NUCLEAR SAFETY AND RELIABILITY

WEEK 1

[Author's Note: These notes were used in lectures at the University of New Brunswick, in the period from 1984 to 1991. As a result, some of the references, methods, and discussions are somewhat dated. For example, the LOCA analysis methods described in Week 8 and Week 9 notes have been superseded by superior analytical models and experimental results. The notes as presented do, however, give a simple guide to the models in use today. – November 2003.]

TABLE OF CONTENTS - WEEK 1

1. Objectives of Course

This course is intended as an introduction to radiation protection principles and practice in commercial nuclear power plants. Emphasis is placed on those aspects of radiological safety that relate to accidents rather than to normal operation of the plant. The course material assumes a fundamental knowledge of low-energy nuclear physics, nuclear reactor physics, heat transfer, and fluid dynamics. ChE 3804, Introduction to Nuclear Engineering, or its equivalent gives a useful background to this course, but these background facts can be collected with a reasonable amount of extra work during the course. Reliability concepts are used to introduce the concept of industrial plant safety evaluation within the framework of risk, defined as the product of accident probability and consequence. Safety characteristics are described for the various types of nuclear reactor now in commercial use. The framework of the Canadian reactor licensing system is discussed in the context of risk management. Risk comparisons are drawn between nuclear plants and other industrial enterprises.

On completion of the course, the student should:

1. Understand the nature of nuclear plant risks and the means used to control and evaluate those risks.
2. Know the basic concepts of reliability analysis, fault tree construction, and event sequence diagrams used in analysis of the public risks of industrial plant operation.
3. Have an understanding of the essential elements of safety practice in large industrial systems.
4. Have a semi-quantitative grasp of the event sequences following rupture of a primary coolant pipe in a CANDU reactor, and of the way in which the plant safety systems would respond to mitigate the consequences of such an accident.
5. Recognize the methods available for evaluating risks and benefits of alternate technologies for power generation and for comparing these risks with those involved with other activities of the society.

It is expected that preparatory reading will be done before certain lectures; identification of such reading material is given in these notes. Tutorial hours will be devoted to investigation of accidents that already have occurred, with a brief outline of the plant involved, the sequence of events, consequences, causes of the accident, and lessons learned. Information will be gathered from handouts and library materials. Students will work in groups of two or three; each group will present their findings to the class on completion of each accident investigation. Tutorial marking will be based on individual performance rather than group performance.

2. References

- John R. LaMarsh, “Introduction to Nuclear Engineering”, Addison-Wesley Publications, 1983
- Karl O. Ott and John F. Marchaterre, “Statistical Evaluation of Design-Error Related Nuclear Reactor Accidents”, Nuclear Technology, Vol. 52, pp179-188, Feb. 1981
- Norman J. McCormick, “Reliability and Risk Analysis”, -- *Methods and Nuclear Power Applications*, Academic Press, 1981
- H-J Hoffmann, L. Oedekoven, K.O.Ott, “Statistical Trend Analysis of Dam Failures Since 1850”, Report, Kernforschungsanlage Jülich GmbH, Feb 1984
- Mary Douglas and Aaron Wildavsky, “Risk and Culture”- An Essay on the Selection of Technological and Environmental Dangers, University of California Press, 1983
- Aaron Wildavsky, “No Risk is the Highest Risk of All”, American Scientist **67**, 32 (1979)
- D.A. Meneley, “A Reactor Cannot Explode Like a Nuclear Bomb”, URL http://canteach.candu.org under AECL-contributed technical notes.
- Kaplan and Garrick, "On the Use of Bayesian Reasoning in Safety and Reliability - Three Examples", Nuclear Technology 44, 231, 1979
- "PRA Procedures Guide", NUREG/CR-2300, Vol.1
- Ralph Fullwood and Robert Hall, “Probabilistic Risk Assessment in the Nuclear Power Industry”, Pergamon Press, 1988, ISBN 0-08-36362-8

3. Hazards

The word "hazard" is used here to describe a potentially dangerous material or condition. If there is no chance of the hazard producing actual harm, the danger level or risk is zero.

There are many hazardous materials and situations present in any operating nuclear plant. Industrial safety practices must address all of these potential risks. This course is, however, concerned only with radiological hazards and their chance of producing harm. The largest part of the course is concerned with only part of this topic - radiological risks during and after accidents.

Radiological safety comes down to protection of the human body from those effects of ionizing radiation which are caused by plant operation. There is no requirement here to provide protection against radiation from natural or medical sources; in the former case it is presumed to be inevitable, and in the latter the benefits are presumed to greatly outweigh the risks. (an interesting exception for natural radiation is developing with regard to high radon content in tightly closed buildings). It is presumed that protection of humans is sufficient to cover the effects on all other life. It is further presumed that damage incurred is directly proportional to the incremental amount of radiation received and is independent of the rate of damage (in spite of some evidence to the contrary).

Radiation damage occurs at the molecular level inside live cells. Ionizing radiation can be thought of as a stream of "bullets" emitted by radioactive materials; these "bullets" remain dangerous only as long as they have enough energy to produce ionization of other atoms and molecules. Direct knock-on collisions with atoms within DNA molecules can affect the functioning of the molecule. Transmutation of one element into another can affect the cell chemistry. The most important damage mechanism is production of ions within a cell which then react chemically with other cell components.

This disruption of the cell's chemistry has an effect which is identical to that produced by many non-radioactive materials; it can result in cell death or in activation of cancerous reproduction. These effects show up mostly during cell division (mitosis) so that cells which are dividing are most sensitive to damage. Some tissues (e.g. bone marrow, intestinal lining, thyroid, gonads) are more sensitive to radiation than others due to their high rate of cell division. Rapid cell division due to growth is the reason that fetuses and young children are more sensitive to radiation than are adults.

Radiation damage to cells will be addressed in Week 2. **Preparation - read Chapter 9 of LaMarsh, Section 9.1 to 9.5.**

4. Radioactivity and the Fission Products

The fission or splitting of fuel nuclei is the key event by which heat is produced in nuclear fission reactors. Atoms of the fuel combine with an incident neutron to form unstable nuclei which then achieve relative stability by splitting into two or more fragments, generally of comparable size. One to five additional neutrons are released at the same time; these are necessary to the operation of the reactor since they initiate the further fissions which produce a

continuous chain reaction. The fission products (elements near the middle of the periodic table created by fission of the fuel) constitute the major and characteristic hazard of the fission reactor. The reason is that some of them are still unstable, and so decay with unique half-lives until they eventually become stable atoms. During decay they emit various kinds of ionizing radiation (the "bullets" referred to in the previous section). The major emissions from fission products are beta particles (electrons) and gamma rays.

Damage to an individual's body can be reduced by any of three mechanisms: (1) **shielding** - that is, by placing ordinary material between the radiation source and the body so as to stop most of the "bullets", (2) increasing the **distance** between the radiation source and the body - to reduce the number of "bullets" hitting the body per second, and (3) by reducing the **time** spent in the path of the radiation - to reduce the total number of "bullets" which strike the body.

The risk to operating staff from fission products retained within the reactor, and from the fraction of the neutrons and gamma rays generated in fission which escape from the reactor core, is reduced to acceptable levels by suitable shielding. The principal risk is then from fission products which may be released to the environment. Other radioactive species apart from fission products also are formed in fission reactors by the absorption of neutrons in fluids and structural materials outside the fuel. These so-called activation products usually represent a much smaller risk than the fission products; therefore, they are mentioned only briefly in the following material.

Radiological safety is achieved if the plant delivers zero incremental amount of radiation to humans during normal operation or following accidents. This ultimate goal cannot be met absolutely. (Radiological safety is demanding in this respect because even very low concentrations of radioactive materials can be measured easily - for many chemicals it is the detection threshold that determines the acceptance limit). Therefore, standards must be established which determine maximum permissible radiation **dose levels**. Normal operation standards are based directly on limits derived by the International Commission on Radiological Protection (ICRP); in the case of potential accidents one must specify not only the permissible level, but also the permissible accident frequency corresponding to this level. In Canada, permissible values are set in the AECB Siting Guide. We are concerned here with those isotopes that might contribute to these dose levels.

Fission product isotopes accumulate inside solid fuel elements during operation, in an amount proportional to the energy produced. These fuel elements are removed periodically during refueling operations; nevertheless, at any time there is an "equilibrium" quantity that might be released to the environment following an accident.

The basic safety strategy in a solid-fuel reactor such as CANDU is to keep the radioactive materials formed in the fuel inside the fuel sheath, inside the primary pressure boundary, and inside the containment envelope. This strategy localizes the dominant source to the greatest possible degree depending on the accident circumstances. All engineered safety systems' design is based on this strategy. This localization strategy is extended to peripheral devices that contain some radioactive materials (e.g. ion exchange [IX] columns). The secondary strategy is to place shielding materials between radiation sources (e.g. core, primary piping, IX

columns) and people. The exclusion zone surrounding the plant, in which no permanent habitation is permitted, is really an extension of this shielding strategy - it provides an "air shield" as well as distance between people and direct radiation which might be emitted after a major accident. The exclusion zone also decreases the concentration of any radioactive materials emitted in these circumstances through mixing with ambient air. This airborne pathway is the largest potential contributor to public radiation dose; liquid pathways also must be considered in safety systems design.

Given that there are several barriers to release of radioactive materials and the "zoo" of isotopes known to be present in the fuel (about 200 different isotopes of nearly 40 different chemical elements), accident analysis is concerned with how these barriers might be penetrated by these isotopes during various reactor accident sequences. It is helpful to first examine the relative importance of each isotope to various accident sequences. Safety design then can concentrate on fewer potentially dangerous contributors.

The principal factors determining importance of any single isotope are (a) its total core inventory (dependent on fission yield, decay chains, half-life, and reactor operating history), (b) its physical properties, including the nature of its radioactive emissions and its volatility, (c) its chemical properties, (d) its transport through the atmosphere (adsorption and deposition properties), and (e) its biological impact (uptake rate, biological half-life, specific effect on organs, etc.). Some isotopes may not be important dose contributors themselves but might produce biologically important isotopes on decay after emission.

Some of these importance-determining factors are inherent (e.g. fission yield, half-life, radiation emission characteristics). Other factors depend on the features of the plant and the particular accident sequence - such as release temperature and flow pathway before release - because these details will determine the chemical and physical forms of many isotopes. It follows that the importance of each isotope depends to some extent on the accident scenario under investigation. However, in general, the most important radionuclides in severe accident sequences are iodine, cesium, the noble gases, and tellurium. Iodine has been considered the dominant single isotope. Figure 1.1 gives a rough sketch of the pathways expected following CANDU loss of coolant accidents. The form of compounds existing in the fuel depends on pre-accident conditions and, to some degree, fuel overheating prior to fuel sheath failure.

FIGURE 1.1
FISSION PRODUCT BEHAVIOUR UNDER ACCIDENT CONDITIONS

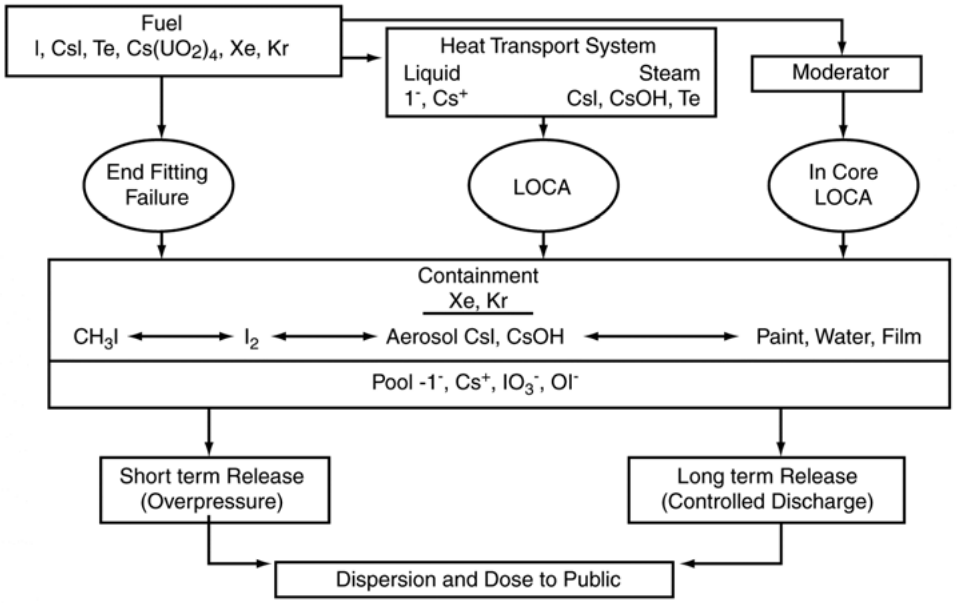

The noble gases are, of course, non-reactive and will behave as gases. The other compounds will either deposit on metal surfaces or react with the water-steam mixture in the channel. The CsI is of particular interest because of its low volatility and strong affinity for the liquid phase - but it would react to CsO if there were free oxygen available, and thereby release elemental iodine. Another iodine compound with relatively high volatility is CH_3I, which can be formed by reaction with organics in the containment space, either in the atmosphere or the water pool. The actual distribution of radioactive species prior to release depends strongly on a number of parameters; two of the most important are oxidation potential and pH.

5. Accident Initiating Mechanisms

Given the primary strategy of isolating the major source of radioactive materials, a sufficient method of doing this is to keep the fuel sheaths intact. Fuel sheaths can fail under nominal operating conditions due to a number of slow degradation mechanisms such as manufacturing error, pellet-sheath interaction, stress-corrosion cracking, etc. Low failure rates are desirable because a contaminated HT system defeats the first barrier to release of activity to the public, increases maintenance staff dose levels, and decreases the sensitivity of fuel failure detection equipment. Experience in LWR reactors has been variable, but recent CANDU experience has been excellent. On-line fueling gives CANDU an inherent advantage in that leaking fuel bundles can be removed without loss of generation (provided they can be located). In LWR's there is a strong motivation to continue operating until the next scheduled fueling outage – so these reactors hold, on the average, a higher fission product concentration in their primary coolant.

Considering off-normal (accident) situations, fuel failures can result from three main mechanisms: excess heat production leading to fuel melting and sheath rupture, insufficient cooling leading to sheath dryout - leading to overheating and rupture, and mechanical damage. These physical mechanisms result in several different potential accident classes being defined for accident analysis. Each of these classes is separated into a number of categories depending on the assumed progress of the presumed accident sequences. Transient analyses are carried out for each category to arrive at consequence estimates. The frequency of each event is estimated by examining the probability of various initiating mechanisms for the sequence as well as the probabilities implied by the assumptions regarding the progress of the sequence.

In licensing analysis, the calculated accident progression is distorted by the limitations imposed by the Siting Guide (e.g. no mitigation by the reactor regulating system [RRS]) and by the conservative models forced upon the analysis itself. Therefore, the licensing analysis does not represent a best estimate of either the frequency or the consequences of any particular sequence. In addition, the assumed sequence is unlikely to be the same as any real sequence observed during operation. The purpose of licensing analysis is to demonstrate, conservatively, that the safety systems are capable of mitigating accident consequences for a broad range of initiating events.

If the expected behavior of the system is required for operator training or any other purpose, a separate analysis should be done using design center values of parameters.

Following is a brief description of each of the major accident classes.

Transient overpower (TOP)

This accident class includes rod ejection in a PWR, turbine stopvalve closure without bypass in a BWR, and loss of regulation in an FBR or CANDU. It does not include the LOCA, because in that case overpower occurs (if the void reactivity is positive) as a consequence of the initiating event. The TOP transient can, of course, result in HT overpressure and eventually lead to a loss of coolant [LOC] through rupture of the heat transport [HT] system pressure boundary. An extreme case of TOP transient effects can be found in the recent Chernobyl accident in the Soviet Union; TOP's must be prevented with a high degree of reliability because of their potential for endangering all of the natural and engineered barriers to release of dangerous radioactive materials to the public.

Fuel sheath failure can result either from melt-induced rupture or dryout- induced overheating. In either case, the concurrent event is coolant voiding; in a CANDU this results in further power increase due to positive void reactivity. In pressurized water reactors [PWR] or boiling water reactors [BWR] the void reactivity is negative, so voiding leads to power decrease. Positive void reactivity is one of the main reasons for inclusion of two independent shutdown systems in CANDU, versus one in the LWR. Current fast breeder reactor [FBR] designs also use two shutdown systems because the coolant void reactivity is positive in that case as well. The major safety weaknesses of the FBR are illustrated by TOP sequences, the short prompt neutron lifetime can result in very rapid fuel energy accumulation, and fuel movement after melting is likely to result in one or more further TOP pulses. Gas cooled reactors [GCR] are very resilient to TOP events because of the high heat capacity of the graphite moderator. The

Soviet RBMK reactor is very susceptible to TOP events, so that the designers must take careful measures to ensure rapid termination of this class of accidents.

In CANDU designs the TOP accident without shutdown has been made extremely improbable, so is not analyzed in licensing submissions. If an unprotected TOP occurs the first inherent negative feedback would come from moderator voiding, because fuel temperature feedback is either zero or slightly positive at high temperatures. Moderator voiding is caused by massive melting, pressure tube/calandria tube rupture, and injection of molten fuel. Potential vapor explosions pose an overpressure challenge to containment in addition to completely destroying the core. This gives an indication of the reason for paying close attention to reliability and capability of shutdown systems in CANDU. On the other hand, if rapid shutdown is accomplished, the existence of moderator heat removal mechanisms almost entirely precludes fuel melting in CANDU. Relative to LWR systems, CANDU relies more heavily on successful shutdown systems action and less heavily on emergency coolant injection for mitigation of the consequences of major accidents. Since shutdown systems are much simpler and easier to test than are injection systems, the net safety advantage lies with CANDU. Containment is important to all reactor concepts; however, any design that effectively eliminates massive fuel melting from consideration greatly simplifies the requirements placed on containment design.

Loss of flow (LOF)

LOF accidents can be local due to flow blockage or global due to loss of pumping power. Reduced heat removal capability necessitates rapid reduction of power so as to prevent sheath dryout and overheating failure. In LWR or FBR systems a local flow blockage can propagate laterally to other fuel channels and produce large-scale fuel melting. In the FBR system it is relatively simple to install exit temperature monitors at the outlet of each fuel channel; the exit coolant is normally well below boiling, so there is margin for detection and shutdown action.

In CANDU at full power, an inlet flow blockage sufficient to produce dryout rapidly results in fuel melting in the single channel involved, pressure tube/calandria tube failure, and ejection of molten fuel from the channel into the moderator. Shutdown signals are not received until the blocked channel fails. Experiments and analysis show that no failure propagation to adjacent channels will occur. However, severe damage might be suffered by core internals (e.g. control rod drives). The economic penalty of such an accident would be very high - the prime reason for regular testing of channel flowrates. At present the only way to test flowrate in a boiling channel is to reduce power so that the exit coolant is subcooled; flow is then inferred from coolant temperature rise from channel inlet to outlet.

This accident illustrates one important safety weakness of CANDU. Placement of the fuel immediately adjacent to the pressure boundary gives a second vital reason for preventing fuel overheating when the HT is pressurized - to prevent failure of the pressure boundary.

A global LOF in LWR or CANDU must be followed immediately by reactor trip to restore the balance of heat production and removal. In the BWR, loss of flow results in immediate negative feedback to the power due to the increased core void; in PWR this is somewhat delayed because the exit coolant is below saturation temperature. Thermosyphoning is sufficient to remove decay heat provided that HT inventory is maintained. Redundant power

supplies to HT feed pumps are required to prevent eventual depressurization and loss of circulation. Modern FBR designs have sufficient natural circulation and negative reactor power feedback to sustain a complete LOF without shutdown action.

Loss of heat sink (LOHS)

This transient can be initiated by secondary side failure (e.g. loss of feedwater, steam main failure). Fuel failure can be initiated by LOC produced by HT overpressure and rupture, by failure of steam generator tubes leading to LOC, or by loss of inventory through HT relief and subsequent loss of HT pumping power resulting from turbine trip consequent to the secondary side event. Redundant heat sinks and power supplies must be provided to terminate this sequence successfully after reactor trip.

This accident class is very important to LWR reactors due to their lack of full-pressure residual heat removal systems and due to the possible core meltdown consequences to containment.

Loss of primary coolant (LOC)

This class of accident sequences defines many of the limiting requirements for safety systems. They range from leaks just beyond the capability of the HT feed system to maintain pressure control, up to failure of the largest piping in the HT loop. The sequences are particularly important to any reactor with positive coolant void reactivity, because they set the detection and performance requirements of shutdown systems. Pressurized HT systems are particularly vulnerable because of the rapid vaporization of depressurizing coolant, especially around the fuel. For breaks large enough to result in rapid voiding around the fuel it is extremely difficult to prove that fuel sheath failure will not occur to some degree. In contrast, the atmospheric pressure HT system in FBR's virtually eliminates the LOC class of accident, particularly in "pot" designs.

In PWR and BWR reactors, large piping is used to circulate coolant to and from the core. Potential breaks in this piping would lead to rapid draining of the fuel from the core, with consequent fuel heatup. Emergency coolant must be added quickly to prevent fuel melting, with consequent massive release of fission products and a potential for melting through of the reactor vessel. The Three Mile Island Unit 2 accident came very close to this extreme condition.

The distribution of piping lengths in CANDU designs leads to the general conclusion that about 98% of all LOC sequences will result in leakage rates less than or equal to rates typical of rupture of the largest feeder pipe. This conclusion is reinforced by the frequent opening and closing of the HT boundary during fueling. Potential malfunction of the fuel channel closure system is a significant contributor to the overall LOC frequency in CANDU.

Fuel bundles could suffer mechanical damage during a LOC event in CANDU because they are held loosely in the channel by the flow force, with some space between the first bundle and the upstream shield plug. The abrupt flow reversal that follows channel inlet breaks could drive the bundles against the upstream shield plug and rupture the sheaths. Limited testing has shown that damage is unlikely unless the upstream space is quite large.

A subset of LOC that is particularly important to public safety is HT auxiliary piping failure outside containment. Isolation of this piping at the containment boundary is the means for terminating this sequence.

End-fitting failure (EFF) - CANDU only

CANDU and RBMK are the only reactor systems in which fuel can be ejected from the HT system during operation. Fuel failure occurs through direct mechanical damage. Recent work indicates that the frequency of occurrence of fuel ejection is quite low; of the same order as the frequency of large LOC. The sequence poses some difficult problems to containment design because of the rapid release of fission products (albeit in limited quantity) very early in the containment isolation sequence.

Accident Precursors

We tend to think of the above (or some longer list) as the beginning of accidents. Looking deeper, we try to find out the cause of the accident (a slight contradiction in terms), or even further the cause of the cause and so on back to the beginning of time. This is an ancient philosophical argument that has no solution. It can be cut shorter in our case to the first reactor that started up on December 9, 1942. Did we think about the safety of these plants in the most effective way? Was the design concept properly framed? Was the project funded sufficiently well? Was the staff organized effectively? Did the designers do their homework in the detailed design phase? Did they really look for potential failure modes, or did they sometimes assume that systems they designed would always work right? Was the design critically reviewed? Did the manufacturers build conscientiously to the specifications? Was the plant constructed properly? Was the maintenance done thoroughly? Did the designers tell the operating staff all relevant facts about their machine? Did the scientists tell the designers all relevant facts about the raw materials of their engineered product? Were the people running the unit properly trained? Were they presented with the necessary information by their instruments? Did they pay attention to the information and interpret it properly? Did operations' management encourage good performance and support those on the front line, or did they only criticize mistakes? - and so on into the night.

This is the stuff of the nuclear safety business. It is much like any other safety business in that we try to examine causal mechanisms and either eliminate or account for them. All failures are human failures in one sense or the other - nuclear accidents get more attention than average just at the moment. Those on the "front line", the operators, tend to gather more than their share of the blame for failures because the root causes are often indefinite. Success is no news. The people who achieve it get few rewards, even though they may richly deserve them.

There is a positive fact that goes back to the philosophical question of free will; that is, most errors (which occur) can be detected and corrected before they lead to consequences. **Since it is obvious that no radiological accident can occur before the plant begins to run**, the people directly responsible, who live "at the pointed end of the ship", are the operating staff. Acceptance of the plant for operation transfers full responsibility and authority for safety, including correction of all potential design and construction errors, to the operating staff. If they

take this responsibility seriously and develop a healthy "safety culture", a very high degree of safety can be achieved.

Summary of potential hazards

The previous listing does not include potential releases to the environment through peripheral systems. Such events have lesser potential for major consequences than do sequences involving fuel failures. On the other hand, relatively high frequency and low consequence events may well dominate the real operating risk spectrum of nuclear plants (with one big "if" - if the low-dose health consequence is really linear).

Accident analysis to consider all of these sequences, and combinations, has become a major activity in plant design offices in recent years. A large institutional framework has been created to deal with this work as well as with regulatory review of the work. The quality and depth of accident analysis has been improved considerably over the past ten years, as has the understanding of post-accident behavior. Unfortunately, the institutional structure threatens to be self-perpetuating beyond the range of its usefulness to public safety. A critical future task is to establish the reasonable bounds of risk investigation and to set firm acceptance criteria.

By any reasonable measure the equipment, procedures, and people now in place at CANDU stations are more than adequate to meet the original goal - that nuclear-electric generation should be safer than similar industrial activities. However, few would claim that the industry is now in a strong position in our society. The present demand for even better plant safety is driven by public fear and mistrust. Little can be done to improve the climate of licensing until this attitude is changed to one of acceptance and, hopefully, active public support.

6. Experience with Failure Trends

Ott & Marchaterre published, in 1981, a small paper that could have a large effect on our perception and management of risk. This paper addressed seven historical reactor accidents in which fission products were released from fuel. This work initiated studies of accident histories in various industrial facilities (at KfK Jülich). The general concept is indicated in Figure 1.2.

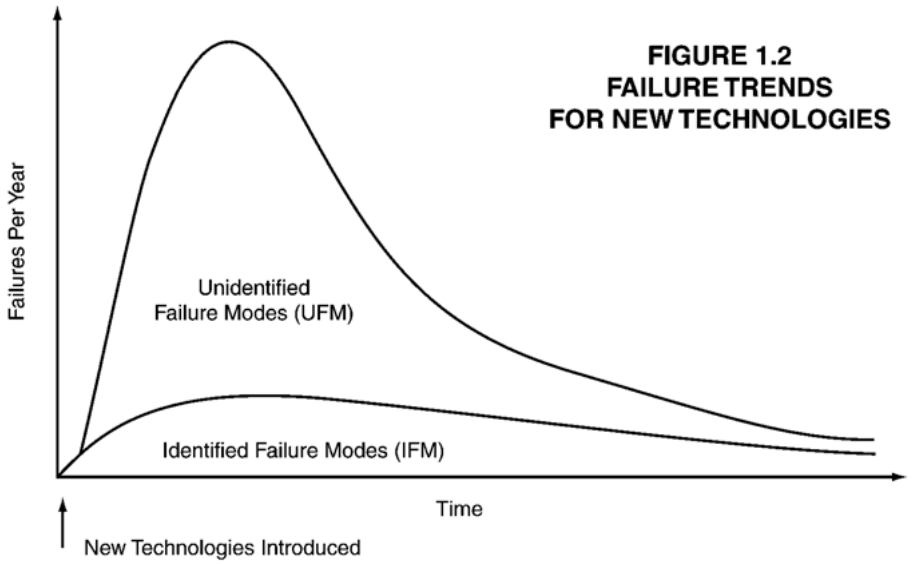

After a new technology is introduced, the failures per year increase more or less linearly in the beginning (assuming the number of units in-service increases linearly). "Failure" could mean death if we are talking about aircraft crashes, rupture if the subject is pressure vessels, or fission product release from fuel in the case of the nuclear industry. As the number of units becomes significant a reaction sets in - the failures either create a nuisance, they cost too much to be tolerable, or public pressure forces the proponents to clean up their act. An example of the last kind is boiler explosions. According to a recent article printed by Scientific American in their "100 years ago today" column, in the late 1800's about 90 people were killed each year by boiler explosions of various kinds. The reaction led directly to creation of the ASME boiler and pressure vessel code. As for steamship boiler explosions it was Lloyd's of London who applied the pressure - obviously in order to save on insurance claims. Commercial competition can also be helpful; no one will buy garbage bags that break easily unless the better product is outrageously expensive. This underlines an important concept in safety design -- the customer gets what he is willing to pay for. In our particular case it seems that the public is willing to pay almost any price for "safety" unless there is a cheaper electricity generation alternative. It is not so clear that people appreciate relative levels of safety between technologies. It is a very important subject to discuss with them before our technology gets priced out of the market by excessive safety precautions.

Returning to Figure 1.2, when a new technology is introduced there is usually some level of appreciation of how it might fail; that is, some failure modes have been identified. The actual failures should (but don't automatically) provide the lessons necessary to avoid recurrence. The stock of Identified Failure Modes (IFM) increases. Nevertheless, it is in the nature of accidents that most of them contain some surprises (e.g. the Pickering pressure tube failure, the Comet 1 disasters); therefore, in the early days the overall failure rate is dominated by Unidentified Failure Modes (UFM). With hindsight, evidence usually can be found that someone considered the particular critical failure mode but his knowledge was never put into practice. A crucial issue - when we talk about IFM and UFM we must clearly understand: identified or unidentified by whom? Obviously, the information must be in the hands of those who need it - designers, builders, QA inspectors, or especially operating staff.

Coming back again to Figure 1.2, in a well managed industry such as commercial airlines or nuclear energy the failure rate eventually decreases even though the population of units increases. The fraction of IFM in the total failure rate increases with understanding and experience. In the long run the total failure rate becomes constant and is accepted by the public (assuming that the number of units in service is now constant). This is the only practical definition of "acceptable risk" of a technology. Safety management encompasses all the actions required to get the risk to this level and to keep it there.

Another useful example from Ott's work on technological risk is the history of large embankment dam failure in the US, as shown in Figure 1.3. This is an isotonic regression plot of the specific failure rate - failures per dam-year, plotted versus the year of failure. The vertical lines represent actual dam failures (not all shown); the horizontal bars indicate the regression estimate of the failure rate based on experience up to that time.

There was a rash of failures in the early 1900's that apparently drew a reaction from either the designers or the regulators (it would be interesting to study the history of this. Perhaps those were exceptionally wet years.) Whatever the cause, the failure rate decreased markedly and steadily after that time, until by 1979, the last year of the study, the rate estimate was down to about .00015 per dam-year. The Grand Teton dam failure is at the end of this plot. This apparently was a classic case of shaving the margins and ignoring some basic rules of dam design and construction. It doesn't affect the average because of the large number of existing embankment dams, but this would do little to compensate the families of those killed or to excuse negligence on the project. There must be a lesson here. It seems to be that generalities are fine when we are examining the overall level of safety of any technology, but it is the specifics that are important at any given installation.

There is another general lesson in this Figure. If the expected failure rate per dam-year had been obtained simply by averaging the annual rate, the expected rate in 1979 would have been much higher than .00015. This procedure ignores the effect of learning the fact that some failure modes have been "taken out" of the system by design, by regulation, or by sound operation.

This procedure can be applied at any level of detail by proper definition of the "failure" to be investigated.

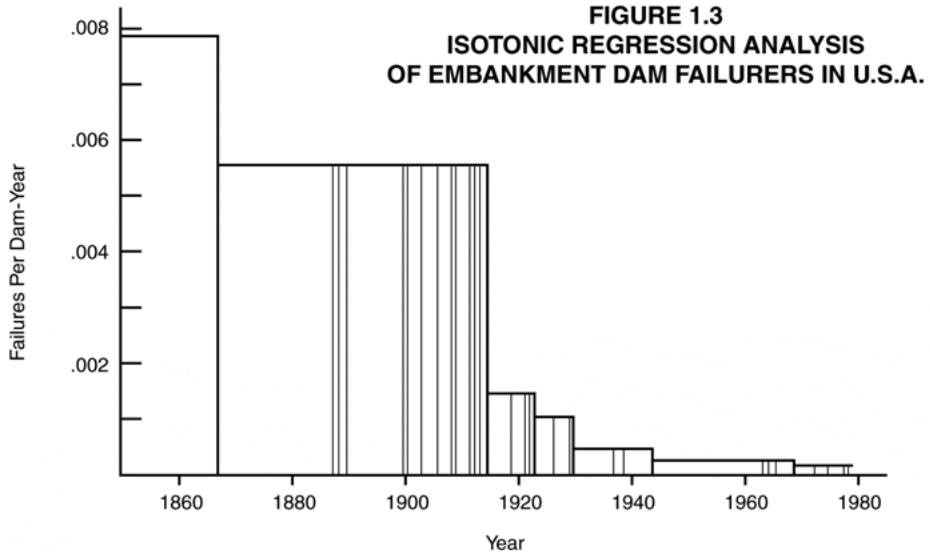

Figure 1.4 is another general sketch of recent reactor operating experience, with an impression of how the specific failure rate curve should behave in the future (it is still too early to tell). It is fair to assume that all the analysis done on the TMI accident and examination of related systems has done some good and has enlightened the subject, so that UFM rates can be expected to go down. Most of these should not show up in the IFM plot because they have been eliminated by design. It is entirely possible, however, that the changes made in the name of safety have introduced new UFM sequences which are now "lurking" in the system. It is important to remember that everyone operating a plant needs to learn these lessons.

The other possibility is that the IFM-related event frequency might increase in the future because of relaxation of good operating practices or wear-out of equipment. In examining the performance of a system (component, plant, or whole technology) there are some simple tests of the procedure that can be used to test for significance of apparent trend reversals. These are only triggers for investigation to determine the cause. They can be used as a management tool.

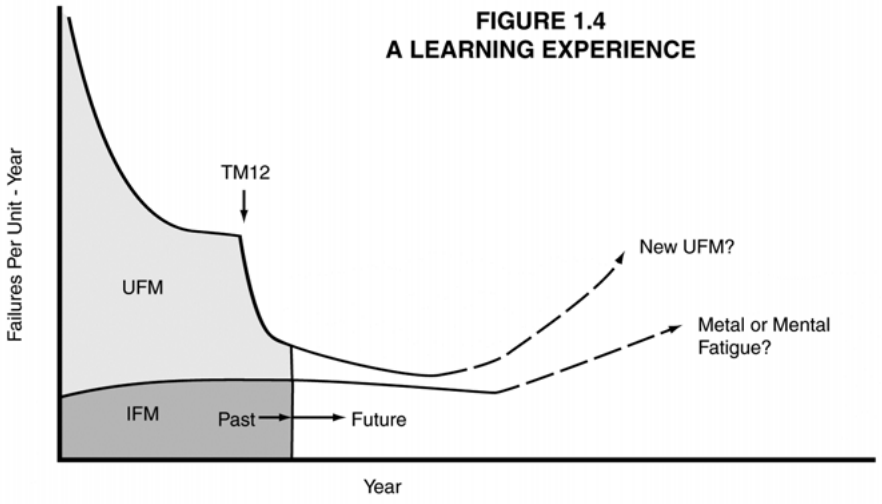

7. Prevention of Fission Product Release from Fuel

The "serious" phase of any accident sequence in a CANDU reactor with respect to public safety begins with failure of the HT pressure boundary. Though there are many potential initiating sequences that can eventually lead to this phase, public safety consequences are minimal unless HT failure actually occurs. Pressure boundary failure may be postulated (LOC sequences) or it may be the result of other initiating events (e.g. malfunction of feed/bleed/relief system or fueling machine operation). Fuel sheath failures are either nonexistent or unimportant unless this phase is reached, because an intact pressure boundary is an effective container for any fission products released to the coolant

The prevention category of accidents includes all common-cause failures such as earthquake, fire, flood, aircraft impact, and external explosions as well as a number of sequences such as loss of reactivity control, loss of forced circulation, loss of one HT pump, feedwater and steamline breaks, etc. Under the assumption that overheated fuel might lead to failure of pressure tubes, the operating power limit criterion of "no centerline melting" is defined. Reactor trip parameters and setpoints are selected so that this limit is not exceeded for a wide range of initiating events. None of these events leads to significant release of radioactive materials to the environment. The residual probability of fission product release represents UFM sequences.

One major emphasis of Canadian regulatory actions over the past few years has been on the prevention side. This may be a source of irritation to the designer, to whom it represents additional capital cost, and who might consider that the investment made in special safety systems is more than adequate to the purpose of the Regulatory Agency's role - auditing of public protection capability. This emphasis on prevention is, on the other hand, at least in the same direction as the Operating Company's goal of plant protection. There might well be differences as to degree of protection required, but there should be no disagreement on the objective. Since none of the accident sequences covered by prevention actually lead to significant public consequences there might also be grounds for agreement on the measure of adequacy - the actuarial risk of plant shutdown or damage. There are many bridges to cross before such an agreement could be contemplated.

Conduct of accident analysis directed toward prevention must use probabilistic methods either implicitly or explicitly - there are no absolutes. The use of fault tree and event tree analysis, along with the accompanying consequence analysis, merely provides a logical framework to aid the decision as to how much is good enough.

8. Mitigation of Consequences

Shutdown is by far the most important mitigating action in CANDU; if the reactor is shut down promptly after a process failure, then the possible consequences of the failure to public health are essentially nil. Fortunately, shutdown is the easiest function to design for high reliability - it requires comprehensive fault detection and reliable activation mechanisms, but then involves only a change of state from "poised" to "fired".

The superficial structure of accident analysis for Loss Of Coolant (LOC) sequences (which are about the only IFM accidents requiring mitigating action other than shutdown) is clear

in the Siting Guide - single failure and dual failure analysis with impairments of four or five sub-systems for each of emergency coolant injection and containment. The simplicity ends here. Given that the primary pressure boundary is presumed to have ruptured, the objective is to contain the fission products in the fuel sheath or inside the containment boundary. In the former case one is interested in cooling each element in each of several thousand fuel bundles so that the rupture threshold is not exceeded, by injecting water into a hot piping system with a hole in it! No one knows at time zero where the hole is or how big it is. If the break is large, some of the elements are steam-blanketed at time zero (effectively) and so are heating up more or less adiabatically. Huge amounts of money have been expended around the world in trying to design a cooling system for the large LOCA and to prove its effectiveness. A very simple analysis done in 1960 and reported by W.G. Morison describes the essence of the problem and recommends a solution.

The most difficult aspect of emergency cooling systems design has been deciding what the performance objectives should be. The most easily understood set of objectives was prepared by M. Milicev as part of a submission to the Advisory Committee on Nuclear Safety [ACNS] in August 1981. The ACNS evaluated all the submissions received and prepared a report for the President of AECB. Aside from some additional experiments, which are now almost complete, the Committee supported Ontario Hydro's position on the issue. This seems to have settled the definition of performance requirements, though the issue is still considered live by some AECB staff members. A large amount of R&D money is still being expended on this topic.

The question of how to predict the actual performance of an ECI system is largely a matter of judgment as to the reliability of the thermal-hydraulics models under various conditions. For the vast majority of break sizes and locations these models give quite reliable answers. However (remembering to concentrate on the worst case) there are some conditions under which some channels will reach stagnation flow conditions. There are two choices - either take the existing model as the best available and use it in spite of its weaknesses, or devise a model that can be confidently supported. The so-called limit consequence analysis was devised under conditions pertaining to Bruce A at that time: a gravity-driven ECI system. The limit-consequence analysis shows that channel integrity is maintained with no fuel melting for the full range of breaks, even when injection occurs very late. Adoption of high-pressure injection changed the expected performance of the ECI system considerably; at the same time there remains considerable doubt that it could actually prevent fuel sheath failure after a large LOC event. With these systems installed the objectives established by Ontario Hydro can be met easily. Work is currently underway to improve the capability for predicting the injection system's performance over the whole range. Realistically, it is unlikely that this will be done to the high standards of proof required by licensing analysis within the next 20 years.

Containment performance analysis is somewhat more straightforward and reliable, except perhaps during the very early pressure transient before any significant amount of fission products have been released into the building. The key question in containment analysis (other than peak pressure) is the space and time distribution of fission products released from the fuel. Until recently, very conservative assumptions were necessary. Recent work on fission product chemistry gives hope that containment release predictions for any given condition will decrease by two or three orders of magnitude.

Figure 1.5 shows the effect of preventive and mitigative design actions on the probability of fission product release to the environment. It is not possible to reduce the residual frequency of UFM events to zero, by definition.

FIGURE 1.5
EFFECT OF DESIGN ON ACCIDENT PROBABILITY

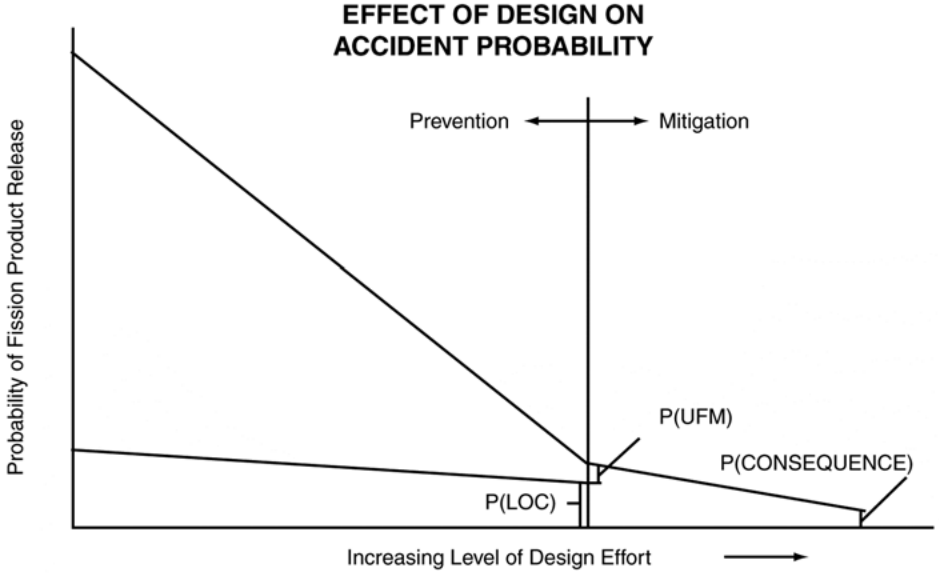

(This category includes initiating mechanisms such as unexpected operator actions, unanalyzed equipment failure sequences, etc.) Precursors of UFM sequences can, however, be recognized and corrected during operation through careful analysis of abnormal operating events.

9. Summary

The only public hazard from operation of nuclear power reactors is radioactive materials. Almost all of these materials are contained in the fuel during normal operation. This is positive from the perspective of safety because the normal release to the environment during operation is zero. However, it must be recognized that the dangerous fission products are located in the same place where the fission heat is released – in the fuel. If heat production is greater than heat removal for any significant length of time the fuel will heat up, materials will reach their temperature limits, and fission products will be released from the fuel. Safety design requires introduction of reliable barriers to prevent fission product release and damage to human lives.

One key aspect of safety that must be emphasized is that humans make **all** of the mistakes - machines are much too stupid to be creative, or to take responsibility for their own failures. The human element is an integral part of any safety system, and must be examined along with the hardware in order to arrive at a correct conclusion.