Reliability Engineering and System Safety **41** (1993) 25–47

Context and human reliability analysis

Ed Dougherty*

Science Applications International Corporation (SAIC), 19353 US Highway 19 North Clearwater, Florida 34624, USA

(Received 3 July 1992; accepted 8 November 1992)

Unfinished business related to human reliability assessment includes the identification and specification of cognitive (diagnostic and decision making) error potential and context. This relates to the so-called NRC commission error issue and is a recognized omission from the recent efforts in IPEs.

By reviewing notable instances of cognitive errors or near misses, by carefully characterizing the environment and situations in which such errors will arise and by borrowing on a scattering of partial techniques, a systematic approach to cognitive context can be developed. This paper takes a stab at gathering the various pieces and suggesting how such a method might proceed.

INTRODUCTION

As Probabilistic Risk Assessment (PRA) enters the post-IPE era for the nuclear power community there seems to be a diminishing future for further Human Reliability Analysis (HRA) developments, while the recent activities associated with performing Individual Plant Examinations (IPEs) have demonstrated a significant technical need. There obviously are several arenas in which lessons learned from the nuclear concern for assessing hazards may have a portability, e.g. the aerospace, the space, or the process industries. But whatever the nuclear HRA future, the purpose of this paper is to redirect. The redirection is intended to apply to some unfinished business but may inadvertently amount to a new direction entirely.

For bounding purposes, the analytical setting of this paper is a nuclear power plant In which the human milieu might be characterized as follows: there is a *crew* of operators plus considerable supporting personnel with relatively clearly stated (although potentially conflicting) *goals* operating in a highly *proceduralized,* i.e. emergency operating procedure (EOP), environment, in which 'events' are detected almost solely by means of a complex technical system of *instrumentation* and alarms.

THE CONTEXT OF CONTEXT

All human action is performed within a specific context, i.e. conditions that are situational (such as cues from plant instrumentation) or environmental (such as the time available in which to perform an action). However, the reliability of an action is not necessarily an obvious function of this context or may only be partially dependent on it.

Swain, of course, writes more than most on the various influences, the reliability context of human performance (see Chapter 3 of Ref. 1). His categories of influences,[1] or as he terms them, performance shaping factors (PSFs), include situational characteristics (relating both to workplace and environment), task instructions (e.g. procedures), task characteristics (e.g. complexity), organismic (i.e. human) factors, and the 'stressors' that impinge on the psyche and soma as a result of these influences. Of course, even an HRA method whose analysis resorts solely to the management of such influences, e.g. Success Likelihood Index Methodology (SLIM),[2] its variant, the Failure Likelihood Index Methodology (FLIM),[3] or a late incarnation,[4] incorporates few of these potential influences, presumably because they are not really that influential.

However, taking Swain's lead, it is fair to assume that the reliability of human performance, particularly the kind that is more knowledge-based (following Rasmussen),[5] or cognitive (following Hollnagel),[6] is a function of several *dimensions.* Figure 1 reinvokes the ancient protagonist to cognitivism, the SOR (stimulus/organism/response) paradigm, to provide an otherwise useless partition of these dimensions.

It is clear that what one is after, whether called goals or values or purposes, heavily influences one's responses to events, even the perception of events (the stimulus itself). Then events, particularly their pace as they unfold, 'dictate' in the above setting the

*Present address: SAIC, 655 Metro Place South, Suite 745, Dublin, OH 43017 USA.

Reliability Engineering and System Safety 0951-8320/93/$06.00

Fig. 1. SOR as a starting point rather than a sore point.

kind and quality of decisions and actions in response to them. Within us (organism) we hold innumerable beliefs and attitudes, as well as emotions and other affective if not cognitive furniture that temper, even to contradiction, our goals and perceptions of the evolution of events. Lastly, we respond, at least in the specified analytical setting, according to procedure (as much as feasible), coordinating the response in terms called crews.

Hence, SOR is not the arch-villain that cognitivists would have it but merely the obvious elevated to an icon. The concern crucial to HRA is not whether SOR is correct but rather how to handle the immense richness of the 'O' in SOR that makes human versus machine performance so interesting. This is why it is important to delimit the setting as, for example, presented above. It is not everyday life we seek to explain, which is more varied and hence less predictable from the response point of view but less hazardous (except maybe for transportation) from the stimulus side.

THE NEED TO MODEL

HRA modeling approaches now appear to fall into four categories (see Table 1): procedural, temporal, influential, and contextual. Nuclear power plant risk assessments have used the first two methods extensively and the third occasionally. So-called contextual HRA (a phrase of Hollnagel)[6] is a newcomer to the scene and is, as yet, not associated with a quantification method, which is both a requisite and the bane of human reliability analysis in IPE.

The contextual approach insists that the human reliability analyst be allowed something to do, i.e. neither the stark holism of the time reliability correlation (TRC) temporal approach nor the quantitative holism of the influential approach is satisfying, even were they sufficient. (Note that the holism that may accompany the time reliability correlation variant on temporal approaches[7] is often tempered by the influential or procedural approaches.[8] Note that what is meant by 'holism' in this sense is the idea that all human performance can be reflected in one (or a few) 'lumped' parameters, e.g. available time or a success likelihood index.) But the contextual approach's reductionism is in the direction of *breadth,* i.e. involving factors or influences, along with or maybe as substitute for *depth,* i.e. a hierarchy of reliability units such as a reduction into subtasks.

Notice that the influential and contextual approaches may find themselves indistinguishable at the quantification stage because of the paucity of actual data. However, there is much more task and situational analysis associated with the contextual approach (as proffered by Hollnagel)[6] than has been exhibited in the variants on SLIM, for example. This distinction might merely be the product of analyst style and it is easy to foresee that the influential and the contextual may merge into a single approach.

The linear metaphor for human performance modeling arises from an observational, purely behavioral viewpoint. A task, once performed, is 'laid out before you' as a temporal, linear order, i.e. action 1 precedes action 2 and action 203 follows action 202. Clearly, one can argue over the 'optimality' or the normative goodness of the order of the task 'elements,' e.g. whether action 202 must precede action 203. This is the root of procedure development. And it is also probably true that the learning of a task by taking such an optimal approach is made easier or is more 'cost -effective.'

The issue, however, is whether a linear output so learned when applied in an actual setting, i.e. the performance of a task at hand, has anything at all to do with the (human) *reliability* of that performance. The *procedural prototype* (as termed by Hollnagel)[6] assumes that there is a (nearly) one-to-one correspondence between a task's reduction into actions and the reliability of the performance of the actions. Swain's Technique for Human Error Rate Prediction (THERP)[1] is the exemplar of this notion but there are

Table 1. Four types of HRA modeling approaches

Model type	Analysis type	Output metaphor
Procedural	Reductionist (to subtask elements)	Linear (activity)
Temporal	Holistic	Linear (in time only)
Influential	Holistic (at quantification)	Nonlinear
Contextual	Reductionist (but not simply subtasks)	Nonlinear

other less obvious examples, e.g. Hollnagel's example is Rasmussen's step ladder model.[5] However, a notorious shortcoming of the procedural prototype is the fact that there is no objective 'stopping rule' for the reduction process, i.e. there is no definition of a 'reliability unit.'

One solution may be to create more sophisticated linear models[9] in hopes that the additional complexity of the modeling will allow the emergence of a complexity commensurate with that believed to be characteristic of human performance. However, the alternate approach is to give up on complexity.

This leads to the other popular HRA approach, the TRC approach,[8] which has many shortcomings.[10,11] The basic flaw is that the TRC approach involves no apparent analysis (at least in the pristine form in which it is sometimes practiced). Because a human performance scenario quantified using a TRC is not always accompanied by a task analysis or even a PSF analysis, much more difficulty arises in trying to glean *qualitative* insights from the HRA results than with the straightforward THERP. TRC abuse occurs because it is easy and takes on a facade of objectivity; and because of this bareness in modeling, there is typically a need to 'adjust' the quantifier to account for PSFs—or something.

Another glaring weakness in the TRC approach arises from the temporal characteristics that make it a TRC.

(1) *Sometimes there is too much available time.* The safety margin concept underlying the TRC, e.g. the HCR's normalized time,[12] is exponentially sensitive to time. Hence, any right positive time distribution has a tail that forever decreases in order of magnitude. When the safety margin is greater than 3–5 (depending on the distribution type), then a time failure estimate using the TRC is incredibly low. That is why Wreathall[7] truncated his TRC on probability, e.g. declaring that, say, 0·000 01 was a least credible value, and Dougherty and Fragola[8] often truncated on available time, say one or two hours.

(2) *Sometimes there is too little time.* Any TRC has a median expected response time. For example, Swain's was about 4 min.[8] Since performance is chance (50:50) at the median time, any shorter time is pretty much irrelevant for HRA/IPE purposes. This might lead one to declaring a safety margin minimum of, say, 2:1. However, for some tasks, e.g. verifying that a safety system actuates immediately following reactor trip, short available time should not preclude success.

Ad hoc fixes around these distributional 'tail' properties of TRCs include (1) declaring a scenario time-independent and substituting a procedural approach or (2) adjusting the input parameters[8] to reflect different behavior types, hence producing *families* of TRCs. The fix by tweaking parameters is notoriously unstable.[13] A more reasonable approach is to more tightly bound the variance on response, maybe by using a normal distribution rather than a lognormal one or maybe by tracking the cue evolution more accurately and tacking on a time-independent approach that moves in time (if you will). The result of this tactic would resemble a dynamic modeling approach such as the Dynamic Event Tree Analysis Methodology (DETAM).[9]

Another flaw in the TRC approach is that there is no real way to accomodate anticipation. For example, operators are not surprised (one would hope) by the alarm indicating the need to transfer to sump recirculation in a Pressurized Water Reactor (PWR) Loss of Collant Accident (LOCA) scenario since they know that initial cooling water sources are finite and have been trained (and even simulated) to be aware of this function. So a model that 'starts the clock' at a time when the alarm arrives is wrong. But without modification according to the cue pattern (which has been proposed as part of the new cue modeling in the human cognitive reliability model (HCR),[14] starting the clock at the initiator is just as incorrect. These objections are technical; global objections concerning simulator fidelity, non-response versus failure, etc., further mean that the TRC approach is tenuous.

The net result, it seems, is that neither the procedural nor the temporal HRA approaches are in themselves sufficient to handle cognitive (or any other) context. The procedural approach too readily gets lost in its own deetails of never ending subtasks promoting the belief that task logic alone reflects nominal performance and PSFs used to reflect context are adjoined to the modeling where most convenient to the analyst. The TRC approach, on the other hand, is virtually impervious to context, while being refreshingly simple.

The influence-oriented approaches have been ignored in this casting of stones simply because:

(1) they have not been nearly as popular in IPEs to date, and
(2) they will most likely merge with context-oriented approaches when the latter mature.

This is intended less to be a slighting of, for example, SLIM, which has sparse face validity and suffers from a host of technical problems, but more to provide substance for analysts who must contend with selling operators and engineers on the insights obtained. As David Gertman notes, many of the faults associated with SLIM may be attributed to its past implementations rather than the method.

UNSAFE ACTS

However, the author does not intend to bury old HRA methods (although a previous plea for a second generation appears to have been all but unheard;[10] the intention is to talk of unfinished business.

The US Nuclear Regulatory Commission (NRC) has the onus for the nuclear industry of regulating safety. (With the issue of the Office of Safety and Health Administration (OSHA) rule, CFR 1910,[15] the chemical process industry has entered the wonderful world of risk assessment or something like it.) As a result, what may be termed *unsafe acts,* should be a primary leitmotiv for HRA. Unsafe acts are errors or deliberate actions made in a hazardous environment.[16] The fundamental taxonomy of unsafe acts is as in Fig. 2 (adopted from Fig. 7.7 in Ref. 16). Unsafe acts may be intended or unintended. (That is, the action is intended; the consequence usually is not.)

An unintended action may amount to a *slip,* a failure to pay enough attention; or put more neutrally, an attentional scheme is used that is suboptimal. (This phrase is not intended to be overly jargon-laden. Our attention paying in the everyday world is seldom, or at least not always, the province of conscious decision but rather a feature of learned and now virtually autonomous tasks.) *Lapses* are momentary (or the reflection at a moment of permanent) memory problems. Each slip or lapse amounts to a process, or implementation, error; one that holds little cognitive meaning, although all action can be said to be performed under cognitive *control.*[17] Unintended actions either surprise us, i.e. we notice their impact immediately, or they lie dormant, or as Reason would say, latent, amounting to, as it were, a 'pathogen in the system', awaiting a 'trajectory of accident opportunity' such as a weakness in a safety barrier.[17] When this occurs, it is the situation that surprises us—or as Taylor maintains,[18] all accidents are 'truly meaningless events'. As a result, we often do not detect the latent errors at the root of it all.

Unintentional errors, however, do not carry the emotional baggage that mistakes and violations do. They seem more the province of error tolerant design rather than HRA or its applications to training, etc. This does not stop the utilities from claiming to 'counsel' the offender in hopes of eliminating what is a natural and fundamental feature of human performance. The utilities have also taken to putting the phrase 'cognitive error' in their Licensee Event Reports (LERs). Just what they mean by it is anyone's guess. But in the framework of Fig. 2, we will assume that what is referred to as cognitive errors is what is designated as *intended but unsafe acts.*

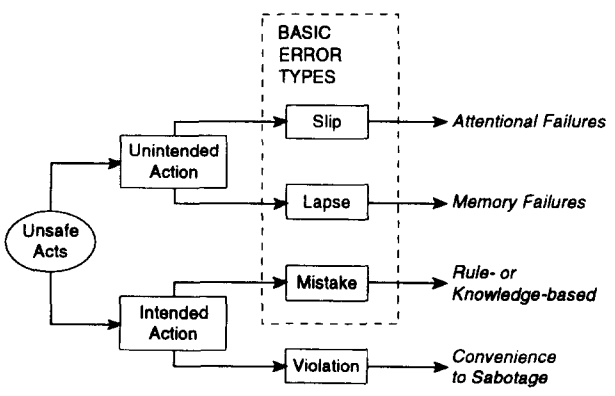

Fig. 2. Reason's taxonomy of unsafe acts.

Most such errors, which are termed *mistakes,* are simply inadequate diagnoses or planning failures. The goals were well perceived; the intention well-formulated; but cognitive performance, which is sometimes woefully fallible, was imperfect. Sometimes failures occur in defiance of having well-formed procedures (rules) and sometimes the failures occur when *ad hoc,* realtime decision making is required but not good enough. However, as Reason's taxonomy allows, some unsafe actions are not errors in any accurate sense. Short of sabotage, which amounts to adopting a goal different from what is generally held, *violations* are deliberate actions that pursue the proper goals but which defy some standard, procedure, or practice. Yet, the actions may be based on beliefs strongly held, although too often with little support, which turn out inappropriate according to the context of the moment.

A vivid event was the procedural violation made by operators at the Davis-Besse nuclear power plant,[19] which is discussed at length later. The action violated procedure (a procedure, by the way, that had never been implemented at any nuclear plant previously nor since). However, the action was based on a belief by the operators concerning the efficacy of the procedural instruction, ameliorated by an ongoing contingency plan that, were it successful, would have made the procedure moot. The contingency in fact was successful, but its occurrence was late relative to the cue for invoking the procedure. Hence, the Davis-Besse operators formally (although maybe not quite deliberately) defied procedure while pursuing a plan that was better according to (some of) their beliefs.

It is interesting to note that the utility subsequently changed the procedure to fit the violation, since it turned out to be the safer way; hence, in hindsight, the action was *correct.* Of course, the regulator shut down the plant for 14 months for the violation, which they judged an unsafe act. Was the action an error? Clearly not; it was even the optimal action. Was it an unsafe act? Yes, the hazard (or a reactor meltdown) existed and the operators took it on themselves to avoid the hazard by following a plan that *they* thought was correct for the specific situation rather than what

others had anticipated in the form of a general procedure. This example will be further detailed later.

As this and many other events attest, *context is everything.*

For clarity's sake, the following definitions are tendered. Notice that they are deliberately bare-bones, since the solution to the problem of HRA context does not seem to be by way of taxonomy.

- *Commission error:* An *action,* rather than an inaction, that produces an effect *not intended* by the actor or that is *inappropriate* considering the situation in hindsight.
- *Cognitive error:* An *action or inaction* that is based on a *decision* (which in turn may be based on a diagnosis, plan, etc., i.e. the error's causation has a 'high level' cognitive content) that produces an effect *not intended* by the actor or that is *inappropriate* considering the situation in hindsight.

Hence, a commission error is potentially but not necessarily a cognitive error (Swain's commission slip is the counterexample) and a cognitive error may not be a commission error—but a cognitive omission, which is not a commission, or a violation, which is not an error. Reason's definition of an unsafe act is thus modified to:

- *Unsafe act:* cognitive error or a (willful) violation.

This definition ignore slips and lapses, the other error modes that Reason calls unsafe acts, because they seem to be more unfortunate acts made in an unsafe or hazardous environment rather than an unsafe act, laden with intent. This means that the issue of wrong unit/wrong train[20] may be relegated to THERP technology, since most such errors amount to technicians or maintenance personnel unintentionally exercising the wrong equipment.

Hence, for us, to commit an unsafe act one must go knowingly, if not willingly, into the fray.

The bottom line of all this is that the NRC's issue over so-called commission errors is probably most fruitfully interpreted as a concern over unsafe acts, as defined above to exclude slips and lapses. Then the distinction of commission/omission and error/violation are semantically moot, while still potent and interesting as technical parameters.

THE EOP CONTEXT

One of the accident mitigation enhancements mandated after the Three Mile Island (TMI) accident was to develop functional symptomatic emergency operating procedures. The motivation of this style of procedure was the fact that a commission error was made during a situation involving multiple failures. The phrase that arose to characterize the new EOPs, 'symptom-based procedures,' is a misnomer, since all procedures are conceived as responses to some symptoms. It was the *type* of symptom that was at issue when the post-TMI requirement appeared.

Partly because of the difference in machines, the Boiling Water Reactor (BWR) and PWR approaches to creating symptomatic EOPs are radically different. However, both are symptomatic: BWRs rely on *parametric* symptoms (instrument readings), e.g. reactor vessel level, and PWRs rely on *functional* symptoms (abstract safety functions), e.g. loss of heat sink (which involves multiple systems and, hence potentially multiple parameters). It is an unresolved issue in HRA as to whether one EOP style might be 'better' than another, or whether their advantages address important but different operational aspects.

It is clear that in such a procedure-dominated environment as a nuclear power plant, errors or violations must be relative to the procedures. Hence, Swain is fundamentally correct to assume that a task, i.e. in this case a procedure analysis, is crucial to an identification of error likely situations. Put alternatively, the *morphology* of cognitive error must lie among the *logic* and *chronologic* of the EOPs. The controversies related to THERP lie among its details rather than at this abstract level.

To illustrate how cognitive error situations might be identified in this heavily proceduralized environment, the EOP system developed by an unnamed utility based on the Combustion Engineering EOP guidelines is described. Figure 3 depicts how this system is to function.

There are seven numbered EOPs. EOP-00 is to be implemented anytime the reactor trips or is judged to require manual trip. This procedure instructs the operators to verify the variety of safety and supporting system equipment that might be needed in *any* offnormal condition. Notice that this would have meant that the operators at TMI would have probably not committed error #1 (see the next section) and maybe would have avoided the misinterpretation error #2. (Of course, it was precisely these errors that led to the new EOPs.) EOP-00 is committed to memory and is part of almost all simulator exercises in training.

EOPs-01 through -07 are *event-oriented* procedures. If a single event has caused the reactor trip and it can be clearly diagnosed, then these procedures would allow the operators to optimally respond to the specific event. (This is why Westinghouse sometimes calls their event procedures Optimal Response Procedures.) The design basis accident, a loss of coolant accident (LOCA), would be treated with EOP-03, for example.

EOP-20 is the *Functional Recovery Procedure,*

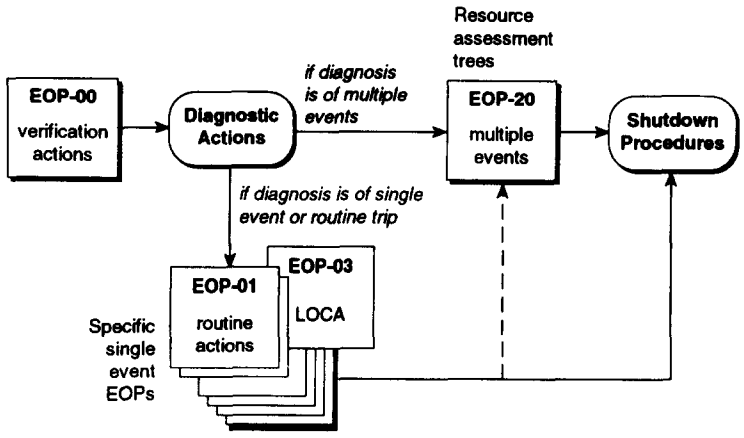

Fig. 3. One variety (of four) of symptom-based EOPs.

which is the hallmark of this functional symptomatic approach. If there are multiple events ongoing or if the operators are not certain of their diagnosis of the specific event, then EOP-20 takes precedence and is not left until the plant can be brought to a safe, stable condition.

EOP-20 includes logic decision trees called Resource Assessment Trees to assist the operators in managing a complicated set of contingencies in a designated priority. The priority concerns six critical safety functions (CSFs) that are to be maintained:

- reactivity control,
- maintenance of vital auxiliaries (ac, dc, instrument air, component cooling),
- Reactor Collant System (RCS) inventory control,
- RCS pressure control,
- RCS and core heat removal, and
- containment integrity.

Each of these CSFs has at least two alternative success paths described in the EOP, any of which would assure the maintenance of the CSF. All CSFs must either already be maintained or one success path must be implemented according to the priority of the CSF and the subsequent priority of the success paths. The purpose of this EOP is clearly to combat fixating on an incorrect diagnosis, by reducing the crew's diagnostic role to that of symptom set pattern matching, as well as to combat the pursuit of lesser important failures, by introducing a goal hierarchy of CSFs and success paths.

A downside to this prioritization tactic is that if there are multiple events but, say, only one 'matters,' and the lesser important event comes first in the priority, then the operators still (are supposed to) attend the less important fault first even if they correctly assess the situation. For example, a recent IPE had a situation in which loss of offsite ac power (LOSP) could accompany any other trip, e.g. a steam generator tube rupture (SGTR), with non-trivial frequency. Since this scenario leads to the diagnosis of multiple events and since ac power is prior to reactor coolant system (RCS) pressure control, the electrical fault would be attended first. The sole recovery option for the particular LOSP scenario of significance included a time-consuming action that would take place outside the control room. An SGTR event is one in which operators should take care of depressurizing the RCS and bottling up the leak as soon as possible. But this latter, more significant aspect of the scenario (in this case) takes a back seat to the LOSP in the EOP-20 priority scheme. It should be noted that the EOPs are 'validated' against a wide range of scenarios, including multiple events. It is not possible to guarantee any single scenario will be handled optimally, just satisfactorily.

Another feature of the EOP system is that the transfer into any of the EOPs beyond the first is accomplished by (literally) a flowchart attached to EOP-00, called *Diagnostic Actions,* to assist in identifying the single event or recognizing the presence of multiple events. (Notice that this distinction is somewhat fuzzy, since, for example, if vital dc power is a single fault, EOP-20 is invoked anyway, whereas were an SGTR to occur followed by faults associated with isolating the leak, the specific EOP for SGTR, EOP-04, is not left for EOP-20.) Notice that *Diagnostic Actions* is a prototype for what is referred to as rule-based behavior.[16] Each EOP also allows for formal rediagnosis and upon specific conditions may transfer control to another EOP; the exception is EOP-20, which once entered is not exited. This transfering feature is indicated by a dashed line on Fig. 3.

One more feature of the EOP system includes the Floating Steps. These are subprocedures that are cued by parameters much like the BWR system, e.g. level in the emergency feedwater storage tank. They 'float' on the back of one or more EOP and are to be implemented whenever their symptom set is observed.

One of the important functions of the Shift Technical Advisor (STA) is to monitor these steps. The Floating Steps take priority over any other steps in a particular EOP, but their action set would typically be implemented in parallel with the ongoing EOP steps.

The crew allocation is another factor in the context of offnormal operation. Figure 4 indicates the minimal crew members in the control room on any shift. Auxiliary operators who are trained to manipulate equipment may be found in the Turbine Generator Building or Auxiliary Building. Technicians and other maintenance personnel may be available particularly during the day shifts. The Technical Support Center (TSC) includes other operators training or operations management staff and would be in operation a half hour or so into the incident under nominal conditions. (Note that this many crew are greater than apparently proposed for advanced reactors.)

The role of this crew is as follows. The reactor operators (ROs), sometimes called board operators or simply licensed operators, are the hands and eyes of the crew. Typically they are the only ones who would maniplate controls on the front control boards and are likely to be the only ones who can read some of the instruments and alarm indicators. The ROs receive instructions from an transmit instrument readings to the Senior Reactor Operator (SRO) who is the 'procedure reader.' The SRO manages the EOPs and other procedures, monitors the evolution of the event, and ensures that the actions are performed, hence acting as the controller of the event. Because of the complexity of the EOPs, however, it is difficult to give much credit for the redundancy provided by the SRO relative to the ROs. (In one simulator, the communication scheme was for the RO to repeat any instruction of the SRO and then perform it. Often, however, the RO anticipated the next step and performed it prior to the SRO's command and then merely aped the instruction callout. Presumably, there was no redundancy in this teaming arrangement but this might only have been a simulator phenomenon.[21])

The Shift Supervisor (SS), who is also an SRO, attends to the NRC notification and event emergency

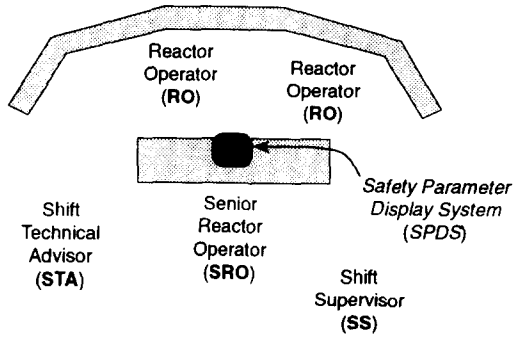

Fig. 4. A typical crew in a horseshoe control panel control room.

classification activities for up to an hour after the initial plant upset. The SS quite literally will be on the phone much of this time and he probably cannot provide the senior advice to the SRO/ROs as might be desirable. The Senior Technical Advisor (STA) is a non-operator (although at some plants they maintain an operating license) whose function is to monitor the CSF maintenance, the Floating Steps, the Safety Parameter Display System (SPDS), and to perform other safety monitoring. He is designated to 'stand back' from the board operators and take in the big picture. The incorporation of the STA has not been smooth at all utilities, so the social and professional status of the STA is always a part of the general context.

As can be inferred from this brief description of the crew and the EOP system, any particular evolution of events will find the operators winding their way through a complex manifold of instructions and cues, using different personnel differently. Human performance in this setting is a concert conducted at the pace defined by the evolution of events under a competence dependent upon the accuracy of the operators' situation assessment and their uncertainty while attending their actions, i.e. their cognitive skills. This effusion of activity is time-embedded but not necessarily simply time-dependent; rule-influenced but not always rule-based; cognitive but in a distributed system of crew and computers. This is the milieu of a nuclear accident and its management.

THE COMMISSION ERROR ISSUE—REVISITING TMI

The event at TMI (see Refs 22–25 for descriptions) led to a popularization of, and a confusion with the term 'commission error.' Table 2 indicates some of the major events at TMI, including 'errors' of the crew.

The accident was a multiple event scenario: a loss of heat sink, i.e. all secondary cooling, which then induced a LOCA when a pilot-operated relief valve (PORV) that opened to relieve the pressure of the lost heat sink failed to reclose as designed. The operators overlooked the immediate signs of the LOCA, such as the increase in the rupture tank level, partly because this information was on a back panel of the control room and partly due to the fixation on the inexplicably rising level in the pressurizer, an unanticipated phenomenon resulting from the loss of heat sink. The apparent fact that this level was rising meant that 'going solid' was imminent, which was of more consequence and importance to operating a nuclear submarine than a power plant. Many hours into the accident, final mitigation of the event was delayed because the thermal-hydraulic conditions of

Table 2. The evolution at TMI

Time	Event	Comment
0	Loss of all feedwater; emergency feedwater (EFW) does not provide flow	A latent error caused the failure of EFW; *TMI#1*—operators do not recognize that EFW is not effective
3 s	PORV opens on high RCS pressure	Due to loss of heat sink for RCS
8 s	Reactor trips	Automatic on high RCS pressure
12 s	PORV fails to close automatically; water from the PORV empties into the drain tank	Piping temperature sensors indicate open valve; tank level indicators not on front boards; *TMI#2*—operators do not diagnose this failure
4 min	Operators isolate (stop flow from) HPI pumps	*TMI#3*—due to a falsely high indication of pressurizer level these steps could have indicated the ongoing LOCA
13 min and 38 min	Operators stop LPI pumps Operators stop sump pump	
1 h 13 min	Operators stop first two Reactor Coolant Pumps (RCPs)	Pumps vibrate due to unknown steam binding
1 h 40 min	Operators stop second two RCPs	As above
2 h 22 min	Operators block PORV	Stops LOCA
11 h	Operators *et al.* stabilize plant	*TMI#4*—hydrogen bubble scare delays actions

the reactor were well outside any knowledge in the industry at the time.

According to the table, there were (at least) four major errors (indicated as *TMI#n*) made during the event evolution (or post-initiator, as PRA terms it). All are cognitive errors and one is a commission error—the infamous one that became the source of the NRC's so-called commission error issue. From the perspective of phenotype, i.e. what happened, the commission error occurred when the TMI operators turned off a safety system (this may have been the sole decision of the crew leader rather than a team decision)—the high pressure injection system (HPI)—and guaranteed that an unusual transient became the only instance in the US of a melted core.

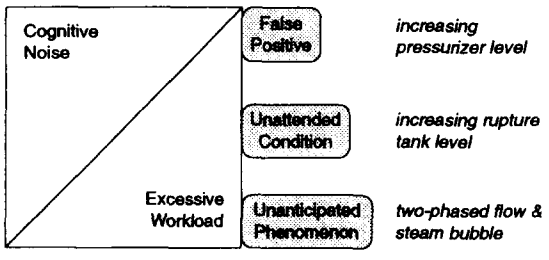

Fig. 5. Part of the cognitive context of the TMI event.

From a perspective of genotype, i.e. why it happened, this commission error was not simple. As Figure 5 tries to indicate, the workload and cognitive noise during the early and mid-range stages of the accident were considerable. For example, seven significant indications arrived in the *28 seconds* following the unknown opening of the pilot-operated relief valve. In short, the operators had too much to do and were overloaded by an alarm system designed to assist them. Hearing a description by one of the TMI board operators of the experience is vivid, leaving no likelihood of volunteers for such a circumstance.[26] To paraphrase Faust, 'I didn't even know whether I was alone or not in the control room for the first ten minutes after all the alarms sounded.'

Hence, the commission 'error' at TMI it seems cannot be fairly attributed to the operators themselves, but, as Table 3 depicts, were induced by problems and limitations of the nuclear industry as a whole. The TMI event was much the analogy to the early failures of the B-727 airplane, which changed airline safety; we in the nuclear industry simply took too much for granted and were enveloped in considerable uncertainty concerning the fundamentals of safe reactor operation.

The nuclear industry, with much consternation, has

Table 3. The multiple layers of 'error' at TMI

All Utilities	The event-based procedures and their commensurate training for operators were inadequate for the occurrence of multiple events along with misleading symptoms
The Industry	There was industry-wide ignorance of the importance of person–system interface; ignorance existed of the impact of two-phased flow and the possibilities related to post-core heat up phenomena

implemented numerous enhancements to the person–system interface in each nuclear plant at a cost of millions of dollars per plant. Yet not one enhancement has been *demonstrated* to be an improvement in safety or human performance. Nor have many of the technical issues been resolved. An indicator of this is that the so-called commission error issue is still unresolved and, although there have been very recent efforts to provide some assessment methodology,[27,28] these are far from final nor are they presently promising.

(The curious common feature between these competing preliminary methods is their lack of consideration of context in a setting that demands the influence of context. Note, however, as Gertman points out, that the use of INTENT assumes that a cognitive task analysis or other supporting effort has priorly determined that decision-based errors are possible. Also, the PSF assessment, as in the case of SLIM-like methods, is *contextful.*)

SIX NOTABLE EVENTS

The nuclear industry has a history of being reticent to analyze events across the industry for HRA and human factors implications (see Refs 29 and 30 for exceptions). Along with TMI, at least six events stand out as being interesting in relation to cognitive and/or commission errors (see Table 4 for a list). TMI occurred in 1979 but two interesting events occurred prior to it, and the other three are post-TMI. The event in 1985 at Davis-Besse is made more interesting because most of the morphology of this event was a replication of TMI's event evolution six years later.

Browns Ferry

On 22 March 1975,[31] a worker ignited a fire in the Unit 1 cable spreading room while using a candle to check for penetration leaks. Ironically the air flow from a leak caused the flame used to find the leak to be drawn into the polyurethane foam sealant. The fire was fought unsuccessfully with CO_2 and dry chemical extinguishers for about 15 min until an evacuation alarm sounded. The operators announced the presence of the fire at that time.

Smoke and inaccessibility hampered the fire fighting. The operators and plant staff decided not to use water, fearing that water might not extinguish the fire and might cause further damage to instruments and controls (I&C) or create an electrical shock hazard. Plant I&C began to degrade 33 min into the fire, and the fire ultimately affected instruments or controls of the residual heat removal, high-pressure and low-pressure injection systems, and radiation monitors. The local fire department was called 30 min into the fire but not used. The I&C problems that began at 33 min left only control rod drive water as a high-pressure water source. At 70 min, the operators began to depressurize to allow use of the condensate booster pumps but the relief valves failed and pressure increased at 5·5 h. At 6·5 h, the shift supervisor approved the use of water to extinguish the fire, which was declared 'out' 45 min later. At 9·5 h the RCS was sufficiently depressurized to allow for long-term cooling through condensate makeup.

Rancho Seco

On 20 March 1978,[32] a plant technician dropped a display light bulb behind an instrument panel at Rancho Seco, shorting out dc power to nonnuclear instrumentation bus Y. This caused the loss in indication for steam generator (SG) level, pressurizer level, and RCS temperature and other equipment. Blind to the secondary side, the operators initiated feed and bleed (FAB) in 1–7 min, using the computer indication of pressurizer level. A false alarm due to the bus loss had tripped main feedwater. Auxiliary feedwater was inhibited by closed inlet valves due to the loss of SG actuation signals. An uncontrolled drift in SG A level 'allowed' the FW valve to receive an open signal. Seventy-five minutes into the accident, power to NNI-Y was restored. RCS pressure and

Table 4. Six notable events

Plant	Error no.	Year	Type	Effect	Mode	Result
Browns Ferry	—	1975	post	omission	mistake	suboptimal
Rancho Seco	—	1978	pre	commission	slip	suboptimal
TMI	#1	1979	post	omission	mistake	failure
	#2			omission	mistake	failure
	#3			commission	mistake	failure
	#4			omission	mistake	suboptimal
Sequoyah	—	1981	pre	commission	mistake	failure
Ginna	—	1982	post	omission	violation	suboptimal
Davis-Besse	#1	1985	post	commission	mistake	failure
	#2			omission	violation	success

temperature were stabilized 30 min later using pressurizer spray and reactor coolant pumps.

Sequoyah

On 11 February 1981, a path from the RCS out of the containment spray header was inadvertently opened by an auxiliary operator (AUO) at Sequoyah Unit 1. This created an LOCA, with 180 000 litres of primary system water and 300 000 litres of refueling water storage tank (RWST) water being sprayed into containment. A major cause of the accident was the need for remote communication. The AUO had been dispatched to open two B loop residual heat removal (RHR) valves and to *verify* that the interconnect valves between the RHR system and containment spray were closed. The AUO arrived at the interconnect valves first and telephoned back to the unit operator (UO), who told the AUO to open the two valves. No mention was made between the two operators as to which of the valves were involved. The AUO *opened* these two valves, which were the ones he was only supposed to verify, and proceeded on the other two valves. When there, the AUO attempted to telephone the UO but this phone was inoperative. The AUO then also opened these valves creating a LOCA path through the spray system. Forty-three minutes later, the AUO along with a second UO pieced together the cause of the LOCA to which the operators had already responded correctly.

Ginna

On 25 January 1982, a steam generator tube ruptured at Ginna. The operators recognized high air ejector radiation alarms and low pressurizer pressure and diagnosed a tube rupture. They isolated the affected SG in 4 min and began cooldown. At 26 min, the operators blocked the SG PORV locally to further isolate the affected SG. At 39·5 min, the operators began to depressurize the RCS through a PORV, which was detected as stuck open, at 41 min. The PORV was blocked a minute later. (The crew reported that they had thought of TMI's PORV problem prior to Ginna's PORV problem.)

In the interim, a steam bubble had formed in the reactor vessel driving the pressurizer level above 100% almost exactly when the procedure step for terminating safety injection (SI) was reached. Apparently,[33] a concern arose in the control room that the steam bubble might change the conditions for terminating SI (*déjà vu* from the TMI accident). However, the crew and personnel in the Technical Support Center (TSC) debated whether to terminate SI per procedure, the shift supervisor's recommendation, or to continue SI to try to collapse the steam bubble, that of the TSC. Debate continued for 17 min from 42 min into the incident. It turned out that the control room crew were correct in hindsight, since it was subsequently shown that the bubble did not present a major risk and the SI termination conditions were appropriate in this case. However, when the TSC convinced the crew to restart SI, the bubble did collapse.

During the debate, an SG B safety valve lifted and the operators had to regulate auxiliary feedwater flow to SG A. At 69 min, the operator overruled the TSC and terminated SI. At 83 min, sump alarms sounded (from a minor leak in the letdown system) and at 99 min the TSC, concerned about the steam bubble, requested and got the restart of SI. At nearly 2 h into the incident, the steam bubble collapsed and SI was stopped again. The system stabilized at 3 h.

Davis-Besse

On 9 June 1985, Davis-Besse experienced a loss of all feedwater flow to its steam generators. As at TMI, auxiliary feedwater (AFW) did work as intended. As part of the routine trip response, an operator went to a back panel to attempt to manually start AFW using the Steam Feedwater Rupture Control System (SFRCS) actuation system. This is a complicated 10-button system; there is no simple AFW start button. Figure 6 shows the actuation controls schematically (adapted from Ref. 19). The operator pushed the two top buttons instead of pushing the

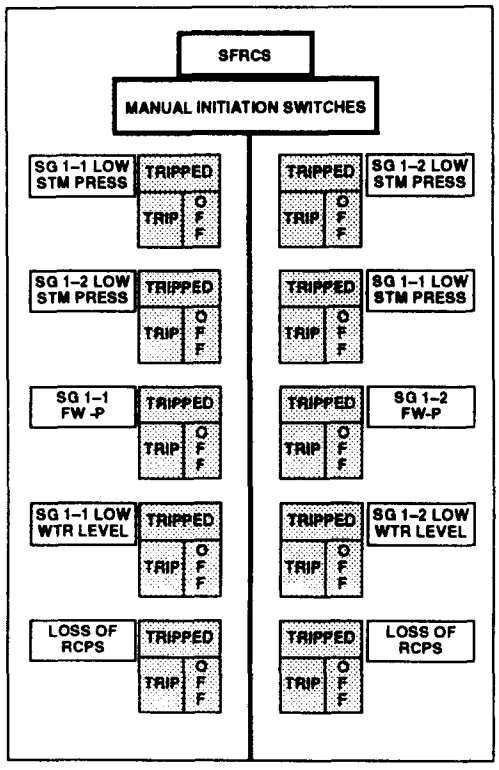

Fig. 6. The auxiliary feedwater initiation control panel at Davis-Besse.

level buttons, the fourth pair from the top. By pushing the buttons for pressure, the SFRCS automatically reacted as if there were a streamline break and isolated the steam generators, i.e. flow to the steam generators remained unavailable. The NRC team assumed that this was 'inadvertent' but noted that the button arrangement contributed to the error. Ironically, the utility had previously advised the NRC that they were planning to change the actuation controller to accommodate better human factors but they had not done so yet.

More significantly, the operators had to recognize the conditions that were present in order to push the correct pair of buttons. Thus, not only was this a commissions error, it might have been a cognitive error as well. (Reason classifies the error as a slip, but it would plausibly seem a *rule-based mistake*).[16] The operator admitted to the NRC team that little training had been provided for this action and that he had never attempted it in the plant or in simulation. Davis-Besse operators had also had to use a generic B & W simulator rather than a plant-specific one at that time and the generic simulator did not have the SFRCS actuation control. It must be noted that the significant effect of this commission error was that of simply not starting AFW, i.e. an omission error effect. The Assistant Shift Supervisor recognized the error and tried to restart AFW within 3 min, which also failed to work. This ended the initial phase of the event.

Having given up on starting AFW from the control room, operators were dispatched to try to start AFW locally (about 9 min into the incident). Without secondary cooling, the steam generators would dry out and stop producing steam. Since all of the AFW pumps were steam driven, there was little likelihood that there would be enough steam left to restart the steam driven pumps. The assistant shift supervisor then decided to use the startup feedwater pump (SUFP), which was electrically driven, to feed water to the steam generators so that the steam produced then could be used to restart the AFW pumps. He took about 5 min to do this, although previous (and subsequent) walkdowns had taken much longer. In the meanwhile, other equipment operators at other locations within the plant opened the valves needed to restart the AFW pumps.

During the course of trying to restore AFW, the steam generators went dry. This was the cue by procedure to use the feed and bleed option, called makeup/high pressure injection (MU/HPI) cooling at Davis-Besse. This is clearly an unwanted option except in the direst conditions, since it will induce a LOCA (the bleed part) and drain primary water into the containment sump. The result will be a long shutdown for cleanup. The shift supervisor, in telephone consultation with the operations superintendent, was 'influenced by a reluctance to release reactor coolant into the containment because of the cleanup and extended shutdown associated with it.'[19] Since the assistant shift supervisor had a viable recovery strategy (the SUFP) and since the core had in no way reached dangerously high temperatures, the shift supervisor delayed the MU/HPI option awaiting the attempt of the assistant shift supervisor. This indeed worked and the utility has since changed the emergency procedures to reflect this strategy. However, the NRC kept Davis-Besse off line for 14 months because of the incident.

It is not insignificant to note that a PORV lifted and did not reclose fully following its third lift, creating a small LOCA similar to that at TMI. The operators routinely used the block valve and closed the path but apparently attributed the resulting depressurization to the use of the pressurizer sprays at about the same time. They did not know that the PORV had stuck open. © Cognitive aspects of the events.

A different synopsis of these six events, more from their cognitive aspects, is made in Table 5. Notice that the initial commission error of the Rancho Seco event, the dropping of the light bulb, was (probably) not a cognitive error. However, the subsequent response of the operators, which was error-less, had considerable cognitive context. This demonstrates the significant reliance in high technology environments on 'artificial' perception, e.g. instrumentation. It also may indicate

Table 5. Cognitive aspects of the six events

Event	Cognitive problem	Category	Recovery
Browns Ferry	Belief state	Error	Down to last resort
Rancho Seco	(Gross) cue underspecification	None	Innovative actions
TMI-1	Unattended cue(s)	Error	Later, after the noise
TMI-2	Unnoticed cue(s); mindset	Error	Never really did
TMI-3	Strong but wrong interpretation	Error	Too late to avoid melt
TMI-4	Unknown phenomena	Error	After much delay
Sequoyah	Unknown cause	None	'Routine'
Ginna	Unanticipated phenomena	Violation	Finally OK
Davis-Besse-1	Unfamiliar action and conditions	Error	Alternate found
Davis-Besse-2	Reluctance in prescribed option	Violation	Ultimately 'correct'

that gross underspecification of an event's cue may lead to increased cognitive tension and thence lead to opportunistic (Hollnagel's term, TBP)[6] but successful behavior. On the other hand, subtle underspecification might not arouse the tension necessary to break a mindset. (Note that all of this synopsis is speculative, i.e. leading to hypotheses, but is hardly a scientific data-driven analysis.) The initial error in the Sequoyah event, going to the interconnect valves first, was also probably a non-cognitive (wrong train) commission error. However, the communication failure between the AUO and the UO was a cognitive error, again a failure in 'remote' perception. Rancho Seco and Sequoyah show that the obvious necessary requirement for proper cognitive performance is reliable information and that obtaining this information has both technical and social dimensions.

The Sequoyah event also demonstrates that even when the *why* of an event is not known, with good procedures the *what* can be accomplished. Moreover, the fact that the 'command and control' of the event was distributed meant that when the control (the AUO) made an error, then the redundancy of the EOP system could influence the command (the UO) and not create an unsafe mindset. Even when the control (the Davis-Besse RO) makes an error in the control room, the effect may only be to delay or redirect the command (the shift superviser). TMI, however, shows the contrary for poor procedures and Ginna shows that procedures probably can never cover all contingencies; that there will always be interpretation required, particularly when conditions depart from the anticipated (as presented in training) or the corporately known (learned from other events at other plants). Davis-Besse shows that procedures cannot negate the strong (even if not wrong) beliefs of operators, that procedures do not make a person an automaton. In a distributed decision-making situation,[34] as in a nuclear power plant, a hesitancy in situation assessment can arise from social dynamics. This volatile potential of a combination of unanticipated conditions along with distributed distributed decision making (DM), as indicated by these few events, must call to question the efficacy of accident management.

The events at Browns Ferry, TMI, Ginna and Davis-Besse demonstrate that ultimately it is the 'belief states' of operators and other personnel that direct their actions unless they are entirely overcome by uncertainty. At Browns Ferry, the operators believed that water was not the means to put out the fire and did not use a readily available resource. At Three Mile Island, the operators believed going solid was the most important aspect of the panoply of cues presented them and acted appropriately with respect to this belief. At Ginna, a steam bubble led to a debate between the control room crew and the TSC staff over terminating or continuing SI, the delay of which was a violation of procedure. And at Davis-Besse, the belief that secondary restoration was imminent and that feed and bleed would prove costly (and might not be effective) led to a procedure violation. Hence, although a nuclear plant is a procedure-laden environment, operators and other personnel still (and should!) operate under their own judgment, which means that the pejorative term of *violation* may only be an indicator of the operators' doing the right thing.

THERE MAY BE ENOUGH MADNESS TO FORM A METHOD

Current developments toward an HRA method for cognitive error modeling[27,28] amount to variations on the 'blackbox,' 'give-me-a-number' school of risk assessment. This is an unfair assessment (of course) since the EPRI method is in part an analysis of simulator data and the Idaho National Engineering Laboratory (INEL) work considered actual events as documented in LERs. However, the leitmotiv of each method seems to be to provide a 'database,' analogous to the numbers of Chapter 20 in THERP,[1] which are intended to be applicable under quite general circumstances, that is to say, *sans* context.

This effort is remindful in intent (sic) and depth of the various mathematical approaches to common cause failures generally, e.g. the 'multi-greek letter' approach.[35] Such methods provide much less than meets the eye. Although providing a way to quantify, no way is provided to analyze, i.e. *qualify*, and hence, the answer cannot be input usefully into a utility risk *management* program: what cannot be described, cannot be managed.

However, there is no reason as yet to give up hope of a model, or at least a framework within which to identify and specify cognitive error likely situations among risk-significant scenarios. The tools of decision analysis used to analyze actual and simulated, i.e. past, events[36] can be modified to provide structure to this search. The techniques developed for knowledge acquisition in artificial intelligence research can be used to supplement the structural analysis.

To develop a cognitive error framework, the cognitive paradigm in a starting place (see Fig. 7). Cognition is a combination of cognitive processes (our competencies, capabilities, and skills related to knowledge and its use in control) and the brain structures that amount to our knowledge, beliefs, and prejudices. This interactive (and possibly inseparable) blend of process and content is influenced both from within, by *affectors*, and from without, by *situational* signs. The most obvious, i.e. objective, of these are the external influences. These amount to a situation's

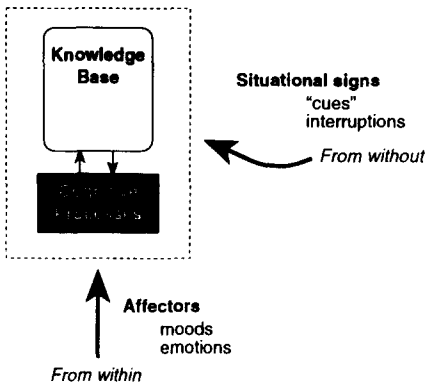

Fig. 7. Oh no, not another engineer's model of cognition!

morphology, represented by a prototypical or optimal response along with potential interruptions that lead the response away from the optimal, forcing the reality of satisficing,[37] leading to paths, i.e. action sets, that are possibly initially suboptimal to those that might ultimately be failure (see Fig. 8). An HRA method that handles cognitive error must address the possibility that a sequence of suboptimal responses might lead to failure (or when do two 'oks' make a 'bad'?).

PRA structures are already robust enough to handle the scenario transition effects of such non-binary possibilities.[38] HRA techniques already have means to transform the sequence and system oriented PRA structures into useful human factors representations, such as timelines, Murphy diagrams, link analyses, etc. The need is to be able to identify (i.e. predict) the interruptions.

The knowledge base (KB) is also accessible. In the nuclear setting, the EOP context described previously represents the nominal if not normative foundation of how operators will come cognitively to a scenario. Elicitation techniques can get at more, as long as we are mindful of the unreliability potential of elicitee and elicitor alike. Another caveat is that the KB is distributed (and peculiar from individual to individual) as is the decision process.

The cognitive processes that people use to exercise the KB are much less accessible. Cognitive psychology has made inroads into the error forms[16] of cognition. These error forms amount to biases (to use the most deprecating term) or heuristics (to use a more neutral one) that are a side effect of the amazing cognitive prowess humans possess. As Reason[39] puts it, there is a cognitive balance sheet on which for each assest, there is a debit that amounts to a source of error.

Table 6 lists and indicates the meaning of several cognitive error forms. These forms have been inferred from observations of performance and protocol analyses. The error forms, like the under trimmings of the other developing cognitive error methods, are good *retrospective* analysis; but they lend themselves minimally to *predictive* analysis. The table will be allowed to stand alone; a source document for each concept is provided for further description.

The fuzziest element in the cognitive paradigm of Fig. 7 are the affectors, i.e. emotions, moods, etc. that at a given time may dominate one's behavior. Cognitive psychologists associated with HRA have systematically avoided modeling these. Swain and many others have listed them; however, these elements seem to have the least predictive potential of all of the influences human.

The desired full cognitive program would be able to identify the kinds from each category of factors and pre-specify, i.e. qualitatively predict them. Table 7 provides an initial prediction as to how far this program may proceed; an answer of yes for all factors would facilely lead to a quantitative model, considering the proclivity of risk analysts to guess at even the mysterious. An example restricted to HRA is SLIM.[2] At least under these ideal conditions, any requisite guessing will be restricted to numerics and be constrained by a fully specified situational assessment (unlike what any operator will be privy to in an actual event). Table 7 claims that the full program can only partially be achieved, but this may be enough to close some unfinished business.

From the foregoing discussion it seems that Swain's original thesis mentioned above may indeed carry the day, but it may be modified so much as to make it unrecognizable. First, a task analysis should be replaced by a cognitive task analysis. Here, the focus

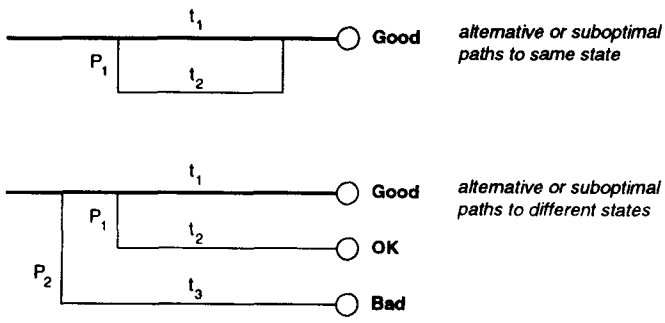

Fig. 8. Satisficing as it applies to performance measure.

Table 6. Cognitive error forms

Form	Definition	References
Availability heuristic	Events are judged likely if it is easy to imagine them or they are remembered to have occurred or occurred recently	34
Cognitive hysteresis (functional fixation, mindset)	Tendency to stick with a decision due to bias to search only for confirming evidence, the danger of partial explanations, and the similarity between actual and perceived events	40
Cognitive overload	Simple fact that cognition is finite	39
Confirmation bias	Only noticing or seeking data confirming belief	39
False consensus	Erroneous belief that others share one's belief	34
Frequency bias	Use of most popular 'solution'	39
Groupthink	Social phenomenon of adhering to shared beliefs	41
Halo effect	Bias toward oversimplification	42
Hesitancy (decisional trauma)	Cognitive lockup because of goal competition or the presence of only undesirable options	43
Matching bias	False positive, matching partially, ignoring rest	39
Overconfidence	Thinking too much of one's 'expertise'	39
Pluralistic ignorance	Erroneous belief that one's at odds with all others	
Risky shift	Group's tendency to adopt more extreme position than that of any individual	
Satisficing	Seeking 'good enough' rather than optimal	37
Strong but wrong	Errors from autonomous processes one is used to accepting	39
Theorizing	Filling in, i.e. imagining, the gaps of a hypothesis and blending fact and theory indistinguishably	44

Table 7. HRA context—the chances of prediction

Area	Can context be identified?	Can context be predicted?
Knowledge base	Yes	Somewhat
Cognitive processes	Indirectly, retrospectively	Hardly, if at all
Effectors	Yes	Yes
Affectors	Generically only	Seems impossible

is on decisions, diagnoses, and plans, rather than an ever-cascading series of motor-perceptual subtask elements. These latter 'atoms' of performance may have some reliability impact, particularly in routine tasks, but are not likely to play a significant role in cognition during dynamic, offnormal events in a nuclear power plant. The Davis-Besse event shows that slips do change the 'boundary conditions' of the cognitive context, but they are themselves eminently recoverable (at least when made in the control room).

The investigations into cognitive error alluded to above suggest that there is no danger of there not being a response at all when it comes to diagnosis. People actively seek stimulus and will invent it when there are gaps in time or logic.[44] Hence, the evolution from a cognitive error abstractly looks like Fig. 9 (adopted with modification from Fig. 2 in Ref. 45), where an early diagnosis is flawed but the action proceeds 'successfully' according to the misdiagnosis. This phenomenon can be extended to the planning aspects of cognition, e.g. in particular to flawed or incomplete or inappropriate procedures (Ref. 46 for a glimpse into planning errors).

Under stress, cognitive hysteresis is the major risk, i.e. the team pursues its early, possibly incomplete beliefs irrespective of future contrary cues, possibly taking on a social dynamic similar to groupthink,[41]† but more likely simply because a completed plan 'confers order and reduces anxiety.'[47] This suggests that the key role of symptomatic EOPs are not to provide the right symptoms but to provide strong enough cues to combat the possibility of cognitive hysteresis. It also means that the crediting of operators merely for their likely presence in the control room, as does the so-called dependency model of THERP (Chapter 10 in Ref. 1), must be re-examined. The issue is whether or not more operators really provide greater human redundancy.

† This behavior is literally irrational but normative psychology is not the issue here; it is not abnormal but highly likely behavior under certain circumstances.

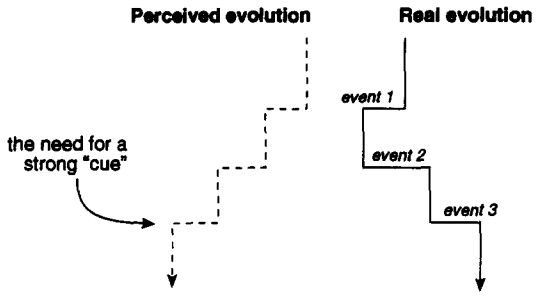

Fig. 9. The role of symptomatic EOPs: how strong are strong cues?

Many cues (too many) were available to TMI operators. The new EOPs assure that a formal mechanism exists to chose among them and find the correct cues. But will the human mechanism that assures the proper use of the formal procedures surface?

The foregoing discussions also assure that we know abstractly where cognitive errors and violations will occur in the scheme of HRA types of events (Fig. 10). At every major decision point, they will occur up front, attached, as it were, to the next goal in the goal matrix of the scenario, either as a gross omission of the goal, e.g. a violation that amounts to deferring or rejecting the goal, or an action that is improper relative to achieving the goal, i.e. the plan or procedure is inappropriate. THERP made the claim that these up front events—'the occasional error of decision-making'—were within the uncertainty bands applied to the associated response failures, i.e. the slips and suboptimal responses.[1] It does not appear that the notable events synopsized above lend credence to this claim.

FEED AND BLEED—THE NEVER-ENDING CASE STUDY

TMI introduced the option of 'feed and bleed' (FAB) to the PWR community. This is a contingency action to a loss of all heat sink that amounts to *deliberately* creating a LOCA by *bleeding* off reactor coolant system water through one or more pilot/power operated relief valves (PORVs) attached to the pressurizer and *feeding* the RCS, i.e. making up the water lost to containment, using the safety injection (SI) equipment as would be used in any LOCA. (Some plants use the term 'bleed and feed' because the bleed option needs to precede feed in order for the pumps to work.) In the terminology at some CE plants, FAB provides a once-through-cooling of the core, primary cooling, that replaces secondary cooling, the complex thermal-hydraulics of the reactor coolant pumps (RCPs), the steam generators, and the secondary plant, i.e. the normal *heat sink*. However, since there are multiple options for accomplishing the normal heat sink, the loss of all of this capability is quite incredible to operators (although we in PRA, as well as a few notorious actual events, seem facile at conjuring it up).

However, and this is significant, FAB does *create* a LOCA, which is precisely the type of event that the whole of reactor safety design bases were meant to avoid and mitigate. At the very best, the plant faces a large cleanup and significant downtime as a result. Of course, as the Davis-Besse event reminds us, the plant may face a long shutdown (14 months in the case of Davis-Besse) if they do not feed and bleed!

Hence, FAB presents the following quandary, a classical *goal conflict:*

- FAB is clearly not a preferred option—it is a last resort undesired by all plant operators.
- Yet, FAB is quite easy to accomplish (from the actual implementation point of view).
- However, no operators really believe that FAB will be needed because of the reliability and diversity of secondary cooling.
- Further, operators, according to procedure, will focus on secondary cooling restoration prior to the arrival to FAB criteria.
- Unfortunately, FAB often (for some plant types) has a short window of opportunity for success.

Secondary options are numerous: some are generic, e.g. main feedwater (MFW), and others are peculiar to the particular plant; some options may be easily determined to be unavailable (which depends on the cause of the loss of secondary cooling and its ability to be diagnosed easily) and other options may take

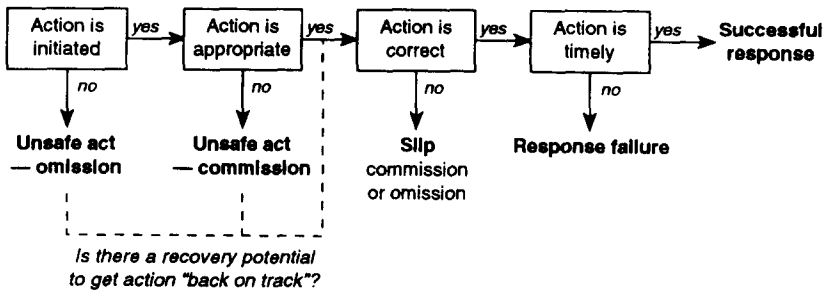

Fig. 10. The general paradigm for unsafe acts.

considerable time to implement, e.g. 'blowing down' the secondary side to use available condensate pumps. Hence, FAB resides in a goal matrix as presented in Fig. 11. The goal quandary above is transmitted to operators as the caution: 'DO NOT . . . unless . . .' Of course, the official line is that FAB is viable and when the cue arrives for FAB, i.e. the FAB criterion is met, then 'you will follow procedure and implement it.'

In reality, operators are thinking adults, highly skilled and trained but knowledgeable of the larger picture as much as anyone. Their decision is not 'just "pushed" by premises but "pulled" by the goal . . .'[48] The larger picture includes the four floating goals in Fig. 11: avoidance of plant trip, avoidance of long shutdown, avoidance of core melt, and avoidance of radioactive release. One of these is already lost, since the loss of heat sink always leads to tripping the plant.

Notice that failure of these goals has progressively more detrimental consequences: loss of revenue, major loss of revenue (e.g. replacement costs of electricity for customers), permanent shutdown of the plant, and the possibility of deaths (to neighbors, no less!). However, the plausibility (not only to operators, of course) of any such consequence also decreases, becoming not so much more abstract, but rather less credible. At least prior to the Davis-Besse event, the failure of the goal to restore secondary cooling implies the failure of the goal to avoid a long shutdown, whereas achievement of primary cooling fails that significant goal while only being one way to achieve core melt avoidance. Hence, there is a natural bias in favor of restoring secondary cooling. Also recall that the most likely failures of multiple train systems such as auxiliary or emergency feedwater (A/EFW) are faults in actuation signals that can be corrected by simple remote manual actuation. So, based on the implausibility of core melt alone (or more concretely, the incredibility of the loss of all heat sink options), it is easy to see why operators (at least pre-DB) might opt to defer implementing FAB beyond what procedure designers felt was the last minute. (Such a delay is officially an instance of the availability heuristic with a reversed gambler's fallacy twist: it has not happened so far; it will not happen today.)

Now that the goal level of the situation is specified, Fig. 12 shows the morphology of the feed and bleed situation from a functional level. Seven branch nodes, four of which involve decisions by the operators denoted by darker diamonds and the others of which are results oriented, define the likely 'flow paths' through the situation. Two 'collectors' exist: *A*, which is where the core melts; and *B*, which is where the operators re-establish a heat sink in time to avoid core melt. Notice that the game is never really up, since the operators would continue to try to restore heat sink even beyond the onset of core damage. Hence, the definition of success in consequence relative to success in performance are not necessarily the same. Also, success is goal-dependent. One might fail to save the core but assure no major release of radiation, which is a failure of incident response but a success of accident management. (It still leads to the permanent shutdown of the plant, however.)

Were the probabilities and times indicated on the figure known, a simulation could accumulate statistics on the relative probability of filling the two collectors, yielding an overall failure probability. Notice that, although the TRC concept has foundational problems as it is conceived as a global reliability indicator, there

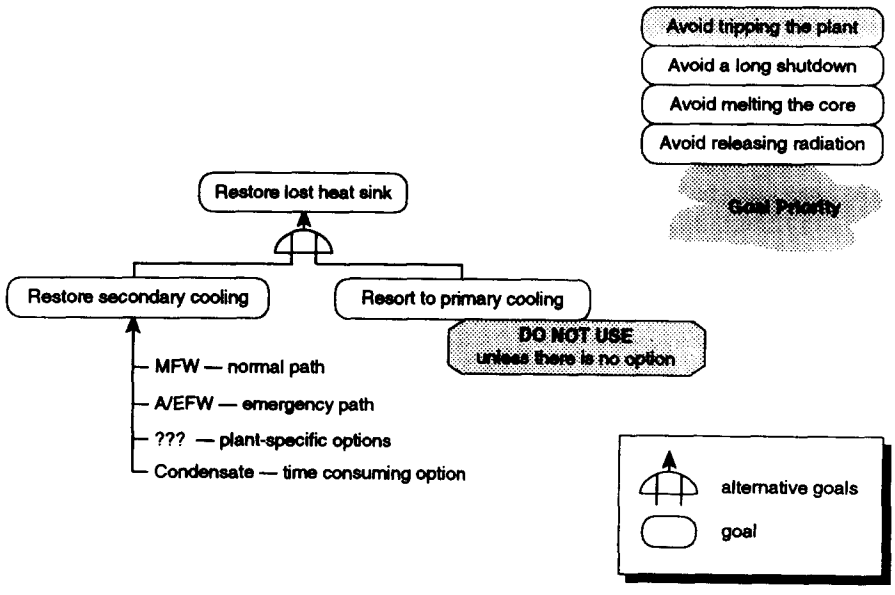

Fig. 11. The goal matrix of feed and bleed.

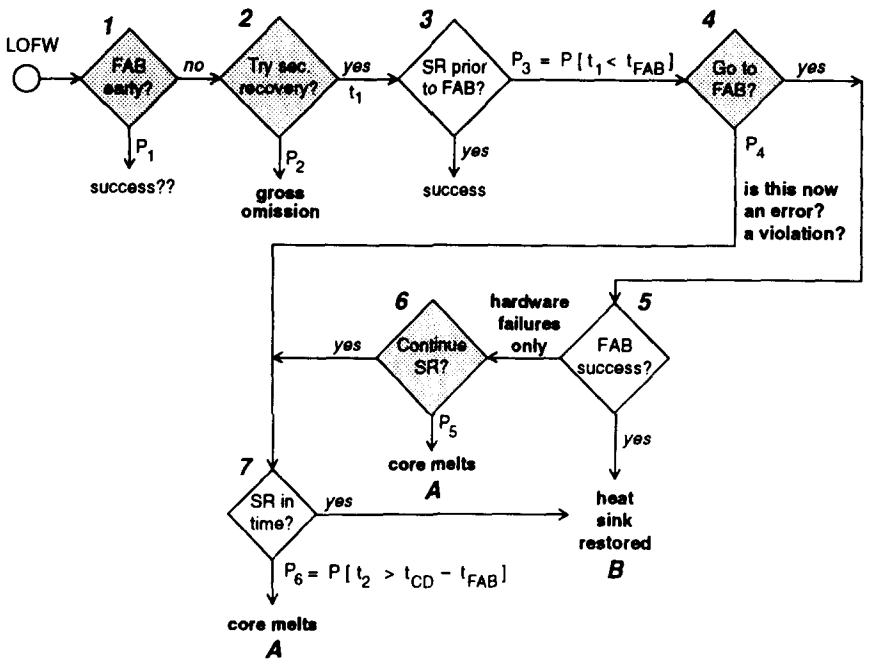

Fig. 12. The functional morphology of feed and bleed.

is clearly a time element in the functional description of the situation, since failure has a time factor in its definition, i.e. the operators do not fail *until* the core melts. The major advantage of such a simulation is of course, its facility in changing the parameters that produce the probabilities or the time distributions to directly reflect the various elements of cognitive context, e.g. those listed in Table 6. In this way a simulation can become an important determinant of the *consequences* of the analyst's assumptions, which is the true merit of codes such as the Cognitive Environment System (CES),[49] or the Dynamic Logical Analytical Methodology (DYLAM),[50] or even the dynamic event tree concept.[9]

Figure 12 allows a search for cognitive errors to proceed systematically as in Table 8. This is because there is no evidence that human error phenomena is a 'random' process, although people can become so

Table 8. Specification of feed and bleed unsafe acts

Node	Unsafe act	Potential causes	Result
1	Commission	Overtrained to loss of FW; short window of opportunity (not generic); distrust of FAB	Actually a success!
	Omission	None	—
	Violation	None	—
2	Commission	Incorrect priority; work on one option too long; think finished and move on	Suboptimal Possible failure New or persisting cues
	Omission	None	—
	Violation	Opt not to try condensate because of duration of effort	Suboptimal
4	Commission	None	—
	Omission	Can only be violation	—
	Violation	The basic quandary at Davis-Besse has not been removed, maybe reduced in likelihood	Core melts
6	Commission	None	—
	Omission	Might get involved with recovering FAB faults	Core melts
	Violation	None	—

Table 9. Diagnostic load in feed and bleed

Required diagnosis	Credibility
Recognize loss of heat sink	Cannot be a problem
Troubleshoot secondary cooling	May be difficult depending on the kinds of fault but the basic fact that there is no heat sink will not change
Recognize FAB criterion	Not as long as parametric symptoms are used

panicked that their behavior is essentially unpredictable, a kind of behavior 'control' that might be called stochastic.[6] The table would require considerable explanation, but first notice the results. First, *vis-à-vis* decision making, there are only four general opportunities for cognitive error. This limitation arises from the assumption—bolstered by common sense, observations of simulator exercises, and reviews of actual events—that operators are fundamentally procedure-following, that human performance in the early stages of an offnormal situation (no matter how severe it may seem or the workload of the cue demand) is not random (but *opportunistic* in Hollnagel's classification of control) and that operators are well in tune with the basic safety goals as implied in Fig. 11. Hence, we must be allowed to ignore those who continue to yell 'commission error' but who come to the table with nothing but their alarm.

Second, a protest may arise that the possibility of misdiagnosis (what was in a loose interpretation the source of the TMI commission error) has not been addressed. Recall first that 'there is no one, single diagnosis task in dynamic worlds; rather there is a continuing need for situation assessment. . .'[45] This is true and will be accommodated next. Recall next that the implementation of the new symptomatic EOPs was motivated by the concern that a complex event could not be identified with reliability. TMI was blamed (in part) on the event-oriented procedures of the day.

A PRA performed post-TMI but pre-symptomatic-EOP used the concept of a confusion matrix to reflect the potential for confusing one event for another.[51] The prototype event considered at that time was a steam generator tube rupture, whose salient cues are: secondary radiation alarms, LOCA signs in the RCS (decreasing pressurizer level and decreasing RCS pressure), and non-LOCA signs in the containment (i.e. water is exiting the ruptured generator and not filling the containment sump). It was postulated that the radiation alarms might be missed and that the operators would fixate on the *proximal* cues for LOCA, pressure and level in the RCS, while ignoring the fact that the sump level was not also increasing (e.g. because of a confirmation bias). Note that the sump level is a *distal* cue, since it may not arrive until a significant time after the RCS LOCA signs are apparent and relates to a 'down stream' effect rather than an immediate correlate to a loss of inventory.

In the post-TMI EOPs, however, *re*-diagnosis is highly supported (e.g. the transfers of Fig. 3) and the symptom sets are prioritized according to safety function priorities, which is at least directed toward combatting the odd human capacity to be enthralled by the least significant thing ongoing. Further, as Fig. 4 indicates, there is a safety function cue 'checker,' the STA, who because of his non-licensed status (which might be a minus in some ways) may not be as susceptible to the mindsets of the SRO and the ROs (however, his authority is low). Hence, without an accompanying failure of the secondary radiation monitors, it is difficult to credibly foresee a persistent misdiagnosis of an SGTR event, even though initially LOCA and SGTR mitigation is the same and is performed by automatic systems rather than human actions.

Returning to the FAB example, is there a possibility for misdiagnosis there? The Integrated Reliability Evaluation Program (IREP) Arkansas Nuclear One (ANO) PRA assessed the potential for failing in the FAB option[52] to be predominantly due to misdiagnosis and the remainder to failure of implementation (slips). This, of course, is consistent with the THERP assumption that decision making is in the uncertainty of the modeling. But, let us see what this really means. The plant in question is the same type of PWR as was TMI. The new Babcock & Wilcox (B & W) EOPs are highly focused on heat sink. The crew is generally as in Fig. 4. There simply is no credible argument that the operators would not notice the relevant alarms. Moreover, the EOP system does not (initially) require an alarm response: any time the plant trips means that the initial verification steps in the EOPs are implemented. Thanks to TMI, if operators at a PWR will notice anything, it will be a loss of heat sink (although they may not accept this belief).

Table 9 seems to be the complete diagnostic load that a 'vanilla'† loss of all heat sink scenario will put on the operators. Recognition of lost heat sink cannot

† Woods has criticized the HRA community for focusing too much on single event, i.e. vanilla or textbook, scenarios. He is correct from a performance perspective, but may not be from a reliability perspective. Instrumentation failures, which would make a scenario more 'interesting' from an HRA perspective, merely lower an already low frequency of occurrence and hence mean the scenario is not interesting from a risk perspective.

be a problem considering the diagnostic context. As the operators attempt to restore secondary cooling (if time permits), then the diagnosis of the root causes of the faults, i.e. trouble shooting, may be difficult, but the failure to find the fault is easy to realize: the heat sink simply is not re-established. Finally, the FAB criterion, or the cue to implement FAB, is the only remaining diagnostic requirement. One of the purported problems at Davis-Besse was that the operators did not realize that the SGs had dried out and hence they did not sharply implement FAB. This is not likely. The actual cue is on a specific level reading from the SG level instruments. The term 'dry out' is a euphemism for this parametric cue but is not confusing as a cue, since it is not literally used as one. The delay arose because of the decisional quandary outlined above plus the important reality that the assistant shift supervisor (the second in command, so to speak) had relayed that he was almost ready with an option to restore secondary cooling. The operators knew that the SGs were dry; it was not the significant element in their belief structure.

Hence, the IREP analysis is off base. The key to the FAB situation is the decision and its time complex. What is needed for the HRA of the situation is a qualitative scheme to direct quantification. (Note that at least one TRC method has such a method, using the concept called burden.)[8] will proceed one more level toward that.

The formal structure of the FAB situation, the EOPs, is another potential clue to the possibility of cognitive errors. One question is whether the events of TMI and Davis-Besse, which are at the root of the feed and bleed issue, are, at least formally, accommodated by the new EOPs. Figure 13 shows the trace of the planned use of the CE EOPs for the loss of heat sink situation. The first 10–15 min will be spent assuring all critical safety and supporting equipment are operating as needed. This will include the attempt to restart MFW and start AFW, the failure of which is a direct cue of the loss of heat sink. However, at this time nothing else will be done toward *restoration* of the heat sink.

At the end of EOP-00, the operators will use the Diagnostic Actions flowchart. This will confirm the loss of heat sink and indicate that multiple failures are not diagnosed. Note that this is an interesting semantic, since all of main and auxiliary feedwater must be faulted, to make heat sink loss apply: clearly multiple failures. The non-multiplicity refers to the fact that no other CSF is affected. If this is the case, the symptom set of EOP-06 is met and the operators will transfer to that EOP.

EOP-06 directs the operators to attempt restoration of secondary cooling, trying the four options in the priority of Fig. 11. During this time (and previously), the levels in the steam generators have been decreasing, reflecting the loss of feedwater. At the instrumented low levels, and when the RCS temperature begins to rise (again a function of the loss of heat sink) the criteria for implementing FAB will be met. (Note that, however, this criterion is slightly different at some other plants.) It is at this point that the operators are supposed to transfer to EOP-20 and implement feed and bleed. From this point on, the instructions are found solely in this procedure and they amount to treating the scenario like a LOCA, which is caused by opting FAB.

The foals of Fig. 11—the four 'floating' goals and the explicit scenario oriented ones—can be matched to the major EOP steps and the nodes of Fig. 12. The abbreviations can be inferred from Figs 3, 11 and 12 and are less important than the fact that the goal matrix changes over time as the functions are implemented using the EOP sections. This results in Table 10, which is similar to the goal-switching timeline developed by Reinartz and Reinartz.[53] What the table shows (if anything) is the 'conflict' between the crucial goals of avoiding a long shutdown (LS) and avoiding core melt (CM). It seems reasonable to assess that both goals are 'active' at one of the EOP nodes and two of the functional nodes and that this is precisely where the conflict exits. (It also seems plausible that this is an analytical gimmick to aid in our understanding and that nothing like this goes on in the brains of operators.) Whether this is a measure of what is referred to as decisional *burden*[8] or not, there seems to be a mechanism to capture a dimension

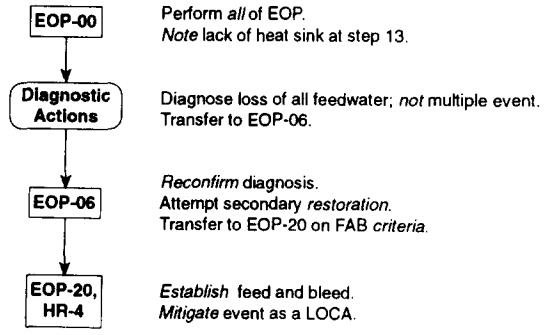

Fig. 13. An EOP trace of feed and bleed.

Table 10. Goal switching in feed and bleed

	EOP node				Functional node						
Goal	00	DA	06	20	1	2	3	4	5	6	7
AT											
LS	×		×		×	×	×				
CM		×	×	×		×	×	×	×	×	×
RR										×	
HS		×	×	×	×	×	×	×	×	×	×
SC			×		×	×	×			×	×
PC				×				×	×		

of cognitive overload, namely when several goals can be active.

As one final representation of the feed and bleed situation, the goal-oriented perspective of Fig. 11 is merged with the EOP trace of Fig. 13 to produce a *Hollnagel diagram* (Fig. 14). (Hollnagel would call this a goal-means task analysis.)[6] The formalism recognizes that goals may have preconditions, which amount to subgoals. Each (sub)goal has one or more tasks that are needed to meet the goal. Fig. 14 shows information about the procedure logic (namely EOP-06 and EOP-20 in Fig. 3), the task location, and the task complexity (at least its basic logic). It does not, as in Fig. 12, show the chronology (the chronologic) of the situation.

The potential for cognitive problems shows up as the fact that a precondition for one task, establishing FAB, is the *negation* of the principal precondition of the main goal, restoring the lost heat sink. This amounts to a 'do it—do not do it' form which could lead to trouble without any of the other elements of the decisional quandary of feed and bleed. Another problem the Hollnagel diagram can be used to point out is that a precondition for restarting the preferred option, secondary cooling, is that the equipment be able to be started from the control room. Barring that, auxiliary operators will be sent to try to start pumps or open valves locally and this raises the spector of distributed decision making (a key causal contributor in the Davis-Besse event) and a reduction in control for the control room crew. This physical distance produces a cognitive distance that depends on communication and patience on the part of the operators in the control room. If you ever want to frustrate a simulator crew to distraction, emulate an auxiliary operator being requested to locally return a piece of equipment to service and do not report back to the control room crew for a while.[21]

It should be noted that from Fig. 14 there is a potential for misinterpreting the establishment of FAB as more difficult than restoring secondary cooling. However, this is merely because the various activities that accompany the four options for secondary cooling listed in Fig. 11 are not modeled in the diagram. It is also likely that for most plants, the concern over A/EFW will dominate the restoration activities in the precious little time prior to the arrival of the FAB criterion. By partitioning the varieties of a scenario into what PRA calls cutsets, then a variation on Fig. 14 might explicate more of the primary goal's activities.

A final use of a Hollnagel diagram as in Fig. 14 is that it is the foundation of a link analysis, i.e. the assessment of the crew as resources. This example

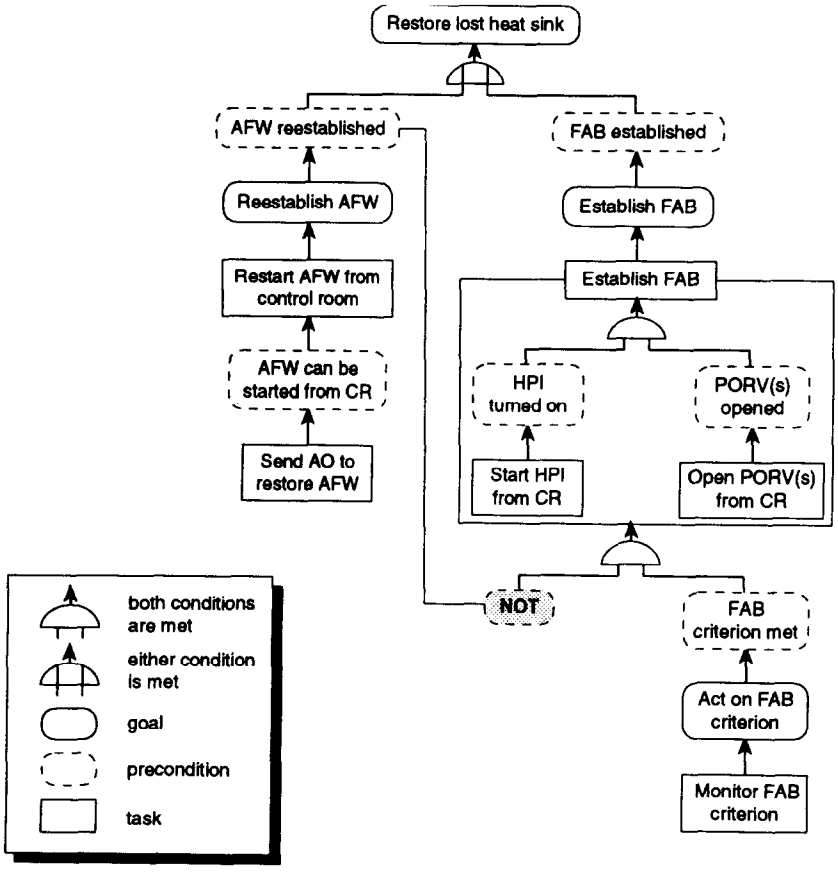

Fig. 14. A Hollnagel diagram of feed and bleed.

shows that the SRO has a role, namely directing the activity; that the STA has a role in monitoring the safety goals, particularly the FAB criterion; that two ROs are busy, one for the secondary equipment and one for the primary equipment; and at least one auxiliary operator (and there may be no others depending on the time of the day) is busy with remote restoration activities. So it would seem that coordination and communication is a major part of the context of this situation and should impact the reliability of the activities.

CONCLUSIONS

It is too early to specify to an *n*th degree a cognitive error HRA method. We are still too much like a conceptual marketplace rather than a profession.[54] However, the above discussions seem to suggest that it is possible to provide a systematic means to identify and, to a lesser degree, specify the potential for cognitive errors in an EOP-dominant environment.

For example, after performing a recent HRA involving the EOP system of Fig. 3, the strategy of Table 11 suggested itself. EOP-00 is essentially contextless and standard HRA techniques can be used to identify and quantify events associated with it. EOPs-01 through -07 are the single event procedures and cue timelines, symptom complexity, and the potential for cognitive hysteresis should be considered. Finally, EOP-20 is a complicated, long procedure which has a fixed chronology that will seldom fit the ongoing event in a way that seems optimal to the operators (or the analysts). Here cognitive overload or distraction might be a problem and even the potential for looping, i.e. getting lost, is a real possibility.

It is also clear that more investigatory use of plant simulators is in order.[21,27] Although the hesitancy that the decision quandary of FAB suggests should exist was indeed observed in simulators[55] early into the implementation of the symptomatic EOPs, that phenomenon has apparently disappeared as the nuclear safety culture ingrains the option into operators more and more. (Operators will never fail NRC requalification exams for 'erring toward safety,' as going early or promptly to feed and bleed in a simulator may be interpreted.) However, there may be vestiges of this reluctance that can be elicted by more variety and greater effort, neither of which has been attempted to date.

This paper has attempted to indicate that *analysis* should remain a fundamental part of human reliability analysis. It has done so by examining cases of cognitive problems, if not errors, and by attempting to provide a system to bound, if not explicate, cognitive context. This system should include four major efforts:

(1) Identify the goal matrix for the situation and any possible goal conflicts, e.g. Fig. 11 and Table 10. This will give justification to assessment of cognitive hysteresis, burden, and other global contextual influences.
(2) Develop a functional chronology of the situation, e.g. Fig. 12. This will allow the analyst to
    (a) specify the time matrix of the actions and decisions to be made,
    (b) provide a structure to identify personnel links and the impact of distributed decision making on the reliability of the task(s),
    (c) estimate cognitive workloads, and
    (d) identify and locate the likely cognitive error modes and forms, e.g. Table 8.
(3) Perform a cognitive task analysis, e.g. the Hollnagel diagrams as in Fig. 14, and a procedure analysis, e.g. as sketched in Fig. 13. Then task and complexity indices may be estimated, diagnostic demands identified, e.g. Table 9, and an assessment of the quality of procedures and overall preparation may be made.
(4) Perform knowledge acquisition, e.g. interviews of operators, simulations, and walkdowns, to 'flesh out' the morphology generated above.

When this morphology is produced, then the

Table 11. Context index of EOPs

EOP	Context	Possible HRA strategies
−00	Contextless	Use THERP's annunciator model or simply an omission slip model
−01 to −07	Rule-based cue matching	Might need to audit cue reliability *à la* confusion matrix; consider cue timeline as source of cognitive hysteresis
−20	Potentially rule-based; but priority might not 'fit'	Cognitive overload or distraction or rushing to perceived most important action or consider potential for getting lost

assessment of subtler influences such as in Table 6 or the types of control suggested by Hollnagel[6] is made possible. From that basis (and only until then) quantification *à la* SLIM or whatever will have substantial justification and the 'results' will be translatable to useful insights for further risk management.

Then with this approach, maybe a sufficient story can be developed in order to 'guestimate' what is now described as a generic, hence contextless, event such as 'circumvent procedure with potentially catastrophic consequences' (Table 1, p. 131 in Ref. 28), replacing it with a specific event such as 'delay FAB because the AFW pumps have a history of failing on start and because the plant personnel are concerned that FAB will not be effective.'

The opportunity to progress is imminent; now if only the opportunity to try it out were to arise.

ACKNOWLEDGMENTS

The author is indebted to Erik Hollnagel of Computer Resources International, who provided much of the grist and motivation of this paper, but who is in no way responsible for its specific features or quirky ideas. The author also indebted to the two referees of the early draft of this paper for their fine suggestions, most of which were accepted. I know that one of them was David Gertman of Idaho National Engineering Laboratory and thank him personally.

REFERENCES

1. Swain, A. D. & Guttmann, H. E., Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications. NUREG/CR-1278, USNRC, Washington DC, August 1983.
2. Embrey, D. E., Humphreys, P., Rosa, E. A., Kirwan, B. & Rea, K. SLIM-MAUD: An approach to assessing human error probabilities using structured expert judgment, NUREG/CR-3518, USNRC, Washington DC, March 1984.
3. Chien, S. H., Dykes, A. A., Stetkar, J. W. & Bley, O. C. Quantification of human error rates using a SLIM-based approach. Conference Record for 1988 IEEE Fourth Conference on Human Factors and Power Plants, 88CH2576-7, Institute of Electrical and Electronics Engineers, 5–9 June 1988, pp. 297–302.
4. Zamanali, J. H., Hubbard, F. R., Mosleh, A. & Waller, M. A., Evolutionary enhancement of the SLIM-MAUD method of estimating human error rates. *Trans. American Nuclear Society*, 7–12 June 1992, Boston, TANSAO 65 1-580 (1992), vol. 65, 1992.
5. Rasmussen, J., *Information Processing and Human–Machine Interaction*. North-Holland, New York, 1986.
6. Hollnagel, E., *Reliability of Cognition*, Academic Press, (in press).
7. Hall, R. E., Fragola, J. R. & Wreathall, J., Post Event Human Decision Errors: Operator Action Tree/Time Reliability Correlation, NUREG/CR-3010, USNRC, Washington DC, November 1982.
8. Dougherty, E. M. & Fragola, J. R., *Human Reliability Analysis: A Systems Engineering Approach With Nuclear Power Plant Applications*. Wiley Interscience, New York, 1988.
9. Acosta, C. & Siu, N., Dynamic event trees in accident sequence analysis: application to steam generator tube rupture. *Reliability Engineering & System Safety*, April 1992.
10. Dougherty, E. M., HRA—Where shouldst thou turn? *Reliability Engineering & System Safety*, **29** (3) (1990) 283–99.
11. Reason, J., The cognitive worm at the core of the TRC apple, unpublished.
12. Hannaman, G. W., Spurgin, A. J. & Lukie, Y., *Human Cognitive Reliability Model for PRA Analysis*, Draft NUS-4531, NUS Corporation, San Diego, CA, December 1984.
13. Kelly, D. L., On the human error probability for injecting boron during ATWS at a BWR. *Reliability Engineering & System Safety*, **35** (3) (1992) 253–5.
14. Spurgin, A. J., Moieni, P., Toksimovich, V., Luna, C. J., Parry, G. W. & Lydell, B. O. Y. A Human Reliability Analysis Approach Using Measurements for Individual Plant Examinations. NP-6560-L, Electric Power Research Institute, Palo Alto, December 1989.
15. Occupational Safety and Health Administration, Process Safety Management of Highly Hazardous Chemicals; Explosives and Blasting Agents. 29 CFR Part 1910, Washington DC, 26 May 1992.
16. Reason, J., *Human Error*. Cambridge University Press, Cambridge, 1990.
17. Reason, J. T., Absent-mindedness and cognitive control In *Everyday Memory, Actions and Absent-Mindedness*, ed. J. Harris & P. Morris. Academic Press, London, 1984, pp. 113–32.
18. Taylor, D. H., The hermeneutics of accidents and safety. In *The New Technology and Human Error*, ed. J. Rasmussen, K. Duncan & J. Leplat. John Wiley & Sons, Chichester, 1987, pp. 31–41.
19. USNRC, Loss of Main and Auxiliary Feedwater Event at the Davis-Besse Plant on June 9, 1985. NUREG-1154, USNRC, Washington DC, 1985.
20. Office of Nuclear Reactor Regulation, Potentially Significant Problems Resulting From Human Error Involving Wrong Unit, Wrong Train, or Wrong Component Events. NRC Information Notice No. 87-25, USNRC, Washington DC, 11 June 1987.
21. Dougherty, E. M., Implications from observing simulator exercises. SAIC internal report, June 1989.
22. Nuclear Safety Analysis Center, Analysis of Three Mile Island-Unit 2 Accident, NSAC-1, Electric Power Research Institute, July 1979.
23. Sugarman, R., Nuclear power and the public risk. *IEEE Spectrum*, November 1979, pp. 59–69.
24. EPRI, "Prelude: The accident at Three Mile Island. *EPRI Journal*, Electric Power Research Institute, June (1980) 7–13.
25. Lombardo, T. G., TMI Plus 2. *IEEE Spectrum*, Institute of Electrical and Electronics Engineers, April 1981.
26. Faust, C., conversations at the UCLA Workshop on Severe Accident Management for BWRs, 26–28 September 1990.
26. Laughery, K. R., Task network modeling of human operators in nuclear power plant control rooms. Proceedings of the Topical Meeting on Advances in Human Factors Research on Man/Computer Interactions: Nuclear and Beyond, June 10–14, 1990, Nashville, TN, American Nuclear Society, pp. 90–5.

27. Beare, A. N., Gaddy, C. D., Parry, G. W. & Singh, A. An approach for assessment of the reliability of cognitive response for nuclear power plant operating crews. In *Probabilistic Safety Assessment and Management,* ed. G. Apostolakis. Elsevier, New York, 1991, pp. 827–32.
28. Gertman, D. I., Blackman, H. S., Haney, L. N., Seidler, K. S. & Hahn, H. A. INTENT: A method for estimating human error probabilities for decisionbased errors. *Reliability Engineering & System Safety,* **35** (2) (1992) 127–36.
29. Pew, R. W., Miller, D. C. & Fehrer, C. E., Evaluation of Control Room Improvements Through the Analysis of Critical Operator Decision, EPRI NP-1982, Electric Power Research Institute, 1981.
30. Ballard, G. M. Reactor events involving misinterpretation/misunderstanding of plant status by plant staff. In *Proc. Int. Conf. Man–Machine Interface in the Nuclear Industry,* IAEA-CN-49/82, Tokyo, 15–19 February 1988.
31. Hanauer, S. H. *et al.,* Recommendations Related to Browns Ferry Fire, NUREG-0050, USNRC, February 1976.
32. Chexal, B. & Wycoff, H., Instrument and control bus power loss at Ranch Seco on March 20, 1978. NSAC-13, Nuclear Safety Analysis Center, November 1980.
33. Woods, D., Operator decision behavior during the steam generator tube rupture at the Ginna Nuclear Power Station. In *Analsis of Steam Generator Tube Rupture Events at Oconee and Ginna,* ed. W. Brown & R. Wyrick. INPO 82–030, Institute for Nuclear Power Operations, 1982.
34. Fischoff, B., Decision making in complex systems. In *Intelligent Decision Support in Process Environments,* ed. E. Hollnagel, G. Mancini & D. D. Woods. Springer-Verlag, Berlin, 1986, pp. 61–85.
35. Fleming, K. N. & Mosleh, A., Classification and Analysis of Reactor Operating Experience Involving Dependent Events, NP-3967, Electric Power Research Institute, June 1985.
36. Hollnagel, E., Pedersen, O. M. & Rasmussen, J., Notes on Human Performance Analysis, Risø-M-2285, Risø National Laboratory, Denmark, June 1981.
37. Simon, H. A., *The Sciences of the Artificial,* 2nd edn. The MIT Press, Cambridge, 1981 [originally, 1969].
38. Apostolakis, G. & Chu, T. L., Time-dependent accident sequences including human actions. *Nuclear Technology,* **64** (1984) 115–26.
39. Reason, J., GEMS: a framework for locating common human error forms. In *New Technology and Human Error,* ed. J. Rassmussen, K. Duncan & J. Leplat. John Wiley & Sons, Chichester, 1987, pp. 63–83.
40. Norman, D. A., New views of information processing: implications for intelligent decision support systems. In *Intelligent Decision Support in Process Environments,* ed. E. Hollnagel, G. Mancini & D. D. Woods. Springer-Verlag, Berlin, 1986, pp. 123–36.
41. Janis, I. L., *Victims of Groupthink.* Houghton Mifflin, Boston, 1972.
42. Reason, J., Recurrent errors in process environments: some implications for the design of intelligent decision support systems. In *Intelligent Decision Support in Process Environments,* ed. E. Hollnagel, G. Mancini & D. D. Woods. Springer-Verlag, Berlin, 1986, pp. 255–70.
43. Janis, I. L. & Mann, L., *Decision Making.* Free Press, New York, 1977.
44. Gazzaniga, M. S., *The Social Brain.* Basic Books, New York, 1985.
45. Woods, D. D., Coping with complexity: the psychology of human behaviour in complex system. In *Tasks, Errors and Mental Models,* ed. L. P. Goodstein, H. B. Andersen & S. E. Olsen. Taylor & Francis, London, 1988, pp. 128–48.
46. Reason, J., Collective planning and its failures. In *New Technology and Human Error,* ed. J. Rassmussen, K. Duncan & J. Leplat. John Wiley & Sons, Chichester, 1987, pp. 121–4.
47. Reason, J., The psychology of mistakes: a brief review of planning failures. In *New Technology and Human Error,* ed. J. Rassmussen, K. Duncan & J. Leplat. John Wiley & Sons, Chichester, 1987, pp. 45–52.
48. Volta, G., Time and decision. In *Intelligent Decision Support in Process Environments,* ed. E. Hollnagel, G. Macini & D. D. Woods. Springer-Verlag, Berlin, 1986, pp. 45–57.
49. Woods, D. D., Roth, E. M. & Hanes, L. F. Models of Cognitive Behavior in Nuclear Power Plant Personnel. NUREG/CR-4532, USNRC, Washington DC, 1986.
50. Amendola, A., Bersini, U., Cacciabue, P. C. & Mancini, G., Modelling operators in accident conditions: advances and perspectives on a cognitive model. In *Cognitive Engineering in Complex Dynamic Worlds,* ed. E. Hollnagel, G. Mancini and D. D. Woods. London, Academic Press, 1988, pp. 145–58.
51. Potash, L. M., Stewart, M., Dietz, P. E., Lewis, C. M. & Dougherty, E. M., Experience in integrating the operator contributions in the PRA in actual operating plants. In *Proc. ANS/ENS Topical Meeting on Probabilistic Risk Assessment,* Port Chester, NY, September 1981, pp. 1054–63.
52. Kolb, G. J. *et al.,* Interim Reliability Evaluation Program—Analysis of Arkansas Nuclear One, Unit One Nuclear Power Plant. NUREG/CR-2787, USNRC, Washington DC, June 1982.
53. Reinartz, S. J. & Reinartz, G., Verbal communications in collective control of simulated nuclear power plant incidents. *Reliability Engineering & System Safety,* **36** (3) (1992) 245–51.
54. Rasmussen, J., Cognitive engineering, a new profession? In *Tasks, Errors and Mental Models,* ed. L. P. Goodstein, H. B. Andersen & S. E. Olsen. Taylor & Francis, London, 1988, pp. 325–34.
55. Woods, D. D., Wise, J. A. & Hanes, L. F., Evaluation of Safety Parameter Display Concepts. NP-2239, Electric Power Research Institute, February 1982.