THE EVOLUTION AND DESIGN OF SAFETY MANAGEMENT SYSTEMS

The application of Cybernetics and the Management Oversight and Risk Tree (MORT) paradigm to the analysis of safety management.

JOHN CHRISTOPHER KINGSTON-HOWLETT
Doctor of Philosophy

University of Aston in Birmingham

September 1996

SUMMARY

Research in safety management has been inhibited by lack of consensus as to the definitions of the terms with which it is concerned and, in general, the lack of an agreed theoretical framework within which to collate and contrast empirical findings.

This thesis sets out definitions of key terms (hazard, risk, accident, incident and safety) and provides a theoretical framework. This framework has been informed by many sources but especially the Management Oversight and Risk Tree (MORT), cybernetics and the Viable System Model (VSM).

Fieldwork designs are proposed for the empirical development of an analytical framework and its use to assist study of the development of safety management in organisations.

DEDICATION

In memory of my father, Maurice Stanley Howlett, to whom these ideas would have appealed.

ACKNOWLEDGEMENTS

Professor Richard Booth of Aston University, without whom this work would not have been started let alone finished.

The staff of the Safety Systems Development Centre of the Idaho National Engineering laboratory and of the Technical Research Advice Center, Scientech, Inc., Idaho, for their generous provision of MORT information, training and unstinting hospitality.

Dr Robert J Nertney, whose energy, imagination and delight in ideas are an inspiration to all who know him. Mr William G. Johnson and Professor W. Ross Ashby, who inspire still through the greatness of their work.

Miranda, my family and my friends. Thank you for your support, belief and above all patience.

List of Contents

List of Tables and Figures

Preface

This thesis turns out not to be the thesis I originally thought I would produce although the subject and title are the same. I had thought that I would produce a piece of research consisting of a discussion of the relevant topics and previous work, an explanation of the methods to be used for field work, an account of its conduct, a description and discussion of its findings and a conclusion summarising what had been found and the need for further work. At least, this is the pattern that BSc and MSc programmes had given me to expect.

What I had in mind was to study how management systems develop in new worksites/plants. The reason for this interest was a hypothesis that the *development* of a system to manage safety might be more predictive of long-term safety performance than the snap-shots revealed by periodic audit during the operational phase of an industrial system. There is a good deal written on this subject but more about how such systems *ought* to develop and mostly concerned with the classical life-cycle of development (that is, conceptual, design, commissioning, operational and disposal phases). However, despite intensive efforts to identify case histories, there do not appear to be documented accounts of how systems actually developed and, in particular, how a safety management system developed within a new system.

Implicit in this is a secure foundation of theory about safety management, in essence, a theoretically supported model of safety management of general application within which to collect and make sense of data obtained through fieldwork. In other words, I had assumed that square one was a secure place from which to launch predominantly empirical research. However, taking the H&S literature as the primary source, it became evident that the solidity of our theoretical knowledge was far from secure.

At the outset of the work I (obviously) had not expected "square one" to be the bounded territory of this thesis because I thought that I already had a model of safety management in the Management Oversight & Risk Tree (MORT) which after a some modifications would admirably serve my purposes. After all what initially attracted me to MORT was that it was big and complicated and so is real-life and also that it claimed to represent the "idealised safety system model" (Knox & Eicher, 1992).

However, little time passed before it became apparent that complex and valuable though MORT is, it is *not* a model of a system as systems are commonly understood to be: dynamic, interactive and purposive. Nor was it at all clear how MORT relates to organisational structures - particularly as represented by the organisational chart. In summary, I was making a mistake that apparently many people had made: looking to MORT as *the answer* and coming away with more questions. The suggestion here is to re-examine the question I was originally asking: "is safety performance predictable not merely from the particular configuration of hardware and activities but from the process by which this configuration was designed"? In other words can we predict long-term safety performance by examining the features of the system that designed the safety system and, for that matter, is this a better method?

On the face of it, the question is a good one but it is rather full of assumptions. Firstly, the process of designing a safety management system is not likely to be a one-off effort but a constant process of refinement and adaptive change, if only at the level of personnel. Second, we have grown perhaps too familiar with the term safety management system... *What is* a safety management system? Given a particular organisation, can one distinguish the elements that are from the elements which aren't part of the management of safety? If the identity of these elements is entailed by their relationship to safety, are we in a position to say, in our current state of knowledge, what these relationships are or even (dare I say it) what safety *is*?

It is possible that I could have quietly stopped thinking quite those thoughts and simply chosen whichever definitions suited my purpose. However, given that I wished to study the development of safety management systems and thereby understand better how safety performance is determined; the need to define the safety management system in general yet rigorous terms is irresistible. In the absence of given definitions of *safety management systems* which "suited my purpose" I have had to construct my own.

To summarise: In order to begin the process of studying the development of safety management systems I needed a framework within which to collect and analyse the data. To develop this framework I first needed a model of what it was I aimed to study: safety management systems. However, implicit in this is study of the systems *which gives rise to such systems* and so my model would need to incorporate these antecedents also. This begs a further question: where does the

"safety management system" begin and the system that designed it end?, and in general, can I define a safety management system with sufficient clarity as to identify elements which belong to it from those which do not?

The same questions are raised, implicitly, by Rasmussen:

> "An obvious research topic would be to explore the possibility of integration of the organisation needed for normal operation and for MORT-type analysis of operating experiences. It would be interesting to see, whether such an integration could lead to an operational simplification when implemented in a particular system, supported by the analyses and functions already in action for normal operation and planning"
>
> Rasmussen & Batstone, 1989, paragraph 88

In a similar vein and somewhat more transparently, the third report of the ASCNI Human Factors Study Group (HSC, 1993) asserts that: "the best safety standards can arguably only be achieved by a programme which has a scope well beyond the traditional pattern of safety management functions" a conclusion they reach from, inter alia, the USNRC research which implicates communication and organisational learning as the two prime indicators of safety performance.

The connection between Rasmussen's suggestion and the argument proposed by the ACSNI group is the embedding and synergy between the wider functions of "normal operation" and those functions conventionally associated with safety management. Thus the question is what are the boundaries of, and what are the connections between, the management of safety and the management of the enterprise as a whole?

I do not claim to have answered this question, indeed I do not believe there is a final answer to it. However, I believe that exploring the issues it embraces has radically improved the objective of "square one": a theoretical framework with which to study safety management and which I look forward to using in practice.

Thus, unexpectedly, I have stayed in "square one" and this thesis contains the results of exploring its varied terrain.

1 Introduction

Summary

The aim of this thesis is to explore the structure of safety management systems and the issues associated therewith. To achieve this it is necessary to first consider what health and safety performance is in more exact terms. The more exact definitions then allow H&S performance to be rendered as special instances of more general phenomena and this permits the evaluation of control methodologies not generally associated with H&S which may provide effective strategies in the management of risk.

It is argued that only a wholly integrated approach to the management H&S can be effective. In other words, H&S is an inseparable aspect of general managerial control just as hazards are an inseparable part of the industrial process that is managed. Within this perspective, the general issue of control in systems is examined. The discipline of cybernetics (established by Wiener, 1948) is devoted to the aim of identifying the mechanisms by which control is asserted in systems generally. Hence cybernetics and other "systemic" approaches are reviewed for their applicability to health and safety management.

In contrast to the foregoing, but central to the issue of H&S, is consideration of ethics and it will shown in chapter 3 why ethical considerations are vital: in defining adverse events within organisations; to the concept of *severity* or *magnitude* of consequences in risk; and in defining social and organisational policy. The first two matters are considered in the light of the assertion that "where a genuine moral choice is involved no accumulation of non-moral considerations is of decisive importance" (Weldon, 1962, page 231). The impact of ethics upon decision-making is discussed as a direct input in developing a general framework for H&S management.

Having developed the two principle topics of control and ethics, the Management Oversight and Risk Tree (MORT) is evaluated (readers who are unfamiliar with MORT are directed to appendix 1 which describes the MORT system). This evaluation considers how the virtues of MORT (ie, its language, logical rigour and internal consistency) can be retained whilst introducing dimensions which allow organisational control to be modelled. At the moment, MORT provides a static "linear" model of isolated barriers between energy and targets (people and

objects) - a vertical and very narrow "slice" through an organisation described in purely functional terms. MORT analysts often aggregate a number of such barriers together when performing root cause analysis (of upstream or "management system" factors) and the fact that this can be done at all relies upon the analysts supplying this extra dimensionality. The reason for this is that analysts must per force meet the axiom "Every good regulator of a system must be a model of that system" (Conant & Ashby, 1970).

Whilst the current approaches to modelling in H&S appear to be limited, other literatures contain radical approaches which may be of use. The general topic of cybernetic modelling is considered and the contribution of Beer (1966; 1979; 1981; 1985) is given special attention in chapter 4.

The themes of control, ethics and modelling (including MORT) are brought together in chapter 5 where a general view H&S modelling is developed. In chapter 6, research designs for the empirical development of models are presented. However, it is noted that whilst empirical validity is desirable, it is not the sole criterion. It is perhaps more reasonable to judge models for their *usefulness* rather than their absolute truth: the map of the London Underground is a model of outstanding usefulness to travellers on the "Tube" but is a remarkably poor representation of the capital's geography.

1.1 Discussion

I have written the remainder of this text in the first person. This is first to assist the reader's apprehension of sometimes abstruse ideas without the additional burden of third-person sentence construction; and, second, to permit easier discrimination of my ideas from those of others to whose work I refer.

In the Preface I spoke of *exploring* because I have made various excursions from the safety literature seeking alternative models and unifying concepts with which to inform this work. The subject of safety is by its nature, multidisciplinary and has been throughout industrial history. However, more recently, the safety literature has increased its theoretical base in response to the growing perception of management involvement in determining H&S performance. Previously, the safety literature was pre-eminently concerned with the technological control of hazards and individual behaviour (such as accident proneness).

I am not aware of any marked transition from the technological to the social-scientific although there are obvious milestones such as the report by the Robens Committee "Safety and Health at Work" (1972). Instead there has been a gradual change in how H&S is considered theoretically and approached in industry. Ordinarily, this might be labelled as a "paradigm shift" (Kuhn, 1962) but this would presuppose a more coherent body of literature and approaches to research to shift from/to than I consider to exist. Instead, what appears to have changed is the legitimacy of management as an arena for study in relation to occupational safety.

The results of the managerial orientation are difficult to gauge but it is early days. I suggest, however, that the progress towards higher standards of safety performance that we hope to gain through this are inhibited by a lack of consensus on what counts as fact in relation to H&S in industrial systems. In previous times when the H&S literature and practice were dominated by a technological/human-error focus, it was the phenomena themselves that provided a tacit definition to the boundaries of the subject and, with it, an implicit paradigm in which to work. However, the breakthrough into the socio-technical view requires better resolution of the issues central to H&S to act as a frame of reference. Without this, it is difficult to determine, firstly, what we know and, secondly, to discriminate the assumptions qualifying the evidence that underpins this knowledge. My reading of the literature is that there is a great deal of mileage in what we know but inconsistency of terms, definitions and concepts belies this. For these reasons I have sought for areas of convergence about the epistemology of systems and of safety within systems; this topic is picked up in chapter 3.

The thesis develops three themes; the first of these is the view of H&S as a systemic control problem. Ultimately, any system, no matter how excellent its control "architecture" must specify the parameters of acceptable performance, that is, what standards pertain to which aspects of behaviour. Here, the logic of system control requires closure from without. What is "good" cannot be specified within a system and requires reference to the system that contains it and thereby constrains it. Further, this "higher" or *meta*system would in turn look above itself to its metasystem. This conundrum, which centres on the question of system closure, has been debated for centuries under various guises: consciousness (eg. Churchland, 1988; Hofstader, 1985) in mathematics (eg Gödel, 1992[1] and Cantor,

[1] Gödel - Formally undecidable propositions

1915[2]) and in mainstream philosophy (eg. Hegel, 1874[3]). In general, the notion is given by the aphorism *quis custodiet ipso custodes* — who guards the guardians?

In H&S, until recently, the issue of closure was imperfectly dealt with by prescriptive H&S legislation: The law stated what should be and the state provided the closure. However, whilst the state succeeds in this logically, the law fails in practical terms. One cannot legislate for all possibilities and, even if it were possible to address all eventualities, communicating this giant to employers would not be feasible[4]. Furthermore, while there have been notable successes in H&S law (eg., The Power Presses Regulations 1965), in seeking *general* application, prescriptive law is not likely to provide the optimum solution in any *particular* workplace. Operated in this way, the law constrains the freedom of management to arrive at better solutions of the same issues (because they have local knowledge) and disallows their judgement of more cost-effective solutions to the variety of hazards in their firm. The state has, as it were, assumed the role of management, but aside from the question of its technical competence, how can the state enforce its requirements? It seems axiomatic that the state can not dispose the resources that an organisation routinely does to the same problem - systemic control. The Robens Committee recognised this:

> "The most fundamental conclusion to which our investigations have led us is this. There are severe practical limits on the extent to which progressively better standards of safety and health at work can be brought about through negative regulation by external agencies. We need a more effectively self-regulating system."
>
> Robens Committee, 1972 (paragraph 41)

The success of the Health and Safety at Work Act (1974) and the establishment of the HSC and HSE in delivering an effective self-regulating system has been, by the most optimistic reading, partial. Some of the possible explanations are examined in Chapters 3, 4 and 5.

A further quotation from the Robens report (Ibid.) provides the connection to the second theme of this thesis. Robens gives the following quotation from Webb (1910) that he recommends as distilling the problems associated with "piecemeal legislation".

---

[2] Cantor's paradox of non-exhaustive sets

[3] Hegel - the theory of internal relations

[4] Although, advances in IT (especially in regard to the accessibility and affordability of the technology) suggest that the communication aspect is fast becoming feasible.

> "This century of experiment in factory legislation affords a typical example of English practical empiricism. We began with no abstract theory of social justice or the rights of man. We seem always to have been incapable even of taking a general view of the subject we were legislating upon. Each successive statute aimed at remedying a single ascertained evil. It was in vain that objectors urged that other evils, no more defensible, existed in other trades or classes, or with persons of ages other than those to which the particular Bill applied. Neither logic nor consistency, neither the over-nice consideration of even-handed justice nor the quixotic appeal of general humanitarianism, was permitted to stand in the way of a practical remedy for a proved wrong".

For my part, it is Webb's references to a theory of social justice and the "rights of man" that are compelling. It is hard to pin down the possible reasons, but amongst workers in the area of health and safety there is a tendency to fight shy of moral arguments[5]. Johnson points out that "moral concepts have slight effect on changing practical behaviour" (Johnson, 1980, page 146) and yet it is upon the assertion of rights and reciprocal duties that employee and public H&S ultimately rests. In chapter 3, I discuss the matter of ethics in relation to H&S and, in chapter 4, look at the logical argument for the role of ethics in generic control systems.

These two themes, ethics and control, are finally reconciled in Chapter 5 in which the third theme (modelling of safety management) is developed and a new model presented. Considered in total the thesis aims to produce a conceptual framework within which to study organisations from the perspective of H&S.

[5] Perhaps in seeking to speak the language of cost-obsessed managers, the absence of words like "rights", "well-being", "justice" in the non-moral vocabulary has rendered these words taboo. Perhaps H&S has succumbed to the post-modern world-view in which nothing can be said to possess moral authority.

1.1.1 Connection with previous work undertaken by the author

The orientation of the thesis originated in 1991 with work I undertook as an MSc (Work design and Ergonomics) research project involving BNFL'S THORP project. This was a collaboration between Birmingham University and BNFL.

The rationale behind the work was that risk assessment methods tend to:

1. suffer from omissions in hazard identification;
2. contain data, such as component failure frequencies, which may be falsified in practice, and that
3. even if the probabilistic risk analysis (PRA) was exhaustive in its treatment of failure modes and equipped with valid data it would still not guarantee the safe operation of the plant. This is because changes occurring in or impinging upon the organisation may violate the PRA assumptions upon which safe operation was predicated.

Strongly implicated in the shortfall between what PRA delivers and the assurance of safe operation, is the management of safety. Can we do for management what PRA does for operations, that is, produce a predictive method that indicates the riskiness of the management of the operations?

A predominant technique in risk assessment is Fault Tree Analysis (FTA) in which the results of techniques such as HAZOP, Task Analysis, SLIM and FMEA can be arranged to estimate the likelihood of specified top-events. The resulting fault trees can then be used with other techniques to find the most cost effective routes to reducing the top-event frequency. FTA forms a substantial component in the production of safety cases and the resolution of design options in high risk industries such as those coming within the definitions of the CIMAH Regulations.

MORT can be considered as an FTA of management where the top event is stated in the following general terms "Injuries, damage, other costs. Performance lost or degraded. Program public impact. Future undesired events." (MORT User's Manual, Knox & Eicher, 1992). On the face of it, this is the "predictive method which indicates the riskiness of the management of the operations".

1.1.2 MORT as the management analogue of operational fault trees

The argument of the previous paragraph is not unreasonable and provided the underpinnings for the MSc work. The objective of this project was twofold: (1) to develop an audit tool with which to assess the operational readiness of the safety system at a pre-operational stage of development and (2) to pilot this tool at THORP. Whilst the project was academically successful and yielded useful information about the safety effort at THORP, I was left with severe misgivings about the tool I had developed. At the time, these concerns concentrated upon the hit-and-miss nature of the audit-tool: the hits testified to the fact that MORT asks the "right" questions but the particular protocol developed was something of a blunderbuss. I loaded the blunderbuss with questions derived from MORT, used the organisational chart to locate knowledgeable individuals and pulled the trigger. Unfortunately, this process was conducted in the dark: I had inadequate appreciation of the structure I was seeking to explore using the MORT structure as a map and MORT questions as a flashlight. I ought to add that the effect of the MORT blunderbuss on individuals was enlightening: interviewees were very impressed by the interview, many volunteering the view that the process had given them further insight and areas to explore in their work.

1.1.3 The present

Notwithstanding my misgivings about the method used in the MSc project, the idea of examining the "longer-lived" influences on safety and, in particular how management systems are designed and established, still seemed valid. In taking this idea forward to the present work, this basic notion required better resolution. Further, the idea and the acquisition of empirical support are necessarily connected through a grounding in theory, in this case, a theoretical model of H&S management.

With hindsight I can see that, amongst the many shortcomings of my approach at THORP, the most fundamental was MORT as a map of organisational safety - it really did not seem to perform at all well. This left me with two non-exclusive hypotheses: (1) my understanding of MORT was inadequate; (2) MORT is inadequate for this use. However, I had exhausted the available literature and expertise and with them my ability to resolve these problems. In essence, I could not judge the matter because I could not gauge my own ignorance. Moreover, I did not have the conceptual structure or language with which to phrase the

problem by correspondence. Enlightenment required the richness of face-to-face communication with the developers of MORT (sadly with the exception of the late Bill Johnson).

To this end, arrangements were made to meet with the staff of SSDC (the Safety System Development Centre, a branch of EG&G Idaho Inc.) who supplied (until 1995) specialist safety expertise, including MORT approaches, to the US Dept of Energy. These meetings took place in March 1992 and were prefaced by two weeks of MORT training in San Francisco[6], accompanied by considerable discussion with fellow trainees (plant managers, safety specialists, NRC inspectors and engineers) which gave me a much improved confidence with the MORT language.

A week of conversation with Dr Nertney of SSDC greatly improved my conceptual grasp and it emerged that we had independently identified the same concern: MORT does not model the dynamics of system processes (either equilibrial or developmental). In MORT terms the problem can be stated: how do products of "M" branch processes arrive in the "S" Branch as built system configuration? Further, what time dependencies operate in this process (the question of time lags being illustrated nicely in Tom Ryan's NRC work (Ryan, 1991), and were a dominant theme in my discussions with Dr Ryan who by this time was employed at EG&G's Idaho National Engineering Laboratory).

Returning to the "two hypotheses" of a few paragraphs ago, it appeared that I did understand MORT and it also seemed true that MORT has important shortcomings as a model. However, I was still unable to specify exactly what these were, although there was some moral support from the MORT team at EG&G who were scratching their heads about the same issues.

1.1.4 Specifying the limitations of MORT

Part of my confusion stemmed from a rather unscientific impression: Johnson really seemed to have understood the management of safety, what he and his team wrote during the same period as the report of the Robens committee was published (1970-72) remains some of the best literature in the field - yet the distillation of his work (MORT) does not seem to address the interactive dynamism

[6] For which I am greatly indebted to the generosity of the SSDC staff and Dr Robert Nertney in particular.

of real systems. This and the fact that MORT has survived a quarter century of wide use belies the idea that it suffers from serious shortcomings. These ruminations suggested that perhaps I was trying to make MORT do things it had not been designed to do. So what *was* MORT designed to do?

1.1.5 The uses of MORT

From my reading of Johnson's developmental work (as reported in SAN 821-2, 1973) it appears that in the course of developing the "superlative safety program" MORT was initially invented as a method of investigating accidents with scientific rigour... something that had not been done before to the standards required by the nuclear industry. Moreover, MORT investigations produced an acceptance by senior management of their role in determining safety that was quite extraordinary. Consider the following quotation from SAN 821-2 (Ibid.):

> "Interestingly, and helpfully, MORT quickly displayed a capacity to gain management confidence despite the fact that MORT puts the accidents on management's doorstep. Two senior vice-presidents of a large research and development organisation said such things as: "Interesting, valuable, very provocative, and certainly opened my eyes to a lot of things," and "the first scholarly, in-depth approach to safety I've ever heard in industry or research."
>
> Johnson, 1973 (SAN 821-2) page 136

I suspect that MORT, a tool within a kit of tools, was promoted to flagship status because of its evident acceptability to senior managements. This promotion was not solely, or even mainly, the will of its creators but impelled by the enthusiasm of the investigators who used the first generation of MORT and later by the appraisal staff who recognised the *Success* Tree in the investigators' *Fault* Tree (Nertney, 1994).

However, MORT while a creation of a philosophy, expresses this philosophy in a limited way. As stated, the MORT diagram was designed to investigate accidents and not to model organisation for safety (as a device to assist the modelling of adverse H&S events in an organisation rather than the model of the organisation). The managerial tools, approaches and risk assessment techniques that comprise the MORT "tool kit" are also rendered coherent within the guiding philosophy of Johnson et al of which the MORT Chart is the most representative artefact.

From the perspective of Johnson and his co-workers on the project (Jack Clark, Jack Ford and Bob Nertney) I believe the acceptability of MORT qualified it as the vehicle by which to transport the whole system safety philosophy that they had conceived: the means to an end but not the end in itself. An unexpected outcome of MORT's promotion is the revelation of its limitations: if the management of the system investigated share the philosophy from which MORT arises, it serves as an instrument in the practical enactment of that philosophy. Conversely, if the management hold dissonant beliefs relative to this philosophy, MORT answers questions no-one is asking in a language unlikely to be understood about a world which the management do not see as their own.

In other words, MORT does not itself contain requisite capacity to convey meaning. The extra information has to be supplied from without and that presupposes requisite knowledge of the organisation in question. A graphical analogy is given in Figure 1.1. Looking first at the hexagonal plan of the figure, a large family of three-dimensional figures are consistent with it but one needs information from a third dimension to determine which member of the family is being depicted on any given occasion. Relating this back to MORT, if one is concerned with organisations in general, then this *extra* information is not available and MORT sits dead on the page: it lacks vital dimensions in which to map organisations generally or in a particular instance.

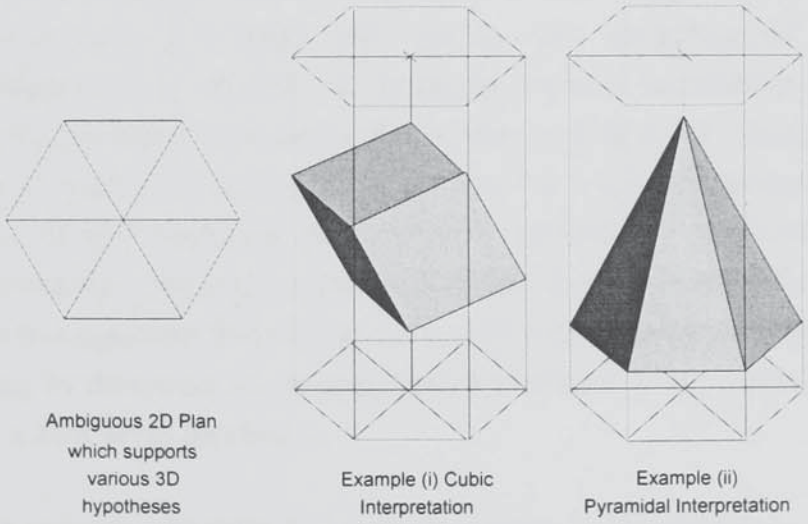

Figure 1.1 Illustration of ambiguity requiring outside reference.

In my opinion Johnson et al did not develop MORT to capture these "extra dimensions" simply because the need to model them was not perceived. However,

the acceptance of the MORT tree was, in my view, due to MORT users automatically *augmenting* MORT with their own mental models of their organisation (just as you or I might perceive the 3-*D* cube in the 2-*D* pattern of lines, because the drawing is augmented by our general knowledge of rectilinear objects).

1.1.6 Expanding the dimensions "flattened" in the MORT representation

Throughout this and later discussions of MORT, It is useful to remember that the MORT literature *does* attempt to describe vital mechanisms within a dynamic safety management system. However, it falls short of actually describing a model of (i) how the management of safety is embedded in a dynamic industrial system or (ii) how the various safety functions interact. These two descriptions are, in my view, inextricable and it is the attempt to treat them separately that has dogged both the theoretical and practical progress of safety. This is not a failing peculiar to workers in the safety field but an inheritance of a reductionist paradigm on the one-hand and the social trend towards specialisation on the other. It is perhaps yet more apparent in safety owing to the *perceived* disparity between the aims of industry and the aims of the safety profession, more of which in Chapter 3.

As mentioned, the MORT literature does not capture the place of safety functions within the "whole" but it does faithfully acknowledge the purposes of the organisation generally and regards safety goals as congruous with them. The upshot of this is that MORT and its literature provides a logic which is compatible within organisations but does not specify the mapping of MORT onto organisational systems. As stated, the mapping requires extra dimensions. Part of this dimensionality reflects an early realisation in this work - that some of this missing dimensionality is recursive in character (c.f. the discussion of closure at page 16, ante). MORT recognises the recursive operation in organisations, but not explicitly. For example, taking the instance of the *controller's controller*, in the MORT logic tree one has operator behaviour (SD5 - b3) supported by a supervisor (SD1 - SD5) according to direction via management services (MA2 - a6) as directed by policy (MA1): a four level hierarchy.

It is clear that any use of MORT for mapping H&S management in organisations must untangle the recursive dimensions currently "flattened" in the logic tree. Furthermore, as this dimensionality was not part of the MORT concept at inception,

it is quite likely that (i) that closure to the hierarchy has not been specified[7] and (ii) the logic of recursion is not consistently contained in the current logic.

Concerning (ii), the matter of closure, the *controller's controller* example illustrates that the addition of new "higher" recursive levels can carry on infinitely. Strictly speaking, it is incorrect to close a recursive system arbitrarily. However, management has little time for notions of the infinite and a practical answer must be found to this problem.

In the early stages of this work, I found that pondering the control hierarchy would lead, with frustrating inevitability, to the inclusion of the state as part of the system. The frustration was due to my desire to find a way of making closure possible within the organisational system - otherwise the method of closure takes us back to the problem previously discussed: the prescription of conditions backed-up by the unarguable force of law. This is analogous to a parent saying to a child do "*x*" in relation to "*y*" because "***I say so***". Given that organisations thus dictated to may both know better what to do about "*y*" and also have "*a, b & c*" problems whose solution may be more beneficial to its workers (say) but for which the state has no regard. Constraining freedom in this way is both oppressive to management and unlikely to result in satisfactory H&S performance overall.

I believe it was this situation that Webb (Ibid.) saw as the outcome of legislating in the absence of "a theory of social justice" or, less grandly, "a general view of the subject". This subject is pursued in Chapter 3. For the moment I hope it is clear that to model the management of safety it is necessary to recapture or introduce the logical hierarchies functioning within organisations and, at each level of the hierarchy, to model the functions required to manage risks to H&S. The list of requirements is a matter handled by MORT as it stands but it is the modelling of functions interacting in a logical hierarchy that can provide the model with descriptive and, hopefully, diagnostic power.

[7] In the example given above, "Policy" is the highest expression of control. In MORT "Policy" is a base-event [MA1] sharing a logic gate with 'Implementation' [MA2] and 'Risk Assessment System' [MA3]. It appears that the logical operation requires a falsification of the facts because it must be that Policy directs risk assessment which in-turn provides information at a lower level of policy. In a similar fashion, the argument can be operated backwards - Initial policy *must* be determined a-priori by a set of factors not available in the MORT tree: incompleteness.

1.1.7 Cohesiveness in the recursive MORT structure

Given that it is possible to tease out the recursions in the MORT tree, the question arises as to how to give the structure cohesion; in other words, how to tie the various levels of recursion together. In his work on recursive organisational structures, Beer (1979) identifies *planning* as the "recursive glue". Contrary to the conventional view of planning as a finite process with finished plans as a product, Beer's view of planning is of a constant process carried out at all levels of an organisation in the light of improving estimates of what the future will be. Hence plans are made and discarded on a moment-by-moment basis as part of the normal operation of the firm. I do not think that this is as radical as a first reading may suggest: the submission by divisions of five-year plans to their parent company is an argument for the allocation of resources which may form the basis for accountability but not the letter-by-letter contract of what will actually happen. The operation of authority in an organisation is essentially to exercise veto within prescribed limits and as Jaques (1990) points out, each managerial layer within an hierarchical organisation is characterised by a time-horizon within which to abort the extant plan in favour of one which better approximates the needs of the future at the limit of the time-horizon.

This iterative planning process is in keeping with the general views of Johnson et al and their "dynamic" model of safety assurance. MORT contains particular references to Hazard Analysis Process (HAP) triggers[8] (ie, monitoring in respect of detected unplanned change), Change Analysis[9] (as the precursor of planned change), Change Review[10] (ie, hardware design and as-built configuration control), Task Safety Analysis[11] (ie, semi-formal task-based risk assessment) as well as other planning and review functions. These functions all speak of planning as an iterative and dynamic process but the "flatness" of MORT obscures the correspondence between the ideas of Johnson et al and the views of Beer.

Another vantage point on this issue is provided in the USNRC NUREG studies of organisations which identified four principal predictive indicators of safety performance (USNRC, 1991). The first two (in rank order of correlation with measures of safety) are *Communication* and *Organisational learning*. These two factors can be considered in isolation, but certainly appear interrelated in the

[8] MORT Event SD1-a4
[9] MORT Event MB3-a1-c4. Further described in Bullock, M.G. (1981)
[10] MORT Event MB3-a2-c45
[11] MORT Event SD5-c9 Further described in Nertney, R.J. (1987a)

iterative planning process: Beer's recursive glue. The fourth-ranked of the USNRC indicators, external factors such as the *Impact of the parent company and regulatory bodies*, can also be seen as higher-order organisational recursions which operate upon the lower-order system-in-focus insofar as they exert control (or regulatory) actions. Again, taken with the first two indicators, this indicator fits into the emerging picture of hierarchical control and information exchange between the levels of the hierarchy. The third-ranked indicator - *Organisational Focus* (that is, focus on safety) is interesting insofar as it is ranked beneath the more general organisational factors in magnitude of correlation to safety outcomes.

As noted in the third report of the ACSNI Working party (HSC, 1993), the factors identified in the NUREG studies (Organisational Focus aside) "do not concern safety directly". The authors of ACSNI III suggest that "it follows that to make managers manage safety better it is necessary to make them more effective managers". I would emphasise that it is the *organisational context* which largely enfranchises the efforts of people, which allows managers to "add value" to the efforts of subordinates and which provides the value system within which to make decisions. As Senge (1992) suggests: "When placed in the same system, people, however different, tend to produce similar results" (page 42).

An important factor in the enfranchisement of individuals within organisations is the flow of information in which relation Nertney writes

> "We [the US DOE] have perhaps neglected the explicit design of communication to the extent that many information systems are in gridlock".
>
> Nertney (1993)

Although, the design of computer-based information systems has advanced considerably in terms of speed, usability and connectivity, I believe that there is more to the problem than is likely to be solved by the provision of databases containing the what/why/who/when/how of safety in an organisation and label it a "Management Information System" or "Decision Support Software". The fact is there is ***too much*** data and, moreover, even if valid abstractions of the data could be produced in manageable form how, ultimately, could the chief executive produce the appropriate complex of control actions required to enact a particular decision?

As Nertney suggests, information and communication systems need to be *designed* and I suggest that this design should be undertaken in recognition of Ashby's law of requisite variety (Ashby, 1956). If left entirely to self-organisation, Ashby's law will assert itself at the cost of whatever index one might choose to assign value to (eg. efficiency, quality, health & safety etc.). In essence, the law of requisite variety is a more fundamental rendering of the quotation offered previously at page 14 (ante): *Every good regulator of a system must be a model of that system.* Every distinct consequential state the system is capable of taking (the set of possible states = the *variety* of the system) must have a matching control action - hence *requisite* variety of the controller. This is true whether the controller is the operator of a machine, a supervisor or the Chief Executive. Information systems which are not the subject of design to RV, leave controllers with little option but to simply discard huge amounts of systemic variety and handle merely the residual that *seems* most pressive and significant. In passing, the third of the NUREG indicators "organisational focus" shows some degree of independence from the two main indicators "communication" and "organisational learning". This lends support to the view that even in organisations characterised by a positive concern for safety, poor information handling is associated with poor safety performance.

In summary, then, it seems that there is a reasonable correspondence between the views of Johnson and Beer and that there is an overlap between the required features of the safety management model and the organisational factors found empirically associated with safety performance. These issues appear united within the concept of organisational cohesion which is characterised by the factors of (a) communication and (b) organisational learning. Cohesion, Beer argues, is achieved through the iterative process of planning and this acts as the glue which binds the various layers of the organisation together. Further, this structure must be *designed* to meet Ashby's law of requisite variety which requires the variety of the lower order system to be matched by the variety of the higher order system if control is to be maintained. The main elements involved in achieving this are:

(i) the competence of the model of the lower-order system held in the higher order system;
(ii) the design of the communication interfaces between the layers in the hierarchy[12].

[12] For example, the driver of a car (higher and lower order respectively) may have a highly developed mental model of the car, perhaps he is an "AA Man". However, if the oil-pressure warning light is inoperative, it is unlikely that he will be able to act in time to prevent damage to the engine by the time other non-designed indications declare there to be a problem (the engine seizes,

1.2 General management and management of H&S performance

The discussion so far has seen that MORT requires a fundamental reconsideration if it is to perform as a research framework or diagnostic model. In the main, these changes are those which a hierarchical view of control requires[13]. Communication and information are matters that MORT does not appear to handle particularly well although useful ideas are contained there. Again this is perhaps inevitable given the "flattened" representation and the previous section has argued the importance of these matters to building a model of the dimensionality required. Further, the matter of designing the information channels both within and between layers of the hierarchy so as to allow requisite variety in the management system are matters which need to be imported into MORT. A point to note here is that the management system is seen as a hierarchy which would certainly include operative staff and even automated equipment. What is paramount is the control of system behaviour rather than conventions as to the "who" and "what" that are the medium through which control is achieved.

A cursory reading of the MORT chart clearly removes any doubt that its subject is the control of H&S phenomena but, in principle, the functions represented are compatible with the general control of unfavourable states that a given system may be capable of expressing. As mentioned previously, Johnson (1980) was strongly motivated by the aim of achieving a congruence of safety aims and methods with production aims and methods and so this apparent generality of application is intentional. This brings the discussion to a central issue: the relationship between the control of H&S phenomena and the control of non-H&S phenomena. Are these commensurate or does H&S require different treatment organisationally? The purpose of this question is to examine whether there are requirements for the control of H&S phenomena which cannot reasonably be accommodated within conventional management approaches to the control of non-H&S phenomena. Similarly, as the USNRC research previously mentioned suggests, is it the case that H&S phenomena are more *testing* of management control in general than other phenomena expressed in organisations?

---

temperature warnings etc). On the other-hand, a naive driver may well notice but ignore an operating oil-pressure warning. the result is the same in both cases.

[13] I think it is worth repeating that 'hierarchy' is used to denote the logical structure of control rather than the institutional structures represented in the familiar organisational charts. It may be that, on occasion, the interpersonal relationships expressed in such charts do correspond to the functional 'reality' of control but, even where this is the case, the chart is particular to the organisation and a gross simplification of the 'true' pattern of interconnectedness that prevails in a community of people.

In Chapter 2, I examine what H&S performance is in terms of the phenomena which characterise this class of organisational behaviour. At this level of analysis there do not seem to be *striking* differences between H&S and other organisational behaviours but more differences of degree. From the discussion in chapter 2 what characterises H&S phenomena is (i) the level of probability and (ii) the subjective class of adversity to which they are allocated. Because of this, I would tend to favour the hypothesis that "H&S phenomena are more *testing* of management control in general".

1.2.1 Conventional management control and low probability states

My own experience of organisations is that, regardless of apparent sophistication, they mostly muddle through and that their deficiencies, taken as whole systems, are compensated by the efforts and ingenuity of individuals. Using the language of variety previously introduced; RV is preserved in hierarchical control systems by individuals "pumping-in" the marginal variety missing in extant organisational arrangements. This view is illustrated by the effects of "working-to-rule" disputes and by the inadequacy of formal procedures generally.

However, whilst the "pumping-in" of extra variety may be successful for high probability states in such systems or states which are easily predicted, low probability states (or those whose potential is not obvious) are characterised by poor reliability of decision-making. There are numerous reasons for this, amongst the most evident are limitations imposed by the means by which people acquire a mental model of the local system under their control (learning from experience) and the biases operating in the controller perception of probabilities (eg. the use of heuristics in the special meaning of Tversky & Kahneman, 1974). Thus, the control of low probability states requires greater institutional support of individual efforts and, therefore, a greater sophistication than "muddling through". Because of the arguments already sketched out with regard to requisite variety, approaching this problem by institutionalising a "risk-assessment think-tank" to provide reliable solutions to these low-probability issues (other than where this is appropriate to pre-operational design) is not merely impractical but impossible (too much information, amongst other reasons). Instead, what is required is the organisation's sponsorship of individuals' understanding of the local system under their control and the means whereby unfavourable states recognised through this enhancement can be communicated with the appropriate part of the

organisation for resolution[14]. Insofar as H&S phenomena are, in general, low-probability this argument applies as much to them as to other unfavourable low-probability states.

1.2.2 Management and hierarchy of values.

Whereas certain classes of *value* are conventionally accepted within organisations (such as productivity, profitability, legal compliance etc.) others are more equivocal. H&S is in many organisations part of the latter category, especially where the potential consequences are perceived to be small in measurable cash terms. Whereas organisations implicitly rely upon the individual "pumping-in" the required extra variety they do not, in general, permit individuals to work to private agenda in the sense referred to. Organisations confer upon themselves the role of arbiter when it comes to determining what is "good" and what is "bad". How this is achieved is considered in Chapter 3 but the position of organisation-to-employee is analogous to the position of State-to-citizen in the definition of good. However, whereas in States like the UK and the US, citizens may exercise their "Protestant conscience" and still remain citizens (in the extreme case, behind bars) the corresponding act by employees may lead to their expulsion from the organisation. For this and other reasons, employees generally leave their personal ethics at home. Further, as Shaw and Barry (1989) point out, moral accountability is so diffused in modern corporations (or, in the case of smaller companies, diffused by "industry practice") as to bring in to question their competence in making moral decisions.

This argument suggests that H&S goals have a double handicap: first that conventional managerial control relies, as stated above, upon individuals to "pump-in" the extra variety to achieve control in general and that this approach is unreliable for low probability phenomena and, second, that organisations tend to disable individual moral responsibility but fail to provide clear (and therefore accessible to external critique) priorities about values[15].

Lastly, the point about external accessibility is not an adjunct to this debate but is quite central to it. What is at issue here is merely the next recursion up in the control hierarchy which happens to traverse the conventionally perceived

---

[14] See the argument presented in chapter 4 (page 112, ante) concerning double-loop learning.
[15] I say *priorities* because values can presumably be ranked (least, more, most) in an ordinal arrangement. However, the rank assigned to a given value may well be unstable and governed by situational forces.

boundaries of the organisation. The logic of the situation argued for the interior of the organisational boundary is the same when the boundaries of the H&S "system" are redrawn to include supra-organisational systems including the state and non-governmental organisations such as the CBI.

Sadly, the key amplifier of feedback which operates in regard to the impact of organisations upon citizens (in their dual role as employees) is the Press. I say "sadly" for two reasons. First, the Press is an inconsistent and partial amplifier of information which is to say it often manages to perform as an attenuator by garbling the facts of what it does report (amplifying noise rather than information) and failing to report at all that which is not "news-worthy". Secondly, Robens recognised the value of the press as an amplifier and the conduit by which H&S data could be accessed by the press was built into the Health and Safety at Work Act (1974) at section 79(2). This effectively required H&S information to be made publicly accessible (via company reports) and, although the extent of the information was limited, this section of the Act was never enacted.

There have been various efforts to remedy this state of affairs, most recently, in a bill proposed by Jeff Rooker MP which sought to amend the Companies Act 1985. The proposed amendments required the disclosure of (i) details of notices served, (ii) convictions of any employees for offences under H&S and environmental legislation, (iii) common law offences, (iv) details of notifications made by the company under RIDDOR (1985), (v) details, including full costs, of payments made in respect of personal injury compensation. The Bill did not receive a second reading in January 1992 "as it was blocked by the Government" (Rooker, 1994). Obviously this state of affairs also blocks this information from company shareholders who, in the light of (ii) & (v) especially, and if only from a position of self-interest, may choose to apply pressure to the board.

1.2.3 Impact of the previous discussion upon the recursive model

Enhancing control of low probability states is, I consider, achieved through institutionalising the functionality of the MORT chart at appropriate levels of recursion. Effectively, this means augmenting the knowledge of individuals in terms of the variety of their regulatory models (as per the Conant-Ashby maxim *every good regulator of a system must be a model of that system)* and the information available to them, matters which must be sustained at each level ascent of the hierarchy. At each level of the system, the metasystem belonging to it must aim to

support the process of maintaining RV and this is predicated on the design of adequate communication.

However, in the absence of appropriate value attached to H&S performance, this comes to nothing for the same reason that an accredited BS5750[16] system will consistently produce the same shoddy goods if the quality specification is set low. Thus there is the logical requirement for the input of values into the control system. The business of this thesis is not to prescribe a probability/consequence function but to describe how it may be communicated and used as a criterion for the various measures appropriate in technology and hierarchical placement within the organisation. It is easy enough to fall into the subtle trap of confusing description and prescription and perhaps it is in this area that we have to tread warily. This said, I do agree with Weldon's statement that *"where a genuine moral choice is involved no accumulation of non-moral considerations is of decisive importance"* (Ibid.). In other words, whilst I strongly agree with proposals for cost-effectiveness in the allocation of resources, I do not see that cost-benefit arguments should have a determining role in the standards of safety prevailing in an enterprise. The role of CBA should be an informing one (to produce more choices and provide one means of choosing between them).

Whatver models might be developed of the regulatory process, they rely upon an organisation deciding what it wants by way of safety and, insofar as the HSE can dispose RV to their task of regulation, that this is always as good as society wants. This is to say that the matter of safety regulation is closed within organisations by the internal ratification of a probability/consequence function acting as the minimal performance standard open to inspection.

1.3 Industrial scope of modelling

MORT was originally designed to meet the needs of large high-hazard industries and interest in the MORT system has largely been restricted to this sector. However, within the US DOE complex of organisations, MORT is applied to contractor businesses which, whilst serving such organisations are themselves characterised by risks of limited consequence (eg construction contractors). Whilst the MORT literature may specify sophisticated management approaches, the functions thus addressed are logically identical regardless of the level of sophistication with which they are applied. Hence to manage risks there needs to

[16] British Standard BS5750: quality systems

be an identification of the hazards; an estimation of the risks associated; implementation of control measures as appropriate; monitoring to assure the efficacy of these arrangements and for change occuring in the system which may render them ineffective. Whether this takes place in a small farm or in a nuclear chemical plant makes little difference to the logic of these requirements although the rigour and sophistication applied in serving this logic will of course bear little resemblance one to the other.

In keeping with this, I do not see it as necessary to restrict the application of the modelling framework advocated here to the high-hazard industries alone. In practice, however, it is unreasonable to expect a small firm to consider supporting sophisticated methods which larger companies might take for granted. However, if like Robens, we ask "what is wrong with the system?" (Ibid.) the answer "some companies are too small to be safe" is not satisfactory. A better answer might be to use the logic of recursive organisation to regard small businesses as though they were divisions of one large business for the purposes of H&S and this thought is explored in Chapter 6.

2 Concepts & Definitions

"To examine ideas which are in such common currency in one's life that they are seldom if ever reflected upon can be most puzzling"

Soltis, J.F. (1968)

Introduction

As stated in Chapter 1, exploration of the issues under the head "H&S Management" has taken me back to first principals in a number of areas. Also, as noted by other researchers, I have presented the view that progress in H&S has been hampered by the lack of an agreed conceptual framework or even a hotly disputed framework that might at least serve to stimulate pursuance of the former. This chapter considers the need for definitions in H&S and presents argued definitions of the basic terms: Accidents, Risk, Hazard and Safety.

It is perhaps surprising to find that attitudes are passionately held about the apparently neutral matter of definitions. Consider two authoritative sources:

> "There is... rather too much talk about definitions. A definition, strictly speaking, is nothing but an abbreviation in which the user of the term defined may please himself..."
>
> Lord Pollock, 1931[17]

> "The definition of anything is the statement of its essence: what makes it that, and not something else... it is a matter of great importance in science. Things belonging to one genus will be studied together: and the aim of our study will be to discover all the general propositions that can be made about them. But there may be some statements that will apply to everything contained within the genus, others will only be true of a portion"
>
> Joseph, 1916, page 115

In considering this matter, and indeed justifying to myself the need for this chapter, I have used an informal test: *is the absence of a clear, unified, definition of*

[17] Cited in Brazier, 1993. This opinion has much in common with Mill's doctrine of *Nominalism* (Warnock, 1975) which maintains that things called by the same name have only the name in common.

H&S an obstacle to progress in the study or practice in the field?. The answer from the academic standpoint, bearing in mind Joseph's observation, appears to be yes.

From the practical standpoint, the need is not so immediately obvious, but it is precisely the preoccupation with practicalities that Webb (Ibid.) complains of as "a typical example of English practical empiricism". However, for Webb, the preoccupation with "practical remedies" was to the occlusion of a wider vision of social justice. Whilst I do not disagree with him, more pertinent to this debate is that the preoccupation with "practical remedies" has left H&S largely ungrounded conceptually[18]. In the absence of delineated boundaries, which a definition should strive to provide, the connections with other literatures and professions cannot be easily specified and maintained. Similarly, failure to agree what counts as fact about events such as "accidents" does not permit the coherent accumulation of knowledge. For example Hale & Hale (1972) conclude that:

> "Many authors have assumed that accidents form a homogenous group of events with common causes. However, none of the published research has demonstrated satisfactorily that different types of accident have the same pattern of causes... Because people have failed to realise the differences between different kinds of accident, they have generalised the results of particular studies too widely"
>
> Hale & Hale, 1972, page 79

Perhaps matters are more serious in the practical expression of this problem. The effect of a predominantly practical literature has been to allow the partition of H&S from the mainstream of academic work, management training and managerial practice[19]. Of particular detriment is that the separation of the management of safety from management in general fosters precisely the state of mind in management we know to be inimical to safety: the lack of ownership of safety issues.

Can these problems all be for want of so abstract a thing as a definition... I do not honestly believe so. However, I submit the definitions derived later as devices for improving this state of affairs.

I ask the reader to bear in mind that the purpose of this chapter is to define ***concepts*** within H&S. Thus our overriding interest is to agree what these concepts

[18] This observation returns us to another distinction made, more on the basis of convention, than logic: the distinction between the academic and practical approaches to H&S.
[19] H&S is not alone in this, the explosion of specialisms of the last 30 years has provided many similar casualties.

are, how these concepts relate and to analyse the concepts themselves. My interest in the words themselves is secondary: as symbols in common currency, overuse and misuse can lead to blurring of the concepts symbolised. It is hoped that having clarified the concepts, that their interelations can be considered on a rational rather than merely conventional basis.

2.1 The construction of definitions

My approach to deriving definitions has been as open as possible and as rigorous as practical. The practical limitation on defining qualities such as "safety" or phenomena such as accidents, is that the resulting definition should not do unnecessary violence to the ordinary meaning attached to the word. I have found H&S to be a minefield of terms upon which the process of definition operates like a detonator to the explosive pressure of contradiction, conflation and abbreviation within; thus violence has not been altogether avoidable.

As alluded to previously, my intention here is to clarify what subjects should be included under the head of *Safety, Accident* etc., and thereby make explicit the relevance of knowledge that, by convention, has not been routinely brought to bear on H&S management. Additionally, the purpose of definition is to declare the natural relation between different things so they can be considered together rather than separated by convention.

2.1.1 Logical division

This question of *natural* relatedness is very difficult to deal with, especially with regard to the broad range of phenomena called accidents. The process of logical division is driven by a particular principal, or *fundementum divisionis,* which should operate upon a given *genus.* For instance, botany classifies plants (a genus) using their structure and, in particular, their flower structure as the fundementum divisionis. In the case of accidents, one typically sees a process of cross-division (where more than one fundementum divisionis is employed), the result is not helpful when the bases used are the accident consequences themselves. The upshot of this is that accident causes are clearly preferable to consequences. However, proceeding by *cause* results in logical errors, particularly the occurrence of the same event in more than one species (because of multicausality). Now, if neither cause nor consequence provides a logically acceptable scheme of definition we may reasonably suspect the essential basis of the genus is ill-conceived. Thus one has to exercise great care in specifying what is essential to the term *accident* if

logical division is to result in anything that assists clear thinking about the subject rather than propagating the confusion.

For a given definition of genus (ie, a description of the essential characteristics that make the subject distinct from other things) it should be possible to see what makes any particular instance a member of this genus as opposed to any other. Hence any instance of the subject should be contained within the concept of the genus and the attribute which distinguishes the particular instance from other instances. The utility of the resulting *family tree* is not necessarily entailed by the rigour of its derivation. For example, once the essential characteristics of the genus "accident" are decided upon, I might then further divide on the basis of the object damaged, say, things and people. Then, in relation to people, I could continue this by the part of their anatomy that has been damaged (eg. limb, torso, neck and head) and so on for each of these parts. Now, there may be instances where this scheme proves useful, but it is most unlikely to serve all or the majority of interests.

2.1.2 Classification

The process of classification (very much related to logical division) provides an alternative. This method moves from the consideration of particulars which are grouped at increasingly general levels. This method is made yet more attractive by statistical methods such as *Cluster Analysis* which provides a computational procedure for grouping the attributes of particulars using data obtained empirically. Thus one might develop a classification of accidents by studying their causes and using these as attributes in an analysis. The expectation of this procedure is the grouping of accidents into species of increasing generality (that is from sub-species to species to genera to summum genus). The resulting *family tree* (or *dendrogram*, SPSS Inc., 1988) would represent an empirical definition of accidents. However, such results (of classification) regardless of the process, are obviously dependent on the attributes included in the scheme and, in turn, the choice of attributes is guided by some prior conceptions about the essential nature of the genus to be studied. Thus whilst there is much to recommend classification it does not avoid the central problem of defining the essential characteristics of the genus to be studied.

Whichever approach is used, the definitions need to be rigorous. For the purposes of this thesis I shall endeavour to conform to the six requirements proposed by Joseph (1916). These require (paraphrasing) that a definition:-

(1) must give the essence of that which is to be defined;
(2) must per genus et differentiam (clearly identify the feature that distinguishes the instance from the general case);
(3) must be commensurate with that which is to be defined;
(4) must not, directly or indirectly, define the subject by itself;
(5) must not be in negative where it can be in positive terms; and
(6) should not be expressed in obscure or figurative language[20].

2.2 The Essence of *Accident*

Many writers (eg. Bird & Germain, 1986 and Johnson, 1980) have distinguished accident from incident on the basis that the former involves loss and the latter does not. However, the two terms are common in that they describe a class of event that is not desirable. This suggests a genus "undesirable event" that includes the two species incident and accident. However, it seems clear that what is undesirable about an accident is the entailment of loss but what then is undesirable about an incident which is definitively loss-free?

Apart from the practical matter that incidents *do* tend to incur *minor* loss (disruptions to process, inconvenience to staff etc) the fact is that "*under slightly different circumstances*" (as Bird puts it) loss would have occurred. Thus a more rigorous, if clumsy, definition of incident is *an event caused by factors likely to cause other events resulting in loss.* Thus the undesirable element is not in the loss but in the prior "factors" whose identity is in turn defined by relatedness to undesirable consequences.

2.2.1 Qualifying use of the term *Event*

It seems to me that the difficulty indicated in the foregoing paragraphs lies in the term "event" which connotes a singularity whereas an accident requires at least two events (1) the occurrence of factors likely to cause loss (2) the damage actually occurring. A similar observation is made by Hale & Hale (1972).

> "Many researchers have also considered injury and accident to be synonymous. They fail to recognise that the accident, the injury and the reporting of the injury are three successive behavioural events and that different factors can have an effect at different stages in the chain."
>
> Hale & Hale, 1972, page 79

[20] or Latin!

For instance, if I cross the road but fail to notice the bus that is bearing down on me; a particular state prevails - the conjunction of myself and trajectory of bus. If either the bus or I manoeuvre with sufficient alacrity, that is, act to cause a new state to prevail then I might appear unscathed. But if time does not allow or the action taken to change matters is otherwise deficient, the mechanical energy of the bus will cause me to change state from one of well-being to one of injury.

Thus the description *event* is misleading given that on a micro time-scale there is a succession of states each of which, though short-lived, are distinct. The upshot of this is that what defines an event is the *subjective placement of boundaries in time and in other dimensions of interest* (bus, myself, their vectors, relative proximity etc). An event, then, is a concatenation of a succession of states causally connected to a state, or future state, of adversity (characterised by loss).

The definitions provided by the HSE vary, but one of the more recent (Successful Health And Safety Management, HSE, 1991) states:

> "Accident includes any undesired circumstances which give rise to ill health or injury; damage to property, plant, products or the environment; production losses or increased liabilities".
>
> *and*
>
> "Incident includes all undesired circumstances and near misses which have the potential to cause accidents"
>
> *Successful Health And Safety Management*, HSE, 1991, page 66.

The HSE definition of incident seems the more immediately problematic by the separation of "undesired circumstances" from "near misses". It is not at all clear what distinction is being drawn between these terms. Other than this, the definitions seem reasonably in accord with the arguments I have presented given the proviso that "circumstances" are synonymous with "states". However, a weakness in these definitions and in the argument that I have so far developed is that neither capture the characteristic element of *rapidity* commonly identified with accidents and incidents. We need to address the *"it all happened so fast"* characteristic.

2.2.2 Time and rate of change

It is in this regard that my adoption of the terminology of systems and states[21] of systems has some advantage. The causation of occupational illness is generally regarded as operating on a longer time scale than the trauma entailed in accidents. There are other distinctions that might be drawn between occupational illness and accident trauma, for instance the agent of harm. However, this is also bound up in the variable of time. For example, pathogenic agents generally require greater time to overcome the body's defences than, say, voltage or mechanical force beyond the threshold of damage to body tissues.

Using the system terminology, we can view the accident event as the bounded epoch in which a system goes from a desirable or neutral state through various transition states towards undesired states. Beer (1981) suggests that this change of state (eg. from state a to state m) can be said to have a *trajectory*. I think it permissible to extend this language further: As well as a trajectory $a \rightarrow m$ one might also introduce the rate of this change by considering the time elapsing between $a \rightarrow b$, $b \rightarrow c$, $c \rightarrow d \ldots$ $n \rightarrow m$. The apparent feature of "rapidity" in accident

---

[21] The term "state" is of course much older than systems theory. The earliest definitive uses of the term I can find go back to Aristotle's doctrine of categories in the *Metaphysics*. The categories may be described as a list of predicates, one or other of which declares the mode of its essential being belonging to any subject which exists. Aristotle's scheme contains ten:

1)	*Substance*	corporeality
2)	*Quantity*	size, weight, number of
3)	*Quality*	colour, loudness, justice, virtue
4)	*Relation*	heavier, before, above, adjacent
5)	*Place*	here, there,
6)	*Time*	a date
7)	*Situation*	horizontal, sitting, lying. Presupposes the distinction of whole and part as well as the categories *place* and *relation*.
8)	*State*	Something which characterises a whole through the condition of its parts (whereas a condition needs to present throughout a given subject). Hence state is more complex than quality. *State* presupposes the distinction of whole or part, which in material things at least, implies the category of *quantity*, and it presupposes the categories of *Action, being acted on* and *quality*. For a whole is in a certain state through the interaction of parts having certain qualities, as when the body is well or ill; or through some done to certain parts of it, as when a computer is off-line or a car is red.
9)	*Action*	running, flying, pressing
10)	*Passion*	(being acted on) pressed, pulled, lifted, electrocuted.

Aristotle often sets *state* and *situation*, apart as they are derivative categories. However, whilst derivative they are distinct and contain something not in the notions from which they are derived. Kant (1937) objected to the inclusion of derivative at the same level as the underived, or pure, categories. But, for Kant, the Aristotelian scheme classified the products of perception and it was the processes of perception which provided Kant with a different scheme.

events may be due to the speed of all or the latter state transitions. Transitions from positive to adverse states which approach or outstrip the response time for system controls might fairly be called accidents (when one of the later transition states involves actual damage). This need not be on the microsecond scale as the hour-scale of Three Mile Island accident makes clear; it is the **match or mismatch** of control reaction to system change, not subjective labels of time that count.

In instances where the system moves slowly (especially on the scale of days or years) between normal and detectable adverse states in persons, this might more usually be referred to as the aetiology of an illness. From the control point of view the situation is not different given the grounds already provided. What is different is inconspicuity of cause and the obscurity of linkage between cause and effect. One has to know that the agent is harmful (knowledge of cause-effect[22]) and know it is present (at all or above a threshold). However, the necessary stages are common to both: a change of state in the system is matched by a change of state in the controller of that system. This view is in accordance with the Hale-Glendon (1987) model[23] of accident causation albeit in different and, I believe, more general terms.

To summarise; so far we have seen that:

- the term "accident" contains a degree of subjectivity concerning the boundaries in time and other dimensions which I have described as the parenthesis of event;
- this event may be comprised of several distinct states of which those termed adverse comprise the set of states which are of immediate interest to us;
- the adoption of the system-theoretic term *state* is profitable because, inter alia, it allows us to discuss accidents, incidents and ill-health in the same language;

[22] I would not like the reader to run away with the notion that relating cause with effect is a routinely successful cognition in more rapid sequences. In logic, this is known as the post hoc fallacy (after this, therefore on account of this). I came across a tale which serves to illustrate this: in the late 1940's a mother and her sons, Adam and Ben, were travelling on a train. In the course of the journey the mother gave the boys a banana each; neither boy had eaten a banana before. Just as Adam swallowed his first bite the train entered a tunnel and they were enveloped in the darkness. Adam stopped Ben taking his first bite saying "don't eat it, they make you go blind"!

[23] The recognition of "danger" is itself a change in state of the system controller which matches a prior change in state in the system to be regulated.

- in particular, the idea of change trajectory and rate of change, permits us reasonably to distinguish accidents and incidents from ill-health in terms of rapidity and;
- the control actions are common in principle to both accidents and the aetiology of ill-health.

Underwriting all of this is the notion of *adversity* of which little has been said so far. There really is too much to be said about this aspect to deal with it here and so the matter has been made the subject of chapter 3. For the moment suffice it to say that adverse states are those which are identified as *those antagonistic to the purposes of the system* as perceived by the various entities (ie, people and groups of people) that have an interest in the system.

The reader may care to note that my adoption of *"adversity"* is coincidental to the use of the same term in the introduction to the Royal Society's report on Risk. Fortunately, my *use* of adverse does not conflict with theirs. However, the Society's definitions imply rather than state what adverse can be predicated of, and the implication is much more restricted than I think it ought to be. This is considered in Chapter 3.

2.2.3 Accidents and intention

If a person intends a state to exist that state cannot be said to be accidental; it is an objective. If one intends a state to exist but the system configuration designed to bring it about involves unintended states, then all of these states can be regarded as "accidental" in the logical philosophic sense. A sub-set of these will meet subjective criteria of adversity (and the other criteria previously mentioned) and these will be accidents by definition. This is unproblematic so far but if these accidental states have been **identified** in the process of developing or overseeing the evolution of a particular system configuration, is acceptance of their possibility tantamount to their being intended and, therefore, not accidents? In other words, if you expect it to occur and it does occur - can the occurrence be called an accident?

This is a problem. If I design a system to produce a desired set of behaviours and in doing so identify, albeit at a low probability, other behaviours that fall under the category of adverse which are somehow entailed by the system itself - I can either design them out in an absolute sense or live with their possibility (albeit on a reasoned basis of tolerable risk). Even if I do not want them but accept them

because I accept the system then, just as I intend the system to be, I must also intend these states to occur.

Within MORT, adverse states of the this type are bracketed under the head of "assumed risk". The processes underlying the assumption of risk are discussed at length in the MORT literature, however, their role as a species of event is not explored. Nevertheless, the MORT tree logic renders accident and assumed risk as distinct causes of loss. This is shown in Figure 2.1

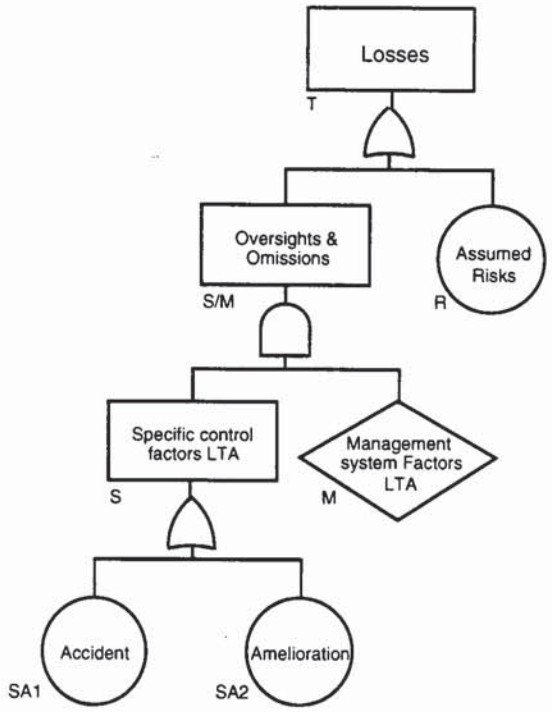

Figure 2.1 MORT Tree at SA Limit of Resolution

In practical terms, the assumed risk (albeit at a low probability, say, 1×10^{-5}) cause of losses will be dealt with in a different way than the unidentified accident event. In the case of the latter, the hazard identification and risk assessment machinery of system control is brought into question (hence "oversight and omission") whereas the former has no such connotation. Thus the distinction is more than academic.

Within the present context, we have to decide whether (1) a lack of intention is essential to the character of *accident* or (2) it is a logical-accident in the definition. A logical-accident, as a term in the technical vocabulary of logic, is described by Joseph (1916) as "an attribute which equally may or may not belong to a subject" (page 75). For example, if we were defining the term organism: it is a

logical-accident in the definition of an organism that it be used as food (because it can be used this way or not used this way). On this basis, I think that the matter of *intention* is a logical-accident in the definition of *accident* because the essentials are already commensurate given the terms (i) rate of change (ii) duration of causal sequence (iii) mismatch between control and the system behaviour and (iv) attribution of adversity. Thus the matter of *intention* is made subordinate to these issues and where it "belongs" will be considered in chapter 3.

2.2.4 A working definition of *Accident*

On the basis of the discussion so far, I suggest the following definition for the term *Accident*:

> ACCIDENT: An event of *short duration* during which a sequence of changes of state occur in a *system* that regulatory mechanisms fail to control before a state of *adversity* is reached.

The emphasis added to the definition is to alert the reader to the subjective nature of the key terms *short duration, system* and *adversity*. This may be uncomfortable, I find it so and perhaps it is empiricist conditioning that is responsible, but it is unavoidable.

The reader may note that this definition applies to *any* event of the type. It is super-ordinate to considerations about the kind of adversity and, therefore, to the mechanisms particular to it. Hopefully this position allows accidents under the H&S head of adversity to be considered with accidents involving other kinds of adversity *in terms of the objective elements of the definition:* the regulatory mechanisms aimed to regulate system states.

2.3 The Essence of *Incident*

I do not propose to repeat all the ground covered in section 2.2, and, as I see it, there is only one issue that requires expansion. This concerns the general acceptance of an incident as an accident in all respects except the realisation of states of adversity.

It seems to me that there are two types of incident. The first has all the attributes of accident but is curtailed by timely action by the controller. Nevertheless this timely action intervenes at a transition state whose probability of changing to an adverse state is considerably beyond the threshold of acceptable risk. For instance, a wrench is dropped from a scaffold. During the transition state (wrench falling and people underneath) the person who dropped it shouts a warning (intervention or an act of regulation) this changes the state of the system by causing the would-be target to move out of the path of the wrench which embeds itself in the ground.

The second incident type again has all the attributes of an accident but no regulatory interventions and terminates without attaining an adverse state. As before, what is essential to the definition is that the probability of the various adverse states that can be reached from the same antecedent state (eg. a wrench in free-fall) are beyond the threshold of acceptable risk. What is characteristic of this type of incident is that chance alone precludes adversity rather than adventitious action as in the previous type.

The first type of incident is a testimony (that being characterised by intervention) to the fact that accidents are as rare events as they are observed to be. As much as workers in the field of safety wring their hands about human error, the hands of people in systems more often than not catch systems in the act of falling. Whilst we should persist in our efforts to produce error tolerant designs, reducing human intervention in systems (and thus the opportunity for error) needs to be balanced against the control resource that people embody. I believe that promoting this ability is at least as important as reducing the need for its exercise in earnest.

Perhaps unexpectedly, it seems that *incident* is a more complicated notion than *accident*. Whereas the character of *accident* is in the adverse state(s) that occur during the event, incident requires acknowledgement of probabilities of different states within the event (particularly the probability of the end-state).

> INCIDENT: An event of *short duration* during which a sequence of changes of state occur in a *system* that regulatory mechanisms fail to control before the probability of an *adverse end-state* exceeds a *threshold of acceptability*.

As before, the italicised phrases indicate subjective elements in the definition and the definition is stated at a super-ordinate level permitting any event to be identified regardless of the kind of adversity attributed (eg. threats to life and limb, financial, reputation, customer relations, etc).

2.4 *Hazard* and *Risk*

A straightforward approach suggests that risk is nothing more than the probability associated with a given adverse state of a system. However, embedded in *risk* is the notion of *hazard* which, so far as I am able to discover, is particular to the H&S head of adversity and this differentiates hazard from the other concepts we have so far considered.

2.4.1 The essence of *Hazard*

The term hazard appears to confuse many people and those it does not confuse seem hard pressed to define the term precisely. This is of particular moment because of the fundamental importance of the hazard identification stage in risk assessment. Consider the following quotations:

> "We define Hazard as - The potential in an activity (or condition or situation) for sequence(s) of errors, oversights, changes, and stresses to result in unwanted transfer of energy with resultant damage to persons, objects, or processes".
>
> Johnson, 1980, page 247

> "Hazard - something that can cause significant harm".
>
> Stephenson, 1991, page 8

> "Hazard means the potential to cause harm, including ill health or injury; damage to property, plant, products or the environment; production losses, or increased liabilities".
>
> HSE, 1991, page 66

> "Hazard - A source or a situation with a potential for harm in terms of human injury or ill-health, damage to property, damage to the environment, or a combination of these".
>
> British Standard 8800, 1996, paragraph. 3.4

> "Hazard: a situation that could occur during the lifetime of a product, system or plant that has the potential for human injury, damage to property, damage to the environment, or economic loss".
>
> Royal Society, 1992, page 4

Of these five, Stephenson's definition is set apart and the difference is made emphatically clear by Levens and Krikorian (1970) who assert that "*Things* are not hazards or potential hazards. Events or a sequence of events are hazards". It appears that Stephenson is identifying *hazard* as the agent[24] that does the damage and this approach leads to difficulties when the term hazard is pressed into practical usage. For instance is a wall a hazard?... not unless I walk into it and demonstrate Newton's 3rd law of motion. Is a 20kv line a hazard?... not unless I provide a route to earth for the energy it conveys. Hazards arise out of interactions between elements in a system. Again, the concept of system states, of state-change trajectory, of transition states and adverse states developed earlier in relation to accident and incident is of assistance here.

Taking a power press as an example and labelling the salient attributes of states as they unfold during an accident event on a "micro" time-scale, we might arrive at a primary event line (using the terminology of Events and Causal Factors Analysis, Kingston-Howlett & Nelson, 1995) resembling Figure 2.2 below.

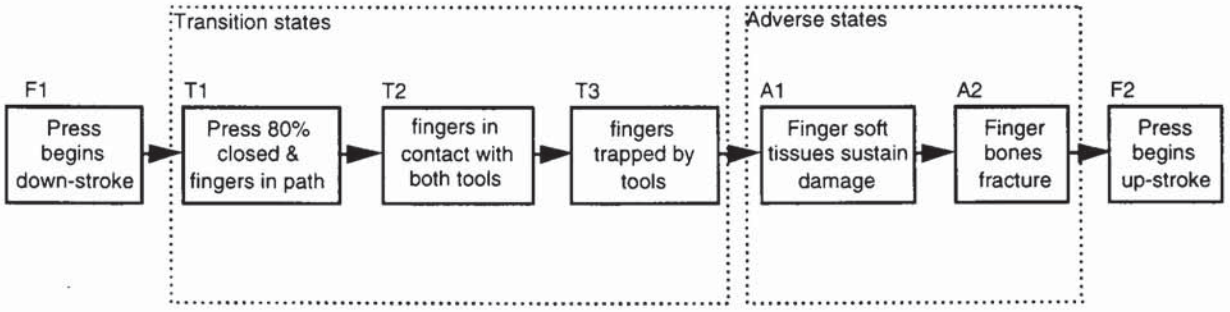

Figure 2.2 Hypothetical power press accident represented as a primary event line illustrating sequence of state changes. The annotations "F", "T" & "A" are Favourable, Transition and Adverse respectively[25].

[24] "Agent" in the sense described in section 2.2.2 (page 39, ante).

[25] Figure 2.2 represents a highly simplified cut-set and is suggestive, therefore, of a fully determined system. It is only reasonable to suppose that the kind of system involved in the

First there is a favourable state, the press closing in the normal way. The next state shows the press to have completed 80% of its down-stroke (an arbitrary value somewhat less than that required to trap the fingers) and the operator's fingers in its path. At T_2 the press has closed to the extent of contact between the tools and the operator's fingers. At T_3 the down-stroke continues and exerts a greater force than at T_2, entrapping the operator but still below a threshold of tissue damage. Were it the case that the sequence could be stopped at this point: no physical damage would be sustained by the operator. The operator would have changed his or her state to one of alarm and pain (perhaps) but the state of tissue damage has yet to be reached at A_1[26].

Returning to the matter of *hazard*, for Stephenson the press *is* the hazard because a press, as the example illustrates, is "something that can cause significant harm". However this lacks the dimensionality of time and interaction necessary for precision. For the HSE it is clear that the press has "the potential" for harm but this too vague, as I trust the "walking into the wall" example illustrates. Johnson's definition seems to capture more of the necessary detail but still seems imprecise.

Using the state description, it seems to me simpler and more exact to describe hazard as a transition state at which the probability of a (consequent) adverse state becomes unacceptable. When the probability of an adverse state is very high (ie, approaching unity) we might call this transition state the point of no return[27]

This approach reflects the ordinary usage of hazard: a tripping hazard for instance describes the ultimate transition state during which (whilst falling) the unfortunate is not injured but at which the probability of an adverse state is unacceptably high or near unity. This allows the possibility of a number of adverse end-states, distinguishable one from the other by nature and degree.

---

example is Markovian; that is a network of distinguishable states whose transition one-to-the-other is given by a probability. An Event Tree is an essentially Markovian model.

[26] I have chosen this level of change to the operator as the first adverse state, although I might have put it earlier in the sequence, say, at the point where discomfort became noticeable... please excuse the sangfroid, this is only a hypothetical example!

[27] the HSE (1996) in "Generic terms and concepts in the assessment and regulation of industrial risks" refer to this situation as "peril". However, their definition of hazard "the disposition of a thing, condition or situation to produce injury" is not discussed here as it does not, in my view, contribute anything positive to the debate.

This conception of hazard is perhaps a little radical. It has lost the common-sense notion of that which *causes harm* in favour of that which *makes harm foreseeable:* It is not the fall that kills you, it is hitting the ground.

However, *hazard* as the postulated "state of no return" seems in essence a matter of control, or rather, the lack of it. Thus the need to qualify the probability associated with consequent adverse end-states is, perhaps, unnecessary. A simpler approach might be to define hazard as a state from which one *cannot control* consequent states, thus leaving the matter of the adverse end-state to stochastic factors beyond control. Further, the point at which "cannot control" arises, is movable by making changes to the system. For instance, if one is working at a height and loses purchase - this is the point at which control is lost. However, if a safety harness is worn, control is retained at the state labelled "loses purchase". However, if the harness breaks under the strain, the moment of its failure becomes the state at which control is lost. Lastly, if the worker loses purchase but some hypothetical energy dissipating material covers the landing area: control is never truly lost because there is no path by which this system can reach a state of adversity, hence there is no hazard[28].

This developing definition appears to agree well with my prior definitions of accident and incident, it also seems to agree with hazards in the arena of occupational ill-health. Here the fact of lengthy separation of exposure to effect, or gradual cumulative effects or dose-response effects are all rendered on a common basis. The only effective difference is that the system states leading to an end-state characterised by accident trauma are all external to the victim (but internal to the system as a whole). In the typical aetiology of occupational ill-health, some of the system states leading to the adverse end-state are internal to the victim. This is not a difficulty for the definition of hazard so long as it is remembered that the system must minimally contain the agent of harm, the target of damage and their immediate surroundings.

[28] Although it is hard to envisage an HSE inspector being easily satisfied by an approach to site safety that is characterised by construction workers raining down from heights.

2.4.2 A working definition of *Hazard*

Given the above, I propose the following definition of hazard:

> HAZARD: A non-equilibrial *system* state with trajectories to *adverse* states characterised by transition probabilities greater than those *acceptable* to *system* regulators.

This can be informally rendered as: *a situation likely to lead to harm or damage*.

2.4.3 A Working Definition of *Risk*

One of the aims of the definition of hazard is clear the path for the definition of risk. Risk is often confused with hazard, probably because the concept of hazard is so intimately connected with the *cause* of harm. This confusion is apparent in the definition of *risk* provided by the Royal Society (and, thereby, in BS 4778, upon which it is based) which is

> "Risk: a combination of the probability, or frequency, of occurrence of a defined hazard and the magnitude of the consequences of the occurrence"
>
> Royal Society, 1992, page 4

This is difficult to reconcile with their own definition of hazard as a "situation... that has the potential for human injury..." unless the probability of a given adverse state is unity, given the occurrence of the hazard state. As argued in section 2.4.1, cases where the probability of a defined adverse state is <1, given the pre-existence of hazard state, are not only possible but more likely to be the norm. However, as I hope the foregoing has made clear, the hazard state is characterised by the loss of control of an agent of harm and someone/thing who is thereby exposed to it. Risk concerns the probability of this whole set of states resulting in a given adverse state (monetary loss, damage to property, injury of people etc.). Thus, simply –

> RISK: The probability of a specified *adverse* state in a specified *system*.[29]

[29] The need to specify a time period, which ordinarily features in rigorous definitions of *risk*, is here subsumed within the phrase "specified system".

As previously, the italicised words indicate subjective terms in this definition which robs "simply" of application to anything accept the wording of the definition. Both "adverse" and "system" are the subject of further discussion in Chapter 3, but suffice it to say for now that because two of the major terms in the definition are subjective, the objective rigour suggested by "probability" (even if we were sure of our data) is illusory.

What my working definition leaves out and that the Royal Society's contains is "*the* ***magnitude*** *of the consequences*". The matter of consequences is already included as "adverse state" but the *magnitude* of adversity is a further dimension contained when risk is used as an index of acceptability. I do not see *magnitude* as essential to the concept of risk, but it *is* essential to the concept of risk acceptability if different adverse states and sets of states are to rendered on a common scale of acceptability. This topic will be further considered in Chapter 3.

2.5 *Safety*

The work of this section is two-fold. First, as with the definitions worked through previously, the term *safety* will be considered conceptually and a technical definition provided. Once derived this provides the basis for a definition of safety management. Throughout this section, *safety* is seen as containing occupational health. This is not because there are no distinctions to be drawn between the professions, but because the current debate is not affected by these distinctions.

2.5.1 The concept of *Safety*

The literal meaning of safety can be given as "freedom from adversity" and this is a satisfactory definition of the ideal. In practice, however, it hardly needs saying that risk of adversity is generally present even if the probability associated is negligible. Hence, one has little choice to opt for a definition where the absolute term "freedom" is qualified if the term safety is to be retained in the technical (ie, practical) vocabulary. An example of where this has been done is provided by the Royal Society (1992) and their use of the phrase "unacceptable risk"

> "Safety: the freedom from unacceptable risks of personal harm"
>
> The Royal Society (1992)

Leaving aside the restriction of the definition to "personal harm", from the previous section it seems apparent that *unacceptable* is predicated of both (i) the nature of adversity and (ii) the probability of a system entering a state characterised by that adversity. Let us suppose that there is consensus by all affected parties on what states are adverse and about the probabilities pertaining to such states given a particular system. Let us further suppose that the system has consequently been modified such that the probabilities of adverse events are all below the threshold of acceptability. Now, the nature of probability is such that no matter how small the value, the adverse state may occur now as at any time. This fact does not easily square with the quality of freedom, and whilst I grant that most freedoms are relative in practice, talk of "the relative freedom from unacceptable risk" would be unnecessarily obscure. However, if risks are divided into two sets - those that are and those that are not acceptable - what is relative is *the placing of the boundary* between the sets. Once the boundary is established through policy and implemented in practice, what we have is an unqualified "freedom from unacceptable harm" because any harms that do then occur are acceptable (given that the implementation is perfect).

To recapitulate, the definition of safety given by the Royal Society includes the notion of freedom; in this case, *freedom from unacceptable risks of personal harm*. Thus safety is defined as the freedom from unacceptable risk but **not** the freedom from acceptable risk. This implication arises because of the need[30] to qualify freedom from its absolute ideal to its relative reality: the relativity arising from the distinction between acceptable and unacceptable risk of personal harm. Hence, without introducing new terms into the definition but restating it positively (the fifth of Joseph's requirements for definitions) we obtain:

SAFETY: the exposure to acceptable risks of personal harm

This result is interesting insofar that an attempt to state a positive yet rigorous definition of safety results in a definition somewhat alien to the normal sense in which the word safety is employed. Nonetheless, it does present a qualitative definition capable of quantitative rendering if we allow "harm" to be a dimensional quantity. It may be made quantitative by describing probability (or estimated frequency) as a function of the harm caused (expressed a some measure - a monetary value is often used).

[30] Otherwise the definition would be "freedom from risk of personal harm" which, given the extreme rarity of a system in which the probability of any harm is zero, could not be generally applied.

This definition may however require further qualification because the term "acceptable risks" is relied upon and thus begs the question *of acceptable to whom?* This definition then makes starkly apparent the major philosophical concern of safety - *how* is the partition between acceptable and unacceptable risk to be placed, *who* or *what* has authority in this matter.

For the moment, it suffices to say that we have to include this "who or what" into consideration whenever safety is used in technical rather than purely ideational sense. In effect, the practical usage of the term requires us to also nominate the authority relied upon to make the judgement of acceptability. Convenience suggests that I should name this entity and the first noun that recommends itself is *regulator*. However, for reasons that shall I hope become clear, there are good reasons to reserve *regulator* for special use later. Instead let me ascribe this role to an entity called an *arbiter*. Whatever or whoever this arbiter is in practice, and regardless of the sagacity they bring to bear, the placement of the line between acceptable and unacceptable is subjective... if only because the matter of *harm* (and adversity in general) is subjective.

2.5.2 Does *safety* apply only to the harm of people?

Armed with the caution to remember the implicit arbiter, we now need to see whether the definition of safety can be broadened (beyond the particular application to *personal harm*) so as to be consistent with the associated definitions provided earlier. As argued at section 2.5.3 below, I do not see any necessity to restrict safety to exclusive concern with personal harm and, indeed it is generally acceptable not to restrict it in this way. Safety is ordinarily predicated of financial investments or of delivery of merchandise (just as it as one can talk about the risks associated with either of these examples). It may be that the primary concern of health and safety professionals *is* the prevention of harm to people but that is not of particular moment here.

If this line of argument is accepted then our definition of safety as *the exposure to acceptable risks of personal harm* can be reduced to *the exposure to acceptable risks,* where a risk is "the probability of a specified *adverse* state in a specified *system*".

Now something is not quite right with this. Writing the definition out in full makes this clear:

> *SAFETY*: the exposure to acceptable [probability of a specified adverse state in a specified system]

At first sight the problem is simply one requiring the risk definition be stated in the plural. The problem goes a little deeper: In my definition of risk I had drawn a line around an assembly and nominated it a system; the vantage point was above the system viewing, as it were, the totality of probable states. The Royal Society definition of safety, on the other hand, does not imply a particular system or the relations of the target of harm to the source of the risk. This is not a flaw in the definition of safety in its ideational sense but it does weaken the definition it its technical sense. The reason being that the subjective judgement of the arbiter (of acceptability) will be conditional on the drawing of the system boundary and, in particular, the relation of the target of harm to the source of the risk.

For example, in the UK the arbiter of nuclear risk acceptability appears to be the HSE. Drawing the system boundary to include only workers, the arbiter produces an annual probability of 10^{-3} for death as "the dividing line between what is just tolerable and what is intolerable" (HSE, 1988)[31]. However, drawing the system boundary so as to include members of the public, the same arbiter produces a annual probability of 10^{-4} for death. It appears then that the matter of specifying the system is a matter of consequence and this is justification for including the notion in the definition of safety.

Returning to the definition, the issue just addressed comes down to the word *exposed* in the definition, as it imposes the perspective of the target of harm rather than the wider system, the identity of which is of consequence in determining risk acceptability.

Taking a system as a whole, the definition above can be transformed to:

[31] The tolerability of risk from nuclear power stations. HSE, 1988 (page 23). In the simplified scheme of an arbiter drawing a "line between acceptable and unacceptable risk" one might expect that if 10^{-4} is intolerable then an immediately smaller probability is acceptable... in fact, this is not the case. 10^{-6} is intolerable in this context and the reasons for a "line thickness" of two orders of magnitude will be considered later.

> SAFETY: a condition of a specified *system* in which the probabilities of *adverse* states[32] are equal to, or less than, those *acceptable*[33].

In this sense "Safe" describes a system so regulated as to realise adverse states with a frequency acceptable to the arbiter of the system's behaviour. It needs be said before passing on that this definition, whilst fit for technical use does seem to have lost something along the way - one can hardly imagine this definition as the rallying cry by which the serried ranks of the "safety first" movement were galvanised. Having dissected the body the heart was mislaid.

2.5.3 The scope of *Safety*

As noted above in regard to the definition of safety at 2.5.1, health and safety traditionally includes all injury incurred through accidents and all ill-health which is occupational in origin. In the case of accidents, some incur damage to property and process as well as injury to people. Therefore, the scope of safety is widened to include both subjects of adversity. By including property and process-centred adverse states, there is a fuzzy boundary between accidents which damage (1) both people, things & process and (2) things and process alone. The latter case has departed from the traditional scope of "Health and Safety" but is connected via mutuality of cause and generality of effect (adverse states).

As time moves on, the scope of safety appears to be ever broadening and the classification boundaries mentioned, crumbling. For example, the range of the HSE definition of hazard quoted earlier "...harm, including ill health or injury; damage to property, plant, products or the environment; production losses, or increased liabilities" and the frequency that one encounters *safety* and *quality* in the same phrase or organisational title.

The question arises as to how far the scope can be widened before safety as a genus bursts at the semantic seams and yields to another term entirely. At the narrow extreme, the scope is defined by adverse states of systems the adversity of

---

[32] See also the discussion in chapter 4 (page 132, ante) concerning the distinction between individual adverse outcomes and outcomes considered sum over time.

[33] As before, the previous definitions, the italicised words indicate subjective terms

which is chiefly in the human elements of the system (ie, occupational injury and ill-health). This would fit with the more traditional view of safety (current until challenged in the 60's in the US with the genesis of "system safety" approaches; (Johnson, 1970; Stephenson, 1991). The other extreme includes all adverse states into the scope but drops the term *safety* as an arbitrary classification - an extreme form of the view reported with reference to "Japanisation" (HSC, 1993).

The rationale of the people-only scope for safety would appear to be that owners of systems have a right to risk the damage and destruction of all things that they *own*. One cannot *own* a person, and this provides a "natural" basis for considering safety in this way. However, as argued, the causes of adversity for people may sometimes be common to adversity located in objects or process etc. Hence, the narrow treatment of safety results in somewhat artificial classification boundaries.

On the basis of the above, there appears to be a compelling case for the widest definition of safety in terms of the adverse system states to be included. As the quotation from Joseph earlier states *"Things belonging to one genus will be studied together"* and viewing safety as essentially a matter of control of adverse system states provides the unifying concept. However, in doing so *safety* becomes subordinated to a wider concept of systemic control.

2.5.4 A working definition of the *management of safety*

If safety is "a condition of a specified *system* in which the probabilities of *adverse* states are equal to, or less than, those *acceptable*" then the management of safety is to ensure that this condition is perpetuated from moment to moment throughout the life of this system.

> THE MANAGEMENT OF SAFETY: the control of a *system* such that the actual probabilities associated with *adverse states* of that *system* are equal to, or less than, those *acceptable*.

2.6 Summary

The endeavour of this chapter has been to analyse the basic terms of safety in order to deliver the conceptual building blocks for use in the remainder of this thesis. I have argued that the terms accident, incident and hazard are best

described in the language of systems theory, and this approach yields the concepts of control and adversity and acceptability as essential to safety.

TERM	DEFINITION
ACCIDENT	An event of *short duration* during which a sequence of changes of state occur in a *system* that regulatory mechanisms fail to control before a state of *adversity* is reached
INCIDENT	An event of *short duration* during which a sequence of changes of state occur in a *system* that regulatory mechanisms fail to control before the probability of an *adverse end-state* exceeds a *threshold of acceptability*
HAZARD	A non-equilibrial *system* state with trajectories to *adverse* states characterised by transition probabilities greater than those *acceptable* to *system* regulators
RISK	The probability of a specified *adverse* state in a specified *system*
SAFETY	a condition of a specified *system* in which the probabilities of *adverse* states are equal to, or less than, those *acceptable*

Table 1 Summary of definitions

3 Systems & Adversity

"No thoughtful man versed in the methods of natural enquiry can fail to be reminded at every moment of the ultimate and universal dependence of every one group of phenomena upon every other."

T.C. Allbutt (1896)[34]

INTRODUCTION

In Chapter 2, I suggested that the disunity of the safety literature was in part due to a failure to agree terms clearly. In the course of deriving definitions for some of these terms, deeper issues became apparent: it is one thing to disagree about the meaning of a given term but quite a different order of complication is engendered by failing to achieve a coherent philosophy which allows whatever terms to mean anything in particular. This is ordinarily captured within the notion of epistemology: a branch of philosophy "concerned with the origin, structure, acquisition, and validity of knowledge" (Klir, 1991). Without travelling deeper into this rather profound issue it is, I hope, sufficient to reduce this to the commonplace observation that whenever we talk about something we inevitably rely on a set of associations. Hence, any given term is distinct because it enjoys a *unique set of relations* to other terms.

For the present purposes what is required is to choose, on a rational basis, the pattern of relations which allows us to be consistent in our treatment of the various subjects that need to be included. Clearly, the choice of subjects will have a considerable impact upon the overall network of relations that is obtained. In the physical sciences, the process of discovery can be seen as the making of connections which, as time goes by, produces an increasingly complex body of knowledge. What renders this body of knowledge accessible is the recognition that whilst there is a unique set of relations associated with any one node in this complex body, the *pattern* of these relations is sometimes repeated for other nodes.

---

[34] Allbutt, T.C. (1896) "A system of Medicine"... *cited* in Hale & Hale (1972)

These regular patterns once recognised can then be mathematically codified as principles, laws and axioms.

Whereas empirical science generally reduces the complex interconnections of natural phenomena by artificial constraint, research into health and safety management does not ordinarily have this option[35]. I intend no criticism of the calibre of research done in this field by drawing comparisons with the physical sciences. The purpose of the comparison is to highlight the epistemological shortcomings in the area which, from the argument above, are summarily:

I. By failing to agree basic terms, the implied network of association between terms becomes very fuzzy;
II. The lack of resolution in the associational network undermines the consensus as to;
   A. the meanings that can be ascribed to terms, and;
   B. the relationships which might render the subject coherent;
III. This compromises the reduction of the total number of relationships to a smaller number of principles.

This is illustrated below in Figure 3.1, where the left hand block diagram represents the ideal scheme and, the right hand diagram, the situation suggested at (1) to (3) above.

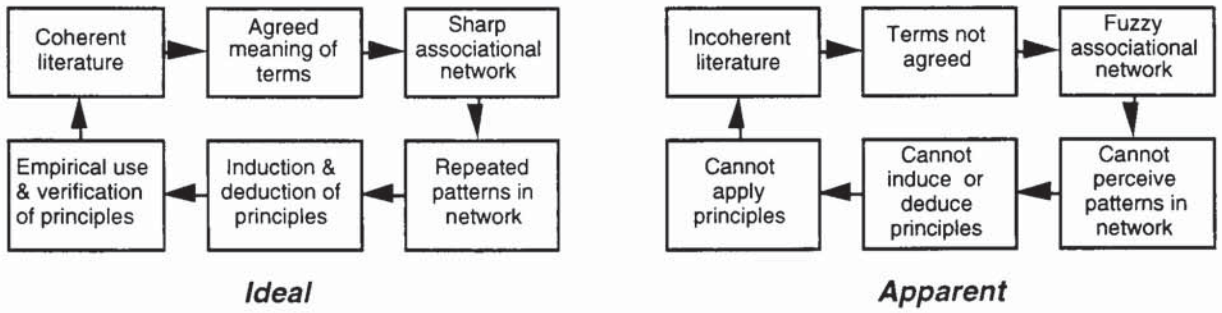

Figure 3.1 Ramifications of terminological inconsistency.

The arguments developed in the remainder of this chapter concern how we might better conceive of the management of safety and introduces concepts drawn from, inter-alia, the systems science and ethics literatures. It may seem a little inconsistent, but I do not suggest that there is a single "truthful" structure of

[35] The other mode, common to both physical science (eg. chemistry) and areas such as safety management research, is statistical of which more later (section 3.5.4, page 80 ante)

knowledge which can be obtained in the area. Indeed, a degree of "fuzziness" can be argued to be a valuable impetus to research - a spur to search for better resolution and with it the opportunity of adventitious discovery en route. The best we can hope for is greater uniformity of approach and description which might allow us to move more easily from one viewpoint to another and thereby allow judgements based, not on absolute truthfulness, but relative utility for the purposes of the researcher.

3.1 Researching safety management systems - the need for *consensibility*

The current problem in health and safety research is not the lack of useful models but the lack of correspondence between them and also the lack of correspondence between the models and the entities they are ultimately applied to - organisations. Hence one might consider the graphical depictions of say[36], the Hale-Glendon accident causation model, HS(G)65 and the loss causation model of Bird and Loftus and although they are all concerned with the same topic (accident prevention/causation) it is not clear how to integrate them. This observation does not necessarily support the argument that one or more of these models is wrong as there is support for all. As true for models generally, they are more or less useful depending on the needs of the person employing them and the situation in which they are being used.

However, whilst models are helpful to communicate ideas, unless there is a clear lineage[37] back to an agreed set of principals and evidence, the danger of failing to produce a *consensible* body of knowledge is very real. Consensible here means that the statements comprising this body of knowledge are "not so obscure or ambiguous that the recipient is unable either to give it a whole-hearted assent or to offer well-founded objections" (Ziman, 1991, page 6). For reasons mentioned earlier, and others mentioned later, I am ***not*** arguing that knowledge about health and safety management must satisfy the all criteria suggested for the physical sciences[38] but merely that consensibility is a minimum criterion[39].

---

[36] Oddly, these models are dealing with very different things; Hale-Glendon is pre-eminently concerned with the non-observable elements of the system (perception and decision) whereas Bird and Loftus is very much more concerned with the observable aspects (and at a much higher level of description). HS(G)65 on the other hand adopts a much more integrated view.

[37] A further issue here is that of isomorphic and homomorphic mapping - chapter 5

[38] Which taken as a whole, mostly falls foul of some of the more stringent criteria such as falsifiability and consensuality but still remains highly predictive of observable events.

[39] By way of example, my colleagues Florus Koornneef (University of Delft, Netherlands) and Philippe Schallier (European Space Agency) and I are in the process of establishing a network of accident investigators and interested academics in which MORT was advanced as a useful

From this, there seem to be two main areas of consideration. Insofar as safety management coheres within organisations, we require:

1. A theoretical model of organisations so as to inform us as to their general structure and the nature of regulation within this structure in terms of general performance;

2. A theoretical model of safety management which is, in principal, compatible with (1) above, and assists the unification of the literature so as to better allow knowledge in the area of H&S to be better applied.

Hence the "general view of the subject" suggested by Webb (c.f. chapter 1, page 17) finds relevance here. As has been mentioned previously, systems theory aims to provide a general perspective and, indeed, is implicit in the conventional parlance of the safety and general management literatures. The business of this chapter is to examine systems concepts for explanatory usefulness; explicitly in relation to (2) above and implicitly in relation to (1)[40].

3.2 Informal use of the term *System*

System is a much used word and embraces an even more widely used set of concepts. Some authors have argued that the word should be reserved for special use and have suggested replacement terms so as to make their special meaning unambiguous[41]. Others have been very generous in what kind of thing the word system can be applied to: scissors, use of information, businesses, geometry etc.

For reasons I hope to explain, whilst it is grammatically acceptable to use "system" as part of a noun phrase (in other words, I can refer to something in the world and say "that is a such-and-such system") it is a sort of *dodge*. Hence when we speak about a "filing-system" we should be understood because most people share a stereotyped idea about the set of operations upon the kind of things and the sort of purpose that are evoked by the phrase (eg. the sorting of documents by

---

approach. However, our principal motivation was not the utility of MORT as an analytical method but as MORT as a vocabulary to enable unambiguous communication - a necessary condition of consensibility.

[40] Chapter 4 will place emphasis on (1) as the cybernetics view of management is discussed.

[41] For example "Org" (Mesarovic, M (1970), "Holon" (Koestler and Smythies, 1969) and "integron" (Jacob, 1974)

a consistent method of classification into filing cabinets for the purpose of being able to find them easily at some time in the future). The dodge I mentioned is twofold: first, the word "system" serves as a token by which the whole set of ideas can be communicated and; second the word serves to demarcate instantly a set of objects and relations from the context on which they depend. In the example of a filing system, no real crime has been committed by using the dodge just as it is innocuous to spare oneself proving the whole of Euclid when saying "use pythagoras, $a^2=b^2+c^2$". However, when the token noun-phrase denotes a more variable set (of objects, operations and purposes) or a more complicated context (with which the set is associated), the assumption of shared meaning becomes less robust.

In view of this, when we talk of systems which are more heterogeneous (that is, less stereotypic than the filing system example) much greater caution must be taken if misunderstanding and vagueness are to be avoided[42]. However, there is some reason for optimism. In the filing system example, I mentioned operations, objects, purposes and context. Purpose is distinctive (in respect of operations, objects and context) because it provides the basic rules for what gets included (ie, operations and objects) and excluded (ie, the dependent context). Hence, if people in conversation can agree on the purpose[43] of the system they are discussing, they introduce a powerful constraint on the variety of what is included in the system (and hence reduce at a stroke, the probability of misunderstanding[44]). Once we can agree on the purpose being served we can be more confident about knowing what kinds of things are required to serve it.

Effectively, the method I am using here is the same as employed when deriving the definitions in chapter 2[45] (a definition must give the essence of that which is to be defined) in this case **the essence of a system is its purpose**. Chapter 2 argued that the essence of safety management systems was "the control of a *system* such that the actual probabilities associated with *adverse states* of that *system* are equal to, or less than, those *acceptable*". So, a safety management system is a system

---

[42] As an aside, just the inclusion of "system" into a noun-phrase carries with it a connotation of qualities such as rigour and consistency which may not warranted in all cases. This may be another kind of dodge entirely - a sort of rhetorical sleight of hand.

[43] An issue for later treatment is that agreeing as to the purpose of a system is far from straightforward.

[44] Logically, the smaller the variety in the category, the more probable a meaning will be agreed.

[45] I hope also the present discussion serves to further justify the importance of deriving the definitions beyond "mere" semantics. Once we have agreed meaning, what counts as proof and the distinction between relevant and irrelevant should fall in line.

which controls a system. However, reviewing Joseph's (Ibid.) requirements, this is suspect because a definition "must not, directly or indirectly, define the subject by itself". Hence, *a system which controls a system*[46] is not an answer but a starting place, the implicit question being what it is about systems which requires me to consider one system as defining another?

Before attempting to answer that implicit question (which is the business of the remainder of this chapter), let me mention another reason for optimism. Agreeing a purpose allows us to define the system in terms of what is included in it and what, outside, it depends on. However, the contents may still add up to an unmanageably large set of operations and *operands* (things operated upon). However these classes (operations and operands) may well be further divided into manageable subsets. In the case of the filing system, this "division[47]" is the basis for classifying the operands (documents). Similarly the operations performed upon them may be classified and thus reduce the number of distinct operations needed to describe or model the filing system *without loss of essential detail* This appears highly relevant to the ideal "reduction of the total number of relationships to a smaller number of principles" that was discussed in the introduction to this chapter.

Relating this back to safety management systems - agreeing purpose provides a way of deciding the membership for the set of operations and operands. It also provides a way of distinguishing between this set and the context upon which it depends and alerting us to the need to specify the dependency[48]. In other words a means of specifying the contents of a model of safety systems and the interface between the contents and the "external" consideration upon which they rely. Hence, whilst a comprehensive survey of safety management systems would reveal an unmanageable number of operations and operands it is possible, in principal, to reduce the dimensions of this problem by classification (in which respect, MORT is of considerable value to the present work).

---

[46] a machine built to achieve a purpose (safety) by constraining other machines built to serve other purposes

[47] Again in keeping with Joseph's requirements for a particular *fundementum divisionis*

[48] ie, the transaction of energy, materials and information between them.

3.3 Systems Theory, "System-Safety" and Safety Systems

The fact that the phrase "safety management system" is ubiquitous certainly suggests that system is an important idea within the management of safety. However, the question arises as to the extent to which the ideas within the literature of system theory are shared within the literature of safety management. As discussed below, whilst there are linkages between the two domains, these are far from rich.

The systems movement appears to have emerged during the second-world war[49]. Bertalanffy (1968) suggests that it arose from the wartime need to co-ordinate the various heterogeneous disciplines involved in producing sophisticated armaments such as self-steering missiles and the atomic bomb. What was required was an approach which provided common ground but was not dominated by any one discipline and, further, one which permitted discussion and identification of issues with precision without the hindrance of the jargon and conventions of any one discipline. More recently, Senge (1992) makes a similar observation:

> "This is why systems thinking is the fifth discipline. It is the discipline that integrates the disciplines, fusing them into a coherent body of theory and practice."
>
> Senge, 1992, page 12

At the end of the second-world war, the concepts characteristic of the U.S. "system safety" approach appeared in publications, although the term "system safety" was coined six years later by Miller in 1954 (Currie, R., 1968). The system safety approach, it should be noted, arose in the same group of industries as mentioned by Bertalanffy and may be seen as something of a post-war zeitgeist. In support of this view, Bertalanffy, placing his own contribution (General System Theory - GST) in contemporary context, observes:

> "It was again one of the coincidences occurring when ideas are in the air that three fundamental contributions appeared at about the same time: Wiener's Cybernetics (1948), Shannon and Weaver's information theory (1949) and von Neumann and Morgenstern's game theory (1947)"
>
> Bertalanffy, 1968, page 15

[49] Which is not to say that people did not think this way before as there is much evidence in various literatures that they did. It is simply that the uniting of disciplines in this purposeful fashion had no precedent in this century.

However, whilst the word *system* is associated with these trends (ie, pragmatic safety on the one hand and theoretical systems approaches on the other) the sharing of ideas appears to have been limited to only the highest level (united by the concept of man-made systems as complex wholes consisting of highly integrated parts). Hence, whilst "system-safety" drew upon the system concept it did not in practice utilise the approaches developed contemporaneously by researchers such as Weiner, Bertalanffy and the others within the "systems movement" (Klir, 1991).

Johnson (1973) captures something of the distinction made above and it is part of his argument for a more comprehensive embrace of systems thinking and systemic approach to safety:

> "The term 'system safety' as used in the aerospace industries was an aspect of system engineering, which was usually a separable project or contract approach, and system safety was sometimes a separate contract. This gives rise to some semantic problems, in that system safety, as thus applied, is an early, finite phase in a project, rather than an on-going management effort characteristic of occupational safety. A system safety effort, despite its considerable virtues, was seemingly set apart from the ultimate on-going operations.
>
> Safety professionals have had difficulty in seeing how and where they could use system safety techniques. Further, the contract organisational form seemed to divorce system safety from the ultimate operational activities, even though the system safety tasks included operational requirements.
>
> Further complicating understanding of system safety by safety professionals was the substantial expense and new, sometimes complex, analytic techniques developed by system safety. Neither of these is a requirement for using a systems approach... We use the term "safety system" to describe the kind of on-going development which now appears needed in occupational safety."
>
> (Johnson, 1973, page 101)

In summary, there is a linkage between systems theory and systems safety but the connection is somewhat superficial - an effect of the zeitgeist in the organisational machines of sophisticated warfare. The contribution of Johnson and his co-workers is to integrate systems thinking (an aim shared by the systems movement generally) with the rigorous but limited systems safety paradigm.

Whilst the highly sophisticated safety management approach set out by Johnson (in SAN 821-2) is consonant with the thinking of the systems movement, his only

citations[50] from "their" core literature is "Cybernetics" (Wiener, 1948) and "The Human use of Human Beings" (Wiener, 1954). Thus the connections, whilst strong at the level of practical exposition[51], are tenuous at the level of formal theory and this is hardly surprising given the practical orientation of Johnson's work.

3.4 Limitations of MORT as a systemic model of safety management

The work of Johnson and his co-workers, as distilled in SAN-821-2, remains a valuable resource despite the passage of 23 years since publication. However, even when considered with the remainder of the MORT literature which is continuously updated and supplemented, this body of information still lacks unity. Paradoxically, whilst this lack of integration may be evident on paper the reverse is true of highly experienced MORT analysts among whom one finds abundant evidence of coherence. During the period of this research, this has been a constant source of optimism that the task that I had set myself was possible - if the men and women who use MORT daily have an internal representation (ie, a construction in the mind) of safety management which is manifestly coherent then this, in principal, can be externalised and explained.

To make this rather peculiar observation clear, one must first recognise that MORT conventionally refers to two quite distinct things:

> "The MORT principle has two meanings:
>
> 1. A total safety programme concept (viewed as a specialised management subsystem) focused upon programmatic control of industrial safety hazards, and
>
> 2. The actual logic diagram which displays the structured set of interrelated safety programme elements and concepts comprising the ideal management programme model called MORT."
>
> MORT Users Manual (Knox & Eicher, 1992, page 4)

However, I suggest MORT has not two but *three* meanings:

---

[50] Additionally, Johnson makes much use of the analytical approaches advocated by Kepner and Tregoe (1965) who cite the work of Simon (1960) which is generally associated with the "systems movement".

[51] A clear point of correspondence between Johnson and system theory is his repeated emphasis of "feedback" as essential to safety management.

a) The MORT logic diagram which can be seen as a classification system, that is, a taxonomy of functions required to assure safety performance;
b) MORT as a set of literature referenced to the MORT diagram *plus* a set of information systems (notably those maintained for the US Dept. of Energy by Scientech Inc.)
c) MORT as a mental model existing, in reasonably[52] isomorphic forms, among experienced MORT analysts but which has no coherent external (objective) analogue other than reduction to the MORT diagram.

MORT analysts distinguish between meaning (a) and (b) using the phrase "MORT in the *programmatic sense*"[53]. Meaning (c) is manifest when discussing a concept or observation whereupon the analyst will use MORT as a child might use a set of lego: picking appropriate pieces and joining them in order to capture something of the network of relationships necessarily excluded from a fault tree representation. Hence, when an analyst looks at a MORT chart they "see" a multi-layered (ie, extra layers of depth) chart, each layer joined to each other by a web of connections, and each layer criss-crossed by interconnections. As with perception in general, sense data are supplemented with the models/theories in the mind of the observer.

The perceptual need to supply "something extra" to descriptions and models of safety management systems is certainly not exclusive to MORT although it is more clearly apparent because MORT is uncommonly explicit (and hence are its limitations). Indeed, the principal utility of MORT is that it provides a defined (by its technical literature) vocabulary by which to communicate messages approximating the real-world complexity of the systems under discussion[54].

Concerning the analyst's contribution to MORT, set out below is a suggested classification where, items (1) and (2) can be seen as information relating to

---

[52] isomorphic literally means identical in structure. Hence, I certainly could not say with authority that these mental models are isomorphic. However, the fact that MORT analysts do understand one another suggests that there must be good correspondence.

[53] It is conceivable that a model derived using the approach set out in this thesis might be "validated" by studying experienced MORT users undertaking accident analysis. The augmentation of the MORT chart representation suggests an internal representation of a different order of complexity to the fault tree (which is still required as an aide memoire by even the most experienced analysts). Whether this "internal representation" is unvarying even within the conception of a given analyst is doubtful, however, the instability (ie, variation) may itself be lawful (ie, not a random feature owing to the limitations of human information processing)..

[54] As discussed later in regard to information theory - this observation is more than an interesting aside: there is a formal relationship between messages about a system and the structure of the system in question.

normal operation, and (3) as information relating to the development of the organisation in a primary sense.

I. the connectivity **between elements in the diagram** both as

   A. transacting elements (eg. Policy [MA1] "converses" with the Risk Assessment System [MA3], and;
   B. logical dependencies (eg. maintenance activities [a1-SD3] are determined by maintenance plans [a2-SD3] which in turn depend upon design [a2-MB3];

II. the connectivity **across organisational boundaries** both as

   A. within the same organisation
      1. transactions between the senior management, line management, supervisors and operative levels;
      2. transactions between departments;
   B. between organisations
      1. vertical transactions, such as those between contractors and sub-contractors
      2. horizontal transactions between "peer" companies.

III. Transactions of a **developmental type** such as those which lead an organisation from greenfield site to operational readiness.

Notwithstanding the consideration of coherence (as described through relations) the MORT literature represents a uniquely consistent and well-researched body of the safety management information. However, in order to render the knowledge available in the MORT literature amenable to the aims set out in section 3.1, a common basis must be found which allows MORT to mapped onto organisational structures[55]. As the systems literature is definitively concerned with the relations in identifiable wholes (such as organisations and organisms) it recommended itself as a likely source of unifying ideas.

Before embarking on a description of the main systems-theoretical concepts, another area which requires discussion is "safety culture" which is relevant

[55] An obvious extra requirement being a model of organisation which allows this mapping to be achieved.

insofar as it also speaks to the need for the "something extra" mentioned previously.

3.4.1 Parallels between "safety culture" and the present work

Safety culture, as defined by the HSE, is "...that set of attitudes and attributes in individuals and organisations which ensure that safety issues, as an overriding priority, receive the attention they warrant" (HSC, 1993)

In preface to this discussion, I should say that the treatment here of safety culture is relatively cursory for the reason that a deeper analysis would serve to confuse rather than clarify the systemic concepts which underlie both "safety culture" and safety management systems. More particularly, whilst "safety culture" does undoubtedly distinguish a relevant set of variables (such as motivation, risk perception, communication etc.) it does not address the internal relations of this set in manner susceptible to analysis[56]. In fact, put at its most extreme, the nebulous quality of safety cultural approaches to explanation can perhaps be misused to fill the voids in our theoretical picture of safety management systems - not by showing relationships between disparate elements but by obscuring them. Irrespective of these limitations, "safety culture" serves two very valuable purposes:

1. it elevates the importance of informal relations to that of formalised orthodox structures (ie, those explicitly specified and sanctioned by the organisation);
2. it provides those who use it with a vocabulary and legitimacy to express ideas and concerns that could not be articulated within the pre-existing set of conventions.

The second of these points is particularly powerful in its implications as it allows safety practitioners from widely different organisations, and also safety researchers, to communicate their experiences and perceptions independent of a particular context. It is entirely conceivable that safety advisors or managers from technologically diverse settings (eg. a nuclear power station and a dairy) could

[56] This is not to say that one can not make a list of, say, the attitudinal variables ascribed to safety culture (such as the detailed list presented in the third report of the ASCNI Human Factors Study Group [HSC, 1993] or Ostrom et al, 1993) and develop from it a psychometrically reliable instrument. Using methods such as principal components analysis (Factor Analysis) one might obtain a statistical view of the relations but not an explanation the mechanisms at work.

share useful information using the language of safety culture, an exchange that would have been exceedingly difficult to mediate otherwise.

When one examines the use of the term "safety culture", one usage that attracted my attention (generally noted in conversations between safety practitioners) is the separation of safety management system and "safety culture" in a manner suggesting the two concepts are related but distinct. An obvious analogy here is the distinction between motivation and competence as cofactors of individual performance. Here it could be argued that motivation is the necessary predecessor of competence as without the drive to acquire expertise, competence cannot be acquired. However, before the requisite level of competence is acquired no surfeit of motivation can deliver the required level of performance. This use of safety culture as the organisational analogue of motivation is undoubtedly useful and this is how it seems to be employed in HS(G)65. For example:

> "Organisations which achieve high health and safety standards are structured and operated so as to put their health and safety policies into effective practice. This is helped by the creation of a positive culture which secures involvement and participation at all levels. It is sustained by effective communications and the promotion of competence which enables all employees to make a responsible and informed contribution to the health and safety effort."
>
> (HSE, 1991, page 2)

> "A positive health and safety culture needs to be developed in which health and safety objectives are regarded by all as aligned to other business goals".
>
> (Ibid., page 10)

However, as well as the motivational sense of "safety culture", there are various other meanings. In the following quotation, culture as the antecedent of competence is present but so is competence as the concomitant of culture:

> "Control is the foundation of a positive health and safety culture and the management techniques used by those in positions of control are considered in more detail in chapters 4 to 6. All four elements [control, co-operation, communication and competence] are, however, inter-related and interdependent so that, for example, action taken to achieve consistent activity in each area is necessary to promote a **climate** in which a positive health and safety culture can develop and targets can be achieved. ...
>
> In organisations achieving success in health and safety, control is achieved by securing the commitment of employees to clear health and

> safety objectives. Managers take full responsibility for controlling all those factors which could lead to ill health, injury or loss. They provide clear direction and take responsibility for the working environment in which accidents, ill health and incidents could occur. This creates a positive atmosphere and encourages a creative and learning culture in which the emphasis is on a collective effort to **develop** and **maintain systems of control** before the event rather than on blaming individuals for failures afterwards."
>
> Ibid., page 16, *(emphasis added)*

Whilst the variety of different constituents and properties implicitly ascribed leave it unclear what makes "positive safety culture" distinct, some important notions are conveyed in the quotation.

1. The reference to "climate" as (I infer) the complex of control, co-operation, communication and competence from which emerges a positive safety culture (somewhat akin in character to the doctrine of vitalism). Seemingly, *climate* is a sort of *proto-safety culture*, all the ingredients but requiring time for them to "consolidate" (page 16, Ibid.) the authors mention a period of 5 to 10 years for this consolidation.

2. The emergent safety culture serves as a complex from which emerge the practical control measures which prevent accidents. It is not clear whether the practical control measures are part of the culture or an output from it. The IAEA (No. 75-INSAG-4, 1991) definition of safety culture embraces the control measures into the safety culture concept and the authors of HS(G)65 give no reason to suppose that this does not apply. Thus, part and parcel of "safety culture" is the competence of risk control and thus distinctions between organisational motivation and competence are no longer sustained.

3. The phrases "creative and learning culture" and, "develop and maintain systems" convey something self-sustaining - perhaps akin to Robens conception of a "self-regulating system" (Robens Committee, 1972).

What is distinguished here are two distinct realms of consideration:

a) item (3) above, concerns the mechanisms underlying the perpetuation of the safety management system once developed: **maintenance of the steady state**. The authors of HS(G)65 put considerable effort in describing some of the sub-processes (eg. supervision, training,

accident/incident investigation and auditing) but do not provide a coherent view of how these tie together (other than a simple schematic diagram)[57].

b) items (1) and (2) above, concern the **development** of an increasingly organised system. This view invests the word "consolidate" with quite profound meaning as it denotes the whole process by which the various constituents are organised into a coherent whole. This remainder of HS(G)65 has various references to this process but, in the final analysis, it is left behind a veil of mystery.

As previously asserted, safety-culture, whilst a useful notion in many respects appears to have little to offer the pursuit of an analytical framework in which to study the development and maintenance of safety management in organisational systems. However, what we take forward from this discussion, is the need to consider both the development and maintenance of safety management as determinants of safety in organisations. Additionally, the importance to safety performance of information exchange, outwith the formalised and orthodox organisational structure, is considerable and given inadequate weight within the MORT literature.

3.5 General systems concepts

In section 3.2 the concept of system was introduced in an informal way, chiefly to set the scene for the discussion in sections 3.3 to 3.4 but also to underline the danger of imprecision often lurking behind the informal use of the term. This danger is manifest in safety management systems in a number of ways. Firstly, safety management is a complex matter both in the number of processes and the interrelations of these processes. Secondly, as argued in chapter 2, certain fundamental concepts in safety are equivocal and, in some cases (such as acceptability of risk) are subjective and hence variable. In this section, systems theory is introduced more formally and, in section 3.6, safety management is reconsidered using these ideas.

[57] It has to be said that this criticism HS(G)65 is most unfair given that its primary aim is to communicate a large set of ideas to industry. Hence simple schematics are there to give a flavour of the relations between this set of ideas rather than to convey a deeper (and necessarily more complex) view of those relations. As is observed by the ASCNI working party on human factors

3.5.1 Closed and open systems

Classical science required an approach of reducing the complexity of nature to limits compatible with its analytical means[58]. One method of achieving this is to draw limits around the objects of the inquiry and to regard the result as absolutely isolated from its surrounds. This isolated set of objects in interaction is referred to as a *closed-system*. The closed-system approach and, more generally, the reductionist scientific paradigm with which it is associated remains the dominant influence in scientific study. Many researchers have spoken of the need for alternative approaches, most recently, Professor Levens[59] addressed the 1996 Edinburgh International Science Festival in the following terms:

> "...intellectual barriers to solving health, agricultural and environmental problems stemmed from the reductionist strategy of Euro-North American science, which chose the smallest possible object as the "problem", and then divided this into its smallest parts for analysis. This approach was historically justified in the struggle for scientific objectivity and could be valuable in research, but as the dominant research strategy it was responsible for the failure of many projects... All research had to make distinctions and recognise different kinds of processes and causes, but science often stopped there, without putting back together what it had separated. False dichotomies such as heredity versus the environment and thinking versus feeling had wrought havoc with scientific analysis, forcing choice between alternatives that were not mutually exclusive".
>
> The Times Higher Education Supplement, 18 April 1996

An early example, often cited in the systems' literature (eg, von Bertalanffy, 1968) of the limitations of the closed-system perspective, are the Driesch experiments in at the end of the nineteenth century. Driesch wished to render embryonic development accessible to the rigour and method of physics. At this time, it was known that if a four-celled embryo had some of its cells destroyed, the development of the embryo would be one sided (that is only the left, or right side would develop). In a series of experiments starting in 1892, Driesch took sea-urchin embryos (2-8 cells) and using the closed system-approach separated these cells, his expectation being that each isolated cell would develop into the particular body part destined for it. This hypothesis is very much characteristic of

---

(HSC, 1993) models which seek to embrace something of the real complexity of the safety management task "do so at the expense of comprehensibility" (page 18).

[58] My Aston colleague, Dr Mark Cooper, entertainingly refers to this as redefining problems in the "far-too-difficult" category.

[59] Professor of population sciences at the Harvard School of Public Health.

the closed-system view which regards the ***final state of the system to be wholly predictable from its starting conditions***[60]. Hence, in a two-celled embryo, one cell when isolated will result in the left side alone and the other the right-side alone. In practice this process obtained, in many of the cells, a normally developing embryo (ie, whole) rather than just a part. Driesch could not explain this result in terms of the physics and mathematics of the day and instead suggested the agency of a vital force "entelechy" which he extended into the philosophy of "vitalism"[61] (Gilbert, 1994).

The Driesch results were incorporated into a mechanistic philosophy called "holistic organicism" (sic) which, although over a century old, has much in common with the contemporary perspective of **open systems** regarded as axiomatic within disciplines such as cybernetics (and its sister disciplines within what is now referred to as systems science - Klir, 1991). As summarised by Gilbert, *holistic organicism*:

> "...refers to the views that (1) the properties of the whole cannot be predicted solely from the properties of the parts, and (2) the properties of the parts are informed by their relationship to the whole."
>
> Gilbert, 1994, page 580

This importance of this notion is further explained by Bertalanffy:

> "The meaning of the somewhat mystical expression, "the whole is more than the sum of the parts" is simply that constitutive characteristics are not explainable from the characteristics of the isolated parts. The characteristics of the complex, therefore, compared to those of the elements, appear as "new "or "emergent". If, however, we know the total of parts contained in a system and the relations between them, the behaviour of the system may be derived from the behaviour of the parts". We can also say: While we can conceive of a sum as being composed gradually, a system as [a] total of parts with interrelations has to be conceived of as being composed instantly."
>
> Bertalanffy, 1968, page 55.

[60] Parallels between this paradigmatic approach in biology and the system-safety approach are apparent: in the prediction of industrial system behaviour from a very detailed specification of its starting conditions.

[61] This theme is subject to further treatment at section 4.2.5.1, page 134, post.

The open-systems perspective deals extremely well with biological phenomena and, in particular, with the property manifest in living organisms – the *increase* in order, not the decrease. This is of no real threat to the second-law of thermodynamics because the long-term trend towards maximum disorder is preserved (living things are not immortal). From the perspective of physics, the biological tendency towards a greater level of order is achieved through the energetic interaction of the organism with its environment (importing ordered, low entropy matter and exporting disordered, high entropy matter). Thus, the organism can be seen as an *open-system* which attains a *steady state,* as contrasted with the closed-system which attains dynamic equilibrium (such as the uniform distribution of energy in a closed system).

Whilst caution must be exercised in making analogies, the embrace of the open-system-biological perspective within the organisational literature has allowed the consideration of organisational systems as analogues of biological organisms. Morgan's (1986) review of the organisational literature places the emphasis thus:

> "We began this chapter with the invitation to view organizations as organisms. And we have ended up with a review of some of the central ideas of modern organization theory. This is because most modern organization theorists have looked to nature to understand organizations and organizational life. The ideas identified provide an excellent illustration of how metaphor can open our minds to a systematic and novel way of thinking. By exploring the parallels between organisms and organizations in terms of organic functioning, relations with the environment, relations between species, and a wider ecology, it has been possible to produce different theories and explanations that have very practical implications for organization and management".
>
> Morgan, 1986, page 71

As I hope to demonstrate in chapter 4, cybernetics offers a means of formalising the basis of such comparisons to allow a firmer foundation than metaphor alone.

Although the closed-system and open-system perspectives have radically different expectations of the phenomena so rendered, it is (importantly) *not* the case that one is the *opposite* of the other. As Goguen and Varela (1979) point out, any open-system can be reduced further to obtain a closed system. Hence, the reductionist approach and the holistic approaches are concerned with the same things but at a higher or lower level of description (which when examining systems is typically referred to as a higher or lower level of *recursion*). However, it seems equally valid to apply the argument in reverse; an open system, which by definition

acknowledges rather than disregards the environment, may also be treated as a closed system at the next recursive level up. The new perspective obtains a closed system containing the previously nominated system and that which was previously nominated as its environment. This property of recursiveness and observer relativity is nicely illustrated by Goguen and Varela (1979) reproduced as Figure 3.2, below. In this figure, the observer's perspective is denoted by the symbol ✠.

Figure 3.2 Alternative system perspectives and configurations (from Goguen & Varela, 1979, page 33)

The narrative provided by Goguen and Varela, is (paraphrased) as follows: Item **1** of Figure 3.2, shows control of a system S_i by its environment E_i from the perspective of an observer shown by the mark; item **2** shows the autonomy of system S_i in its environment E_i; item **3** shows control of a subsystem S_i in a system S_{i+1}; item **4** shows the autonomy of a subsystem S_i in a system S_{i+1}; item **5** illustrates feedback control of system S_i by system $S'_{i,}$; item **6** shows communication between (co-ordination of) two subsystems S_i and S'_i within system S_{i+1}; and item **7** shows the co-ordination (ecology) of subsystems S_{i+1}, S'_{i+1} and S''_{i+1}.

The foregoing supports, I trust, the assertion made passim in chapter 2; that a system is a subjective entity. When we define *a* system it is clearly incumbent upon the perceiver to make clear their perspective. I believe that researchers in safety management have been labouring for many years with the implicit burden

imposed by the complexity of the situations they observed. This complexity is not merely the complication of the number of elements and their interconnections but also the unlimited number of viewpoints by which these elements can be perceptually grouped[62].

3.5.2 Systems and recursive logic

Concerning the definition of systems, Bertalanffy provides various forms of words (as well as various mathematical formulations) for example: a system is a set "of elements standing in interrelation" (Bertalanffy, 1968, page 36). Beer (1979) puts matters as follows:

> "A System consists of a group of elements, dynamically related in time according to some coherent pattern. That much seems to be essential. And it is not clear that we can say much more. The point that I find that I am most anxious to add is that this System has a **purpose**... It is you the observer of the System who recognises its purpose."
>
> (Beer, 1979, page 7-8).

Whilst Beer and many other authors are content to allow the term "system" to rest there, others like Koestler appear anxious to ensure that the recursive properties of open-systems are explicit. Koestler (1969) focusing on the issue of element or part writes:

> "A Part, as we generally use the word, means something fragmentary and incomplete, which, by itself would have no legitimate existence. On the other hand, there is a tendency among holists to use the word "whole" or "Gestalt" as something complete in itself which needs no further explanation. But wholes and parts in the absolute sense do not exist anywhere, either in the domain of living organisms or of social organisations. What we find are intermediary structures on a series of levels in ascending order of complexity, each of which has two faces looking in opposite directions: the face turned towards the lower levels is that of an autonomous whole, the one turned upward that of a dependent part. I have elsewhere proposed the word "holon" for these Janus-faced sub-assemblies – from the Greek *holos* – whole, with the suffix *on* (cf. neutr*on*, prot*on*) suggesting a particle or part."
>
> (Koestler and Smythies, 1969, page 64)

Throughout the systems literature, the convention used to convey the distinction emphasised by Koestler is as follows. System is the main focus of concern: the

[62] c.f. section 3.5.4 *Systems and problems of "organised complexity"*, page 80, post.

unitary whole which is the subject of analysis. Above system there is the Metasystem and, below the System, is the Subsystem. If the focus of our enquiry was an organisation we might nominate that as the System. When we consider its behaviour in relation to the industry of which it is a part, the industry might be regarded as its Metasystem. Similarly, if we were concerned with the internal behaviour of the organisation, we might consider the departments within it as sub-systems. Some authors have supplemented this simple 3-level approach with other epistemological schemes (eg. Klir, 1983 and Klir & Rozehnal, 1996).

3.5.3 Internal connectedness and *Variety*

Having dealt briefly with the issue of recursivity, the matter of internal structure requires further elaboration. A course grained view of internal structure can be obtained by considering the degree to which the elements are joined together in a system. As the study of systems is pre-eminently concerned with the relations between things, connections are determined by information flow rather than any physical medium. Hence two systems which happen to be physically joined are not necessarily connected in terms of information (eg. a party wall between two workshops).

The fully joined system. Imagine, if you will, a system as a set of things connected together to form a network. The things in this example are pressure vessels and their connections are pipes. Supposing our regulatory interest is to ensure a uniform given pressure throughout the network, this is automatically entailed by controlling the heat flow to just one vessel. This system is very simple to regulate. So long as we could control the heat flow to one vessel and have information about the pressure inside it, we could control the pressure in the entire system of vessels.

The partly joined system. Let us now suppose the system is slightly different: a pressure relief valve has been placed into each pipe. If the valves are all set at a very low value (ie, greatly less than the average in the system generally) the system remains very simple in its regulation - as all the valves are open. If, however the valve settings are changed, say to the mean pressure, the variation of pressure amongst the various boilers will become highly irregular as the tiniest variance in the setting of each valve will cause them to open at different times. If we freeze time at any point the network will reveal a pattern of closed and open

valves and vessels at different pressures. The fully joined (all valves open) system has become a number of more or less isolated subsystems (some valves closed).

We can rightly expect this system to settle down *eventually* so long as we assume no heat loss nor pressure loss from any vessel in the system to the environment. As soon as we allow greater variations in the vessels or valve settings we have gone from a simple whole (ie, simple in regulation) to a complex system of interacting subsystems.

Ashby (1956) suggests that the measure of complexity of a system is the number of distinguishable states that can be expressed in that system. This quantity he terms ***variety***. For example, throwing a dice can have 6 distinguishable outcomes, its variety is 6. Using the pressure vessel system, if this had 4 vessels connected in a way which allowed all possible combinations of directional flow of pressure among the four elements (ie, the valves are bi-directional), the number of distinguishable states would be 2 (steam flow or no flow) raised to the power of the number of possible directional connections (*n* elements x *n-1*); $2^{4(4-1)} = 4{,}096$.

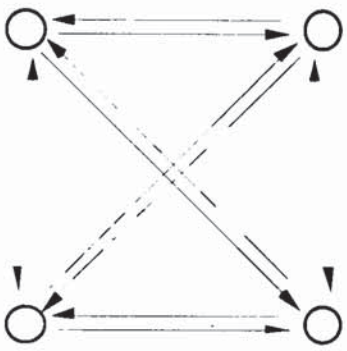

Figure 3.3 Fully joined four-element system

However, the complexity of the system quickly becomes much greater as the number of vessels increases: 5 vessels = $2^{5(5-1)} = 1.05$ million states; 6 vessels = $2^{6(6-1)} = 1.1$ billion states; 7 vessels = $2^{7(7-1)} = 4.4 \times 10^{12}$ states; 8 vessels = $2^{8(8-1)} = 7.2 \times 10^{16}$ states. In fact with just 30 vessels, the number of distinguishable states exceeds Eddington's estimate of the number of atoms in the universe.

Ordinarily, a large proportion of the maximum number of states of open-systems are in fact never realised and, of the remainder, the possible states do not have equal probability of being realised in practice. Hence, whilst the tendency is to define systems as wholes consisting of richly interconnected parts - only a fraction of the maximum possible number (as demonstrated in the scheme above) are expressed. This partial interconnection gives rise to localised properties - where certain behaviours of the localised parts do not affect other parts within the same system. This partiality of connection, or *localisation*, therefore acts as a set of *constraints* to the system as a whole. In any system that might be studied the action of constraint is that which makes it predictable, hence the laws of nature are

constraints upon the behaviour of natural systems. This conception, clearly has implications for how research into such systems might be undertaken.

3.5.4 Systems and problems of "organised complexity"

To ascribe to a new and radical analytical approach, the grand title of "paradigm shift" (after Kuhn, 1962) has become something of a cliché. However, taken as a whole, the range of approaches falling under the umbrella of General Systems Theory, do seem to qualify for Kuhn's description. Given below is a quotation from Weaver (1948) who elegantly captures the essence of this debate: the fundamental difference between the problems amenable to "classical" scientific approaches, those amenable to statistical approaches and, most importantly, the shortfall between.

> "The classical dynamics of the nineteenth century was well suited for analysing and predicting the motion of a single ivory ball as it moves about on a billiard table. In fact, the relationship between positions forms a typical nineteenth-century problem of simplicity. One can, but with a surprising increase in difficulty, analyze the motion of two or even three balls on a billiard table"... "But as soon as one tries to analyze the motion of ten or fifteen balls on the table at once, as in pool, the problem becomes unmanageable, not because there is any theoretical difficulty, but just because the actual labor of dealing in specific detail with so many variables turns out to be impracticable. Imagine, however, a large billiard table with millions of balls rolling over its surface, colliding with one another and with the side rails. The great surprise is that the problem becomes easier, for the methods of statistical mechanics are applicable. To be sure the detailed history of one ball cannot be traced, but certain important questions can be answered with useful precision, such as: On the average how many balls per second hit a given stretch of rail?" ... "This new method of dealing with disorganised complexity, so powerful an advance over earlier two-variable methods, leaves a great field untouched. One is tempted to oversimplify, and say that scientific methodology went from one extreme to another - from two variables to an astronomical number - and left untouched a great middle region"...
>
> "The really important characteristic of the problems of this middle region, which science has as yet little explored or conquered, lies in the fact that these problems, as contrasted with the disorganised situations with which statistics can cope, show the essential feature of *organisation*. In fact, one can refer to this group of problems as those of *organised complexity*"..."problems which involve dealing with simultaneously with a *sizable number of factors which are interrelated into an organic whole.*"
>
> Weaver, 1948, pages 541-543 (his emphasis)

A similar point is made by Senge (1992) who refers to organised complexity as dynamic complexity:

> "...*dynamic complexity*, situations where cause and effect are subtle, and where the effects over time of interventions are not obvious. Conventional forecasting, planning and analysis methods are not equipped to deal with dynamic complexity."
>
> Senge, 1992, page 72 (his emphasis)

The problems of organised/dynamic complexity seem to me to capture the nature of research into safety management. There is a considerable motivation in the field towards multivariate statistical research as this at least seems to address the fact that there are a great many variables and allows a mechanism to be implied. Hence the USNRC work referred to in chapter 1 provides us with "proof" of relationships that we thought intuitively to be connected (eg. communication is correlated with safety performance). However, whilst this is interesting (and a relief, as if such a relationship could not be found statistically, this would be a cause of great concern) it still begs the question of the mechanism between the variables: the linkage, direct, circuitous or coenetic[63], which mediates the statistical relationship. The point here is that system theoretical approaches aim to identify the mechanisms by which variables are connected and, where possible, to reduce these to mathematical form. Hence, the adoption of system theoretical means in this research is not to the exclusion of multivariate methods, far from it. However, given that this the domain of this work is located in "middle region", as Weaver would have it, our orientation needs to contemplate the mechanisms of organised relations as a priority above their statistical traces.

3.6 Safety management reconsidered through systems concepts

A theme in this text has been the need to supplement safety management descriptions with extra information: this was evident in the discussion of safety culture and more so in the discussion of MORT. The extra information is characterised by the interrelations of elements within identifiable wholes. Evidently, systems theory, which attempts to discover *generalised* principles underlying relations in systems is of assistance. Now, having introduced some of the basic concepts of system theory, safety management is reconsidered.

---

[63] *coentic* (pron. *sennetic*) from Sommerhoff (1950) meaning the common cause which determines the behaviour of two (or more) otherwise uncorrelated variables. A notion drummed-into users of correlational statistical techniques is that *correlation does not equal causation* - the possibility of coentic linkage is a very common alternative hypothesis.

3.6.1 Circularity and safety management system definition

In section 3.2 (page 63) the definition of a safety management system produced a circularity: a safety management system is [paraphrasing] a "system that controls a system". For Joseph (Ibid.) as a good logician, circularity is a sign of poor logic. However, the usefulness of logical division is a limiting case in which time has been reduced to zero; a linguistic equivalent of differentiating an equation with respect to time. The requirement of a systems viewpoint, particularly when accounting for the steady-state of an open system, is the inclusion of dynamic relations (system - environment, or between constituent subsystems). For example, adopting the scheme shown as No. 5 in Figure 3.2, I shall nominate S'_i as the safety management system (SMS) the purpose of which is to control an operational system S_i.

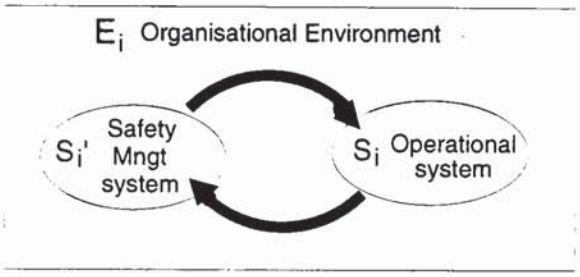

Figure 3.3 Cyclic relationship between a safety management system and an operational system.

The SMS seeks to constrain the behaviour of the operational system (OS) and this requires the transmission of information along the top arrow. However, the message transmitted must in part be a response to the behaviour of the OS as received via the bottom arrow of Figure 3.3. While the viewpoint of the SMS suggests that it has authority in the matter, its internal state is driven, to some extent, by the state of the OS. This scheme is a limited view of a steady state. When one then considers the internal organisation of both the SMS and the OS, this again is circular. As the OS develops and changes with time (for example, re-tooling to permit the manufacture of a different kind of component), if the SMS is to remain effective in its purpose it must change its structure. Otherwise the SMS would be unable to "perceive" the new set of states of the OS and equally unable to regulate the states of the OS to those acceptable.

The conclusion of this is that we cannot define a system using the conventions of logical division. This does not require us to *break* the rules but simply to acknowledge the limits of the realm to which the rules apply. The import of the open system/recursive system view is that circularity *necessarily* follows: the structure of a safety management system is *entailed* by the pattern of connectedness with the systems it seeks to regulate. **This concept, so far as I am able to assess, whilst implicit in much of the safety literature as has not been explicitly recognised before.** The emphasis is changed from Johnson's desire for integrated safety management to the recognition that it always has been. What is disputed, however, is the extent to which the configuration of relations by which this integration exists at all are optimum in the particular case. The hope of this research is to make some contribution to showing what these patterns of relations are in the general case.

3.6.2 Purpose and safety management system definition

As argued passim, the *purpose* of a safety management system is essential in its identification within the wider organisational system —it provides boundaries as to the wider system of which it is a part and suggests the likely processes and relations required in the service of its purpose. This simple yet far reaching systems concept provides an important contribution to the basic problem identified in the preface to this thesis:

> "To summarise: In order to begin the process of studying the development of safety management systems I needed a framework within which to collect and analyse the data. To develop this framework I first needed a model of what it was I aimed to study: safety management systems. However, implicit in this is study of the systems *which gives rise to such systems* and so my model would need to incorporate these antecedents also. This begs a further question: where does the "safety management system" begin and the system that designed it end?, and in general, can I define a safety management system with sufficient clarity as to identify elements which belong to it from those which do not?"
>
> Page 11, ante.

It seems at least that the imputation of purpose does provide the *basic rationale* for unravelling the organisational components and patterns of interrelation, in keeping with what I have suggested to be a basic criterion for research work in safety management: "a clear lineage back to an agreed set of principals" (page 60).

3.6.3 Recursiveness and safety management system definition

In chapter 1, I reported that my early attempts to analyse safety management led "with frustrating inevitability, to the inclusion of the state as part of the system" (page 24). Evidently, if it is granted that systems are holonic, that one should hunt from level to recursive level *is indeed* inevitable[64]. The "frustration" aspect is also to be expected because the general view requires not a single cognitive perspective (eg. a synoptic view at the level of the nation state, looking "down") but several viewpoints maintained in parallel. The fact is that, genius and insanity aside, maintaining several parallel viewpoints simultaneously is a cognitively tall order. Ordinarily just two simultaneous viewpoints is accompanied by the unpleasant stress of cognitive dissonance and results in the sufferer escaping by embrace of one and denial of the other competing view. What the systems concepts provide (and special models such as the Viable System Model offer in particular) is a means of describing the different epistemological levels (Klir & Rozehnal, 1996) of systems that allows one to *move* to different viewpoints whilst specifying the relations between them. By attempting to discover the principals underlying relations in systems, the emphasis is on *integration* rather than *competition* between different perspectives.

Recognition that all organised systems may be considered recursive provides a further theoretical guide to our expectations of safety management systems. This consideration, together with the points at 3.6.1 and 3.6.2, suggests that the view of safety management as "a specialised management subsystem" (MORT users manual, page 4) *must be regarded as very limited.* There undoubtedly are instances where organisations believe they have a management subsystem dedicated to health & safety and it certainly is not for me to deny that state of affairs. However, whilst such a subsystem may be a focus of activity it does not alone achieve the purposes of safety management - one cannot, for example, audit safety *into* a work system anymore than one can inspect quality *into* a manufacturing system. The safety subsystem contributes but relies ultimately on, inter alia, the organisational apparatus of control, resource allocation, and decision-making distributed throughout the whole enterprise. Hence in organisations without a safety subsystem, H&S purposes are still served albeit in a sub-optimum fashion by the service of other more explicit goals and by instincts of self-preservation.

---

[64] The fact is that this 'hunting' (in the engineering sense of the word) is extremely baffling until one realises the epistemological rules of the game. In effect, by following a line of reasoning starting at one level one does not gradually obtain a view of the next level, one switches spontaneously from the level one started at to the next.

Interestingly, Robens' view of the system whilst hierarchical is not necessarily recursive: that is, whilst he acknowledges the need to embrace both the state and industry (and all levels within a given industrial organisation) the regulatory properties of recursive systems is not used to advantage. This rather gnomic statement is explained in detail in chapter 4. For the moment, consider a simplistic scheme of system regulation as depicted below in Figure 3.4.

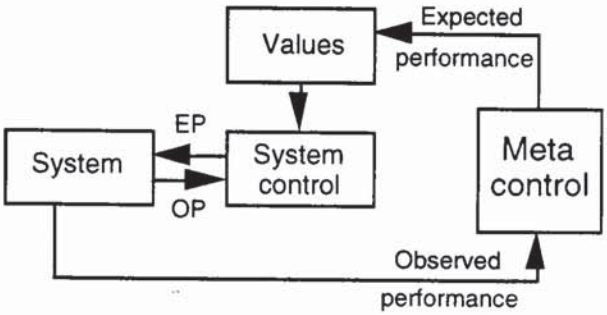

Figure 3.4 Simplified instance of Regulation at system level

The system's behaviour (ie, output) is controlled by the system controller. However, if all there was to this scheme was a system and controller the question arises as to the values of the output which are desired: how is the controller to know? They would have to be told (or calibrated if the controller was an automaton) at least once. The question then is who or what provides this instruction? In this scheme I have named the person or automaton that provides these values the *metacontrol.* In this simple scheme two levels of feedback are provided. The first, or lowest, is the feedback loop operating between system and system control: as the system behaviour deviates from the values required this state of affairs is transmitted to system control (the line marked "OP" meaning observed performance) which modifies the inputs to the system (by transmitting information along the line "EP" - expected performance) to bring its behaviour back within tolerances of these values. The second level of feedback involves metacontrol, this transmits information concerning changes in expected performance (values) to system control and assures that this transmission has succeeded by comparing expected performance with the observed performance of the system in question.

In keeping with the recursive concept, we can move one level up to include the meta-metacontrol as illustrated in Figure 3.5, below.

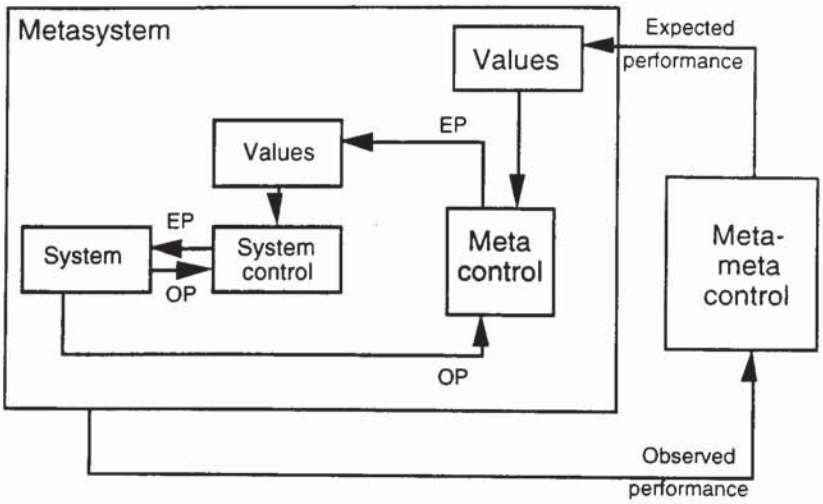

Figure 3.5 Simplified instance of Regulation at Meta-system level

Other than illustrating recursive control in the abstract, the importance of metacontrol is that it exchanges information with other systems *at its own level of recursion* which do not directly impinge upon system-control. Thus whilst the arrangement seems unduly complex (ie, why does meta-metacontrol require an intermediary, when it could transmit direct) this is because (a) there are perhaps many systems alongside but not featured in the diagram each of which are sending information and requiring regulatory inputs and (b) there are other metasystems with an interest in the system that need to be mediated and (c) the meta-metacontrol may be wishing to influence the behaviour of the systems regulated by metacontrol but this influence needs to be translated into an appropriate language as suggested in Figure 3.6 below.

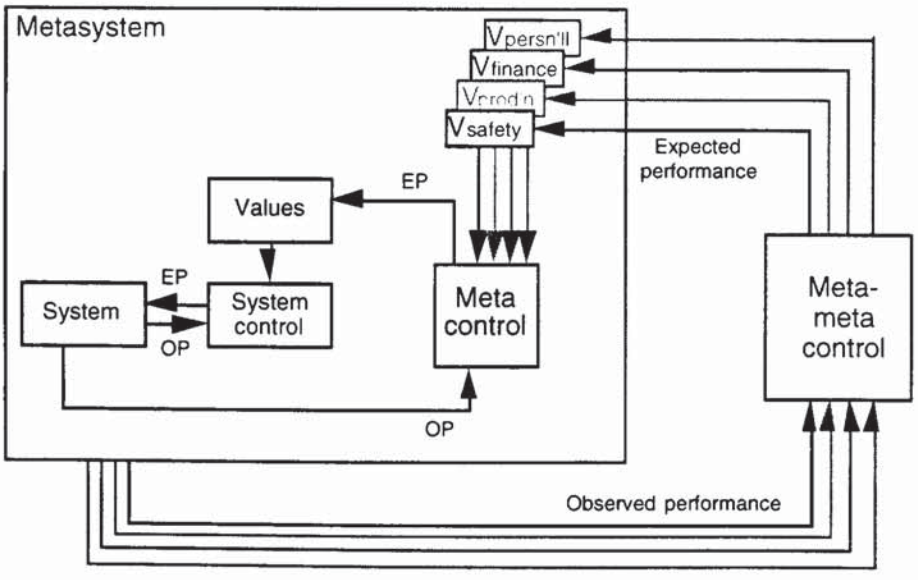

Figure 3.6 Regulation at Meta-system level showing the integrative role of Metacontrol

Here we see Koestler's "Janus-faced sub-assemblies": metacontrol looks "down" at the systems it regulates but can look "up" towards the assembly above it which provides the set of values it needs to obtain through regulation. The question now arises as to the origins and authority of these values.

3.6.4 Recursiveness and the definition of *Adversity*

Given the importance of *purpose* to system identity it is quite reasonable to suppose that the behaviour of the enclosed subsystems are likely to be regulated in keeping with the overall purpose of the whole. Hence, what is good is that which serves the purposes of the whole and what is bad is that which is inimical to those purposes. With recursion aforethought, what is true at the highest organisational level is also true at the next level down. Hence, each of the sub-systems within the organisational boundary will themselves have purposes. These subsystems, however, exist only to serve the purposes of the organisation, there would be no reason to create or maintain them otherwise. However, there comes a point in this unfolding when one encounters sub-systems which do not exist solely to serve the purposes of the organisation: these are called people. We will return to them shortly.

In the previous paragraph we started with the organisation as the highest level of system, this was, if you like, the "system in focus" (after Beer, 1979). In Figure 3.5 and Figure 3.6 we saw metasystemic inputs (observed performance) and outputs (expected performance or values) to the system in focus. The question now is what is metasystemic to the organisational system and the answer is highly variable as it depends upon the interests of the observer of the system. As argued previously (eg. in section 3.5.1, page 76) the purposes of a system are not objectively given but depend upon the interest of the observer: hence a Marxist may see organisations as instruments through which the state oppresses the people, an economist may see a financial institution creating wealth for the nation, a psychologist may see organisations as a means of reducing the anxiety of its members (eg. Jaques, 1955). For each of these observers, the metasystem which includes the organisation will be different: the economist may see the international economic system, the Marxist might see the nation state or, perhaps, a more global bourgeoisie. The point is, whilst these observers will certainly disagree about most things they may each be correct from their own viewpoint - what is material here is that whilst an organisation can be regarded as having an independent (ie, material) existence, that existence is manifest in many different dimensions of relations. Beer (1979)

uses the phrase "recursive dimension" to characterise this rather perplexing notion. In the health and safety dimension, and with the corporate level of the firm considered as the system in focus, the metasystem is the state as evident[65] in the following quotation:

> "HSE is very unlike a conventional Government Department. Working to HSC, it is a regulatory authority with detailed power and influence over the way industry works. Its powers need to be exercised expertly to help, not to burden, the process of creating wealth and welfare. Health and Safety are dimensions of every industrial process and to carry out their functions of protecting workers and the public over the whole range of industry, HSE has to be a scientific and technological body, first and foremost"
>
> (Rimington, 1992, page 6)

As before, the question arises as to the purposes of this system (ie, the state) which, taken broadly, is contained in this quotation as the creation of "wealth and welfare". The curious aspect of the HSC/E is, as Rimington intimates, it is unconventional insofar as it has considerable autonomy - thus whilst operating in recognition of the metasystemic purpose (ie, "to help, not to burden, the process of creating wealth and welfare") it has its own distinctive purpose. In the quotation above this is mentioned as "protecting workers and the public" and from a more recent document:

> "The aims of the Health and Safety Commission and Executive... are to protect the health, safety and welfare of employees, and to safeguard others, principally the public, who may be exposed to risks from industrial activity"
>
> HSC (1994)

Industrial organisations then, in pursuing their own purposes, receive metasystemic inputs from the state about the behavioural values to be maintained. The nature of metasystemic inputs is that they are of a different logical order to those generated at the level of the system. Values at the system level are discussible, open to enquiry and occasionally battled out by chief executives, directors, shareholders and worker representatives. Metasystemic values are, contrastingly, not discussible but absolute. The obvious point here is that we do not live in a dictatorship which might better be characterised by the issue of

[65] However, the matter is more complicated because as Rimington et al (1992) put it "HSE is only one of numerous actors in the 'industrial safety system'... In addition, insurers are particularly active as certifiers and inspectors of plant...". For the moment, we will confine ourselves to the recursive dimension primary in determining the values used.

absolute decree or ukase but in a democracy which is definitively based on discussion and, so far as the democratic apparatus allows, consensus.

The apparatus for health and safety is twofold: first, there is the democratic process in the ordinary way and then there is the specialised machine specified by the HSW Act 1974 (these are illustrated in Figure 3.7 on page 91). In the latter case, those subject to metacontrol have a role in determining the values applied to the regulated systems. In other words, industry has a role in the metasystem but is bound by whatever is finally agreed through the wider machineries of state. In this way, the apparent contradiction between metasystemic autocracy and systemic democracy is dissolved. In chapters, 4 and 5, the multiple recursive roles of individuals and groups will become an increasing feature of the emerging picture of safety management.

The view so far presented is of adversity as defined through the purpose of the system in question. What is *good* is what tends to bring about or *favours the purposes of the system*, what is *bad* is that which works *against those purposes*. The connection between systems theory and ethics (ie, moral philosophy) is quite apparent in this respect, as the notion of teleological[66] systems of ethics has been long established in philosophy. These are also referred to as *consequentialist* theories in recognition of the fact that they judge the moral correctness of an act by its consequences. The most widely known (and used) of the teleological theories is utilitarianism (eg. Mill, 1975) in which a morally right action is one that promotes the greatest good for the greatest number of people. Implicit in Mill's view is that the purpose of the state as system is to maximise the good of society and this is consonant with the view expressed by Rimington, above (implicitly, the purpose of the state is the creation of wealth and welfare).

At first sight, the teleological principal in ethics appears to be so strongly compatible with the systems viewpoint that one might be forgiven for not looking further. However, the main alternative to teleological theories of ethics are termed *deontological*[67] or *nonconsequentialist*. As one might expect from the name, non-consequentialist theories hold that, whilst consequences are a material

[66] Teleology is defined as: "The doctrine of final causes, esp. that natural and historic processes are determined not only by causality but also by their ultimate purposes, eg. attainment of the kingdom of heaven, human welfare etc. [fr. Mod. L. *Teleologia* fr. Gk *teleos*, end + *logos*, word]" Longman Modern English Dictionary, 1976.

[67] Deontology is defined as: "the science of duty or moral obligation [fr. Gk *deon (deontos)*, obligation + *logos*, discourse]" Longman Modern English Dictionary, (Watson, 1976).

consideration in determining actual behaviour, the essence of morality resides in the nature of the act. Hence there are certain acts or, more broadly, certain classes of behaviour, which exist as a duty. The most extreme deontological theories are those of Kant, which are succinctly summarised in the categorical imperative:

> "There is therefore but one categorical imperative, namely, this: Act only on that maxim whereby thou canst at the same time will that it should become a universal law."
>
> Kant, 1785. *Abbott translation.*

The distinction between the teleological and deontological is provided by Kant earlier within the same work:

> "In the meantime it may be discerned beforehand that the categorical imperative alone has the purport of a practical law; all the rest may indeed be called principles of the will but not laws, since whatever is only necessary for the attainment of some arbitrary purpose may be considered as in itself contingent, and we can at any time be free from the precept if we give up the purpose; on the contrary, the unconditional command leaves the will no liberty to choose the opposite; consequently it alone carries with it that necessity which we require in a law".
>
> Kant, 1785. *Abbott translation.*

Referring this back to the metasystemic-systemic arguments previously stated, it can be seen that, in general terms, the deontological formulation of "good" is consistent with the system's perspective of the metasystem input to it. In other words, whilst the teleological formulation of values (ie, the definition of good behaviour or performance) fits with the purposeful essence of systemhood[68], deontological formulation is consistent with the receipt of values from the metasystem above.

[68] a term coined by Rosen, 1986.

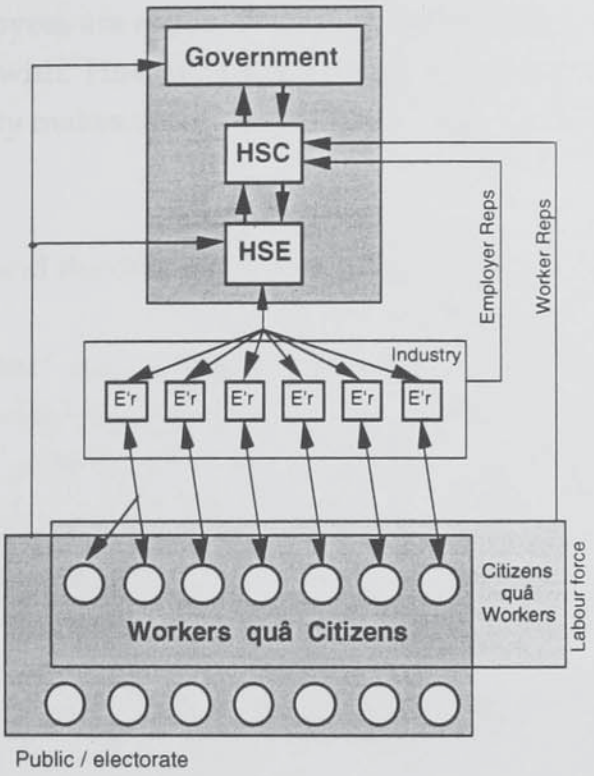

Figure 3.7 Schematic showing roles and information flow determining UK H&S regulation (E'r meaning employer]

In terms of Figure 3.7 above, the HSC is the focus of a value formulating system which through discussion and agreement of purpose (teleological) new policy and regulatory values are developed. These are received in the "industry" box (consisting of individual employers within the purview of the HSE) as deontological values. This is to say they need not be conceived of in terms of the purposes of the organisation but are categorical in nature - a matter of *prescribed duty*. Generally, neither people nor institutions gladly accept metasystemic inputs especially those arising from a system of values outwith their own set of purposes.

For this reason, the HSC collaborative method of value formulation appears successful. As Rimington puts it

> "This close involvement of all the parties explains the high degree of "stickability" of health and safety legislation introduced since 1974, and helps give industry confidence in the standards to which they are to work"
>
> Rimington, 1992, page 15.

What is clear is that a definition of adversity is implicit in the HSW Act and its nature is deontological: there exists a duty upon employers to ensure "the health, safety and welfare at work" of employees. In system language, the health, safety

and welfare of employees are *essential variables* (Ashby, 1956) that the organisation must concern itself with. However, the qualification of this duty by the test of *reasonable practicability* makes the *values* of these essential variables a complicated issue.

3.6.5 Recursiveness and the derivation of H&S values

The Kantian "hard-line" deontology described earlier was not arbitrary, but the product of a sophisticated argument. Kant's position is essentially that all rational beings are conscious of being under an obligation to act in particular ways. The reciprocal of this is that rational beings have the right to expect treatment from other rational beings in particular ways. The outcome of this is contained in Kant's statement:

> "Accordingly the practical imperative will be as follows: So act as to treat humanity, whether in thine own person or in that of any other, in every case as an end withal, never as means only."
>
> Kant, 1785. *Abbott translation.*

I include this, not because Kant's argument is irrefutable (despite the attractiveness of its humanitarian outcome), but simply because its shines a strong light on this debate. Looking at Figure 3.7, the intersecting lower boxes are labelled "Workers quâ Citizens" and "Citizens quâ Workers". By *Workers quâ Citizens* I mean people who are being considered as citizens within a democracy. On the other hand, *Citizens quâ Workers*, places the emphasis in the reverse: people considered primarily in their organisational context.

Using the Kantian distinction, *Workers quâ Citizens* are **ends in themselves** within a democratic state. In contrast, *Citizens quâ Workers*, are **means to organisational ends**.

The regulatory problem for the Government is how to achieve a practical balance between the good of the individual and society whilst serving the purpose of the "creation of wealth and welfare" for the state as a whole through business organisations. This observation is explained further by Dworkin[69] who, speaking in regard to the behaviour of Government and role of the Judiciary (judicial review in particular), comments:

---

[69] Ronald Dworkin, Professor of Law, Oxford and New York Universities. The quotation is from an interview he gave to the BBC, broadcast on Radio 4 in February 1996.

> "I do think there's an important distinction between two kinds of questions that a political community must somehow answer. The first are questions of policy. What economic institutions should be established? What tax policy should be put in place? What kinds of business organisations, what sort of redistribution of wealth in the interest of efficiency should take place? These are the kind of issues that I believe that in a decent running democracy should be left to a representative institution like Parliament. Now there's a second kind of issue that really cuts across this, and that's the issue of what **means** a government in all fairness may use to achieve these ambitions of policy. These are issues about, as it were, **how far the interests or position or dignity of an individual can be sacrificed to these overall goals.** And these are questions of principle. And it's these questions that I think are essentially adjudicative questions though they're also moral questions. Judges aren't the only people who should be asked to think about these. I have to be absolutely clear about this. I mean obviously a government should be thinking about what's fair and right and decent too and so should we all."
>
> Dworkin, 1996 (*emphasis added*)

To some extent, the regulatory values for H&S are driven politically after the occurrence of major accidents[70]. The nominal situation, however, is very much more utilitarian:

> "143 The whole **philosophy and structure** of British health and safety legislation is designed to ensure that the costs which are imposed on business are proportionate to the benefits obtained. The main duties placed on employers under the Health and Safety at Work etc Act 1974, and under much of the secondary legislation made under it, include the proviso that risks should be reduced "so far as is reasonably practicable" in the specific circumstances.
>
> 144 Most of the important recent legislation, especially that implementing EC requirements, has been based on the related concept of **risk assessment and control**, which is again designed to focus control measures on areas of greatest risk."
>
> HSC, 1994 (*emphasis added*)

The good of society is achieved (in terms of creation of wealth and welfare) but at the cost of individual citizen's lives or quality of life. Thus, in itself, the death or major injury of a relative few is not necessarily falling foul of the utilitarian purposes on the state, indeed that the number of casualties is not larger might be

[70] the report on Piper Alpha disaster called for "goal-setting" regulations and a corresponding increase in the regulatory expertise and resources of the inspectorate (paragraph 22.12 Piper Alpha Report, Dept of Energy [UK], 1990)

viewed as a quite in keeping with the greatest good of the greatest number argument.

A curious outcome of the UK position is that the ALARP[71] principal, which underlies the standard of the duty of care for employer of employees, whilst providing an elegant basis for uniform **risk acceptance**, creates a de facto *pluralistic* standard of **risk exposure**. It appears inevitable that in practice what is reasonably practicable for one employer may not be reasonably practicable for another - creating a diverse profile of risk exposure across industry[72]. When discussing this with the former Director General of the HSE, Mr Rimington presented an interesting justification of ALARP. His explanation was that the state, by fixing an upper limit on tolerability (an individual risk of fatal injury of 10^{-4} per annum), can rely upon the "geometry" of accidental injury ratios to effectively extend the tolerability threshold across the scale of harms. ALARP ensures that the outcome of this arrangement is generally bettered and seldom made worse. However, the premises of this argument are suspect.

The first area of weakness concerns the "geometry" of accident injury ratios which are highly variable. All other things being equal, firms with steep gradients (ie, ratios such as 1:56, the numbers of 3-day lost time accident to the number of minor injuries[73]) have a higher tolerability function than those with shallow gradients (eg. 1: 10)[74]. Which gradient is the basis of regulatory acceptability? If it was the hospital (1:10) then the construction firm (1:56), whilst perhaps meeting the 10^{-4} goal for risk of accidental fatal injury, exceeds tolerability for "lesser" accidental consequences.

The second area of weakness, is the assumption that the accident ratio represents a relationship between accident frequency and severity where this vector is homogenous. The fact is that while accident severity can be considered as a continuous variable, accidents can not. As Hale & Hale (1972) point out, there is no evidence which supports a view of accidents as causally homogenous. The upshot of this is that whilst accident ratios can be easily obtained they are in fact *accidental injury ratios.* By placing the emphasis on the fact that it is the outcome of

---

[71] ALARP principal: that risks should be reduced *As Low As Reasonably Practicable.* This is conveyed by, inter alia, Section 2(1) of the HSW Act by the inclusion of the phrase "so far as is reasonably practicable".

[72] In conversation with enforcement officers, concern has often been expressed about this

[73] The Construction case study in HSE (1993)

[74] The Hospital case study in HSE (1993)

accidents and not the causal factors which is represented, the dubious nature of the "ratio" statistic becomes more obvious assuming as it does two factors, neither supported by evidence: (a) that the ratio is stable with respect to time and (b) that the injuries are causally identical irrespective of severity. Concerning (b), Johnson (1980) makes the point succinctly:

> "There is widespread acceptance of the premise that prevention of minor accidents prevents major accidents. Yet as can be seen in Figure 7-4, and in many other tests, this can be true only for some homogeneous class of accidents."
>
> "Figure 7-4 [*Right*] Log-log plot of frequency and severity used to determine slope of risk line"
>
> Johnson (1980) pages 96-97

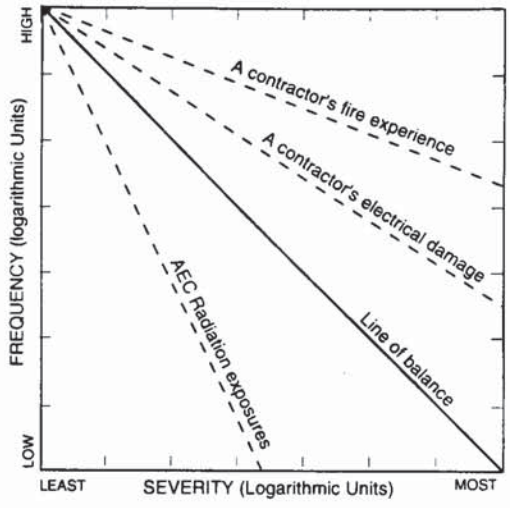

It occurs to me that whilst ALARP is a very useful test to be applied in the teleological determination of health and safety values at the organisational level, the complaints voiced above suggest that the single-point tolerability figure developed at state level provides a metasystemic standard for that point only. In fact, this deontological nature of this single-point figure seems aberrant with regard to the rest of the regulatory approach to H&S in the UK. Pursuing this argument, the organisational application of CBA (to determine teleologically the de facto level of risk exposure for risks that are anything less than intolerable and anything more than negligible) provides a rational method of determining de-facto levels of risk but cannot "square the circle" of diverse risk exposure across industry. Thus, the observation first made at page 32, ante, is, repeated here: the role of CBA is clearly useful to *inform* the choices of decision-makers but, deontologically, cannot be relied upon to *determine* such decisions. Lastly, if it is the case the State has delegated to organisations (in Britain, via S.2(1) HSW) the authority to determine the level of industrial risk to which citizens are exposed, then it is to the State that the accountability for such decisions[75] refers. As argued in the next section, the propriety of this arrangement is open to question.

[75] if properly reached through a transparent method such as CBA or equivalent teleological calculi. If an organisation were to fail in this respect, presumably they would incur criminal liability.

3.6.5.1 People as *ends* and People as *means*

The foregoing discussion allows us to consider the ethical issues raised in chapter one more simply. The systems theoretical approach and the key dichotomisation of normative approaches in ethics allow H&S values to be considered thus:

> **Deontological health and safety values:** duty of care to absolute standards regardless of other purposes which impinge upon the duty bearer.
>
> **Teleological health and safety values:** duty of care to a relative standard as influenced by other purposes which impinge upon the duty bearer.

It should be noted that whilst these two approaches are distinct, the results of their operation are not necessarily radically different: the deontological approach *emphasises* people as ends, and the teleological *allows* people as means. In this perspective, the three quotations offered in chapter 1 can be considered again:

- When Johnson (1980) said "moral concepts have slight effect on changing practical behaviour" (page 17, ante) he was speaking from within a teleological framework about the concept of deontological values. Organisational purposes, Johnson recognised, are definitive in the matter and arguments based on standards unrecognised corporately have little meaning;

- When Weldon (1962) asserts that "where a genuine moral choice is involved no accumulation of non-moral considerations is of decisive importance" (page 13, ante) he is arguing from a deontological basis - people as ends in themselves. From the H&S perspective, the truth of Weldon's assertion depends upon the highest authority impinging upon the organisation. Generally this means which nation state one's organisation is subject to; and

- When Webb asserts of the 1900s "This century of experiment in factory legislation affords a typical example of English practical empiricism. We began with no abstract theory of social justice or the rights of man" (page 17, ante) he is contrasting the teleology inherent in *practical empiricism* with the deontological concepts of human rights (and the fact that we English have a tradition evident in our common law and our moral philosophy which sets us apart from the rest of Europe).

Of the many issues raised in this thesis, this is both the least tractable and, arguably, the most fundamental. What is at question here, ultimately, is the moral authority of the state in relation to deciding the value of human life. It is disappointing to find myself still in the grip of the same English practical empiricism that Webb complained of eighty years ago, but the service of this research requires me to remain thus bound for a little while yet. Whilst I cannot pursue this topic further within the scope of this thesis, I would at least like to register my assessment of this situation: simply that this matter will not be long in becoming a focus of public debate. Consider three contemporary legal opinions in relation to H&S law, the law of tort and the law considered generally:

> "The standard of reasonable practicability is not found in the Directives, though it occurs in places in the implementing Regulations. Where the use of that standard (in implementing or pre-existing legislation) would diminish the standard of care required by the directives in an identical situation, the lower standard is open to challenge. The preamble to the Framework Directive specifically provides that health and safety at work "is an objective which should not be subordinated to purely economic considerations", a principle which is inconsistent with the element of cost present in the assessment of what is "reasonably practicable".
>
> Hendy & Ford, 1993, page lxxxiii.

> "Tort is concerned with wrongs. Gradually the common law judges seem at last to be equally concerned with rights. The story of the next decade may be one of judicial incorporation of the European Convention of Human Rights at least in part via the medium of the law of torts".
>
> Brazier, 1993, page 13

> "I think it is perfectly true, and in my case very welcome, that English judges have been breathing new life into our unwritten constitution. And they've been doing it through the common law. And they have been looking into the common law, which traditionally has been at least ethically aimless... and injecting ethics into the common law. Principles like the principle of free speech or the principle of equal treatment without unfair discrimination".
>
> Lord Lester, 1996

3.6.6 The translation of values into actions - regulation and self-regulation

The question of the regulatory values pertaining to a given system, whilst made vexatious through the duality of ethical standpoints, are certainly not the sole deciding factor in determining the *actual* level of risk within an organisation. In other words, one cannot legislate the actual level of risk, merely the desired level. Addressing this point, Johnson (1980) observes[76]:

> "Some say safety is compromised by value conflicts when safety motivations are of insufficient strength. However, case histories do not pinpoint managerial factors in terms of values. In the best companies major factors are weaknesses in safety analysis, failure to provide successive, alternative countermeasures. ... Defined criteria which management ought[77] to use in assessing safety may also measure the strength and nature of beliefs in practical ways."
>
> Johnson, 1980, pages 224-225

Johnson's point here seems to distinguishing between espoused-values and values-in-use (Argyris, 1988), specifically, identifying implicit organisational values through the explicit behaviour of their agents (decision-makers). This may look rather like putting the cart before the horse but is very much in keeping with the systems viewpoint which strongly emphasises the circularities apparent in open systems. Although it is customary to consider metasystemic values (ie, criteria) as being prior to the system to which they are applied it need not be so. As Bertalanffy (page 74, ante) points out, when taking a systems perspective assuming simultaneity (of system elements and relations) is an aid to thought. A related point is made by Beer (1979, page 9) "both the nature and the purpose of a System are recognised by an observer with his perception of **what the system does**" (his emphasis).

---

[76] This does rather beg the question of why these practical weaknesses are tolerated at the most senior level (ie, the question of values again). One can only hang onto a hypothesis of organisational incompetence so long as it extends to a level of incompetence in the acquisition and communication of "bad news" to the most senior level in the organisation (as was argued by the directors of Townsend Thoresen as part of their defence in the legal action following the Zeebrugge disaster).

[77] Johnson, whilst seemingly eager to shelve the issue of ethics and concentrate on the practical matter of how to deliver safety performance, here provides a demonstration of "Hume's Law". Hume (1740) noted that when reasoning on issues of a moral nature, authors tend to suddenly use the word "ought", thus relying on extra information than that presented in their preceding argument. Johnson does something of this kind when saying that management ought to use defined criteria... in other words a general basis for judgement which has decisive authority in some sense beyond the level of the management who use it. The question of values, and where human welfare is concerned- moral values, cannot be divorced from the practicalities of obtaining system performance which meets these criteria.

Before proceeding I would like to summarise the points made in the foregoing, presenting as they do a rather entangled set of considerations.

1. The boundary of a system is entailed by its purposes.
2. The teleological definition of good and bad follow from the purposes of a system. These can be considered as the *essential variables* of the system.
3. The purposes of subsystems are entailed by the greater purpose of the system in which they cohere.
4. Depending on their interest, different observers will perceive different purposes and, hence, different systems.
5. Given the recursive structure of systems, the different purposes will yield different sets for subsystem membership and imply different metasystems. These have been referred to as recursive dimensions.

When we consider a given physical manifestation (such as Aston University) the constituent elements can be argued to be contained within the bounds of the campus. If, however, we want to understand the relations between the constituent elements we need to impute a purpose. If we chose research as the purpose, a large number of the geographically bounded constituents would appear redundant or simply disappear from consideration (eg. staff who do not contribute to the research effort, several thousand students on taught courses). Equally, we might be mystified by the sudden disappearance of researchers with large sums of university funds (eg. Aston researchers working overseas). The essential variables for this system might include, the number of papers published, the research funding attracted, the number of research degrees awarded, and perhaps less tangible variables such as the contentedness of the research staff.

If instead we chose teaching as the purpose we might be puzzled or infuriated by the sizeable number of staff who spend considerable effort on tasks other than teaching - after all, spending resources on other than serving the teaching purpose, constitutes an adverse state. Again, there would be puzzling absences and mysterious episodes of cash arriving (eg. teaching and training off campus). The essential variables for this system might include, the number of students successfully obtaining degrees, the number of applicants seeking places here, the speed by which graduates obtain relevant employment, quality measures of the teaching provided, productivity of teaching staff.

Turning to metasystem identity. The research view of purpose will reveal metasystems such as the research funding councils (eg. ESRC, EPSRC, NERC etc), the various editorial boards of learned journals might also be considered as part of a broadly drawn metasystem. As for teaching at undergraduate level, the metasystem would certainly include the HEFC but also various professional institutions (eg. the British Computer Society, the Institute of Mechanical Engineering etc).

Given that the two purposes are served by some mutual elements there has to be some apportionment of resources which optimises both (although probably maximising neither). There are probably a great number of purposes which could legitimately be perceived as shared the same people, bricks and mortar - each of these has to entered into the optimisation equation. Hence further considerations need to be added to those enumerated above:

6. For each recursive dimension pertaining to a given set of elements, a separate purpose can be imputed.
7. These interdependent purposes all yield sets of essential variables.
8. These sets of essential variables need to be balanced one with the other.
9. The task of management can be characterised, therefore, as primarily a task of optimising a multivariate equation.
10. Within this multivariate equation, some of the weightings of the variables will be set by the relevant metasystem thus removing degrees of freedom from the weightings assigned to the other variables.

The points above are quite in keeping with Johnson's further discussion of the interaction between establishing and assuring H&S values (quoting Schroeder, 1970) he continues:

> ""It is not simply a question of the value of safety... It is rather a question of the nature of this belief... 'Immaturity'... is the level or complexity of conceptual structure for processing information. The initial question is to determine the number of independent classes (or scales) of information the person selects as being relevant. Weighting should also be assessed."
>
> Such a method would also test the limits of values by examining a list of alternatives. When safety clashes with budget and schedule constraints, the trade-off criteria and mechanisms are weak.... Management must require confrontation between alternative solutions in its bases for choices and decisions."
>
> Johnson, 1980, pages 224-225

Schroeder's assertion is that perceptual[78] sophistication operationalises a set of values. In other words, it is a fine thing to believe safety to be important but another to have the a conceptual structure able to make judgements as to the salience and importance of safety related information. This idea is strongly related to Simon's concept of "bounded rationality" (Simon, 1962 and 1982). Simon argues that organisations can never be fully rational as their human elements, labouring under the multiple burdens of limited cognitive ability, time, information and indistinct values cannot arrive at optimal or fully rational decisions or sets of policy.

As well the congruence between the ideas of Simon and Schroeder, the latter's metaphor of *immaturity* is strongly reminiscent of Rousseau's *sleep of reason* (Grimsley, 1973). Rousseau used this as a simile for the essential condition of childhood before the conceptual sophistication of adulthood gives rise to a capacity for reasoning and responsibility. In keeping with Rousseau, *awaking* from the sleep of reason may serve to encapsulate Robens'[79] intentions for the an industrial move to a "more effectively self-regulating system" (Robens, 1972, paragraph 41) from what Booth (1994) has described as a condition of "unthinking compliance".

Whilst Robens does not define what he means by self-regulation, what is apparent in the report is a heavy emphasis on **responsibility** at **all levels** of industry. *All levels*, he makes abundantly clear[80] includes all employees and not merely those beyond a certain level of seniority. With responsibility and involvement as the main ingredients of Robens' recipe for self-regulation, the regulatory environment provided by the state was very much in terms of goal-setting. Relating this back to the views of Schroeder and Simon, it is clear that there is more to self-regulation than simply assigning responsibility for H&S to everyone involved[81].

---

[78] That is, conceptual richness logically underpins perception. It is taken as axiomatic that an observer cannot perceive matters for which he has no theory (ie, set of concepts) nor make distinctions beyond those available through the current theory of what is observed.

[79] To pursue this metaphor further: The alarm clock designed by Robens and built by the HSW Act (1974) did not go off until the MHSW Regulations (1992). Even now, much of industry sleeps on.

[80] Paragraphs 59 to 71 of the report (Robens Committee, 1972) are devoted to the subject of the involvement of workpeople in the process of self-regulation.

[81] Whilst I have never been a fan of safety posters, one for which I reserve a special distaste runs "Safety is everybody's business" which manages, with an élan characteristic of the genre, to translate a vital concept into a very hollow slogan.

3.6.6.1 Self-regulation and the Three Senses of Responsibility

Authors always run the danger, when taking a first-principles approach, of discovering something that whilst new to them is quite commonplace to everyone else. Thus at the risk of demonstrating obtuseness, I have to admit to my delight in discovering the felicity with which the word "responsibility" opens up the issue of self-regulation.

According to Shaw and Barry (1989) Moral responsibility can be considered in three ways; these are given below in Table 2. I do not see any argument for restricting the scheme they propose to moral issues only. From the foregoing discussion of the definition of adversity it is, I trust, evident that the principles involved generalise across different instances and classes of adversity.

Sense One **Responsible *as in* accountable for past actions**	***Sense Two*** **Responsible *as in* having a duty (responsible for somebody or something of value)**	***Sense Three*** **Responsible *as in* able to make informed decisions**
To behave within given constraints or beyond some minimum standard (as given by law, the organisation, and society).	To have a role derived from the specific requirements placed upon the actor by higher authority	In this sense, responsibility requires requisite: (1) knowledge of salient facts (2) cognitive ability and (3) the ability to act.

Table 2 The three senses of *Responsibility*

Within this scheme, the various senses of responsibility are each interrelated; Sense two requires sense one to provide some basis for judgement (ie, sense two without the qualification of sense one might be referred to as "notional responsibility"). However, the real importance of this is that, whilst each sense of "responsibility" is distinct, the scheme is **dominated by sense three**. If not responsible in sense three, one cannot be considered responsible in sense one and two.

As outlined in Table 2, to be responsible in Sense three requires various minimum criteria to be met. Should any of these criteria be absent or deficient the agent in question cannot be responsible in sense three or, because of the domination of sense three, responsible in either sense one or sense two. The criteria for sense

three, enumerated below, concern: the availability of information to the agent; the agent's mental model; the cognitive ability to integrate the information with the model; a set of values with which to judge the situation, and; the ability to act.

3.6.6.2 Criteria for Sense three of *Responsibility*:

I. availability of the "salient" facts. That is, information which is relevant to their role;

II. a mental model of the situation. This is strongly associated with Schroeder's idea of "conceptual structure" and Simon's "bounded rationality". This mental model:

   A. provides the definition of "salience" for information (theory provides the basic cue to what is relevant and irrelevant among the data available to the senses)
   B. must match the variety of situation (otherwise the agent will be unable to discriminate between different and possibly consequential states of the situation)
   C. provides the basic configuration for actions (a similar notion to the skill– and rule– levels of Rasmussen's SRK scheme. (Rasmussen, 1980);

III. the cognitive ability to manipulate the model so as to derive alternative decisions. This is more at the knowledge-based performance level of Rasmussen's SRK scheme;

IV. a set of values to permit judgement of
   A. the need for action
   B. the optimum result of action;

V. the ability to implement the decision - *to act*. This concerns the agent's ability to affect the situation for which s/he has responsibility.

Whereas the domain of individual moral responsibility places the emphasis on the individual to acquit themselves against the criteria given above, the domain of the organisation must assure that these criteria are met. This is the rationale behind my suggestion in chapter one, repeated below for ease of reference:

> "The authors of ACSNI III suggest that "it follows that to make managers manage safety better it is necessary to make them more effective managers"... "I would emphasise that it is the *organisational context* which largely enfranchises the efforts of people, which allows managers to "add value" to the efforts of subordinates and which provides the value system within which to make decisions".
>
> section 1.2.7, page 26, ante

3.6.6.3 The three senses of *responsibility* considered organisationally

sense one: accountability against *agreed* measures of past performance or current behaviour (which might include measures of productivity, behavioural policies such as equality of treatment and indeed measures of risk or H&S audit variables);

sense two: the assignment of an operational process, geographic area or organisational function to an individual or group (the nominated responsible agent or controller of that process, area, etc.) by a higher authority - the *metacontrol* introduced in section 3.6.3;

sense three is specific to the particular organisational situation that the "responsible" agent is assigned to control but the general principles are the 5 criteria listed on the previous page. Summarily:

1. Relevant information from monitoring systems or directly from machine interfaces;

2. a mental model of the system acquired through education, training and experience - much of which provided through the resources of the organisation;

3. the cognitive ability to manipulate the model to deal with novel or unusual system states (this would include problem solving skills);

4. a set of values - performance criteria including quality standards, risk acceptance criteria or, at a higher level, policy and directives (all of which arrive from the metasystem);

5. the ability to implement the decision - *to act.* This combines competence with organisational resources made available to the agent by the metasystem... *power* commensurate with *responsibility* .

What becomes clear from this account is the involvement of the next level "up" both logically (as a recursion) and organisationally (as seniority). Equally, it should be appreciated that this recursive pattern continues until the board or owner are included thus binding the organisation together though line and staff responsibility. This organisationally-rendered view of the "three senses" is illustrated at Figure 3.8 below.

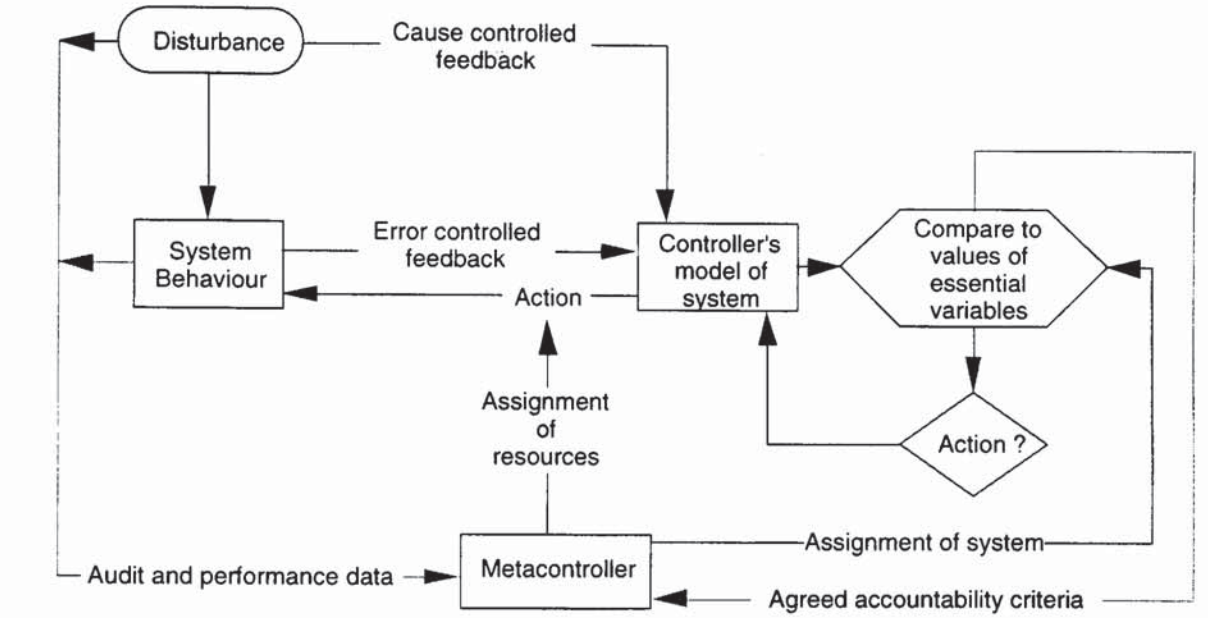

■ Responsibility as accountability

■ Responsibility as rationality

■ Responsibility as assigned duty

Essential variables - Those system variables that must be maintained within a pre-set range

Model of system - The mental or computational model of the system maintained by the controller

Error-controlled - data about change in the values of the essential variables

Cause-controlled - data about change in the parameters of the system environment

Figure 3.8 The 3 senses of responsibility in organisational perspective

Figure 3.8 contains certain terminology that will be explained fully in the next chapter.

In passing, another feature of Figure 3.8 not mentioned in the foregoing is the box marked "disturbance"[82]. This is included to alert the reader to the fact that whilst our natural inclination may be to concentrate on the states unfolding in the "system assigned", some of the system's behaviour will undoubtedly be influenced by the impact of "environmental" variables upon it. *Environment* in this instance simply means everything that is not contained within the boundaries of "system assigned" and would certainly include other "systems assigned" belonging to other regulators within the organisation.

Within this wider perspective, it is reasonable to regard the various systems assigned to various controllers (ie, all the peer systems not shown in Figure 3.8) as one sub-system in a wider array of subsystems with which it is linked. These linkages or *couplings* may be either *loose* or *tight* to use the language of process industries (an example of the latter being where two-subsystems are coupled through the exchange of product in process, an example of the former being where the are linked merely though the mutually competition for organisational resources).

The Metacontroller then has an overview of each of the subsystems which gives an advantage of cognisance of events in one sub-system which will impact on another sub-system. With this position comes some extra-requirements particularly the requirement to dampen down oscillations which may set in when each sub-system is regulated in ignorance of the others (ie, the regulating agents merely react to the error messages arising from their own area. A reaction which, because they are unable to perceive the wider scene, may have the effect of amplifying the disturbance to the connected sub-systems). To some extent, this oscillation is offset by either the explicit design (or more frequently, the informal evolution) of an extra-circuit of information. This is shown as the attenuated input from "disturbance" to the "model of system assigned". In practice this may be nothing more than telephone calls between supervisors or conversations in the canteen.

There are many examples available of the impact of these oscillations on H&S performance. Some may be gross such as the TMI-2 accident in which various erroneous control actions interacted to produce a near disaster. Others may be subtle such as oscillating standards of prevention effort or "safety behaviour".

---

[82] Which is shown sending information under attenuation along lines labelled "cause controlled". Chapter 4 will develop ideas about the merits and uses of cause and error controlled feedbacks.

3.7 Self-regulating safety management and managerial cybernetics

This chapter has covered a great deal of ground and, I hope, has succeeded in the terms set out at the start: attempting to draw on unifying principles and introduce systems theoretical ideas to allow a coherent and sufficiently broad view of the subject of safety management.

Much of what has been said has been supported diagrammatically and, indeed, the recurrent idea of circular relationships is certainly made easier to conceive in that format. By way of summarising the central ideas of purpose as definitive of systems, purpose as the progenitor of teleological values and purpose as entailed by what the system does; a sketch (Figure 3.9 below) made early-on in this research may be helpful.

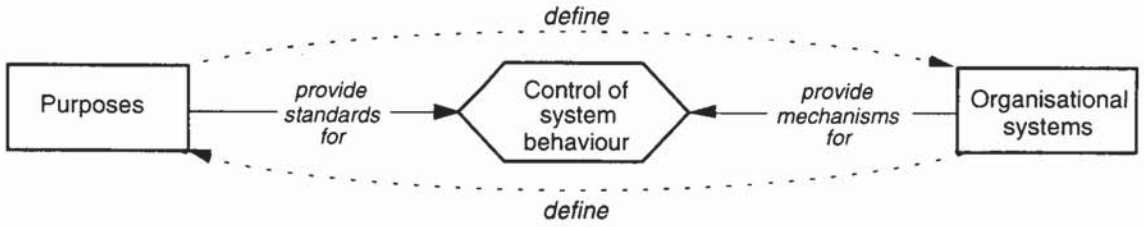

Figure 3.9 Circular relations between system purposes and system behaviour

This chapter has focused on the importance of self-regulation within the management of health and safety and is quite in agreement with Robens in this. However, whilst Robens emphasises the importance of self-regulation and some of the key ingredients within organisational systems that characterise it, he does not specify the supporting organisational machinery. Attempts such as that in HS(G)65 (HSE, 1991) place much emphasis on "positive safety culture" and this may be seen as a contemporary synonym for H&S self-regulation. As with Robens, whilst the ingredients are unexceptionable, the principles underlying self-regulating behaviour are left mostly untouched.

I have made extensive use of Johnson (1970, 1973 and 1980) and the rest of the MORT literature as this represents, in my belief, the most extensive and semantically unified body of technical literature within health & safety management. Whilst I have criticised this literature for lacking overall cohesion, its considerable virtues make it possible to provide this through the application of

systems-theoretical principles without fundamental reconsideration of the MORT literature itself. One of the reasons for this is that MORT is based, first and foremost, on an information-processing view of safety management and information is very much the currency of the systems theoretical approach to understanding organisations. Within the group of systems theoretical disciplines, it is cybernetics which recommends itself to the task of this thesis. As Morgan (1986) puts it:

> "From a cybernetic perspective, everything can be understood in informational terms. It is no accident that the word "information" contains the word "form", for the cyberneticians believe that form rests in information, or difference". "...its basic insight is that systems in the natural and social world can be understood as changing patterns of information has major epistemological implications".
>
> Morgan, 1986, page 357

In the next chapter, I shall present a review of cybernetics with particular emphasis on the sub-discipline of managerial cybernetics. The aim is to present the principles underlying self-regulation and build on the current chapter's introduction to mechanisms through which these principles are enacted. Whilst my emphasis will be at the level of the organisation, various aspects of the wider regulatory picture will be addressed as it is the regulatory environment (as Robens termed it) which provides the conditions in which self-regulation can occur.

4 Cybernetics & Modelling

" This, I believe, is the ultimate precept a theory of organisation can give: not a manual for dictators of any denomination more efficiently to subjugate human beings by the scientific application of Iron Laws, but a warning that the Leviathan of organisation must not swallow the individual without sealing its own inevitable doom."

Ludwig von Bertalanffy (1968)

Introduction

In this chapter, I shall introduce the cybernetics concepts which have formed the basis of this work and demonstrate their importance to the task of regulating health and safety performance. Here, *regulating performance* is considered in a broad sense: as concerning the systematic selection and maintenance of the operating parameters of a system or subsystem.

I have attempted to balance the need to set out the principles of cybernetics with the need to demonstrate the relevance and appropriateness of cybernetics to this thesis. This is not a simple task as cybernetics is a richly interconnected literature – following one line of enquiry tends to encounter the whole weave. With this in mind, the approach I have taken is to provide sufficient coverage and explanation of the main ideas for the limited purposes of this thesis whilst providing key references to the wider cybernetics literature upon which these ideas are founded. The result is perhaps more favourable to the exploitation of cybernetics ideas in health and safety than it is to demonstrating the intrinsic richness and rigour of the cybernetics literature.

In the previous chapter, fundamental systems concepts were reviewed and their application to theoretical issues in safety management considered. The chapter concluded with a discussion of self-regulation in the light of a reasonably straightforward analysis of the term *responsibility*. This showed how the moral philosophic analysis reveals three senses; summarily: responsibility as accountability for past actions; responsibility as having a duty and; responsibility as rationality. Responsibility as rationality dominates the other two senses - one cannot be held accountable nor regarded as able to discharge a duty in the absence of rationality. My further analysis reveals that responsibility as rationality requires 5 criteria[83]:

1) Information about the current state of the system
2) A mental model of the system
3) The computational ability to manipulate this model
4) A set of values against which to judge the performance of the system
5) The ability to alter the behaviour of the system in keeping with (4)

It is reasonable to consider the question of rationality described here as being a relative rather than absolute matter. In other words, instead of rational vs. irrational there is a continuum between perfect rationality and irrationality. Where the responsible agent is found on this continuum at any moment, depends upon the adequacy of each of the five criteria. This is entirely in keeping with Simon's concept of bounded rationality (Simon, 1962).

The fact that each of the aforementioned criteria involve the transaction and processing of information, provides strong linkage between the ideas of Simon (1962, 1982), Johnson (1973 and 1980), Schroeder (1970) and the cybernetics literature. Indeed information is the essential subject matter of cybernetics and, as asserted in the following quotation, essential also to the matter of regulation in systems:

> "In defining cybernetics in the first instance, the great Norbert Wiener used the famous phrase, "the science of communication and control in the animal and the machine". In so doing, he was trying to emphasise two discoveries. **The first was that communication and control were virtually synonymous. Regulation is an informational process**".
>
> Beer, 1983, page 1 (*emphasis added*)

[83] Each of these criteria are themselves reliant upon further considerations which will be examined later in this chapter.

In keeping with systems science generally (as mentioned in the opening paragraph of section 3.6) the aim of cybernetics is to elucidate the general principles of regulation found in all systems. In this relation, Beer continues:

> "The second discovery was that the classical dichotomy, inherited from the ancient Greeks, between the animate and inanimate was delusory. **There are invariant laws of regulation that apply**, like gravity, **to everything**".
>
> Ibid. (*emphasis added*)

The hyperbolic style of these quotations might suggest exaggeration. However, whilst these are bold claims they do accurately reflect the beliefs of the cybernetics movement. Also, whilst cybernetics has obtained many remarkable insights into the nature of regulation it does not make similarly bold claims with respect to designing or discovering a full set of actual regulatory devices to match the range of theoretical mechanisms and principle it has elucidated. This is not a serious criticism but merely a warning that whilst cybernetics can aid the design and analysis of regulatory systems it does not have all the answers to the practical solution of such problems as may be revealed. Like logic, cybernetics is an aid to analysis and, like logic, it provides a rigorous means of stating problems to which answers are not always available. As Ashby puts it:

> "Cybernetics... takes as its subject-matter the domain of "all possible machines" and is only secondarily interested if informed that some of them have not yet been made, either by Man or by Nature. What cybernetics offers is the framework on which all individual machines may be ordered, related and understood".
>
> Ashby, 1956, page 2

The word Cybernetics is derived from the Greek κυβερνητης (pronounced *Kybernetes*) meaning steersman as in the steersman of a ship. This underlying notion is well exemplified in Von Foerster's observation that:

> "cybernetics arises when effectors, say, a motor, an engine, our muscles, etc. are connected to a sensory organ which, in turn, acts with its signals upon the effectors. It is this circular organization which sets cybernetic systems apart from others that are not so organized."
>
> Von Foerster, 1994, page 3

This description of "circular organisation" is seen by Morgan (1986) as typified by four principles:

> "First, that systems must have the capacity to sense, monitor and scan significant aspects of their environment. Second, that they must be able to relate this information to the operating norms that guide system behaviour. Third, that they must be able to detect significant deviations from these norms. And fourth, they must be able to initiate corrective action when discrepancies are detected".
>
> Morgan, 1986, page 87

However, whilst the four principles enunciated by Morgan describe the basics of self-regulation[84], cybernetics also deals with the question of how the "operating norms" are established and changed. This aspect is sometimes referred to as *second-order* cybernetics (eg. Umpleby, 1979., Foerster, 1994) and is identical in its basic concerns to the scheme set out in section 3.6.3 (page 83, ante). There, the involvement of the metasystem was shown to be important in determining the operating norms of the system. Similar ideas have been developed outside of cybernetics, in particular, Argyris (1994) identifies the same set of concerns under the title of *single* and *double-loop learning*. In this model, single-loop learning concerns the system's ability to refine its performance, for example the speed and efficiency with which its behaviour is brought back to normative values when perturbed by a change in the environment. Double-loop learning, on the other hand, introduces the reconsideration of these operating norms or values.

As may already be apparent, Ashby's conception of cybernetics as a generalised descriptive and analytical framework, holds up well even in relation to the apparently more complex realms and less well-defined machines that this thesis concerns - organisations. This opens cybernetics to the charge of being perhaps *too* general and, speaking personally, I have a suspicion of *theories of everything*. After nearly three years reasoning upon the cybernetics literature I have to confess that this suspicion remains with me but, as Scottish law would allow, my considered verdict is *case not proven*. There are many alternative views open to a critical reader, the more relevant amongst those published will be presented in section 4.3 (page 150, post) as they relate to Beer's Viable System Model.

---

[84] In this relation, Ashby (1960) provides a detailed exposition of these principles even to the extent to demonstrating his argument with a working electromechanical device - *The Homeostat*

The stance I have adopted has two aspects:

First, I accept the notion that any system is in fact a subjective phenomenon (or "a mental construct artificially delineated from the unity of the universe" as Beer[85] puts it). In this perspective, what cybernetics provides is a highly refined set of principles and techniques with which to analyse the regulation of such systems. In this respect, cybernetics embodies the most highly developed approach available.

Second, that the perspectives provided by cybernetic consideration provide insights into systemic regulation only partially addressed by alternative theoretical standpoints such as the cultural (eg. Kunda, 1992), economic (eg. Wildavsky, 1989) and psychoanalytic perspectives (eg. De Board, 1978). Each of these has its own distinctive virtues and would shed light on aspects both within and outside the interest of the cybernetician. For example an anthropologically-derived description of an organisation (eg. Kunda, 1992) would better capture and communicate to a reader the symbolism and meanings within that organisation than cybernetics. It would not however be able to identify incompetent regulatory arrangements, analyse them nor identify appropriate principles from which competent arrangements could be derived. Hence, I do not believe it is true that the cybernetics approach is inherently superior to any other. What it does provide is a view of regulation which is more fundamental and onto which *other theoretical views can be mapped more readily than they can be mapped onto each other*. This difficulty of moving between different theoretical views has been termed, at the extreme, paradigm incommensurability (Kuhn, 1970):

> "the proponents of competing paradigms practice their trades in different worlds... two groups of scientists see different things when they look from the same point and in the same direction".
>
> Kuhn, 1970, page 150

Given what I have already set out in chapters 2 and 3 concerning the need to find a common conceptual basis and lingua franca for health and safety, that the "transdisciplinary" (Beer, 1983) character of cybernetics might be coupled with the rigour of MORT is a very attractive possibility.

---

[85] Beer, S. 1983, page 3

4.1 Information and Variety

Given the central importance of information to regulation, as suggested by Beer in the quotation given earlier and the recurrent theme of variety throughout this thesis, it is important to be clear about the interrelation of these terms. Taking information first, whilst one generally thinks of information as being part and parcel of its communication (for example, one ordinarily thinks of a message as *containing* information) this conception is flawed. Ashby (1956) offers a useful example:

> "Two soldiers are taken prisoner by two enemy countries A and B, one by each; and their two wives later each receive an identical message "I am well". It is known, however, that country A allows the prisoner a choice from
>
> *I am well,*
> *I am slightly Ill,*
> *I am seriously ill,*
>
> while country B allows only the message
>
> *I am well,*
>
> meaning "I am alive". (Also in the set in the possibility of "no message"). The two wives will certainly be aware that though each has received the same phrase, the informations they have received are by no means identical."
>
> Ashby, 1956, page 124

What Ashby demonstrates here is that the information conveyed by a message is defined to a large extent by the set of possible messages arising from the transmitting system. Hence, the variety of a system is strongly implicated in the notion of information. Further the definition of *variety* as the number of distinguishable states of a system, subtly introduces the *observer* into the estimation of variety. For example, a simple system consisting of a light that may be green, red or off has three states for an observer with normal vision, two states for an observer with red/green colour blindness, and a variety of zero for someone totally blind.

This conception leads to a number of questions. First, for those system states that are distinguishable, some channel of communication must exist between the system and the observer. The degree to which this channel preserves information is, therefore, an important consideration. The second issue concerns how the

changing state of the system is encoded into a form compatible with the medium of the channel. Again, the degree to which this encoding stage preserves information is of consequence to the transmission from system to observer. Within cybernetics the mechanism which encodes (or decodes) such information is called a *transducer*[86]. It is apparent then that the information received by the observer depends not just on the variety of the system observed but also on the characteristics of channel and transducers which mediate the communication.

If we suppose that neither the encoding or decoding transducers nor the channel itself loses or otherwise corrupts messages being transmitted between the system and the observer, this is still insufficient reason to suppose that information has been received by the observer. *Something* will have been received by the observer, but whether that is disregarded as noise or double-Dutch or misinterpreted by the observer is determined by factors outside of these considerations — What is in question here is the *semantics* of information transfer.

4.1.1 Conceptual or mental models as determinants of information

Aside from the transmission of data and the various mechanisms entailed, the informational content for the recipient is importantly derived from their conceptual model of the transmitting system. Again, the topic of perception arises and what is said in chapter 3 (page 67, ante) applies here. Perception is a creative process involving, not the passive import of data into the brain via the senses, but the active engagement of theories of the system in question as a conceptual analogue or *mental model*[87] of the perceived system. Hence perception can be regarded as another transduction, in this case not of translation between one code and another (eg. as electrical code is transduced into acoustic code by a telephone receiver) but as the mapping of information about a system onto the conceptual model of that system in the mind of the perceiver. From this perceptual perspective, it follows that information is obtained from data insofar as the data available to the senses corresponds or *maps* onto the perceiver's conceptual model of the system from which those data arise. Essentially then, information can be distinguished from data on the basis that the former is (1) assimilable within an existing conceptual model and (2) changes that model in the sense of aligning it more closely to the current state of the system that it represents. This reasoning fits

---

[86] From the Latin *transducere* - to **lead** across

[87] A term first coined by Craik, K.J.W (1943) but subject to renewed interest since the 1970's (see, for example, Gentner and Stevens, 1983; Johnson-Laird, 1983; Norman, D.A., 1988).

with the definition provided by Beer (1985) of information as subjectively *"that which changes us"*.

The foregoing does not necessarily mean that where there are discrepancies between the conceptual model and the system so represented that those data which do not fit will simply be disregarded as noise, although this is a possibility (as the TMI-2 accident demonstrates[88]). The other possibilities, for purposes of illustration rather than exhaustive listing, include:

a) **The data are simply not registered in consciousness** - they have been cognitively filtered out as noise in the signal. For example, the stub axle in my car was worn (in fact near to failure) - I was wholly unaware of anything untoward until a friend borrowed the car and within two minutes behind the steering wheel suggested to me that *"That* doesn't sound too clever". I still could not hear *that* at all;

b) **The data are mapped using a many-one transformation which does not preserve variety**. Here, perhaps consequential distinctions of system state do not exist within the conceptual model although a coarse representation does. For example, a junior doctor may pick up a heart murmur during an examination. In contrast, an experienced cardiologist may be able to specifically diagnose the precise heart condition even though the data available to both doctors is identical;

c) **The data are mapped using a many-one transformation which whilst not preserving variety initiates a process of restoring requisite variety to the model.** An example here is when I worked as a picture framer. The job in question consisted chiefly of mass-producing framed prints in runs of about 500 per day, where I would cut the mitres using a double circular saw. Because I was using this saw for about 20 hours a week, I developed a sensitive knowledge of the nuances of vibration, smell, sound associated with the machine. On the occasion in question, the "feel" of the machine changed and although the quality of cut was unaffected. Nevertheless I was concerned (the machine figured largely in my livelihood at the time and I had respect for its

[88] During the incident, operators disregarded certain instrument readings as these did not fit with their model of the system condition (an example of what is sometimes referred to as *mindset*). Fortune and Peters (1995) provide many similar examples of people disregarding apparently clear signals because they do not fit in with prior expectations and one need only consult the signal detection literature (eg. Wickens, 1992 and Helstrom, 1995) to find myriad examples of this and of theories which seek to predict the behaviour.

ferocity) and began stripping the machine down to find the cause of the change. This turned out to be a worn drive-belt. Thus, at the start of this chain of events, whilst the variety of my conceptual model of the machine was something less than requisite, part of this model was "the feel of the machine is a vector composed of x, y, z variables at α, β, χ values respectively" a change in one of these values alerted me to my own lack of variety as a regulator of that machine. The search for the cause was as much initiated and led-by my current model as led-to a new model of greater variety.

Common to all of these examples is the distinction between data and information. Also present in these examples is a fundamental concept of information theory - that information can be regarded as the reduction of an observer's uncertainty about the state of a system.

4.1.2 Information as defined within information theory

The theoretical information literature, whilst strongly associated with cybernetics, has its own independent existence. The principal difference between the two disciplines is that cybernetics is concerned with regulation in systems, which is certainly informed by information-theoretical notions. In particular, as regulation involves the passing of information between regulator and system, our interest is both with this exchange and with the varieties of the communicating parties.

We saw with Ashby (page 114, ante) that information is importantly determined by the variety of the transmitting system and information theory is quite in agreement with this. This is illustrated by Shannon's formula (Shannon and Weaver, 1949) for the information content of a message:

$$H(x) = \sum_{i=1}^{i=N} p(i) \log_2 p(i)$$

where H is the information content each symbol within the message[89] (X) in units using logarithms to the base 2, or *bits;* p(i) is the probability of symbol (X_i) occurring; and N is the total number of different symbols available. Hence, the equation characterises the ensemble of all possible messages (for a given system) and this determines the information content of a particular message. In this way,

[89] which consists of a collection of symbols ($X_1, X_2, \ldots X_i \ldots X_n$)

the *ensemble of all possible messages* is formally equivalent to variety defined as the *total number of distinguishable states.*

If we consider a symbol (X_i) as signifying that a state (i) has occurred within the transmitting system, it becomes apparent that more information is conveyed by symbols occurring with observed probabilities *significantly different* than is expected of them (ie, given by a knowledge of the relative frequencies of the symbols within the ensemble). In other words data which registers that a system is behaving wholly as predicted is not informational, as (Atlan, 1983) puts it:

> "From this point of view, the more a given event is unexpected, the more we learn from it when it actually happens. If no a priori uncertainty exists as to its occurring, the very fact that it occurs does not teach us anything: its information content is zero. On the other hand, the more uncertainty about its occurring, the more information content characterises this event. Thus, the measure of information is reduced to that of uncertainty, which in itself is given by a probability. This is how, within these limits, we can understand in a natural way that the information content is as high as its uncertainty is high, ie, its a priori probability is low. "
>
> Atlan, 1983, page 11

As an aside, this conception serves to further demonstrate one of the problems noted in chapters 2 and 3 (ante) concerning the dangers of treating accidents as homogenous events. This danger is that we forget that we are talking about an event and, instead, talk about its set. Hence an accident happening may be treated as having little information because accidents happen often enough. Alternatively, the probabilities of the various system states subsumed within the accident "message" may well be significantly different from those predicted and this, according to the Shannon definition, has a higher information content.

The limitations of information theory are apparent in the foregoing, as section 4.1.1 (page 115, ante) makes clear, the meaning or semantic content of information is untouched by the Shannon conventions. However, what information theory does provide is a rigorous treatment of the various attributes of communication from the transmitting system up to and including the decoding of messages by the observer as suggested in Figure 4.1 below:

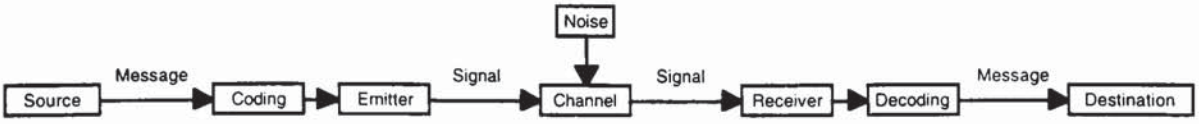

Figure 4.1 Communication linkage between source and receiver (after Atlan, 1983)

4.2 Regulation, Variety and Information

As the quotation from Morgan (page 112, ante) makes clear, whilst in the previous section we have considered the communication of information in one direction, regulation minimally requires the exchange of information in both directions.

This is all quite unremarkable and intuitive, however, the cybernetic (and, more particularly, Ashbian) conception of regulation has profound implications for the fundamental concepts upon which MORT is founded: *change* and *energy flows*. Before these can be discussed, some further description of the cybernetic principles of regulation is required.

4.2.1 The Ashby conventions for describing regulation.

In chapter 3, the concept of ***essential variables*** was introduced (section 3.6.4, page 92 and thereafter). Essential variables are those aspects of system behaviour which are regarded as important and, in an organism, those which are essential to viability. For each of these essential variables, there will be a range within which values must be maintained, although there will be some variables the values of which are absolute (ie, a value less than or more than this would be lethal). In complex systems there may well be a large number of these essential variables and for simplicity, I will refer to them as a set.

Given what has already been said (in chapter 3) concerning open systems - exchange between system and the environment is essential to viability. However, it is a matter of the nature and degree of this exchange. Those exchanges, the action of which are deleterious with respect to the essential variables of the system, are regarded as a ***disturbance***.

In the ordinary way of things, organisms fit with respect to the demands of their environment, have evolved means of handling the transmission of entropy from their environment to, as it were, their essential variables. Thus the beneficial portion of environmental entropy (such as food) is incorporated and the hazardous residual (such as mechanical insult) is blocked. This selective transmission and blocking of entropy can be seen as fitting within[90] the Shannon conception sketched out at 4.1.2. Thus, whether the medium is gas exchange at the

[90] Indeed, the relation between entropy and information is formalised to the extent where, as Atlan (1983) points out (citing the work of Brillouin (1959)) statistical thermodynamics can be treated as a special treatment of information theory.

lungs, consumption of food or a conversation, these transfers can be properly treated in terms of information. With this in mind, regulation can be seen as the selective blocking of information. This is most easily be imagined as a physical blocking, such as a pair of suitably specified ear-defenders or an insulated protective suit for working in heat. Also, it may be a blocking effect where the information impinging upon the system is opposed by information generated by the regulator.

Ashby (1956) provides an illustrative example - a thermostatically controlled water bath. Here, the essential variable is the temperature of the water for which a value (eg. 25°C) has to be constantly maintained regardless of fluctuations of temperature in the room which houses the bath. In the example, Ashby supposes that an experimenter wishes to choose between 2 such baths on the basis of the essential variable already mentioned. Figure 4.2, below shows the temperature of the room and the two baths as recorded over a 24 hour period

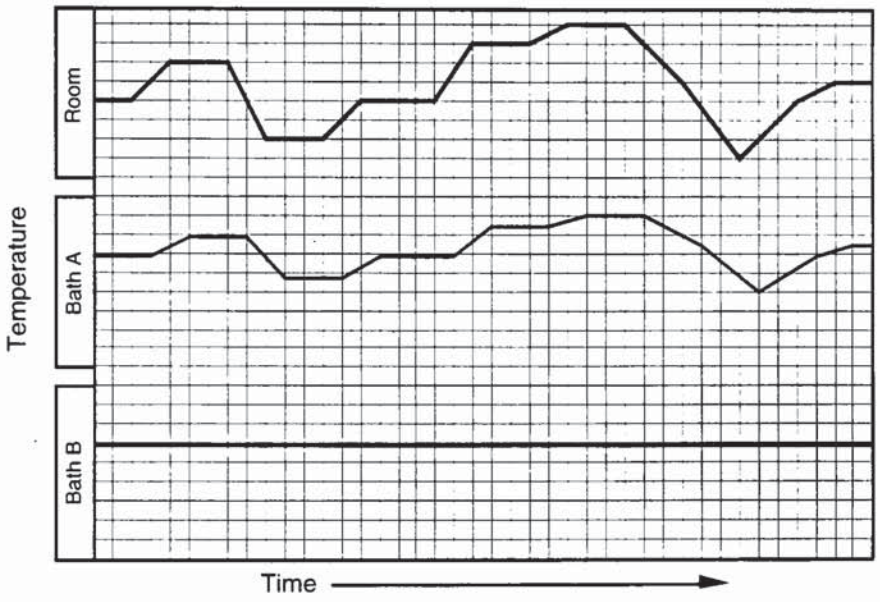

Figure 4.2 Thermostatic temperature regulation in two water baths relative to a fluctuating ambient temperature.

If the changing temperature is thought of as a series of messages transmitted to the water bath system one can see that the information is not received in Bath B but is (albeit slightly phase-shifted and attenuated) in Bath A. Hence the effect of the thermostat is to block information flow from the room to the essential variables of the system (the water temperature). Another example might be the action of a defender in basketball to the basket shooting attempts of an opposing member of the team: each duck and dive of the would-be scorer is matched by the

defender blocking the transmission of the ball to the basket. Similarly, the ear-defenders mentioned above which achieve their effectiveness by a mixture of reflection and absorption of acoustic energy have their equivalent in the use of anti-phase acoustic techniques. Here, the acoustic signal arriving from the environment is simultaneously received and a new signal (the phase of which has been adjusted so as to cancel out the original signal) sent to the system.

Implicit in all of the examples provided (particularly evident in the last case) is the need for the regulator to match *message for message* the environmental disturbance that would otherwise adversely impinge upon the essential variables of the system. Hence current anti-phase technology is extremely effective when the noise is highly predictable but has problems when noise is unpredictable in any of its parameters (Fitzgerald, 1996). Similarly, for the basketball players; if the would-be scorer is 7 feet tall and the defender is a mere 6'6", the effectiveness of the latter as a regulator is seriously compromised. Similarly, even if both players are the same height, if the scorer is more agile or simply more skilful than the defender, the degree of regulation (moves and shots blocked) will be partial.

4.2.2 Regulation and requisite variety

In terms of *variety*, the examples provided illustrate the need for the regulator to possess *at least as much variety* as the source of disturbance. As the example of the water bath illustrated, perfect regulation is achieved when the variety transmitted from disturbance to essential variables is zero - the unvarying water temperature in bath B of Figure 4.2. This is the regulatory principle of requisite variety. As Ashby states it:

> "only variety in R [the regulator] can force down the variety [in the essential variables of the system] due to D [the disturbance]; **only variety can destroy variety".**
>
> Ashby, 1956, page 207

Whilst the basic principle enunciated above is common to all regulation, there are two main forms, **cause-controlled** regulation and **error-controlled** regulation. The first of these is illustrated below at Figure 4.3 as a directed graph (after Ashby, 1956).

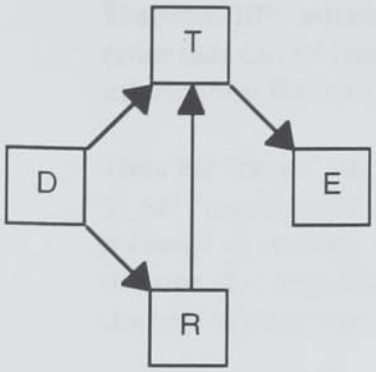

Figure 4.3 Cause-controlled regulation.

The conventions here are *D* (Disturbance), *R* (Regulator), *E* (Essential variables) and *T* stands for "Transformation[91] Table". Concerning the last of these, the reason for the rather unfamiliar term is that Ashby is describing *functions* rather than *real* objects. The function of *T* (which might be a machine or an organism or an ecosystem) is the transduction of variety from the environment (including *D*) and transmission to *E*, a process made selective through the functioning of *R*.

In the water bath example, *D* would correspond to the fluctuating ambient temperature of the room; *R* would be the thermostat; *T* would be the water bath and its contents and *E* would be the temperature of the water (of which of the whole set of values, η is the subset of those acceptable). Because the Ashby conventions are drawn in functional terms we have the freedom to redraw the functional boundaries between the real-world objects and attributes represented. In this example, we might place the boundary between the water bath and regulator slightly differently to include any insulation on the body of the bath as part of the regulator - *R* because the function of insulation is regulatory.

The cause-controlled regulator is, in principal, the superior design as it does not have to wait for any effect to occur in *T* or be made manifest in the essential variables before acting. The information route shown in Figure 4.3 can be redrawn to make this point clearer.

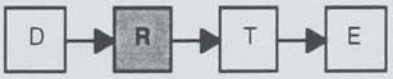

The cause-controlled regulation scheme affords *R* the possibility of perfectly blocking the flow of information *D* to *E* (subject to *R's* variety relative to *D* as mentioned previously). This presentation also makes clear the relation between Ashby's Law of requisite variety and the work of Shannon.

> "The law of requisite variety says that R's capacity as a regulator cannot exceed R's capacity as a channel of communication. In the form just given, the law of requisite variety can be shown in exact relation to Shannon's

[91] The term "transformation" is in the normal control engineering sense and the box "T" can be regarded as conforming to the conventional principals of transfer function analysis (for example, Raven, 1995).

> Theorem 10[92], which says that if noise appears in a message, the amount of noise that can be removed by a correction channel is limited to the amount of information that can be carried by that channel.
>
> Thus his "noise" corresponds to our "disturbance", his "correction channel" to our "regulator R", and his "message of entropy H" becomes, in our case, a message of entropy zero, for it is *constancy* that is to be "transmitted". Thus the use of a regulator to achieve homeostasis and the use of a correction channel to suppress noise are homologous".
>
> Ashby, 1956, page 211[93]

Error controlled regulation is required when the regulator cannot react directly to a disturbance. In the case of the water-bath, for example, the thermostat reacts not to the change in the ambient temperature (*D*) but to the change of value in the essential variable (*E*) - the water temperature.

In general, using the conventions of Figure 4.3, there are two options for error controlled regulation. These are shown as directed graphs in figure 4.3b below.

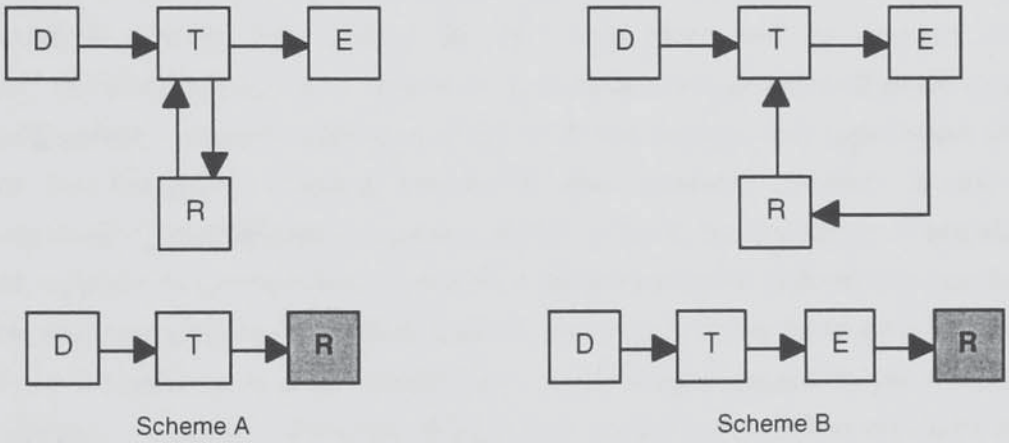

Figure 4.4 Two generic forms of error-controlled feedback

Under scheme A, as the lower portion shows, the regulator obtains its information by way of *T*. This means that the disturbance is already acting on *T* before any regulatory action can be initiated. Additionally, information or variety may not be preserved by its transduction through *T*. For these reasons, the error controlled regulation can not be made perfect in the sense that cause-controlled regulation can. Were the regulator intelligent (eg. human) this scheme suggests that *R* can

[92] Shannon Theorem 10 (also known as "the noisy channel theorem") . Atlan (1983) provides a proof of the formal correspondence between this theorem and requisite variety.

[93] I cannot help but feel that in the context of Ashby's normally moderate prose, these paragraphs have something of the flourish with which a magician produces a rabbit from a top hat - complete with a fanfare

only make sense of what is going on (the action and nature of *D*) through, as it were, the eyes of *T*.

Scheme B, takes matters yet further. Because *R* is receiving its information via *E*, successful regulation (eg. through trial and error) means that, as the water bath example showed, the channel of information from *D* is blocked to some extent. As Ashby puts it: *"the more successful R is in keeping E constant, the more does R block the channel by which it is receiving its necessary information. Clearly, any success by R can at best be partial"* (1956, page 224). However, so long as it is true that the effect of the disturbance is not so dramatic as to force the *E* value(s) outside of their viable range before *R* can act, the scheme can prove highly effective. Small errors are compensated by regulatory action thus providing adaptation to the *T-E-R* complex overall with respect to the environmental variations *D*.

A further refinement in this picture concerns the error-controlled regulatory situation when both *T* and *R* are recognised as probabilistic rather than deterministic in their internal arrangements. Deterministic meaning a system whose state is wholly determined by its initial state and its current input - ordinarily, simple electrical and mechanical machines conform to this description. Health and safety certainly concerns itself with the design and regulation of such machines, but the issue of health and safety management certainly is one more characteristically probabilistic in nature. Ashby's basic formulation of regulation, as set out, applies to probabilistic systems without resort to subsidiary conditions. However, the regulation of a probabilistic system (also referred to as a Markovian machine) by a regulator is characterised by a non-unique trajectory (ie, from initial state to desired / equilibrial states). Whilst this seems a much less efficient means of achieving regulation it is both simple (because a detailed algorithm does not need to be prepared - every state of *D* matched by a state of R) and within limits can work with less advance information about *D* . *R* behaviour can be characterised overall as never reaching equilibrium until the values of the essential variables are within the set η. Ashby (1956) refers to this as regulation by *vetoer*[94].

[94] An issue here is the likely behaviour of several Markovian machines linked together (ie, sub-systems of a larger system). In such an assembly, all or some perform as regulators of their neighbour sub-systems and the possibility of continuous vetoing of the states of neighbouring leading to incessant oscillations is very real. In control engineering, such behaviours are well understood as is the generic term for their remedy - damping. As discussed later, organisational systems suffer from the internal sub-systemic oscillation and the cybernetic approach to damping these is considered in section 4.4.2.2, page 178, post).

4.2.3 The relations between general regulation and the regulation of H&S.

Having now introduced the various conventions, Ashby's statements concerning the general form of regulation can be considered. Ashby (1956) provides a general statement of prerequisites for regulation:

> "In practice the question of regulation usually arises in this way: The essential variables *E* are given, and also given is the set of states η in which they must be maintained if the organism is to survive (or the industrial plant to run satisfactorily) These two must be given before all else. *Before any regulation can be undertaken or even discussed, we must know what is important and what is wanted."*
>
> Ashby, 1956, page 219

Reconciling this with H&S regulation: the nomination of essential variables and the acceptable values pertaining to them, has largely been dealt with at the philosophical level in section 3.6 of chapter 3 (ante). In terms of MORT, these are largely decided by the M-branch functions as shown in Figure 4.5 on page 126. Policy determines the essential variables and, to some extent[95], the maximal risk exposures, The risk assessment system determines the risk levels pertaining to the essential variables (Ashby's η set)[96]. This rather cursory treatment will be deepened later as it depends on the more technical matters discussed next.

Ashby's formulation of the general problem of regulation is:

> "*Given* E, η, T, *and* D, *to form the mechanism* R *so that* R *and* T, *coupled, act to keep* E *within* η".
>
> Ashby, 1956, page 220

As before, can we reconcile this viewpoint with H&S regulation? I believe we can. Consider the definition of accident from the MORT literature:

> "An unwanted transfer of energy or an environmental condition which, due to absence or failure of barriers and/or controls, produces injury to persons, property or process"
>
> Horman, 1992, page 134.

---

[95] The matter is complicated by the cost-benefit trade-offs in the risk assessment system (MA3) and by the prior determination of certain risk criteria at a yet higher metasystemic level (such as state prescribes a maximum tolerability for risk of death and absolute limits exposures to certain harmful substances). These are considered in proper detail in chapter 5.

[96] An argument for later development concerns the relations between the η set and the "Assumed risks" (R branch) of MORT.

This is a statement of the MORT diagram at SA1 (see appendix 1, post) which for ease of exposition is given in summarised and positive form as Figure 4.5 below.

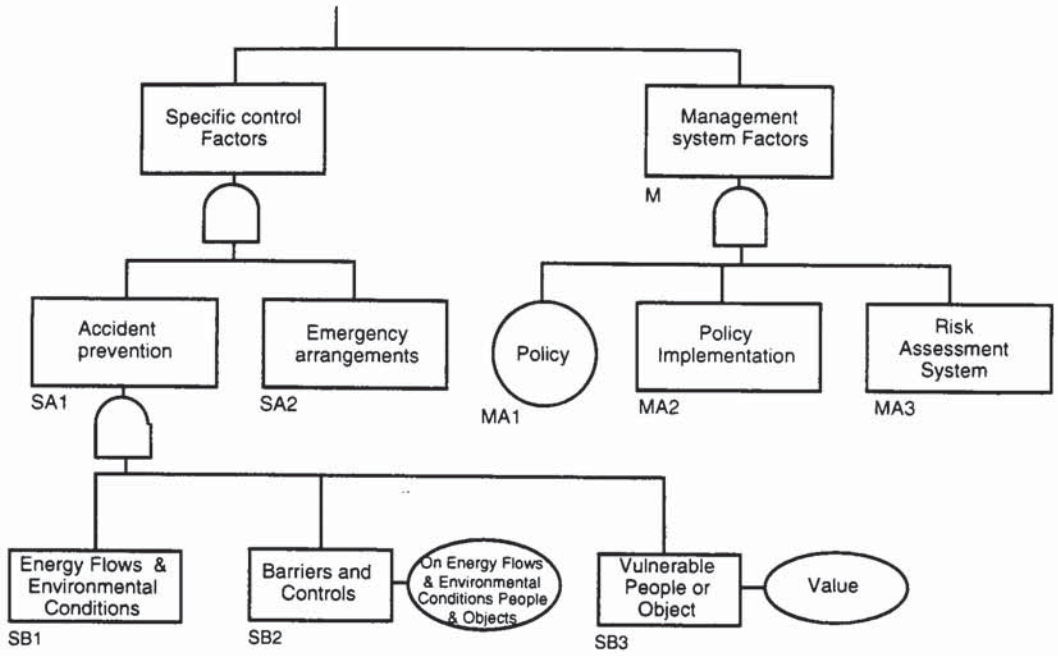

Figure 4.5 MORT elements S/M to SB/MA level in success-tree format

Using the MORT conventions, the *functional* equivalence between the regulatory accounts of Ashby and MORT at the operational level are summarised below in Table 3.

ASHBY FUNCTIONS		**MORT FUNCTIONS**	
Terminology	***Symbol***	***Terminology***	***Symbol***
Essential variables	*E*	Vulnerable people or objects	SB3
Set of values for *E*	η	Assumed risks	R_n[97]
Table of transformations	*T*	Operational system	PPP[98]
Disturbance	*D*	Unwanted energy flows	SB1
Regulator	*R*	Barriers & Controls	SB2

Table 3 Equivalence between cybernetic and MORT regulatory functions

[97] As per Johnson (1980) page 90. This is subject to later discussion.

[98] PPP configuration - a MORT term meaning the operational system as defined by the Plant itself, the People who staff and operate it according to the operating Procedures.

Using the conventions of Figure 4.3 (page 122, ante), the equivalencies set out in Table 3, are illustrated graphically below in Figure 4.6. Here, I have used the cause-controlled paradigm but the choice is arbitrary.

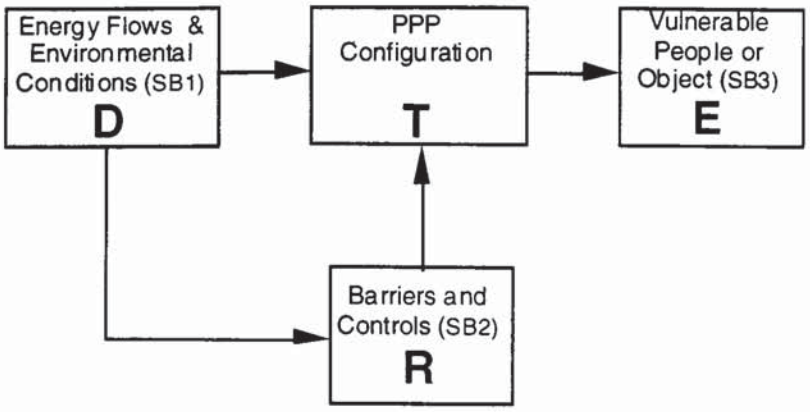

Figure 4.6 Equivalence cybernetic and MORT regulatory functions

When I first considered harmful energy flows in the mode of Ashby's *D* function, I was concerned as to how far I was stretching assumptions. The first verification I made was to reconsider the ideas informing the MORT concept of accidents as summarised by Horman (1992) above. This is based on, amongst others, McFarland who summarises his thesis as:

> "All accidental Injuries (and damage) result (1) from the application of specific forms of energy in amounts exceeding the resistance of the tissues (or structures) upon which they impinge, or (2) when there is interference in the normal exchange of energy between the organism and the environment (eg., as in suffocation by drowning). Thus, the various forms of energy... constitute the direct causes of injuries in accidents. Also, prevention of injuries can often be achieved by controlling the source on the energy, or the vehicles or carriers through which the energy reaches the body."
>
> McFarland, 1967, page 146

McFarland's account certainly fits with those of Ashby (his *Disturbance*) and Shannon (his *Message*). The MORT concept of barriers and controls distinguishes them on the basis that *barriers* describe those devices (procedural and/or physical) designed only with the control of unwanted energy flows. *Controls*, on the other hand, whilst offering similar protection are designed principally to control desired energy flows. Thus, the idea of barriers and controls as regulatory devices within a system are quite in keeping with MORT.

The matter of which physical objects and attributes correspond to Ashby's *D* and *T* was dealt with above (page 122, ante) where it was noted that the boundary between *D* and *T* is arbitrary - in other words we can please ourselves so long as once the boundary is placed for a given analysis we are consistent. When using MORT for accident investigation, the analyst is faced with similar decisions and often has to adopt a series[99] of different *D* | *T* boundaries. Also, where the boundaries are placed is determined by functional rather than physical considerations. For example, a car colliding with a tree carrying an unrestrained driver may be regarded as having aspects of *D*, *T*, *E* and *R* sharing the same physical system (the driver himself): *D* includes the driver's own velocity with reference to the local elements of the car (the steering wheel/ dashboard etc), *T* includes his body tissues, *E* includes his blood volume, blood pressure and the functioning of his internal organs and, *R* as his decisions not to wear a seat-belt, to drive too fast etc. The question of *D* | *T* | *R* boundary placements is further considered in section 4.2.4.

Having reflected on the relationship and tested the correspondence (MORT and Ashbian regulation), the MORT SB1-3 scheme appears homologous to Ashby's, and thereby with Shannon's communication theory. Further, as mentioned on page 119, as information can be formally treated as entropy within statistical thermodynamics, there exists the possibility of formalising (ie, mathematically) the set of relations between regulation, information, thermodynamics and the McFarland (1967) energetics. However, this is beyond the scope of this work (and this author) although, the principles will be further considered insofar as they inform the arguments later presented.

---

[99] Note that this is a *series* - not a parallel set of boundaries. The serial element arises because the may be several energy flows to account for. For example, a traffic collision might require the consideration of:

lorry A as a kinetic energy flow with respect to the barrier (say a red traffic light) with car B as the target;
then B as a kinetic energy flow with brakes and crumple zones as controls/barriers with the occupant of B as the target;
then the occupant of car B both as a kinetic energy flow and the target with the crash restraint system (eg. seat belt) as the barrier;
and, if perhaps, driver B sustained trauma due to the seat belt itself...
then the kinetic energy flow of driver B (forward momentum) as concentrated through the small surface area of the belt (during rapid deceleration), with his/her internal organs as the target and the design and installation of the seat belt as the barrier.

4.2.4 Regulation of the large system

In the foregoing section, I introduced the relations between MORT and the fundamentals of cybernetic regulation. For the purposes of explanation, I chose fairly straightforward and easily imagined examples. However, can this simple scheme cope with the regulation in a large system?

First of all, the adjective *large* needs qualification: the solar system is large but as a system of massive bodies in relative motion, it is highly predictable. Thus whilst spatially vast, in respect of its variety as an orbital system, it is relatively small: largeness, for our purposes, is defined by the variety to be matched by the regulator.

As mentioned in chapter 1 (page 33, ante), whilst the subject matter of this thesis might be conventionally regarded as belonging and restricted to the province of the so-called high-risk industries, a sustaining belief during this work has been the conviction that it applies as much to a self-employed plumber as it does to a nuclear reprocessing installation. The rationale is this: the logic of regulation which pertains to the management of H&S is the same although the sophistication of method and practical difficulties may be vastly different. From the regulatory standpoint the plumber's occupation involves much less variety than the operation of the nuclear reprocessing plant and, as such, presents less of a challenge analytically. Yet both of these must balance their regulatory variety to the systemic and environmental varieties that threaten them. Whilst I have confidence in the simple argument that the principle of requisite variety applies in any instance of regulation, it is quite apparent that the variety problems faced in the regulation of the large system are not generally solved by having a single comparably large regulator. The perspective which assists here is one that views the large system not as singular high variety "box" to be regulated, but as a set of nested boxes each in interaction with the others. Thus, while it is convenient to think of D as arising in the environment outside the large system and impinging upon it, a more realistic view is of a set of single boxes being acted upon by not only D-variety arriving from the environment but also the variety transmitted by the neighbouring boxes to which they are coupled. For example in Figure 4.7 below, sub-system T_5 is subject to variety transmitted from D and from T_3 and, indirectly from T_4. Further some of T's variety is exchanged with D via T_1 which may have the effect of increasing the variety of the former.

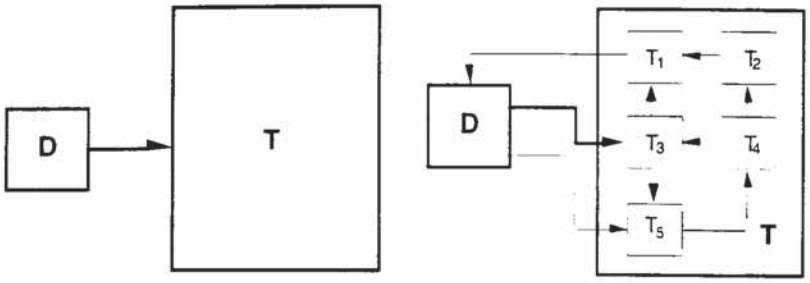

Figure 4.7 Sub-systemic interaction contributes to Systemic Disturbance

Ashby describes this matter as follows:

> "The size of the dynamic system that embodies *T* tends to be correlated with the variety in *D* for several reasons. If *T* is made of many parts, and there is uncertainty about the initial state of any part, then the variety will be allocated to *D*; so in general, other things being equal, the greater the number of parts the greater the variety in *D*. Secondly, if each part is not completely isolated from the world around, each part's input will contribute some variety which is allocated to *D*; so in general, the greater the number of parts, the greater the number of components in *D*; and, therefore, if the components have some independence, the greater the variety in *D*."
>
> Ashby, 1956, page 245

There are four options which can be exploited in the regulation of a system whether the system be large or small:

1. Allow a wider range of values in η;
2. Reduce the variety of the Disturbance;
3. Exploit constraints in the variety of Disturbance[100];
4. Increase the variety of the Regulator.

Each of these four are considered in turn in sections 4.2.4.1 to 4.2.4.4, below.

4.2.4.1 *Allow a wider range of values in* η

This, in essence, is to lower the standards For example, the regulatory variety required to control the goal-scoring variety of Manchester United Football club is

[100] eg, if the numerical variety of the disturbances is, say, 100 but only 10 are frequent and the remainder rare — the disturbance is showing some form of constraint. Hence, a low-variety regulator such as the cough reflex will clear the most frequent portion of disturbance variety whilst leaving rare occurrences such as impaling fishbones unregulated.

roughly equivalent to the goal-preventing variety of Aston Villa F.C.[101]. Aston Villa might decide to sell all their current players and employ a less skilled team. Perhaps the club had run up large debts or need extra funds to finance a new stadium. However, the new Villa team will be running a considerable variety deficit when they next play Manchester United and would likely be unable to contain the latter's goal-scoring variety. The same argument reversed would suggest that the "average" regulatory variety disposed by Du Pont sites must be higher than that of comparable ICI sites in order to achieve the order of magnitude superiority in accident rates associated with the former. The ethical issues raised in chapter three (ante) are clearly implicated here[102] in relation to H&S but, as Rasmussen (1996) and Senge (1992) note, this is a generalised problem (i.e. across different categories of adversity). Further, both authors note that this phenomenon is more commonly experienced as gradual change that is unopposed (as *entropic drift* for Rasmussen, and *eroding goals* for Senge) rather than an explicit decision to accept a lower standard.

4.2.4.2 *Reduce the variety of the Disturbance*

This is often the first option within H&S regulation and, insofar, as we are addressing man-made systems it is quite apparent that to "remove or reduce the energy" is the top of the safety control hierarchy. However, there are three limits to this approach. The first limit is that operational effectiveness will require a certain quantity of variety and will seek to preserve it[103]. The second limit is related and concerns the power of an entity seeking to reduce the variety of the disturbance. For example, the regulatory approach prior to the HSW act was to limit or proscribe certain processes and substances[104]. Aside from the political difficulties inherent in this approach (that is, the operational attempts to preserve variety by political means), the state as a regulator has serious practical limitations on its own variety (both acquiring adequate information about industry and handling information through so constrained a channel as the parliamentary

---

[101] Non-British readers should note the personal peril engendered by the author by using such an example

[102] Indeed, the HSE have little option than to allow this in the face of the huge industrial variety they seek to regulate given their limited variety as a regulator. This is the subject of further discussion.

[103] This argument is the old chestnut of "safety vs. production" somewhat disguised or, indeed, revealed less emotively by this analysis. Indeed, violations may be seen as the unilateral expansion of variety within the operational system.

[104] Of which the regulatory attempts to proscribe the use of Asbestos in Europe and the US is a prime example.

process). The third limit concerns the portion of *D* which truly arises from the environment beyond the control of the system designers.

4.2.4.3 *Exploit constraints in the variety of Disturbance*

The point here is that whilst the variety in the components that comprise *D* may add to a huge total variety, *D* may nonetheless show some form of constraint. In other words, the actual variety of *D* is less than it might otherwise seem[105]. The argument here is that if the variety of *R* cannot be increased, nor *D* artificially decreased, we may have little option other than to regulate against the repetitive, or constrained, variety of D. The perspective here is of the long-run, regulating the repetitive disturbances and living with the effects of the rare non-repetitive disturbances. In relation to H&S, this option raises some of the ethical concerns of chapter three (ante). Ashby provides an engaging example:

> "It should be noticed that what is "Good" in the Grand Outcome [the long-run] does not necessarily follow from what is "good" (η) in the individual outcomes; it must be defined anew. Thus if I go in for a lottery and have three tickets, a win on one (and consequent loss on the other two) naturally counts as "Good" in the Grand Outcome; so here 1 good + 2 bad = Good. On the other hand, if I am tried three times for murder and am not found guilty for one, the individual results are still 1 good + 2 bad, but in this case the Grand Outcome must naturally count as Bad[106]. In the case when the individual disturbances each threaten the organism with death, Good in the Grand Outcome must naturally correspond to "good in every one of the individual outcomes"."
>
> Ashby, 1956, page 249

The ethical issue here is the acceptability of a "bad" individual outcomes (people quâ citizens, as ends in themselves) as contrasted with a "good" grand-outcome for the organisation (people quâ workforce or local residents, as means to organisational ends).[107]

---

[105] The logic here is much the same as the effect of AND gates in fault trees - they imply constraint and, heuristically, the greater the proportion of AND gates to OR gates the more constrained the system in question.

[106] The death penalty for murder was still an option open to UK courts at the time Ashby wrote these words.

[107] c.f. Section 3.6.4 "**Recursiveness and the definition of *Adversity***" (page 87, ante).

4.2.4.4 *Increase the variety of the Regulator*

This one might have expected to be dealt with first as it appears less open to ethical compromise than the options so far discussed. However, increasing the capacity of a given regulator is often difficult in practice. As an extreme example, my personal regulatory capacity to match the goal-scoring variety of Manchester United F.C. could not (within the rules of the game at least) be increased to anything approaching requisite variety. Whilst this is certainly ridiculous it does open the debate of somewhat more serious suggestions. Let us suppose we wished to halve the annual number of reportable accidents occurring in a large firm and we looked to the safety department as the regulator of this performance to achieve this. What increase in the resources of the dept would it require to achieve this - double the number of staff? Treble? I doubt that even a quadrupling would do greatly more than dent the reportable accident rate, as the variety deficit would still be huge. Now, the safety department if asked would I am sure pronounce this example to be no less ridiculous than the previous one. Nevertheless it makes the point: *increasing the variety of a given regulator is likely to meet practical limitations long before it meets requisite variety*. This is not to say that we should abandon this option as maximising the amount of regulation a given regulator can provide must surely be a meritorious option.

4.2.5 Regulator design and amplified regulation

The reason why some of the examples of the preceding section were so ridiculous was because they so blatantly ignore the principle of *regulatory amplification*. Here, instead of attempting to increase the variety of a *particular* regulator we attempt to increase the variety of *overall* regulation deployed against disturbance. In simple terms regulatory amplification is achieved in much the same way as power amplification. The law of conservation of energy requires that this amplification is, in fact, a supplementation achieved in, minimally, *two stages*: the energy provided by the input and the supplementation of the this energy from some abundant source to produce a more powerful output. Similarly, the law of requisite variety requires that a low-variety input is supplemented using an abundant source of variety to produce a higher (ie, requisite) variety output. Hence a regulator with a variety of say 10 units, facing a disturbance of 1000 units must find an amplification of minimally 100x to reach requisite variety. Returning to the example of the previous section, let us take the safety department as currently able to dispose 10 units of variety (in this case the regulatory variety of the safety advisor) to the task of regulating the 1000 variety safety management of the firm.

Increasing the staff complement of the safety department with 3 equally experienced others increases the regulatory variety from 10 to 40 Units. If, on the other hand, the current regulator uses its resources to obtain a second stage source of variety, the option of multiplying its own variety becomes available[108]. This statement, whilst purely illustrative provides the necessary change of focus, namely, the view of regulatory amplification as involving design: the obtaining of a second (or indeed, third, fourth, etc.) stage of regulation.

4.2.5.1 Design and Information Theory

The process of design can be seen as a process that determines the final machine from an arbitrarily large number (or high *variety set*) of possible machines. Although design is in practice a creative process (eg. Carroll & Rosson, 1985), it is essentially a process of *selection*. Selection in this way can be seen as the reduction of the variety of possibilities to a variety of 1 - the final design itself. Hence, when we consider the role of the designer, it can be seen as a process of selecting one state from a set of states. Returning to the discussion of section 4.2 (ante), this set of possible machines is a special case of Shannon's ensemble of messages and, thereby, the transmission of a particular message can be brought into correspondence with the selection/design of a particular machine. Ashby summarises this as:

> The act of "designing" or "making" a machine is essentially an act of communication from Maker to Made, and the principles of communication theory apply to it.
>
> Ashby, 1956, page 253.

The change of focus referred to on the previous page is considerable because the impact of Ashby's just quoted statement is profound. In recognition of this, and before continuing on this line of reasoning, a brief three-point summary is given below.

I. We have already seen how regulation is an informational process and that the formal relations between the cybernetic view of regulation, information theory and McFarland's energetics allow transfer of this conceptualisation into health and safety management via MORT conventions (SB1-SB3 acting as our gateway).

---

[108] Often, in safety practice, this second stage is frequently identified with supervision. However, whilst this observation is noted for the sake of illustration, I am jumping the gun somewhat.

II. The regulatory requirements of the high-variety system require us to consider how to possess the regulator of requisite variety and we have seen four options (at 4.3.4 above). Each of the options (a) to (d) above are either limited in their effectiveness or lead to untenable compromises[109], or both. In particular, the obvious route of increasing the variety of the regulator will often be found subject to practical limitations

III. This leads us to consider how to *amplify* regulatory variety which, as discussed, involves the design of a secondary regulator possessed of greater variety. The process of bringing this secondary regulator into practice has various noteworthy attributes:

A. the act of design can be seen as an act of communication from designer or maker to made;
B. we can therefore apply everything discussed about information theory to the study of this process;

The change of focus allows the following set of conclusions:

I. Design, seen as the process of selecting from a large variety set a single design option (which conforms to certain design criteria), *is homologous to the regulation of a system with respect to the values of the chosen essential variables* (ie, the criteria to be met in the final design = a vector of essential variables at given values, ie. η).

II. Therefore, the process of design *is itself* an act of regulation and, as such, all that has recently been said of regulation and regulatory amplification can be applied to the process of designing regulators.

A. Thus the amplification of a regulator (R_1) via a second stage regulator (R_2) to achieve RV in a regulatory process, corresponds to -
B. the amplification of selection by a designer (of limited variety) via a second stage selector (of greater variety) to achieve RV in a design process

[109] Which is not to say that these options should not be exploited, merely that we can not expect these alone to provide a solution to the regulatory requirements of the system.

III. From this we can see that regulators may be designed by other regulators and that:

A. the regulation of the large system may be accomplished by the inception of a regulator R_1 which makes another regulator R_2;

B. the task of R_2 may be the direct regulation of a system if RV is obtained or the design of another regulator R_3 to obtain amplification to RV;

C. this scheme (whether selection in stages or amplification in stages) may be repeated as many times as it requires to achieve RV in the regulatory task;

D. The regulation and regulatory design processes are both reducible to the same analytical form. In *principle*, these are reducible to measurable quantities should we wish to know the variety sums involved.

Probably one of the best studied class of phenomena, in which this cascading regulation and design is expressed, belong to developmental biology, the sub-discipline dealing with embryonic development. Consider the following quotation from Gilbert (1994):

> "Organs are complex structures composed of numerous types of tissues. If one considers an organ such as the vertebrate eye, for example, one finds that light is transmitted through the transparent corneal tissue; it is focused by the lens tissue, the diameter of which is controlled by muscle tissue; and it eventually impinges upon the tissue of the neural retina. The precise arrangement of tissues in this organ cannot be disturbed without damaging its function. Such co-ordination in the construction of organs is accomplished by one group of cells changing the behaviour of an adjacent set of cells, thereby causing them to change their shape, mitotic rate, or differentiation. This action at close range, sometimes called **proximate interaction** or **secondary induction** enables one group of cells to respond to a second group of cells and, in changing, often to become able to alter a third set of cells".
>
> Gilbert, 1994, page 647[110]

Put simply, the inducing cells act as a regulator, selecting the appropriate sets of genes in the responding cells. Concerning the latter, given the requirement of abundant source of variety for regulatory amplification, the DNA is the source of

[110] An elegant demonstration of these principals can be found in Nüsslein-Volhard (1996)

variety in the responding cells[111]. Further, as implied at the end of the quotation just given, the construction of an organism involves perhaps many cascading acts of amplified selection. The reason why the process of design often occurs in stages involving supplementation is twofold. First, the genetic code of "higher organisms[112]" does not contain sufficient information to construct a phenome of requisite variety to regulate its essential variables with respect to the environment. Second, the task of construction is too complex to achieve in one stage and must use the method of successive supplementation so as to meet the requirements of the law of requisite variety at each stage and overall.

Given the concern of this thesis with the design and evolution of safety management systems, the notion of design by supplementation and regulation through amplification are of obvious importance and application. The evolutionary mechanism is also of application although it is more subtle and complex. For the current purpose we will consider two aspects: Firstly, a process of selection which has been operating for some considerable time before the genesis of the system of interest; secondly, as a method of regulatory design within the extant system (whatever its level of development).

4.2.5.2 A-priori Selection - *Inheritance*

When we consider the development of a particular organism - the result of the evolutionary process is relevant whilst the *process* itself is not. It has, as it were, paused whilst the organism battles against the selection forces rallied against it drawing upon its genetic inheritance as well as acquired regulatory characteristics.

---

[111] Shannon's information theory would predict that the regulation (as selection) of the inducing cell would be bound by the genetic variety of the regulated cell (ie, the genetic variety = Shannon's ensemble of messages). Empirical justification of this position has been available since the Spemann experiments of the 1930's. In one experiment, Spemann introduced embryonic salamander cells into a frog gastrula (both the salamander and frog gastrula were sufficiently developed to allow approximate identification of the host and donor regions). The result of this was a frog tadpole with a salamander mouth. As Spemann puts it "The ectoderm says to the inducer, 'you tell me to make a mouth; alright, I'll do so, but I can't make your kind of mouth; I can make my own and I'll do that' " (quoted in Gilbert, 1994).

[112] It should be noted that regulative development is one mode of development found to a greater or lesser extent in all species. The other mechanism of development is called "mosaic" and, essentially, this involves certain cells differentiating autonomously. Thus, at an early stage of development, regions of the embryo will already be dedicated to a particular course of development and will fulfil this fate even if separated from the host embryo. It seems apparent that whilst the mosaic form of development is more commonly found in the development of "lower organisms" and regulative development more characteristic of the "higher organisms", both types are used in all organisms.

Similarly, if we consider the development of an organisational system at whichever level of recursion, it can be seen to have an *inheritance* which pre-existed it. If we were to consider the THORP facility now operational at BNFL's Sellafield site, whilst many aspects of the plant are unique to it, it is nevertheless the inheritor of some portion of the sum of nuclear engineering knowledge, BNFL's policy structure and operational experience, IAEA and NII experience which pre-existed it. If any of this history was different, THORP would be a different system. Likewise, the regulatory system of THORP is, in part, determined by these antecedents which owe nothing to its technical particularities. In one plant I am familiar with, the tendering policy of the commissioning organisation resulted in a much a larger variety of replacement parts (functionally identical but physically different in various aspects) and machine variances than might have occurred without this policy. Similarly, the chemical process at the ill-fated Nypro works at Flixborough was designed to mix liquids rather than vapours. The mixing of liquids is a less efficient process than the modern method of mixing vapours (less molecular contact between the reagents per volume). Thus a greater volume of reagents is required to produce a given volume of product and, therefore, a greater volume of in-process inventories were required (Kletz, 1991).

What is common to both these examples is that the plant designers made selections from what technology was extant and even their innovations were predicated on available knowledge. Similarly, the design criteria (the essential variables) were in large part the *inheritance* of a-priori selections. The principal here is that, even when innovating new methods of regulation, the designer is predisposed to make certain selections and not others because of technological and systemic[113] heritage.

4.2.5.3 Evolutionary Regulator Design

In some cases, a regulator-designer will create a deterministic regulatory machine, that is, one which is entirely predictable given a knowledge of its starting state. This is likely to be the preferred option if the regulatory design is well within the variety constraints of the designer. However, when the regulator design is of higher variety than the designer can bring to bear upon the process of its design another option needs to be found. One option is to supplement the design process using another source of variety as has been discussed extensively above. Another

[113] That is, everything that impinges directly or indirectly - whether legal, social, organisational or personal.

method exploits the method described in section 4.2.2 (page 124, and footnote 84) regulation or, in this case, design by *vetoer*. The advantage here is that instead of specifying an algorithm for the to-be-designed regulator (which achieves RV by matching every consequential state of the system to be regulated by a specified state of the new regulator) a simpler *heuristic* method is used. The design process in this case is not dominated by the need to specify an algorithm to be embodied in the new regulator but to build a Markovian machine with plenty of states (ie, greater variety than the system ultimately to be regulated) but which has been constrained such that its only regions of stability correspond to those desired states of the system (the η set). The description offered by Ashby is of assistance here:

> Let the regulator R be built as follows. Let it have an input that can take two values β and γ. When its input is β (for "Bad") let *no* state be one of equilibrium, and when its input is γ (for "Good") let them all be equilibrial. Now couple it to T so that all the states in η are transformed, at R's input, to the value γ, and all the others to the value β. Let the whole now follow some trajectory. The only states of equilibrium that the whole can go to are those that have R at a state of equilibrium; but this implies that R's input must be at γ, and this implies that T's state must be one of η. Thus the construction of R makes it a vetoer of all states of equilibrium in T save those in η. The whole is thus regulatory; and as T and R are here Markovian, the whole will seem to be hunting for a "desirable" state and will stick to it when found. R might be regarded as directing T's hunting.
>
> Ashby, 1956, page 233

This approach is noteworthy in many aspects. Firstly, from the vantage point of an observer T can be seen to be teaching R and, if R *has a memory*, its slow hunt-and-stick method used on the first occasion of T taking a novel state outside of the η set, will result in a learned regulatory response the next time T moves to this state. This is the basic rationale underlying the design of computational neural networks and is a reasonable model of organismic learning. Second, so long as R has appropriate transducers[114] , the extra variety required to regulate novel states of T arises from the source of the disturbance itself. In this way, one can see that it is mice which teach kittens how to catch mice thus supplementing the information lacking in the gene-pattern of the kitten. This is explained in detail in section 4.2.7 and its subsections (starting at page 142, post).

---

[114] ie, to allow T variety input to R - for example, if I am operating a machine the imminent failure of which is heralded by a terrible noise but at 50KHz, I will be none-the-wiser on the next occasion as on the first - I cannot learn because my transducers (hearing) cannot preserve this variety.

4.2.6 Modelling and regulatory efficiency

Taking up the theme of the previous paragraph, I should now like to again draw attention to the Conant-Ashby maxim (first mentioned on page 14, ante) as the cybernetic principle that has been implicit throughout section 4.2, namely, that *any regulator of a system must contain a model of that system*. Whether we consider cause-controlled or error-controlled regulation (c.f. page 121, ante) we have taken it as axiomatic that the regulator must possess at least as much variety as the system regulated. Thus if a machine T can take up to 6 distinguishable states but regulator R only has 5 states available, regulation will be imperfect (if the sixth system state is outside of the acceptable η-set of states). For the purposes of illustration, Figure 4.8 shows an error-controlled regulator in schematic form.

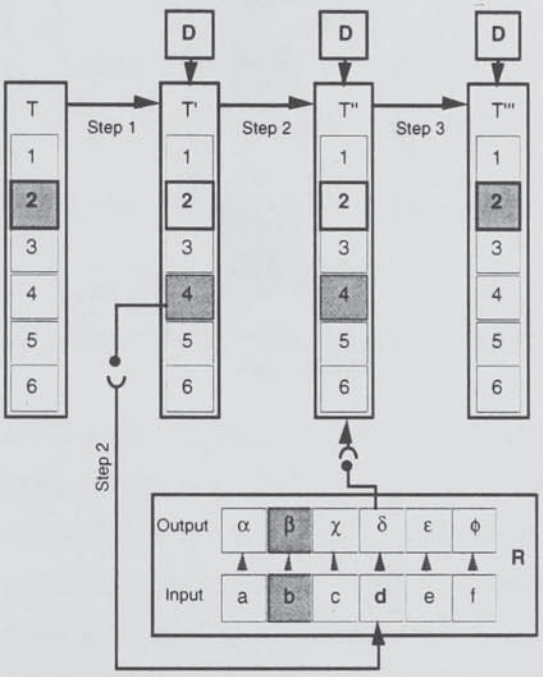

Figure 4.8 Schematic error-controlled regulator showing mapping of T to R (input), and R input-output relation.

In this figure, four successive states of the $R - T$ complex are shown. At T (the left hand column) the system is at state #2 which is the only acceptable state for this system (ie, the set η contains only one member, 2). A disturbance D acts upon the system such as to change the system to state #4 in one step. This is shown as T'. At step 2, this change of state is registered in R as d (T-R uses the mapping 1→a, 2→b, 3→c, etc.). R's response to input d is to transmit output δ to T''. The effect of regulatory input δ is to bring T back to state #2 at step 3. (An example is given at footnote [115], below.)

[115] For example, if a car (T) hits a large stone in the road (D) which has the effect of jerking the wheels from straight-on (state #2) to right-turn (state #4). This is transmitted to the driver (R) through various inputs including the clockwise rotation of the steering wheel (as input state d) who initiates output δ (an anti-clockwise rotation of the steering wheel corresponding to the speed and angle denoted by δ). This returns (in this simplified example) the car (T) to the state #2 - that is, a direction of travel parallel to the line of the road.
NOTE: In the scheme of Figure 4.8, I have shown every input to R as having a unique regulatory output - it need not be so... for example there are a large variety of different states whilst driving to which the low variety regulatory response of "braking" will be adequate.

The importance of this is that, whilst a variety of 6 available to both system and regulator is a necessary condition for successful regulation, it is not sufficient. There must also be a formal correspondence between the varieties of system and regulator such that the states of the system map onto the states of the regulator. Taking Figure 4.9, below, it can be seen that of the six states available to *T*, a mapping exists only for *T* states 1, 2, 4, 5 and 6, although *R* also has 6 states. *T*-state 3 has no corresponding state in *R*, hence if state 3 is expressed in *T* no change occurs in *R* and no regulatory response will occur until *T* reaches another state.

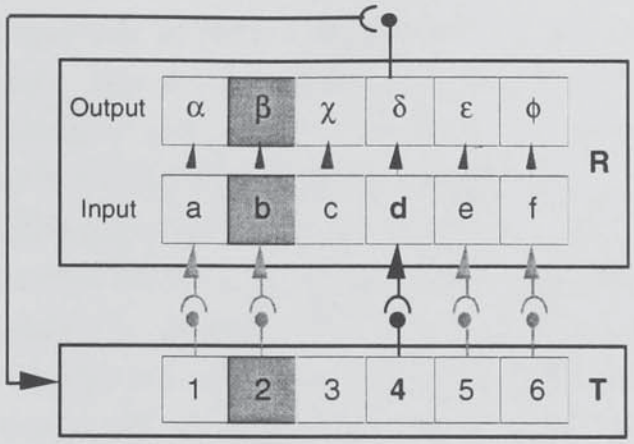

Figure 4.9 Partial mapping between regulator (R) and system (T)

Thus, the regulation by *R* will be sufficient for all *T* states except #3. Were it the case that all *T* states mapped onto unique states in *R*, the mapping would be *isomorphic*, identical. Hence, under some transformation, *R* can be regarded as a *model* of *T*. It is perhaps worth recalling at this point that the conventions used are functional (c.f. page 122). Thus, when I say that "*R* is a model of *T*" this refers only to those aspects of *T* which are material to the task of regulation for which *R* exists. In the same way a model car will preserve certain characteristics under an appropriate transformation (such as 1/30th scaling of physical dimensions) whilst not preserving the variety of other characteristics such as the engine configuration, materials used etc. In the case of regulatory modelling, the attributes of the set of real-world objects and relations to be included depend upon the purposes of the model maker (ie, the designer of the regulator).

Thus the argument is brought full circle to the discussion of chapter 3 (in particular, section 3.6 on) and allows a broad conclusion: *A regulator, whatever its material embodiment, is functionally a model of a system as defined by the purposes perceived in the system by the designer of the regulator.* Further, the regulation of a system can only be as competent as the regulatory model of it is complete. From the managerial perspective, this means that the management of a process will be bounded by the competence of the model of that process. States of that process which do not correspond to the regulatory model of that process cannot be

recognised nor, therefore, acted upon. In this way the Ashby-Conant maxim (Ibid.) can be seen to correspond to Simon's concept of bounded rationality (Simon, 1962) and, with it, the related ideas of, Johnson (1973 and 1980) and Schroeder (1970) as noted on page 110.

The question of how close the mapping in *R* of *T* must be, is an intriguing one. The best answer I can provide owes little to technical prowess: the mapping should be as close as provides the level of regulation judged acceptable and *no more*. As a general principal we will always wish to have a regulator which is as simple as possible to avoid unnecessary labour in its construction, maintenance and especially to reduce the variety required in its meta-regulator. Thus, for the most part, regulatory models will be some many-to-one (homomorphic) mapping of the system's states rather than a one-to-one (isomorphic) mapping. In effect, homomorphic models may treat distinct states of the system as identical - thus different system states will evoke the same regulatory response. Insofar as this arrangement achieves acceptable regulation of the essential variables, requisite variety is de-facto established in the regulator. For example, I do not need a detailed knowledge of electronics to regulate the behaviour of a radio set: merely a model which relates the outputs to the inputs (knobs, switches and batteries). My regulatory model has RV.

4.2.7 Construction and acquisition of regulatory models

An immediate question concerns how, in principle, the regulatory model is established in a regulator. We have already touched on this in discussion of the adaptive error-controlled regulator (section 4.2.5.3, page 138, ante) but the matter requires a more even-handed treatment to include the other regulatory paradigms and to provide an introduction to the managerial cybernetics of Beer (1959; 1979; 1981; and 1985).

4.2.7.1 Regulatory modelling and the cause-controlled paradigm

In the cause-controlled case (also known as *feedforward*), the regulator is clearly a model of the system it regulates; as Conant and Ashby (1970) put it (paraphrasing[116]): *R* must be a homo- or isomorph of *T* (since it has the same input as *T* and a mapping related output).

[116] The paraphrasing is due to a difference between the notations used here and in Conant and Ashby (1970), the latter having adopted the notation used by Sommerhoff, 1950 (op.cit. Footnote 63, page 81).

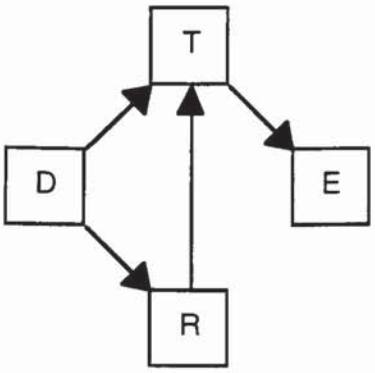

Figure 4.3 Cause-controlled regulation.

However, inspection of Figure 4.3 (page 122, and reproduced adjacent for ease of reference) reveals no means within the data flows between *D, T, E* & *R*, by which *R* might acquire a model. What is required then is a designing agent somehow *outside* the diagram. This designer, as discussed in section 4.2.5 (page 133, ante), must possess a variety equal to the variety of *T* (as given by the particular attributes which affect the essential variables *E*), the variety of the disturbance *D* and the variety of the set of acceptable states (η)[117].

The limitations of this approach are: (i) only states of *D* or *T* which affect *E* foreseen by the designer will be regulated against and therefore; (ii) any change in *D, T* or *E* must be subject to the designer's modification of *R*. Thus the cause-controlled regulator, whilst "perfect" in the sense described in section 4.2.2 (page 121, ante) has the drawback of being unable to adapt to changes in *T* and *D*. In the language of GST, such an arrangement is stable (delivers "perfect" regulation of *E*) but not *ultrastable* (delivering regulation of *E* even in circumstances not envisaged by the designer of *R*).

In the case where *R* is human, and assigned the task of regulating a system *T*, the designer implicitly provides a model of *T* through the provision of procedures. Given appropriate information sources (displays of the state of *D* - *stimuli*), controls (input devices to affect appropriate changes in the state of *T* - *responses*) and sufficient alacrity and reliability in the association of stimuli and responses, this may result in stable performance (ie, as defined in η) of the system. However, if a state of *D* or of *T* occurs outwith the implicit model provided by the designer, the human regulator will be unable to stabilise the resulting change of state in *E*. This example is noteworthy in a number of respects. Firstly, such an arrangement looks "fragile" as it depends for its success on no consequential changes in any of the functional elements. Not least amongst these elements is the human regulator who is required in this cause-controlled regime to behave as a deterministic machine (neither likely nor an appropriate use of human cognitive resources). Secondly, the scheme will keep the human regulator firmly within the skill- and

[117] This is to say that attributes of the real-world objects comprising *T* need not be modelled - the system T is defined through the purposes imputed to it by the designer of *R*.

rule-based behavioural domains of Rasmussen's SRK scheme (Rasmussen, 1980) which whilst blocking-off the opportunity for them to acquire a mental model through which they may enter the knowledge-based domain in the event of unforeseen disturbances or system states.

Such a regulatory system may work in three types of situations:

- a system characterised by no unforeseen states of consequence to η (by the designer) in neither T nor D, and in which there is a large investment in both the constant maintenance of the human regulator (ie, the provision of training to maintain stimulus-response contiguity to the required level of reliability) and displays/controls;

- a system which is subject to changes in D or T, where the procedures used by R are instantly modified by the designer. What this means is that the designer must always possess both perfect information (of changes in D and T) and the variety to compute and effect (instantly) corresponding changes to all the regulators that depend on his/her action (which, of course, presupposes equal variety to the actual task of regulating T);

- a system in which the η-set is generous. As encountered in section 4.2.4.1 (page 130, ante) a wider range of states in η may allow the system to tolerate sufficient variability in E to allow the designer to effect the necessary changes in procedures (the "grand-outcome" approach buys time for the designer). However, whether this option exists or not takes us to a yet higher recursion - at which the regulator of the designer is brought into consideration.

From the above it appears that the cause-controlled paradigm offers the opportunity of *perfect* regulation but does so under very strict cybernetic assumptions (as well as dubious psychological ones) that are unlikely to be met in practice. Not least of these is that the effective denial of the self-organising abilities of the human regulator (ie, learning) means that the full variety of the regulatory task must be borne by the designer. Thus, in many settings, the cause-controlled paradigm has serious practical limitations.

4.2.7.2 Regulatory modelling and the error-controlled paradigm

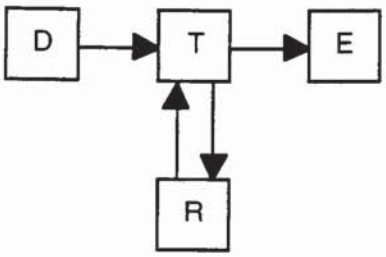

In section 4.2.2 (page 121, ante) two forms of error-controlled regulation were illustrated (schemes A and B of Figure 4.4). Scheme A, (shown left) which involved the regulator receiving information about *D* via the changes in *T*, is similar to the cause-controlled paradigm insofar as *R* receives no information about *E*. Thus, much of what is said at 4.2.7.1 applies here also, and, in fact is different in two respects only.

The first difference is that the response speed for *R* may be slower[118] than in the cause-controlled (as the changes effected by *D* in *T* must be manifest before *R* can act). The second difference is that *R* now has some scope for error-recovery insofar as it receives feedback about *T* except its essential variables. This is akin to a procedure where a sequence of control actions is recorded and available for review. For example, if a control is pushed and, in doing so, illuminates a light; the regulator at least has the reassurance that the control has been activated. This, it should be noted, does not necessarily mean that the essential variables have been regulated. It merely means that the control action *envisaged by the designer as associated with regulation of the essential variables* has been correctly carried through[119]. Thus, this mode of regulation relies upon the model implicit in the designer and the caveats given at 4.2.7.1 apply here.

[118] It will be slower under conditions where the change in *T* registered in *R* simultaneously corresponds to a change in *E*. However, certain states in *T* may herald a trajectory of change which acts to change *E* only after a number of steps. In this case, *R*'s action upon *T* may intervene during this trajectory, returning *T* to a stable state before *E* is influenced by D.

[119] Whilst this appears a rather bizarre approach it is evidenced in the design of certain displays in reactor plant 2 at Three Mile Island. One of the events in the sequence leading to the loss of coolant incident was the failure of a Pilot Operated Relief Valve (PORV) to close. The purpose of this PORV in the design was to provide relief of over-pressurisation in the primary cooling system. Part of this, naturally enough, is to shut when the pressure had been reduced to normal. Owing to a history of PORV closure failures, an annunciator light had been provided to show when PORV was closed. However, this indicator was not activated by sensors measuring the action of PORV (eg. sensors registering flow through PORV) but merely the presence of a signal "ordering" PORV to close. (Data was available to the operators which would have supported, by inference, that PORV had not closed... but this hypothesis was not entertained at the time and complicated by many other factors). This account is drawn from Perrow (1984), Woods (1987), Kletz (1988) and Reason (1990).

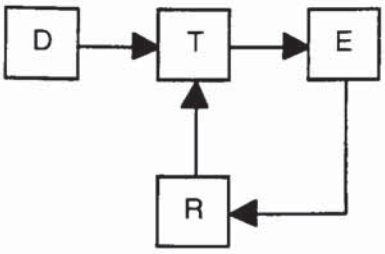

Scheme B (shown left) of Figure 4.4, illustrates the regulation of a system *T* using information about *D* as registered in the essential variables *E*. As was noted on page 124 (ante), information reaching the regulator about *D* has been filtered by *T*, and may well have lost information in the process. *R's* success in regulation has the effect of further reducing the information it receives about *D*[120].

Despite the paradoxical effect that *R's* success causes loss in the information by which it regulates, in practice, the residual variety reaching *E* is enough to allow *R* to maintain adaptability to small changes in either *D* or *T*. Thus, so long as the effects of a novel state in *D* upon *T* are not so great as to cause a lethal variation in *E*, this error-controlled paradigm provides the possibility of adaptive regulation. That is, to respond to changes in *D* or *T* not foreseen by the designers of *T* or *R*. However, this requires the further conditions described in section 4.2.5.3 (page 138, ante): that *R* be Markovian; that *R* have a large variety of states that it may occupy; that *R* be equipped with the ability to distinguish states of *E* that are within and outside of the set η. In principal then, the designer of *R* need not specify a model of *T except* in respect of the values essential variables which are acceptable — *R* cannot logically acquire these at its own level of operation, the η-set must be given by the designer.

As for the development of a model of *T* in *R*, the prescription above requires two additions. The current provisions will allow *R* to go into a *hunting* mode as soon as the value(s) of the essential variables leave the set η; this means *R* will *veto* all states of *T* and will only *stick* when a state of *T* is reached which returns the *E* values within η. Any change of state in *D* which displaces *E* from η, no matter how predictable, will cause *R* to begin hunting. In light of this, the first addition required to the design of *R* is a *memory* to allow it to recall the regulatory response it obtained via hunt-and-stick to a particular change of vector in *E*. Thus, should this pattern of inputs (*E* values) be repeated at a future time it may use a learned response. Over time, this would allow a repertoire of such responses to be

[120] To recapitulate: *R's* task is to reduce variation in the values of the essential variables of the system to those within the η-set. Thus, *R* can be said to be perfectly regulating if no information is communicated from *D* to *E* through *T*. In this way, the better *R* performs, the less information is available to it from *E*. However, as *R* can only *react* to change in *E*, in practice it will always receive some information arising from *D* and thus never be a perfect regulator.

developed. Whilst this in principle will reduce the time that R might otherwise spend in a *hunting* response to previously encountered states of D, this has not addressed the time it may take on the first occasion. Ashby (1962) and Beer (1966; 1979) both demonstrate that a *random* process of hunting will, under the majority of circumstances (typically where there are a large number of essential variables to maintained and T is large), take too long to achieve adaptation for practical purposes. The solution to this problem is complex and the subject of extensive discussion in Ashby (1962) and Pask (1975). For the present purposes it is sufficient, in my view, to note that a *random* process of selecting an appropriate regulatory action may lead to interminable hunting but that a *directed* process may reduce the time taken to feasible limits.

Returning then to the provision of a model of T in R, whilst the error-controlled feedback can in principal develop its own model by associating outputs states of T (ie, E states) with regulatory inputs made to T, this is likely to take an impractical length of time. There seems the need, then, for the designer to provide R with a low variety model which predisposes it to sets of *hunting* experiments upon T rather than a random undirected experimentation. In this way, the regulator may develop a model of T of greater variety than that provided in the first instance by the designer. This scheme is illustrated below

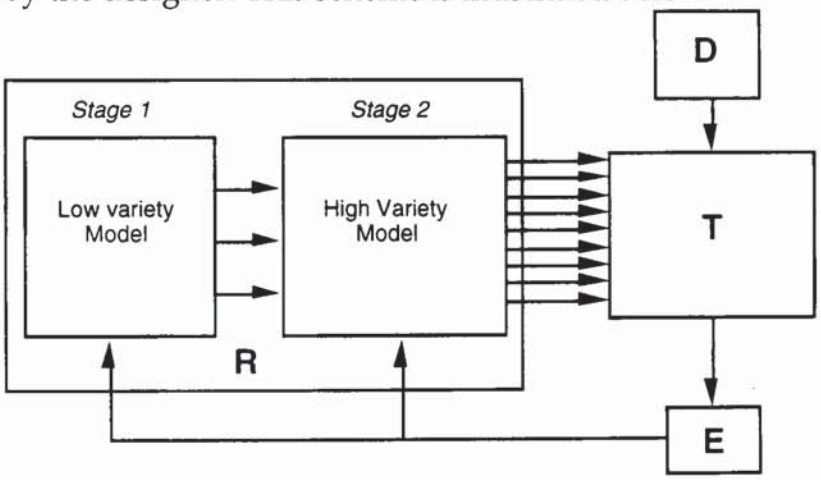

Figure 4.10 Two-stage adaptive regulation

The provision in R of a low variety model plus the mechanisms of adaptation does, I trust, show clear correspondence with the discussion at section 4.2.5.2 (page 137, ante) where the notion of *inheritance* was considered in relation to design selections. Further, the means suggested in the foregoing show, where R is human, how the knowledge-based realm of performance (Rasmussen, 1980) is introduced; the production of a high variety of regulatory actions from a limited variety mental model.

As with the other regulatory paradigms, this model must minimally include the η-set of acceptable states. These requirements will give rise to regulator which can adapt to unforeseen states in *T*. In terms of variety, the regulator described may develop a model of *T* which is not equivalent to the full variety of *T*, but just that proportion which is induced by the action of limited variety disturbances. However, the regulator is still a hostage to fortune insofar as its state of information about *D* remains much reduced - thus whilst the new arrangements allow the possibility of effective regulation in the face of small perturbations of *E* by *D*, where the classes of disturbance has been foreseen by the designer of the model; large perturbations by novel classes of disturbance remain effectively unregulated.

4.2.7.3 Regulatory modelling and the full-information paradigm.

As the name suggests, the full-information paradigm (Klir, 1991) seeks to combine the cause-controlled and error-controlled regulatory schemes. It may be recalled that the cause-controlled paradigm relied upon the designer of the regulator to provide a model which has requisite variety to the task of regulating the system *T*. However, aside from the problems of changes not foreseen by the designer, there is the more immediate problem that this scheme presupposes RV in the designer. As has already been argued, one reason for designing a regulator is to amplify the variety of the designer when he is seeking to amplify his variety.

The adaptive error-controlled scheme, using feedback from the essential variables permits the regulator to learn from small errors and thereby develop a higher variety model of system *T*. In order to allow *R* to develop responses in a timely fashion, it was suggested that *R* be equipped with a low variety model, the action of which is to predispose its experimental inputs. This assumes a foreknowledge of the likely classes of disturbance on the part of designer as well as a mapping between these classes of disturbance and classes of regulatory input (the idea being that *R*, thus equipped, experiments not at random but within classes[121]). However, should a disturbance act upon the system of a type not foreseen by the designer, this regulatory arrangement breaks-down: the regulator will go into interminable hunting meanwhile the disturbed *E* values will leave the η-set region of viability before the regulator can return *T* to an acceptable state.

[121] This is homologous to the amplification of design variety as set out in section 4.2.5.1 (page 134, ante)

What is required to offset the exigency described, is a means of updating the low variety model "originally" supplied by the designer to R in the light of new information about the possible states of D before these send the E values outside of the η-set. In other words, the low-variety model "originally" provided by the designer to the regulator was based on the designer's best estimate of what the future might hold at the time of the regulator's inception. If at a later time the designer is in possession of better information about the future, it should modify the low variety model provided in the light of this information. In this way, the designer may better equip the adaptability of the regulator.

Ashby (1962) illustrates this debate in relation to genetics:

> "What concerns us in this book is the fact that the active defences can be *direct* or *indirect*. The direct... include all the regulatory mechanisms that are specified in detail by the gene-pattern. They are adapted because the conditions that insisted on them have been constant over many generations.
>
> The earlier forms of gene-pattern adapted in this way only. The later forms, however, have developed a specialisation that can give them a defence against a class of disturbances to which the earlier were vulnerable. This class consists of those disturbances that, though not constant over a span of many generations (and thus not adaptable to by the gene-pattern, for the change is too rapid) are none the less constant over a span of a single generation. When disturbances of this class are frequent, there is advantage in the development of an adapting mechanism that is (1) controlled in its outlines by the gene-pattern (for the same outlines are wanted over many generations), and (2) controlled in details by the details applicable to that particular generation.
>
> This is the learning mechanism. Its peculiarity is that the gene-pattern delegates part of its control over the organism to the environment. Thus, it does not specify in detail how a kitten shall catch a mouse, but provides a learning mechanism and a tendency to play, so that it is *the mouse* which teaches the kitten the finer points of how to catch mice.
>
> This is regulation, or adaptation, by the indirect method. The gene-pattern does not, as it were, dictate[122], but puts the kitten into the way of being able to form its own adaptation, guided in detail by the environment".
>
> Ashby, 1962, page 234

[122] The same argument, transformed into organisational terms, is presented by Beer (1979), see page 161, post.

The intriguing realisation in this is that the process by which an adaptive regulator is to be made effective (that is, to obey the law of requisite variety) is through amplification of models by stages. In Ashby's example, there are two stages: the low variety model provided through the genome and a higher variety model developed in the phenome (utilising a plentiful supply of states in the regulator assigned, as it were, by harnessing the variety provided through disturbance).

As Ashby (Ibid.) later observes, there is no reason why the genome should accomplish this indirect regulation at one remove, many recursive stages may be used, each stage making selections of the next stage, each successive stage amplifying the variety of its precursor.

4.3 Adaptive regulation and managerial cybernetics[123]

Whereas the biological examples of Ashby have the genome as the highest recursive level and an evolutionary rate of change, the same underlying cybernetic principles have been demonstrated for organisations somewhat up-tempo for modern day industrial society. In this application, the work of Beer (1966; 1979; 1981; 1985) provides an substantial contribution to the cybernetic and management literatures (Jackson, 1988 & 1993; Espejo, 1989; De Raadt, 1991) and in which interest has been growing since the publication[124] of his various texts.

Within this thesis (especially in chapter 5), I place considerable reliance on Beer's Viable System Model (VSM). In view of this, I have expended considerable effort in the foregoing pages to lay the foundations from first principles and using the sources that Beer himself cites. In part this validation was a matter of academic prudence (ie, verifying that Beer's conclusions are supported on the strength of his analysis). Additionally, as the quotation below suggests, this validation was a necessary exercise to ensure that I had a firm grasp of the issues myself.

---

[123] Whilst an unusual request in a PhD thesis, I suggest that the reader omits this section on the first reading and progresses to section 4.4 (page 157, post) - this allows a direct connection between the "conventional" cybernetic principals rendered in section 4.2 and its subsections and their equivalent as rendered in the terminology used by Beer (eg, 1979). The current section serves as a preface to the account of Beer's viable system model.

[124] Inspection of both the science citation and social science citation indices (1981-1995) reveal increasing numbers of citations for all of Beer's publications. Taking Beer (1979) as an example; the social science citation index reveals 157 citations between the years 1981-95. Of these 157: 48 occur between 1981-85; 47 between 1986-90, and 62 between 1991-95.

Anderton (1989) observes:

> "Stafford Beer's work on the principles of 'viable systems' and their application to problems of organizational design constitutes material of great interest and importance, both practically and theoretically. However, in spite of extensive and brilliantly written publications, a large, diverse and enthusiastic following, and a significant number of attempts, often very successful, at application, it remains true that many people, including highly intelligent ones, attracted by the ideas, find them exceedingly difficult to grasp; they understand them superficially and find them cogent but when it comes to the point of detailed practical use they seem to slip away ... Solution of the problem which the above suggests would be made easier if a more explicit version of the underlying theory were to be propounded".
>
> Anderton, 1989, pages 40-41

As indicated previously (pages 24 and 84), before encountering the cybernetics literature I had found myself wrestling with the issue of systemic recursion and had come to some "natural" understanding of the general problem. This had the virtue of preparing me somewhat for what is probably one of the more difficult aspects of Beer's theorising; the constant demand to maintain *at least* two levels of recursion in mind at all times. Of further assistance was an adventitious approach to the VSM via systems theory[125]. However, it is only with hindsight, having studied various of Ashby's writings *after* many re-readings of Beer's texts, that Ashby provides the optimum *preparation* for comprehending the VSM. Thus whilst Anderton's (Ibid.) argument for a formal proof of the VSM speaks to a real need, the "proof" is, I suggest, already present (albeit camouflaged by an exceedingly demanding style of delivery) in Beer's main texts and the wider cybernetics literature from which it arises. Having said this, I must confess that there remain aspects of Beer's exposition that remain quite opaque to me. Where these aspects appear consequential to the present research I have noted them as requiring further work (in chapter 6). I have proceeded on the basis of adopting from the VSM, only those functional conventions and terminology, that I have been able to vouchsafe through the wider cybernetics literature and the process of my own reasoning. Given this last statement, I should like to make clear that I am *not* suggesting of Beer what Eysenck (1985) asserted of Freud: "All that is good in Freud is not new, and all that is new in Freud is not good". My contention is simply that some of what is new in Beer, whilst apparently not essential to the

[125] Starting with von Bertalanffy (1968), Weiner (1948) and Klir (1991). A further headstart was afforded by a grounding in psychophysics and neurophysiology which Beer (1981) uses to explain the VSM.

VSM, requires deeper consideration as to its validity and application in the wider gamut of managerial cybernetics.

Beer (1989), on his account, observes:

> "What was perhaps novel, for the record, was the *recognition* that the VSM homeostats requisite variety applies in three distinct ways: the blocks of variety homeostatically related, to the channels carrying information between them, and to the transducers relaying information across boundaries".
>
> Beer, 1989, pages 18-19 (*emphasis added*)

Turning now to critique of the VSM, whilst there is considerable controversy surrounding it this is of the paradigmatic type. This is to say that the critics of the model (notably Checkland, 1986) object to it on terms other than those on which the model itself operates. As stated earlier, a cultural approach will certainly be better able to capture and convey more of the "social reality" as perceived by members of an organisation. In these terms the VSM is simplistic indeed. For example, the attenuation of information flowing from a managerial to an operational sphere will be noted as a problem within the VSM. However, the cultural approach may reveal a long history of distrust as the reason why communication has become denatured. The question is - does this constitute a problem in the present context? Given the basic rationale - to develop an analytical framework with which to study the development of safety management systems - I do not see that it does. My empirical interest is to discover what *is* done in *functional* terms, in reference to an abstract framework of what logically is required, to deliver acceptable risk. Thus, the attainment of a clear analytical view in these terms, must be at the cost of a richer phenomenological description (the argument at page 113, ante, deals with this).

The criticisms of the VSM fall into three classes. The first of these, the comprehensibility of its exposition by Beer, has already been discussed. The second concerns the extent to which the VSM identifies the true determinants of system performance. The word *true* in this context reveals again the essentially epistemological nature of the debate (c.f. page 113, ante). In the VSM, Beer (Ibid.) claims to model the necessary and sufficient conditions for any system to be viable, that is, to be able to support an independent existence. As has been stressed in the foregoing, these conditions are ultimately concerned with cybernetic regulation - the principal currency of which is the timely exchange of information within a system and between a system and environment. As will be seen during

the description of the VSM; this process of information exchange is context-free[126]. It is because of this singular *virtue*, in my view, that the VSM has been the subject of considerable criticism from organisational theorists operating within the "cultural metaphor" (Jackson, 1989). Theorists such as Checkland (1986; 1991) stress the importance of the interpretive approach: understanding the social reality of systems as represented in the viewpoints of the organisation's members. Hence, from this perspective, an approach such as the VSM is apt to be seen as restricted to a view of organisations as machines[127] (Jackson, 1989; Checkland, 1980) rather than as definitively involving social groups, constructed, populated and continuously thought into being by their human members.

One of the curiosities of the debate, is the fact that there is no real contradiction except that which arises when the gear-teeth of two paradigms fail to mesh (as conveyed earlier in the quotation from Kuhn, 1970; page 113, ante). One aspect of this conflict arises from the purposes of the modellers. As indicated earlier (page 14, ante) the *goodness* of a model can not be judged by the extent to which it codifies reality but merely the degree to which it succeeds in serving the purposes of the modeller. For interpretivist theorists (as modellers) such as Checkland (Jackson, 1993) their goal is to understand organisations from different perspectives, therefore embracing the need to model a plurality of viewpoints:

> "It [the VSM] assumes that they [organisations] are instrumentalities. This is one legitimate view of an organisation, certainly. An argument for viewing an organisation as an instrumentality can always be mounted. But many other views are possible. In order to grapple with the problematical aspects of organisations it is frequently useful to them to be - as well as instrumentalities - social processes in which appreciative settings are formed and modified, and political processes in which accommodations between permanently conflicting interests are continually established and re-established."
>
> Checkland, 1986, page 270

What this adds up to is an indictment of limited usefulness of the VSM within protocols such as the Soft Systems Methodology (Checkland, 1991) in which it

---

[126] To the extent of applying to any viable system - including those of biology. This should come as no surprise given the roots of the viable system model in Ashby (eg, 1956;1962) and Sommerhoff (1950): a psychiatrist and a biologist respectively. Thus the context is irrelevant in the general case, just as mathematics can be used irrespective of what real world objects are represented by the notation and numbers.

[127] Which is a further demonstration of the difficulty found by many commentators in finding a convenient epistemological "pigeon-hole" for the VSM.

would be one model amongst a repertoire of models to be drawn upon as the circumstances on as selected by the circumstances of the particular study.

Checkland's (Ibid.) attribution of *instrumentality* is also worthy of note and brings into view the third class of criticism voiced in respect of the VSM: that it is a methodological device which may be used by organisations to refine their ability to oppress their members. I take it as a fact that the VSM, or any organisational application of cybernetics is aimed at better understanding regulatory effectiveness within the system. Regulatory effectiveness by definition has as its aim the reduction of variability in whatever aspects of performance are judged to be important. If it is the case that "what is important" is defined oligarchically and, by its assertion, has a negative impact upon the well-being and freedom of individual members, then there is little in-built defence[128] within the VSM to prevent it. However, this all reduces to the questions already discussed at section 3.6.5 (page 92, ante) of people in their dual roles as *means* and *ends;* of deontological and teleological ethics and the correspondence therewith of recursive regulatory structures. Also, the notion that organisational change based on VSM analysis is necessarily more disposed to deliver oppressive regimes than, say, cultural engineering approaches, is open to question. Consider the following quotation:

> The rhetoric of culture, however, indicates a shift in managerial sensibilities to a different form, one that Etzioni [1961] refers to as normative control. Normative control is the attempt to elicit and direct the required efforts of members by controlling the underlying experiences, thoughts, and feelings that guide their actions. Under normative control, members acts in the best interest of the company not because they are physically coerced, nor purely from an instrumental concern with economic rewards and sanctions...membership is founded not only on the behavioural or economic transaction traditionally associated with work organisations, but, more crucially, on an experiential transaction, one in which symbolic rewards are exchanged for a moral orientation to the organisation. In this transaction a member role is fashioned and imposed that includes not only behavioural rules but articulated guidelines for experience. In short, under normative control it is the employee's *self*—that ineffable source of subjective experience—that is claimed in the name of the corporate interest.
>
> Kunda, 1992, page 11.

[128] Other than Beer's constant entreaties that autocracy will tend to deliver non-viable systems and insistence (eg. Beer, 1989b and 1993) that democratic principles to foster viability as much for cybernetic reasons as psychological and political ones.

Whilst it makes for lively argument, I do not see the cybernetics and phenomenological approaches as *thesis* and *antithesis*. Wherefore then, *synthesis*[129]? One route to synthesis is the model-utility argument, the strength of which owes as much to pragmatism as to the logical premise that there is no final objective truth (eg. the most *truthful* model of *x* is *x* itself, as this is unusable: a *good* model is one which best represents the purposes imputed to the system by the modeller). Another route to synthesis lies in the work of Aulin (1982; 1989) and his cybernetic analysis of social structures. The essentials of his argument have already been presented in sections 4.2.5 to 4.2.7, in particular the establishment of requisite variety through recursive stages of regulatory amplification. Aulin (1989) refers to this as the *law of requisite hierarchy*. The "height" of this hierarchy is determined by the variety disposed at each regulatory level - the less variety disposed at each level the greater the number of amplifying stages and, hence, the "taller" the hierarchy required to achieve RV. This is not to say that extra hierarchical levels may not be established, *it is* to say that beyond the minimum number which ensures RV, these will be redundant with resulting inefficiency overall. Within each hierarchical level the individuals or groups of individuals are both *steered* (by the assimilation of beliefs, values and the development of norms[130]) and are *self-steering*[131] (as the individual and/or group continuously modify their beliefs, values and norms in the light of the empirical data they acquire through experience of the world). Recalling the need for an adaptive regulator to have "plenty of states" (page 139, ante) and the fact that the human brain has a superabundance of states, it is clear that the self-steering regulator can improve its regulatory performance[132]. If this improvement in self-steering regulatory capacity is sufficient then the pre-existing hierarchical depth may become inefficient. Further, this inefficiency whilst notionally redundant will tend, if insisted upon[133], to restrict the growth of self-steering capacity in the system. Given the implications of self-steering to personal freedom, creativity and locus of control, an oppressive social regime would be predicted.

---

[129] As per the Hegelian dialectic concept (Rinaldi, 1992) . For Hegel, this logical process goes beyond words to the social arena (it is *socially enacted*, as Weick (1995) would have it).: a thesis becomes an actuality which in time develops an opposition - an antithesis. This in turn will yield and may yield to a formulation - a synthesis - which unites the thesis with its antithesis.

[130] Corresponding to the feedforward selections made by the designer-regulator (section 4.2.7.3, page 148, ante).

[131] The adaptive self-organisation described in sections 4.2.7.2 (page 145) and 4.2.7.3 (page 148, ante)

[132] To limits imposed chiefly through transduction (ie, information handling) capabilities.

[133] Which in Beer's terminology is "pathological autopoiesis" - *autopoiesis* meaning self-production (Varela et al, 1974).

What Aulin's analysis clearly demonstrates is both the necessity of regulatory hierarchy and of individual creativity and personal growth at any level within it. Similarly, the presence of redundant levels in the hierarchy is shown both to add nothing to the regulatory capacity of the organisation whilst pathologically limiting the freedom and growth of individuals within it. Similarly, the restriction and manipulation of information available to self-steering elements (eg. through methods of normative control) provides a means of sustaining the existence of the highest levels of a hierarchy which otherwise would be redundant[134].

The matter of redundant hierarchical levels and the corresponding redundant capacity in individuals at each level is recognised in Senge's thesis of the "learning organisation" (Senge, 1992 and 1994). As Senge (1992) points out, the unused capacity of individuals in organisations may be an outcome of wider social conditioning which acts as a motor, perpetuating poor learning at both individual or organisational levels.

> Most people have grown up in an authoritarian environment. As children, their parents had "the answers". As students, their teachers had the answers. Naturally, when they enter organisations, they assume that "the boss" must have the answers. They are convinced deep down that people above them know what is going on, or at least they ought to know if they are competent. This mentality weakens them as individuals, and the organisation as a whole. At some level it absolves them of responsibility in the organisation's learning. It also predisposes them to cynicism when events naturally reveal that the people at the top did not have all the answers.
>
> Senge, 1992, page 282

Within Aulin's concept of requisite hierarchy, the cultural and cybernetic viewpoints cohere within one theoretical framework - a *synthesis*. If an organisation succeeds in obtaining acceptable performance, whether it be through the engineering of culture, the application of "scientific management" or by complete automation of process - it cannot do so in violation of the law of requisite variety upon which the viable system model is predicated.

---

[134] A rather clear demonstration of information as power and that normative control (eg. the engineering of culture) is one method of preserving redundant levels of hierarchy. This reasoning in the justification for the views of von Bertalanffy quoted at the head of this chapter. The cybernetics of the situation show the informational mechanisms by which a system is regulated within its environment. The same reasoning also shows that, in social systems, that the powerful must be obedient to these principals. Should they act with disregard for them, the system will either fail to regulate its essential variables (and thus be doomed) or ultimately be dissolved through a state of revolution (and thus be doomed).

4.4 The Viable System Model

Beer has presented his VSM using three different approaches: a mathematical account (Beer, 1959); A neurophysiological account (Beer, 1981); a management science account (Beer, 1979 and 1985). These later expositions contain, as Beer (1989) puts it:

> "a topological version of the original set-theoretic algebra that it seemed no-one would study properly. The drawings were now rigorous mathematics in themselves in that they offered explicit homomorphic mappings of any one VSM recursion onto the next..."
>
> Beer, 1989, page 13

As I am unqualified to comment on the set-theoretical account, the topological method of presentation will be used.

Scheme B of Figure 4.4 (page 123, ante) shows a representation of an error-controlled regulator. This is shown below, rotated clockwise through 90°, as Figure 4.11.

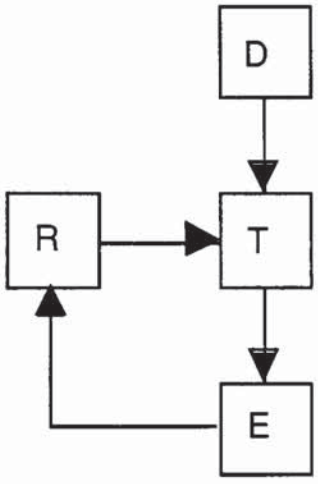

Figure 4.11 Error-controlled regulation (using feedback from essential variables)

Beer (1979) points out that the transformation table - *T*, can be regarded, using the nomenclature of electrical engineering, as a *Black Box*. The purpose of doing so is to stress that a representation of the internal configuration of black boxes, which is to say, the transformations they perform upon inputs to them, can be understood by associating the *inputs* to with the outputs from the box. In this way, a black box may be rendered "muddy". Beer (Ibid.) stresses that whilst a deterministic box of moderate variety may be rendered wholly transparent, a high variety Markovian box (such as that which contains human elements) may not.

In this way, Beer (Ibid.) arrives at his two "regulatory aphorisms":

> "The First Regulatory Aphorism: It is not necessary to enter the black box to understand the nature of the function its performs."
>
> (Ibid., page 40)

> "The Second Regulatory Aphorism: It is not necessary to enter the black box to calculate the variety that it potentially may generate."
>
> (Ibid., page 47)

Concerning the second aphorism, the calculation of variety depends on a sufficient period of observation to exhaust the variety of inputs which act as parameters to the box (including sufficient trials to reveal the statistics of the box from its outputs across the variety of parameters). It may be that the time involved is great and so it is reasonable to expect, in the majority of cases, that the model of the box will be at best muddy rather than transparent (ie, isomorphic).

The "aphorisms" are used by Beer to demonstrate that the contents of the muddy box cannot be managed by direct managerial action, that is, the manager responsible for a particular operational system lacks requisite variety for this task. This is not to deny that managers may have insight to *some* areas of the box (ie, certain areas are transparent) but overall they cannot be expected to have a mental model equal in variety to the box configuration. Thus, and with a certain implied humour, Beer (Ibid.) provides a schematic reproduced below as Figure 4.12.

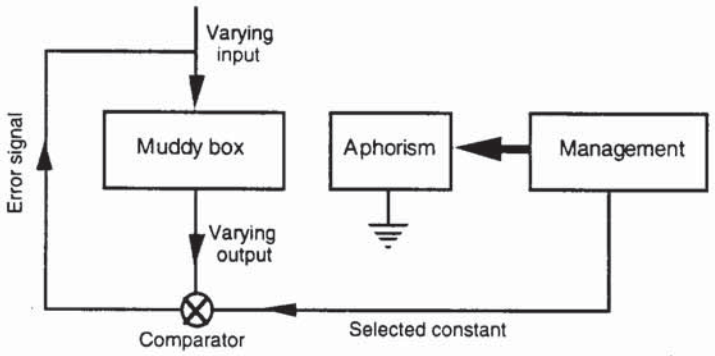

Figure 4.12 Managerial tendency to act against the "First Regulatory Aphorism"

"The management is shown as moving off [bold line] in the general direction of the muddy box that it is due to manage. It runs into the First Regulatory Aphorism[135]..."

Beer, 1979, page 59.

Other features of note in Figure 4.12, are the "Selected constants" (which corresponds to the of essential variable values and the ⊗ comparator symbol which compares value(s) of the current output from the box to the η-set selected by management. The comparator is a regulatory device shown previously as the lower portion of Figure 4.8 and Figure 4.9 (pages 140 and 141, ante, respectively).

Even at this elementary level, Beer's account shows considerable insight. One of the commonplace complaints voiced in the safety literature is the tendency for accidents to result in remedial measures developed by managerial "diving-in" to the muddy box without requisite variety and devising new controls based on their low-variety models. Thus the tendency for such low-variety solutions to fail to

[135] I have to confess to laughing aloud when I first read Beer's account; recognising in it a parsimonious account of my own managerial shortcomings in various organisations.

work in the face of high variety situations and to be either circumvented in practice or to create yet further problems[136].

The scheme at Figure 4.12 is still deficient insofar as management, restrained by the first aphorism, has no means of amplifying its variety to the task of regulating the outputs of the muddy box. Figure 4.10 (page 147, ante) shows this as a two stage regulatory arrangement and this notion is preserved in Beer's account as shown below in Figure 4.13.

Figure 4.13 Two stage adaptive regulation using terms of Beer (1979, page 64)

Recalling the discussion presented at section 4.2.7.2 (pages 145 to 148, ante), the role of the *feedback adjuster* is to transform the difference between the observed output value of the muddy box (T) and the desired value (ie, within the set η) into an appropriate regulatory input[137]. Similarly, the role of the *adjuster organiser* is to direct the hunt-and-stick operations of the *feedback organiser* to which end it also

[136] One is spoilt for choice of examples, but here is one from my own experience. An accident had occurred in which a firefighter had injured his back whilst raising a roller-shutter on an appliance rear locker. The immediate problem was that a fire beater (a long wooden pole with a piece of leather or rubber fastened to one end - used for beating out grass fires) had fallen in the locker, fouling the shutter. In response to this, the senior officer responding to the accident ordered that beaters should be stowed forthwith in another locker. What this missed is the facts that the roller shutter mechanisms had been found generally problematic for a variety of reasons and the alternative location that he had ordered to be used for the beaters (which are not always doused after use) also contained petrol vapours (from a petrol driven portable pump stowed in this locker).

[137] As noted at footnote 118 (page 145, ante): part of this transformation may need to incorporate the amount of time taken between D's action upon T and the information reaching the regulator via E.

receives feedback via comparator "B". However, as noted in the third paragraph of section 4.2.7.3 (page 149, ante), the success of this arrangement depends upon the design of the adjuster organiser being itself subject to revision in the light of improved information about the disturbances likely to act on the muddy box, *T*. As was described, this is achieved by invoking a higher level of recursion (Ashby's example of the gene pattern of the species as the higher recursion relative to the adaptive regulator embodied in an individual of that species). Corresponding to the regulator/designer operating from the higher recursion (ie, the **metasystem**), Beer (Ibid.) invokes a function which he calls the *Organisational landscape*[138].
Beer (Ibid.) in his commentary states that:

> "... The species does not 'tell' the individual what to do. And in real management, we do not observe the ordering about that the naive view of the organisation projects. ... On the contrary, our enquiry suggests that the role of the 'higher' level is to express a perception of the scene observed by the 'lower' level, that is actually inaccessible to this 'lower' level. That is why the 'higher' level is able to formulate the epigenetic landscape. Moreover: the 'higher' level has a language to talk which is of a different *logical* order from the language of the lower level. This partly because it has a way of observing the impact of external disturbances on the 'lower' muddy box, and accounting for what is happening in its own language – whereas the only words for these shocks that are available in the 'lower' language are such as 'OUCH' and 'HELP'. And it is partly because it is the recipient of a landscape from *its* 'higher' order, which (naturally) cannot be transmitted to its 'lower' order – because it would have no meaning there. ...
>
> A better account of the 'higher' level is to call it METASYSTEMIC to the 'lower' level. 'Meta' means 'over and beyond', referring to the perception and the logic, and not to seniority."
>
> Beer, 1979, pages 68-69.

This is shown below at

[138] Adapted from the term "Epigenetic landscape" coined by Waddington (1957) and cited by Beer. Waddington's notion, as explained by Beer, is that evolution cannot occur wholly through random mutation as this would be too slow to account for the pace of evolutionary change, therefore, the gene-pattern must be subject to non-random mutation. Hence a "flat landscape" corresponds to random mutation (or experimentation) whereas a epigenetic landscape corresponds to a non-random or predisposed mutation.

Figure 4.14, the bold box demarcating a "Management Unit" and, thereby, a single recursion. Over leaf, Figure 4.15 shows a recursive hierarchy of these management units.

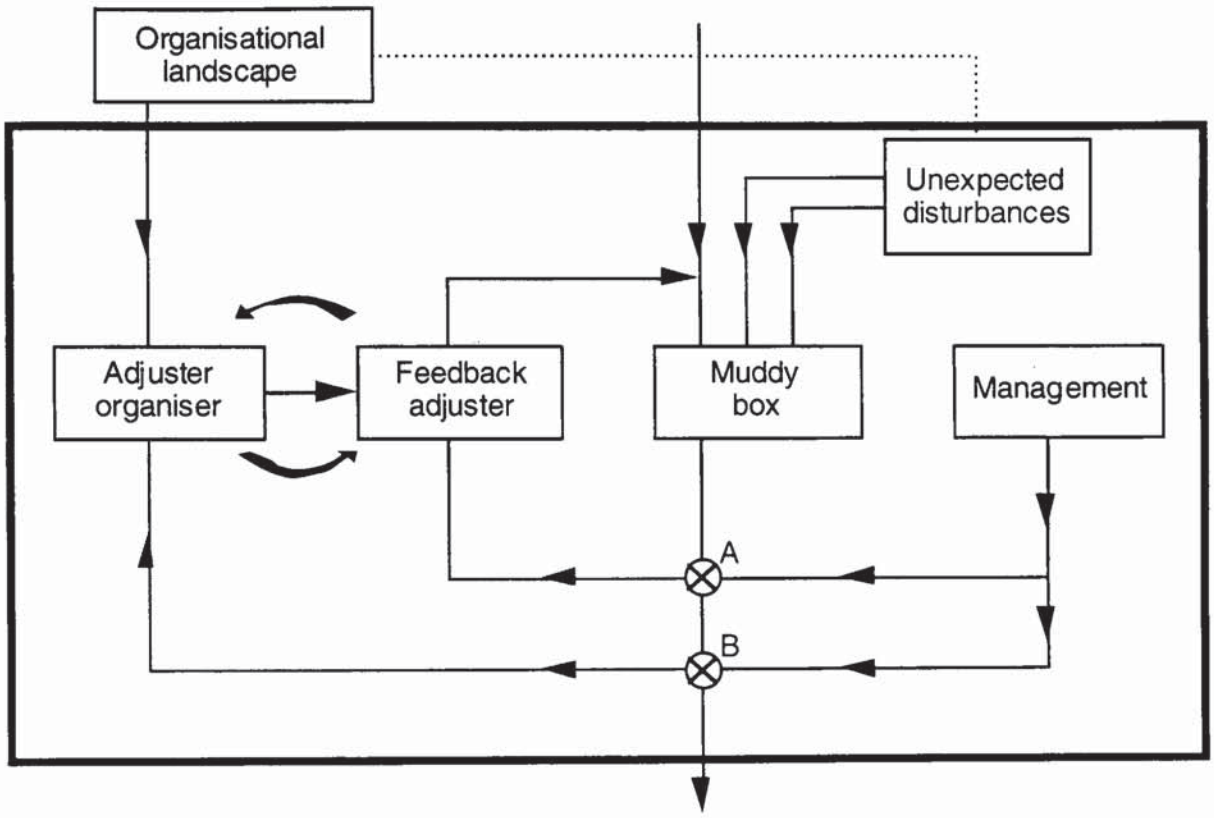

Figure 4.14 Adaptive regulation showing metasystemic input

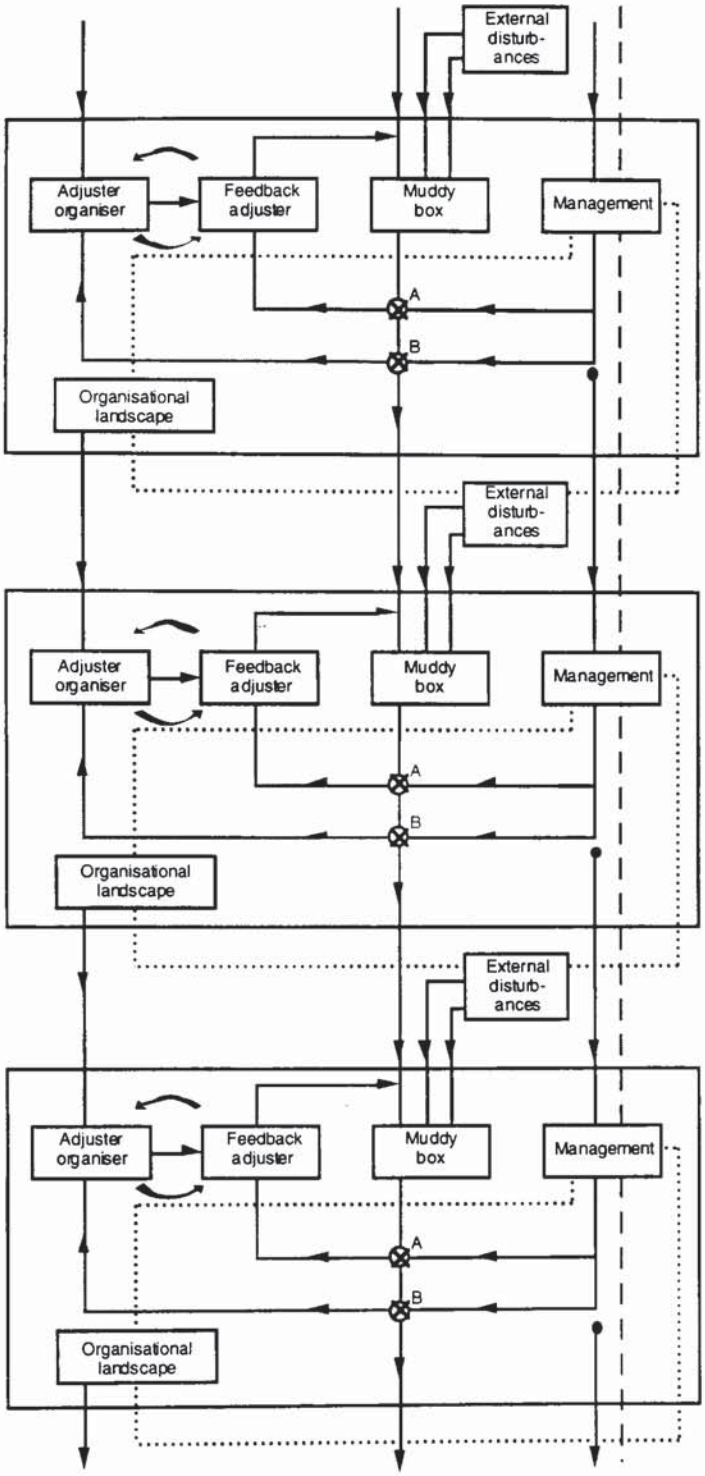

Figure 4.15 A recursive hierarchy of Management Units (After Beer, 1979, page 71)

The meaning of the discontinuous line (broken by a ●) between the hierarchical "Management" boxes on the right-hand side of Figure 4.15 is that only some of the contents of the η-set are selected by the metasystem, others are set at system level. The continuous (dashed) line connecting the "Management" boxes denotes the personal dialogue between them.

Before moving on to the next set of concepts, it needs to be made clear what the muddy box corresponds to. Whereas the *T* function, in the discussions previously given, corresponded to a physical system; in Beer's account, the *muddy-box* denotes a *mental model*[139] held or shared by management *of* a physical system. Thus the managerial cognition using muddy box representation must be transduced into the *real world* objects, actions and attributes to which this representation corresponds. In other words, management act *through* their model of the world and receive their information of the world as mapped onto this model (as discussed in relation to perception on pages 67 and 101, ante). The importance of the mental models to the performance of management cannot be over-emphasised whether in Beer and elsewhere. Senge (1992), for example, makes the explicit consideration of managerial mental models one of the necessary conditions for a learning (i.e. adaptive) organisation.

In terms of regulatory variety, each successive (lower) level of the recursive hierarchy acts as a variety amplifier for the level above. Equally, thinking now from the design perspective (as discussed at section 4.2.5.1, page 134, ante), each selection made by the higher level of the lower represents a potentially massive curtailment in the variety of the latter. For example, imagining a corporate HQ and its division, if the former says to the later "we require a five percent reduction in accident rates" this is a very low variety statement which will be massively amplified by the consequent risk assessment and risk reduction activities (both high variety) undertaken by the division. If, on the other hand, HQ introduces an procedure (say, for manual handling) this will tend to reduce the variety of the receiving system (both in the way things are done at the operational level, and in range of options open to local management in workplace layout, how they allocate resources etc). Given that the primary purpose of each recursive level is amplificatory, then it becomes obvious that the meta- level must be judicious in how it delimits system-level variety (1) because it is curtailing the degree amplification that can be provided and (2) because it runs the risk of making such selections from a lower than requisite variety model.

The situation described suggests something of a fine balance; this is true both cybernetically and psychologically. Indeed, engineering this balance between amplifying and attenuating variety, can be said to characterise the primary concerns of the VSM. In relation to H&S management, this balance is, if anything more difficult to strike: on the one hand, it is widely accepted that there be

[139] Or computational model in a non-human regulatory device.

"ownership" of the problems at local level (ie, adaptive change characterised by feedback), whilst on the other hand, prescriptive (and even safety-case style) regulation by the state fosters a top-down (a massively feedforward approach characterised by "fragility"). This does not mean that it is mistaken for the state to require organisations to demonstrate control (as prescribed conditions or processes), merely that what counts as substantive "proof" arises from a reductionist and non-systemic viewpoint. Further, even where the State requires demonstration of practices consistent with adaptive ability (such as training, consultation[140], review of risk assessment, etc) how is it to judge whether such functions provide requisite variety or are merely a "laundry list" of requirements included to jump through the regulatory hoop?

4.4.1 VSM elemental systems

In order to obtain an orientation from which to explain, consider the following imaginary example: three shops in the high street of a small town. The first shop "Fixit" offers a repair service for domestic electrical appliances. The second shop "Machine-Mart" is a retail outlet for a wide variety of domestic electrical goods. The third shop "Kitchen-Co" provides fitted kitchens. Each of these is privately owned and run by individuals who know each other well and have made a practice of directing customers towards each other.

After some time, all three business are prospering but each of the owners are facing the problems of being large enough (in terms of trade) to create considerable difficulties in terms of paperwork but not quite large enough to justify hiring someone to manage this side of the business. The three owners, whilst discussing their common experiences hit upon the idea of becoming partners - as together they can afford to hire the extra staff that alone they could not justify financially. They also realise that the businesses can be combined synergistically: "Fixit" can provide warranty cover for the goods provided by Machine-Mart and those installed by "Kitchen-Co"; "Machine-Mart" can supply "Kitchen-Co" with its appliances taking advantage of the greater discount offered by wholesalers for the larger volume of orders. For these reasons and also tax advantages, the three partners form a new company "FKM Ltd", but decide to keep their old names on the shop fronts for the time being so as to maintain their pre-existing good image with their customers.

---

[140] European Community Framework Directive 89/391/EEC as implemented by "The Health an Safety (Consultation with Employees) Regulations 1996".

From the vantage point of the VSM each of these three businesses is a *viable* system - not merely could they separate in the future and maintain an independent existence but cohere together in a single *viable* system (FKM Ltd). Looking at the whole (ie, at the level of FKM) each shop is termed an **elemental subsystem** (or **element**) and the three elemental subsystems taken together are **System One** of FKM Ltd. It should be noted that three elemental subsystems taken together do *not* constitute FKM as there are also the various corporate aspects of FKM which are extra to the three elements. Thus in addition to *System One* (the collective term for the three elemental subsystems), there are corporate functions which, together with System One, "add-up" to the viable system called "FKM". Summarily, these corporate functions concern the synergistic operation of the three elements (corporately seen as the single System One); the forward planning and market research activities of the corporation as a whole; and the policy-formulating partnership (consisting of the three partners). Together, these corporate functions are the **metasystem** of System One.

If we keep our system-in-focus as FKM, but consider its subsidiary Kitchen-Co, the managerial interactions between it and FKM (its metasystem) are said to take place on the **vertical** axis of the VSM. This is in contrast to the interactions at the level of Kitchen-Co (its local management, operations and the environment served by its operations) which are said the take place on the **horizontal** axis of the VSM.

4.4.1.1 Elemental subsystems and the horizontal axis of the VSM

In the following diagrams, the management units shown singly in Figure 4.14, correspond to the square box below at Figure 4.16.

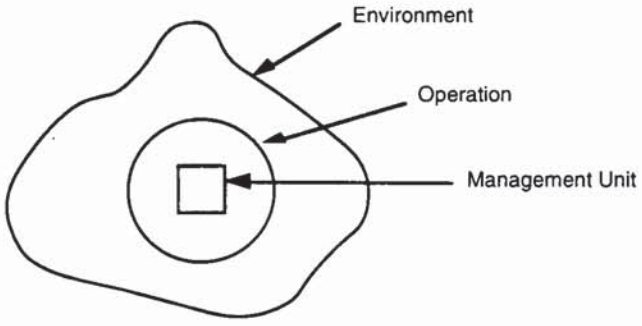

Figure 4.16 VSM graphical conventions: Environment, Operation and Management Unit

Beer (eg. 1979) uses the conventions shown in Figure 4.16, to represent a management system embedded in its operation, itself, embedded within the environment. The sizes may be taken as token indications of the variety each domain possesses: ie, *Management Unit* < *Operation* < *Environment.*

Given what has already been demonstrated, the natural variety imbalances between these functional elements will tend to assert themselves as recognised in the law of requisite variety. As has been argued (section 4.2.4) whilst there are other options to be exploited to address this imbalance, chief among them is variety amplification: thus the regulation of the operation by its management must be amplified as must the regulation of the environment by the operation. It follows from this that the feedback reciprocal to this amplification must be attenuated: otherwise the limited variety of the receiver will be swamped. To give a simple example: we (the Health and Safety Unit) recently decided to amplify our variety to the task of tutoring our students (especially distance learners) using information technology. In brief, we created multiple choice questionnaires to allow students to evaluate their performance at each stage of the courses involved using the medium of the Internet. Thus amplified, a very large number of students can now sit tests at their convenience. However, without further amplification we would be swamped by the completed questionnaires requiring to be marked. Again using IT we developed automated means of marking these tests and giving feedback based on the pattern of responses in *each* case (thus maintain requisite variety).

Now we faced being swamped by copies of the results (which we needed to update our model of *each* student's performance). The results required attenuation that, whilst reducing the amount of data to manageable proportions, *preserved* variety. The variety in question is equal to the variety of residual states from the automated testing loop: that is, students gaining very low marks, students performing satisfactorily, students not sitting tests by the dates suggested. Thus, we needed a range of marks and a schedule (for progress through the course syllabus) which, taken together, provide the criteria of stability for *each* student. Only indications of instability (low marks or not attempting MCQ tests) result in a message being transmitted to the tutoring staff, updating their model of *each* student's state from "progressing satisfactorily" to "there are problems".

In summary, the students' perception was of feedback tailored to their *individual* performance: whether through automated amplification or enquiries initiated by the tutor (again based on their *individual* performance). Hence by a combination of amplification and attenuation *every* student's requirements (large variety) could be balanced by the limited variety resources available to the HSU[141].

[141] I should add that whilst this method goes a long way in meeting requisite variety, other methods are used to meet the margin - especially when Beer's "Ouch!", "Help!" messages are received.

Another example, this time emphasising attenuation of environmental variety, is the action of the Fire Service through the inspection of buildings: the aim being to both ensure adherence to the various Buildings Acts (ie, reduce the variety of non-conformance to standard fire precautions and thereby attenuate the number and severity of incidents they are require to attend).

The concepts and symbols for amplification (⊲) and attenuation (⋀) adopted by Beer (eg. 1979) from electrical engineering are shown at Figure 4.17. The grey line between the management unit and the environment is added for completeness although it is given less status within the VSM. Beer (1979) says of these

> "They exist, even though managers who explore their environments may perforce work through their operational models at the psychological level."
>
> Beer, Ibid., page 127

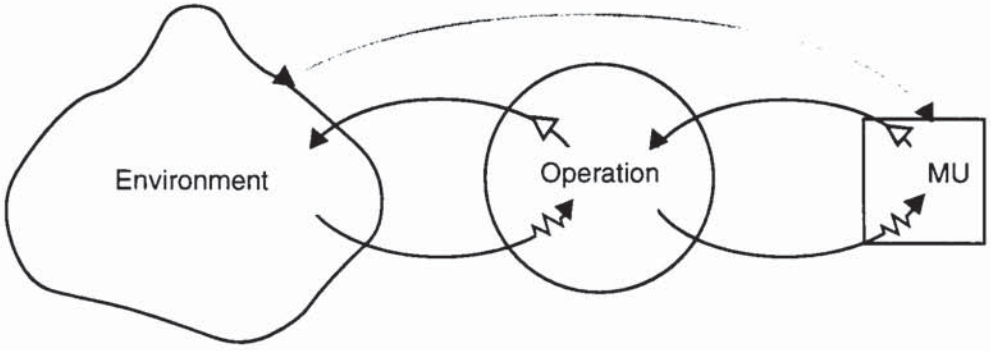

Figure 4.17 Variety Amplification and attenuation on the horizontal axis

Beer (1979) summarises his account of amplification and attenuation between these three functional elements in his *first principle of organisation*:

> "Managerial, operational and environmental varieties, diffusing through an institutional system, tend to equate; they should be *designed* to do so with minimal damage to people and to cost".
>
> Beer, 1979, page 97.

In keeping with information theory, (as summarised in Figure 4.1, page 118, ante) each looping arrow of Figure 4.17 above, assumes the **channel capacity** and **transduction** arrangements as discussed in section 4.1 (pages 114-119, ante).

Transducers, as described earlier, should neither amplify nor attenuate variety but merely transform one code into another without loss. For example, the erstwhile British Rail "Rule-book" was intended by its management to amplify the low variety message "act within the rules" into a high variety set of IF-THEN routines for staff within the operational circle. Thus, in appearance at least, the condition for managerial amplification to its operation has been satisfied. Leaving aside the requisite variety implications of the rules themselves: what of channel capacity and transduction in this example? Channel capacity has been ensured by publication arrangements, however the transduction across the operational circle boundary (ie, decoding from the printed word [reading] and the further transformation through and of the reader's conceptual model [comprehension]) was found to be problematic. A crude measure of this kind of transduction is *reading-age*. In the case of the people within the operational circle, the *mean* reading-age was found to be **nine** years old. This, in itself, is not a problem so long as the encoding of the rule-book was compatible with the variety of the decoders possessed by the receiving staff. However, when investigated, this was found to assume a reading age of **21** years. Thus, for the majority of staff there was no means of preserving the variety of distinctions (encoded through vocabulary and syntax) in the message. The result is that whilst information *defined as the book of messages* has apparently been received, information, as *defined as changing the recipient* in the intended way has not.

An example of channel capacity problems, this time in the opposite direction (ie, messages transmitted from the operational circle to the management unit), might be an audit report. One objective of such reports may be to enrich the variety of the conceptual model (the muddy box) that management relies on for regulation. The great danger is that whilst the recipient may well have the necessary decoding rules[142] (ie, special vocabulary, command of language generally) the transduction takes *time* to achieve. Given that there are likely to be a great number of messages competing for managerial attention, the result is a much constrained channel. If the manager is unwilling or unable to make such time as is needed available - the result is much the same as the previous example of limited literacy[143].

---

142 I have seen examples where this is not true in relation to risk assessment reports. On these occasions managers could not transduce the variety of the reports although their subsequent decisions were supposedly informed by them.

143 Another danger, such as that noted by Eisner and Leger (1988), is that time-pressured managers receiving an ISRS (International Safety Rating System) audit report may well rely upon the scoring system - which, in the case of ISRS reduces attenuates a massive variety to 10 states (five points on the standard, and five points on the advanced scheme) and possibly even less as it is attenuated to a 3 state "better, worse or the same as before".

An example of variety lost during transduction into the management unit, arises from personal experience of accident reporting. Despite earnest attempts to preserve variety in a brief (ie, attenuated) analysis and presentation would often be met by "but, who was to *blame*": even though my words were understood, the conceptual model could not distinguish as many states as the variety of the situation presented.

Concerning **channel capacity**, this needs to be as large as the greatest demand placed on it. This obviously has implications for the handling of *unforeseen* abnormal situations (foreseen abnormal situations *should* have an attenuating protocol... a fire alarm, for example, represents a 1-bit variety but will activate a high variety plan). An example of this condition not being met was the computer system relaying information to the controllers during the Three Mile Island accident (see footnote 119, page 145, ante) which was so constricted as to lag message delivery by a matter of hours. Similarly, the recommendation by Hidden (1989) of two-way radio communication for train crew is an attempt to provide requisite channel capacity.

The question of channel capacity and transduction are summarised by Beer (1979) as the second and third "principles of organisation":

> **"The Second Principle Organisation:** The four directional channels carrying information between the management unit, the operation and the environment must each have a higher capacity to transmit a given amount of information relevant to variety selection in a given time than the originating sub-system has to generate it in that time."
>
> Beer, 1979, page 99

> **"The Third Principle Organisation:** Wherever the information carried on a channel capable of distinguishing a given variety crosses a boundary, it undergoes transduction; and the variety of the transducer must be at least equivalent to the variety of the channel."
>
> Beer, Ibid., page 101

Beer's ***fourth* principle**, is that the operation of first three principles stated above must be "maintained through time without hiatus or lags" (Beer, Ibid., page 258). This is, to some extent, both obvious and unrealistic - does there exist anywhere an organisation that might survive the test of the fourth principle? In my view, the fourth principle is made realistic when one acknowledges that the communication involves transmission of **difference**, thus it is only when change occurs that

information can be transmitted. Further, the parameters of change must also be a factor - if *all* change is transmitted on *all* communication lines the result is unavoidably gridlock. Similarly, if the gain on the line is adjusted too low (ie, only report of huge change transmitted) then two things can be expected: firstly, small changes heralding the large change will not have resulted in the regulatory adaptation that might otherwise have maintained regulatory requisite variety (as per the discussion at page 146, ante). Second, if a communication channel continuously registers a single value (eg. no change) the transducers associated with it tend to denature[144].

When Beer's fourth principle is viewed alongside the various concepts summarised (ie, *information-as-difference*, adaptive regulator dependence upon reliable registration of change, variable gain, maintenance of transducers) it can be seen to be both realistic and imperative. Lastly, it should not be forgotten that the options laid out in sections 4.2.4.1 to 4.2.4.3 (pages 130 - 132, ante) are always available: if the practicalities of communication are beyond what can be achieved to deliver the regulation of essential variables as defined with the set η, then the firm may be forced to accept a more forgiving η-set. Beer's contention is, in my view, that once an organisation has *decided* what is important and what is wanted (as per the requirements of Ashby (1956), page 125, ante) it has little option to deal with the realities consequent to such choices.

---

[144] In behaviourist psychology the phenomenon is referred to as "habituation" to a constant stimulus. Organismic defences to the denaturing of transducer function are various. The best example I know of it is implied by the Gate Control Theory of pain perception (Melzack and Wall, 1965). Without going into the intricacies of GCT, it provides explanation for a range of phenomena such as traumatic analgesia, chronic pain and simple phenomena such as the tendency for vigorous rubbing of a painful area to give relief from pain. What is thought to occur (and, over the years, has been both supported histologically and empirically through treatment) is that pain information will be relayed to the brain only when the *ratio* of nerve impulses relating to afferent (normal sensory information: temperature, pressure, etc) and nociceptic (abnormal sensory information corresponding to physical trauma) changes. In other words, there are always some nociceptic messages being sent, even in the absence of external stimuli - but these are ordinarily balanced by afferent information at the relay "gates" in the spine. Whilst this is conjecture, it seems to be logical that the survival value of nociception is so great as to explain the evolutionary selection of a system which maintains *tone* in the channels concerned. Further, the value of the ratio at which nociceptic messages are relayed up the spine to the brain can be varied by "descending control" mediated by the control of neurotransmitter release into cerebro-spinal fluid. Thus the gain of the system (number of messages relayed) can be varied by control at a higher neural level - even to the extent of inhibiting spinal reflexes (Kandel and Schwartz, 1991).

4.4.1.2 Multiple elemental subsystems and the horizontal axis of the VSM

Having now introduced the basic conventions for a single elemental subsystem and the operation of variety exchanges within the subsystem and its environment, the instance where System One is composed of multiple elements is considered. It should be noted that the elemental subsystems depicted are *peers,* just as within the FKM example, "Machine-Mart", "Kitchen-Co" and "Fixit" were peers regardless of their relative turnovers.
The occurrence of multiple *peer* subsystems (ie, on an equal footing at a given level of recursion) creates the need for further types of communication linkages to be brought into consideration. Taking the illustration of a single sub-system given in Figure 4.17, the choices available for the graphical depiction of multiple peer systems is limited. Beer and other users of the VSM have opted to use the convention illustrated below in Figure 4.18.

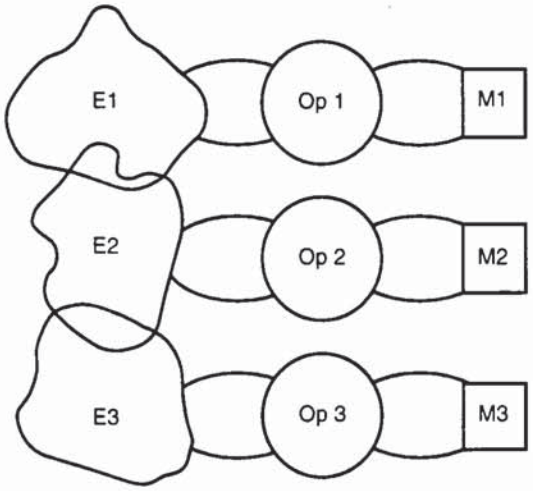

Figure 4.18 VSM Conventions for depicting multiple peer subsystems

The danger here is of misconstruing the relationship between the peer management units as hierarchical: it is not. An attempt to convey this is given at Figure 4.19. Here the large cube indicates the metasystem (eg. FKM Ltd) shared by three peer subsystems (eg. "Kitchen-Co", "Machine-Mart" and "Fixit").

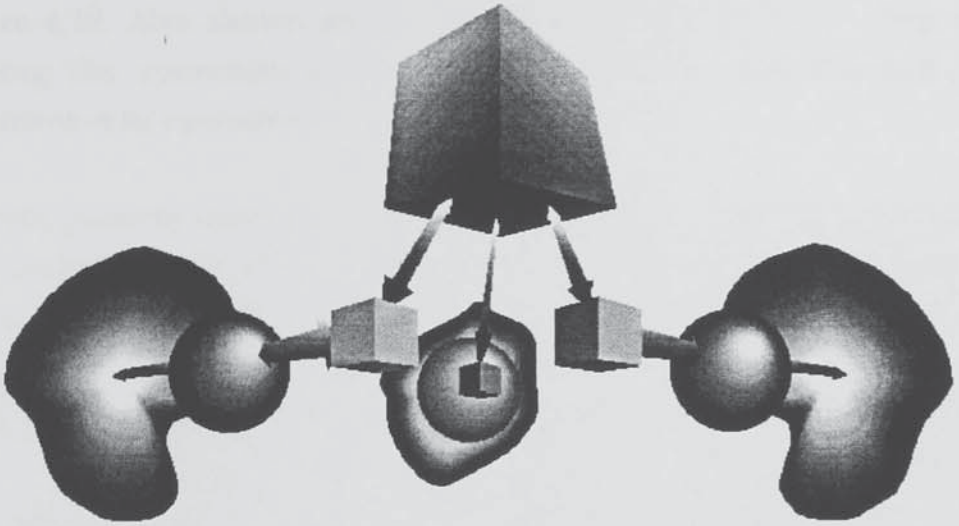

Figure 4.19 A "three dimensional" representation of peer subsystems

It follows then, from the perspective of the metasystem that these subsystems aggregate to a single system: **System One,** where each autonomous subsystem is an **element** of System One. However, some of this autonomy must be sacrificed to the cause of the coherence of the whole. The expectation is, therefore that these elements will, ordinarily, be interconnected. As shown in below, the connections occur within three **domains** - those within the environmental domain, those within the operational domain and those with the managerial domain.

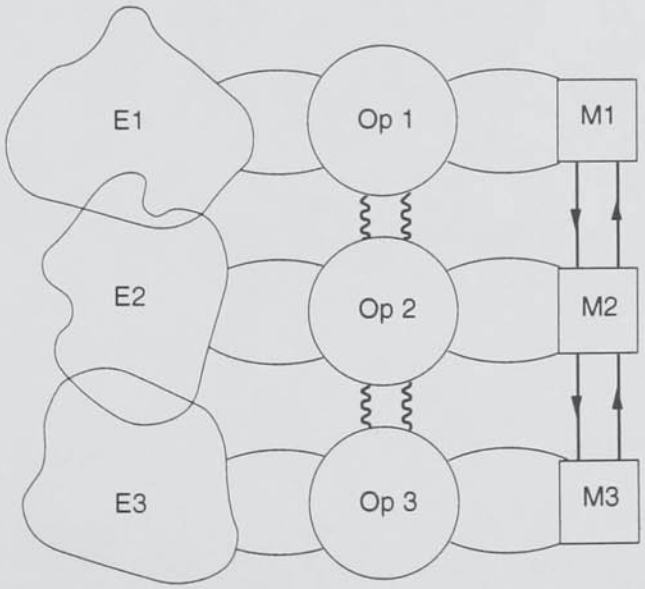

Figure 4.20 Elemental interactions in the environmental, operational and managerial domains

The connections shown between subsystem 1 and subsystem 2 are implied also between subsystem 1 and subsystem 3. To make this point clear, Figure 4.21 shows these domain connections from the synoptic perspective of the metasystem

in Figure 4.19. Also shown are the sub-systemic connections within elements (including the connections between the each management unit and the environment of its operation).

It is worth pausing momentarily to reflect on the complexity unfolding at this level of analysis. In Figure 4.21 there are three elemental subsystems, each of these has three bi-directional channels (shown as loops) making 9 channels on the horizontal axis. Then there are the domain connections (linking each environment, each operation and each management unit); again 9 bi-directional channels. The tally so far is 18 to which can be added the three bi-directional channels linking the elemental management units to the metasystem, making 21 in all. For each of these the first three principles of organisation[145] need to be assured meaning, at this stage, with three elements we have a list of 63 cybernetic checks[146].

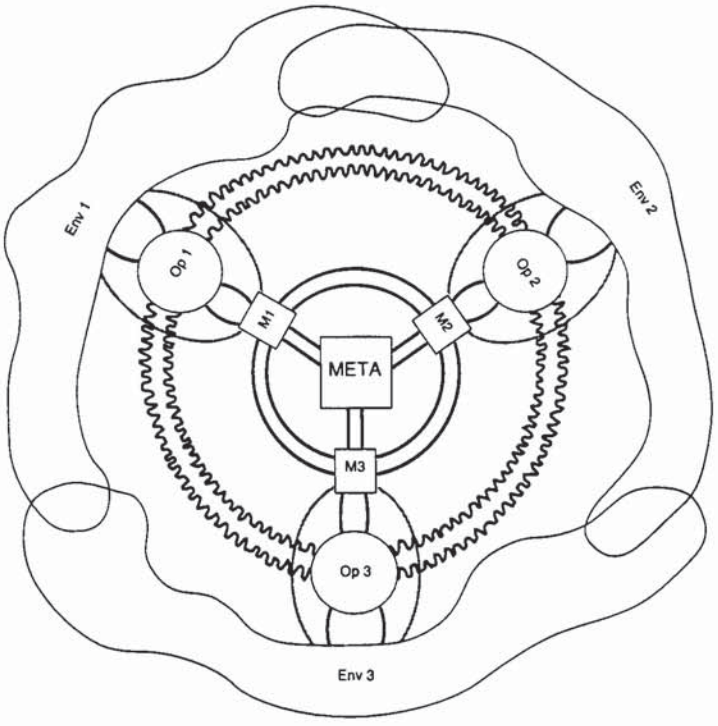

Figure 4.21 Synoptic view of interconnection of the elements of System One

[145] Requisite variety, channel capacity and transduction integrity, respectively

[146] The calculation is achieved by adding the sum of **horizontal** connections to the sum of **domain** interconnections. Each elemental subsystem has three bi-directional channels *plus* one channel linking it to the metasystem, thus *four* channels *per* subsystem. If n equals the number of elements the $4n$ gives the number of **horizontal** channels. Each elemental subsystem has three domains. The number of connections between n objects is given by $n(n-1)/2$. Therefore the total number of domain connections is $3(n(n-1)/2)$. Therefore, the total number of channels (Σ) is given by

$$\sum_{Channels} = (4n) + 3\frac{n(n-1)}{2}$$

Were there five elemental subsystems then 50 channels would exist (150, cybernetic checks).

Whilst the "squiggly line" (sic) connections exist between operations, they may vary considerably in their strength. In a chemical plant, these might represent physical inter-process connections. Alternatively they may be weak to the extent of *appearing* absent. Beer makes an illuminating observation concerning apparently weak linkage:

> "...the separate divisions are at the least competing for capital. They may also be competing for managerial talent, trained within the metasystem. It is within these circumstances, whatever they happen to be, that the three principles need investigation. ... It has to be remarked that organised labour often appears (to me) to act with an instinctive understanding of cybernetic principles that is lacking in the boardroom. To account for the observation is not so hard. It is **inordinately easy to annihilate variety in the boardroom** – it can be done by clearing the throat. **On the shop floor it is much more difficult**."
>
> Beer, 1979, page 127 (*emphasis added*)

Concerning the strong inter-operational linkage:

> "I have often noticed that that the variety-exchangers at the operational linkage [ie, the squiggly line connections] are far richer, more accurate and speedier, than the bureaucratic variety-exchanges between the managerial units – which are supposedly controlling the action. Of course, the reason is that the operations are their own representation... they furnish their own requisite variety. That variety is necessarily attenuated in the office representation of reality."
>
> Beer, 1979, page 128

This quotation is supportive of the point made in chapter 1 (section 1.2.1, page 29. ante) where the notion of individuals "pumping-in" the marginal variety not provided by extant organisational arrangements. It is not the case that such arrangements are a sign of poor organisation, indeed, Galbraith (1973; 1994) sets out the case for maximising organisational effectiveness by optimising these "lateral" arrangements. Contrarily, in H&S, the situation is often encountered where the managerial organisation is acting as though the squiggly-line self-organisation does not exist. Thus, the operational generation of regulatory variety becomes a kind of "black-economy" where procedures, originated metasystemically and without requisite variety in the generating models, are adjusted to fit with operational demands. Taking a strict viewpoint *all* such

adjustments constitute violations although it is only when they result in visible failure might they be called as such[147].

4.4.2 Metasystemic implications and System One

As was suggested in section 4.2.4 (page 129, ante) the disturbance to be regulated against in the large system is likely to be composed of the variety arriving from the environment and the variety generated by enclosed subsystems, where the latter have some degree of independence. As suggested in the section 4.4.1.2, the independence of the elemental subsystems is likely to be large given that they are, so far as the cohesiveness of the whole System One allows, autonomous. The degree of autonomy possessed by each elemental subsystem is decided by two factors: (a) as suggested, the cohesiveness of the whole enterprise as overseen by the metasystem (considered in section 4.4.2.1) and (b) the effectiveness of the operational domains, which includes the operational interactions of all the elemental subsystems comprising System One (section 4.4.2.2, page 178, post).

4.4.2.1 Metasystemic and elemental subsystem linkage.

In order to reduce clutter in the VSM diagrams, the loops are replaced by lines which are bi-directional unless otherwise indicated. The solid circle (•) terminating the lines in Figure 4.22 denotes a transducer.

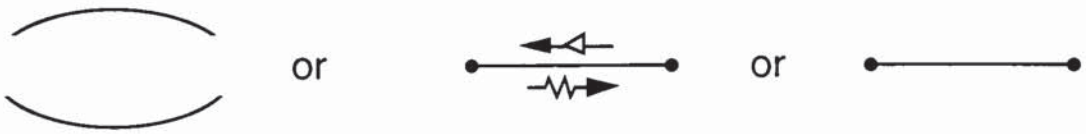

Figure 4.22 simplifying graphical conventions for channels

[147] The Fire Service has a special brand of exceptional violation, when such adjustments *succeed* during operations they are called acts of *heroism* and are *praised*. When the *same* adjustments to operating procedures *fail*, they are *violations* and result in *disiplinary* action.

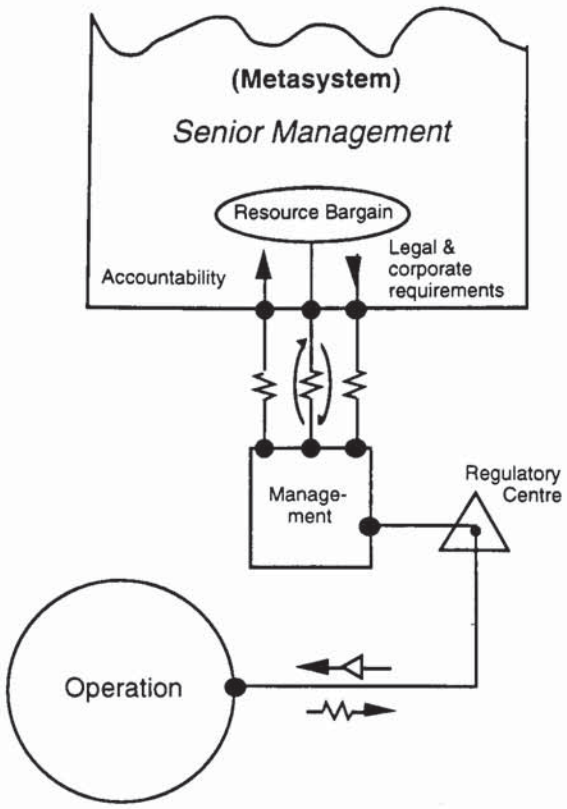

Figure 4.23 Fundamental linkages between senior and local management

Using these conventions, Figure 4.23, below shows a single elemental subsystem and its metasystem. The three lines between the management unit and its senior management (part of the metasystem) one can see that the *net* effect of all is attenuation of the management unit's variety. Beer (1985) states:

> "... senior management often assumes - and likes to exercise - the power to poke around in the intimate managerial detail of its subsidiaries... But *think*: the so-called prerogative to intervene indiscriminately does not have requisite variety. It cannot be competently done."
>
> Beer, 1985, page 38

As previously noted (page 163, ante) the metasystemic interventions to system level management have to find the balance that permits the latter to provide amplification to requisite variety (in the regulation of the operations within its purview) whilst attenuating only that managerial variety which sends the essential variables (those determined within the metasystem) outside of the η-set.

Taking first, the channel labelled ***legal and corporate requirements***, this would include statutory requirements, policy requirements as well as, perhaps, the specification of the technology to be used in the operational circle "belonging" to the local management. Elsewhere (eg. Beer, 1979; 1981) this channel is given the more general title of the "command channel".

The ***resource bargain*** is elegantly portrayed by Beer in the quotation provided below. It is worth emphasising that the resource bargain is a *function* rather than an *event* - and this function is a very frequent rather than occasional operation. Perhaps resource bargain*ing* might be a more appropriate term.

> "*The Resource Bargain* is the 'deal' by which some degree of autonomy is agreed between the Senior Management and its junior counterparts. The

> bargain declares: out of these activities that System One elements might undertake, *these* will be tackled (and not *those*), and the resources negotiated to these ends will be provided. The homeostatic loop sketched into the diagram properly indicates that a **dynamic** process is involved. It is essentially attenuative because it excludes a huge range of alternatives. This is not to say that the senior management never provides variety amplification to the junior enterprise within the attenuating scheme: it may, by superior knowledge or through unexpected financing, open up opportunities not conceived by System One on its own initiative."
>
> Beer, 1985, pages 38-39.

The ***accountability channel*** is closely associated with the resource bargain homeostat, insofar as the bargain sets the attenuated parameters for accountability (much in common with management by objectives).

> "...think about this responsibility for resources provided... in terms of variety engineering. Can you possibly itemize every single thing that the subsidiary does, demand a report on it, and expect a justification? Obviously not. Therefore accountability is an **attenuation** of high variety happenings."
>
> Ibid., page 40.

The ***regulatory centre*** (marked by the triangular symbol on the right of Figure 4.23) is the means by which management co-ordinate the action of their operation[148]. In essence, this is amplification of the low-variety agreement reached in the resource bargaining process to a higher variety description (eg. plans and procedures) which allows the operational implementation of the agreement. For example, if the agreement concerned a programme of training for operational staff this would need to be amplified into a schedule which released staff for training without interfering intolerably with production. Additionally, the return loop to the regulatory centre from operations carries **monitoring** information relative to the amplified resource bargain. This allows local management to ensure that it abiding by its agreement and to (under further attenuation) provide reports back along the accountability channel to senior management. From the operational perspective, the regulatory centre is the collection of service functions which enable and co-ordinate its work.

---

[148] There is rather more to the regulatory centre than clarity permits explanation of here - in particular its anti-oscillatory role at the *next recursion down* which will be discussed in the following section)

4.4.2.2 System One and anti-oscillatory metasystemic regulation

Much of the interaction in the horizontal domains (eg. as mediated by the so-called *squiggly line* channels) may be necessary and desirable. However, some of this inter-elemental exchange may allow System One as a whole to go into oscillation.

Using the FKM example, let us suppose that "Machine-Mart" has built-up stocks of a particular brand of dish-washer and wishes to clear them to make way for new models which have appeared on the market. To assist this it advertises a special offer and manages to clear its stock. "Kitchen-Co" has five fitted kitchen installations already "on the books" to be carried out this week: part of the specification for these kitchens is the model of dishwasher no longer stocked by "Machine-Mart". So as not to delay the installation, "Kitchen-Co" agrees to use the new type of dishwasher on the understanding that "Machine-Mart" absorb the extra cost of these machines (£70 per unit). However, this agreement failed to establish that this absorption concerned not merely the five imminent installations but all of the orders based on the old price (25 orders, or £1,750). Meanwhile, the Kitchen-Co plumbers and joiners have prepared the sites and the engineer from "Fixit" arrives to find that (a) the hot and cold feed laid-on by the plumber does not correspond to the new machine (which only has a cold feed and an unusual type of connector) and (b) is 4 cm deeper than the old machine and sits proud of the carpentered framework (c) is faced by a householder who points out that the colour mismatch between the dishwasher and the other appliances!

This example could be elaborated to account for the disruptions to the schedules of "Fixit" and "Kitchen-Co", the pricing structures used by "Kitchen-Co" and the finances of "Machine-Mart" to name but some. However, it is sufficient to demonstrate the general principle that the mutual adjustments of every element require *metasystemic* oversight to prevent such *oscillations* within System One. This functional type of regulation (ie, anti-oscillatory) is provided by **System Two** of the VSM, and is denoted by the triangular symbols and connections shown in Figure 4.24, below.

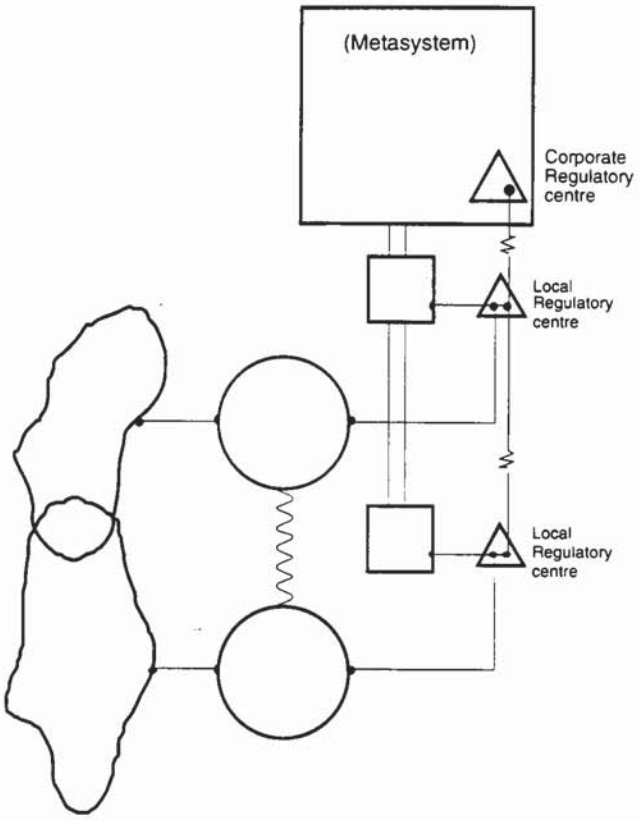

Figure 4.24 System Two of the VSM

Beer's various publications are united in their emphasis of System Two. It is seen only in terms of a **service** to System One and only serving the objective of preventing **oscillations** between the elemental subsystems of System One. Probably the most important prescription Beer offers for System Two is that its design should be in the hands of the System One managers themselves with the role of the metasystem kept to assistance rather than dictation of its design. Beer (1979) conveys much of this in the following quotation:

> "Contrast the following two statements:
>
> 'I am the boss, and therefore in the nature of things I am entitled to tell you what to do about your System Two. But I am also an enlightened boss, and therefore I shall allow you to have a say in the matter.'... and
>
> 'You need a System Two as a service, and it falls to you to design it. However, I have views on the matter, since I belong as an operational element to the next higher level of recursion. Moreover, and because of this, I can make expertise available to your work which you do not have'
>
> These are radically different approaches to the matter, and they do not at all reflect the orthodox division of approaches. For this accepts the authoritarian hierarchical model in advance of seeking participation in that predetermined authority."
>
> Beer, 1979, pages 182-183.

System Two, whilst a service to System One, nevertheless achieves this by reducing the operational variety of the elemental subsystems that comprise System One. Hence, there is likely to be confusion between variety attenuation in the cause of anti-oscillation and with interventions on the central "command" axis between senior management and local (elemental) management. In relation to the latter, these metasystemic interventions are:

1. those which arise from the recursion beyond metasystem (such as statutory requirements)
2. those which arise from the *purposes* of the metasystem (the synergistic purposes of System One)

These are in contrast with System Two activities which reflect but *do not make* the rules and policies implied by (1) and (2) above. System Two activities operate within the framework established by (1) and (2) to promote the harmonious conduct of each elemental operation **insofar as it relates to the others within System One**. Thus one would expect System Two to embody such services as:

Plant

- Maintenance Scheduling (eg. reliability-centred)
- Design guidelines (eg. standardisation of components)
- Information technology specifications

People

- Code of conduct
- Selection and appraisal procedures

Procedures

- Site wide emergency plans and alarm design
- Accounting practices
- Production scheduling (involving multiple elemental systems)
- Reporting procedures (for all corporately agreed essential variables)
- Risk assessment protocols

More generally, System Two involves both the promulgation of various activities *"downward"* to System One as well as the creation of common reporting conventions for the *"upward"* transmission of monitoring information

corresponding to these activities. As previously, the upward transmission necessitates attenuation. To permit this, System Two, must either speak all the languages of the elemental subsystems or require (by agreement) a common language, or both. The suggestion of "both" seems rather perverse until it is realised that System Two includes a variety of different classes of activity.

Lastly, Beer (1979) notes that the corporate regulatory centre (ie, the metasystemic component of System Two) needs to be carefully constructed to preserve channel capacity to the elemental subsystems:

> "The impression of variety overload normally experienced by an operational element assaulted by twelve different Systems Two is more due to the unnecessary diversity in presentation that to the inevitable diversity of content. After all: the divisional [ie, elemental subsystem] manager knows perfectly well that the twelve matters... are indeed aspects of his managerial responsibility, because he handles them all the time. What is oppressive to him is the fact that he is required to handle each matter with System Two using utterly different conventions: different terms for the same thing, different expressions of the same measure, different codes, different formats, different categories of praise, blame and excuse."
>
> Beer, 1979, page 473

4.4.2.3 System One and operational monitoring by the metasystem

In the basic regulatory schemes introduced at sections 4.2.2 (page 121, ante) and discussed further in section 4.2.7 (pages 142-150, ante) the need for an anti-oscillatory control was not explicit. This was because T was treated as unitary. However, the large system contains many independent subsystems each contributing to the total disturbance to be managed in the system as a whole. Therefore, Beer's requirement for System Two is quite justified.

Similarly, the requirement for variety amplification in stages, in this case, the senior (metasystemic) management by the elemental managements is necessary and requires *minimal intervention* by the metasystem if it is to succeed.

However, thus far, the metasystemic regulation of System One lacks requisite variety relative to: the needs imposed by its accountability to *its* metasystem; its competence in assuring coherence among the elemental subsystems of System One, and; its ability to provide services to assist the regulatory amplification achieved in the horizontal domains. Whilst the granting of maximal autonomy consistent with coherence to subsystem management is achieved in the first

instance by the resource bargain, the metasystem only has accountability information arising from the management units themselves and the routine (anti-oscillatory) System Two channels. Therefore some degree of check on the *actual* (rather than reported) state of affairs in the elemental subsystems is required. The immediate questions concern (a) how is the metasystem to dispose requisite variety to this task and (b) how to do this without massively curtailing the variety of System One required to ensure requisite variety in the horizontal domain.

Taking (b) first, Beer (1979) notes that many organisation fall foul of the variety implications and develop "cancerous activities" on the command axis of the VSM (ie, the direct channels between local and senior management). As previously quoted Beer's rhetorical question "Can you possibly itemize every single thing that the subsidiary does, demand a report on it, and expect a justification?" (page 177, ante) puts the matter plainly[149]. The approach advocated by Beer (and implicit in the Ashby formulations, already presented) is that high variety enquiries of this type should be directed to the operational domains of System One with the cognisance of the subsystem managers.

Turning to (a): the variety problems noted are to some extent addressed by the fact that whilst the enquiries involve high variety this is offset by the requirement that they be: (i) sporadic, not constant; (ii) each such enquiry be narrowly scoped to a particular class of operational activity or process, and; (iii) attenuated for input into the metasystem.

The question remains as to how to act upon such problems as may be revealed by enquiries of this type. The minimal intervention principle rules out a reflexive bloc of recommendations as commands delivered down the command channel between senior and local management. The suggestion is that, where necessary to the cohesion of the system as a whole, this can be achieved by the transmission of few and low variety rules[150] to be *amplified* by local management into regulatory requisite variety in the horizontal axis. Alternatively, the metasystem might make available expertise to the local management in the service of helping it improve its methods of regulation on the horizontal axis.

---

[149] Beer (Ibid.) notes that information technology can be abused to this end as it makes it simple for senior management to demand far more data than it can actually transduce into information but without the mountains of paper reports that might signal the futility of the exercise and the needless burden placed upon System One.

[150] Which is logically dependent upon the design of the attenuator which transduces the high variety information from the enquiry into the metasystem.

Beer (Ibid.) regards the *audit* to be an excellent example of conventional management practice that conforms to cybernetic principles. Indeed, in much of the VSM literature, the channel we are considering is labelled "Sporadic Audit". Whether this is true of all auditing is a moot point and I am not confident that "safety auditing" is sufficiently classified to hold it as an example. Within H&S, the other prime example is accident investigation: both positively as the cybernetically valid acquisition of information (if conducted with appropriate rigour - eg. Johnson, 1985) as well as negatively – the most extreme flouting of the *minimal intervention* principal (for example the long list of high variety commands contained in the Kings Cross Inquiry Report (Fennell, 1988). Whilst "Sporadic Audit" is a descriptive pseudonym, the given name of this function within the VSM is *Three star*, or, 3*, as shown in Figure 4.25.

4.4.3 Metasystem components

The need for the metasystem involvement in the work of the System One has already been introduced and some of the metasystem components introduced: System Two and System Three-Star. The purpose of these was to "look after" System One both in ways that the elemental managements could not, as well by providing the subsystems with services and expertise outside of their own resources. This is a reasonable summary of the purpose of the metasystem as a whole, and how it achieves this will be considered in the following subsections. To orient the reader, Figure 4.25 provides a simplified view of the VSM.

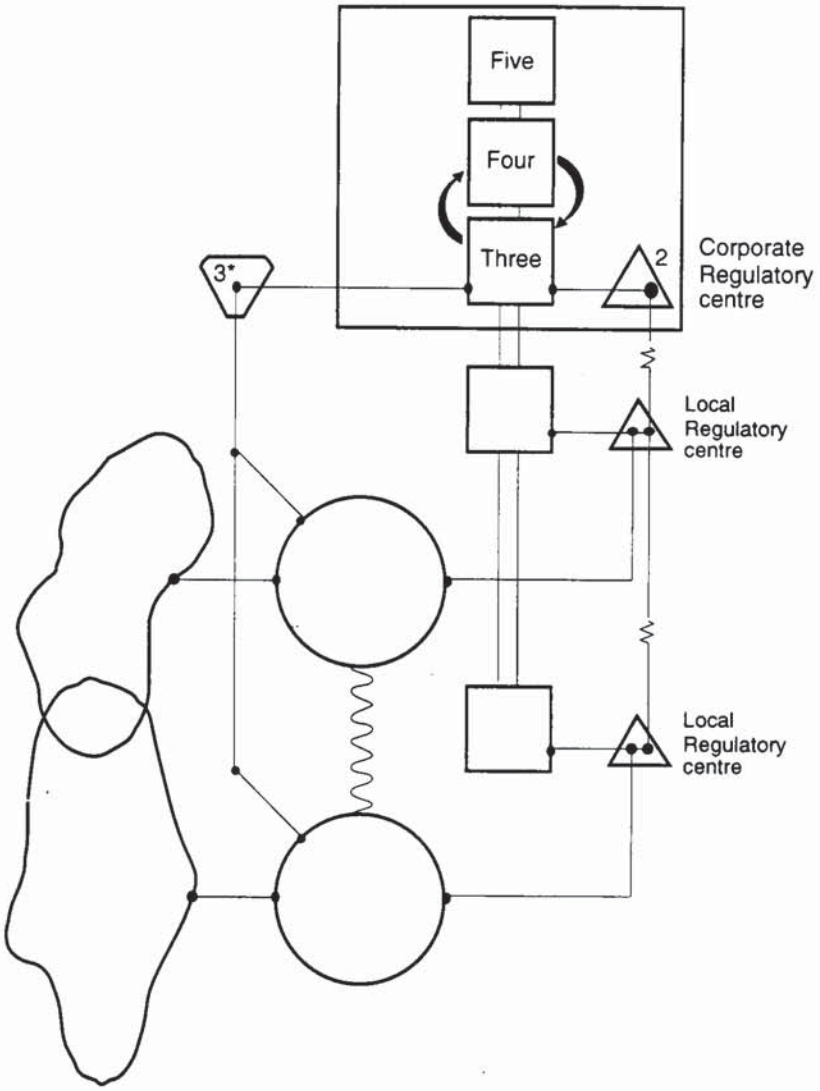

Figure 4.25 The VSM (simplified view)

It is a timely moment to reinforce the point made first in section 4.2.2 (page 122, ante): the various components of the VSM are identified *functionally*, which is to say that they may correspond to a physical bounded area (eg. the operational circles may correspond to a particular work site or part of a process) but they equally may be distributed geographically and personally. As we turn to consider the metasystem in more detail, it should be borne in mind that the functions concerned may be served by the same individual, as it were, wearing different hats.

4.4.3.1 VSM System Three

System Three is part of the "senior management" alluded to passim. Its principal function is to attend to the running of the elemental subsystems but to do so as

governed by the principle of minimal intervention. Thus far we have seen how System Three approaches this: via System Two to damp oscillations between subsystems that otherwise threaten the efficiency (or even survival) of System One; via the resource bargain struck with the management of each element which aims to select from a universe of possibilities a subset both agreeable to the purposes of the subsystem and agreeable to the purposes of the system as a *whole*; lastly, via the command channel along which such rules as must be recognised (such as statutory requirements and corporate rules) are transmitted. Beer summarises this:

> "Thus the system Three role, conceived as fundamentally synergistic, offers a powerful (to the firm) yet minimal (to System One) intervention in elemental autonomy. It is worth repeating that this intervention could reasonably be made by System One itself, acting in concert, and adopting a System Three role. This fact suggests that the loss of freedom entailed by adopting synergistic policies in the cause of corporate viability is no particular hardship, providing the reason for it is fully understood at the elemental level. Then that fact in turn suggests that the most effective managerial mode in which System Three may apply synergistic policies, is to co-opt the head of the managerial units in System One as members of System Three *to that end*. Those words are italicised , because there is no good reason to co-opt System One into System Three for any other purpose."
>
> Beer, 1979, page 207.

In addition to the direction of System One to obtain synergy, System Three also has an oversight role which has already been discussed in relation to System Three-Star. This high-variety monitoring role obviously depends upon appropriate expertise being available in System Three (engineers, accountants) or available to it (ie, commissioning consultants as required). Similarly, other high-variety excursions into the operational domain of System One via System Three-Star: Engineering services (eg, design), IT specialists, Safety Specialists, are all examples of functional resources that reside in System Three.

In summary, System Three attends to the immediate smooth functioning of System One by identifying and overseeing the implementation of synergistic plans and discharging the obligations (eg. legal and corporate) of the metasystem. Thus, at the corporate level: System Three is the guardian of Argyris' Single-loop (page 112, ante) forever dedicated to improving the "here and now" aspects of system one.

4.4.3.2 VSM System Four

If the organisation was limited only to those functions up to and including System Three, we might conclude from a "spot inspection" that it is well run. However, whilst System One is well maintained and running smoothly, its ability to *adapt* is limited to the *hunt-and-stick* mode. This is to say that minor or repetitive disturbances, not fully absorbed through the elemental operations in their interaction with the environment, can be coped with by System Three as these can be regulated by small changes in Three's regulatory model.

Looking again at
Figure 4.14 (page 161, ante), Systems Four and Three correspond to the *Adjuster organiser* and the *Feedback adjuster* respectively. Similarly, the *muddy box,* is the regulatory model of System One held in System Three. Without the *Adjuster organiser* the *Feedback adjuster,* may well enter the *hunt-and-stick* mode (experimentally trying to adapt to the disturbances) interminably. Thus the job of System Four (as identified with the *Adjuster organiser*) is to *regulate change* in System Three (and hence to System One). In organisational terms, then, how is System Four to determine what change and when? To some extent this requires information about the disturbances arising from the environment (eg. trends and the anticipation of change likely to impact upon the system). Equally, the model of System Three in System Four must have requisite variety which implies communication between them. Lastly, as was evident in
Figure 4.14, there must also be communication from the "organisational landscape" as formulated meta- to System Four (heralding discussion of the role of System Five). To set the scene for this discussion, a quotation is given below:

> "They [human beings and firms] expand variety by contemplating rather than creating alternatives. They reduce variety by the mental elimination of those alternatives. Thus, I do not say to my children: 'quick – run across the road... oh, too bad, they didn't survive', and then replace them by further breeding. Nor do we say to the firm: 'quick – here's a money-making opportunity... oh, too bad, we are bankrupt'. In both cases, we hope to acquire degrees of freedom needed to promote mutation, learning, adaptation and evolution (in a word survival-worthiness, or in another word VIABILITY) by *simulating* the amplification and attenuation of variety."
>
> Beer, 1979, page 230.

This quotation calls attention to the fact that this kind of simulation must deal with the essential variables of the firm as all of these must be maintained in the η

set of acceptability which at the extreme means viability (ie, not all negative happenings are lethal but some negative happenings are). Thus these simulations must take account of the full set of essential variables as the final decisions resulting must promote a state located within the η set.

In H&S, a popular complaint is that H&S essential variables are not integrated corporation-wide process of business development. However, of the applications of the VSM I have read (eg. those reported in Espejo and Harnden, 1989) most suggested that this complaint is not limited to H&S but is generalised to the other classes of essential variables. The root cause of this appears to be that whilst Systems One and Three conform to conventional management approaches[151], the need for System Four is poorly recognised because its function is obscured by conventional practice. What this means is that the traditional functional divides (along such lines as product design, market research, production technology, finance, human resources, etc) are not adequately bridged. Just as System Three must consider the synergistic convergence of the elemental subsystems, System Four must consider the optimisation of several classes of essential variables. Each new class further constrains the variety of choices (just as was discussed earlier in relation to design in general) and all classes must be considered in parallel if the emerging design is to be anything other than prejudicial to the class(es) that have been omitted.

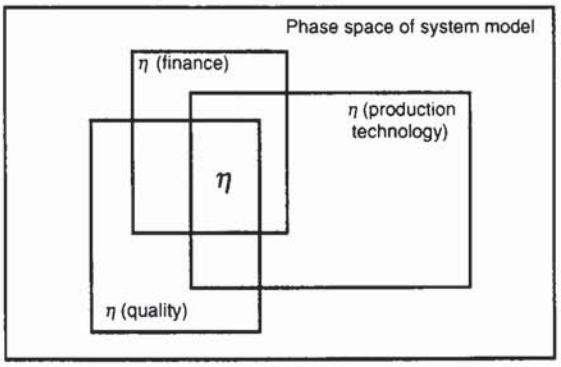

Figure 4.26 Degrees of freedom without inclusion of H&S essential variables

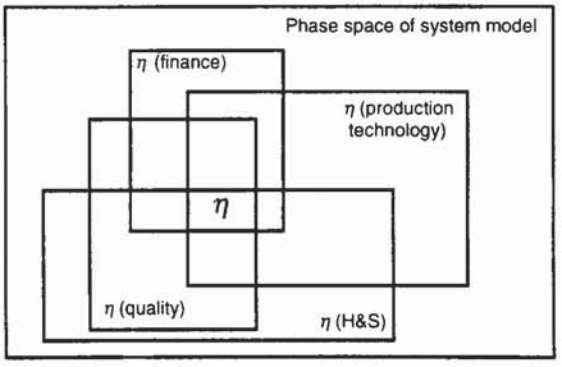

Figure 4.27 Degrees of freedom *with* inclusion of H&S essential variables

Figure 4.26 and Figure 4.27 are illustrations of the point argued. If Figure 4.26, is the determining model which becomes material through change mediated by System Three's engineering, later insistence upon the H&S performance with η, becomes a costly and difficult matter of redesign. In the Figure 4.27 situation,

[151] It should be recalled that the VSM claims to be a model of what *is* as much as a model of cybernetically what *needs* to be.

whilst the variety of options has been reduced greatly, one should expect there to remain considerable room to manoeuvre.

The modelling process in System Four is characterised by Beer (perhaps over-simply) as "Outside and Then" as contrasted with System Three as "Inside and Now". In some discussions the System Four function is referred to as "intelligence", emphasising the enquiries made by System Four in the wider environment of the whole system at this level of recursion - a matter that System Three is not designed to consider. However, I have found it difficult to obtain a clear view of System Four in isolation, it only seems to make sense when considered in dynamic interaction with System Three[152]: comparing the model of the present with models of the future to determine the direction and nature of change. An essential requirement is, therefore, that System Four contains a model of System Three in parallel with a model of itself. Beer (1979; 1985) makes great play of this, as does Harnden (1989), as it gives rise to the paradox of infinite regression and thereby opens up various debates (centred on self-awareness and consciousness). Rather than enter into this debate, suffice it to say that this is has bearing upon Argyris' double-loop (page 112, ante): the capacity for reflection upon what should be done differently to bring about an acceptable future, rather than merely to accept a fatalistic delivery of the present. Beer in all of his main texts, draws attention to the need for a physical location for System Three-Four interaction where the emphasis is availability of information (preferably graphical) and the fostering of open debate and providing, as he calls it, an ***environment of decision***. This has clear parallels[153] with the suggestions made by Johnson (1980); contrast the two quotations offered below:

> "An Operations Room, considered as the physical manifestation of our focus - in which in particular the kernel of the System Four model of itself is displayed - might taken on any form. But outstandingly it must be an ergonomically viable locale. The people who use this club-house are human beings constrained by their own neurophysiological limitations. Therefore it is absurd to overload them with data they cannot ingest, absurd to present those data in ways that actually suppress their inherent patterns, and especially absurd to maintain data that are out of date."...
>
> "Returning to the analysis of the Three-Four loops, it seems evident that the easiest way to meet the requirements of all *four* management principles is to extend the notion of the Operations Room. Let it become the Three-Four

[152] In the Ashby/Beer terminology, a homeostat, where the variety of Three and Four act reciprocally on one another.

[153] An understatement - at the danger of overstatement, these are functionally the same.

clubhouse and not merely a development-oriented place." ... "Here, System Three and System Four would exhibit themselves to each other, in a continuous mode, and absorb each other's variety"

Beer, 1979, pages 243 and 258.

"The "War Room" at Aerojet during the MORT trials was conceived as a development phase to better understand the information system which was evolving. The War Room quickly displayed two additional advantages:

1. It was a great place to explain plans and show people how objectives might be attained.
2. As data developed, it became a great place to hold managerial decision meetings.

Safety Program improvement projects (SPIPs) described thus far in the MORT trials became numerous and increasingly difficult to control and assess. To maintain visibility for the analysis and results, a War Room display of the working papers was organised on a blank wall some 20 ft in length. Although the War Room was a working room and not a polished display, it served its purposes: (1) a working focus for the MORT team, (2) a briefing room which had constructive effects on participants, and (3) management information. ...A prototype display for the whole corporation was prepared to show summary progress along the same lines in each division of the company and each branch of the safety division. ...

If an organisation plans to learn new systems to reach higher goals, the War Room approach is worthy of consideration."

Johnson, 1980, pages 448-449.

4.4.3.3 VSM System Five

In the previous section, mention was made of double-loop learning and the System Three-Four interaction appears to be involved with this insofar as it provides the basis for the debate of future actions as well as a focus of self-awareness for the corporation. However, the matter of how the essential variables are defined and the η set established and reviewed (the prerequisites of regulation as mentioned on page 125, ante) has not been settled.

Figure 4.14, showed an input from the "organisational landscape" to what we have latterly called the Three-Four homeostat. In biological terms this was the predisposition of genetic mutation, Waddington's "epigenetic landscape" (c.f. footnote 138, page 160, ante). In psychoanalytic terms, or so it seems to me, the same role can be ascribed to parental norms, which according to Freud (1915), precondition the formation of the *super-ego* - notionally the "moral sense" of the individual. This super-ego then regulates the functioning of the ego (which certainly fits the self-awareness function of the Three-Four homeostat)[154]. In this way, System Five conforms to the role of the super-ego: the metasystemic arbiter of good. In this relation and considering the role of System Four, Beer writes:

> "*Can* it review everything, however cursorily? Somehow or other, surely, it has to acquire criteria of relevance.... I think the rules come from System Five: not so much by stating them firmly, as by creating a corporate ethos - an atmosphere." ...
>
> "But the point about the *ethos* concept is that it is a **variety sponge** of gigantic capacity. Try to think of a way-out idea in your organisation — so way-out that certainly no-one has ever considered it, although it is not manifestly daft. *How would the board react to that?* The betting is that you know the answer exactly. No-one has put the idea forward because the answer is self evident. This is not to say that the answer is correct." ...
>
> In terms of the VSM, what we are discussing is the *intervention by System Five* in the balancing activity of the *Three-Four homeostat.*
>
> Beer, 1985, pages 125-127.

Within this scheme, policy-making is unconventionally represented and is evidently a much less clearly defined matter than ordinarily expected. As with

[154] Pursuing the analogy suggests that the "muddy box" must equate to the Id... a notion that again bears scrutiny, but I have already arrived at the destination intended.

much of Beer's writing, this has considerable verisimilitude; in this case, whilst there may be written policy (in the sense of corporate intentions and objectives) there is a perhaps larger "volume" of *unwritten meta-policy* (ie, the deeper set of values in which particular issues are judged). This may be both good and bad thing. Cybernetically it is "good" in the sense that the variety involved is large (Beer's variety sponge). Analytically, it is "bad" in the sense that it leaves much open to interpretation by Systems Four, Three and One as to what the essential variables are and how the η set is defined *proactively*.

Hence, if Beer's description is correct, we would expect Three-Four people to be in a guessing-game, deducing the overlap between classes of essential variable distributions that defines the η-set. The danger here is paradoxical and stems from the same reasons set out on page 124 ante: *"the more successful R is in keeping E constant, the more does R block the channel by which it is receiving its necessary information. Clearly, any success by R can at best be partial"*. In this case, the *more stable* the operation of the Three-Four homeostat, the *less exceptions* must be presented to the highest authority for resolution[155]. As illustrated below in Figure 4.28, the selections performed by Five in response to exceptions presented from Three-Four are a means by which the latter gains information (and hence builds and updates a model of Five). Thus, whilst the Three-Four arrangement is an excellent regulatory device - one might expect matters drift away from the region of the η-set where values are least emphasised by the corporate ethos (as perceived in Three-Four).

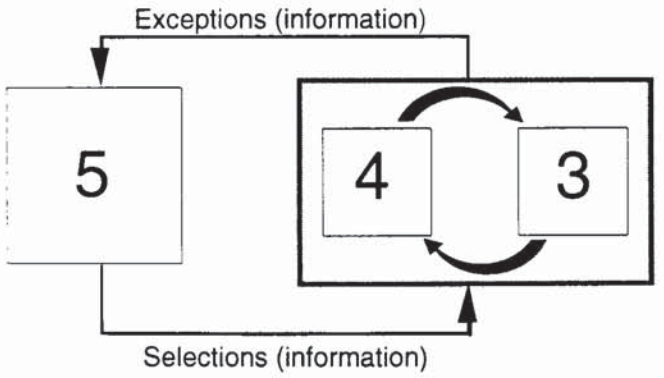

Figure 4.28 Transactions between System Five and the Three-Four Homeostat

[155] A possible explanation of the adage that "good government is boring government".

There are two approaches to controlling this effect. The first, is an additional VSM mechanism explicitly provided in recognition of the problem which Beer (1979; 1981) terms the *algedonic* channel. This has two notable features. The first is that it is unique in the property of transmitting information directly from System One to System Five without attenuation by the intervening Systems Two, Three and Four. The second property is that it is non-analytical, that is, it corresponds to a generalised *alarm* signal (shown as the dotted line branching from System One to System Five in Figure 4.29, page 194, post) and is initiated when indications of instability are registered not by Three-Four imbalance but instabilities registered in the behaviour of System One.[156] In this relation Beer offers a provocative observation:

> "These approaches effectively repudiate the **causal** view of the world on which existing managerial measurement is based.... yet our culture to this day continues to propagate the belief that 'every event has a cause'. Thus, when it comes to the management of very large systems, we still look for a **unique cause** of systemic failure - and this is not at all the appropriate methodology. Complicated systems fail because they are potentially unstable, and because some concatenation of circumstances has made the potentiality actual. No unique event is *the* cause; and when we look for one it often seems as if the total system had been in a different state, that event would not have lead to disaster. Even when the event is intrinsically disastrous of itself, it is not feasible to isolate it from is system milieu."
>
> Beer, 1979, page 290.

> "The causal model in a complex, probabilistic system does not have requisite variety to predict the future." ... "if, on the other hand, we say that the behaviour of the system, as evidenced by real time data is incipiently unstable, and that therefore we shall take action to increase the probability in favour of stability, we have changed ourselves (acquired information) which may make the future other than it would otherwise have been."
>
> Ibid., page 376.

The second check against systemic drift out of the η-set, arises from the logic of the VSM itself and from the discussion of ethics in chapter 3 (section 3.6.4, pages 87-96, ante). In chapter 3, the question of defining the η-set (and, reciprocally, *adversity*) was seen to have two main approaches which whilst distinct in their underlying philosophy were logically compatible: *deontological* (duty-based and metasystemic) and *teleological* (purpose-based and systemic).

---

[156] The best example I have come across concerns the use of Extreme Value Projection. This is described in chapter 5.

(1) **Teleologically** - the values of System Five are calibrated by **System One** *because*:

> purposes define what is good, *and*...
> purposes are defined through the perception of what the system does, *and*...
> what the system *does*, as Beer (1985) points out, "is *done* by System One".

(2) **Deontologically** - the values of System Five are calibrated by the **metasystem** at the next level of recursion (of which it is System One) *because*:

> the purposes of this System are subject to the cohesiveness of the purposes as formulated at the next level of recursion.

(3) The system is ***closed*** by the fact that, at the higher recursion, the same argument of points (1) and (2) is repeated.

In relation to the VSM description, this argument underlines the need for:

- System Four to be organisationally powerful (so as to champion adaptive change against the grain of established ideas);
- System One must be adequately represented at System Five level (Beer suggests this should be in person(s)).

4.4.4 Multiple recursion and the VSM

The foregoing has illustrated the notion of recursive levels and that the same logical structure can repeated for as many levels of recursion as is desired in analysis. This is conveyed by the Recursive System Theorem:

> "In a recursive organisational structure, any viable system contains, and is contained in, a viable system."
>
> Beer, 1979, page 118.

This provides a basic rationale to be further explored in the next chapter - that in relation to the management of health and safety, the functions which regulate the level of risk within this class of adversity are, and *must be*, regarded as distributed throughout recursive systems. The embedded viable systems within the system in focus with which we have been dealing are shown in Figure 4.29, overleaf.

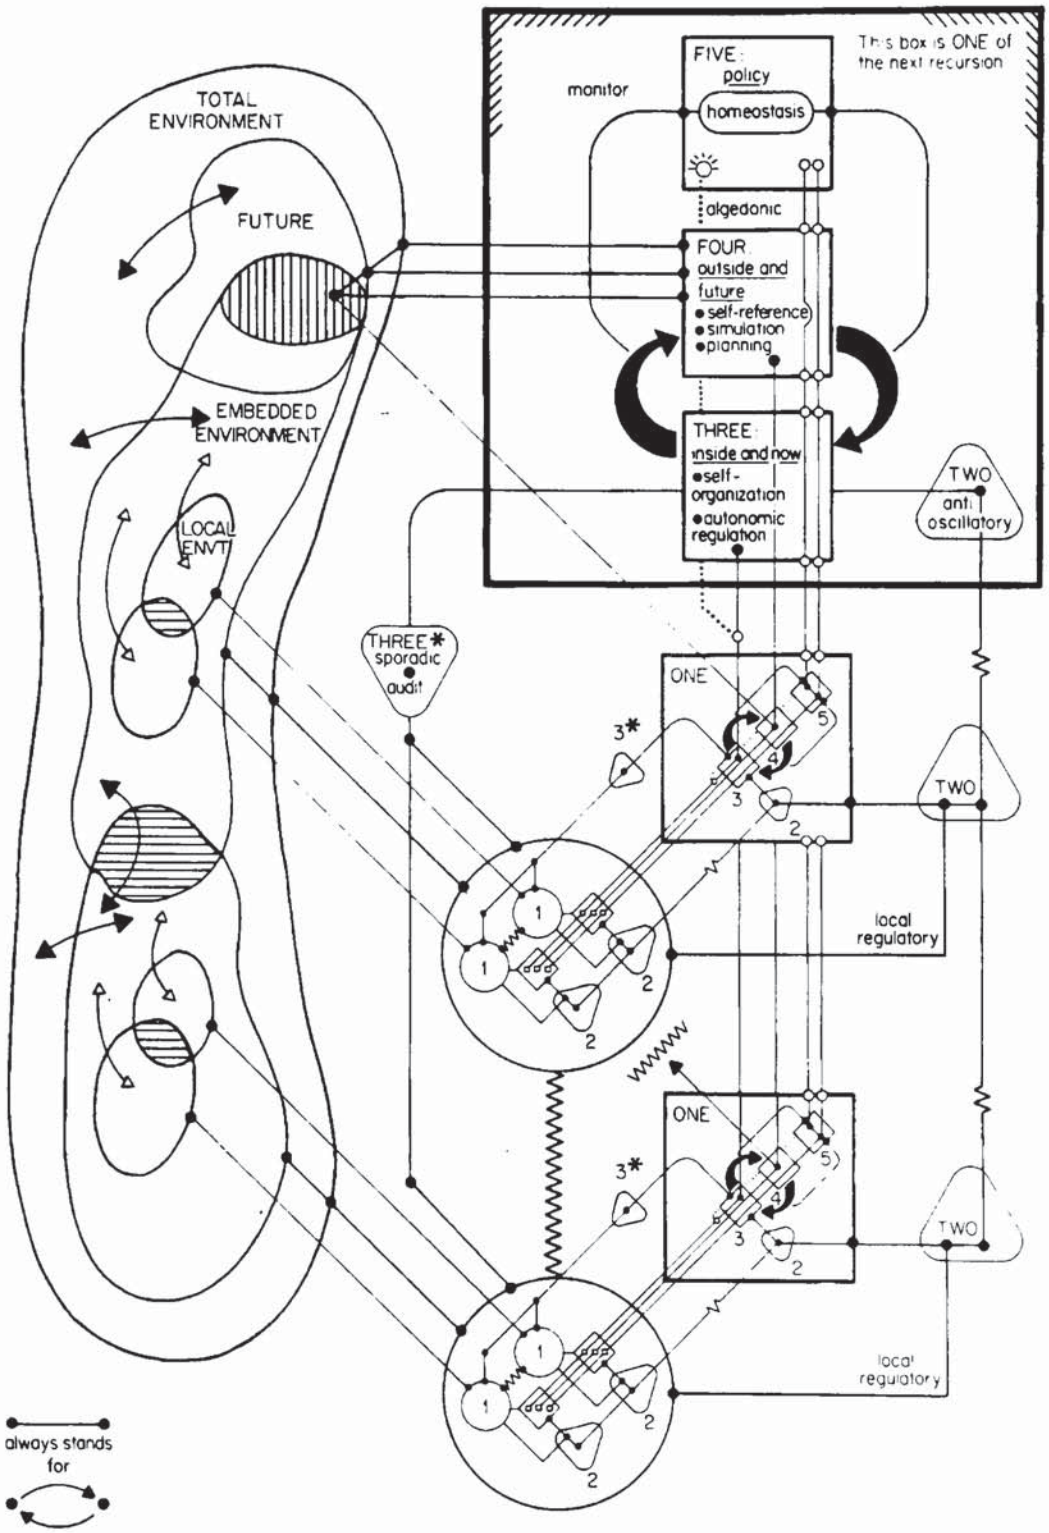

Figure 4.29 The viable system (after Beer, 1985)

5 Modelling H&S Management

"Half the wrong conclusions at which mankind arrive are reached by the abuse of metaphors, and by mistaking general resemblance or imaginary similarity for real identity."

Palmerston[157]

5.1 Introduction

As set out in the preface to this dissertation, the purpose of this work is to lay foundations for the study of particular instances of safety management systems. Should we wish to apply to one organisation what we have learned from another, some method of achieving the transfer of knowledge must be in place. Further, should we wish to say something about safety management systems in general on the basis of empirical studies then we are making a bolder statement: that there are invariances which permit such generalisation.

In seeking a set of principles from which to generalise, chapters two, three and four, framed the issue of H&S management in *systems* terms - gaining access thereby to the unifying concepts of system theory and cybernetics. H&S, recast, is considered to be amenable to the logic and principles of regulation which pertain to systems in general. This, as argued from various perspectives, reduces to a matter of information: both in respect of the channelling of information in the organisational system and in the models used by the regulators of the system.

The contention of chapter 4 is that, whilst the control of a particular system must obviously reflect the empirical nature of the system itself[158], the nature of the regulatory task relies upon common principles. These principles when rendered in such a way as to be universal are, however, of limited practical usefulness. In other words, the basic schemes set out in sections 4.2.2 and 4.2.3 (pages 121-129,

---

[157] Palmerston, H.J.T., Private correspondence with Sir George Villiers (1835), cited by Weldon (1962)

[158] eg. the control arrangements to effect safe performance in a dairy will bear scant resemblance to the control arrangements required in a nuclear power plant: the energies, toxins and biological pathogens are radically different and the practical control of these will reflect this.

ante) provide only general formulations from which more sophisticated models of regulation may be constructed.

I believe that it is misleading to view the development from universal principles to actual instances of large organisational systems as implying a continuum of complexity. This is not to say that such a view is disallowed, merely that it is unhelpful because it suggests that there is a *given* model of regulation that can act as a template for organisational design. As suggested in the quotation given on page 200 (post), there is no such holy grail.

What is more plausible, is the empirical pursuit of descriptions of H&S regulation relating to particular systems based on, and reducible to, a limited (and hence usable) set of modelling conventions. Further that these conventions are based upon a yet smaller set of rigorous principles. The rationale for this is twofold.

- *First*, as set out in the introductions of chapters 2 and 3 - a set of conventions allied to deeper principles permits exchange of information. Further, that these principles and conventions can be challenged in the light of empirical experience. This is operating within the requirements for consensibility as discussed at 3.1 (page 60, ante).

- *Second*, a reasonable expectation is that the process of empirical description will discover some aspects of regulation to show invariance across organisations. However, it is also to be expected that what general models *gain* in their *level of description* they *lose* in *range of applicability.* Rasmussen (1994) suggests that there may be classes of organisational system: *within each class* it may be possible to develop models at a lower level of description (ie, more detailed) than is permissible across classes (c.f. section 5.3, page 210, post)

General models, even with the caveats stated, are by no means easily obtained. What is advocated here is the development of methods of generative modelling in the first instance. Once these are developed, the empirical process of obtaining *general* models of H&S regulation can be taken forward.

5.1.1 General models of H&S management and the need for integration

Although mentioned before (section 3.1, page 60, ante), it perhaps needs repeating that such models as do exist are not suitable for the purposes of studying the design and evolution of safety management within organisations. Models, such as those implied in HS(G)65 and BS8800 are aimed at communicating *good* practices dedicated to the aim of securing H&S performance. They assume that such practices can naturally be integrated into the organisational structure and also assume that the exhortations to develop a "positive safety culture" will allow these to be amplified through the existing regulatory system. However, it is by no means established that a regulatory system that delivers system performance relating to the core activities of the firm (eg. serving the purposes defined by what the system does) necessarily is competent to deliver the control of hazardous states. As indicated in chapter 1 (pages 28-29, ante) it may be that the category of organisational performance called H&S outcomes, present a greater challenge to the regulatory system. As noted in chapter 4 (particularly pages 130-132, ante), the regulatory system may dispose sufficient variety to regulate against the effects of repetitive disturbances (and to some extent have engineered the system to show this constraint). Even so, it may be that the regulator does not possess sufficient information handling capability to deal with **less frequent** disturbances (and, correspondingly, less probable states of the system).

For example, messages transmitted by the State to industry such as "The costs of accidents at work" (HSE, 1993) *assume* that the set of system states associated with acceptable H&S performance is naturally convergent with the set of system states conducive to profitability. Various case studies are presented showing that accidents are a major source of uninsured loss[159]. Thus, improving H&S control improves profitability. However, what is far from supported is that organisations are able to obtain the savings that a simple argument of accident prevention suggests. The traditional view is that better motivation and application of H&S "best practices" will naturally lead to improved performance and the savings that HSE indicate. However, a tenable counter-argument may be put. It may be that the financial savings suggested by the HSE study, involve an increase in the information handling arrangements of the firm that is beyond its capacity. In other words, the cash loss may be an overhead. A colleague in the insurance industry suggested a similar formulation (used in the persuasion of otherwise

[159] Similarly, Veltri (1990) presents the argument that conventional methods of accounting under-represent the size of losses incurred through accidents.

unresponsive senior managers): what organisations insure is, in fact, "*the margin of their managerial incompetence*".[160]

Given the arguments put in chapter 4, it may be that the information handling abilities implied by given organisational structures simply cannot be ramped-up by increasing the capacity of the given regulators distributed in the system. As discussed in section 4.2.4 (page 129, ante) a given regulatory structure (ie, interconnection of regulating elements) may alter its performance by:

1. Allowing a wider range of acceptable system states (ie, tolerate lower levels performance - perhaps by trading-off one class of essential variable at the expense of another);
2. Reducing the variety of the Disturbance (eg. engineer to reduce variation at subsystem level or, reduce the variety in the environment, else move to one which is less hostile);
3. Exploiting constraints in the variety of Disturbance (ie, identify the most frequent disturbances and regulate against those - tolerate the less probable and higher variety residual);
4. Increasing the variety of the regulators (eg. train to distinguish more states in D or T, increase capacity of channels and transducers, increase control variety).

However, these options each and collectively have limitations. To go beyond these limitations, requires amplification (as introduced in section 4.2.5, page 133) and this requires supplementation in stages. This presupposes change in regulatory structure (which *may* not[161] necessarily entail change in staff but does entail change in how they are deployed[162]).

It seems to me that there is an *urgent* need to study the relations between H&S management and the regulation achieved in organisations generally. The need has

---

[160] The colleague in question recently emigrated to New Zealand.

[161] For example: A supervisor is responsible for the regulation of Section A. If her staff (x, y and z) have no discretion and must bring all matters to her to receive instructions, the amount of regulation achieved in Section A is limited to the capacity of the supervisor. If however, x, y and z, are granted some degree of discretion, the amount of regulation achieved by the same four people will be increased)

[162] The introduction of IT into organisations has been noted on many occasions (eg. Bjørn-Andersen et al, 1986, and Eason 1988) to result in sometimes dramatic change in regulatory structure without apparent alteration in staff structure as shown in the organisational chart).

always been there but the *urgency* stems from the increasing financial instability of the market-place. Organisations, striving to maintain competitiveness - financial viability - are required to become adapted to the market in shorter time frames than hitherto. Further, the same forces have resulted in high levels of redundancy and acceptance of the notion that slack resources are wasted resources. A view that this research has strengthened is that the regulatory resource presented by the workforce at large provides the amplification for the limited resources of management.

The regulatory resource of the workforce concerns the capacity to select and maintain the operating parameters of the system (considered at the operational level). It is not that this resource is unused – far from it, as mentioned in chapters One (page 29) and Four (page 174) this is a vital commodity in any organisation. What is in issue is the extent to which the resource is under-utilised and inadequately recognised by management. Whilst the "safety culture" concept stresses the involvement of all, it does little to challenge the Tayloristic assumption that regulation is something that managers do *to* the workforce, as contrasted with something achieved *by* the workforce and *steered* by their management.

Stata[163], in discussion with Senge (1992), puts a similar view:

> "The "scientific management" revolution of Frederick Taylor took the traditional division of labor, between workers and managers, and gave us the "thinkers" and the "doers." The doers were basically prohibited from thinking. I believe our fundamental challenge is tapping the intellectual capacity of people at all levels, both as individuals and as groups."
>
> Senge, 1992, page 350

What the cybernetic viewpoint emphasises is that the regulatory resources of an organisation, whilst self-organising to some extent, are amenable to optimisation. Perhaps more seriously, because the Tayloristic viewpoint discounts the importance of the regulatory capacity of the workforce, the negative affect of staffing reductions and restructurings upon risk may also be discounted.

For these reasons, improved understanding of organisational regulation and the relations of this with the processes, practices and motivations of H&S are required. Whilst there is good reason to give H&S special consideration (as highlighted in

[163]Mr Ray Stata, President and CEO, Analog Devices, Inc.

the ethical considerations of chapter 3), improvements in understanding about organisational regulation generally is a pre-eminent concern. As Jaques puts it:

> "What we need is not some new kind of organization. What we need is managerial hierarchy that understands its own nature and purpose. Hierarchy is the best structure for getting work done in big organizations. Trying to raise efficiency and morale without first setting this structure to rights is like trying to lay bricks without mortar. No amount of exhortation, attitudinal engineering, incentive planning, or even leadership will have any permanent effect unless we understand what hierarchy is and why and how it works. We need to stop casting about fruitlessly for organisational holy grails and settle down to the hard work of putting our managerial hierarchies in order."
>
> Jaques, 1990, page 133.

5.2 Generative models

A well known marketing aphorism is that "the medium *is* the message"[164] whereas our concern is diametrically opposite: the use of models (a) as part of a methodology for obtaining a description of a given system and (b) as a way of communicating what was found in this system.

Seen in relation to (a), both MORT and the VSM are generative schemes for developing models: MORT reflecting the particular considerations of H&S management (and notably, as was Johnson's intention, including requirements bearing on management generally); the VSM reflecting information handling within organisations as grouped within a five-fold functional framework. In relation to (b) the VSM is often used to communicate the findings of a study (although it will not necessarily look like the "classical" diagram of Figure 4.29 (page 194, ante) and will often be used as the reference diagram for more detailed descriptions rendered in different ways). However, as discussed further in section 5.2.1 below, the MORT diagram is not ordinarily used as the means of communicating the results of a study[165].

The ideas I am attempting to communicate here are elaborated by Espejo and Harden (1989):

---

[164] In the systems literature the antithetical notion is conveyed by "the map is not the territory".

[165] The notable exception being the use of MORT as a supportive medium in the communication between experienced MORT users (as discussed on page 67) . This, in my view, is one of the reasons that people who use MORT become somewhat emotionally attached to it.

> "One way of looking at any model, including the VSM, is as a means of 'gathering' descriptions that might themselves concern non-intersecting phenomenal domains, under an umbrella of intersection. Quite simply, models enable diverse people with quite different mindscapes to have conversation about diverse matters. A model does this by providing a context or mood which directs discourses in particular paths, upon acceptance of such a model as a common *convention*.
>
> A model is expected to provide a setting, a common frame – in other words, it is expected to *make visible a set of constraints*, within which certain problems can be enunciated in a particular way, and certain problems solved.
>
> Let us be clear about this. A model is a *convention – a way of talking* about something in a manner that is understandable and useful in a community of observers. It is not a description of reality, but a tool in terms of which a group of observers in a society handle the reality they find themselves interacting with. ... But whatsoever, an individual may never *communicate* what is accessed to another individual except in terms of models. This is not a limitation, but is precisely the motor *for the generation of a consensual domain*. A consensual domain is none other than the play of a particular set of interacting models."
>
> Espejo and Harden, 1989, pages 445-446

Neither MORT nor the VSM are answers in themselves and require reference to the particular case (or the particular class) of system examined. However, both come with a considerable body of insights gathered by the experience of their progenitors. In my view, MORT and the VSM address the same issues but in complementary ways. The question is, can we combine these to produce a scheme to generate models that describe the integration of H&S and general regulation achieved in the firm?

5.2.1 MORT and generative modelling

The foregoing chapters have been concerned with the foundations upon which generalisations may be constructed. In regard to the Management Oversight & Risk Tree, this achieves its general applicability by identifying the elements (further considered below) required in the regulation of specific hazard states: essentially those concerned with the manipulation of probabilities associated with such states to a occur within pre-determined level. However, the justification of MORT analysis is that, whilst it formalises the practice of identifying the causes associated with a given accident, the causes are held to be more general in their effects than their association with the event in question. For example, if

inadequate training is identified as an issue in relation to the event, the hypothesis must be entertained that training may be *generally* inadequate. Similarly, if the evaluation of the hazard state involved in an accident showed that it was more probable than the the original risk assessment had anticipated then both the reliability and validity of the *whole* risk assessment process maybe in question.

Each barrier/control failure examined using MORT identifies a subset of elements found to be less-than-adequate (LTA). When an accident sequence is thus examined, the results can be associated together in an events and causal factors chart, showing causal relations between MORT elements and events, events and consequent events, and common modes (eg. both events and MORT elements). Such a result provides an account of the accident as a model of the processes and interrelations of processes as revealed by the sequence of changes unfolding in time[166].

What these observations point to is that MORT provides a generalisable list of "programme elements" (and the properties required, in general, of such elements) but leaves the temporal and structural (ie, arrangement of processes within an organisational structure) to be addressed on a case-by-case basis. To attempt to model the temporal and structural aspects relating to the MORT diagram elements themselves (and with the same rigour) would involve immense variety. However, this is not a serious loss when the products of MORT analysis are applied and restricted to a given organisation (so long as its structural pattern is stable with respect to time) because the organisation itself provides the information concerning structure[167] (ie, it is its own model). Similarly, analysts may discover that certain MORT elements are problematic across a number of organisations (a common example is the set of processes concerned with developing and maintaining effective procedures). This may be an indication that certain processes are inherently difficult and/or operate upon generalised weaknesses across different organisational settings.

---

[166] This model is informative to an organisation because it challenges the assumption that the various processes revealed were functioning satisfactorily at the material time.

[167] This brings in to view an alternative viewpoint provided through cybernetics: that regulation has common principles irrespective of the system in question. In other words, the regulation achieved in all systems can be reduced to a set of universal principles. In this way, Beer produces his account of organisations which claims that, irrespective of the particular structure it possesses (eg. as described in various ways by physical arrangement, organisational chart etc.), patterned within it is a recursive five-fold structure of regulatory *functions* and their interrelations. How particular organisations serve these functions (eg. the processes/methods used) will vary but the five functions are necessary and sufficient conditions of regulation in any viable system.

5.2.1.1 Characterisation of MORT elements

One of the great advantages of MORT, evident in its utility as a *general* tool in generating *specific* accident models, is that the nominated MORT elements are generically named: we have in MORT a meta-syntax[168] for the various elements involved in H&S management.

What then can be generally predicated of the *elements* as classified in MORT? Again, let us approach this from the retrospective viewpoint of accident investigation. In the MORT User's manual[169] (Knox and Eicher, 1992) the various elements shown in the MORT diagram are referred to in the question set, as *events*. To some extent this is to maintain parity with Fault Tree conventions, however it also points to the notion that (from the temporal perspective of the accident) the action of an element is, logically, an event - it acted in past time. If the investigation finds (in present time) this past act to have been less-than-adequate then the hypothesis is that the *present* actions (and contemporaneous actions) of this element may also be suspect.

For the moment let us suppose that MORT elements are *processes*[170]. A process is taken, simply, as something which *changes* whatever it is that is subjected to the process. In the language of ECFA these episodes of process action are called *conditions*. In other words the action of a process *Y* upon *X* in past time is to predispose *X* in present time (ie, present with respect to the accident sequence) to behave in particular ways. Thus, the behaviour of *X* (whatever this actually refers to) is *conditioned* by process *Y*. For example, *maintenance*, is a process - it acts to condition the behaviour of a machine. Similarly, *training* is a process - it acts to

---

[168] Whilst valuable, this may not be without cost. As noted in a letter to Dr Nertney "...at the level of problem identification, the level of implementation and at the level of post-implementation review, mismatches may exist between the abstract representation and the system. In other words the translation of a given system into MORT for Root-Cause analysis is a vital analytical stage but information is lost in the process. Equally, translation of the abstracted system back into the actual system for implementation needs an amplification stage to recapture the detail lost originally. Mismatches between systems and maps of those systems could quite easily go unchallenged. A special issue here is the constitution of the Investigation board: does it contain members of sufficient organisational overview to assist the translation process (they speak, as it were, the metalanguage of MORT and can translate this without loss to and from the language of their own organisation).." (Kingston-Howlett, 1994)

[169] See Appendix 1

[170] Inspection of MORT suggests that in the majority of cases, the events are processes or imply processes. In a smaller number of cases the events are attributes of processes.

change the behaviour of people. What of *inspection*, is this a process and, if so, what does it change? It seems to me that an inspection process may *result* in change to, say, the machine inspected but not immediately. What it may do is communicate to the maintenance process, which is to say, *change* the maintenance process either temporarily (eg. this machine, or its components, needs to be changed) or enduringly (by altering the schedule of maintenance to include periodic adjustments to this machine). Thus, inspection is a process - it acts to change the maintenance process[171]. On this basis, there is evidence to suppose that we have two levels of process:

- those which change operational elements in the system (such as people through training, and machines through maintenance); and,

- those which change processes (which we might call higher order or *meta*-processes).

[171] In MORT terms, inspection at SD4 will communicate with MB3-b10 to achieve a change in the maintenance schedule.

The formulation suggested above (process *Y* changes the behaviour of subject *X)* needs adjustment. The accident investigation view ordinarily[172] has all events modelled at a probability of 1 - these events took place. However, the causal factors which conditioned the occurrence of these events are in general not deterministic but probabilistic in their operation. In other words processes such as maintenance and training condition the probabilities of the behaviour of their subjects with regard to the transformations effected by those subjects. Hence, training changes the probabilities of behaviours expressed in response to various inputs, machine maintenance changes the probabilities associated with the various operations performed by the machine upon its inputs. This is to say that the first-order processes act by altering the probabilities associated with the transformations achieved by people and machines of the inputs provided to them into the outputs required of them within the work process. The second-order processes influence these probabilities at one or more remove and so forth. In general, then, processes are operations upon the probabilities of system states and are characteristically regulatory and Markovian.

Also considered in MORT is the *mediation* between the process and its subject. As discussed in chapter 4, *to change* equates with *to inform* and communication theory applies to this. Johnson's (1973; 1980) concern with such issues is strongly apparent and, as noted passim, he made considerable use of these ideas in the development of the MORT programme. Perforce, a fault-tree representation permits only a weak representation of the communication aspect, but nevertheless, the representation is present. Thus in addition to MORT events as processes, there are MORT elements explicitly concerned with communication (eg. the SD1 branch). However, greatly more consideration of communication is implied than is stated (for example the various MORT events concerning supervisory observation presuppose the passage of information).

Can we further characterise these processes? In general they are doing one of two things - changing regulatory models (eg. altering the models and/or altering the probabilities associated with transition states in the models) or changing the system itself (altering the system and/or probabilities associated with transition states in the system). In general, the effect of upward flow of information is to change probabilities in regulatory models (generally, this will be attenuated in stages). Similarly, the effect of downward flow is to change probabilities in the

[172] Exceptionally, evidence will be equivocal and in consequence events may be stated with a probability <1.

system. However, such change will be often be mediated through changes in the intermediate regulatory levels (ie, subject to their amplification). Therefore downward flow may be characterised as making selections in regulators and thence to the physical system itself[173] (ie, design - probably in stages).

Lastly, in addition to processes and mediation of processes, the remaining MORT events suggest *attributes* of processes. For example, MORT events SD5-e26 & e27 (see Appendix 1) suggest that the supervisory and task design processes should positively reinforce "correct performance" on the part of operative staff.

5.2.2 Integrating MORT and the VSM for generative modelling

Whereas MORT provides a set of processes logically necessary to the regulation of H&S performance, the VSM provides a recursive structure and functional classification which may be used to locate and interrelate these processes. The expectation being that the MORT processes should map onto the five functions of the VSM. As indicated earlier, MORT analysis relies upon the organisation itself as a means of providing structure for the interaction of various processes. Therefore, the advantage of mapping MORT into the VSM is to provide a means of modelling the organisational structure in which these processes cohere.

What is required then is to reconcile MORT (a many - one model of safety management processes) with the VSM (a many - one model of organisations) to obtain a generative model which shows the relations between the processes and, further, provides a basis for examining the information handling within and between these.

5.2.2.1 Expectations of a MORT-VSM mapping

As stated, the purpose of a mapping study is to model how an organisational system structures the various processes which *explicitly effect* safety performance alongside those which *implicitly affect* safety performance. A cursory mapping of limited aspects of MORT (eg. the MA1-POLICY and MA2 - IMPLEMENTATION branches) was attempted as a desktop exercise. The results, enumerated below, were instructive.

---

[173] Klir (1991) refers to this as the *data-level system*. The distinction is useful but has not been adopted here to avoid confusion.

1. Although these questions are concerned with policy implementation, very limited traffic was required on the "corporate intervention" channel, relying instead upon systems 3* and 2 and the resource bargain.

2. Phrase-by-phrase, the requirements described for each process in the MORT User's manual tend to involve multiple levels of recursion. Generally this is *implicit* rather than explicit. In many cases, it is exceedingly difficult to maintain a focus on a given level of system (the notion of a system-in-focus was exemplified on page 165, ante);

3. The processes as described in the manual often require more than one metasystem element (ie, System 2, 3*, 3, 4 or 5) and more than one channel of communication. This reflects the different functions served within each process. For example MORT event MA2-a1. The MORT User's manual poses the following questions in relation to implementation of policy:

> "Is there a comprehensive set of criteria used for assessing the short and long-term impact of the methods on safety for the desired results? Does management demand that adequate analyses be performed and alternative countermeasures examined, or are criteria simplistic and therefore LTA?"
>
> (Knox & Eicher, 1992, page 42).

These rather innocent sounding questions are actually very sophisticated. Firstly, *policy implementation* refers to the planning and enaction of any new initiative developed at managerial level. Thus, we need to consider the VSM 3-4 interaction in terms of the changes initiated by 4 and the analyses it performs with respect to the information it obtains from 3 (and from outside information sources). Similarly, 3 may well have to initiate inquiries via 3* into the operational circles themselves to obtain data to assist the debate with 4 in the choice of how to proceed with operational change. Further, the question calls for examination of the methods used (ie, for risk evaluation, CBA and cost-effectiveness analysis). Meanwhile, one hopes that 5 has given or can provide sufficient guidance to 4 & 3 as to what purposes and values should inform these analysis (eg. with regard to acceptable and tolerable risks). These requirements are illustrated below in Figure 5.1. Nor should it be overlooked, that depending on the scope of the decision, more than one level of recursion may be involved and the interaction between these examined.

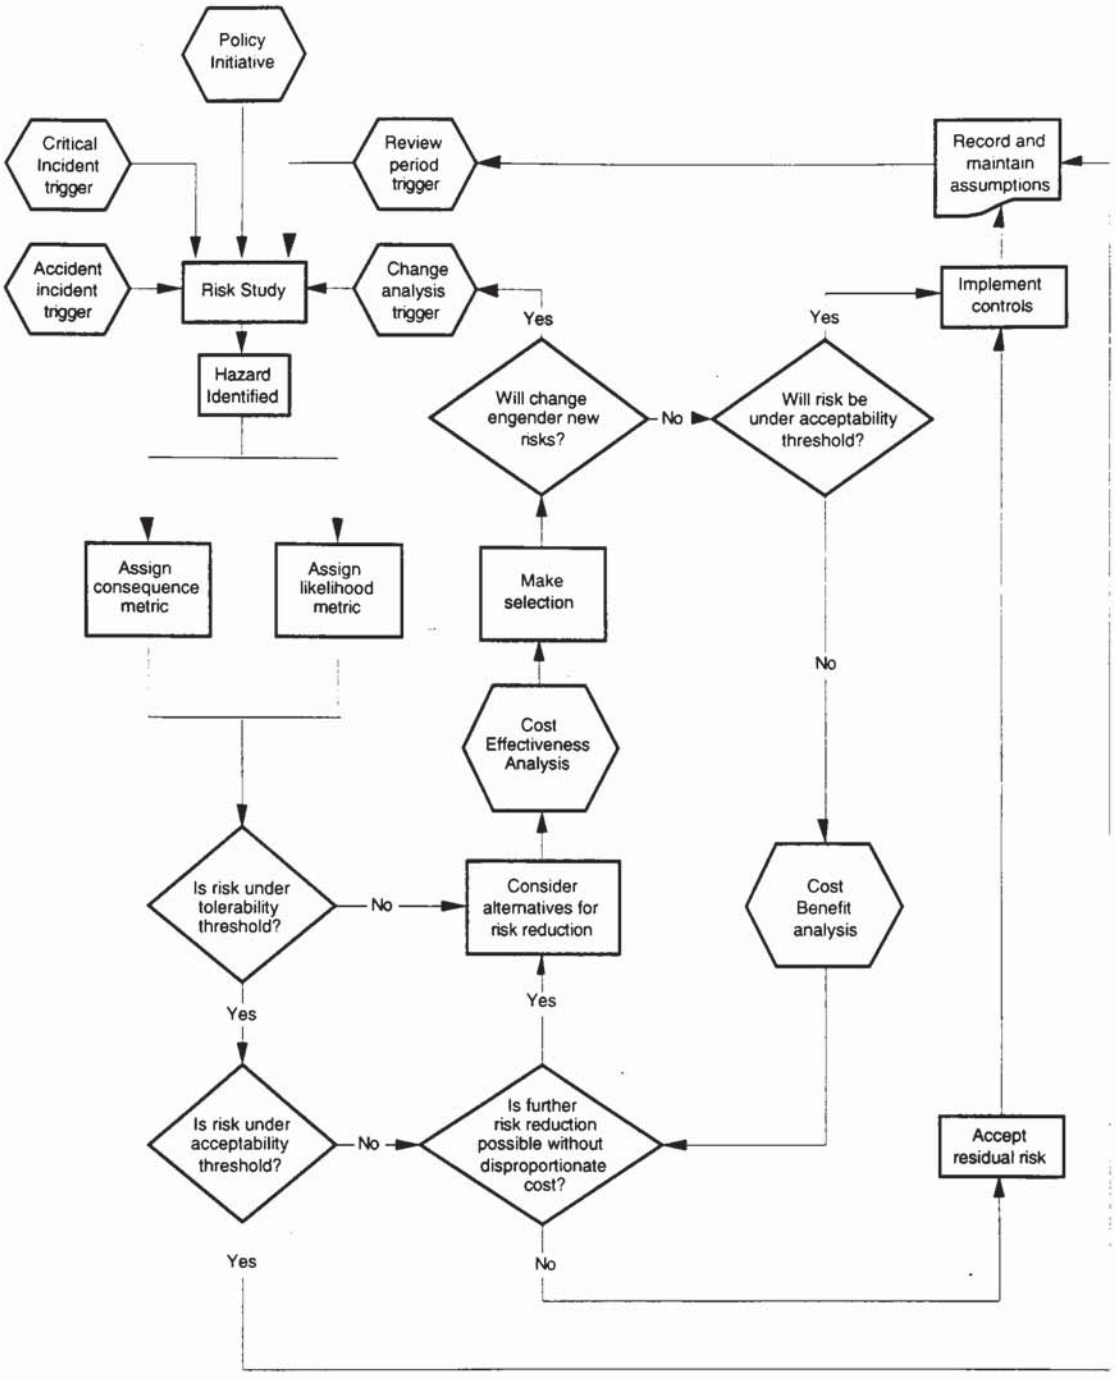

Figure 5.1 Elaboration of issues within MORT event MA2-a1

The fact that MORT events are clearly complex is a matter that this tentative exploration of VSM mapping made clear, as indicated by the involvement multiple VSM elements and, more problematically, multiple recursions. It is fair to say that event MA2 is at a high level in the MORT tree and that was the reason for its choice. MA2, implies many processes that are addressed specifically elsewhere in the MORT diagram. However, the experiment indicates that the mapping task needs to be performed *using a case-study approach.* In other words it is possible but

exceedingly difficult to approach this task without an external frame of reference in the first instance.

In line with the distinction of models as generative schemes (methodologies for modelling), the first objective of such a case study would be the development of protocols for modelling work of this kind. In other words, the first stage is to *discover and resolve the problems entailed in conducting the analysis*. The second objective addresses the requirement for *models as communication media for encoding the results of analysis*[174]. But for the scale of the task, this might better be referred to a pilot study.

In brief (described in detail in chapter 6), a case study would consist of identifying the basic regulatory arrangement of the organisation using the VSM. Further, I would advocate that this analysis be performed in the first instance with no special emphasis on *explicit* safety processes and requirements (indeed *de-emphasis* so long as this does not lead to excessive practical difficulties). After the initial ("non-safety") VSM description is completed, the analysis should be repeated, this time identifying the processes and requirements of MORT and mapping these onto the organisational description. There are :

a) to provide empirically an expansion and clarification of the recursive tangles evident in MORT;

b) to identify MORT processes which require modelling of recursive levels outside of the system-in-focus;

c) to highlight discontinuities between the processing of information which is *explicitly* concerned with H&S and the processing of information relating to other major essential variables (eg. productivity, quality, human resource development);

d) to highlight areas in the VSM which do not allow mapping (ie, incommensurabilities between organisations as conceived in MORT and as conceived in the VSM).

[174] By analogy, just as an events and causal factors chart encodes the results of MORT (and other) analyses into a model of the occurrence investigated, this part of the study would aim to producing a general format for encoding the results of MORT-VSM analysis.

5.3 Scope of modelling required of the H&S regulatory process

This question concerns the different requirements of different players in and inquirers of the H&S regulatory system. Here, the "system" is as defined by Robens:

> "By 'system' we mean here the whole complex of arrangements and activities, whether of a statutory or voluntary nature, which seek to protect and promote the safety and health of people at work, and to protect the public from hazards of industrial origin. The system can be seen as comprising two very broad elements: regulation and supervision by the state, and industrial self-regulation and self-help".
>
> Robens, 1972, paragraph 15.

As indicated in section 5.2 (page 200, ante) and the foregoing chapters, the philosophy of the approach taken in this thesis is to assist the process of debate and sharing of knowledge amongst all inquirers of the regulatory problem. My adoption of the systems approach and of cybernetics is because these seek to include rather than exclude any given discipline or level of inquiry. Given the aim of consensual debate writ large, the development of broadly-scoped analyses which attempt to describe the recursive linkages from the level of the State to the last decision element at the "coal-face" operating level, seems necessary.

Rasmussen (1996) reaches a similar conclusion which he illustrates in Figure 5.2. The following commentary is provided:

> "Many nested levels of decision making are involved in risk management and regulatory rule making to control hazardous processes. This social organization is subject to severe environmental pressure in a dynamic, competitive society. Low risk operation depends on proper coordination of decision making at all levels. However, each of the levels are often studied separately within different academic disciplines".
>
> Rasmussen, 1996, page 2.

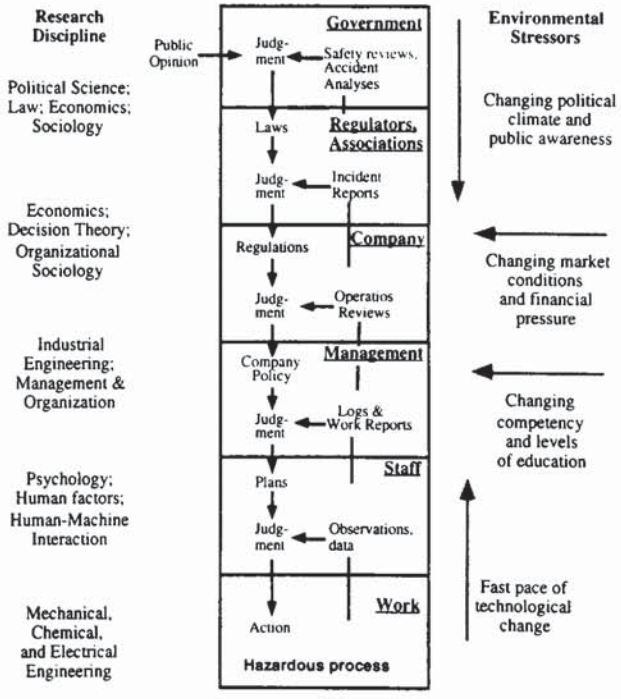

Figure 5.2 The regulatory hierarchy (from Rasmussen, 1996)

5.3.1 Classes of system (Rasmussen)

Rasmussen (1994) suggests that the modelling performed at the level of the operating systems themselves will need to reflect the different natures of the technologies involved:

> "Modelling risk management in a dynamic society in which all actors continuously strive to adapt to changes, we have to apply a closed loop, feedback control point of view based on a functional abstraction to derive an information flow representation of the control function. The information flow network involved in control must be mapped for the entire system, including the technical core at the bottom. The control function and its stability **cannot be studied at a general level across systems**, but must be analyzed for a particular system, or **type** of systems. Therefore, modelling risk management must be based on a categorization of hazard sources according to their control requirements."
>
> Rasmussen, 1996, page 14 (*emphasis added*)

The categorisation suggested by Rasmussen (1994) is based upon the dominant mode of control relative to the adverse states of greatest consequence in the systems in question. For example, process industries typically involve marshalling concentrated sources of energy or chemicals with associated high consequence potentials. Hence, the control asserted upon such systems needs a much greater level of reliability than is required in systems utilising lesser amounts of energy or chemicals. He suggests a three-way classification (illustrated in Figure 5.3) by: rate of technological change; amount of energy, and; by concentration of energy.

<table>
<tr><td colspan="2" rowspan="3"></td><td colspan="4">Hazard (that which causes harm)</td></tr>
<tr><td colspan="2">Concentrated</td><td colspan="2">Distributed</td></tr>
<tr><td>Low</td><td>High</td><td>Low</td><td>High</td></tr>
<tr><td rowspan="2">Pace of technological change</td><td>Slow</td><td>Empty</td><td>Concent-rated sources of high hazard; slow place of change
(eg. railways, aircraft, shipping)</td><td rowspan="2">Distributed sources of low hazard
(eg. factories, construction sites)</td><td></td></tr>
<tr><td>Fast</td><td>Empty</td><td>High hazard systems in Fast pace of change
(eg. nuclear & chemical process plants)</td><td></td></tr>
</table>

Figure 5.3 Industrial system classification

Rasmussen's scheme arises from the consideration of the type of control paradigm required by the nature of "the accident creating process" (Ibid.). He distinguishes two main types of control, *empirical control by objective* and *analytical safety control by design and plans,* which may be the dominant form in a particular system.

Rasmussen's description of ***empirical control by objective,*** has a metasystemic focus (ie, at a level above the operating system whose self-organising properties are relied upon by the meta-regulator):

> "One possible strategy is an empirical control aimed at maintaining an acceptable level of safety by reacting properly to a continuous measurement of safety as, for instance, defined by the observed rate of fatal and lost-time-injuries. This approach is frequently discussed in terms of an acceptable 'safety culture', that is, general efforts serving to minimize the frequency of initiating events and/or serving to control the propagation of accidental courses of events. This strategy basically depends on reactive feedback control..."
>
> Rasmussen, 1994, page 3.

Analytical safety control by design and plans is described in the following terms:

> "Another possible strategy is the analytical, open loop, feedforward strategy, aimed at a focused control of the particular accident creating processes of the particular system. This strategy exercises safety control through a proper design of the system and a safe plan for its operation. It depends on a predictive risk analysis, based on an understanding of the mechanisms underlying the hazards of the system and the accident creating processes. This requires a reliable predictive model of system function, a stable system and stable operating conditions, and a pre-planned risk management with reference to the predictive risk analysis".
>
> Ibid., page 4.

In the terms established in chapter 4, the principal distinction is between the *direct* and *indirect* modes of regulator design (ie, design by the metasystem) as discussed on page 149, ante. In keeping with this, the variety held in the metasystemic model of the system will greatly differ between these two forms. The "control by objective" being the variety required by a muddy-box level of system perception; whereas the "analytical safety" requiring the much higher variety of transparency in the "boxes" comprising the system.

A third mode of control may overlay whichever is the dominant form present in a particular system: ***Control by evolutionary choice.***

> "Even when systems are planned for open loop, analytical control, feedback correction will take place following accidents. It is basically a matter of time scales whether feedforward or feedback control is the typical control mode. Often, basic safety features of the individual system are controlled analytically by design whereas the level of safety across the industry is subject to empirical control by evolutionary choice. That is, the predictive model and the plans for operation will be updated for future systems and their operation in reaction to accidents in present systems. This control strategy depends on stable conditions, that is, past experience is relevant for future operation, and on reliable generalization from particular events to type of relationship"
>
> Ibid., page 4.

An interesting feature of Rasmussen's scheme is that it implies the importance of modelling recursions. Thus the pace of technological change at the meta-level, whilst remaining a highly relevant concern, needs also to be accompanied by a consideration of the rate of change in work-practices when the system level is taken as the system-in-focus. Similarly, the level of consequence (ie, high vs. low hazard in Rasmussen's scheme) when dropped a level of recursion needs to be accompanied by a more detailed consideration of the consequences at this level. This opens-up the right hand cells of Figure 5.3 to enquiry.

With regard to the class of system nominated as "concentrated sources of hazard - slow pace of change" the distinction between *direct* and *indirect* modes of regulatory design is revealing. In his discussion, Rasmussen notes that such industries typically rely on the detailed specification of the system and its local regulation, the *direct* approach in Ashby's terminology. However, this is not achieved analytically but by the accretion of designer competence over many generations of system. So long as the variety of the disturbances to be regulated against is slow-changing (relative to the "life-span" of systems established in this way) this is likely to be effective. Just as industrial process plants require their designer's assumptions to be maintained, so too do these systems. However, as Rasmussen points out, the assumptions of design are not maintained but instead are increasingly falsified. In short, the regulatory systems (at the system and meta-levels) are not adaptive to their changing environment: regulatory requisite variety is not maintained. As there are limits as to how much of the missing variety that can be "pumped-in" by individuals at the operational level, accidents become inevitable.

In conclusion, Rasmussen's classification appears compatible with the terms set out in chapter 4. Therefore, it may provide a scheme for identifying systems in the empirical work discussed in section 5.2 (page 200, ante) and in chapter 6.

5.3.2 Supra-organisational regulation

Whilst much of our interest is necessarily focused upon the organisations that create the risks to which workers and the public are exposed (and indeed the creation of wealth upon which society relies), particular emphasis needs to be given to the regulatory apparatus at the level of the State, at the level of industries, and their interaction with each other and the organisations "beneath". It seems to me that the VSM (and general cybernetic formulation of the situation) is of clear application here. General indications for further work in this area are given in the next chapter.

As a general point of concern, the fact that cybernetics in general, and the VSM also, provides no defence against the employment of values (and consequent definition of the η-set) which are decidedly antithetical to the well-being of individuals. The VSM requires the well-being of *functional* elements and, on occasions when the functional element corresponds to an individual, this seems fair. However, this speaks only of *statistical* persons: people qûa workers, the loss of one suffers only the inconvenience of replacement with another. However, this rather bleak appearance both acknowledges the amoral excesses of some organisations and ignores the basic issue of the State at the nth recursion. The State has the key position to rein-in the excesses of amoral organisations - because only at this level are people truly ends in themselves. Wildavsky (1989a; 1989b) provides the antithesis to this argument by insisting that the trend is for wealthy societies to be safer - as people become richer within a society they effectively buy themselves less risk, market forces thus driving up the threshold of risk acceptance. Obviously, a common attribute to these "wealthy" nations is that, without exception, the countries in question all possess H&S regulatory agencies and are all democracies (people as ends). However, whilst average wealth may seem high - modern day industrial societies are typified by increasing polarisation of income. Hence, richer may be safer *on the average.* This takes us back to the problems as originally stated in chapter 3 - if a State can legitimately pursue a policy of the grand-outcome which must discount the rights of the individual to the statistical person, then so be it. If on the other hand this legitimacy is challenged at the supra-State level or through the operation of the democratic

process, then the net-benefit approach cannot be sustained. A market-forces argument speaks only of means and not of ends, if a society wishes to formulate its operation in terms of ends, then it has little option than to regulate to achieve them. The basic requirement of this, as we have seen, is to decide what is important and what counts as success. Thereafter, through feedforward and feedback regulation in a requisite hierarchy, may what is defined as success be achieved in practice.

Rasmussen (1996) citing the conclusions of Breyer (1993) notes that the success of the state level regulator depends upon a number factors:

> "He [Breyer] finds that success will depend on five factors:
>
> 1. the group must have specific risk related missions, together with
> 2. the inter-agency jurisdiction needed to transfer resources. The group must also
> 3. have a degree of political insulation to be able to withstand pressure and
> 4. have *prestige* to be able to attract a highly capable staff *that understands science, some economics, administration, and law.* It must have the ability to communicate in a sophisticated way with experts in all these fields. Finally, the group must
> 5. have the *authority* that will give it a practical ability to achieve results."
>
> Rasmussen, 1996, page 25.

In the UK, we are unusually fortunate to have in the HSE/HSC, a regulatory agency which conforms to at least four of these five requirements (the requirement for inter-departmental co-operation[175] is a matter which, I understand, could be improved upon). However, just as Rimington (1992; 1993) makes clear the strength of the HSC/HSE under the first, fourth and fifth quoted requirements, the documents themselves attest to a political situation which is weakening the "insulation" suggested as necessary in the third. If there is a general conclusion to my own reading of the regulatory requirements for H&S it is that, whilst organisations themselves are the vital amplifier between policy set by the State and the actual probabilities of harm to workers and the public, self-regulation must be considered at the level of the whole. Thus the State must ultimately arbitrate what is good and regulatory agencies such as the HSE must dispose requisite variety to the task of ensuring it.

[175] ie, between HSE and other civil service departments.

6 Conclusions & Further Work

" 'I'm very glad,' said Pooh happily, 'that I thought of giving you a Useful Pot to put things in.'
'I'm very glad,' said Piglet happily, 'that I thought of giving you Something to put in a Useful Pot.' "

Milne, A.A.

Introduction

As set out in the Preface, this thesis has been concerned with ways and means of approaching the study of H&S regulation. In pursuing this, fundamental questions have been asked and, in most cases, the best that has been achieved is to reformulate the questions so that answers might be found. In H&S and other areas[176], the response to complex problems seems too often to be oversimple answers. These are enthusiastically received and acted upon until the complex reality intercedes: as Jaques notes with some acerbity: "Encouraged by gimmicks and fads masquerading as insights, we have burdened our managerial systems with a makeshift scaffolding of inept structures and attitudes" (Jaques, 1990, page 128). What is advocated in this dissertation is an approach to gaining a better understanding of the nature of the H&S regulatory task, and a means of facilitating the sharing of knowledge between the various inquirers in and of the H&S regulatory system. There is no final right answer nor any inherently supreme right view.

6.1 Conclusions

Set out at (a) to (e) below are summary conclusions together, where applicable, with suggested areas of further work.

[176] As Senge (1992) puts it "Because we see the world in simple obvious terms, we come to believe in simple, obvious solutions. This leads to the frenzied search for simple "fixes", a task that preoccupies the time of many managers." (page 267).

(a) *It has been suggested that the study and practice of safety management has been hampered by a poorly developed conceptual structure.*

This conclusion was reached by a semantic examination of vocabulary. This revealed highly variable meanings for even the most basic terms used in the area of safety management. It is suggested that adoption of a systems theoretical perspective would assist by providing a unifying set of concepts.

(b) *Ethics has been argued to be a necessary aspect of the study of safety management. Further, the cybernetic view of regulation provides a unifying basis for dialogue between ethics and safety management.*

Many of the issues within safety management impinge upon wider debates of ethics in general and business ethics in particular. It was noted that the safety literature generally avoids "moral arguments". It is hoped that the compatibility demonstrated between the cybernetic view of regulation and ethical perspectives (page 90) may serve as an example of bridging the gulf between the technical and moral domains.

(c) *MORT has been shown to be compatible with a cybernetic view of regulation. This provides a means of bringing MORT methods and literature into correspondence with those of systems theory.*

It is suggested that further work be undertaken to formalise the correspondence between MORT and Systems theory. As noted at page 128, this concerns the set of relations between regulation, information, thermodynamics and the McFarland (1967) energetics.

(d) *In its present form, MORT is not an appropriate means of modelling safety management systems. This is because MORT, whilst strong in describing* ***detail complexity****, does not address the structure or* ***dynamic complexity*** *of organisational systems. However, it is suggested that cybernetics provides a means of addressing these systemic factors and that the VSM may provide a framework to achieve this end.*

Further work is required to develop a modelling methodology based on MORT and the VSM and this is described at 6.2. It should be noted that whilst the VSM is advocated as a means of describing and abstracting the systemic aspects of regulation, there are other approaches that could be considered for

use in this way. Examples of alternatives include the Soft System Methodology (Checkland and Scholes, 1988) and the "tools for systems thinking" provided by Senge (1992; 1994). Future work should review these for application in the context of safety management research.

(e) *It is argued that the structure of any safety management system is entailed by the structure of the organisational system from which it is abstracted. Thus, efforts towards the development of a "template for safety management system design" for general application, are very likely misdirected.*

It has been argued that there are certain regulatory principles that invariably apply to safety management in organisational systems. However, this will only be true at a high level of description. At a low, that is, very detailed level of description, each system will vary from each other. However, as Rasmussen (1994) suggests, it may be that *within* certain classes of organisational system there is less variability than between classes. If so, some generalised models may be developed at an intermediate level of description. The question remains, and further work needs to address, how to develop classification schemes for this purpose.

6.2 Development of a modelling methodology

The justification for this work has been stated in chapter 5. In summary, a theoretical approach was identified as involving an unfeasible level of difficulty: both MORT and the VSM are generative techniques for modelling aspects of actual rather than theoretical organisations. Therefore, the development of a modelling technique to describe the regulatory arrangements for H&S in organisations will be facilitated by working with an organisation as an external frame of reference. This will allow the correlation of (a) the requirements of H&S management (as classified in MORT) and (b) the general regulatory arrangements of the organisation (as codified by the VSM).

The primary objective of the work will be the development of a first generation methodology for use in later modelling (although even at this stage one would hope collateral benefits to be obtained - immediate results of use to the host organisation and the production of hypotheses concerning the relation between H&S regulation and "general" regulation to be examined in later studies).

The intermediate goals of this fieldwork were given on page 209 (ante) but are repeated here for ease of reference:

a) to provide an expansion and clarification of the recursive structures implied in MORT;
b) to identify MORT processes which require modelling of recursive levels outside of the system-in-focus (eg, at the level of the state or industry);
c) to highlight discontinuities between the handling of information which is *explicitly* concerned with H&S and the handling of information relating to other major essential variables (eg. productivity, quality, human resource development);
d) to highlight areas in the VSM which do not allow mapping (ie, incommensurabilities between organisations as conceived in MORT and as conceived in the VSM).

In sections 6.2.1 and 6.2.2, the requirements for this empirical work are elaborated. However, the purpose of this elaboration is *not* to specify in detail exactly how this work must be conducted but to indicate the general form that it is likely to require.

6.2.1 Choice of host organisation

A concern here is the type of organisation that might best serve these aims. As discussed in chapter 5 and at 6.1(e) above, the variations in regulatory systems *between* types of organisations (such as chemical process, railway transportation, and general manufacturing) may be greater than the *within* each class. However, the typology suggested by Rasmussen (1994) is characterised by archetypes - extremes of reliance upon analytical risk management, or extreme reliance upon established technology or empirical (reactive) management. Given that the research aim at this stage is to develop a basic modelling methodology, a reasonable choice of organisational host would be one with a mixture of concentrated and distributed energy sources. Also, given the requirement of the Management of Health and Safety at Work regulations (1992), the majority of organisations will be using analytical approaches as part of their regulatory effort. Lastly, the organisation should be one which is known to be reasonably competent in its management of H&S.

The scope of the analysis would be expected to encompass the level of individual operatives through to the level of senior executives. In the example of chapter 4, the analysis would consider the operative level of FKM's three businesses and FKM

itself. The level of detail required is difficult to state with precision. It should be sufficient to resolve the responsibilities and characterise the roles of the individuals; and to identify the communication network which connects them within the organisation. One might expect this requirement to have some bearing upon the size of organisation chosen for the study. However, it seems to me that the dimension of importance is not the absolute number of staff, but the diversity of operations. For example, in British Telecommunications plc, the variety of different types of task performed by its engineers is smaller than the number of engineers: the same types of task being repeated in a large number of different geographical areas.

6.2.2 Indicative format for the study

There are three main phases. The first requires a VSM description[177] of the host organisation with no emphasis upon its safety management. The second phase involves mapping MORT requirements onto the VSM description obtained. The aim of this is to produce a VSM augmented with explicit regard to H&S management. The third phase involves applying this augmented VSM to the organisation with explicit consideration to the regulation of H&S performance.

First Phase: VSM description (*explicit regulation of essential variables other than safety performance*)

1. As in any application of the VSM, particular attention should be given to **correctly identifying** System One and its subsystems[178] starting at what Johnson (1980) refers to as field level (eg. Kitchen-Co and its peers as the elemental subsystems in FKM's system One (page 164, ante).

    a) Throughout the analysis, efforts should be made to follow the advice of Nertney (1992) to "*hang it on the people*" - in other words, identify

[177] Beer (1985) provides an accessible description of how VSM analysis may be undertaken. Similarly, Espejo (1989) provides a useful and concise source of reference.

[178] As Beer (1989) points out, some applications of the VSM have fallen at this hurdle by including such matters as maintenance as an operational subsystem within System One. System One includes *only* those activities that produce the wealth of the organisation: they are what the Viable System does.

at all stages of analysis the people involved and the various different "functional hats" they are wearing[179]

2. The horizontal domain loops should be identified at this stage.

   a) This includes those between elemental management, operation and environment, and minimally[180] the squiggly line loops connecting the operations.

   b) For each loop examined here and later, a "Homeostatic Loop Analysis (HLA)[181]" should be attempted (as described and exemplified in Beer, 1985).

3. Following this, the various co-ordinating efforts of the metasystem served via System Two should be identified.

   a) How thorough (ie, finely/coarsely graded) the analysis is required to be at this stage is a moot point: generally enough for the analyst to get a *feel* for the emerging communication network and media).

   b) The general requirements for System One's design of System Two needs also to be explored as this provides a means of beginning the exploration of the System Three-Two relationship.

   c) The same rationale applies to the identification of processes "belonging" to the Three-Star function (ie, three-star should be approached both from the operational *end* and the System Three *end*).

4. System Three should then be approached from the perspective of the elemental System One managers, using:

   a) the resource bargaining process (which includes accountability arrangements), as well as;

---

[179] An additional agenda for analysis is the level of dissonance serving these different functional roles inflicts and the methods used by these people to ameliorate this (such as satisficing, Simon, 1956)..

[180] That is, the environmental loops, unless of obvious significance, should be ignored at this stage.

[181] This concerns the input and output transducers, the mediating amplifiers and attenuators and the criteria of stability used by local management to ensure the operation of the loop under the "test" of requisite variety.

b) the operation of the intervention channel descending from System Three.

c) The examination of System Three should focus on the resource allocation process. This includes

i) the synergistic co-operation and communication between the components of System Three (identification of the System Three components will be well under way following investigation of the parties involved in the resource bargain, and the various processes and activities identified for Systems Three-Star and Two).

d) The *n*-way homeostatic interactions of System Three components is likely to be extremely informative both in regard to

i) how resources are synergistically organised (ie, as services to System One);

ii) how such *rules* as are required are formulated, and;

iii) strategies developed to implement (mostly via System Two) the requirements of the higher recursive level[182].

5. The methods that System Three employs to develop its strategies are important ways, in my view, to begin enquiries into the activities of System Four and the relations (ie, information exchange) between System Three and Four.

a) Where did the methods in use in System Three originate,

b) what determined the technology in use in System One

c) what is directing future-oriented change in System Three.

6. Hopefully, answers to questions of the kind exemplified above should indicate the existence and components of System Four (I say hopefully, because experience suggests that System Four may not be developed in the

[182] There may be several of these in different recursive dimensions. If the System-in-focus is the highest level of system within the organisation (as was FKM Ltd, in the example of Chapter Four) then these recursive dimensions may include the industrial bodies relevant to this sector, the HSE and other statutory organisations, and professional bodies such as the IEEE, IMechE.

organisation and the role of System Four is of great importance to risk-related decisions at this level of the organisation[183]).

a) Attention should be focused on the information channels between System Four and the environment (as the information it obtains a crucial factor in the future development of the organisation)

b) Again HLA should be performed to gain an appreciation of this function.

c) The other important aspect of analysis of System Four concerns the trade-offs it makes between the various essential variables of the organisation...

d) comparing the η-set thus derived with the current η-set as apparent in the functioning of System Three (this is, as it were, a change engine - propelled by the difference between "*where we are*" and "*where we need to be*"). As indicated in Chapter 4, this *change engine* involves considerable exchange of information (the Three-Four homeostat) as the twin needs of (a) resources to maintain and operate System One and its services (3-2-1) are balanced against (b) the resources needed to change System 3-2-1 to maintain adaptation with the environment (and the requirements of the next recursion - *see footnote 182*).

[183] System Four (and the Three-Four homeostat) is of particular moment because it is upon this that notions such as *reasonable practicability* rely, and hence to actual risk levels acceptable in relation to various identified adverse states. Cost-benefit analysis, for example, is very much a Three-Four process: Three is providing details of risk assessment studies conducted via System Three-Star, Four will be considering these data in relation to predicted losses (eg. via future-discounted cash flows). Further, Four will also be taking cognisance of the practice of its peer systems (ie, competitors), particularly, figures for value-of-life, BATNEEC, etc. Whilst this process may in theory lead to a *ratcheting-up* of risk acceptance thresholds (i.e. towards greater margins of safety), of greater concern is a *levelling-down* effect (ie, higher levels of risks becoming accepted with the passage of time). For example, parcel delivery services impose schedules upon their drivers which force them to drive unsafely (examinations of the schedules and routes reveal that they cannot be met within the speed limits). The reason for this is that extremely fierce competition opposes the obvious solution of more vans and drivers (the creation of slack resources) as these cannot be provided without increasing the cost of delivery to customers. No one company dare do this because of the loss of their competitive position in the market. (Royal Society for the Prevention of Accidents, 1996)

e) In practice, one might expect the Three-Four homeostat to be a potentially unstable device as the desire for constancy and the desire for change struggle for convergence (or indeed supremacy if a power struggle is observed). Again, an HLA needs to be prepared to describe the Three-Four homeostat including the practical matter of the physical embodiment and time constraints which pertain.

7. The outcome of this last aspect of the analysis should indicate System Five and its role in:

a) monitoring the stability of the Three-Four homeostat;

b) providing closure in endorsing/selecting the η-set definition arising from the operation of the Three-Four homeostat, part of which is;

c) the production and communication (ie, in writing and in person[184]) of "statements of intent" and the more general notion of the corporate ethos.

d) Lastly, in relation to System Five, is its communication with System One. Here we need to examine:

i) What algedonic "alarms" arise from System One, how is the gain set on this channel (eg. how timid is System One, how interested is System Five, how in general is the gain calibrated;

ii) How are the algedonic channels maintained[185];

iii) which essential variables are fitted with "alarms";

iv) the adequacy of these channels (by HLA) and

v) verification (another HLA) of the channels assumed by (c) above.

---

[184] For example, the manager of a local (Coventry) chemical plant makes a habit of floor-walking the various areas of the site. He, for some reason, is particularly motivated to see good manual handling practices which he personally enforces. This leaves no one in any doubt that this and matters associated with it, are part of the real η-set (as opposed to the espoused η-set) of this organisation. Incidentally, this plant has maintained a productivity level of 175% for the last year which demonstrates both questionable metasystemic measurement techniques and, perhaps the general effectiveness of this management style across a range of essential variables.

[185] Perhaps fractional dead-time (FDT) concepts can be applied to such channels?

Second Phase: mapping MORT into the VSM

This stage is aimed at producing a VSM description *augmented* for examination of *explicit* safety regulation. This, as indicated at (a) to (c) has been produced preserving the logic of the VSM and to incorporate the MORT logic of what is generically required in the management of H&S. Another way of looking at this is that the intermediate stage aims to translate MORT into VSM conventions.

8. The MORT processes are mapped onto the VSM description of the organisation obtained in the first phase. This is done on the strength of the information obtained so far. The expectation here is that certain MORT elements will:

   a) **fracture** into sub-process each allocated to a different VSM function or channel (a hypertext database will need to be set up to handle the resulting cross-referencing of MORT codes and VSM codes) this, it is suggested, is highly desirable as it represents the augmentation of MORT;

   b) be **replicated** at different recursive levels not yet modelled (these will need to be recorded for subsequent VSM description of the adjacent recursions);

   c) **fail to readily map** into the VSM (these should be noted for more fundamental cybernetic analysis - the question being *what* regulatory function do these elements perform? If it is subsequently found that the MORT is cybernetically valid yet does not conform to VSM conventions, this may suggest problems with the VSM logic itself).

Third Phase: MORT augmented VSM description of the organisation (*explicit regulation of H&S essential variables*)

The rationale is that the augmented VSM will be used to examine the difference between the regulation of H&S essential variables and the regulation of essential variables *not explicitly* related to H&S (as revealed in the first phase) by re-examination of the organisation in a second-pass.

9. This second-pass will be much accelerated by the information gathered in the first. This second pass will:

   a) explicitly consider H&S regulation in the system-in-focus;

   b) consider the relations of the H&S regulation to regulation of the other essential variables in the system-in-focus;

   c) note the recursions required by MORT but not yet modelled at 8(b) above.

The **differences** are expected to be instructive:

- Some differences will indicate that H&S regulation *requires* unusual information routing in the organisation;
- Some differences will indicate that H&S regulation is not handled well by arrangements that adequately regulate other essential variables;
- Some differences will indicate that H&S regulation is handled better than the regulation of other essential variables. For example, personal experience of joint committee work in the Fire Service revealed that communication and decision-making for safety-related change was often handled with greater ease and efficiency than other categories of performance.

The **similarities** are expected to be instructive:

- Some similarities will indicate that H&S regulation is accomplished via the same routes and regulatory functions;
- Some similarities will indicate that the organisation achieves poor regulation of all (or not merely H&S) its essential variables.

6.3 Specific theoretical areas of MORT and the VSM requiring further study

6.3.1 Planning

As mentioned in chapter 1, Planning is the communication process which ties the adjoining recursive levels of the VSM together. The description offered by Beer (1985) for applying the VSM is quite shallow in this respect; planning is confined to his discussion of the "resource bargain" although it is implied throughout. In contrast, the treatment by Beer (1979) is more abstract and considerably more detailed (pages 335-356 give explicit consideration). Whilst understanding and

applying the VSM can be achieved without this more complex treatment of planning, the topic is an important one and further work is required to evaluate the approach described by Beer (1979) .

6.3.2 Measurement of instability

Measurement is an area of clear synergy between Beer's work and H&S management. Beer's notions of measurement are based *not* in the notion of *what gets measured gets done* - which certainly invokes the functions of the "resource bargain" and its associated accountability feedback, but, moreover - *what doesn't get measured is also happening*. To some extent systems 2 and 3* are dealing with this but in a feedforward or feedback fashion - both predicated on existing understanding of cause and effect in the system. Beer's measures are metasystemic in their perception because, at this level, management "...does not deal with the *stuff* of the system at all - it deals only with the managerial consequences of what the system does" (Beer, 1979, page 289). Hence, metasystemic measurement is concerned with identifying *instability* - the early signs that the system is not maintaining control to satisfactory levels.

The ideas presented by Beer have much in common with forms of statistical analysis such as "extreme value" methods (Briscoe, 1982). The perception of instability does not require an analytical understanding of the system involved. The definition of hazard in chapter 2, suggested that the term best describes a unstable state in a system in which the probability of the system then moving to one or more adverse states is greater than that acceptable to the system's regulators. In essence, *hazards* are states of the system **outside of the margins of control**. As we move up recursive levels, the occurrence of incidents and accidents in the various operating level systems, can be used to statistically describe these margins of control. It is suggested that H&S regulation could make wider use of such techniques and others developed along similar lines to those indicated by Beer (1966 and 1979). Particular emphasis should be given both to (a) deriving measures of more populous events than accidents and (b) utilising these and related data at successive levels in the recursive control hierarchy. An example would be *near-miss* occurrences although it should be noted that some method of classification would be needed to structure these data (e.g. by class of hazard involved as in the example from Johnson (1980) cited on page 95, ante).

6.3.3 Risk Assumption and individual discretion

Allied to the both the issues mentioned above is the *communication* of risk acceptance criteria and, with it, the *discretion* for accepting risk. As stated in the MORT User's manual (see Appendix 1):

> "Assumed Risks (R) ...
>
> R factors are defined as only those risks that have been analysed and accepted by the proper level of management; unanalysed or unknown risks are not considered to be Assumed Risks" (page 9) ...
>
> "Was there a specific decision to assume each risk? Was it made by a person who had management delegated authority to assume the risk?" (page 55).

Further work is needed to address the process by which the "proper" level of management is defined. Research questions include (taking the lead offered by "delegated authority"): what conditions should be associated with this delegation; how are such boundaries of discretion defined and communicated; to be defined; what is the process for upward referral?

6.3.4 Algedonic alarms

Lastly, a particular research issue concerns how appropriate algedonic channels be designed and maintained. A popular belief amongst experienced accident investigators is that "there were always warnings" before the accident in question. How then do we design these channels and ensure that they are responsive enough to provide timely "alarms" to senior management but without a false-alarm record likely to breed their demise?

6.4 Study of statutory regulation

In addition to studies of particular organisations and the development of such generalisations as empirical findings allow, it is suggested that research should also be aimed at exploring the statutory regulatory task.

The basic problem faced by regulatory agencies such as the HSE is the variety in the organisations to be regulated is large compared to the variety that the HSE can dispose to the task of regulation. The argument that "more inspectors" is not necessarily an answer to the task of regulating industry has already been

introduced in section 4.2.5 (page 133, ante). This is not to say that there are already too many (!) or not enough. We have repeatedly encountered the method of amplifying regulation by stages. Further work should be directed to considering whether this principle can be employed in the service of statutory regulation.

Another side to this argument concerns the perception of amplified regulation as design-by-stages: each regulator above making selections of the regulator below. Recasting the problem of statutory regulation in this light suggests that the HSE has a role in designing the regulators at the industrial level. There are three subsidiary questions here:

- What sources of supplementation can be utilised;
- How to achieve this regulatory design whilst observing the "minimal intervention" principle (as much a political as cybernetic requirement);
- How to achieve this affordably.

This is certainly a complex problem but not an insoluble one. For example, in chapter 3 (page 91, ante) the collaboration between the State (as embodied in the HSC) and industrial bodies such as the CBI and TUC was noted to be an effective partnership in devising H&S policy. What is suggested is that this be pursued down recursive levels and equipped at the supra-organisational level to provide a service to the organisations beneath as well as gathering data (very much in the mode of the VSM's Systems Two and Three-Star. To be politically acceptable as well as cybernetically valid, the further proviso is that this regulatory function be composed, so far as possible, by people from the industry in question and financed jointly by the businesses themselves, insurers, as well as by the State. I can well imagine the reactions of many to this suggestion (eg. "Cloud-cuckoo-land idealism!", "monstrous bureaucracy!", etc). However, Aulin's notion of "requisite hierarchy" (page 155, ante) suggests that intervening levels between the organisation and the state may be necessary. A hierarchy that lacks the necessary layers simply cannot dispose adequate variety to the task - if the research is undertaken and establishes the need for change: are we to shoot the researchers and burn their notes? Whilst I do not have blind faith in the computer as the answer to everything: *imaginative* use of IT may allow such change as is revealed necessary, to be accomplished without major upheaval and expense (c.f. footnote 162, page 198, ante).

Considerable value would be added to such research by undertaking parallel studies in other EU Member States. Germany, for example, makes use of local level organisations (chambers of commerce) as well as Insurance companies for improving the reliability of feedback data (accident reporting).

List of References

Anderton, R. (1989). The need for formal development of the VSM. In: The viable system model: interpretations and applications of Stafford-Beer's VSM. Edited by Espejo. R., and Harnden. R., John Wiley & Sons, Chichester. pp. 39-50.

Argyris, C. (1988). Problems in producing usable knowledge for implementing liberating alternatives. In: Decision making: descriptive, normative and prescriptive interactions. Edited by Bell, D.E., Raiffa, H., and Tversky, A., Cambridge, Cambridge University Press.

Argyris, C. (1994). Good communication that blocks learning. Harvard Business Review, July-August 1994, pp. 77-85.

Ashby, W.R. (1956). Introduction to cybernetics. London, Chapman and Hall.

Ashby, W.R. (1960). Design for a brain. 2nd edition, London, Chapman and Hall.

Atlan, H. (1983). Information theory. In: Cybernetics Theory and Applications. Edited by Trappl, R., London, Hemisphere(1983).

Aulin, A. (1982). The cybernetic laws of social progress, towards a critical social philosophy and a criticism of Marxism. Oxford, Pergamon Press.

Aulin, A. (1989). Foundations of mathematical system dynamics, the fundamental theory of causal recursion and its application to social science and economics. Oxford, Pergamon Press.

Beer, S. (1959). Cybernetics and management. English Universities Press.

Beer, S. (1966). Decision & control. John Wiley & Sons, Chichester.

Beer, S. (1979). The heart of enterprise. John Wiley & Sons, Chichester.

Beer, S. (1981). Brain of the firm. John Wiley & Sons, Chichester.

Beer, S. (1983). Introduction: questions of quest. In: Cybernetics Theory and Applications. Edited by Trappl, R., London, Hemisphere

Beer, S. (1985). Diagnosing the system for organisations. John Wiley & Sons, Chichester.

Beer, S. (1989). The viable system model: its provenance, development, methodology and pathology. In: The viable system model: interpretations and applications of Stafford-Beer's VSM. Edited by Espejo. R., and Harnden. R., John Wiley & Sons, Chichester. pp. 39-50.

Beer, S. (1989b). National government: disseminated regulation in real time, or 'how to run a country'. In: In: The viable system model: interpretations and applications of Stafford-Beer's VSM. Edited by Espejo. R., and Harnden. R., John Wiley & Sons, Chichester. pp. 39-50.

Beer, S. (1993). World in torment: a time whose idea must come. Kybernetes, Vol. 22, No. 6., pp. 15-43.

Bird, F.E., and Germain, G. L. (1986). Practical loss control leadership. Institute Publishing, Loganville, Georgia.

Brazier, M. (1993). Street on Torts. Ninth Edition. Butterworths, London.

Briscoe, G.J. (1982). Risk management guide. DOE-76-45/11, SSDC-11, EG&G Idaho, Idaho Falls, USA.

British Standards Institution (1996). BS 8800: Guide to health & safety management systems. London, BSI.

Bjørn-Andersen, N., Eason, K., and Robey, D. (1986). Managing computer impact, an international study of management and organizations. Norwood, N.J., Ablex.

Bullock, M.G. (1981). Change control and analysis. DOE-76-45/21, SSDC-21, EG&G Idaho, Idaho Falls, USA.

Cantor,G. (1915). Contributions to the founding of the theory of transfinite numbers. Dover Publications New York, reprinted 1955.

Carroll, J.M., and Rosson, M. (1985). Usability specifications as a tool in iterative development. In Harton, R (Ed). *Advances in HCI 1*. Ablex.

Checkland, P.B. (1980). Are organisations machines? Futures, Vol. 12, pp. 421-424.

Checkland, P.B. (1986). Review of "Diagnosing the system". European Journal of Operational Research. Vol. 23, pp. 269-270.

Checkland, P.B. (1988). Soft systems methodology: an overview. Journal of Applied Systems Analysis, 15, pp. 27-30.

Churchland, P.S. (1988). The significance of neuroscience for philosophy. Trends In Neurosciences, 1988, Vol. 11, No.7, pp. 304-307.

Committee on the Financial Aspects of Corporate Governance (1992). The financial aspects of corporate governance. (The Cadbury report). London, The Committee and Gee and Co.

Conant, R. and Ashby, W.R. (1970). Every good regulator of a system must be a model of that system. International Journal of Systems Science, 1970, Vol. 1, No. 2, pp. 89-97.

Craik, K.J.W. (1943). The nature of explanation. Cambridge, University Press.

Currie, R. (1968). System safety and industrial management. National Safety Council, Chicago, Illinois.

De Board, R. (1978). The psychoanalysis of organizations, a psychoanalytic approach to behaviour in groups and organizations. London, Tavistock Publications.

De Raadt, J.D.R (1991). Information and managerial wisdom. Paradigm Publications, Idaho.

Department of Energy (1990). The public inquiry into the Piper Alpha disaster. HMSO.

Dworkin, R. (1996). In: Analysis: Power to the Judges. BBC News and Current Affairs Department, Radio 4. Transcript of broadcast on 1.2.96. BBC Tape reference TLN605/96VT1005.

Eason, K.D. (1988). Information Technology and Organisational Change. Basingstoke, Taylor & Francis.

Eisner, H.S. and Leger, J.P. (1988). The International Safety Rating System in South African mines. Journal of Occupational Accidents, Vol. 10, pp. 141-160.

Espejo, R. (1989). A cybernetic method to study organisations. In: The viable system model: interpretations and applications of Stafford-Beer's VSM. Edited by Espejo. R., and Harnden. R., John Wiley & Sons, Chichester. pp. 361-382.

Espejo, R., and Harnden, R.J. (1989). The VSM: an ongoing conversation. In: The viable system model: interpretations and applications of Stafford-Beer's VSM. Edited by Espejo. R., and Harnden. R., John Wiley & Sons, Chichester. pp. 39-50.

Etzioni, A. (1961). A comparative analysis of complex organizations. New York, Free Press.

Eysenck, H.J. (1985). Decline and fall of the Freudian empire. Harmondsworth, Viking.

Fennell, D. (1988). Investigation into the King's Cross underground fire. Department of Transport., HMSO.

Fitzgerald, R.M. (1996). Matched-phase noise-reduction. Journal Of The Acoustical Society Of America, 1996, Vol. 99, No. 3.

Fortune, J. and Peters, G. (1995). Learning from failure - the systems approach. Wiley, Chichester, 1995.

Freud, S. (1915). Instincts and their vicissitudes. In: The standard edition of the complete psychological works of Sigmund Freud. Edited by Strachey, J. and Freud, A., Vol. 14. London:Hogarth (1957).

Galbraith, J. (1973). Designing complex organisations. Reading (Mass), London, Addison-Wesley.

Galbraith, J. R. (1994). Competing with flexible lateral organisations. 2nd Edition. Addison-Wesley.

Gentner, D. and Stevens, A.L. (1983). Mental models. Hillsdale, N. J., London, Laurence Erlbaum Associates.

Gilbert, S.F. (1994). Developmental biology. 4th Edition, pub. Sinauer Associates, Sunderland.

Gödel, K. (1992). On formally undecidable propositions of Principia mathematica and related systems. translated by B. Meltzer. Dover Publications, New York.

Goguen, J.A., and Varela, F.J. (1979). Systems and distinctions: duality and complementarity. International Journal of General Systems, 5 (1), pp. 31-43.

Grimsley, R. (1973). The philosophy of Rousseau. Pub. Oxford University Press, London.

Hale A.R. and Glendon A.I. (1987). Individual behaviour in the control of danger. Amsterdam, Oxford, Elsevier.

Hale, A.R. and Hale, M. (1972). A review of the industrial accident literature. HSMO, London.

Harnden, R. (1989). Outside and then: An interpretive approach to the VSM. In: The viable system model: interpretations and applications of Stafford-Beer's VSM. Edited by Espejo. R., and Harnden. R., John Wiley & Sons, Chichester. pp. 39-50.

Hegel, G.W.F (1874). The logic of Hegel: translated from The encyclopaedia of the philosophical sciences with prolegomena. Edited by Wallace, W., Clarendon Press, Oxford.

Helstrom, C.W. (1995). Elements of signal detection and estimation. Pub. Englewood Cliffs, N. J., Prentice-Hall International, London.

Hendy, J., Ford, J., and Brodie, D. (1993). Redgrave Fife & Machin, Health and Safety. 2nd Edition, Butterworths, London.

Hidden, A. (1989). Investigation into the Clapham Junction railway accident. Department of Transport. HMSO.

Hofstadter, D. (1985). Gödel, Escher, Bach: An eternal golden braid. Penguin. England.

Horman, R.L. (1992) Glossary of SSDC terms and acronyms. DOE-76-45/28, SSDC-28. EG&G Idaho, Idaho Falls, USA.

HSC (1993). ACSNI Human Factors Study Group. Third report: Organising for safety. HSE Books.

HSC (1994). Review of health & safety regulation - Main report. HSE Books.

HSE (1988). The tolerability of risk from nuclear power stations. HSMO, London.

HSE (1991). Successful health & safety management. HS (G) 65. HSE Books.

HSE (1993). The costs of accidents at work. HS (G) 96. London, HMSO.

HSE (1996). Generic terms and concepts in the assessment and regulation of industrial risks. Discussion document. HSE Books.

IAEA (1991). Safety culture: A report by the international nuclear safety advisory group. Report: Safety Series No. 75-INSAG-4.

Jackson, M.C. (1988). An appreciation of Stafford Beer's 'viable system' viewpoint on managerial practice. Journal of Management Studies, Vol. 25, No. 6, pp. 557-573.

Jackson, M.C. (1989). Evaluating the managerial significance of the VSM. In: The viable system model: interpretations and applications of Stafford-Beer's VSM. Edited by Espejo. R., and Harnden. R., John Wiley & Sons, Chichester. pp. 39-50.

Jackson, M.C. (1993). Signposts to critical systems thinking and practice. Kybernetes, Vol. 22, No. 5, pp. 11-21.

Jacob, F. (1974). The logic of life, a history of heredity. London, Allen Lane.

Jaques, E. (1955). Social systems as a defence against persecutory and depressive anxiety. In: New directions in psychoanalysis. Edited by Klein, M., Heinmann, P., and Money-Kyrle. T., London: Tavistock publications.

Jaques, E. (1990). In praise of hierarchy. Harvard Business Review, Jan-Feb., Vol. 68. pp. 127-33.

Johnson, W.G. (1970). New approaches to safety in industry. London, Industrial & Commercial Techniques.

Johnson, W.G. (1973). MORT - The Management Oversight and Risk Tree. SAN 821-2. US Atomic Energy Commission.

Johnson, W.G. (1980). MORT safety assurance systems. New York, Marcel Dekker.

Johnson, W.G. (1985). Accident/incident investigation manual. 2nd edition, DOE-76-45/27, SSDC-27, EG&G Idaho, Idaho Falls, USA.

Johnson-Laird, P.N. (1983). Mental models, towards a cognitive science of language, inference, and consciousness. Cambridge, Cambridge University Press.

Joseph, H.W.B. (1916). An introduction to logic. Oxford, Clarendon Press.

Kandel, E.R. and Schartz, J.H. (1991). Principles of neural science. 3rd edition. London, Prentice-Hall International Inc.

Kant, I. (1937). Fundamental principles of the metaphysic of ethics: translated by Abbott, T.K., 10th edition. Longmans, Green, London ; New York

Ostrom, L., Wilhelmsen, C., and Kaplan, B. (1993). Assessing safety culture. Nuclear Safety, Vol. 34, No. 2, pp. 163-172..

Kepner, C.H., and Tregoe, B.B. (1965). The rational manager - a systematic approach to problem solving and decision-making. Kepner-Tregoe, Inc.

Kingston-Howlett, J.C. (1994). Observations on "The accident investigation process" (published by SSDC). Personal communication with Dr R. J. Nertney (12 July 1994).

Kingston-Howlett, J.C., and Nelson, H.K. (1995). Events and causal factors analysis. SCIE-DOE-01-TRAC-14-95, Scientech Inc., Idaho Falls, USA.

Kingston-Howlett, J.C., Nelson, H. K. and Nertney, R. J. (1995). Barrier analysis. SCIE-DOE-01-TRAC-28-95, Scientech Inc., Idaho Falls, USA.

Kletz, T.A. (1988). Learning from accidents in industry. London, Butterworths.

Kletz, T.A. (1991). Plant design for safety: a user-friendly approach. London, Hemisphere Publishing.

Klir, G.J. (1983). General systems concepts. In: Cybernetics Theory and Applications. Edited by Trappl, R., London, Hemisphere.

Klir, G.J. (1991). Facets of systems science. Plenum Press, New York.

Klir, G.J. and Rozehnal, I. (1996). Epistemological categories of systems: A overview. Intern. J. of General Systems, 24(1-2), pp. 207-224.

Knox, N.W. and Eicher, R.W. (1992). MORT user's manual. DOE-76-45/4, SSDC-4, EG&G Idaho, Idaho Falls, USA.

Koestler, A., and Smythies, J.R. (1969). Beyond reductionism: new perspectives in the life sciences. Hutchinson, London.

Kuhn, T.S. (1970). The structure of scientific revolutions. 2nd edition. Chicago, University of Chicago Press.

Kunda, G. (1992). Engineering culture: control and commitment in a high-tech corporation. Philadelphia, Temple University Press.

Levens, E. and Krikorian, M. (1970). Search 1. ASSE Journal, January, 1970.

Lord Lester (1996). In: Analysis: Power to the Judges. BBC News and Current Affairs Department, Radio 4. Transcript of broadcast on 1. 2. 96. BBC Tape reference TLN605/96VT1005.

McFarland, W.R. (1967). Application of human factors engineering to safety engineering problems. National Safety Congress Transactions.

Melzack, R., and Wall, P.D. (1965). Pain mechanisms: A new theory. Science., Vol. 150, pp. 971-979.

Mesarovic, M. (1970). Theory of hierarchical, multilevel systems. New York, London, Academic Press.

Mill, J.S. (1975). Utilitarianism: Introduction by Warnock, M., Collins/Fontana, Glasgow.

Milne, A.A. (1926). Winnie the Pooh. London, Methuen & Co.

Morgan, G. (1986). Images of organization. Sage Publication, Delhi.

Nertney, R.J. (1987a). Process operational readiness and operational readiness follow-on. DOE-76-45/39, SSDC-39, EG&G Idaho, Idaho Falls, USA.

Nertney, R.J. (1993). Personal communication with the author (10 July 1993).

Nertney, R.J. (1993). Personal communication with the author (16 September 1993).

Nertney, R.J. (1994). Personal communication with the author (7 October 1994).

Nertney, R.J., Clark, J. L. and Eicher, R. W. (1975). Occupancy-use readiness manual - Safety considerations. ERDA-76-45-1, SSDC-1, EG&G Idaho, Idaho Falls, USA.

Norman. D.A. (1988). The psychology of everyday things. New York, Basic Books.

Nüsslein-Volhard, C. (1996). Gradients that organise embryo development. Scientific American, Vol. 275, No. 2, pp. 38-43.

Pask, G. (1975). The cybernetics of human learning and performance, a guide to theory and research. London, Hutchinson Educational.

Perrow, C. (1984). Normal accidents: living with high-risk technologies. New York, Basic Books.

Rasmussen, J. (1980). What can be learned from human error reports? In: Changes in Working Life. Edited by Duncan, K.D., Gruneberg, M.M., and Wallis, D.J., Wiley, Chichester, 1980.

Rasmussen, J., and Batstone, R (1989). Why do complex organisational systems fail? Environment working paper No. 20, The World Bank, Washington, D. C.

Rasmussen, J. (1994). Safety control: some basic distinctions and research issues in high hazard low risk operation. In: Control of Safety. Edited by Bremer, B. and Reason, J., Hove, UK: Lawrence Erlbaum Associates.

Rasmussen, J. (1996). Risk management in a Dynamic Society: A Modelling Problem. Key-note address: Conference on human interaction with complex systems. Dayton, Ohio, August 1996.

Raven, F.H. (1995). Automatic control engineering. 5th edition. New York, London, McGraw-Hill.

Reason, J. (1990). The contribution of latent human failures to the breakdown of complex systems. Phil. Trans. R. Soc. London. Vol. 327, No. 1241, pp. 475-484.

Rimington, J. (1992). Market testing of HSE regulatory functions - Feasibility study. Report to Secretary of State (employment). In Confidence.

Rimington, J. (1993). Market testing of HSE regulatory functions 2 - HSE's Scientific services. Report to Secretary of State (employment). In Confidence.

Rinaldi, G. (1992). A history and interpretation of the logic of Hegel. Lewiston, N.Y., Lampeter, Edwin Mellen Press.

Rooker, J. (1994). Personal communication with the author (30 August 1994).

Rosen, R. (1986). Some comments on systems and system theory. International Journal of General Systems, 13(1), pp. 1-3.

Royal Society for the Prevention of Accidents (1996). Personal communication with the author (Bibbings - Kingston-Howlett). March, 1996.

Ryan, T. (1991). Organisational factors research - lessons learned and findings. US NRC.

Senge, P.M. (1992). The Fifth Discipline, the art and practice of the learning organization. London, Century Business.

Senge, P.M. (1994). The Fifth Discipline fieldbook, strategies and tools for building a learning organisation. London, Nicholas Brealey Publishing Ltd.

Shannon, C. and Weaver, W. (1949). The mathematical theory of communication. Urbana [Illinois], London, University of Illinois Press.

Shaw, W.H. And Barry, V. (1989). Moral issues in business. 4th edition. Belmont, California, Wadsworth Pub. Co.

Shroeder, H.M. (1970). Safety performance measurement. Symposium, National Safety Council.

Simon, H.A. (1956). Rational choice and the structure of the environment. Psychological Review, Vol. 63, pp. 129-138.

Simon, H.A. (1960). The new science of management decision. Harper & Row, New York.

Simon, H.A. (1962). The architecture of complexity. Proceedings of the American Philosophical Society. Vol. 106, pp 467-482.

Simon, H.A. (1982). Models of bounded rationality: Vol. 1. Economic analysis and public policy. Cambridge, Mass., London, MIT Press.

Soltis, J.F. (1968). An introduction to the analysis of educational concepts. Addison Wesley.

Sommerhoff, G. (1950). Analytical biology. Oxford University Press, Oxford.

SPSS Inc. (1988). SPSS-x™ User's guide. 3rd Edition. SPSS, Chicago.

Stephenson, J. (1991). System safety 2000: a practical guide for planning, managing and conducting system safety programs. Van Nostrand-Reinhold, New York.

The Royal Society. (1992). Risk analysis, perception and management. The Royal Society, London.

The Times Higher Education Supplement (1996). Call to reduce role of reductionist strategy in science: report by Scottish Editor Wojtas, O., 18 April 1996.

Tversky, A., and Kahneman, D. (1974). Judgement under uncertainty: heuristics and biases. Science, Vol. 185, pp 1124-1131.

U.S. Nuclear Regulatory Commission (1991). Influence of organisational factors and performance reliability. NUREG CR-5538.

Umpleby, S. (1979). Heinz Von Foerster, a second order cybernetician. Cybernetics Forum, Vol. IX, No. 3.

Varela, F.G., Maturana, H R. and Uribe, R. (1974). Autopoiesis: The organization of living systems, its characterisation and a model. Biosystems, Vol. 5(4), pp. 187-196.

Veltri, A (1990). An accident cost impact model: the direct cost component. Journal of Safety Research, Vol. 21(2), pp. 67-73.

Von Bertalanffy, L. (1950). The theory of open systems in physics and biology. Science, Vol. 111, January 1950, pp. 23-29.

Von Bertalanffy, L. (1968). General systems theory. George Braziller, New York.

Von Foerster, H. (1994). Ethics and second-order cybernetics. The Stanford Electronic Humanities Review. Vol. 4, Issue 2.

Waddington, C.H. (1957). The strategy of the genes, a discussion of some aspects of theoretical biology. London, Allen & Unwin.

Watson, O.C. (Ed). (1976). Longman modern English dictionary. Longman Group Limited, London.

Weaver, W. (1948). Science & complexity. American scientist Vol. 36, pp. 536-544.

Weick, K.E. (1995) Sensemaking in Organisations. Thousand Oaks, CA. Sage.

Weldon. T.D. (1962). States and morals - A study in political conflicts. London, John Murray.

Wickens, C. (1992). Engineering psychology and human performance. 2nd edition. New York, NY, Harper Collins Publishers.

Wiener, N. (1948). Cybernetics. New York, Wiley.

Wiener, N. (1954). The human use of human beings: cybernetics and society. 2nd edition. London, Eyre and Spottiswoode.

Wildavsky, A. (1989a). The secret of safety lies in danger. Society Vol. 27 (2). 1989. pp. 3-5.

Wildavsky, A. (1989b). Thanks for commentary: replies to critics and critiques. Society Vol. 27 (2). 1989. pp. 28-31.

Woods, D.D. (1987). Technology alone is not enough: reducing the potential for disaster in risky technologies. In: Human Reliability in Nuclear Power. Proceedings of a two-day conference, Regent Crest Hotel, 22-23 October 1987.

Appendix 1

MORT - The Management Oversight & Risk Tree

The title MORT conveys two meanings: the MORT programme (as developed at Aerojet Nuclear - Johnson, 1973) and the expression of this programme as a logic tree. The latter is the subject of this appendix. This contains a copy of the MORT question set together with a the MORT diagram reproduced as a series of cutsets[i]. Perusal of the cutsets in tandem with the corresponding entries in the manual indicates the rich interconnection that, perforce, is lost in a fault tree representation. As inspection of the diagrams and corresponding text shows, MORT cannot be properly understood from study of the fault tree representation alone - hence the inclusion of the questions here.

MORT is a generalised fault tree which is to say that it has been designed for application to any occupational accident. The events modelled in the MORT diagram fall into two classes. The first class of events relate to a particular barrier or control failure judged to be causal in an accident sequence[ii]. The objective of this part of the analysis is to establish the causes of the failure in question which appear to be under local (managerial and, generally, geographical) control. These events are located in the *specific control factors* (or "S") branch of the tree. In the manual, the questions relating to these events are in the past tense as it is assumed that these will be consulted for incident/accident investigation. The tree is traversed from left to right and top to bottom. Each event represents a hypothesis and a cue for the analyst to make a judgement of the event as: *irrelevant* on this occasion; *less than adequate (LTA); satisfactory,* or; as *requiring more information* before a judgement can be made. In practice, this process is accompanied by colour coding the chart (irrelevant -

[i] At the time of publication, the most recent edition of the MORT diagram is that produced by the author. Copies of this edition are available from the Health & Safety Unit at Aston University or from the Technical Information Service of Scientech, Inc., 1690 International Way, Idaho Falls, ID 83402, USA.

[ii] It should be noted (and this is conveyed by MORT event SB4) that the use of MORT in accident investigation relies upon a preliminary analysis of the sequence of accident events. This is required to identify the number and sequence of energy-barrier-target transactions. This is achieved via "Barrier Analysis" (Kingston-Howlett, Nelson and Nertney, 1995). For each barrier judged LTA (less-than-adequate) at this stage of analysis, a MORT analysis will be performed.

crossed out in black, LTA - *red*, adequate - *green* and, more information - *blue*[iii]) and annotating the chart with notes justifying the evaluation.

The second class of events, is found in the *management system factors* (or "M") branch of MORT. These events are intended both for accident investigation and for appraisal (ie, proactive) analyses. Because of this, the questions are stated in the present tense. The two branches of MORT are usually presented as concerned with "what happened?" (the "S" branch) and, in relation to this, "why?" (the "M" branch). This is perhaps misleading to those unfamiliar with MORT, as both branches are predominantly concerned with failures in the human element of control. This is evidenced in the "S" branch from the SD level downwards, the explicit cross-referrals from the "S" to the "M" branch and, yet more so, in the referrals implied by the questions and commentary of the user's manual.

The MORT diagram is, at first sight, rather daunting. Although the tree is extremely extensive, its detail arises from a *limited span* of high events as indicated in the truncated version below.

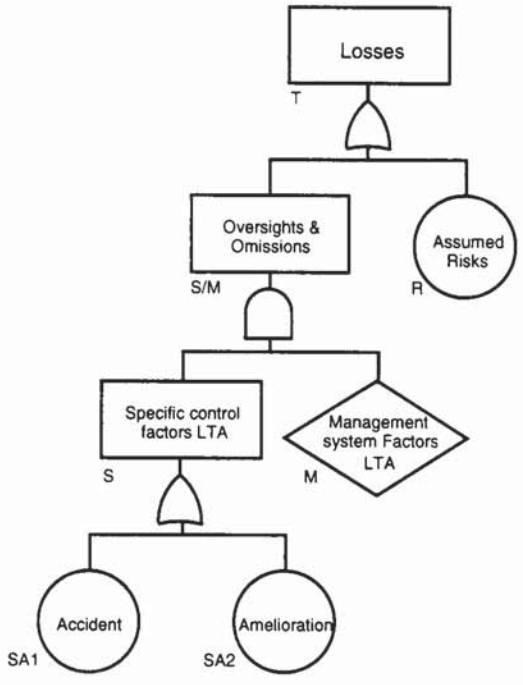

The top event is a generalised statement of the losses incurred. Descending through the OR gate, two hypothetical alternatives are presented: either the accident was due to faults in the system (Oversights & Omissions) *or* the accident was the outcome of a risk correctly identified and accepted (an *assumed risk*). If

[iii] One of the most useful aspects of MORT analysis is the prompting of further enquiries - hence a chart with many blue-coded events is an indication of the usefulness of the exercise.

the latter event is true, the analysis terminates here. In such circumstances the occurrence of the accident is an indictment of the correct functioning of the system. However, it is the "oversights & omissions" hypothesis which is true in the vast majority of cases.

If the Oversights & Omissions hypothesis is entertained, an AND gate is traversed. This conveys the logic that for the particulars of the work-place/work-practice ("S" branch) to have caused an accident then *it is always true* that this is due to problems in the management system. Returning to the specific side of the tree, the argument is presented that losses incurred through the accident were either due to the accident itself *or* (non-exclusively[iv]) the emergency action associated with it. This is to say that a minor accident can be made greatly worse by inappropriate amelioration measures (eg. attempting to extinguish a chip-pan fire with water). The remaining events in MORT develop this argument into the full MORT tree (some 1500 events including the transfer elements) additional to which are the analytical trees relating to operational readiness (Nertney, 1987; Nertney, Clark and Eicher, 1975).

MORT has been slow to gain acceptance in Europe and this is thought to be a product of the limited training available here as well as the perception of MORT as difficult and cumbersome. However, interest in MORT as a tool for accident investigation is now growing - particularly in the United Kingdom. This is perhaps a reflection of the change in regulatory emphasis from reliance upon prescriptive legislation towards H&S regulation through risk management[v].

[iv] All OR gates used in the MORT chart are non-exclusive (meaning "AND/OR").
[v] As well as the fact that MORT and its supporting literature is written in English.

MORT User's Manual Page Number **13**

The following question set is reproduced from the most recent edition of the MORT User's Manual (Knox and Eicher, 1992). The page numbering (**top right hand corner**) preserves the pagination of the original as this is a useful source of cross reference in the MORT diagram. The MORT diagram has been reproduced as a series of cutsets at the end of this question set. The relevant User Manual page number is shown in **bold** face to the left of each event MORT in the trees. Another feature of importance is the page referencing to Johnson (1973) which indicate further criteria to assist the judgement of the adequacy of a given MORT event during the process of investigation (see, for example, event SD1-d1 on page 22 - MORT manual numbering).

QUESTIONS FOR THE MORT ANALYST

The following questions may be used in conjunction with the MORT diagram. The reader is reminded that the Specific Control Factors (S) branch questions presume an accident, whereas the General Management Factors (M) branch questions are phrased from the assumption of a safety system evaluation. Whatever is the actual case, the user should have no great difficulty in making the mental adjustment.

T. Fundamental Questions (the Top event)

What happened?
Why?
What were the losses? (Specify the number and type of injuries, the amount of property damage, production downtime, product degradation, reduction in employee morale, programme impact, negative publicity, or any other type of loss.) [The construction layout shows an alternative top event connected to the diagram by a dashed line. This unique method is employed to show the duality of MORT application. When using MORT as an appraisal tool, the analyst views the Top event statement as future potential losses that may result from an Assumed Risk or from an Oversight/Omission existent in the safety system evaluated.]

S/M. Oversights and Omissions

[The tree structure depicts two fundamental causes of the adverse consequences listed by the Top event: (1) Management Oversights and Omissions, or (2) Assumed Risks. All contributing factors in the accident sequence are seen as Specific Oversights and Omissions until such time as they are transferred to Assumed Risks. Further discussion of Assumed Risks is provided under its appropriate heading. Input to the Oversights and Omissions event is through an AND logic symbol, because MORT experience, to date, shows the S and M branches to be mutually inclusive.]

S
SA1
SA2

S. Specific Control Factors

What were the specific control factors of the management system that were overlooked or omitted? [Detailed understanding of the incident/accident sequence leads naturally to: (1) consideration of the Management System Factors, and (2) judgement whether the fault (failure potential) was an Assumed Risk.]

SA1. Accident

Describe what happened.
[MORT conceives the accident occurred when an unwanted energy flow or environmental condition that results in adverse consequences reaches persons and/or objects. MORT combines this concept and others into a functional accident definition as follows: An unwanted transfer of energy or environmental condition because of lack or inadequate barriers and/or controls, producing injury to persons and/or damage to property or the process.]

SA2. Amelioration LTA

Once an accident has occurred, was there adequate amelioration on the part of all concerned parties?

[Amelioration can only be considered and evaluated after an accident, thus the "Accident Occurrence" constraint on the gate leading to lower branch elements. The intent of amelioration is to limit the consequences of what has immediately occurred and to reduce the sensitivity of those consequences whenever possible. When evaluating Amelioration from an overall management system standpoint, consider the following: (1) Are all of the amelioration functions pre-planned (as opposed to the possibility of having them occur fortuitously at the time of a particular accident)? (2) Does the plan adequately scope the types and severity of accidents which it intends to cover? (3) Are adequate resources allocated to properly execute the plan? and (4) Is management aware of any residual risk beyond the scope of the plan?]

a1. Prevention of Second Accident LTA:
Was prevention of a Second Accident or further losses adequate? Through the efforts of individuals at the accident scene and those who arrived later, were steps taken to prevent a second accident caused directly or indirectly as a consequence of the first?

b1. Plan LTA:
If properly executed, was the plan adequate to accomplish the intended function? Was the plan provided to those who needed it?

b2. Execution LTA:
Was the plan executed as was intended?

c1. Practice LTA:
Was there sufficient practice of various plan assignments? Was the practice realistic?

c2. Personnel and/or Equipment Changes:
Were there personnel or equipment changes that caused the execution of the plan to be LTA? Were trained personnel free of any recent physical or mental changes? Was the equipment familiar to the users and free of defects or modifications?

d1. D/N Counter-change:
Had appropriate counter-changes been considered and introduced where applicable for changes in personnel or equipment?

c3. Task Performance Errors:
Was the plan executed properly through successful completion of all steps?
[Note the transfer in of other lower tier events from SD5-b3.]

a2. Emergency Action (Fire Fighting, Etc.) LTA:
Was the emergency response prompt and adequate?
Which emergency response teams were required? Were they notified and did they respond?
[Include local facility fire brigade, health physics team, fire department, bomb squad, and other speciality teams. Be sure to consider delays or problems in both notification and response.]
[Note other lower tier events included by transfer from a1.]

a3. Rescue LTA:
Were trapped or immobilised victims satisfactorily removed to a safe area? Before entering a hazardous area, did rescuers consider the risk of injury to themselves versus the ability to lessen the severity of injuries to victims? Include the evacuation of employees or the public from potentially hazardous areas.
[Note other lower tier events included by transfer from a1.]

a4. Medical Services LTA:
Was adequate medical service available?

b3. First Aid LTA:
Was adequate first aid immediately available at the scene? Was it used properly to prevent immediate injuries from becoming more severe?

b4. Transport LTA:
Was mobile service available to transport medical personnel and equipment to the accident scene and/or to transport injured to medical facilities? Was transport executed properly?

c4. Plan:
Was there a medical service plan? Was it distributed to appropriate personnel?
[Consider such things as: (1) how to make a notification, (2) training of medical personnel and drivers and when they are available, and (3) who and what equipment will respond.]

c5. Notice:
Was notification made in an adequate time and manner? Were employees instructed on how to notify medical services?
[Consider whether notification process was easy to do, especially during the stress of an emergency.]

c6. Personnel and Equipment:
Did the personnel use the equipment correctly? Did the equipment function properly? Did the medical and transport personnel have all the equipment necessary to properly perform the jobs expected of them? Were the personnel adequately trained relative to the postulated needs?
[Consider whether equipment could be operated easily during the stress of an emergency.]

c7. Distance:
Was there a significant distance between medical services and the area to which the service responded?
[If the distance is great, response time is increased.] [Note the event is flagged with R1 assumed risks. Top management must assume distance/time response risk.]

b5. Medical Treatment LTA:
Was there adequate medical treatment en route and at the medical facilities?

a5. Rehabilitation LTA:
Was rehabilitation of persons and objects made after the accident?

b6. Persons:
If the injury was disabling, could its overall disabling effect have been reduced and/or the individual made more functional? If such rehabilitating activity was possible, was it done?

b7. Objects:
Was damaged equipment, buildings, or other property expeditiously repaired, salvaged, or replaced?

a6. Relations LTA:
Was there a management plan outlining the protocol to be followed and steps to be taken subsequent to a significant accident? Was the accident news disseminated to all concerned parties in a proper and timely manner?

b8. Employee:
Did the relatives of the injured employee first hear about the accident from a responsible, tactful individual within the organisation? Were the other employees in the organisation notified firsthand about the accident with some assurance that significant corrective action would be taken?

b9. Officials:
Were the facts about the accident given accurately and in a timely manner to the proper officials of: (1) the organisation, (2) the customer, (3) the local municipality, (4) the state, and (5) other governmental agencies as appropriate?

b10. Public and

b11. Media:
Were the news media (and thereby the public) given the accident facts and assurance that significant corrective actions were being taken? Was a specific point of contact within the organisation provided as the source of additional information?

SB1. Potentially Harmful Energy Flow or Environmental Condition

What was the energy flow or environmental condition that resulted in the accident. [SB1 denotes an energy flow or environmental condition which could result in harm if barriers and controls are inadequate and a vulnerable person or object is exposed.]

a1. Nonfunctional:

Was the energy flow or environmental condition causing the harm a functional part of or product of the system?

b1. Was there adequate control of nonfunctional energy flows and environmental conditions?

b2. Was such control practicable?

[Note that the event is flagged with R4 assumed risk symbol. Proper management level must assume responsibility for this decision.]

a2. Functional:

Consider the lower tier elements below this only if the harmful energy flow or environmental condition was a functional part of or a product of the system. Given a failure of the barrier system, the following questions should be considered:

b3. Were the administrative controls adequate to prevent the harmful energy flows or environmental conditions from reaching vulnerable persons or objects?

SB2
SB3

b4. Diversion LTA:

c1. Was there adequate diversion of harmful energy flows or environmental conditions?

c2. Was diversion impractical?

[Note that this event is flagged with R5 assumed risk symbol. An appropriate management level should assume risk responsibility for this decision.]

SB2. Barriers and Controls LTA (Incident)

Were adequate barriers and controls in place to prevent vulnerable persons and objects from being exposed to harmful energy flows and/or environmental conditions?

Note: The constraint placed on SB2 is intended as a device to prevent oversight. It is designed primarily to draw attention to barriers and control related to harmful energy flows or environmental conditions and those controls designed to control movement of target persons or objects.

Both types of barriers should be considered but rigorous and proper classification is not necessary to the analytical processes, provided that all barriers are considered.

Were the barriers and controls designed to prevent harmful energy flows or environmental conditions from reaching vulnerable people and objects LTA? [Refer to SC1 and SC2 for further development.]

Were barriers and controls designed to prevent vulnerable people and objects from encountering harmful energy flows and environmental conditions LTA? [Refer to SC1 and SC2 for further development.]

SB3. Vulnerable People or Objects

Note: The constraint "value" in place here. An accident is defined in terms of loss of something of "value."

What vulnerable people and/or objects of value were exposed to the harmful energy flow or environmental condition?

a1. Nonfunctional:

Was the person or object performing a functional role in operation of the system?

b1. Was there adequate control of nonfunctional persons and objects?

SB4
SC1
SC2

b2. Was such control practicable?

[Note that the event is flagged with R4 assumed risk symbol. Proper management level must assume risk responsibility for this decision.]

a2. Functional:

Consider the lower tier elements below this only if the person or object was performing a functional role in operation of the system. Given a failure of the barrier system consider the following: **Note**: The constraint in place here. An accident can only occur if the barriers were LTA.

b3. Were the administrative controls adequate to prevent persons or objects from being exposed to the harmful energy flow or environmental condition?

b4. Evasive Action LTA:

c1. Was there adequate evasive action for vulnerable persons or objects?

c2. Was evasion impractical?

[Note that this event is flagged with R5 assumed risk symbol. An appropriate management level should assume risk responsibility for this decision.]

SB4. Events and Energy Flows Leading to Accident-Incident

What were the events and energy flows leading to conversion of hazards to actual accident-incidents. (Analyse as appropriate to the accident events.) [Refer to SC3 and SC4 for further development relating to precursor events and energy flows.]

Note: Energy-barrier analysis (1) and events and causal factor analysis (2) should be used as appropriate to the situation.

SC1. Control LTA

Were there inadequacies in the control system that was established to prevent vulnerable people and objects from interacting with harmful energy flows or environmental conditions? [Analyse the specific nature of these failures in terms of SD1 through SD6.]

SC2. Barriers LTA

Were there failures of barrier systems provided to prevent interactions between vulnerable people and objects and harmful energy flows or environmental conditions?

Were there adequate barriers? What were the specific barriers? [The 12 barriers listed on page 33 of the MORT text and their order of listing should be reviewed carefully. MORT treats the first five listed barriers as a function of concept and design. The four "intermediate" barriers (i.e., source-target) are listed by MORT as specific inputs to this Barriers event. The final three "target" barriers are treated elsewhere on the MORT diagram. The example of grinding wheel safety practices, Figures 2-3 on page 35 of the MORT text, is particularly helpful in illustrating the barrier concept.]

Note: The following breakdown (a1, a2, a3, a4) is intended as a device to prevent oversight. All barrier types should be considered.

Rigorous and proper classification in terms of a1, a2, a3, a4 is not necessary to the analytical process provided that all barriers are considered. A supplementary barrier analysis form available from SSDC may be used in recording this information.

a1. Were there barriers on the energy source?

[Note other lower tier events included by transfer from a3.]

a2. Were there barriers between the energy source and the injured person/damaged equipment?

[Note other lower tier events included by transfer from a3.]

a3. Were there barriers on persons and/or objects?

[Note all lower tier development under this event also transfers to a1., a2., and a4.]

b1. None Possible:

[Note use of the Diamond event symbol, indicating termination of fault sequence because of the lack of solution. Note also the event is flagged with R2 assumed risk symbol. Appropriate management must assume risk for design where no barriers were possible.]

b2. Barrier Failed:

Did the barrier function as intended?

b3. D/N Use:

Were barriers used?

c1. D/N Provide:

Were barriers provided where possible?

SC3
SC4
SD1

[Note the event is flagged with R3 assumed risk symbol. An appropriate level of management must assume risk for failure to provide barriers, e.g., failure to provide safety glasses.]

c2. Task Performance Errors:

Were the provided barriers used properly? (e.g., Were available safety glasses improperly used?) [Note that all the lower tier development under event SD5-b3 transfers to this event also.]

a4. Were there "barriers" of time or space which separated the energy and the person or object?

[The term "barrier" has the connotation of physical intervention; however, the barrier may be a "paper barrier." Separation by time or space in particular may be accomplished by written procedure or some other type of administrative control.]

[Note other lower tier events included by transfer from a3.]

SC3. Barriers and Controls LTA

Were barriers and controls on energy transfers and other events leading to conversion of a hazard to an actual accident less than adequate? [Refer to SC4 for description of these events and to SB2 for further development relating to barriers and controls on preceding events.]

SC4. Events and Energy Flows

What were the precursor events and energy flows that resulted in conversion of a hazard to an actual accident? [Refer to SB1 for further development relating to these preceding events.]

Note: Energy-barrier analysis (1) and events and causal factor analysis (2) should be used as appropriate in the SC3 and SC4 subjects.

SD1. Technical Information Systems LTA

Was the technical information system adequate (with respect to the unwanted energy flow)? (349)

[Complex work flow processes must be supported by complete technical information systems. It is axiomatic that complex systems will depart from plans and procedures to some degree. Therefore, information systems need to detect deviations, determine rates and trends, initiate corrections, and, in general, assure that goals are attained. MORT conceives a technical information system as consisting of "research" persons, "program" persons, and "action" persons obtaining, handling, and providing technical information relevant to the work flow process in a "communication" network.]

a1. Technical Information LTA:
Was there adequate technical information relevant to the work flow process? [Often relevant information exists but is not available to the "action" persons associated with the process. Possible reasons are investigated by the following series of questions.]

b1. Knowledge LTA:
Was knowledge of the work flow process adequate? [The question is investigated by subdividing into known and unknown precedent.]

c1. Based upon known precedent (i.e., for the prevention of the unwanted energy flow):

d1. Was application of knowledge obtainable from codes and manuals adequate? (260)

d2. Was the list of experts (to contact for knowledge) adequate?

d3. Was any existing but unwritten precedent relevant to the work flow process (i.e., part of the supervisor's regular practice) known to the "action" person?

d4. Were there studies directed to the solution of known work flow process problems? Was the effort being spent in the search for their solution reasonable and adequate? (265)

c2. If there was no known precedent:

d5. Were there investigation and analysis (i.e., risk analysis) of prior similar accidents/incidents or the work flow process accident potential? Was the investigation adequate? (95)

d6. Was there research directed to the obtaining of knowledge about the work flow process? Was the research effort reasonable and adequate? (97)

b2. Communication LTA:
Was the exchange or transmittal of knowledge adequate (relative to the potential unwanted energy transfer)?

c3. Was the internal communication adequate? (391)

d7. Was the definition of the internal communication network adequate?

d8. Was operation of the internal network adequate?

c4. Was the external communication adequate? (411)
[The query relates to the interface between the in-house (internal) information system and national information systems, such as the National Safety Council, NASA, NSIC, and others.]

d9. Was the definition of the external communication network adequate?

d10. Was the operation of the external communication network adequate? Was the method of searching, retrieving, and processing relevant information adequate?

a2. Monitoring Systems LTA:
Was the monitoring system adequate? Were the principal elements of a good monitoring system present? (351)
[Highly complex work flow processes require a high order of excellence of the monitoring subsystem of the technical information system. The management Risk Assessment system must be closed-loop to maintain the process "in control." Triggers for fast action fixes and data for achieving long-range hazard reduction goals are generated by the Monitoring System and transmitted by the Technical Information System to the Hazard Analysis Process (HAP) portion of the Risk Assessment System.]

b3. Was the safety observation plan (employed by work flow process supervision) adequate?

b4. Was there a planned independent search out effort for high potential hazards by a safety professional? Was the safety inspection search out effort adequate?

b5. Was incident/accident information, relative to prior incidents/accidents in similar processes, recorded and reviewed?

b6. Was there a planned Reported Significant Observation (RSO) system? Was it operative? (116, 361) MORT uses this term for a special safety study rather than the better known "critical incident" for semantic reasons. ("Critical" and "criticality" have very different and specific meaning in the nuclear energy field.) The RSO concept relates to the study of near-miss incidents observed and reported by line supervision and work level personnel.]

b7. Error Sampling System LTA:
Was there an error sampling plan? Was it operating adequately? (357)
[Error sampling is a specific management plan whereby staff personnel systematically sample for operating errors, using prepared checklists and definitions.]

b8. Were the routine field safety work site inspections made? Were they adequate?

b9. Was the audit of "upstream" work flow processes conducted in an adequate manner? (114)

[MORT separates the general work flow process into: (1) worksite operations; and (2) upstream work flow processes such as design, construction, selection and training, etc. Each segment must be examined relative to the three basic work ingredients - hardware, procedures, and people.]

b10. Was the general health monitoring of work flow process personnel adequate? (373)

a3. Data Collection and Analysis LTA:
Were the data collection and analysis procedures adequate? Were there analyses (i.e., measurement techniques) made of the data? Did the analyses provide the proper risk assessment information to the decision maker responsible for the risk assumption? (415, 372)

b11. Was there a priority problem list? Had it been updated to be a current list? (375)
[Management should, at all times, know what its most significant assumed risks are thought to be. Any delay in corrective action for budget reasons becomes an assumed risk for the present.]

b12. Were the available status and predictive statistics adequate? (415)

b13. Was the diagnostic statistics analysis adequate?

b14. Was the risk projection analysis adequate?

b15. Was the "War Room" status display of current problems, analyses, and results adequate? (439, 97)

a4. HAP Triggers (Fix Control Initiators) LTA:
Were triggers (stimuli) for the initiation of the Hazard Analysis Process (HAP) adequate? Were they utilised to obtain early safety anticipation and review in planned or unplanned changes? (233) [MORT postulates HAP triggers as part of the HAP portion of the Risk Assessment System, but originating from the Monitoring Subsystem of the Technical Information System.]

b16. One-On-One Fixes LTA:
Was the information from the technical information system adequate to trigger the HAP preventive action plan for individual problems? (397)

b17. Priority Problem Fixes LTA:
Was the information from the technical information system adequate to provide a continuous trigger to the HAP Priority Problem Lists? (234)

b18. Planned Change Controls LTA:
Were HAP triggers from planned changes in the work process adequately recognised? Were they used? (233)

b19. Unplanned Change Controls LTA:
Were HAP triggers from unplanned changes in the work process adequately recognised? Were they used? (233)

b20. New Information Use LTA:
Were HAP triggers from research, new standards, etc., detected and used? (234)

a5. Independent Audit and Appraisal LTA:
Was there a recent appraisal of the total safety system (or audits of parts thereof)? Were the audits and appraisals conducted in a truly independent manner? Was the appraisal plan adequate? (371, 399)

SD2. Facility Functional Operability LTA

In accident/incident analysis, the accident is considered to be prima facie evidence that the system was not operationally ready. MORT, therefore, deals with "operational problems" by evaluating operational readiness in terms of readiness of the three major system elements: "plant/hardware," "procedures/management controls," and "personnel." In MORT analysis, the "operational problems" are analysed in terms of these three elements in two steps: (1) Operational readiness trees are used to better define and localise the functional inadequacies; and (2) Conventional MORT analyses are used to define causal factor chains associated with the inadequacies (including definition of root causes).

Was the facility and process operationally ready? Were the necessary supplementary operations supportive to the main process ready? (293) [This branch probes the status of "upstream processes" (design, training, etc.) which supports the ingredients of the work process (hardware, procedures, and people). The ingredients used at the worksite are obtained from two major upstream sub processes: (1) the original design, construction, test, and qualification plus documents defining operating limits and performance specification, and (2) modification projects to the facility. All "upstream processes," including the Hazard Analysis Process, are susceptible to constructive analysis as "work processes" in themselves. Each upstream process can be analysed as to hardware, procedures, and personnel.]

a1. Verification of Occupancy-Use Readiness LTA:
Was verification of the facility and/or work process adequate? [Two publications of the System Safety Development Center, SSDC-1, "Occupancy-Use Readiness Manual," September 1975, and SSDC-39, "Process Operational Readiness and Operational Readiness Follow-On," February 1987, provide detailed criteria for this major functional branch.]

SD3
SD4

b1. Was the conduct of an operational readiness review specified?

b2. Were the criteria used for determining the facility or process readiness adequate?

b3. Was the required procedure for determining occupancy-use readiness followed?

b4. Were the personnel who made the decision on occupancy-use readiness adequately skilled and experienced?

b5. Was the follow up of action items from occupancy-use readiness review adequate? Were all outstanding action items resolved prior to start-up of the work flow process?

a2. Organisational and Functional Relations LTA:
Was there adequate technical support furnished to the work flow process, particularly at the worksite? Were the organisational versus functional relationships adequate to assure the required level of operability? [Highly complex processes need close field liaison by scientific and engineering personnel.]

a3. Interface Between Operations and Maintenance and Testing Activities LTA:
Was the interface between operations personnel and testing and maintenance personnel adequate? Were administrative procedures well-planned to preclude misunderstanding of operational status due to a breakdown of communication?

a4. General Design Process LTA:
Was the actual physical arrangement or configuration identical with that required by latest drawings, specifications, and procedures? Were the configuration and documentation of modification to the facility or process adequately controlled? Was the general design process adequate to assure functional operability?

SD3. Maintenance LTA (Basic logic same as SD4, Inspection LTA)

SD4. Inspection LTA

Was there adequate maintenance (or inspection) of equipment, processes, utilities, operations, etc?

a1. Plan LTA:
Was the plan scope broad enough to include all the areas that should be maintained (or inspected)? Was management aware of those areas not included in the plan?

b1. D/N Analyse Failure for Cause:
Did the plan require that any failed item be analysed for cause of failure? Were the analysis results required to be acted upon by an appropriate individual or group?

(Note: Items b1 and b2 of this section were interchanged in the Maintenance branch.)

b2. D/N Specify:
Were maintainability (inspectability) requirements specified by the design or procurement documents? If not, are they provided adequately by operations plans? (311)

c1. Maintainability (Inspectability) LTA:
Did the plan address methods for minimising problems with equipment, processes, utilities, operations, etc. when they are undergoing maintenance (or being inspected)?

c2. Schedule LTA:
Was there a schedule? Did the plan schedule maintenance (inspections) frequently enough to prevent or detect (as appropriate) undesired changes? Was the schedule readily available to the maintenance (inspection) personnel? Was the schedule co-ordinated with operations to minimise conflicts? (313)

c3. Competence LTA:
Did the plan specify minimum requirements for the competence and training of individuals used in the program?

a2. Execution LTA:
Was there adequate execution of the maintenance (or inspection) plan?

b3. Task Performance Errors:
Were the individual tasks (as set forth by the plan) performed properly? [Note other lower tier events included by transfer SD5-b3.]

b4. D/N Maintain "Point-of-Operation" Log:
Was there a log of maintenance (inspections) kept at the point-of-operation on the piece of equipment, process, etc.? (311) [This is distinct from other logs that may be kept in a control room, back at the main office, or in someone's desk or file. Familiar examples include the periodic inspection tags found on fire extinguishers and in elevators.]

b5. Caused Failure:
Was maintenance (inspection) of the work flow process performed without the maintenance (inspection) activity itself causing a failure or degradation of the process?

b6. Time LTA:
Was the time specified in the plan's schedule sufficient to adequately perform the task at each station? Was the time budgeted for personnel adequate to fulfil the schedule? Was the time actually provided?

SD5. Supervision LTA

Was the worksite supervision adequate? Were the necessary supportive services adequate? (297)
[MORT identifies the first line supervisor as a "key man" in worksite safety, as he unquestionably is. However, if the supervisor is to adequately fulfil his responsibilities, he must have competent and useful advice and support from several kinds of supportive services. The adequacy of site supervision is, therefore, examined by MORT in this broader context. In particular, MORT tries to assess management's role in support and service to the supervisor. The emphasis throughout is to discuss what in the management system failed - not who.]

a1. Help and Training LTA:
Were the help and assistance given to supervisors adequate to enable them to fulfil their roles? Was the feedback of information to the supervisor adequate? Was it furnished in a form usable by the supervisor? What training had the supervisor been given in general supervision? What training had the supervisor been given in safety? Has the supervisor training programme been evaluated? (300)

a2. Time LTA:
Did the supervisor have sufficient time to thoroughly examine the job?

a3. Supervisor Transfer Plan LTA:
Were there any gaps or overlaps in the supervisory assignments related to the event? If the supervisor was recently transferred to the job, was there protocol for orderly transfer of safety information from the old to the new supervisor? (303)

a4. D/N Detect/Correct Hazards:
Were the supervisor's efforts adequate in detection and correction of hazards? [Knowledge of hazards is often available from the work force. The supervisor must be receptive and accessible and must display vigour in acting on suggestions, if he is to gain access to that knowledge.]

b1. D/N Detect Hazards:
When did the supervisor last make an inspection of the area? Was any unsafe condition present in this accident/incident also present at the time of inspection? Was the condition detected? (303)
[Note that if the condition was detected but not corrected, the analysis shifts to D/N Correct Hazards.]

c1. Knowledge (Checklists) LTA:
Was a checklist specific to the process available? Was it used? Was the supervisor considered generally competent to assess safety aspects of his area of work?

c2. Detection Plan LTA:
Was there an overall detection plan for uncovering hazardous conditions?

d1. Logs/Schematics LTA:
Was the point-of-operation posting of warnings, emergency procedures, etc., provided for in a general detection plan? Were maintenance and inspection logs at the point of operation adequate? Were work schematics adequate? Were equipment change tags used? (304)

d2. Supervisor Monitor Plan LTA:
What guidance was given to the supervisor relative to inspecting and monitoring status of the process ingredients (i.e., equipment, procedures, and personnel)? Did he use the guidance? Was he given guidance on detection of individual personnel problems, such as alcoholism, drug use, personal problems?

d3. D/N Review Changes:
Was guidance given on review methods and change detection? Were the changes involved known to the supervisor? What counter-changes were made for the known changes?

d4. D/N Relate to Prior Errors:
If there were any known prior errors afflicting the process, was the supervisor told they might correlate with safety errors? Had he made an effort to correlate them? Was he aware of other signs or warnings that the process was moving out of control?

c3. Time:
Did the supervisor have adequate time to detect the hazards?

b2. D/N Correct Hazards:
Was an effort made to correct the detected hazard? (305)
[Some facts about non-correction of hazards were dealt with under non-detection. There are some basic factors of non-correction still to be examined.]

c4. Interdepartmental Co-ordination LTA:
If the accident/incident involved two or more departments, was there sufficient and unambiguous co-ordination of interdepartmental activities?
[Interdepartmental co-ordination is a key responsibility of the first line supervisor. It should not be left to work level personnel.]

c5. Delayed:
Was the decision to delay correction of the hazard assumed by the supervisor on behalf of management? Was the level of risk one the supervisor had authority to assume? Was there precedent for the supervisor assuming this level of risk (as then understood by him)?
[Note a decision to delay correction of the hazard may or may not transfer to the Assumed Risk branch. It was an assumed risk only if it was a specific named event, analysed, calculated where possible, evaluated, and subsequently accepted by the supervisor who was properly exercising management-delegated, decision-making authority.]

d5. Was the decision to delay hazard correction made on the basis of limited authority to stop the process?

d6. Was the decision made because of budget considerations?

d7. Was the decision made because of time considerations?

c6. Programme Housekeeping LTA:
Was the housekeeping of the ongoing programme adequate? Was the storage plan for unused equipment adequate?
[The true role of housekeeping in the accident experience is usually unclear.]

c7. Supervisory Judgement:
Was the judgement exercised by the supervisor to not correct the detected hazard adequate considering the level of risk involved? If there were previously established supervisor authority limitations, were the supervisor's actions generally in accord with those limitations? [Evaluation of the performance of a supervisor in a given situation is, of course, retrospective and must be fairly considered. If the authority limitations of the supervisor have been defined (as they should be), then the adequacy of his performance is more easily measured.]

a5. Performance Errors:
Was the work activity at the worksite free of performance errors by work level personnel? (331)
[The MORT analysis separates performance errors into task, non task, and emergency shutoff errors. Worksite activity can be viewed as usually proceeding in a normal manner to attainment of performance goals. If the ongoing activity enters an abnormal phase requiring work process shutoff, it is described as an emergency, and is analysed in the light of the additional stress associated with emergency action. The analysis proceeds more easily with these considerations. It should be pointed out that the kinds of questions raised by MORT are directed at systemic and procedural problems. The experience, to date, shows there are few "unsafe acts" in the sense of blameworthy work level employee failures. Assignment of "unsafe act" responsibility to a work level employee should not be made unless or until the preventive steps of: (1) hazard analysis, (2) management or supervisory detection, and (3) procedures safety review have been shown to be adequate.]

b3. Task Performance Errors:
Was the task-related work activity free of hazards caused by performance error?

c8. Task Assignment LTA:

Was the task assignment properly scoped with steps and objectives clearly defined? Was the task assignment one the supervisor should have made? (332)

c9. Task Safety Analysis (e.g., JSA) Not Performed:
Was any form of task safety analysis performed as part of the work process? (316)
[Effort directed to task Safety Analysis should be scaled to fit the magnitude of the safety hazard posed by the work task. The safety analysis effort applied to work processes having high energy potential or high hazard potential is usually highly formalised. The analysis results are implemented by a written procedure developed by the task supervisor and a small group of his most skilled craftsmen, and will usually be subjected to independent review. An example of this kind of task safety analysis is Job Safety Analysis (JSA), a process used by many large industrial companies. At the other end of the spectrum of Task Safety Analysis is the informal, oral review of task safety measures by the task supervisor, before work level personnel start to work the task. This latter level of safety analysis is applied to tasks having relatively low energy or low hazard potential. It is used most often with tasks related to routine maintenance and repair activity and will usually not have been independently reviewed.

The task safety analysis level of effort actually applied will range somewhere between the extremes described. MORT uses the concept of Pre-Job Analysis by which is meant that nearly every task must be surveyed step-by-step to determine the level of effort of Task Safety Analysis that should be applied to the work task to be performed. The MORT diagram analysis proceeds with the premise a Pre-Job Analysis should always be made for tasks assessed as having significantly high hazard potential.]

d8. High Potential:
Was an analysis performed for a work task involving a high potential for error, injury, damage, or for encountering an unwanted energy flow?

e1. Pre-Job-Analysis Not Required:
Did the operations management require a pre-job-analysis to scale the magnitude of task safety analysis to be performed?

e2. If required, was the pre-job-analysis, as performed, adequate to scale the magnitude of task safety analysis to be performed?

e3. Pre-Job-Analysis Not Made:
Was a pre-job-analysis required but not made?

f1. Was it not made because of lack of authority?

f2. Was it because of budget reasons?

f3. Was it because of schedules?

f4. Was it because of a decision by the line supervisor?

d9. Low Potential:
Was the work task assessed as one involving low potential? Was this a reasonable assessment? Was the decision to not perform a task safety analysis properly delegated to the supervisor?
[Note the event is flagged with R6 assumed risk symbol. If the criteria for risk identification and assessment were properly met, this event transfers to the Assumed Risk branch.]

c10. Pre-Task Briefing LTA:
Was the work force given a pre-task briefing (prior to task performance)? Was it adequate? Did the pre-task briefing adequately consider the net effect of recent changes, maintenance, new hazards, etc.?

c11. Task Safety Analysis (e.g., JSA) LTA:
Was the task safety analysis adequate? Was the task safety analysis scaled properly for the hazards involved?

d10. Preparation LTA:
Was the preparation (and content) of the task safety analysis adequate?

e4. Selection LTA:
Were the safety hazards associated with the work task adequately identified and selected? (318)

f5. Were the criteria used adequate?

f6. Were the methods used in prioritising the identified hazards adequate?

e5. Knowledge LTA:
Was the knowledge input to the task safety analysis adequate?

f7. Employee Suggestions and Inputs LTA: Was consideration of employee-developed suggestions and inputs adequate?

g1. JSA Programme LTA:
Was a JSA programme used to obtain work level employee participation? Was the process of accomplishing the JSA programme adequately defined and staffed?

g2. General System LTA:
Was the general management system for collecting and utilising other employee suggestions and inputs adequate?

g3. RSO Study LTA:
Were "Reported Significant Observation" (RSO) studies used to gather employee inputs? Were these RSO's readily accessed? (116)

g4. D/N Use Suggestion:
Were employee suggestions and inputs (made through JSA, RSO, and other processes) used in the task safety analysis?

f8. Technical Information LTA:
Was the technical information (with respect to the preparation of the task safety analysis) adequate? [Technical information relevant to safety aspects of the work process often exists but is not available to the "action" persons associated with the process. Possible reasons are investigated by a series of lower tier questions. Analysis of these lower tier events is shown under SD1-a1. Note the analysis transfers to this event also.]

e6. Development LTA:
Was the development of the specific task safety analysis by the first line supervisor adequate? If judged to have been inadequate, what were the true underlying causes for the inadequacy? [An honest assessment should be made of what could reasonably be expected of the supervisor, taking into account existing time and budget restrictions placed upon him by higher supervision.]

f9. Time LTA:
Was there time for an adequate development of task safety analysis?

f10. Budget LTA:
Was there sufficient departmental budget?

f11. Scope LTA:
Were the scope and depth of the task safety analysis development sufficient to cover all related hazards?

f12. Professional Skill LTA:
Were the experience and skill of the supervisor and other participants adequate to accomplish the required work task safety analysis?

d11. Safety Analysis Recommended Controls LTA:
Were adequate worksite controls placed on the work process, facility, equipment, and personnel by the task safety analysis?

e7. Organisation and Clarity LTA:
Were the organisation and clarity of presentation of the task safety analysis recommendations adequate to permit their easy use and understanding?

e8. Programmatic Conflict:
Were the recommended controls free of conflict with the overall project goals and requirements?

e9. Control Testing LTA:
Were recommended controls tested at the worksite for feasibility before being directed for use?

e10. Directive For Use LTA:
Was the management directive for use of the task safety analysis recommended controls adequate? Was it explicit and not subject to possible misunderstanding?

e11. Availability LTA:
Did the management information system make knowledge of the recommended controls available to the worksite personnel?

e12. Adaptability LTA:
Were the recommended safety controls made in a form which allowed them to be adequately adapted to the varying situations?

c12. D/N Use Safety Analysis Recommended Controls:
Were the safety controls recommended by the task safety analysis used?

d12. Use Not Mandatory:
Was use of the recommended safety controls mandatory?
[If use of the recommended safety controls was not mandatory, failure to use them is either an Assumed Risk or a management system failure.]

d13. Deviant Performance:
If use of the recommended safety controls was mandatory, were they actually used?
[If use was mandatory, failure to use them is a deviant performance on the part of the line supervisor.]

c13. Task Procedure D/N Agree With Functional Situation:
Did the work task completion procedure, as directed by oral or written instruction, agree with the actual requirements of the work task?
[Direction or requirements, as defined by specifications, operating procedures, equipment manuals, etc., may conflict with actual work task requirements.]

c14. Personnel Performance Discrepancy:
Did the individuals assigned to the work task perform their individual task assignments properly? [Possible causes of performance discrepancy should be considered for each individual whose performance was judged to be discrepant.]

d14. Personnel Selection LTA:
Were the methods of personnel selection adequate? (325)

e13. Criteria LTA:
Were the safety-related job requirements adequately defined so as to select an individual with desired characteristics?

e14. Testing LTA:
Did the individual meet the standards established for the task? Had the assigned individual been recently re-examined to the standards established for the task?

d15. Training LTA:
Was the training of personnel adequate? (327)

e15. None:
Was the individual trained for the task he or she performed?

e16. Criteria LTA:
Were the criteria used to establish the training programme adequate in scope, depth, and detail?

e17. Methods LTA:
Were the methods used in training adequate to the training requirements? [Consider methods such as realistic simulation, programmed self-instruction, and other special training in addition to basic indoctrination, plant familiarisation, etc.]

e18. Professional Skills LTA:
Was the basic professional skill of the trainers adequate to implement the prescribed training program?

e19. Verification LTA:
Was the verification of the person's current trained status adequate? Were retraining and re-qualification requirements of the task defined and enforced?

d16. Consideration of Deviations LTA:
Was adequate consideration shown by the supervisor for the need to observe deviant personnel performance? (334) [The analysis shows contributions to Deviations from both Normal Variability and Changes. Normal personnel performance variability is viewed as manageable through appropriate equipment design, good planning, training, and application of human factors. Change is more the characteristics of illness, fatigue, personal problems, etc., which results in individual performance outside the normal range of variability.]

e20. Normal Variability:
Was the deviation in personnel performance within the range of normal variability?
[The Scroll event symbol is used to show that some degree of variability is normal and expected.]

e21. Changes:
Was the deviation in personnel performance significantly different than the performance standard needed for the task?
[The Scroll event symbol is used to show that some degree of change is normally expected to occur.]

e22. D/N Observe:
Was the deviation (i.e., extreme variability or significant change) observed by the line supervisor?

e23. D/N Correct:
Did the supervisor act to correct the observed personnel performance deviations?

f13. D/N Re-instruct:
Did the supervisor re-instruct the person observed as to the correct performance?

f14. D/N Enforce:
Did the supervisor enforce established correct rules and procedures? Were disciplinary measures taken against personnel who wilfully and habitually disregarded rules and procedures?

d17. Employee Motivation LTA:
Were the employee motivation, participation, and acceptance adequate? (337)
[Employee motivation plays a significant role in personnel performance in accomplishment of the work task. Various aspects of employee motivation are analysed by lower tier events.]

e24. Management Concern, Vigour, and Example LTA:
Was management concern for safety displayed by direct vigorous personal action on the part of top executives? (200)

e25. Schedule Pressure:
Were task schedule pressures (as experienced by the individual) held to an acceptable level?

e26. Performance Is Punishing:
Was the employee fairly treated by management for performing as supervision desired?
[From the viewpoint of the employee, sometimes there is an undesirable consequence to the person doing a good job.]

e27. Non-Performance Is Rewarding:
Did the employee find the consequence of doing the job incorrectly more favourable than doing the job as directed?
[Obstructive behaviour may be more rewarding to the individual than facilitating behaviour.]

e28. Job Interest Building LTA:
Does performing the task well <u>really</u> matter to the individual performing it?
[Perhaps the performing individual believes the consequence is the same to him whether he does the task right or some other way. Good performance should be followed at least periodically by an event considered favourable by the individual.]

e29. Group Norms Conflict:
Are the actions and attitudes of the individual's peer groups in harmony with the task requirements and the goals of the larger organisation?

f15. Worker Participation LTA:
Was there adequate opportunity for the worker to participate in analysis, training, or monitoring systems (e.g., JSA and RSO studies)?

f16. Innovation Diffusion LTA:
Was there adequate use of management motivational programmes to develop desired behavioural change in individuals (i.e., application of innovation diffusion techniques)? [Appendix H of the MORT text.]

e30. Obstacles Prevent Performance:
Were obstacles that might prevent task performance reduced to an acceptable level? [Often a task would get done more efficiently if conditions were changed. If performance discrepancies appear not to be because of lack of skill or motivation, one thing to look for is an obstacle.]

e31. Personal Conflict:
Are individual personal conflicts, which may have a negative relationship to task safety, adequately resolved in the individual? Does the individual have good standards of judgement?

f17. W/Supervisor:
Were employee and supervisor personalities compatible in the work environment?

f18. With Others:
Was the employee's personality compatible with other workers in the work environment?

f19. Deviant:
Were the psychological traits exhibited by the individual judged acceptable when rated against the task safety requirements? Individuals exhibiting abnormally high levels of social maladjustment, emotional instability, and conflict with authority produce more than their share of accidents. The decision to employ an individual in a given task ultimately rests with the line supervisor.

If the organisation has maximised its contribution in the areas of management concern, safeguarding environment, good job safety procedures, good job training, sound human relations, etc., the use of an individual with known deviant performance characteristics in a high potential energy task becomes an assumed risk. [Note the event is flagged with R7 assumed risk symbol. If the criteria for risk identification and assessment were properly met, this event transfers to the assumed risk branch.]

e32. General Motivation Programme LTA:
Was there a general motivation programme on safety, employed by management, to adequately motivate employees to perform correctly and safely? [Slogans, posters, leaflets, and contests are a highly visible part of many safety programs. Their true value is difficult to ascertain. These programmes do play a supporting role, however and the adequacy of the safety programme in these regards should be evaluated.]

b4. Non-Task Performance Errors:
Was the performance of non-task work free of performance errors? [A "non-task" is one not assigned by a supervisor.]

c15. Peripheral:
Was the work peripheral to the principle task performed error-free?
[Examples are going to or from work on the premises, authorised work break, etc. The activity was not in conflict with the rules.]

c16. Unrelated:
Were all activities unrelated to the authorised work activity performed error-free?
[Examples are going to lunch, recreational programs.]

c17. Prohibited:
Were all performed activities permitted? If not, were the prohibited activities performed error-free?
[Activity in violation of rules, horseplay, etc., is defined as prohibited activity.]

b5. Emergency Shutoff Errors:
If there was an emergency shutdown of some activity from its normal operating mode, was it done error-free?
[Emergency situations usually are a time of rapid change and high stress. The emergency may evolve from a planned task (an in-process work activity) or from a non-task activity. Note the use of the Constraint event symbol requires an off-normal initiating anomaly to have occurred.]

SD6

c18. Task Performance Errors:
Was there an emergency shutoff? Was the execution of a planned shutdown sequence accomplished error-free?
[The entire MORT analysis accomplished under SD5-b3 transfers into this event.]

c19. Non-Task Performance Errors:
If there was an emergency situation arising with a non-task activity (i.e., one not assigned by a supervisor), was it free of performance errors?
[See the classification and explanation of non-task performance errors provided under SD5-b4.]

SD6. Higher Supervision Services LTA

Did upper level management provide the type of supportive services and guidance needed at lower organisation levels for adequate control of unwanted work process energy flow?

a1. Research and Fact Finding LTA:
Was necessary information, which was not otherwise readily available, sought out through established research and fact finding techniques?

a2. Information Exchange LTA:
Was there an accessible, open line of communications which permitted transmittal of needed information in both directions between upper and lower levels? Was study of a problem a shared responsibility? Were results provided to users?

a3. Standards and Directives LTA:
In cases where the organisation and external sources of codes, standards, and regulations did not cover a particular situation, did management develop (or have developed) adequate standards and issue appropriate directives?

a4. Resources LTA:
Did management have the resources derived from standards and directives it needed to perform the supportive services?

b1. Training LTA:
Was there sufficient training to update and improve needed supervisory skills?

b2. Technical Assistance LTA:
Did supervisors have their own technical staff or access to such individuals? Was technical support of the right discipline(s) sufficient for the needs of supervisory programmes and review functions?

b3. Programme Aids LTA:
Did management have available, for support of its programs, such aids as: useful analytic forms, training materials, reproduction services, audio-visuals, capable speakers, meeting time and rooms, technical information, monographs, etc.?

b4. Measure of Performance LTA:
Were there established methods for measuring performance which permitted the effectiveness of supervisory programmes to be evaluated?

b5. Co-ordination LTA:
Were other management programmes and activities co-ordinated with the groups and individuals who interfaced with the programme participants? Did this co-ordination eliminate conflicts which could have reduced programme effectiveness?

a5. Deployment of Resources LTA:
Were the available resources used effectively and to the greatest advantage of supervisory efforts?

a6. Referred Risk Response LTA:
Was management responsive to risks referred from lower levels? Was there an established system for analysing and acting upon such risks in a timely manner? Was there a fast action cycle to process imminent hazard/high risks?

"Higher Supervision Services" serves as a mechanism to transfer management intent to field activities. In addition to the specific questions dealt with above, a number of questions should also be asked relative to the more general effects of management on employee performance. These include five questions for example:*

- To what degree do such considerations as pay and benefits, working conditions, job security, rewards, and recognition combine to produce adequate worker job satisfaction?
- To what degree do work ethics and practices, concern for safety, concern for quality team orientation combine to produce adequate worker motivation?
- To what degree do loyalty, sense of ownership, pride, respect for radiation, morale, sense of accountability result in adequate work attitudes?
- To what degree do management activities in controlling person-machine interfaces and person-environment interfaces satisfy human factors requirements in an adequate manner?
- To what degree do management activities relating to fitness for duty, testing programs, and discipline result in adequate human reliability?

*Derived from S. B. Haber and D. P. Thurmond, "A Preliminary Identification of the Human Performance Issues Within the Department of Energy," Task 1 Report prepared for the U.S. Department of Energy (1989) and applied in Human Performance Appraisal in June 1990.

M
MA1
MA2

M. Management System Factors LTA

Are all the factors of the management system necessary, sufficient, and organised in such a manner as to assure that the overall programme will be "as advertised" to the customer, to the public, to the organisation itself, and to other groups as appropriate?
[In the event-by-event review which follows, the questions are phrased in the present tense. Assume the diagram is being used for evaluation of an existing safety system. For accident investigation, rephrase questions to past tense.]

MA1. Policy LTA

Is there a written, up-to-date policy with a broad enough scope to address major problems likely to be encountered? Is it also sufficiently comprehensive to include the major motivations (e.g., humane, cost, efficiency, legal compliance)? Can it be implemented without conflict? (175, 183)

MA2. Implementation LTA

Does the overall programme represent the intended fulfilment of the policy statement? If there are problems encountered in carrying out the policy, are these relayed back to the policy makers? Is the implementation a continuous, balanced effort designed to correct systemic failures, and generally proactive rather than reactive? (185)

a1. Methods, Criteria, Analyses LTA:
Are selective methods used for management implementation and for improving human performance? Is there a comprehensive set of criteria used for assessing the short and long-term impact of the methods on safety for the desired results? Does management demand that adequate analyses be performed and alternative countermeasures examined, or are criteria simplistic and therefore LTA? (185)

a2. Line Responsibility LTA:
Is there a clear, written statement of safety responsibility of the line organisation, from the top individual through the first line foreman to the individual employee? Is this statement distributed and understood throughout the organisation? Is it implemented? (190)

a3. Staff Responsibility LTA:
Are there provisions for assigning and implementing specific safety functions to staff departments (e.g., safety, personnel and training, engineering, maintenance, purchasing, transportation, etc.)?

a4. Information Flow LTA:
Has management specified the types of information it needs and established efficient methods by which such information is to be transmitted up through the organisation? Has management, in turn, supported this process by providing the information needed in lower organisation levels? (198)

a5. Directives LTA:
Is safety policy implemented by directives which emphasise methods and functions of hazard review, monitoring, etc., rather than specific rules for kinds of hazards? Are directives published in a style conducive to understanding and without interface gaps? (193)

a6. Management Services LTA:
Has management provided the type of supportive services and guidance needed at the lower organisation levels? Is there a formal training programme for all management personnel which addresses: (1) general aspects of management and supervision, (2) specific technologies, (3) human relations/communications, and (4) safety? (195)
[Note the transfer in of all the lower tier event analysis from SD6.]

As indicated earlier (SD6), Higher Supervision Services is a mechanism by which management intent as described on the "M" portion of the MORT is transferred to specific activities dealt with on the "S" side of the MORT. In addition to explicit analysis of the management functions themselves, one should also ask more general questions relating to management mechanisms for transfer of management intent to the specific organisational activities.*

- To what degree do guidance and direction, management attention, manager presence, promotion of goals and unity, and appropriate recognition of constituencies/customers combine to provide adequate leadership?

- To what degree do strategic/long-range planning, work planning, programme co-ordination and resource allocation lead to adequate overall planning and scheduling?

- To what degree do information collection and processing, upward reporting systems and information dissemination lead to adequate information management?

- To what degree do work processes, work documentation, quality control, authority and use of procedures combine to result in adequate work controls and practices?

- To what degree do problem identification/reporting processes, root cause analysis, corrective action processes, use of lessons learned combine to result in solution of field level problems?

*Derived from S. B. Haber and D. P. Thurmond, "A Preliminary Identification of the Human Performance Issues Within the Department of Energy," Task 1 Report prepared for the U.S. Department of Energy (1989) and applied in Human Performance Appraisal in June 1990.

- To what degree do problem assessment/trending, self appraisal, independent oversight, independent performance evaluation and performance improvement programmes combine to provide an adequate performance assessment program?

- To what degree do safety and environment policy and program, safety and environment requirements definition, and compliance monitoring and reporting result in transfer of environment and safety considerations to field activities in an adequate manner?

- To what degree do quality policy and program, quality requirements definition, and compliance monitoring and reporting result in transfer of quality considerations to specific activities in an adequate manner?

- To what degree do levels of decision making, assignments of responsibility and accountability, and influence of safety concerns result in effective risk based decision making related to specific activities?

a7. Budgets LTA:
Is the budget adequate not only for the safety group but also for related safety programme aspects for which other groups in the organisation have responsibility? (189)

a8. Delays:
Are safety programme elements implemented in a timely manner? Are solutions to safety problems introduced early in the life cycle phases of projects? (189)
[Delays can and should be made known to management. If this is done and delay is a practical need, the delay becomes an assumed risk.]

a9. Accountability LTA:
Is line management held accountable for safety functions under their jurisdiction? If so, are there methods for measuring their performance? (198)

a10. Vigour and Example LTA:
Have top management individuals demonstrated an interest in lower level programme activities through personal involvement? Is their concern known, respected, and reflected at all management and employee levels? (200)
[Do people tell stories of a manager's vigour in support of safety? If not, the manager's example may be LTA.]

MA3. Risk Assessment System LTA

Does the risk assessment system provide management with the information it needs to assess residual risk and to take appropriate action if the residual risk is found unacceptable? Does the system also provide: (1) comparative evaluation of two or more systems; and (2) development and evaluation of methods supporting the hazard analysis process? (205)

MB1
MB2
MB3

MB1. Goals LTA

Are there high goals for policy and implementation criteria as well as specific goals for projects? Are the goals non-conflicting, sufficiently challenging, and consistent with policy and the customer's goals? (206)

MB2. Technical Information System LTA

Is the technical information system adequate to support the needs of the risk assessment system?
[Note other lower tier events included by transfer from SD1. Refer to SD1 section of this outline for write-up.]

MB3. Hazard Analysis Process LTA

Is the Hazard Analysis Process (HAP) properly conceptualised, defined, and executed? (225, 215, 234)

a1. Concepts and Requirements LTA:
Are the concepts and requirements of the HAP adequately defined? (237)

b1. Definition of Goals and Tolerable Risks LTA:
Have goals and tolerable risks been defined for both safety and performance and any conflicts between the two resolved? (237)

c1. Safety Goals and Risks Not Defined:
Do the goals state what degree of safety excellence should be attained and when? Are tolerable direct and indirect safety risks defined and actual risks quantified?

c2. Performance Goals and Risks Not Defined:
Have goals been set for performance efficiency and productivity? Have tolerable risks for lost efficiency and productivity been established and actual risks quantified?
[Such goals complement safety goals by requiring greater assurance of error-free performance.]

b2. Safety Analysis Criteria LTA:
Have the necessary criteria been specified and elements defined to adequately support the safety analysis program?

c3. Plan LTA:
Has a system safety plan been developed that describes "who does what and when" in analysis, study, and development? (238)

c4. Change Analysis LTA:
Has a specific change-based analytic method been established to review form, fit, or function of components and subsystems (including interfaces) upwards in a review process until no change is demonstrated? (59)

c5. Other Analytical Methods LTA:
Are other appropriate analytical skills available in the organisation (or from a consultant) and are they used (e.g., Hazard Identification, Failure Modes and Effects, Fault Tree, MORT, Nertney Wheel, Failure Analysis, Human Factors Review, etc.)? (223, 228, 248)

c6. Scaling Mechanism LTA:
Has some reasonably clear-cut mechanism been established for scaling the seriousness/severity of prior events. Is there a mechanism to project past events to a scaled effort to evaluate current processes? (238)

c7. Required Alternatives LTA:
Does management require confrontation between alternative solutions in its bases for choices and decisions? (186, 208)

c8. Safety Precedence Sequence LTA:
Is the preference for safety solutions prioritised as: (1) Design, (2) Safety Devices, (3) Warning Devices, (4) Human Factors Review, (5) Procedures, 6) Personnel, and (7) Acceptance of Residual Risks (after considering the preceding six items)? (98, 225)

b3. Procedures Criteria LTA:
Are engineers and designers made aware of their limitations in writing procedures for operating personnel, for the need for selection and training criteria for operators, and of supervisory problems? (315, Appendix F)

b4. Specification of Safety Requirements LTA:
Have all applicable and appropriate safety requirements been specified, made available, and used? (260)
Consider whether the following documents have been adequately called out to the extent they are applicable:

c9. DOE (customer) requirements developed in-house.

c10. OSHA regulations that are law.

c11. Other Federal and National Codes by agencies other than the customer and OSHA.

c12. State and Local Codes applicable to the geographical area where the work is to be performed.

c13. Internal Standards developed within the organisation to cover situations not addressed by outside requirements.

b5. Information Search LTA:
Is an adequate information search required? (262)

c14. Nature of Search LTA:
Does the nature of the search include incident files; codes, standards, and regulations; change and counter-change data; related previous analyses; and other comments and suggestions?

c15. Scope of Search LTA:
Is the search scoped in a manner that would seek information on problems from conceptual design, through construction and use, to final disposal?

b6. Life Cycle Analysis LTA:
Is there an adequate safety analysis which starts with planning and continues through design, purchasing, fabrication, construction, operation, maintenance, and disposal? (263, 225)

c16. Scope LTA:
Does the scope include not only the prime mission equipment, but also checkout and test equipment and procedures, facilities and operations, procedures for operation, selection of personnel, training equipment and procedures, maintenance facilities, equipment and procedures, and support equipment?

c17. Analysis of Environmental Impact LTA:
Is the life cycle analysis scoped to include an analysis of environmental impact which complies with all applicable requirements? (259)

c18. Requirement for Life Cycle Analysis LTA:
Is the requirement for Life Cycle Analyses (LCA) rigid enough to assure that a thorough LCA will be initiated during the planning stage?

c19. Extended Use Factors LTA:
Has sufficient consideration been given to special requirements, new problems, and other factors to be encountered if the facility/operation is extended beyond its original intended life?

a2. Design and Development Plan LTA:
Does the development phase provide for the use of the major safety results of the Concepts and Requirements Phase (MB3-a1)? Is the design a true representation of the developed criteria, definitions, specifications, and requirements? (267)
[Note that barriers and amelioration, analysed separately in accident investigation, are part of the design process.]

b7. Energy Control Procedures LTA:
Is there an attempt, whether by design or procedure, to control energy to only that which is needed for the operation and to contain its interactions to the intended function?

c20. D/N Substitute Safer Energy:
Does the design use the safest form of energy that will perform the desired function?

c21. D/N Limit Energy:
Is the amount of available energy limited to that which will perform the operation without any unnecessary excess energy?

c22. Automatic Controls LTA:
Are there devices to automatically control the flow of energy and to maintain it in its operating mode? Is use of redundant design adequately employed? (267)

c23. Warnings LTA:
Are there clear, concise warnings for all situations where persons or objects might unintentionally interface with an energy flow? (268)

c24. Manual Controls LTA:
Are there manually-operated controls to maintain the proper energy flow during the normal mode or as a manual override of automatic controls?

c25. Safe Energy Release LTA:
In the event that the energy containment fails through normal flow channels, is there a designed-in route through which the energy can be safely released?

c26. Barriers and Controls LTA:
Are there adequate barriers included as part of the design, plan, or procedure? Do they separate energies and/or protect people and objects? (33, 268)
[Note other lower tier events included by transfer from SB2.]

b8. Human Factors Review LTA:
Has consideration been given in design, plan, and procedures to human characteristics as they compete and interface with machine and environmental characteristics? (273)

c27. Professional Skills LTA:
Is the minimum level of human factors capability, needed for evaluation of an operation, available and will it be used? (275)

c28. D/N Describe Tasks:
For each step of a task, is the operator told: When to act? What to do? When the step is finished? What to do next? (276)

c29. Allocation Man-Machine Tasks LTA:
Has a determination been made (and applied) of tasks that humans excel in versus those tasks at which machines excel?

c30. D/N Establish Man-Task Requirements:
Does the review determine special characteristics or capabilities required of operators and machines?

d1. D/N Define Users:
Is available knowledge about would-be users defined and incorporated in design?

d2. Use of Stereotypes LTA:
Are checklists of stereotypes (typical, normal, expected behaviour) used in design? (e.g., Is a control turned right to move a device to the right?) Are controls coded by size, colour, or shape?

d3. Displays LTA:
Are displays used which can be interpreted in short time with high reliability?

d4. Mediation LTA:
Is consideration given to delays and reliability of interpretation/action cycles?

d5. Controls LTA:
Are controls used which can be operated in short times with high reliability?

c31. D/N Predict Errors:
Is there an attempt made to predict all the ways and frequencies with which human errors may occur, and thereby determine corrective action to reduce the overall error rate?

d6. Incorrect Act:
Have all the potential incorrect acts associated with a task been considered and appropriate changes made?

d7. Act Out of Sequence:
Has the consequence of performing steps of a task in the wrong order been considered and has appropriate corrective measures been made?

d8. F/T Act:
Is there an attempt to reduce the likelihood of operators omitting steps or acts which are required by procedure?

d9. Act Not Required:
Are all the steps that are needed to accomplish a task required in the procedures? Are only those steps in the procedure?

d10. Malevolence:
Are deliberate errors and other acts of malevolence anticipated and steps taken to prevent them or reduce their effect?

b9. Maintenance Plan LTA:
Is maintenance of an operation/facility given consideration during the conceptual phase and on through the rest of the life cycle? Is there an adequate maintenance plan? (311)
[Note other lower tier events included by transfer from SD3-a1.]

b10. Inspection Plan LTA:
Is inspection of an operation/facility given consideration during the conceptual phase and on through the rest of the life cycle? Is there an adequate inspection plan? (312)
[Note other lower tier events included by transfer from SD4-a1.]

b11. Arrangement LTA:
Does the design consider problems associated with space, proximity, crowding, convenience, order, freedom from interruption, enclosures, work flow, storage, etc.?

b12. Environment LTA:
Are people and objects free from physical stresses caused by: (1) facility physical conditions, (2) conditions generated by the operation, or (3) interactions of one operation with another?

b13. Operational Specifications LTA:
Are there adequate operational specifications for all phases of the system operation? (269)

c32. Test and Qualification LTA:
Is there a "dry run" or demonstration to prove out all associated hardware and procedures and to check for oversights, adjust for the final arrangement, and provide for some first "hands-on" participation?

c33. Supervision LTA:
Are there guidelines for the amount of supervision required, minimum supervisory capabilities needed, and responsibilities of operating supervisors?

c34. Task Procedures D/N Meet Criteria:
Do the procedures for each task meet selection and training criteria and applicable operating criteria? Are the procedures responsive to supervisory problems that can be addressed in written procedures?
(315, 269, Appendix F)

d11. D/N Fit With Hardware Change:
Are procedures revised, if necessary, to agree with changes in plant or equipment?

d12. Clarity and Adequacy LTA:
Does the writing style of the procedures give consideration to variations in reading skills and intelligence of intended users? Are procedures sufficiently scoped to cover all steps of a task and is enough information given about each step?

d13. D/N Verify Accuracy:
Are procedures rechecked with applicable criteria and tested for correctness under "dry run" operating conditions?

d14. Emergency Provisions LTA:
Do procedures give users clear instructions for all anticipated emergency conditions? Are instructions easy to perform under the stress of an emergency?

d15. Cautions and Warnings LTA:
Are dynamic and static warnings used when appropriate? Are they located at point-of-operation as well as in procedures? Is their meaning unambiguous?

d16. Event Sequence LTA:
Do procedures have steps performed in a sequence: (1) according to criteria; (2) that is safe; and (3) that is sufficient?

d17. Lockouts LTA:
Are lockouts called for where hazardous situations are encountered or created through use of procedures?

d18. Communication Interfaces LTA:
Do the procedures adequately convey their intended message? If procedures call for contact between users and other individuals, are these interfaces clear?

d19. D/N Specify Personnel Environment:
Do procedures specify maximum permissible levels of physical stresses imposed on the users?

c35. Personnel Selection LTA:
Are personnel selected on the basis of the capability (both physical and mental) which is necessary and sufficient to perform the operation? (327)

c36. Personnel Training and Qualification LTA:
Are personnel given all the training they need for the equipment and procedures they will be using? Do they demonstrate through "hands-on" use that they know how to apply the training properly? (327)
[Personnel training and qualification factors are considered in detail under SD5-d15.]

c37. Personnel Motivation LTA:
Do personnel want to perform their assigned work task operations correctly? (337)
[Personnel motivation factors are considered in detail under SD5-d17.]

c38. Monitor Points LTA:
Are there sufficient checkpoints in written procedures during an operation to assure that steps are performed correctly? (351)

b14. Emergency Provisions LTA:
Does the design of plant and equipment provide for safe shutdown and safety of persons and objects during all anticipated emergencies? (306)

b15. Disposal Plan LTA:
Is the design such that disposal problems and hazards are minimised when the facility or operation has served its useful life? (270)

b16. Independent Review Method and Content LTA:
Is provision made for thorough and independent safety review at pre-established points (e.g., milestones) in the life cycle process? Are the risk-reduction trade-offs documented? Is the technical competence of Review Board members properly scaled to the level of technology involved? (283, Tree of Exhibit 8)

b17. Configuration Control LTA:
Is there a formal programme to assure adequate configuration control throughout the entire life cycle of the facility? Does the programme allow for easy access for review of modified procedures, drawings, and other documentation? (270, Tree of Exhibit 3)

b18. Documentation LTA:
Are all types of documentation complete, up-to-date, and accessible to users? (271)

b19. Fast Action Expedient Cycle LTA:
Is there an existing method to bypass the usual delays in order to get an immediate correction for an imminent hazard or problem of significant consequences?

b20. General Design Process LTA:
Are commonly recognised good engineering practices, including safety, reliability, and quality engineering practices, adequately incorporated into the general design process? (281)

c39. Code Compliance Procedures LTA:
Are there written procedures to assure compliance with applicable engineering and design codes?

c40. Engineering Studies LTA:
Where codes, standards, regulations, and state-of-the-art knowledge cannot furnish required design data, are engineering studies conducted to obtain the needed information?

c41. Standardisation of Parts LTA:
Is there an attempt to use proven existing standardised parts where possible, or to design so as to encourage their use?

c42. Design Description LTA:
Does the design description provide all the information needed by its users in a clear and concise manner?

c43. Acceptance Criteria LTA:
Are acceptance criteria stringent enough to assure operability/maintainability and compliance with original design?

c44. Development and Qualification Testing LTA:
Is there adequate testing during development of a new design to demonstrate that it will serve its intended function? Does qualification testing assure that non-standard components satisfy the acceptance criteria?

c45. Change Review Procedure LTA:
Does change review cover form, fit, and function on up the part-component-subsystem chain to a point where no change is demonstrated? Are there change dockets on drawings and at points-of-operation?

c46. Reliability and Quality Assurance (R&QA) LTA:
Is there an adequate reliability and quality assurance programme integrated into the general design process? (282)
[In some organisations, the reliability and quality assurance functions are very specifically separated; other organisations combine them. Whether combined or separated, R&QA is a strong complement to safety. Close mutual support between safety and R&QA should be evident throughout the general design process.]

MB4. Programme Review LTA

Do the Environment, Safety and Quality (ESQ) programme reviews assure a planned and measured program, with low cost/high volume services, professional growth, and use of modern methods? (445)

a1. Definition of Ideals and Policy LTA:
Are there adequate ESQ policy statements and are the ideals of the SEQ programmes articulated? Do these summarise what management should know (and require) the ESQ process? Do the ideals provide a base that measures the programme and projects improvements?

a2. Description and Schematics LTA:
Are programme ideals documented in operating manuals and schematics? Are programme operating data available and evaluated? Are there outlines, steps, and criteria that substantially describe the ESQ programs?

a3. Monitoring, Audit, and Comparison LTA:
Is there a formal measurement system that compares actual performance with ESQ programme ideals and objectives? (446)

a4. ESQ Programme Organisation LTA:
Are the programmes organised with the necessary and adequate elements? (449)

b1. Professional Staff LTA:
Do ESQ personnel rate well by both ESQ and management criteria? Are they effective in both technical and behavioural aspects? Do they have good organisational status and are they educated, experienced, and promotable? (454)

b2. Management Peer Committees LTA:
Are special-purpose and ongoing committees and boards used to improve ESQ understanding and attitudes within scientific and engineering groups? Do these ongoing groups have a positive, action orientation toward real-life problems?

b3. Scope LTA:
Does the ESQ programme scope address all forms of hazards, including anticipated hazards associated with advanced technological development and research?

b4. Integration LTA:
Is the staff support for ESQ integrated in one major unit rather than scattered in several places?

R

b5. Organisation for Improvement LTA:
Is the ESQ programme organised adequately to achieve the desired pace of ESQ improvement? (119, 209, 457) [Achievement of a breakthrough goal in accident reduction by an ESQ programme requires clear goal definition and distinctive organisational effort, particularly by staff personnel.]

a5. Block Function and Work Schematics LTA:
Are charts and drawings of the full array of ESQ-related processes and functions adequate and reviewable?
[This may include provision of ESQ equipment, delivery of other safety reviews to point of need, and other safety-related functions, plus the schematics of various "upstream processes" to be audited or monitored.]

b6. Not Up-To-Date:
Are charts and drawings kept up-to-date?

b7. Incomplete:
Are all items that are needed for review included in the charts and drawings?

b8. Completion Criteria LTA:
Are criteria clear and specific as to what should be included in drawings and when they should be finished and revised?

a6. ESQ Programme Services LTA:
Does management provide the supportive services and guidance needed at the lower organisational levels for an adequate ESQ programme review?
[Note the transfer in of all lower tier event analysis from SD6.]

R. <u>Assumed Risk</u>

What are the assumed risks? Are they specific, named events? Are they analysed and, where possible, calculated (quantified)? Was there a specific decision to assume each risk? Was it made by a person who had management delegated authority to assume the risk?
[The specific risk may be: (1) tolerably low (minor) in frequency or consequence, (2) high in consequence but impossible to eliminate, (i.e., hurricane), or (3) simply too expensive to correct when weighed against the risk consequences. The assumed risk events are shown elsewhere on the MORT diagram and flagged with a numbered "R" symbol.]

MORT CUTSET: **S - Branch (SA To SC Level)**

The disadvantage of cutsets is that it is easy to become confused with multiple pages. It is therefore essential to work in a disciplined and systematic way, carefully labelling the sheets used with the information requested in the spaces at the top of each page. Where possible use one normal MORT chart per barrier/control analysis and use the cutsets for branch repeats (for example, a given barrier or control failure may require many repeats of the b3-SD5 sub-tree when operation, maintenance and inspection failures have occurred).

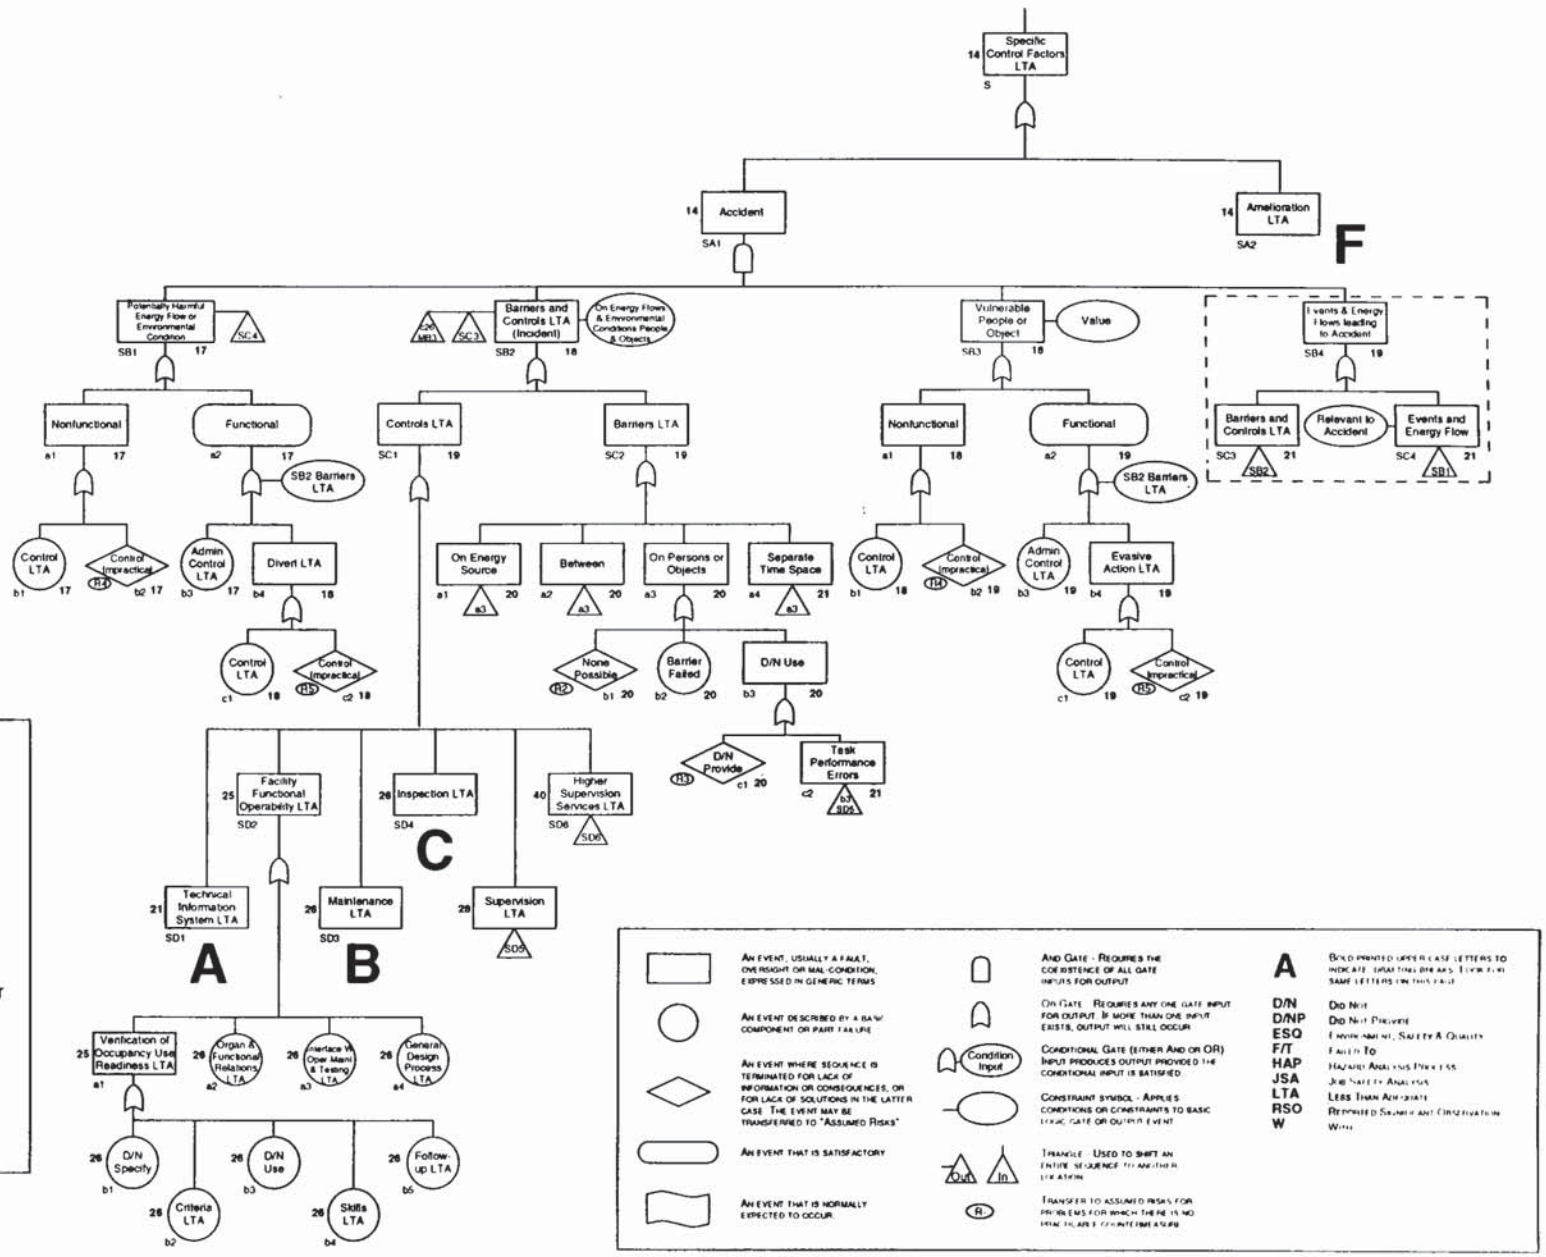

MORT CUTSET: **"A"(SD1 & MB2) Technical Information System**

The "A" branch plugs into two places in the MORT chart. The SD1 "A" sequence focusses on the provision of information to, and gathering of information from, the operational interface containing the barrier in question. The MB2 "A" sequence considers the provision of information into the formal risk assessment cycle.

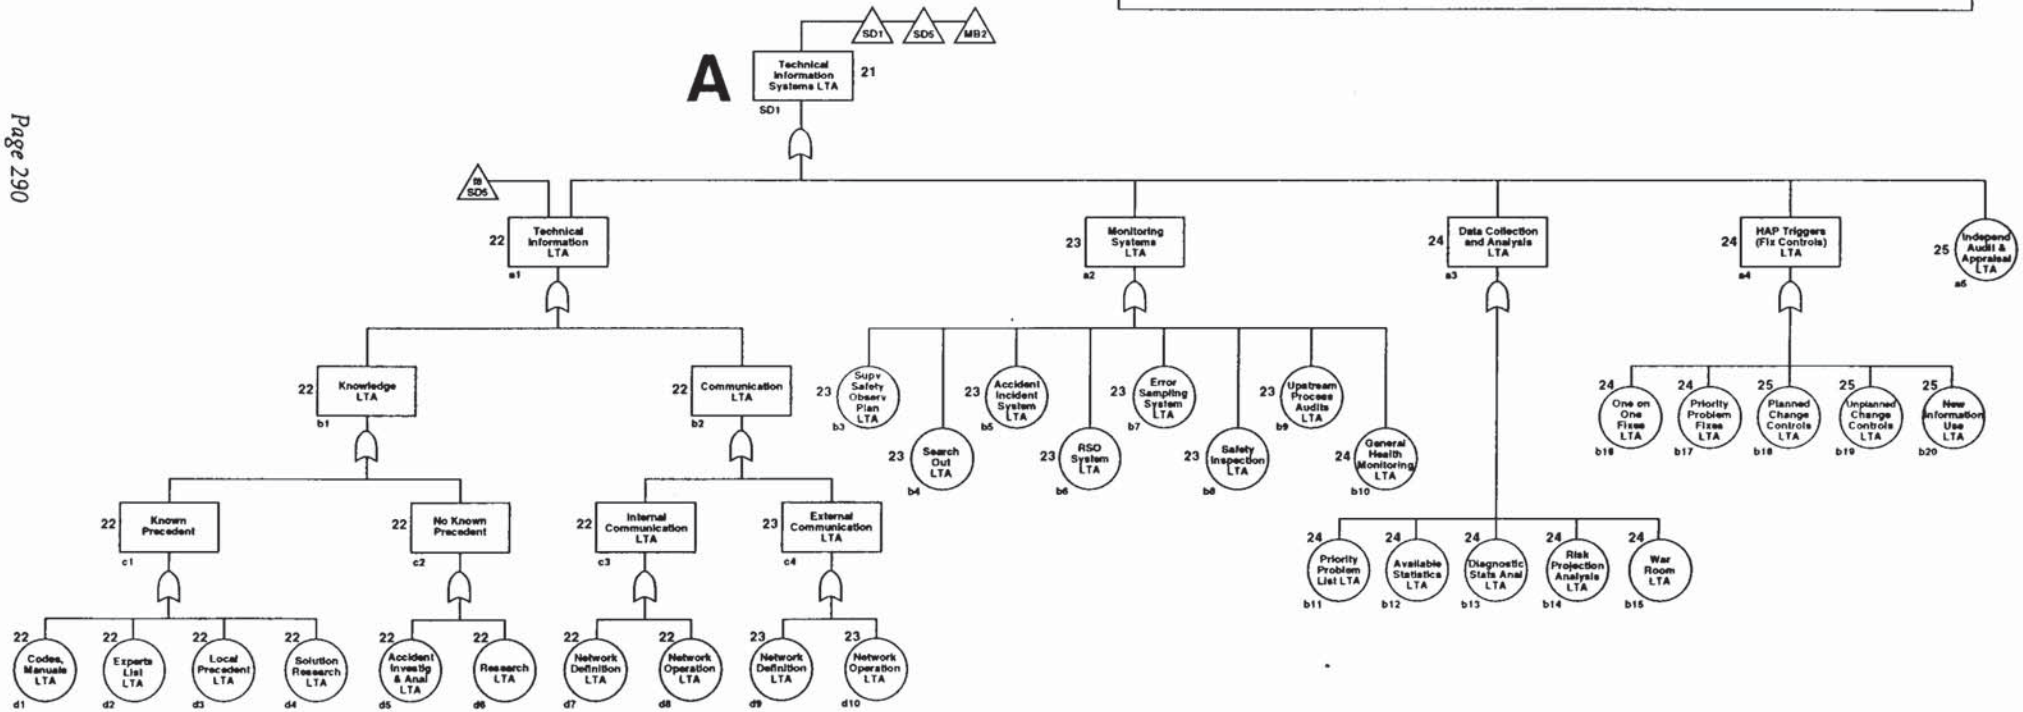

MORT CUTSETS: **"B"& "C"** (**SD3, SD4. *Also* MB3-b9 & b10) Maintenance and Inspection.**

These branches plug into the events underlying SC1 "Control". They are also partially used via referrals from the "MB3" management side of the tree

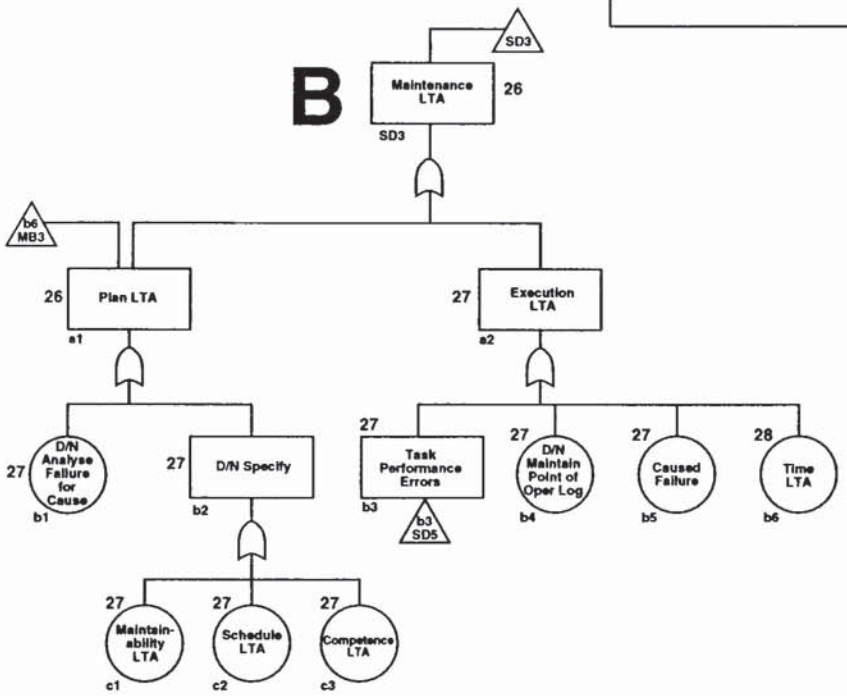

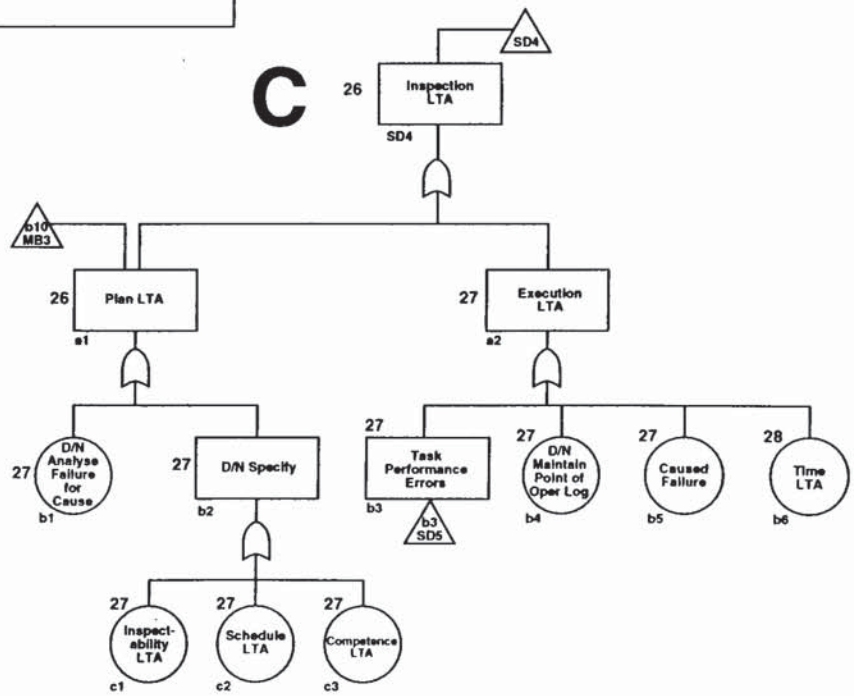

MORT CUTSET: **SD5**

The full MORT tree will guide the process which will (1) consider the whole of SD5 from the perspective of SC1: the assurance the barrier in question; (2) The lower events of SC1 (such as Inspection) refer the analyst to consider only the elements underlying b3 "Task Performance Errors"; and (3) if the problem was that the barrier failed because it wasn't used, the analyst considers the actions of the person who failed to use it.

lastly there are 3 referrals from the Amelioration branch

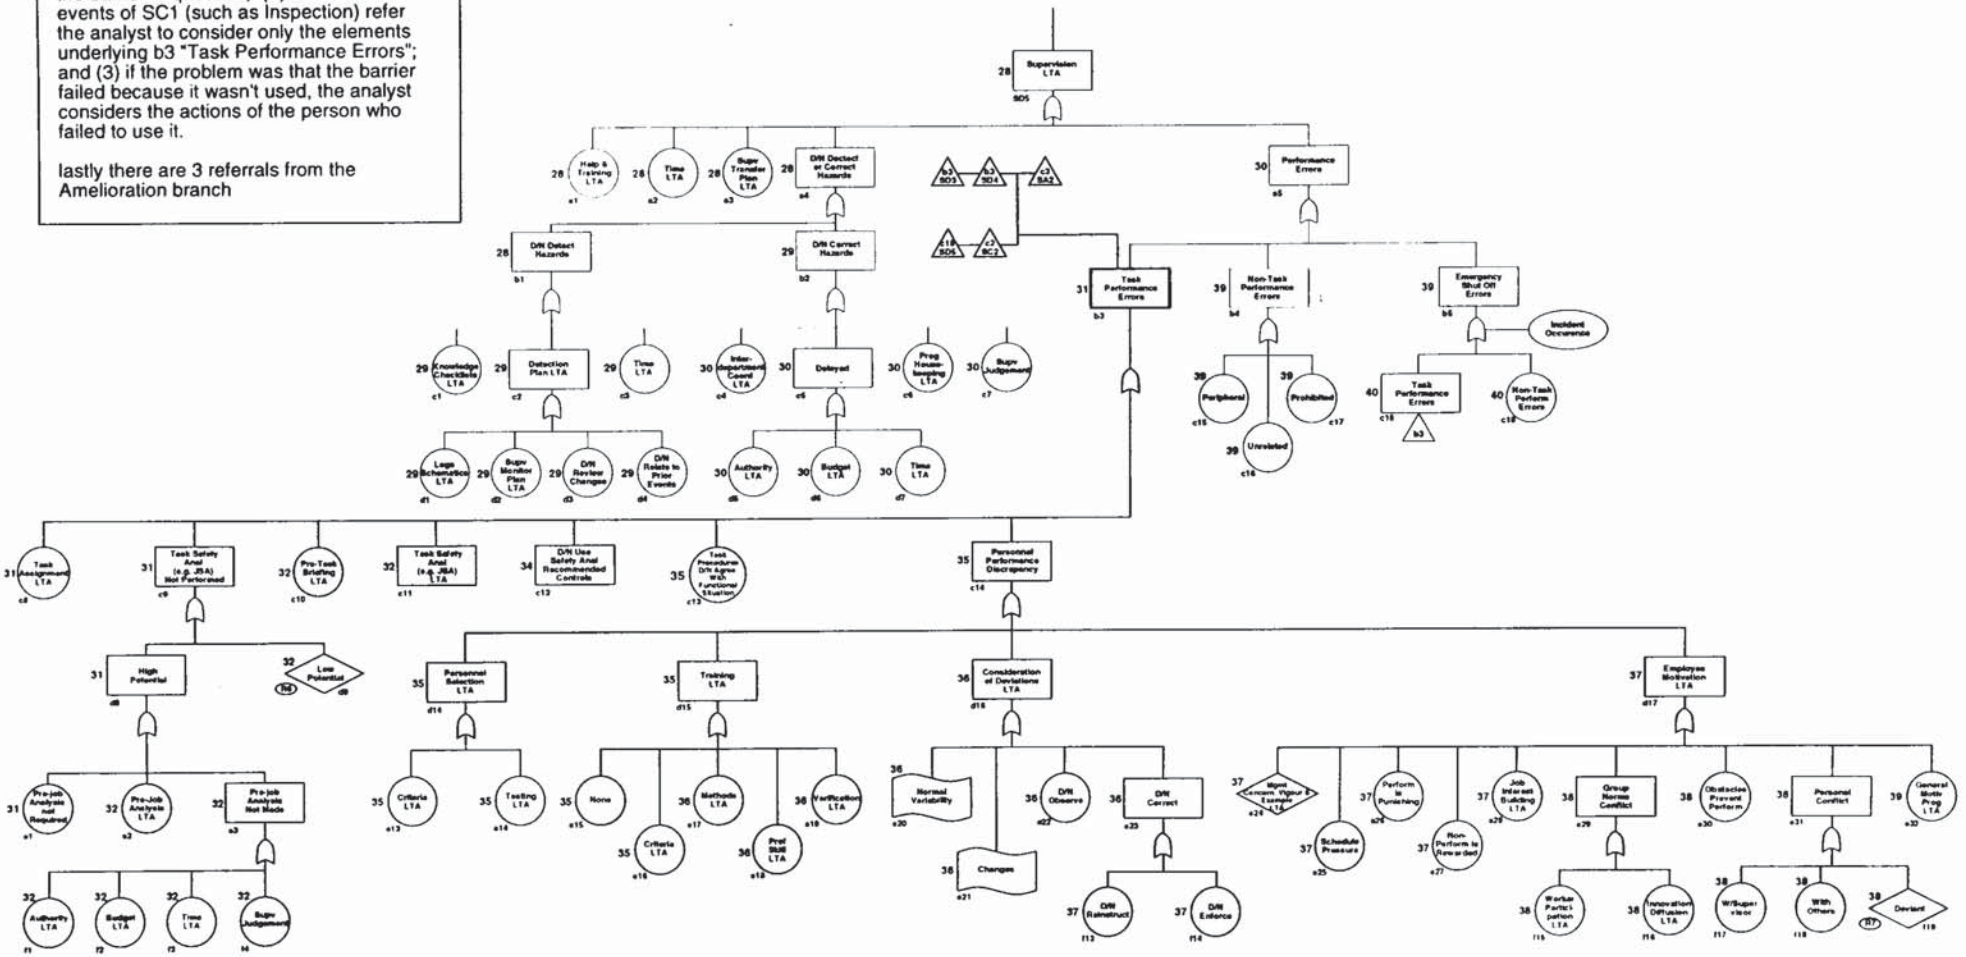

MORT CUTSET: "**E**" (**SD6, MB4 & MA2**) **Services Branch**

The "A" branch plugs into three places in the MORT chart. The SD1 "A" sequence focusses on the provision of information to, and gathering of information from, the operational interface containing the barrier in question. The MB2 "A" sequence considers the provision of information into the formal risk assessment cycle.

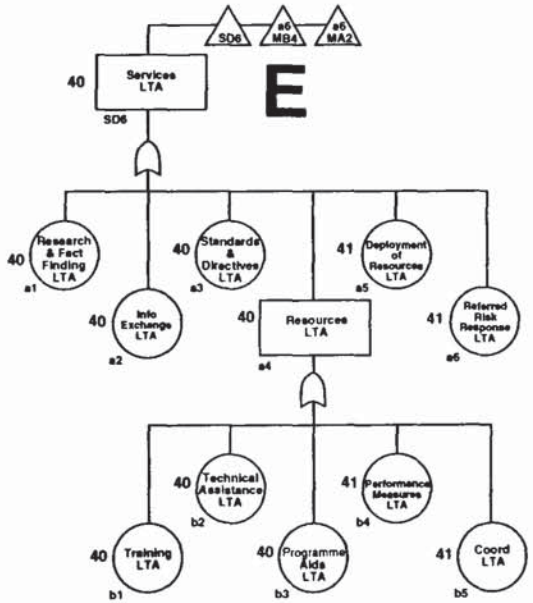

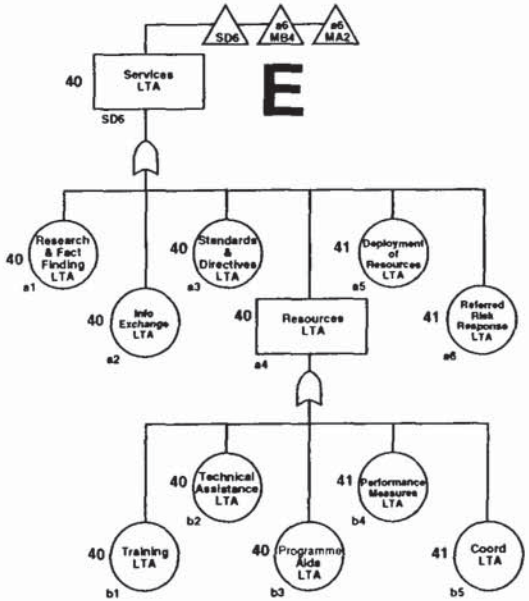

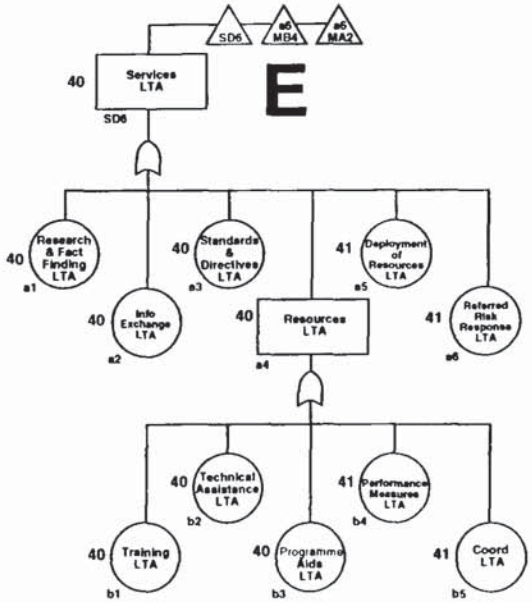

MORT CUTSET: **"F" Amelioration Branch**

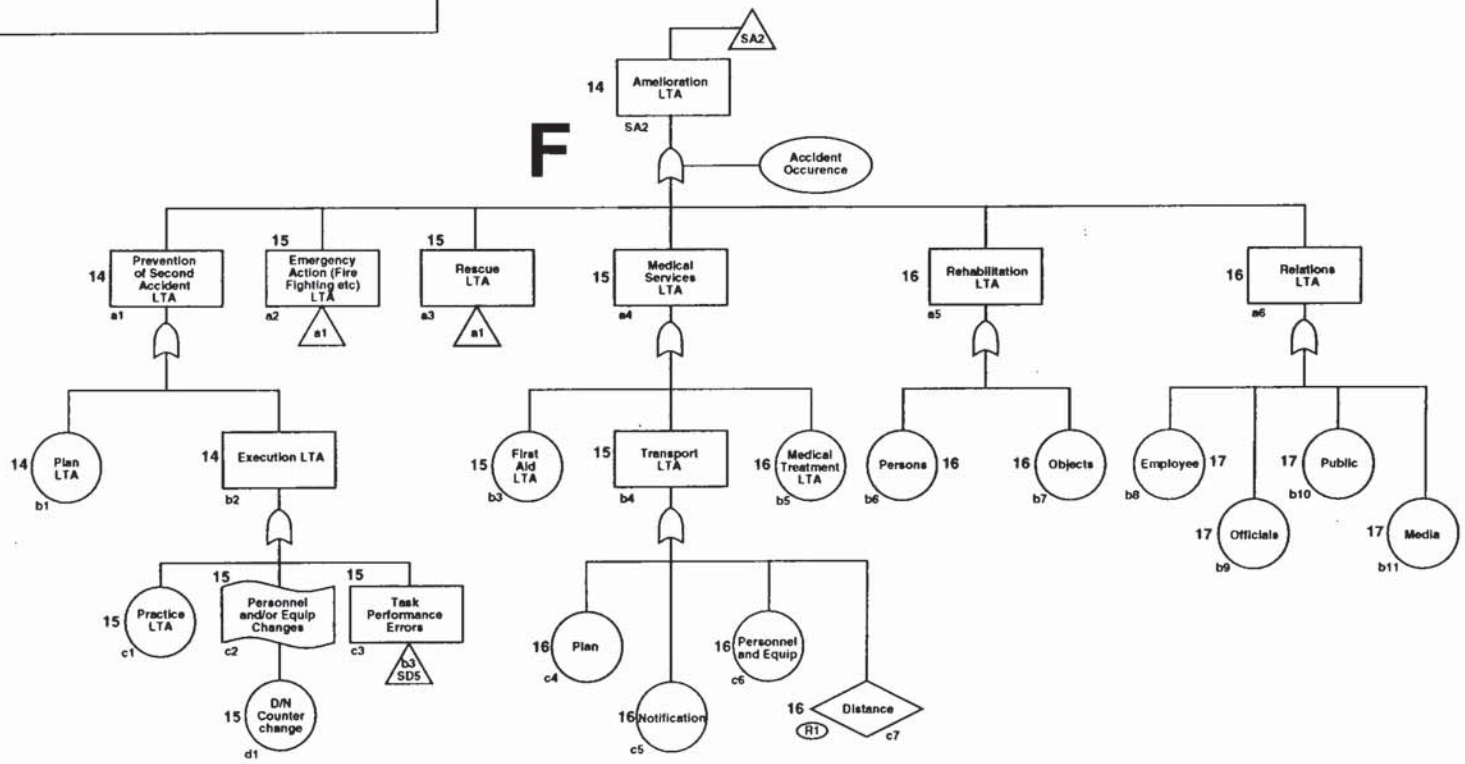

MORT CUTSET: **"G" Human Factors Branch**

The HF Branch should be considered when analysing the "M" side of the tree (under MB3-b8), that is, when the full "S" branch analysis is completed. However, you might also consider the G branch when analysing task peformance errors (e.g. under SD5-e3 "Obstacles prevent peformance"). If so, don't forget the barrier name and referral information

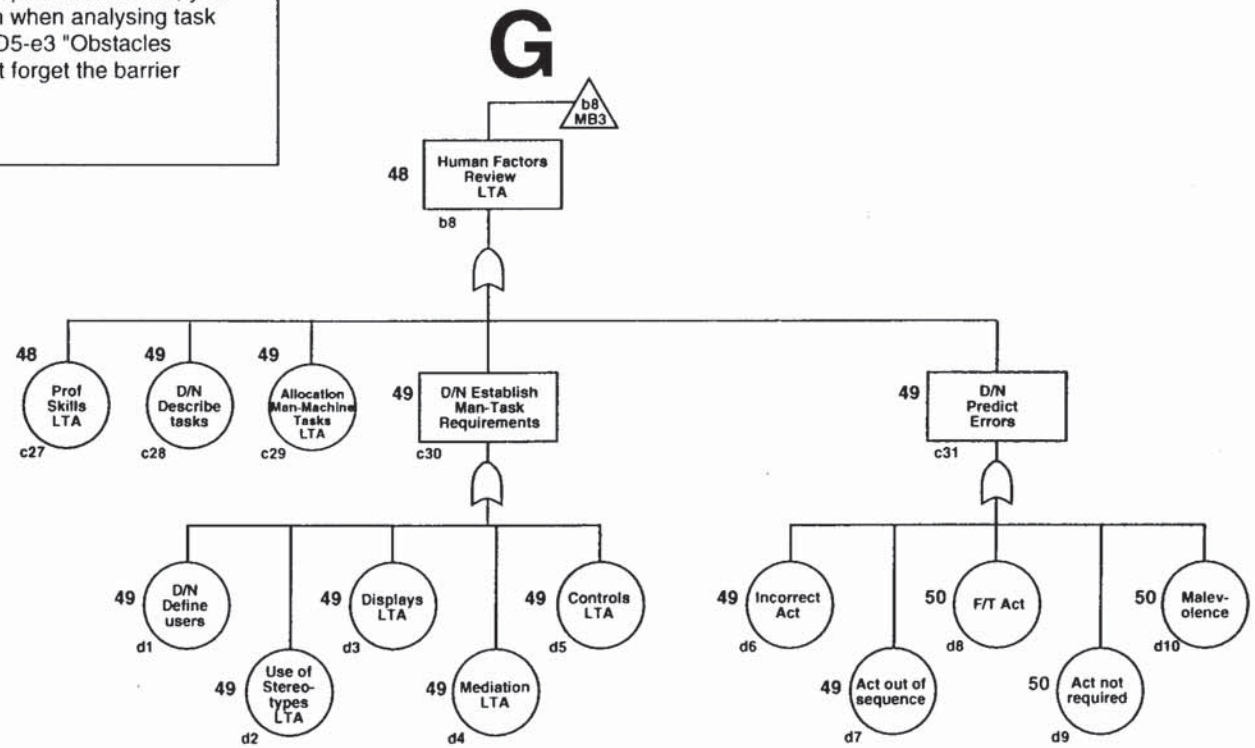

MORT CUTSET: **S BRANCH**

As event SB4 suggests; there may well be multiple energy flows to consider (e.g. electrical arcing --> insulation fire --> appliance fire --> area fire & smoke) in order to capture all the barriers and points of control. Each *significant* barrier/energy flow must be considered.Significance is for you to define.

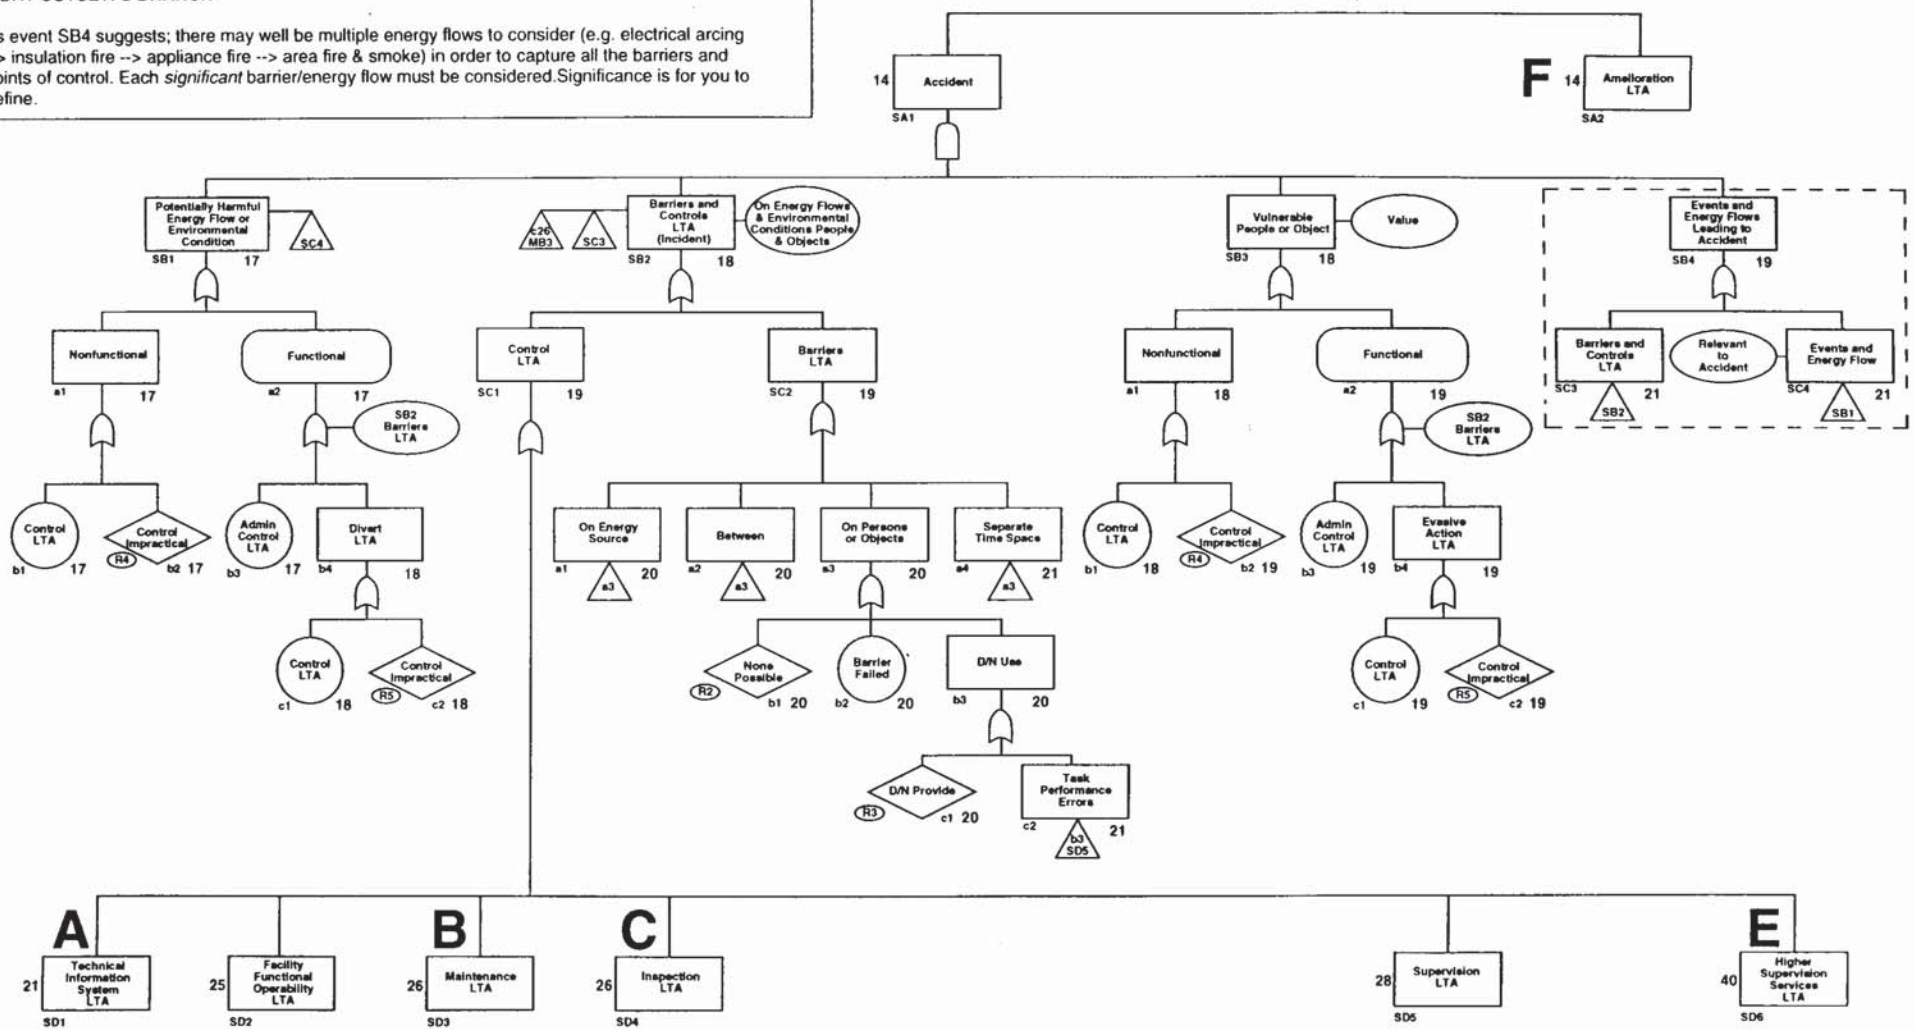

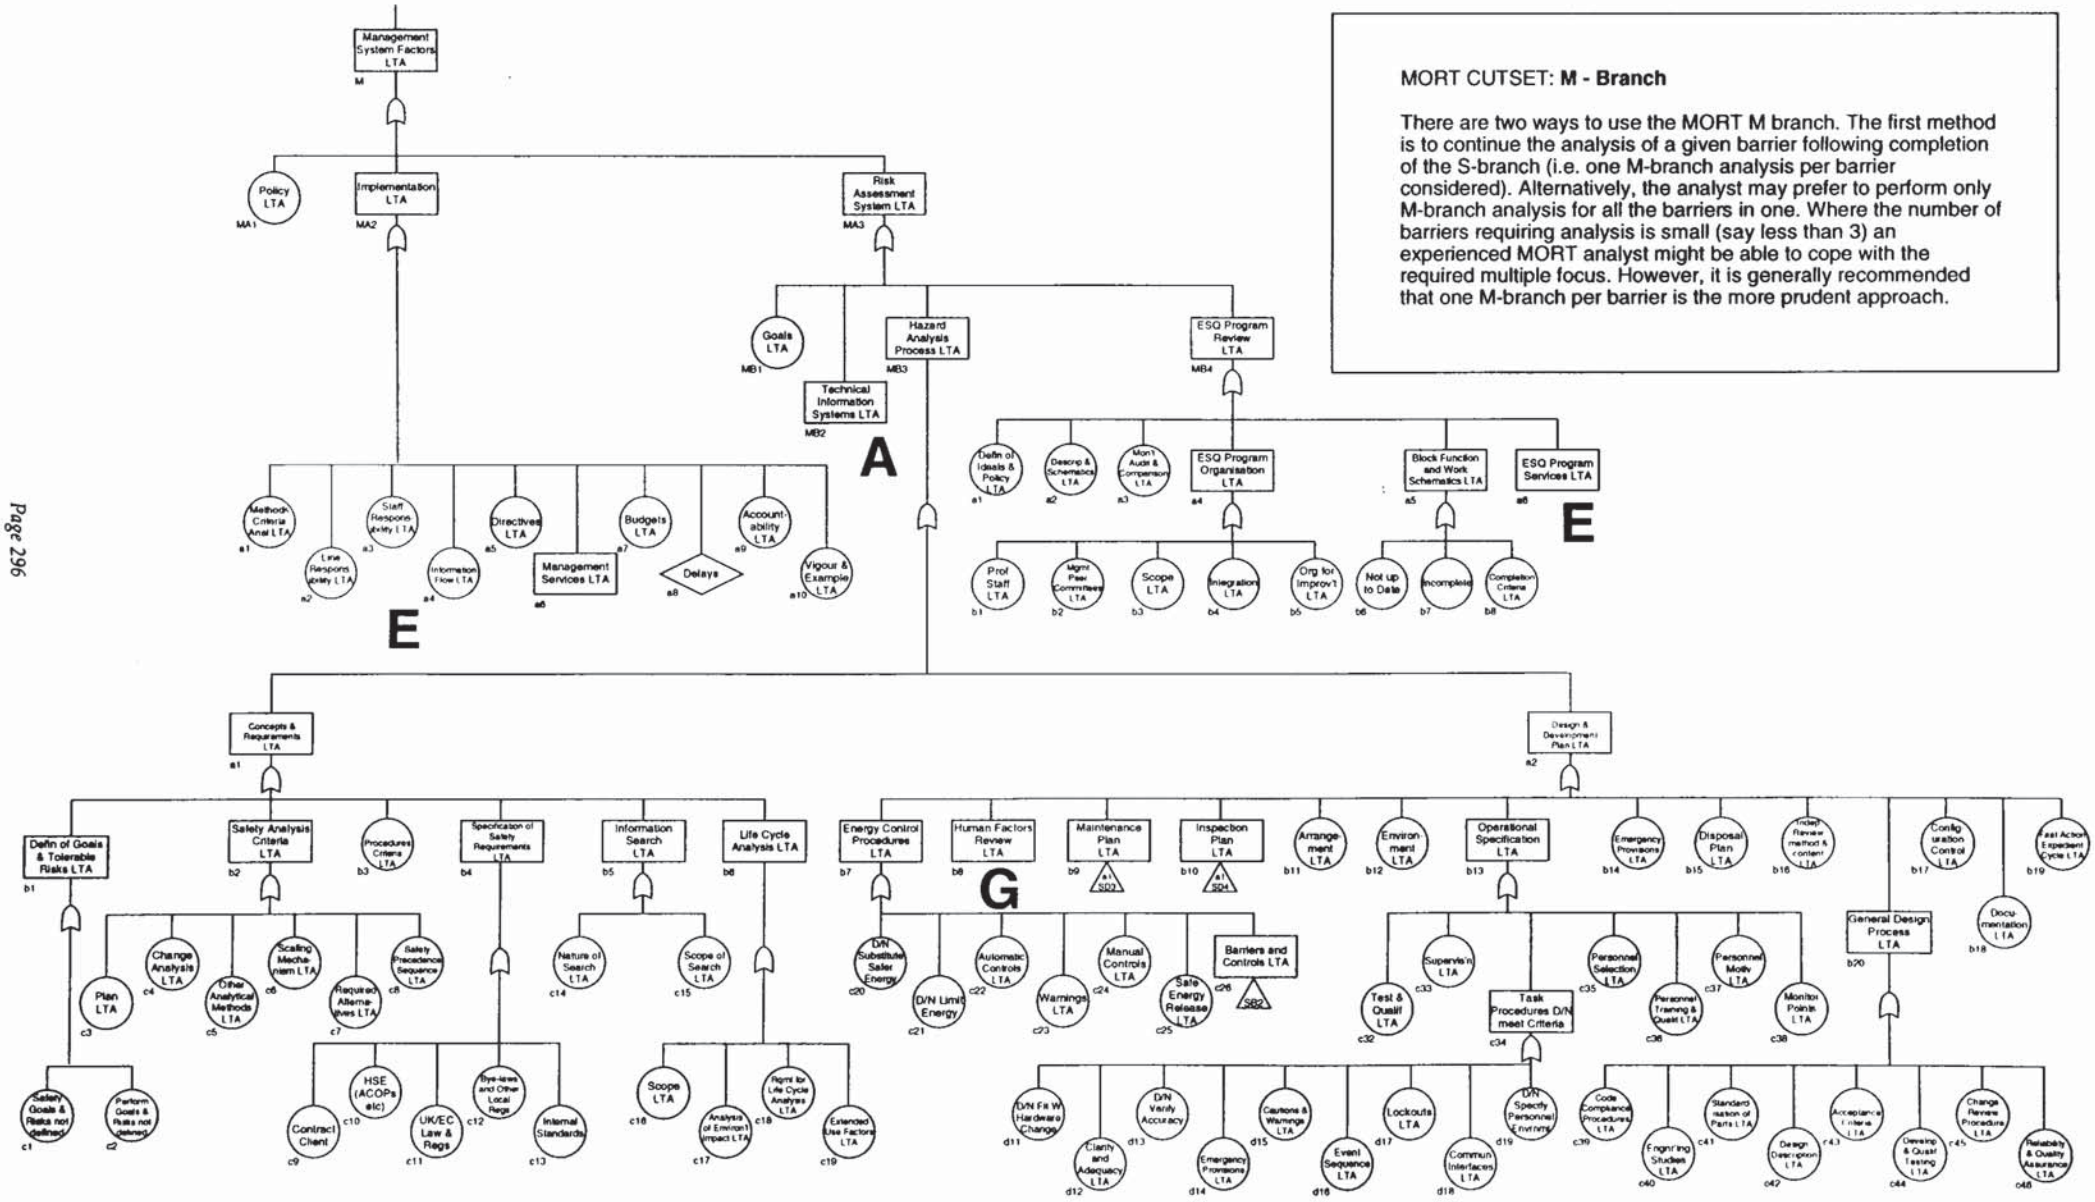

MORT CUTSET: **M - Branch**

There are two ways to use the MORT M branch. The first method is to continue the analysis of a given barrier following completion of the S-branch (i.e. one M-branch analysis per barrier considered). Alternatively, the analyst may prefer to perform only M-branch analysis for all the barriers in one. Where the number of barriers requiring analysis is small (say less than 3) an experienced MORT analyst might be able to cope with the required multiple focus. However, it is generally recommended that one M-branch per barrier is the more prudent approach.