14

Nuclear Reactor Safety

14.1 General Considerations in Reactor Safety

14.1.1 Assessments of Commercial Reactor Safety

The historical record of nuclear reactor performance can be interpreted as showing that they are very safe or that they are very dangerous. The former conclusion follows if one limits consideration to plants outside the former Soviet Union (FSU). The latter conclusion follows if one focuses on the Chernobyl accident and takes it as a broadly applicable indicator.

For commercial reactors in the non-Soviet world, which account for the largest part of the reactor experience, the safety record is excellent. As of the end of 2003, these reactors had a cumulative operating experience of about 10,100 reactor-years, of which about 2870 reactor-years were logged by U.S. reactors.[1] There has been no accident in any of these reactors, including the 1979 Three Mile Island (TMI) accident, that has caused the known death of any nuclear plant worker from radiation exposure or that has exposed any member of the general public to a substantial radiation dose.

If one goes beyond Western commercial reactors, there are three exceptions to this excellent record. Two involved reactors built for military purposes and are sometimes overlooked—the 1957 Windscale accident in a British plutonium-producing reactor that led to some significant exposures and the 1961 SL-1 accident in the United States in which three army technicians died. (These accidents are described briefly in Section 15.1.) The third was much

[1] The number of reactor-years is extrapolated from December 31, 2002 data [1, Table 7]. It includes the contribution from commercial reactors that are no longer in operation.

greater in impact and has received far more attention: the Chernobyl accident in the Soviet Union in 1986 (see Section 15.3).

The Chernobyl reactor was graphite moderated, with a number of unusual design features, and the circumstances of that accident could not be repeated in the standard LWRs and HWRs used outside the FSU. Nonetheless, no reactor has a truly zero chance of an accident and Chernobyl demonstrated that a major reactor accident could potentially impact hundreds of thousands of people. For this reason, high importance is attached to issues of reactor safety by proponents and opponents of nuclear power alike.

Assessments of reactor safety involve estimates of both the probability and severity of accidents. In the remainder of this chapter, we will explore some of the general issues involved in achieving and evaluating nuclear safety. In the following chapter, we will look at the failures, cases where accidents did, in fact, occur. We will be interested in both their causes and consequences.

14.1.2 The Nature of Reactor Risks

Categories of Reactor Accidents

There is a large spectrum of possible consequences from a nuclear reactor accident. The most serious accident is one in which there is a large external release of radionuclides, as was the case at Chernobyl. Less harmful, but still serious, are accidents in which there is damage to the reactor core, but with no appreciable release of radionuclides to the outside environment, as at Three Mile Island. The spectrum can be extended downward to include everything from near misses to harmless breakdowns that have no actual or likely adverse consequences. These lesser mishaps are of interest primarily because of the cost of the remedial measures and lost time, and for the light they shed on the probability of more serious accidents (see Section 14.4.3 on precursor analyses).

Potential major nuclear reactor accidents fall into two main categories, each illustrated by one of the two major past accidents in power reactors, the Chernobyl and Three Mile Island accidents:[2]

- *Criticality accidents.* These are accidents in which the chain reaction builds up in an uncontrolled manner, within at least part of the fuel. In an LWR of normal design, such accidents are highly improbable, due to negative feedbacks and shutdown mechanisms. They are less unlikely in some other types of reactor, given sufficient design flaws. The 1986 Chernobyl accident was a criticality accident, although much of the energy release was from a steam explosion following the disruption of the core.
- *Loss-of-coolant accidents.* When the chain reaction is stopped, which can be accomplished quickly in the case of an accident by inserting control

[2] The Windscale accident does not fit into either of these categories (see Chapter 15).

rods, there will be a continued heat output due to radioactivity in the reactor core. Unless adequate cooling is maintained, the fuel temperature will rise sufficiently for the fuel cladding and the fuel to melt, followed by the possible escape of radioactive materials from the reactor pressure vessel and perhaps from the outer reactor containment. The TMI accident was a loss-of-coolant accident. There was substantial core melting, but no large escape of radioactive material from the containment.

With appropriate precautions, such as the assurance of intrinsic negative feedback and the capability for rapid insertion of control rods, a criticality accident is virtually impossible in a well-designed reactor. Therefore, almost all of the attention to reactor accidents in the United States and elsewhere is directed to the more demanding task of avoiding a loss-of-coolant accident.

In the light of possible misapprehensions, it is worth noting that a bomblike nuclear explosion cannot occur in a nuclear reactor. In a bomb, a critical mass of almost pure fissile material (^{235}U or ^{239}Pu) is brought together violently and compressed by the force of a chemical explosion, and the chain reaction develops fully within one-millionth of a second—quickly enough for much of the fuel to fission before the mass is disassembled (see Section 17.2.3). In a reactor, most of the mass is not fissile. Even in the fuel, the fissile mass is small compared to the ^{238}U mass.[3] A reactor also contains a great deal of other nonfissile material in the form of coolant, moderator (if there is a moderator other than the coolant), fuel cladding, and metal support structures.

The presence of the nonfissile material has two consequences that are pertinent to the issue of explosions: (1) The multiplication factor k in a reactor is close to unity, whereas in a bomb it approaches 2, and (2) the average time between fission generations (the mean neutron lifetime l) is greater in a reactor than in a bomb, because the most frequent neutron reactions in a reactor are elastic or inelastic scattering, not fission. As a result, the chain reaction builds up much more slowly in a reactor than in a bomb [see Eq. (7.15)].

Overall, the first "line of defense" against an explosion in a reactor is the negative feedback that prevents criticality accidents. This should suffice. However, if there are mistakes in the design or operation of the reactor and the chain reaction reaches too high a power level, there is time for the ultimate "negative feedback" to come into play—the partial disassembly of the reactor core, which stops the chain reaction after only a relatively small amount of energy has been produced (i.e., only a small fraction of the nuclei have fissioned). This is what happened in the Chernobyl accident, where most of the energy of the explosion came from chemical reactions, including steam interacting with hot metal (see Section 15.3.2). Such an accident can be very serious, but the consequences are not on the scale of the consequences of a nuclear explosion.

[3] The fission cross section for neutrons colliding with ^{238}U is small for neutron energies below 2 MeV and is negligible below 1 MeV (see Section 6.2.3).

Aftermath of a Reactor Accident

Nuclear accidents pose particular problems because of the persistent effects of radioactivity. The heat output immediately after the reactor is shut down is about 7% of the thermal output of an operating reactor (see Section 14.2.2). Although the activity and energy release fall rapidly with time, a serious accident can occur if this heat is not removed by the cooling system. At Three Mile Island, this continued production of heat led to the fear that the accident might progress further, with the release or ejection of radioactive material from the reactor containment. At Chernobyl, there was a very large release of radioactive material, and the dispersed debris has created problems that will last for many years.

This may be contrasted with the situation in many other sorts of accident (with important exceptions, such the Bhopal accident in India in 1984). Once a dam breaks or a natural gas facility explodes, the damage is done and society feels moderately secure in coping with the aftermath. There may be more immediate fatalities than in a nuclear accident,[4] but when the accident is over, it is usually deemed to be over, and there is little investigation of possible lingering consequences. With nuclear accidents, serious consequences may persist for a long period of time—in particular, cancers caused by both the initial exposure and the continuing exposures due to radionuclides deposited on the ground.

These factors, plus less well-defined but widely held fears, put nuclear accidents in a special category of societal concern and make it particularly urgent that they be avoided. There can be debates as to the effect of an accident on the health of the public. There is no doubt, however, that each nuclear accident has been something of a disaster for the nuclear industry.

14.1.3 Means of Achieving Reactor Safety

General Requirements

Underlying the approach to safety, for any sort of equipment, are high standards in design, construction, and the reliability of components. In nuclear reactors, concern about possible accidents has led to particularly intense efforts to achieve high standards. Individual components of the reactor and associated equipment must be of a codified high quality. As described in an OECD report:

> In the early years of water reactor development in the USA, a tremendous effort was put into development of very detailed codes and

[4] For example, explosions in liquid-natural-gas tanks and the associated fires killed 130 people in Cleveland, Ohio, in 1944 and 40 workers on Staten Island in New York, in 1973 [2, p. 162]; in each case, the casualties exceeded the prompt fatalities at Chernobyl (see Section 15.3).

> standards for nuclear plants, and these were widely adopted by other countries where nuclear plants were initially built under US licenses. [3, p. 62]

The efforts of the United States have since been supplemented by parallel efforts by other countries and the International Atomic Energy Agency (IAEA). In parallel, a nuclear reactor safety philosophy has developed which includes a number of special features, as summarized in the succeeding subsections.[5]

Passive or Inherent Safety

A distinction is made between *active* and *passive* safety systems. An active safety system is one that depends on the proper operation of reactor equipment, such as pumps or valves. For example, active safety systems include the pumps and valves that control the water supply for emergency core cooling and the motors used to insert control rods in emergency shutdowns. Passive safety features are aspects of the system that are arranged to come into play automatically, without the action either of the operators or of mechanical devices that might fail. The gravity-driven fall of a control rod is a passive feature, although its purely passive character would be compromised if the release of the rods is initiated by an active system.

The terms "passive safety" and "inherent safety" are often used interchangeably, although some authors may intend a difference in meaning or nuance. These terms suggest that the safety of the reactor will depend on immutable physical phenomena rather than on the proper performance of individual components or correct actions by reactor operators. For example, if the thermal expansion of the reactor core provides a negative feedback, the expansion provides an inherent safety feature. In the extreme version of the concept, in a passively safe reactor all operators could become incapacitated and all external electricity and water could be shut off, and still the reactor would turn itself off in the case of an accident and gradually cool with no damage.

This terminology is widely used but has also been criticized. The objections have had several strands:

- Inherent or passive safety is a matter of degree, rather than a totally new departure. A negative temperature coefficient or a negative void coefficient is a passive safety feature and, therefore, most existing reactors already have passive safety features.
- The terms are misleading in that they seem to suggest that an accident would be *totally* impossible, whereas, in fact, one can find circumstances in which any given reactor might fail if arbitrarily improbable scenarios are permitted.
- The terms could appear to have a prejudicial aspect because they could seem to suggest that existing reactors are *not* safe.

[5] The discussion loosely follows the organization used in Ref. [4, pp. 9ff].

The criticisms have had some force, and to defuse them alternative words have sometimes been suggested [5]. However, whatever words are used or caveats included, the concept is clear: It is safer to rely on basic physical phenomena (e.g., gravity or thermal expansion) rather than on the consistently good performance of equipment and operators.

Redundancy

The likelihood of any sort of accident can be reduced by redundancy, which can be achieved in a number of ways:

- *Identical units of the same type.* Often, more than one pump or motor is provided to perform a given safety task, although it is only necessary that one of these operates properly. It is particularly important in such cases to avoid common-mode failures, in which one failure could simultaneously disable all of the units. To achieve this, among other demands, there must be adequate physical separation between the units and between the control systems for them.[6]
- *Diverse types of systems.* An example of diversity in reactor safety design is the provision of different types of emergency core-cooling systems, which act independently.

Defense-in-Depth

A special kind of redundancy is sometimes singled out as being "at the heart of nuclear safety." This is the reliance on *multiple barriers* or *defense-in-depth*, which is described as "a hierarchically ordered set of different independent levels of protection" [6, p. 109]. The principle of defense-in-depth is seen in considering the barriers that prevent or minimize exposures due to the release of radioactivity from a reactor:

- The UO_2 fuel pellets retain most radionuclides, although some gaseous fission products (the noble gases and, at elevated temperatures, iodine and cesium) may escape.
- The zircaloy cladding of the fuel pins traps most or all of the gases that escape from the fuel pellets.
- The pressure vessel and closed primary cooling loop retain nuclides that escape from the fuel pins due either to defects in individual pins or, in the case of an accident, overheating of the cladding.

[6] After the Browns Ferry fire in 1975, it was recognized that multiple wiring systems, intended for redundancy, were carried in the same cable trays and, therefore, were all disabled at the same time. A simple solution is to use different paths for redundant cabling.

- The heavy outer reactor containment, with its associated safety systems, is designed to retain radionuclides that escape through the cooling system or, in the case of a very severe accident, from the pressure vessel.
- If these systems all fail and there is a significant release of activity to the outside environment, the population can be partially protected through evacuation. However, if radiation escapes the containment, then the system has been defeated even if evacuation reduces the damage.

Steps taken to avoid the overheating of the fuel—in particular, the standard and emergency cooling systems—as well as systems to suppress overpressurization of the containment can also be considered to be part of the defense-in-depth.

These barriers against radiation exposures have been put to a severe test in only two instances. In the TMI accident, the reactor containment was highly successful. In the Chernobyl accident there was no containment, as the term is understood in Western design practice, and there was a massive release of radioactive material to the outside surroundings. Subsequent emergency evacuations reduced the exposure of people in the evacuation zone, but there was substantial exposure of the public nonetheless (see Section 15.3).

Defense-in-depth and the various forms of safety redundancy represent a sophisticated version of the view that although it is likely that *something* will go wrong, it is highly unlikely that *everything* will go wrong. Illustrating the value of redundancy (reiterating a point made in Section 12.3.2), if the causes of the failures are uncorrelated, three independent barriers that each have a 1% chance of failure provide a system in which there is only one chance in one million of overall failure.[7]

14.1.4 Measures of Harm and Risk in Reactor Accidents

The most fundamental harm in reactor accidents is that caused by radiation exposures. The extent of the harm can be alternatively measured in terms of individual radiation exposures, the collective population exposure, the number of prompt fatalities caused by intense exposures, or the number of latent cancers caused by lower radiation doses. Of these, prompt fatalities represent the most dramatic and least ambiguous effect. However, the greatest predicted health consequence is latent cancer fatalities (i.e., the eventual cancer deaths expected to occur due to radiation exposures). The doses might be received

[7] Another way of looking at reactor safety, also sometimes termed "defense-in-depth," is to divide it into phases of accident avoidance, accident correction or protection, and accident mitigation (e.g., Ref. [7, p. 339]). Avoidance is achieved by proper design, maintenance, and operation. Accident correction is achieved by reliable safety systems that, for example, shut the reactor down promptly and alert the operators. Accident mitigation is achieved by, for example, restoration of lost cooling, an effective containment system, and, as a last ditch measure, evacuation of the immediately surrounding population.

mostly in the first few days or in the first year following the accident, but the cancer fatalities would appear over many decades, generally starting after a latent period of 10 years.

Other harm includes physical damage to the reactor plant and contamination of the surrounding environment that may force the evacuation of large regions. Plant damage was clearly the most important direct consequence of the TMI accident and ground contamination was a major, perhaps in the end *the* major, consequence of Chernobyl.

Reactor accident risks are often analyzed in terms of the probability of two defining aspects of reactor accidents. One is the probability of reactor *core damage*—in particular, the melting of part of the core. The other is the probability of a *large radiation release*, stemming from the failure of the barriers provided by the reactor pressure vessel and the reactor containment.

The distinction between these consequences is illustrated by the TMI accident where, as discussed in Chapter 15, there was surprisingly little release of radioactive material to the environment outside the reactor containment although the damage to the reactor fuel assemblies was great. This focused attention on the accident *source term*—the inventory of radionuclides released to the outside environment, as distinct from the inventory of radionuclides in the fuel.

For ^{131}I and other iodine isotopes, the source term at Chernobyl was essentially the total initial core inventory. At TMI, it was close to zero: 18 Ci out of 64 million Ci, as reported in an American Nuclear Society study [8, pp. 1–12]. The low release of iodine at TMI was the result of the fact that there was much more cesium than iodine in the core inventory and the iodine predominantly formed cesium iodide (CsI), rather than the volatile gas I_2 (see, e.g., Ref. [8, pp. 8–9]).[8] The CsI was then trapped by dissolution in water or deposition on surfaces.

A crucial question is whether the very good performance of the containment system is generic to all LWRs or was peculiar to TMI. There have been a number of studies of this matter—for example, studies carried out under the auspices of the American Nuclear Society (ANS) [8] and the American Physical Society (APS) [9]. These studies concluded that, in most cases, the source term will be substantially less than the core inventory. If the source term is sufficiently low, then there is no "large release" of radionuclides.

Although there is no single indicator of reactor safety, in practice the most significant measure may be the core damage probability. Any instance of core damage at least raises the possibility of a significant radiation release and would inevitably deeply concern the public. Further, even with no release of activity outside the reactor containment, the cleanup expense after the core is damaged would be punitively expensive for the utility. Thus, much of the

[8] The differences in abundances results from the continual decay of ^{131}I ($T = 8.02$ days) during the months of reactor operation, whereas ^{137}Cs ($T = 30$ yr) kept increasing in amount.

efforts in assessing and reducing reactor risks focuses on the possibility of core damage.

14.2 Accidents and their Avoidance

14.2.1 Criticality Accidents and Feedback Mechanisms

General

In normal operation of a thermal reactor, prompt criticality is avoided. The reactivity of the system is kept low enough to make delayed neutrons crucial for criticality. Thus, even if the reactivity rises, the rates of increase of the neutron flux and of the power output are relatively slow. The magnitude of any power excursion is limited in an appropriately designed reactor by inherent negative feedbacks that come into play automatically. This gives time for the insertion of control rods, which have high neutron-absorption cross sections and will terminate the chain reaction. We consider below two major feedback mechanisms that enhance reactor safety.[9] Unless otherwise indicated, it will be assumed that the reactor considered is a standard LWR.

Fuel Temperature Feedback: Doppler Broadening

Although we have been tacitly treating the nuclei of the fuel as motionless targets undergoing bombardment by neutrons, this is not a precise description. The uranium nuclei are in thermal motion, with an average speed that increases as the temperature increases. The result is to increase the effective cross section for neutron absorption in ^{238}U if the temperature of the fuel rises, through the Doppler broadening of the absorption resonances (see Section 5.2.3). The number of neutrons available for fission is reduced, and the reactivity and the reactor power output decrease.[10] This negative feedback comes into play quickly, reversing the rise in power output as soon as the fuel temperature rises.

However, the fuel temperature feedback is not automatically negative in all types of reactors. If a fuel has relatively little ^{238}U and is primarily made of fissile material, then the main effect of Doppler broadening is to increase the rate of fission at nonthermal energies, giving a positive feedback. Thus, to keep the fuel temperature feedback negative, the fraction of fissile fuel in liquid-metal fast breeder reactors is kept below 30% [7, p. 146].

[9] This is not intended as a full listing of feedback mechanisms. Additional ones exist, both positive and negative (see, e.g., Ref. [7, pp. 145ff]), and must be taken into account in reactor design.

[10] In terms of the four-factor formula [Eq. (7.5)], the resonance escape probability, p, is reduced.

Void Coefficients

In an LWR, water is essential for moderating the reaction. If the water is removed (e.g., if there is a pipe break and insufficient replacement water is provided), the moderation will be inadequate and the reactivity will drop, because with less thermalization, there will be more loss of neutrons through absorption in ^{238}U. More voids also mean a greater escape of neutrons from the reactor. Loss of water in the reactor vessel is the limiting case of a "void."

The term *void coefficient* is usually applied to the replacement of liquid coolant by bubbles. The void coefficient is defined as the ratio of the change in the reactivity to the change in the void fraction. A negative void coefficient means that the reactivity decreases as the volume of steam bubbles increases (i.e., the void fraction increases). The loss of water leads to two effects which contribute to a negative void coefficient: (1) less effective moderation (i.e., relatively less elastic scattering of neutrons by hydrogen) and therefore increased resonance absorption of neutrons in ^{238}U and (2) more leakage of neutrons from the reactor.[11] A negative void coefficient corresponds to a negative feedback in accident situations, because the void fraction rises as the power level rises.

However, water also acts as an absorber of slow neutrons, and too much water leads to too much absorption (a low thermal utilization factor f). Were this the dominant effect, then the void coefficient would be positive. Thus, there is a competition between the moderating and absorbing roles of water, with opposite feedback signs. When the void coefficient is negative, the reactor is *undermoderated*; when it is positive, the reactor is *overmoderated*.

For BWRs, in which steam and water are both present, an increase in the steam content corresponds to less water. The moderating role is more important than the absorbing role, and an increase in steam content decreases the reactivity. Thus, the void coefficient is always negative for BWRs. In PWRs, there is usually no direct void coefficient, but thermal expansion of water has the same general effect of reducing moderation and providing a negative feedback.

The situation is more complicated for water-cooled graphite-moderated reactors, and the sign of the feedback can go either way depending on the relative amounts of water and graphite. The role of the water as moderator is less important, and the main effect of the water (aside from the intended function of cooling) can be to absorb neutrons. Loss of this water, by conversion to steam or otherwise, can increase the reactivity (i.e., the void coefficient is positive). This was the situation at Chernobyl. However, this is not intrinsic to all water-cooled graphite reactors. In particular, the N reactor formerly operating at Hanford had a negative void coefficient.

In sodium-cooled fast breeder reactors, the sodium plays only a small role as a moderator, but this moderation acts to lower the reactivity, because

[11] Referring to the five-factor formula [Eq. (7.4)], these feedbacks correspond to a lower resonance escape probability p and a lower nonleakage probability P_L.

the fission cross section increases with energy for neutron energies in the neighborhood of 1 MeV. Thus, the thermal expansion of the sodium or the development of bubbles reduces the moderation, increases the average energy in the neutron spectrum, and increases the reactivity. At the same time, with less sodium in the path of a potentially escaping neutron, more neutrons can escape from the reactor, reducing the reactivity. Overall, these competing effects may leave a sodium-cooled reactor with a positive void coefficient and it is important that there be counterbalancing negative feedbacks.

14.2.2 Heat Removal and Loss-of-Coolant Accidents

Decay Heat from Radioactivity

The central problem in loss-of-coolant accidents arises from the need to remove the heat produced by radioactivity during the period after reactor shutdown. The magnitude of the initial rate of heat generation can be understood in terms of the total energy release in fission, as discussed in Section 6.4.2. On average, for each fission event, about 7.8 MeV is released in beta decay and 6.8 MeV in accompanying gamma decay, for a total of 14.6 MeV out of about 200 MeV (i.e., approximately 7% of the total energy release). Strictly speaking, this result is applicable only when equilibrium has been reached between the production of radionuclides and their radioactive decay. However, the initial activity is dominated by short-lived radionuclides with half-lives of several days or less. Thus, if a reactor has been operating at full power for, say, a month, the total activity reaches a value close to its equilibrium level.

The activity just after shutdown is the same as the activity just before shutdown (treating shutdown as essentially instantaneous), and the initial thermal output from radioactive decay is 7% of the thermal output of the reactor during normal operation, or about 20% of the electric output. Thus, at shutdown of a 1000-MWe reactor, the heat output is initially about 200 MW. It drops to about 16 MW after 1 day and about 9 MW after 5 days [10, p. S23]. Without cooling, these heat production rates are sufficient to melt the fuel.

Core-Cooling Systems

During normal operation, reactor cooling is maintained by the flow of a large volume of water through the pressure vessel. This flow can be disrupted by a break in a pipe, failure of valves or pumps, or, in PWRs, a failure of heat removal in the steam generators. Such accidental disruptions of the normal cooling system are generically termed loss-of-coolant accidents (LOCAs). To guard against the overheating of the fuel in a LOCA, light water reactors have elaborate emergency core-cooling systems intended to maintain water flow to the reactor core.

A distinction is sometimes made between large and small LOCAs. The prototypical large LOCA is a break in the pipes carrying the primary cooling water to the reactor. In a large break, the pressure in the reactor vessel will be lost and a large amount of water will escape. The emergency core-cooling system (ECCS) then comes into play. Initially, replacement water is delivered from "accumulators" driven by nitrogen gas under pressure. Later, low-pressure pumps can provide additional water from external supplies. A large LOCA would be a dramatic event, and much of the early concern about reactor safety focused on preventing such an accident and, if prevention failed, assuring an effective and independent ECCS.

A small LOCA may occur from a leak in the primary cooling loop or, as was the case for the initiating event in the TMI accident, from a problem in the secondary cooling loop. Loss of secondary flow means that heat cannot be removed in the heat exchanger from the primary loop. In such an event, the pressure in the reactor vessel may not be relieved, and it may be difficult to establish the flow of replacement water in the complex hydraulic environment created by the mixture of steam and water at high pressures. To cope with such circumstances, the ECCS has a high-pressure injection system to provide replacement water to the reactor vessel.

The effectiveness of the ECCS for both large and small LOCAs has been the subject of many studies, starting before and intensifying after the TMI accident. In addition to calculations and theoretical analyses, there have been extensive tests, particularly the loss-of-fluid test (LOFT) program at the Idaho National Engineering Laboratory. This program was carried out from 1978 to 1985 and involved simulated accidents on a specially built 50-MWt test reactor. This was an NRC facility, but tests were also carried out there for the Nuclear Energy Agency of the OECD. Analyses of the results of these tests of system performance under simulated accident conditions have led to improvements in equipment and procedures (see, e.g., Ref. [3, pp. 39–42]).

Release of Radionuclides from Hot Fuel

If either the normal or emergency core-cooling system operates properly, there will be no damage to the reactor core in case of a reactor malfunction and no concern about release of radionuclides. However, if the cooling system fails to keep the cladding temperatures low enough to avoid melting, radionuclides will escape into the pressure vessel and into the primary cooling system.

The radionuclides include both fission products and actinides. They can be grouped according to differences in their volatility. The most volatile are the noble gases. These can diffuse out of the fuel into the fuel pins even at normal fuel temperatures. As the fuel temperature rises, damage to the fuel and the cladding causes release of additional elements, the most volatile of which are iodine and cesium. Some other radionuclides, in contrast, are quite refractory and are not released in substantial amounts even under extreme circumstances.

Thus, in one hypothetical accident, presented as an example in an NRC study, the median fission product release from the fuel rods was close to 100% for the noble gases, 6% for iodine, 1% for cesium, and 0.05% for strontium [11, p. A-34]. Although these particular values cannot be taken as precise measures of what would happen in a specific actual accident, they illustrate the main trends.

Although we have emphasized transport through the cooling system as the main avenue for radionuclide release, as was the case at TMI, there are other possibilities. In one extreme case, molten reactor fuel might settle in the bottom of the reactor vessel, melt through the vessel wall, and penetrate into the concrete base below. This scenario is sometimes referred to as the "China syndrome." The main consequence is not as extreme as the name might suggest. It comes from the generation of gases (such as CO_2 and others) in the interaction between the molten fuel and the concrete. This could produce an aerosol that carries nonvolatile radionuclides out of the fuel and into the atmosphere of the containment.

If radionuclides escape from the cooling system or from the reactor vessel, the next barrier is the containment structure. The integrity of the containment can be compromised by overpressure, most likely from the buildup of steam. To avoid this, there are containment cooling systems, either passive or active, intended to condense the steam. For example, PWRs commonly have spray systems for condensation, and BWRs have pools of water for pressure suppression. Some units also have refrigeration units. It is also possible, in the case of an excessive buildup of pressure, to release gas from the containment through valves, with filters to remove radionuclides.

14.3 Estimating Accident Risks

14.3.1 Deterministic Safety Assessment

One approach to establishing and evaluating reactor safety is to establish strict criteria for reactor design and construction and to analyze the behavior of the resulting system for a variety of postulated failures. The more demanding of these failure scenarios are termed *design basis accidents*. The reactor performance is studied through experiments and computational models to investigate whether the safety systems are adequate to cope with a design basis accident. For example, one can postulate a break in a cooling system pipe and then examine whether the emergency core-cooling systems will provide alternative cooling.

This straightforward approach is called *deterministic safety assessment* and it is useful in establishing and verifying design criteria for the reactor. A limitation of the approach is that it does not address the question of likelihoods. In particular, it does not consider the probability that the design basis accident will occur or the probability that the safety system will work as

intended. Obviously, a particular sequence of events is more serious if the initiating problems are relatively probable and the safety systems have a relatively high probability of failing.

14.3.2 Probabilistic Risk Assessment

PRA Implementation: Reactor Safety Study, WASH-1400

Estimating risk probabilities is not an easy matter in the case of nuclear reactors. For automotive safety, by contrast, it is relatively easy to answer questions about the chances of a fatal accident. One merely has to look at the annual fatality rate, subdivided, if one wishes, by type of car, road conditions, driver, and so forth. There are ample data on auto fatalities, and these lend themselves to extensive analysis. Thus, there is reasonable quantitative knowledge of the safety of automobiles and roads.

With no fatal accidents and with no major accidents of any sort in light water reactors other than the Browns Ferry (1975) and Three Mile Island (1979) accidents, overall reactor safety cannot be determined from direct accident experience.[12] (Of course, it would be unacceptable to have enough accidents to provide meaningful statistics.) Instead, it is necessary to rely on calculations or assessments. An early effort in this direction was an 1957 Atomic Energy Commission study (known as WASH-740) on the possible consequences of an accident, but for many years, there was no careful estimate of the *probability* of an accident. A major expansion of nuclear power was expected in the 1970s in the United States and throughout the world, but although there were many intuitions as to the level of risk, there was no defensible quantitative analysis.

To address this issue, the Atomic Energy Commission sponsored an extensive study under the direction of Norman Rasmussen of the Massachusetts Institute of Technology. This study was issued in draft form in 1974 and in final form in 1975 under the institutional sponsorship of the Nuclear Regulatory Commission—which by this time had assumed the regulatory functions of the disbanded AEC. The study is variously referred to as the Rasmussen report, the Reactor Safety Study (RSS), and WASH-1400 [12]. It was the first major study to combine in one analysis the probability and consequences of accidents, in order to assess the *risk* associated with reactor accidents. It was a limited study in that only one PWR and one BWR were analyzed in detail, although the results were often taken to be representative of the situation for other PWRs and BWRs.

The RSS was controversial from the moment the first draft appeared, and the controversies were never fully resolved. However, it is generally agreed that the study made a very important contribution in pioneering the application of methods of *probabilistic risk assessment* (PRA) to the analysis of

[12] See Chapter 15 for a discussion of these accidents.

nuclear reactor safety. In later terminology, especially in international usage, this approach has also been called *probabilistic safety assessment* (PSA). The terms are often used interchangeably.[13]

In principle, this approach permits an objective estimate of the *absolute* risk of accidents, although, at present, it is widely believed that the absolute PRA numbers have large uncertainties. However, even if the data and analyses fail to establish the absolute risks precisely, they can be useful in suggesting the *relative* risks of different configurations and in pinpointing weaknesses. There are some who argue that the chief value at present of probabilistic risk assessments is in identifying places where safety improvements are needed. In this view, PRAs are more useful for improving reactor safety than for estimating it.

Although improvements in reactor equipment and advances in analysis techniques have made the detailed numerical results of the RSS obsolete, they remain of historical significance, and the report itself remains a historic milestone. Subsequent to the TMI accident, numerous steps have been taken to improve reactor safety as well as to refine the analyses, with separate analyses carried out for individual reactors. The general methodology employed in the RSS has been retained.

Event Trees and Fault Trees

The PRA tools used in the RSS were event-tree analyses and fault-tree analyses. In an event-tree analysis, one imagines the occurrence of some initiating event and traces the possible consequences. We illustrate in Figure 14.1 the event tree for studying the consequences of a major pipe break, following which the emergency core-cooling system (ECCS) must operate successfully for damage to be avoided [12, Main Report, p. 55]. The worst case in this example would be the electric power failing to operate, the ECCS not functioning, the fission product removal systems within the containment not operating, and the containment integrity being breached.

The probability that everything goes wrong in this sequence is shown in the bottom leg of the "basic tree" in Figure 14.1. It is the product of five individual failure probabilities. In the "reduced tree," shown in the bottom part of Figure 14.1, cognizance is taken of the possibility that the probabilities are not independent. In particular, the bottom leg of the reduced tree, which bypasses three steps, is based on the assumption that without electrical power, the other systems will also fail and the accident will proceed to the breaching of the containment.[14]

[13] For example, the NRC describes the analysis in its study NUREG-1150 as a PRA [11], whereas the (American) chairman of INSAG terms this study a PSA [13, p. 50].

[14] Figure 14.1 is a simplified version of the event trees that are actually used and is shown for illustrative purposes.

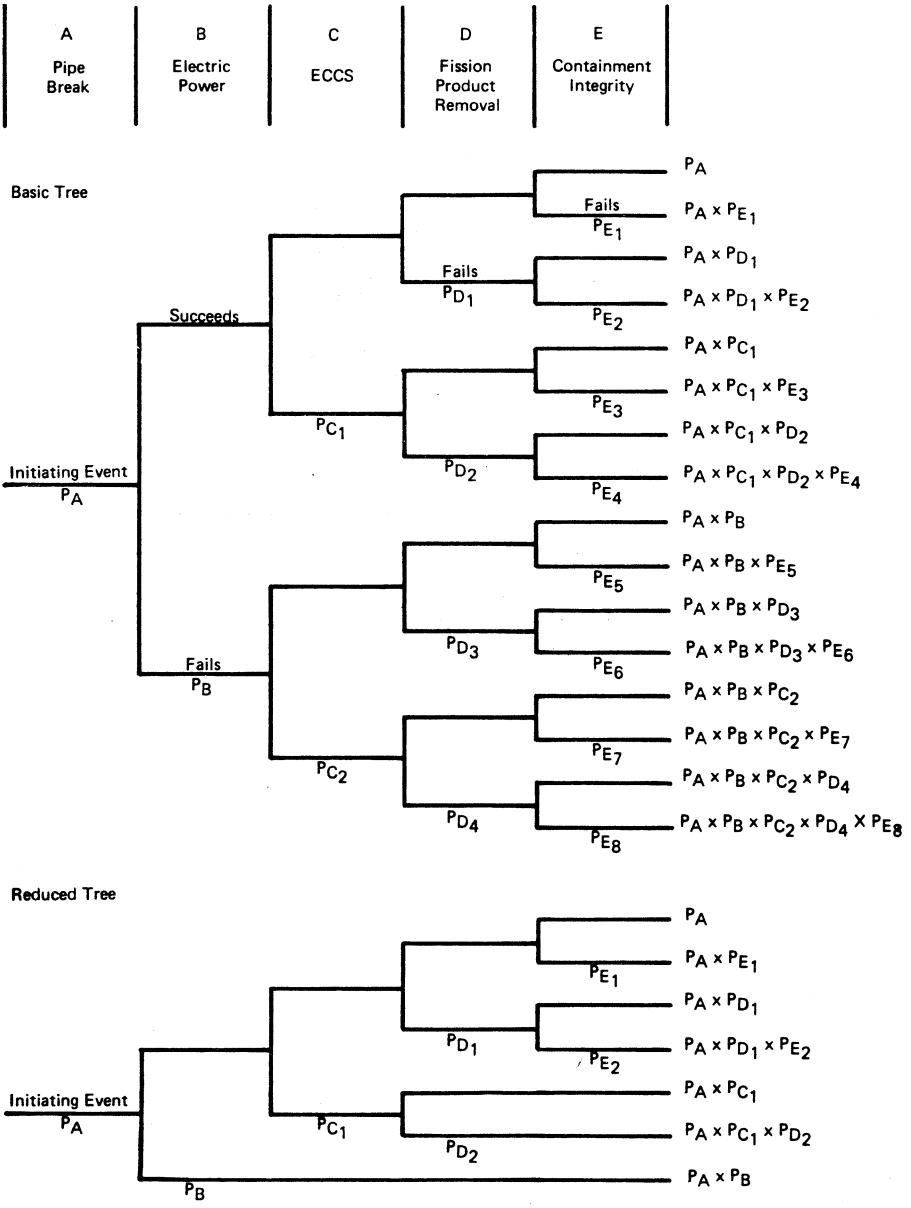

Fig. 14.1. Simplified event trees for a large loss-of-coolant accident. [For this diagram, it is assumed that the failure probabilities (P_f) are small and therefore factors of the form $1 - P_f$ are not explicitly indicated.] (From Ref. [12, p. 55].)

This is an example of a *common-mode* failure (i.e., a case in which individual failures are causally related). Such scenarios could, at least in principle, greatly increase the chance of a serious accident. It is therefore necessary, but not necessarily easy, to identify sequences in which the failure of one system enhances the likelihood of the failure of others.

What is the probability that the electric power will fail, as assumed for the event tree of Figure 14.1? That question is answered in principle by a fault-tree analysis, diagrammed in Figure 14.2. For the electric power to fail, there must be a loss of *both* the off-site AC power (the standard source) and the on-site AC power (one or more emergency generators). The loss of AC power *or* the loss of DC power (required in this case to control the AC system) would mean that the safety systems would not operate.

In many cases, the individual ingredients for the event-tree and fault-tree analyses come from an extensive database (e.g., the rate of failure of a given type of valve or motor that may be widely used outside of the nuclear power

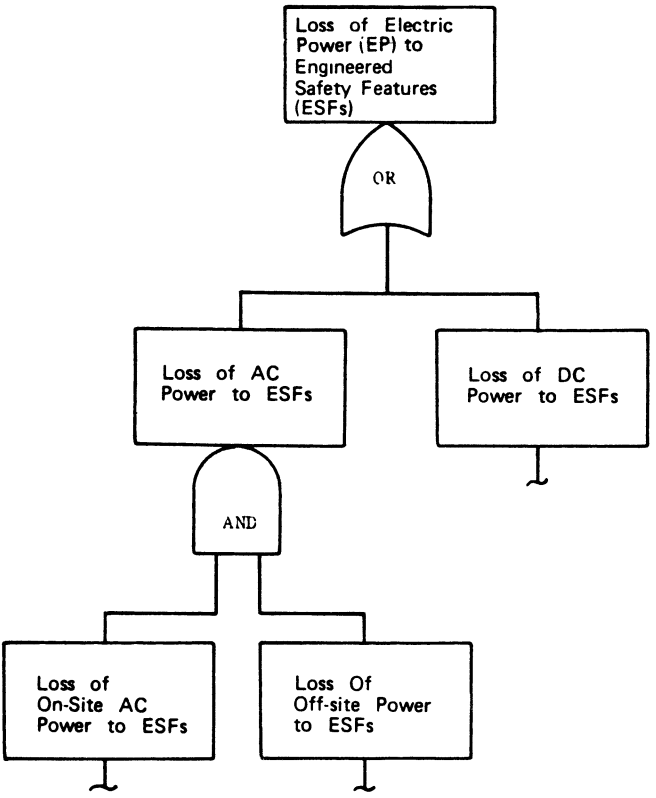

Fig. 14.2. Fault tree for loss of electric power. (From Ref. [12, p. 56].)

industry). In other cases (e.g., the probability of human error), the input numbers are likely to be only rough surmises.

Combining the outcomes of the event-tree analyses and the fault-tree analyses gives the probability for an accident scenario. Some scenarios will represent accidents with large releases of radioactivity to the environment; others will represent small releases. The overall results of the study can be embodied in graphs or tables in which the probability of an accident of a given or greater severity is displayed as a function of the severity of the accident.

The Role of Probabilistic Risk Assessment

In the original Reactor Safety Study, considerable emphasis was put on the absolute magnitude of the reactor accident risks. Uncertainties in the analysis were explicitly indicated, but there were criticisms that these had been underestimated. In NUREG-1150, a later PRA study, the issue of uncertainties was featured more prominently (see Section 14.4.2).

Despite the difficulty of making precise estimates of reactor risk with PRA techniques and the uncertainties that surround their results, they appear to offer the best available approach to risk estimation. As analysis methods are improved and input data on failure rates becomes more extensive, there can be increasing confidence in the applicability of the results. However, ambivalence remains, as reflected in comments made in a 1993 report prepared by the Nuclear Energy Agency of the OECD:

> *Probabilistic safety assessment* (PSA) is a powerful technique for providing a numerical assessment of safety. It is being increasingly used as a guide for comparing levels of safety. As such it complements the deterministic approach to safety assessment, but it is not considered as an absolute measure of safety for regulatory purposes. . . .
>
> But the importance of PSA is not so much in the final answer that it gives for the chance of accidents. Its main value lies in the insights that are obtained in the process of the analysis. It will highlight those elements in a chain of events which contribute significantly to the probability of serious accidents—the weak links—and which if strengthened will therefore give a significant improvement in overall safety. [3, p. 63]

There appears to be little dissent from the view that PRA (or PSA) studies give useful information on relative risks and on the identification of "weak links." However, the uncertainties in the PRA estimates of absolute risk magnitude of the risks may be large, and the policies on the use of PRAs for regulatory purposes by agencies such as the NRC appears to be still evolving (see Section 14.5.1).

14.3.3 Results of the Reactor Safety Study

Summary of Results

The results of the RSS included estimates of the probability distributions for a variety of forms of harm: early fatalities, early illness, latent cancer fatalities, thyroid nodules, genetic effects, property damage, and magnitude of the area in which relocation and decontamination would be required. These results were presented in the form of graphs of the probability of occurrence as a function of the magnitude of the harm. Thus, for instance, the calculated probability of an event that would cause more than 1 latent cancer death per year was about 3×10^{-5}/reactor-year (RY); the probability dropped to 2×10^{-6}/RY for more than 100 latent cancer deaths per year [12, p. 97]. Large uncertainties were indicated for both the probabilities of the events and the resulting number of cancer fatalities.

The probability of a core melt was estimated to be 5×10^{-5}/RY, with an upper bound of 3×10^{-4}/RY, or about 1 per 3000 reactor-years of operation [12, p. 135]. The most probable cause of a core melt was found to be not a break in the large pipes providing the main cooling water but rather an accumulation of smaller failures. This was surprising in view of prior prevailing beliefs.

To provide perspective, the RSS also compared the risks from reactor accidents to those from other sorts of accidents or natural mishaps. For these other accidents, there are few data on latent effects. Perhaps the trauma of a nonfatal airplane accident increases one's chance of dying 30 years later, but this is not customarily included as a fatal consequence of airplane accidents. Thus, a direct comparison between nuclear power and other hazards is made simpler, although incomplete, if consideration is restricted to early fatalities. For a nuclear reactor accident, these would be primarily caused by very high early radiation exposures. In Figure 14.3, the annual risks from 100 reactors, as estimated in the RSS, are compared with the annual risks from other causes, such as airplane accidents and dam failures. For example, Figure 14.3 indicates an average of 1 airplane accident causing 100 or more fatalities every 3 years, whereas a nuclear reactor accident with this early toll was predicted to occur only once every 80,000 years.[15]

Responses to the Reactor Safety Study

The RSS was received very differently by different groups. Nuclear power advocates greeted it enthusiastically as a vindication of their belief in nuclear safety. It was possible to draw all sorts of dramatic comparisons from it, and these were gleefully put forth; for example, that there was less chance of

[15] It should be noted that with accidents of this magnitude, the consequences other than early fatalities are likely to be much more severe for the reactor accident than for the airplane accident.

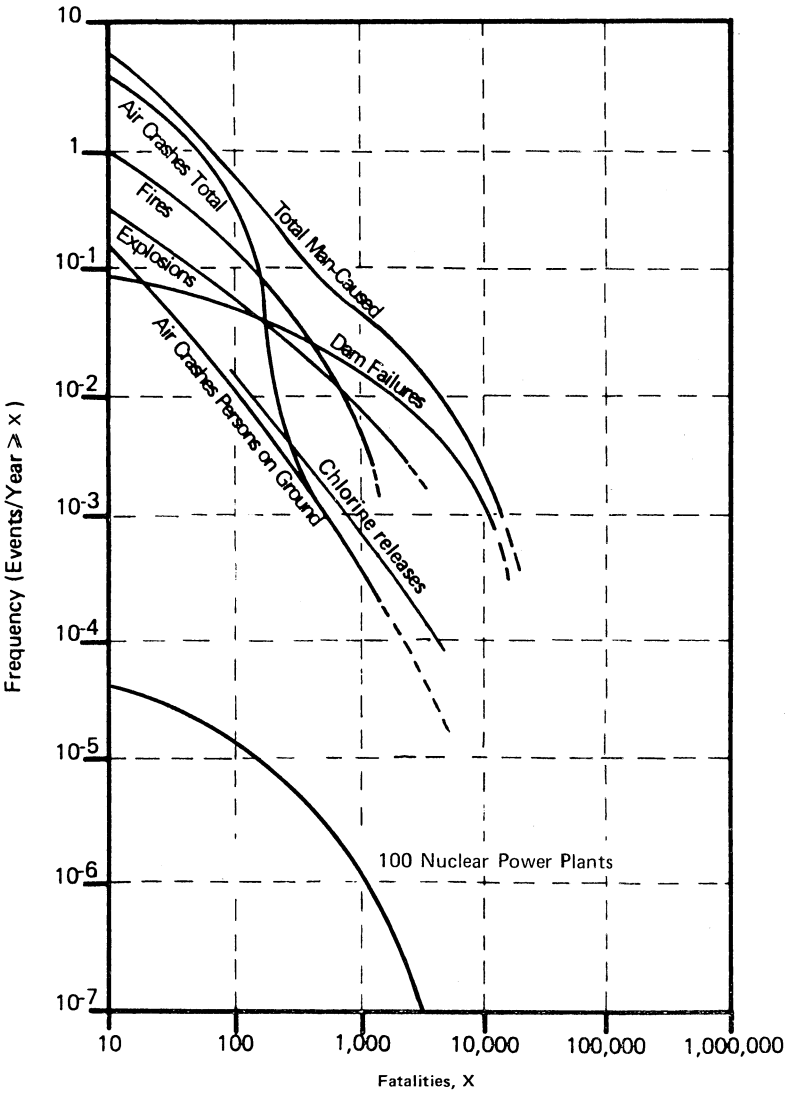

Fig. 14.3. RSS comparison of annual probabilities of accidents causing x or more (early) fatalities: 100 nuclear reactors compared to other "man-caused" events. (From Ref. [12, p. 119].)

being killed by an accident in a nearby nuclear power plant than by an errant automobile, even if you were neither in a car nor crossing a street yourself. Nuclear opponents greeted the RSS with strong criticism and even scorn. It was not surprising, in their view, that a study sponsored and carried out by the "nuclear establishment" would conclude that nuclear power was safe.

An influential critique of the RSS was done by a special review group, commissioned by the Nuclear Regulatory Commission and chaired by Harold Lewis of the University of California at Santa Barbara [14]. The main conclusions of the Lewis report were (1) the methodology used in the RSS was basically sound, (2) significant mistakes had been made, for example, in some of the statistical methods, (3) it was difficult to balance the instances of conservatism and nonconservatism, (4) the uncertainties were much greater than those quoted in the RSS, (5) the executive summary was misleading, and (6) the review group could not conclude whether the probabilities of a reactor core melt were higher or lower than those quoted in the RSS.

The Lewis report was regarded by some as a "repudiation" of the RSS, and the NRC backed away from using it as a guide for regulatory decisions. However, Lewis himself took a consistently "pro-nuclear" position in congressional testimony, stating that he felt "the plants are actually safer than stated in the Rasmussen report."[16] Lewis made this last statement, which reiterated earlier statements by him in the same vein, in May 1979, shortly *after* the TMI accident. However, TMI made such studies at least temporarily irrelevant. Quite apart from the merits and demerits of studies by academic scientists and engineers such as Rasmussen and Lewis, a significant fraction of the public concluded after TMI, and all the more after Chernobyl, that nuclear reactors were not safe enough for nearby siting. That conclusion has had a profound influence on the subsequent pace of nuclear power development.

Hindsight on RSS Predictions

It is tempting to look back with the benefit of hindsight on the RSS estimate of a core damage probability of 5×10^{-5}/RY with an upper bound of 3×10^{-4}/RY. As of the end of 2003, there had been about 2870 reactor-years of LWR operation in the United States.[17] If one uses the RSS, the predicted number of core melts through 2003 would be 0.14 with an upper bound of 0.9. The actual number of core melts was one (TMI), so the actual experience does not differ greatly from the predicted upper bound.

This comparison is not quite appropriate, however. The RSS was specific to reactors as they existed in the 1970s. Therefore, its results should be compared to reactor performance in that period before the post-TMI improvements were made. As of the end of 1979, there had been under 600 reactor-years of LWR operation in the United States, and the above "anticipated" accident rate should be reduced by about a factor of 5. Thus, one could infer that the average estimates of core melt probabilities given in the RSS were underestimates.

If one goes beyond average estimates, another interesting viewpoint emerges. As indicated earlier, the RSS studied certain reactors in detail and

[16] References and further quotations are given in Ref. [15].

[17] The RSS was for United States LWRs, and, therefore, it is appropriate to restrict comparison to their record.

these were taken as representative of all LWRs. In particular, the PWR analysis was based on the Surry 1 power plant, manufactured by Westinghouse. According to the RSS, using this as prototypical for all PWRs would "tend to overestimate, rather than underestimate the risk," because this was a relatively old plant and newer ones would, on average, be safer. However, as discussed in a subsequent study of the implications of the TMI accident, conducted under the auspices of the American Physical Society:

> The first reaction of many observers to the accident was that the Reactor Safety Study methodology was completely wrong because it had not predicted that type of accident would soon occur. The particular sequence...was calculated for the Surry facility...to have a frequency of once in 10^5 years. Yet...if the RSS procedures had been applied to a Babcock and Wilcox reactor like TMI-2, the methodology would have predicted a frequency of occurrence of one in 300 years. Babcock and Wilcox reactors had an operating history of about 30 reactor years. The differences stemmed from: (a) the pressure relief valve settings that caused the valve to be released before reactor scram and (b) the fact that the steam generators had a small heat capacity and dried out in ten minutes, compared with a time of about an hour calculated for the Westinghouse reactors such as Surry...if the methodology had been applied to the reactor at Three Mile Island, the plant-specific scenario differences might have been noted, modifications might have been made, and the accident perhaps avoided. [9, p. S11]

In short, the greatest mistake with the RSS analyses was the failure to apply the analyses to all reactors, individually.

14.4 Post-TMI Safety Developments

14.4.1 Institutional Responses

The TMI accident showed that the prior efforts of government agencies and reactor manufacturers to achieve safety had been insufficient. Although the consequences were limited, in that there was no large release of radioactive material to the outside environment, it was clear that further measures were needed to improve reactor safety. This was obviously felt by the public and was recognized by the nuclear industry and the U.S. Nuclear Regulatory Commission (NRC). The intensity of their subsequent efforts was increased by the fear that any accident would bring discredit to all of nuclear power. This concern is encapsulated in the phrase: "An accident anywhere is an accident everywhere."

More specifically, the accident revealed defects in physical components of the cooling system and in the systems that provided the operators with information about the status of the reactor. It also showed the need for improve-

ments in operator training and in communication among utilities. Remedial steps were taken in all of these areas, including, in some cases, expensive and time-consuming retrofitting of existing reactors and modifications of reactors under construction. The first impact of these measures in the United States was a pause in the licensing of new reactors and even in the operation of some reactors. In consequence, the output of nuclear electricity dropped in the 3 years following TMI, before beginning a substantial rise in the mid-1980s. The last new nuclear reactor in the United States went into operation in 1996, but despite some subsequent shutdowns of older reactors, total nuclear output continued to rise during the 1990s and reached a new high in 2002.

At first, the rise was attributable mainly to additional reactors coming on line. However, since about 1990, it has been due almost entirely to an increase in the capacity factors of the reactors (see Section 2.4.2). The capacity factor provides a good overall measure of reactor performance and its rise reflects the positive impact of post-TMI improvements in equipment and maintenance procedures. As discussed in Section 14.5.3, the rise in capacity factors was accompanied by other, more direct, indicators of increased safety.

On an institutional basis, in the United States the nuclear industry established the Institute of Nuclear Power Operations (INPO) to exchange information and coordinate and monitor efforts to make reactor operation more reliable and safer. At the same time, the NRC intensified its watchdog role in setting standards and monitoring performance.

Internationally, there are long-standing reactor safety programs operated by the International Atomic Energy Agency [e.g., the IAEA's International Nuclear Safety Advisory Group (INSAG)] and by the Nuclear Energy Agency of the OECD. In addition, the World Association of Nuclear Operators (WANO) was established after the Chernobyl accident (1986) as an international counterpart of INPO [3, p. 27]. It now plays an active role in reviewing the performance of individual plants throughout the world and in facilitating the exchange of information on reactor safety.

14.4.2 1990 NRC Analysis: NUREG-1150

Analysis Procedure in NUREG-1150

A further step in the development of reactor safety analysis methods in the United States was marked by the publication in 1990 of the NRC report *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants*, also known as NUREG-1150 [11]. Five LWRs were analyzed in detail for this study. These are in some sense typical of LWRs in the United States, but the reported results are specific to the individual reactors.

In NUREG-1150, an explicit distinction was made between *internal* and *external* events. Internal events are those due to the malfunctioning of components of the reactor, including its control systems. External events are those

initiated by things that happen outside the reactor (e.g., earthquakes).[18] External events were analyzed for only two of the five reactors.

The NUREG-1150 analysis was divided into several separate stages:

- *Accident frequency.* The goal here is to estimate the probability that the reactor core is damaged. The starting point is to identify possible initiating events and assess their frequency. For events due to internal system failures, accident probabilities were determined through a combination of event-tree and fault-tree analyses. Human error and "dependent failures" (also known as common-mode failures) were included. For events due to external hazards—including earthquakes, fire, and aircraft impacts—somewhat analogous procedures were used, but the database for the initiating events is not as good and there is a greater chance of a simultaneous failure of several components.
- *Accident progression.* Given damage to the reactor core, it is important to know what further damage occurs. Thus, probabilities were estimated for the breaching of the reactor vessel and for either breaching of the concrete containment or leaks through it.
- *Transport of radioactive material.* Given damage to the fuel and to the reactor vessel or cooling system, there will be a transfer of radionuclides to the reactor building. Release to the environment then depends on whether or not the containment fails. The noble gases are the most likely to be released, with virtually all escaping given sufficient fuel damage. For other radionuclides, the release rates depend on the volatility of the element, ranging from high for iodine and cesium and low for ruthenium and strontium. The aggregate of total releases to the environment is the source term.
- *Off-site consequences.* The radiation doses received by people outside the reactor depend on the magnitude of the source term, the movement of the plume of radioactive material in the air, and the details of the pathways by which radiation exposures occur. Doses and health consequences were calculated for a variety of assumptions as to the evacuation of the surrounding population.
- *Integrated risk analysis.* An overall integrated risk is found from the array of probabilities for each of the various stages.

As discussed in connection with the RSS, the ultimate result of this analysis is a probability distribution for the risk of occurrence versus the magnitude of the consequence, for each adverse consequence of interest. Thus, the result might be the probability distribution for exceeding various levels of population dose or of latent cancer fatalities. Although such a probability distribution cannot be fully represented by a single number, both medians and mean values are given in NUREG-1150 to provide an easily encapsulated overall

[18] The distinction is not clean and, customarily, loss of power from off-site sources is included as an internal event, whereas floods and fires within the plant are termed external events [11, p. 2-4].

Table 14.1. Estimated mean probabilities per reactor-year of core damage and other effects of reactor accidents, for reactors studied in NUREG-1150.

	Reactor Studied				
	Surry 1	Zion 1	Sequoyah 1	Peach Bottom 2	Grand Gulf 1
Reactor type	PWR	PWR	PWR	BWR	BWR
State located	VA	IL	TN	PA	MS
Capacity (MWe)	781	1040	1148	1100	1142
Commercial operation (year started)	1972	1973	1981	1974	1985
Internal events					
Core damage[a]	4×10^{-5}	6×10^{-5}	6×10^{-5}	4.5×10^{-6}	4×10^{-6}
Early containment failure[b]	4×10^{-6}	6×10^{-6}	7×10^{-6}	2×10^{-6}	1×10^{-6}
Individual early fatality[c]	2×10^{-8}	3×10^{-9}	1×10^{-8}	5×10^{-11}	3×10^{-11}
Individual latent cancer[c]	2×10^{-9}	3×10^{-9}	1×10^{-8}	4×10^{-10}	3×10^{-10}
External events, core damage[d]					
Seismic events, LLNL	1.2×10^{-4}			8×10^{-5}	
Seismic events, EPRI	2.5×10^{-5}			3×10^{-6}	
Fires	1.1×10^{-5}			2×10^{-5}	

[a] Number for Zion reactor reflects plant modifications after study was initiated [11, p. 7-4].
[b] A containment failure here includes both breaks in the containment structure and bypass of it. It may lead to large early release of radionuclides.
[c] The individuals considered are those within 1 mile and 10 miles of the reactor boundary for early and latent fatalities, respectively.
[d] Separate results are given for studies from the Lawrence Livermore National Laboratory (LLNL) and the Electric Power Research Institute (EPRI).
Sources: Capacity data and commercial operation dates are from Ref. [16]. Core damage data are from Ref. [11, pp. 3-4, 4-4, 5-4, 6-5, and 7-4]. Large early release data are from Ref. [11, p. 9-6]. Individual risk data are from Ref. [11, p. 12-3].

perspective.[19] Some results of the NUREG-1150 analysis are summarized in Table 14.1, for the five reactors studied.

The NUREG-1150 study is more pertinent to the present situation than the original Reactor Safety Study, because it used more advanced analysis techniques and considered reactors as they were after a period of considerable upgrading. This is a continuing process, however, and the analyses were spe-

[19] The mean is, in general, higher than the median, because the probability distribution for a given consequence generally has a tail extending to high magnitudes.

cific to the situation at the time they were made (in the later 1980s). Since then, conditions may be better due to further modifications in equipment and operating procedures or worse due to aging. Continual reactor-by-reactor monitoring is necessary. As discussed in Sections 14.4.3 and 14.4.4, other indicators suggest continuing gains.

Core Damage Probabilities

The mean calculated probability of core damage from internal events varied from 4×10^{-6}/RY in the best case to 6×10^{-5}/RY in the worst case, with a rough (arithmetic) average of about 3×10^{-5}/RY (see Table 14.1). The probability of core damage was considerably lower for the two BWRs than for the PWRs, although the study cautioned that it would be "inappropriate" to conclude that this was true in all cases [11, p. 8-11]. Nonetheless, some advantages of BWRs were pointed out, particularly more redundancy in the emergency core-cooling systems.

The main causes of core damage differed among the reactors [11, p. 8-3].[20] For two (Zion and Sequoyah), events with loss-of-coolant were the most important factor. For two (Surry and Grand Gulf), loss of power (station blackout) was the main factor. For one (Peach Bottom), roughly equal responsibility was placed on loss of power and failure of control rod insertion during transient disturbances.

Core damage due to external causes was considered for only two of the reactors, Surry 1 and Peach Bottom. In both cases, seismic events and fires were the only significant external sources of risk. The data of Table 14.1 might suggest that the external risks are greater than the risks due to internal failure. However, such a conclusion may be premature. For one, the risk for seismic events is highly uncertain, with two analyses considered in NUREG-1150 differing substantially. Further, the seismic risk distributions are very broad and are skewed so that the median risks are considerably lower than the mean risks [11, p. 8-6].[21]

Early Containment Failure

Radionuclides can escape into the environment due either to a breaching of the containment structure or due to valve failures, through cooling system pipes that bypass the containment and vent outside it. An emphasis is put on early failures because if the release of radionuclides is delayed, "mitigative features within the plant can substantially limit the release that occurs" (i.e.,

[20] A compact summary is given in Ref. [17, p. 3–7].

[21] For the LLNL analyses, which give the higher core damage probabilities, the median risk is of the order of one-tenth the mean risk: 1.5×10^{-5}/RY for Surry 1 and 4×10^{-6}/RY for Peach Bottom 2.

radionuclides may be retained inside the containment despite the containment failure [11, p. 9-5]). In each case, the early containment failure rate is estimated at under 10^{-5}/RY.

Consequences of Accidents for Human Health

The key potential health effects of a reactor accident are early fatalities, due to very high radiation doses, and latent cancers, due to the long-term effects of smaller doses. The NRC has put forth "quantitative risk objectives" for these accident consequences (see Section 14.5.1), and the NUREG-1150 results in Table 14.1 are couched in terms that permit a direct comparison with the NRC objectives. Average early fatality risks are calculated for individuals within 1 mile of the reactor, and latent fatality risks for those within 10 miles of the reactor. Of course, the doses decrease substantially with distance, and the risks are higher in these regions than in broader surrounding areas.

The mean early fatality risk for individuals is calculated to range from under 10^{-10}/RY to 2×10^{-8}/RY. Even for the greater of these numbers, the risk is small—1 chance in 50 million per year. This is only 4% of the NRC objective of 5×10^{-7}/RY. For latent cancers, the highest of the calculated risks (1×10^{-8}/RY) is only 0.5% of the NRC objective (2×10^{-6}/RY). Therefore, unless the calculated results are greatly in error, the NRC's safety goals for individuals are satisfied by a large margin.

Seismic Risk

A striking aspect of Table 14.1 is the relative importance of seismic risks in the tabulated core damage frequencies.[22] Determination of the seismic core damage probabilities involves estimating both the probability of earthquakes of various magnitudes at the reactor site, the so-called seismic hazard, and the ability of the reactor to withstand the resulting ground accelerations.[23] The NUREG-1150 calculation used seismic hazard assessments from both Livermore and EPRI. There is no conclusive method for predicting earthquake probabilities, and both the Livermore and EPRI studies relied upon an array of expert evaluations.[24]

[22] The large differences between the EPRI and LLNL results prompted the NRC to commission a study on methods to be used in carrying out a probabilistic study of seismic hazards [18]. No new comprehensive analysis of seismic hazards incorporating these recommendations has been published as yet.

[23] The methods used in these analyses are described in detail in Refs. [19] and [11, Section C11].

[24] The EPRI and Livermore studies were both part of a major program undertaken in the late 1980s to assess seismic hazards, in the region of the United States to the east of the Rocky Mountains, where the large majority of the reactors are situated. These hazards are ultimately couched in terms of site-specific probability distributions for ground acceleration.

Differences in these evaluations in the two studies led to substantially different mean results for the core damage probability. The associated probability distributions are very broad. For example, for the Livermore seismic hazards, the 5th and 95th percentiles in the core damage frequency distribution differ by more than a factor of 1000 for the Surry plant, and the median probability is only about one-eighth of the mean probability [11, p. 8-6]. The large width of the distribution makes any "average" result a poorly established quantity. Subsequent to the publication of NUREG-1150, a new Livermore study of seismic hazards was carried out, which, in an overall sense, tended to move the Livermore results in the direction of the EPRI results (i.e., to lower the estimated risk [20]).

Despite the emphasis here on probabilistic risk assessments for evaluating reactor earthquake risks, this is not the chief approach adopted by the NRC. Instead, *seismic margin* methodology is being used for some present and all future reactors. This method is somewhat less demanding in terms of the analysis required and may be the most reasonable approach given the large uncertainties in estimating earthquake probabilities. The starting point is the specification of the so-called *safe shutdown earthquake* (SSE).[25] This is an earthquake whose magnitude is based on the "maximum earthquake potential" in the vicinity of the site. The reactor must be designed to shut down safely should this largest expected earthquake occur.

If a reactor can withstand an earthquake more severe than the SSE, then the reactor has a "seismic margin." The extent of the seismic margin is based on a reference earthquake, more severe than the SSE, for which there is a "high confidence of a low probability of failure" (HCLPF).[26] For example, if the SSE corresponds to a peak ground acceleration of 0.3 g at the reactor site, the seismic margin condition might be established by demonstrating fulfillment of the HCLPF condition for an acceleration of 0.45 g. The NRC requirement for new reactors will be that they demonstrate an adequate seismic margin.

Although the safety study NUREG-1150 included estimates of seismic core damage probabilities for two of the five reactors, no estimates were given of large release probabilities for seismic effects. The rationale was that if an earthquake is severe enough to damage a reactor, there will be damage to other structures such as buildings and dams, with consequences that are expected to be more severe than those from possible reactor releases [11, p. 1-4]. In the absence of meaningful estimates of the effects of this other damage on the surrounding population, the NRC lacked a reference point for considering the

[25] The definition and use of the safe shutdown earthquake is discussed in Appendix A to Part 100 of Ref [21]. The SSE had previously called this the *design basis earthquake*.

[26] The HCLPF criterion can alternatively be established by deterministic or probabilistic determination of failure modes. In the latter case, it is assumed to correspond to a greater than 95% confidence that the failure probability is less than 5% [19, p. 5-4].

significance of the off-site effects of nuclear reactor containment failures and did not calculate their probability for NUREG-1150.

Overall, although considerable effort has gone into making nuclear reactors "safe" against earthquakes, there has been difficulty in quantifying the level of safety. In a quite different approach, some designs for new reactors incorporate seismic isolation between the reactor and the surrounding ground, with the goal of decoupling the reactor from possible ground motion.

14.4.3 Predictions of Core Damage and Precursor Analyses

The Record Since TMI

Estimates of core damage frequency based on PRAs cannot be checked against actual experience because there has not been an LWR accident that has caused core damage since the TMI accident in 1979. In this period, there have been roughly 2300 reactor-years of operation in the United States, a record that can be taken to suggest that the core damage frequency during this period was probably less than 5×10^{-4}/RY. However, such an upper limit considerably exceeds both what is acceptable and what is predicted and, therefore, is not very useful.

The database, of course, substantially increases if one looks at all LWRs throughout the world. Again, there has been no core damage. However, it is not fully appropriate to compare worldwide experience to estimates made in NRC studies of individual U.S. reactors because, although the guiding principles are the same, differences in regulatory, construction, and operating practices may make reactors in other countries more safe or less safe than U.S. reactors.

Analysis of Precursor Events

There are continuing malfunctions of reactor equipment short of core damage, spanning a wide range of severity. These malfunctions can be viewed as *precursor* events (i.e., potentially the initial first stage in a chain of failures that, if they all occurred, could lead to core damage). Analyses of precursors provides a powerful tool for inferring the expected core damage frequency and—perhaps even more valuable—gauging progress in reactor safety.

The precursor events are identified from "licensee event reports" that reactor operators are required to submit to the NRC for each significant malfunction in reactor performance. It is possible through the PRA methodology to then calculate a "conditional core damage probability" (CCDP) (i.e., the probability that core damage would result given this first mishap). These calculations have been made in the NRC's Accident Sequence Precursor Program. An index of reactor performance is given by the sum of the CCDPs for all failures in a year divided by the number of operating reactors. This index has been variously called the "inferred mean core damage probabil-

ity" [22], the "annual core damage index" [23], or—in the terminology now used by the NRC—the "accident sequence precursor (ASP) index" [24]. It provides an overall indication of progress in reducing the likelihood of core damage, although the ASP index is not a complete predictor of core damage frequency [24].[27] For example, it does not include external events. Further, year-to-year results for the index are not strictly comparable, because the detailed analysis methods and range of events included have changed with time [25].

Despite these caveats, this index provides a valuable indication of progress in reactor safety. The reported magnitudes of the ASP index for the period from 1969 through 2000 are shown in Figure 14.4 [23, 24]. The ASP index is dominated by a few events of relatively great severity (high conditional core damage probability) rather than by a large number of relatively minor events. For example, the high value of the index in 1979 was due to the TMI

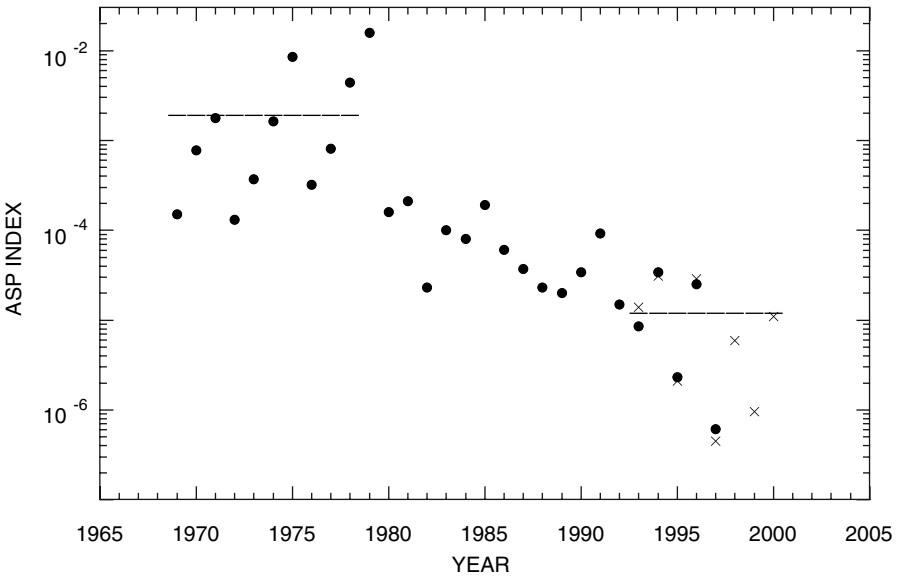

Fig. 14.4. Summary of precursor analyses: the ASP index as a function of time. The horizontal lines are averages over extended periods. [Solid circles: 1969 to 1997 (calendar years), data from Ref. [23]; crosses: 1993 to 2000 (fiscal years), data from Ref. [24].]

[27] The numerical values of the "mean core damage frequency" reported by Thomas Murley in 1990 [22], the "core damage index" reported by Murley in 1999 [23], and the ASP index of the recent NRC documents [24] are in good agreement, confirming that the differences are more in terminology than in basic meaning (see Figure 14.4).

accident[28] and the high excursion in 2000 was due to a single event at a two-reactor station.[29] Such events lead to large year-to-year fluctuations, although not enough to mask the overall downward trend, exhibited in averages of the ASP index taken over a number of years. This average dropped from about 2×10^{-3}/RY for the 1969–1978 period (the 10 years preceding the TMI accident) to about 1.2×10^{-5}/RY for the 1993–2000 period—an improvement of better than a factor of 100.

The retrospectively calculated pre-TMI ASP index is considerably higher than even the upper bound on core damage probability of 3×10^{-4}/RY estimated in the 1975 Reactor Safety Study. Therefore, taking the ASP index to represent an approximate core damage probability, it would appear that the reactors were less safe than then thought. Subsequent changes in equipment and operating procedures have greatly improved matters (e.g., see Ref. [26]). However, the improvement in the ASP index seems to have ended in the early 1990s, although this is difficult to interpret due to the large fluctuations during the 1990s (see Figure 14.4), in part caused by the dominance of only a few events in each year. Preliminary NRC study of the experience for 2001 and 2002, prior to the completion of ASP analyses for those years, does not suggest any marked changes, although the analysis of the Davis–Besse corrosion problem was continuing as of early 2003 (see the next section) [27].

14.4.4 Other Indications of Performance

Following the Three Mile Island nuclear accident, the U.S. nuclear industry established the Institute for Nuclear Power Operations (INPO) designed to coordinate the industry's efforts to remedy existing problems and achieve safer and more economical reactor operation. As part of this activity, INPO monitors and reports upon various performance indicators. Results comparing 2001 to early years include the following [28]:

- *Unplanned scrams.* Automatic shutdowns of a reactor, initiated by a failure in one of the reactor components, are called scrams. The scrams are a safety measure to avoid the development of serious accidents and the occurrence of a scram is part of proper operation of the safety systems. However, just as it is desirable to avoid the need for household circuit breakers to trip, it is desirable to minimize the need for scrams, and the number of scrams per year suggests how well a plant is operating. Measured in terms of the

[28] For 1979, when there were 69 reactors, the contribution to the ASP index from TMI alone (CCDP =1) was 1/69 = 0.0145/RY; the total calculated ASP index was 0.0157/RY.

[29] The CCDP for this event was 4.5×10^{-4}. Summing for the two reactors gives, for 103 reactors, a contribution of 0.9×10^{-5} to the average for all reactors of 1.1×10^{-5}/RY for 2000 [24].

number of scrams per 7000 h of reactor operation,[30] the rates were 7.3 in 1980, 1.2 in 1990, and under 0.1 in 2001. This more reliable operation was one of the reasons for the improved capacity factors.

- *Radiation exposure of workers.* In the absence of a serious reactor accident, the reactor workers are the only individuals at potential risk from exposure to nuclear radiation. Some radiation exposure is inevitable among workers in a nuclear plant, but it is desirable to keep this exposure "as low as reasonably achievable." The collective exposure per reactor (i.e., the sum of individual exposures of all the workers at that unit) dropped by more a factor of 6 from 1980 to 2001.[31]
- *Industrial accident rate.* The rate of industrial accidents provides an indication of the overall safety of working conditions, although it has little to do with nuclear safety per se. The rate of accidents at nuclear reactors that led to lost or restricted work time or to fatalities dropped by a factor of 9 from 1980 to 2001. The rate per 200,000 worker-hours in 2001 was 0.24 for nuclear reactors compared to an average of 4.0 for U.S. manufacturing industries as a whole.

These measures of reactor performance indicate that there have been substantial safety improvements over the past 20 or 25 years. However, this does not mean that the nuclear industry or the NRC can afford to be complacent. This was forcefully brought home by the discovery in March 2002 at the Davis–Besse plant of deep corrosion in the reactor vessel head (i.e., the top portion of the reactor pressure vessel).[32] The corrosion was discovered during an inspection conducted by plant personnel while the reactor was shut down for refueling. It was caused by boric acid that leaked through small cracks in the nozzles that allow control rods to move up and down in the reactor.[33]

Both boric acid leaks and the corrosion caused by boric acid were familiar matters in the nuclear industry, but the NRC thought that it was on a scale that did not constitute a safety problem, as long as it was monitored. In fact, the inspections at Davis–Besse and other reactors were specifically required by the NRC in recognition of past boric acid leaks. However, the magnitude of the corrosion was unexpected. The cavity was about 4 in. by 5 in. and extended to a depth of approximately 6 in. This means it penetrated almost all the way through the top cover of the pressure vessel. The extent of the corrosion suggests that it might have been discovered during an earlier shutdown, given sufficient vigilance. Following this discovery the NRC required opera-

[30] This is equivalent to one reactor-year for a 1000-MWe reactor operating at an 80% capacity factor.

[31] Collective exposures per reactor in 2001 averaged 0.68 person-Sv for PWRs and 1.49 person-Sv for BWRs.

[32] This account of the corrosion problem is based primarily on Ref. [29].

[33] Boric acid is introduced into the reactor cooling water to adjust the reactor's reactivity through neutron absorption in boron-10 (^{10}B).

tors of other PWRs to report anew on boric acid leaks and corrosion at their plants. No comparable problems were found at other reactors. To decrease the possibility of similar future occurrences, the NRC is imposing strengthened requirements for the inspection of reactor vessel heads and is also making a broader examination of ways to assure that other reactor corrosion problems will be avoided [30].

It is of interest to determine how "near" a miss this was. There was no release of radionuclides and no damage of any sort except at the location of the leak. The NRC had not yet completed its analysis of the accident when its March 2003 report on the Accident Precursor Program was released [27]. This analysis will eventually provide an estimate of the likelihood that the accident might have led to core damage. There might have been no core damage even if the corrosion had penetrated all the way through the wall. Some reactor cooling water would have been ejected into the containment building as the pressure inside the reactor vessel was suddenly relieved, but it may have been possible to maintain cooling and avoid core damage. Pending the NRC report, it seems reasonable to conclude that in the case of the Davis–Besse incident, the overall safety system, including the inspection regime, in the end worked, and even had the inspection not come in time to avoid a reactor vessel breach, there probably would have been no major release of radionuclides to the outside environment. Nonetheless, the failure to detect and correct the corrosion promptly showed serious weaknesses in the monitoring procedures of the reactor operator and the NRC. This single event does not negate the very good and improving record of nuclear reactor performance, but should serve as a reminder against complacency.

14.5 Reactor Safety Standards

14.5.1 U.S. Nuclear Regulatory Commission Position

General Approach

Although the above discussion has emphasized probabilistic measures for specifying reactor safety, the NRC—the agency responsible for U.S. reactor safety—has been reluctant to adopt probabilistic criteria in setting reactor licensing requirements. Nuclear reactor safety is regulated using a "deterministic" approach, which, in recent years, has been modified to "risk inform" the application of the deterministic criteria. As described by the NRC in 1995:

> The NRC established its regulatory requirements to ensure that a licensed facility is designed, constructed, and operated without undue risk to the health and safety of the public. These requirements are largely based on deterministic engineering criteria. Simply stated this deterministic approach establishes requirements for engineering margin and for quality assurance in design, manufacture, and con-

> struction. In addition, it assumes that adverse conditions can exist (e.g., equipment failures and human errors) and establishes a specific set of design-basis accidents. It then requires that the licensed facility design include safety systems capable of preventing and/or mitigating the consequences of those design-basis events to protect the public health and safety. [31, §IIIA]

Probabilistic risk assessment methods are used to supplement the deterministic approach by identifying both vulnerable and overprotected parts of the system:

> A natural result of the increased use of PRA methods would be the focusing of regulations on those items most important to safety. Where appropriate, PRA can be used to eliminate unnecessary conservatism and to support additional regulatory requirements. Deterministic-based regulations have been successful in protecting the public health and safety and PRA techniques are most valuable when they serve to focus the traditional, deterministic-based, regulations and support the defense-in-depth philosophy. [31, §IIIA]

This is a limited use of PRA methods. The NRC has resisted suggestions to take the further step of setting probabilistic limits as part of the regulatory criteria (e.g., setting limits on the calculated core damage frequency).

Health Effects Criteria

The NRC's primary safety goals were set forth in 1983 in the report *Safety Goals for Nuclear Power Plant Operation* and were published in 1986 in a slightly revised form in the *Federal Register* [32]. This policy statement has remained the primary guide for NRC regulatory activities. The NRC staff in 2001 recommended modifications to the statement, but these were disapproved by the NRC commissioners who indicated that a change was not then timely given broader efforts underway to further "risk inform" NRC's safety regulations and the press of other demands upon the NRC [33, 34]. The 1986 statement put forth two "qualitative safety goals":

> Individual members of the public should be provided a level of protection from the consequences of nuclear power plant operation such that individuals bear no significant additional risk to life and health.
>
> Societal risks to life and health from nuclear power plant operation should be comparable to or less than the risks of generating electricity by viable competing technologies and should not be a significant addition to other societal risks. [32, § II]

The NRC also presented "quantitative risk objectives" to give a specific numerical meaning to the term "significant additional risk," but it prefaced the quantitative criteria with the cautionary comment:

> The Commission wants to make clear. . .that no death attributable to nuclear power will ever be "acceptable" in the sense that the Commission would regard it as a routine or permissible event. We are discussing acceptable risks, not acceptable deaths. [32, § IIIB].

Essentially, the NRC indication of objectives interpreted "significant addition" to mean an increase of 0.1% over the risk from non-nuclear sources [32, § IIIC]:

- *Prompt fatalities.* The risk to an average individual living within one mile of a plant should not exceed 0.1% of the combined average risk from all other accidents. Taking the latter risk to be 5×10^{-4} per year, the individual risk limit for a prompt fatality accident translates to 5×10^{-7} per year, as indicated in NUREG-1150 [11, p. 12-3].
- *Cancer fatalities.* The risk to the population living within 10 miles of the plant should not exceed 0.1% of the risk from all cancers. Taking the average annual cancer rate to be 19 per 10,000, the individual risk limit for cancers attributable to releases from the nuclear plant translates to 2×10^{-6} per year (again as in NUREG-1150 [11, p. 12-3]).

These were specified to be "quantitative objectives." Although they do not carry the legal force of regulations, they were deemed by the NRC to provide a "useful tool by which the adequacy of regulations. . .can be judged" [32, §V].

These objectives have remained unchanged in subsequent years and have been commonly referred to in safety evaluations. They were cited as "safety goal[s]" in the NRC document NUREG-1150 (e.g., Ref. [11, p. 12.3]) and were, for example, reiterated in 1998 in the context of an analysis of a proposed new reactor, the AP600 (see Section 16.3.2), with the statement: "The Commission approves the use of the qualitative safety goals, including use of the quantitative health effects objectives, in the regulatory decision making process" [35, p. 19-4].

Accident Frequency Criteria

As discussed earlier, the NRC has no official standard that specifies reactor safety requirements in terms of the probabilities for core damage or a large early release, as estimated through PRAs. However, many NRC documents have discussed the matter. Thus, a 1983 document on safety criteria stated that "the Commission has selected the following design objective" [36, p. 14]:

> The likelihood of a nuclear reactor accident that results in a large-scale core melt should normally be less than one in 10,000 per year of reactor operation. [36, p. 14]

The document went on to indicate the importance of mitigating the effects of such an accident through "containment, siting in less populated areas, and emergency planning as integral parts of the defense-in-depth concept."

It should be noted that these are "design objectives," but they were not necessarily to be incorporated in "the regulatory framework" [36, p. 15]. The terminology and formal status of this objective has been the subject of continual discussion in subsequent years. In a 2001 recommendation to the commissioners, the NRC staff suggested adopting "useful subsidiary benchmark[s]" of 10^{-4}/RY for the core damage frequency and 10^{-5}/RY for the large early release frequency (LERF) [33, §IIID]. However, as already mentioned, the commissioners declined to accept that recommendation at the time.

However, the NRC is thinking in terms of a more ambitious goal for new reactors. As early as 1986, the Commission had recommended consideration of a performance guideline for a LERF of less than 10^{-6}/RY [32, § V]. More recently, in the design certification documents for the ABWR and AP600 reactors, the NRC compared the expected performance to the "Commission's goal" of a LERF of under 10^{-6}/RY [35, p. 19-180; 37, p. 19-39]. A further indication of NRC thinking on desirable safety levels is provided in Regulatory Guide 1.174, a 1998 document discussing its approach to handling requests by a reactor operator for changes in the reactor [38]. In brief, it calls the calculated increased risk "very small" if it is less than 10^{-6}/RY for the core damage frequency and is less than 10^{-7}/RY for the large release probability.

Of course, limiting an increase to 10^{-6}/RY is less stringent than limiting the magnitude to 10^{-6}/RY. Nonetheless, this stipulation suggests that if the NRC eventually adopts a formal limit for the core damage frequency in new reactors, the number selected is likely to be more demanding than the value of 10^{-4}/RY that has appeared in past discussions.

Cost Considerations

In considering reactor safety, there is a certain temptation to ignore cost issues and say that society should "spare no cost" in its efforts to reduce the probability of a reactor accident. However, costs are not ignored for other activities (e.g., in preventive medicine, highway safety, and airplane construction, to cite a few examples) and they are not ignored for nuclear power plants. In all of these cases, one eventually reaches the point of diminishing returns.

The issue is explicitly addressed by the NRC in the spirit of cost–benefit analysis. For example, in its review of the application for design approval of the Westinghouse AP600 reactor, the NRC considered a number of possible changes beyond those already implemented in the design and compared the cost of each change to an assumed benefit of $5000 per person-rem of averted exposure [35, p. 19-255].[34] The costs for the 14 design alternatives that were considered ranged from $19,762 to $14,679,500 per person-rem [35, p. 19-270] and, therefore, none of these design changes was required.

[34] This figure is based on an an assigned benefit of $2000 per person-rem for health effects and $3000 per person-rem "to account for offsite property damage" [35, p. 19-251].

The number of significant figures in these numbers is not to be taken literally. The costs are somewhat uncertain and the averted risks (in person-rem) are probably considerably more uncertain. In addition, the assumed benefit of $5000 per person-rem can be disputed as being too high or too low. Nonetheless, these calculations address an essential consideration in seeking to balance the needs for safety and economy. The challenge for reactor developers is to achieve designs that will be acceptably safe and economically affordable. Cost–benefit comparisons can be a guide in considering modifications to the designs.

14.5.2 Standards Adopted by Other Bodies

Other organizations have been more explicit than the U.S. NRC in setting forth criteria for core damage frequency. The International Atomic Energy Agency has established a group specifically charged with considering matters of nuclear safety: the International Nuclear Safety Advisory Group (INSAG). In a 1999 report on safety principles, INSAG indicated a severe core damage frequency target of under 10^{-4}/RY for existing reactors [39, p. 11]. It further suggested that the application of appropriate safety principles and objectives to future plants "could lead to the achievement of an improved goal of not more that 10^{-5} severe core damage events per plant operating year."

The U.S. utility industry, through a "Requirements Document" issued by the Electric Power Research Institute in 1990, also adopted a core damage frequency limit of 10^{-5}/RY for future light water reactors [40, p. 94].

The INSAG document, in addition, suggested as an objective for future reactors the "practical elimination of accident sequences that could lead to large early radioactive releases." No quantitative meaning was assigned to the term "practical elimination" and no serious analyst will ever claim, or talk in terms of, "zero risk." However, one can speculate that if one tries to interpret the words in terms of a quantitative criterion, "elimination" might mean at least a factor of 100 beyond the core damage frequency, which would correspond to a LERF of less than 10^{-7}/RY. This is perhaps as close to zero as can be meaningfully considered.

14.5.3 Standards for Future Reactors: How Safe Is Safe Enough?

There is no universal answer to the question of "how safe is safe enough?" The acceptability of a given risk depends on circumstances, including the risks involved in alternative options. Many auxiliary factors enter, including—but not limited to—whether the risk is created by one's own actions, the actions of external institutions, or the actions of nature. Experience suggests that, at any level of numerically calculated danger, risks associated with nuclear energy are far less acceptable to the public and to most policy makers than are many other existing risks we encounter (e.g., those

from automobile travel and the chemical emissions from coal-burning power plants).[35]

With these attitudes in mind, it is reasonable to conclude that the risk levels of the "useful subsidiary benchmarks" for present reactors suggested to the NRC (but not adopted by it)—of 10^{-4}/RY for core damage and 10^{-5}/RY for a large release of activity—would not meet a socially acceptable criterion of "safe enough" for a future world with, say, 4000 reactors. Taken literally, these numbers would imply a TMI-type accident every 2 or 3 years and a Chernobyl-type accident every 25 years. No matter what number of casualties is assumed for such an accident, it would not be acceptable to have a Chernobyl every few decades. The fact that the world accommodates to more severe tragedies from natural events and small-scale wars is probably irrelevant in terms of public response to nuclear power accidents.

On the other hand, a core damage frequency of 10^{-6}/RY to 10^{-5}/RY and a large early-release frequency of 10^{-7}/RY to 10^{-6}/RY might be satisfactory. Were such criteria met at the start, there would be only a small chance of either type of accident during the first decades of a large nuclear energy buildup. As discussed in Chapter 16, it may be possible to meet and exceed such standards with the new reactors that are now becoming available. It is pointless to attempt to estimate safety levels beyond a few decades, because continuing changes in nuclear reactor design—presumably with further safety improvements—would accompany a major revival of nuclear plant construction.

References

1. International Atomic Energy Agency, *Nuclear Power Reactors in the World*, Reference Data Series No. 2, April 2003 edition (Vienna: IAEA, 2003).
2. U.S. Environmental Protection Agency, *Accidents and Unscheduled Events Associated with Non-nuclear Energy Resources and Technology*, Report EPA-600/7-77-016 (Washington, DC: EPA, 1977).
3. Organization for Economic Cooperation and Development, Nuclear Energy Agency, *Achieving Nuclear Safety: Improvements in Reactor Safety Design and Operation* (Paris: OECD, 1993).
4. Uranium Institute, *The Safety of Nuclear Power Plants: An Assessment by an International Group of Senior Nuclear Safety Experts* (London: The Uranium Institute, 1988).
5. C. W. Forsberg and A. M. Weinberg, "Advanced Reactors, Passive Safety, and Acceptance of Nuclear Energy," *Annual Review of Energy* 15, 1990: 133–152.
6. International Atomic Energy Agency, *The Safety of Nuclear Power: Strategy for the Future* (Vienna: IAEA, 1992).

[35] Illustrating the disparity in public reactions, the death of five people due to a natural gas explosion and fire in Philadelphia on May 11, 1979, less than 2 months after TMI and in the same state, was virtually unnoticed [41, p. 930].

7. Ronald Allen Knief, *Nuclear Engineering: Theory and Technology of Commercial Nuclear Power*, 2nd edition (Washington, DC: Hemisphere Publishing Company, 1992).
8. American Nuclear Society, *Report of the Special Committee on Source Terms* (La Grange Park, IL: ANS, 1984.)
9. "Report to the APS of the Study Group on Radionuclide Release from Severe Accidents at Nuclear Power Plants," Richard Wilson, Chairman, *Reviews of Modern Physics* 57, no. 3, part II, 1985.
10. "Report to the APS by the Study Group on Light-water Reactor Safety," H. W. Lewis, Chairman, *Reviews of Modern Physics* 47, Supplement 1, 1975.
11. U.S. Nuclear Regulatory Commission, *Severe Accident Risks: An Assessment for Five U.S. Nuclear Power Plants*, Final Summary Report, Report NUREG-1150, vols. 1 and 2 (Washington, DC: NRC, 1990).
12. U.S. Nuclear Regulatory Commission, *Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants*, Report WASH-1400 (NUREG 75/014) (Washington, DC: NRC, 1975).
13. H. Kouts, "The Safety of Nuclear Power," in *The Safety of Nuclear Power: Strategy for the Future* (Vienna: IAEA, 1992), pp. 47–54.
14. U.S. Nuclear Regulatory Commission, *Risk Assessment Review Group Report to the U.S. Nuclear Regulatory Commission*, H.W. Lewis, Chairman, NUREG/CR-0400 (Washington, DC: NRC, 1978).
15. David Bodansky, "Risk Assessment and Nuclear Power," *Journal of Contemporary Studies* 5, no. 1, 1982: 5–27.
16. "World List of Nuclear Power Plants," *Nuclear News* 37, no. 3, March 1994: 43–62.
17. American Nuclear Society, *Report of the Special Committee on NUREG-1150, The NRC's Study of Severe Accident Risks* (La Grange Park, IL: ANS, 1990).
18. Senior Seismic Hazard Analysis Committee, R. J. Budnitz, Chairman, *Recommendations for Probabilistic Seismic Hazard Analysis: Guidance on Uncertainty and Use of Experts*, Report NUREG/CR-6372, UCRL-ID-122160 (Livermore, CA: Lawrence Livermore National Laboratory, 1997).
19. Electric Power Research Institute, *Use of Probabilistic Seismic Hazard Results: General Decision Making, the Charleston Earthquake Issue, and Severe Accident Evaluations*, EPRI Report TR-103126, prepared by Risk Engineering, Inc. (Palo Alto, CA: EPRI, 1993).
20. U.S. Nuclear Regulatory Commission, *Revised Livermore Seismic Hazard Estimates for 69 Nuclear Power Plant Sites East of the Rocky Mountains*, Draft Report NUREG-1488 (Washington, DC: NRC, 1993).
21. *Energy, U.S. Code of Federal Regulations*, Title 10 (1993).
22. T.E. Murley, "Developments in Nuclear Safety," *Nuclear Safety* 31 no. 1, 1990: 1–9.
23. T.E. Murley, "Safety Culture Indicators," MIT Safety Course (July, 1999), unpublished.
24. William D. Travers, *Status of Accident Sequence Precursor and SPAR Model Development Programs*, SECY-02-0041 (Washington, DC: U.S. Nuclear Regulatory Commission, 2002).
25. R. J. Belles, et al., *Precursors to Potential Severe Core Damage Accidents: 1997*, Report NUREG/CR-4674, ORNL/NOAC-232, Vol. 26 (Oak Ridge, TN: ORNL, 1998).

26. "Changes in Probability of Core Damage Accidents Inferred on the Basis of Actual Events," NRC staff report (forwarded to the Chairman of the NRC by James M. Taylor, April 24, 1992).
27. William D. Travers, *Status of the Accident Sequence Precursor (ASP) and the Development of Standardized Plant Analysis Risk (SPAR) Models*, SECY-03-0049 (Washington, DC: U.S. Nuclear Regulatory Commission, 2003).
28. "Performance Indicators: Another Successful Year in Performance, Safety," *Nuclear News* 45, no. 6, May 2002: pp. 28–30.
29. U.S. Nuclear Regulatory Commission, *NRC Update: Davis–Besse Reactor Head Damage* (November 2002).
30. "The Nuclear News Interview. The NRC's Brian Sheron: On Reactor Vessel Degradation," *Nuclear News* 46, no. 7, June 2003: 29–33.
31. U.S. Nuclear Regulatory Commission, "Use of Probabilistic Risk Assessment Methods in Nuclear Regulatory Activities; Final Policy Statement, *Federal Register* 60, no. 158, August 1995: 42622–42629.
32. U.S. Nuclear Regulatory Commission, "10CFR Part 50, Safety Goals for the Operation of Nuclear Power Plants; Policy Statement; Correction and Republication, *Federal Register* 51, no. 162, August 1986: 30028–30033.
33. William D. Travers, *Modified Reactor Safety Goal Policy Statement*, SECY-01-0009 (Washington, DC: Nuclear Regulatory Commission, 2001).
34. U.S. Nuclear Regulatory Commission, *Committee Voting Record, Modified Reactor Safety Policy Goal Statement* (Washington, DC: NRC, April 16, 2001).
35. U.S. Nuclear Regulatory Commission, *Final Safety Evaluation Report Related to Certification of the AP600 Standard Design,*" NUREG-1512 (Washington, DC: NRC, 1998).
36. U.S. Nuclear Regulatory Commission, *Safety Goals for Nuclear Power Plant Operation*, NUREG-0880 REV 1 (Washington, DC: NRC, 1983).
37. U.S. Nuclear Regulatory Commission, *Final Safety Evaluation Report Related to the Certification of the Advanced Boiling Water Reactor*, Report NUREG-1503 (Washington, DC: NRC, 1994).
38. U.S. Nuclear Regulatory Commission, *An Approach for Using Probabilistic Risk Assessment in Risk-Informed Decisions on Plant-Specific Changes to the Licensing Basis*, Regulatory Guide 1.174 (Washington, DC: NRC, 1998).
39. International Atomic Energy Agency, *Basic Safety Principles for Nuclear Power Plants, 75-INSAG-3, Rev. 1*, International Nuclear Safety Group Report INSAG-12 (Vienna: IAEA, 1999).
40. National Research Council, *Nuclear Power, Technical and Institutional Options for the Future*, Report of the Committee on Future Nuclear Power Development, John F. Ahearne, Chairman (Washington, DC: National Academy Press, 1992).
41. *The World Almanac and Book of Facts 1980* (New York: Newspaper Enterprise Association, 1979).