing the original pump designation, looked up the details in a catalogue so did not consult the mechanical designer and did not find out about the change. The error is banal, but shows how simple mis-predictions about a dependent designer's strategy can lead to basic inconsistencies in design:

> piping design to centrifugal pump had to be modified
>
> > *because* mechanical engineer had requisitioned different pump from catalogue design
> >
> > > *because* mechanical engineer attempted to reduce purchase cost by 'buy-down'
> > >
> > > > *because* organisation has incentive to minimise procurement cost
> >
> > <u>*because* mechanical engineer failed to consult piping designer before requisitioning different pump</u>
> >
> > > <u>*because* mechanical engineer expected piping designer to consult him to get detailed design data</u>
> > >
> > > > <u>*because* mechanical engineer failed to realise piping design would work from catalogue design</u>
> >
> > *because* requisitioned pump had different outlet orientation from catalogue design

As Table 3 indicates, there were several 'cueing' errors. Cueing matters in distributed problem solving because one participant often depends on another for partial results, because one participant happens to come across facts that another needs to know, because participants rely on each other when their memories are fallible, and so on. Cueing is particularly important when information emerges fitfully, when the environment provides a lot of interruptions, and when collaborations are temporary. In some cases a designer failed to attend to a cue, whereas in some the cue was absent. In the case that follows there was no cue from one designer to another that pockets should be avoided in lines connected to a vacuum unit. The designer had failed to provide this cue because he or she had made an assumption about the location of the equipment that was not warranted. This assumption had probably been made because, in the absence of any cue, it was normal for the location to be at grade. The example illustrates both the importance of cues, and why they are sometimes missing:

vacuum unit piping requires extensive modification on site
- *because* piping design included pocket
  - *because* piping had to be supported over very long span
  - *because* piping staff had not realised pockets were unacceptable
    - *because* need for vacuum unit unusual
    - *because* equipment designer had not specified unacceptability of pockets
      - *because* equipment designer had assumed vacuum unit located at grade
      - *because* unit located above grade

Generally, Table 3 demonstrates the richness of the knowledge participants have to have of one another in a distributed design process. To solve design problems jointly you need to know who the other designers should be, what they can do, what is an exception to them (not necessarily to you), what they need, what they might be waiting for (which could depend on something you are waiting for, from them), what they will do when you act in a certain way, how capable they are, what state they are in, what they do not know they do not know but you do, the history of the thinking that lies behind their design, how they attribute discrepancies when they verify your results, and so on. Deficiencies in knowledge of all these kinds lead to error. Distributed cognition in design therefore requires much more than a logical task decomposition and the coordination of results, and the implication of errors like these is that participants in the distributed cognition are sometimes insensitive to the richness of what they need to know about their collaborators.

Sometimes the ‘collaborators’ were past designers who provided partial solutions that could be used by a current designer. But the current designer still needs knowledge of the past designer, not just his or her design. One of the most obvious needs is to know, if one is about to improve on someone else’s past design, what blind alleys they have already explored if one is not to try, and fail, at doing the same things again. This is a theme that has not been tackled in prior work on distributed cognition in tasks other than design. For example, Hutchins[33] talks about cultures passing down partial solutions from one generation to the next to aid their cognitive tasks—but does not talk about the value of passing down failed searches for solutions.

3.2 *Failures involving interactions with designs*

These are failings in which participants err in their dealings with a design, typically because they misread the intentions behind someone else's design, or because they do not make plain their intentions in a design they give to others. Table 4 shows the types of interaction with a design that led to error.

As Table 4 indicates, many of the failures had to do with cues. In the example that follows, one designer had became used to taking another designer's specification of boundary conditions in his or her design as a reminder to incorporate a device that would detect operation of the equipment at its limits. When the specification was omitted from the other person's design (for unknown reasons) the cue was absent, the device was forgotten and the equipment failed in service:

auxiliary steam generating boiler failed
- *because* carryover of caustic substance into boiler seams
  - *because* steam flow at higher than expected velocity
    - *because* operating pressure lower than expected
      - *because* attempted to draw greater mass of steam than expected
        - *because* boiler design overlooked steam requirement of de-aerator
      - <u>*because* no failsafe device to shut down operation</u>
        - <u>*because* boiler manufacturer specified no acceptable pressure limit</u>
        - <u>*because* designer used limit specification as cue to incorporate failsafe device</u>

The absence of cues was also problematic when people did not expect them. For example, one designer frequently supplied dependent designers

Table 4 Failures in interactions of participants with designs

Type of interaction	*Nature of failures*
Cueing of action	Not being told what is the unstable element in a 'typical' design, not noticing an unusual need in an otherwise familiar specification, not associating an exception with the need to re-examine assumptions, another's design failing to provide a cue that is depended on to include an associated device
Legitimation of use	Omitting to specify an unacceptable but predictable misuse of a design feature
Prediction of treatment	Failing to take account of idiosyncrasy in a reused design, failing to fully understand the rationale of another's design, wrongly believing a hazardous item in trace quantities to be insignificant
Knowledge of state	Not knowing that the design is inconsistent with its replicas

No. cases = 19

with 'typical' rather than definitive designs when there was insufficient information to make the designs definitive. These 'typical' designs failed to provide cues that told the dependent designers which elements of the design were tentative and likely to change, and the dependent designers generally proceeded as though the design were definitive.

As Table 4 also indicates, there were failures in designers predicting how they should treat designs. For example, a designer forgot to include a particular kind of feature (air dry-up connections on a process plant) because he had reused an earlier design which happened to have been developed on a project on which there was a convention *not* to show air dry-up connections. This convention was unusual. It could be argued in hindsight that the earlier designer should have indicated this idiosyncrasy in some way, because it was predictable that others would reuse the design. But the earlier designer might have believed that the idiosyncrasy—the absence of air dry-up connections—was indication enough:

> start-up design review revealed design omitted air dry-up connections
> > *because* air dry-up connections had been removed from specific design that was being reused
> > > *because* design being reused came from project where practice not to show dry-up connections

The real problem with these cases is not so much that designers have problems interacting with designs, but that designs are an imperfect medium for designers interacting with others (including other designers). Designs may be complete, internally consistent, and completely objective, but as the table indicates this is not enough. It is not enough to reveal the intentions of the designer, in many cases, and it is not enough in a world that is characterised by delay and tentativeness.

3.3 Failures involving interactions with tools

Since Hutchins'[33] work on distributed cognition in navigation, and probably earlier, it has been profitable to see artefacts of various kinds as participants in distributed cognition. There was a relatively small number of errors involving interactions with tools, shown in Table 5.

Again there were problems with cues. One arose when designers forgot to compress graphics files, found tool performance becoming degraded, and assumed the fault lay with the hardware. Another arose when a change administration system provided too few cues to help an occasional user do what he or she intended to do. There were also problems with feedback. For example, in one case a designer repeatedly copied the same elements into an invisible layer thinking the copying operation had failed on each

Table 5 Failures in interactions of participants with tools

Type of interaction	*Nature of failures*
Constraint on inconsistent action	No constraint on inappropriate adaptation of modifiable tool
Cueing of required action	No prompting of discretionary house-keeping operation, no guidance on means–end relations for occasional users
Feedback on earlier action	No feedback on consequence of user's earlier action for later operations, no feedback on consequences of simplifying representation for later operations

No. cases = 5

occasion. The tool provided no indication that the operation had succeeded when the outcome was invisible.

These tool-using errors seem to be characterised by two elements. The first is something about the tool that is ill-matched to the information needs of the user. The second is attribution bias on the part of the user: that is, attributing failures to some property of the tool rather than the user's own prior actions. This bias obviously gets in the way of learning how to use the tool, so this kind of error-making is not as productive as it should be in the learning process.

3.4 Failures involving interactions with the organisation

There are various ways in which the formal organisation contributes to the problem solving of individual designers, perhaps most obviously in providing conventions, codes and standards which constrain the solutions the designer can use, or providing partial design solutions themselves. Table 6 shows the errors that involved interactions of the designer with the organisation.

Some of the failed interactions are similar to those that have already been described under *Failings among participants*, such as failures to co-opt appropriate participants, and failures in the verification of one person's results by another. In this case, the failing is with the organisational mechanisms that are meant to meet these needs, rather than with individuals. For example, the failure involving co-option arose from the absence of a representative of a particular discipline in design reviews generally—not just in a single instance. Obviously, attributing failures of this kind either to the organisation or to individuals is problematic, because the boundary between an organisation's and individual's responsibilities is a matter of taste. But generally, where a failure is systematic, there is an argument that it should be an organisational concern.

Table 6 Failures in interactions of participants with the organisation

Type of interaction	*Nature of failures*
Allocation of work	Allocated enlarged task packages without realising effect on schedules for information exchange
Conventions for activity	Standards were inconsistent with each other, obsolescent standards remained extant, arbitrariness of standards led people to mis-predict them, variation of standards between organisations meant people applying those they were used to applied the wrong ones
Co-option of disciplines	Absence of participants with appropriate disciplinary knowledge in formal routine
Cueing of action	Absence of prompts to maintain consistency of replicated data, latency in mechanisms alerting people to changing state of information, absence of reminders to follow new drafting convention, absence of reminders to revisit tentative assumptions
Goals imposed	Imposition of intermediate deadlines requiring use of incomplete information, failure to state goal of robustness to uncertainty and change to new joiners
Verification structures	Absence of planned synchronisation of activity precludes clash checking, checking activity fails to test external consistency of design, absence of checking mechanisms for transcription

No. cases = 19

As indicated in the table, several errors involved conventions. Some involved conventions which the designers did not know of, partly because they had not been told them, and partly because they were by nature arbitrary and therefore unpredictable. For example, one involved the convention of modelling unresolved branches in pipe designs by adding stubs:

project manager believed that piping isometric design was more advanced than it was
- *because* piping system report showed unmodelled branch lines as modelled
  - *because* piping designer had modelled stubs in place of unresolved branch lines
    - *because* piping designer wanted to extract isometrics without error reports
    - *because* piping designer followed practice at previous employer
      - *because* piping designer did not know local practice for dealing with unresolved branch line

Other errors in conventions involved inconsistencies and obsolescence in the conventions themselves. For example, a formal design standard called for labels that turned out to be illegible and had to be changed, on the constructed plant, at considerable cost. On previous projects, the standard had not been used because labelling had been the subject of ad hoc negotiations. It had therefore not been tested. One suspects, also, that because it was a formal standard the designer who resorted to it did not inspect the assumptions it was based on. Generally speaking, conventions of one kind or another were important in the design processes studied here, and in many ways made it more robust. But, at the same time, they introduced new possibilities for error.

Again there were errors in cueing. Often these occurred in connection with change, where cues were needed to tell designers that some constraint or requirement which affected them had been modified. Errors obviously occur when such cues are absent, typically because it has not been predicted that a certain kind of change affects a certain designer. Errors also occur if there is delay in providing these cues—for example where change administration systems queue their workload. And errors occur where a cue is needed to test whether a change is required in the first place—typically when a designer makes a provisional assumption about constraints which he or she needs to revisit at some later time, but forgets. This forgetting is fairly predictable, given that information emerges over finite periods, designers frequently switch between tasks, and the environment provides frequent interruptions.

Generally, the kinds of error shown in Table 6 reflect commonplace problems with organisations, and similar errors probably occur in technical tasks other than engineering design. It is debatable whether they could therefore be labelled 'design errors'. But these errors do indicate that design activity in organisations is not just a question of designing in groups. It is not the fact that design is collective that is the problem, but that design takes place in a setting where there are formal structures and routines that have arisen from historical events. Tasks are decomposed and allocated in certain ways, conventions come into being for certain reasons, individuals acquire goals that arise from certain roles they are required to adopt, and so on. And people rely on organisational mechanisms to provide information processing. They rely on the organisation's information processing for cues, for solution methods, and even for partial solutions themselves. Errors of various kinds arise when this organisational information processing is flawed.

3.5 *Failures involving interactions with the environment*

There was a small number of errors involving interactions with the environment beyond the design organisation, shown in Table 7.

One failure arose because a member of a collaborative project adopted goals that were inconsistent with this collaboration. The organisation in question had become habituated to an environment where there were fewer inter-dependencies among different design disciplines, and it attempted to optimise its sub-design:

Table 7 Failures in interactions of participants with the environment

Type of interaction	*Nature of failures*
Cueing of action	Absence of cues indicating that organisation did not know what it did not know about client requirements, absence of cues that collaborator's design was changing state, external interruptions leading to lapses when switched tasks
Feedback on solution	Absence of collaborator feedback on proposed solution
Goal for collaboration	Failure to state goal of robustness to uncertainty and change to new collaborators

No. cases = 5

rework required to structural steelwork design
- *because* contractor attempted to optimise steelwork design for material content
  - *because* contractor believed in client's interest to optimise steelwork for material content
    - *because* contractor had little experience of working in collaborative setting
  - *because* required duty loadings stated were conservative
    - *because* relevant loading information such as pipe thrusts remained undetermined
      - *because* steelwork design had to start with provisional information
        - *because* schedule adherence required early start to steelwork design

Another error arose when one organisation proceeded with a sub-design that was incompatible with another's, but received no feedback saying so. And, once more, errors arose from absent or inappropriate cues. For example, an error occurred where an external intervention (by a client organisation) caused a designer to switch tasks in the middle of making a set of modifications. He or she forgot to resume the modifications and so the design was unknowingly left in an inconsistent state. 'Lapses'—where an intention is not followed by the required action—are obviously common and predictable human errors, especially where the environment provides interruptions. Design activity is vulnerable to lapses because it involves relatively complex operations which, if not taken to completion, lead to a design being in an inconsistent state.

4. *Conclusion*

4.1 *Variety in the modes of failure*

There is considerable variety in the types of failure described in the *Findings*, which points to the importance of studying errors in reasonably large numbers—but plainly requires a great deal of effort when each case has to be analysed in detail. But this variety is perhaps unsurprising, because in an established organisation the errors someone is most likely to observe are those whose lack of a systematic connection with earlier errors has made it less likely that the organisation will have learned to avoid them. Obviously it is quite probable that there will be systematic limitations in the learning process that mean errors cannot all be attributed to chance or circumstances, but it seems reasonable to expect that the residue of errors, after learning, will contain a large non-systematic component.

4.2 Cueing problems

This said, there were two common themes that linked many of the cases: cues and norms. Cues are occurrences which participants in the distributed cognition use to determine when to act and how to act. Norms are rules of some sort, spoken or unspoken, that help make the participants' subtasks consistent with each other.

Cueing is important in collective design that extends over long periods and long distances. It is not just that one participant's tasks have to be co-ordinated with others', and with the needs of the external world, but that participants have to interpret what is needed of them. They have to know, for instance, when they are required to tell someone else what loading they have decided someone else's structural subdesign will have to withstand. They have to know that when a client uses a feedstock with a high carbon dioxide content they will need to specify particular types of material and this will affect loadings on the supporting structure. It is not surprising that cues are problematic: the process is complex so it is a difficult organisational problem to ensure cues are there when they are needed. And the question of how designers will attend to and interpret cues is as much a function of their own experience as the nature of the cues. Some of the errors in this study, for instance, occurred where designers did not realise that some characteristic of the design problem was a cue to adopt a particular kind of solution, and they did not realise this because their prior experience happened not to have involved a situation where this cue was relevant.

One important source of cues is the design itself, or at least the sub-designs that are exchanged between the different design disciplines. In one error, a sub-design did not show limiting conditions so a designer used to using limiting conditions as a cue to specify detection devices forgot to do so. Designs were also an important source of cues when a designer wanted to adopt or adapt the solutions of a predecessor. But, as became clear in several errors, the cueing function of a design is compromised by the difficulty designers sometimes have in interpreting its rationale.

Cues also cause problems when the design is passed to production, installation, operating and maintenance staff. Construction staff, for example, sometimes take signs of familiarity (like a design which outwardly looks like a previous one) as cues to avoid wasting time in detailed study of the design, proceeding on the assumption that design is just like the previous one when in some cases it is not. Designers, in such cases, seem to conclude that construction staff are in error and do not change their design practices. But construction often operates under very heavy time pressure (sometimes during maintenance shutdowns) and it is predictable that construction staff will positively seek cues that tell them to avoid unnecessary activity. One could well argue that designers need to be attentive to the cues the design provides, and where these cues are inadequate or misleading to arrange better ones.

However, the cases where the obvious providers of cues failed to provide them were instructive about why cues could be missing. People who could have drawn designers' attention to exceptional constraints sometimes did not because they believed the constraints to be trivial. They believed them to be trivial because they were, to *them*. Similarly, people who could have provided data on particular variables failed to do so because *their* way of doing an analysis did not need it—even though, as it turned out, other's analyses did need it. The upshot is that differences in practice and perspective among members of the design process limit their ability to provide cues to one another. Moreover, the provision of one cue sometimes relies on receiving other cues, in a kind of nested structure. In one of the examples given in the *Findings*, one designer omitted to cue another to the need for a particular feature because the first assumed circumstances did not require the feature. He or she made this assumption because he or she, in turn, received no cue that circumstances did in fact require it. It is important that participants in the design process are aware of the cues they need to provide to others, and of the fact that their own experience may not help them predict what these cues are.

In the context of real-time operations, in contrast to design activity, the exchange of cues is again very important but again often not consciously registered. Norman, for instance, points to the role of central throttle levers in aircraft cockpits for ensuring that both pilot and co-pilot know their current state—and suggests that task analyses typically fail to document this social communication[36]. Hutchins' studies of how a team of people develops a collective ad hoc problem solving algorithm during a navigation emergency shows how this can take place without a deliberate attempt to design such an algorithm[37]. It seems that the use of cues is both an adaptive

36 Norman, D A 'Design principles for cognitive artefacts' *Research in Engineering Design* Vol 4 (1992) pp 43–50
37 Hutchins, E 'Organizing work by adaptation' *Organization Science* Vol 2 (1991) pp 14–39

strength, in allowing this kind of unplanned problem solving to emerge, but also a weakness in that it is associated with certain kinds of error.

4.3 Problems with norms

Hutchins[33] writes about 'culture' as passing partial solutions from one generation to the next. In this study, it was a specific aspect of culture—the norms which designers shared—which provided much of this solution passing, and played a part in many errors. Norms were often seen in fact as things that had been implemented in the aftermath of past errors, in order to avoid their recurrence. They ranged from informal understandings (how, for example, structural and piping design was demarcated), through explicit procedures (for example, how to represent unmodelled piping branches), to codes that specified necessary properties of an artefact (which could be as simple as the height of handrail). These norms were especially important when different people's work had to be consistent, but when in isolation they would have found it hard to predict what their colleagues' actions would have been. Whether handrails are 70 or 75 cm high, for instance, matters much less than that all designers choose the same height. If they do not, a vessel designer will design fittings at a height which clashes with the structural designers' design of handrails. Another way of looking at this is that norms are important when participants find it hard to predict the effect of their actions on others. If they can follow a set of rules that others know about they do not need to put themselves in the minds of others to work out how they could confound their designs.

Unfortunately, norms in various guises are also the source of errors. Ironically, one type arises for the same reason that norms are introduced. When they are arbitrary (and function by making people's work consistent, not because there is only one best way of working) they are most essential. Yet when they are arbitrary they are also most vulnerable because people who are unaware of them are least likely to predict them correctly. One error, for example, arose when a designer did not realise there was a convention in the organisation about how to represent unmodelled pipe branches. He followed the convention he had adopted with a previous employer. An obvious remedy is to brief people better, constrain them more, and make norms even more prescriptive. But given the number of norms a designer follows, and the difficulty even for experienced people to remember them all, this does not look promising. Some of the errors found in this study also showed how readily people will violate conventions if they believe they are unnecessary and they are working under pressure. And it is quite easy to become conditioned to believe that norms are unnecessary if one happens to have had experiences only of events where the norms were not needed. The only robust approach to reducing errors in applying norms is to help designers understand the extent to which they use norms and help them develop a sensitivity to their misapplication.

Ultimately, the question of norms is closely related to that of cues, because norms help participants proceed when cues are missing or unreliable. If a mechanical designer repeatedly forgets to give nozzle loading requirements to a piping designer one remedy is to adopt a convention that nozzle loadings will always be the same. Equally, cues are needed to avoid errors with norms—like cues that suggest to the designer a potentially outdated norm needs to have its assumptions inspected. Both cues and norms become basic elements in organised, collective design, and both play a central role in errors.

4.4 Assumptions and not knowing

It is easy to believe that assumptions are always in the designers' heads. In one error in this study, the designer assumed the water table was so low that concrete casting of sumps would be straightforward and a special coating would not be problematic. This tuned out not to be the case. But, as far as we know, the designer never said to himself or anyone else 'the water table is low in this instance'. He either followed his previous practice or the previous practice of others who had done similar designs. The 'assumption' was contained in the fact that previous designs happened to have been done for locations where the water table was low, so no-one in the design process discovered, in the process of making an error, that the solution principle would not work when the water table was higher. If they had, then cues—like the fact that the location for the plant was coastal—would indicate that the solution might fail. Thus assumptions like this are not part of the designer's mental state before the design activity, a fact which plays an important role in theories of situated action[38]. Therefore, when designing, the process of testing whether one's assumptions are valid is not so much about asking what one knows and might be fallacious, but what one does not know because one's experience is a biased sample of all possible experiences. This is not perhaps a test that most designers seem to apply in any obvious way.

Some of the errors in this study involved not knowing what needed to be known. One designer did not ask another whether a client had an idiosyncratic requirement, even though the other designer undoubtedly did know, because he did not know that clients had idiosyncratic requirements (of this particular type). He did not know what he did not know. It is very much like the problem of assumptions 'being in the world', because it was a chance property of the designer's experience that this experience had not been completely representative. An organisation can compensate for this

38 Suchman, L A *Plans and situated actions* Cambridge University Press, Cambridge UK (1987) kind of problem because it can provide cues and norms. It can cue individuals to find out bits of information that they might not have planned to obtain, and it can establish norms such as consulting people with broad experience at important points in the design process (like design reviews). However such remedies give no guarantee that people who do not know what they do not know will find out, so errors are always likely to arise because assumptions exist in the world and people's experience can lead them to be insensitive to this.

4.5 How appropriate is the distributed cognition model?

People diagnose errors as 'communication problems' very often, but this is so commonplace as to be a rather empty categorisation, and it suggests too limited a model of how designers jointly do designs. The failures that there are arise not so much because X does not communicate Z to Y, but that Y interprets Z in the light of some prior experience which X does not know about so fails to make allowances for, and Y does not realise X does not know this because Y thinks both that his experiences are representative and that X's are, and so on. The model of distributed cognition helps to reveal this kind of nested structure, its dependence on the inferences of the participants, and the root of these inferences in their historical experience of earlier episodes of similar kinds.

The model also has some potential usefulness to design managers because failings in each of the different types of interaction described in the *Findings* point to various heuristics for running design organisations. For example, failures which involve the interaction of current and past participants suggest organisational omissions in the preservation and dissemination of knowledge. Specific failures suggest such remedies as having roving design consultants, lessons-learned meetings that draw out the systemic causes of failure, project programmes that allocate resources to doing research into what problems have been encountered with similar projects in the past, and so on. Similarly, failures which involve interactions of design tools with designers point to remedies like incorporating systems support staff as part of design project teams, and focusing support and training efforts on specific individuals whom other designers regularly consult. One advantage of the distributed cognition model is that it is systemic in nature, drawing attention to the way articles of knowledge and bits of processing need to come together from multiple sources—rather than isolating particular elements (like the individual designer) and suggesting that remedies lie in improving those isolated elements (like changing designers' attitudes).

The main limitation in the model is that it does not help very much to explain errors that arise from the individual designer. So understanding error could never be purely an exercise in applying the distributed cognition model. There is also the question of whether it is really distributed cognition when the multiple participants are not self-consciously part of same problem-solving effort. For example, the designers and the users of an artefact are not obviously part of the same process, so if designers have misconceptions about the needs of users, and users have misconceptions about the intentions of designers, is this really a failure of distributed cognition? A somewhat similar view of the collaborative design process has been proposed, in terms of shared memory[39], and possibly this would compensate for the some of the shortcomings of distributed cognition. But the reality seems to be that neither concept is all-embracing, and possibly that there is no one concept that could fully explain all there is to know about the design process.

4.6 Limitations

Finally, there are several elements of this work that set limits to its wider applicability. It relied on individuals' reports of historical events which are open to biases to do with recall and self-presentation. There was no independent verification of the inferences made from informants' accounts, so there is no information on inter-rater reliability. And no attempt was made to find second reports on the errors described by each individual so there is no information on whether alternative perspectives would have led to different accounts. All these factors suggest that particular cases and their interpretation are not especially dependable. The suggestion is also that, because of the possibility of bias, quantitative inferences would be suspect. But demonstrating the existence (if not the extent) of error-producing characteristics of cues and norms in the design process is less vulnerable to reliability limitations. And this work can still serve as an exploratory basis on which to base confirmatory work. The challenge is to find research methods which gain reliability without sacrificing validity when the subject in question is one that traditionally makes research problematic.

Acknowledgments

Many thanks are due to the staff of Stone and Webster Engineering Ltd, particular to Mike Welsh and Janet Nicholson for making arrangements. Thanks are also due to the anonymous referees for comments on an earlier manuscript. The work was funded by Stone and Webster and the EPSRC.

39 **Konda, S, Monarch, I, Sargent, P and Subrahmanian, E** 'Shared memory in design: a unifying theme for research and practice' *Research in Engineering Design* Vol 4 (1992) pp 23–42