The International Yearbook of ORGANIZATION STUDIES

Edited by David Dunkerley and Graeme Salaman

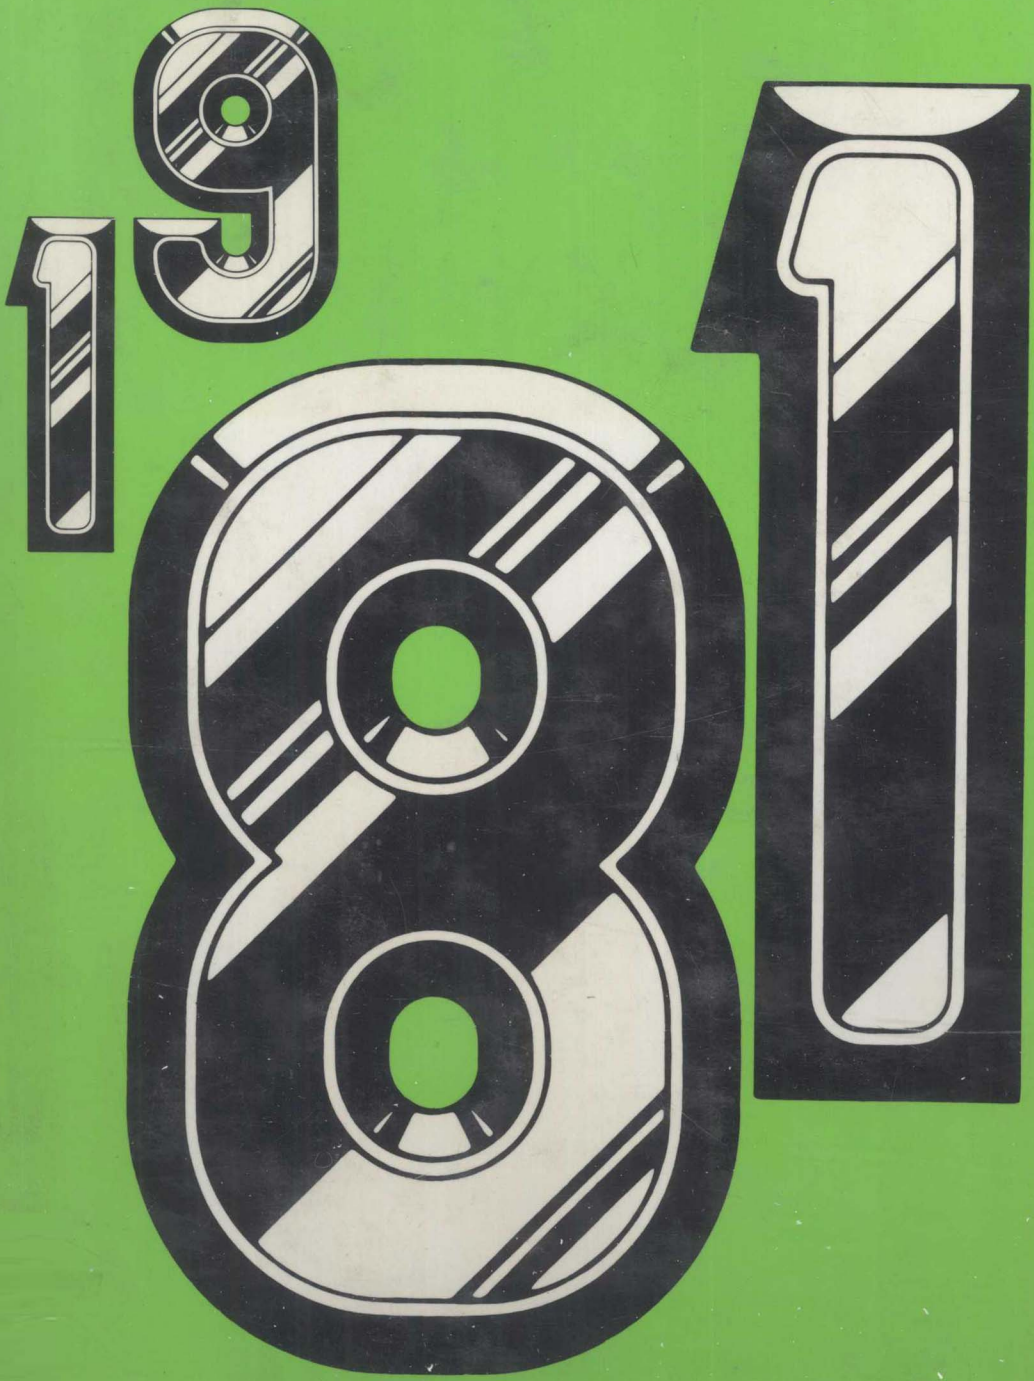

The internationai yearbook of organization studies 1981

edited by
David Dunkerley
Department of Sociology and Politics
Plymouth Polytechnic
and
Graeme Salaman
Faculty of Social Sciences
The Open University, Milton Keynes

Routledge & Kegan Paul
London, Boston and Henley

First published in 1982
by Routledge & Kegan Paul Ltd
39 Store Street,
London WC1E 7DD,
9 Park Street,
Boston, Mass. 02108, USA, and
Broadway House,
Newtown Road,
Henley-on-Thames,
Oxon RG9 1EN
Printed in Great Britain by
The Thetford Press Ltd,
Thetford, Norfolk

British Library Cataloguing in Publication Data

The International yearbook of organization studies. — 1981
1. Organisation — Periodicals
302.3'5'05 HM131

ISBN 0-7100-0996-8

Contents

1 Three Mile Island: a normal accident

Charles Perrow

Accidents will happen, including ones in nuclear plants. But, by and large, we believe accidents can be prevented through better training, equipment or design, or their effects can be localized and minimized through safety systems. The accident at Three Mile Island is being assessed in this fashion. The industry started a new training program, the equipment at the Babcock and Wilcox (B&W) plants is being improved, the design has been modified, the utility chastised - all useful, if minor steps. Furthermore, to nuclear proponents, such as Edward Teller, the accident proved that the effects can be localized and minimized. It is safe. No one has died as a direct result of radiation injuries in all the years of commercial nuclear plant operation.

But the accident at TMI was not a preventable accident, and the radiation vented into the atmosphere could easily have been much larger, and the core might have melted, rather than just being damaged. TMI was a 'normal accident'; these are bound to occur at some plant some time, and bound to occur again, even in the best of plants. It was preceded by at least sixteen other serious accidents, or near accidents in the short life of nuclear energy in the USA (Webb, 1976; 'New York Times' 31 March 1979, p. B.31), and we should expect about sixteen more in the next five years of operation - that is, in industry time, the next four hundred years of operation in the plants existing now and scheduled to come on stream.

Normal accidents emerge from the characteristics of the systems themselves; that is why I call them normal, and why they cannot be prevented. They are unanticipated; it is not feasible to train, design, or build in such a way as to anticipate all eventualities in complex systems where the parts are tightly coupled. They are incomprehensible when they occur. That is why operators usually assume something else is happening, something that they understand, and act accordingly. Being incomprehensible, they are partially uncontrollable. That is why operator intervention is often irrelevant. Safety systems, back-up systems, quality equipment, and good training all help prevent accidents and minimize the catastrophe, but the complexity of systems outruns all controls.

I will take the example of the accident at TMI as a prototype of

the normal accident and consider it closely later on, and also several other nuclear accidents and several non-nuclear ones in high-risk industries. We should become familiar with several types of accidents, especially normal ones, because we are more prone to them in an advancing industrial society.

The normal accident has four noteworthy characteristics: signals which provide warnings only in retrospect, making prevention difficult; multiple design and equipment failures, which are unavoidable since nothing is perfect; some operator error, which may be gross since operators are not perfect either, but generally is not even considered error until the logic of the accident is finally understood; and 'negative synergy,' wherein the sum of equipment, design and operator errors is far greater than the consequences of each singly. The normal accident generally occurs in systems where the parts are highly interactive, or 'tightly coupled,' and the interaction amplifies the effects in incomprehensible, unpredictable, unanticipated, and unpreventable ways. When it occurs in high-risk systems, such as those dealing with toxic chemicals, radioactive materials, microwaves, recombinent DNA, transportation systems, and military adventures, the consequences can be catastrophic. Even a fairly benign system, such as electrical power distribution, can cause considerable harm.

No one who owns or runs a high-risk system wants to consider a classification scheme for accidents that includes a normal accident category. It would be an admission of liability, and for some unknown but finite period of time, an admission of inevitable disaster. The category takes on more meaning when we contrast it to preferred ones. I will consider three major categories, though there are others. The best type of accident, for owners and managers, is the 'unique' accident, such as the collapse of a building in a giant earthquake, or simultaneous heart attacks for the pilot and co-pilot of an airliner or a bomber near a city. No reasonable protection is possible against freak accidents or Acts of God, so no liability can be assigned. They are so rare we need not fear them, and more important, even reasonable expenditures will not produce a significant reduction in risk. Otherwise, we would build no dams, buildings, airplanes, or armies.

Nevertheless, the unique accident is sometimes contemplated for high-risk, long-lived systems. About halfway into the nuclear power age it was required that the new plants be built to withstand large earthquakes, and the impact of a jet airliner. But even here it was only the reactor building that was so guarded; the auxiliary buildings and the pipelines to it from the reactor building, essential for using radioactive liquids to cool the reactor core in an emergency, are generally not protected. It is easy to imagine the loss of both main power system and the back-up one during an earthquake or even a storm. The designs, of course, have not been given destructive testing by actual earthquakes or falling planes. We missed a chance a few years ago when a Strategic Air Command Bomber, flying directly at a nuclear power plant in Michigan, crashed in a stupendous explosion just two miles short of the plant. The pilots at the nearby SAC base were routinely warned not to fly near or over the plant, though they routinely did, at 1000 feet, suggesting it would not have been a unique accident after

all, if it had occurred two seconds later (Webb, 1976, pp. 194-5).

Because liability cannot be assigned, owners and managers cry 'unique' when they can. The first explanation for the massive power failure in New York City in 1977 by Consolidated Edison spokesmen was 'an Act of God'. Lightning struck twice on the same power line. Mayor Beame's office just as immediately blamed the event on the gross incompetency of the utility, a type of accident we won't consider here. Subsequent investigation disclosed that while incompetency was rather extensive in the utility, and the act of nature unusual, it was the nature of the complex system that sealed the fate of the island and led to damages estimated to be at least $310 million. There was multiple equipment failure, leading to incomprehensible readings. A key operator decided that one line must have automatically reopened as it was supposed to (it didn't). Since it normally carried little or no current, the zero reading was more plausible as a reading of normality than the highly implausible alternative of a shut-down line. Thus he ignored what he probably viewed as panicky suggestions to shed load from the office that controlled the state power grid. In the few remaining minutes, each step that was taken made it worse, but each step, given his understanding of what should have been happening in such an emergency, was plausible. It was a normal accident (Wilson and Zarakas, 1978a, pp. 39-45; 1978b, pp.994-6).

Con Ed abandoned the 'unique' label and moved to the next most desirable one. This is the 'discrete' accident - there was an equipment failure, but it could be corrected and it won't happen again. Generally discrete accidents - which do occur - indeed are very plentiful in all human-machine systems, and nature itself - involve the failure of one of the two pieces of equipment (a relay failing to reclose after the lightning strikes), a limited design error, or an operator error. In a discrete accident the system responds to that source of error without any significant synergistic developments. Back-up systems and isolation devices come into play. Con Ed was wrong about the blackout because there were multiple failures of all sorts, and incomprehensible interaction effects for those at the controls, but they tried to treat it as a discrete accident. While liability can be assigned (nothing should fail, no matter how complex the system), it is generally limited (things will fail, nothing is perfect). More important, the label of a discrete accident is comforting because the system will not be abandoned; it can easily or conceivably be fixed. It will even be 'safer' afterwards than before, as with the nuclear power industry after each publicized accident.

At the press conference two months after TMI, Babcock and Wilcox, which built the reactor, argued that this was a discrete accident. There had been an instance of equipment failure, the pilot-operated relief valve (PORV), but it was the only instance of this and the system contained planned means to rectify the failure. The actual cause of the accident was the failure of the operators to follow correct procedures after the failure, they argued. If the operators had been on their toes it would have been a trivial event (Babcock and Wilcox, 1979, pp. 82-90). As we shall see, there were multiple equipment failures, a major design error, and the operators did just what at least some of the experts, some

months before, had said they should do. And the event was 'mysterious' and 'incomprehensible' even to B&W experts at the site. But management prefers the discrete label to the one that suggests the complexity of the system is at fault.

Similarly, the first reports of the cause of the Plain Bag Breeder accident in Michigan, where toxic chemicals were mistakenly fed to livestock, indicated that workmen made a simple, discrete error and mixed up bags of chemical with bags of feed. Such an event could be avoided in the future. Subsequent investigation disclosed other failures, akin to equipment failures, and leads to a presumption that organizational pressures were involved. The company ran out of the proper bags; instead of curtailing sales and deliveries, they substituted plain bags for the red ones used for the poison, and did not label them with the contents. The similarities of the names ('Firemaster' for the PBB mixture, used as a fire retardant, and 'Nutrimaster' used for the feed), and the stenciling and color of the bags led to an easy substitution. Operator error was not gross and was abetted by the equivalent of equipment failure (lack of supplies, in this case) and, possibly, work pressures. I suspect that further investigation might disclose warning signals about inventories, labeling problems, handling problems, work pressure, and so on in the firm, moving it into the category of a normal accident.

Discrete accidents allow for operator intervention; the accident itself is comprehensible - someone made a mistake; the equipment failed; the design did not allow for this eventuality - so something can be done. They are also preventable (to the extent that accidents ever are) by noting warning signals, by using back-up or safety systems, and, of course, by rectifying the problem after the accident. Liability can be assessed, but our system of governance and our judicial system is lenient in this regard; 'It won't happen again, sir.'

Discrete accidents are convenient for risk assessment. The Rasmussen Report on the safety of nuclear plants, which found that the risks were comparable to being hit by a falling meteor, pursued a 'fault-tree' analysis which in essence looked at a long chain of possible discrete events stemming from a discrete failure. The possibility of multiple failures at any juncture in the tree was not, critics noted, adequately investigated (Union of Concerned Scientists, 1977, pp. 10-16). But it is hard to see how it could be; the complexity of the analysis would increase exponentially if multiple failures were included. Nevertheless, as a regular consultant to utilities using nuclear energy, Rasmussen was no doubt prone to thinking of accidents as either Acts of God (being hit by a meteor) or discrete.

The most troublesome category of accidents, both for owners and managers and for the theorist, is the 'calculated risk' accident. Liability, where risk is calculated, could easily be assigned, so owners and managers avoid any admissions of calculation and prefer the category of unique, or failing that, discrete accident. Theorists have troubles too since on the one hand there is a sense in which a calculation is made of every known risk, making the category vacuous, and on the other hand, there are presumably many unknown risks in complex systems, so calculations are not possible,

again rendering the category vacuous. Between these two extremes (more could be done to prevent it since calculations are made and nothing could be done to prevent it since some things will be incalculable) is a messy but useful area.

Consider the reasonable, state-of-the-art efforts that organizations continually make to calculate risk and weigh the benefits and costs of safety measures. Documents from the Ford Motor Company clearly indicate that they decided the few dollars (variously reported as $12 or $16) required to protect the gas tank from exploding when the Pinto was struck from the rear at a low speed would cost more over the production run of the model than the expected suits for vehicle damage, personal injury, and deaths. Reportedly, the fire that killed the astronauts on the launch pad was considered to be possible, but the level of safety deemed acceptable was below the level of this possibility, so the risk was run. Once it happened, the system was redesigned, perhaps because of the unfortunate publicity rather than a reassessment of the risk calculations, just as the still unburned Pintos were recalled once the government intervened in the private calcuation of risk. However, our country was built on risk, as we are hearing lately in the USA. Thus, a stadium built in the flight path of O'Hare Airport in Chicago had its roof mysteriously collapse, perhaps as a result of shock waves from the jets passing closely overhead. Undaunted, the owners are rebuilding, calculating that the next roof will be stronger, and calculating that the chance of a plane crash is extremely remote - at the busiest airport in the country.

Not surprisingly, the calculation of risks rarely becomes public knowledge, so we often do not know whether an accident falls in this category or some other. When the nuclear submarine, the USS 'Thresher' went to the bottom off the Continental Shelf some years ago, it could have been a unique accident (unlikely, but possible), or the result of a single piece of equipment failure or an operator error (more likely), or a normal accident (quite likely, given the complexity of the system and problems experienced in previous test dives). But most likely it was a calculated risk accident. After a year of sea trials the vessel was in dry dock and Admiral Rickhover (notorious for his insistence upon quality construction and crew training) had a small sample of the welds tested. It was found that 14 per cent would not pass inspection. (Shades of the weld inspection in the movie, 'China Syndrome.') Then the vessel was taken on a deep test-dive, even though the emergency deballasting system had never been tested in any nuclear submarine. After the fatal accident, the system was tested in port on another submarine and it failed ('Newsweek,' 1965; 'US News and World Report,' 1963, pp.38-9, 1965, pp. 10-16; Curtis and Hogan, 1969, pp.72-5). The last intelligible message from the 'Thresher' suggested that the emergency deballasting system was not working. Water apparently came in and damaged the reactor which was shut down leaving the ship without the power to surface. The case is interesting, because with the extreme safety consciousness of the Admiral, and the resources of the Navy, we would not expect to find either a normal accident or a calculated risk one.

Nuclear proponents are fond of saying that all imaginable risks have been calculated; indeed, this is cited as a major reason for the escalation in plant costs by proponents. However, substantial risks

that are considered too high to run in new nuclear plants, and thus must be designed out of them, are left to simmer in old plants. In an important decision in 1973 the Atomic Energy Commission (AEC) ruled that it would not be necessary to 'retrofit' existing plants and those to be licensed for the next four years with a back-up SCRAM system (an emergency system to halt reactivity). It was economically prohibitive. The TMI plant lacked the back-up emergency core cooling system (ECCS) that is required of newer plants, just as the early ones such as one at Indian Point, New York, required none. As we shall see, however, it probably would not have made any difference in the TMI transient. (A 'transient,' incidentally, is a technical term, and does not imply an insignificant degree of danger, as some might think. It means, in this connection a steady loss of coolant.)

Westinghouse, a major reactor builder, and the Phillipine government, wish to build a nuclear plant in Batan that would service a US military base and a growing industrial area that is of interest to US corporations. Opponents in both countries argue that the plant would involve an unacceptable calculated risk. It would be built in an area that has had severe earthquakes, tidal waves, and is within one hundred miles of four active volcanoes (whose ash would foul the plant filters, causing severe dangers) and ten miles from a long dormant volcano. The Nuclear Regulatory Commission (NRC) had the plan under study for over a year, and finally ruled against it. If it is built, as still seems quite possible, and were there to be an accident as a result of an earthquake, tidal wave or volcanic activity, it would undoubtedly be attributed by Westinghouse and the NRC to an Act of God, a unique accident, while opponents would consider it a calculated risk of unacceptable proportions.

As suggested, this category is a messy one, open to debate after the fact and hidden from view before it. In any case, the tendency is to classify accidents as unique events, or discrete accidents, rather than calculated risks. Calculated risk accidents that we are able to learn about are generally cataclysmic (that is why we know of them), and thus, like unique accidents, operator intervention is negligible and synergistic effects are irrelevant, though probably present in those few seconds of disaster.

WARNINGS

Complex human-machine systems abound in warnings - signs in red letters, flashing lights, horns sounding, italicized passages in training manuals and operating instructions, decals on equipment, analyses of faults in technical reports, and a light snowfall of circulars and alerts. Warnings are embedded in construction codes, testing apparatus, maintenance checks, and, of course, fire drills. All personnel are expected to be only indifferently attentive and concerned, so training, drills, reports and alarms are repetitive, numbing, essential parts of these systems.

Warnings work; there is no question about that. But not all. Virtually all of the major, serious accidents I have examined in detail are replete with warnings which are unheeded. We should not be surprised; the very volume of warning devices testifies to this

likelihood. If warnings were heeded, we would need only a few modest and tasteful ones rather than a steady drill of admonitions punctuated by alarms and lights.

Yet we stand incredulous when confronted with, for example, the same engine on the same DC-10 aircraft failing twice within a few months (one fatality - a passenger sucked out of the plane), or the cargo doors of DC-10 aircraft, after repeated warnings, blowing open three times (the third a fully loaded plane crashed and all died), or an accident at Three Mile Island that seemed to be almost a simulation of two previous accidents at other plants and fulfilled the predictions of a hypothetical analysis of an engineer. Why are warnings not always heeded? There are many reasons, and when we consider the over-population of complex, high-risk systems that someone has decided we cannot live without, they are disturbing.

Consider three categories of warnings. First, there are deviations from steady-state conditions that do not activate significant alarms. We will consider a rather long list of these from the Three Mile Island plant later, when we look at equipment and operator failures. Each one individually is considered trivial or interpreted in a routine framework. Only retrospective judgement, hindsight, discloses the meaning of these deviations. Second, there are alarms, such as flashing lights or circuit breaker trips or dials reading in the red zone. But operators are accustomed to reinterpreting these alarms as insignificant when they have a conception of the problem which is violated by the alarm. Or, if there is as yet no conception of the problem by the operators, the alarm may be attributed to faulty alarm equipment. Since dials sometime give faulty readings or breakers trip for no good reason even under routine conditions, and since disturbed conditions can create misleading alarms through malfunctioning or complex interactions, the operators may be correct. Alarms, like deviations, always outnumber actual accidents; warnings are in greater supply than actual malfunctions. 'If we shut down for every little thing...' the reasoning goes.

Past accidents, mute predictors of future ones, is the third category of warnings. But history is no guide for highly infrequent events; they are not expected to occur again. Generally, they don't. Or, there may be compelling economic reasons for continuing to run the risk. It would have been quite expensive to redesign the DC-10 so that the controls either were not ducted through the cabin floor, which tended to collapse when there was a loss of air pressure, or to strengthen the floor. A cargo door that opened out, instead of in, and had a poorly designed locking system, tended to blow open as the cabin was pressurized. Rather than redesign the door and strengthen the floor, a venting system was used instead to relieve pressure after two occasions when the floor collapsed and severed many of the controls. Untested and problematical, it did not prevent a third, fatal collapse near Paris, when the door again blew open (Godson, 1975).

In addition to believing it will not happen again , and compelling economic reasons, past accidents fail as warnings if the warning is available to only one part of the system, and that part is only loosely connected to the other parts. This was a major problem at TMI.

Any single plant with a complex technology is likely to be tightly

coupled; a disturbance in one part will reverberate quickly to the other parts. But the plant may be only loosely coupled (Weick, 1976, pp. 1-19) with other parts of its system. Warnings from another plant may not reach it; the mechanisms for transmitting such warnings in the case of nuclear power plants are reasonably redundant and plentiful - the NRC, the reactor builders, numerous institutes, university centers, and industry bodies all function in this capacity. They have a strong interest in preventing accidents. Indeed, in a crisis, the system comes together tightly; it responded exceptionally well to the TMI transient. They knew that the future of nuclear power was at stake. But under normal conditions they have an interest in minimizing the dangers that exist, avoiding costly shutdowns, and carrying out their separate organizational concerns. These interests buffer the part of the system that experiences a disturbance from the other parts, unless the disturbance is very large and widely publicized. In such a manner TMI was buffered from a technical report prepared by an engineer at another utility, a somewhat similiar accident in Europe, and a very similar accident in an adjacent state. All constituted unheeded warnings.

The technical report was prepared by Caryle Michelson, an engineer with the Tennessee Valley Authority (TVA) which was considering the purchase of a reactor from B&W, one quite similar to the two reactors at TMI (Report of the President's Commission, 1979; Hearings, 19 July 1979). Michelson wrote a long memo raising a number of concerns. Among these concerns was a remarkably prescient description of the dynamics of the TMI accident: A LOCA (a loss of coolant accident) occurs, a high pressure injection (HPI) goes on to maintain pressure in one part of the system, the pressurizer. The pressure rises there, but falls in the reactor core for complex reasons. The operators fear over-pressurizing the pressurizer, because it might 'go solid' (become saturated with water and/or steam). Going solid is to be avoided, since it means the reactor must be shut down if it isn't already (SCRAM, or inserting graphite control rods to stop the fission process) and even if it is already, it takes a long time to get it back in operation after going solid, and the utility loses money because it must buy electricity rather than make it. So they 'throttle back' on the HPI, but this will mean less cooling of the reactor core and could lead, in minutes, to damage to the core and even a melt-down.

Michelson's report was sent to the NRC in November 1977, and a reply acknowledged they understood the problem, but they kept it to themselves. In April 1978, eleven months before the accident, it was sent to the vendors, B&W. There it received normal handling. The engineers read it, considered it, and wrote a reply nine months later, two months before TMI, stating that these matters had all been considered. We do not know what happened to it at the NRC; it seems to have disappeared in their vast files.

Meanwhile, on 24 September, 1977, a LOCA occurred at the Davis-Besse plant near Toledo, Ohio. The operators throttled back on the HPI when they saw the pressure in the pressurizer rising, even though it was falling in the core. Fortunately, the plant was operating at only about 9 per cent capacity, and in a short time they discovered the cause of the accident - a faulty PORV, and bypassed it before any damage to the core occurred. An engineer

from B&W, Mr Kelley, was sent to the plant to investigate the accident. Returning to B&W, he gave a seminar on the transient, warning about the improper operator action of throttling the HPI system prematurely, and then wrote a memo suggesting that all units using this kind of equipment be warned about this improper action.

Mr Kelley's superior, Mr Dunn, took up the matter and had his memo sent around B&W. Only one engineer responded, and he misunderstood it and dismissed it. Dunn persisted, and the memo, now fathered by a Mr Novack, made a slow ascent. It was sent over to that division of B&W concerned with customer services, to Mr Karrasch. He said he gave it to two subordinates, but they dc not recall ever seeing it. It was sent there because customer service is traditionally concerned about anything that might unduly interrupt service, and since going solid would, they should review it. (Kelley-Dunn-Novack were concerned about the far more dangerous matter of core damage and melt down.) No word came from Karrasch, so Novack kept calling. Months went by, and still no answer as to whether they should alert all utilities to this danger. Meanwhile, the training department had assured Kelley-Dunn-Novack that operators were, indeed, instructed to not throttle back on HPI in a LOCA, even though they had at Davis-Besse.

Finally, a Mr Walters met Karrasch at the water cooler and asked about the memo from his people on the engineering side. Karrasch replied, off-handedly, something to the effect that 'It's okay, no problem.' Mr Walters pondered the reply as Mr Karrasch hurried off to a meeting - did it mean there was no problem of going solid, or no problem of uncovering the core, or what? Irresolutely, he left the matter hanging. It all came out after the operators at TMI throttled back on the HPI and made a serious accident even more serious. Nineteen months had transpired since Kelley first wrote his memo. B&W then quickly sent out the Kelley-Dunn-Novack memo to all units using this equipment.

To the members of the President's Commission on the Accident at Three Mile Island this was the familiar curse of a failure in communication, the phlogiston of organizational problems and of many disasters. Warnings were not made available to the proper people; Karrasch, at the least, had failed to communicate with the engineers. Mr Karrasch was more perceptive, if aggressively defensive. There was no failure in communication he insisted; the matter was simply one of low priority. He then went on to suggest the several obviously high priority matters his office was dealing with, ones forced upon them the implication runs, by new and pressing NRC safety standards (Hearings, 1979, pp. 249-51). He was right. Everyone at B&W did what they were supposed to do, with both the Michelson and Kelley memos. Only in retrospect had they assigned the wrong priority. In retrospect we often do.

How many warnings can one heed? The best set of warnings lie among the 2,OOO Licensee Event Reports (LERs) that are sent by the utilities to the NRC every year. These are required by law, and report significant events that might effect safety. The NRC has gagged on them; no reasonable system for analyzing them exists. The utilities dutifully report these and they sink into an enormous file. Perhaps the reports hope the investigator of some future catastrophe will find them and say, 'See, we were warned!' What would the

operators, even if they were college-trained engineers, do with a steady stream of reports, memos, instructions, analyses that they would be required to remember for years on end, use rarely, and recall instantly in a complex emergency? Would the memo on the HPI have made any difference at TMI? Only if it had been remembered, along with all the other instructions that continually change, and more important, only if they had known that it was this type of accident that they were in. As we shall see, they didn't. Even the experts that were quickly on the scene did not know soon enough.

It is not clear that the system should be more tightly coupled so that warnings, for one thing, should travel faster and create their intended 'perturbances.' Were the TVA, NRC, Battelle Institute, Brookhaven Labs, university departments, Electric Power Research Institute, Oak Ridge Laboratories, Westinghouse, Combustion Engineering, B&W, Davis-Besse and TMI and some seventy other plants all wired together into one low resistance circuit, the number of untoward events and immense complexities lying in the nuclear industry would drown them all in signals. Loosely-coupled systems have slack, reserve time, and resources. One part of a system can be made to withstand the brunt of a disturbance and protect the others from incessant shocks. Parts can be isolated and even left to fend for themselves. Information is absorbed, summarized, compacted into bits of information in one part that can be sent to the others without inundating them. Centralization is avoided; innovation encouraged.

Such loosely-coupled systems are resistant to change from the outside, however. By focusing upon TMI, the President's Commission unwittingly reinforced the survival values of loosely-coupled systems - the utility was segregated from the industry, and reprimanded. Indian Point, with its old equipment grandfathered from safety requirements, perched up wind of the millions in the New York metropolitan area, is buffered. Better equipment and training and management at TMI will supposedly take care of the problem, along with a single-headed rather than a hydra-headed NRC and some new 'attitudes' there. Operators will be flooded with new warnings. But it is normal for the systems to have accidents; warnings cannot affect the normal accident.

Tight coupling encourages normal accidents, with their highly interdependent synergistic aspects, but loose coupling muffles warnings. The PBB accident in Michigan was no doubt abetted by the loose ties that linked chemical plants to farmer co-operatives to farmers buying animal feed, though the tight coupling within the chemical plant allowed poison to be packaged in similar ways and right next to animal feeds. The loose ties that the drug company Hoffman-LaRoche had with its plant in Seveso, Italy, no doubt contributed to that disaster. The parent firm, in clean white Switzerland, was extremely careful of Dioxin, a byproduct of an insecticide, in its spotless laboratories. They had no plants built in Switzerland. It had the chemical plant that produced the insecticide built in dirty northern Italy, and it is likely that the workers and even many of the managers were not aware of what they were handling. The plant was poorly run and leaked Dioxin so frequently that farmers who brought in dead animals were routinely compensated for them - so much a chicken, so much a hog. The production crew left a batch of chemicals in a reactor over the weekend, and while the plant was

unattended, the batch unaccountably heated up and exploded, drenching the surrounding community with the deadly poison. Warnings that a very toxic matter was being handled (and quite casually) had not reached the plant from Switzerland (Whiteside, 1979). Loose coupling had its advantages; Hoffman-LaRoche was buffered. It was several hours after the explosion before anyone even bothered to warn the nearby residents, and then only to stay indoors and not eat fresh fruits and vegetables for a while.

Whether systems are loosely or tightly coupled, they all face another problem with warnings - the signal to noise ratio. Only after the event, when we construct imaginative (and frequently dubious) explanations of what went wrong, does some of the noise reveal itself as a signal. The operators at TMI had to literally turn off alarms; so many of them were sounding and blinking that signals passed into noise. The extremely detailed log of the accident (accurate to the tenths of a second), put out by B&W, performs this merciful winnowing task for us now, selecting out the noise and giving us the signal, with the unspoken admonition 'see this reading; *that* was significant.'

Noisy systems illustrate the banality of the normal accident. Prior to the attack upon Pearl Harbor, there were a dozen or so bits of information that warned that the Japanese were about to attack. But there were thousands of bits of information among these dozen that indicated an attack was not forthcoming, and was even impossible (Wohlstetter, 1962). Even when a few of the warnings were singled out as they passed up the hierarchy, they were, properly, discounted. No one in their right strategic mind would have anticipated that the Japanese would be so foolhardy as to steam thousands of miles across the ocean in a large attack fleet, through shipping lanes and within reach of patrol planes, and expect to be undetected. Had the commanders at Pearl Harbor used limited and precious fuel to keep planes aloft and ships at sea looking for an unimaginable event, they would have been cashiered, because all intelligence indicated the resources would be needed for expected attacks in South-East Asia (which did occur). The problem was not a failure of 'intelligence' (in the military sense of the term), but the routines in which the signals were embedded, and the strategic 'mind-set' that was legitimately operative.

Complex systems are simply not responsive to warnings of unimaginable or highly unlikely accidents. Because they are complex, organizational routines must be carefully followed and off-standard events reinterpreted in routine frameworks. Fortuitous events are always more plentiful than unfortuitous ones, Murphy's law (if anything can go wrong it will) notwithstanding. Most things that go wrong don't matter; the redundancies are plentiful. The 'mind-set' that the Commissioners referred to so often in their discussions with witnesses allows organizations to go forth without an agony of choice over every contingency. The phrase 'I'll believe it when I see it' is misleading, an organizational theorist, Karl Weick notes; it is equally true that 'I'll see it when I believe it' (Weick, 1976). The warning of an incomprehensible and unimaginable event cannot be seen, because it cannot be believed. But since it is inconceivable that there were not warnings, investigators, congressional committees and the superiors of hapless operators dig among

the wreckage until they find what can pass for an unheeded warning. But the normal accident is unforeseeable; its 'warnings' are socially constructed.

DESIGN AND EQUIPMENT FAILURE

It is obvious that designs cannot be perfect or fail-safe, nor can equipment. Everything dangerous would be far too expensive to build and maintain if we required maximum state-of-the-art efforts in equipment and design. Some risk must be run if we wish to have nuclear plants, rail and air transportation, chemical fertilizers, large buildings, military raids and so on. Even nearly fanatic efforts to reduce risks, such as those reported of Admiral Rickhover, are insufficient. Recall that 14 per cent of the tested welds on his submarine were not up to standard, and the 'Thresher' failed its first deep water dive. Given the robustness of most industrial systems, equipment and design failures are not likely to be catastrophic; though they are obviously heavily involved in the 5,000 or so industrial accident deaths we produce in the USA each year. Failures might be catastrophic in high-risk industries, such as the nuclear power industry, especially when the failures are multiple and interacting. Multiple and interacting equipment and design failures abounded in the case of the TMI incident, and several other nuclear accidents or near accidents.

The major piece of equipment failure at TMI was the PORV. It stuck open. The event was not without prior warnings. There were at least eleven other failures of this key valve at other plants before TMI. The same valve, but made by a different subcontractor, stuck open at the Davis-Besse plant causing that accident less than a year before TMI. Fortunately, Davis-Besse was only operating at about 9 per cent capacity, rather than TMI's 98 per cent. The valve had failed once before in TMI Unit 2, and some corrections had been made, but they were obviously insufficient. Furthermore, prior to that failure, it was not possible for the control room operator to easily determine whether the valve was open or closed. After the initial failure, a parsimonious step was taken. A signal was installed, but it only indicated whether a signal was sent to the valve to open or close it, not whether it was actually open or closed. In the March 1979 accident, the indicator said it was closed, while in actuality it was open. Furthermore, the valve had been leaking for some weeks, making check readings from the drain pipe attached to the valve unreliable.

The valve is a particularly crucial one in the pressurized water reactor design of B&W, since the steam generators may boil dry very rapidly - two or three minutes - rather than slowly as in the boiling water reactor designs built by other firms (15 minutes in one design, and 30 in another). This instance of tight coupling makes core uncovery more likely, though B&W officials argue that it also provides advantages in other kinds of accidents. It also has the distinct commercial advantage of allowing the reactor to continue operating even if the turbine shuts down, thus minimizing expensive downtime.

This advantage was removed after TMI when B&W, following

discussions with the NRC, reduced dependence upon this critical valve by having the reactor shut down whenever the turbine tripped. In testimony, a B&W official was reluctant to say that this corrective action signified a design problem in the original B&W equipment, but it would appear to indicate quite a significant one. Thus, there were several warnings, insufficient corrective action, a major failure, and only then, a design change in the system (not the valve).

There were other equipment failures during the transient. Paper jammed in the computer printout, and to get the printout operating considerable data logging had to be sacrificed. The computer was presumably not designed to handle the volume of a major accident and was one and one half hours behind in its printout at one point. There was an error in the instrumentation for the level indicator in the miscellaneous waste holding tank. A check valve was faulty and it let water into the condensate polisher system; this had been noticed before, but the attempt at correcting it had not succeeded. This particular failure probably started the whole transient but in normal accidents the particular trigger is relatively insignificant; the interaction is significant.

There were serious leaks - the source of which was still unknown some weeks after the accident - in the venting system, allowing unintended radioactive releases to the atmosphere. A safety system was not used because it was not safe; it could easily leak. This was the normal back-up system for cooling the reactor by returning liguid from the auxiliary building. Because it could not be trusted, poisonous gas was vented directly into the auxiliary building (and then went to the atmosphere) in a controversial decision which produced the large radioactive puff. Several people (including a utility official from Metropolitan Edison) testified that leaks in this 'safety' system made it a dangerous procedure. That a safety system would be too dangerous to use suggests both a design and equipment failure of some magnitude.

The following items were not working at the time of the transient or had failed in the recent weeks: the auxiliary building sump tank had blown a rupture disc some weeks prior to the accident and operators were by-passing the tank (there are no regulations that prohibit this). It complicated the intervention efforts. One operator testified that the plant had tripped twice before in connection with the condensate polishers. In addition, two weeks before the accident there had been a 'sizeable leak' in the air lines going to the polisher. A pump came on 'inadvertently' about a month before the accident, was bypassed, and was still awaiting repairs at the time of the accident. Three auxiliary feedwater pumps had been taken out of commission two weeks before the accident and left out, in violation of Federal Regulations.

This review indicates that there was not just a single piece of equipment failure that might have been bypassed, but equipment failure (and design problems) on a level that should cause concern even in a less deadly, non-nuclear plant, and the presence of warning signals that were not heeded. But the important point is not that Metropolitan Edison was particularly derelict, but that such a state of affairs is fairly normal in complex industrial and military systems. Ammonia plants, a mature part of the chemical