Core-Physics Aspects of Safety Analysis

by

B. Rouben
Manager, Reactor Core Physics Branch
AECL

Presented at Chulalongkorn University
Bangkok, Thailand, 1997 December

Abstract

To illustrate the physics input to reactor safety analysis, a discussion of the physics analysis for various accident scenarios is presented. One of the most-important design-basis accidents considered in the safety analysis for CANDU is the large loss-of-coolant accident (LOCA). The basic role of reactor physics is to examine the neutronics of the core for all postulated events, and provide the evolution of fundamental quantities such as the core reactivity and the 3-dimensional neutron flux and power distributions. The discussion presented here illustrates the analysis of large and small LOCAs and of the Regional Overpower Protection (ROP) system design. It covers the major physics considerations, methods, and models.

Table of Contents

Figures

1. Introduction

Reactor physics provides essential input into any reactor safety analysis. Essentially, the safety analysis postulates various types of accidents and aims to demonstrate the following:

- the accident is quickly terminated,
- the reactor is brought to a safe state and remains safe,
- the power is reduced to small values everywhere in the reactor,
- the heat added to the fuel does not result in unacceptable consequences.

Some accident consequences which the safety analysis would like to show are avoided are fuel break-up (fragmentation), fuel centreline melting, changes in the fuel-bundle geometry which would restrict coolant flow (e.g., ballooning of fuel sheath to contact with the pressure tube), and compromise of fuel-channel integrity. The ultimate criteria to be met are those which limit the radiation exposure of the public.

The basic role of reactor physics is to examine the neutronics of the core for the situation of interest, determine the neutron balance, show that the fission chain reaction is shut down and remains shut down, and calculate the rates of the various neutron reactions in space and time through the postulated accident. The fundamental quantities in which the reactor physicist id interested are the following:

- the core reactivity,
- the neutron flux distribution,
- the reactor bulk power and the spatial power distribution, and
- the energy added to the fuel (i.e., the fuel enthalpy).

The present discussion will describe how the physics analysis is carried out for typical accident scenarios for CANDU reactors. One of the particular cases which we shall consider is the postulated large-Loss-of-Coolant accident (LOCA), an important design-basis event. The LOCA analysis illustrates well the important parameters which the CANDU physicist must always consider, and the methods and models in current use. The discussion will also cover the major conservative assumptions used in the analysis.

To perform credible reactor-physics analysis, it is of course essential to first have tools (computer programs) which properly capture and model the phenomena at play. In particular, for safety analysis, the computer programs should be able to reliably calculate the core parameters of importance, those listed above. It is also crucial to be able to understand and model the action of the reactor protective systems, since it is their mandate to recognize accident situations, actuate the shutdown system(s), and terminate the event.

The physics analysis cannot be done in isolation from the other components of the safety analysis. Reactor physics provides input to the thermalhydraulics, fuel, fuel-channel, and radiation-dose analyses. In turn, any data which has an impact on the geometry, configuration, or physical properties of the nuclear lattice should be input into the physics analysis. An important example is thermalhydraulics data, such as coolant density and temperature.

2. CANDU Shutdown Systems

The CANDU reactor is equipped with two independent shutdown systems, SDS-1 and SDS-2. By design, these are physically, logically, and functionally separate. The CANDU-6 systems will be presented here, although they are very similar in other CANDU designs. The required differences between SDS-1 and SDS-2 are achieved by using vertically oriented mechanical shutoff rods in one system and horizontally oriented liquid poison injection nozzles in the second system.

According to system-separation criteria, each shutdown system is to be fully capable, acting on its own, to shut the reactor down from any postulated accident condition. Thus the full safety analysis must demonstrate the capability of each shutdown system.

Each shutdown system must have an availability factor of 0.999, to be demonstrated by periodic testing. The presence of two shutdown systems, each with such availability, means that, if both CANDU shutdown systems are called upon, there is an incredible probability of non-shutdown (10^{-6}).

Each shutdown system can be actuated by a number of means. Two such means, which rely on neutronic parameters, are out-of-core ion chambers and in-core detectors.

2.1 Shutoff Rods (SDS1)

The shutoff rods are tubes consisting of a cadmium sheet sandwiched between two concentric steel cylinders. The rods are inserted vertically into perforated circular guide tubes which are permanently fixed in the core. The locations of these rods in the CANDU 6 are shown in Figure 2.1. The diameter of the rods is the maximum that can be physically accommodated in the space between the calandria tubes (about 113 mm), when space for the guide tubes and appropriate clearances are considered. The outermost four rods are about 4.4 m long, while the rest are about 5.4 m long. The rods are normally fully withdrawn from the core and are held in position by an electromagnetic clutch. When a signal for shutdown is received, the clutch releases and the rods are initially accelerated by a spring and then fall by gravity into the core.

2.2 Liquid-Poison Injection System (SDS2)

This consists of a system of high-speed injection of a solution of gadolinium in heavy water into the calandria. This is accomplished by opening high-speed valves which are normally closed and retain the solution at high pressure in a vessel outside of the calandria. When the valves open, the liquid poison is injected into the reactor moderator through six horizontally oriented nozzles that span the core and are located in positions shown in Figure 2.2. The nozzles are designed to inject the poison in four different directions in the form of a large number of individual jets. This disperses the poison rapidly throughout a large fraction of the core. The gadolinium solution is held in the pressure vessel at a concentration, typically, of about 8000 g of gadolinium per Mg of heavy water.

3. Neutronic Protection Systems

CANDU reactors are equipped with protection systems which detect an emergency situation and actuate the safety system(s) discussed in the previous Section. The CANDU 6 neutronic protection systems are described here. It is important to note however that protection is also provided by monitoring of non-neutronic, i.e., process, parameters.

The neutronic protection system is in fact a double system. That is, there is a **separate** neutronic protection system for each of the two shutdown systems. Each protection system is **triplicated** and consists of out-of-core ion chambers and self-powered in-core detectors. Triplication means that there are three separate "logic" channels for each protection system. These channels are labelled D, E, and F for SDS-1 and G, H, and J for SDS-2.

There are three ion chambers in each protection system, one per logic channel. They are located at the outside surface of the calandria (see Figure 3.1). Each ion chamber is designed to "trip" its logic channel when the measured rate of change of flux (ϕ), or more precisely the quantity

$$\frac{d\,\ell n\,\phi}{dt}$$

exceeds a pre-determined setpoint (e.g., in the CANDU 6, 10% per second [$0.10\ s^{-1}$] for SDS1, and $0.25\ s^{-1}$ or $0.15\ s^{-1}$ for SDS-2).

There are also a number of fast-responding (platinum or Inconel) in-core detectors in each protection system. The system of in-core detectors is designed to identify and protect against high local or regional flux, and is therefore called the regional-overpower-protection (ROP) system. The CANDU 6 reactor has 34 SDS-1 detectors, arrayed in vertical assemblies, and 24 SDS-2 detectors, arrayed in horizontal assemblies. Figures 3.2 and 3.3 show the location of some of these in-core detectors. The detectors are assigned variously to the individual logic channels, so that channels D, E and F contain 11 or 12 detectors each, while channels G, H, and J contain eight each. The detectors trip their logic channels on high neutron flux: when the reading of any one detector reaches a pre-determined setpoint, the logic channel to which it is connected is tripped.

The triplicated logic which governs the "tripping" (actuation) of the shutdown systems is shown schematically in Figure 3.4. It is designed as follows:

- In a given logic channel of either system, if the ion chamber or any in-core detector in that logic channel reaches its trip setpoint, the channel is tripped;
- For either shutdown system, when any 2 of the 3 corresponding logic channels are tripped, the shutdown system is actuated.

This logic achieves simultaneously several objectives:

- provide very reliable protection in a genuinely abnormal situation,
- allow individual logic channels to be temporarily "removed from the system" (taken out of action) for testing, and
- minimize spurious trips.

Note that, although the tripping of only 2 out of 3 logic channels is sufficient to actuate the corresponding shutdown system conservatism, the safety analysis normally assumes, for conservatism, that one of the triplicated logic channels is in testing and is therefore unavailable. The most effective logic channel, i.e., the earliest to trip, is ignored, so that each shutdown system is assumed to be actuated only when the corresponding two least effective logic channels trip.

4. Physics Analysis for Regional Overpower Protection (ROP) against Loss of Regulation

A loss of regulation (LOR) is an event in which the Reactor Regulating System (RRS) loses control of the global or local power. If a power rise - even a slow loss of regulation - is left unchecked, it may lead to damaging overpowers in the fuel, at least in some locations. The CANDU ROP system is provided to address such scenarios. Since the loss of regulation may lead to an unacceptable increase in power in any part of the core, the ROP system is designed so that no fuel channel reaches its "critical" channel power. The critical channel power is currently defined as the power which leads to onset of fuel dryout.

The role of the ROP analysis is then to determine appropriate trip setpoints for the in-core detectors such that, for anticipated flux shapes in the reactor, under both normal and off-normal operation, fuel dryout will be avoided if there is an increase in the local or global power.

Note nonetheless that the in-core ROP system is useful not only in loss-of-regulation situations. It provides the welcome possibility of neutronic trips from in-core signals in any situation leading to a global or local power transient, such as a loss-of-coolant accident (LOCA), examined in the next Section.

The ROP analysis is based on a large number - hundreds - of calculated flux shapes. These hundreds of flux shapes are selected to span a wide range of flux distributions, corresponding to:

- different possible reactivity-device positions,

- power manoeuvres,
- power recovery following a reactor shutdown,
- xenon transients and oscillations,
- possible scenarios initiating a loss of regulation, such as the accidental draining of one or more zone-control compartments, etc.

The role of reactor physics in the ROP analysis is to calculate these flux distributions and provide to the ROP code, for each flux shape, the complete 3-dimensional flux and power distributions. The computer program which designs the ROP system must then use the totality of this data to

- select appropriate detector positions within the detector assemblies,
- assign detectors to the various logic channels, and
- define detector setpoints.

The basic idea or principle to follow so as to achieve a good ROP-system design is to place detectors so that they "see" perturbations very effectively. There must be (at least) one detector in each logic channel to protect against an LOR from that perturbed flux shape. That is, the detector must be placed so that it would reach its setpoint (to be chosen judiciously) before any channel reaches its critical power.

The detector positions, channelization, and setpoints must, consistent with the triplicated logic described in the previous Section, protect the reactor against a loss of regulation while providing sufficient operating margin to allow normal operation without undue spurious trips. Of course, the challenge is to design an effective ROP system with a reasonably small number of in-core detectors.

The flux shapes needed for the ROP analysis are calculated with the finite-core code RFSP (Reactor Fuelling Simulation Program). These calculations start with the time-average core configuration as basis, as it is not practical to compute hundreds of flux shapes for each snapshot in a reactor's operating history. The time-average picture, however, does not include the flux "ripple" due to daily refuelling, a measure of which is the Channel Power Peaking Factor (CPPF), defined as the highest value of the ratio of instantaneous (snapshot) to time-average channel power (in a defined, high-power region of the core). To compensate for the absence of fuelling ripple in the calculated ROP shapes, the in-core detectors are calibrated daily (at site) upwards from their full-power reading, by a factor equal to the CPPF. For instance, if the detector trip setpoint is (arbitrary numbers, for illustration only) 1.22 and the current value of CPPF is 1.08, then the effective setpoint (i.e., the effective margin to trip) is 1.22/1.08 = 1.13.

5. Neutron Kinetics

Fast neutronic transients, where large changes in power occur over intervals of seconds, are analyzed using neutron-kinetics computer codes. This section describes features of importance in neutron kinetics.

(Note: Simulations over intervals of seconds need not address $^{135}I/^{135}Xe$ kinetics, which becomes important on a time scale of minutes or hours. These effects must, however, be taken into account in other types of analyses, such as those for slow transients, which may be followed for several minutes or longer.)

The driving term in neutron kinetics is the system reactivity ρ, a measure of the imbalance between the rates of neutron production and loss (absorption or leakage). If we start from the reactor multiplication constant

$$k_{eff} = \frac{\text{Rate of neutron production}}{\text{Rate of neutron loss (by absorption and leakage)}} \tag{1a}$$

the system reactivity ρ is defined by

$$\rho = 1 - \frac{1}{k_{eff}} \tag{1b}$$

A zero value of ρ (i.e., $k_{eff} = 1$) denotes a critical reactor. The fission chain reaction is just self-sustaining, and if the reactivity remains nil, the neutron population will be steady in time. A positive value of ρ denotes a supercritical reactor (with an increasing neutron population), and a negative ρ denotes a subcritical reactor (with a decreasing neutron population). Changes in the reactor power are in the same direction as changes in neutron population.

While reactivity (as also k_{eff}) is a pure number and therefore has no real units, values of reactivity are generally small (much smaller than unity), and a unit often used (at least in Canada) for reactivity is the milli-k:

$$1 \text{ milli-k} = 0.001$$

A reactivity of 1 milli-k may seem small, but its effect may not be inconsequential, depending on other nuclear characteristics of the core.

Any imbalance between neutron production and loss (i.e., a non-zero value of reactivity) causes the neutron population to increase (or decrease) from one generation to the next. It is therefore natural to think that the *rate* at which the neutron population (and, consequently, the power) will change will depend on the mean generation time T, the average time interval between successive neutron generations. In a simplistic treatment of kinetics, in fact, the power varies exponentially with reactivity and with time in units of T:

$$P = P_0 \exp\left(\frac{\rho t}{T}\right) \tag{2}$$

However, this treatment is much too simplistic because it does not account for the important fact that some neutrons are delayed. Neutrons produced in fission are either prompt or delayed. Prompt neutrons are produced essentially at the same instant as fission. If all fission neutrons were prompt, the mean generation time T would be identical to the ***neutron lifetime*** Λ (more frequently known as the prompt-neutron lifetime, although it applies to all fission neutrons), defined as the average time interval between the birth of a neutron in the reactor and its absorption in a subsequent fission reaction. In the CANDU lattice, Λ has a value of approximately 0.9 millisecond. In LWRs, Λ is about 30 times shorter.

With a mean generation time of 0.9 ms, a reactivity of 1 milli-k would lead (via Eq. 2) to a power increase by a factor of 3 in 1 second, a factor of 9 in 2 s, etc. Such a fast rate of change would be extremely difficult to control. (In LWRs, the rate of change of power would be even 30 times as great for the same reactivity.)

Delayed neutrons, however, reduce the rate of power change considerably. This is so even though the number of delayed neutrons is a very small fraction (~0.6%) of all neutrons from fission. Most delayed neutrons are produced in the beta decay of fission products, the delayed-neutron precursors. For practical purposes, delayed-neutron precursors from any one fissionable nuclide can be subdivided into six distinct groups, with beta half-lives ranging from 0.2 s to 50 s. Table 5.1 shows the typical decay constants and half-lives of the delayed-neutron precursors. Table 5.2 shows the fractions of neutrons that are delayed, when considering fissions from the various fissile nuclides. The total number of delayed neutrons is lower in plutonium than in uranium, so that the delayed-neutron fraction decreases with fuel irradiation.

In CANDU, there are additional delayed neutrons, produced by the photodisintegration of the deuterium in heavy water. These photoneutrons appear from another nine distinct precursor groups, with even longer time constants, in the hundreds to tens of thousands of seconds. Table 5.3 shows the time constants and delayed fractions for the photoneutron groups.

When analyzing transients over a time scale of a few seconds, all delayed-neutron groups, including the photoneutron groups, are "collapsed" onto a set of 6 "effective" groups. The delayed-neutron data for these collapsed groups depend on the exact method used for the collapsing, and on the average fuel irradiation (burnup). Table 5.4 shows typical delayed-neutron constants for the collapsed groups for a CANDU equilibrium core.

These delayed-neutron and photoneutron time constants are such that, in spite of the small delayed fraction, we can see that the "effective" (weighted-average) mean generation time would be much longer than the prompt-neutron lifetime. In fact, if we calculate the weighted-average mean generation time for CANDU, we find a value of the order of a tenth of a second, about 100 times the ~1-ms prompt-neutron lifetime. From Equation 2, we can see that a reactivity of 1 milli-k would then lead to an increase in power of a factor of only about 1.01 s^{-1}, compared to 3 s^{-1} without delayed neutrons. Figure 5.1 illustrates schematically the large influence of delayed neutrons on the evolution of power in transients. It is clear as well, from the above, that delayed neutrons facilitate reactor control considerably.

The positive influence of delayed neutrons on reactor transients disappears however in the regime where the chain reaction can be sustained by prompt neutrons alone, i.e. when $\rho > \beta$. When this is the case, the reactor is said to be prompt supercritical, and delayed neutrons can for all practical purposes be neglected. In this regime, the generation time becomes equal to the prompt-neutron lifetime Λ, and power evolves on the correspondingly much faster time scale.

In summary, in the regime of reactivity far below prompt criticality, the mean generation time is governed by the delayed-neutron-precursor lifetimes and is not sensitive to the value of the prompt-neutron lifetime. But for transients close to or above the prompt-critical regime, the value of the neutron lifetime Λ governs the response rate of the reactor to the perturbation. This is where transients evolve much more slowly in CANDU than in light-water reactors. Figure 5.2 illustrates this by showing the sensitivity of the "reactor period" (the time scale for the increase in power by a factor e) to reactivity for various prompt-neutron lifetimes.

6. Kinetics Methods

For a proper treatment of delayed-neutron effects in fast transients, it is neither accurate nor sufficient to simply define a mean generation time and use Equation 2. A rigorous approach demands a proper mathematical treatment of the rates of production of prompt neutrons and of the production and decay of the various delayed-neutron precursors. This leads to a set of differential equations coupling the neutron flux ϕ and the delayed-neutron-precursor concentrations c_i:

$$\frac{d\phi}{dt} = \frac{\rho - \beta}{\Lambda}\phi(t) + \sum_{i=1}^{N} \lambda_i c_i(t)$$

$$\frac{dc_i}{dt} = \frac{\beta_i}{\Lambda}\phi(t) - \lambda_i c_i(t) \qquad i = 1, \ldots, N \tag{3}$$

where

N is the number of delayed-neutron-precursor groups,

β_i and λ_i are the partial delayed neutron fractions and decay constants for the various precursor groups, and

β is the total delayed-neutron fraction (the sum of the partial delayed fractions).

6.1 Point Kinetics

In point kinetics, the reactor core is treated as a single point. The premise is that, predominantly, only the spatially uniform component of the power change need be examined and that spatial variations of the response can be ignored. Core-average values would then be used for the parameters in the set of Equations 3. The method determines the time variation of the global (average) values of power and delayed-neutron-precursor concentrations. In this approximation, this variation is then superimposed on the pre-event power shape.

There are many self-standing point-kinetics codes. At AECL, the program PTK is currently used. Also, a number of thermalhydraulics codes incorporate subroutines to solve the point-kinetics equations. Examples of such thermalhydraulics codes with point-kinetics capability are SOPHT, FIREBIRD, and CATHENA.

6.2 Inadequacies of Point Kinetics

The point-kinetics approximation is largely inadequate for the purposes of modern safety analysis. Its weakness originates in the presence of spatially non-homogeneous effects. The following subsections illustrate two very important sources of inhomogeneity.

6.2.1 Voiding Transient

The coolant voiding in a large LOCA is certainly not uniform. For instance, in the CANDU 6, the heat transport system is subdivided into two side-by-side loops (see Figure 6.1), each servicing one half of the cylindrical reactor. The two loops are isolated from one another in a large LOCA. Thus, the break will induce a side-to-side asymmetry in the core coolant density, leading to a side-to-side asymmetry in the ensuing power pulse.

It is also to be noted that a postulated pre-accident side-to-side asymmetry in the power distribution, with the tilt in the same direction as that due to the void asymmetry, will further accentuate the non-uniformity in the power pulse (see Section 7.3.2).

These effects cannot be adequately represented by point kinetics.

6.2.2 Shutdown-System Coverage

The safety analysis is done with conservative assumptions. In particular, it is assumed that the shutdown system is not fully operational, i.e., that some of its components do not act. For instance, two of the shutoff rods in SDS-1 (e.g., 2 of 28 in the CANDU 6) are assumed not to be functioning (see Section 7.7). This results in non-uniformity in the shutdown-system spatial coverage. Once again, point kinetics cannot hope to properly deal with this situation.

6.3 Spatial Kinetics

The previous section illustrates the reasons why the capability to model spatial effects is important. This has led to the development of spatial-kinetics methods, which are used for the detailed analysis of fast transients over time scale of a few seconds. Point-kinetics methods are still useful as a means to continue the analysis over very long times, for instance to provide input to long-term thermalhydraulics simulations.

The development of spatial-kinetics codes for CANDU has followed two different routes, the nodal and modal methods. These have resulted in the codes CERBERUS and SMOKIN respectively.

6.3.1 Improved Quasi-Static Method (CERBERUS)

CERBERUS solves the time-dependent neutron-diffusion equation in three spatial dimensions and two neutron energy groups (fast and thermal). The methodology includes the space-and-time-dependent delayed-neutron precursors. The equation is solved in its finite-difference form, using a finite-difference nodal model (see Figures 6.2 and 6.3) with tens of thousands of mesh points. The flux is written as the product of a space-independent amplitude and a space-and-time-dependent flux-shape function:

$$\Phi(\vec{r},t) = A(t)\Psi(\vec{r},t) \tag{4}$$

The idea of the Improved Quasi-Static (IQS) method is to cast the major time dependence into the amplitude. This is achieved simply by constraining a core integral of the flux to be constant in time [all core integrals in the IQS method use the steady-state adjoint flux as a weighting function]. Equations are then derived for the amplitude, the flux shape, and the precursor concentrations.

The equation for the amplitude $A(t)$ is space-independent, but is coupled to equations involving integrals of the precursor concentrations. This set of coupled equations has **exactly** the same form as the point-kinetics equations, except that the parameters within the equations [e.g., reactivity ρ, prompt-neutron lifetime Λ, and effective total delayed fraction β, are all evaluated as core-integrated quantities. It is here that the choice for the adjoint (also known as the importance function) as the weighting function in the integrals is recognized as crucial. This is so because the adjoint has the property that it makes the all-important value of reactivity a stationary function of the flux shape, thus minimizing its sensitivity to possible numerical (e.g., round-off) errors in calculating the flux.

Since they are not functions of space, the point-kinetics-like equations in the amplitude and integrated precursor concentrations can be solved with little numerical effort.

On the other hand, the equation that is obtained for the flux-shape function $\Psi(r,t)$ is similar to the time-independent diffusion equation, except that it has additional terms in the amplitude and the precursor spatial concentrations. Because of the number of unknowns (tens of thousands of flux values), this equation must be solved using iterative methods. Consequently, this part of the calculation is by far the most time consuming. However, the choice of the form of solution permits the use of relatively large time steps between flux-shape calculations, while retaining good accuracy.

Most recently, the CERBERUS methodology has been integrated within the major core-physics code, RFSP, so that the most sophisticated, best-validated core models can now be used in the analysis of fast transients.

6.3.2 Modal Method (SMOKIN)

In SMOKIN, the one-energy-group (thermal) neutron flux is expanded in a finite series of pre-calculated flux modes. These modes are normally the flux "harmonics" of the time-independent neutron diffusion equation. Physically, they represent the three-dimensional "global" flux shapes which are expected to be most "excited" (or promoted) by core perturbations. The flux harmonics are calculated by a code such as RFSP. Some 10 to 20 harmonics can practically be computed. Examples of the harmonic flux shapes obtained are sketched in Figure 6.4. They represent azimuthal (e.g., side-to-side), radial, and axial flux perturbations of various orders. Note that these modes reflect global shape changes and cannot represent fine details or very localized perturbations.

In SMOKIN, the mode amplitudes are the unknown variables of the problem. The time-dependent neutron diffusion equation can be re-cast in terms of a small number of linear differential equations in the mode amplitudes. The equations incorporate terms which are various core integrals of products of the harmonics and the nuclear cross sections. These modal "weights" can be pre-calculated since the modes are fixed. The numerical solution of the equations can thus be computed very quickly.

6.3.3 Pros and Cons

Few would argue with the assertion that the Improved Quasi-Static method, based on the finite-difference diffusion-theory methodology, has a more rigorous basis and contains fewer approximations than the modal method. Thus the consensus is that of the two codes, CERBERUS is the more accurate. On the other hand, the great advantage of SMOKIN is the relatively small numerical effort required to analyze even long transients; safety analysis and scoping studies can therefore be performed in a shorter time frame.

Another very useful feature of SMOKIN is that it can model the response of the Reactor Regulating System (RRS) to core perturbations, so that reactivity-device movements, in particular, can be predicted. Modelling of the RRS, using the IQS methodology, has recently been added to RFSP as well; however SMOKIN retains the advantage of computational speed.

SMOKIN has not been in great favour at AECL. It has seen extensive use mostly by analysts at Ontario Hydro, the electric utility.

6.4 Kinetics Effects on Flux Shape and Reactivity

Delayed-neutron precursors have a marked influence on the reactivity worth of shutdown systems because of their flux-shape-change retardation effect. Consider an event in which a shutdown system has been actuated. In the pre-event steady state, the shape of the delayed-neutron-precursor distribution is the same as that of the fission (prompt-neutron) source (curve labelled 4 in Figure 6.5).

On shutdown-system actuation, the system reactivity becomes negative, i.e., the system is subcritical, and the prompt source drops quickly in regions covered by the shutdown system. If there were no delayed neutrons, the flux shape would be close to that calculated in a hypothetical steady state (critical reactor) with the shutdown system inserted. The flux values would be very low in regions directly covered by the shutdown system (curve labelled 1 & 3 in Figure 6.5). However, the presence of the delayed source results in an overall flux shape which is in fact less depressed in the shutdown-system regions (curve labelled 2 in Figure 6.5). This has the effect of increasing the neutronic importance, and therefore the reactivity worth, of the shutdown system (in other words, the delayed source in the region of the shutdown system increases the effectiveness of the latter).

A typical example: the "static" reactivity worth of 26 (out of 28) shutoff rods in the CANDU 6, calculated with the *time-independent* diffusion equation in a *hypothetical, non-physical steady state* with the reactor critical and 26 rods inserted, is approximately -55 milli-k. The "dynamic" reactivity worth of the same rods is approximately -80 milli-k; it is calculated with the *time-dependent* diffusion equation and the proper *time-dependent flux shape* appropriate to the scenario considered.

7. Schematic of a Physics Analysis for a Large LOCA

Coolant voiding in CANDU introduces positive reactivity and promotes a power rise. The power pulse arising from a large LOCA is a fast transient. The power rises quickly (time frame of a fraction of a second), shutdown-system action is initiated typically within a second, and the power has been turned around and reduced to small values within a few seconds (if the shutdown-system design is adequate). Thus, a detailed (3-d) physics simulation of a large LOCA usually extends to a few (~5) seconds after the postulated time of the break.

The root cause of the positive void effect in CANDU lies in the pressure-tube configuration, with the coolant separate from the moderator. In light-water reactors (LWR), the same liquid serves as both coolant and moderator, and a loss of coolant is also a loss of moderator, leading to a less self-sustainable chain reaction, i.e., a *decrease* in reactivity. In CANDU, however, the role of the coolant in moderating neutrons is relatively very small. Thus, the loss of coolant does not imply a significant reduction in moderation. On the other hand, the loss of coolant does result in changes in the neutron spectrum (distribution of neutron energy) which:

- go in the direction of reducing the probability of neutron absorption in the fuel resonances, and
- increase the contribution of fast-neutron-induced fission.

These effects are the main components of the reactivity *increase* on coolant voiding. For irradiated fuel, in which plutonium is present, the change in neutron spectrum gives also a *negative* component in the reactivity change, due to a reduction in thermal fission in plutonium, but the *net* reactivity change on coolant voiding is still positive (but smaller than for fresh fuel).

A large loss of coolant is in fact the accident which presents the greatest challenge to CANDU shutdown systems in terms of the rate of positive reactivity insertion. A large LOCA is caused by the rupture of a large pipe such as a Reactor Inlet Header (RIH), Reactor Outlet Header (ROH), or Pump-Suction pipe (see Figure 7.1). Such a rupture has the capability to lead to a sudden power surge (power pulse) beyond the capability of the Reactor Regulating System to control. The manner in which the shutdown systems act (separately) to terminate the power excursion must therefore be carefully studied.

Because of its importance, we will deal mostly with the large LOCA in the next few sections. For other types of accidents, the analysis may use different computer codes and assumptions as appropriate, but it will still be concerned with the same quantities, i.e. reactivity, flux, and power generation.

Figure 7.2 presents a schematic of the physics analysis for a large LOCA. The main steps in the physics analysis are seen to be:

- the simulation of the pre-accident reactor configuration,
- the modelling of the postulated perturbation (in this case the LOCA, modelled in conjunction with a thermalhydraulics calculation),
- the simulation of the early part of the accident, prior to shutdown-system actuation,
- the calculation of the shutdown-system actuation time, sometimes loosely referred to as the reactor "trip time" (the shutdown systems are actuated either by process trips or by neutronic trips. The role of the physics analysis is to determine the actuation time resulting from neutronic signals.),
- the simulation of shutdown-system action,
- the simulation of the combined effects of the perturbation and the mitigating shutdown-system response.

The quantitative results of the analysis, in terms of reactivity, bulk, channel, and bundle powers, integrated powers, and peak fuel enthalpy, are then available to assess the consequences of the accident.

While the positive reactivity theoretically available to be inserted in full-core voiding is of the order of 10-15 milli-k, heat-transport-system subdivision and/or practical considerations limit the amount of coolant loss that is possible in the first few seconds after the break. Thus, a typical large LOCA may insert of the order of 4-5 milli-k of positive reactivity within 0.5 s, while the shutdown-system response will counter with negative reactivity of the order of 50-100 milli-k in 1-2 seconds. Typically, the resulting power pulse will then be as illustrated schematically in Figure 7.3.

8. Analysis Methods, Models, and Assumptions

Many modelling inputs must be assembled for a LOCA calculation. System-parameter values must be chosen. Many conservative assumptions are usually made in a safety analysis. This section presents some of the methods, modelling details and inputs needed.

8.1 Neutronics and Thermalhydraulics Models for LOCA Analysis

With the evolution of computer capacity and performance, and the analyst's desire to better and better capture the physical phenomena at play, the models used for reactor analysis have increased substantially in size and complexity over the last few years.

The evolution of physics models for LOCA analysis has gone hand in hand with that of thermalhydraulics models. Twenty years ago, a LOCA calculation would use a single coolant density transient over the entire broken PHTS loop. Present-day LOCA calculations feature different density transients in different parts of the core. The critical pass in the broken loop is modelled with 5-10 different thermalhydraulics channel groups (see Figure 8.1) representing channels with different conditions, instead of a single "average" group. Non-critical passes, where the voiding is much slower, can be modelled by one channel group.

Large-LOCA analysis is now done with coupled neutronics and thermalhydraulics codes, so that the greatest benefit can be garnered from the evolution of the models. The coupling can be done in either of two ways:

1. Cycle between the thermalhydraulics and neutronics calculations over the entire LOCA simulation interval, starting from a "guessed" power transient, and repeating until convergence is achieved. This method has been used mostly with the SMOKIN/TUF combination.
2. "Walk" through the transient only once, in small time steps (e.g., the CERBERUS flux-shape time step), sequencing the thermalhydraulics and neutronics calculations at each

step. This method has been used mostly with the CERBERUS/FIREBIRD and CERBERUS/CATHENA combinations.

8.2 Coolant-Void Reactivity

Void reactivity enters the calculation via the lattice parameters (nuclear cross sections), computed with a cell code such as POWDERPUFS-V, resident within RFSP. As the LOCA proceeds, these lattice parameters will change in regions of voiding, inserting positive reactivity.

The void reactivity increases as the isotopic purity of the heavy-water coolant decreases, i.e., as the H_2O content of the coolant increases. Consequently, for conservatism, the safety analysis of a LOCA is usually done assuming the minimum coolant purity allowed by the station Operating Policies and Procedures (OP&P). This may vary from plant to plant, but is usually in the range 97-99 atom % D_2O.

In order to take into account a possible underestimation of the positive void effect by the cell code, a further allowance is made in the direction of artificially increasing the void reactivity. This allowance can be introduced by artificially degrading the D_2O coolant purity below the minimum operational value. The full-core void reactivity increases by about 0.6-1 milli-k per percent reduction in the coolant isotopic purity.

8.3 Reactor Pre-Accident Configuration

Since it is required to demonstrate adequacy of the shutdown system(s) under any credible situation, the safety analysis should be done for a variety of pre-accident configurations. In particular, it is important to try to identify and analyze those pre-accident states which can be expected to increase the severity of the power pulse. This section describes two such configurations.

8.3.1 Poison in Moderator

The presence of poison (boron or gadolinium) in the moderator increases void reactivity. The reason is that, on coolant voiding, there is a redistribution of neutron flux in the lattice cell: the flux in the moderator region decreases, causing a reduction in the rate of neutron absorption in the poison. The full-core void reactivity increases by approximately 0.6 milli-k for each additional ppm of boron in the moderator.

Configurations which induce larger void reactivity will thus be those in which the moderator-poison concentration is high, for example

- the "young" reactor core, from initial criticality to first refuelling,
- after a long reactor shutdown, when ^{135}Xe and other saturating fission products have decayed away,.
- in periods of intentional overfuelling in anticipation of planned fuelling-machine maintenance

A critical core following a long shutdown at the plutonium peak will feature the highest poison concentration.

8.3.2 Flux Tilts

Pre-accident flux tilts may increase the power pulse after a LOCA. For instance, a top-to-bottom tilt with the high flux at the bottom would tend to reduce the effectiveness of dropping shutoff rods, since they will take longer to reach the region of higher flux.

Also, in the CANDU 6, a pre-existing side-to-side flux tilt will increase the neutronic importance of the void (and therefore the void reactivity) when the coolant loss is on the high-flux side. In fact, this pre-accident configuration has been found to induce the most severe power pulses in the CANDU 6.

8.3.3 Pressure-Tube Creep

One of the effects of reactor aging is pressure-tube radial creep under neutron bombardment. The increased pressure-tube radius results in a greater volume of coolant in the core. With the greater coolant volume is associated a larger void reactivity. Predicted values of the radial creep, consistent with the actual or anticipated age of the pressure tubes, should therefore be used in analyses of "mature" reactors.

Axial creep also needs to be taken into account in certain reactors. In normal operation the fuel string in a channel is pushed by the force of the flow to the coolant-outlet end of the channel. Elongation of the pressure tube due to axial creep means therefore that the fuel string may be partly out of the core at the channel outlet. In an RIH break, pressure differentials may push the fuel string back into the core. If the refuelling scheme in the reactor is against the direction of coolant flow (such as in the Bruce reactors), the irradiation distribution of bundles in the channels is such that the movement of the fuel string back into the core ("fuel-string relocation") would introduce positive reactivity *in addition to* the void reactivity. This must be taken into consideration in the LOCA analysis. The magnitude of the effect depends on the length of the gap in the channels, which changes with the pressure-tube age (axial creep).

8.4 Fuel-Temperature Reactivity Feedback

As the power rises during the LOCA, fuel temperatures will increase in consequence. Rising fuel temperature will increase neutron absorption due to Doppler broadening of the uranium resonances. This phenomenon adds negative reactivity in the early stage of the LOCA, when temperatures are higher than normal, but can compensate only in part for the void reactivity. Fuel temperatures are typically calculated in the thermalhydraulics code (e.g., CATHENA), and input into the kinetics code along with the coolant-density transient.

8.5 Protection-System and Detector Modelling

The physics simulation determines the actuation time of the shutdown system in response to the LOCA. In order to calculate the actuation time, the code needs:

- the position, channelization, and trip setpoints of the in-core ROP detectors and out-of-core ion chambers,
- the delayed-response characteristics of the in-core detectors,
- the characteristics of the electronics (amplifiers, compensators, etc.) to which the detectors are connected,

8.6 Shutdown-System Configuration

For additional conservatism, the safety analysis is performed assuming that a part of the shutdown system does not function. For instance, for SDS-1, it is assumed that two of the shutoff rods do not drop into the core. Analysis is needed to select the two missing rods in such a way that the remaining rod configuration is the least effective. The missing rods are usually adjacent (see Figure 8.2), leaving an uncovered region in which the power pulse will be higher.

The safety analysis also makes a conservative assumption regarding the speed of insertion of the shutoff rods. The insertion characteristic should not be faster than demonstrated in field tests of shutoff-rod drop.

In simulations of SDS-2, conservative assumptions are made for the pressure in the injection tanks and the poison concentration in the tanks. In addition, one of the (6 or 7, depending on the reactor) poison tanks is usually assumed to be non-functional.

8.7 Decay Heat

In the calculation of the power distribution, it is necessary to remember that the thermal power produced in the reactor has two components:

- the "prompt", or neutronic, component, which appears very quickly following fission, and
- the decay heat, which is produced in the decay of fission products and which appears delayed (seconds/minutes to weeks/months) following fission.

In steady-state operation, the decay heat is approximately 7% of the total thermal energy generated. In a transient situation, the decay power has a time variation which is very different from that of the prompt (neutronic) power. While the prompt power increases quickly and is reduced quickly (within seconds), the decay power decreases very slowly. This is illustrated schematically in Figure 8.3.

In addition to the decay heat from fission products present before the accident, there is a decay-heat component originating in fission products newly created in the power pulse. This component should also be taken into account.

8.8 Calculation of Peak Fuel Enthalpy

The fuel enthalpy can be evaluated for any fuel bundle by integrating the sum of prompt and decay components of bundle power during the power pulse, and adding to the pre-accident enthalpy (consistent with the fuel temperature).

To cater to the possibility that the LOCA may occur at a time when a bundle power is at the license limit, whereas that situation may not be found in the *calculated* pre-accident configuration, a conservative assumption is made. The fuel bundle in which the energy generation during the LOCA is the largest is identified. The *relative* power pulse for that bundle (i.e., the normalized ratio of instantaneous to initial power) is then applied to a bundle at the license limit. In fact it is applied to the hottest fuel element in such a bundle, and the fuel enthalpy calculated for this element is taken as the peak enthalpy for the LOCA.

The margin to fuel fragmentation is obtained by comparing this peak fuel enthalpy to a conservatively low limit for fuel break-up. This value, derived from experiment, is in the range of 200-250 cal/g UO_2.

It is now increasingly acknowledged that fuel break-up (fragmentation) is not a credible consequence of LOCA in CANDU, because the rate of energy addition is too slow. Instead, fuel-centreline melting and pressure-tube integrity must be examined. This requires further analysis with thermalhydraulics and fuel-channel codes.

9. In-Core LOCA

The scenario for an in-core LOCA consists of a slow loss of coolant resulting from a single-channel event such as a pressure-tube rupture or a stagnation-feeder break. The coolant lost through the break displaces moderator volume. A number of assumptions worsen the consequences of the slow LOCA:

- the moderator is assumed to contain a certain concentration of soluble poison (boron or gadolinium) before the event; the poison is diluted by the discharging coolant, thereby inserting positive reactivity in addition to the positive void effect;
- the pre-event poison concentration is maximized by assuming, for instance, that the accident occurs during reactor restart following a long shutdown; in this scenario saturating fission products such as ^{135}Xe have decayed away and reactor criticality requires a compensating concentration of moderator poison;

- if the reactor under consideration is still relatively "young", the moderator poison concentration is further increased by assuming the accident occurs at the plutonium peak, i.e., at maximum lattice reactivity, requiring extra compensating poison;
- if the reactor is not young, coolant-void reactivity is assumed increased by pressure-tube radial creep;
- pipe whip by the ruptured pressure tube is assumed to damage a number of shutoff-rod guide tubes, preventing these shutoff rods from falling (either completely or even partially) into the core.

The physics analysis must assess the combined reactivity effect of all these assumptions, and demonstrate that the (reduced) SDS-1 can arrest the fission chain reaction and maintain subcriticality until such time at which the reactor operator can be assumed to intervene - usually taken as 15 minutes after an "unambiguous indication of the event" in the control room.

The physics analysis in this accident scenario thus reduces essentially to an accurate assessment of the evolution with time of the overall system reactivity, i.e., of the balance between the negative reactivity of the diminished SDS-1 and the positive reactivity effects assumed. This assessment will again require neutron-diffusion calculations in three dimensions, using as accurate a modelling of the various assumptions as possible. Some components of the modelling may of course be difficult to pin down:

- a precise assessment of the expected damage to shutoff-rod guide tubes may be difficult to obtain, and it may be difficult to verify the prediction; and
- the precise manner in which the discharging coolant displaces moderator may also be open to question: the "piston" and "instantaneous uniform mixing" models are at opposite extremes of possible modelling, and recent consensus seems to favour the more reasonable "delayed-mixing" model.

From the above discussion, it is quite clear that the physics analysis of the in-core LOCA is not isolated, but indeed tightly coupled to the assumptions and conclusions of other disciplines.

Note that it is only the effectiveness of SDS-1 which is assumed degraded in an in-core LOCA. The continued effectiveness of SDS-2 is not in question, since high-pressure injection of poison into the moderator can take place even if the SDS-2 nozzles are damaged by pipe whip, and since the reactivity depth of SDS-2, at several hundred (negative) milli-k, is much larger than that of SDS-1.

10. Summary

The physics analysis is essential to the quantitative understanding of the behaviour of the core following a hypothetical accident. Loss of coolant in CANDU introduces positive reactivity and promotes a power rise. The large LOCA is the accident which presents the greatest challenge to CANDU shutdown systems in terms of rate of positive reactivity insertion.

Many conservative assumptions are made in typical physics simulations of hypothetical accidents. The physics component provides important quantitative information used in the rest of the safety analysis.

Neutronics methods and models for accident analysis have greatly evolved over the last twenty years.

TABLE 1.2

PHOTONEUTRON DATA

PRECURSOR GROUP g	FRACTIONAL GROUP YIELD a_g^{pn}	DECAY CONSTANT λ_g^{pn} (s^{-1})	TOTAL PHOTONEUTRON YIELD PER FISSION Y^{pn}	TOTAL PHOTONEUTRON DELAYED FRACTION
1	.0005	6.26×10^{-7}	8.53×10^{-4}	U-235: 3.5×10^{-4}
2	.00102	3.63×10^{-6}		Pu-239: 3.0×10^{-4}
3	.00320	4.37×10^{-5}		Pu-241: 2.9×10^{-4}
4	.0232	1.17×10^{-4}		U-238: 3.0×10^{-4}
5	.0205	4.28×10^{-4}		
6	.0333	1.50×10^{-3}		
7	.0695	4.81×10^{-3}		
8	.2025	1.69×10^{-2}		
9	.6462	2.77×10^{-1}		

TABLE 1.3 CANDU-6 DELAYED-NEUTRON DATA COLLAPSED TO 6 GROUPS

GROUP	FRESH FUEL		EQUILIBRIUM FUEL (1.8 n/kb)	
	β_g	$\lambda_g\ (s^{-1})$	β_g	$\lambda_g\ (s^{-1})$
1	.000385	.000733	.000295	.000612
2	.001526	.03173	.001165	.03155
3	.001385	.1172	.001033	.1218
4	.003258	.3128	.002350	.3175
5	.001049	1.402	.000780	1.389
6	.000240	3.912	.000197	3.784
β_{tot}	.007843		.005819	
l^*	0.962×10^{-3} s		0.902×10^{-3} s	
V_{th}	2.725×10^5 cm/s		2.726×10^5 cm/s	

Table 5.1
Decay Constants for the Precursors of Delayed Neutrons from the Fuel

Group	λ (s^{-1}) (decay constant)	Half-life (s)
1	0.0129	53.73
2	0.0311	22.29
3	0.134	5.17
4	0.331	2.09
5	1.26	0.55
6	3.21	0.22

Table 5.2
Fraction of Delayed Neutrons from the Fuel, β_{ki}

Group	^{235}U	^{239}Pu	^{241}Pu	^{238}U
1	0.000251	0.000087	0.000066	0.000206
2	0.001545	0.000639	0.001234	0.002174
3	0.001476	0.000493	0.000932	0.002570
4	0.002663	0.000747	0.002102	0.006156
5	0.000756	0.000235	0.000981	0.003570
6	0.000293	0.000080	0.000086	0.001190
Total β	0.006984	0.002281	0.005389	0.015866

Table 5.3
Photoneutron Time Constants and Delayed Fractions

Group k	$T_{k,1/2}$	λ_k (s^{-1})	β_k
7	12.8 d	6.26×10^{-7}	1.653×10^{-7}
8	2.21 d	3.63×10^{-6}	3.372×10^{-7}
9	4.41 h	4.37×10^{-5}	1.058×10^{-6}
10	1.65 h	1.17×10^{-4}	7.670×10^{-6}
11	27 min	4.28×10^{-4}	6.779×10^{-6}
12	7.7 min	1.50×10^{-3}	1.101×10^{-5}
13	2.4 min	4.81×10^{-3}	2.298×10^{-5}
14	41 s	1.69×10^{-2}	6.694×10^{-5}
15	2.5 s	2.77×10^{-1}	2.136×10^{-4}
TOTAL (photoneutrons)	16.7 min*	6.92×10^{-4}	3.306×10^{-4}

* Average values: $\lambda_{ph} = \Sigma\ (\beta_k/\lambda_k)/\ \Sigma\ \beta_k$ and $T_{1/2} = \ln 2 \cdot \lambda_{ph}$

Table 5.4
Typical Delayed-Neutron Data for CANDU Equilibrium Core
(6 Collapsed Groups)

Group	**λ_k (s^{-1}) (Decay Constant)**	**Delayed Fraction β_k**
1	0.000608	0.000291
2	0.03154	0.001150
3	0.1221	0.001019
4	0.3181	0.002314
5	1.389	0.000773
6	3.778	0.000196
Total Delayed Fraction β	-	0.00574

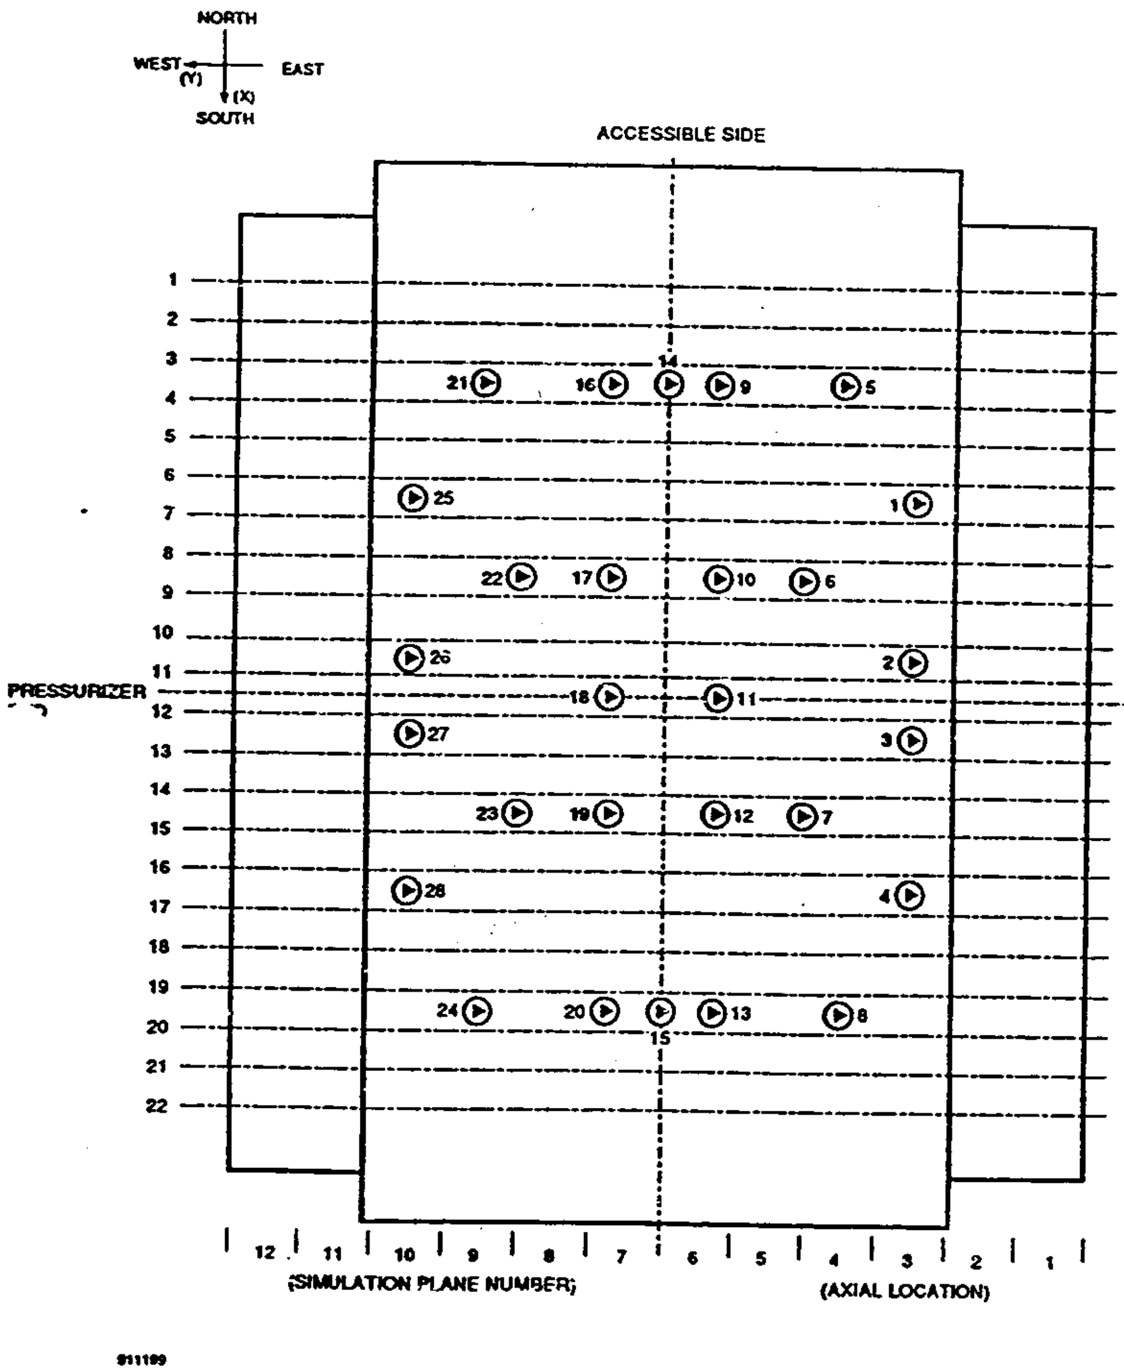

Figure ~~5-17~~ 2.1 Shutoff Rods Numbering

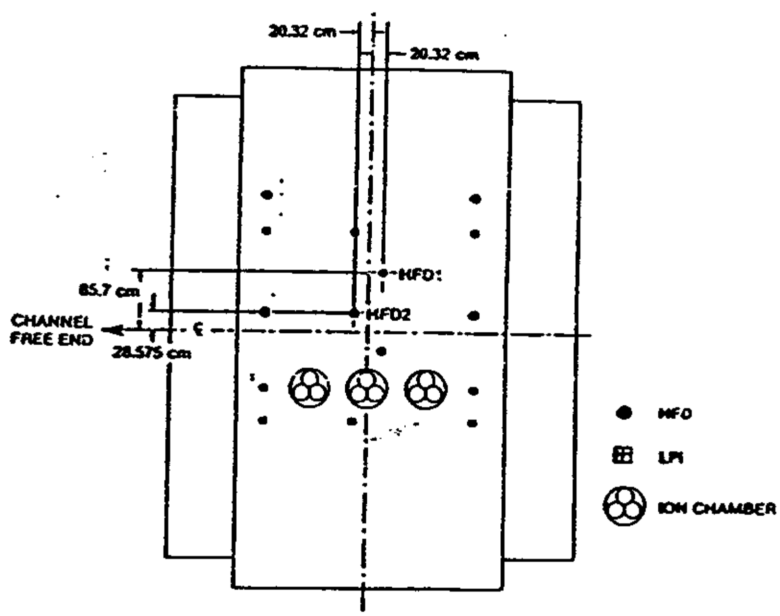

2.2

Figure ~~5.19~~ Calandria Shell Side Elevation View of the CANDU 6 Reactor Showing Positions of the Horizontal Flux Detector Assemblies, Liquid Poison Injection Nozzles and the Ion Chambers

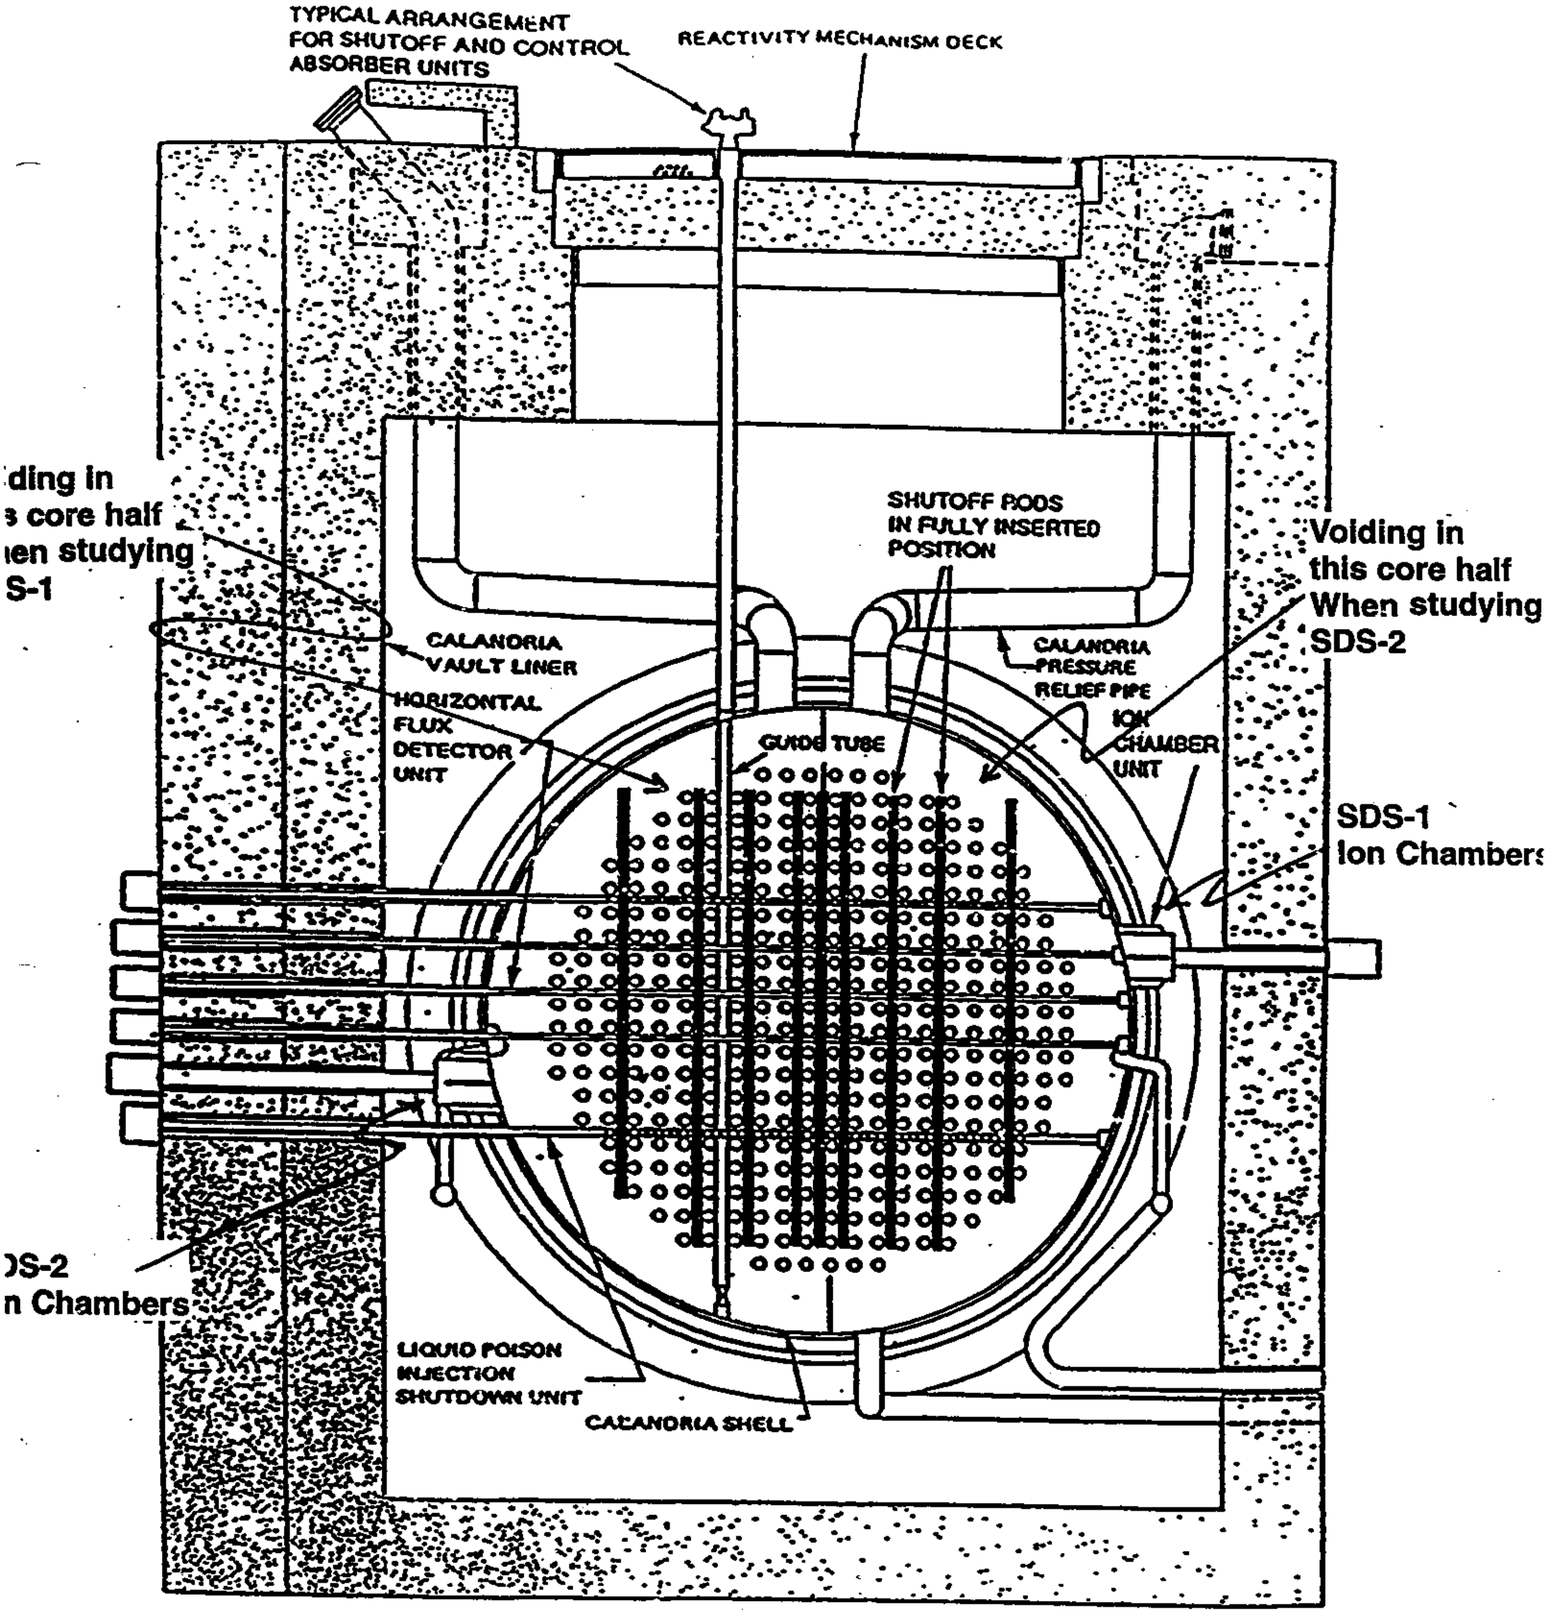

Figure 3.1 a,b
Location of SDS1 and SDS2 Ion Chambers in the CANDU 6

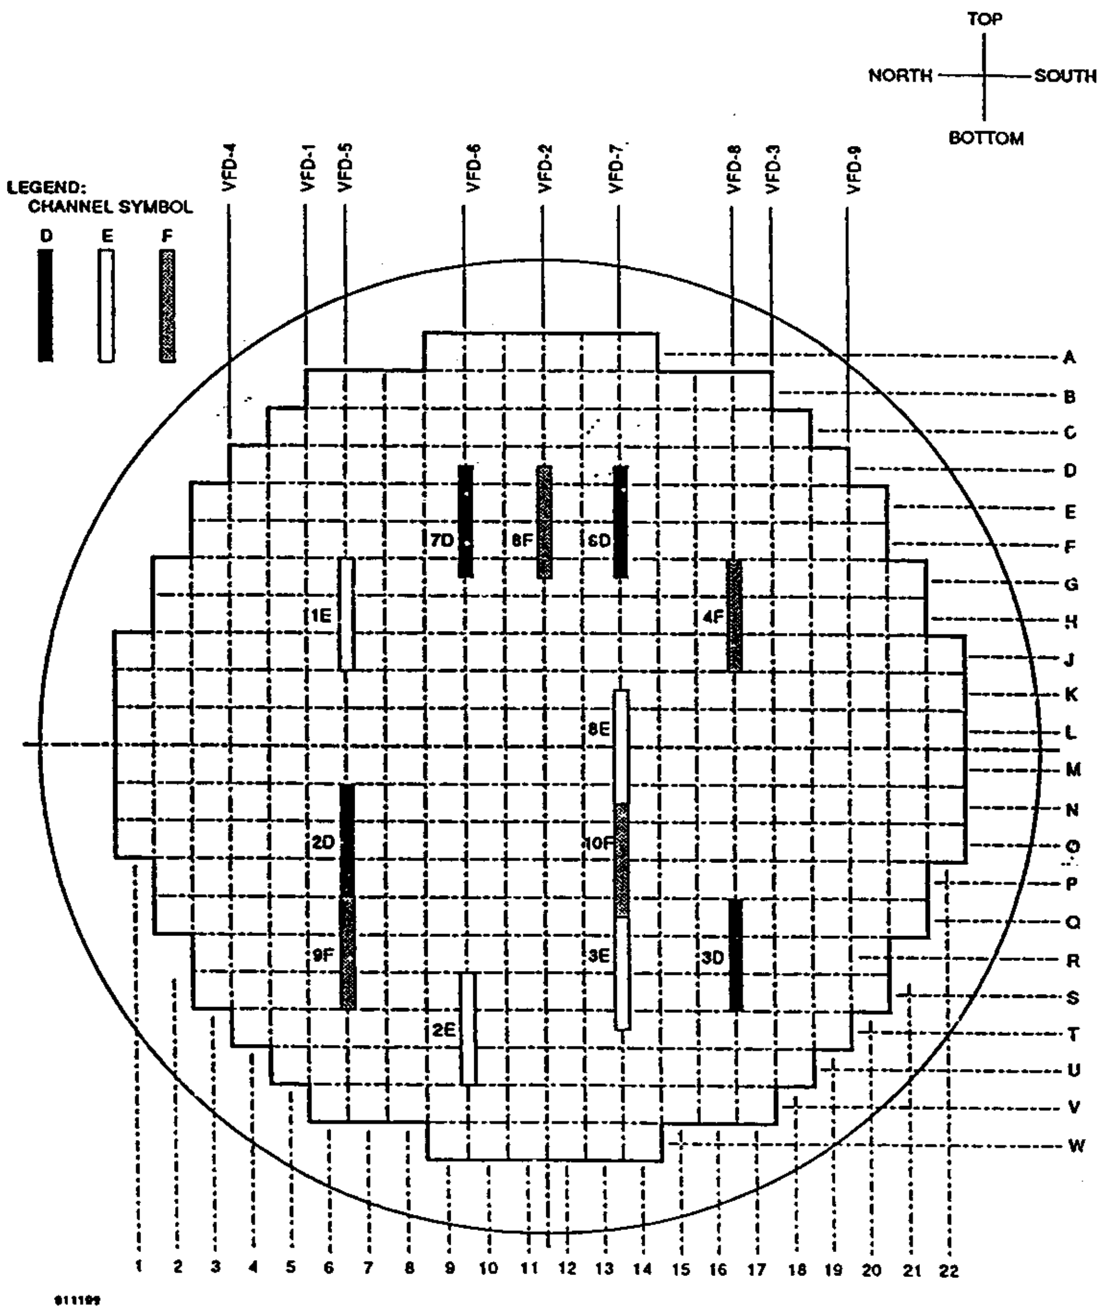

Figure ~~1(a)~~ 3.2
Location of Some Vertical In-Core ROP Detectors in CANDU 6

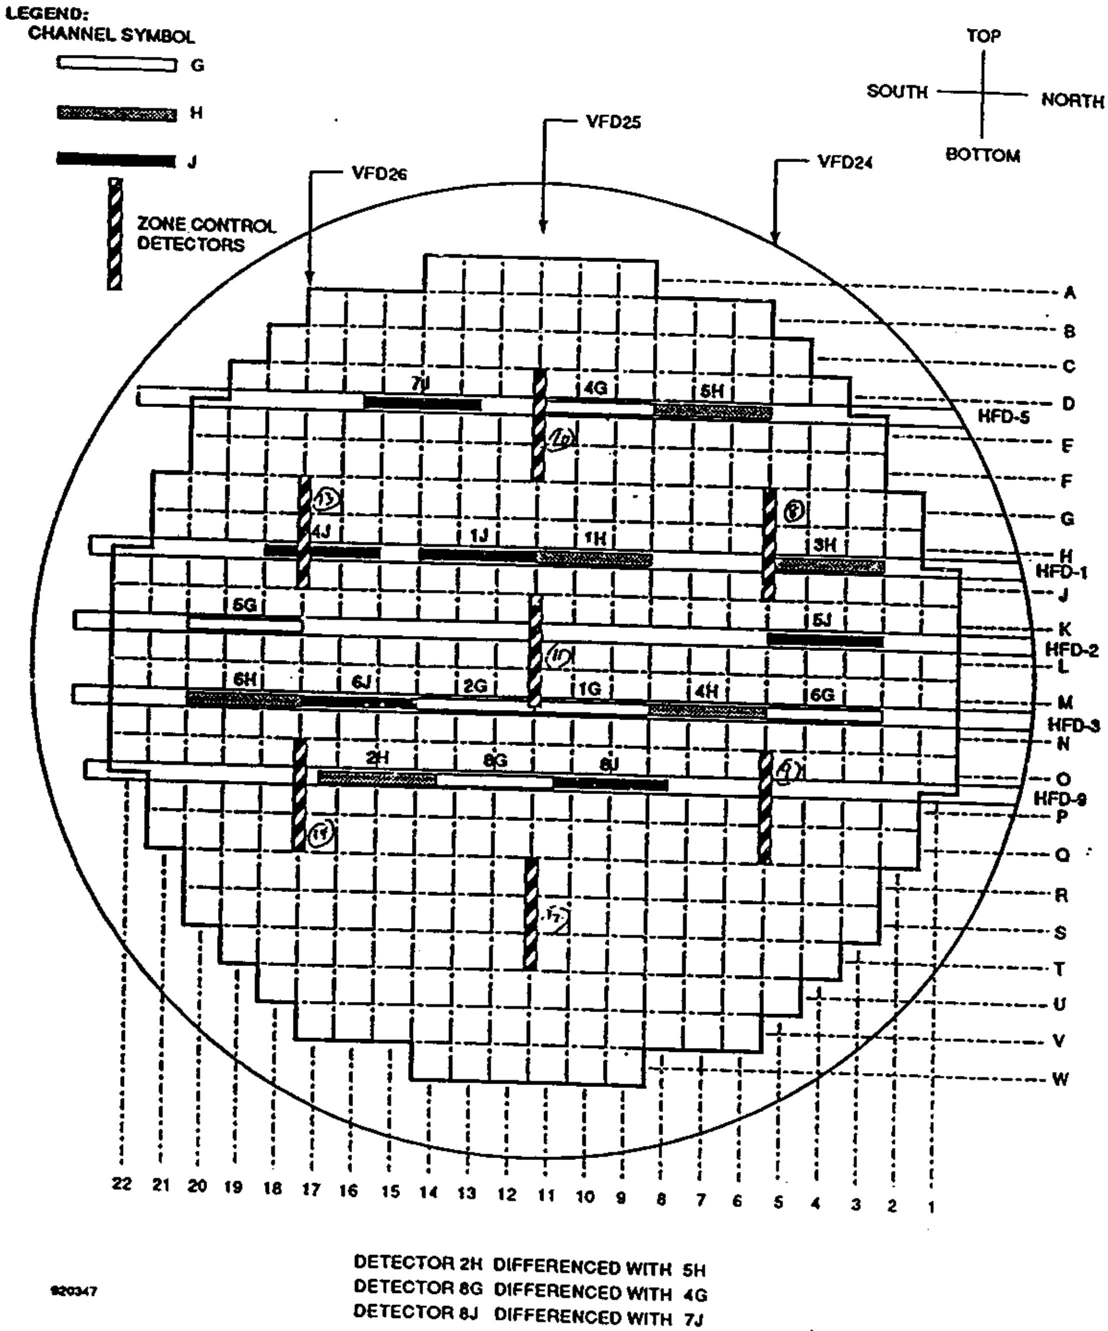

Figure ~~1(5)~~ 3.3
Location of Some Horizontal In-Core ROP Detectors in CANDU 6

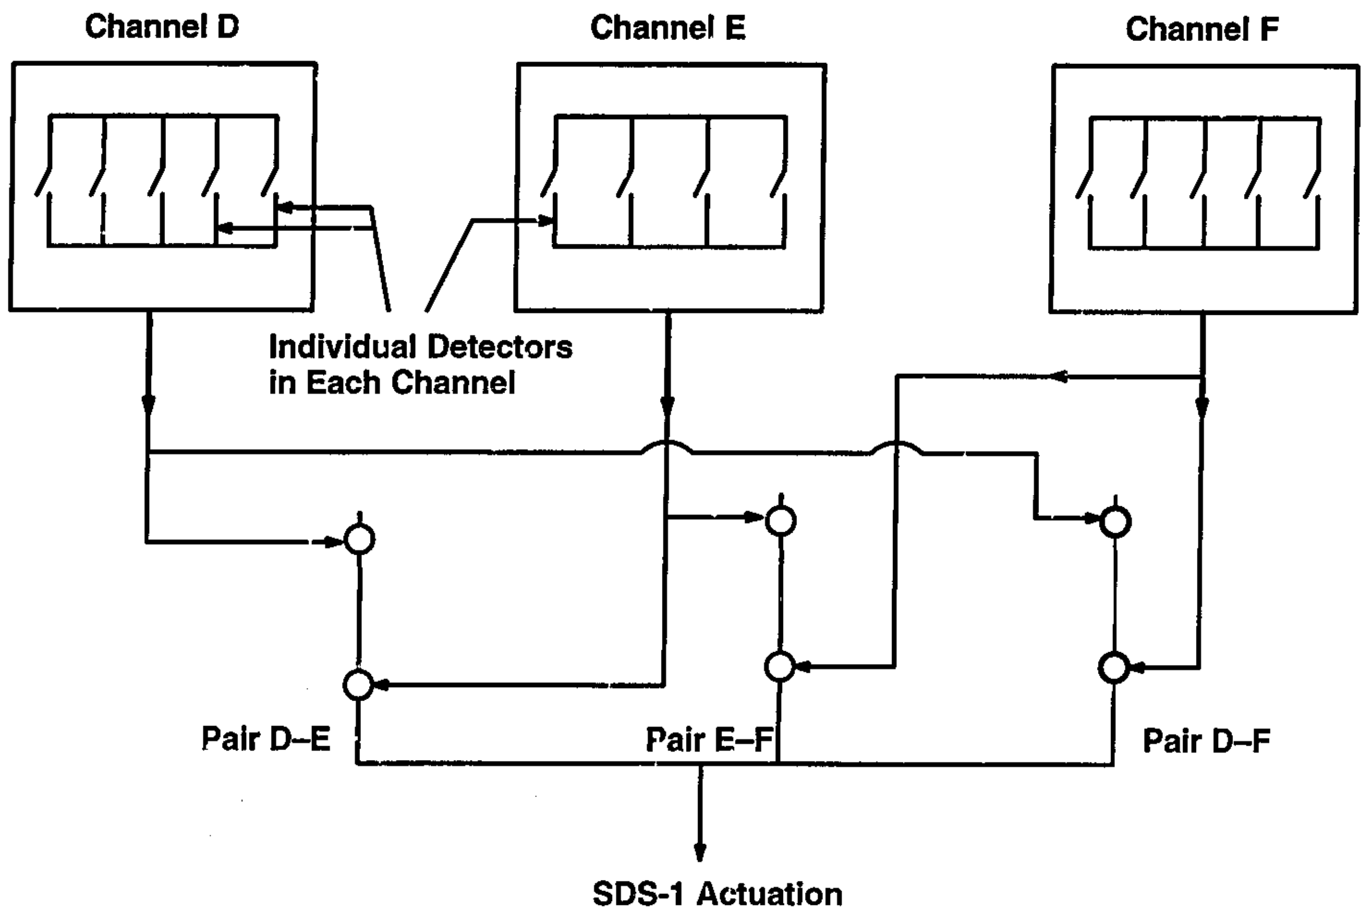

Figure 3.4
Triplicated Logic for Shutdown-System-1

u65 / bused / ben
figs_overheads / circuit_logic
dbb

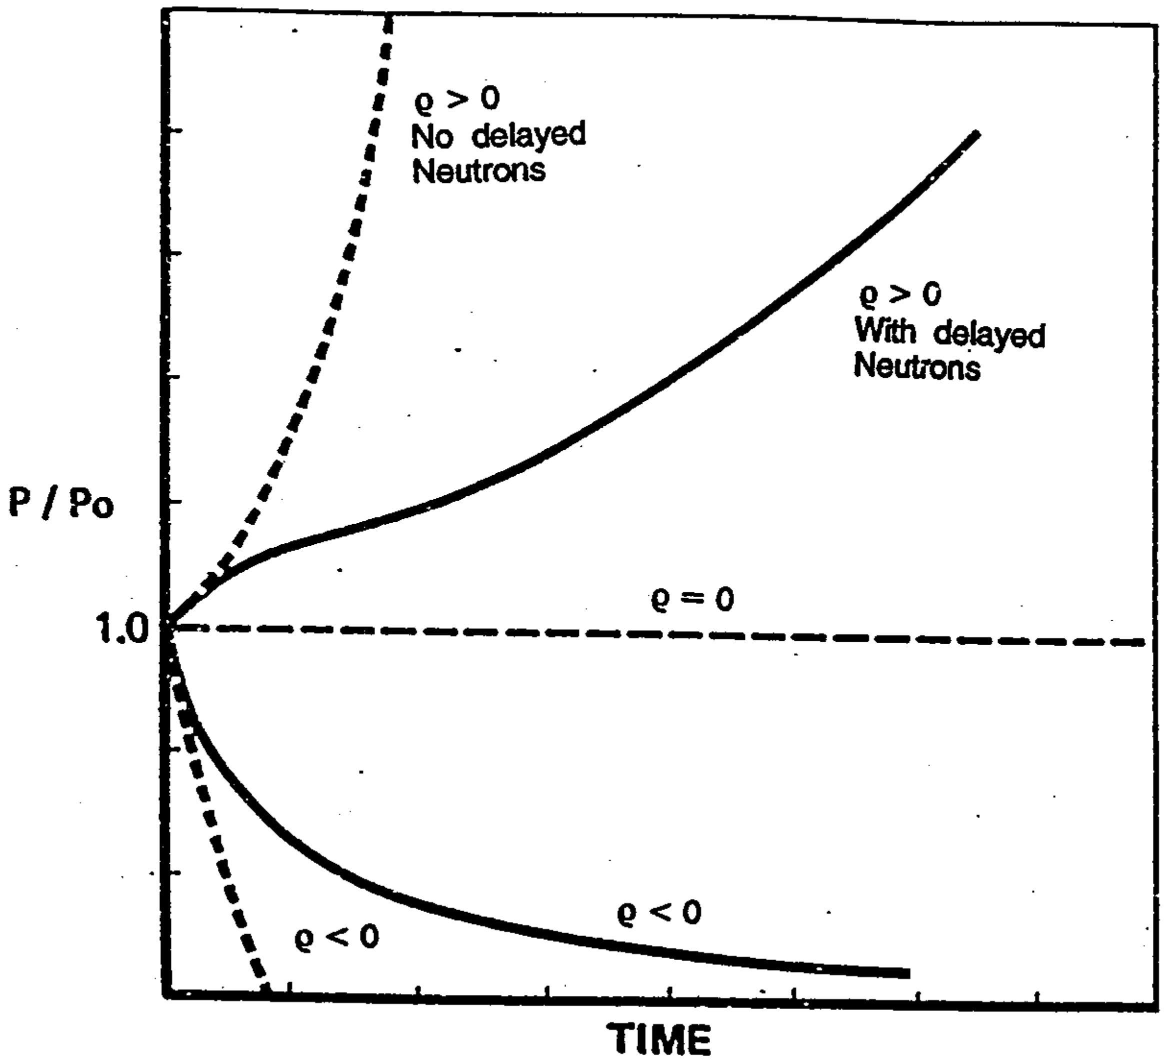

Figure 5-1
Influence of Delayed Neutrons on Power Transients

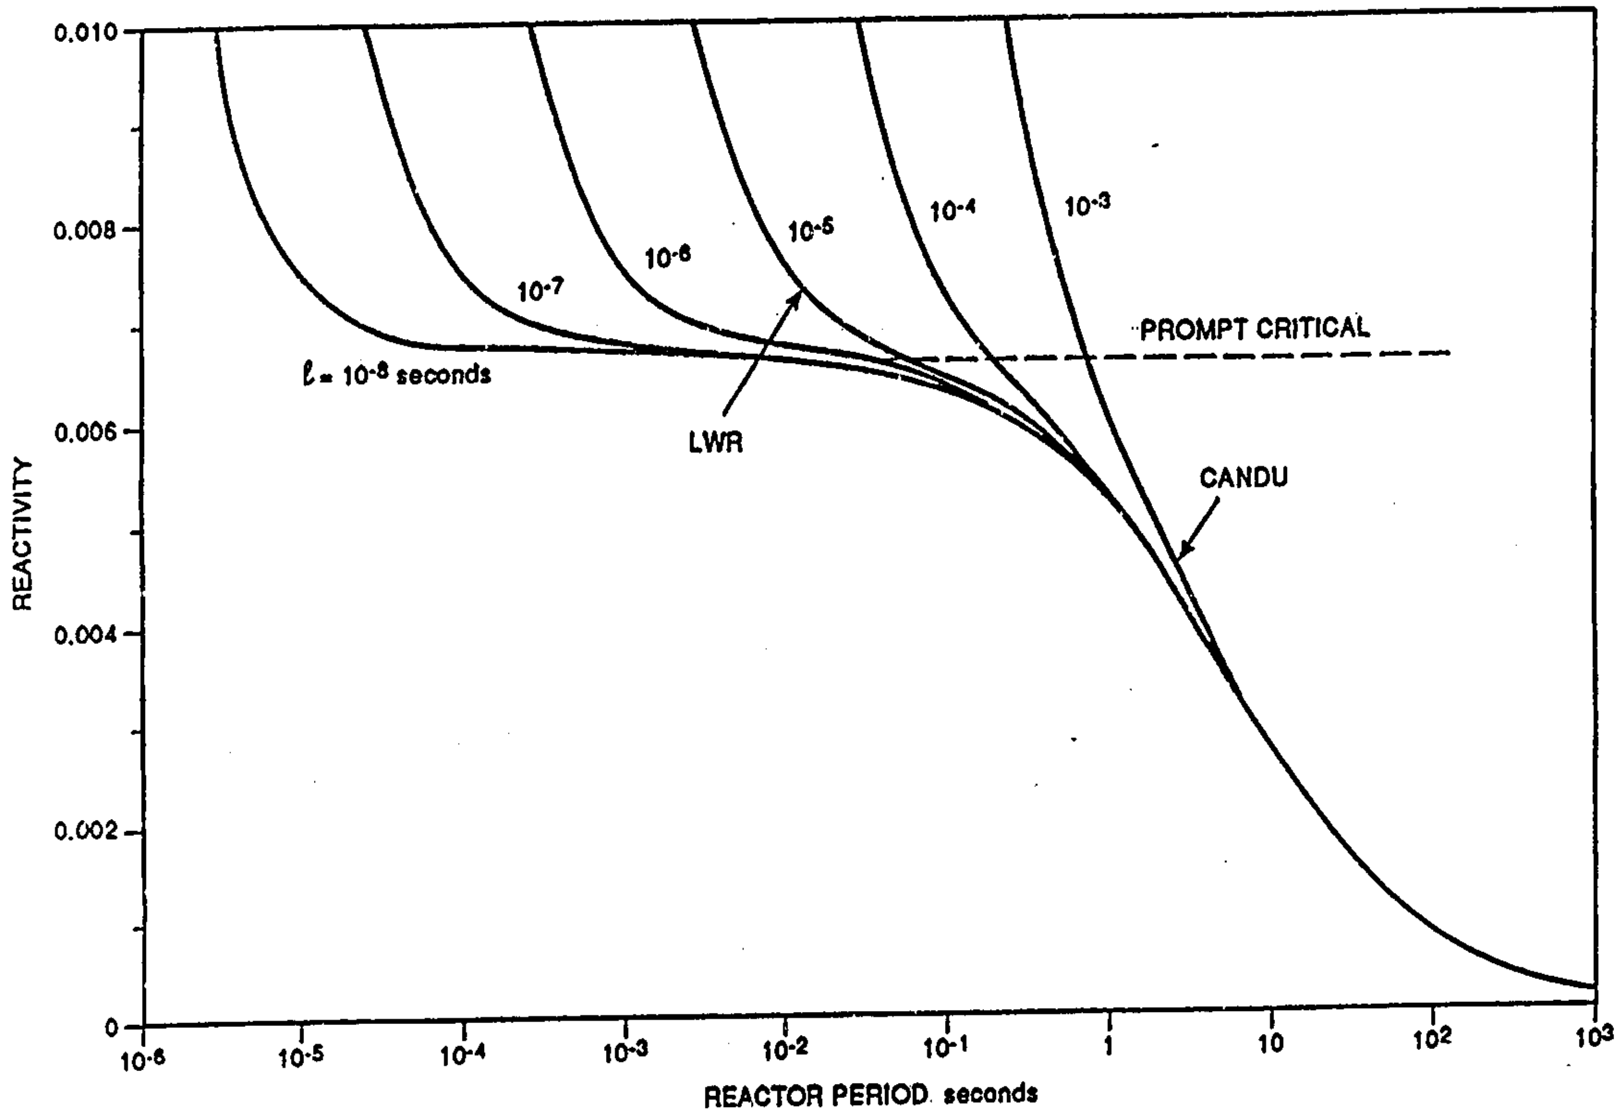

Relationship Between Reactor Period and Reactivity for Various Neutron Generation Times

FIG. 5.2

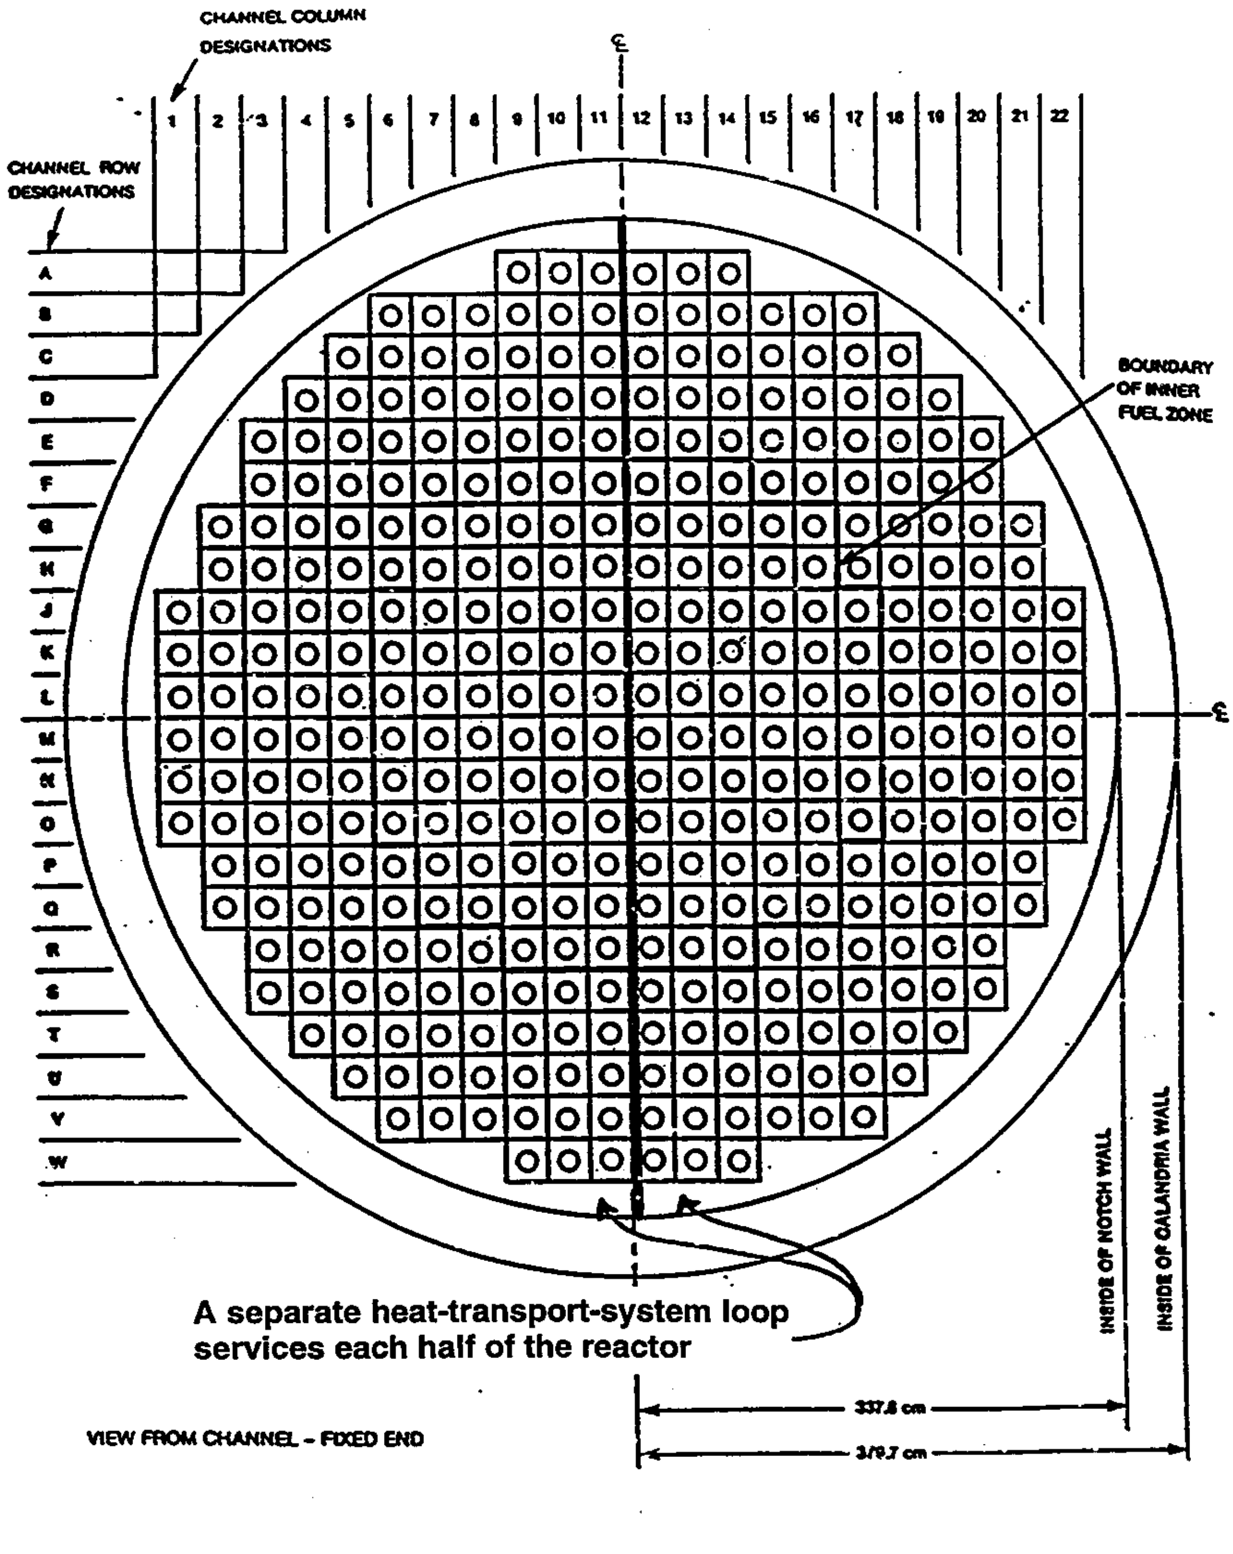

Figure 7 6.1
Subdivision of Heat-Transport-System Loops In CANDU 6

MESH ARRAY = 44 x 36 x 24
LATTICE ARRAY = 22 x 22 x 12
ORIGIN OF LATTICE ARRAY: X = 68.525
Y = 68.525
Z = 3.945

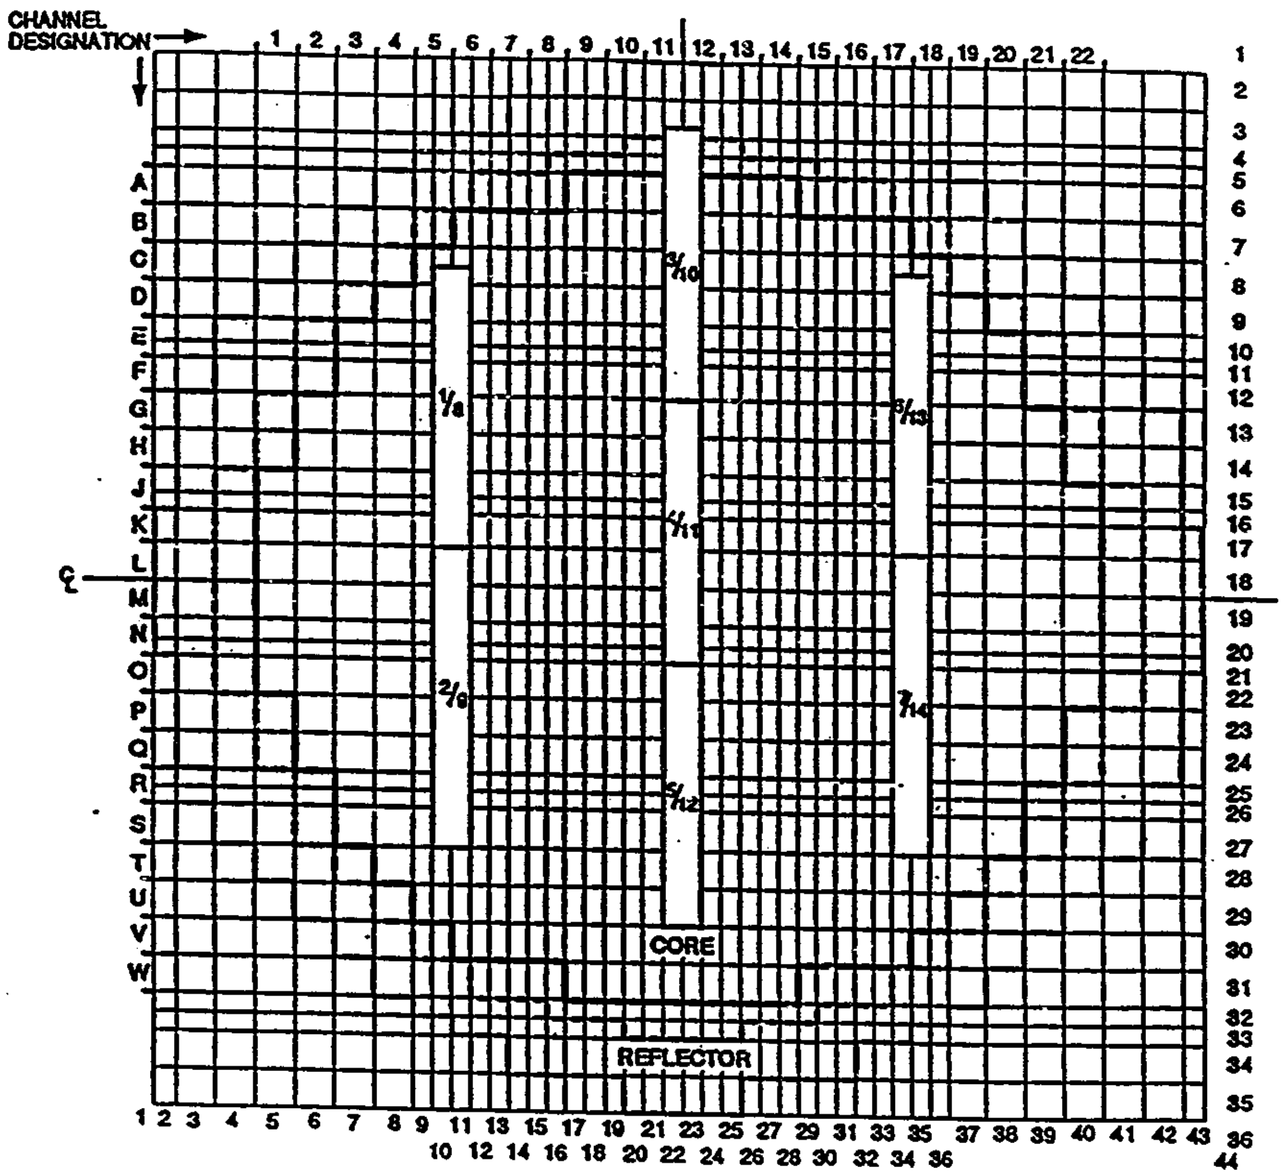

Figure ~~8(a)~~ 6.2
Typical Finite-Difference Neutron-Diffusion Model
Face View

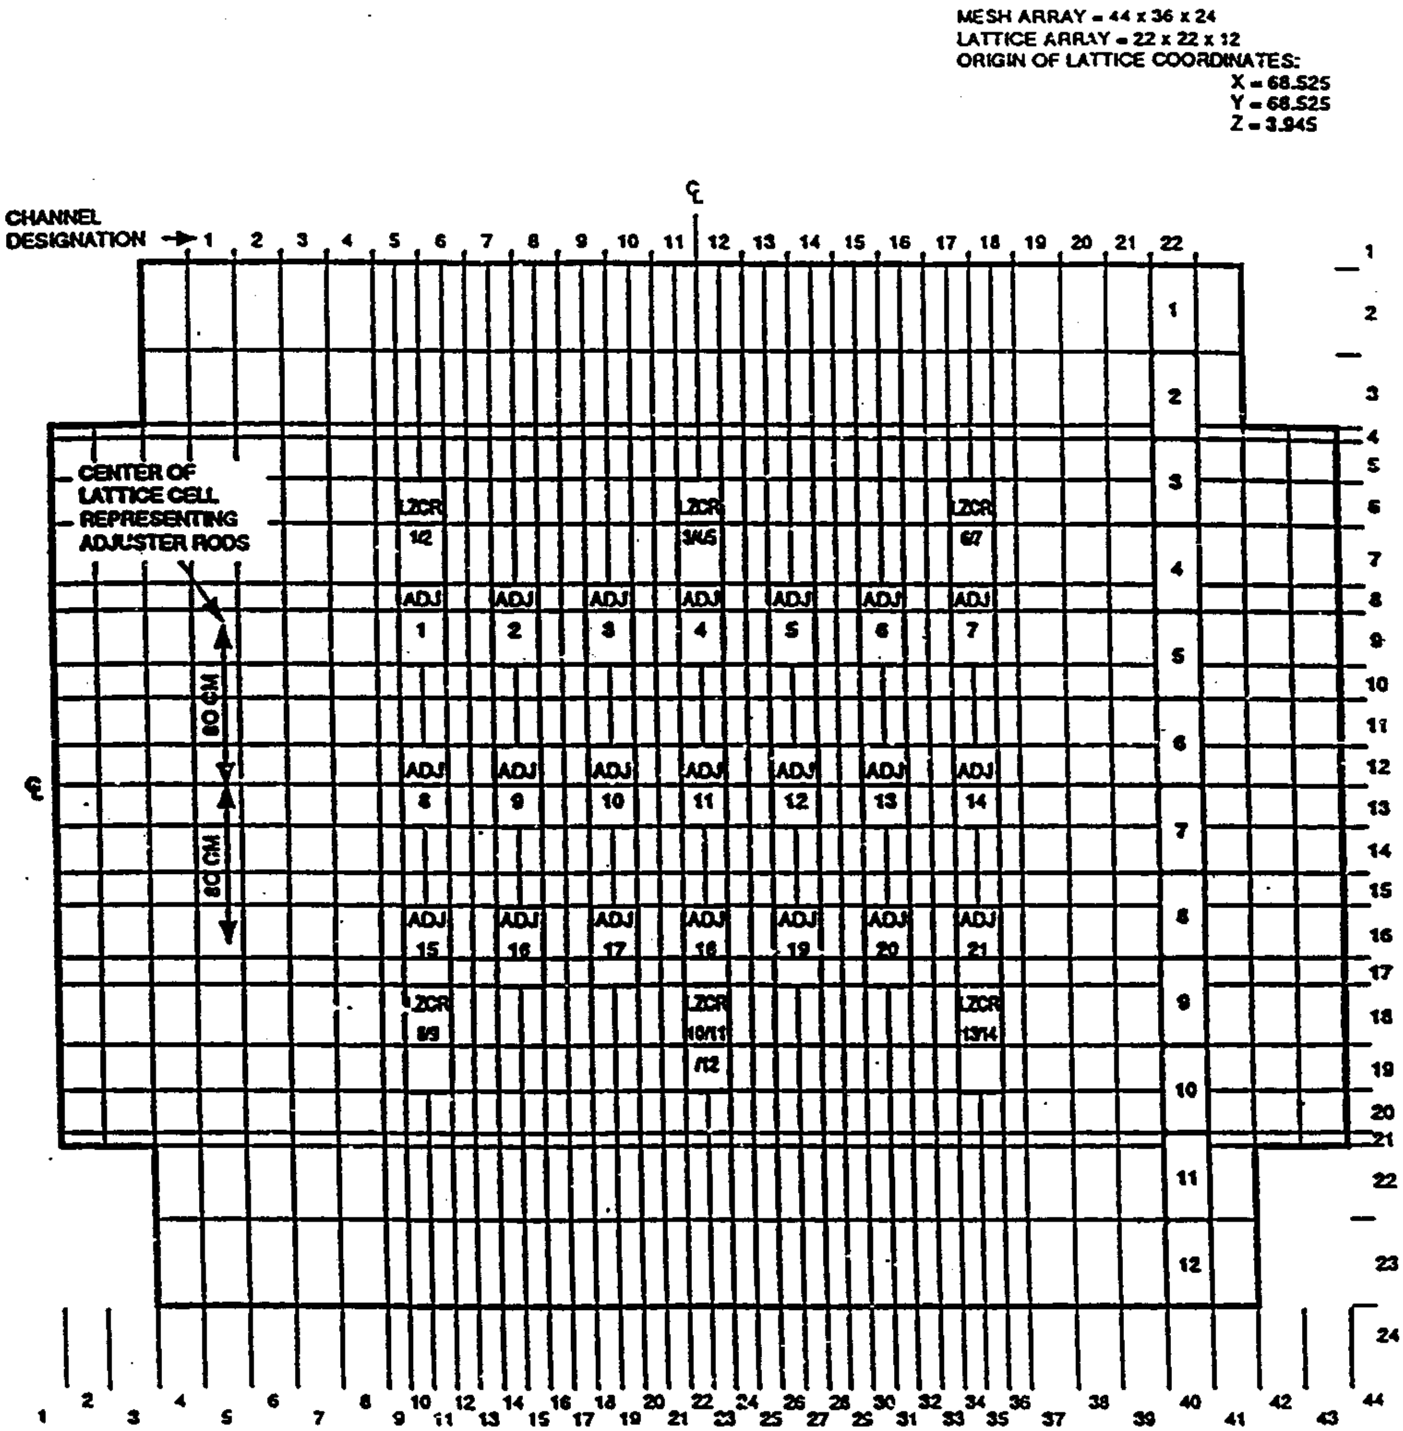

Figure 8~~(b)~~ 6.3
Typical Finite-Difference Neutron-Diffusion Model
Top View

MODE NUMBER	DESIGNATION	SUBCRITICALITY mk	MODE SCHEMATIC (IDEALIZED)
0	FUNDAMENTAL	0	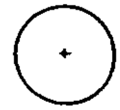
1	FIRST AZIMUTHAL-A	16.2	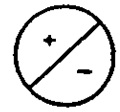
2	FIRST AZIMUTHAL-B	16.9	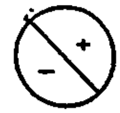
3	FIRST AXIAL	27.1	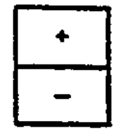
4	SECOND AZIMUTHAL-A	44.0	
5	SECOND AZIMUTHAL-B	47.0	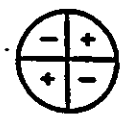
6	FIRST AZIMUTHAL-A X FIRST AXIAL	46.9	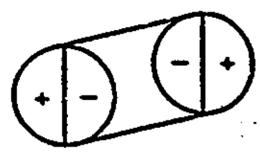
7	FIRST AZIMUTHAL-B X FIRST AXIAL	47.7	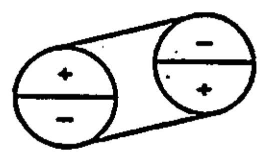
8	FIRST RADIAL X SECOND AXIAL-A	66.3	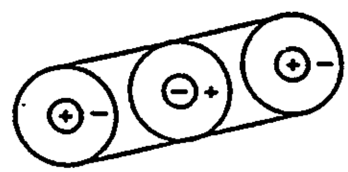
9	FIRST RADIAL X SECOND AXIAL-B	80.6	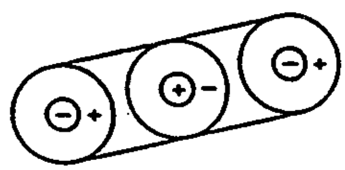

Figure Ø 6.4
Schematic of Harmonic Flux Shapes

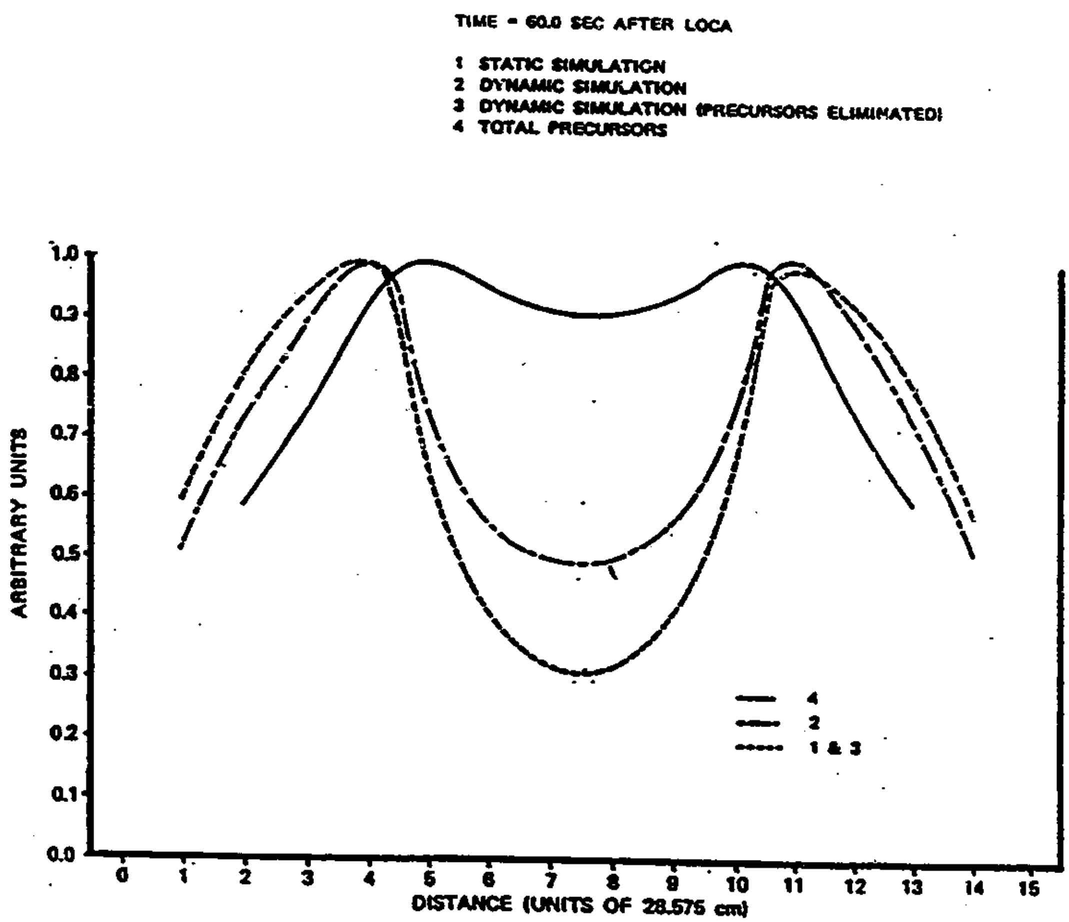

Figure ~~10~~ 6.5
Effect of Delayed-Neutron Precursors on a Flux Shape

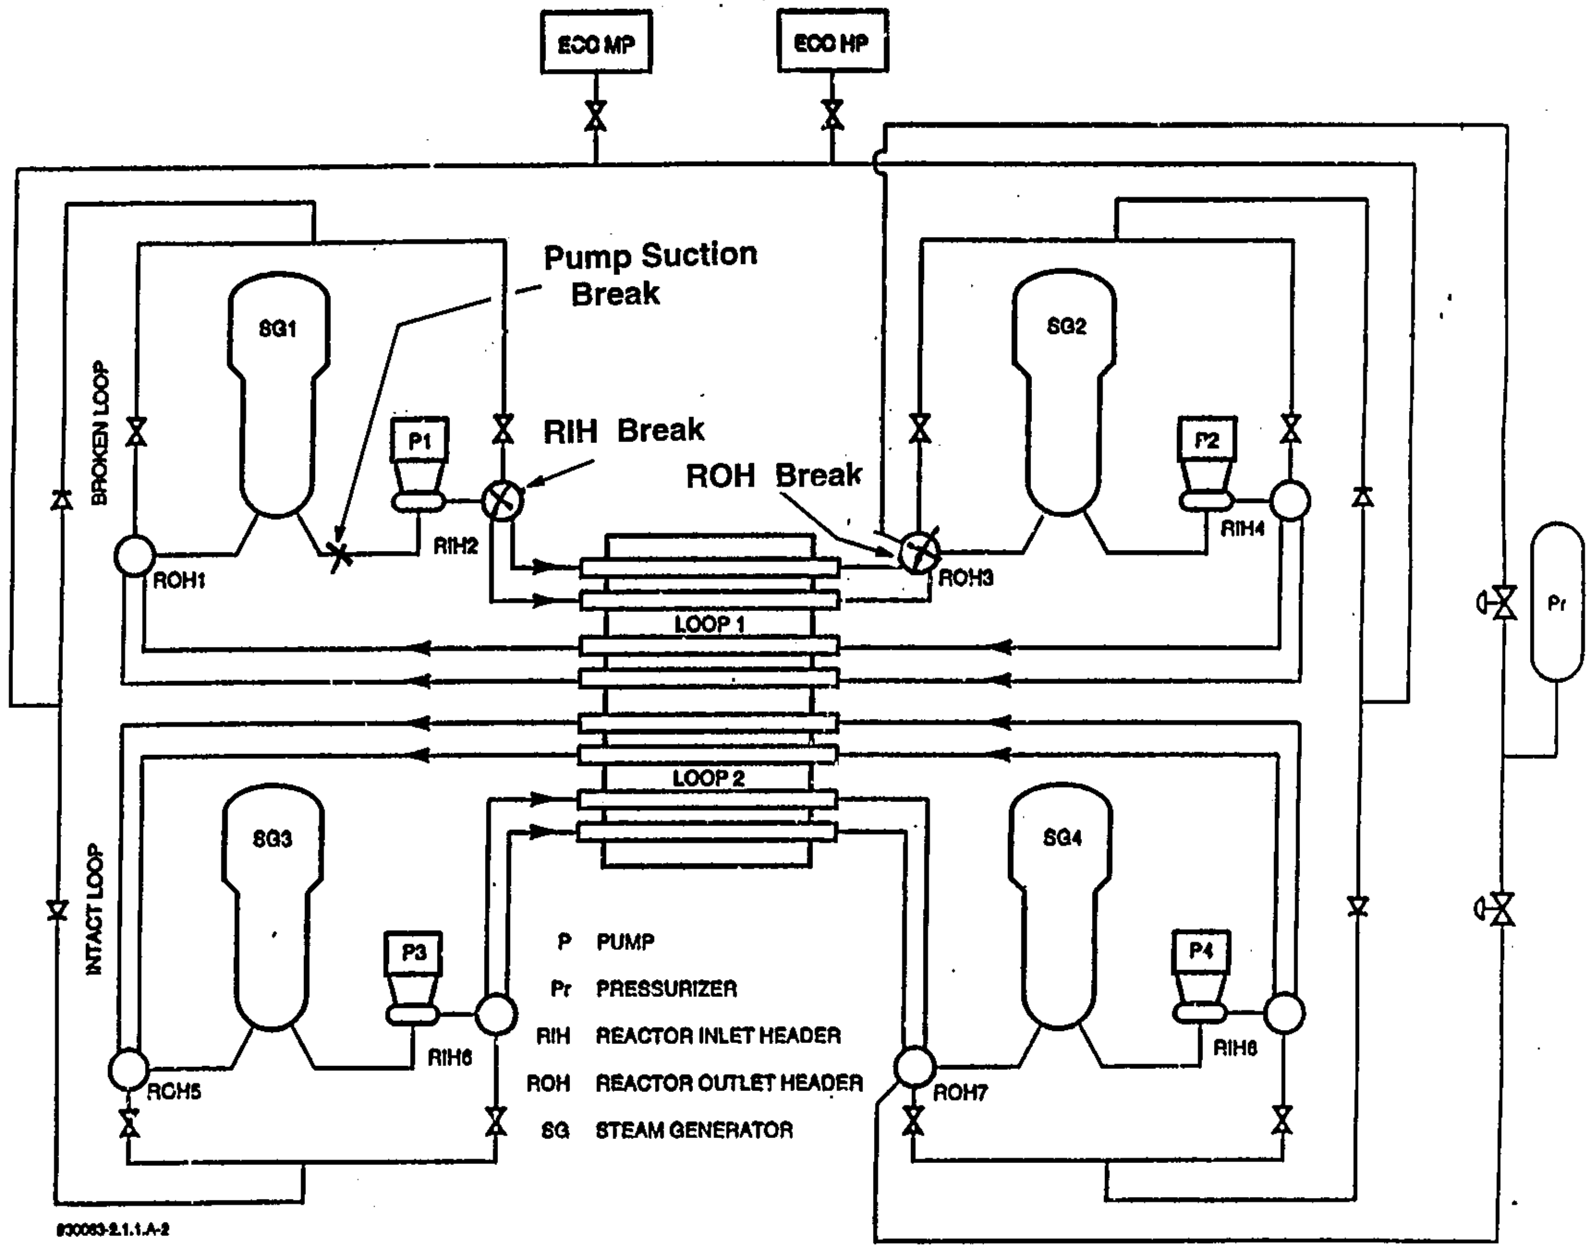

Figure 3 7.1
Examples of Break Locations Giving Rise to a Large LOCA

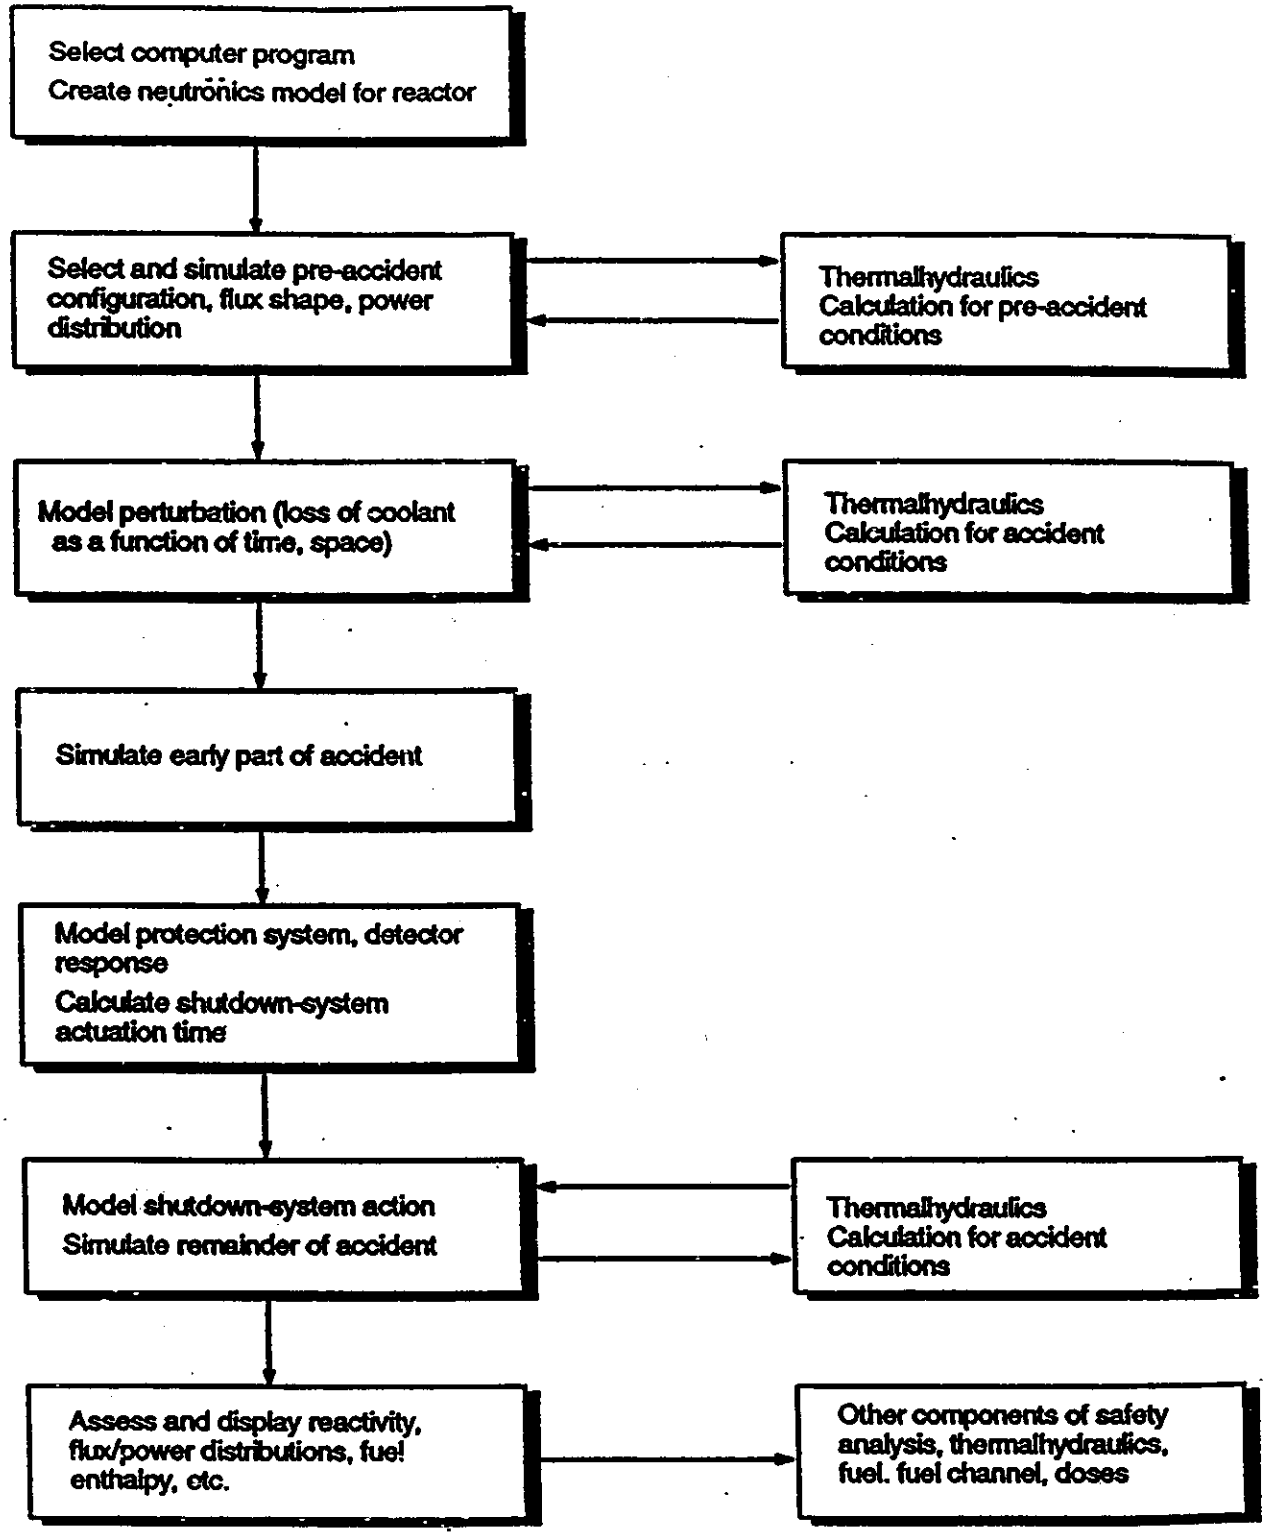

Figure 4 7.2

Schematic of a Physics Analysis for a Large LOCA

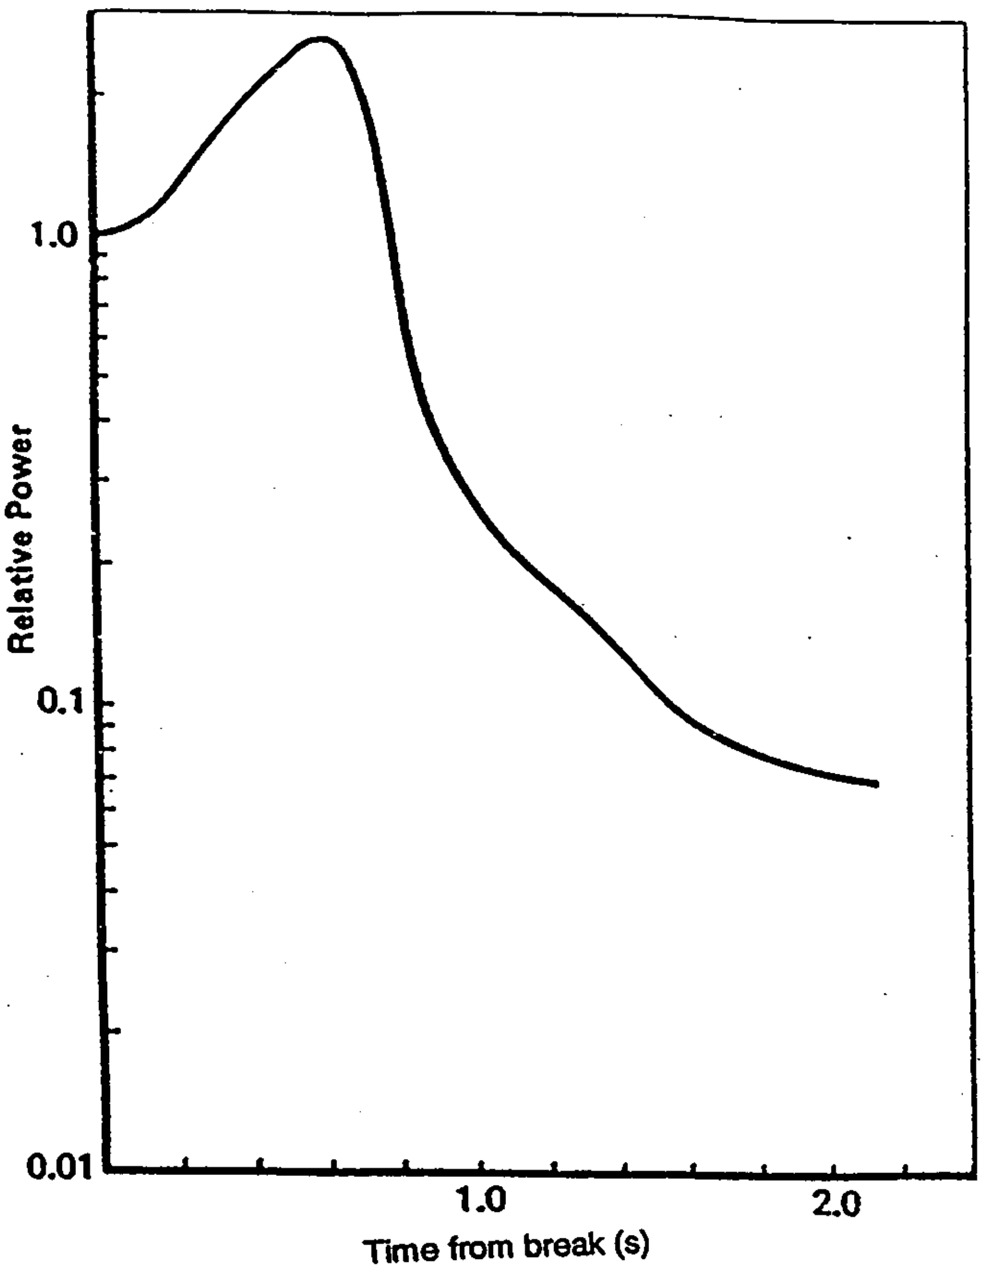

Figure 6 7-3
Schematic of a Power Transient Terminated by Shutdown-System Action

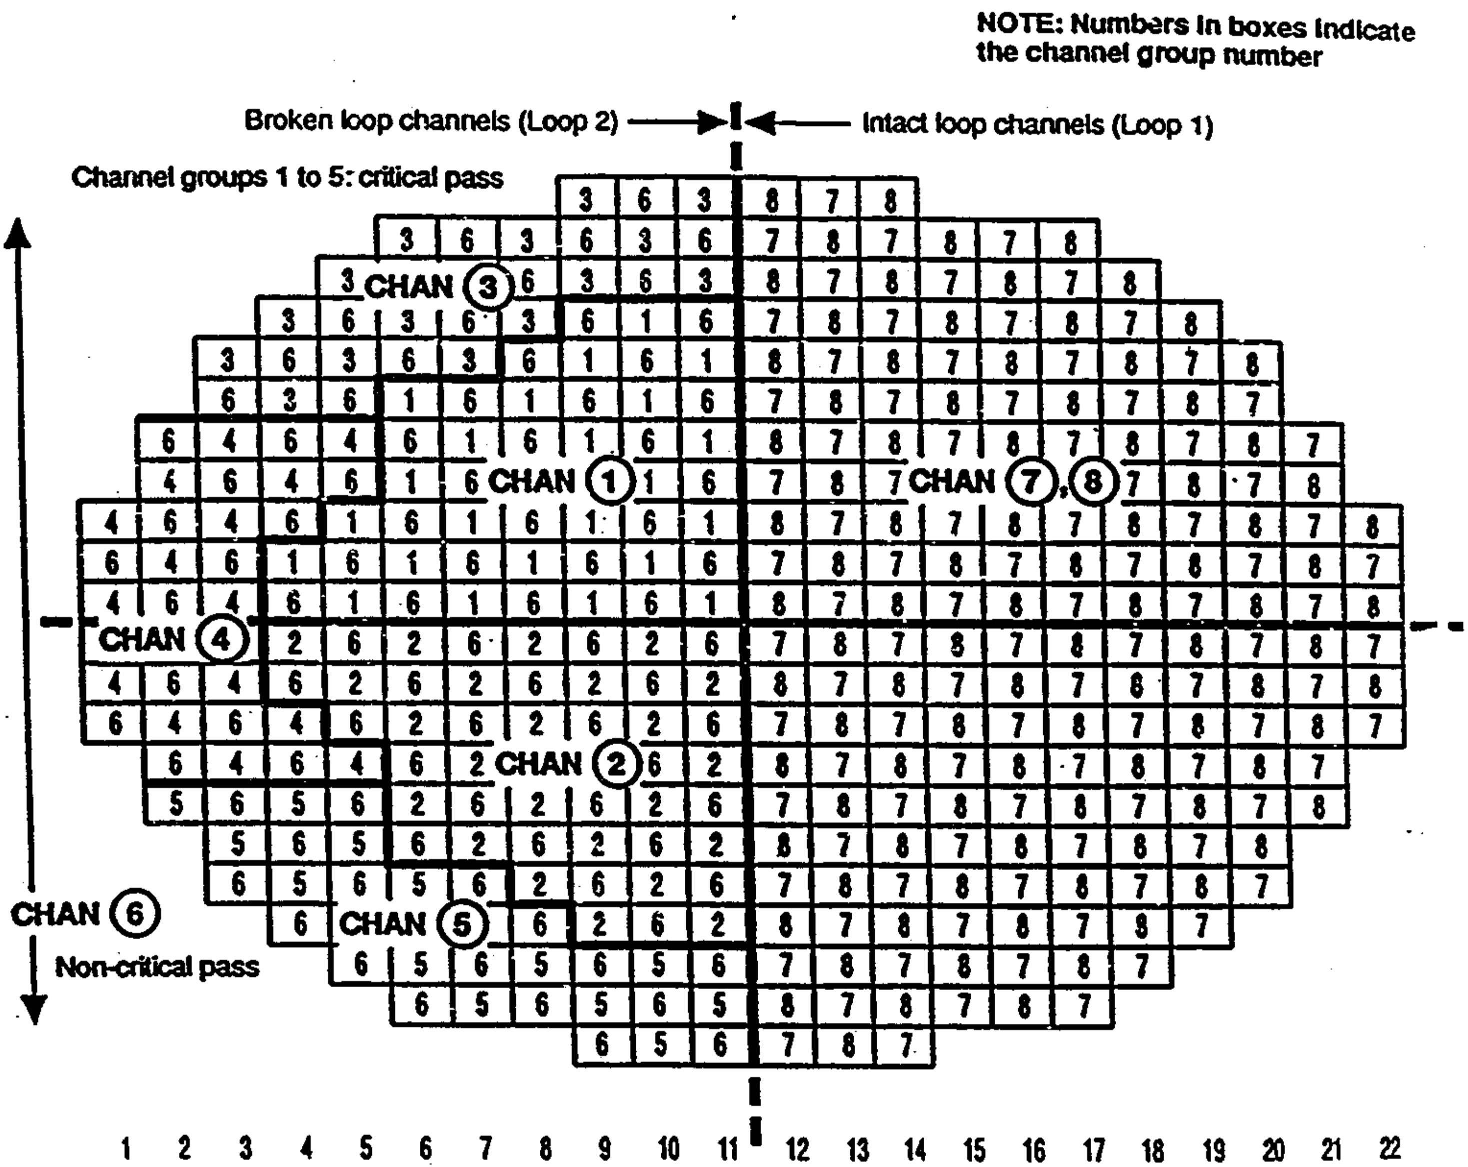

Figure ~~14~~ 8.1
Thermalhydraulic Channel Groups for a LOCA Calculation

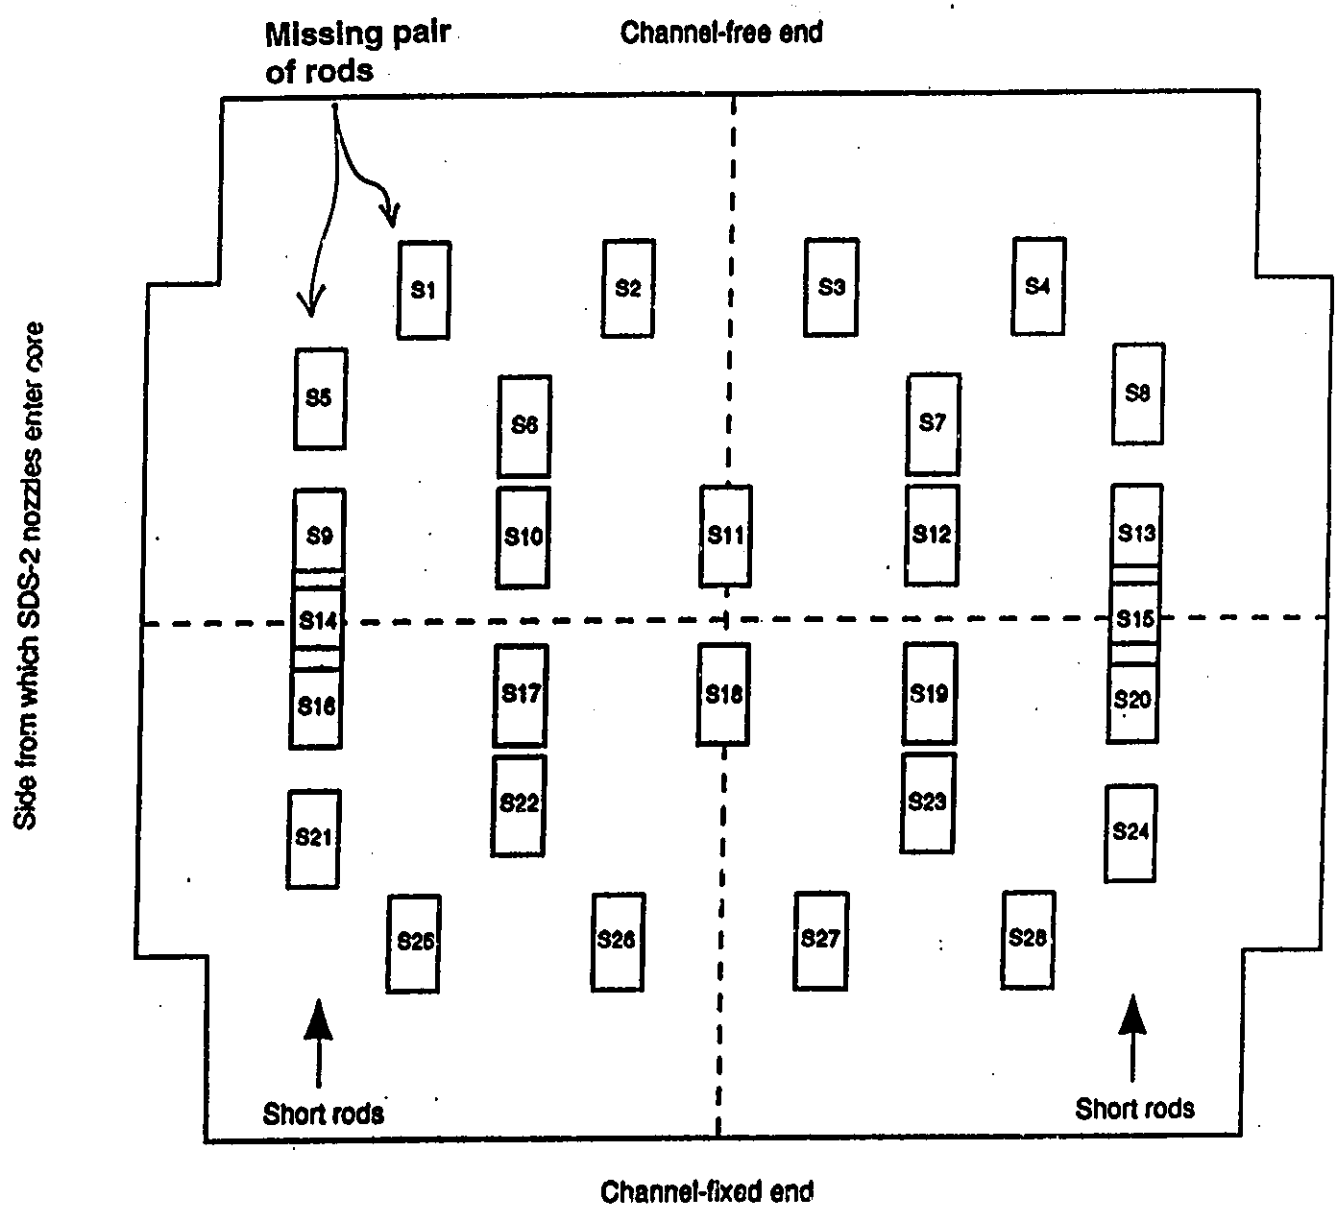

Figure ~~16~~ 8.2
Top View of Reactor Showing Shutoff Rod Location

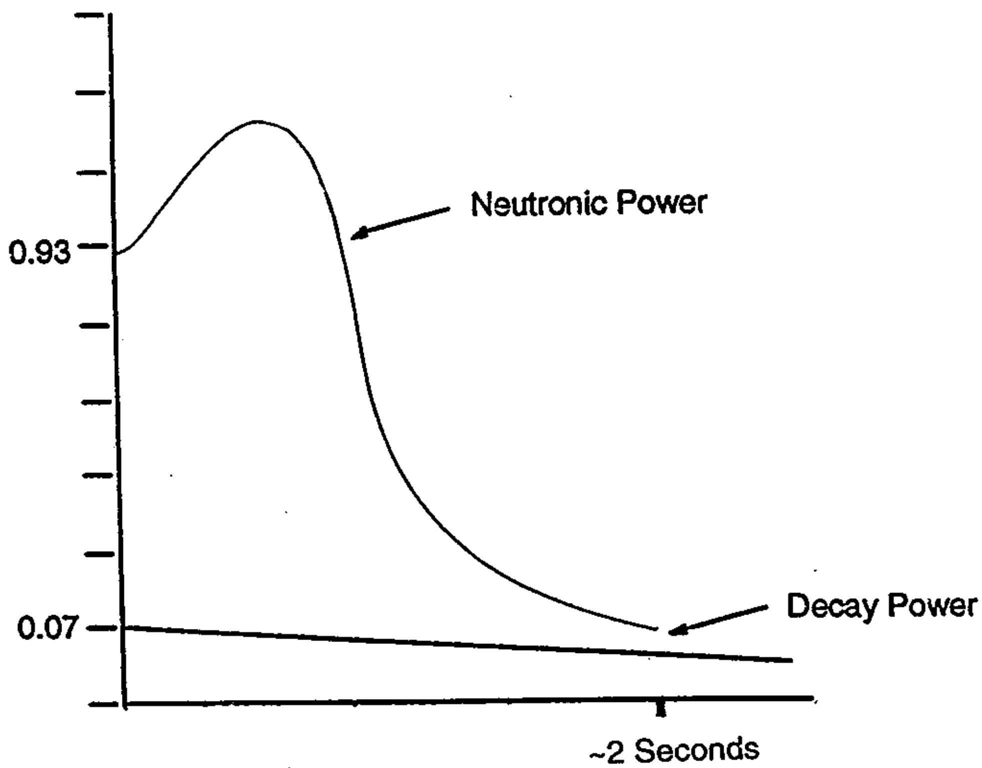

Figure ~~17~~ 8-3
Time Variation of Neutronic and Decay Powers

Function	Device	Total Reactivity worth (mk)	Maximum Reactivity rate (mk/s)
Control	14 Zone Controllers	7	± 0.14
Control	21 Adjusters	15	± 0.1
Control	4 Mechanical Control Absorbers	10	± 0.075 (driving) - 3.5 (dropping)
Control	Moderator Poison	—	- 0.01 (extracting)
Safety	28 Shutoff Units	80	- 50
Safety	6 Poison Injection Nozzles	>300	- 50

Figure 5.8 Reactivity Devices, Worths and Maximum Rates