

Chapter 1 - Introduction

The Double-Edged Sword

Nuclear reactors are inherently dangerous. So are hydro dams, and fossil fuel electrical generating stations. So is life, for that matter. The nature of the hazard in each case is, however, quite different. Hazards can be sudden (acute) or delayed. For hydro dams, one of the hazards (to humans) is rupture of the dam and massive floods downstream; and build-up of toxic mercury in the water behind the dam due to leaching from the rocks. For natural gas plants, there is a local hazard due to explosion, and a world hazard due to climate change from the release of combustion products (greenhouse gases) to the atmosphere. Coal plants are likewise a major source of greenhouse gases, and in addition can cause cancer from the combustion and release of chemicals in the coal. They also pose a small hazard due to release of radioactivity; depending on the source of the coal, some coal plants emit more radioactivity to atmosphere in normal operation than a nuclear power plant. For nuclear power, as discussed later, the hazard of most interest is the release of radioactivity in accidents.

We 'accept' hazards of technologies when they have a benefit which is perceived to offset the risk. Sometimes this decision is made on an individual basis: You may go sky-diving (an activity so objectively risky that you cannot get insurance coverage for it) because you believe the unique thrill is worth the risk. You accept the hazards of electrical shock for the convenience of using electric lights, etc. Few things that we do on a day-to-day basis are as risky as hurtling down a narrow strip of levelled ground at 100 km/hr in a thin metal container containing 60 litres of explosive liquid towards someone else in a similar device, using a painted strip as a guide to avoid collision. Yet we do it. Presumably, we judge that the benefit is worth the risk.

Sometimes the decision is made on a societal basis: if you live in a city, you cannot really choose to accept or reject risks such as: being hit by a car (even if you choose not to drive one); breathing polluted air; or getting mugged. Activities which pose an involuntary risk are often, but not always, regulated by law; in our three examples, the regulators would be traffic laws; emission controls on cars, industries, and fossil fuel generating stations; and the criminal laws enforced by the police.

The acceptability of risk is a social issue somewhat beyond the scope of this course, although we cover aspects of it (safety goals) in Chapter 6. However quantification of risk is a major part of the course, since it is the job of experts to provide the facts to decision-makers so that their decisions are not arbitrary or counterproductive.

The benefits of nuclear power include tangible benefits (production of electricity) and intangible ones (avoidance of greenhouse gases and carcinogenic chemical releases). (An intangible benefit is simply one that has not *yet* been costed, of course). In Ontario, about half of the electricity comes from nuclear power; in countries such as France, as much as 80%. Less attention is paid to the non-power aspects of nuclear technology, which include medical and industrial applications, insect control, environmental protection, and scientific research. In the U.S. it has been estimated that in 1991, the *non-power* use of radioactive materials (most of which are produced in nuclear reactors) were responsible for US\$257 billion in total industry sales, and 3.7 million jobs¹. Using the 10% rule, we can estimate the Canadian figures at a tenth of these. Nuclear technology is big business, because there are significant benefits in the use of nuclear technology, benefits that are sufficiently large to pursue in spite of the potential for accidents. Most of you know that Canada has until recently been the source of over half of the world's production of medical isotopes, largely originating from the NRU reactor at Chalk River; the current intent (as of this writing, in September 2009) is to repair and restart NRU.

It is not the purpose of this course to “sell” nuclear power. But a consideration only of risk without an acknowledgement of benefits does not lead to sensible decision-making.

Note that this course is mostly concerned with risk to humans. An important aspect of any technology is also its risk to other living creatures (referred to as ‘non-human biota’). For nuclear technology, it has been generally believed that if radiation risk to humans is made acceptable, the risk to non-human biota will also be acceptable, because in general they are less susceptible to radiation (e.g., they do not live as long (so do not develop cancer as easily) or are inherently more resistant to radiation damage (e.g., insects)). However radioactive elements and compounds can be concentrated up the food chain, so the pathways have to be modelled to provide a scientific basis for the claim that humans are limiting. There are also non-radiological routine environmental aspects to nuclear power generation (such as discharge of warmed water - also common to fossil plants) which are beyond the scope of this course.

Once we have decided to employ a technology, the job at hand is to minimize the risk, minimize the cost, and maximize the benefit. These objectives are usually competing and ensure that the job of the designer is challenging. It is essential to note that tradeoffs are inherent in the nature of the problem. It is not acceptable to require absolute safety at all costs. In fact, it is nonsense to require absolute safety. Nothing is absolutely safe. And we do not have infinite resources. The unrestrained pursuit of additional safety at some point incorrectly and unjustly diverts resources (time, people, money, natural resources ...) away from other important programmes (health care, education, transportation, ...). *Optimization* in the face of conflicting objectives, as opposed to *maximization or minimization* of any one, is the essence of good engineering and is not unique to nuclear power.

We need a methodology, then, to quantify risk, safety, benefit, etc., and to permit design, construction

and operation to take place on a rational and justifiable basis. Part of this course is an attempt to elucidate that methodology, a methodology employed by the nuclear industry and other industries such as the space and aircraft industries. It is called “probabilistic safety analysis”. However safety is more than just numbers: history is replete with apparently ‘incredible’ events that actually happened, in many technologies. Experience is a powerful teacher in nuclear safety, and use of experience is embedded in something called “deterministic safety requirements” - in simple terms, this means “provide some protection against event x no matter how unlikely you think it is”. Both deterministic and probabilistic techniques have their advantages and disadvantages; most nuclear designers now use both.

What is the Hazard?

The hazard most people think of when nuclear power is mentioned is ‘radiation’. This is correct, but not completely correct. Let’s go through possible hazards systematically and see how nuclear power stacks up.

A hazard in a broad sense can be *physical, chemical, biological or radiological*.

Nuclear power plants do not pose a *physical* hazard - there is no risk of offsite injury due to explosion or debris, something which is not true of other energy technologies such as natural gas or hydro power. Sometimes people confuse nuclear energy with nuclear bombs, and worry that a nuclear power plant could somehow explode like a nuclear weapon. Not so - despite the ‘explosion’ at Chernobyl, which was a steam explosion with no offsite *physical* consequences (the release of radioactivity off-site as a result of the destruction of the reactor core was another matter, of course). We will come back to this later on. For the time being here is a simple comparison.

A bomb works by making a mass of fissile material supercritical, *and holding it together long enough to reach very large energies*. The hard part is holding it together, which requires three things:

- banging two sub-critical masses together very fast, so that the supercritical mass formed does not disintegrate as the pieces approach each other and heat up; and
- ensuring that the source of neutrons that initiates the explosion is located at the centre, and is triggered at the right time, and
- using pure fissile material - U^{235} or Pu^{239} - so that the mass goes critical on *fast* neutrons. Fast neutrons have very short lifetimes. The basic time unit that bomb-designers use is a ‘shake’, or 10^{-8} seconds^a. It takes only about 50 chain-reaction generations of neutrons to produce the enormous nuclear energies in the few shakes before the mass blows apart and the chain

^a From “in a few shakes of a lamb’s tail”

reaction stops.

Most power reactors, however, slow down the neutrons to thermal energies, and thermal neutrons have lifetimes of milliseconds. This alone is not enough for a safe plant: we shall learn later on that a power plant is critical on *delayed thermal neutrons*, with lifetimes of the order of tenths of seconds to several seconds. Thus if you somehow make a power reactor (e.g., a CANDU) supercritical, the energy doubling time is of the order of hundreds of milliseconds. This is slow enough that you can stop it with mechanical or hydraulic devices; but if these fail, the thermal energy buildup destroys the fuel and the reactor geometry before the energy level gets above perhaps ten times normal power, and that ends the chain reaction. The result is not minor (cf. Chernobyl) but is not a nuclear bomb.

Most people do not think of a nuclear power plant as posing a *chemical hazard* - but thermal power plants need a large supply of cooling water. About ½ to ⅓ the energy produced by any thermal power plant is wasted, courtesy of the second law of thermodynamics; the waste energy is rejected to a lake, river, sea, or atmosphere. This water is used, for example, in once-through mode in the main condenser and in many plants (fossil as well as nuclear) is chlorinated to avoid growth of biological material in the plant equipment, such as zebra mussels. It follows that such plants have relatively large tanks of chlorine somewhere on site. The consequences of rupture of these tanks could be severe off-site (cf. the Mississauga train derailment in 1979). Because this hazard is ‘conventional’ (a word which really means ‘we are used to it’), it does not attract much attention. In nuclear plant safety design, e.g. CANDU, it *is* considered by providing a self-contained secondary control area away from the main control area, so that in case of such a release, the operators can shut down and maintain the safety of the nuclear plant without being incapacitated.

There is also no *biological hazard* associated with a nuclear plant because they do not contain or produce bacteria or viruses^b.

That brings us to the *radiological hazard*. The hazards of radiation are well-known. The effect can be *somatic* - affecting a living individual - or *genetic* - appearing in the yet-to-be conceived offspring of the person irradiated, or in later generations. (The rather dry terminology here and in the next few paragraphs is worth remembering so when others use it, it makes sense to you; and you can also spot when it is being used incorrectly).

Let's first deal with *somatic* effects. Large doses to an individual can cause illness or death (*acute*, or

^b There could be a small biological hazard if the plant uses cooling towers and doesn't keep them clean - they could become a source of bacterial growth.

prompt, or *early*, or *non-stochastic* effects - they all refer to the same concept); smaller doses can increase one's risk of contracting cancer several years later (*delayed* or *latent* or *stochastic* effects). The word *stochastic* means random, and reflects the fact that if a large number of individuals is exposed to a moderately 'high' dose of radiation (above about 0.2 Sv each - see below), one can predict the number of such individuals who will one day get cancer as a result of the exposure, but one cannot predict *which* individuals will be affected. A third classification of somatic effect is called *teratogenic*: once a woman is pregnant, the foetus could be damaged by radiation.

The second type of hazard is *genetic* - effects on children or later generations due to irradiation of the father or mother **before the children were conceived**. While such an effect has been observed in animals, it has not - despite all the cartoons - been observed in people².

In fact radiation is but one of many sources which damage the DNA in our cells; others are chemicals, and the natural error rate produced in DNA when cells divide. Nor is the challenge from radiation (from nuclear power) unique - we are born in, live in, and die in a bath of cosmic radiation. Had our cells not evolved a highly effective repair mechanism, I would not be writing these words nor would you be reading them.

Effects of Radiation

To put this in a more quantitative framework:

Radiation from a health physics point of view consists of energetic particles, which retain rather quaint names from the days before people knew what they were:

- alpha rays, or helium nuclei
- beta rays, or electrons
- gamma rays, or photons (X-rays are low-energy gamma rays)

These are characteristic of radiation emitted by the radioactive fission fragments of split uranium and plutonium nuclei (fission products).

A nuclear reactor can also be a source of neutrons, and the moving fission products themselves have energy. Neutrons are not normally a concern to the public in reactor accidents, as they slow down very rapidly in the reactor structure; however they can be a concern to workers if they are near a shutdown reactor which inadvertently goes critical, or in a fuel reprocessing criticality accident as happened in Japan.

The mechanism of damage from radiation is through deposition of energy in the cells of the body, via ionization of the molecules - hence the term ionizing radiation, to distinguish it from, for example, solar radiation (sunshine)^c. A measure of the effect of radiation is the energy per unit mass absorbed in the material through which it passes. The first such unit³ was called the “Roentgen”^d (R) and is:

The **Roentgen** is that quantity of X- or gamma-rays which deposits 87.7 ergs in one gram of air at Standard Temperature & Pressure (STP).

The difficulty with this definition is that the dose depends on the material - for example the same Roentgen produces about 97 ergs/gram in soft body tissue. Thus the “rad” was defined, applicable to any type of radiation and any material:

The **rad** is the unit of radiation dose which produces 100 ergs/gram of absorbed energy.

Note that 1R is about the same as 1 rad in body tissue.

These units do not however measure the amount of damage that different types of radiation cause: for example alpha particles are more effective in causing cell damage than beta particles, even for the same dose in rads. This effect is incorporated by specifying a Relative Biological Effectiveness (**RBE**) which compares the cell damage from all forms of ionizing radiation to that induced by gamma rays, as follows:

Radiation	RBE
X-, γ -rays, β -particles	1
Thermal neutrons	3
α -particles, fast neutrons	10
Heavy recoil nuclei (fission fragments)	20

^cThere are other ‘rays’ from nuclear reactions. Neutrinos are produced from nuclear reactions such as beta-decay but their chance of interacting with material as they pass through it is very small indeed (which is why they are so hard to detect) - so they are not significant in terms of damage to humans.

^dThe abbreviation is capitalized as the unit is named after the discoverer of X-rays.

A more relevant measure of biological effect is therefore obtained by multiplying the dose in rads by the RBE:

$$\text{Dose in rem (Roentgen-Equivalent-Man)} = \text{dose in rads} \times \text{RBE}$$

While you will often see this unit used in older texts, or by older health physicists, the current SI unit of dose is the Sievert, abbreviated Sv, defined simply as:

$$1 \text{ Sievert} = 100 \text{ rem}$$

The dose in Sv can now be related to health effects. Figure 1.1 puts this in perspective.

The boundary between stochastic and non-stochastic effects is about 1 Sv. Doses of that magnitude make you sick early on, although you will recover. Doses above 5 Sv have an increasing probability of early death, approaching 100% near 10 Sv.

For low doses, the effect is only stochastic. Since an individual dose is not a predictor of whether or not *that individual* will get cancer, the measure of hazard used is *collective dose* - i.e., the sum over a large number of people of the dose each individual receives. The unit is therefore person-Sv, although sometimes Sv alone is used if it is clear that collective dose is being discussed.

Paradoxically, radiation is not a very effective way of inducing cancer (compared to other carcinogens), and can even be used medically to cure cancer. Much of the data on how much radiation causes how many cancer cases comes from survivors of the atomic bombs which were dropped on Japan - since these provide the large numbers of people exposed to the relatively high doses required to distinguish a small effect. This data is supplemented by data from animals, on which one can do experiments on small doses over long periods of time. Even so, it is difficult to see any effect below an average dose of about 0.1 Sv.

For large doses in the stochastic range, one can predict the following effect⁴:

100 person-Sv will produce about 5 fatal cancers in the exposed (general) population

The effect would occur over a period of 10-30 years, which makes it even harder to detect against the large number of fatal cancers that occur 'normally' (your chance of dying from cancer in North America is about 25%).

Because there is little observable effect at low doses, this relationship - derived from high doses - is assumed to be linear with dose - that is, to apply whatever the dose *rate* in the stochastic regime. This so-called *linear dose-effect hypothesis* is just that - a hypothesis - but because it is believed to overestimate the effect, is used to set dose limits for workers and the public, in all nuclear endeavours - e.g., X-ray technicians, nuclear power workers. In other industries, with a lack of knowledge of the behaviour of toxic chemicals at very low concentrations, this approach is usually not used. Instead a “threshold” value of exposure or dose is postulated above which no harm is observed; and (with some safety margin) exposure to the toxin at levels below the threshold is assumed to be ‘safe’. If that

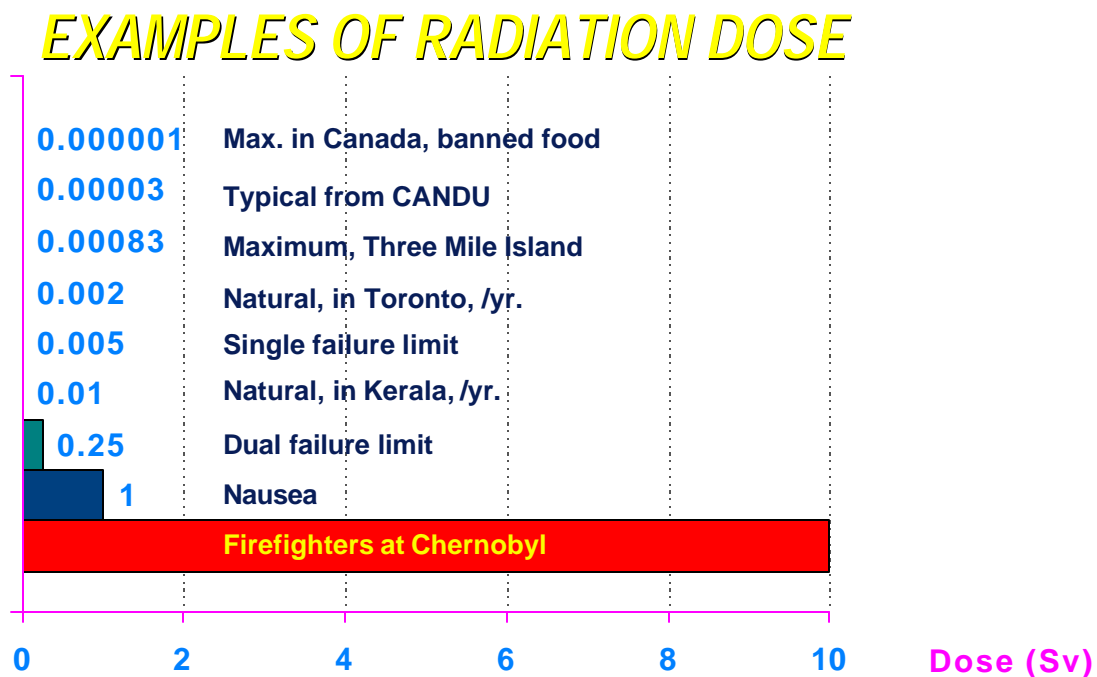


Figure 1.1

approach were followed for nuclear energy, allowed doses to the public would go up by a factor of ten or more. In fact an increasing body of minority scientific opinion now holds that doses in the range of 0.01 Sv are beneficial to you - an effect called *hormesis*.

Taken literally, the linear hypothesis has some interesting consequences. For example, we are all exposed to ‘natural’ background radiation due to cosmic rays, radioactivity in the soil and rocks, and radon gas. The dose per year varies over the globe, but is in the range of 0.001 - 0.002 Sv / year to every one of us. If you take Canada’s population of 30 million people, this results in an annual collective

dose of about 60,000 Sv. According to the linear dose-effect hypothesis, such a dose would produce 3,000 cases of fatal cancer a year. According to the hormesis theory, such a dose is beneficial.

Appendix A summarizes a little more of the science behind the response of cells to radiation dose.

Returning to Figure 1.1 above, and reading from the highest doses to the lowest: The firefighters who stood over the burning reactor at Chernobyl received doses of the order of 10 Sieverts, and most died. Doses of 1 Sv, as noted, produce nausea but one recovers. A dose of 0.25 Sv has been historically used in the US and Canada as a dose limit for a severe accident (“dual failure” in Canadian terms); such an accident has not occurred in Canada, but the reactors were designed such that if it does occur, the dose will be below that limit. For more frequent accidents (“single failures” - those that might occur rarely in the plant lifetime), the regulator in Canada has historically set a dose limit of 5 mSv, corresponding at the time to the international limit for the annual dose to a member of the public from normal operation.

These values have recently been revised as part of the development of high-level design safety requirements in Canada⁵. Single failures have been subdivided into Anticipated Operational Occurrences (AOOs), which are likely to occur at least once during plant operation, and less likely Design Basis Accidents (DBAs), with dose limits of 0.5 mSv and 20mSv respectively. (Severe accidents are now handled through a combination of specific design requirements and numerical safety goals - more on this later.)

The natural background individual radiation dose in Toronto is, as noted, about 0.002 Sv / year (medical radiation adds on average 0.001 Sv/ year); however there are places in the world with higher doses: in Kerala, India, where the rocks have a lot of thorium in them, the background dose is 0.01 Sv / year. This would seem to be a good place to study the effects of radiation dose - as has been done - but no relationship was found. In general, poverty is a much more important indicator of life expectancy than background radiation dose.

The maximum radiation dose to an individual member of the public from the Three Mile Island accident was very low (0.00083 Sv) despite the fact that much of the core melted; this was largely due to trapping fission product iodine in water inside the containment building, a phenomenon we shall cover later on. The typical annual individual dose nearby a normally-operating CANDU is 0.00003 Sv, or 1.5 % of the natural background radiation. Finally you may recall that during the Chernobyl accident, some food grown in Canada which was in the path of the fallout cloud was banned. Had you bravely eaten such food anyway, your additional dose would have been 0.000001 Sv, or about 5 hours' worth of natural background radiation.




Figure 1.2 - Measures of Radiation Including Historical Weapons Fallout


Figure 1.2 is similar to Figure 1.1 but shows the time evolution of fallout (mostly from nuclear weapons testing) compared to other sources of radiation.

How Radioactivity Can Escape

Since the course will now focus on the radiological hazard of nuclear power plants, we need to know where the radioactivity is normally, and how it can escape.

Most of the radioactivity (fission products) is of course in the *fuel* in the core. Nuclear power plants also have spent fuel on site, in either spent fuel storage pools or in dry shielded storage (concrete containers). We'll come back to fuel, but there are a few other sources of radioactivity we'll discuss now.

Classic CANDU reactors use heavy water for coolant and moderator, and this becomes activated via neutron bombardment as follows, with the deuterium atom capturing a neutron to become *tritium*:



Tritium is radioactive with a half-life of about 12 years, decaying back to deuterium and hydrogen, with emission of an electron. It is hazardous if inhaled, ingested or if it comes in contact with skin, but you need little shielding to protect yourself - the beta particle can be stopped by a sheet of plastic. If you work in an area where there is a tritiated water or vapour hazard, a plastic suit and a respirator are sufficient protection. Note that tritium oxide (T_2O) is far more hazardous than tritium gas (T_2) because of the ease with which it can be absorbed by the body, in which it behaves (chemically) like water.




Figure 1.3 - Cross Section of a Fuel Element with Sheath

The moderator is also a source of *Carbon-14*, produced by neutron bombardment of dissolved nitrogen. Normally the C^{14} is removed onto the ion-exchange columns which continually purify the moderator, and the issue is one of long-term waste disposal rather than acute exposure to C^{14} from an

accident.

Back now to the fuel. In normal operation, the radioactivity in the fuel consists of:

- fission products trapped within the ceramic UO_2
- fission product gases in bubbles or interlinked spaces within the fuel ceramic or free between the fuel and the sheath

in the ratio of about 90:10 for the highest-powered fuel element in a CANDU reactor as shown in the photo (Figure 1.3 - the cross-section of a Zircaloy-clad CANDU fuel element).

Thus accidents which damage the fuel sheath (but which don't damage the UO_2) have the potential to release something less than 10% of the gaseous fission products only. Sheaths can be damaged mechanically (fuel handling accidents), or by overheating: if the sheath overheats from its normal temperature of 300C to about 600-800C, it will plastically deform because of the pressure of the fission product gases it contains, and eventually rupture. To drive out the remaining gaseous fission products and the solid fission products such as caesium and strontium, the fuel temperature has to be raised to close to the melting point (2840C) or the fuel itself heavily oxidized by direct exposure to air or steam.

Thus accidents which release significant amounts of radioactive material are initiated by:

- overheating the fuel in the core via power/cooling mismatch
- leaks or pipe breaks in the coolant or moderator
- mechanical damage to the fuel
- overheating the spent fuel in storage via power/cooling mismatch

The use of low enriched fuel in the ACR-1000 design, and of course in all LWRs, gives a subcategory of the last item - can we get inadvertent criticality in the new fuel storage area due to e.g. flooding by light water? Or in the spent fuel bay by incorrect loading of the spent fuel racks? Or in core collapse following a severe accident? Generally for criticality accidents the design approach is to make them impossible via inherent means (geometry, absorbers), and we shall not cover criticality analysis in this course.

All accident analysis reduces to these categories of failures.

First, however, we need some tools to determine not just consequences but risk.

That's Incredible

Given a design, the basic methodology can be stated quite succinctly:

Show that the frequency and consequences of possible accidents are within acceptable limits

or

Show that the frequency of an accident is too small to consider.

Acceptable limits are defined with respect to the event frequency. For example, frequent occurrences (minor faults such as a loss of electrical power) should not stress the system, damage fuel or invoke protective systems. Very infrequent events, like a large loss of coolant, are permitted to push the physical systems into plastic deformation or damage fuel, but not allow a radioactive release beyond a prescribed limit.

Below some accident frequency, say one in a million reactor-years, in the past one normally did not provide further *explicit* design defences, but relied on the excess capability of the systems designed for more frequent events. (The term “reactor-year” means: if you have n_i reactors each running for i years, the cumulative number of reactor-years of experience is:

$$N = \sum_i i \times n_i$$

This implies that the more reactors that are built, the safer they have to be.) In recent years there has been increased focus on rare severe accidents and for new designs one provides “complementary features” such as core flooding systems, containment atmosphere cooling and a robust containment structure that can collectively cool debris and contain the releases until accident management measures stabilize the situation.

Anything above an “incredible” frequency typically gives rise to varying degrees of concern as shown in Table 1.1.

So, safety, or its negative counterpart, risk, is a function of the frequency of occurrence of an event and the consequence of that event.

Table 1.1 - Acceptability of Risk

Annual individual fatality risk level from an accident, per year	Conclusion
10 ⁻³	This level is unacceptable to everyone. Accidents providing hazard at this level are difficult to find. When risk approaches this level, immediate action is taken to reduce the hazard.
10 ⁻⁴	People are willing to spend public money to control a hazard (traffic signs/control and fire departments). Safety slogans popularized for accidents in this category show an element of fear, i.e., “the life you save may be your own”.
10 ⁻⁵	People still recognize these as of concern. People warn children about these hazards (drowning, firearms, poisoning). People accept inconvenience to avoid them, such as avoiding air travel. Safety slogans have a precautionary ring: “never swim alone”, “never point a gun”, “never leave medicine within a child’s reach”.
10 ⁻⁶	Not of great concern to the average person. People are aware of these accidents but feel that they can’t happen to them. Phrases associated with these hazards have an element of resignation: “lightning never strikes twice”, “an act of G-d”.
	Extracted from H. L. Otway and R. C. Erdmann ⁶

Risk

Safety concerns are ultimately expressed in terms of *risk*. The risk of a system, which must be specified (e.g., of a component failure, of an activity, of a nuclear reactor, of the nuclear fuel cycle, etc.) is customarily defined as:

$$\mathbf{Risk} = \sum_i \mathbf{expected\ frequency\ of\ event}_i \times \mathbf{expected\ consequence}_i \quad (1)$$

Risk is a summation of events over the chosen system, and will increase when either the number of events or the magnitude of the events increase. This is by no means a unique definition; for instance, if one wanted to amplify the importance of events with large consequences (risk aversion), risk could be defined as:

$$\text{Risk} = \sum_i \text{expected frequency of event}_i \times (\text{expected consequence}_i)^k \quad (2)$$

where $k > 1$

We seek to optimize risk, not to minimize risk. We could start by choosing the least risky path to achieve the desired goal. But lowering risk is usually expensive and, since we have finite resources, we need to balance the cost versus the benefit. This can be done by setting quantitative risk targets^e. The




Figure 1.4 Cost Versus Risk

target levels of acceptable risk are set with respect to the alternative ways of achieving the same goals.

^eRisk can be quantitatively optimized, usually within a limited set of alternatives, by using a technique called Benefit Cost Analysis (BCA).

For instance, acceptable levels of risk for nuclear power plants should ideally be set at levels comparable to the level of risk inherent in coal and oil fired plants if we allow all three to be used for power generation. Because of many factors (newness of technology^f, fear of radiation), we find that the acceptable level of risk for nuclear power has been set substantially below that of most alternative means of large scale power production (natural gas is safer in terms of the overall fuel cycle because the safety of nuclear generation is offset by the conventional hazards of uranium mining). This has ensured that nuclear power is safer than most alternatives (and indeed safer than most human activities), but this safety has come at a significant social cost. One can argue that the funds spent on the extra safety should have been spent elsewhere.

Figure 1.4 illustrates that dealing with risk (i.e. providing safety) becomes more and more expensive as the risks become smaller - a form of diminishing returns on our efforts to make the world a safer place to live. Conversely, the social cost increases as the risk level increases. We seek to minimize the total cost (assuming that the true cost can be properly quantified). Starting from the right side of Figure 1.4, the high social cost of very risky things and the relatively low cost of implementing safer systems leads society to invest wisely in these safer systems (example: car seat belts). As we progress to consider endeavours of lower and lower risk, the increasing cost of implementation of safer systems begins to outweigh the benefits derived from the safer systems. At some point, we have to say “enough”. But how do we know when enough is really enough?

Three Approaches to Design

Quantification of “enough” implies quantifying the consequences and quantifying the frequencies of possible events. In short, we need to analyse the safety aspects of the endeavour in question. There has always been a recognition of the role of probability and consequence in determining the risk of a design even if it was not explicitly stated. Thus for boiler design in the early 1900s, because our analysis capability was limited and because failure data was not readily available, risk could either be *accepted* (boilers were expected to explode occasionally) or *reduced* by over-design. The latter approach reduces risk but increases cost. Furthermore the increased costs are not all easy to identify. A 10-ton automobile might offer increased safety and the increased cost of manufacture may be well-defined, but how does one cost the increased effects on the environment (due to increased fuel consumption), and to occupants of lighter vehicles in a collision? Further, where analysis capability was limited, improvements often occurred more as a result of “leaning by mistakes” than as a result of pre-

^fThe fact that nuclear technology is relatively new compared to, say, coal mining means that people are not “used” to a poor level of safety, as in the latter; and also that the technology is not as entrenched.

production design and analysis. This may be acceptable for products that can be exhaustively tested to failure (like automobiles) but it is not acceptable for the nuclear industry or similar industries where it is usually not financially possible nor socially acceptable to test complete systems to failure in anger^g.

Consequently, prudent engineering required a more deterministic approach: i.e., ensure protection against prescribed events. It is only recently that failure rate data has become more available, enabling safety optimization through the probabilistic approach.

In summary, there are three overlapping approaches:

- 1) **design by probabilistic safety analysis** - i.e., design according to the predicted frequency and consequences of failures, optimizing to deal with the high-risk contributors.
- 2) **design by deterministic safety analysis** - i.e., design according to a prescribed list of failures based on past experience and judgement. Sometimes these are called ‘design basis accidents’.
- 3) **design by rule** - e.g., use the ASME code for pressure vessel design. It is implied that following the Code or the Standard reduces the likelihood of failure of the material to a very low level. This is largely based on long experience and test and, more recently, analysis. Special versions of these Standards are used for pipes and vessels in the nuclear industry. In some cases the rule is considered to be ‘conservative’ enough that one does not have to consider failure in the design if the rule is followed, pressure vessels being a case in point. This is a disadvantage in some ways, as it gets the designer ‘off the hook’ as long as he follows the rules.

We expand on these three methods..

Design by Probabilistic Safety Analysis

The probabilistic approach provides a rational framework and it is useful to cast our study of safety design in those terms first. Probabilistic Safety Analysis (PSA) seeks to categorize each event by

^gThis is somewhat oversimplified - much of the information on how a reactor behaves under a large power rise, and what happens to the fuel, comes from the SPERT/BORAX series of destructive tests on real reactors in the U.S. Public safety was assured by doing them in remote desert areas.

probability of occurrence and then demonstrate that certain criteria are met.

PSAs therefore proceed using the following methodology:

- define the acceptance criteria,
- generate a set of accidents to consider,
- predict the frequency and consequences of the event,
- show that the appropriate risk-based criteria are met.

Acceptance Criteria

Each event or collection of events is associated with criteria against which they are to be judged. The nuclear industry uses two general types of acceptance criteria for PSAs: *Binning* and *Averaging*:

- *Binning* techniques are based on limiting the consequences for each event based on its individual frequency. An example (sort of) is the Canadian Nuclear Regulatory Commission (CNSC) Consultative document C-6 discussed later on in this course.
- *Averaging* techniques are based on setting a limit on the frequency of a given outcome, which we will call a “safety goal”: for example, that the expected frequency of the release of X TBq^h of radioactivity be less than 10^{-6} events/year; or that the core damage frequency be less than 10^{-5} events/year. This methodology also requires the summation of the frequency of all events that exceed the stated criteria.

To use either criteria we need the PSA methodology described in this course.

Note that the consequences are usually expressed in terms of radioactive releases (since these can be directly related to health effects, which is what you want to limit). They may be worked backwards to define subsidiary acceptance criteria for design and safety analysis as discussed in detail in Chapter 7. The subsidiary criteria can be *probabilistic* (the likelihood of consequence x occurring at frequency y must be less than z) or *deterministic* (the consequence of accident x has an upper limit y). Usually the criteria for probabilistic safety analysis are probabilistic, but sometimes deterministic criteria are used for simplification. For example instead of calculating the frequency and consequences of core-wide fuel damage in a small LOCA with ECC, one could simply say that fuel damage must not occur at all in a

^h 1 Becquerel (Bq) of radioactivity is that quantity which has a disintegration rate of 1 nucleus per second. It's a very small quantity. A litre of milk contains about 40 Bq of naturally radioactive K⁴⁰. In terms of old units (Curies), 1 Bq = 2.7×10^{-11} Ci. 1 T(era)Bq = 10^{12} Bq.

small LOCA with ECC, regardless of consequences.

Accident set

The task here is to define all the *initiating* events that are deemed necessary to analyze (predict the consequences and the frequency thereof). The discussion in Chapter 4 summarizes methods used to ensure that all event initiators have been captured, and how event sequences are built up from event initiators. Suffice to say at this point that there is no way of *proving* that all events have been captured properly using probabilistic methods. That is one argument in favour of using deterministic methods in design, in a complementary fashion, particularly for novel technologies.

Predict Frequency & Consequences

Since events are classified by the frequency of occurrence, the numerical reliability of systems has to be measured or analyzed. The frequency of an accident is built up from the frequency of the initiating event (e.g., pipe break in the primary coolant system, or failure in the reactivity control system), and the reliability of each of the safety-related mitigating systems called upon after the accident to stop it or contain its consequences (e.g., other control subsystems, makeup systems, shutdown systems, emergency core cooling system). Fault trees (FT) are the tool used to determine the reliability or the failure rate of a system; event trees (ET) are the tool used to link the initiating event frequency with the reliability of the mitigating systems.

Compare to Criteria

The result (in terms of frequency and consequence) is then compared to the acceptance criteria, and the design is changed if they are not met. Usually there is a low-frequency cutoff below which further mitigation is not considered justified (in shorthand, the event is ‘incredible’ - which the author believes is an imprecise term which should not be used). Often it is not so simple, and benefit-cost analyses are used to see if a design change really would reduce risk by a significant amount.

Design by Deterministic Safety Analysis

Historically the approach to accidents did *not* use PSA, for two reasons: the PSA tools were not well developed, and there was not enough experience to confidently support the frequencies and reliabilities which PSA needs. *Common cause failures* were a particular concern. For example if one had three separate emergency heat removal systems, each with a failure probability of one in 100 demands, then one could be tempted to deduce that the probability that *all* three systems failed simultaneously on demand would be one in a million. But if the systems are all maintained by the same crew, or use

equipment from the same manufacturer, or are subject to a common environment after an accident, or all rely on a single source of cooling water or electrical power - then the combined failure probability is much higher. It was not until PSA tools were developed to quantify these common-cause failures that the PSA methodology became more widely accepted.

Thus at first a different methodology was used. In Deterministic Safety Analysis, a set of stylized accidents - called Design Basis Accidents - is defined based on past experience, knowledge of the plant, and engineering judgement. Each accident sequence is chosen to be severe enough that the consequences of a 'real' accident would be less; thus only a small subset of possible accidents need to be analyzed. Sometimes unphysical assumptions are used; sometimes variables are set at the most pessimistic limit. The consequences of these stylized accidents are predicted and compared against acceptance criteria. Such acceptance criteria are very loosely based on frequency. For example, in Canada, two broad classes of accidents were historically defined, along with dose limits for each class (single failures, and dual failures); within each class, however, the 'real' frequency of an accident could vary by three orders of magnitude. In the U.S., a severe release of fission products *into* containment was *prescribed* as a basis for the containment design, regardless of the actual reactor design within.

Design Basis Accidents, or DBAs, are discussed in detail in Chapter 2. The use of DBAs is double-edged, being both inclusive and exclusive: they define *a priori* a list of accidents against which the designer must provide a defence; and accidents beyond this set were considered to be sufficiently rare that no specific design provisions need to be made.

“Design Basis Accident” is however a poor term and, by itself, a weak concept. Accidents can be subtle or may evolve in unusual ways, and the creator of the Design Basis Accident list can too easily dismiss them; ‘rare’ accidents can indeed occur; and restricting one’s defences to a pre-ordained list can lead to a lack of robustness in the safety design, and a lack of questioning attitude on the part of the designer. Conversely some Design Basis Accidents are indeed very rare (sudden large LOCA) and much money has been wasted in performing and justifying sophisticated analysis of them, and providing equipment to mitigate them.

It is telling that the Three Mile Island accident was not in the Light Water Reactor Design Basis Accident set. In response to that accident, LWRs began to investigate severe accidents (sometimes called, unimaginatively, “Beyond Design Basis Accidents”) to ensure the plant retained some residual defences, notably that the containment would not be damaged early. To Canada’s credit, some severe accidents were always part of the Design Basis for CANDU. The current requirements in Canada and elsewhere, as noted, place much more emphasis on handling severe accidents.

Design By Rule

Finally design by rule is still used for cases where the frequency is very low and/or indeterminate, or where the consequences of failure are unacceptable - notably in the design of large pressure vessels. Handbooks (Standards) list the rules that must be followed.

The interplay between probabilistic and deterministic analysis is illustrated in Figure 1.5.

Safety Analysis

For each accident, whether from each branch of the event tree that is “credible”, i.e., has a frequency higher than a predefined cutoff, or from a deterministic list, safety analysis must be performed, usually by computation and experimentation, to determine if the consequences are within acceptable limits or not. Safety analyses are very complex and require extensive knowledge of an event. Since very few significant accidents have occurred, there is strong reliance on computer codes with the associated challenge of showing their predictions are correct (i.e. the codes are “validated”). The details of these analyses are largely beyond the scope of this course. However Chapters 7 and 8 cover the elements of safety analysis (also called ‘accident analysis’ depending on whether you view the glass as half full or half empty). If the limits are not exceeded no further action is required. If they are, something has to be done to mitigate the issue. That something is (re)-design, refinement of the safety analysis, or demonstration that design changes do not materially alter the risk.

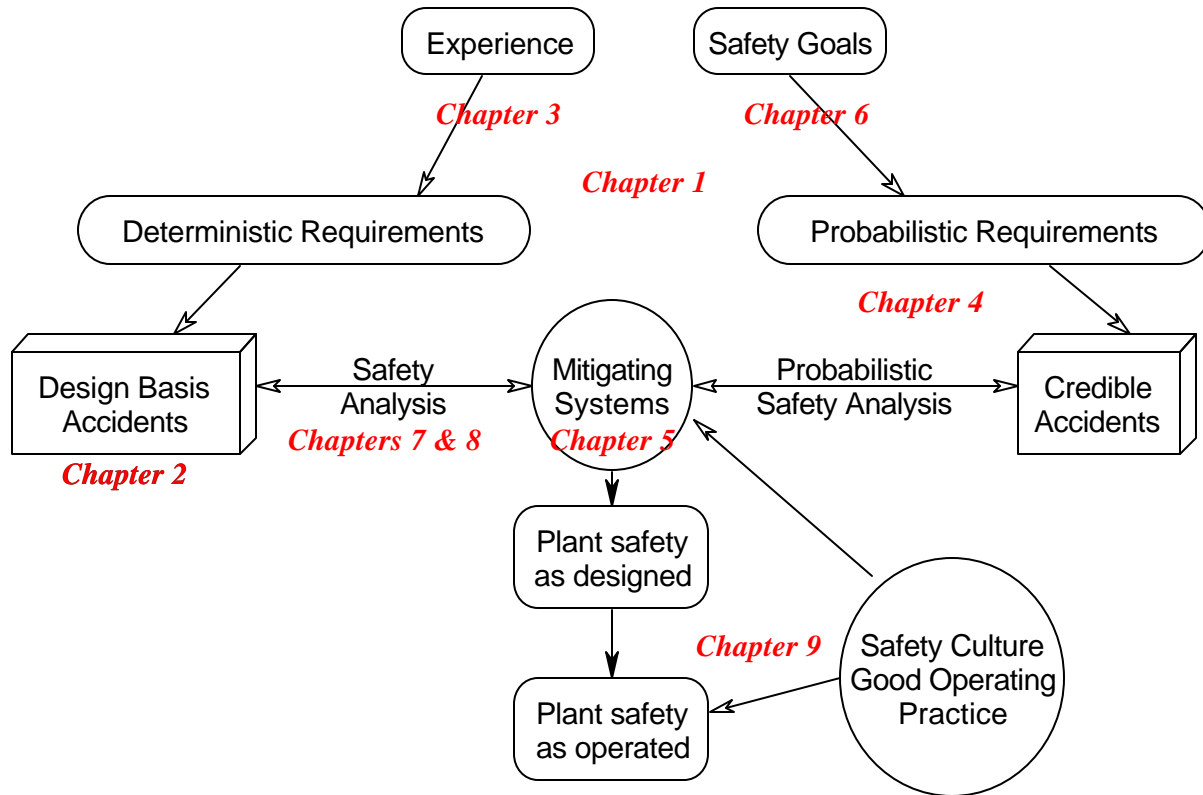


Figure 1.5 Probabilistic and Deterministic Analyses

How is Safety Design *Really* Done?

Most jurisdictions nowadays use *both* Deterministic and Probabilistic Safety Analysis for designing and licensing nuclear reactors, for reasons which should be obvious from the above discussion. But the real safety design process is more complex than setting criteria and meeting them. As one analyzes a given design, weaknesses and areas for improvement show up. We will find that reactors with negative void coefficients of reactivity are not necessarily safer than those with positive coefficients. We will likely find that most equipment faults of consequence are caused by secondary and supportive systems, not the reactor and reactor heat transport system proper. We will find that most accidents are caused, and often cured, by human error, not machine error. We might find that all designs, even passively safe ones, have failure modes (like loss of reactor power control) that are not passively safe. However, we won't find anything unless we look and we can't judge what we find unless we are able to quantify our findings.

The subject of “safety design” is a combination of safety system design and safety analysis. Design is the process by which a system is engineered to perform its intended function. Ideally, we would like to be able to work forwards from the design criteria to define the actual design: that is, from a performance specification to a system specification to a component specification (geometry, materials and operating parameters). Certainly design requirements are written but they are not sufficient to determine a design. One learns mathematics by learning theorems and finding that, lo and behold, they have useful consequences. But in reality, the theorems came from generalizing experience and examples, not the other way around. So in design, we use past experience and accepted practices to conceive of an initial design and proceed to analyze that design to see if it meets the performance specifications. Obviously this is an iterative process.

In the nuclear industry, practical design relies heavily on previous designs. New designs tend to be evolutionary rather than revolutionary for at least two reasons: cost, and performance assurance. It has been estimated that the overall cost of taking a reactor concept from paper to a commissioned prototype power reactor is about \$1-2 billion, of which the design cost alone is now about \$400 million. This alone biases the design process to lean heavily on past designs. But apart from the cost, overall operating and safety performance is a strong function of accumulated operating experience and laboratory testing. Utilities, who after all buy the design, tend to be conservative; in fact the only thing that makes a utility more nervous than buying a new design, is being the *only* utility to buy a new design.

As a part of the design, safety principles are declared and must be shown to be met. They likewise did not come from academic investigation, but from real accidents and some of the hard lessons learned in the early days of nuclear reactors. Because accidents are relatively rare, they can be enormously instructive when they occur: the Three Mile Island accident caused a large shift not just in the way LWR designers approached safety design, but also forced the U.S. nuclear regulator (the U.S. Nuclear Regulatory Commission, or USNRC) to reassess how it operated. By the same token, the fundamental shutdown system design philosophy of CANDU came from an accident in the NRX research reactor in 1952.

Figure 1.6 is an overview of the design process from a very generic stance. Can you see where the PSA and the deterministic assessment fit in?

Figure 1.6 is but one way to view the whole process. We'll see other views as well, such as that of the International Atomic Energy Agency (IAEA) in Chapter 9 and the CNSC in Chapter 2. We shall see that the views are complementary. All views revolve around the common sense approach that is inherent in good engineering practice: start with a good design, follow established safety and design practices, and provide protection against the risks.

Key CANDU system designs result from this type of process, and are discussed in Chapter 5.

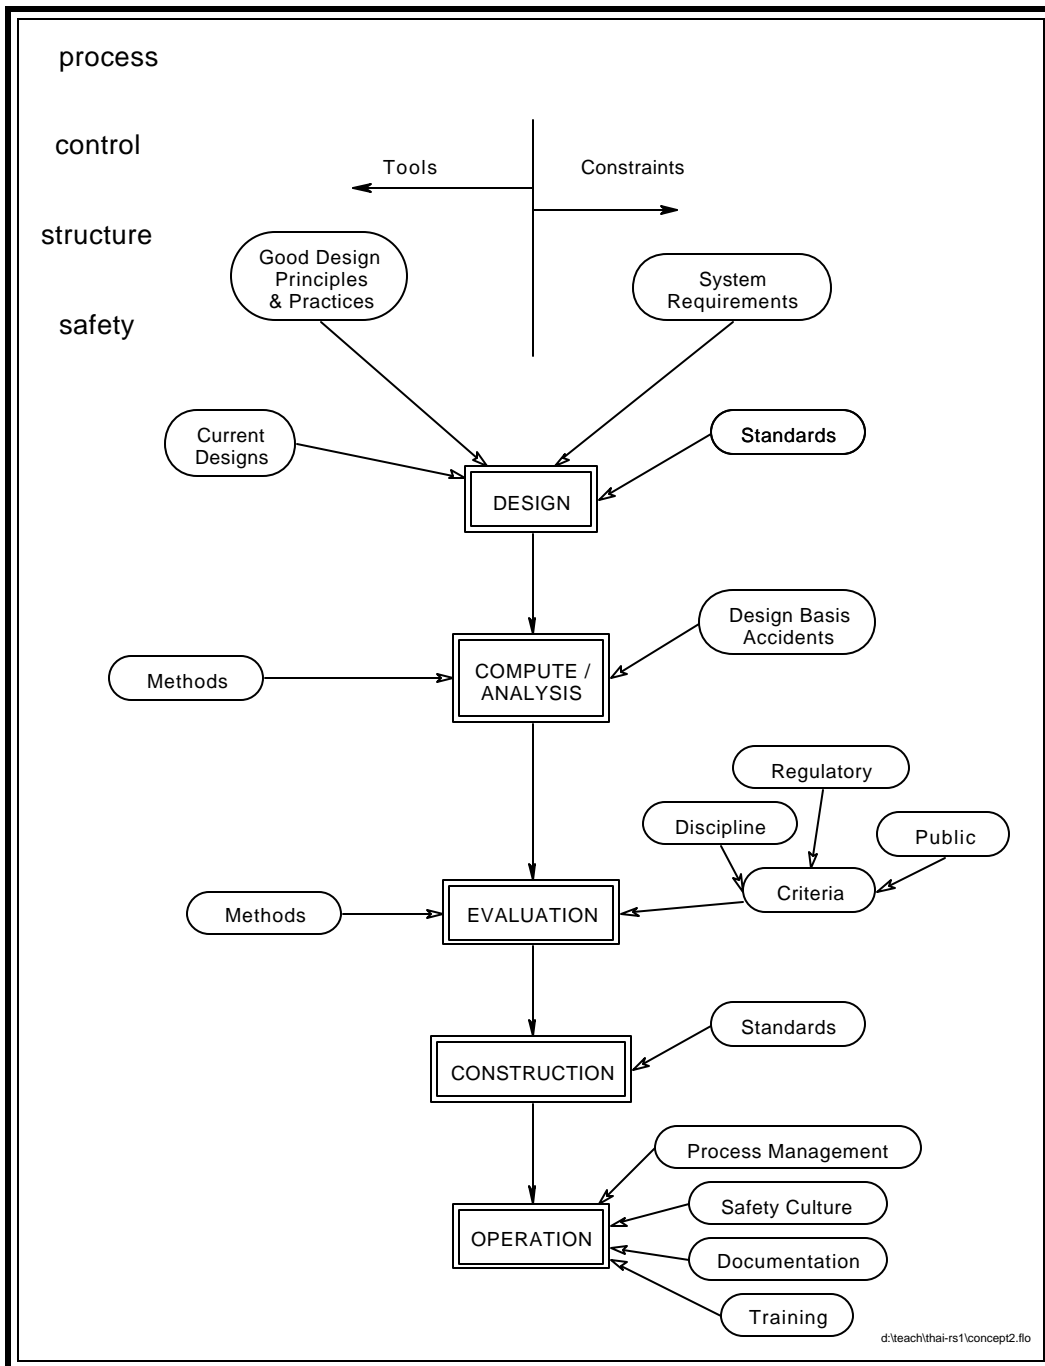
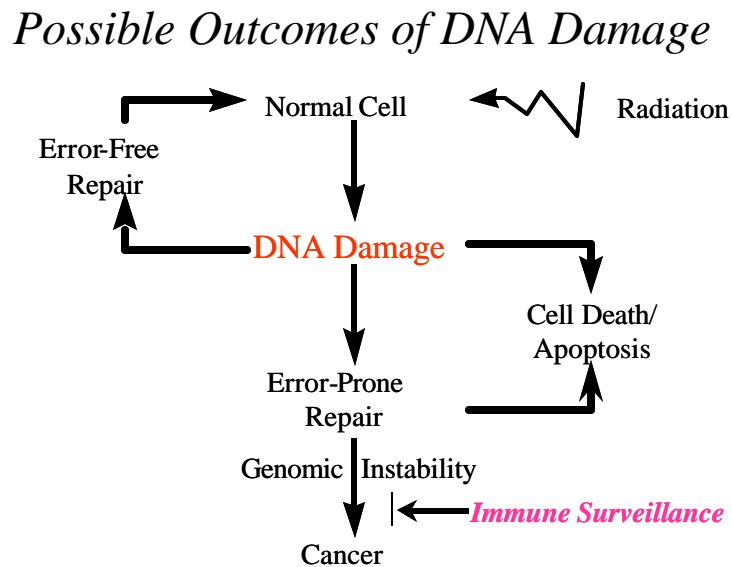


Figure 1.6 Overview of the design process.

Appendix A - Evolving Views on the Linear Dose-Effect Hypothesis

This summary is taken from a lecture in 2007 at AECL by Dr. Shu-Zheng Liu, Professor of Radiation Biology, Jilin University, a medical doctor and cancer expert with 367 scientific papers to his credit.

Figure 1.7 below shows the possible outcomes of DNA damage to a cell:



Ron Mitchel, 2004, modified 2007


Figure 1.7 Outcomes of DNA cell damage

Significant DNA damage from radiation (e.g. double-strand break of the DNA) normally results in error-free repair or the death of the cell. An example of how cells repair themselves with time after exposure is shown in Figure 1.8 below.

If the repair is error-prone, the body's immune system can destroy the damaged cell (labelled "Immune Surveillance" in Figure 1.7). Only if all these mechanisms fail can a cancerous cell develop.

Dr. Liu's work has shown that low-doses of radiation can actually *stimulate* these overall defence mechanisms.

Repair of γ H2AX foci as a function of time and dose



Rothkamm and Löbrich 2003

Figure 1.8 Cell repair as a function of time and radiation dose

Given the basic problem that the effects of low doses are small, regulators initially tended towards using the linear dose/effect hypothesis. This has been increasingly challenged (as being too conservative) in recent years, although it remains the basis for regulation in most countries.. The table below summarizes the history.

Table 1.2 - Debate on Dose/Effect Relationship

Year	Expert group Recommendations
1972	BEIR I report recommended using a linear model for estimating radiation risks
	UNSCEAR VI questioned the validity of using a linear model for estimating radiation risks
2001	NCRP report concluded that the LNT Model is the best model for risk assessment though there were many sets of scientific data not supporting it.
	The ANS published a position statement emphasizing that below 100 mSv risks of health effects are either too small to be observed or are non-existent.
2005	BEIR VII report insisted on using the LNT model for estimation of risk for low and very low doses though it recognized the uncertainty of such judgement
	Joint report of French Academy of Science and Academy of Medicine stated that cancer is a multi-cellular disease arguing against BEIR VII report
2007	ICRP new recommendations: LNT hypothesis still remains a prudent basis for radiation protection at low doses and low dose rates
2007	DOE LDR Program Update: signalling from non-irradiated cells can actually eliminate damaged cells from a tissue instead of assuming that the single cell is the unit of function.

Exercises

1. A few years ago the Washington sniper(s) terrorized Washington and the surrounding areas for several weeks. Suppose you lived in Washington and had the option of ‘waiting it out’ in Toronto. Calculate and compare the risk to you due to the sniper if you stayed in Washington; and the risk of flying to Toronto and back and avoiding the sniper.[Hint: You’ll need to look up some airline accident statistics]. Discuss the cost incurred of moving to Toronto temporarily, versus any risk reduction. How would you judge the acceptability of this cost/risk tradeoff (i.e. what numerical benchmark would you use)?
2. A fault tree identifies all the failure modes of a piece of equipment and assigns a numerical frequency or demand availability to each one. Do the first part: List the failure modes of an active safety system (a household circuit breaker) and also those of a similar passive one (a household fuse). Make sure you define the term “failure”, recognizing that this is a safety course.
3. Every computer user is always told to back up data. Assume you have just finished your Master’s thesis on your computer and you need a reliability of 999 times out of 1000 that your thesis is readable. You have a computer with a floppy disk drive with a reliability of 0.95 per floppy disk written, and a CD re-writer with a reliability of 0.9 per CD-RW written (they aren’t very good, as you may have noticed). What strategies would you use to get the required reliability to ensure you could recover your data? (Assume your data will all fit on one floppy disk and that you have as many floppy disks and CD-RWs as you need). Discuss the strengths and weaknesses of your chosen approach. Even if your numbers pan out, what could be wrong with your assumptions?
4. Assume a collective dose of 100 person-Sv is given to:
 - a. 1,000,000 people
 - b. 1000 people
 - c. 10 people
 What would be the expected number of cancer cases in each situation? (Why?)
5. Rank the magnitude of the following risks to you as an individual (express the answers numerically and explain your reasoning):
 - a. A one-time dose of 10 Sv

- b. A dose of .33 Sv / year for 30 years
 c. As (b) but the dose is due entirely to heavy nuclei.
 d. A one-time dose of 5 Sv
6. International bodies set limits for the amount of dose an individual should receive from all man-made sources. Nuclear power plants are required to meet these limits on public dose in normal operation (in practice they fall well below). There are a number of issues lying behind this apparently simple statement. Discuss the following four and draw reasoned conclusions:
- a. How should exposure from radiation used for medical purposes be controlled (i.e., what factors should determine whether or not, and how much, radiation should be used)?
- b. Should large power reactors have the same limits as small research reactors such as at McMaster (which also produces medical isotopes)? Why?
- c. You are a nuclear regulator and have been asked to approve two devices: a smoke detector, and an X-ray machine for looking at your feet in a shoe store to make sure your shoes fit (pretend this is in the 1950s, when many shoe stores had these!). Assume (for the sake of this problem) that the smoke detectors will give a dose of 0.01 mSv per year to the whole body of 20,000,000 people in Canada; and that the X-ray machine would give a dose of 1 mSv per year to the feet of 200,000 people. What would your decisions be, and why? (What factors would you look at?)
- d. What should the dose limit be for lifesaving (i.e., your colleague is trapped in a very high radiation field and you are asked to go in and save him)?
7. Many people refused to fly after the attack on the World Trade Centre in September 2001, because of their belief that the risk of death due to flying had increased. Clearly for those who are personally affected by the crash, the impact is disastrous and tragic. However how does it affect a decision to fly in future? Estimate (numerically) the change in risk of death per year to an individual who flies 10 times a year, assuming that four *extra* planes crash each year. How does it compare to hi/her risk of death from other causes? [Hint: You will have to look up flight statistics]
8. Rank the magnitude of the following risks to a group of people (express the answers numerically and explain your reasoning):
- a. A collective dose of 1000 person-Sv given to 1,000,000 people
 b. A collective dose of 1000 person-Sv given to 100,000 people
 c. A collective dose of 1000 person-Sv given to 100 people
 d. A collective dose of 100 person-Sv given to the entire world

9. The attached figure shows a passive water makeup system for a reactor. The water is in a tank pressurized by gas, at a pressure lower than the operating pressure in the reactor. A one-way rupture disk separates the two. It will break only when the pressure in the reactor is 1MPa less than the pressure in the tank. When the pressure in the reactor falls, say due to a loss of coolant, the rupture disk will break, and water will flow into the reactor. No instrumentation or control is needed and no electricity, so this would be classified as a passive system. What are its failure modes?

10. If you had to take one of the following two risks, which risk would you prefer, and why?
 a.1 chance in 1000 of losing \$1
 b.1 chance in 1,000,000 of losing \$1000?

11. If you had to take one of the following two risks, which risk would you prefer and why?
 a.1 chance in 1000 of losing \$1000 or
 b.1 chance in 1,000,000 of losing \$1,000,000?

12. If you had to take one of the following two benefits, which benefit would you prefer and why?
 a.1 chance in 1000 of receiving \$1 or
 b.1 chance in 1,000,000 of receiving \$1,000?

13. If you had to take one of the following two benefits, which benefit would you prefer and why?
 a.1 chance in 1000 of receiving \$1,000 or
 b.1 chance in 1,000,000 of receiving \$1,000,000?

14. Where do your choices fall on the risk plot of Figure 1.8, below? Are you averse to risk with large consequences?

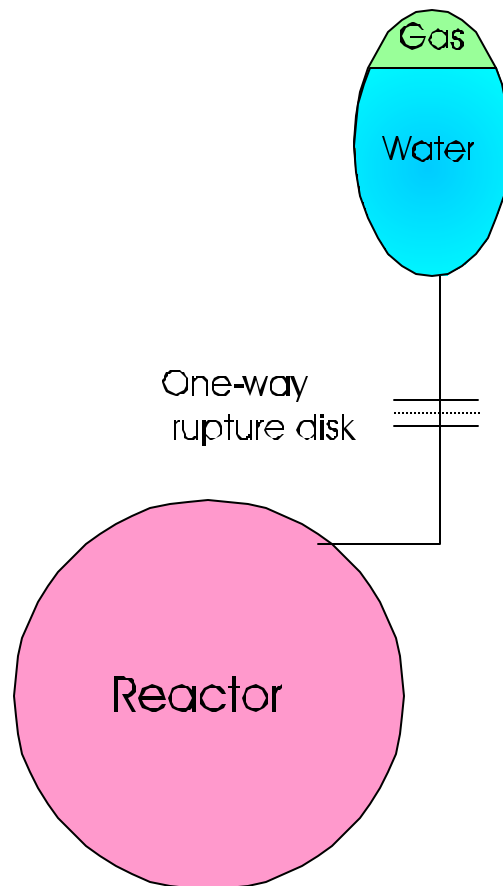


Figure 1.9 - Passive Makeup System

15. A nuclear regulator is considering a high-level safety goal for new nuclear power plants in Canada. He proposes two requirements:
- The risk to an individual close to the nuclear power plant of dying immediately from an accident must be less than 10^{-6} per year
 - The risk to an individual close to the nuclear power plant of getting cancer from an accident must be less than 10^{-5} per year.
- Two nuclear power plants apply for a licence. They have done an accident analysis and the results are as follows:
- For plant 1, there are no significant releases for any accident above a frequency of 10^{-7} per year. However there is an uncontained core melt at that frequency which gives a dose of 10 Sv to each individual in the nearby population.
- For plant 2, two accidents are the major contributors to risk. One causes severe fuel damage but prevents core melt. It occurs at a frequency of 10^{-4} per year and gives a dose of 0.25Sv to each individual in the nearby population. The other is a core melt but it is contained - it occurs at a frequency of 10^{-6} per year and gives a dose of 1 Sv to each individual in the nearby population.
- Determine numerically whether these plants meet either, both, or neither safety goal.
Hint: consider converting average dose to risk.
16. A nuclear designer is trying to optimize his design. He knows of an accident with a frequency of 10^{-7} per year which leads to a contained core melt and causes the following effects:
- Permanent damage to the plant (i.e. cannot be recovered)
 - Evacuation of nearby people (5,000) for three days
 - No prompt fatalities
 - A collective dose to the closest population of 100 Sv
- He can reduce the frequency (but not the consequences) of this accident by a factor of 10, by putting in an extra heat removal system, costing M\$10 in capital costs and an extra \$100,000 per year in maintenance and operating costs. How would you make this decision in an quantitative way?
Hint: Consider expressing accident consequences in terms of dollars.
17. A massive *spontaneous* failure (i.e. not as a result of a core melt) in an LWR pressure vessel would simultaneously breach all the physical barriers which prevent radioactivity from escaping - the fuel, the primary coolant pressure boundary, and the containment. Research and describe the approach taken by LWR designers to show that this is “incredible”.




Figure 1.10 - Example of a Constant Risk Line

References

1. “The Untold Story: Economic and Employment Benefits of the Use of Radioactive Materials”, report prepared for *Organizations United for Responsible Low-Level Radioactive Waste Solutions*, by Management Information Services Inc.; March 1994.
2. A.B. Reynolds, “Bluebells and Nuclear Energy”, Cogito Books, 1996.
3. D.J. Bennet, “The Elements of Nuclear Power”, Longman Group Limited, 1972.
4. International Commission on Radiological Protection, “1990 Recommendations of the ICRP”, report ICRP-60, table B.11.
5. Canadian Nuclear Safety Commission, “Design of New Nuclear Power Plants”, RD-337, November 2008.
6. H. L. Otway and R. C. Erdmann, *Nucl. Eng. Design* **13**, 365 (1970).