Chapter 3 - Case Studies[1]

Introduction

The purpose of this chapter is to examine some real-life accidents as an introduction to the methodology that will take up the next few chapters. As you may have gathered, it is all too easy to present safety analysis as a formalism, with well-defined rules and "right" answers. Real accidents almost never follow the simple assumptions in safety analysis. Most real accidents tend to be very complex, have more than one contributing factor, have a high component of human error and human recovery, and often leave one wondering "How did all these improbable things happen all at once?"

The simple answer is, of course, that if they didn't happen all at once, there would be no accident. It is a poorly-designed nuclear power plant where a single failure causes a catastrophic accident. Most accidents are the end-point of a chain of events, prevention of any one of which would have stopped the accident from happening. The fact that such chains occurred provides powerful lessons learned for future designs and current operation, lessons which are ignored at our peril.

That is not to denigrate the purpose of safety analysis. Traditionally safety analysis is structured so that its answers provide an upper bound to the accident consequences in each class of accidents - "things can not be worse than this". So if you base your design on the safety analysis, it will be quite robust and able to withstand a real accident with some margin - at least as far as public safety is concerned. Safety analysis is not generally concerned with economics, so that a plant owner can take little comfort *alone* from the fact that the public is protected, if his plant is damaged beyond repair in the process. Economics is actually a more powerful influence on design than safety, at least for the more likely accidents - simply because of the huge cost of plant damage. From an economic point of view, it is important to know how the 'real' plant is likely to behave - something traditional safety analysis with its pessimistic assumptions distorts badly. For the same reason, if you are writing abnormal or emergency operating procedures to help an operator manage a real accident, traditional safety analysis is not a good basis, as it tends to compress the time-scale, ignore the (possibly beneficial) effects of non-safety systems, over- *and* underestimate human intelligence and ingenuity, and exaggerate the consequences. A more realistic picture of the plant is obtained from PSA (which gives both a best estimate of the frequency of event chains and normally uses "best estimate" safety analysis.

In general, when any new technology is introduced, its characteristics are not fully known and there is an early peak in the accident rate. As people learn from design and operating mistakes, the accident rate declines, as shown in the sketch below. The decline does not, however, account for organizational forgetting, and if nothing is done to preserve corporate or industry memory, the rate will start to rise again. One of the purposes of case studies is to ensure that past mistakes are not repeated.

What you are expected to get from this chapter then is the ability to dissect an accident - what was its nature? What were the root causes? What were the lessons learned on its design? On operation? The exercise at the end of the chapter gives you examples of an actual incident and asks you to figure out what *really* went wrong.

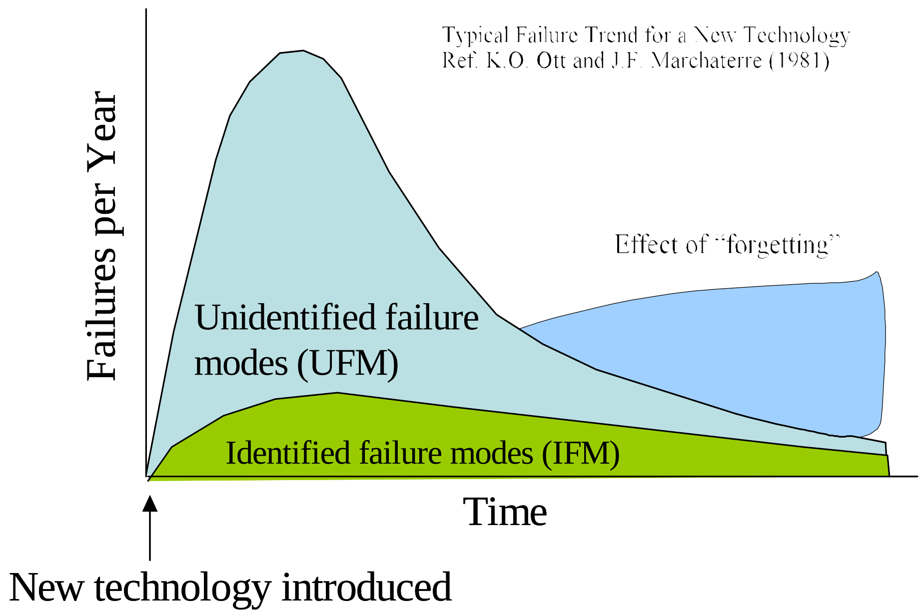

Figure 3-1 - Technology Learning Curve

Reactor Physics[2,3]

Four Factor Formula

It is not possible in the course of a few lectures to cover all reactor technology. However since some understanding of reactor physics is necessary to make sense of some of the case studies in this chapter, and also to avoid a chapter which is all words, we make a brief detour to point kinetics. This is by no means a course in reactor physics and the knowledgeable reader is asked not to cringe at some of the simplifications made to achieve understanding.

Point kinetics is a model of reactor dynamics - that is, how a reactor state can change in time. The point kinetics model is a simplification which treats the reactor, as implied, as a point - in other words, all parts of the reactor are assumed to behave the same way in a transient. (This is not a very good approximation for CANDU but it illustrates the basic concepts). First we look at some simple kinetics of a finite reactor.

Consider a chain reaction with no external sources of neutrons. Define a reproduction constant, k, as the ratio of the number of neutrons in one generation to the number in the preceding generation:

$$k = \frac{\text{Number of neutrons in generation i+1}}{\text{Number of neutrons in generation i}}$$

Clearly if $k<1$, the reactor is subcritical and will shut down; if $k>1$, the reactor is supercritical and power will increase; if $k=1$, the reactor is critical, the chain reaction is stable and the reactor will stay at constant power.

Splitting of U^{235} produces fast neutrons. Let's start off with n such neutrons and track the possibilities between one generation and the next. Some of these neutrons produce fast fission before they slow down. Define the *fast fission factor* ϵ as the number of neutrons slowing down per neutron produced by thermal fission. Clearly $\epsilon > 1$ due to the addition of neutrons produced from fast fission. Thus $n\epsilon$ neutrons are slowing down per n thermal neutrons we start out with. Some will leak out of the reactor before they slow down: define *the non-leakage probability* P_{NLf} as the fraction of fast neutrons which do **not** leak out. That leaves $n\epsilon P_{NLf}$ neutrons left of our original n. Some will get captured in U^{238} resonances[a]; let p be the fraction that escape resonance capture. Now we're down to $n\epsilon p P_{NLf}$ neutrons which have now been thermalized. Some will leak at thermal energies: let the fraction which does not leak be P_{NLth} (the *thermal non-leakage probability*). In order to cause the next generation of fission, they have to be absorbed in the fuel, not in other structures; let the *fraction absorbed in the fuel* be f. Thus of our original n neutrons, $n\epsilon p f P_{NLf} P_{NLth}$ are absorbed in the fuel. About half of these cause fission, producing between 2 and 3 neutrons per fission. Let η (the *fuel utilization factor*, be defined as the average number of fission neutrons produced per neutron absorbed in the fuel. So finally we have $n\epsilon p f \eta P_{NLf} P_{NLth}$ neutrons born in the next generation, and hence

$$k_{eff} = \epsilon p f \eta P_{NLf} P_{NLth}$$

where k_{eff} (effective) is used instead of k for a finite reactor. If the reactor is infinite, there is no leakage, so

$$P_{NLf} = P_{NLth} = 1$$

and the infinite reproduction constant is therefore

[a] This isn't altogether a bad thing because it produces Pu^{239}, which can be fissioned later on. About half the energy in a CANDU comes from fissioning Pu^{239}.

$$k_\infty = \epsilon p f \eta$$

This is the famous “four-factor-formula”. k_∞ can be calculated by calculating each of the four factors on the right. Table 3-1 shows[4] how neutrons are born and die in a CANDU.

Note that about twenty times as many neutrons are produced from thermal fission as from fast fission, which is why one needs a moderator in a water reactor. Think of trying to catch a baseball - it’s much easier if the baseball is going slowly. You can make a reactor work on fast neutrons only, but you have to increase the likelihood of a collision between a neutron and a U^{235} nucleus, and you do that by increasing the number of U^{235} nuclei - hence fast reactors use highly enriched fuel.

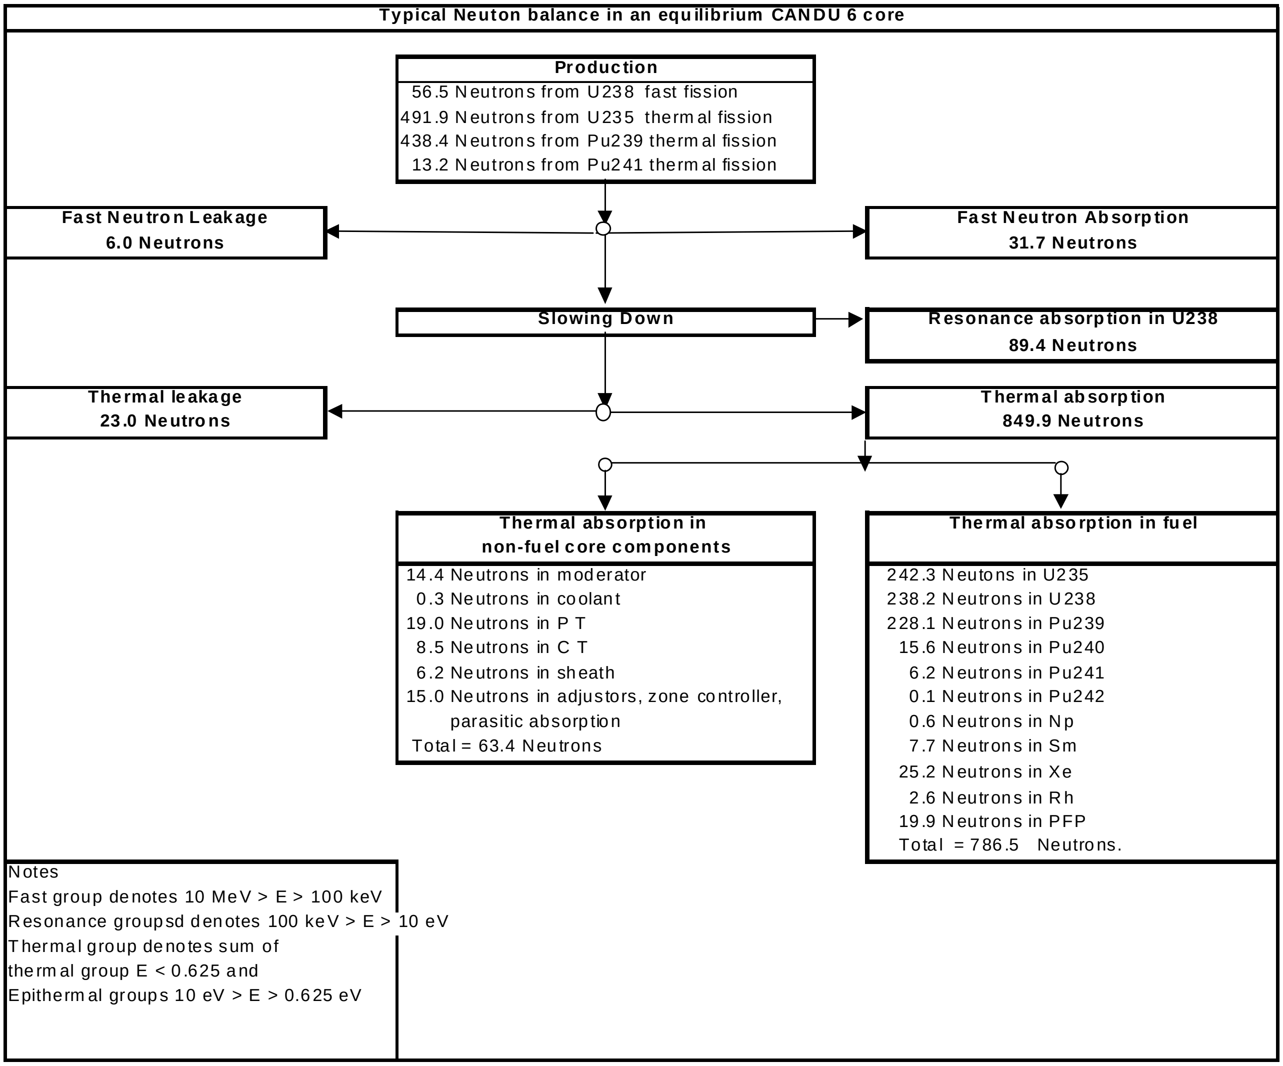

Table 3-1 - How Neutrons are Born and Die in CANDU

Kinetics with No Delayed Neutrons

Consider now an infinite reactor with no delayed neutrons (that is, all the neutrons produced by fission are born instantaneously or are ‘prompt’). Define l_p as the *prompt neutron lifetime*, which is the average time between the ‘birth’ of a prompt neutron and its ‘death’ by absorption somewhere in the reactor. With no delayed neutrons, l_p is equal to the average time between successive generations of neutrons. So if $N_f(t)$ is the number of fissions per unit volume at time t, the number in the next generation is $N_f(t+ l_p)$. But this is the definition of k_∞, so

$$k_\infty = \frac{N_f(t+l_p)}{N_f(t)}$$

but also since

$$N_f(t+l_p) \approx N_f(t) + l_p \frac{dN_f(t)}{dt}$$

we have a differential equation

$$\frac{dN_f(t)}{dt} \approx \frac{k_\infty - 1}{l_p} N_f(t)$$

with the solution

$$N_f(t) = N_f(0)e^{(k_\infty - 1)/l_p t}$$

The "e-folding time" or *period* (the time taken to increase reactor power by a factor *e* or 2.72) is therefore

$$T = \frac{l_p}{k_\infty - 1}$$

For a light-water reactor, $l_p \approx 10^{-4}$ seconds. Assume the reactor is supercritical by 1 part of *k* per thousand - i.e., one *milli-k* or *mk* (much less than the worth of a single control rod in an LWR), i.e., $k_\infty = 1.001$. Then T=0.1 seconds, which means the reactor increases in power by e^{10} in 1 second, or a factor of 22,000 times its initial power. This would make it very difficult to control, to put it mildly. [This is still not a nuclear bomb. You may recall from Chapter 1 that for a nuclear bomb, $l_p \approx 10^{-8}$ seconds. Typically k~1.7 for a small bomb. (Exercise: Calculate the period). A small nuclear bomb releases as much energy inside a baseball-size sphere in 70 **nano**-seconds as a Pickering unit releases in about 15 **hours** - a difference of 12 orders of magnitude.]

Kinetics with Delayed Neutrons

Fortunately, not all neutrons are born at the instant of fission. About 99.4% are - or at least within 10^{-17} seconds of fission. **Nuclear safety depends on the 0.6% of neutrons which are delayed**. This is called the delayed neutron fraction β. We shall see why. First, they arise from the β-decay of certain fission fragments[b]: emission of an electron can lead to a nucleus with excess neutrons, which it sheds at the half-life of the unstable nucleus - which is called a *delayed neutron precursor*. There are a large number of such precursors, each with its characteristic half life. We

[b] Note the unfortunate use of the same symbol β for two concepts which are totally different - the delayed neutron fraction, and the name for an electron emitted in a nuclear decay.

could, if we wished, track each isotope separately; however it is conventional to group the isotopes by half-life into six groups (for ease of calculation). Each group then has an "average" half-life for that group, and contributes a certain fraction β_i to the total delayed neutron fraction β. Half-lives range from tenths to tens of seconds. The fractions and the total delayed neutron yield are particular to each fissionable isotope: thus, for example, for fissioning in U^{235}, the total delayed neutron fraction is 0.0065; whereas for plutonium it is 0.0021. A pure plutonium-fuelled reactor is thus more of a challenge to control than one fuelled with U^{235}, although recall that as the U^{235} burns up, some U^{238} is converted to Pu^{239}.

Table 3-2 shows a typical set of delayed neutron fractions and decay constants[c] for CANDU. Note the first group with a very long lifetime; this is unique to heavy-water reactors and is due to *photoneutrons* - neutrons produced by gamma-ray bombardment of heavy water.

Table 3-2 Delayed Neutron Fraction and Precursor Decay Constants for CANDU

Group	β	λ (sec^{-1})
1	0.295525E-03	0.612382E-03
2	0.116650E-02	0.315532E-01
3	0.103604E-02	0.121850E+00
4	0.235673E-02	0.317638E+00
5	0.783992E-03	0.138922E+01
6	0.198545E-03	0.378577E+01
TOTAL	5.837332E-03	
l (sec)*	0.00091806	

where l^* is the mean neutron lifetime. Let each delayed neutron group I (we shall take 6 groups, as is conventional) have a fraction β_i and a mean lifetime l_i (or t_{mi}), measured from the moment of fission to the 'death' of the delayed neutron[d]. The total delayed neutron fraction is, obviously,

$$\beta = \sum_{i=1}^{6} \beta_i$$

and the mean life of all fission neutrons, prompt and delayed, is the weighted average of each component, namely

---

[c] See Appendix A for a brief discussion of radioactive decay

[d]Note that $\lambda = \frac{1}{\ell}$ - See Appendix A

$$l = (1-\beta)l_p + \sum_{i=1}^{6} \beta_i l_i$$

The term l_i is approximately the same as the mean life of the precursors of the i^{th} group (since the time to produce the delayed neutrons by decay of the precursors is much longer than the time the neutrons exist).

As an exercise, you could show that the period of a reactor with delayed neutrons is at least an order of magnitude larger than one which is assumed to have no delayed neutrons, which makes conventional control possible.

We can now generalize the neutron differential equation presented earlier, by making two corrections: only the fraction (1-β) of all neutrons is prompt; and we have an additional source of neutrons from the decay of the delayed neutron precursors, including photoneutrons. Hence, the rate of production of neutrons is the rate of production of prompt neutrons minus the rate of absorption, plus the net rate of production of delayed neutrons, or

$$\frac{dN_f(t)}{dt} = \frac{k_\infty(1-\beta)-1}{l_p} N_f(t) + \sum_{i=1}^{6} \lambda_i C_i$$

The second term gives the neutrons added by the decay of each delayed neutron precursor: i.e., C_i is the concentration of the ith group and λ_i is the decay constant of the ith group.

We can now write the same equation for each group: the net rate of production of delayed neutron precursors from group i is the rate of formation of the precursor minus the rate of decay, or

$$\frac{dC_i}{dt} = \frac{k_\infty \beta_i N_f(t)}{l_p} - \lambda_i C_i$$

At this point we have seven simultaneous first-order differential equations. They are easy to solve with computers.

We can define the *reactivity* for an infinite reactor as

$$\rho = \frac{k_\infty - 1}{k_\infty}$$

Note that near criticality, $\rho \approx k_\infty - 1$, i.e., the excess reactivity. Recall that

$$k_\infty = \frac{N_f(t+l_p)}{N_f(t)}$$

i.e., the ratio of the number of neutrons from generation to generation. Hence

$$\rho = \frac{N_f(t + l_p) - N_f(t)}{N_f(t + l_p)}$$

You can think of ρ as *the fractional increase(or decrease) in the number of neutrons from generation to generation.*

You will hear this term used often, as ρ is a measure of how close the reactor is to being critical. If $\rho < 0$, the reactor is subcritical; if $\rho = 0$, the reactor is exactly critical; if $\rho > 0$, the reactor is supercritical.

Reactivity is clearly a dimensionless quantity, but to give it meaning, it is expressed in units of:

1. *k*, or more commonly *milli-k*: thus if $k_\infty - 1 = 0.001$, ρ is ~ 1 mk.; or
2. β, or "total delayed neutron fraction". When $\rho = \beta$, the reactor is said to have a reactivity of "one-dollar" or "one β". Since in normal operation of a reactor, reactivities are usually quite small, ρ is also measured in units of 0.01β, or "cents"; or
3. percent Δk, i.e., $\Delta k / k \times 100\%$.
4. *pcm*, "percent milli-rho", defined as $10^{-5}\,\rho$

Thus 2% Δk is about 20 mk or 2000 pcm or (for CANDU reactors) 3.5β or 3.5$ or 350 cents. Note that the relationship between mk and dollars depends on the reactor type.

When we talk of reactivity of a device, we use the term "worth" meaning the net reactivity difference due to the device, as in "the reactivity worth of a control rod is 1 mk". This means that if the control rod goes from fully inserted to fully withdrawn, the change in reactivity of the reactor is +1mk.

To put this in perspective, the normal operating control range of a CANDU Classic[e] is about ±7.5 mk. The maximum worth of a single control rod is about 2 mk. The reactivity worth of the shutoff rods in total is -70mk. The reactivity added by a large loss of coolant accident over a period of two seconds is about +8 mk. Adding one new fuel bundle increases reactivity by about a mk[f].

---

[e]We will use the term "CANDU Classic" when referring specifically to the heavy-water-moderated, heavy-water cooled CANDU. We'll use the generic term "CANDU" to refer to either CANDU Classic or the Advanced CANDU Reactor.

[f]These numbers are highly specific to the reactor type. Small reactors typically have higher device reactivity worths since there is less space to fit rods into. For the McMaster Nuclear Reactor (MNR), the regulating rod is worth 4 to 6 mk, the 5 shutoff rods are worth a total of 90

The kinetics equation therefore becomes:

$$\frac{dN_f(t)}{dt} = \frac{k_\infty(\rho - \beta)}{l_p} N_f(t) + \sum_{i=1}^{6} \lambda_i C_i$$

We can (over)simplify the kinetics equations above by assuming one delayed group with average fraction β and average decay constant λ. Then the equations are:

$$\frac{dN_f(t)}{dt} = \frac{k_\infty(\rho - \beta)}{l_p} N_f(t) + \lambda C$$

and

$$\frac{dC}{dt} = \frac{k_\infty \beta N_f(t)}{l_p} - \lambda C$$

Look at the first equation. If $\rho < \beta$, the first term on the left hand side is negative and reactor is maintained critical on delayed neutrons; i.e., the increase in power is mostly governed by the production rate λC of delayed neutrons, with a time constant of the order of tenths of seconds to seconds. If $\rho = \beta$, the reactor is maintained critical on prompt neutrons alone and is said to be *prompt critical*; the production rate of prompt neutrons is exactly equal to the production rate of delayed neutrons. If $\rho > \beta$, the first term on the right hand side is positive and dominates the second, if l_p is small; the reactor is superprompt critical, and we are back to the case we started out with in the absence of delayed neutrons.

The averaging of delayed neutrons with lifetimes varying by four orders of magnitude does not unfortunately give very accurate numerical answers. Don't design a reactor based on a single delayed group.

Most reactors have available more reactivity than β in their control systems, or generated by accidents. The defence against large insertions of positive reactivity is one or both of the following:

- ensure the insertion *rate* of positive reactivity is limited, so that shutdown systems have a chance to act before short periods are reached
- ensure that there are fast *inherent* negative feedback mechanisms which compensate the inserted reactivity

---

mk or so, and a new fuel assembly is worth about 25mk. In the Maple family of reactors, a control/shutoff rod can be worth up to 50 mk. In the Advanced CANDU Reactor (ACR), the void reactivity is slightly negative (a few mk.).

For a CANDU Classic reactor, l_p is about 0.001 seconds, so there is in fact not a rapid change in period as the reactor goes above prompt criticality. A large loss-of-coolant accident, as noted, adds about 8 mk. of reactivity in the first two seconds. The addition rate is slow enough (limited by the speed at which coolant can physically escape from the fuel channels) that shutdown systems can act within the first second and stop the power rise before it does any serious damage. Even if the reactivity does go slightly above prompt critical, this conclusion still holds.

For a Light-Water Reactor (LWR), each control rod has a much higher worth than in a CANDU (several dollars). They are within the coolant pressure boundary (as opposed to being within the low pressure moderator as in CANDU) and can therefore, in an accident such as failure of the rod housing, be ejected under coolant pressure very rapidly, rendering the reactor prompt critical. The prompt neutron lifetime is more than an order of magnitude smaller, so that no shutdown system can be fast enough to catch the accident. Fortunately the use of enriched fuel allows a large *negative fuel temperature feedback* - i.e., as the fuel heats up, the resonance absorption of neutrons in U^{238} increases (due to the Doppler effect) and there is an offsetting negative reactivity which acts fast enough to stop the power increase - or at least stabilize it at a reasonable level. In CANDU Classic, the fuel temperature feedback coefficient is close to zero and plays little role in safety.

The ACR falls between these cases: the void reactivity is negative but relatively small, and the fuel temperature feedback is likewise negative but smaller than in a LWR. Is the small magnitude a safety benefit or a disadvantage?

Note that the fuel temperature feedback coefficient α_f is defined as the reactivity change per change in average fuel temperature:

$$\alpha_f = \frac{\partial \rho}{\partial T}$$

and can be expressed in terms of mk/°C etc. Similar coefficients can be defined for reactivity effects due to moderator temperature, coolant temperature, etc. As noted, α_f is large and negative for Light Water Reactors, small and negative for ACR, and near zero for CANDU Classic.

Figure 3-2 shows reactor period versus inserted reactivity for various prompt neutron lifetimes. Note that as long as the reactivity is below prompt critical, the period is insensitive to prompt neutron lifetime. For LWRs, the period is highly sensitive to reactivity once the reactivity inserted is above prompt critical; it is much less so for CANDU Classic because of the longer prompt neutron lifetime. ACR behaviour falls part-way in between LWR and CANDU Classic.

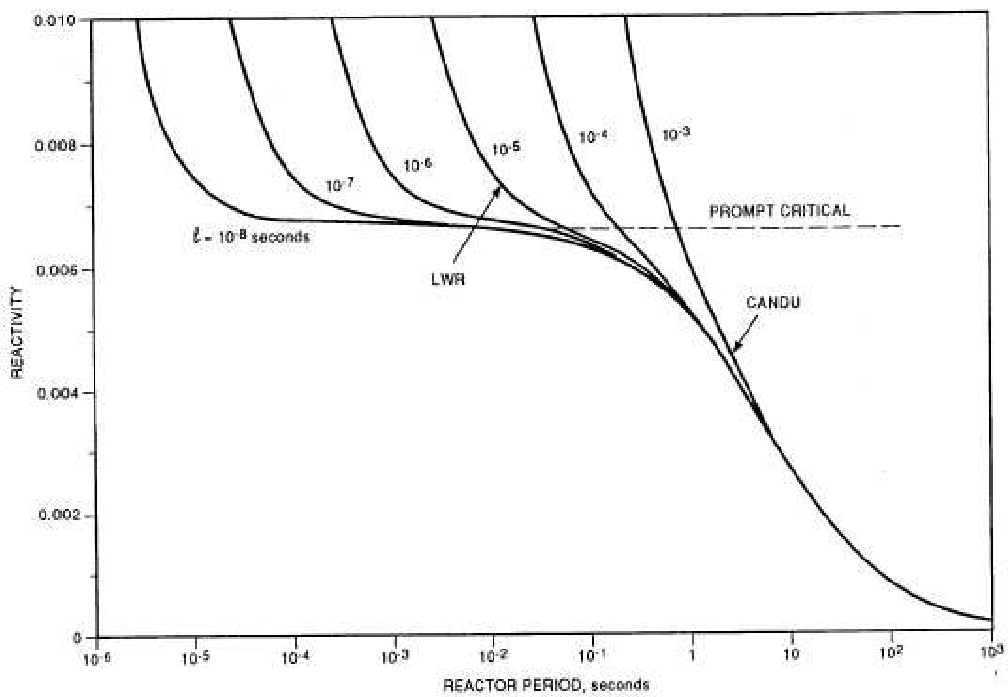

Figure 3-2 Relationship between Reactor Period and Inserted Reactivity for Various Neutron Lifetimes

Figure 3-3 shows the same phenomenon in a different way - it shows the effect of prompt neutron lifetime on reactor period *in a transient*. Two step insertions of reactivity are modelled at time zero: 1 mk. and 6 mk. There is no shutdown. For each one, the prompt neutron lifetime is set at 0.0009 seconds (typical of CANDU Classic) and 0.00009 seconds (still somewhat larger than LWRs). Note that for as long as $\rho \ll \beta$, the prompt neutron lifetime has little effect; but it dominates the behaviour above prompt criticality. The examples do *not* include negative fuel temperature feedback, so the LWR behaviour is artificial: a strong negative fuel temperature feedback would stabilize the transient and is in fact *required* for LWRs.

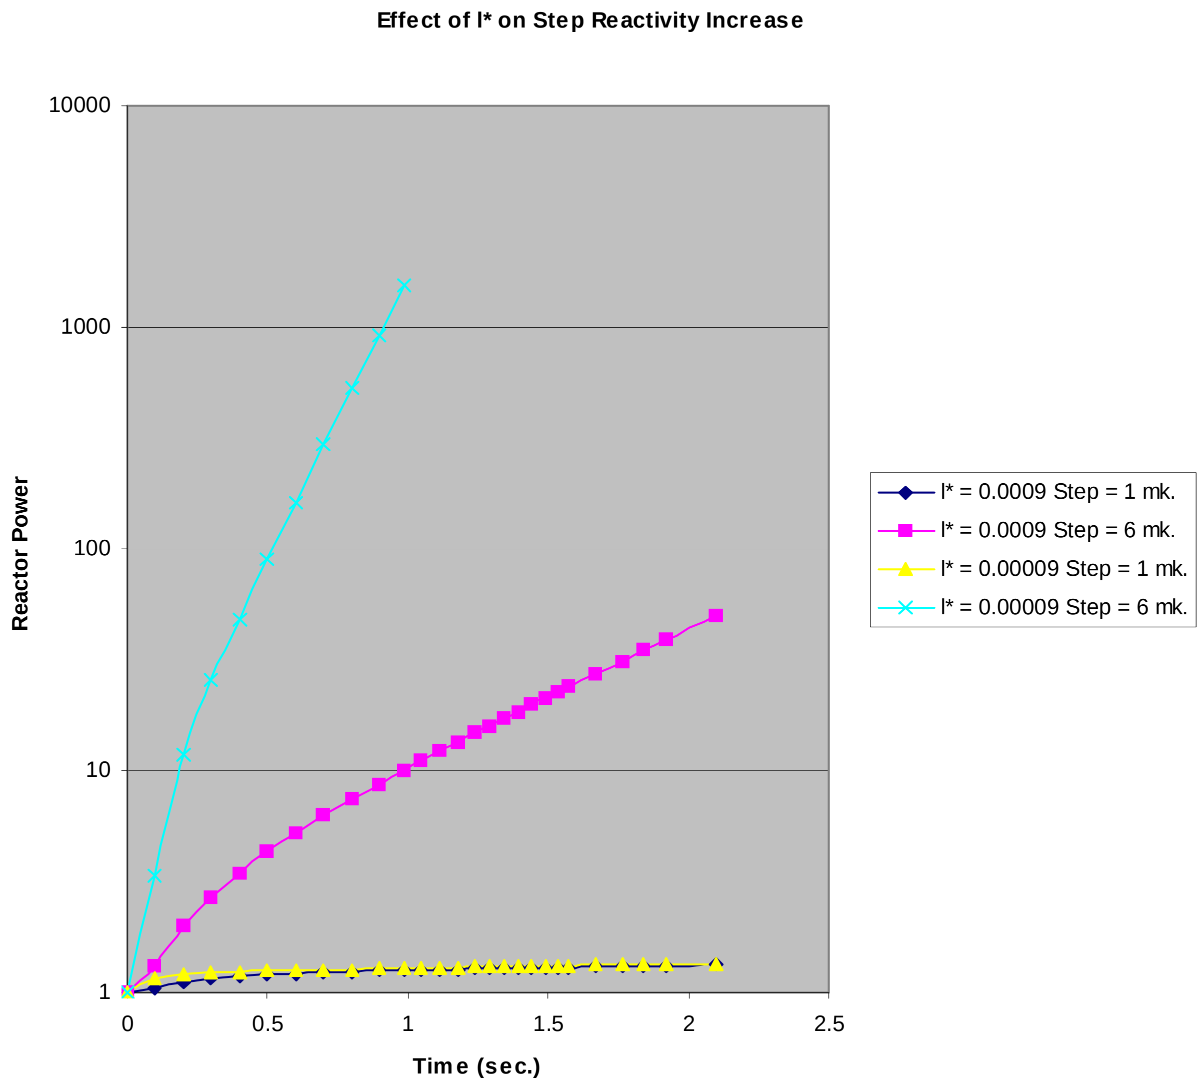

Figure 3-3 - Effect of Prompt Neutron Lifetime for a Step Reactivity Transient

Criticality Accidents and Power Excursions

The earliest accidents in nuclear reactors were criticality accidents. This is in retrospect not too surprising because most early reactors were research reactors, small enough that if the reactor was shutdown, the decay heat could be removed by the water and structures surrounding the core for long periods. Loss of heat sink was not a major issue. Thus the main safety concern was avoiding unwanted criticality, and shutting down quickly if it occurred. A contributing factor was the large worth of some of the reactivity control devices: because the reactors were small, there was not a lot of space for control rods, so that they tended to have high worth.

Finally a number of reactivity accidents have taken place when the reactor was supposed to be shut down. An initially sub-critical reactor is not as safe as it sounds. The power is fairly constant, and proportional to the number of source neutrons divided by the absolute value of the (negative) reactivity - see Appendix B. However it is not under active "control". Detectors are insensitive and indications that the reactor might be near-critical are not very obvious. In some cases the reason for the shutdown is to do maintenance on the control and/or shutdown system so that they are not available. So the safety depends on ensuring an adequate shutdown margin (negative reactivity) via procedural means - i.e. inserting a number of absorber rods, or removing fuel, or adding a liquid absorber to the moderator or coolant, or removing the neutron reflector[g] surrounding the core, and also depends on *maintaining* this margin. This is the reason for the Guaranteed Shutdown State in CANDUs - to have a large enough negative margin that the reactor is resistant to operator mistakes. It is also good practice, as we shall see from the case studies, to have a shutdown system always poised during shutdown.

You may want to think about whether it is safer to do refuelling of a reactor while it is critical or shutdown.

We shall take two examples, SL-1 and NRX.

SL-1[5]

The SL-1 (Stationary Low Power Reactor No. 1) was a natural-recirculation pressurized boiling-water military reactor with a thermal power of 3 MW. It was located at the Atomic Energy Commission National Reactor Testing Station in Idaho Falls, Idaho. Figure 3-4 shows a cutaway of the reactor vessel. At the time of the accident, on January 3 1961, there were 40 fuel assemblies and 5 control rods in the core.

Three operators were engaged in reassembling one of the control rod drive mechanisms while the reactor was shut down. Part of the assembly process required raising the shaft of the rod a limited amount, by hand. This was done; the rod was re-clamped; and the only remaining task was unclamp and to lower the rod. For some unknown reason, the rod was instead *raised* rapidly making the reactor super-prompt critical. The pressure vessel jumped up nine feet, with the chain

---

[g] Many reactors have a reflector which reflects neutrons back into the core - i.e. increases the value of both P_{NLf} and P_{NLth}. In CANDUs it consists of the outer region of the heavy-water moderator; in MAPLE, it is a separate heavy-water jacket surrounding the core. In LWRs it is a region of the coolant outside the core.

reaction being compensated (probably) initially by the negative feedback due to heatup of the fuel and moderator, and by the melting and vaporization of the fuel itself.

Two operators were standing on the lid at the moment of the accident. One was thrown to one side as the vessel rose. The other was impaled on an shield plug ejected from the top of the reactor, and was carried up to the roof of the reactor room where he remained suspended. The third man was killed by radiation and flying debris.

Table 3-2 from Ref. 3 gives a chronology of the events.

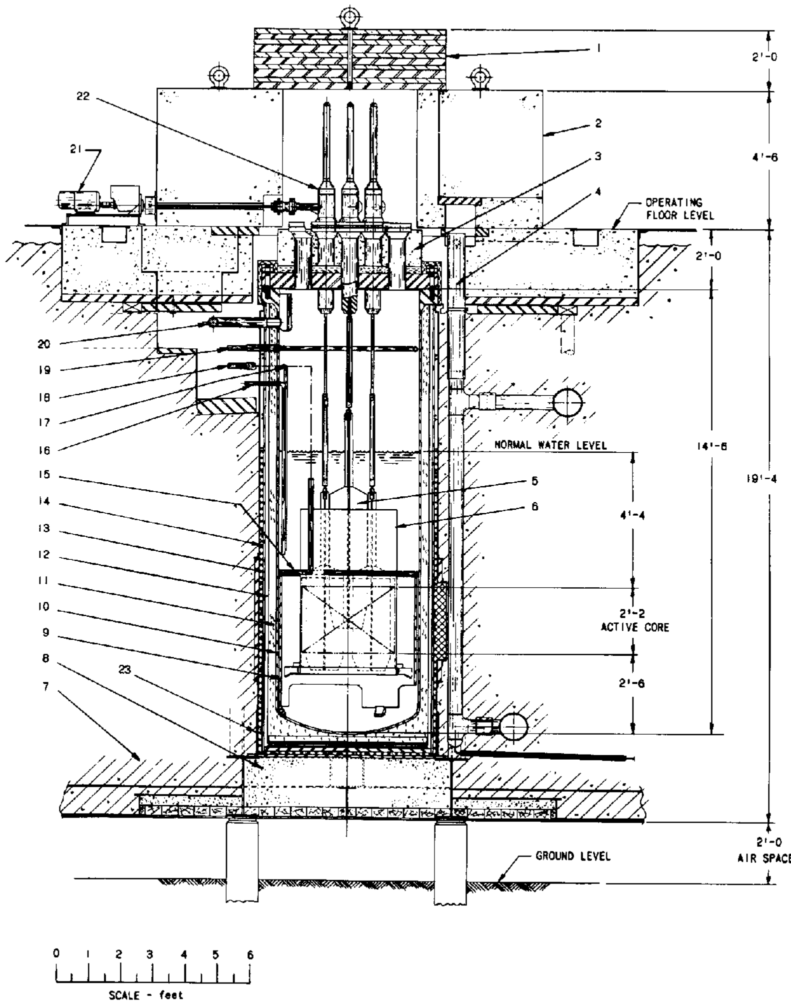

FIG. 3-22 SL-1 reactor installation, vertical section. Numbers indicate principal parts, as follows: 1 - laminated top shield; 2 - concrete shield; 3 - dry shield mixture; 4 - instrument well; 5 - control rod; 6 - core structure; 7 - gravel; 8 - dry shield mixture; 9 - thermal shield; 10 - pressure vessel (4.5 ft outer diam.); 11 - insulation; 12 - air space; 13 - support cylinder; 14 - lead thermal shield and cooling coils; 15 - feedwater spray ring; 16 - separator return line; 17 - purification purge line; 18 - feedwater inlet; 19 - boron spray ring; 20 - steam line; 21 - control rod drive motor; 22 - control rod drive; 23 - Boral shield.

Figure 3-4 - SL1 Cutaway

Table 3-2

Chronology of SL-1 Accident

Time	Event
-500 msec	Central control rod withdrawal starts
-120 msec	Reactor goes critical with rod at 16.7 in. (40.6 cm).
...	Central rod at 20 in. (50.8 cm), period = 3.9 ±0.5 msec, (2.4 ± 0.3)% Δk
0	Peak of power burst. (1.9 ± 0.4) × 10^4 MW.
...	Portion of plates reach vaporization temperature, 2060C (3740F); 5% of centre 16 elements
~2 msec	Prompt nuclear energy release ends; total nuclear energy of excursion = (133±10) MW-sec [+(24 ±10) MW-sec in metal-water reaction]
...	20% of plate area destroyed; centre 16 elements 50% melted; central shroud and control blade ejected from core.
...	Water column above core accelerated by average pressure (500 psi or 35 atm) to velocity 160 ft/sec (49 m/sec).
34 msec	Water slams against lid of vessel.. Maximum pressure ≈10,000 psi (~700 atm).
...	Head shielding ejected. Plugs ejected with velocity of 85 ft/sec (26 m/sec) or less. Vessel rises, shearing connecting pipes. Guide tubes collapse. Nozzles and vessel expand.
160 msec.	First plug hits ceiling.
...	Two-thirds of water expelled. 5-10% fission products expelled.
...	Vessel hits ceiling. Total kinetic energy involved ~1% of total energy released.
...	Insulation ripped from vessel.
2000 - 4000 msec	Vessel comes to rest in support cylinder.

The philosophy of reactivity design is vitally important. The fuel in SL-1 had boron strips attached to the fuel; these were supposed to burn up[h] with the fuel to allow the core reactivity to stay within the control range of the control rods. They had burned up faster than planned, and in addition pieces of boron had fallen off the fuel during operation. The net result was that the core became more reactive, such that it could be made critical with four rods fully in the core and the central one partially withdrawn. The criticality position of the central rod at the time of the accident was 16.7 inches withdrawn; the extra 4 inches that the rod was withdrawn above this value added more than four times β, giving a period of 4 msec. Clearly once the transient had started, no operator action and no shutdown system could stop it in time.

There were many lessons learned from SL-1, but here we summarize the major ones concerning reactivity control:

- It should be impossible for a reactor to be made critical by withdrawing the single most effective rod. Conversely it should always be possible to shut down a reactor with the single most effective rod stuck in its outermost position. This so called "single rod rule" has been followed ever since in safe reactor design.
- It should not be possible to quickly withdraw a high-worth rod. It is still unknown why the operator withdrew the rod. Speculations abound based on second-hand information - was he pulling on a stuck rod to free it? Was it a murder-suicide? Was he "goosed" by one of the other operators who was known to be a prankster? In any case the lesson is that if a reactor requires high-worth rods, there must be mechanical means to prevent their rapid withdrawal.
- In a reactor where large amounts of reactivity can be added in short times, there must be inherent fast negative reactivity feedback via e.g., fuel temperature, such that the transient is stopped short of reactor damage. This lesson was later applied in commercial Light Water Reactors.

NRX[6,7]

NRX was Canada's first large research reactor, built in 1947 with a thermal power of 30MW. The moderator was heavy water, contained in a cylindrical calandria. Coolant was once-through, light-water, taken from, and returning to, the Ottawa River. Passing through the calandria were vertical tubes open to the air at top and bottom[i]. Each fuel rod, made from metallic uranium sheathed in aluminium, had its own cooling jacket (Figure 3-5); each rod, with its cooling jacket, was located

---

[h]Jargon for: being transmuted by neutron absorption to an isotope which (in this case) is less absorptive of neutrons

[i]This describes the reactor at the time of the accident; the design of both the fuel and the fuel channels were changed later on.

in one of the vertical tubes (Figure 3-6). A stream of air passed between the fuel rod and the calandria tube. Twelve boron shutoff rods passed through 12 of these tubes. Startup was by removing half of these rods, whereafter the reactor was controlled by varying moderator level.

The shutoff rods could be raised by compressed air and were then held up by an electromagnet. They could be driven back in using high pressure air, although they would fall more slowly under gravity if released.

The 12 rods were divided into 6 groups, or banks, of rods, ranging from 4 to 1 rods per bank. Normally a "safeguard bank" was held in reserve: that is, bank 1 containing 4 rods was the first to be removed during a startup, and was designed to overcome any reactivity added if the reactor should go inadvertently critical on removal of any of the remaining banks of rods. Removal of the safeguard bank was prohibited by limit switches on each rod unless all other rods were down; however the switches had been giving trouble and had been disabled.

To ensure fast startup after a shutdown (and avoid xenon poison out), it was possible to drive up the first four rods, worth 30 mk., in a few seconds.

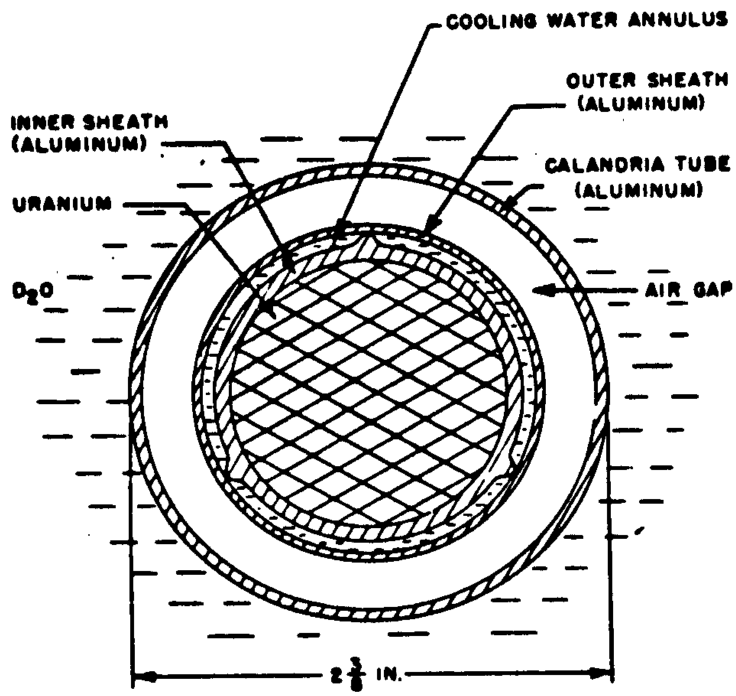

Fig. 1—Cross section of rod structure and calandria tube.

Figure 3-5 NRX Fuel Cross-Section

There were four pushbuttons on the control panel. Pushbutton 4 charged air to the heads of the shutoff rods; release of this air drove the rods down. Pushbutton 3 temporarily increased the current in the electromagnet and ensured that the rods were properly seated. Pushbutton 1 raised the safeguard bank, and Pushbutton 2 raised the remaining banks in automatic sequence. The normal operation was to press Pushbutton 3 along with 4, 1, and 2. If 3 was not pressed with 4, air might leak from the head system; if 3 was not pressed with 1 and 2, the shutoff rods might not be drawn fully home and the safety circuits would prevent startup from proceeding.

The void reactivity was positive and large - that is, removal of all the light water coolant would result in an increase in reactivity of 25 mk.

At the time of the accident, the reactivity of certain fuel rods was being measured at low power. The cooling to these rods was reduced and one was being cooled only by air.

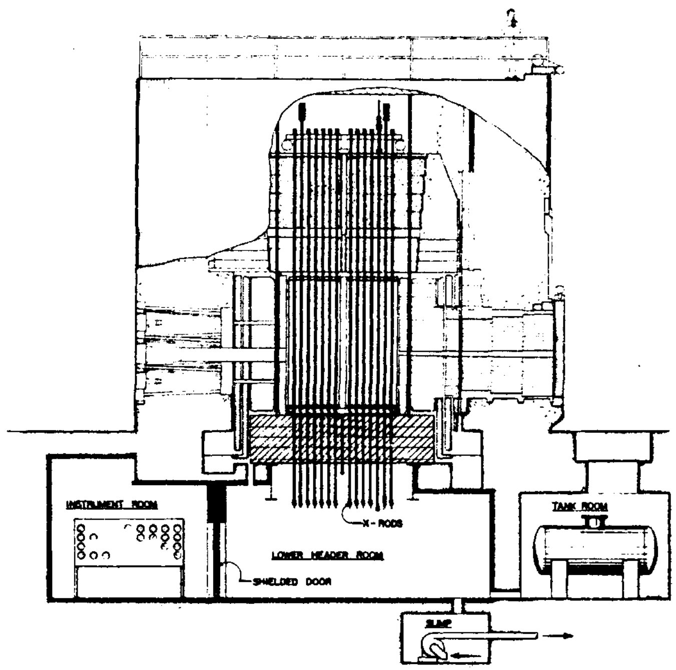

Figure 3-6 NRX - Elevation

The following sequence of events is taken from the presentation by Cross[8]:

“The accident occurred during a start-up procedure. Just as the first group of shut-off rods was about to be removed, an operator in the basement of the building (who had nothing to do with the start-up) mistakenly turned some air-valves which caused several shut-off rods to rise. This was immediately shown by the indicator lights in the control room. The reactor supervisor phoned the operator to stop and went down to the basement himself to make sure that the valves were properly reset. When this was done the rods should have gone down into the reactor. In fact, they went down only part way, but far enough that the lights in the control room indicated that they were down.”

This failure to reinsert is not explained in any of the published reports. Cross continues:

“The supervisor in the basement phoned the control room and told his assistant to press two numbered buttons. He gave the wrong number for one of the buttons and when it was pressed, instead of resetting the air pressure as was intended, it raised four more shut-off rods. If the first group of shut-off rods had been down, as their lights indicated they were, raising four rods was a reasonable thing to do, so the mistake was not recognized."

Specifically, the supervisor asked the assistant to press buttons 4 and 1. This charges the air to the heads and raises the safeguard bank. He had meant to say 4 and 3, which charges the air to the heads and seats the rods. However since button 3 was not pressed, the air to the heads brought up by button 4 leaked away; and the safeguard bank (which appeared to have been up already) was raised. The supervisor realized his mistake but the assistant in the control room had to leave the phone to use both hands, to push the two buttons.

From Cross:

"It was soon apparent from instruments in the control room that the reactor was above critical and the power level was rising. This was surprising but not alarming since the reactor could easily be turned off by dropping the shut-off rods just raised. But when, after 20 seconds, the button was pressed to do this, only one of the 4 rods actually went down. The power level continued to climb and, after some discussion in the control room, it was decided to dump the moderator[j] into a storage tank. Within less than 30 seconds the power-level metres were back on scale and the power dropped rapidly to zero."

The removal of the safeguard bank made the reactor supercritical by about 6 mk, and started a power rise up to about 100 kW. The lone shutoff rod started to fall in at this point, but the power kept rising to about 17MW. Boiling then occurred in some of the temporarily cooled rods, expelling the light water and increasing the reactivity by another 2.5 mk. Power continued to rise to 60-90 MW, when it was stopped by dumping the moderator.

The power surge melted a number of fuel rods, and failed a number of calandria tubes. Eventually in a major operation the reactor calandria was removed and buried; the building was decontaminated; and the reactor was replaced.

The major lesson learned was the importance of the safeguard bank. Thereafter, before a reactor is made critical, the safeguard bank was removed and poised, so that if anything went wrong, it could be quickly reinserted. In CANDU, both shutdown systems 1 and 2 are poised before the reactor is allowed to go critical using the control devices.

A second lesson learned was to make the shutoff rods much simpler. Indeed, after NRX, designers mistrusted shutoff rods as a safety shutdown mechanism; hence NPD and Douglas Point used moderator dump as their emergency shutdown system, and had no safety rods. Pickering had

[j]In the spirit of just-post-war secrecy, the moderator was called "the polymer" to disguise its nature; hence the cry, "Dump the polymer!".

a few safety rods which acted in conjunction with the moderator dump; and it was only by the time of Bruce A that shutoff rods had been restored as a primary shutdown mechanism. They were of course very different from NRX: they consisted of a rod on a cable, which was wound round a pulley and held in place by an electromagnetic clutch. On a shutdown signal, current to the clutch was cut off, and a spring assisted the rod to drop by gravity. It was, and is, a simple reliable design with large clearances.

A third lesson was to separate the control and shutdown systems, with the latter being used only for shutdown and removed from the reactor when it was critical.

Loss of Coolant

TMI

The Three Mile Island reactor was a conventional Pressurized Water Reactor; a flow-sheet is shown in Figure 3-7[9]. The reactor core is inside a large pressure vessel; unlike CANDUs and BWRs, the coolant does not operate in boiling and is kept subcooled by a pressurizer connected to the pressure vessel. A pressure relief valve provides conventional overpressure relief from the pressurizer.

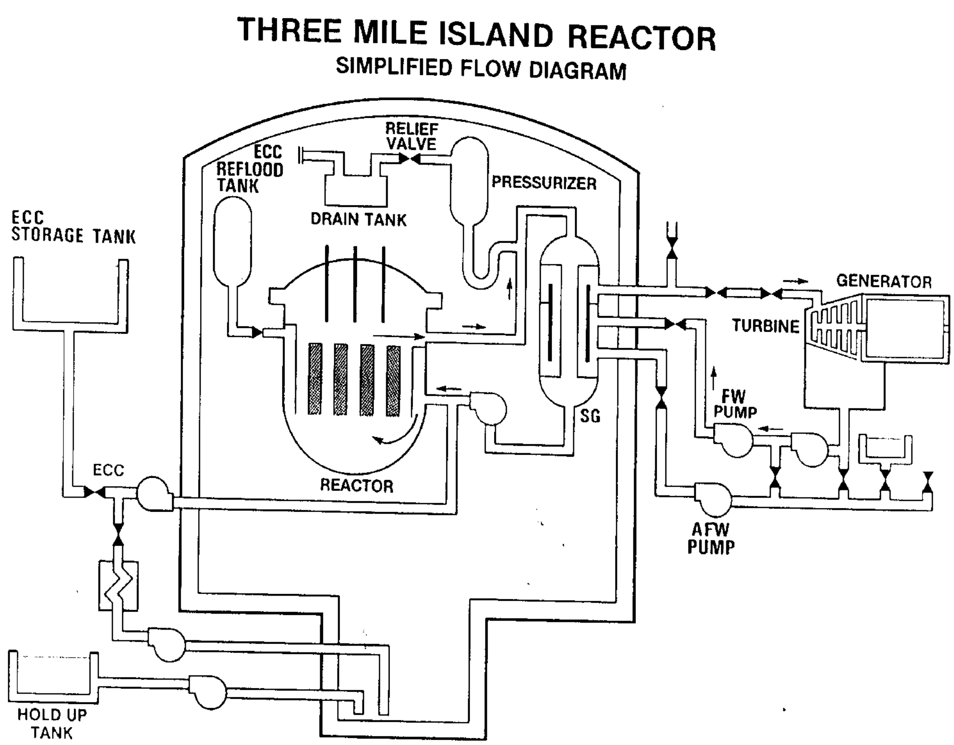

Figure 3-7 Three Mile Island Schematic

On March 28, 1979, a maintenance error resulted in a loss of feedwater to the steam generators when the reactor was at power. Because there was no heat sink for the primary coolant, the coolant pressure began to rise and the pressurizer relief valve opened to limit the pressure, as it was supposed to do; and the reactor tripped (shut down automatically). As the pressure fell, the emergency core cooling system activated to make up for lost water, again as designed; and the pressurizer relief valve should have re-closed, but did not. The net result was a small and continuing loss of coolant through the open valve.

Such a sequence is uncommon but not very rare, and failed-open relief valves, while not treated casually, have occurred several times on CANDU plants. What converted this operating transient into a partial core melt was that the operators did not realize the valve had stuck open, partly because the valve indicator in the control room was derived from the valve actuating signal, not the valve stem position; and partly because the pressurizer level was rising due to the loss of coolant from the top. The operators had been trained when operating a “solid” system (no boiling) to use pressurizer level as a measure of the inventory in the circuit, so they made the erroneous conclusion that there was *too much* water in the system. If you add too much water to a subcooled system, the pressure can rise very fast indeed when the steam space in the pressurizer is filled up, and they were trying to avoid this. Their actions over the next several hours were complex and confused, clearly indicating that they did not understand what was happening, but the net result was that the continual loss of inventory uncovered the core, leading to a partial core melt, hydrogen production from oxidation of the fuel sheaths at very high temperatures, a hydrogen burn inside containment, and a partial melt-through of the pressure vessel, which however remained intact. Finally the failed valve was recognized and then isolated by closing a blocking valve, the main pumps were turned back on to re-pressurize the system and high-pressure makeup water was added to refill it.

The containment did its job but small amounts of noble gases and iodine were released through a leakage recovery system (which could not be isolated because that would also stop primary pump seal cooling). It is in fact a credit to the containment design concept that a partial core melt gave rise to such small releases. It is even more of a credit to water chemistry: fission products such as iodine and caesium combine when released from overheated fuel, to form caesium iodide (an aerosol). However when CsI meets water, it dissociates:

$$CsI \longrightarrow Cs^{+} + I^{-}$$

and is thereafter very difficult to remove (compare this to trying to remove chlorine after dissolving table salt - NaCl - in water). TMI taught that in terms of public safety - “wetter is better”.

The lifetime health effects on the surrounding population were predicted to be 0.7 fatal cancers *total*, and 0.7 nonfatal cancers[10] (this is a prediction, not a measurement, and used the linear dose-effect model we discussed in Chapter 1. More people were predicted to die from the coal burned to replace the electricity lost when TMI was shut down.

The sequence of events has been coherently analyzed by G.L. Brooks and E. Siddall[11] , from which we extract the following:

THREE MILE ISLAND ACCIDENT SEQUENCE

Logic	#	Event
		PHASE A - INITIATING SEQUENCE (A Mishap Within the Normal Range of Viable Operation) (First Minute)
START	1	Wrong detailed design of polisher resin transfer system leading to resin blockage in transfer line.
AND	2	Wrong action by operators in connecting service instrumentation air system to higher pressure water system in trying to clear resin blockage.
RESULT	3	Water in air system caused sudden loss of normal feedwater supply. Turbine tripped. Pilot operated relief valve (PORV) opened correctly.
AND	4	PORV stuck in open position
		PHASE B - SEQUENCE LEADING INTO MAIN ACCIDENT (A Very Important Safety Related Event) (First 50 Mins. After Start)
AND	5	PORV position lights in control room worked from signal and not from valve plug.
AND	6	Previous occasion when 4 had occurred at this same unit (TMI 2) (although for different reasons) had not come to these operators' attention
AND	7	Operators forgot or did not know about 5 and failed to recognize significance of sustained high PORV discharge temperature
RESULT	8	The operators did not realize that the PORV was stuck open.
AND	9	Design engineers and regulatory agency were excessively concerned about avoiding solid (i.e. all liquid) system operation.
AND	10	Operating manuals placed undue emphasis on avoiding solid system arising from 9.

Logic	#	Event
RESULT	11	The operators concentrated on establishing a vapour space in the pressurizer by cutting back emergency injection and maximizing letdown flow despite the fact that the primary coolant was boiling away.
		PHASE C - THE MAIN ACCIDENT (Failure to Take the Last Clear Chance to Avoid the Accident) (From 50 Mins. To 2 Hours after Start)
AND	12	The operating manuals and the operators' training did not deal with the case of a loss of coolant accident resulting from a small leak at the top of the pressurizer, which is what this was. The operators had no guidance or instructions for this contingency.
AND	13	During a period of more than an hour, 2 senior operators, 2 control room operators and the Superintendent of Technical Support all failed to recognize from basic principles that the primary coolant was boiling away, despite clear instrument indications of high primary coolant temperature, low primary coolant pressure, low primary pump flow, high primary pump vibration, and reactivity disturbances.
AND	14	There had been a failure to communicate to the operators or to incorporate in Operating manuals or training the lessons of at least two previous events (Beznau and Davis-Besse) which had given warning of many of the problems here encountered.
RESULT	15	Two hours after the start, much of the primary coolant had boiled away and serious fuel damage started. Despite partial recovery by the operators and other management and engineers who arrived later, the damage continued and adequate cooling was not permanently restored until nearly 16 hours after the start.
		PHASE D - LATER STAGES (A Warning Event for the Future) (during 15 Hours after Start)
AND	16	Primary coolant liquid and vapour, hydrogen and fission product gas and vapour escaped through the open PORV to the reactor coolant (RC) drain tank and then into the containment building.

Logic	#	Event
AND	17	Containment system was wrongly designed for small leaks in primary system. Operators could not isolate (i.e. close off) the containment building without loss of primary coolant pump seal cooling because of the inflexible design of the isolation system. Automatic isolation was even less helpful and was in fact overridden by the operators.
AND	18	Gases and vapour were pumped from the RC drain tank into the waste gas header which had leaks and pressure relief devices - so that gases and vapour escaped to the atmosphere.
RESULT	19	Much of the noble gas fission products and a small amount of radioactive iodine escaped to the atmosphere and caused limited radiation exposure of the public. Much of the radioactive iodine was trapped in water in the auxiliary building and the containment building. Hydrogen, formed from overheated fuel sheaths in contact with steam, burned in the reactor building.

You will note the incredulous tone above that the operators could not see what was happening. It is easy to blame operators after an accident. However analyses of real accidents usually shows a complex interplay between design deficiencies, “safety culture” (see Chapter 9) and operator errors. One cannot fully blame an operator if he is misled by the design.

The TMI accident severely challenged the approach to nuclear safety in the U.S., and to a lesser extent in other countries using similar regulatory philosophies. Some of the lessons learned are obvious, and relate to design of valve position indicators. A lot of effort went afterwards into provision of aids to the operators, so they can understand the state of the plant. One could argue that they already had this information had they chosen to credit it. However once one forms a hypothesis, it is extremely difficult to dislodge it - and the operators were convinced that they had a risk of overfilling (which they had been trained to avoid), not a loss of coolant. The role of the main coolant pumps was strongly debated afterwards: as long as they were kept running, the fuel was cooled even though the inventory was decreasing; and as soon as they were turned off, the core level collapsed and fuel overheating began. Despite this, the final decision was not to keep them running in future.

The main weakness was (much later) conceded to be the prescriptive approach to reactor safety. In the U.S. regulatory approach, the definition of design basis accidents was quite rigid, and there was little incentive to look beyond the “official” list. Most of the safety analysis effort and research money had been spent on the large Loss of Coolant, since it was believed that this would be worse than other loss of coolant accidents. An un-terminated small loss of coolant was not felt to be as important, so there was little training for it and little safety analysis. The formation of hydrogen and its combustion in containment was not anticipated, because regulatory requirements stated that the fuel sheath temperature in all design basis accidents had to be less than 1200°C, which precluded much hydrogen generation. Unfortunately the physics of the TMI accident chose not to obey the regulatory requirements.

The other major changes in direction occurring after TMI were a much increased emphasis on severe accident phenomenology, prediction, mitigation and management; and more use of PSA to enable designer and owner to look “outside” the design basis accident box.

Power Runaway

Chernobyl

Our final example is the well-known accident at the Chernobyl Nuclear Power Station on April 26, 1986.

On that day, the unit 4 reactor suffered a severe accident. The core and much of the building were destroyed; all of the noble gases and several per cent of other fission products were released to the environment. (Note that this is still much less than the ‘catastrophic’ accidents *postulated* in the early days of nuclear power design, which typically assumed *half* the total fission product inventory became airborne - these were supposed to be worst case scenarios).

The reactor design and the accident sequence have been studied extensively since then. The accident evolution is complex and the lessons learned are many. In this chapter we can touch only on the most obvious aspects.

Figure 3-8 shows a cutaway of the reactor core, and Figure 3-9 shows a plant schematic.

Accident Sequence

In the process of performing a safety-related test just prior to a scheduled shutdown, a sequence of events occurred which took the reactor outside the permissible operating range, and at the same time led to the ineffectiveness of emergency shutdown. The combination of operating conditions, control rod configuration, apparent[k] operator violations of procedures, and the inherent core characteristics, led to a large reactivity transient and rapid power rise.

[k]The official Soviet brief laid heavy blame on the operators for not following procedures - some of them were dead and could not defend themselves, and some surviving operators were gaoled afterward. Other experts have claimed that the entire safety culture of reactor operation in the USSR encouraged such violations.

Specifically the test was to show that after a main pump trip, the small amounts of steam generated from the low residual reactor power could provide enough electricity via the turbine-generators to extend the rundown time of the main coolant pumps. Thus the test was done at low power and involved a pump trip. The pump trip at these conditions caused a power rise due to boiling in the fuel channels and the positive void coefficient, exacerbated by the firing of the shutdown system (with the reactor flux shape being highly unusual, the shutdown system worked in reverse and *increased* the power). We shall cover the reasons why later on.

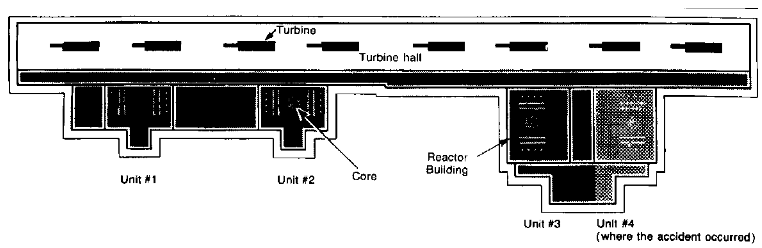

Figure 1 Layout of four reactor units.

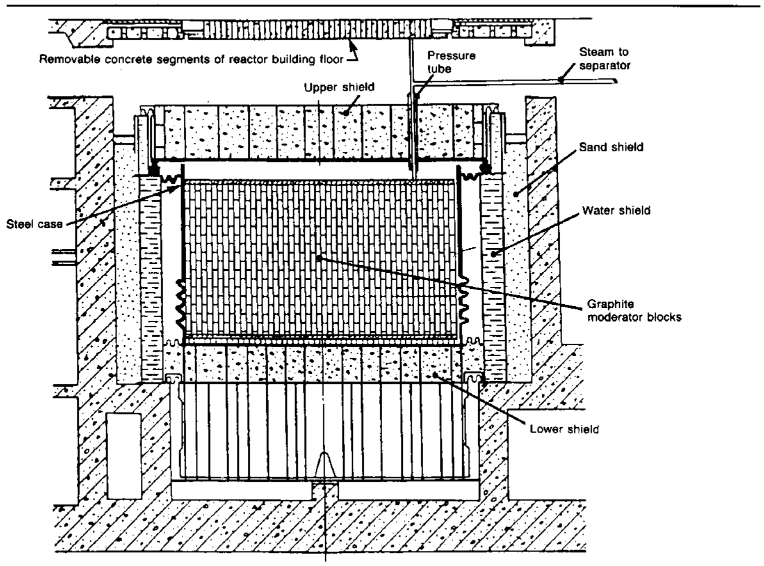

Figure 3-8 Chernobyl Core Cutaway

The fuel energy reached a mechanical breakup level, causing rapid fuel fragmentation in the bottom portion of the core: this resulted in an overpressure in the cooling circuit. Pressure tube failures led to pressurization of the core vessel and loading of the 1,000-tonne reinforced-concrete top shield slab, expelling it from the reactor vault. Burning fragments were ejected from the core, starting 30 fires in the surrounding area.

Immediate Effects of Power Runaway

The core expanded into the surrounding space in the reactor vault (i.e., there was destruction of the radial reflector and the water shield), and dispersal of the fuel and the graphite moderator resulted in the core becoming subcritical. The severing of reactor inlet pipes and outlet pipes and the destruction of the upper portion of the reactor building led to air access to the core. The graphite began to burn locally.

Radioactive Releases

Fragmented fuel and fuel aerosols were expelled in the explosion, and gases and aerosols were taken high (0.8-1 km) into the atmosphere by the thermal plume from the hot core. This continued for several days as the graphite burned and the fuel oxidized, with the rate of release falling as the fuel cooled.

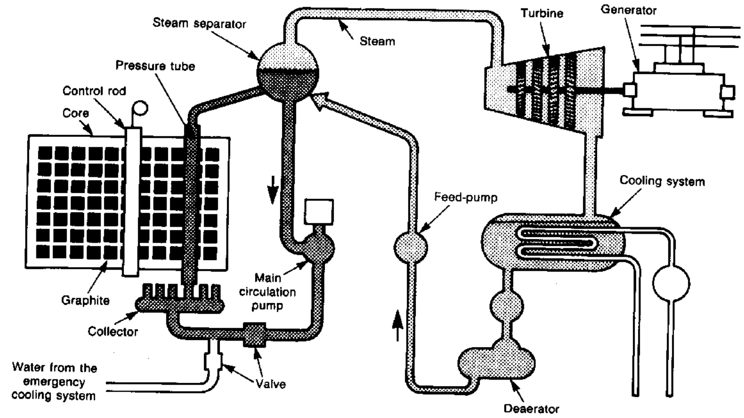

Figure 3-9 - Chernobyl Schematic

To stop the release the Russians dropped about 5,000 tonnes of material, including boron carbide (to ensure shutdown), dolomite (to produce carbon dioxide to try to smother the fire), lead (to absorb heat and provide shielding), and sand and clay (to create a filter bed). This led to a rise in fuel temperature as the convective cooling was cut off. The core reached a hot, oxidizing condition (peaking on May 4), and fission product release rates increased again.

At this stage nitrogen was fed to the bottom of the reactor cavity, cutting off the ingress of oxygen and extinguishing the graphite fire. The fuel temperatures dropped, with a corresponding sharp reduction in releases. The core was believed to then be in a 'stable' air convective cooling mode. Much later, it was discovered that the core cavity was in fact almost empty - the fuel had melted and had run down into the lower parts of the building below the reactor, where it re-solidified. Figure 3-10[1] shows some of the debris which "froze" as it came out of large pipes.

Figure 3-10 - Resolidified molten debris

[1]In the "steam distributing corridor": 1 - fuel lava, 2 - concrete, 3 - steam discharge valve, 4 - capacitor.

Total releases are estimated to be 100% of the noble gases, 20-60% of the volatile fission products (Caesium and Iodine), and approximately 3.5% of the long-lived fission products as shown in the table below[12].

Estimate of Radionuclide Releases During the Chernobyl Accident

(modif. from De95)

Core inventory on 26 April 1986			Total release during the accident	
Nuclide	**Half-life**	**Activity (PBq)**	**Percent of inventory**	**Activity (PBq)**
33Xe	5.3 d	6 500	100	6500
131I	8.0 d	3 200	50 - 60	~1760
134Cs	2.0 y	180	20 - 40	~54
137Cs	30.0 y	280	20 - 40	~85
132Te	78.0 h	2 700	25 - 60	~1150
89Sr	52.0 d	2 300	4 - 6	~115
90Sr	28.0 y	200	4 - 6	~10
140Ba	12.8 d	4 800	4 - 6	~240
95Zr	1.4 h	5 600	3.5	196
99Mo	67.0 h	4 800	>3.5	>168
103Ru	39.6 d	4 800	>3.5	>168
106Ru	1.0 y	2 100	>3.5	>73
141Ce	33.0 d	5 600	3.5	196
144Ce	285.0 d	3 300	3.5	~116
239Np	2.4 d	27 000	3.5	~95
238Pu	86.0 y	1	3.5	0.035
239Pu	24 400.0 y	0.85	3.5	0.03
240Pu	6 580.0 y	1.2	3.5	0.042
241Pu	13.2 y	170	3.5	~6
242Cm	163.0 d	26	3.5	~0.9

Accident Recovery

Firefighting started immediately and external fires were brought under control in four hours. Extensive cleanup and decontamination began. A 'sarcophagus' (reactor burial structure), utilizing a forced-convective air-cooled system with open ventilation and a filtration system, was built around the reactor and turbine hall. The sarcophagus surrounds the reactor and turbine of unit 4.

Radiological Consequences

On-site staff. There were 2 immediate deaths as a result of the accident. Over the next few weeks there were 29 (later 31) fatalities from high radiation doses and burns received by station staff trying to bring the accident under control. The dose distribution for these people was as follows:

Dose (Sv)	*No. Patients*	*No. Deaths*
6 - 16	22	21
4 - 6	23	7
2 - 4	53	1

Offsite - effects of the accident on the surrounding population. Emergency response measures included: 1) distribution of iodine tablets to the population around Pripyat; 2) sheltering for residents of Pripyat before evacuation; and 3) evacuation once the plume shifted towards Pripyat.
The collective dose commitment in the USSR was estimated at the time as:

31×10^4 Sv *external* (over 50 years), and
210×10^4 Sv *internal* (over 70 years).

In both cases the dose commitment is mostly from caesium.

From UNSCEAR[13]: "Doses to the thyroid received in the first few months after the accident were particularly high in those who were children and adolescents at the time in Belarus, Ukraine and the most affected Russian regions and drank milk with high levels of radioactive iodine. By 2002, more than 4,000 thyroid cancer cases had been diagnosed in this group, and it is most likely that a large fraction of these thyroid cancers is attributable to radioiodine intake. It is expected that the increase in thyroid cancer incidence due to the Chernobyl accident will continue for many more years, although the long-term level of risk is difficult to quantify precisely." ... Among Russian recovery operation workers with higher doses there is emerging evidence of some increase in the incidence of leukaemia ... There is no clearly demonstrated increase in the incidence of solid cancers or leukaemia due to radiation in the most affected populations. Neither is there any proof of other non-malignant disorders that are related to ionizing radiation. However, there were widespread psychological reactions to the accident, which were due to fear of the radiation, not to the actual radiation doses". Note that most thyroid cancers are, fortunately, curable.

Design Aspects Relevant to the Accident

This section identifies aspects of the Chernobyl unit 4 design and operation relevant to the accident.

General Design

Chernobyl unit 4 was of the RBMK (roughly translated as ‘large reactor with tubes’) type, and the most recent of the 1,000 MW(e) series. It was a graphite-moderated, boiling-light-water-cooled, vertical pressure tube design, using enriched (2% U^{235}) UO_2 fuel with on-power refuelling. It utilized a direct cycle, to produce electricity from twin turbines (see Figure 3-9).

The reactor core is shown in Figure 3-8. One of the key reactor physics parameters in the equilibrium fuel state is a positive void reactivity with a strong dependence on the operational configuration of the reactor. The design basis called for a maximum void reactivity coefficient of 0.2 mk / % void, whereas at the accident conditions it was reported to be 0.3 mk / % void. (Note that at their *normal* operating conditions, i.e., above 20% full power, the void coefficient is about 0.05 mk / % void.) The overall power coefficient (which includes both the positive void coefficient and the negative effect of fuel temperature increase) is negative under normal high-power conditions, but positive at low power (below ~20%), as was the case just before the accident.

The moderator temperature coefficient is strongly positive for the irradiated core, but because of the slow response characteristics of graphite, it did not play an important role in the accident.

The large core size is noteworthy, since it leads to the potential instability of power distribution, and, in the extreme, to local criticality. In the RBMK reactor a spatial control system is required, primarily for feedback-reactivity-induced spatial instabilities.

The graphite moderator heat capacity is very large, being at least 400 FPS (full power seconds) above ambient at nominal conditions, as compared to that of fuel (11 FPS) and the primary coolant (150 FPS). A distinguishing feature of the RBMK reactor design is the use of the primary circuit as a sink for the moderator heat (5.5% of fission energy). Considerable sophistication has gone into the design of the contact conductance between the pressure tube and moderator, and the conductivity of the moderator cover gas.

With respect to emergency shutdown, the most important features are a slow rate of negative reactivity insertion and *a dependence of that rate on the control rod configuration.* Administrative controls were required to ensure at least 30 equivalent rods were in the core at all times. This heavy reliance on administrative control was traced to early USSR experience in which operators were more reliable than automatic systems.

Thermalhydraulic Design

The RBMK thermalhydraulic design is based on a boiling water, direct-cycle heat transport circuit (see Figure 3-9). Steam mass qualities range from 11 to 22% at nominal conditions. Provision for individual channel flow adjustment is made and is performed manually a few times between channel refuelling, in order to match flow to power.

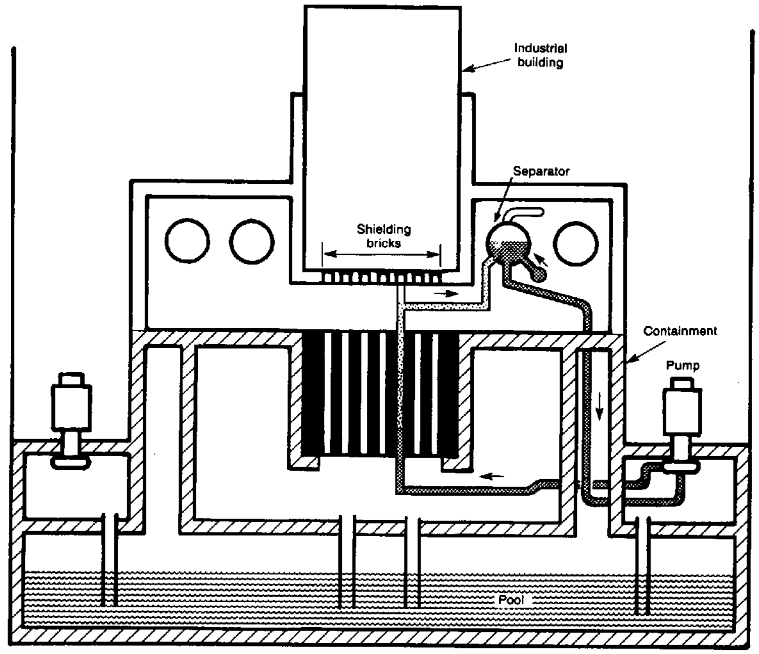

Figure 3-11 Chernobyl Containment /Confinement

There are two normally independent primary circuit loops, which can be interconnected to a single turbine-generator at low power (as at the time of the accident). The primary circuit flow is driven by three pumps per loop. The pumps have significant rotational inertia that permits a transition to thermosyphoning on loss of power without fuel heat transfer concerns. There is a spare pump in each loop that can be started up at power, but because this leads to a reduction of the net positive suction head, it is not normally used.

The condensate from the turbine is returned to the steam separator, and mixing occurs in the drum. Changes in feedwater flow can therefore have a direct feedback on core inlet temperature (separated in time only by a transport delay).

Containment Design

The containment 'localization' system at Chernobyl was a recent RBMK design (see Figure 3-11). In this design the containment was divided into local compartments with distinct design pressures and relief / pressure-suppression to the water-filled 'bubbler pond' in the bottom of the building.

The top portion of the reactor (risers, separators, steam lines, fuelling machine room) was not within a pressure-retaining containment. For small pipe breaks in this group (e.g., a riser tube rupture), it was believed that the large fuelling hall was adequate for the limited discharge rates and low expected levels of radioactivity. Pressure relief for the graphite core vessel was provided by eight 30-cm. pipes connected to the bubbler pool. Relief capacity was stated to be capable of handling a single channel rupture.

In the accident, the steam explosion led to multiple pressure tube failures, which caused a pressure rise in the reactor vault, well beyond design capacity. Thus, the containment localization system played no real role in accommodating the accident. The basic structural integrity of the lower 'containment' compartments was preserved. The upper portions of the building were designed for modest loadings and suffered dramatically from the thermal, and possibly chemical, explosions that occurred.

Variation of Void Reactivity with Reactor Operating State

At Chernobyl, if coolant was lost (voids) from the pressure tubes, there was a positive reactivity addition leading to a rise in power. In fact, the plant was designed to cope adequately with this effect at high power. It was *not* designed to cope with the effect at low power, because the *size* of the void reactivity effect was strongly dependent on reactor operating parameters. Because of the unusual conditions of the reactor just prior to the accident (i.e., low reactor power; only 6-8 control and shutdown rods equivalent in the core, versus 30 required; high coolant flow through the core), there was an abnormally high void reactivity holdup.

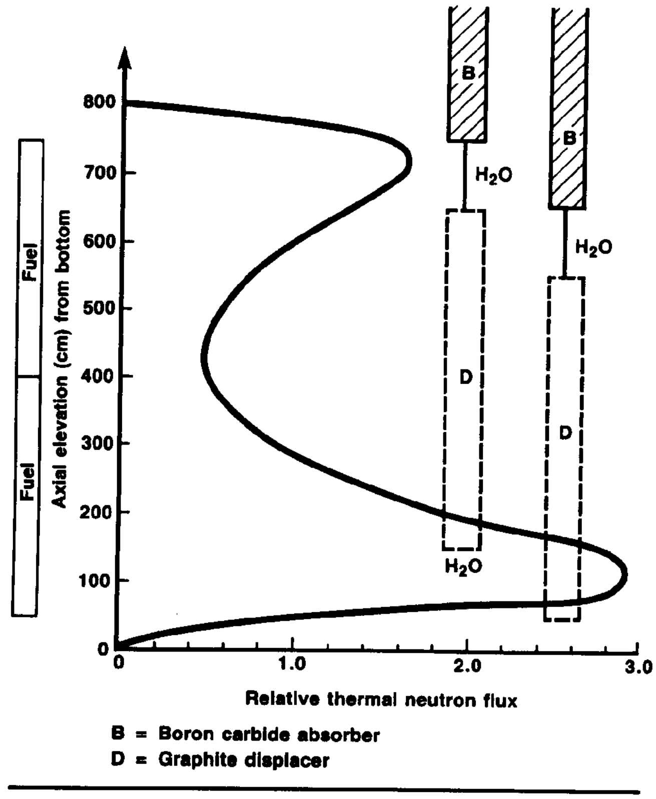

Figure 5 Axial flux distribution preceding accident.

Figure 3-12 Reactivity Effect of Control Rods

Simulations done at AECL and at the U.S. Department of Energy suggest that *positive reactivity was also added by the shutdown system*[14,15]. Normally, the absorber rods are attached to graphite displacers or followers, to increase their reactivity worth. As they are inserted, the absorber rods move into the high-flux region in the centre of the core, which was previously occupied by the graphite, so the absorber rod effectiveness is enhanced (see Figure 3-12). If there were no graphite, the rod would displace water - also an absorber - so the change in reactivity with insertion would not be as great. But in the accident, most of the absorbers were well removed from the core. The flux was peaked at the top and the bottom, where most of the reactor power was being generated. Thus, when insertion of the absorbers first started, the *water* in the high-flux region at the bottom of the core *was first displaced by the graphite follower,* leading to a reactivity increase. Thus operating the plant in an abnormal condition resulted in an unusually large holdup of void reactivity, exacerbated by a deficient shutdown system design, which led to the large power excursion and the resultant core damage.

The size of the system void reactivity in the RBMK-1000 reactor can be controlled to a large extent by operational constraints. The safety of the reactor, therefore, is dependent on the operators' adherence to these constraints. The system void reactivity in this reactor can become significantly higher under abnormal operating conditions. Such conditions include: a) reduction in the number of in-core absorbers with concurrent increase in fuel burnup, which is plausible during loss of refuelling capability; and b) reduction in reactor power level without a matched decrease in coolant flow rate.

In particular, the RBMK-1000 reactor is very sensitive to item b). In order to maintain a near-constant coolant void level in the reactor core, the flow is normally reduced as the power is reduced. However, at low powers (i.e., less than 20% full power), the flow cannot be reduced to match the power, and small changes in coolant conditions can have large effects on coolant void.

Shutdown Systems and Reactor Control

The accident was characterized by a power excursion and an ineffective shutdown; the former, as noted above, may also have been initiated or worsened by the shutdown system design. Reactivity protection (shutdown) and control in the RBMK reactor is complex and requires manual involvement. The control function of the RBMK-1000 reactor is divided into:

1. bulk reactivity control for power manoeuvring and for maintaining criticality in the presence of perturbations caused by absorber rod movement or by feedback reactivity,
2. control of flux and power distribution in the radial plane to limit channel power,
3. emergency reduction of total reactor power to safe power levels when necessary,
4. emergency reduction of local reactor power to safe power levels when necessary, and
5. emergency shutdown of the reactor with the insertion of all absorber rods at their maximum speed.

Demands on the absorber rods are made according to certain rules. The automatic control system attempts to meet these demands. If the operator finds that the automatic control system is insufficient, he inserts or removes 'supplementary' absorbers manually. The number of supplementary absorbers present at any time depends on a combination of factors. Some of these are: 1) the extent of power shaping required, 2) the neutron poison override capability that was required, and 3) the operating value of the coolant void reactivity.

As the demand on the automatic control system increases, supplementary absorbers are driven in or out by the operator to keep the automatically controlled absorbers in their range of travel. However, 24 absorbers are normally kept outside of the core to provide reactivity depth on reactor shutdown.

Required absorber rod positions. A significant feature of this mode of operation is that the maximum negative reactivity rate achieved in an emergency shutdown depends on the *number* of supplementary absorbers present in the core, and in *which locations* they are inserted. For this reason, the equivalent of at least 30 absorber rods are always required to be inserted at least 1.2 m into the core and spread reasonably uniformly over the reactor diameter. This rule was violated prior to the accident.

A significant feature of the rod design, as discussed, is the ingress of water into the bottom of the core that occurs when the absorber and its graphite displacer are pulled out of the reactor.

Emergency shutdown. The emergency protection (shutdown) is designed for both bulk and spatial power excursions. Protection is based on three types of signals:

1. Ion chambers outside of the reflector are used for high flux and high rate trips. One description states that rate is monitored only below 10% full power. Some degree of spatial protection is afforded by tripping if setpoints are exceeded at 2 ion chambers on the same side of the reactor. A total of 8 ion chambers is used by the protection system.

2. Two fission chambers are located near each of the automatic spatial control rods. Both chambers near one rod must exceed their setpoint to initiate protective action

3. One hundred and thirty radially distributed in-core flux detectors (using a silver emitter) are compared to appropriate pre-calculated setpoints, and a partial forced power reduction is initiated by the protection system if the setpoint is exceeded. This system is stated to be effective only above 10% full power. The detectors have a slow response (25-second time constant), so this system would be of no use during a fast excursion in power.

In summary, the ion chambers give only poor spatial protection, but their response is prompt. The fission chambers give better coverage, but there are only a few detectors to cover a large core. Fission chambers are usually also prompt in their response. The in-core detectors give very good coverage, but have a slow response.

The emergency rods are complex devices which can be inserted at various rates, the fastest of which is quite slow (about 10 seconds). This speed limitation is due to the hydraulic drag as the rods are driven or dropped into their water-filled guide tubes. Trips do not appear to be locked in; when a flux reading is no longer high, rod insertion is interrupted. Rods do not appear to be rigidly assigned to the control or protection systems; some appear to serve a dual role.

Physics assessments at AECL show that the Chernobyl reactor is potentially subject to very local, very large flux perturbations[16]. Less than 10% of the core can sustain criticality.

Lessons Learned

Whereas Three Mile Island led to a wholesale re-examination of the approach toward severe accidents, the West learned few new lessons from Chernobyl. Once the accident was understood, a process which took over a year, a thorough check was made of Western designs to see if there was any similar way of disabling the shutdown systems. The answer was negative. In Russia, of course, the effect was far-reaching. Shutdown systems on other RBMKs are reported to have been made faster, and the total void holdup was reduced to $< 1\beta$.

Exercises

Exercise 1

Use Table 3-2 and Appendix A to:

a. Calculate the mean life of each delayed neutron group
b. Calculate the average decay constant of the delayed neutrons (i.e., reduce to one group of delayed neutrons)
c. Estimate the period of a thermal (CANDU) reactor with delayed neutrons. You can use an equivalent "one-group" average delayed neutron model if you wish. You can express the period as a function of the reactivity insertion, or work it out numerically for a few typical values of reactivity. It won't be very accurate.
d. Estimate the period using the same model with no delayed neutrons. You can express the period as a function of the reactivity insertion, or work it out numerically for a few typical values of reactivity.
e. Assume a sudden *negative* reactivity insertion ρ = -50mk. in a CANDU reactor operating at full power. Calculate the power rundown assuming a) no delayed neutrons b) one average group of delayed neutrons

Exercise 2

Analyze the fire in the Narora plant, attached separately, and indicate lessons learned. Describe how you derived these lessons learned. The "official" conclusions from the report have been removed. Your answer will be judged **not** on whether it matches those conclusions but **how you reached it.** You might want to think of dividing your analysis into three types of causes of the event:

- direct causes (immediate precursors which caused the event)
- root causes (causes which if prevented would have prevented the event)
- contributing causes (causes which increased the likelihood of the event)

Exercise 3

A light-water reactor has control rods worth 20 mk each. If one is ejected at zero power, you need to design the reactor so that fuel temperature feedback quickly compensates the positive reactivity, and stabilizes the power at a level (say 130% of full power) which gives the shutdown system time to shut down the reactor. Indicate how you would calculate the required value of the fuel temperature feedback coefficient α_f in mk/°C. Look up typical values of fuel thermal parameters, or leave them as symbols in your solution. You will need to figure out how to calculate the average fuel temperature - express the result symbolically. A steady-state solution is adequate - the solution should not take more than a few lines. [Hint: calculate the stored energy in the fuel].

Exercise 4
Analyze the Point Lepreau spurious douse, attached, and indicate lessons learned. Describe how you derived these lessons learned. The "official" conclusions from the report have been removed. Same caveats as Exercise 2.

Note: There are six headers (large pipes with many spray nozzles) connected to the dousing tank, each header having two valves in series. To operate dousing, both of these valves must be opened. The valves have local air tanks as a backup to the instrument air supply.

Exercise 5
Analyze the McMaster Reactor criticality accident, attached separately, and indicate lessons learned. Describe how you derived these lessons learned. Part of the report has been removed - same caveats as Exercise 2. In addition comment on the robustness of the design and indicate if design changes should be considered (and why).

Exercise 6
The dose distribution for the workers who worked on the (longer term) recovery of the Chernobyl accident was as follows (from UNSCEAR):

Table 16
Estimated effective doses from external irradiation received by recovery operation workers in the 30-km zone during 1986–1987
[I25]

Group	*Number of workers*		*Average dose (mSv)*		*Collective dose (man Sv)*	
	1986	*1987*	*1986*	*1987*	*1986*	*1987*
Staff of nuclear power plant	2 358	4 498	87	15	210	70
Construction workers	21 500	5 376	82	25	1 760	130
Transport, security workers	31 021	32 518	6.5	27	200	870
Military servicemen	61 762	63 751	110	63	6 800	4 000
Workers from other power plants		3 458		9.3		30
Annual total or average	116 641	109 601	77	47	8 970	5 100
Total or average	226 242		62		14 070	

Calculate the expected health effects assuming the linear dose/effect hypotheses.

Point Lepreau Generating Station

Event Report

Inadvertent Opening of Dousing **ER-34310-00.08.21** **(PICA 00-1003)**

This report constitutes:

i) a Detailed Event Report (3)

This Detailed Event Report is believed to be complete ■Yes ■ No

If No, an additional Event Report will be issued as a revision to this report, and will be identified using the same ER number.

ii) an Event Notification Report ■ (3)

iii) an Internal Event Report ■ (4)

PREPARED BY:	xxx Investigator/s	DATE:	00-10-04
APPROVED BY:	____________ Investigation Group Supervisor	DATE:	____________
ACCEPTED BY:	____________ Station Manager*****	DATE:	____________

TRACKED BY: xxx DATE: 00-10-04

*****_The Station Manager's signature indicates his acceptance of the report, and acknowledges his responsibility for implementation of proposed corrective actions._

TABLE OF CONTENTS

1.0 DATE AND TIME OF THE EVENT

2.0 CONDITIONS PRIOR TO THE EVENT

3.0 DESCRIPTION OF THE EVENT

4.0 INVESTIGATION FINDINGS

5.0 REVIEW OF PREVIOUS OPERATING EXPERIENCE

6.0 CAUSE ASSESSMENT AND PROPOSED ACTIONS

7.0 ADDITIONAL OBSERVATIONS

8.0 INTERNATIONAL NUCLEAR EVENT SCALE (INES) RATING

9.0 SAFETY SIGNIFICANCE OF THE EVENT

10.0 REPORTING REQUIREMENTS

FIGURE(S)

Figure 1	Events and Causal Factors Chart

1.0 Date and Time of Event

Event Date

The event occurred on 00-08-21 at 1440 h.

2.0 Conditions Prior to the Event

Details

Prior to the event

- the reactor was in the Guaranteed Shutdown State, and
- air hold testing was being performed on the Dousing System Reserve Air Tanks.

3.0 Description of the Event

Summary

On 00-08-21, the Dousing System Pneumatic Valves (PV) were being test stroked, using the reserve air tanks, in support of the annual test of Loss of Instrument Air to the Dousing PVs. While PV11 was being tested and was open, PV 12 was inadvertently selected Open resulting in a dousing release from one header. Both PVs were immediately selected Closed and a cleanup of the water was initiated. A number of workers were in the Reactor Building at the time of the incident. No serious personnel injuries were experienced as a result of the partial dousing.

4.0 Investigation Findings

Findings

Investigation into the incident revealed

- on 00-08-21, test stroking of the dousing PVs was being performed using Operating Manual Test (OMT) 63431.3, in support of reserve air tank demand test which was being performed using OMT 63431.8. This is a mandatory test scheduled every 52 weeks

Findings (cont'd)

- on the night shift preceding the incident, air hold testing was successfully completed on three reserve air tanks and the six dousing PVs associated with the airtanks. The remaining testing (three air tanks and six associated dousing PVs) was turned over to day shift
- on 00-08-21, at the start of dayshift briefing, two Operators were assigned to progress the remaining testing
- Operator #1 was assigned to the dousing tank platform in the Reactor Building (R/B) to perform the air hold test on the reserve air tanks. He had been assigned the task the previous day and had an opportunity to review the test.
- Operator #2 was assigned to Secondary Control Room (SCR) duties, which included the dousing PV stroke testing as well as other SCR functions. He was assigned the SCR duties upon arrival the morning of the incident. This was his first shift on duty after returning from 2 weeks of time off
- the SCA portion of the testing was intermittent and was performed as required by the air hold test
- the testing also required intermittent support from Main Control Room (MCR) personnel to confirm alarms and indications
- no formal pre-job briefing was conducted, however the 2 Operators were instructed to review the testing information to ensure they understood their individual roles before commencing
- the two Operators stated that although they had performed these tests infrequently, they were familiar with the test procedures, the equipment, and the consequences of inadvertently opening both valves in the same header
- OMT 63431.3 "Manual Operation of the Dousing Valves from the SCR," requires that Operations staff establish communications between the MCR and SCR. Operator #2 elected to use the desk phone in the SCR to establish communications with the MCR.
- the telephone normally used for panel testing in the SCR was out of service with a Work Order Report of Deficiency raised for repairs. Because the cord on a desk phone was not long enough to reach the panel the Operator was required to break communication with the MCR to perform the valve stroke
- the device used to operate the PVs from the SCR is not a typical handswitch, but is a Manual Operation Valve for the actuator air supply 3-way valve
- the Manual Operation Valves are the same for 6 pneumatic dousing valves. The Open/Closed position is not easily discernable by looking at the handle. As well, the identification lamacoids are covered by the valve operator when it is selected in the Open position

Findings (cont'd)

- upon completion of air hold testing on 2 of the air tanks and stroking of the 4 associated dousing valves, the Operators took their dinner break
- before leaving for his dinner break, Operator #2 was requested to perform a test run on the temporary diesel air compressors, which were in place for the Reactor Building Leak Rate Test
- upon returning from his dinner break Operator #2 was asked to support the Instrumentation and Control (I&C) Technicians in the installation of a jumper for Shutdown System 2 (SDS 2) shutter withdrawal, which is supported from the SCR
- Operator #1 returned from dinner and contacted Operator #2 to resume testing. Operator #2 instructed Operator #1 to commence testing, and to contact him when he required the valves to be stroked
- 00-08-21 at 1400 h, Operator #1 completed the air hold test and contacted Operator #2 to have the stroke test performed
- Operator #2 informed him that he was in the MCR preparing for the SDS 2 shutter withdrawal and would contact Operator #1 when he had returned to the SCR and was available to stroke the valves
- 00-08-21 at 1420 h, Operator #2 completed his preparations for the shutter withdrawal and proceeded to the SCR to complete the dousing valve stroke testing. He was expecting a page from the I&C Technicians to assist in the shutter withdrawal
- 00-08-21 at 1435 h, Operator #2 arrived in the SCR and communicated to Operator #1, and the MCR, that he was going to commence valve stroking
- the sequence of subsequent events was as follows:

  - ➢ at 143727.012 h, Operator #2 selected PV 12 to Open from the SCR
  - ➢ after the valve was opened he returned to the phone and confirmed with the MCR that the proper alarms and indications were received. After receiving confirmation he terminated communication with the MCR
  - ➢ at 143812.018 h, Operator #2 selected PV 12 to Close from the SCR
  - ➢ after the valve was closed he returned to the phone and confirmed with the MCR that the proper alarms and indications were received. After receiving confirmation he terminated communication with the MCR
  - ➢ at 143930.686 h, Operator #2 selected PV 11 to Open from the SCR
  - ➢ after the valve was opened he returned to the phone and confirmed with the MCR that the proper alarms and indications were received. After receiving confirmation he terminated communication with the MCR
  - ➢ as he was completing his alarm confirmations with the MCR, Operator #2 was distracted by a telephone page that he believed was from the I&C Technicians inquiring about support for the SDS 2 shutter withdrawal

Findings (cont'd)

- ➢ at 144012.390 h, Operator #2 inadvertently selected PV 12 to Open, rather than the desired action of selecting PV 11 to Close. He immediately realized the error and closed both PVs to isolate the header

- PV 12 was off of its open limit for 25 seconds. It was estimated that 55,000 litres of light water was released to the R/B
- Operator #2 did not follow the Stop-Think-Act-Review (STAR) philosophy before he selecting PV12 to Open
- Following the incident an evacuation of the R/B was initiated. The evacuation was hampered by the fact that the Equipment Airlock was out of service. However, the Equipment Airlock was available to be recalled if required in an emergency situation
- the new security access control tracking and mustering system was a positive aid for Operations Personnel in quickly accounting for personnel in the R/B
- the cleanup was progressed quickly with a large number of station staff volunteering to assist in the cleanup, and
- a technical evaluation was performed on systems affected by the dousing incident, which required inspection of critical equipment to confirm its condition. This information will be documented in a separate technical report.

...end of excerpt

Appendix A - Decay of RadioNuclides

The basic law of radioactive decay is that the decay rate is proportional to the amount of radioactive material remaining - a radioisotope decays at a certain *percent per second of what is left*. Thus if *N(t)* is the number of radioactive atoms at time *t*, then

$$\frac{dN(t)}{dt} = -\lambda N(t)$$

The proportionality constant λ is called the *decay constant*, and the sign is negative since *N(t)* decreases with time.

Transposing and solving,

$$N(t) = N(0)e^{-\lambda t}$$

The *half-life* is defined as the time during which the number of radioactive atoms is reduced by half, or

$$N(t_{1/2}) = \frac{1}{2} N(0)$$

Solving the two equations above, we find

$$t_{1/2} = \frac{\ln 2}{\lambda}$$

Another quantity of interest is the *mean* life, which is just the average:

$$\text{Mean life} = \frac{\sum_{\text{all atoms}} (\text{atoms decaying between } t \text{ and } t + \Delta t) \times t}{\text{All atoms}}$$

Or as $\Delta t \to 0$,

$$t_m = \frac{1}{N(0)} \int_{N(0)}^{0} t\,dN$$

the upper limit of integration being 0 since at the 'end' there are no atoms left. But we know that

$$dN(t) = -\lambda N(t)dt$$

so that

$$t_m = \lambda \int_0^\infty t e^{-\lambda t} dt$$

This is solved by integration by parts[m] giving

$$t_m \equiv \ell_m = \frac{1}{\lambda}$$

Thus the mean life is just the inverse of the decay constant.

[m] $\int v du = uv - \int u dv$

Appendix B - Subcritical Multiplication

The kinetics equations for a subcritical reactor differ from those presented by the addition of a neutron source S_e.

$$\frac{dN_f(t)}{dt} = \frac{(\rho - \beta)}{l_p} N_f(t) + \sum_{i=1}^{6} \lambda_i C_i + \frac{S_e}{l_p}$$

and

$$\frac{dC_i}{dt} = \frac{\beta_i N_f(t)}{l_p} - \lambda_i C_i$$

When $N(t)$ is constant and $\rho<0$:

$$N_f = \frac{S_e}{-\rho}$$

References

1. Some of the material in this section is drawn from lectures by D. Meneley, whose insight I gratefully acknowledge.

2. D.J. Bennet, “The Elements of Nuclear Power”, Longman Group Limited, 1972.

3. J.R. Lamarsh, “Nuclear Reactor Theory”, Addison Wesley, 1966.

4. M.A. Lone, “Fuel Temperature Reactivity Coefficient of a CANDU Lattice”, AECL report RC-2566, December 2000.

5. T.J. Thompson & J.G. Beckerley, “The Technology of Nuclear Reactor Safety”, Vol. 1, M.I.T. Press, 1964

6. W.B. Lewis, “The Accident to the NRX Reactor on December 12, 1952", AECL Report AECL-232, July 1953.

7. D.G. Hurst, “The Accident to the NRX Reactor Part II”, AECL Report AECL-233, October 1953.

8. W.G. Cross, “The Chalk River Accident in 1952”, presented at a symposium on “Historical Perspective on Reactor Accidents”, Seattle, Washington, July 1980.

9. E.M. Yaremy, “A Review of the Incident at the Three Mile Island Nuclear Power Plant”, presented to the Canadian Electrical Association Thermal and Nuclear Power Section meeting, Calgary, Alberta; October 1979.

10. W.P. Dornsife, “The TMI Accident - What Really Caused the Crisis”, Pennsylvania Department of Natural Resources, compilation of talks given by the author, 1979.

11. G.L. Brooks and E. Siddall, “An Analysis of the Three Mile Island Accident”, presented at the CNS First Annual Conference, Montréal, Canada, June 1980.

12. OECD - Nuclear Energy Agency, “Chernobyl - Ten Years On: An Assessment by the NEA Committee on Radiation Protection and Public Health”, November1995.

13. “Exposures and Effects of the Chernobyl Accident”, UNSCEAR Report, Annex J, 2000.

14. P. Chan *et al*., “The Chernobyl Accident: Multidimensional Simulations to Identify the Role of Design and Operating Features of the RBMK-1000", presented at the Conference on Probabilistic Safety Assessment and Risk Management, Zurich, Switzerland; September 1987.

15. U.S. Department of Energy, “Report of the US Department of Energy’s Team Analyses of the Chernobyl-4 Atomic Energy Station Accident Sequence”, DOE/NE-0076, November 1986.

16. P. Gulshani, A. Dastur & B. Chexal, “Stability Analysis of Spatial Power Distribution of the RBMK-1000 Reactor”, presented at the ANS Annual Meeting, Dallas, Texas; June 1987.