Chapter 2
INTRODUCTION TO SAFETY ANALYSIS

V.G. Snell
Manager, Safety Branch
Atomic Energy of Canada Limited - Engineering Company
Mississauga, Ontario

ABSTRACT

This chapter introduces public safety of Nuclear Power Plants from a multi-disciplinary scientific and social point of view. First, the risk is identified: the only risk posed by a Nuclear Power Plant is radiological; that is, related to the escape of radioactive materials. The goal of nuclear safety is to use both engineered and natural defences to prevent or reduce such releases. Almost all of the radioactivity resides in the fuel itself, and can only be released in quantity if the fuel becomes very hot. We prevent such overheating in an accident by stopping the chain reaction and by removing the few percent of residaul heat which is generated after shutdown. Should, however, release from the fuel occur, containment and the chemical behaviour of radioactive elements in water provide very powerful means to contain the radioactivity within the plant. In particular, CANDU plants have certain safety-related features, which we list. Beyond the containment building, natural processes, such as weather, and design features such as an exclusion zone, provide even further dilution. These engineered and natural defences are effective enough that the historical safety record contains very few accidents which challenge the design. Thus, we postulate a spectrum of hypothetical accidents so we have something to design to. We predict the course of these accidents using analytical tools which have been confirmed by experiments, and estimate the dose of radiation to the public. We then compare these estimates with targets set up by our regulatory authority.

2.1 What is the Risk from A Nuclear Power Plant?

The purpose of a Nuclear Power Plant is to produce electricity. In Ontario, it does that at about half the cost of currently available alternative means. Along with this benefit, as with all other forms of making electricity, goes a certain amount of risk. The risk from a Nuclear Power Plant comes from the amounts of radioactivity which are produced as a by-product of the splitting of the uranium atoms. The practice of nuclear safety has been to minimize the risk to a socially acceptable level. In practice, this means keeping the radioactivity where it can be controlled and isolated from contact with people. By comparison, the risk from an oil-fired electrical generating station is chemical rather than radiological and Safety for such plants means minimizing release of disease-causing chemicals to the environment. For hydro-electric plants, the risk is physical and is minimized by ensuring that the chance of a dam failure, coupled with the damage it could do <u>if</u> it failed, is low enough that people accept it. <u>None</u> of these risks is zero.

The behaviour of radioactivity is better understood than that of most environmental pollutants, so that the means of measuring it and handling it have become fairly standard. It can be easily measured by anyone with a hand-held counter. Safe handling means keeping it contained and shielded. During normal operation of a reactor, this is helped by the fact that the radioactive fission products are formed within the uranium bundles and stay there - in fact, only about 5% of the fission products ever leave the solid structure of the uranium dioxide, and then these rest as a gas in the space between the uranium pellets and the sealed metal sheaths (Fig. 2.1).

In the following section, we shall discuss the effect of an accident on this container.

2.2 Potential Release Paths for Radioactivity and Safety Requirements

Since the radioactivity is in the fuel bundle, a safety expert must consider ways in which it can get out. A simple one is by mechanical damage of the bundle. Such incidents have occasionally happened during the loading and unloading of fuel bundles from a CANDU reactor. The result has been ruptures in some of the sheaths of the bundle involved, releases to the reactor building of a small amount of gaseous fission products, and cleanup of the solid bits and pieces.

Things could potentially be worse if bundles in the core somehow get overheated. Normally bundles operate with their sheaths at around 300°C. If the temperature should rise to about 700 or 800°C, the sheaths will get weaker, and if the coolant pressure should fall significantly, they may rupture releasing the gaseous radioactivity contained in the gap between the fuel and the sheath. At sheath temperatures around 1500°C, the corresponding fuel temperatures are such that the radioactivity in the solid fuel can begin to diffuse out of it, and if the fuel melts (2800°C), then most of the volatile contained radioactivity will diffuse out quickly.

By definition, fuel can only get hot if the heat being generated was more than the heat being removed. For example, if the reactor coolant system was operating normally as described in the first lecture, but the power began to rise above its normal power setpoint, eventually the fuel would overheat. For this accident, all that is required is to terminate the fission reaction and this "reactor shutdown" can reduce the power from 100% or more to about 6% in a matter of a couple of seconds.

Conversely, we can imagine a set of accidents where the power remains controlled but the cooling is lost, so that if nothing is done, the fuel once again overheats. Here the need is two-fold. First one must terminate the fisson reaction. As I mentioned before, this still leaves us with 6% power being generated. Six percent is an "after glow" from the spontaneous decay of the radioactive fission fragments. It cannot be turned off but will eventually die out with time as the fission fragments disintegrate into stable elements. Two seconds after shutdown, the afterglow is 6%

of what the power was before shutdown; after about an hour it has decreased to 2%. This may not sound like very much, but for a reactor like one of the ones at Pickering 'A', the decay heat right after shutdown is about 90 MW; enough to heat up each fuel pin (if no cooling is provided) at a rate of 1 to 2°C per second. So our second need besides shutdown is to provide backup systems, for cooling the fuel at decay power levels.

How could such under-cooling accidents arise? Well, if one of the pipes containing the high pressure coolant should break, the coolant would be lost through the hole. We would need a backup system to make up for the loss. Another example is a loss of power to the large pumps which circulate the primary fluid. Here the design of the plant relies more directly on mother nature than on an engineered system: we build so that the reactor core - the heat source - is located at the bottom of the heat transport system, and the boilers - the heat sink - are at the top (Fig. 2.2). The thermosyphoning flow that is set up is enough to take away decay power, and has the added virtue of being passive - we don't have to push a button to get it going and there are no moving parts.

Finally, we can imagine a loss of flow or fluid on the boiler secondary side, so that the heat deposited in the primary side fluid has nowhere to go. The solution is brute force and effective: to provide extra pumps to back up the ones that normally force the secondary flow around, and to provide extra tanks of water to make up for any loss of water. Even without the pumps at all, we can get a flow of water from the tanks to the boiler just due to gravity.

So far I have discussed how radioactivity could be released, and how we respond in general - i.e. shutdown and remove decay heat. The next section covers the latter topics, and others, in more detail.

2.3 Engineered and Natural Defences

We in the industry often get so used to phrases such as high quality construction, installation, inspection of components once installed, and quality assurance, that they become catch phrases and lose meaning. These activites are a major part of the process and are important since by far the best defence against accidents is not the number of systems built to cope with them but our ability to prevent them in the first place.

The electrical utility, the plant owner, has really two safety-related goals. The first is to protect operating personnel and the public adequately against radiation, and the second is to protect its rather large economic investment - typically a billion dollars for a 600 MW plant - against the damage that an accident could cause. Fortunately, economics and safety drive a utility in the same direction - since the measures it takes to protect its investment are also those which assure public safety, particularly for the more probable accidents.

Plants routinely have normal backup systems which can take care of most of the accidents described above. For example, the reactor power control system has separate subsystems which watch for and correct power errors; the primary heat transport system has a makeup system which can replenish the coolant lost through leaks or breaks up to about 50 kg. per second; and automatic diesel generators, which supply about 5 MW, start within a few minutes of loss of internal and offsite power.

However, we have also added an additional layer of systems called safety systems. These are systems whose only goal is to intervene if an accident occurs - they have no role in producing electricity or helping with the day to day operation of the plant. Let's take for example reactor power control. This is not a safety function, and is normally done by a series of in-core tanks filled with ordinary water. Ordinary water absorbs more neutrons than heavy water, so by filling and emptying the tanks, the amount of neutron absorption and thus the power can be controlled. These are augmented by solid absorber rods which move up and down in the core. This is also a normal system, not a safety system. Should the power start to rise beyond its set value due to a fault in the power control system, a monitoring function called setback will fill the light water tanks, and drive in the control absorbers; should the fault persist, another monitoring function called stepback will release the control absorbers and let them drop rapidly into the core shutting down the reactor. Should stepback not recognize the fault, a separate and independent set of solid absorbers with their own detectors, amplifiers and actuators - completely physically independent of the absorbers and the light water tanks - will then drop into the core, shutting it down (Fig. 2.3). This *is* a safety system and is called "shutdown system number one". Actually, in Canada, we have yet another shutdown system which consists of a liquid absorber injected into the moderator through horizontal tubes (Fig. 2.3). This is also independent of the first shutdown system and of the control system. Both shutdown systems are examples of *safety systems* whose purpose is to back up and augment the normal defenses of the plant in the event of an accident. Indeed, the shutdown systems are called upon in *most* accidents to reduce the amount of heat that has to be handled.

Of particular interest to thermohydraulicists in the emergency core cooling system which I mentioned earlier. The idea is simple. If we have a hole in the heat transport system, we must push water in at the same rate that it gets out in order to keep a water cover over the fuel. For a practical design things aren't that simple and the best we can do in terms of speed and flow rate allows temporary voiding of the heat transport system and some fuel uncovering before replenishment of the lost water and refilling of the core.

The emergency core cooling system for CANDU reactors has up to three stages: injection of water at high pressure (4-5 MPa) to overtake the break and refill the circuit; injection of water at lower pressure to maintain inventory; recovery and reinjection of cooled water indefinitely. For example, the emergency core cooling system for the 600 MWe reactors consists of two water tanks connected to a high pressure gas tank which is the driving force. The water tanks have enough water in them to fill a heat transport loop three times over, and the water is driven into the heat transport system at an initial pressure of about 4 MPa. Other CANDU reactors, noteably Pickering, have used pumps for the high pressure stage.

Water is injected into all reactor headers (Fig. 2.4), regardless of break location, and this means we have to supply enough flow so that even allowing for the water which could be wasted through a potentially broken header, the remaining flow is sufficient to refill the core and keep it full. It also means the operator does not have to worry about where the break is.

For a large break, things happen fairly quickly. Initial rates of discharge from the break can be as high as 5000 kg. per second. For all that, the emergency core cooling system has typically overtaken the break and is refilling the system within a couple of minutes of the break.

Since in most cases the break cannot be valved off, we need to supply water continually to the heat transport system for a couple of months to remove decay heat. After the high and medium pressure stages of emergency core cooling are finished, the operator will switch to recovering the spilled water from the floor of the containment building. The water goes through heat exchangers and is cooled, then pumped back into the heat transport system. This recirculating mode can last indefinitely without the need to add further water to the building.

During the period of fuel uncovery before the ECC has refilled the core, some fuel sheaths could overheat and be damaged,and this leads us to the third type of safety system. We have had shutdown systems, emergency core cooling systems; now we consider the containment. The containment is a thick concrete building, surrounding those reactor components from which radioactivity could be released. Its function is to minimize the amount of activity released to the environment. The challenge here is to find a way of coping with the enormous volumes of steam and the accompanying high building pressures that would follow a large pipe break. There are several design approaches to this challenge. One is to build a building so that it is very leak tight at the highest pressure reached if all the coolant should discharge into the containment building. Such buildings have steel liners and design pressures of about 0.4 MPa. This is the approach followed by many builders of pressure vessel reactors. Or one can opt for a system which suppresses the pressure rise by actively condensing the steam. This is done in several ways. In CANDU, a dousing tank stores enormous quantities (3500 m^3) of water which is released in a fine spray over the building volume following a loss of coolant accident. The spray condenses the escaping steam and reduces the pressure and thus reduces the demands on the building structure. If the dousing spray is located in the same building as the reactor, then we call it a single unit containment, as we have built at Gentilly-2 in Quebec and Point Lepreau in New Brunswick, and overseas (Fig. 2.5). We have another type where we put dousing in a separate building, kept at reduced pressure. This is called the vacuum building. It is connected to a number of reactor buildings, and the idea is to suck up any steam produced in a reactor building into the vacuum building and condense it there (Fig. 2.6). This so-called vacuum containment system has economic advantages if there are more than two reactors per site, and is the approach followed at Ontario Hydro plants at Pickering (8 units for one vacuum building) and Bruce, (two plants, of four units and one vacuum building each). Typical leakage rates are: for a single unit containment, 0.1% of its volume per day at design pressure, and for a vacuum containment, 1% of its volume per hour at design pressure. The higher allowable value for a vacuum containment results because the vacuum system provides greater pressure suppression, so that the time the containment pressure is above atmospheric (and hence leakage could occur) is shorter.

In addition to the three "prime" safety systems - shutdown, emergency core cooling and containment - there are a number of safety related systems. These include systems which are seismically qualified to function through the worst earthquake at the plant expected in over a thousand years. Example: emergency power and emergency water systems, which supply both electrical power and cooling water to the station should the regular supplies not work.

I would now like to leave engineering systems for a moment and concentrate on Mother Nature. The same Mother Nature that gave us the gift of radioactivity also gave us the means to handle it and to control its dispersal. Many of the radioactive isotopes which are released in an accident and are of concern from a human health point of view, are very difficult to disperse. Many species decay very rapidly, before they can reach a member of the public. Some are solids. Others, while gaseous, are chemically active. Radioactive iodine, for example, had long been feared as a potentially large health hazard following a major accident, because, in 1957, an accident at a gas-cooled experimental reactor, at Windscale, England, had released quantities of iodine. It happens to be gaseous and therefore releasable to atmosphere, and if inhaled, it concentrates in the thyroid gland and can in sufficient quantities, produce thyroid cancer. However, following Three Mile Island, everyone was asking the same sort of questions as the one Sherlock Holmes asked about the dog in the night. Paraphrasing,

> Holmes asked Watson, "What was so strange about the dog in the night".
>
> Wastson replied, "What dog in the night? I didn't hear it bark".
>
> "Exactly my dear Watson".

What is so strage about iodine is that there probably isn't any. The amount of iodine released at Three Mile Island was negligible (about 10 Curies) in comparison to what some people had expected, but a rather basic knowledge of iodine chemistry quickly explains this. Iodine has an enormous affinity for water, and in the presence of metals such as cesium which are also released in reactor accidents, will quickly form stable salts, i.e. cesium iodide. This is chemically similar to our table salt, sodium choloride, and exhibits similar behaviour in that on contact with water it dissociates into an ionic form and stays in the water. Most of you know that it is much easier to get chloride ions into water by dissolving a teaspoon full of salt in it, then to get them back out again and airborne in the form of chlorine gas. So it is with iodine and we calculate that with a "wet" accident - that is, one with lots of water around - all but a small fraction of the iodine will go to the water and stay there and thus be unavailable for airborne release.

There are other radioactive gases which might be released to atmosphere, but these are less biologically important, Curie for Curie, than iodine. The dose of radiation someone would get would depend on how concentrated the gas is by the time it gets to him. Here the solution uses both engineering and Mother Nature. Engineering in the sense that we surround all our Nuclear Power Plants by a one kilometer exclusion zone, in which no housing is allowed (Fig. 2.7). This allows for considerable dilution (by at least a factor of 10) of radioactivity released from the containment building before it reaches a member of the public. Nature also plays a large role in that further dilution, up to factors of 100 or more, will occur if the weather is "good". Good weather following a nuclear accident is not the same as good weather for a picnic. The best dispersal of radioactivity occurs if the weather is blustery and rainy, and the winds are high. Since we don't know what the weather will be at the time of an accident, in our calculations we assume the most pessimistic weather that prevails ten percent of the time - a light wind blowing uniformly in the direction of the highest population density.

The operator also has the option of controlled releases of the containment contents through filters and he can wait to do this until the weather is "good" and the wind is blowing away from population centres.

2.4 Safety Features of a CANDU Reactor

I would like to switch to hardware for a while and describe some of the safety features of a CANDU reactor. Some are intrinsic to the basic design: others are a matter of careful design choice. I shall group them under the types of accidents we have discussed already, namely loss of coolant in the primary side, loss of pumping on the primary side, loss of coolant on the secondary side, loss of pumping on the secondary side, and finally, over-power accidents.

2.4.1 Loss of Coolant (Primary Side)

A. A loss of coolant allows water in the reactor core to boil. This changes the energy of neutrons entering the uranium so as to increase the fisson rate. Thus, a loss of coolant is accompanied by a rise in reactor power. The rate of rise is modest, being limited by two aspects of basic physics. The first is that water can only escape so fast from the core so that the rate of core voiding is physically restricted; and the second is that the time between successive generations of neutrons from uranium fission in a CANDU reactor is about 30 times longer than for other reactor types. This is intrinsic to CANDU, because of the high moderating ability of the D_20 moderator, plus its lower absorption of neutrons.

So although the power rises, it does so slowly enough that simple mechanical or hydraulic devices can turn the power off before it gets too high. For a large break, required shutdown signal delays are about half a second, and the shutoff rod drop time of about two seconds is fine.

This so-called "positive void effect" is often seen as a disadvantage of CANDU reactors, particularly by those people whose reactors decrease in power on a loss of coolant. However, it is a matter of what one is used to. Even for reactors with a negative void effect, one simply has to look for accidents which increase the coolant density in the core (e.g. a sudden closure of the turbine stop valve on a Boiling Water Reactor). For such cases, the power will rise at the time of the accident and shutdown is required.

B. CANDU reactors separate the fluid which takes heat away from the fuel (the heat transport system) from the moderator which does most of the slowing down of the neutrons. The heat transport system is hot (300°C) and at high pressure (10 MPa) and the moderator is cold (about 65°C) and at low pressure (about atmospheric) (Fig. 2.8). All shutoff and control devices work in the moderator environment, not in the coolant environment. They are unaffected by a major loss of coolant.

This is not true of pressure vessel reactors where their shutoff and control devices experience the high temperatures and the hydraulic forces of a loss of coolant accident.

C. The cool low pressure moderator surrounds each fuel channel (Fig. 2.8). Should we have a loss of coolant, and if the emergency core cooling system does not act, the decay heat from the fuel will flow through the pressure tube and the calandria tube to the moderator. The moderator cannot only take away decay heat, it can also preserve the channel integrity. In addition, no drastic change in mode of operation is required: the moderator is always present and the pumps and heat exchangers continuously running. So we have a backup to the emergency core cooling system, something which no other major water reactor design possesses.

D. On the negative side, the same proximity of the fuel to the pressure tube which allowed us to get rid of decay heat to the moderator, also means that if a channel is partially blocked at full power, the fuel and pressure tube will overheat and the pressure tube may fail. This can only happen if the blockage is more than 90% of the area of the channel. Indeed, the channel design allows monitoring of the flow through each channel and this is routinely done. Nevertheless, the design must be able to handle the failure as noted in (G) below.

E. Because the reactor uses natural uranium fuel, we are particularly stingy about neutrons, so the arrangement of pressure tubes in the reactor is almost at an optimum. In other words, if they are compacted or separated, the power will decrease rather than increase. In addition, criticality is impossible with natural water in the channels, or even diluted heavy water, so that a loss of coolant followed by ECC injection cannot cause recriticality.

F. The pressure tube design means that there are hundreds of small pipes connecting each channel to the collectors or "headers" above the core (Fig. 2.4). This has both advantages and disadvantages. It is a disadvantage in that the likelihood of a small pipe break in a CANDU reactor is much greater than for other reactor types. It is an advantage in that because we are aware of this, the CANDU reactor design has paid particular attention to controlling the small loss of coolant. (e.g. provision of automatic boiler cooldown to reduce the pressure quickly, so that the ECC can be injected, (Fig. 2.9).)

G. The reactor core consists of hundreds of pressure tubes, each containing high pressure coolant and separated by moderator. Failure of one of these tubes because of manufacturing defects, is unlikely; we have done experiments which show that a defect will lead to leakage, which will be detected in the gas between the pressure tube and calandria tube, and alert the operator, long before it becomes large enough to rupture the tube. Actually leaks have been detected on Pickering A and on Bruce A, and the owners of the plant were able to replace the channels before the defects became serious. We acknowledge the possibility of failure in the design by providing relief pipes from the calandria to the containment building; these four pipes allow the high pressure coolant to escape from the calandria without over-stressing it. In addition, we have shown that a pressure tube failure, should it occur, will not propagate across the core to other pressure tubes. The pressure tube rupture several years ago in Swiss reactor at Lucens, a CO_2-cooled, heavy water moderated research reactor, confirms this conclusion.

H. The pressure tube structure of the core means that key components can be replaced without dismantling the reactor. This is not true of pressure vessels, because of their size. Thus we do not have to "put up" with growing concerns. The replacement of the leaky Pickering A pressure tubes mentioned in "G" was, in effect, a correction of a design/installation error, done years after the plant went into operation and without compromising Pickering's position as a world leader in overall capacity factor.

I. We replace fuel on power, at the rate of about 1½ channels per day. We use the fuelling machines as well to remove any defective fuel bundles. This reduces radiation fields during operation, and reduces the amount of contamination to be cleaned up should we get a small spill of coolant.

J. Finally, the cost of heavy water (about $300/kg) makes us quite careful about detecting small leaks early, so we can fix them. Hence the annulus gas system mentioned in G. Hence also our ability to monitor the secondary side water for deuterium, and so pick up evidence of a leaking boiler tube. Prevention is again better than cure, and our boiler tube defect rate (61 defects in 300,600 tubes) is about 1% of that for pressurize water reactors.

2.4.2 Loss of Pumping (Primary Side)

As mentioned, the CANDU reactor places the heat source at the bottom and the heat sink at the top of the circuit. To ensure a smooth transition to thermosyphoning, the main pumps are provided with flywheels so that even if power is lost, they take about 2 to 3 minutes to rundown, and in the meantime provide enough flow to give thermosyphoning a chance to start.

2.4.3 Loss of Coolant (Secondary Side), and

2.4.4 Loss Of Pumping (Secondary Side)

Here we have taken some rather careful design choices. The first is that the boiler vessel on the secondary side by itself holds enough water to take away heat for about half an hour at decay power. This means, for example, if the operator finds he has lost feedwater, he has about half an hour to either restore it or to bring in an alternate heat sink. This allows him time to think and diagnose and reduces the need for fast acting automatic systems. Second, all CANDU reactors have a way of removing decay heat from the primary side (other than through the boilers) at all operating pressures. This is called the shutdown cooling system. It gives the operator another heat sink, should he have a loss of coolant or cooling on the secondary side, without having to worry about reducing primary side pressure.

2.4.5 Loss Of Power Control

Because of the natural uranium fuel, the major means of keeping the fission reaction going in the long-term is on-power fuelling. Thus the control devices do not have to have a large "depth". This means that should the control devices go wrong, they can only raise the power at fairly modest rates and for a limited length of time.

In addition, the normal coolant flows are quite high. This is because the fuel pins are relatively close together, to conserve neutrons. We pay for the high flows in higher pumping power, but gain in that should the power rise so that the sheaths dryout, the heat transfer after dryout is still quite reasonable, causing only modest sheath temperatures and allowing an easy rewet after shutdown.

2.4.6 Destructive Events

To cope with destructive events such as fires, missiles, and earthquakes, we follow three principles: redundancy, separation, and qualifying. By duplicating safety functions (such as shutdown, removal of decay heat, and monitoring), and by physically separating them, we reduce the chance of a fire or a missile knocking them both out. This two-group concept is shown in Fig. 2.10. Each group of systems, by itself, can do the three safety functions mentioned above. For plant-wide events, such as earthquakes, at least *one* of the groups is designed to withstand the earthquake, i.e. is qualified.

2.5 Historical Review of Safety Record of Power Reactors

Most of what I've covered so far is pretty theoretical, so it is worthwhile pausing a moment and reviewing what the actual safety record has been. We look first at research reactors, typically low powered reactors which do not produce electricity, but which are used to gain knowledge so that newer reactors can be well designed. Every country that has built up a nuclear industry has had fairly serious accidents at the research reactor stage. In at least a couple of cases, reactor staff have been killed (the SL1 accident in the United States and the criticality accident in Yugoslavia) and at least another three cases have had detectable releases to the surrounding areas - not sufficient to cause a serious health concern, however - Windscale in England, the NRX accident in Canada, and LUCENS in Switzerland. The importance of such accidents should neither be under or over-estimated. They cannot necessarily be read over to power reactors because the design of the latter is so different. On the other hand, they have provided important lessons which were used in the design of the first generation of power reactors. For example, the NRX accident in Canada, in which a reactor went on a power rise which was terminated not by the shutdown systems, but by rupture of a pressure tube, resulted in improvements to the reliability of shutdown and simplification of the shutdown system design for all future CANDU reactors.

In power reactors, the record has shown a negligible effect from accidents of all the free world's power reactors since they were begun. The best predictions from the Three Mile Island Accident are about one to two fatal cancers. These would be undetectable against the number of cancer cases from other causes (160 per 100,000 people per year) and could be expected to occur perhaps twenty years from now.

The real experience leaves us without a basis for the design of the safety systems. What size break will occur? Where can such a break form? Will there be any warning? To get on with the design, we "postulate" accidents: in other words, build up model accidents in sequences which are expected to be more severe than anything that will happen in reality. For example, since we do not know what size break (if any) will occur in the heat transport system, in analysis we postulate that a break can occur anywhere in the heat transport system and be any size up to that of the largest piping. We then test on paper, the design of the emergency core cooling system against these events, and if it performs adequately on paper, we feel we have a robust design. In addition, we have a fairly large experimental program on ECC performance done for both full scale fuel channels and for scaled down heat transport systems. The results of all this is a design which can cope with any real accident that is likely to occur, even though the exact sequence of events of such a real accident cannot be known beforehand. In this sense, we anticipate acts of God.

2.6 Operations

The other point that power reactor experience has taught us, is that the operator is at least as important as all the hardware one can put in. He is the key to recognizing what is going on in an accident, and in taking measures to control it. If he deduces the wrong thing, he can defeat the action of the automatic safety systems and make the accident worse, as happened at Three Mile Island. To help him recognize the symptoms of an accident, both Ontario Hydro and the United States utilities develop full scale simulators of the reactor control room which when coupled with a computer model of the reactor system, can run through a wide range of plant conditions, and teach the operator rather dramatically of the consequences of his decisions. These computer models use, among other things, thermohydraulics. Their ability to simulate accurate detail, however, is limited by the requirement that they must run in real time, since that is what a simulator does; a development effort is now underway in some countries to get accurate thermohydraulics in real time, computing times.

The operator must also have a pretty basic awareness of thermohydraulics, and the designer, in designing instruments for the operator, must know when these instruments are reliable and when they are not. For example, in Three Mile Island, the loss of coolant was from the top of the pressurizer. This caused the pressurizer to blowdown through the valve at the top and the contents to swell. The level measurement in the pressurizer, which worked on a differential pressure measurement, indicated that the pressurizer was full up. The operator interpreted this as saying that he had lots of water in the system and he blocked the addition of emergency core cooling. This fundamental error resulted from the designer not clarifying when such level measurements would be unreliable and from the operator using that measurement alone and ignoring other conflicting evidence of an emptying system. In particular, he had pressure and temperature measurements which told him that his pressure vessel had gone from being liquid filled to two-phase and that his fuel was overheating.

So there is a whole new field opening up where thermohydraulic knowledge is applied to pre-processing reactor signals to give the operator a comprehensive picture of what is going on.

2.7 Regulatory Philosophy

Up to now we have covered the nature of the risk in a Nuclear Power Plant, some of the design counter-measures that can be taken, and some of the natural processes which protect the public. What we have not discussed is how safe in numerical terms this process makes the Nuclear Power Plant.

There are two aspects to this question:

1) It is possible to take an existing Nuclear Power Plant and analyze all types of accidents above a certain frequency, work out the radiation dose to the public coming from each class, and sum up the total risk. I can represent this mathematically: if F_i is the expected frequency of each class of accident and C_i is the number of deaths or injuries resulting from this type of accident, then the total risk can be written as follows:

$$\text{Risk (harm/year)} = \sum_i F_i \text{ (events/year) } C_i \text{ (harm/event)}$$

A number of people have done this with varying degrees of sophistication and what they end up with in general, is a conclusion that Nuclear Power plants are about as safe as or safer than other means of generating the same electricity; namely, coal, oil, hydroelectricity and natural gas, and for that matter, although this is very controversial, renewable forms of electricity production such as solar power and tidal power.

That work is the subject of another lecture by itself, and I won't go any further into it. However, since Nuclear Power is relatively new, it has not taken the traditional approach used by other industries, namely to build up the industry and then see by experience how safe it is. Rather in every country with a Nuclear Power Industry, government bodies called Regulatory Agencies have been set up which specify in advance of construction of the plant, the safety standards it must meet. These standards are related to the real numerical safety of the as-built plant, but their main purpose is to provide the designers with numerical targets to use in the design. The details of each country's approach differ. In Canada, our Regulatory Agency - the Atomic Energy Control Board (AECB) - sets radiation dose limits in the event of an accident, and the designer must test his plant - by analysis - against a wide range of postulated accidents and show that these limits are usually met even for the worst event in the class. Two very broad classes of accidents are examined. The first class, a so-called single failure, covers those cases where a normal reactor system - one that is used in the day-to-day operation of the plant - is assumed to fail completely. However if the normal plant systems can handle the accident so that relief limits for normal operation are met, then the failure doesn't qualify: a single failure is one which requires safety systems action. The second class consists of dual failures, where a single failure is assumed to occur simultaneously with the unavailability of any one safety system. For example, single failures would include a rupture of a pipe in the primary heat transport system, a rupture of a pipe in the secondary side heat transport system, loss of circulation in the main heat transport system, an uncontrolled power rise, etc. Examples of dual failures would include a loss of primary coolant together with the unavailability of the emergency core cooling system, a loss of primary coolant together with an impairment in the containment building such as failure of the ventilation dampers to close, etc. (Fig. 2.11 gives a summary). However, random coincidental process failures are not in this category. In addition, a peculiarity of Canadian Safety Philosophy, only one shutdown system at a time is credited in any given accident.

Given these two classes of accidents, the designer should restrict single failures to an occurrence rate of less than once in three years. Design calculations establish the frequency and can be verified in reality after a few years of plant operation. The frequency for dual failures will not be observable but can be deduced from a combination of the observed frequency of single failures and the reliability of the safety systems. This reliability is established during frequent on-power testing and it is an AECB requirement that each safety system must work 999 times out of 1,000, or have unavailability of 10^{-3} years/year, i.e., 8 hrs/year. Thus the inferred frequency of dual failures should be less than one in 3000 years. The actual experience on unavailability targets has been generally, that shutdown systems exceed the requirements, and ECC and containment fall short. Thus a problem can be overcome in some cases by just an increased test frequency. There has been a continued effort to fix up the latter two, a successful application of the "test/fail/fix" philosophy.

In addition to frequencies, the designer is also required to show that accidents falling into each of these classes meet radiation dose limits to both an individual presumed to be standing at a plant boundary, and to the neighbouring population. These limits are listed in Figure 2.12. The dose limit increases as the frequency decreases (Fig. 2.13); this is in common with other technological risks such as airplane crashes, where the consequences of the crash are accepted if its probability is low enough relative to, for example, car crashes which occur more often, and kill far greater numbers of people but in small numbers per crash.

The dose limits for nuclear power plants originally evolved from a comparison of the acute risks of nuclear power and coal. They have been refined downwards since then. The effects of radiation on the human body are well enough understood that we can make a translation between dose and harm. The individual doses for both single and dual failures would have no significant early effects; however, a one-time dose of 25 rem whole body would over a period of about 20 years increase the risk of cancer death for an individual by about one-half of one percent. For the population dose, the 10^4 man-rem would on average produce one and a half cases of fatal cancer over the entire exposed population plus one case of curable cancer; it would also cause perhaps three cases of hereditary disease within a factor of 5 either way. The dual failure dose of 10^6 man-rem for the population would have 100 times this effect.

The actual experience in about 100 reactor years has been zero dual failures and a few single failures, all of which produced negligible dose.

The latter result is not surprising; if a manufacturer of cars guarantees his transmissions for 50,000 miles, on an average they must last much longer than that, or he will go bankrupt. Similarly, in order for a designer to claim dose targets are met, there must be, and is, redundancy in the design and conservatisms in the analysis which go above and beyond the minimum required to just meet the target.

As a check of the design, and to assess the robustness of the design in the accident recovery stages, we also analyze event sequences not covered by the single/dual failure categories. These include multiple failures in the ordinary reactor systems, as well as failure of the operator to act within a reasonable time. A lot of the thermosyphoning analysis we do falls into this category, since it usually involves multiple process system failures, such as a loss of coolant plus loss of the internal and external power supply.

The results of these analyses are compared to a "risk" line - i.e. - a line of decreasing (probability×consequence) which passes through the single and dual failure points (Fig. 2.13). We do not design for events if the probability is $< 10^{-7}$ per year. If the probability is greater than this, we compare the consequences to our "risk line" consequences at the same probability; and fix the design if necessary.

8. Conclusions

An ideal safety analyst is an expert in reactor physics, properties of materials, reactor design and operation, control theory, two-phase transient thermohydraulics, chemistry, the biological effects of radiation, atmospheric dispersion, history and philosophy, human engineering, and public policy. Perhaps one day we will find one. Meanwhile the thermohydraulics expert plays a key role in understanding what is happening inside the reactor and also in understanding what is happening inside the containment building, and the discipline forms one of the fundamental building blocks of safety analysis.

9. References

The following document provides a good introduction to safety in general, and the 45 references it contains lists most of the key available publications on CANDU safety; "Safety of CANDU Nuclear Power Stations", by V. G. Snell, Atomic Energy of Canada Limited, Publication AECL-6329, July, 1980.

CANDU FUEL BUNDLE

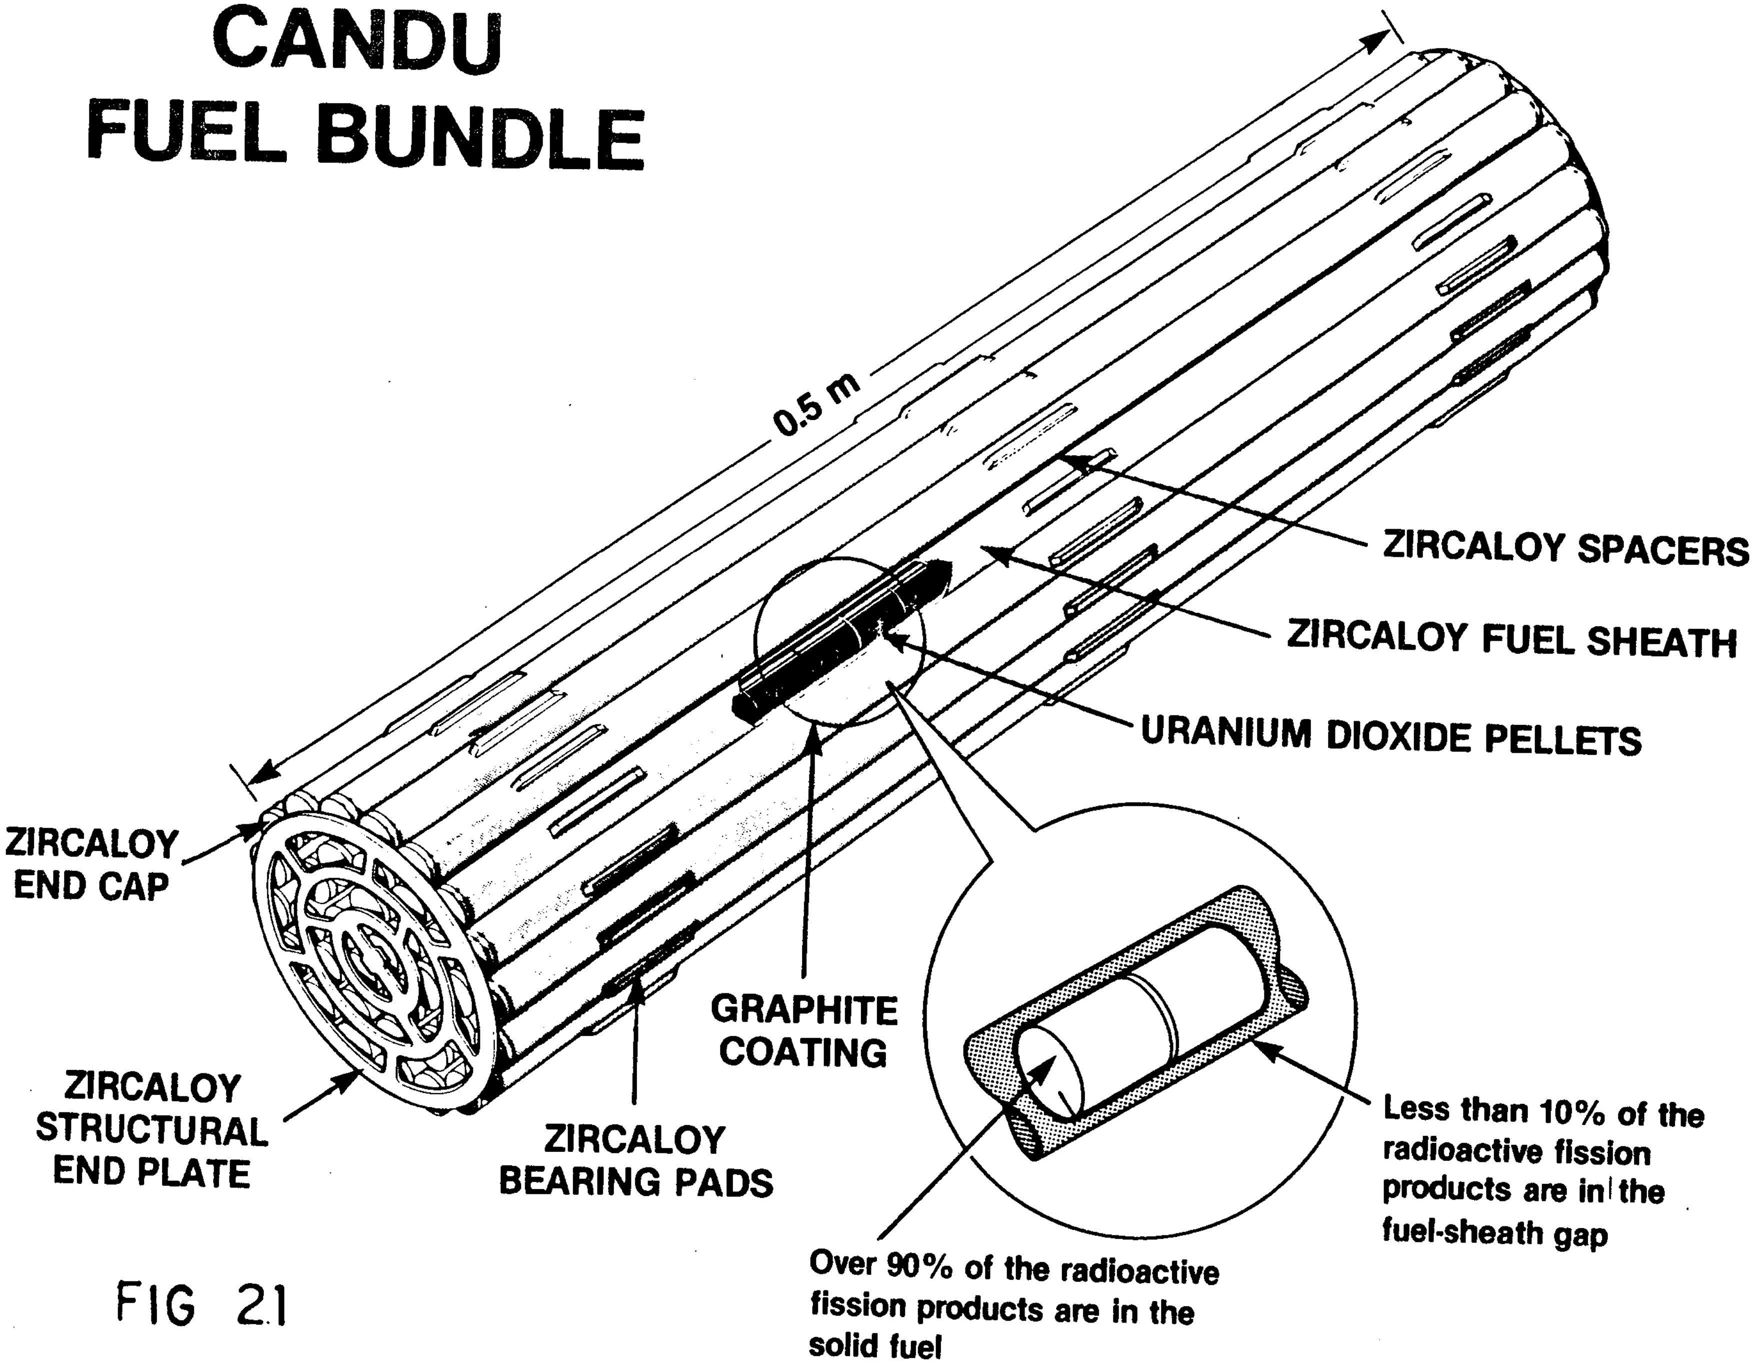
0.5 m
ZIRCALOY SPACERS
ZIRCALOY FUEL SHEATH
URANIUM DIOXIDE PELLETS
ZIRCALOY END CAP
GRAPHITE COATING
ZIRCALOY STRUCTURAL END PLATE
ZIRCALOY BEARING PADS
Less than 10% of the radioactive fission products are in the fuel-sheath gap
Over 90% of the radioactive fission products are in the solid fuel

FIG 21

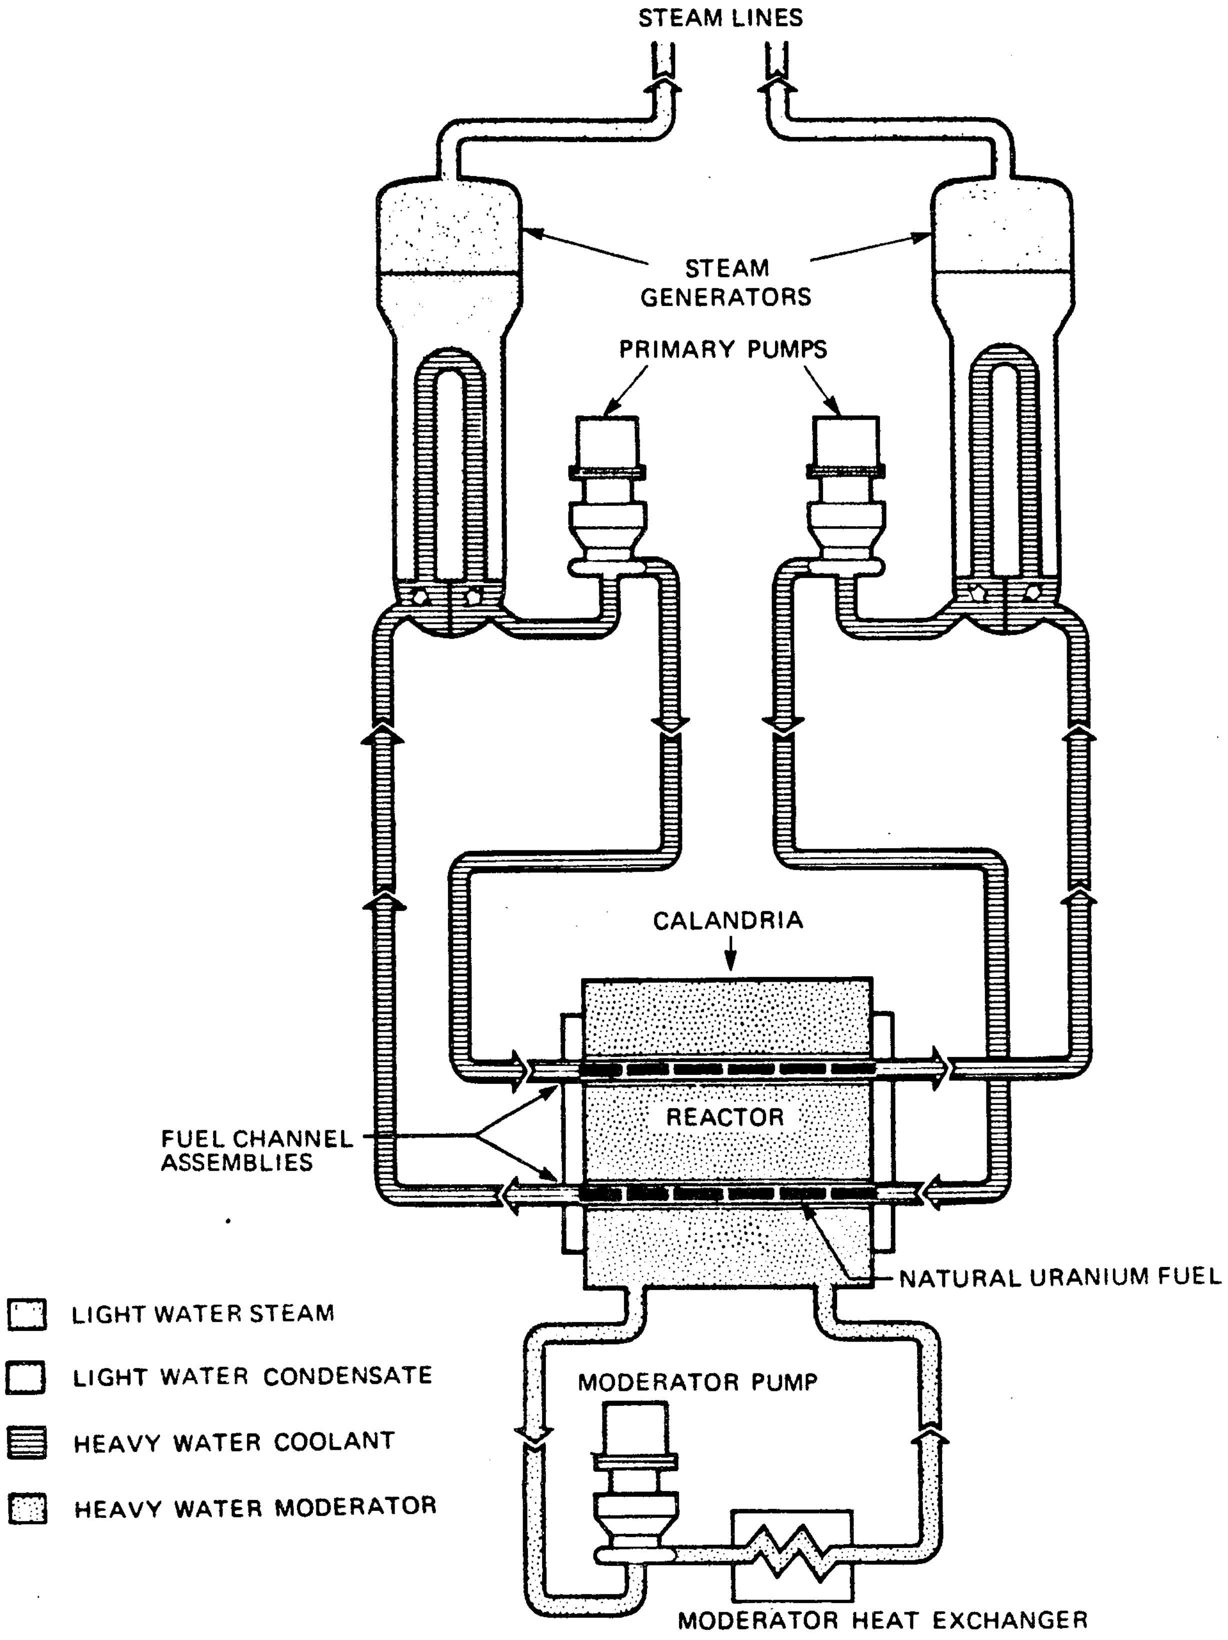

FIGURE 2.2 CANDU NUCLEAR STEAM SUPPLY SYSTEM

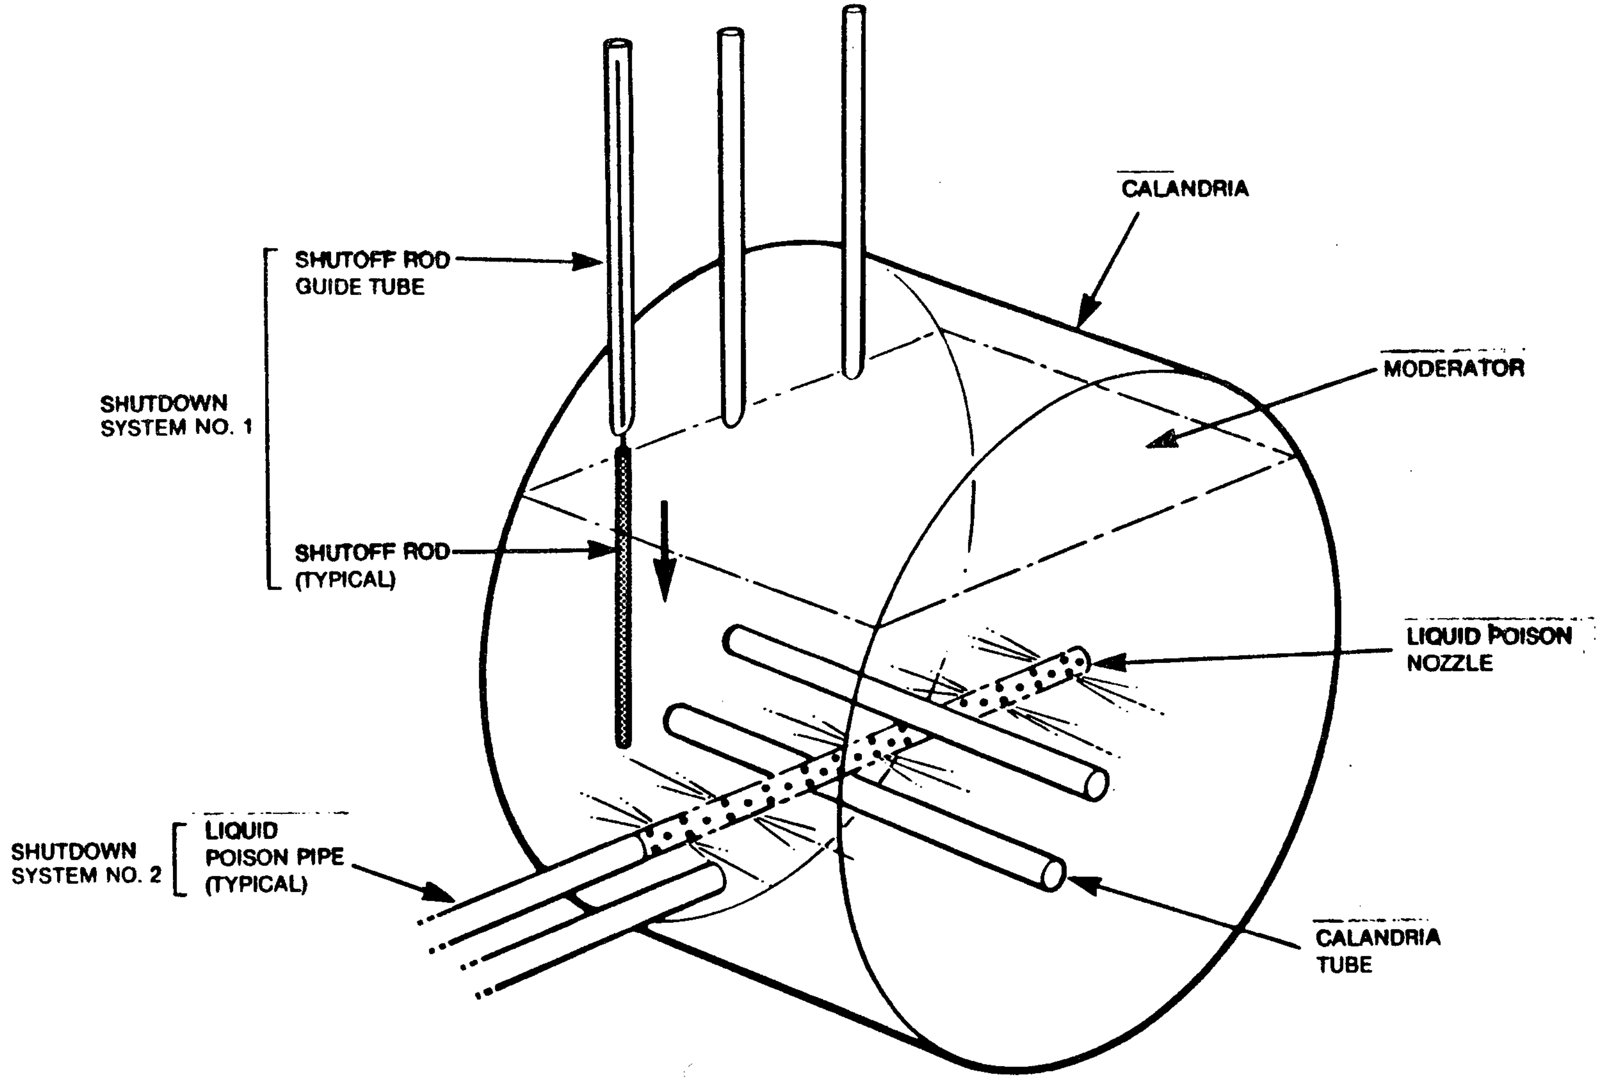

FIGURE 2.3 SHUTDOWN SYSTEMS: SHUTOFF RODS AND LIQUID "POISON" INJECTION

EMERGENCY COOLING

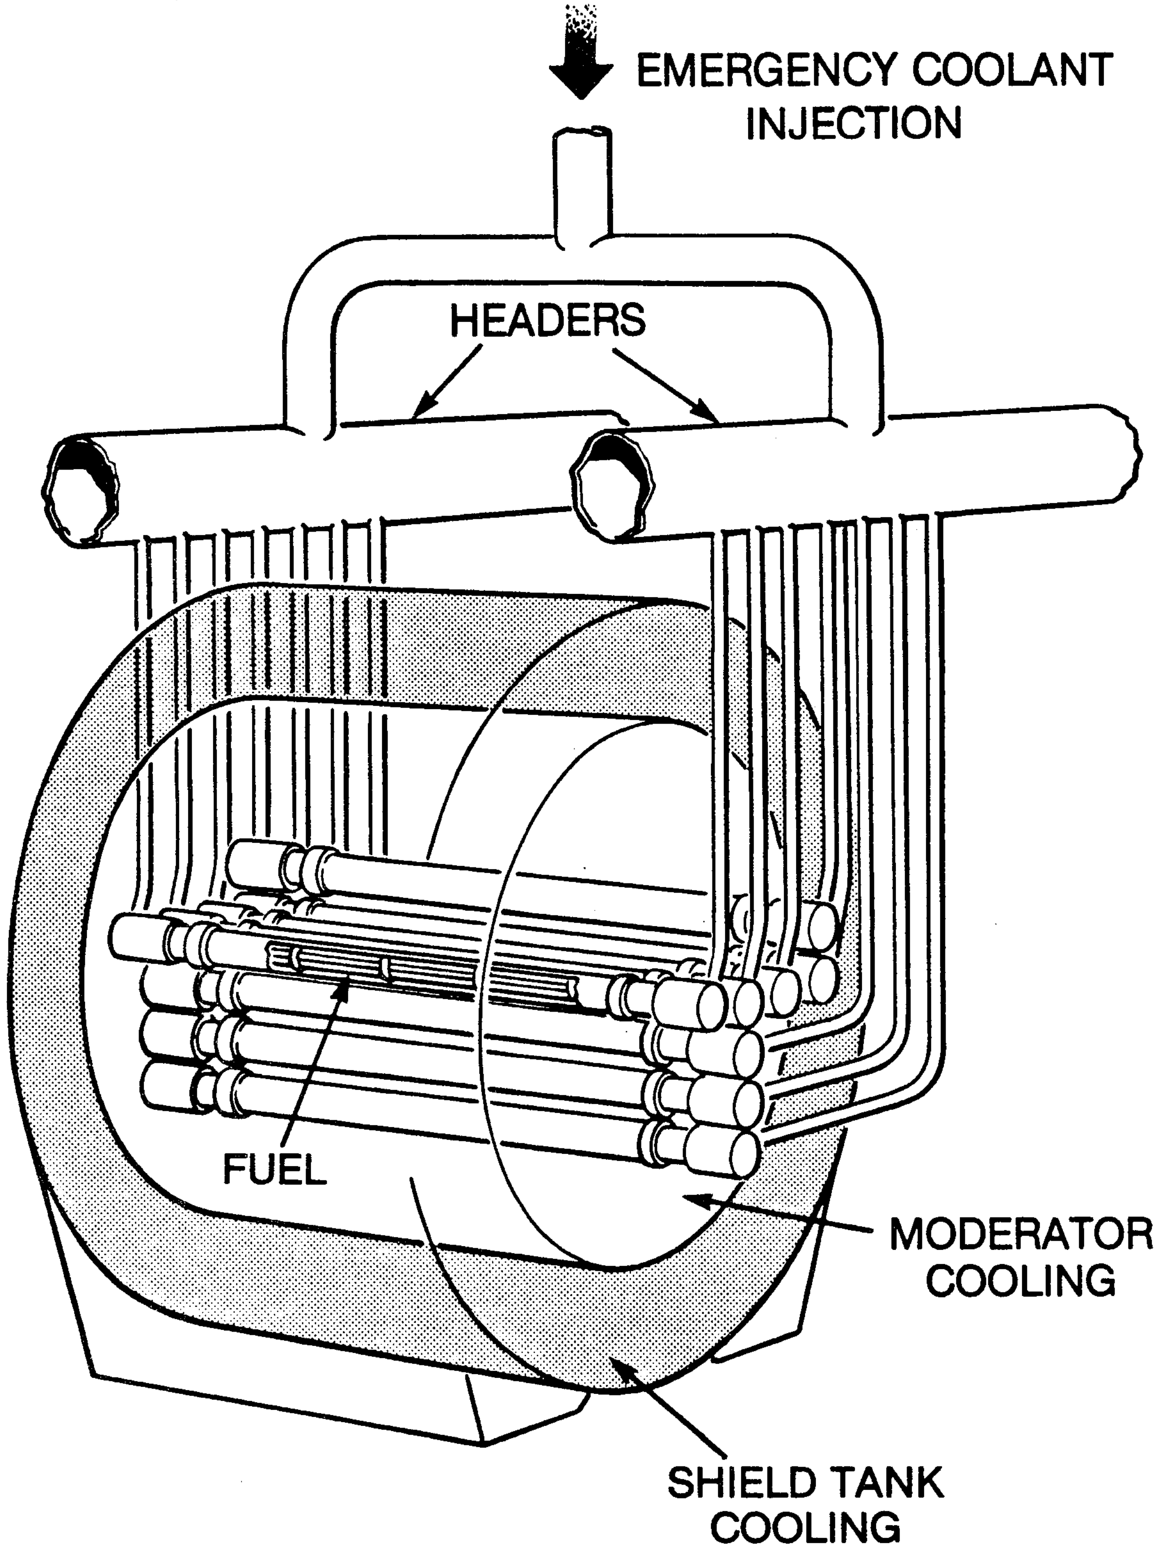

FIG 2.4

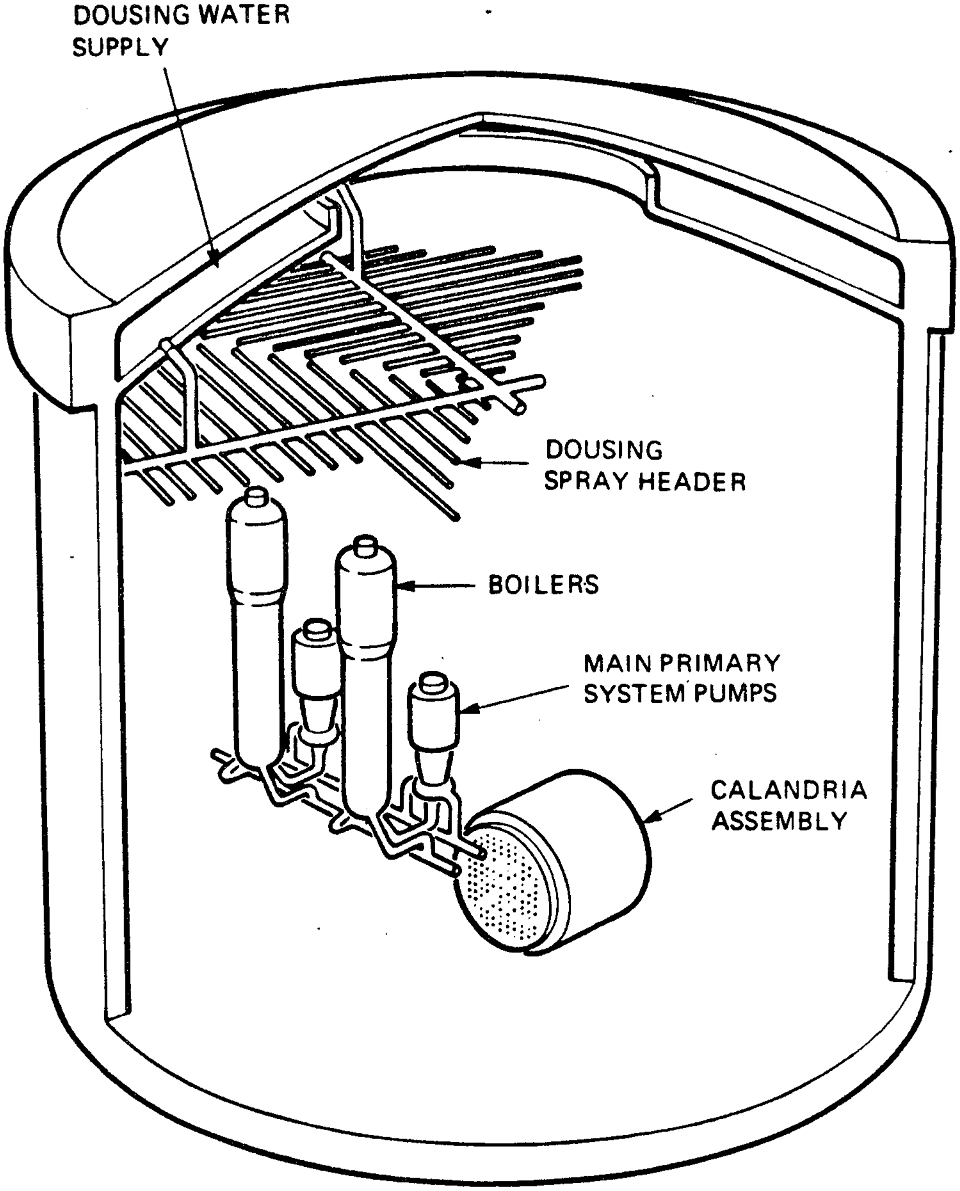

FIGURE 2.5 SINGLE UNIT CONTAINMENT

MULTI-UNIT CONTAINMENT

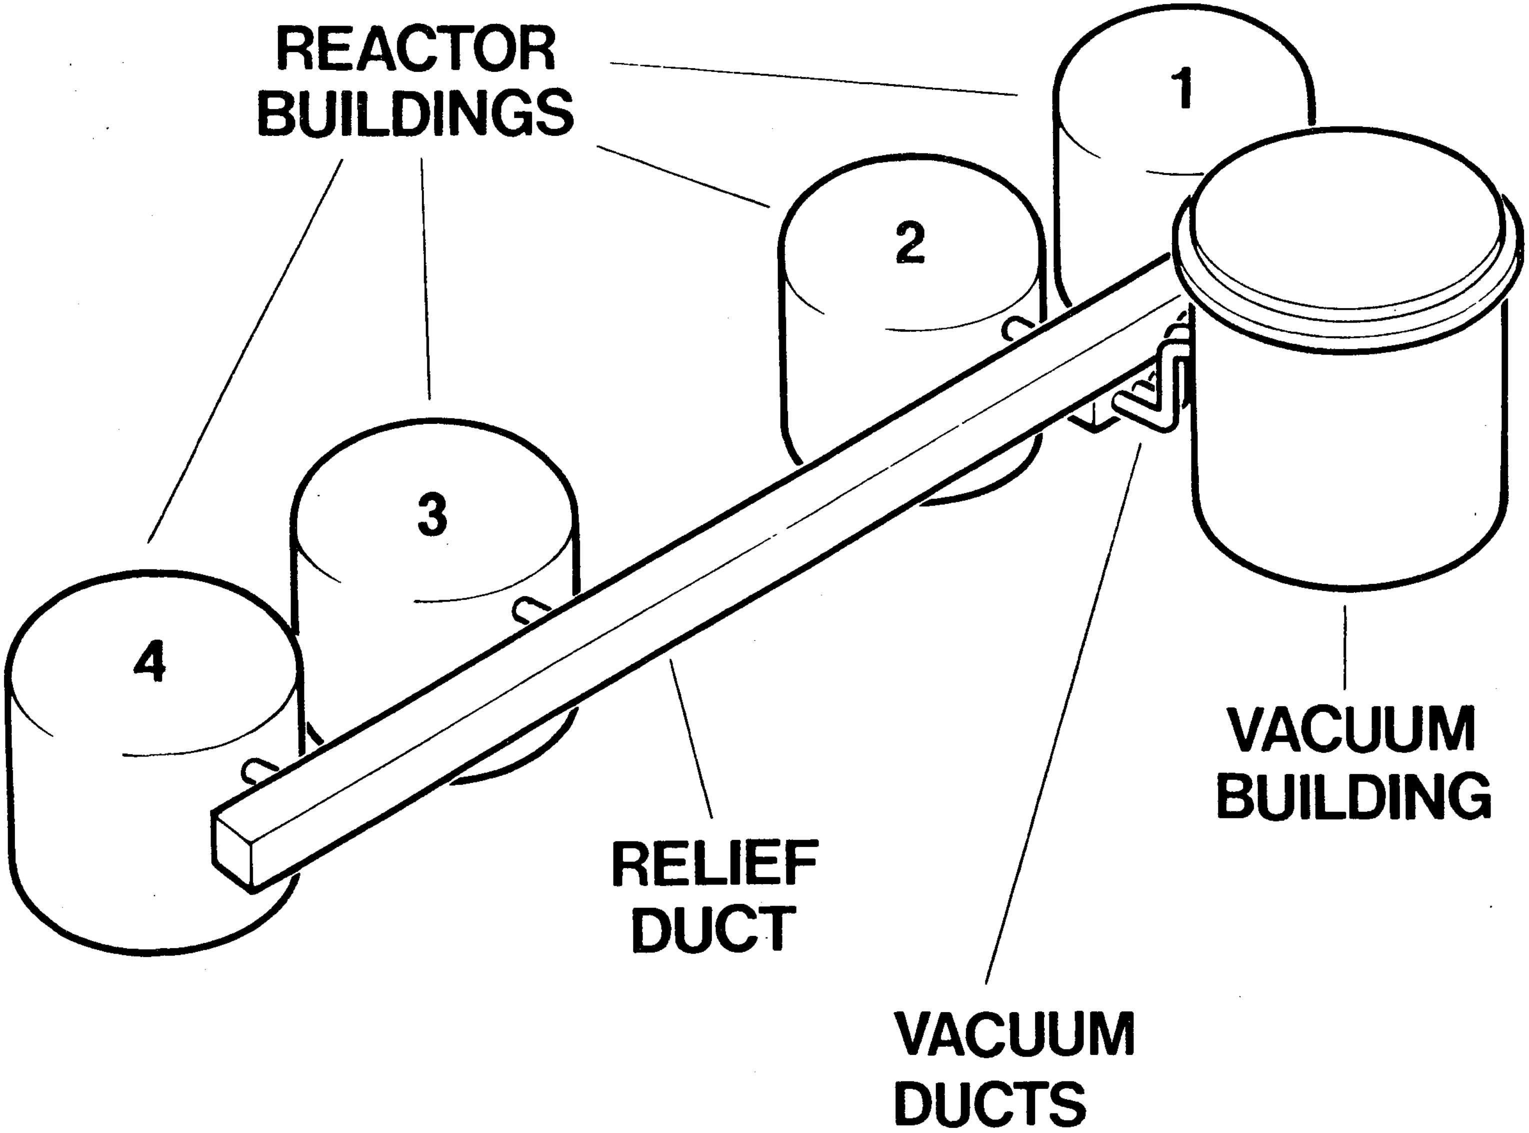

FIG 2.6

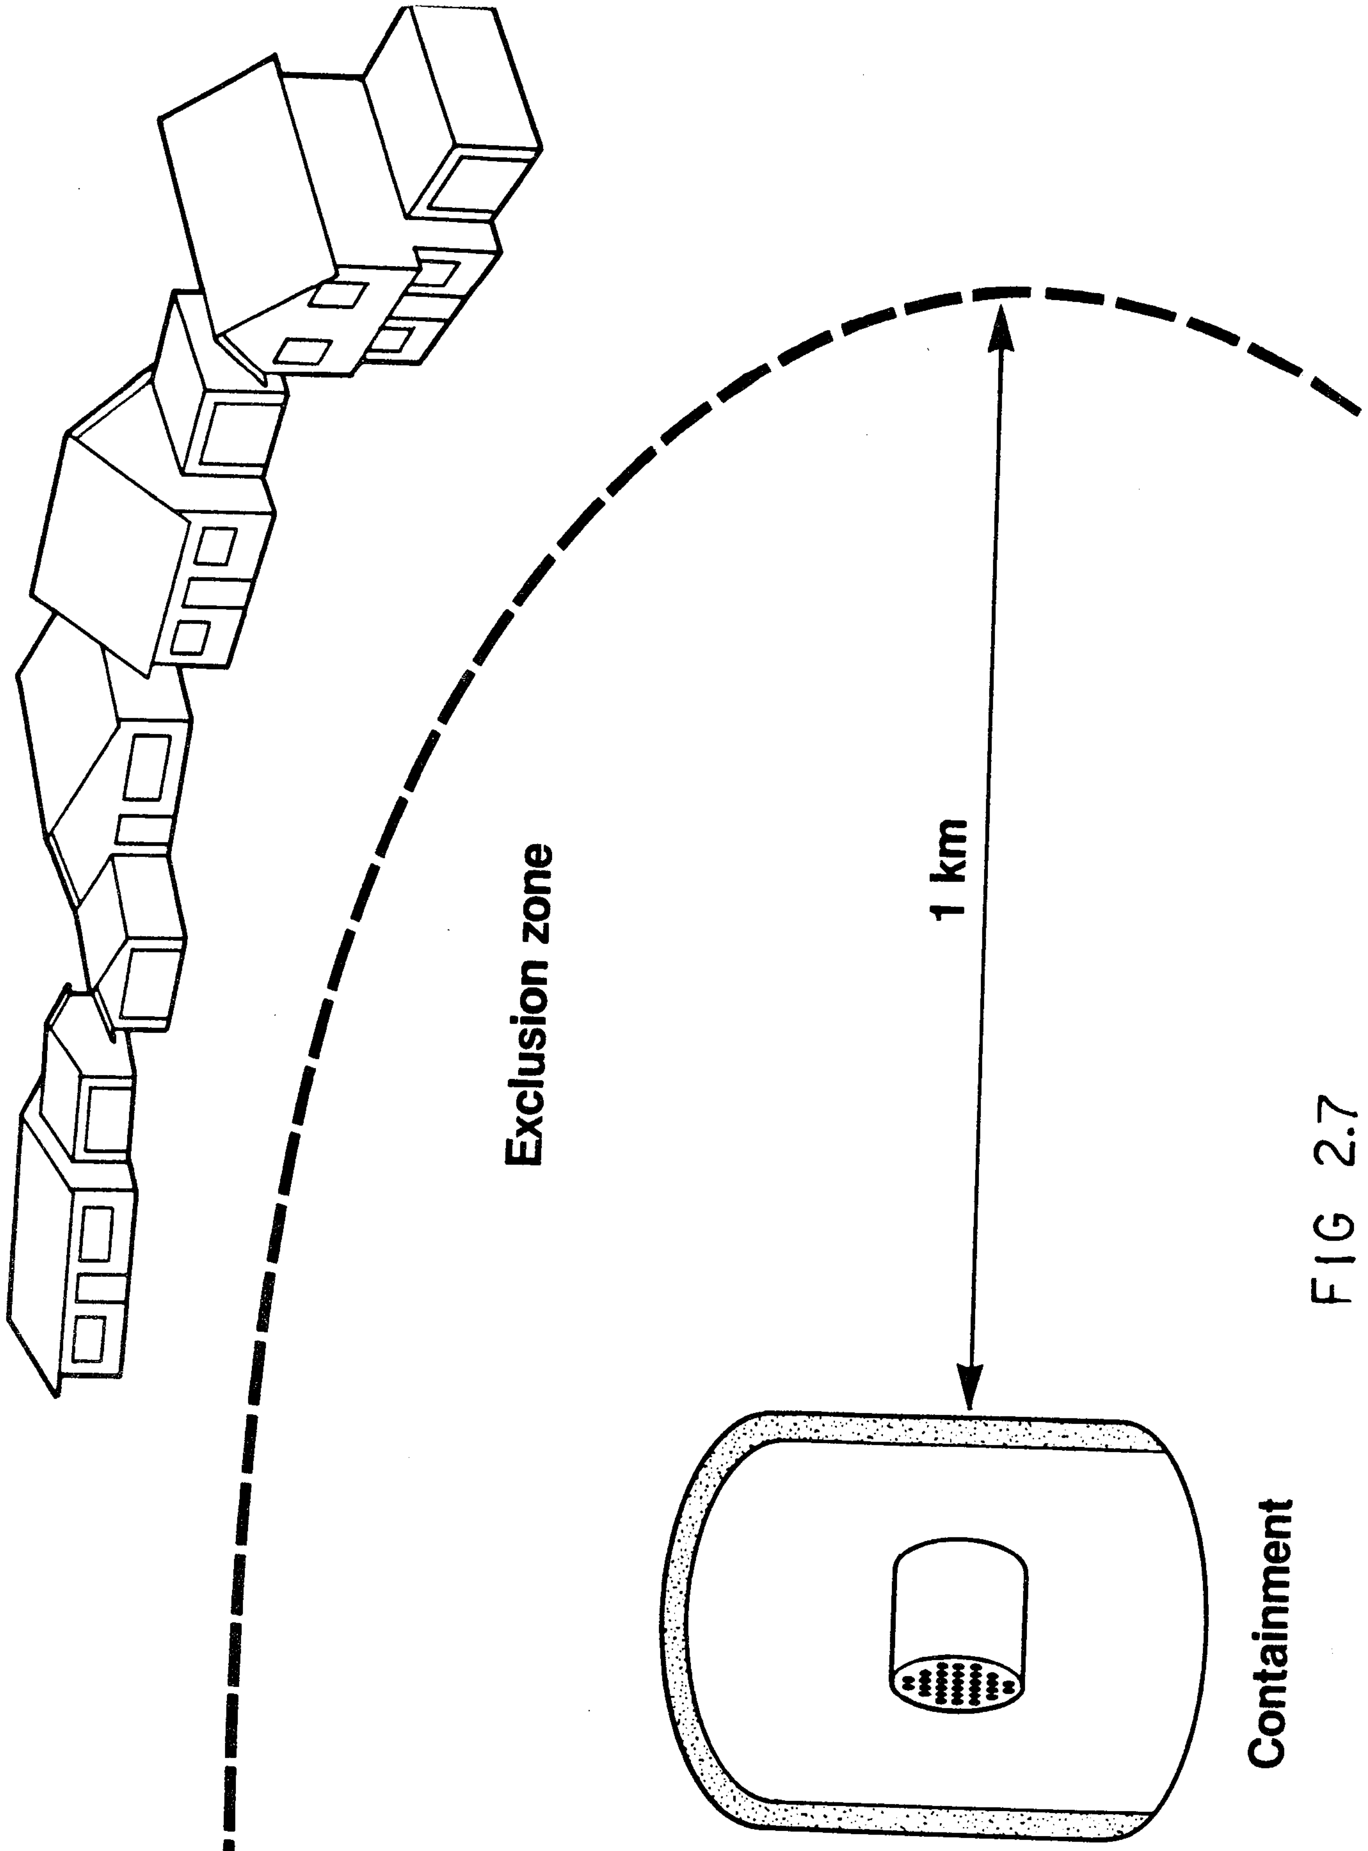

FIG 2.7

SEPARATION OF COOLANT AND MODERATOR

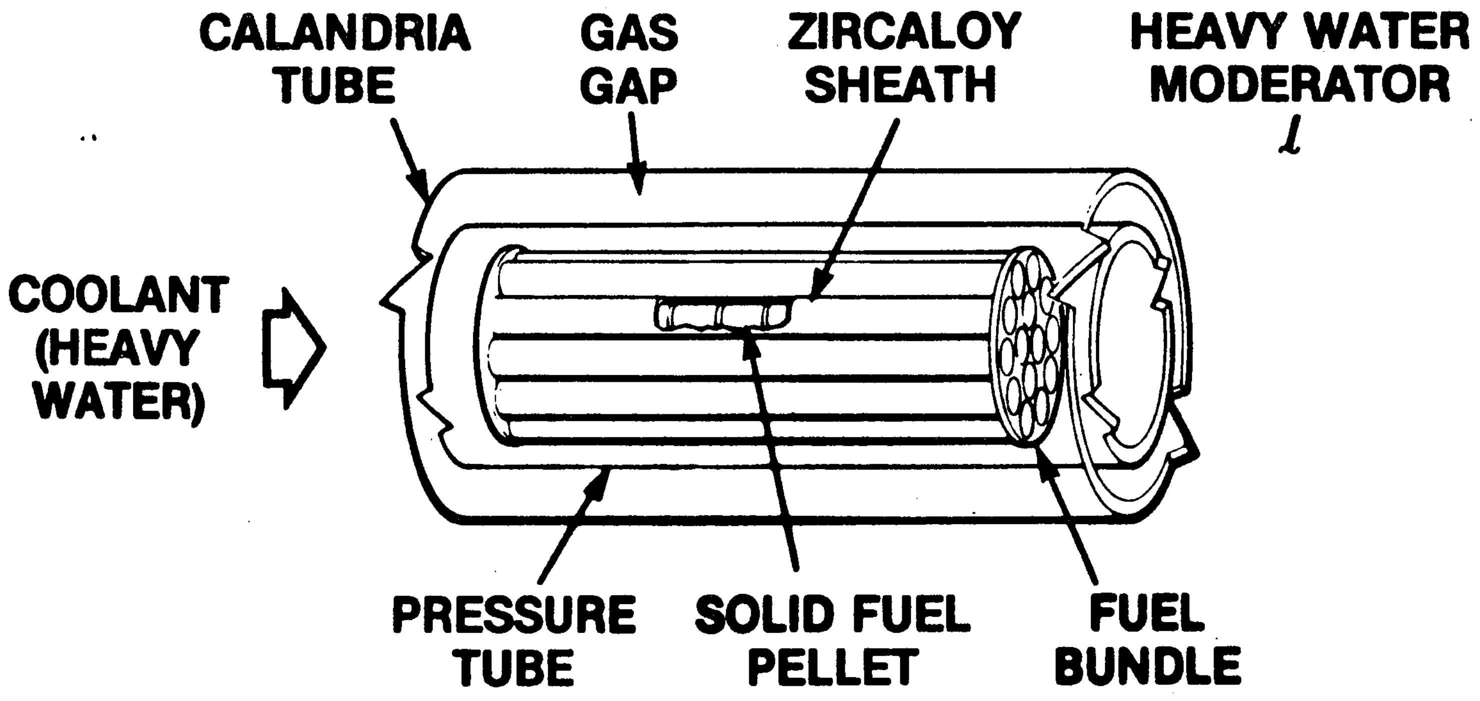

FIG. 2.8

HEAT FLOW PATHS FOR LARGE & SMALL BREAKS

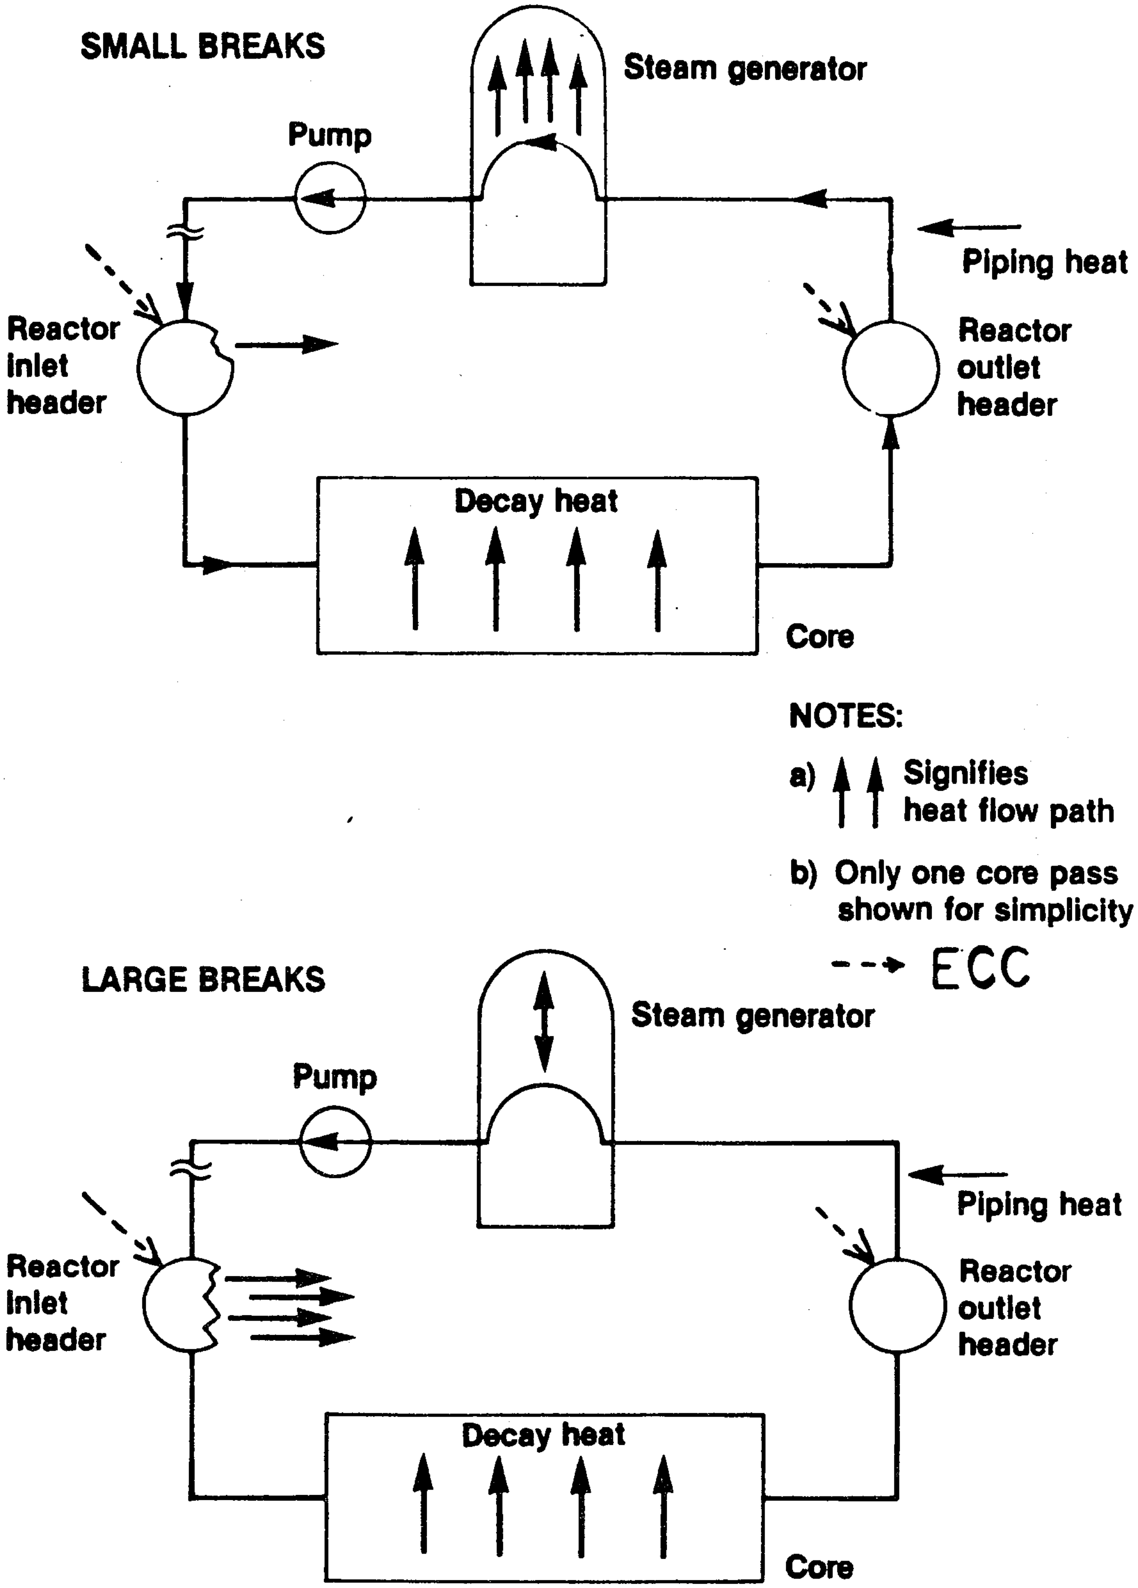

NOTES:

a) ↑↑ Signifies heat flow path

b) Only one core pass shown for simplicity

--→ ECC

FIG. 2.9

TWO GROUP CONCEPT

FUNCTION	GROUP 1	GROUP 2
Shutdown	SDS1	SDS2
Fuel cooling	Normal electrical and water supplies	Emergency power and water supplies
Plant monitoring	Main control room	Secondary control area

FIGURE 2.10

SINGLE AND DUAL FAILURE CONCEPT

SYSTEM CATEGORIES

PROCESS SYSTEMS

- Heat transport
- Reactor control
- Electrical
- Fuel and fuel handling

SAFETY SYSTEMS

- SDS1
- SDS2
- ECCS
- Containment

* Single failure ▶ process failure only

* Dual failure ▶ process and one safety system

FIG 2.11

AECB GUIDELINES FOR ACCIDENTS

EVENT	FREQUENCY	DOSE LIMIT	
		INDIVIDUAL	POPULATION
SINGLE FAILURE	1:3 yr	0.5 rem (body) 3 rem (thyroid)	10^4 man-rem 10^4 thyroid-rem
DUAL FAILURE	1:3000 yr	25 rem (body) 250 rem (thyroid)	10^6 man-rem 10^6 thyroid-rem

FIG 2.12

CANADIAN CRITERIA FOR SINGLE AND DUAL FAILURES

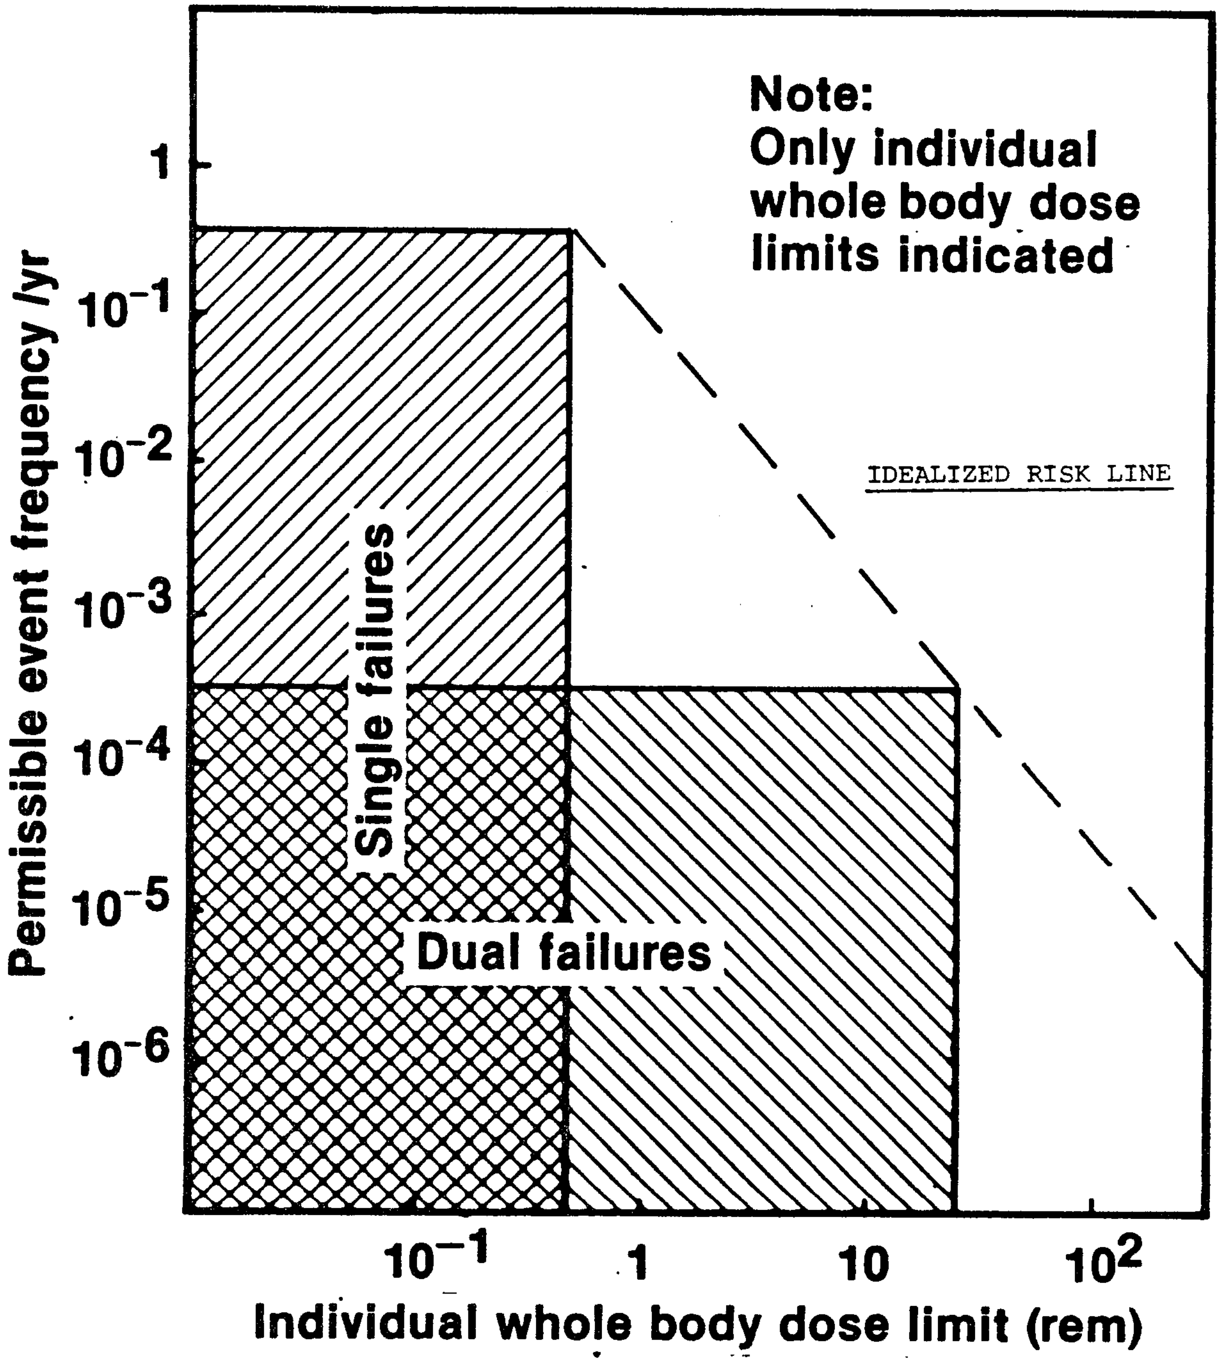

FIG 2.13