Normal Accident at Three Mile Island

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 (TMI) is being assessed in this fashion. The industry started a new training program, the equipment at the Babcock and Wilcox 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 one, and the amount of 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 at 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 United States, 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 of the plants existing now and scheduled to come on stream.

Normal accidents emerge from the characteristics of the systems themselves. 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, backup systems, quality equipment, and good training all help prevent accidents and minimize catastrophe, but the complexity of systems outruns all controls.

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 contrasted to preferred ones. I will consider three major categories, although 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 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 unreasonable ex-

0147-2011/81/0701-0004$03.50/1

penditures 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 the main power and backup systems 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, had it occurred two seconds later.

Because liability cannot be assigned, owners and managers cry "unique" when they can. Failing that, they move to the next most desirable category. 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 piece of equipment, 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. Backup systems and isolation devices come into play. 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.

Normal accidents emerge from the characteristics of the systems themselves.

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, 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. 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 Babcock and Wilcox experts at the site. But management prefers the discrete label to the one that suggests the complexity of the system is at fault.

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 can also be prevented (to the extent that accidents ever can) by noting warning signals, by using backup 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."

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 categories of unique or, failing that, discrete accidents. 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.

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 calculation of risk. However, our country was built on risk, as we are hearing lately.

Nuclear proponents are fond of saying that all imaginable risks have been calculated; indeed, they cite this as a major reason for the escalation in plant costs. 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 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 Three Mile Island plant lacked a backup emergency core

cooling system (ECCS) that is required of newer plants, just as the early ones, such as one at Indian Point, New York require *none*. As we shall see, however, it probably would not have made any difference in the TMI accident.

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; but not all the time. 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-10s, after repeated warnings, blowing open three times (the third time 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 an engineer's hypothetical analysis. Why are warnings not always heeded? There are many reasons, and when we consider the overpopulation 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. There was rather a long list of these at Three Mile Island, to be considered later. Each one individually is considered trivial or interpreted in a routine framework. Only 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 triggered them. Or if the operators have no conception of the problem, the alarm may be attributed to faulty alarm equipment. Since dials sometimes 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, form 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—as with the DC-10 cargo doors prior to the fatal crash near Paris. Past accidents also 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 with other parts of its system.

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.

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 Nuclear Regulatory Commission, the reactor builders, numerous institutes, university centers, and industry bodies all function in this capacity. Indeed, in a crisis, the system comes together tightly; it responded exceptionally well to the TMI accident. 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 shut-downs, 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 similar 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, which was considering the purchase of a reactor from Babcock and Wilcox, one quite similar to the two reactors at TMI. Michelson wrote a long memo raising a number of concerns, including a remarkably prescient description of the dynamics of the TMI accident; a

LOCA occurs (a loss of coolant accident), a high-pressure injection system (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 means less cooling of the reactor core and could lead, in minutes, to damage to the core and even a meltdown.

Michelson's report was sent to the NRC in November 1977, 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, Babcock and Wilcox (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.

Alarms, like deviations, always outnumber actual accidents; warnings are in greater demand than actual malfunctions.

Meanwhile, on September 24, 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 nine percent capacity, and in a short time they discovered the cause of the accident—a faulty Pilot Operated Relief Valve (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 accident, 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 do 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 meltdown.) 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? 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. 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. 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 2000 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 affect safety. The NRC has gagged on them; no reasonable system for analyzing them exists. The utilities dutifully report these and they sink into the enormous file. 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? Only if it had been remembered, along with all the other instructions that continually change, and more important, only if the operators had known it was this type of accident they were experiencing. As we shall see, they did not. Even the

experts who were quickly at 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, Combusion Engineering, Babcock and Wilcox, Davis-Besse and TMI and some seventy other plants all wired together into one low resistance circuit, the number of untoward events

Tight coupling encourages normal accidents, with their highly interdependent synergistic systems, but loose coupling muffles warnings.

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 upwind 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 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.

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 literally to 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 tenth 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.

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 notwithstanding. Most things that go wrong do not 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." 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 are insufficient. 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 5000 or so industrial-accident deaths we produce in the United States 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 pilot operated relief valve (PORV). It stuck open. The event was not without prior warnings. There were at least 11 other failures of this key valve at other plants before TMI, including Davis-Besse. 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 easily to 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—in 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 down-time.

> The operators at Three Mile Island literally had to turn off alarms; so many were sounding and blinking that signals passed into noise.

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 accident. 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 a 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 accident, 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 backup system for cooling the reactor by returning liquid 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.

Numerous items were not working at the time of the accident or had failed in the recent weeks. The auxiliary building sump tank had blown a rupture disc some weeks prior to the accident; operators were bypassing 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 "sizable 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 repair 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.

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 industry, had an average failure rate of 10 to 11 shutdowns per year; 50 days of down-time per year; 1 fire per plant every 11 months. The nuclear power industry is extremely safety-conscious, compared to most industrial concerns, but it will still have problems such as these, as the large number of accidents indicate. Equipment failures, like accidents, are normal, though not frequent.

Operator Error

From the beginning it was widely believed that operator error was the fundamental cause of the TMI accident. B&W flatly stated this, as did the British Secretary of State for Energy, who cited the cause of the accident as "stupid errors." This conclusion was attributed to the Nuclear Regulatory Commission by the press. The President's Commission on Three Mile Island, in their final report, blamed everyone, but most particularly the operators. They twice note that "the major cause of the accident was due to inappropriate actions by those who were operating the plant and supervising that operation," though problems of design, training, and procedures contributed to operator failure. But they also feel "they should have known" that they were in a Loss of Coolant Accident, and "failed to realize" that various problems were due to a LOCA, and were "oblivious" for over four hours to the threat of uncovering the core. (A report prepared by outside consultants, the Essex Corporation,

for the Nuclear Regulatory Commission, came to a different conclusion, one that deliberately distinguished the causes of human errors themselves: "The primary conclusion reached on the basis of this investigation was that the human errors experienced during the TMI incident were not due to operator deficiencies, but rather to inadequacies in equipment design, information presentation, emergency procedures and training."

Complex systems are not responsive to warnings of unimaginable or highly unlikely accidents.

It is not comforting that the most blatant operator error at TMI (though not, it is said, an important cause of the accident) is the one least susceptible to remedial action by educational requirements or training programs. After a routine testing procedure, two valves which were closed for the check were left closed rather than opened. Perhaps some people are more likely to lock themselves out of their houses or cars than others, but educational degrees and training would hardly seem to account for the variations. Such things simply happen. Operating personnel testified that with one or two thousand valves in the plant (making checks on every valve every shift unrealistic) one will expect to find one or two out of position for no good reason at times. One operator testified to personal knowledge of two in the previous year, and about five in the short history of TMI Unit 2.

Some valves are so important that they are locked, and a locked valve book is maintained; but an operator testified that it is "sometimes" not kept up to date. A valve that is checked every shift was once found open despite the check at the beginning of the shift. The problem is aggravated by engineering design where, presumably to save money, indicators do not tell whether the valve is actually open or closed, but only the position of the switch that is supposed to open or close it. Such designs create opportunities for operator error. One minor accident at TMI was caused by an operator inadvertently bumping into some switches while investigating a problem. Not only are there a large number of valves, but frequent testing and maintenance routines require them to be placed in non-normal positions for varying periods of time. Valves in wrong positions have caused and contributed to accidents in other nuclear plants. As long as eternal vigilance is a desideratum rather than a reality, the valve position will continue to cause or greatly complicate nuclear plant accidents.

Errors of judgment by operators are more difficult to analyze (and thus more easy to attribute) because the judgment becomes an error only after the fact. Most cases of operator error in normal accidents are "retrospective errors." Presumably many decisions are made that would be classified as errors according to the books or the training programs, but if they work or cause no problems they will be unnoticed, and thus not lead to a revision of standard procedures. If they work, they are in effect being misclassified as errors by virtue of erroneous procedures in the manuals. The cards are stacked in favor of a declaration of operator error, for operators will not be credited with successful actions which violate procedures, but only charged with those that result in investigated accidents. To aggravate the problem, the system and its procedures are not generally under review, only the operators.

More important, though, is the context of judgment errors. A high pressure spike in the reactor was noted because it automatically brought on a safeguard system. But the operator testified, regarding the spike, that "we kind of wrote it off at the time as possibly instrument malfunction of some sort." This was not an unreasonable conclusion, since instruments *were* malfunctioning. "We did not have a firm conclusion" regarding the spike, he went on, since it appeared and went away with such rapidity. Information about the spike was not widely disseminated at this time because it was neither believed nor understood (though a Senate investigation revealed that at least one person drew the correct conclusion at the time). Such is the common fate of novel signals in normal accidents.

The most significant error, by all accounts, was the failure to maintain the high pressure injection (HPI) system. But consider the context, and the matter of organizational routines and goals. Available readings indicated no problem with the level of coolant in the core; at the time it was not even clear what kind of a Loss of Coolant

The warning of an incomprehensible and unimaginable event cannot be seen, because it cannot be believed.

Action (LOCA) this was. These readings were misleading because of vaporization in the core, but that information was not available. There was no direct way to read the water level in the core, and one B&W official was reluctant to encourage having such an indicator because it would increase other problems (a typical interdependency problem in complex systems). Unaware that there was a danger of uncovering the core, the key danger in nuclear plants, the operators focused upon another danger of considerable magnitude—going "solid" in the pressurizer. They were faced with contradictory indicators. (Even a B&W officer, who blamed the accident upon the operator error of cutting back on the HPI system, said the indicators put on "a mysterious performance.") Pressure was low in the core, but it was

thought to be adequately covered, while pressure in the pressurizing vessel was high and getting higher. As the B&W official put it, the pressurizer level should have been going down, but it was coming up, and high pressure injection only *aggravated* it, he added.

The operators did what the Davis-Besse operators had done in a similar accident; they throttled the HPI system back to about half level to prevent going solid. Going solid can cause serious damage to parts of the system, and can easily be avoided by manually overriding the HPI system, and cutting it back. The fear of the operators was shared by some experts at B&W and at least one in the NRC. The reason the Kellog-Dunn-Novack memo was held up and debated so many months was that B&W experts feared that keeping the HPI system on in almost identical circumstances would result in going solid. Subsequent to the accident, experts changed their mind and released the new instructions.

The key problem remains, however. It is not always possible to know just what kind of a LOCA one is in, and when the memo will apply. As one commission member noted, the decision to cut back on HPI has to be taken *before* one can know that it would be the wrong decision, and the B&W engineer testifying agreed. The new instructions may not solve the problem at all, except that they weigh heavily in favor of a more conservative course of action, risking going solid rather than uncovering the core. (In fact, the new instructions proved to be dangerous when followed at another plant several months later, and were revised back to something closer to pre-TMI instructions.)

There were other errors. Operators thought the complex pathways for radioactive wastes led to one tank, but they in fact led to another, which overflowed. The plumbing is so complex that scientists on the President's Commission could not read the tiny details in the chart when trying to trace out parts of the system. A technician was taking a sample to test for radioactivity. The reason they knew the liquid was not radioactive was that he got some of it on his hands, and then they checked *him* for exposure! The operators read on-line display indications of temperatures of 230°, whereas the computer print-out (delayed an hour or so) indicated a much more serious 285°, which would have led to a different course of action. (It seems likely here that the operators misread the indicator because a higher figure was not congruent with the interpretation they were working under and which made most "sense"—a common attribute of normal accident behavior.) A supervisor testified he believed there was significant core damage the morning of the accident (Wednesday), but he did not mention this to anyone else at the plant when he talked by phone or when he came in the next day; other plant personnel (and the arriving experts) reportedly did not reach this conclusion until late Thursday, early Friday, or even the next week in the case of some Metropolitan Edison officials. The supervisor testified it would not have made any difference if people had been aware of significant damage, but this is hard to believe. Significant events such as the pressure spike and extreme thermocouple readings of core temperatures were not communicated to key personnel—because they were simply not believed.

> In normal accidents the particular trigger is relatively insignificant; the interaction is significant.

The most serious case of possible operator error was the decision to vent radioactive gases (producing the puff and plume over the plant that almost triggered evacuation orders). An NRC officer in the control room at the time believed the venting was the automatic result of excess pressure; the supervisor who ordered it said it was an intended venting, with the concurrence of the NRC officer and prior warnings to civil defense personnel. Most people, including some officials of Met. Ed. and B&W and the NRC, believe that after the "puff" the valve was closed—the pressure having been relieved. But the supervisor testified that after the puff the valve was left open for days, since the level of radioactivity fell off rapidly in the next few minutes.

Defending the venting, the supervisor claimed that he was running low on water being used to cool the system. (A B&W official testified there was plenty of water.) He wanted that water in case the core heated up. (This is puzzling, since he was also sure that the core was stable two days before and had remained stable; indeed, everyone was sure.) He did not trust the back-up safety system since its packing and valves might leak if it had to be used to cool the core should the core happen to become unstable. The pressure in the tank had been relieved several times by the previous shift supervisor by brief ventings, and by the present supervisor himself in his shift. Worried about his water reserve, and about the safety system, he decided to try what can only be called a large vent, and when the radioactivity did not continue at a high rate, it became a permanent vent. (Since there are some ambiguities in the published Staff Reports of the Commission, it seems possible that much more radioactive material was released than has been acknowledged. But this is only a possibility.)

Negative Synergy

Of such complexities are normal accidents made. Even in this most studied and documented piece of complex organizational behavior, the testimony is contradictory and the reasoning, elusive. Safety systems are not considered safe; cores are stable but are not considered stable so radioactive venting is risked in a large dose; the supervisor calls civil defense to alert them to the venting and they think he has said that the island is being evacuated, and so on. The closer normal accidents

are studied, the more they reveal their potentials for even greater disasters. This is why, after close scrutiny, one can always say, no matter how serious the accident, "it was just luck that saved it from being worse."

Synergy is a buzz word in business management circles, indicating that the whole is more than the sum of its parts, or in their congenial familarity, two plus two equals five. But minus two plus minus two can equal minus five thousand in tightly coupled, complex, high-risk military and industrial systems. This article has given repeated examples of negative synergy where complex, unanticipated, unperceived, and incomprehensible interactions of off-standard components (equipment, design, and operator actions) threaten disaster.

The observation is hardly novel; engineers frequently test for multiple failures and design against them. But it is significant that in possibly the most dangerous of all our industries, nuclear power generation, there are just two official categories of accidents, simple and complex, and only the first is used in training, since it is impossible to train for the second.

Operator training for accidents is based upon "design-based accidents," that is, those accidents that are anticipated and guarded against through plant design. If one part of the system fails—emergency feedwater, power outage, etc.—a backup system or a means of isolating faulty equipment and bringing other equipment into service is provided for. (This is akin to a "discrete" accident.) What the industry calls a "worst case" accident is one where there are failures not anticipated in the design, and no obvious or tested emergency procedures are available. (This is akin to a normal accident.) Multiple-failure accidents are generally "worst case" accidents, because design-based accidents generally anticipate only one major cause. One author notes that "practically all of the reactor accidents that have occurred in the past have been multiple-failure accidents." Multiple-failure accidents are not simulated in training. The number of possible multiple-cause accidents is nearly limitless.

"Normal" accidents, in my terminology, are largely multiple failure accidents, or "worst case" accidents. They are infrequent, but far from rare. The more complex the system, the more likely they are to occur. They may, of course, have a single source; it is not the case that there have to be two or more *simultaneous* equipment or operator failures. But the single source leads to further events which are unanticipated and often unimaginable.

An indication of complexity at TMI is provided by this quote from one supervisor: "I think we knew we were experiencing something different, but I think each time we made a decision it was based on something we knew about. For instance: pressure was low, but they had opened the feed valves quickly in the steam generator, and they thought that might have been 'shrink.' There was logic at that time for most of the actions, even though today you can look back and say, well, that wasn't the cause of that, or, that shouldn't have been that long."

All operators and supervisors testified to experiencing a very unusual situation, and there are repeated indications that an attempt was made to force these situations into normal, routine explanations—the kind called for by the emphasis upon "design-based accidents." This is the significance of the widely reported comment by the NRC commissioners that if only they had a simple, understandable thing like a pipe break they would know what to do, and the (presumably joking) remark that perhaps they should arrange for one since there were standard procedures for handling it.

Testimony from operators as to why discharges were in one tank rather than another, involving back flows, spillovers, a previously ruptured and unrepaired disc,

The closer normal accidents are studied, the more they reveal their potentials for even greater disasters.

speculation as to the source of water in the sumps, concerns about whether to keep it in the containment building or the auxiliary building, and so on indicate just one part of the system that was difficult to visualize or conceptualize. This was only one of several parts of the system the operators were attempting to deal with and coordinate.

There is unfortunately a good reason for limiting training to design-based accidents rather than normal ones. As the power and size and complexity of plants increase, the permutations will increase geometrically, and so will the perturbations making protection humanly impossible. The recommendation (by Hans Bethe, for one) to marry fusion processes with fisson processes in order to extend the life of current fisson plants will extend the possible perturbations unimaginably.

The confluence of events is not limited to multiple equipment failures, of course. These will interact with expected operator errors. Equipment or design failures are likely to *elicit* operator errors because they are responding to expected, or routine scenarios, and will misinterpret the unexpected signals. There are several instances of not believing the signals in the TMI accident, and given the occasional unreliability of instrumentation, this is to be expected. When faced with ambiguous events and signals, operators can be expected to construct interpretations around familiar readings, dynamics, and routines, and will have the latitude to discount signals and construct interpretations. But of course, it is the novel interactions, the unexpected, the unimagined, that form the basis of a normal accident. Operators, then, are not conditioned to look for the novel

explanation, and training in design failures reinforces the tendency to avert a glance into the unknown. The permutations referred to above make training for novelty, for the normal accident, exceedingly difficult. Prosaic failures—valves left closed, a valve failing to close though the indicator says it has, water in the condensate polisher system—quickly interact to produce the fourth and final characteristic of normal accidents: synergistic effects which are negative for the system and beyond the reach of training and experience.

Future Alternatives

All sorts of things will reduce the risk of discrete and calculated-risk nuclear accidents—revamping the NRC; better operator training, testing, and qualifications; closing of plants near large cities. A meaningful liability system would help. Financial risks from accidents need not be passed on to the rate payer in higher rates, or the public in general through the Price-Anderson Act which limits the liability of the utility and passes it on to the taxpayer. In addition, warnings help. Some designs are demonstrably safer than others. The Navy may pay more attention to quality control than the private businesses that run most of our utility plants.

But normal accidents, whose origins lie fallow and simmer in the very complexity of the interactive system, waiting upon some failure of equipment, design, or operator action to give them their brief, fierce life, cannot be eliminated. Indeed, they grow with the complexity of the system, including the complexity added by the safety features.

Normal accidents are more likely to be perceived in folk expressions than they are in the technical studies of the labs, the NRC journal *Nuclear Safety*, or the literature on the regulatory process and the sins of the old Atomic Energy Commission. The average person, when she resignedly invokes Murphy's law (if anything can go wrong it will), or notes that "for want of a nail the shoe was lost, for want of shoe the horse was lost...," or mutters the ubiquitous blanket phrase "accidents will happen," is closer to the truth than the experts. The dominant theme for the experts is the accident that can be prevented by design—"design-based accidents" as they are perversely called. This is what the literature covers. A newer, subdominant theme, popularized by the President's Commission, is the "man-machine" interface, and the lack of attention to the "man." But neither theme is responsive to the key characteristic of tightly coupled, complex, high-risk systems that we pride ourselves on. Synergistic effects of a negative nature are bound to occur. Warnings will not prevent them, nor training, nor equipment and design changes, and intervention is limited.

The defenders of nuclear power are correct that "no energy source can be completely safe." The President's Commission agreed; the only thing to do is to make the risk tolerable. But the degree of risk and the level of toleration have not been tested. The industry is young. The catastrophic potentials—e.g., 100,000 immediate deaths, a poisoned land—have not been given a fair chance to be realized. Unanticipated, unpreventable, incomprehensible, uncontrollable accidents in high-risk systems are sufficiently rare to give us another five or ten years of grace. Then we shall see what is tolerable.

There are always alternatives to systems with catastrophic potential. Periodic, low-cost flooding of sparsely settled flood plains adjacent to rivers is less costly and dangerous than huge dams that tower over densely settled flood plains that people come to consider safe. The argument that "we must have dams" is insubstantial; we can live elsewhere. Our energy crisis does not require building Liquified Natural Gas ships that are the length of three football fields, with control panels that have to be as complicated as those in a nuclear plant. Most of the toxic substances we inevitably spew about are not essential to our lives, but only to private profits or war machines. At the very least let us consider including the externalities, real and potential, in the costs of these goods and services. Thus, for example, each propane gas truck in a city should pay insurance to cover the cost of blowing up a few office and apartment buildings when the truck crashes, the gas flows into the sewers for two or three minutes and spreads, and then is ignited by someone's cigarette near the scene of the accident.

Nuclear plants produce about twelve percent of the U.S. electrical power now, but we have over twice as much excess peak-demand power standing by in non-nuclear facilities that could be put to use instead. The potential of both conservation and the various forms of solar energy was completely unexpected by most experts and officials.

Decentralized, loosely coupled, low-risk alternatives abound. In these, a serious normal accident does not lie incubating. We get only the irritating but tolerable foulups that plague our daily life and our organizations.□

Low-risk alternatives to nuclear energy abound; in these, a serious normal accident does not lie incubating.

Charles Perrow is professor of sociology at the State University of New York at Stony Brook. An organizational theorist, he has published four books and numerous articles, including a revised edition of his widely used text, Complex Organizations: A Critical Essay. *He prepared a paper for the President's Commission on the Accident at Three Mile Island on the organizational aspects of the accident.* *The Social Science Research Council, the University Awards Committee of the State University of New York, and the National Science Foundation provided support for the preparation of this article.*