REACTOR PHYSICS

COURSE INTRODUCTION

This Training Manual assumes prior knowledge of Nuclear Theory. It extends this information into a discussion of Reactor Physics, particularly as it relates to CANDU reactors.

The course begins with the general principles of reactor configuration required to maintain a self-sustaining chain reaction. It continues with reactor dynamics (in both the critical and subcritical core), reactivity feedback effects (temperature effects, fission product poisoning, and fuel burnup), and ends with operational considerations (at low and high power).

The material covers four main areas, subdivided into eight major sections as follows:

- The Critical Reactor at Steady Power Output (Section 1)
- The Dynamic Reactor (Sections 2 and 3)
- Reactivity Feedback Effects (Sections 4, 5, and 6)
- Reactor Operations (Sections 7 and 8)

TABLE OF CONTENTS

Objectives 1

The Critical Reactor at Steady Power Output 9

1.0 INTRODUCTION 9

1.1 Fission 9

1.2 Harnessing Fission 15

1.3 Movement of Neutrons Through the CANDU lattice Lattice 18

1.4 The Finite Reactor 26

Response of The Critical Reactor to a Reactivity Change 31

2.0 INTRODUCTION 31

2.1 Exponential Power Rise 31

2.2 Corrections to Exponential Reactor Response 33

2.3 The Effect of Delayed Neutrons 36

2.4 Prompt Criticality 40

2.5 Power Rundown: The Prompt Drop 42

Responsiveness of The Subcritical Reactor 45

3.0 INTRODUCTION 45

3.1 Neutron Flux in a "Shut Down" Reactor 45

3.2 Dynamics in the Subcritical Core 50

3.3 Examples 55

Effects of Temperature and Voiding on Core Reactivity 61

4.0 INTRODUCTION 61

4.1 Feedback—An Introduction to Temperature Effects e 61

4.2 The Physical Basis for Temperature Coefficients 63

4.3 Temperature Coefficients of Reactivity 72

4.4 Reactivity Variation with Temperature 77

4.5 Void Reactivity 79

Effects of Fission Products on Core Reactivity 85

5.0 INTRODUCTION 85

5.1 Xenon and Iodine Buildup 85

5.2 Transient Xenon Behaviour 93

5.3 Xenon Oscillations *103*
5.4 Samarium-149 *106*

Effects of Fuel Irradiation and On-Power Fuelling on Core Reactivity 111
6.0 INTRODUCTION *111*
6.1 On-Power Fuelling *111*
6.2 Fuel Burnup—General *112*
6.3 Transient Reactivity Changes *115*
6.4 Long-Term Reactivity Effects *115*

Reactor Operations at Low Power 125
7.0 INTRODUCTION *125*
7.1 Thermal Power, Neuron Power, AND Fission Power *125*
7.2 Reactor Power Rundown *126*
7.3 The Shutdown State *133*
7.4 Approach to Critical *135*
7.5 Low Power Operation Following Startup *142*

Reactor Operations at High Power 145
8.0 INTRODUCTION *145*
8.1 Flux Flattening *145*
8.2 Flux Shape Details *150*
8.3 High Power Protection *161*

OBJECTIVES

At the end of training, the participants will be able to:

The Critical Reactor at Steady Power Output

List the reaction products of the fission process and for each describe its importance for CANDU operation.

1.0 Describe the characteristics of fission products with respect to yield, stability, radiation hazard, delayed neutron production, and capability to absorb neutrons.

2.0 Identify the different energy contributions that make up the approximately 200 MeV per fission deposited in the reactor.

3.0 Define the terms:
- unit cell multiplication factor,
- effective neutron multiplication factor.

4.0 Describe the neutron life cycle in terms of the following neutron processes:
- fast fission of U-238,
- resonance absorption in U-238,
- absorption in non-fuel core materials,
- thermal fission following absorption in fuel,
- neutron leakage.

5.0 Define reactivity and give its units of measurement.

6.0 Define the terms:
- reactivity worth,
- excess core reactivity,
- nominal core,
- control reactivity.

Response of the Critical Reactor to a Reactivity Change

7.0 Define log rate and reactor period and state their relationship

8.0 Describe the response of a CANDU reactor at low power to a small step insertion of positive reactivity

9.0 Explain the effect of delayed neutrons on reactor control.

10.0 Define the following and state how they arise:
- prompt jump,
- prompt drop.

11.0 Describe the change in reactor power following a positive reactivity insertion large enough to cause a prompt jump and state why power first rises rapidly and then increases more slowly.

12.0 Describe the change in reactor power following a large negative reactivity insertion and state why the power first drops rapidly and then decreases more slowly.

13.0 Define prompt criticality, explain how it arises, and state the approximate reactivity insertion required to cause prompt criticality.

Responsiveness of the Subcritical Reactor

14.0 Define the subcritical multiplication factor.

15.0 Explain how subcritical multiplication of a neutron source in a subcritical core causes:
- an observable, steady power level that is larger than the source,
- a change in power level after a reactivity change that leaves the core subcritical.

- Describe the rate of response to a reactivity change in a subcritical core.

- Indicate how and why the dynamic response changes in going from a deeply subcritical reactor to an almost critical reactor.

Effects of Temperature and Voiding on Core Reactivity

- Define temperature coefficients of reactivity for the:
  - fuel,
  - moderator, and
  - heat transport coolant.

- State the typical operating temperatures for the fuel, moderator, and coolant and the approximate range of temperatures encountered in going from cold shutdown to full power.

- Explain how thermal expansion of the moderator affects:
  - neutron path lengths in the moderator,
  - leakage from the core.

- Explain how increased molecular speeds caused by heating affect the:
  - resonance Absorption in U-238 (Doppler broadening and self shielding),
  - thermal neutron spectrum, and
  - thermal neutron path lengths.

- Describe how changes in the thermal neutron temperature affect absorption in:
  - U-235,
  - Pu-239.

16.0 Explain the effect on reactivity caused by a change in temperature of the:
- fuel,
- moderator,
- coolant.

- Given typical values of moderator, HT coolant and fuel reactivity coefficients, calculate the reactivity change that would occur for typical unit operations, including:
  - HT system warm up from cold to zero power hot,
  - increase in unit power from zero to full power, and
  - decrease in unit power from full power to zero.

- Define Power Coefficient

- Compare the sizes of the reactivity changes due to the moderator, coolant, and fuel for a given change in reactor power with reference to:
  - temperature coefficient,
  - temperature change,
  - time for the effect to show up.

Describe the effect of a typical CANDU power coefficient on:
- normal regulation,
- a power transient following an upset.

- Define the term “void reactivity”.

- Explain how voiding of the coolant simultaneously increases fast fission and decreases resonance capture.

- Describe how the thermal neutron temperature changes on coolant voiding and state the effects on fission rate.

- Explain how void reactivity leads to a lower limit on coolant isotopic.

- Explain the upper limit on heat transport isotopic.

Effects of Fission Products on Core Reactivity

- Define the term, fission product poison.

- State the characteristics of xenon-135 that make it a significant fission product for reactor operations.

- Describe how nuclear processes in the fuel:
  - produce I-135 and Xe-135,
  - remove I-135 and Xe-135.
- Define xenon load and iodine load.
- Explain the given curve shapes for:
  - Iodine load vs. time,
  - Xenon load vs. time.

  for start up after a long shutdown.

17.0 Describe xenon buildup following a trip.

18.0 State the operational problem xenon causes when the reactor trips.

- Explain the following terms:
  - poison out,
  - poison prevent operation,
  - poison override capability,
  - poison override time,
  - decision and action time.

19.0 Explain the following features of the reactivity change following a trip from full power:
- initial rate of xenon buildup,
- the peak,
- eventual decrease in xenon concentration.

20.0 State the approximate time for xenon to peak and the approximate reactivity worth of xenon at the peak for a reactor trip from full power with equilibrium fuel.

21.0 Compare the sizes of the transient peaks following a reactor trip from equilibrium conditions for a trip from full power and a trip from a lower power level.

- Explain the power transient that follows a change in reactor power in the high power range.
- Describe how the reactivity transient is counteracted after a return to power:
  - before a poison out,
  - following a poison out.
- Define the terms:
  - xenon oscillation,
  - flux tilt.
- State why large oscillations are unacceptable
- Describe how oscillations are controlled in CANDU reactors.

- Explain how a small, local, reactivity change can cause a large flux tilt in a reactor operating at high power without adequate spatial control.
- Describe how a flux tilt changes with time if left to itself.
- Explain why an uncontrolled oscillation may continue indefinitely.
- State the characteristics of the samarium-149 that make it a fission product poison.
- Describe how nuclear processes in the fuel:
  - produce Pm-149 and Sm-149,
  - remove Pm-149 and Sm-149.
- Compare the operational effects of samarium with those of xenon with respect to:
  - initial build up,
  - transient following a trip or shutdown,
  - return to equilibrium following a restart,
  - transients on power changes.

Effects of Fuel Irradiation and On-Power Fuelling on Core Reactivity

- List the chief characteristics associated with using on-power fuelling for maintaining core reactivity.
- Define the terms:
  - fresh fuel,
  - fuel burnup, and
  - equilibrium-fuelled reactor.
- State and explain the fuel burnup units

22.0 Describe the transient change in Pu-239 following a shutdown from high power, and following return to high power after a shutdown.

- Define the terms “saturating fission product” and “non-saturating fission product” and compare their long-term effects on the reactivity worth of a fuel bundle.
- Describe the changes that occur in the composition of a fuel bundle as it is exposed to neutron flux in the core.

23.0 Given a graph showing the reactivity change of a fuel bundle with irradiation, explain the shape of the graph in terms of:
- U-235 burnup and Pu-239 growth,
- buildup of Pu-240 and Pu-241,
- increasing fission products.

Reactor Operations at Low Power

24.0 State and explain the reasons for non-linearity between changes in neutron power and changes in reactor thermal power.

Explain the size and duration of the prompt drop given a curve of neutron flux decrease following a trip.

Describe how prompt neutrons, delayed neutrons, photoneutrons, and spontaneous fission neutrons affect the shape of the rundown curve.

25.0 Explain the following differences between the thermal power and neutron power rundown curves following a trip:
- initial rate of drop,
- duration,
- cooling requirements.

26.0 State the approximate value of decay heat at full power, and at 3 minutes and 60 minutes after a trip from full power.

- Identify the reactivity changes that occur in a reactor after it is shutdown from extended operation at power and, for each change give the:
  - sign,
  - approximate size,
  - time scale.

27.0 Describe the variation in neutron flux while the reactor is in a shutdown condition for a long period of time.

28.0 List the parameters specifically related to criticality that are monitored and controlled during the approach to critical.

29.0 List reactivity control mechanisms required during the approach to criticality.

30.0 Explain (for step reactivity increases) the observed changes in:

- stable count rate,
- flux detector time response,

as the core is taken from deeply subcritical to almost critical.

- Predict the power level following a specified change of poison concentration, given the subcritical multiplication formula, the measured power, and the poison concentration relative to criticality.

31.0 Explain why start-up procedures require monitoring of neutron flux during start-up and do not depend solely on criticality predictions.

32.0 List parameters that should be monitored and controlled on reaching criticality.

33.0 Explain in general terms how a reactor, critical at low power, could become subcritical:
- if held in the low power "critical" state for a long time,
- following power manoeuvres at low power.

34.0 Describe the power response of a slightly supercritical reactor at low power.

35.0 Explain the change in power response as power approaches the last decade of reactor power.

- Describe how reactor power is increased to rated power from hot shutdown.

36.0 List the reactivity effects that occur as power rises and state how they are expected to change as power increases.

Reactor Operations at High Power

37.0 Explain how flux-flattening permits increased reactor power output without an unsafe increase in peak bundle or channel power.

38.0 Describe and explain how flux flattening is achieved with respect to:
- the overall flux shape,
- relative zone to zone flux shape,
- control of local flux peaks.

39.0 Explain how a change of reactivity device configuration can cause flux peaks in the core and how this can affect bundle and channel powers.

40.0 Define the terms:
- reactivity device worth,
- differential reactivity device worth.

41.0 Explain the variation in rate of reactivity insertion given a curve of rod worth vs. position.

42.0 Describe the response of the liquid zones while the adjusters are driving from full in to full out.

43.0 Describe the effect on flux shape of replacing high burnup fuel with fresh fuel while operating at power.

44.0 Explain how shutdown fuelling could lead to an unacceptable spatial flux distribution on restart.

45.0 Describe what the "Reference Flux Shape" means and explain why the flux shape in an operating CANDU differs from the reference shape.

46.0 Define the following terms as they apply on an operating CANDU reactor:
- fuelling ripple,
- channel power peaking factor (CPPF).

47.0 State why bulk power is limited when the reactor is operating with adjusters out of core.

48.0 State how the reactor is protected from excessive high power when operating in an unanalyzed flux shape.

49.0 State the approximate power levels above which reactor operation is affected by:
- transient xenon,
- xenon oscillations.

THE CRITICAL REACTOR AT STEADY POWER OUTPUT

Notes:

1.0 INTRODUCTION

The defining characteristic of a nuclear reactor is the fission chain reaction. When an atom such as uranium fissions (splits) it releases energy, and the amount it releases, per reaction, is huge compared to a chemical reaction. This translates into a power source that has relatively low fuel cost. A crucial feature of the fission reaction is that it generates neutrons that can cause fission. When at least one fission neutron per fission survives to cause fission, the chain reaction sustains itself without any help from an outside source of neutrons; it is a self-sustaining chain reaction.

A nuclear reactor is a mixture of nuclear fuel, moderator, coolant, and structural material arranged in a suitable configuration of sufficient size, to support a controllable, self-sustaining chain reaction that releases useful energy. The fission reaction also produces radiation and radioactive by-products that present an immediate hazard. This hazard defines many characteristics of reactor design and operation.

Natural uranium (UO_2) fuel and heavy water (D_2O) moderator characterize CANDU reactors. (The name derives from the words CANada, Deuterium, and Uranium). Requirements related to these design choices resulted in a pressure-tube reactor with on power fuelling: a reactor that is neutron efficient (wastes very few neutrons), with a well-thermalized neutron spectrum. This section reviews the fission process and examines the balance between neutron production and loss in the reactor core that allow for a sustained rate of fission. It then describes how the CANDU characteristics affect neutrons as they travel through the CANDU lattice; effects that help explain features of day-to-day operation described in later sections.

To discuss nuclear processes that go on in a reactor, the trainee must be acquainted with some basic nuclear theory about the structure of the atomic nucleus and the principles governing the release and deposition of energy in nuclear processes such as radioactivity and fission. These notes assume that previous courses have presented the required background information, including the description of atoms, nuclei, atomic particles, mass-energy equivalence, and the interaction of radiation with matter.

1.1 FISSION

Nuclear power production relies on the fact that some nuclei will fission, releasing energy (about 200 MeV per fission) because, in splitting into more tightly bound fragments, some of the nuclear mass becomes energy.

Each watt of power requires about 3.1×10^{10} fissions/s. For a CANDU reactor operating at 2250 MW, the required fission rate is a steady:

$$2250 \times 10^{6} \times 3.1 \times 10^{10} = 7 \times 10^{19} \text{ fissions/s}$$

Notes:

The fission process provides the neutrons it needs to keep itself going. Each fission releases, on average, about 2.5 neutrons. A steady fission rate requires that neutron losses by capture (radiative capture in fuel or other materials) or by leakage (escape from the reactor) not exceed about 1.5 neutrons per fission.

Reaction Products of the Fission Reaction

The particular mode of fission illustrated in Figure 1.1 is one of many possible ways in which the nucleus may divide. The fission fragments Xe-140 and Sr-96 are two of about 300 nuclides that fission produces.

Initially the fission fragments are highly excited and get rid of excitation energy by emitting neutrons and gamma rays. Most of the fission neutrons are emitted almost immediately (~10^{-17} s) after the fission takes place; these are known as prompt fission neutrons, while the gamma rays are known as prompt gammas.

Following de-excitation, a nucleus formed by fission does not have sufficient excitation energy to throw off any more neutrons. It usually decays by beta emission; Figure 1.1 shows an example decay chain. In a few cases, the daughter nucleus created by beta decay has sufficient energy to emit a neutron. This process results in approximately ½% or so of the fission neutrons showing up long after the fission event. These are the delayed fission neutrons.

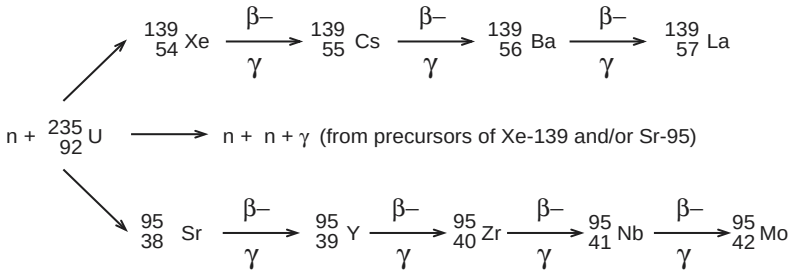

Figure 1.1
Fission Product Decay Chains

The standard notation for the particular fission just described is as follows:

$$n + {}^{235}_{92}U \rightarrow {}^{236}_{92}U^{*} \rightarrow {}^{139}_{54}Xe + {}^{95}_{38}Sr + 2n + \gamma$$

After U-235 absorbs the neutron, the ***compound nucleus*** ($^{236}_{92}U^{*}$) survives for about 10^{-14} s, and then fission occurs about 85% of the time. (Radiative capture occurs about 15% of the time). The equation displays the immediate break up, with prompt neutrons and prompt gammas. Leaving out the compound nucleus simplifies the notation:

Notes:

$$n + {}^{235}_{92}U \rightarrow {}^{139}_{54}Xe + {}^{95}_{38}Sr + 2n + \gamma$$

The energy release per fission, averaged over all fission reactions in a reactor core, is about 200 MeV (corresponding to a mass loss of nearly 0.2 u). This is, coincidently, approximately the energy per fission deposited in the reactor core. Individual fission reactions usually do not deviate from 200 MeV by more than a couple MeV, making it a convenient benchmark for later calculations. About 180 MeV of the energy release occurs at the moment of fission, the moment displayed in the notation just given. The beta emissions that follow release the remaining energy. Table 1.1 shows the distribution of energy released over the reaction products, and shows the energy deposition in the core.

	Energy Released	Energy Deposited in Core	
Energy Source	**Energy (MeV)**	**Energy (MeV)**	**Energy (%)**
Kinetic energy of lighter fission fragment	100 MeV	100 MeV	50%
Kinetic energy of heavier fission fragment	67 MeV	67 MeV	33.5%
Kinetic energy of fission neutrons	5 MeV	5 MeV	2.5%
Energy of prompt gamma rays	6 MeV	6 MeV	3%
Radiative capture of excess fission neutrons	—	8 MeV	4%
Beta particle energy gradually released from fission products	22 MeV (shared by β^- particles, gamma rays and neutrinos)	8 MeV	4%
Gamma ray energy gradually released from fission products		6 MeV	3%
Neutrinos		—	—
Total	200 MeV	200 MeV	100%

Table 1.1
Approximate Energy Distribution Resulting from a Typical U-235 Fission

Notes:

The number of neutrons emitted following fission is variable, but the most probable yields are two or three neutrons per fission. The average number of neutrons emitted per thermal fission is an important quantity in reactor physics. The Greek letter ν (pronounced "new") usually denotes it. Table 1.2 summarizes the values of ν for the fissile isotopes. Fast fission produces marginally more fission neutrons than thermal fission. These notes assume that the average number of neutrons per thermal fission in an equilibrium fuelled reactor is $\nu = 2.5$.

Notes:

Fissile Nucleus	ν
U-235	2.43
Pu-239	2.89
Pu-241	2.93
Equilibrium Fuel	~ 2.5

Table 1.2
Value of ν for thermal fissions

Fission Fragments and Fission Products

The name "fission products" applies to all members of the decay chains—including the original fission fragments. The term fragment describes the initial flying pieces that de-excite before they come to rest as fission products. The majority of fission products have half-lives that range from fractions of a second to about 30 years.

The kinetic energy of the fission fragments accounts for most of the energy released in the fission reaction. The energetic fragments slow and stop in about 10^{-13} s, (after travelling about 10^{-3} mm), heating the fuel by passing their energy to the surrounding fuel atoms.

Each fission product consists of a substantial piece of the original nucleus. They are likely to have mass numbers between 70 and 160, with mass numbers near 95 and 140 the most probable. Symmetrical fission (equal fragments) is rare.

Fission products are a potential radiological hazard. They also absorb neutrons that might otherwise cause fission, limiting the long-term energy output from the fuel. They have one important useful effect; reactor control would not be possible without the delayed neutrons that result from fission product decay. The next few paragraphs discuss these highlighted characteristics one at a time.

The Radiological Hazard from Fission Products

The fission fragments are almost invariably unstable (radioactive). The reason for this is that the neutron/proton ratio of the fragments is about the same as that of the fissioned nucleus, and this is too high for stability at medium mass numbers. Standard beta decay notation records the subsequent beta decays.

For example, following one chain of Figure 1.1:

$$^{139}_{54}\text{Xe} \rightarrow ^{139}_{55}\text{Cs} + \beta^- + \gamma \qquad \left(t_{\frac{1}{2}} = 40 \text{ s}\right)$$

$$^{139}_{55}\text{Cs} \rightarrow ^{139}_{56}\text{Ba} + \beta^- + \gamma \qquad \left(t_{\frac{1}{2}} = 9.3 \text{ m}\right)$$

$$^{139}_{56}\text{Ba} \rightarrow ^{139}_{57}\text{La} + \beta^- + \gamma \qquad \left(t_{\frac{1}{2}} = 1.4 \text{ h}\right)$$

Notes:

The ceramic fuel material (UO_2) and fuel sheathing must encase the fission products to prevent them entering the heat transport system and leaving the reactor core. (There are additional barriers in place to limit release to the public if the piping fails.) As long as the fission products remain in the fuel and the fuel remains adequately shielded, there is no biological risk. Fission products in the heat transport system (or elsewhere outside the core) are a radiation hazard that would prevent equipment access even with the reactor shut down.

Heavy shielding is required around the reactor for protection against the prompt radiation (neutrons and gamma rays) during operation. It also limits exposure to gamma radiation that continues to be emitted by the fission products after shutdown.

Fuel must be replaced remotely; special precautions must be taken in handling, and storing spent fuel.

Neutron Absorption by Fission Products

Some of the fission products have high neutron absorption cross-sections and thereby poison the reactor. A relatively high percentage of fissions produce the two most important poisons, Xe-135 and Sm-149, and these fission products capture a significant number of neutrons. Section 5 examines the effects of fission product poisons.

Delayed and Photoneutrons from Fission Product Decay

Fission product decay sometimes produces delayed neutron precursors, which in decaying produce a small fraction of the neutrons in the reactor core. Energetic gamma rays from decaying fission products may occasionally produce photo-neutrons. Fission neutrons make possible a self-sustaining chain reaction, in which the fission neutrons from one generation of fissions cause the next generation of fissions. Delayed neutrons make reactor control possible (See Section 2). Photoneutrons affect low power operation (See Section 3).

The Chart of Nuclides identifies nearly one hundred delayed neutron precursors. For thermal fission of U-235, the delayed neutrons are only 0.70% of all fission

neutrons. This fraction, represented by the symbol β, is called the delayed neutron fraction. For Pu-239, the delayed neutron fraction is even smaller, at β = 0.23%.

Notes:

For equilibrium CANDU fuel, the effective average delayed neutron fraction is close to β ~ 0.50%. Equilibrium fuel accounts for fissions of all fissile isotopes in the core, as well as fast fission of U-238. The weighted average depends on the particular mix of isotopes in the reactor, a mix that changes with fuel burnup.

A set of six artificial delayed neutron precursors accurately models the behaviour of all the delayed neutrons. These notes make the simplifying assumption of a single group of delayed neutrons with β = 0.5% for equilibrium fuel, and an average delay time t_m = 12.5 s (equivalent to $t_{½}$ = 8.7 s and λ = 0.08 s^{-1}).

1.2 HARNESSING FISSION

A single fission reaction releases about 40 million times as much energy as the release of chemical energy that comes from burning a single carbon atom. However, assembling the reactor materials in a way that allows a continuous, controllable extraction of energy requires careful design. Neutrons from fission are essential to maintain a chain reaction and steady power production in the core. With ν = 2.5, the reactor design must limit neutron losses to 1.5 neutrons per fission. Absorptions that do not cause fission and neutrons that escape into the shielding make it difficult to arrange the survival of at least one neutron from each fission reaction.

For example, consider the possibility of a self-sustaining chain reaction in a large block of natural uranium. Thermal neutrons absorbed in a block of natural uranium have about a 55% chance of causing fission. (The ratio of the fission to absorption microscopic cross-sections for natural uranium is 4.18 b/7.58 b = 0.55). Since fission produces almost 2.5 new neutrons, this would seem to allow for a continuous chain reaction. However, fast neutrons produced by fission are less likely to cause fission than thermal neutrons. (In U-235, the fission cross-section for thermal neutrons is 580 barns, compared to one or two barns for fast fission).

Now consider what happens as a fast neutron slows down by successive scattering reactions in the block of uranium. The energy loss per collision is very small, because the uranium nuclei are so heavy. Before its kinetic energy reduces to thermal energies, the neutron will reach an energy range where the absorption cross-section of U-238 is very high and it will be absorbed uselessly. Even ignoring neutron leakage or absorption in impurities (including absorption in fission products from those fissions that do occur), the resonance capture loss makes it impossible to achieve a chain reaction in pure natural uranium.

One solution is to mix the fuel with a moderator that slows the neutrons quickly through the resonance energy range, without absorbing them once they are thermalized. For a natural uranium reactor using uranium dioxide fuel, the only

possible moderator is heavy water. The choice of the relative volumes of fuel and moderator and their geometric arrangement determines the neutron efficiency.

Notes:

To follow the fate of a neutron as it bounces randomly through a mixture of reactor materials we need a quantitative comparison of the possible reactions. Most convenient is a comparison of reaction rates, which depend on nuclear cross-sections and neutron flux, quantities we will now introduce.

Macroscopic Cross-Section

The rate at which a reaction takes place in bulk material depends on not only the reaction probability with a particular nucleus (represented by the microscopic cross-section) but also on how many target nuclei there are. To account for this we introduce a quantity known as the macroscopic cross-section (Σ—upper case sigma), which combines the microscopic cross-section (σ) of an individual nucleus of a given material with the number density, N**,** (that is, the number of atoms per cm^3) of the material in the region considered.

The macroscopic cross-section is given by:

$$\Sigma = N\sigma$$

The unit for Σ is $cm^2/cm^3 = cm^{-1}$ (taking N in units of cm^{-3} and σ in units of cm^2). It is more difficult to visualize the physical significance of the macroscopic cross-section with its rather odd units than it is for the microscopic cross-section, which is simply an area (although a somewhat artificial one). One way to think of macroscopic cross-section is as the total apparent area of all the nuclear targets in one cm^3 of material. (The blacker the sky is with birds, the less distance a shotgun pellet is likely to travel into the flock before it makes a hit.)

The inverse of the macroscopic cross section for a given material, which has the dimensions of a distance, does have an easily visualized meaning. For example, the quantity $1/\Sigma_a$ equals the average distance that a neutron will travel before being absorbed by the material, and is known as the absorption mean-free-path, (mfp). Similarly, the inverse of the macroscopic scattering cross-section, $1/\Sigma_s$, is equal to the average distance traveled by a neutron between scattering collisions.

Neutron Flux

The neutron flux (ϕ) measures the intensity of neutrons passing through a cubic centimetre of material. It is given by:

$$\phi = nv$$

where n is the density of neutrons (the number of neutrons per cm^3) and v is the speed of the neutrons. The units of flux are:

$$\frac{neutrons}{cm^3} \times \frac{cm}{s} \quad or \quad cm^{-2}s^{-1}$$

Notes:

Unlike a shotgun, which produces a single "pulse" of pellets that travel together in more-or-less straight, parallel lines, neutrons travel continuously through the one cubic centimetre volume in random directions. The chance of a hit depends on each of the factors n and v. Increasing the number of neutrons in the one cm^3 volume makes it more likely that one of them will hit a target. Increasing neutron speeds increases the number of visits in the one cm^3 volume each second.

In physical terms, the quantity ϕ is the total distance traveled in one second by all the neutrons in the one cm^3 volume, since it is obtained by multiplying the number of neutrons in that cm^3 by the speed each is travelling. This is equivalent to the total length of all the neutron tracks laid down in one cm^3 in one second. Visualize it, if you like, as a time exposure of car taillights taken at night; each light shows up as a track.

The expression for neutron flux applies to any neutron energy. When applied to thermal neutrons the product is known as the thermal neutron flux.

Reaction Rate

The reaction rate, R, is the number of reactions per second per cubic centimetre of material. Consider, for example, a thermal neutron absorption reaction with microscopic absorption cross section σ_a representing the apparent absorption size of each target nucleus. The macroscopic cross section ($\Sigma_a = N\sigma_a$) is the total apparent flat area of targets that a neutron "sees" as it approaches the one cm^3 volume. The mean free path to absorption is $mfp = \Sigma^{-1}$. The flux (ϕ) is the total neutron track length laid down in one second in one cm^3, so dividing flux by the length of track required (on average) for one absorption, we get the total number of absorptions, that is:

$$R_a = \frac{total\ \ track\ \ length\ \left(per\ s\ per\ cm^3\right)}{neutron\ \ mean\ \ free\ \ path\ \ to\ absorption}$$

$$so, \qquad R_a = \frac{\phi}{mfp} \quad or \quad R_a = \phi\Sigma_a$$

With ϕ in $cm^{-2}\ s^{-1}$ and Σ_a in cm^{-1}, R has units $cm^{-3}\ s^{-1}$.

Although we used the thermal neutron absorption reaction for illustration, the result is perfectly general, for any specified reaction and for neutrons of any energy:

$$R = \phi\Sigma$$

By far the majority of neutrons in a CANDU core are thermal neutrons, and rates of reactions induced by these neutrons are directly proportional to both the thermal neutron flux and the macroscopic cross section for the reaction. Quoted values of cross-section or flux are always for thermal neutrons unless otherwise stated.

Notes:

Tables of microscopic cross-sections usually give values for neutrons travelling at 2200 m/s, that is, with kinetic energy ($\frac{1}{2}mv^2$) of 0.0253 eV. This is the most likely speed for neutrons thermalized at 20°C. Assuming that the temperatures of CANDU components are all at 20°C is obviously not physically realistic. Corrections to the reaction rates for realistic neutron temperatures become important when we discuss temperature effects in Section 4. As we shall see, neutron temperature does not strongly affect reaction rates in most reactor materials. Reaction rates in the fissile isotopes are a notable exception.

1.3 MOVEMENT OF NEUTRONS THROUGH THE CANDU LATTICE

Figure 1.2 shows a neutron following a zigzag path as its direction of motion changes randomly with each scattering collision. The straight portions of path between two scattering collisions are usually about 1 to 3 cm in length in most materials. For reference, the CANDU lattice pitch (centre-to-centre distance of nearest neighbour channels) is 28.6 cm. The fuel bundle diameter (or pressure tube inside diameter) is about 10 cm, and the bundle length is close to 0.5 metres.

The fate of a neutron depends on the relative probabilities of the scattering and absorption reactions for the various nuclei it encounters as it diffuses through the CANDU lattice. These probabilities depend on the relative abundance of the different nuclei and the cross-sections for particular reactions.

CANDU fuel begins as uranium dioxide with 99.28% U-238 and 0.72% U-235 (the natural abundance). This new fuel is known as fresh fuel**.** An equilibrium-fuelled reactor, in which day-to-day operation includes routine replacement of "burned up" fuel with fresh fuel, contains both uranium and plutonium. Fuel with a mix of isotopes comparable to the equilibrium-fuelled core average is called equilibrium fuel.

In equilibrium, fuel there will be thermal fissions (fissions induced by thermal neutrons) in U-235, Pu-239, and Pu-241, and fast fissions (fissions induced by fast neutrons) in U-238. There is a lot of U-238 in the core (more than 100 tonnes in a large CANDU, compared to less than 1 tonne of the fissile isotopes) but U-238 can fission only if the neutron kinetic energy is greater than about 1.2 MeV (the threshold energy for the reaction). There are far fewer fast neutrons than thermal, so there is relatively little fast fission. (U-238 fast fission contributes less than 3% of CANDU power production).

Neutrons of any energy can initiate fission in the fissile isotopes (this is the meaning of the adjective "fissile"—there is no threshold energy for fission in U-235, Pu-239, or Pu-241), but the neutron flux and cross-sections are both small for high-energy neutrons. We can ignore fissions in the reactor core other than thermal fission of U-235, Pu-239, and Pu-241, and fast fission of U-238.

Notes:

CANDU moderator is high isotopic heavy water, typically 99.8% or more D_2O, with 0.2% or less H_2O impurity. Similarly, the isotopic of the heavy water coolant that surrounds the fuel elements in the channels is near 99.8 % isotopic, but contains slightly more light water impurity than the moderator contains.

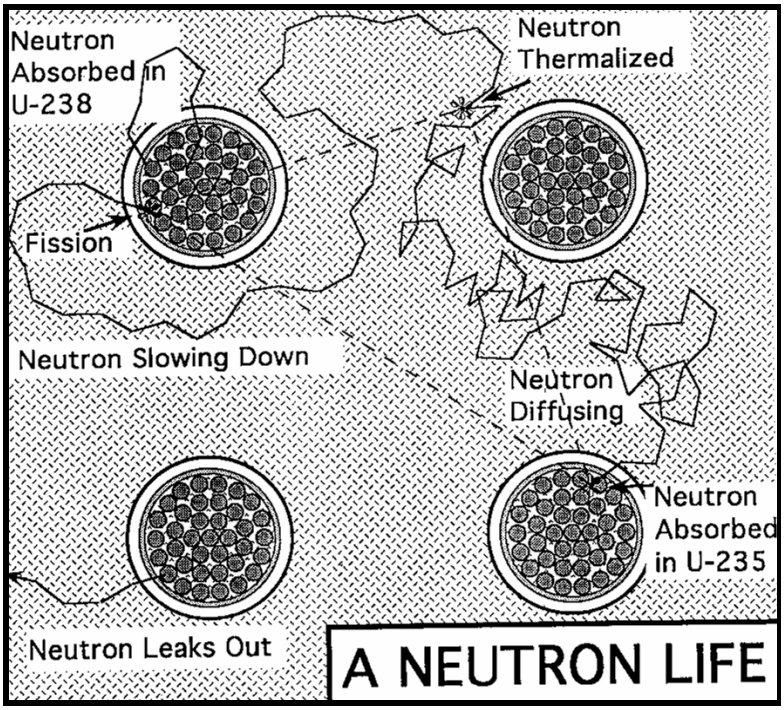

Figure 1.2
Typical Path of a Neutron from Birth to Absorption

In Figure 1.2, elastic collisions with nuclei in the moderator slow a neutron born at the point labelled "fission" until it reaches thermal energy at the point labelled "neutron thermalized". In a CANDU reactor, the average distance between these points is about 25 cm and the time is a few microseconds.

When a fast neutron collides with a nucleus, the fraction of the neutron's kinetic energy transferred to the recoiling target nucleus depends on the mass of the target nucleus and the angle of collision. The average energy transferred on each collision in the moderator is a fixed fraction of the neutron kinetic energy. In an elastic collision between a neutron and a heavy water nucleus, the neutron transfers 40% of its energy, on average, to the target nucleus. Elastic collisions

conserve kinetic energy, so the neutron emerges from each "average" collision with 60% of the energy it started with. To estimate the number of collisions that will thermalize a fission neutron, assume the fission neutron starts with a kinetic energy of 2 MeV, loses 60% of its energy on each collision, and is "thermal" when it has kinetic energy of 0.025 eV. It takes about 36 collisions to thermalize a fast neutron [$0.6^{36} \times (2 \times 10^6) \sim 0.025$].

Notes:

Following thermalization, the neutron diffuses as a thermal neutron until something, usually fuel, absorbs it (unless it reaches the edge of the reactor and leaks into the shielding). The time between thermalization and absorptions accounts for most of the neutron lifetime, about one millisecond (10^{-3} s). The average distance between birth and absorption in Figure 1.2 is almost 40 cm (close to the diagonal distance across the CANDU lattice).

The values quoted are "crow-flight" distances—the paths followed by the neutrons are considerably longer. For example, a thermal neutron moving at 2200 m/s travels a zigzag distance of 220 cm in 10^{-3} s. For a scattering mean free path of 2.9 cm, this implies approximately 75 elastic collisions while diffusing as a thermal neutron (220 cm/2.9 cm).

Figure 1.3 is a flow chart of Figure 1.2 showing four energy intervals during the neutron life cycle. Labels indicate the dominant nuclear reaction in each interval. In a critical reactor, only one neutron per fission completes the cycle, keeping the number of neutrons constant.

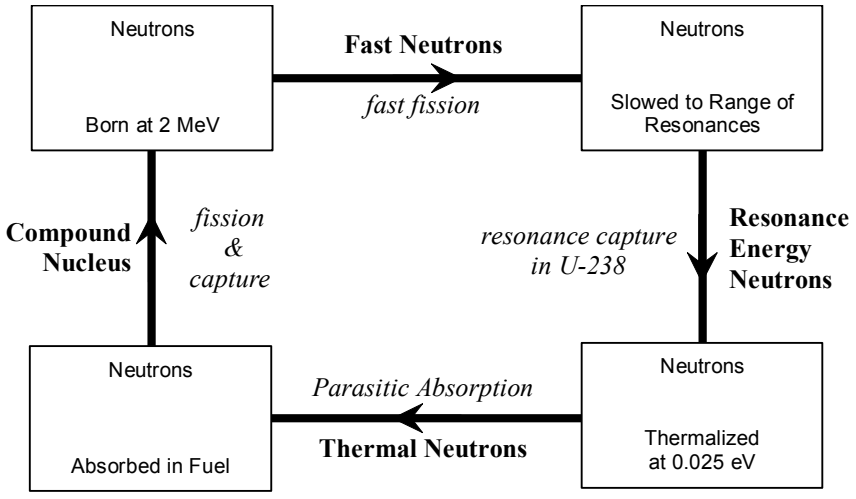

Figure 1.3
Outline of the Neutron Life Cycle

Note the relative importance of the loss mechanisms while the neutron moves through the lattice:

- resonance capture in U-238 accounts for about half of the losses;
- thermal neutron absorption in non-fuel material accounts for over one-third of the losses;
- leakage (not shown in the figure) accounts for the remaining losses.

Notes:

Careful choice of materials and their configuration keeps the net loss of neutrons (fast, intermediate, and thermal) by leakage and capture to about 20% of all the neutrons produced by fission. The remaining 80% are absorbed in the fuel as thermal neutrons. With our choice of fuel, only about half of these cause fission, just enough to return the same number as started the cycle (2.5 neutrons/fission × 40% = 100%).

The parameter "k" (the neutron multiplication factor) characterizes the chain reaction. The value k = 1 describes a reactor in which exactly one neutron from each fission (on average) causes another fission. Such a reactor sustains a steady fission rate and is said to be critical. This subsection develops an operationally useful expression for k by following a group of neutrons around the flow chart and looking at the production and loss mechanisms in each energy range. For simplicity, the discussion begins in the interior of a very large reactor (the "infinite" reactor) so that we can ignore leakage effects. The reactor comprises a large number of cells, each assumed to contain exactly the same mix of fuel, coolant, moderator, and other materials. By specifying identical cells in an infinite reactor, we can assume there is no net transfer of neutrons between cells, reducing the description to a single unit cell: a quadrant of Figure 1.2.

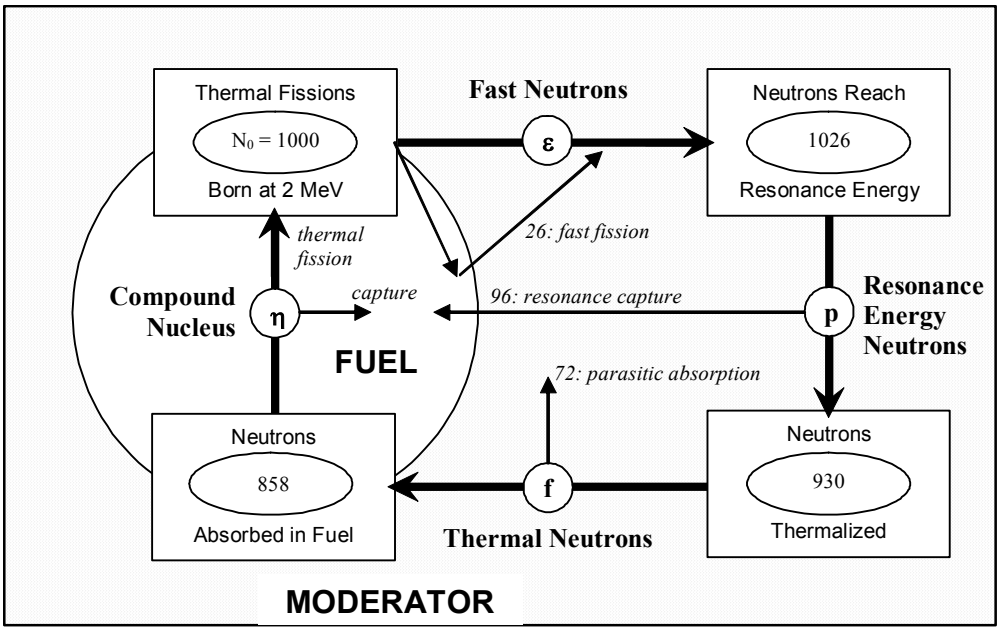

Figure 1.4
The Neutron Life Cycle in the Unit Cell
(Numbers are for a Critical Infinite Reactor, $k_\infty = 1$)

Notes:

We will now follow a group of neutrons as they diffuse through the CANDU lattice, that is, as the neutrons interacts with nuclei in the unit cell. Begin at the top left of Figure 1.4. The initial number of neutrons is $N_0 = 1000$ fast neutrons from thermal fission. Follow the heavy arrows around the closed loop as you read through the following subsections.

Fast Fission

Starting with N_0 fission neutrons, the first thing that happens is that we actually gain some neutrons, because fission neutrons leaving the fuel channel are energetic enough for some of them to cause fast fission in U-238. The U-238 fast fission cross-section is about half a barn for neutrons between 2 and 5 MeV. One or two collisions with coolant molecules reduce the energy of a fission neutron below the fast fission threshold (1.2 MeV). The designers calculate that 1000 fast neutrons escaping the fuel channel generate 26 extra neutrons from fast fission.

The definition of the fast fission factor, denoted by the symbol ε (epsilon), is

$$\varepsilon = 1 + \frac{\text{number of neutrons from fast fission}}{\text{number of neutrons from thermal fission}}$$

For the design numbers in Figure 1.4,

$\varepsilon = 1 + 26/1000 = 1.026$

The number of fast neutrons produced by all fissions (fast and thermal) in the unit cell, starting with N_0 fast neutrons from thermal fission, is εN_0 (= 1026 in our example).

Notes:

The fast fission factor (ε) depends mainly on the relative amounts of fuel and heavy water in the fuel channel. Voiding of coolant, which delays thermalization until the neutrons have left the channel, is the only operational event that significantly affects fast fission. Section 4 examines the effect of voiding the coolant.

Resonance Capture in U-238

Capture of neutrons in the U-238 resonance peaks account for the biggest loss of neutrons. Collisions with coolant nuclei can slow some fast neutrons to the resonance energy range before the neutrons leave the channel. Fast neutrons that escape into the moderator are not immune to resonance capture. Random collisions in the moderator will return some neutrons to the channel of birth or take them to nearby channels before thermalization is complete. Any resonance energy neutron that enters the fuel faces almost certain capture.

The factor that accounts for resonance capture is the resonance escape probability (p). This is the probability that a neutron will not undergo resonance capture in U-238 while slowing down.

$$p = 1 - \frac{\text{neutrons captured in } {}^{238}_{92}\text{U resonances}}{\text{neutrons reaching resonance energy}}$$

For the numbers in Figure 1.4

$p = 1 - 96/1026 = 0.906$

The number of neutrons reaching thermal energy in the unit cell, starting with εN_0 fast neutrons is $p\varepsilon N_0$ (= 930 in our example).

A two-pronged approach limits losses by resonance capture:

- the moderator steps the neutrons quickly through the resonance energy range to reduce the chance of an interaction.
- the fuel is "lumped" to shield its interior from resonance energy neutrons (an effect described as "self-shielding");

In effect, rapid slowing of the neutrons limits the flux of neutrons in the resonance energy range; self-shielding (mentioned again in Section 4) limits the number of

U-238 targets. The combined effect is to limit the rate of resonance capture, making criticality possible even with natural uranium fuel.

Section 4 examines three situations that significantly affect the resonance escape probability:

Notes:

- Increasing the fuel temperature, as occurs during a power increase, quickly and dramatically increases resonance capture.
- Increasing the moderator temperature decreases moderator density. More neutrons reach nearby channels while still in the resonance energy range, so resonance captures increases.
- Voiding of the coolant reduces the number of resonance energy neutrons available for capture, so resonance capture decreases.

Parasitic (Non-fuel) Absorption

Thermal neutrons in the unit cell are absorbed either in the fuel or in other core materials. The thermal utilization (f) is the factor that accounts for non-fuel thermal neutron absorption. It is the fraction of thermal neutrons absorbed by the fuel compared to all the thermal neutrons absorbed in the cell.

$$f = \frac{\phi(\text{fuel})\Sigma_a(\text{fuel})}{\phi(\text{fuel})\Sigma_a(\text{fuel}) + \phi(\text{non fuel})\Sigma_a(\text{non fuel})}$$

ϕ(fuel) and ϕ(non-fuel) are, respectively, the average thermal neutron fluxes in the fuel and away from the fuel. Accurate analysis requires taking into account the variation in flux across the unit cell.

A simplified formula that assumes the flux is the same everywhere suits are purposes here:

$$f = \frac{\Sigma_a(\text{fuel})}{\Sigma_a(\text{fuel}) + \Sigma_a(\text{non}\cdot\text{fuel})}$$

Using the numbers absorbed in fuel and non-fuel in Figure 1.4, we get

f = 858/930 = 0.923

Starting with $p\varepsilon N_0$ thermal neutrons, the number of thermal neutrons absorbed in the fuel is $f p\varepsilon N_0$ (= 858 in our example).

Parasitic absorption is the second most important loss mechanism (after resonance capture). The designers keep the non-fuel thermal neutron absorption to a

minimum by selecting low absorption materials for moderator, pressure tubes, and other core components. The lattice pitch is kept small enough so that thermal neutrons do not spend too much time diffusing, which would increase the probability of their absorption in the moderator.

Notes:

Thermal utilization is also the easiest factor to adjust for short-term control of the fission rate. The designers have provided an array of neutron absorbing devices intended for insertion or removal from the reactor to adjust neutron absorption.

Operating temperatures and coolant voiding affect f. Section 4 examines these effects.

Fuel Absorption: Fission and Capture

All the neutrons remaining in the unit cell at this point are absorbed in the fuel. The term "fuel" as used in this text includes all the content of the fuel bundle (fissile material, U-238, fission products, Pu-240, etc.). Not every neutron absorbed in fuel causes fission. For every neutron absorbed in equilibrium fuel, fission returns 1.2 neutrons. This is the value of the reproduction factor denoted by the symbol η (eta, pronounced "eight-ah"). Note carefully the distinction between η, the number of fission neutrons per thermal neutron absorbed in the fuel and ν, the number of neutrons produced per fission

$$\eta = \nu \times \frac{\Sigma_f(\text{fuel})}{\Sigma_a(\text{fuel})} = \frac{\nu\Sigma_f(\text{fuel})}{\Sigma_f(\text{fuel}) + \Sigma_\gamma(\text{fuel})}$$

The definition of η includes all fission neutrons, prompt and delayed. The fuel cross sections are effective averages over all the contents of the fuel. The number of neutrons per fission (ν) is an effective average over all the fissile isotopes.

Starting with $f p \varepsilon N_0$ thermal neutrons, the number of neutrons returned to start the next cycle is $N = \eta f p \varepsilon N_0$ (= 1000 in our example).

For the cycle in Figure 1.4:

$\eta = 1000/858 = 1.166$

The value of η required for criticality in a finite reactor is higher than the value just given, to offset losses by neutron leakage. In a practical reactor, fuel replacement keeps the inventory of U-235 higher and the inventory of fission products lower, maintaining η near 1.2.

The reproduction factor depends on the fuel composition. Individual fuel bundles change composition as they spend time in the neutron flux, but continual replacement of high burnup bundles with fresh ones (on power refuelling) keeps the core average composition the same from day to day.

Operationally, the rate of replacing fuel, fuel temperature changes, and coolant voiding all affect η. Section 5 examines fuel burnup effects and on power refuelling. Section 4 discusses temperature effects and voiding.

Notes:

The Lattice Cell Multiplication Factor

Summarizing the results above, we started the cycle with N_0 neutrons, and on completing the cycle, the number of neutron returned was $N = \eta f p \varepsilon N_0$. We introduce the symbol k_∞, (k-infinity, the lattice cell multiplication factor) such that $N = k_\infty N_0$ for the infinite lattice with no leakage. This gives:

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

When there is a perfect steady-state balance between neutron production and loss in the infinite lattice, $k_\infty = 1$. This balance occurs in each energy range. The production of fast neutrons (by thermal and fast fission) exactly balances the loss of fast neutrons (by slowing the fast neutrons to lower energy). The production of resonance energy neutrons (by slowing of fast neutrons) exactly balances the losses of resonance energy neutrons (by resonance capture and thermalization). The production of thermal neutrons (by thermalization) exactly balances the loss of thermal neutrons (by absorption in fuel or other reactor core materials). Criticality implies this sort of detailed balance.

Reactor engineers can calculate each of the four factors if they know the composition and geometry of the core. This allows them to compare and optimize configurations and materials. The numerical example carried throughout our discussion used design values of the factors for a critical infinite lattice, but the engineering formulas for calculating the factors do not make any prior assumptions about criticality. The four-factors and k_∞ can be evaluated for any mixture of materials, not just for $k_\infty = 1$.

1.4 THE FINITE REACTOR

Now we will extend the idea of neutron balance across a finite core, where each cell has a different composition and different flux. In the reactor, the random motion of individual neutrons results in an overall movement of neutrons from regions of higher neutron density to regions of lower density. (In the jargon, net neutron transfer depends on the flux gradient). In the finite reactor, the neutron balance in each lattice cell now includes neutrons that transfer across cell boundaries. For example, production of thermal neutrons in a cell would include production by thermalization, as before. It could also include the transfer of thermal neutrons into the cell from a neighbouring cell. Similarly, losses from the cell, in addition to those discussed for the infinite reactor, could include transfer of neutrons out of the cell.

The principle of criticality is the same; a steady-state critical core exists when neutron production balances neutron losses in each cell for neutrons in each

energy range. In the finite reactor, we allow for fuel bundles with different composition in each cell. There are also some cells with extra absorption because some bits of reactor hardware (adjuster rods, measuring instruments etc.) pass through them. There are cells in the reflector where there is no neutron production by fission, only by transfer from neighbouring cells. At the boundary, neutrons transfer into the shielding, but we assume no neutrons come back.

Notes:

For the finite core, each cell links to each neighbouring cell (and indirectly to every other cell in the reactor) because neutrons lost from one cell (in any energy range) become neutron production in a neighbouring cell. The net transfer from a cell is called cell leakage. It could be positive (transfer out of the cell), or negative (transfer in). To achieve a neutron balance within one cell means balancing neutrons in all cells. Anything that affects an individual cell inevitably affects, at least to some degree, every other cell.

The effective multiplication factor (k*)* accounts for possible net imbalance across the entire finite core. Quite clearly, we cannot put together an arbitrary mixture of fuel, moderator, reactor hardware etc. and expect removal and production will be equal. When production is greater than loss, $k > 1$ (the core is supercritical) and when production is less than loss, $k < 1$ (the core is subcritical).

Criticality occurs for a single value $k = 1$ across the core. This involves bringing neutrons in every energy group into balance at every point in the reactor simultaneously. For the situation of perfect balance (the steady-state critical reactor) mathematicians are able to determine steady state fast and thermal neutron fluxes by solving the complicated set of interconnected equations. Section 8 discusses flux shape.

There is no steady state solution for flux when $k \neq 1$. With $k > 1$ flux is increasing. For $k < 1$ flux is decreasing—if there are no neutron sources present. Section 2 introduces power change in the supercritical core and Section 3 introduces the behaviour of the subcritical core.

The Six Factor Formula

The composition of each cell in the finite reactor determines the four factors of k_∞ for each cell. By this we mean the values the four factors would have in an infinite reactor with the composition of the particular cell. Using effective average factors across the whole core, we can write the effective multiplication factor:

$$k = \eta f p \varepsilon \Lambda_t \Lambda_f$$

The factor Λ_t (lambda-t) accounts for the transfer of thermal neutrons out of the cells at the reactor boundary. It is the probability that a thermal neutron will not leak out of the reactor, so its name is thermal non-leakage probability.

The fast non-leakage probability denoted by Λ_f (lambda-f) is the factor that account for the transfer of fast neutrons out of the cells at the reactor boundary. It is the probability that a fast neutron will not leak from the reactor.

Notes:

The fast and thermal non-leakage probabilities depend on:

- How far the neutrons wander as they diffuse, (i.e., on the number of collisions and the path lengths);
- The shape and size of the reactor;
- The reflector.

The first two of these are related. In a small core, too many neutrons will reach the edge and leak out. It is not possible to achieve criticality below a particular size, known as the ***critical size***. The reflector returns some of the leakage neutrons to the core, making criticality possible with a smaller core size.

Reactivity

When k = 1 the neutron flux and the power level remain constant.[1] It is important to note that criticality depends on neutron balance, not on power level; saying that a reactor is critical is not the same as saying it is at high power.

Under normal operating conditions a reactor operates close to criticality (that is, k is nearly equal to one). It is more convenient under these circumstances to talk in terms of the amount by which k differs from one than it is to keep quoting the value of k itself.

The effective multiplication constant is defined in terms of a detailed neutron balance everywhere in the core:

$$\text{(neutron removal)} = \frac{1}{k}\ \text{(neutron production)}$$

The balance exists for k = 1. The amount by which (1/k) differs from one defines reactivity.

$$\text{Reactivity} = 1 - \frac{1}{k} = \frac{k-1}{k}$$

Reactivity is a measure of `the imbalance between neutron production and removal across the entire core.

[1] Steady power occurs with k < 1 when sources are present. We ignore sources in this section.

Because we usually deal with values of k that are close to one, the above expression is nearly equal to k-1 and the following is a satisfactory approximation:

$$\text{Reactivity} \approx k - 1 = \Delta k$$

Notes:

In this course, we usually take this absolute difference between k and one as the reactivity.

Reactor regulation is achieved with very small reactivities, so numerical values of reactivity are usually written in a notation based on the metric prefix "milli-", the milli-k (mk). The following example illustrates the use of the mk to express reactivity = Δk:

Suppose we have $k = 1.002$
then, $\Delta k = k-1 = 1.002 - 1 = 0.002$

In this case, Δk is equal to 2 parts in a thousand, which we call 2 mk:

Δk = 2 mk (Notice that Δk/k = 1.996 mk)

Conversely, for Δk = -1 mk, k = 0.999.

The notation Δk for reactivity is easily misused. The notation "Δ" is used generally to mean any difference between two values of a parameter, but in the equation for reactivity, it is the difference from a fixed reference point, the critical core. (It is k - 1, not $k_2 - k_1$).

The term reactivity worth expresses the difference between two different reactivities, and the term excess core reactivity describes a change in core reactivity from a defined configuration. The following are some common expressions associated with reactivity.

Reactivity worth describes the change in reactivity a device could introduce into the core, whether or not the core is critical. For example, the reactivity worth of the liquid zones between full and empty is about 6 mk.

Excess reactivity is the difference in reactivity between the actual core and the nominal core. The nominal core is a core configuration, defined for each station, in which each reactivity device is in its "normal" operating state. A typical nominal core would have the adjuster rods fully inserted, control absorbers and shutoff rods poised, liquid zone compartments 40% full, with reactor grade moderator (and coolant), free of dissolved poisons, and normal system temperatures.

Control reactivity is the total reactivity worth of the control devices compared to the nominal configuration. To keep the core critical, the control reactivity and the core excess reactivity must be equal and opposite.

For example, suppose the core is critical ($k = 1$, $\Delta k = 0$) in the nominal configuration and the rate of fuelling is increased to allow the fuelling machines to be removed from service. The regulating system holds power by raising the liquid zone levels to keep $k = 1$ during fuelling. Following fuelling the core has excess reactivity, even though $\Delta k = 0$. The difference in zone level from the nominal 40% measures the control reactivity, which exactly offsets the excess core reactivity.

Notes:

RESPONSE OF THE CRITICAL REACTOR TO A REACTIVITY CHANGE

Notes:

2.0 INTRODUCTION

This section deals with a topic of direct operational importance—the way the reactor responds to an imposed change in the effective multiplication constant, that is, a change in core reactivity, which measures the difference between a critical core (k = 1) and a core where neutron production and removal do not balance.

The neutron flux increases exponentially when the reactor operates with a positive reactivity (k > 1). The reactor period, which governs this rate of rise, depends on the reactivity and on the neutron lifetimes in the system. The long lifetimes of the delayed neutrons are crucially important in limiting the rate of power rise, but their effect becomes less important as the size of the reactivity insertion increases.

- For a small reactivity insertion in a reactor core, the sort used in normal regulation, the exponential rise is well controlled.
- Reactor power increases very quickly for a fraction of a second following a larger reactivity insertion. After this prompt jump, the delayed neutrons control a slower stable power rise.
- Still larger reactivity insertions make the reactor prompt critical, causing an uncontrollably fast rise of the sort that occurred at Chernobyl.

The reactor makes a smooth, gradual transition between operating regimes, but the accurate description of reactor response gets more complicated as the reactivity insertions get larger. Consequently, this section uses different models to account for the observed reactor behaviour.

Power reduction (a prompt drop followed by a gradual power rundown) is part of normal operation, but operating staff should never see a prompt jump (or prompt criticality). Reactivity insertions must be limited in size (by design and operation) to limit reactor power increase. Station staff must ensure that changes to systems or operating practice do not violate the built in constraints.

2.1 EXPONENTIAL POWER RISE

December 2, 1942 saw the first controlled, self-sustaining nuclear chain reaction. Enrico Fermi was in charge. Here is an eyewitness account:

"But suddenly, his whole face broke into a broad smile." 'The reaction is self sustaining,' he announced quietly, happily. 'The curve is exponential.'"

Notes:

Exponential power rise characterizes a supercritical nuclear reactor; power rises as follows:

$$P(t) = P_0 e^{\frac{t}{\tau}}$$

P_0 is the power at time t = 0, and power in this equation (and throughout this section) is neutron power; it does not include decay heat. Section 7 will distinguish between neutron power and reactor thermal power.

The parameter that characterizes the rate of power rise is the reactor period, τ (Greek tau). The reactor period, by definition, is the time it takes power to rise by a factor e, the base for natural logarithms (e ≈ 2.718…). A long reactor period corresponds to a slow power rise, a short period to a fast power rise.

Dividing the exponential relationship by P_0 and taking the natural logarithm gives the equivalent relationship:

$$ln\left(P/P_0\right) = t/\tau$$

Figure 2.1 shows this linear power increase with time. We also describe this behaviour by saying that the rate of change of logarithmic power (log rate in station jargon) is constant.

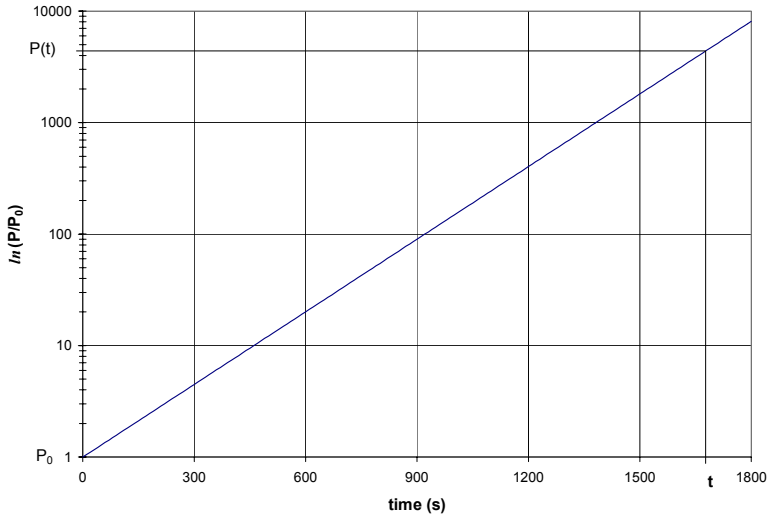

Figure 2.1
Logarithmic Power Increase for an Exponential Power Rise

Log Rate

Notes:

Using the expression above, the average rate of change of the logarithm of power over the time interval t is

$$\log \text{ rate} = \frac{\Delta ln\mathrm{P}}{\Delta \mathrm{t}} = \frac{ln(\mathrm{P}) - ln(\mathrm{P}_0)}{\mathrm{t}} = \frac{1}{\tau}$$

Using the exponential equation, the following relationships give the instantaneous rate of change[2]:

$$\log \text{ rate} = \frac{\mathrm{d}ln\mathrm{P}}{\mathrm{dt}} = \frac{1}{\mathrm{P}}\frac{\mathrm{dP}}{\mathrm{dt}} = \frac{1}{\tau}$$

By either approach, we conclude that

log rate = $1/\tau$

Whenever the time rate of change of the logarithm of some variable equals a constant, (or, equivalently, the time rate of change of the variable is proportional to the variable) integration of the equation demonstrates exponential time dependence.

The log rate represents a fractional increase per second. The rate log meters on CANDU control room panels give the rate log in units of percent present power per second (% P.P. s^{-1})

2.2 CORRECTIONS TO EXPONENTIAL REACTOR RESPONSE

The simple exponential behaviour described in the last section applies only to a slightly supercritical reactor at low power. When the reactor is operating at high power, feedback effects complicate the response to an imposed reactivity. For example, if the control system increases reactivity, power starts to rise, causing changes in the temperatures of various reactor components. This in turn affects the reactivity of the reactor. In this section, we assume that the control system inserts a reactivity Δk, and that nothing changes this original value. This is exactly the kind of behaviour observed during low power operation where feedback reactivity effects do not occur, or are small enough we can ignore them. Later sections examine feedback.

The second deviation from simple exponential behaviour is the response of the reactor to a rapid insertion of a larger reactivity. In practice, an important feature of CANDU operation is that during normal operation the automatic regulating system adjusts reactivity. The computer software cycles through a

[2] the relation $\frac{dlnP}{dt} = \frac{1}{P}\frac{dP}{dt}$ is a calculus identity.

sequence of tasks, and, once per cycle, if necessary, it initiates a small step change to move the measured power level towards the requested power level. The likely variation in reactivity is a ramp change made of a sequence of very small step changes. However, we will illustrate reactor response with a single step change, which is easier to examine than a sequence of steps.

Notes:

The Prompt Jump Approximation

Many texts introduce a model of nuclear dynamics that produces a set of differential equations with an analytic solution. In spite of the relative simplicity of the model, it quite accurately mimics more detailed modeling. The "two group" reactor dynamics theory models the neutron population with a large group of prompt neutrons and a single, small group of delayed neutrons. For a *critical* reactor with initial *steady state* power P_0, the power as a function of time, P(t), after an addition of a step reactivity making $k > 1$, is:

$$\frac{P(t)}{P_0} = \frac{\left\{ e^{\frac{t}{\tau}} - \left(\frac{\rho}{\beta}\right) e^{-\frac{t}{\tau_C}} \right\}}{\left[1 - \frac{\rho}{\beta}\right]}$$

where $\rho = \dfrac{k-1}{k} = 1 - \dfrac{1}{k}$ is the core reactivity

and $\tau = \dfrac{1 - \left(\frac{\rho}{\beta}\right)}{\lambda\left(\frac{\rho}{\beta}\right)}$ is the reactor period.

$\tau_C = \dfrac{\left(\frac{\ell}{\beta}\right)}{\left(1 - \frac{\rho}{\beta}\right)}$, where $\ell \approx 10^{-3}$ s is the time for one prompt neutron cycle.

The "transient time constant" (τ_C), is very much smaller than the reactor period, except when $\rho \approx \beta$; but the equation is not valid for $\rho \approx \beta$ anyway. This equation is a good approximation if the final reactivity (ρ), is not close to the value of β. For small reactivity insertions $k \approx 1$, as it must be for control of a critical core, so $\rho \approx \Delta k$.

In the first formula above, the final term in the curly bracket is a transient term that gives the time dependence of the initial power rise during the first fraction

of a second. Ignoring the time variation in the first fraction of a second and substituting Δk for the reactivity gives us the prompt jump approximation:

Notes:

$$P(t) = \left[\frac{\beta}{\beta - \Delta k}\right] P_0 e^{\frac{t}{\tau}}$$

The reactor period is given by

$$\tau = \left(\frac{\beta - \Delta k}{\lambda \Delta k}\right)$$

Figure 2.2 shows this power increase for a reactivity insertion of Δk = 0.5 mk. Later in this section, we will use a numerical model to show why the reactor behaves this way.

The long-term power increase in Figure 2.2 is generally similar to what we saw in section 2.1. However, on inserting the step Δk = 0.5 mk at time t = 0, the first thing seen is a rapid power increase—just over 10% (in this case).

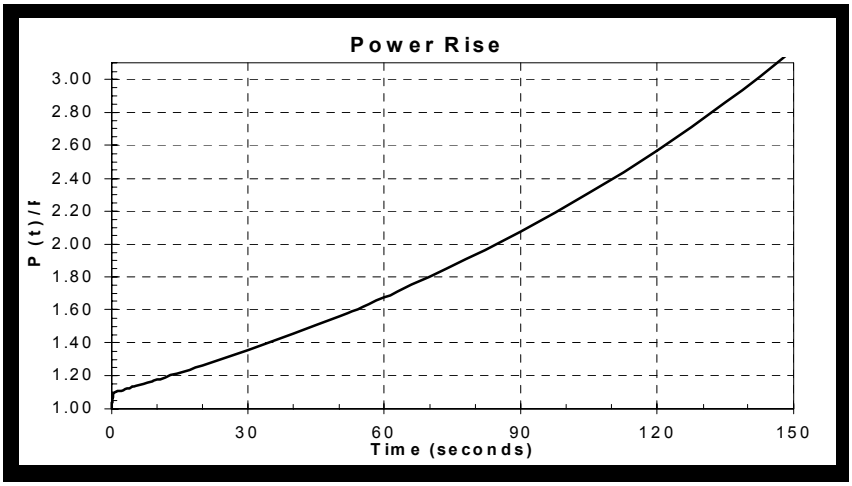

Figure 2.2
Power Rise after a Moderate Reactivity Insertion (Δk = 0.5 mk)

This initial fast rise is the prompt jump, a completely new feature. The more gradual exponential increase following the prompt jump, the stable rise, is nearly the same as the exponential rise in section 2.1.

In the figure, the size of the prompt jump is $\left[\frac{\beta}{\beta - \Delta k}\right] = \frac{0.0050}{0.0050 - 0.0005} = 1.11$

The size of the prompt jump depends on the size of the reactivity insertion (or, more precisely, on the ratio of Δk/β). Notice that the prompt jump is negligible when $\Delta k << \beta$.

because $\left[\frac{\beta}{\beta-\Delta k}\right] \rightarrow 1$ for $\Delta k << \beta$

Notes:

In the figure, the reactor period is $\tau = \left(\frac{\beta-\Delta k}{\lambda \Delta k}\right) = \frac{0.0050-0.0005}{0.08 \times 0.0005} = 112.5$ s

The reactor period also simplifies for small Δk.

$$\tau = \left(\frac{\beta-\Delta k}{\lambda \Delta k}\right) \rightarrow \frac{\beta}{\lambda \Delta k} \text{ for } \Delta k << \beta$$

In other words, when reactivity (Δk) is a very small fraction of the delayed neutron fraction (β), the prompt jump formulas revert to a simple exponential rise. The exponential power rise calculated with this formula for τ agrees with actual power rise for small reactivity additions such as those used by the CANDU regulating system (Δk < 0.1 mk). It is less accurate for step additions a little larger than this. The prompt jump approximation is adequate for larger additions as well as for small reactivity additions.

2.3 THE EFFECT OF DELAYED NEUTRONS

The reactor period, which governs the rate of stable rise, is proportional to the delayed neutron lifetime (1/λ), and also depends on the ratio of the reactivity to the delayed neutron fraction (Δk/β). Reactivity devices can adjust reactivity. Parameters β and λ change with fuel composition, but are very nearly constant in an equilibrium fuelled CANDU. The reasons why reactor control depends crucially on the delayed neutrons is examined next.

The Average Lifetime Model

The simplest model of reactor behaviour (the average lifetime model) ignores the vast lifetime difference between prompt and delayed neutrons. It assumes a neutron generation time (time from one fission to the next) that is a weighted average of the prompt and delayed neutron lifetimes. For equilibrium CANDU fuel, the generation time is approximately:

$99.5\% \times 10^{-3}$ s + 0.5% × 12.5 s = 0.000995 s + 0.0625 s = 0.063 s

The prompt neutrons have such a short lifetime compared to the delayed neutrons that the generation time is calculated accurately enough for this crude model as the product of delayed neutron fraction (β) and the average delayed neutron lifetime (1/λ).

Generation time ≈ β/λ ≈ 0.06 s

In this model, when k > 1, power rises from generation to generation, and the generation time governs the rate of rise. For a short neutron generation time, the increase occurs many times in a second. For a longer generation time, there are fewer cycles in one second and the rate of power increase is less. If all the neutrons were prompt fission neutrons, the generation time would be very short ($\ell \approx 10^{-3}$ s) and the rate of rise would be uncontrollably fast.

Notes:

The power rise during one generation (ΔP) is

$\Delta P = kP - P = P(k-1) = P\Delta k.$

The time for this power increase is the average generation time, β/λ, so the average rate of change over one cycle is

$$\frac{\Delta P}{\Delta t} = \frac{P\Delta k}{(\beta/\lambda)} = \frac{\lambda\Delta k}{\beta}P$$

Equating this average increase in a generation with the instantaneous rate of change, dP/dt, yields the following equation:

$$\frac{1}{P}\frac{dP}{dt} = \log \text{ rate} = \frac{\lambda\Delta k}{\beta} = \frac{1}{\tau} \qquad \text{so} \qquad \tau = \frac{\beta}{\lambda\Delta k}$$

This constant rate log confirms a simple exponential power rise that depends on the generation time, which is determined almost entirely by the delayed neutrons. For this model, the effect of the delayed neutrons is to make the generation time (and the reactor period) much longer than it would be if there were only prompt neutrons.

The Two Group Model

The most noticeable difference between the prompt jump approximation and the neutron average-lifetime approximation is the prompt jump. Treating all the neutrons alike is an oversimplification that fails to predict the power change immediately following the change in reactivity. The prompt power rise is the immediate response of the 99.5% of the neutrons with one-millisecond lifetimes. The prompt power rise ends quickly and a slower, exponential rise continues, controlled by the delayed neutrons.

To understand this behaviour, we must abandon the unrealistic model of a neutron generation based on an average neutron lifetime and consider the time behaviour of the prompt and delayed neutrons separately. Figure 2.3 illustrates how we will do this and introduces the notation.

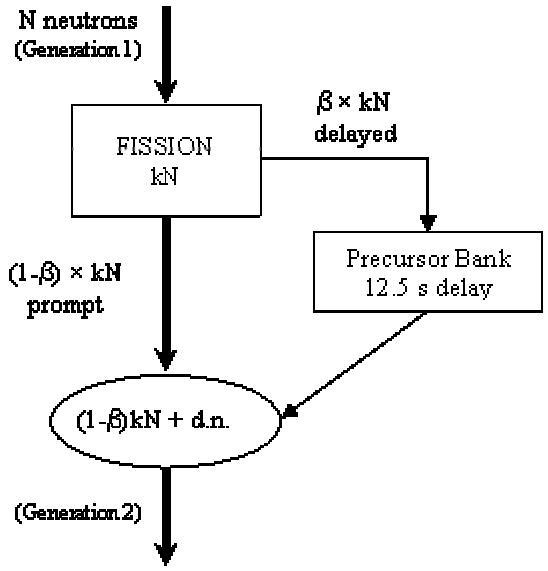

Figure 2.3
A Model for Separate Prompt and Delayed Neutrons

Notes:

The figure shows one prompt neutron cycle (ℓ = 1 ms). The multiplication factor is k, so N neutrons of one "generation" give rise to a total of kN in the next "generation". The majority of the neutrons are prompt, so the prompt fraction, (1 - β) × kN neutrons, appear immediately. The delayed fraction, β × kN, are added to a "bank" of delayed neutron precursors, to be released some time later. The delayed neutrons feed into the fast cycle at a rate controlled by the delayed neutron lifetime and the bank size.

For a steady-state critical reactor, the precursor bank is at an equilibrium concentration where the decay from the bank equals the input. On a step increase in reactivity, the prompt population increases quickly because of an increase in (1 - β) k. The delayed neutron contribution increases slowly, as it depends on the precursor bank.

You may wish to demonstrate this behaviour with a numerical example. Take N = 1,000 initially, and to make the computations easier, use a greatly exaggerated delayed neutron fraction, β = 10% instead of 0.5% and a reactivity insertion Δk = 0.050 (in other words, k goes from 1.000 to 1.050). This simplifies the arithmetic, as there is a significant change from one prompt cycle to the next, but in no way alters the qualitative arguments.

Before the reactivity insertion, N = 1000, and in each cycle 100 of these are stored in the precursor bank, exactly balanced by 100 that enter the cycle by precursor decay. Now, since k is greater than one, the total number of neutrons produced by fission rises to kN = 1050. Ten percent of these are in the form of delayed neutron precursors, so 105 neutrons are "stored" in the precursor bank, to be released later. The prompt neutrons number 945, to which we add 100 neutrons from the precursor bank (stored earlier).

This gives a total of 1045 in the second generation.

If you repeat this calculation for a few more cycles, you should notice two things:

Notes:

Although the input to the precursor bank continues to increase, the output stays constant as the neutrons coming from it reflect a precursor concentration accumulated before the change in reactivity took place. (Each cycle is only 1 ms compared to the "delay time" of 12.5 seconds).

Although the neutron density rises rapidly, the rate of rise falls off (increases of 45, 42, and 40 in successive generations). The reason it falls off is that, at each generation the precursor bank progressively stores more fission neutrons rather than immediately releasing them into the chain reaction.

If you persist and follow, say, fifty cycles, (0.05 seconds after the reactivity insertion) you will see that the increase from one cycle to the next is quite small, and the number of potential neutrons going into the precursor bank each cycle is quite high.

The result is that the rate of increase of the neutron density falls off to zero after a few hundred generations. If we assume that there is no increase in the output of the precursor bank, the situation after 1,000 generations (1 s) would be stable as shown in Figure 2.4.

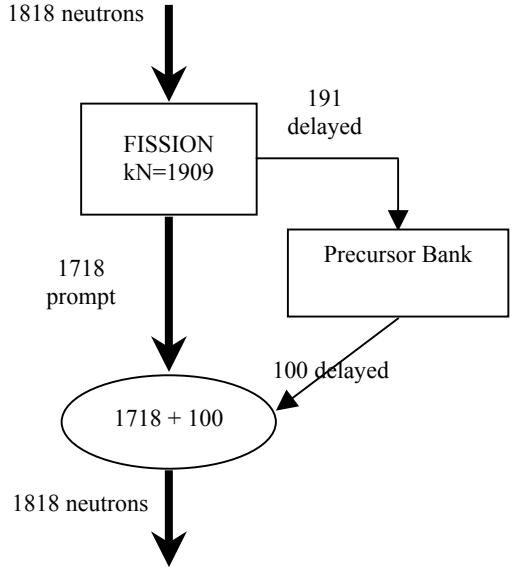

Figure 2.4.
Eventual End of the Prompt Jump (with k= 1.05 and β = 1.10)

By this time, the precursor bank has increased significantly and the number of decays in each cycle will gradually increase from the initial 100 to 101, then 102 etc. The arrival of these delayed neutrons keeps the power increasing.

In summary, a step increase in Δk makes $k > 1$, but $(1-\beta) \times k < 1$. (If no neutrons came out of the precursor bank, the reactor would be subcritical). Prompt multiplication causes the neutron population to rise quickly, but only because the neutron bank continues to "top up" the neutron population at the same rate as before the reactivity addition. Extra precursor decays from the accumulation in the precursor bank drive the stable rise that follows the prompt jump.

Notes:

Example

A few seconds after the insertion of 0.5 mk of positive reactivity (i.e. $\Delta k = 0.0005$), reactor power is observed to rise 100 fold (e.g. from 10^{-4} F.P. to 10^{-2} F.P.). Compare the times of the power rise:

a) For a real reactor with delayed neutrons,

b) If all the neutrons could be prompt.

In each case the power rises exponentially, given by

$$P(t) = P_0 e^{\frac{t}{\tau}}$$

a) The period for a real reactor is
$\tau = (\beta - \Delta k)/\lambda\Delta k = (0.0050 - 0.0005)/(0.08\ s^{-1} \times 0.0005) = 112.5$ s

This gives $t = \tau \times ln(P/P_0)$
$= 112.5\ s \times ln(100) = 518\ s \approx 8½$ minutes.

b) If all the neutrons were prompt, the period would be
$\tau = \ell/\Delta k = 0.001\ s/0.0005 = 2\ s$

This gives $t = \tau \times \ln(P/P_0)$
$= 2\ s \times \ln(100) \approx 10\ s.$

This is almost sixty times faster than actually occurs. Comparing the rates of power increase:

a) Rate log $= 1/\tau = 1/112.5 \approx 0.9\%$ P.P./s

b) Rate log $= 1/\tau = 1/2 = 50\%$ P.P./s

A rate of 0.9% P.P. s^{-1} is a fast power-manoeuvring rate, used only at low power. A 50% P.P. s^{-1} rise is much too fast to control.

2.4 PROMPT CRITICALITY

To this point, we have restricted the reactivity addition Δk to the range $\Delta k << \beta$, or at least to values not very close to β. What happens if we insert a large positive step-change of reactivity, greater than β? Go back to our example, where $\beta = 0.1$, and insert a reactivity step $\Delta k = +0.15$ ($k = 1.150$).

Figure 2.5 shows that starting with 1,000 neutrons in generation 1, the number of prompt fission neutrons in generation 2 is 1035.

Notes:

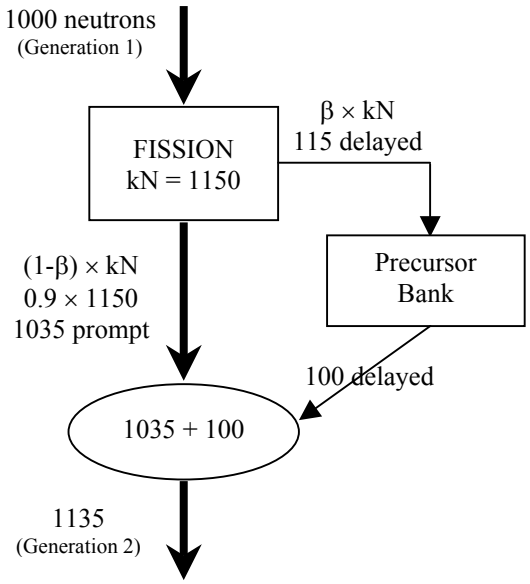

Figure 2.5
Prompt Supercritical

The reactor becomes supercritical on prompt neutrons alone, without having to wait for the delayed neutrons. A reactor with an effective multiplication factor large enough to make the prompt neutron fraction increase enough to make the reactor critical without waiting for the prompt neutrons is prompt critical.

The condition for prompt criticality is $(1-\beta) \times kN = N$

Dividing through by kN and isolating β on one side of the equation gives

$\beta = 1 - 1/k = (k-1)/k$ (= reactivity)

or $\Delta k \approx \beta$

For the equilibrium core, the value of β is approximately 0.005, so prompt critical occurs for $\Delta k \approx 5$ mk. Under these circumstances, the exponential power rise from the fast neutrons alone is similar to that in the example where we assumed all the neutrons were prompt. None of our formulas are valid for $\Delta k \approx \beta$, but accurate numerical analysis shows the reactor period is a little less than 1 second for a prompt critical CANDU reactor, making the rate of power increase (rate log = $1/\tau$) greater than 100% P.P/s.

There is no abrupt change in reactor behaviour as we move through the prompt critical condition. As reactivity $\Delta k \rightarrow \beta$, the reactor becomes less and less dependent on the delayed neutrons so that reactor period (τ) decreases smoothly to shorter and shorter values.

The rate of reactor power increase is unacceptable for reactivity insertions well below $\Delta k = \beta$.

Notes:

Automatic reactor shutdown occurs when instruments detect a high rate of power increase: reactor stepback drops control absorber rods into the core, a trip on shutdown system one (SDS1) drops shutoff rods into the core, or a shutdown system two (SDS2) trip injects neutron-absorbing poisons into the moderator. These systems trigger at rates corresponding to the following reactor periods:

τ = 12.5s (rate log 8% P.P. s^{-1}) for stepback,

τ = 10 s (rate log = 10% P.P. s^{-1}) for SDS1,

τ = 6.7 s (rate log = 15% P.P. s^{-1}) for SDS2.

2.5 POWER RUNDOWN: THE PROMPT DROP

We can use the equations describing the prompt jump to estimate power decrease following a large negative reactivity insertion. In this case, we have a prompt drop as the prompt neutron population collapses, followed by a stable negative period. With a negative reactivity insertion, the prompt population multiplication factor $(1 - \beta)\, k$ is much less than 1 so the prompt population drops. The precursors bank gradually decays as there are not enough fissions to replenish it.

The prompt "jump" approximation gives the ratio of power P immediately after the step insertion of reactivity to the original level P_0, but in this case, Δk is negative. For example, a full stepback may introduce approximately 15 mk of negative reactivity. If Δk = -15 mk is injected into a critical reactor, the power immediately after the step change is,

$$P = P_0 \frac{\beta}{\beta - \Delta k} = P_0 \frac{0.005}{0.005 - (-0.015)} = 0.25 P_0$$

Following the prompt drop, the power settles down to decaying with a negative period. The equation for the reactor period also works for negative reactivity. For Δk = -15 mk

$$\tau = \frac{\beta - \Delta k}{\lambda \Delta k} \quad \text{gives}$$

$$\tau = \frac{0.005 - (-0.015)}{0.08 \times (-0.015)} = -17\text{s}$$

This represents a decreasing exponential that looks like decay. The power continues to drop with a decay constant of approximately $1/17 = 0.06\ s^{-1}$.

Notes:

If $|\Delta k|$ is >> β, the period becomes approximately

$$\tau = \frac{\beta - \Delta k}{\lambda \Delta k} \rightarrow -\frac{1}{\lambda}$$

Thus, the decay constant of the precursors determines the stable period for a large negative reactivity insertion. In this case, power decays following the prompt drop with a decay constant of approximately $1/12.5 = 0.08\ s^{-1}$.

For large negative Δk, the approximation becomes increasingly inaccurate because it models all the delayed neutrons as though they have the same decay constant. Consider the fact that there are actually six precursor groups, the longest-lived having a half-life of 54 s. The power, after the shorter-lived groups have decayed, will decrease with a period (τ) of about 80 s. ($\tau = 1/\lambda$ and $\lambda = 0.693/t_{1/2}$).

As Section 7 describes, the rate of power decrease eventually slows even further because of the presence of photoneutrons, which have longer half-lives than the delayed neutrons.

A final cautionary note: remember that this section deals only with the neutron power. Following a reactor shut down from high power, the thermal power is not at all proportional to the neutron flux, due to the heat contributed by the accumulated fission products.

RESPONSIVENESS OF THE SUBCRITICAL REACTOR

3.0 INTRODUCTION

Notes:

CANDU reactors have two neutron sources (not including neutrons from induced fission). These are photoneutrons and neutrons from spontaneous fission. In this section, we show that the fission process amplifies the source neutrons by a factor (the subcritical multiplication factor) that is inversely proportional to the reactivity.

The nuclear instruments can measure observable flux in a subcritical core, even when the source is quite small, because there is subcritical multiplication of source neutrons. The observed flux increases as k increases in the subcritical core because the source amplification is inversely proportional to reactivity. When deeply subcritical, very large changes in reactivity produce little noticeable change in flux. In contrast, an almost critical reactor behaves, for all practical purposes, nearly like a critical reactor. The essential difference is that for $k > 1$, power continues to rise; for a reactivity increase that leaves $k < 1$ the power stops rising at a level that depends on reactivity.

The section also examines the stabilization time, the time for the power level to establish a new value after a step change in reactivity. The stabilization time gets longer the closer the reactor is to being critical.

Operationally, we can determine the reactivity of the shutdown reactor by inserting or withdrawing a known reactivity worth, and measuring the corresponding change in indicated power. In particular, a power doubling occurs in a subcritical reactor whenever we reduce the shutdown reactivity by half. Procedures use this behaviour to approach criticality cautiously during startup, and to test whether or not a low power reactor is critical.

3.1 NEUTRON FLUX IN A "SHUT DOWN" REACTOR

The observed neutron power (flux) in a subcritical reactor with sources present does not drop quickly to zero. The observed neutron power is considerably higher than can be accounted for by the size of the neutron source itself. To explain the phenomenon, consider a reactor held deeply subcritical by neutron absorbing poisons dissolved in the moderator. Now suppose we inject a single pulse of neutrons into the core. Most of the neutrons are absorbed without causing fission but, inevitably, some cause fission and generate additional neutrons. (The effective multiplication factor of the "shut down" reactor is much less than one, but it is not zero.)

Introducing a pulse of neutrons does not initiate a self-sustaining chain reaction; (less than one neutron per fission survives). However, if a neutron source injects a steady trickle of neutrons, the instruments register more flux than just the source

flux. Moreover, as k gets closer to k = 1, more neutrons survive from each fission that occurs and the flux increases. The flux would drop quickly to zero without source neutrons in the reactor, but when there is a source, there are always more neutrons than just those from the source.

Notes:

The next two subsections address, respectively, the natural sources of neutrons in a CANDU reactor, and the amplification effect due to the reactor.

Neutron Sources

The term "source neutrons" refers to neutrons other than prompt or delayed fission neutrons. A neutron source injects a steady supply of neutrons, independent of power level, temperature, or any controllable parameter. A CANDU has two "built-in" neutron sources, one arising from the nature of the fuel, and the other from the presence of heavy water.

1. ***Spontaneous fission neutrons***

Neutrons from spontaneous fission come mainly from U-238, which constitutes most of the fuel. The contribution to reactor power from fissions caused by the spontaneous fission neutrons is constant, at about 10^{-12} % of full power. This low-level source is important only at first startup (or after a prolonged shutdown) when there are essentially no other neutrons.

NOTE: ***Photoneutrons***

High-energy gamma rays (> 2.2 MeV) emitted during fission product decay generate photoneutrons by interactions with deuterium nuclei. The photoneutron source strength depends on the presence of energetic gamma rays, which in turn depends on how long the reactor has been operating. Following prolonged operation at significant power levels (>10% full power), the decay gamma flux is proportional to the power. The photoneutron source is typically about 0.03% of the power level—less than 10% of the delayed neutron fraction. This source becomes important following the disappearance of the delayed neutrons after shutdown.

Neutrons	Contribution
Prompt Neutrons	99,470,000,000,000
Delayed Neutrons	500,000,000,000
Photoneutrons	30,000,000,000
Spontaneous Fission	1
Total (10^{14} n cm^{-2} s^{-1})	100,000,000,000,000

Table 3.1
Relative Neutron Strengths

Notes:

Table 3.1 shows the relative importance of source neutrons by showing the makeup of the full power flux.

Source Neutron Time Dependence

Spontaneous fission produces a small, constant supply of neutrons (about 130 neutrons per second per CANDU bundle).

Following shutdown, the photoneutron source decreases with the decay of fission products that produce high-energy gamma rays. Power from fissions caused by photoneutrons is typically close to 5×10^{-5} % one day after shutdown. (The precise value depends on the operating power in the weeks before shutdown.)

After the initial rapid drop, the photoneutron source drops slowly, decreasing a little each day. The longest-lived fission product with gamma ray energy above the threshold is Ba-140, which has a half-life of 12.75 days; thus, the photoneutron source persists for many weeks after shutdown.

This course follows the usual convention of treating neutrons from delayed neutron precursors as delayed fission neutrons, and neutrons from photodisintegration of deuterium as source neutrons. The distinction rests on the practical operational effects that result from the different half-lives of these kinds of neutrons.

The longest-lived delayed neutron precursor has a 54-second half-life. When the fission rate changes, it takes a few minutes for the delayed neutron precursor bank to come back into equilibrium. Its slow change controls the rate of change (in both the subcritical and the critical core) of the current fission rate. The current fission rate determines the size of the equilibrium precursor bank.

Photoneutrons, in contrast, have much longer precursor lifetimes. Weeks after a shutdown, fission products from power operation in the weeks before shutdown continue to generate photoneutrons. Following a power change, it takes weeks for the photoneutron fraction to reach its new equilibrium strength.

For neutron balance in a critical steady-state core (Section 1), it is convenient to lump the photoneutrons with the delayed neutrons—they act like delayed neutrons with long lifetimes. Dynamic reactor behaviour (Section 2) depends on a delayed neutron "delay tank"—a neutron source that changes slowly (over many prompt neutron cycles) following a reactivity insertion. The distinction between source neutrons and delayed fission neutrons depends on whether or not the neutron supply is nearly constant over the time interval considered. Operationally, things like power ramps take minutes or hours. From an operational perspective, that makes neutrons from delayed neutron precursors delayed fission neutrons, while photoneutrons are source neutrons.

Notes:

Source Multiplication in a Subcritical Reactor

The previous sub-section stated that the contribution to reactor power from fissions caused by the spontaneous fission neutrons is about 10^{-12} % of full power, and the contribution from photoneutrons is 0.3% of full power. These values are for fissions caused directly by the source neutrons themselves; that is, we ignored the fact that these fissions give rise to additional fission neutrons that, in turn, may cause further fission. The subcritical reactor acts as a multiplier of source neutrons, so that the actual power generated because a source is present is much greater than would be produced by the source neutrons alone. Source amplification in the subcritical reactor depends on its reactivity, as shown below.

Assume that a neutron source is inserting S_0 neutrons in each generation into a subcritical reactor whose effective multiplication factor, k, is less than one. Imagine that you could label the neutrons so that you could identify the origin of each neutron in the reactor. In the neutron population present in the core now, there are S_0 neutrons emitted by the source during the current generation.

During the immediately preceding generation, the source emitted S_0 neutrons and some of these source neutrons caused fission, producing a contribution of kS_0 neutrons to the present population.

The total number of neutrons in the neutron population from this and the immediately preceding generation is therefore:

$$S_0 + kS_0 (kS_0 \text{ is smaller than } S_0).$$

Similarly, one generation earlier, the source emitted S_0 neutrons, generating kS_0 neutrons in the immediately preceding generation, resulting in k^2S_0 neutrons in the present generation. Looking back, the contribution of three generations of source neutrons and the fissions they caused contributes

$$S_0 + kS_0 + k^2S_0$$

If we extend this argument indefinitely (working backwards a large number of generations), we see that the total neutron population (S_∞) is made up as follows

Notes:

$$S_\infty = S_0 + kS_0 + k^2S_0 + k^3S_0 + \ldots$$

$$= (1 + k + k^2 + k^3 + \ldots)\, S_0$$

The series in brackets has an infinite number of terms, but a high-school algebra trick yields a finite sum. For $k < 1$

$$S_\infty = \left[\frac{1}{1-k}\right] \times S_0$$

The fission rate in the subcritical reactor is proportional to the neutron population, so we can write these equations in terms of neutron power rather than neutron population. Thus, the observed power level in the reactor is

$$P_{obs} = \left[\frac{1}{1-k}\right] \times P_{source}$$

where P_{source} is the power that would be generated by the source neutrons themselves in the absence of any multiplication by the fission process (that is, if ν were equal to zero). P_{obs} is the measured (actual) power level.

Variation of Source Multiplication

The factor $[1/(1 – k)]$ is called the subcritical multiplication factor. Notice its effect on a source. Even in a reactor that is well below critical, say $\Delta k = -100$ mk, ($k = 0.9$) the equilibrium observed power level ($P_{observed}$) is 10 times greater than the actual photoneutron source power:

$$P_{obs} = \left[\frac{1}{1-0.9}\right] \cdot P_{source} = 10P_{source}$$

In this example, the subcritical multiplication factor is 10. This means that fission produces the great majority of neutrons in the system and not the source directly.

The amount of subcritical multiplication depends only on the value of k. For example, if the control absorbers fell into the core and the zones filled, the reactivity insertion might be about -20 mk, giving a subcritical multiplication factor of 50.

$$P_{obs} = \left[\frac{1}{1-0.980}\right] \cdot P_{source} = 50P_{source}$$

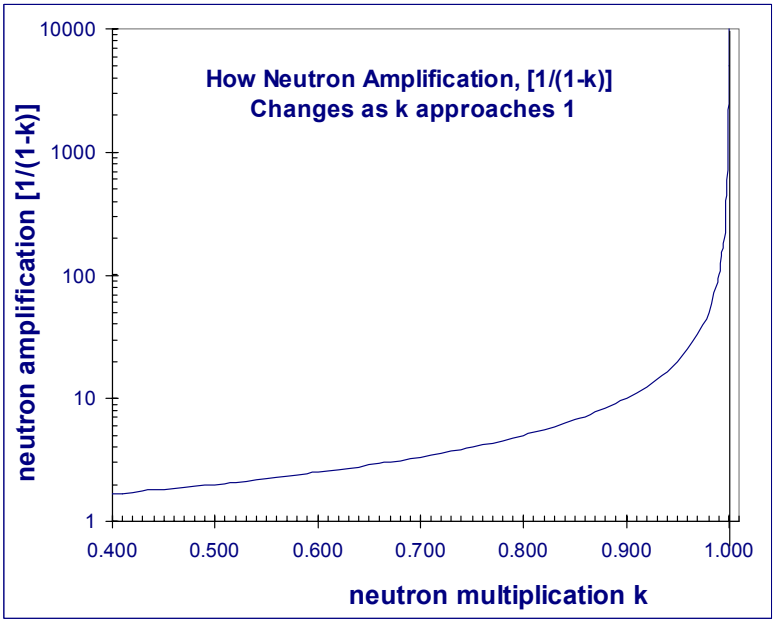

Figure 3.1
The Subcritical Multiplication Factor

Figure 3.1 shows the variation of the subcritical multiplication factor as k changes. When a source is present, most of the neutrons in the reactor (at least for k > 0.5) are not neutrons from the source, but neutrons that originate in fissions generated because there are source neutrons. A glance at the figure shows that adding 100 mk to a reactor with k = 0.400 has less effect than adding, say, 20 mk to a reactor with k = 0.950.

3.2 DYNAMICS IN THE SUBCRITICAL CORE

A positive reactivity insertion into a subcritical reactor increases power to a new equilibrium level. The size of the increase and the time it takes for the power to stabilize at the new value depends on the reactivity inserted and on the final value of k. The closer k is to one, the larger the increase in power for a given reactivity insertion, and the longer the time required to stabilize.

Figure 3.2 illustrates the increasing size and stabilization time for step reactivity insertions that bring the core closer and closer to critical. The power level starts at 1×10^{-5}% with k = 0.900. (The core is 100 mk subcritical). The figure shows nine step reactivity insertions of +10 mk each, with P given time to stabilize after each reactivity addition. (The numbers chosen are for illustration only; making 10 mk steps would be quite impossible in practice.)

Notes:

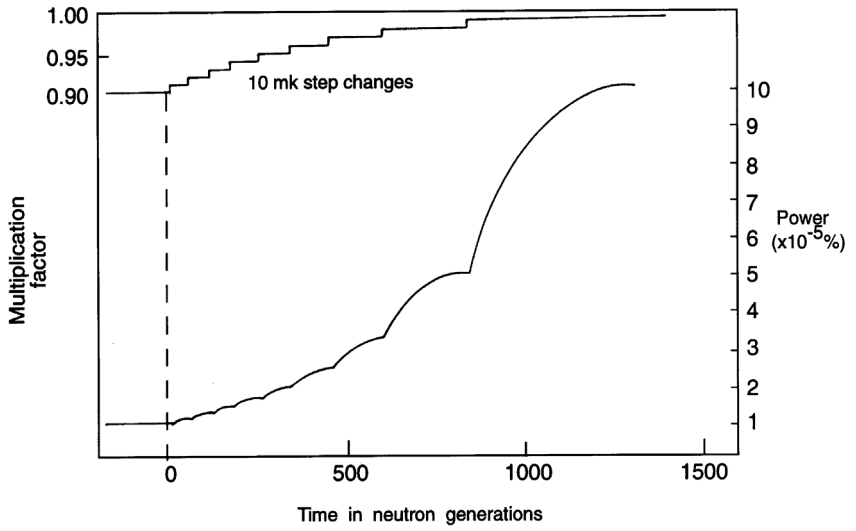

Figure 3.2
Stabilization Time After Change in Multiplication Factor

Notes:

Example:

The following example calculates the power changes in Figure 3.2 for two specific +10 mk reactivity step increases as k approaches 1:

k changes from 0.900 to 0.910:

Power at k = 0.90 is $P_{0.90} = \frac{P_{source}}{0.10} = 10P_{source}$

Power at k = 0.91 is $P_{0.91} = \frac{P_{source}}{0.09} = 11.1P_{source}$

So, $P_{0.91}/P_{0.90} = 1.11$; power increases by 11%

k changes from 0.980 to 0.990:

Power at k = 0.98 is $P_{0.98} = \frac{P_{source}}{0.02} = 50P_{source}$

Power at k = 0.99 is $P_{0.99} = \frac{P_{source}}{0.01} = 100P_{source}$

So $P_{0.99}/P_{0.98} = 2$; power doubles

It is clear from this example that the change in power for a given reactivity addition increases as k tends towards one.

The fact that the stabilization time gets longer as k tends to one is more subtle.

Notes:

During normal subcritical reactor operation, moderator poisons keep the core subcritical. The rate of reactivity addition by poison removal is slow, so we have no need of "two group" kinematics equations. Instead, we will discuss time dependence in the subcritical core using the average neutron generation-time model. While the calculation is only approximate, it shows why the stabilization time gets longer as the subcritical core gets closer to critical.

As shown earlier, the sum of an infinite series gives the total, stable neutron population arising from a source. The series is

$$1 + k + k^2 + \cdots + k^{n-1} + k^n + k^{n+1} + k^{n+2} + \cdots = \frac{1}{1-k}$$

The sum exists because k is less than one, and each successive term is smaller than the previous one. For practical purposes, the series ends after n generations if the value of k^n has become very small compared to the first term (that is, 1).

More precisely, if we pick a value of n that makes $k^n = 0.01$, then the sum of the entire series comes 99% from the terms before the n^{th} term, and 1% from k^n onward. [Writing the series as $1 + k + k^2 + \cdots + k^{n-1} + k^n(1 + k + k^2 + \cdots)$, we see that the series from the n^{th} term onwards is just 0.01 of the entire sum].

We can calculate the number of generations, n, for the power to stabilize following a change in k (that is reach within 1% of its equilibrium value) from the equation:

$k^n = 0.01$

where k is the new value of k after a step insertion of reactivity.

$n \log k = \log (0.01) = -2$, or

$$n = \frac{\log(0.01)}{\log(k)} = \frac{-2}{\log(k)}$$

(Note: for convenience, use log to base 10)

Table 3.2 shows the number of generations to stabilize for four selected values of k. To determine approximate stabilization times, we take the time for one generation to be the average neutron generation time (≈0.07s).

Notes:

k	# of generations n = -2/log(k)	stabilization time
0.800	21	1 s
0.900	44	3 s
0.990	458	30 s
0.999	4603	5 min

Table 3.2
Time to Stabilize

Clearly, as k→1, more neutron generations pass before the power stabilizes, and the stabilization times get longer.

Formulas exist that will calculate the response of a subcritical reactor to insertions of reactivity more accurately. Compare Figures 3.3 and 3.4. Each figure shows an increase in subcritical power of about 25%, that is, about 0.1 decade ($10^{0.1} \approx 1.25$).

The 60 mk step increase in reactivity in a deeply subcritical core produces a power change in much less time than any reactivity device could, in practice, insert the reactivity. In Figure 3.4, the power increase takes about five minutes for a small reactivity insertion (+ 0.3 mk, making k = 0.9988) in an almost critical reactor. In passing, notice that the more accurate analysis plotted in Figure 3.4 shows a prompt jump, followed by a gradual rise. Notice also that the stabilization time is comparable to the estimate in Table 3.2.

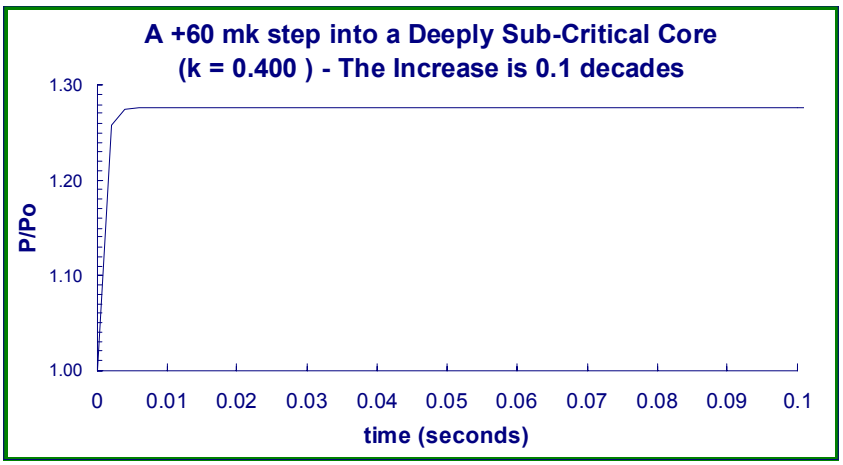

Figure 3.3
A Large Reactivity Insertion in a Deeply Subcritical Core

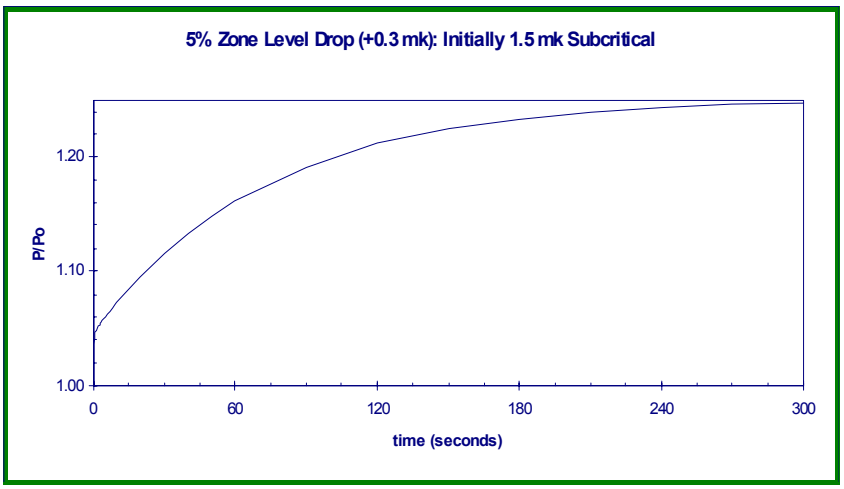

Figure 3.4
A Small Reactivity Insertion in a Nearly Critical Core

As the reactor approaches criticality, each successive reactivity insertion produces a larger power increase, and the time to stabilize gets longer. This produces a smooth transition between the subcritical and critical states. When deeply subcritical the core is very sluggish; even large reactivity increases could go unnoticed. An almost critical reactor behaves a lot like a critical reactor. Figures 3.5 (opposite) illustrate this.

The initial one minute of power rise is very similar in the two graphs of Figure 3.5. However, power continues to rise in the critical core. The power reaches an equilibrium power and stops rising in the subcritical core, where there is a single power level for each $-\Delta k$, given by $P_{obs} = [1/(1-k)]\ P_{source}$

Notes:

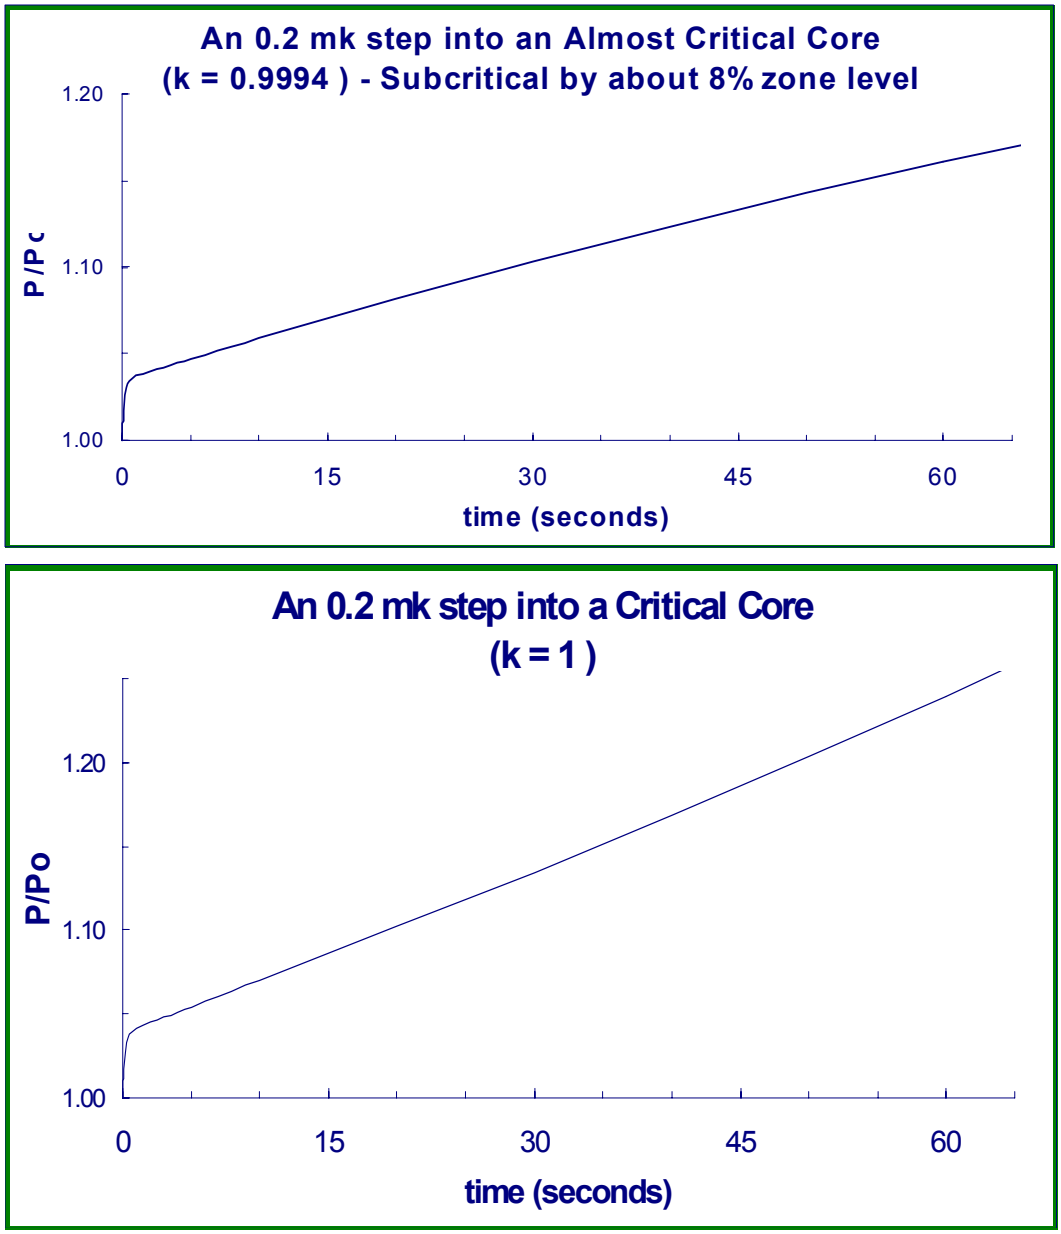

Figure 3.5
Comparison of a Critical Core with One that is Almost Critical

Notes:

3.3 EXAMPLES

Determining Core Reactivity

We can use the subcritical multiplication equation to calculate the reactivity of the core by observing the power before and after a known reactivity change. Usually we do not know the source strength, but we do know it is constant on the short term. We can write the equation for each observed power level and eliminate the unknown source strength.

Conversely, once you determine k, you can find the reactivity worth of a device by observing the power level before and after inserting the device.

Notes:

Suppose a reactor is shut down with a constant indicated power of 2×10^{-5}%. The operator inserts +1 mk by withdrawing an adjuster and power stabilizes at 3×10^{-5}%. Find the original value of k.

$$P_{obs} = \left[\frac{1}{1-k_i}\right] P_{source}$$

(k_i = initial multiplication factor)

Before the reactivity addition:

$$2 \times 10^{-5}\% = \left[\frac{1}{1-k_i}\right] P_{source}$$

After the reactivity addition

$$3 \times 10^{-5}\% = \left[\frac{1}{1-(k_i + 0.001)}\right] P_{source}$$

The source strength is the same in both cases:

$P_{source} = (1 - k_i) \times 2 \times 10^{-5}$%, and

$P_{source} = [1 - (k_i + 0.001)] \times 3 \times 10^{-5}$%

Hence, equating to eliminate P_{source}

$$\begin{aligned} 2(1 - k_i) &= 3\,(0.999 - k_i) \\ 2 - 2k_i &= 2.997 - 3k_i \\ k_i &= 0.997 \end{aligned}$$

(Now, if you want to, you can substitute back and find out that the subcritical multiplication factor initially was 333, and the source strength is 6×10^{-8} %.)

Power Doubling

One example with important practical applications in reactor operation is the power-doubling rule, stated as follows:

When a certain reactivity addition causes power to double in a subcritical reactor, a further addition of the same reactivity would make the reactor critical.

To demonstrate why this is so, suppose we have a reactor that is subcritical with an effective multiplication factor of k_i (i = initial).

The power level (P_i) in the reactor is obtained by multiplying the source power by the subcritical multiplication factor, and is

$$P_i = \left[\frac{1}{1-k_i}\right] \cdot P_{source} = \frac{P_{source}}{-\Delta k_i}$$

Notes:

where $(k_i - 1) = \Delta k_i$ is the amount by which the multiplication factor falls short of criticality.

Now suppose we insert positive reactivity of ½ Δk_i so that the final reactivity is $\Delta k_f = ½ \Delta k_i$. Figure 3.6 shows the reactivities. The equation for the final power level, P_f, is

$$P_f = \frac{P_{source}}{-\Delta k_f} = \frac{P_{source}}{-\frac{1}{2}\Delta k_i} = 2 \cdot \frac{P_{source}}{-\Delta k_i} = 2P_i$$

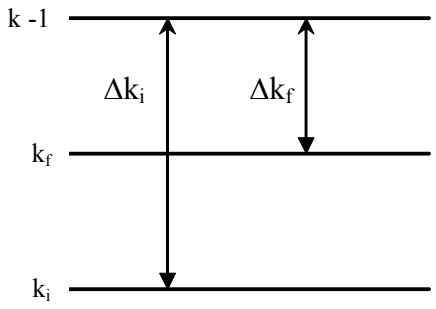

Figure 3.6
Change in Reactivity for a Power Doubling

Removing half the negative reactivity causes the power to double. Said another way, when we observe a power doubling in a subcritical reactor, the amount by which the reactor was subcritical has been cut in half.

Approaching Criticality by Power Doubling

Table 3.3 demonstrates how, in principle, a series of power doublings makes a deeply subcritical reactor approach criticality. Consider the first entry in the table. The reactor is subcritical and initially k = 0.680 (Δk = -320 mk). Measured power is 1.25×10^{-6} of full power ($10^{-5.9}$ decades).

Notes:

	-Δk (mk)	power % F.P.	power (decades)
0	320	1.25×10^{-4}	$10^{-3.9}$
1st	160	2.5×10^{-4}	$10^{-3.6}$
2nd	80	5×10^{-4}	$10^{-3.3}$
3rd	40	1×10^{-3}	$10^{-3.0}$
4th	20	2×10^{-3}	$10^{-2.7}$
5th	10	4×10^{-3}	$10^{-2.4}$
6th	5	8×10^{-3}	$10^{-2.1}$
7th	2.5	1.6×10^{-2}	$10^{-1.8}$
8th	1.25	3.2×10^{-2}	$10^{-1.5}$
9th	0.6125	6.4×10^{-2}	$10^{-1.2}$
10th	0.31	0.128	$10^{-0.9}$

Table 3.3
A Series of Power Doublings

Ten successive doublings raise the power by three orders of magnitude ($2^{10} = 1024 \approx 10^3$). Each doubling leaves half the negative reactivity in the core, so the reactor is not yet critical. It is, however, within a few tenths of a mk from critical (the equivalent of about 5% zone level). The regulating system could respond to a request for a power increase and manoeuvre power at the demanded rate.

Each time power doubles in the subcritical core, noting the change in reactivity that caused the doubling is, in effect, a measurement of k.

Figure 3.7 illustrates the principle of the power doubling. The graph assumes startup instruments monitor count-rate (proportional to neutron power). The graph is a straight line, because the subcritical multiplication formula

$$P_{obs} \propto \frac{1}{\Delta k} \quad \text{implies} \quad \left[\frac{1}{CR}\right] \propto \Delta k$$

The point where the extrapolated line crosses the axis represents criticality. (when $\Delta k = 0$, $1/CR = 0$). The graph demonstrates visually that taking the reactor to criticality by power doubling (i.e., by cutting [1/CR] in half), is a cautious way of approaching criticality.

As long as one doubles power, the reactor gets closer and closer to critical without actually going critical. By plotting the graph, or by keeping track of the reactivity changes that double power, the operator can predict the reactivity worth of the controlling device at criticality.

Notes:

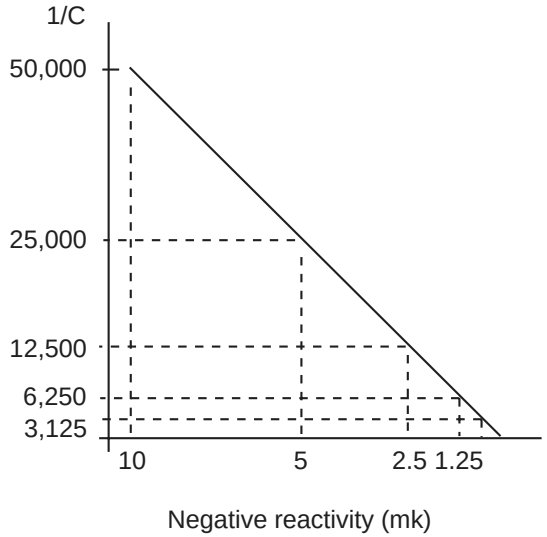

Figure 3.7
Approach to Critical by Power Doubling

Finally, Section 7 examines a commonly used criticality test that applies the principle of power doubling. When the reactor is at low power it is difficult to tell if the regulating system is holding power constant with $k \approx 1$, or if the reactor is subcritical with an observed power level corresponding to a particular $-\Delta k$.

The operator applies the test, typically once per shift at low power, asking the regulating system to double power. If the liquid zones change excessively in attempting to double power, the reactor is definitely subcritical. If a small movement of zone level produces a power doubling, the operator knows the regulating system could manoeuvre power to any demanded level at the demanded rate, and is therefore, in effect, “critical”.

EFFECTS OF TEMPERATURE AND VOIDING ON CORE REACTIVITY

Notes:

4.0 INTRODUCTION

A change in power, for a reactor operating at high power, generally alters the temperatures of the fuel, moderator, and coolant. A change in temperature of any of these components causes a change in reactivity that, in turn, affects reactor operation (a feedback effect). The temperature coefficients of reactivity (the change in reactivity per unit change in temperature) determine the size and direction of the reactivity change.

This section reviews the principal physical mechanisms that give rise to the temperature coefficients. These are:

- Doppler broadening of the U-238 resonances;
- Changes in the thermal neutron energy spectrum;
- Density changes of the coolant and moderator.

Changes in each factor in the six-factor formula for k may contribute to the temperature coefficient of a given reactor component (for instance, fuel). We describe how the physical mechanisms above determine the sign of each relevant contribution to both fuel and moderator temperature coefficients.

This section also looks at some operational implications of temperature coefficients, including the changes of reactivity that occur on startup and shutdown, and the importance of the Doppler Effect in providing negative feedback to simplify normal regulation and to limit positive power transients.

Finally, the effects of void formation are an important source of feedback in power transients. Voiding of the heavy water coolant from the fuel channels produces positive reactivity. The practical implications of this include the need for a high coolant isotopic.

4.1 FEEDBACK—AN INTRODUCTION TO TEMPERATURE EFFECTS

At low power, when the fission rate is too low to generate much heat, addition of positive reactivity produces an exponential rise, (that is, a constant log rate increase). At power levels where there is significant heating, temperature changes in the fuel, moderator, and coolant change reactivity—an effect that occurs quite quickly. The following description demonstrates this reactivity feedback effect.

In 1949, a controlled experiment at AECL, Chalk River, allowed the NRX reactor to “run away”. NRX was a heavy water moderated reactor with control rods for reactor regulation. The heavy water level was set somewhat above the height at which the reactor would be critical at low power with the rods withdrawn. The

experiment allowed reactor power in the supercritical reactor to increase "unchecked".

Notes:

The manner in which the power changed was not what you might expect from Section 2. At low power, the power did increase exponentially with a rate log of 3% P.P s^{-1} ($\tau \approx 33$ s, $\Delta k \approx +1.6$mk). However, as the temperature of the fuel increased, the reactivity decreased and caused the rate of power increase to slow (see Figure 4.1). Later, the reactivity decreased even more, as the heavy water became warmer. The total decrease in reactivity was enough, in this case, to make the reactor subcritical, and so power reached a maximum and began dropping.

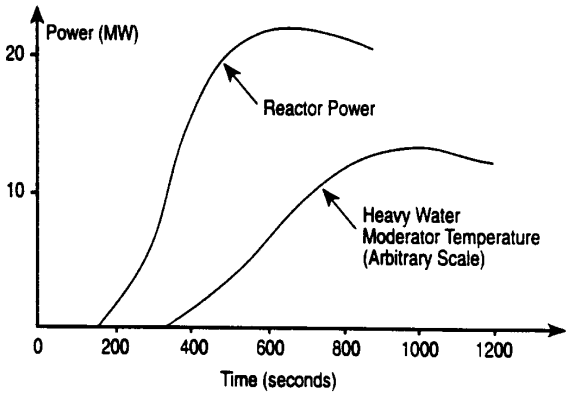

Figure 4.1
Self-Regulation Characteristics of the NRX Reactor

The reactor behaviour was self-regulating since temperature changes induced by a power increase reduced reactivity and so prevented power from increasing indefinitely. Of course, the initial excess reactivity in this experiment was quite small; if it had been larger, the power likely would have continued rising.

The NRX experiment illustrates the effect of a negative temperature coefficient. By definition, the temperature coefficient of reactivity is the change in reactivity per unit increase in temperature. Its units are mk/°C or μk/°C (1 μk = 10^{-3} mk). A temperature coefficient may be positive or negative, but in the example, a temperature increase caused a loss of reactivity so the coefficient was negative. Temperature changes occur in the fuel, coolant, and moderator more or less independently so there is a temperature coefficient of reactivity associated with each.

It is very desirable for the overall temperature coefficient of a reactor to be negative to provide the self-regulating feature illustrated by NRX. It is particularly helpful if the fuel temperature coefficient is negative because in a transient the fuel heats up more rapidly than the other core components do.

The reactivity effect of a temperature change is the product of the temperature coefficient and the temperature change. It is helpful to know typical temperatures of the fuel moderator and coolant in different operating states.

Table 4.1 shows the effective average temperatures of fuel, coolant, and moderator assumed in these notes. Temperatures in particular stations may differ somewhat from these values.

Notes:

Component	Cold Shutdown	Hot Shutdown	Full Power
Fuel	25	290	790
Coolant	25	265	290
Moderator	25	66	69

Table 4.1
Typical Temperatures of Reactor Components (°C)

As stated, the temperature coefficient is the change in reactivity per unit (°C) change in temperature. The overall temperature coefficient is the slope of a graph of reactivity (Δk/k) vs. temperature at the normal operating point. We will introduce such graphs later in the section for moderator, coolant, and fuel.

Introducing $\frac{\alpha}{} = \frac{d(\ln k)}{dT}$ for the slope, where $k = \eta f p \varepsilon \Lambda_f \Lambda_t$, allows us to write the temperature coefficient as a sum:

$$\frac{1}{\varepsilon}\frac{d\varepsilon}{dT} + \frac{1}{p}\frac{dp}{dT} + \frac{1}{f}\frac{df}{dT} + \frac{1}{\eta}\frac{d\eta}{dT} + \frac{1}{\Lambda_f}\frac{d\Lambda_f}{dT} + \frac{1}{\Lambda_t}\frac{d\Lambda_t}{dT}$$

To see how the physical effects of a temperature change influence the temperature coefficient we need to recall the neutron life histories in the reactor and determine the effect of temperature change on factors in the six-factor formula.

4.2 THE PHYSICAL BASIS FOR TEMPERATURE COEFFICIENTS

Later in this section, we look at the individual temperature coefficients of the fuel, moderator, and coolant, and describe the effects of each temperature change on each relevant factor in the six-factor formula. Before we do this, it is useful to consider generally how temperature affects these factors. Three physical processes produce the observed effects.

A temperature increase in a substance increases atomic and molecular motion. It also (usually) causes thermal expansion. Each of these effects influences the factors in the equation for reaction rate (R).

$$R = \phi\Sigma = N\sigma\phi$$

where σ refers to the fission, absorption or scattering cross-section, as appropriate.

The three most important effects, in order of increasing importance, are:

- Thermal expansion decreases density and this directly affects N, the number density of the nuclides present in the reactor. A subsection below explains the effect in detail.
- Higher temperatures in the neutron's environment change the thermal neutron spectrum. This affects the thermal neutron cross-sections for various scattering, absorption, and fission processes. A subsection below explains these effects in detail.
- Increased molecular motion of U-238 increases the resonance capture of neutrons. The motion increases the apparent width of the resonances in the microscopic cross section, an effect known as Doppler Broadening. A subsection below explains Doppler Broadening and its effects.

Notes:

Density Changes

Thermal expansion of the moderator as moderator temperature rises, decreases moderator density. Since the number of atoms per unit volume (N) decreases, neutrons travel further between interactions, so they diffuse over a wider range, both while they thermalize, and as thermal neutrons.

A fuel density change does not have a large effect. Fuel bundles accommodate the expansion of the fuel pellets so a change in fuel density may cause the reaction rate per cubic centimetre of fuel to decrease, but the rate over the whole core does not. By contrast, thermal expansion of heavy water displaces it from the core so the amount of heavy water between the fuel channels decreases.

The effect of longer path lengths in the moderator and coolant, leading to increased overall range of the neutrons, is an increased chance that neutrons will leak out of the reactor. The fast and thermal non-leakage probabilities decrease, tending to lower the reactivity, but the effect is not big.

A more important effect of longer path lengths in the moderator is that neutrons are more likely to reach a neighbouring channel before the moderator has thermalized them. Resonance capture increases so resonance escape probability decreases.

On the other hand, the reduction in atomic density lowers the macroscopic absorption cross-sections in the moderator or coolant, which increases the thermal utilization (f), thus tending to increase reactivity. The effect on f of the reduction in atomic density is particularly strong in the moderator if it contains an appreciable quantity of poison.

The Neutron Energy Spectrum

We usually assume all thermal neutrons have energy of 0.025 eV and travel at 2200 m/s. In fact, thermal neutrons, by elastic collisions, distribute themselves over a range of energies. On average, the neutrons share the kinetic energy of the atoms and molecules with which they interact, but any particular elastic collision can increase or decrease the neutron kinetic energy. A rather complicated mathematical expression known as the thermal neutron spectrum shows the variation in neutron density as a function of thermal neutron energy. We associate a neutron temperature with each such distribution; the temperature depends on the temperature of the surroundings.

Notes:

A temperature change of any component of the reactor affects the neutron spectrum. This, in turn, changes fission and absorption rates in the fuel and absorption rates throughout the reactor, because cross sections vary with neutron energy. The following subsections provide a step-by-step guide to the effects of spectrum changes on reactivity.

The Thermal Neutron Spectrum

Figure 4.2 shows thermal neutron spectra for 20°C and 300°C. Thermal neutrons in a CANDU core will always have spectra somewhere between the distributions shown. (At full power, the thermal neutron temperature of neutrons interacting with CANDU fuel is about 200 °C. When shut down, the neutron temperature will be hotter than 20°C.) Figure 4.2 is a Maxwellian distribution of neutron speeds, plotted as a function of neutron energy, with a logarithmic scale that allows easy comparison with cross-section graphs.

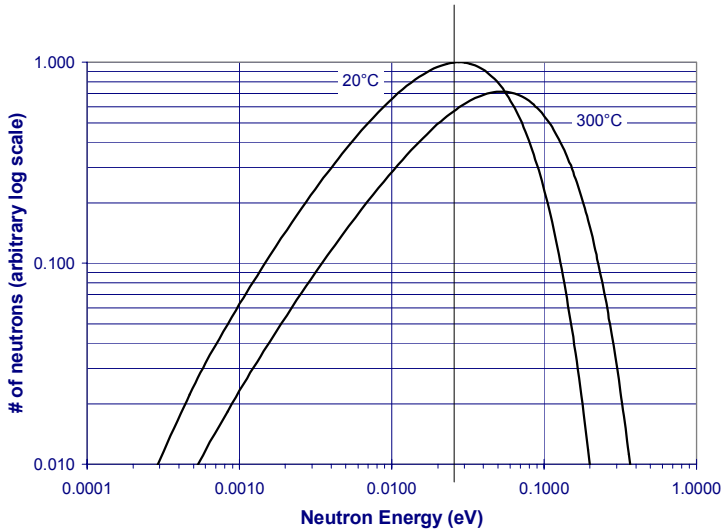

Figure 4.2
Neutron Energy Shift with Changing Temperature

At room temperature, the most probable value for thermal neutron energy (the peak of this curve) is 0.0253 eV, corresponding to a speed of 2200 m/s. The average neutron energy is proportional to the ***absolute*** temperature, so the shift in the spectrum for a temperature change from 20°C to 300°C (293 K to 573 K) almost doubles the energy at the peak in Figure 4.2. The shift of the neutron spectrum to higher temperatures chiefly affects the reproduction factor (η), as discussed below, because the main effect of changing the neutron spectrum is to change the rates of thermal neutron absorption in the fissile isotopes.

Notes:

Reaction Rates in Fissile Isotopes

Many of the materials in the reactor have absorption cross-sections that vary inversely with the neutron speed. As neutron speeds increase, the number of thermal neutrons passing through a region each second increases ($\phi = nv$) but the increased neutron speed reduces the target size ($\sigma_a \propto 1/v$). This means that the absorption rate, $R_a = N\sigma_a\phi$, does not change for 1/v cross-sections.

If all materials exhibited the same 1/v dependence on neutron speed, changes in neutron temperature would not affect reaction rates and there would be no significant reactivity effect to discuss. However, the absorption cross-sections of the fissile isotopes deviate significantly from l/v behaviour.

Figure 4.3 illustrates this with a plot the absorption cross-sections of U-235 and Pu-239 relative to that of U-238. The U-238 cross-section is very nearly l/v, so the rate of absorption in U-238, as for all 1/v absorbers, hardly changes as the temperature of the thermal neutrons change. The ratios plotted show how the absorption rates in U-235 and Pu-239 change with temperature.

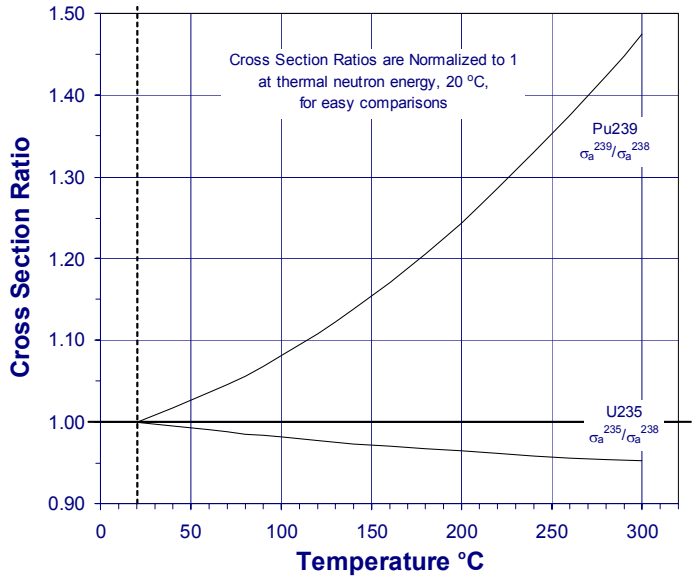

Figure 4.3
Deviation of Absorption Cross-Sections of Pu-239 and U-235 from 1/v

Figure 4.3 shows that, by comparing to U-238, the thermal neutron absorption rate in U-235 drops a little as the neutron temperature increases. Looking at $R = N\sigma\phi$, the absorption cross-section for U-235 drops off a little faster than 1/v, while the increase in thermal flux is proportional to v ($\phi = nv$).

For Pu-239, on the other hand, an increase in neutron temperature causes a very strong increase in the absorption rate in Pu-239. Although the absorption cross section is dropping (see Figure 4.4), it is not decreasing nearly as much as 1/v.

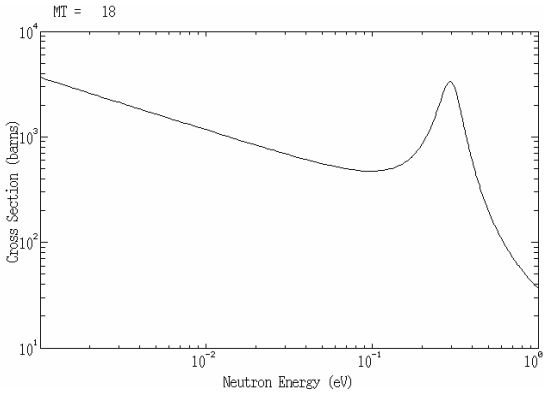

Figure 4.4
Pu-239 Fission Cross-Section

Notes:

For Pu-239, the absorption rate in a CANDU increases with neutron temperature mainly because the cross-section does not drop enough in the thermal energy range. (The bulk of the thermal neutrons are in the energy range below about 0.1 eV.) The large resonance centred at 0.3 eV enhances the absorption rate, because a few percent of the neutrons at the high end of the thermal spectrum move into the resonance region as the temperature rises.

Notes:

In summary, as the neutrons get hotter, the rate of absorption, both fission and capture increases strongly in Pu-239, and they drop a little in U-235.

Variation of η with neutron energy

Unfortunately, knowing how the absorption rates for the fissile isotopes change with neutron temperature does not tell us, directly, how η (a ratio of cross-sections) changes for equilibrium CANDU fuel, which is a mixture of isotopes.

The reproduction factor for CANDU fuel is the ratio of an effective flux weighted average over all the fuel contents:

$$\eta = \nu \frac{\Sigma_f(\text{fuel})}{\Sigma_a(\text{fuel})}$$

The cross sections for both U-235 and Pu-239 affect the numerator and denominator. Each changes differently with neutron temperature, so we need to re-write this expression to show the contributions of each fissile nuclide. It turns out that reproduction factor for CANDU fuel is a weighted average of the reproduction factors for U-235 and Pu-239, as follows.

$$\eta = \frac{\eta_5 \Sigma_a^5 + \eta_9 \Sigma_a^9}{\Sigma_a(\text{fuel})}$$

where the labels 5 and 9 refer to U-235 and Pu-239 respectively. (We've ignored Pu-241).

The reproduction factors η_5 and η_9 are for the pure isotopes—$\nu \times (\sigma_f/\sigma_a)$. The temperature dependence of the microscopic cross-sections and the reproduction factors for the pure isotopes is well known and many textbooks give them. Figure 4.5 shows how η_5 and η_9 change as temperature changes. The upper line in the figure is a horizontal line, to make clear that η_5 decreases a little with temperature. The reproduction factor for Pu-239 drops even more than does η_5, but even for Pu-239, η decreases only about 7% between 20°C and 300°C.

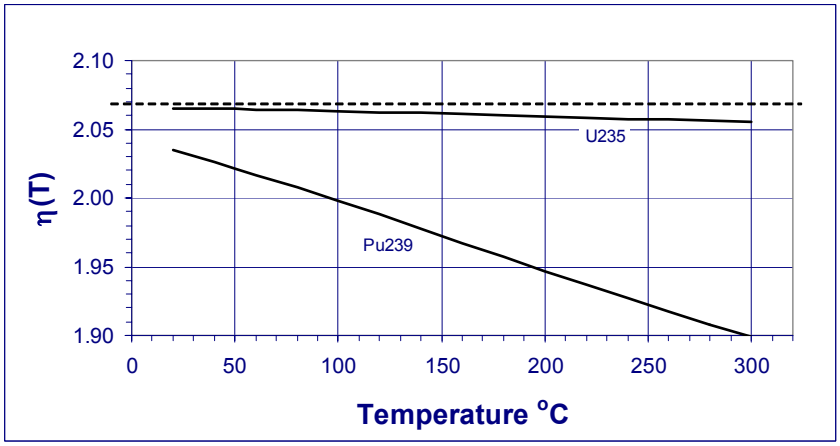

Figure 4.5
Reproduction factors of U-235 (η_5) and Pu-239 (η_9)

Notes:

We can now assemble the pieces to see how η for CANDU fuel varies as the neutron temperature increases. For fresh fuel bundles with only U-235 in the fuel, the situation is straightforward. As the thermal neutron temperature increases, there is a decrease in both the neutron absorption rate in U-235 (Figure 4.4) and the number of neutrons returned per absorption in U-235 (η_5, Figure 4.5). Neither of these variations is very pronounced. It is clear that the contribution of U-235 to the temperature coefficient is negative, but it is not very big.

For fuel with some Pu-239 present, an increase in neutron temperature greatly increases the absorption rate in Pu-239. This outweighs the small decrease in η_9 that occurs at the same time. The contribution of Pu-239 to the overall temperature coefficient is strongly positive. The effect is so large that for an equilibrium-fuelled reactor, it easily dominates the negative U-235 contribution, so:

$\frac{1}{\eta}\frac{d\eta}{dT}$ is positive for equilibrium CANDU fuel.

Other Spectrum Effects

The only factors other than η that the thermal neutron spectrum might affect directly are the thermal utilization factor (f) and thermal non-leakage probability factor Λ_t. (Other factors in the six-factor formula describe fast and intermediate energy neutrons.)

Section 1 gives the thermal utilization factor:

$$f = \frac{\Sigma_a(\text{fuel})}{\Sigma_a(\text{fuel}) + \Sigma_a(\text{non}\cdot\text{fuel})}$$

As we have seen, the Pu-239 content of equilibrium fuel causes a strong increase in neutron absorption as the neutron temperature increases. The non-fuel absorption is almost independent of temperature, so the contribution to reactivity of thermal utilization is positive for equilibrium CANDU fuel. The size of this reactivity increase is not nearly as big as the contribution from η.

Notes:

Thermal neutron leakage increases a little for warmer thermal neutrons. All absorption cross sections drop as neutron energy increases (whether or not the decrease varies as 1/v). Consequently, the mean free path to absorption for thermal neutrons increases (similar to the effect of a density decrease, but restricted to the thermal neutrons). They wander a bit farther from their starting points and are a bit more likely to leak out.

Doppler Broadening in the Fuel

The Doppler Effect arises directly from a temperature change in the fuel. Resonance capture in U-238 increases as the fuel gets hotter for the following reason. The absorption cross-section of U-238 in the resonance region consists of a set of sharp peaks of the sort shown in Figure 4.13. In the resonance region, the probability that a stationary U-238 nucleus absorbs a neutron depends on whether or not the neutron kinetic energy is at the exact energy of a resonance. Both the nucleus and neutron are moving, so the crucial variable in determining the probability of absorption is the speed of the neutron relative to the U-238 nucleus. Heating the fuel makes the atoms of U-238 move more vigorously, so the relative speeds of the neutrons and the U-238 nuclei change.

To demonstrate the effect, consider the stationary nucleus shown at the left of Figure 4.6. A neutron with a speed corresponding to the peak of the resonance has a high probability of being absorbed, while the stationary nucleus is unlikely to capture a neutron travelling slightly slower or faster than this.

Now think about what happens when the fuel is heated and the U-238 nuclei are moving vigorously. A neutron whose speed is such that it previously lay well outside the peak may encounter a U-238 nucleus that is moving at that instant in such a way that the speed of the neutron relative to the nucleus coincides with the peak.

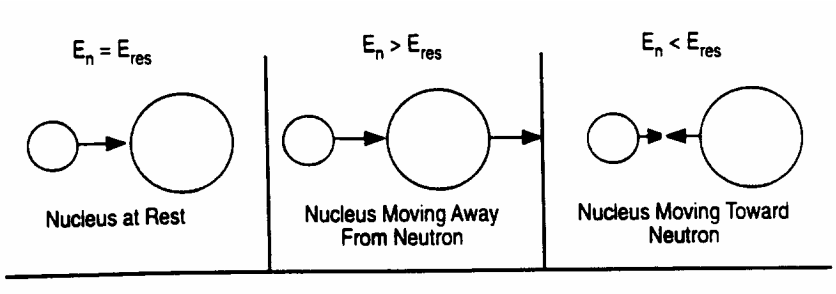

Figure 4.6
Mechanism of Doppler Broadening

Here E_n represents the kinetic energy of the neutron and E_{res} the energy at the peak of the resonance. All three neutrons appear to the nucleus as having the resonance energy.

The centre diagram in Figure 4.6 shows a neutron with speed higher than in the left-hand diagram. However, the neutron encounters a nucleus moving away at a speed that makes the neutron's speed of approach the same as in the first diagram. To the nucleus, this neutron appears to be at the resonance peak energy. The same thing can happen to a neutron with speed below the peak of a nucleus at rest, but which happens to encounter a "hot" uranium nucleus moving towards it at exactly the right speed (right-hand diagram).

Heating the fuel effectively "broadens" the resonances as shown in Figure 4.7. At the same time, increased motion of the U-238 atoms reduces the height of the peak because neutrons moving at exactly the correct speed originally are no longer moving at that speed relative to the moving nuclei. The overall absorption might not change very much if the reduced rate of absorption at the peak were to offset the increased absorption away from the peak.

Notes:

However, the U-238 cross-section is so high in any case that a neutron within the overall peak region is almost certain to be absorbed (although it will likely travel further into the fuel first.)

Notes:

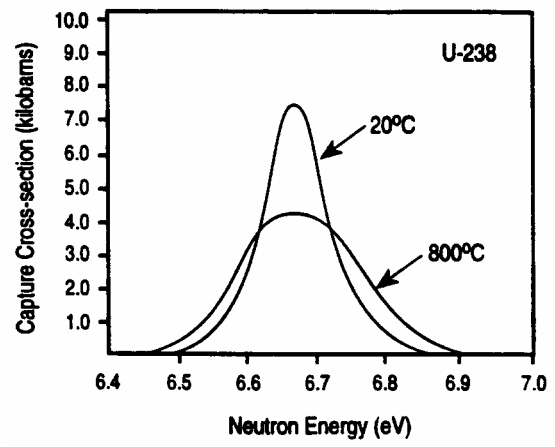

Figure 4.7
Doppler Broadening of U-238 Resonance

Put another way, the peaks are sharper when the fuel is not as hot. With a sharp peak, neutrons at the resonance energy are absorbed close to the fuel surface, (mfp $\approx$ 0.05 mm in UO_2) and neutrons away from the peak energy are likely to pass through the fuel to continue slowing in the moderator. (The surface absorption self-screens the bulk of the fuel from resonance energy neutrons). Hotter fuel exposes a large volume of U-238 to neutrons with energies slightly different from resonance peak, and many of these are absorbed in addition to the absorption of the resonant energy neutrons.

The net effect of heating, however one describes it, is to broaden the range of neutron energies (i.e., the numbers of neutrons) with a high probability of resonance capture. The reduced resonance escape probability (p) results in reduced reactivity.

4.3 TEMPERATURE COEFFICIENTS OF REACTIVITY

We have now introduced the physical effects needed to discuss each of the reactivity coefficients in turn.

The Moderator Temperature Coefficient of Reactivity

Figure 4.8 shows the change in core reactivity as moderator temperature changes. (The curves shown embrace a range of possible equilibrium fuel conditions. Section 6 introduces the burnup unit, n/kb.) The reactivity change in going from 25°C to an operating temperature of 75°C is about 4 mk, for an average temperature coefficient of 4 mk/50°C = +80 μk/°C.

The slope near the normal operating temperature is a little less than this average value, about 75 μk/°C.

Notes:

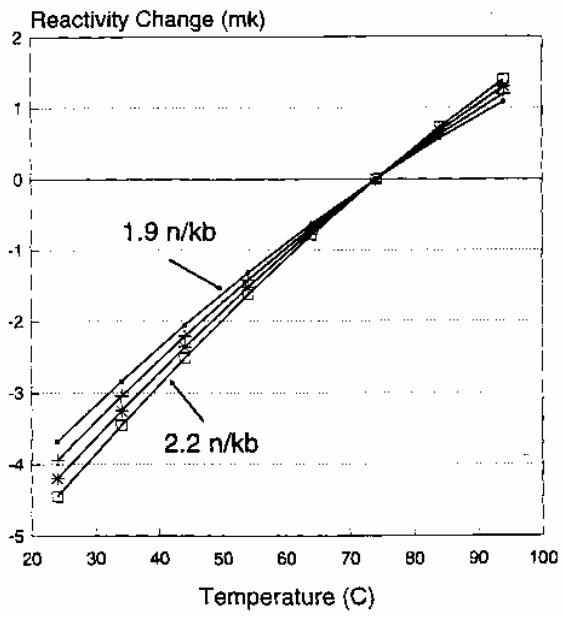

Figure 4.8
Change in Reactivity with Moderator Temperature

One practical effect of the positive moderator temperature coefficient is the possibility of increasing reactivity by adjusting the moderator temperature set point, so that the Moderator Temperature Control program drives the temperature up. For example, suppose the moderator temperature is normally 60°C and the temperature is raised to 65°C. For a temperature coefficient of 75 μk/°C, this produces a reactivity increase of 0.38 mk (5°C × 75 μk/°C = 0.38 mk). This amount of positive reactivity is equivalent to a day or so of fuelling, and is a possible response to unexpected fuelling machine unavailability.

There are few circumstances where it would be possible to gain more than this, as the temperature of the moderator is limited, usually below about 70°C (as measured at the outlet) for safety reasons. Safety Analysis credits the moderator as a "last ditch" heat sink—heat removal by moderator water would limit fuel failures in an accident where no other cooling is available. The analysis requires an upper limit on moderator temperature so that cooling will be effective if it is required.

The explanation of the increase in reactivity with moderator temperature is as follows. Heating the moderator produces two important physical effects: (a) a decrease in moderator density and, (b) an increase in the average neutron energy.

The temperature of the moderator has a big effect on the neutron temperature, so we might expect the reproduction factor (η) to dominate the moderator temperature coefficient. We have seen that the reproduction factor makes a positive contribution to the reactivity as neutron temperature increases.

Notes:

When moderator density decreases, the resonance escape probability (p) decreases. The increase in distance between moderator atoms increases the distance a neutron travels while slowing down; so more neutrons reach neighbouring channels before the moderator has completely thermalized them. (Longer neutron path lengths also increase fast and thermal leakage, but the decrease in the non-leakage probabilities is quite small).

Lower moderator density decreases thermal neutron absorption in the moderator. This results in a positive reactivity contribution from f. The changes in f and p tend to cancel each other but there is a net increase in reactivity that adds to the increase in reactivity from η. The reason the increase in f outweighs the decrease in p is because CANDU reactors are overmoderated. Adequate thermalization occurs even with fewer moderator atoms between the fuel channels, so the main effect of reducing the amount of moderator is a decreased chance of absorption.

The reduction in absorption in the moderator is particularly pronounced when control poisons are present. In this case, there is a much larger reactivity change with temperature than is shown in Figure 4.8.

The Coolant Temperature Coefficient of Reactivity

Notes:

Figure 4.9 shows the effect of changing the coolant temperature. For equilibrium fuel, the coefficient is positive throughout the whole temperature range. Its value near the normal operating temperature (the slope of the curve at that point) is typically about 40 μk/°C. The control system does not adjust coolant temperature independently of the reactor power, but coolant temperature changes a little as the fuel temperature changes.

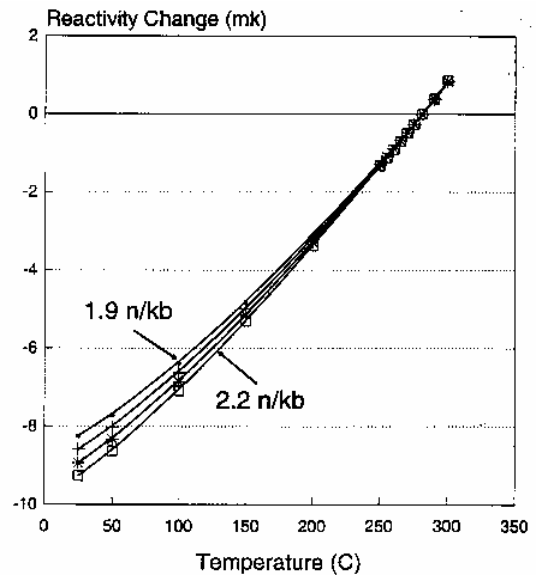

Figure 4.9
Change in Reactivity with Coolant Temperature

From the point of view of the neutrons, the coolant is just more moderator. (The pressure tubes and calandria tubes are nearly transparent to neutrons). The same kinds of reactivity effects that occur in the moderator also occur in the coolant. The sizes of the various contributions are different from the moderator contributions because of the much smaller volume of coolant, the geometry (coolant close to the fuel) and higher operating temperature.

The Fuel Temperature Coefficient of Reactivity

Figure 4.10 shows the effect of changing fuel temperature, as occurs on any power manoeuvre. The average value of the fuel temperature coefficient between hot shutdown and full power is about - 8 μk/°C. Near the normal operating temperature, the slope of the graph is about - 4 μk/°C.

The fuel temperature coefficient arises principally from two factors: one is the Doppler Effect and the other is the neutron spectrum effect. The Doppler Effect reduces the resonance escape probability and this effect on reactivity is very much larger than all other reactivity effects in the fuel. The neutron spectrum shifts to a higher temperature because of the hotter fuel, so for equilibrium fuel, the positive η coefficient partly offset the strong negative Doppler contribution.

Notes:

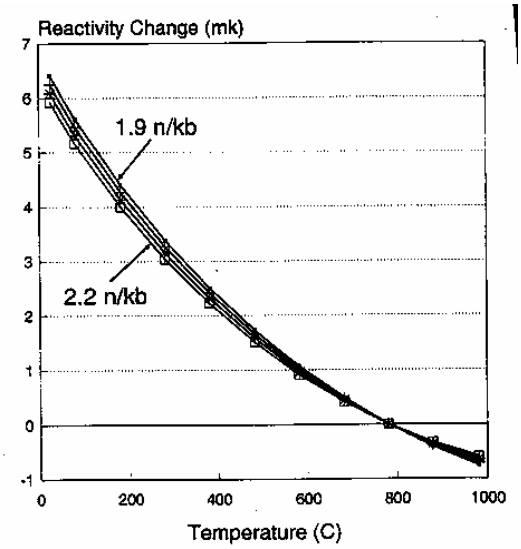

Figure 4.10
Change in Reactivity with Fuel Temperature.

4.4 REACTIVITY VARIATION WITH TEMPERATURE

Notes:

Figure 4.11 shows the variations in overall reactivity in a CANDU 600 as the power changes from cold shutdown to full power. (The figure gives curves for both fresh and equilibrium fuel.). With equilibrium fuel, the net reactivity change in going from cold shutdown to full power is less than +2 mk.

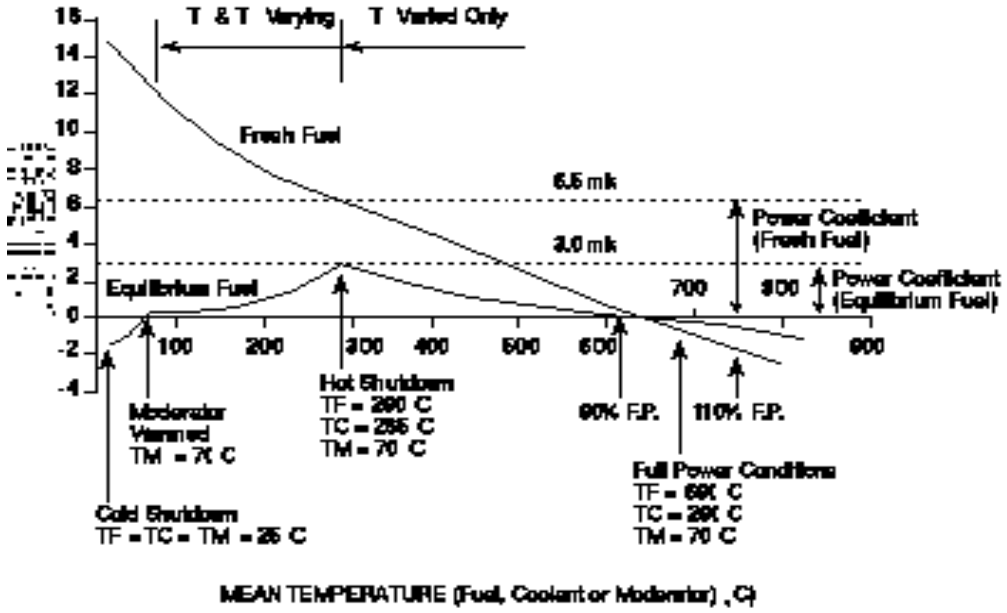

Figure 4.11
Reactivity Changes with Temperature of Reactor Components

The Power Coefficient

Of particular interest in Figure 4.11 is the reactivity loss of about 3 mk in going from the hot shutdown condition to full power. The change in reactivity in moving from the hot shutdown to the 100% power condition is known, somewhat confusingly, as the power coefficient. This "coefficient" is the total reactivity change between the two states rather than the reactivity change per °C, as was the case for the coefficients discussed earlier. The CANDU power coefficient is typically about -2mk to -4mk for equilibrium fuel.

To increase power from hot shutdown to full power, the regulation system, in addition to providing positive reactivity to increase power, must increase reactivity by +3 mk to offset the loss due to heating. The regulating system will insert much smaller amounts at regular intervals to ramp the power up. As the power rises, the loss of reactivity from heating tends to offsets the reactivity insertion. A reactivity insertion big enough to start the power rise is soon overcome by the negative power coefficient, so, for a moderate rate of power increase, the regulating system is continuously playing "catch up" to keep the power rising at the demanded rate.

There is little risk of overshooting the target power, and the regulating system achieves smooth control without complicated feedback circuits.

Notes:

Conversely, small negative reactivity insertions do not shut the reactor down. Suppose a rod falls accidentally into the core, inserting one mk of negative reactivity. Ignoring the regulating system response to this, the loss of 1 mk "uses up" about 1/3 of the power coefficient, so the power drop will slow and stop somewhere below 70% full power, with the reactor critical.

Example

We can calculate the overall reactivity change in going from one condition to another, given the temperature coefficients and temperature changes for each component. Let us estimate the loss of reactivity in going from hot shutdown to full power (the power coefficient). The values of the temperature coefficients involved, averaged over the appropriate temperature ranges, are:

Fuel coefficient: -8 μk/°C

Coolant coefficient: +40 μk/°C

Moderator coefficient: +80 μk/°C

Using the temperature values in Table 4.1, we expect effective average fuel temperature to rise about 500°C from hot shutdown to full power. The coolant temperature may rise about 25°C, (except for the Pickering reactors, which ramp boiler pressure down as power rises, to keep the average coolant temperature nearly constant). Moderator temperature control keeps the moderator temperature nearly constant. (Typically, moderator temperature control keeps the outlet temperature constant, but the bulk temperature rises a little.) Any change in moderator temperature that does occur is relatively slow because of the large mass of moderator water.

The expected reactivity change in moving from hot shutdown to full power is the sum of the contributions from Fuel, Moderator, and Coolant.

Fuel: -8 μk/°C × 500 °C = -4 mk

Coolant: +40 μk/°C × 25°C = +1.0 mk

Moderator: +80 μk/°C × 0°C = 0 mk

TOTAL -3 mk

Self Regulation

To see whether the fuel, moderator, and coolant temperature coefficients, working together, provide the sort of "self regulation" described in Section 4.1, we must consider:

- The size of the three temperature coefficients;
- The size of each temperature change when a power change occurs;
- The time for each component to heat up.

Notes:

In a fast power transient, the fuel temperature increase is much larger and occurs more immediately than any coolant temperature increase. The coolant temperature is likely to lag the fuel temperature by several seconds. Therefore, the negative fuel temperature coefficient provides some degree of the desired self-regulation. This effect arises because of Doppler broadening of the U-238 resonances in the fuel as the fuel heats.

Ignoring the response of the regulating and safety systems, after an upset that inserts a small amount of positive reactivity into a critical core at low power, power will not rise indefinitely. For a large reactivity upsets, the power may continue to rise beyond full power, but the power coefficient reduces the rate of rise, giving time for automatic safety system response.

4.5 VOID REACTIVITY

Voids form if the moderator or the heat-transport system fluid boils. Void formation is more likely to occur in the pressurized coolant than in the moderator, so we will restrict our discussion to the effects of loss of coolant.

Reactivity increases with loss of heavy water coolant, and this produces a power pulse just as heat removal capability is deteriorating. The large loss of coolant accident (large LOCA) defines many of the requirements of the emergency safety systems.

The overall reactivity effect of voiding the coolant is the insertion of positive reactivity. The total reactivity change for full core voiding is typically in the range of +7 to +13 mk, depending on the model of CANDU. Notice that the size of this reactivity is more than enough for prompt criticality.

Fortunately, it takes time for all the coolant to flash to steam through a rupture; long enough for emergency shutdown instruments to detect the power increase and trigger a reactor trip. High rate log, measured by shutdown system ion-chambers, or reactor power high, measured by in-core detectors, are the most likely trips.

Possible causes of coolant boiling are:

- Low pressure (pipe rupture, pressurization system failure);
- Low flow (blockage, pipe rupture, pump failure);
- Excess power (flux distortion, regulating system failure).

Notes:

Under these circumstances, steam gradually displaces the liquid coolant and eventually the channel(s) become totally depleted of liquid.

Voiding fuel channels significantly affects:

- The fast fission factor (ε);
- The resonance escape probability (p);
- The reproduction factor (η); and
- The thermal utilization factor (f).

Of these, the first two are the main contributors to the positive void reactivity.

Voiding causes a decrease in the moderation of neutrons in the immediate neighbourhood of the fuel elements. Looking at Figure 4.12, you can see that a neutron born in one fuel element (for instance, element A) normally passes through some coolant before reaching the next fuel element (element B); coolant molecules usually begin the moderating process. The first elastic collision of a fission neutron with a deuterium nucleus reduces the neutron energy to below the fast fission threshold. A few collisions could reduce the energy into the resonance energy range. (Figure 4.13 shows the radiative capture and fission cross-sections of U-238.) Without coolant present, neutrons in the fuel channel are more likely to be above 1.2 MeV and less likely to be below 10 keV than if coolant is present.

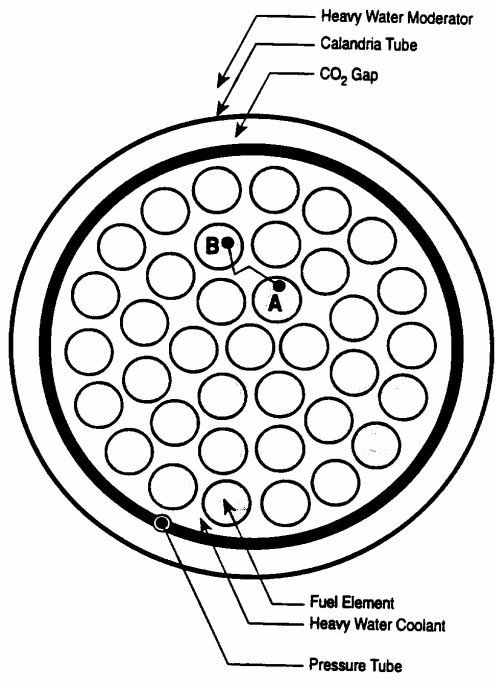

Figure 4.12
Disposition of Moderator and Coolant around the Fuel Bundle

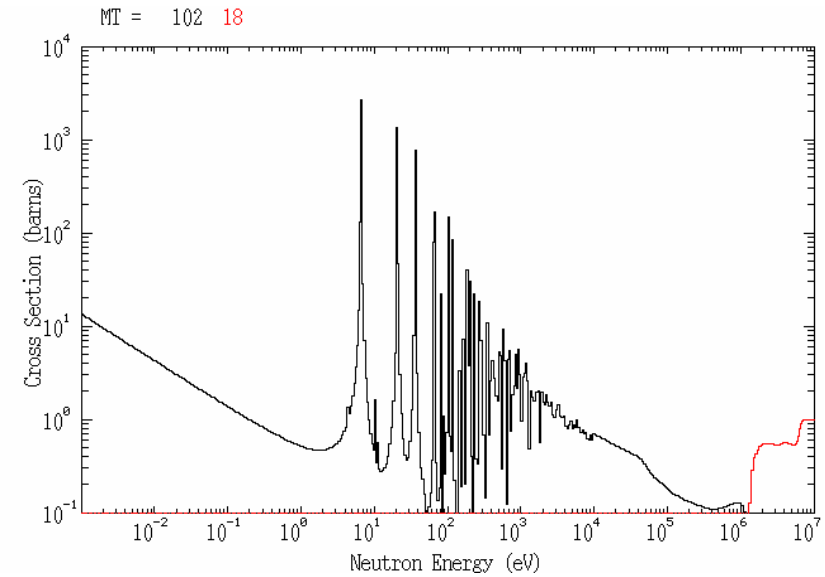

Figure 4.13
Capture and Fission Cross-Sections of U-238

Notes:

The removal of coolant produces the following effects:

- An increase in the fast fission factor (ε), since more fission neutrons have a chance of interacting with U-238 while still above the threshold energy for fission;
- An increase in the resonance-escape probability (p), since fewer fission neutrons reach resonance energy before escaping from the fuel channel.

Notes:

Normally, thermal neutrons drifting back into the fuel channel from the moderator often interact with coolant molecules before the fuel absorbs them. This "re-warming" of the neutrons by the coolant hardens the spectrum. Voiding removes hot coolant from the channel and this reduces the thermal neutron temperature. Without coolant present, the thermal neutrons absorbed by the fuel are "cool" neutrons from the moderator that have not been "re-warmed" by interactions with the coolant. This thermal neutron spectrum change causes a change in η that is not quite as large as the reactivity effect from ε or from p.

- In this case, since we are dealing with a decrease in neutron temperature, the reproduction factor (η) decreases with voiding for equilibrium fuel.
- Voiding the coolant decreases the amount of absorbing material in the reactor (mainly light water impurity in the coolant), increasing the thermal utilization (f). Keeping the coolant isotopic high limits the increase in f on voiding.

License Limits on Coolant Isotopic

Downgrading the moderator water (increasing its light-water impurity) quickly results in the reactor going subcritical because of the high neutron absorption. Low coolant isotopic does not physically prevent reactor operation since relatively fewer of the neutron collisions occur in the coolant. Consequently, administrative procedures enforce a lower coolant isotopic limit, imposed by the operating licence. The limit is usually in the range of 97.5%. If the isotopic is near this license limit, the size of the contribution of f to the void reactivity is comparable to the other contributions. If the isotopic decreases below the allowed limit, f can quickly become the largest contributor to an unacceptably large, positive void coefficient.

An excessively large positive void reactivity causes a large power surge during the void formation. This is likely to cause severe damage to the reactor if the protective systems cannot respond adequately. Trip set points, amount of negative reactivity, and rate of insertion are safety shutdown system parameters designed to accommodate the worst-case accidents; analysis of such accidents determines the operating limits on parameters, such as coolant isotopic.

This is a convenient place to mention the upper limit on coolant isotopic. Accident analysis demonstrates the effectiveness of backup trip parameters for each analysed accident. One of the analysed accidents is the rupture of a pressure tube, causing calandria tube failure such that coolant (at high pressure) displaces moderator water. If this occurs during startup with a heavy poison load in the moderator, the displacement of neutron absorbing moderator water with clean coolant causes a large reactivity increase. This should result in a trip on high rate of power increase. It is also necessary to ensure a timely backup trip. By keeping the coolant isotopic lower than that of the moderator (or setting an upper limit on coolant isotopic), the rate of positive reactivity addition is decreased slightly in the accident just described, allowing time for a process trip.

Notes:

The process of upgrading moderator and coolant heavy water tends to favour upgrading moderator water, simply because the saving in fuel cost is higher for upgraded moderator water. Even in the days before there was an upper limit on coolant isotopic, the moderator isotopic was usually higher than the coolant isotopic. For coolant isotopic lower than moderator isotopic, the reactivity effect of displacing the poisoned moderator, while still high, is not quite as high. Accident analysis credits this slight reduction in reactivity in demonstrating that the backup trips occur in a timely fashion, so the Operating Policies and Principles include a requirement to keep coolant isotopic below moderator isotopic (or specifies an upper limit on coolant isotopic).

EFFECTS OF FISSION PRODUCTS ON CORE REACTIVITY

Notes:

5.0 INTRODUCTION

All fission products absorb neutrons to some extent, so are known as reactor poisons. Most fission product poisons simply build up slowly as the fuel burns up and are accounted for as a long-term reactivity. The neutron absorbing fission products xenon-135 and samarium-149 have particular operational importance. Their concentrations can change quickly, producing major changes in neutron absorption on a relatively short time scale. Each arises from the decay of a precursor fission product, which controls their production rate, but, because they have large absorption cross-sections, their removal changes quickly with changes in flux.

We start by considering the mechanisms for creation and destruction of xenon and its precursor, iodine-135. This allows us to derive expressions for the steady state concentrations of I-135 and Xe-135. It also enables us to analyze the sequence of events following a ***s***hutdown after prolonged operation at power, and to understand why this leads to a rapid increase in Xe-135 concentration.

Any power manoeuvre produces a transient change in xenon concentration. Reactivity changes caused by xenon concentration changes are not as immediate as the reactivity changes caused by a change in fuel temperature, but the size of the effect can be much larger. Unlike temperature feedback, feedback from xenon is positive. For example, a power increase causes an increase in reactivity that causes a further power increase. This is important because the reactivity effect can exceed the capability of the automatic control system to compensate for it.

Another important consequence of the presence of Xe-135 in a CANDU reactor is that it leads to the possibility of xenon oscillations. These can cause reactor power to rise and fall with a period of 15-30 hours, with the possibility of over-rating the fuel. We will describe the process by which xenon oscillations can occur, and see why they necessitate continuous flux monitoring at a number of points in the reactor.

Finally, we will look at the effects of Sm-149. Although it does not create as dramatic an effect as Xe-135, it contributes a significant reactivity at equilibrium, and builds up to a considerably larger load after shutdown. Unlike the xenon reactivity, this does not subsequently disappear, because Sm-149 is a stable isotope.

5.1 XENON AND IODINE BUILDUP

Xenon-135 (often simply referred to just as xenon) is the most important fission product poison. It has a very large absorption cross-section and high production rate.

Notes:

Fission produces Xe-135 in two ways:

Directly from fission—about 0.6% of all fissions in equilibrium CANDU fuel produce Xe-135.

Indirectly from the decay of I-135—fission produces I-135 directly, or as a fission product daughter of tellurium-135. About 6.4% of all fissions in CANDU fuel produce either I-135 or Te-135.

The entire decay chain is as follows:

$$^{135}_{52}\mathrm{Te} \rightarrow ^{135}_{53}\mathrm{I} + \beta^- + \gamma \quad t_{½} = 19\ \mathrm{s}$$

$$^{135}_{53}\mathrm{I} \rightarrow ^{135}_{54}\mathrm{Xe} + \beta^- + \gamma \quad t_{½} = 6.6\ \mathrm{h}$$

$$^{135}_{54}\mathrm{Xe} \rightarrow ^{135}_{55}\mathrm{Cs} + \beta^- + \gamma \quad t_{½} = 9.1\ \mathrm{h}$$

Due to the short half-life of Te-135, we normally consider the entire 6.4% fission product yield to be I-135.

The decay of I-135 to Xe-135 is, in practice, the only way in which iodine is lost. (Removal of I-135 by radiative capture is negligible compared to its decay because I-135 has a very small absorption cross section).

There are two loss mechanisms for xenon; both burnup and decay are important removal processes. At high power, neutron-capture removes much more xenon than does beta decay. An important point is that burnup changes immediately when flux changes, while the 9.1-hour beta decay half-life governs its decay rate.

Xenon Reactivity Effect

Xenon is a strong neutron absorber so its presence in the fuel creates a large negative reactivity in the core. The reactivity worth of the Xe-135 is known as the xenon load. At full power, the reactivity worth of equilibrium xenon in a large CANDU is about -28 mk. This reactivity changes by only a couple of mk above 50% of full power.

It is also common practice to express the concentration of iodine as iodine load (in mk). It is important to realize that iodine alone is not a significant poison; there is no appreciable reactivity associated with it.

The definition of iodine load is the reactivity that it would insert into the reactor if all the iodine present suddenly changes into xenon. Note this point carefully and

understand that we are not talking about a real reactivity that exists in the system, but a reserve bank of potential reactivity that will gradually, and inevitably, become neutron-absorbing Xe-135. For a large CANDU, the equilibrium iodine load at full power (the iodine inventory, or reserve "bank") is worth about -320 mk. The value for any particular reactor depends on the full power flux.

Notes:

Balancing Xenon and Iodine Production and Removal

We now look at the processes of production and loss that determine the buildup of I-135 and Xe-135 when reactor power increases. You may find it useful to refer to Figure 5.1, a block diagram that depicts the production and loss rates of the two isotopes.

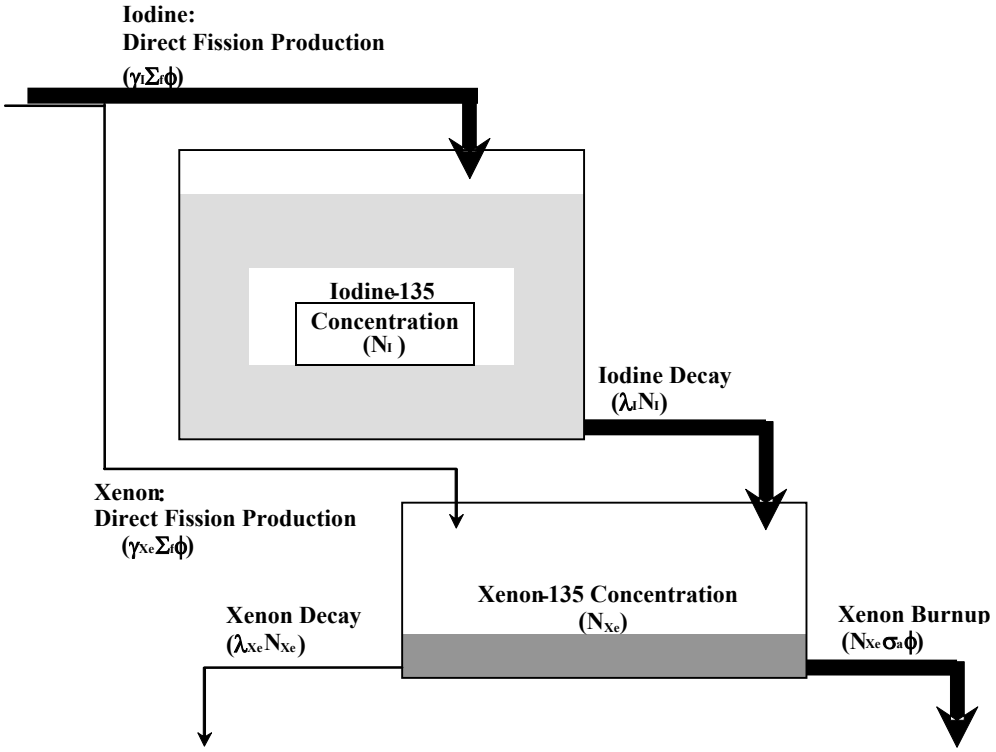

Figure 5.1
Production and Loss Mechanisms for I-135 and Xe-135

The rates of production depend on the fission rate per unit volume:

$R_f = \phi\Sigma_f$

where ϕ is the average thermal neutron flux in the fuel and Σ_f is the macroscopic fission cross-section. Typical values of flux and macroscopic fission cross section in the central region of a large CANDU at high power are $\phi \approx 10^{14}$ n cm^{-2} s^{-1} and $\Sigma_f \approx 0.1$ cm^{-1}.

Table 5.1 shows the microscopic cross-sections, yields, and half-lives for xenon and iodine.

Notes:

The rate of production per unit volume of I-135 from fission $R_f(\text{I-135}) = \gamma_I \Sigma_f \phi$ where γ_I is the fission product yield of I-135 (that is, the fraction of fissions that gives rise to I-135).

Depletion of I-135 occurs by its decay to Xe-135. The decay rate = $\lambda_I N_I$, where λ_I is the decay constant of I-135, and N_I is its concentration in atoms per cm^3.

Subtracting the I-135 loss rate from its production rate gives the net rate of change of iodine concentration at any time. We therefore have (see the upper block in Figure 5.1)

$$\frac{d}{dt}(N_I) = \gamma_I \Sigma_f \phi - \lambda_I N_I$$

For xenon-135, the equation is a little more complicated. We have two production terms, since fission produces xenon directly and as a fission product daughter. We also have two loss mechanisms; both burnup and decay are important removal processes for xenon.

The rates of Xe-135 production (per unit volume) for each production method are:

Direct fission production
$R_f(\text{Xe-135}) = \gamma_{xe} \Sigma_f \phi$ where γ_{xe} is the Xe-135fission product yield.

Iodine decay = $\lambda_I N_I$.

The rates of removal of Xe-135 (per unit volume) for each loss mechanism are:

Burnup
$R_a(\text{Xe-135}) = \sigma_a^{Xe} \phi N_{Xe}$

Xenon decay = $\lambda_{xe} N_{xe}$

where N_{Xe} is the concentration of Xe-135 in atoms per cm^3, λ_{Xe} is its decay constant, and σ_a^{Xe} is its microscopic absorption cross-section.

	Iodine-135	Xenon-135
σ_a (barns)	negligible	3.5×10^6
fission yield	6.4%	0.6%
$t_{½}$ (hours)	6.6	9.1
λ (s^{-1})	2.93×10^{-5}	2.1×10^{-5}

Table 5.1
Properties of Iodine and Xenon

Notes:

Neither Xe-136, formed when Xe-135 captures a neutron, nor Cs-135, from Xe-135 decay, is a strong neutron absorber, so xenon removal reduces neutron absorption in the fuel.

Subtracting the Xe-135 loss rate from its production rate gives the net rate of change of xenon concentration at any time. We therefore have (see the lower block in Figure 5.1)

$$\frac{d}{dt}(N_{Xe}) = \left[\gamma_{Xe}\Sigma_f\phi + \lambda_I N_I\right] - \left[\sigma_a^{Xe}\phi N_{Xe} + \lambda_{Xe}N_{Xe}\right]$$

Iodine and Xenon: Equilibrium Buildup

Looking at the I-135 equation, we see that if we start a reactor (raise the flux to a high, constant value) with no I-135 present in the fuel, the second term on the right hand side is initially zero. It remains negligible compared with the steady production rate as long as the I-135 concentration is low. Thus, I-135 initially builds up rapidly and then, as N_I gradually increases, (so that $\lambda_I N_I$ increases) the net rate of growth of iodine falls (see Figure 5.2).

Eventually, N_I increases to a high enough value that the decay rate is as high as the steady fission production rate, and the net rate of growth becomes zero. I-135 is then said to have reached its equilibrium concentration. In Figure 5.1, the outflow from the upper box matches the inflow.

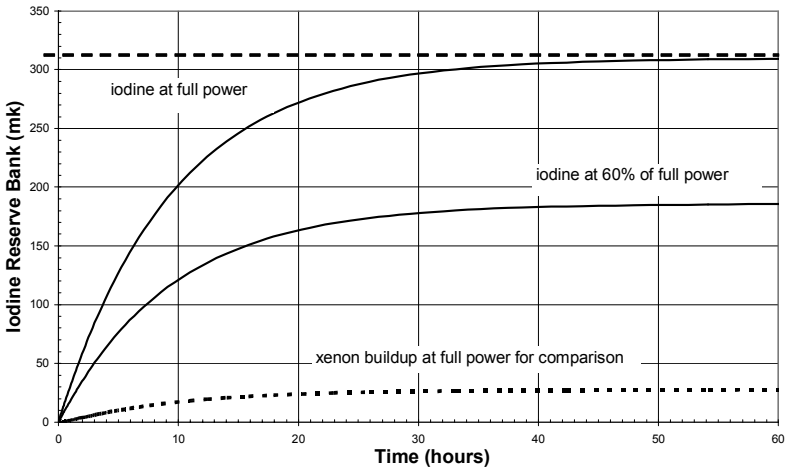

Figure 5.2
Buildup of I-135 to Equilibrium Concentration

Notes:

Under these conditions, the iodine equation becomes $0 = \gamma_I \Sigma_f \phi - \lambda_I N_I(eq)$ where $N_I(eq)$ is the equilibrium concentration of I-135. Hence, re-arranging

$$N_I(eq) = \frac{\gamma_I \Sigma_f}{\lambda_I} \phi$$

This equation gives iodine concentration in the fuel in atoms/cm^3. It is a little easier to understand this expression if we convert it to iodine load in mk. For a large CANDU this is: $N_I(eq) = -320 \text{ mk} \times P$ where P, the neutron power, varies between 0 and 1. The equilibrium concentration is proportional to flux (or power) so, for example, the equilibrium iodine load is -320 mk at full power and at 50% full power the equilibrium iodine load for the same reactor would be -160 mk.

Figure 5.2 shows the approach to this equilibrium value after reactor start-up. The concentration reaches within 2% of its equilibrium value after 40 hours of reactor operation. Note that the equilibrium level of I-135 is directly proportional to the thermal neutron flux ϕ.

The equation for the buildup of I-135 as a function of time is well known

$$N_I(t) = N_I(eq) \cdot \left(1 - e^{-\lambda_I t}\right)$$

Because the exponential term in this equation is ½, ¼ etc. for successive half-lives of iodine, the buildup reaches half way to equilibrium in one iodine half-life, three-quarters of the way in two half-lives, seven-eighths in three, etc.

Now we turn our attention to xenon buildup, which is more complicated than iodine buildup. The initial xenon production rate, directly from fission, is low. As soon as there is appreciable iodine, iodine decay increases the production rate, and then xenon buildup tracks iodine buildup to equilibrium. The short initial delay in xenon buildup is known as "hold-up".

Notes:

At equilibrium, xenon production equals loss (outflow matches inflow in the lower box in Figure 5.1), so the condition for equilibrium is

$$0 = \left[\gamma_{Xe}\Sigma_f\phi + \lambda_I N_I\right] - \left[\sigma_a^{Xe}\phi + \lambda_{Xe}\right]N_{Xe}(eq)$$

where $N_{Xe}(eq)$ is the equilibrium concentration of Xe-135.

Xenon, produced mainly by I-135 decay, cannot reach equilibrium concentration while the I-135 net production rate is changing. Since iodine reaches equilibrium first, we can substitute the equilibrium concentration of I-135 into the xenon equation for equilibrium. With this substitution, we find

$$0 = \left[\gamma_{Xe} + \gamma_I\right]\Sigma_f\phi - \left[\sigma_a^{Xe}\phi + \lambda_{Xe}\right]N_{Xe}(eq)$$

Hence, re-arranging,

$$N_{Xe}(eq) = \frac{(\gamma_{Xe} + \gamma_I)\Sigma_f}{\sigma_a^{Xe}\phi + \lambda_{Xe}}\phi$$

Again, we can convert this equation to xenon load in mk:

$$N_{Xe}(eq) = \frac{-28\ \text{mk} \times P}{0.94P + 0.06}$$

Figure 5.3 shows the buildup of Xe-135 to this equilibrium value of -28 mk over about 40 hours; the upper curve is for the buildup at full power (P = 1), the lower curve is for buildup at 60% full power (P = 0.6). Like I-135, Xe-135 reaches within 2% of its equilibrium value after 40 hours of reactor operation.

The equilibrium Xe-135 concentration shown in Figure 5.3 is nearly the same at 60% and 100% full power. This is quite different than we noted earlier for the equilibrium I-135 concentration, which is directly proportional to the power level. Figure 5.4 shows the dependence of equilibrium xenon on reactor power.

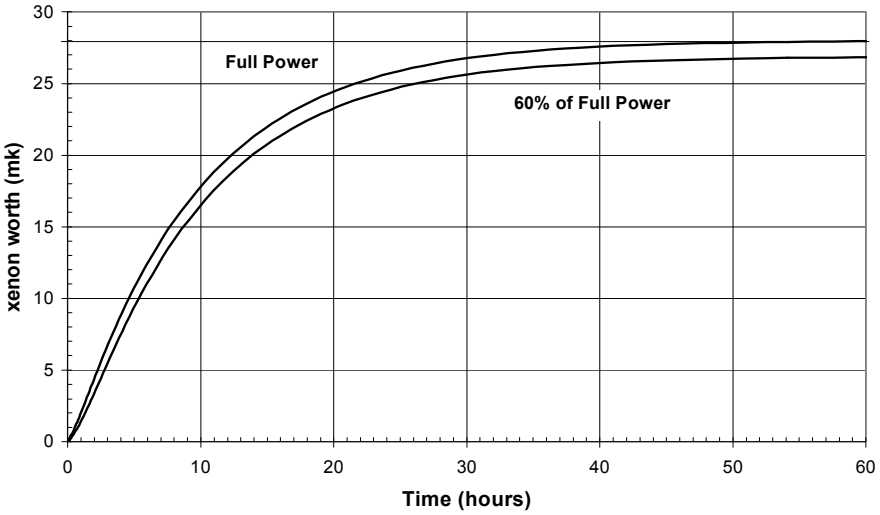

Figure 5.3
Buildup of Xe-135 to Equilibrium Concentration

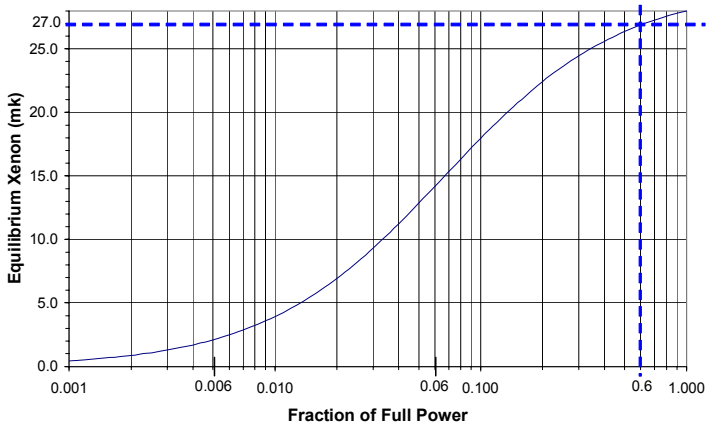

Figure 5.4
Equilibrium Xenon Load as a Function of Reactor Power

We can see why equilibrium xenon is nearly independent of flux for high flux by examining the equation for $N_{xe}(eq)$. Taking the full-power flux for a large CANDU reactor to be about $\phi = 10^{14}$ n cm^{-2} s^{-1}, the sizes of the two terms on the bottom line of the expression are:

$$\sigma_a^{Xe}\phi = 3.5\times10^{6}\times10^{-24}\times10^{14} = 35\times10^{-5}\,s^{-1} \quad \text{and} \quad \lambda_{Xe} = 2.12\times10^{-5}\,s^{-1}$$

At high flux, where the absorption rate is much bigger than the decay rate, we can ignore the second term compared to the first. When we drop the decay constant out of the denominator of $N_{Xe}(eq)$, flux cancels out and the equation becomes independent of flux.

Notes:

This condition applies within the range of about 60% to 100% of full power (see Figure 5.4).

Xenon Simulation

Notes:

The reactor design gives it enough excess positive reactivity to overcome the negative reactivity arising from equilibrium xenon. This means there could be a large excess positive reactivity in the system when there is no xenon present, which will be the case before the reactor has started to operate or on a restart after a long shutdown—when the xenon has decayed. We must be able to compensate for this excess reactivity. The usual way of offsetting the excess reactivity is to dissolve a neutron absorbing poison (boron or gadolinium) in the moderator and removing it as xenon builds up. The addition of poison to the moderator to compensate for xenon that is "missing" is called xenon simulation.

5.2 TRANSIENT XENON BEHAVIOUR

Although the reactor has sufficient excess reactivity to offset equilibrium xenon, a problem occurs following a shutdown after operation at power. Xenon reactivity rises rapidly, peaking in about 10 hours. It takes nearly two days for xenon decay to return the xenon load to near the full power equilibrium value. During this time, there is not enough excess reactivity available to make the reactor critical so it remains shut down.

The explanation of this effect, in broad outline, is not very difficult. Just before shutdown, xenon removal (primarily by neutron absorption) matches the xenon production (mainly from iodine decay). The large reservoir of iodine continues to decay after shutdown, but xenon removal nearly stops. This imbalance produces excess xenon.

To see in detail how this occurs, let us look back at the xenon equation. What we have done below is to indicate underneath each term what its relative magnitude is for a CANDU that has been running steadily at full power long enough to establish equilibrium conditions (See Figure 5.1).

$$\frac{d}{dt}(N_{Xe}) = [\lambda_I N_I + \gamma_{Xe}\Sigma_f\phi] - [\sigma_a^{Xe}\phi N_{Xe} + \lambda_{Xe}N_{Xe}]$$

90% 10% 90% 10%

At equilibrium, the production rate of xenon from iodine decay ($\lambda_I N_I = \gamma_I \Sigma_f \phi$) is much larger than its production rate directly as a fission product ($\gamma_{Xe}\Sigma_f\phi$). At equilibrium, the ratio between production from iodine decay and total production—cancelling $\Sigma_f\phi$—is just $\gamma_I/(\gamma_{Xe} + \gamma_I) = 0.064/0.070 = 0.91$, hence the 90% to 10% (approximate) production ratio.

Consider the xenon loss terms. The removal by burnout depends on the value of the full power flux, so will be somewhat different for different CANDU reactors. We have already established that under full power equilibrium conditions for a large CANDU, the term $\lambda_{xe}N_{Xe}$ is less than one-tenth of the term $\sigma_a^{xe}\phi N_{Xe}$. ($\sigma_a^{xe}\phi N_{Xe} \approx 10 \times \lambda_{xe}N_{Xe}$ for $\phi = 0.6$ of the full power flux.) Consequently, burnup by neutron capture accounts, very approximately, for 90% of the loss of xenon, while its own radioactive decay accounts for only about 10% of the loss.

Notes:

Consider now what happens following a reactor shut down after a long period of operation at full power. The flux ϕ drops to a near-zero value within a minute or so, removing direct xenon production almost immediately, but xenon production by iodine decay continues. Net production remains near 90% of the equilibrium value. On the loss side, we lose 90% of xenon removal, as burnup drops to zero, leaving only radioactive decay. The result is that the xenon concentration starts to rise quite rapidly, fed by the decay of iodine, as shown in Figure 5.5.

This cannot continue indefinitely, because there is a limited quantity of iodine in the core and iodine production by fission stops following reactor shut down. Iodine decays with its characteristic half live; half the iodine load decays in the first half-life (6.7 hours), half of the remaining in the next 6.7 hours, etc.

This generates a lot of xenon immediately after the trip, but less as time goes on (see Figure 5.5). The xenon therefore reaches a peak value, (about 10 hours after shutdown). Thereafter the xenon load gradually decreases, as the reduced rate of iodine decay cannot keep up with the enhanced xenon decay.

Notes:

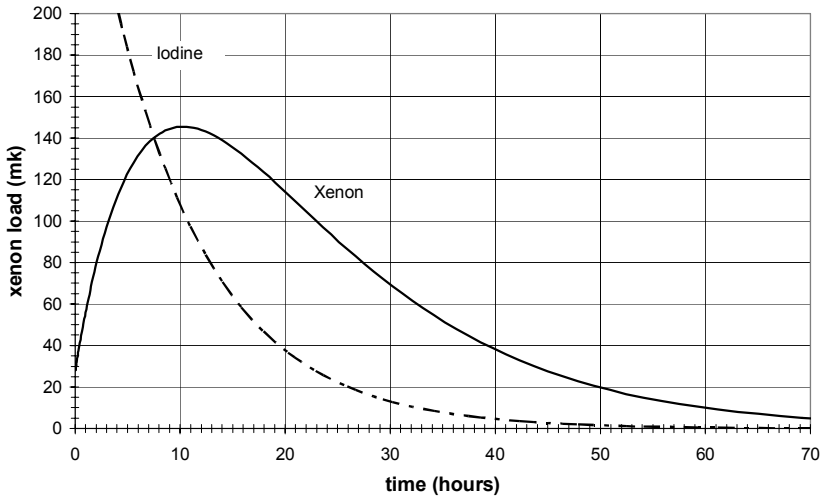

Figure 5.5
Change in Iodine and Xenon Concentrations
Following a Trip from Full Power

Poison Prevent and Poison Override

Before we look at transients that are more complicated than a complete shutdown from full power equilibrium conditions, we will explain the terminology shown in Figure 5.6.

One function of the adjuster rods is to provide excess reactivity to override xenon transients. Withdrawing adjuster rods from the reactor core contributes positive reactivity, up to a maximum of 15 or 20 mk depending on the particular reactor. If the negative reactivity due to xenon exceeds the adjuster reactivity worth, the reactor is sub-critical with no way to restart it. We say it is poisoned out. Figure 5.6 shows that on a trip from full power, a reactor that is poisoned out cannot be restarted until 35 or 40 hours after the trip, when xenon has decayed to near the –28 mk equilibrium level.

Holding reactor power near 60% (or higher) allows sufficient xenon burnout to prevent a poison out, (as Figure 5.8 shows). It is important to realize that on a turbine trip it may make economic sense to keep the reactor operating and to exhaust steam to a condenser (or the atmosphere). We call this mode of operation poison prevent.

Thirty minutes or so after a trip the negative reactivity from xenon exceeds the reactivity worth of the adjusters (the poison override capability). (See Figure 5.6 again.) A reactor restart and return to high power during the 30 minutes or so of poison override time (the time from the trip to a poison out), rapidly burns out the xenon and may prevent a poison out.

Notes:

The adjusters drive out slowly, bank by bank, so, to avoid a poison out, the operators must make the decision to try to restart and begin adjuster withdrawal within about 20 minutes of the trip (the decision and action time).

Poison override is possible in principle, and is part of the reactor design, but is often not practical. Before restarting the reactor following a trip, it is important to find the cause of the trip and eliminate the fault. The operators must make a number of checks before judging the trip "spurious" (that is, a trip that occurs without an actual fault). Checks following a trip (or repairs) usually take longer than the decision and action time, so operating procedures often do not allow the operators to try to "beat the poison out", removing the temptation to take a short cut.

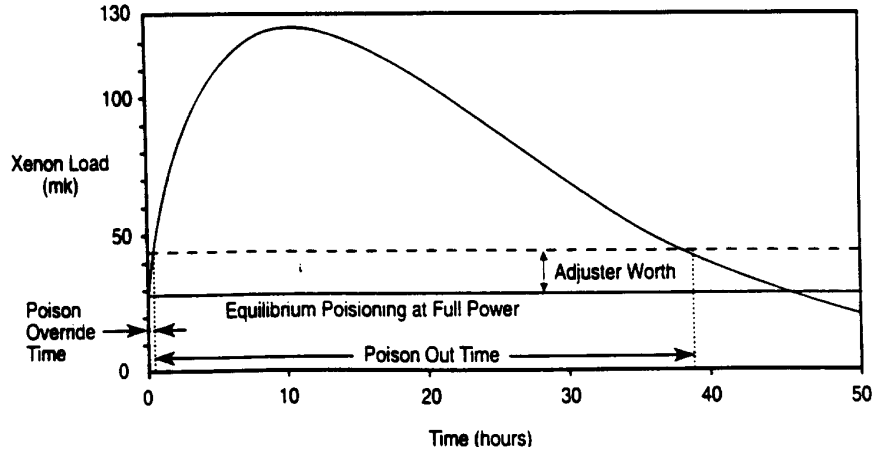

Figure 5.5
Xenon Reactivity Transient for a CANDU 600

The rate of rise of the xenon load after a trip is a function of the equilibrium conditions before the trip. In CANDU reactors, the xenon load increase at about 0.5 mk per minute following a trip from full power. This number and the available reactivity for poison override determine the poison override time. For example, if a particular reactor has a maximum available reactivity of 16 mk, it must be brought back to high power within (16 mk)/(0.5 mk/min) = 32 minutes.

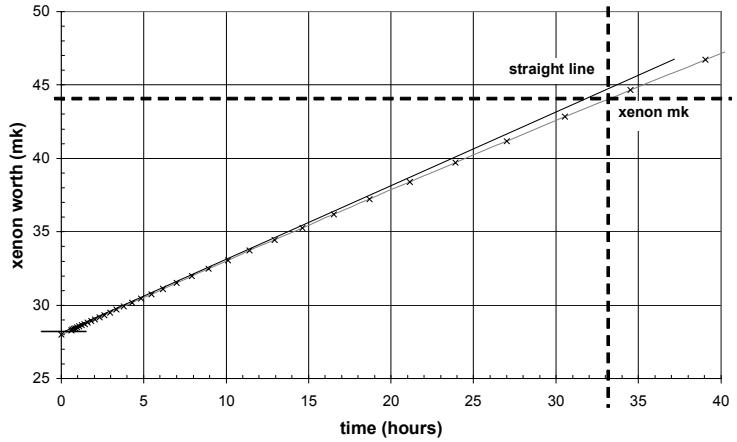

Figure 5.6
The Initial Increase in Xenon Load Following a Trip

Notes:

Example: Estimate the rate of xenon load increase following a trip for a reactor that has an equilibrium iodine load of 300 mk.

The net rate of xenon production immediately following the trip is:

$\lambda_I N_I(eq) - \lambda_{Xe} N_{Xe}(eq)$

Using 300 mk and 28 mk for equilibrium iodine and xenon concentration gives

$2.93 \times 10^{-5}\ s^{-1} \times 300\ mk - 2.12 \times 10^{-5}\ s^{-1} \times 28\ mk$
$= 8.79 \times 10^{-3} - 0.59 \times 10^{-3} = 8.2 \times 10^{-3}$ mk/s
= 0.5 mk/minute

Inevitably, this initial rate falls off, because the increase in xenon load makes the product $\lambda_{Xe} N_{Xe}(eq)$ bigger, while the depletion of the iodine load makes $\lambda_I N_I(eq)$ smaller. Figure 5.6 shows the first 40 minutes of increasing xenon load following a trip to demonstrate that, during the override time, the rate of rise is nearly constant.

Trips from Lower Power

Figure 5.7 shows trips from various power levels compared to a full power trip. In every case, we are assuming equilibrium xenon and iodine loads before the trip. The size of the xenon peak following a trip is nearly proportional to the iodine load before the trip, and the equilibrium iodine concentration, as shown earlier, is proportional to the pre-shutdown flux, so the height of the peak of the xenon load is strongly dependent on flux level before the trip.

Thus, for a trip from full power, the Pickering reactors have a xenon peak of about 80 mk above the equilibrium xenon level, while the Bruce and Darlington reactors have a peak 115 mk or so above equilibrium xenon. (All CANDU reactors have nearly the same -28 mk worth of equilibrium xenon.)

Notes:

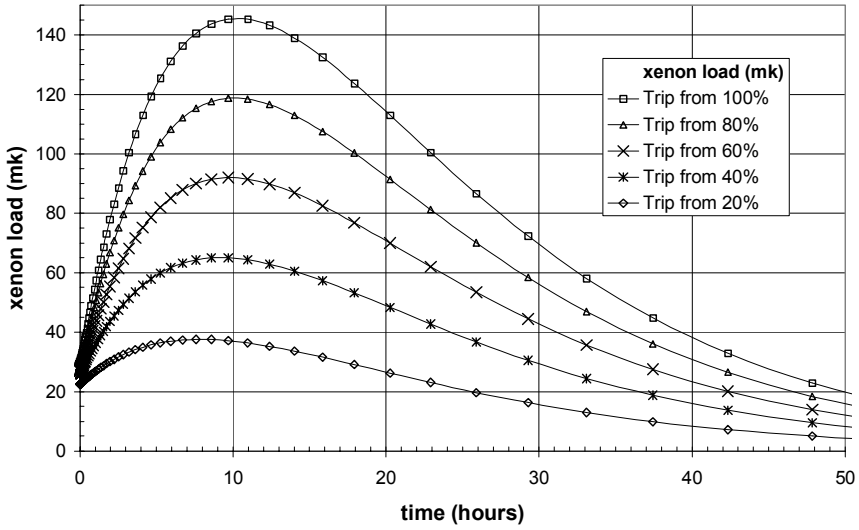

Figure 5.7
Xenon Transient Following Trips
from Various Equilibrium Power Levels

The time to the xenon peak is close to 10 hours for all high power trips, that is, trips that occur with the initial equilibrium xenon close to 28 mk. Peaks following trips from low power occur sooner.

After the peak passes, the xenon decay rate increasingly dominates the decrease in xenon. (Twenty hours after the trip—about three iodine half-lives—the iodine bank is down to about one-eighth of its pre-trip value.) As the iodine drops to insignificant levels, the xenon decays with its half-life ($t_{½}$ = 9.1hours).

As you would expect, Figure 5.7 shows that a smaller initial iodine load makes the poison override time longer and the poison out time shorter for trips from less than full power.

Xenon Transients Following Power Changes

So far, we have discussed xenon transients occurring after a shutdown. The flux goes quickly to zero, and only iodine decay and xenon decay remain as xenon production and loss terms. In practical reactor operation, we are also interested in the transients after a step reduction or a step increase in power. Solving the corresponding xenon equations is a laborious chore and computer codes are normally used.

Nevertheless, you can explain the general shape of the curves by comparing the sizes of the production and loss terms for xenon, just as we did for a trip from full power. Any change in reactor power from the equilibrium steady-state condition produces a transient change in xenon load because of the temporary mismatch between the xenon loss and production terms. You can visualize the changes in concentration as changes in "tank level" when the steady state "flows" are changed in Figure 5.1

Notes:

Figures 5.8 through 5.11 show the results of computer simulations of xenon for a large CANDU reactor following various step changes in reactor power.

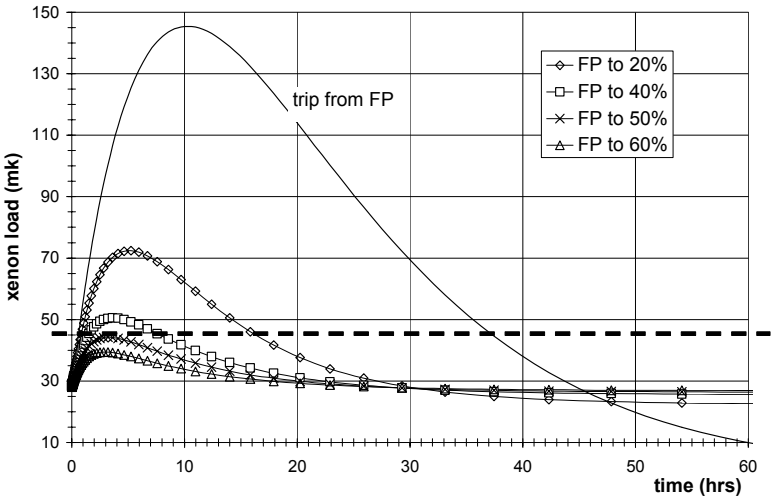

Figure 5.8
Xenon Transients Following Step Power Reductions from Equilibrium Full Power

Figure 5.8 shows the transients for 40, 50, 60, 80, and 100% power reductions from initial full power. For a reduction of, say, 40% (that is, from full power to 60% F.P) the xenon removal by neutron capture also decreases by 40% from its full-power value, but because xenon is still being removed by burnout, the transient will not reach its shutdown peak. Looking at the figure, you will see that for a reduction to 60% F.P., an available excess reactivity of ~15 mk would be more than enough to override the transient altogether.

Ultimately, iodine reaches a new equilibrium at 60% of its full power value, and the xenon reaches equilibrium with the xenon load a little less than -28 mk.

Figure 5.8 also shows that the rate of xenon buildup is less for a 60% reduction (to 40% F.P.) than for a trip. Reactivity peaks at a lower level a little sooner, and the poison override time is longer.

Figure 5.9 illustrates this in more detail. For a fixed amount of excess reactivity, the poison override time depends on the size of the power reduction. For example, the curve shows that the reactor will poison out in 50 minutes following a power reduction to 30% full power, while a reduction to 40% full power extends the override time to more than an hour.

Notes:

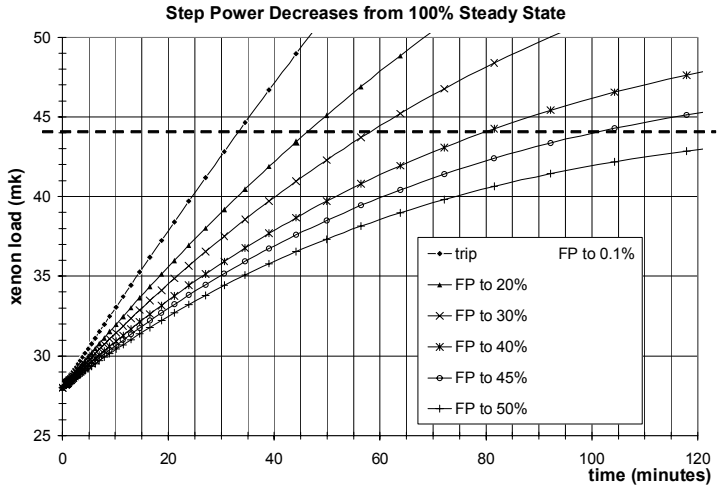

Figure 5.9
Time to Poison Out vs. Size of Step Reduction from Full Power

The converse of these curves also applies. Figure 5.10 compares the change in xenon reactivity when power changes between 60% F.P. and 80% F.P. For example, if the reactor is running at 60% F.P. (at equilibrium) and is taken to 80% F.P., the immediate effect is a gain in reactivity due to increased xenon burnup. At the same time, the higher fission rate produces more iodine, but this does not increase xenon production until later. Consequently, the xenon curve runs through a minimum, and then xenon production rises because of the increasing reservoir of decaying iodine.

Notes:

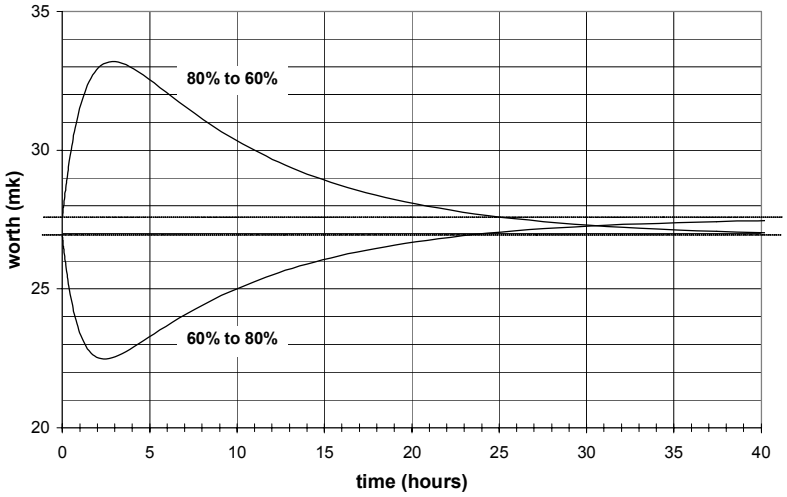

Figure 5.10
Xenon Load Following a Power Increase to 80% from 60% FP and a Decrease from 80% to 60%

The xenon stabilizes at 80% equilibrium load after the iodine has reached equilibrium. The iodine reaches close to its 80% equilibrium value from its initial 60% following the usual half-life behaviour. It gets half way there (to 70%) in one half-life, ¾ of the way (75%) in two half-lives. After 4 half-lives it is nearly at 79%. The iodine bank always proceeds smoothly in this way towards its new equilibrium value. The xenon load reaches its new equilibrium value only after a transient change in the opposite direction. This is caused by the immediate change in removal by xenon burnout, coupled with the delayed change in production.

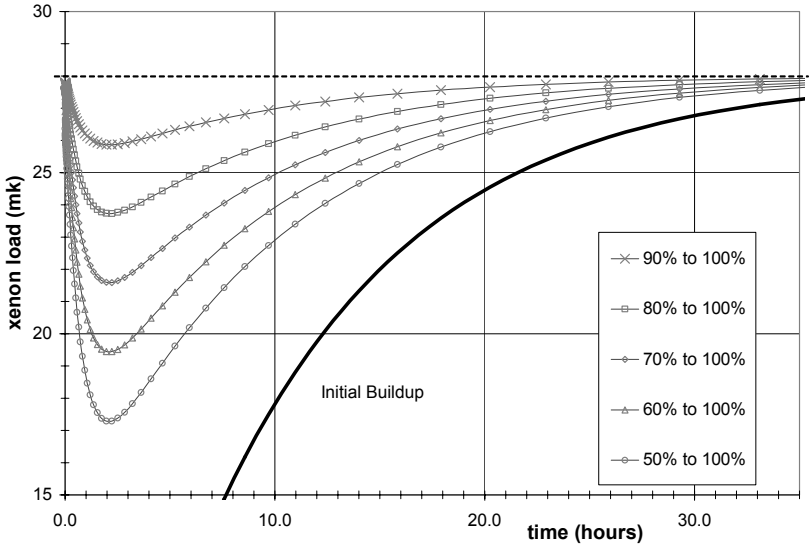

Figure 5.11
Variation in Xenon Load Following Increase in Power to 100% FP

Notes:

Figure 5.11 shows the change in xenon reactivity following an increase in power to 100% from equilibrium operation at various lower powers. Normally, xenon transients on power increases do not present particular operational problems, as xenon simulation with moderator poison offsets the excess reactivity.

Reactor Restart with Xenon Present

If, following a trip, the reactor power increases within the poison override time, (by adjuster removal) the burnout of xenon increases immediately. Consequently, removal exceeds production and the xenon load peaks as power rises and then comes down. Recall that the rate of iodine decay exceeds the rate of xenon decay by about 0.5 mk/minute during the override time. Adding burnout to the xenon decay brings the xenon back near equilibrium quickly. The burnout rate at full power is, as we saw earlier, more than ten times the decay rate, so the "half-time" for burnout is less than an hour.

Figure 5.12 shows the return to equilibrium following restart within the poison override time. The curve assumes an instantaneous return to power, which makes the peak unrealistically sharp. As the xenon burns down to the normal -28 mk load, the regulating system will call for adjuster re-insertion, one bank at a time, whenever the zones reach a high level.

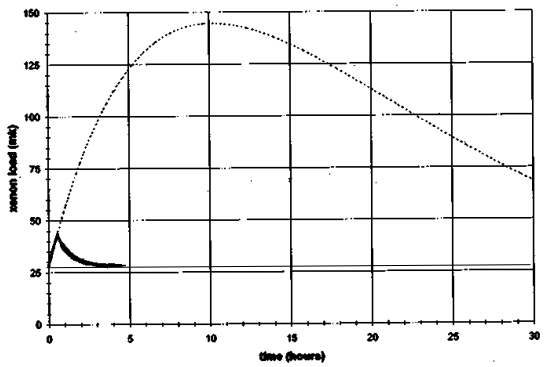

Figure 5.12
Restart During Poison Override

Notes:

Return to power following a poison out (Figure 5.13) is a more usual operating circumstance. Restarting the reactor after the xenon load has returned to about -28 mk but with almost all the iodine gone, produces a xenon removal rate (by decay and burnout) that far exceeds the almost zero production rate.

Xenon simulation offsets the excess reactivity resulting from the rapid decrease in xenon load. The transient positive reactivity looks like the Figure 5.11 curves following power increase, with xenon returning to equilibrium following the return of iodine to equilibrium.

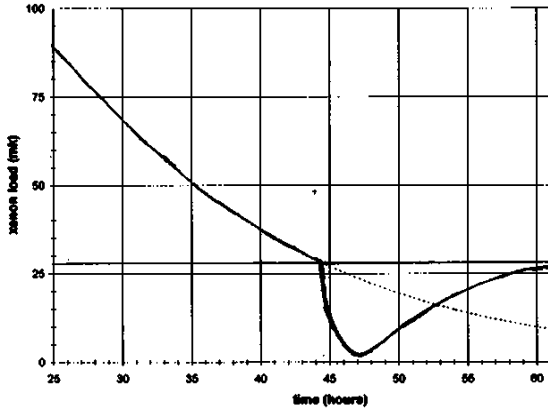

Figure 5.13
Restart After a Poison Out

5.3 XENON OSCILLATIONS

Change in local flux causes a rapid change in xenon load, while the matching change in xenon production is delayed by iodine buildup and decay. This can cause a reactor to undergo periodic (i.e., repetitive) oscillations in flux level. A suitably designed control system limits the size of the resulting flux peaks and stabilizes the flux.

For the moment, ignore temperature reactivity feedback in the reactor and look at xenon in isolation. Consider a small reactor operating at full power and suppose that a small fluctuation takes place in the power with the effect of slightly increasing the flux throughout the system. Because the flux increases, xenon burnup increases and this increases the reactivity of the core, tending to produce a further increase in the flux. There is therefore a positive feedback effect, which pushes the power higher and higher unless some compensating action reduces reactivity to counteract the reactivity increase.

Notes:

In practice, the control system intervenes to prevent this continual rise. Even if the control system is not entirely effective, the upwards drift would eventually stop because iodine load gradually increases, producing more xenon and, eventually, a decrease in reactivity. This would reduce the flux and therefore the burnup of the xenon, magnifying the negative reactivity. The flux would then start to drop and only recover when the reduced production of iodine at lower power resulted in a lower input of xenon. In this way, the reactor power could continue to oscillate in amplitude with a period of several hours between peaks. A control system that reacts promptly to changes in power and maintains it at a more or less constant level easily prevents this kind of bulk power oscillations. Temperature negative reactivity feedback helps damp out small fluctuations.

For a small reactor, careful monitoring of bulk reactor power permits remedial reactor power regulation and prevents oscillations. For a large reactor, however, such as a CANDU, monitoring the overall power level is inadequate, because local xenon oscillations can drive power up in one region of the reactor and down in another, even though the total power remains constant.

Consider how localized oscillations occur. Suppose that the reactor is operating at high power with a uniform power distribution, but the liquid zones in the top of the reactor are a little above their nominal level, while at the bottom they are a little below their nominal level. Although the control system is functioning as intended to keep the flux distribution flat, this spatial configuration deviation from the desired nominal configuration is sometimes referred to as a tilt. It may occur, for example, after refuelling a few channels at the top without refuelling a similar number of bottom channels.

Now, suppose that the regulating system continues to keep the total power output of the reactor constant, but spatial control is lost and the zone levels return to nominal. The flux increases a little in top of the core where the zone levels drop, and simultaneously decreases in the bottom. This deviation of flux shape from the normal flux distribution called a flux tilt. (In this particular example, it is a top-to-bottom flux tilt). In the region of increased flux, xenon now burns out more rapidly than it did before the change and its concentration decreases. The decrease in xenon concentration leads to a higher flux, which again results in increased local xenon burnup, increased local reactivity, increased flux and so on.

Meanwhile, in the region of decreased flux, the xenon concentration increases due to its reduced burnup and to the continued decay of the existing iodine produced in the original, higher flux. This increased xenon concentration decreases reactivity in this region, which reduces the flux, and in turn increases the xenon concentration, and so on. The thermal flux, and hence the power density, decreases in this region and increases in the other, while the total power of the reactor remains constant.

Notes:

These local power excursions do not continue in the same direction forever. At the same time as the increased flux is causing xenon to burn out more rapidly in the high-flux region, it is also increasing the production of iodine. The decay of this enhanced iodine bank eventually leads to an increase in xenon concentration, reducing reactivity and thus the flux and power in that region. Likewise, in the region of reduced flux, the lowered production of iodine combined with the decay of accumulated xenon increases the local reactivity and reverses the flux and power transient in that region.

In this way, unless the regulation system responds to control them, the flux and power may oscillate between different regions (top to bottom, end-to-end, or side to side) indefinitely. Calculations show that ***xenon spatial oscillations*** have a peak-to-peak cycle time of about 15 to 30 hours, and the height of the peaks may increase from cycle to cycle.

Operationally, it is important to recognize that as the peaks turn into valleys and valleys into peaks, there is a point in the cycle where the flux shape is normal. Even though the xenon reactivity is (temporarily) the same in the two regions, the oscillation is not under control. In the region where power is dropping, the iodine load is near its peak and continues to force power down. Conversely, in the region with rising power iodine load is at its lowest, so power continues increasing.

Conditions for Spatial Oscillations

This type of localized xenon oscillation can take place only in a large reactor, a reactor whose spatial dimensions are large compared to the diffusion length of the neutrons. With a small core, a disturbance started in one region affects other regions because neutrons from the affected region spread the disturbance across the core. As mentioned earlier, a regulating system that controls bulk power adequately prevents oscillations in a small reactor.

When the dimensions of the reactor greatly exceed the distance travelled by the thermal neutrons during their lifetime (which is the case in a large CANDU), a disturbance that begins in one place does not spread its influence to a distant part of the core, so the various regions act much more independently. Thus, if a flux increase occurs in one region due to a fuel change, for example, a control system based on maintaining the overall power constant will reduce the flux a little throughout the core to compensate. This would set up a xenon oscillation in the second region exactly out of phase with the one in the first region.

The other condition that must pertain before spatial xenon oscillations can occur is that the reactor operates at high power. When the flux increases at some point in the reactor, the immediate increase in xenon burnup initiates the oscillation. For a marked change in xenon concentration, xenon burnup must predominate over xenon decay. We have already seen that this is the case for a large CANDU, where xenon burnup at full power is at least a factor of 10 larger than xenon decay. Typically, a spatial disturbance cannot become an oscillation in a CANDU unless power is above 25% F.P.

Notes:

The CANDU, then, like several other types of power reactor, satisfies the two conditions required for spatial xenon oscillations to occur. Oscillations can occur with constant (overall) reactor power, so they can continue unnoticed unless instruments monitor the flux and/or power at several points throughout the reactor, and localized absorbers respond to adjust local reactivity to counteract any flux tilts.

Limiting Spatial Oscillations

Douglas Point at 200 MW(e) was the first CANDU to have regional absorbers to control xenon oscillations; it was equipped with four absorbers, which controlled four regions or zones in the reactor.

Pickering A at 540 MW(e) is quite a bit larger and so it is divided into 14 zones, each with its own liquid zone control compartment and a pair of in-core flux detectors. The regulating system uses the detector output to adjust the flow of light water into the zone compartments. All later CANDU designs (Bruce, Darlington, and Pickering B units and the CANDU 600) continue the use of 14 pairs of in-core detectors coupled with 14 liquid zone compartments.

Even with a zone control system, a severe xenon oscillation could risk significant damage to the fuel. The size of the flux tilt depends on the size of the reactivity upset that initiates the oscillation, and on the amount of compensating reactivity that the zone control system provides. A large oscillation could drive one or more liquid zones to their operating limits and spatial control would be lost in those zones. Continued operation with a flux oscillation of such magnitude could lead, at least, to a reactor trip or, more seriously, to dangerously high local fuel temperatures or even fuel meltdown. Even without such severe consequences, xenon oscillations burden the core materials with unnecessary temperature cycling that could lead to premature materials failure.

5.4 SAMARIUM-149

Samarium-149 (often simply referred to just as samarium) has a large absorption cross-section (42,000 barns) and a high production rate (total fission yield is 1.2%). Neither the cross section nor the yield is as big for this nuclide as for Xe-135, so the samarium reactivity effects are much smaller.

Fission does not produce Sm-149 directly. Sm-149 is the fission product daughter of neodymium-149 and promethium-149. The decay sequence is as follows:

Notes:

$^{149}_{60}\mathrm{Nd} \rightarrow {}^{149}_{61}\mathrm{Pm} + \beta^- + \gamma$ $\quad t_{½} = 1.7$ h

$^{149}_{61}\mathrm{Pm} \rightarrow {}^{149}_{62}\mathrm{Sm} + \beta^- + \gamma$ $\quad t_{½} = 53$ h

Because of the short half-life of Nd-149 (compared to the half-life of Pm-149), we consider the entire 1.2 % fission product yield to be promethium.

Like I-135, Pm-149 does not absorb neutrons; only decay can remove it. An important difference between Xe-135 and Sm-149 is that samarium-149 is a stable isotope, and therefore remains in the core after shutdown. Because it is stable, only the process of neutron capture, with the reactor at power, can remove it.

Buildup of Pm-149 and Sm-149

We can write the relationships for rate of change of Pm-149 and Sm-149 concentrations in the same way we did for I-135 and Xe-135. You can easily modify Figure 5.1 to apply to promethium and samarium.

Just like I-135, fission produces Pm-149 and decay removes it. The Pm-149 absorption cross-section is low enough that its rate of burnup (like 1-135) is negligible compared to its rate of loss by decay. Consequently, the upper box in Figure 5.1, with an appropriate change of labels (I $\rightarrow$ Pm) applies to Pm-149, and the equation for the net rate of change of Pm-149 is identical in form to that for I-135.

Promethium buildup, therefore, is similar to iodine buildup (½ way to equilibrium in one half-life, ¾ in two half lives, etc.). Because of its long half-life, it takes about 300 hours to reach equilibrium, compared to 40 hours for I-135. As with iodine, equilibrium promethium concentration is proportional to flux.

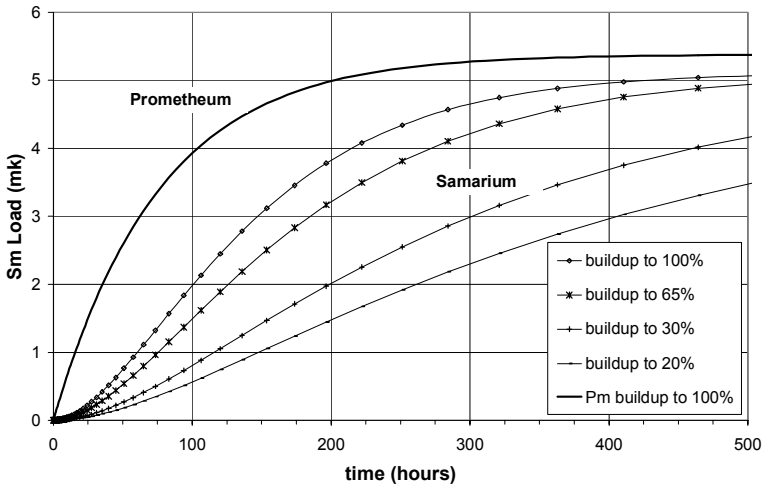

Figure 5.14
Buildup of Pm-149 and Sm-149 in Fresh Fuel Following Startup

Notes:

The expression for Sm-149 buildup is simpler than for Xe-135, because there is no direct production of samarium from fission and no loss by decay. In re-labelling the lower box of Figure 5.1 to apply to samarium, you need to drop the two lightweight lines, the ones for direct production and decay.

Figure 5.14 shows the buildup to equilibrium of samarium and promethium in a high flux CANDU reactor. The upper curve is for promethium, which builds up to an equilibrium load (i.e., reserve inventory) of about -5.5 mk.

The hold-up for samarium is very much greater than we observed in Figure 3 for xenon, and hold-up is more significant at low power. At high power, samarium follows promethium buildup, reaching close to its equilibrium load of about -5 mk in about 300 hours of operation.

Unlike xenon, the time required to reach equilibrium is a function of the flux level, but the equilibrium samarium concentration is independent of the flux (for all power levels). Another important difference from xenon is that samarium does not decay after shutdown, so these samarium buildup curves (Figure 5.14), with no initial samarium, apply only to fresh fuel inserted in the reactor.

Samarium Load After Shutdown

There is a transient rise in samarium concentration after a shutdown because promethium decay continues, but burnup by neutron capture ceases when the flux disappears. The maximum samarium load after shutdown depends on the promethium load before shutdown, which depends on the reactor flux. For CANDU reactors, the maximum samarium load after shutdown is between about 9 mk and 12 mk. Figure 5.15 shows the buildup. Note that the vertical axis starts at the equilibrium samarium load.

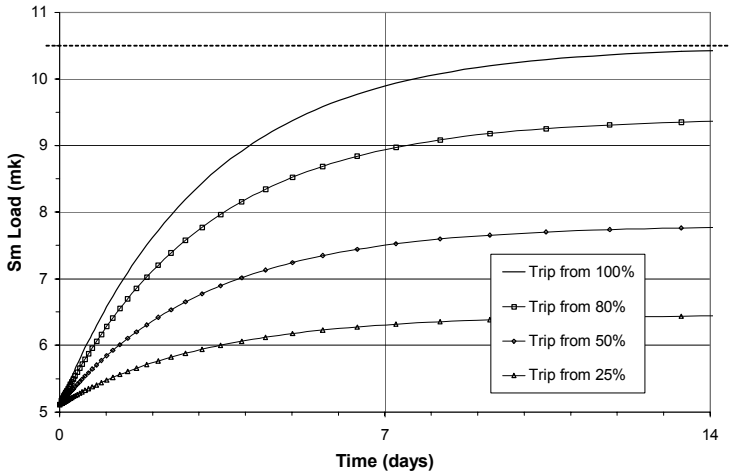

Figure 5.15
Increase in Samarium Load After Shutdown

Notes:

Although reactor design must allow for the equilibrium samarium load, the shutdown load does not cause operational problems. There are two reasons for this:

By looking at the time scale of Figure 5.15, you will realize that the maximum samarium load appears long after the xenon peak decays. There will be lots of reactivity available to deal with the samarium buildup when it occurs. The transient rise in samarium is negligible during the xenon poison override-time, so this does not present a problem either.

In addition to the excess reactivity available following xenon decay, there is a plutonium-239 transient buildup that increases reactivity by a similar amount and at about the same rate as samarium decreases it. The rate at which samarium is formed after shutdown is governed by the Pm-149 half-life of 53 hours which, by coincidence, is almost the same as the half-life of Np-239 (56 hours). Np-239 is the parent of the Pu-239 that adds reactivity as it builds up in the fuel. After shutdown, the Pu-239 starts to increase above its pre-shutdown value because Np-239 decay continues, but Pu-239 burnup stops.

It turns out that the increased reactivity gained from the Pu-239 buildup more than offsets the reactivity loss due to the increased Sm-149. The net result may be a small reactivity gain of a couple of mk.

Samarium Load on Return to Power

Although samarium does not decay during shutdown, it will burn back to equilibrium following a return to power. On restart, the promethium load builds to equilibrium over 300 hours just as described earlier. The burnup rate of samarium is significantly faster than this. It is also faster than the reduction of excess Pu-239 to equilibrium.

The net effect is that, after xenon returns to equilibrium there is an excess reactivity of several mk that disappears in a few more days. This does not cause operational problems, unless normal fuelling stops (in the belief that there is enough excess reactivity). If this happens, the core will require rapid refuelling when the excess reactivity vanishes with the excess plutonium.

Similarly, transient changes in samarium on power level changes are very small compared to xenon, and change very slowly over a week or so. In each case, the samarium level returns to the same equilibrium value and the liquid zones easily correct for small deviations from this value.

Notes:

EFFECTS OF FUEL IRRADIATION AND ON-POWER FUELLING ON CORE REACTIVITY

Notes:

6.0 INTRODUCTION

During reactor operation, neutron reactions continuously change the composition of the fuel. U-235 is burnt up, fissile Pu-239 and Pu-241 are produced (and burnt), while neutron absorbing Pu-240 and fission product poisons build up. In CANDU reactors, on-power fuelling compensates for the gradual decrease in net reactivity that occurs as the fuel is exposed to neutron flux.

This section discusses the principal changes that occur in the fuel and examines their magnitude and the time scale on which they are taking place. First, we outline some characteristics of on-power fuelling. Then we define a unit that specifies the degree of fuel burnup—the neutron per kilobarn. We will then look specifically at the rates of U-235 consumption and Pu-239 build up at full power. These studies enable us to consider how the changing composition of a fuel bundle affects its contribution to the overall reactivity as irradiation proceeds. This determines the fuelling rate that will keep the core critical.

6.1 ON-POWER FUELLING

CANDU reactors use on-power refuelling to maintain adequate long-term reactivity. This keeps the amount of fissile material nearly constant by replacing irradiated fuel with fresh fuel more or less continuously. While there are costs associated with on-power fuelling (capital cost and maintenance costs of the fuelling machines), this system has several distinct advantages over batch refuelling:

- No downtime for refuelling;
- Easy removal of failed fuel, without a shutdown;
- Better average fuel burnup;
- Better flux shaping;
- It avoids the very large poison shim that is required if batch refuelling is used.

In an equilibrium fuelled CANDU reactor, there are small local variations in flux (peaks and valleys) associated with the state of burnup of the fuel locally, but, in contrast to systems with batch refuelling, the overall reactivity of the core hardly changes from day to day and month to month. The locations of the peaks and valleys gradually change as daily fuelling proceeds, but the bulk properties of the core remain the same. Section 8 describes how systematic fuelling controls local flux peaks.

6.2 FUEL BURNUP—GENERAL

Notes:

The initial fuel load of a new reactor is entirely fresh fuel, that is, fuel with no plutonium or fission products present. Poison in the moderator compensates for the excess reactivity of this fuel in the first few months of operation. Exposure to neutron flux gradually changes the composition of the fuel, a process known as fuel burnup. After four to six months of operation at high power, core reactivity drops to a level where a poison shim is no longer required and routine replacement of fuel becomes necessary to maintain core reactivity. The reactor is then described as equilibrium-fuelled. Fuel is replaced on a daily basis (between 8 and 18 bundles per day) to add reactivity at a rate equal to its rate of loss from burnup.

These notes do not describe reactor operation or characteristics of the fresh-fuelled core during the first months of operation. However, we do look at how the composition of individual fuel bundles change during their time in the core.

In a freshly fuelled bundle, the only fissile material is U-235, which constitutes 0.72% of the natural uranium. Exposure to neutron flux gradually depletes the U-235, decreasing reactivity. Buildup of the fission products, especially those with significant absorption cross-sections for thermal neutrons further reduces reactivity. These losses are only partly compensated by the buildup of fissile Pu-239 following neutron radiative capture in U-238 (producing U-239, which subsequently decays to Np-239, followed by a second beta decay to Pu-239). Similarly, the net effect of Pu-240 and Pu-241 buildup is a net loss in reactivity. Eventually, the gradual change in fuel composition requires replacement of the high burnup fuel with fresh fuel to maintain the core critical.

Before reviewing the effects produced by fuel burnup, we should look at how to track the reactivity contribution of an individual bundle, and at how to specify the degree of burnup of fuel. The following two subsections describe these underlying concepts.

Burnup Units

Three different units are in common use for describing the state of the fuel. Each is described below.

Equivalent Full Power Days (EFPD)

Perhaps the simplest way to specify the burnup of a given fuel bundle is by the number of equivalent full-power days (EFPD) it has resided in the core. This is the number of days of exposure to full power flux. (A bundle exposed to 50% of full power flux for two days would have a burnup of one EFPD). Documents for non-specialists, such as public relations publications and management memos, often use this measurement.

Energy Extracted per Unit Mass (MWh/kgU)

Each watt of power production requires about 3.1×10^{10} fissions per second (see Section 1.1). One megawatt-day of heat energy production requires the neutrons to fission about one gram of fissile material. One way to specify burnup, then, is in terms of the total cumulative heat energy extracted from the fuel. The unit is megawatt-hours per kilogram uranium (MWh/kgU). Accountants or people involved in fuel purchasing are most likely to use this measurement. Note that the MWh here are thermal, not electrical, energy.

Notes:

Total Neutron Exposure (n/kb)

Since reaction rate is given by the product $R = \phi\Sigma$ (Section 1.2), the *rate* of burnup is proportional to the neutron flux. The accumulated burnup over a specific period of time (t) is therefore proportional to the product of flux and time (ϕt). This product is known as the total neutron exposure of the fuel; it is the fuel burnup measurement of choice for reactor physicists, the fuelling engineer, and other specialists. The units of neutron exposure are:

$$\phi t \rightarrow \frac{\text{neutrons}}{\text{cm}^2\text{s}} \times \text{s} = \text{neutrons}/\text{cm}^2$$

The total neutron exposures in these units are very big numbers, so a common modification of the unit neutron/cm^2 is the neutron per kilobarn, defined by changing the unit of area to the kilobarn:

$$1\ \text{kb} = 10^3\ \text{b} = 10^3 \times 10^{-24}\ \text{cm}^2 = 10^{-21}\ \text{cm}^2$$

The relation between n/kb and n/cm^2 is then:

$$1\ \text{n/kb} = \frac{1\,\text{neutron}}{10^{-21}\ \text{cm}^2} = 10^{21}\ \text{neutrons/cm}^2$$

The typical average exit burnup for fuel bundles discharged from a CANDU reactor is around 1.8 n/kb. Individual bundles in the core at any given moment have exposures that range from 0 to 2 n/kb or so.

One way to look at this rather arcane unit of exposure is to say that 1 n/kb is equivalent to 10^{21} n-cm per cm^3, that is, the flux exposure measures the accumulated track length of all neutrons that have passed through a unit volume of the fuel during its time in the core.

The approximate relationship among the burnup units is:
100 MWh/kgU $\approx$1 n/kb =10^{21} n/cm^2 $\approx$ 115 EFPD

The units for energy extracted is not exactly proportional to the unit for neutron exposure, as the energy extracted by a given flux exposure is slightly different for

fresh fuel and highly irradiated fuel. The conversion between EFPD and the other burnup units above assumes a CANDU reactor with a full-power average flux of 10^{14} n cm $^{-2}$s $^{-1}$.

Notes:

Tracking Fuel Burnup

Once a reactor has operated for some time, there is a wide variation in the composition of the fuel bundles, since each will have a different degree of burnup of U-235 and buildup of plutonium isotopes and fission product poisons. The fuelling engineer keeps track of changes in individual bundles with a computer simulation that calculates $k_\infty = \eta f p \varepsilon$ for each fuel bundle.

The fuel bundle reactivity is the difference $(1 - 1/k_\infty)$. It is really the reactivity of an infinite reactor in which each unit cell contains a bundle with the same composition as the bundle being tracked.

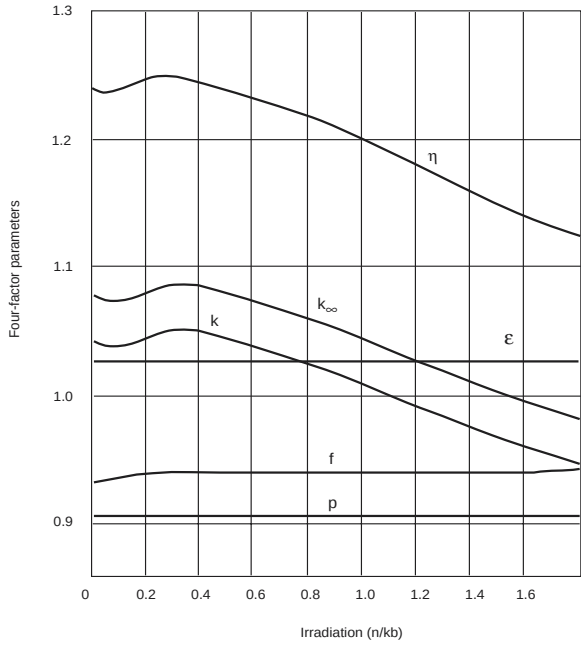

Figure 6.1
Variation with Burnup of the
Four-Factor Parameters, k∞ and k

The leakage factors (Λ_t and Λ_f) depend on the size and shape of the finite reactor. The relation between k_∞ and the effective multiplication factor k, is:

$$k = k_\infty \Lambda_f \Lambda_t$$

The k value in Figure 6.1 is about 30 mk below k_∞, indicating a value $\Lambda_f \Lambda_t = 0.97$.

The equilibrium fuelled critical core with k = 1 consists of a large number of bundles with different k_∞ values determined by their individual degrees of burnup. The lower k_∞ values of the more highly irradiated bundles are compensated by the higher k_∞ values of the less irradiated ones. Each bundle can be labelled with its k_∞ value and its fuel bundle reactivity—i.e. its lattice cell reactivity.

Notes:

6.3 TRANSIENT REACTIVITY CHANGES

The concentrations of Pu-239 and its precursor Np-239 change transiently following shutdown and restart (see Section 5.5). Steady operation produces an equilibrium inventory of the precursor Np-239, which continues to decays following shutdown. The decay to Pu-239 with a 56-hour half-life increases the reactivity attributable to Pu-239 content by 12 mk over about 300 hours. On restart, the excess Pu-239 burns back to equilibrium during the time that Np-239 is rebuilding to its equilibrium value.

This section ignores all such transient effects associated with power changes or shutdowns. We assume that once the fuelling machine inserts a fresh fuel bundle into the core, the flux continuously irradiates the bundle at a steady high level. Transient effects would appear "on top" of the long, steady changes in fuel composition described here.

This section describes the initial buildup of plutonium isotopes and fission products in the fuel, except for the initial buildup of xenon, which is not included here. The initial 40-hour buildup of xenon produces a sharp 28 mk drop in reactivity during the first 0.01 n/kb of flux exposure.

6.4 LONG-TERM REACTIVITY EFFECTS

The following subsections review the rate of burnup of U-235, the rate of buildup of Pu-239, production of Pu-240 and Pu-241, the buildup of neutron absorbing fission products, and the overall rate of reactivity change in the reactor due to the changing composition of the fuel.

Burnup Rate of U-235

The fuel loses one U-235 atom whenever a U-235 nucleus absorbs a neutron. The rate of absorption decreases as the U-235 concentration drops and this produces a characteristic exponential drop in U-235, as Figure 6.2 shows. The removal rate (slope of the curve) is highest when there is the most U-235, and decreases as U-235 decreases.

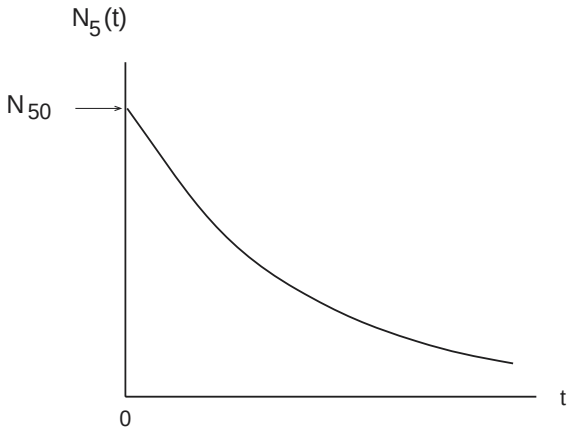

Figure 6.2
Exponential Decrease in U-235 Concentration

Notes:

We can see how this behaviour occurs by looking at the removal process. Neutron absorption removes U-235 from the fuel at the rate $R_a = \phi\Sigma_a^5 = \phi N_5 \sigma_a^5$ We write this as:

$$\frac{dN_5}{dt} = -\left(\sigma_a^5 \phi\right) N_5$$

N_5 is the number of U-235 nuclei per cm^3,

σ_a^5 is the microscopic absorption cross-section of U-235 (in cm^2),

ϕ is the neutron flux in n/cm^2 s.

The negative sign on the right hand side accounts for the fact that N_5 is decreasing.

This equation is of the same form as the radioactive decay equation:

$$\frac{dN}{dt} = -\lambda N$$

where the product ($\sigma_a^5 \phi$) in the burnup equation is equivalent to the decay constant λ in the decay equation, provided ϕ is constant. The U-235 concentration therefore decreases in exponential fashion, as Figure 6.2 shows. It typically drops to about ¼ of its initial value of 0.72% by the time the fuel is discarded.

The equation for U-235 concentration is

$$N_5(t) = N_{50} e^{-\left(\sigma_a^5 \phi\right) t}$$

where N_{50} is the value of N_5 at time t = 0.

Buildup Rate of Pu-239

Each neutron capture in a U-238 nucleus soon produces a Pu-239 nucleus. The percentage change of the U-238 is very small (it constitutes about 99% of the fuel regardless of whether it has been irradiated or not). While the loss of a small amount of U-238 has no noticeable effect on reactivity, the production of a small amount of fissile Pu-239 has a large effect.

Notes:

The buildup of Pu-239, starting with fresh fuel, is quite similar to the buildup of iodine. Figure 6.3 shows that it approaches an equilibrium value, reaching half way to equilibrium in a certain interval, and then to ¾ in the same time interval, etc. We can see how this behaviour occurs by looking at the production and removal processes.

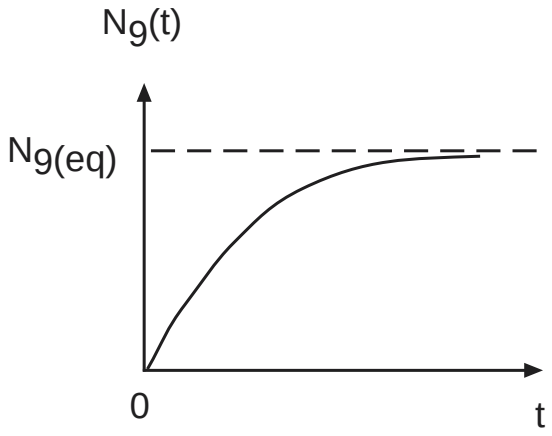

Figure 6.3
Buildup of Pu-239

The steady state rate at which Pu-239 forms (ignoring hold-up—the time it takes the two precursors to decay) equals the capture rate in U-238; or

$$R_a = \Sigma_a^8 \phi$$

where $\Sigma_a^8 = N_8 \sigma_a^8$ is the U-238 macroscopic absorption cross section. The U-238 concentration decreases only a little during irradiation so the macroscopic absorption cross-section is almost constant. (The small thermal neutron absorption cross-section σ_a^8 of 2.7 b causes a low burnout rate. Resonance capture is limited by self-screening.)

Neutron absorption destroys Pu-239, resulting in either Pu-239 fission or capture producing Pu-240. We can ignore Pu-239 removal by alpha decay because its half-life is so long (over 24,000 years). The rate of loss of Pu-239 then equals:

$$R_a = \left(\sigma_a^9 \phi\right) N_9$$

The net change in Pu-239 concentration per cm^3 per sec is:

$$\frac{dN_9}{dt} = \left(\Sigma_a^8 \phi\right) - \left(\sigma_a^9 \phi\right) N_9$$

Notes:

This equation is of the same form as the equation for iodine buildup, encountered in Section 5. Initially, when there is no Pu-239 present, there is a constant production rate. Once there is some Pu-239 present, the net production rate falls off, because of Pu-239 burnout.

The buildup equation for Pu-239 is: $N_9(t) = N_9(eq)\left[1 - e^{-\left(\sigma_a^9 \phi\right)t}\right]$

Given sufficient time, the Pu-239 would build up to an equilibrium concentration.

$$0 = \left(\Sigma_a^8 \phi\right) - \left(\sigma_a^9 \phi\right) N_9(eq)$$

We can calculate this equilibrium value as a fraction of the U-238 content if we replace Σ_a^8 with $N_8 \sigma_a^8$ to give

$$\frac{N_9(eq)}{N_8} = \frac{\sigma_a^8}{\sigma_a^9}$$

Substituting σ_a^8 = 2.7 b and σ_a^9 = 1013 b gives $N_{9(eq)}/N_8$ = 0.27%, which compares well with data presented in Table 6.1. The agreement is better than you might expect, as we have ignored Pu-239 production by resonance capture in U-238 and have ignored the strong temperature dependence of thermal neutron absorption in Pu-239.

In practice, due to the overall loss of reactivity from burnup, it is necessary to start replacing the fuel in the core well before the Pu-239 has reached equilibrium.

Pu-240 and Pu-241 Buildup

As noted earlier, some of the neutron absorptions in Pu-239 (about 27%) produce Pu-240, which is non-fissile and has a large absorption cross-section for thermal neutrons (290 barns). As significant amounts of Pu-239 begin to build up in the fuel, Pu-240 also starts to appear. In principle, Pu-240 also eventually reaches an equilibrium value, but actual fuel irradiation times are much too short for this to occur. As far as we are concerned, Pu-240 is a poison that increases at a nearly steady rate as the fuel burns up.

Pu-240 has one useful characteristic; neutron capture in Pu-240 creates fissile Pu-241, with similar properties to Pu-239. This provides only inadequate compensation for overall reactivity loss because only a small quantity of Pu-241 is produced during the fuel irradiation time.

Fission Product Buildup

The rate at which any given fission product builds up depends on the size of its absorption cross-section. Through an argument similar to the one for Pu-239, we can demonstrate that each fission product approaches an equilibrium concentration at a rate that depends on its absorption cross-section.

Notes:

The high cross-section fission products approach equilibrium relatively quickly. They are known as saturating fission products. The two most important of these are Xe-135 and Sm-149. The reactor design must include enough excess reactivity to allow operation with saturating fission products present. The weaker absorbers, the non-saturating fission products, build up gradually over long periods. Their neutron absorption significantly affects the requirement to replace high burnup fuel while it still contains significant amounts of fissile material.

Changes in Reactivity with Burnup

Table 6.1 lists the concentrations of the fissile isotopes U-235, Pu-239, and Pu-241 as a function burnup. Figure 6.4 plots these concentrations as a function of burnup. The figure indicates the approximate exit irradiation for Point Lepreau, Pickering A, and Bruce A fuel. The higher burnup achieved by the Bruce A reactor is because it does not have adjuster rods in the core.

Burnup		Concentrations (g/kgU)		
n/kb	MWh/kgU	U-235	Pu-239	Pu-241
0	0	7.20	0	0
0.2	19	6.37	0.60	0.002
0.4	39	5.62	1.10	0.009
0.6	59	4.90	1.48	0.025
0.8	79	4.30	1.77	0.049
1.0	100	3.76	1.98	0.078
1.2	120	3.32	2.14	0.107
1.4	140	2.90	2.25	0.145
1.6	159	2.56	2.33	0.177
1.8	179	2.26	2.39	0.211
2.0	198	1.98	2.43	0.245
2.2	216	1.74	2.46	0.278
2.4	235	1.54	2.48	0.309
2.6	253	1.35	2.49	0.338
2.8	271	1.18	2.50	0.366
3.0	289	1.03	2.50	0.393
*** Strictly speaking, the values shown in this table apply only to the Pickering reactors, but they will be correct to within a percent or so for all natural uranium, D_2O moderated reactors.**				

Table 6.1
Burnup Data*

Notes:

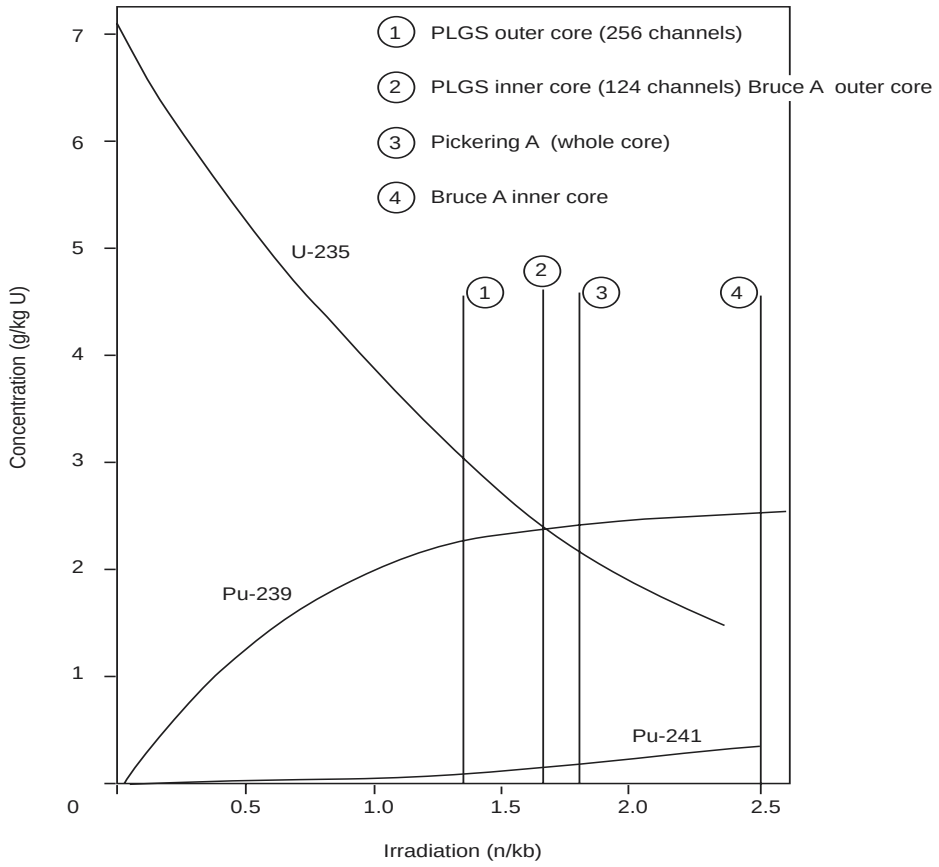

Figure 6.4
Fissile Isotope Concentrations as a Function of Burnup

Notes:

Figure 6.5 shows the more important components of the change of fuel bundle reactivity (that is, the reactivity of the unit cell) with irradiation. Each fuel bundle goes through the composition changes associated with the concentration changes shown in Figure 6.4. Initially, as indicated by the combined curve (U-235 + Pu-239), the positive reactivity contribution of the Pu-239 overcomes the negative contribution due to burnup of U-235, so that the net effect is an increase of reactivity. Initially, only eight Pu-239 atoms are produced for every ten U-235 destroyed, but the higher fission cross-section of the Pu-239 (742 b compared to 580 b for U-235) more than compensates for this ratio being less than one.

Eventually, the Pu-239 buildup rate slows. The fall-off in net Pu-239 production means that, even though its concentration continues to increase, it can no longer compensate for the continuing U-235 burnup, and the curve turns over.

The creation of Pu-240 produces a steady reactivity decrease, shown in the figure. The buildup of fissile Pu-241 partly compensates for this reactivity loss. The combined effect is a relatively steady decrease in reactivity.

The neutron-absorbing fission products cause a negative reactivity that builds up continuously as irradiation proceeds. The initial steep part of the curve labelled "fission products" in Figure 6.5 is due to these, and particularly to Sm-149. Note that the effect of the very strongly absorbing fission product xenon-135 is ***not*** included in Figure 6.5.

Notes:

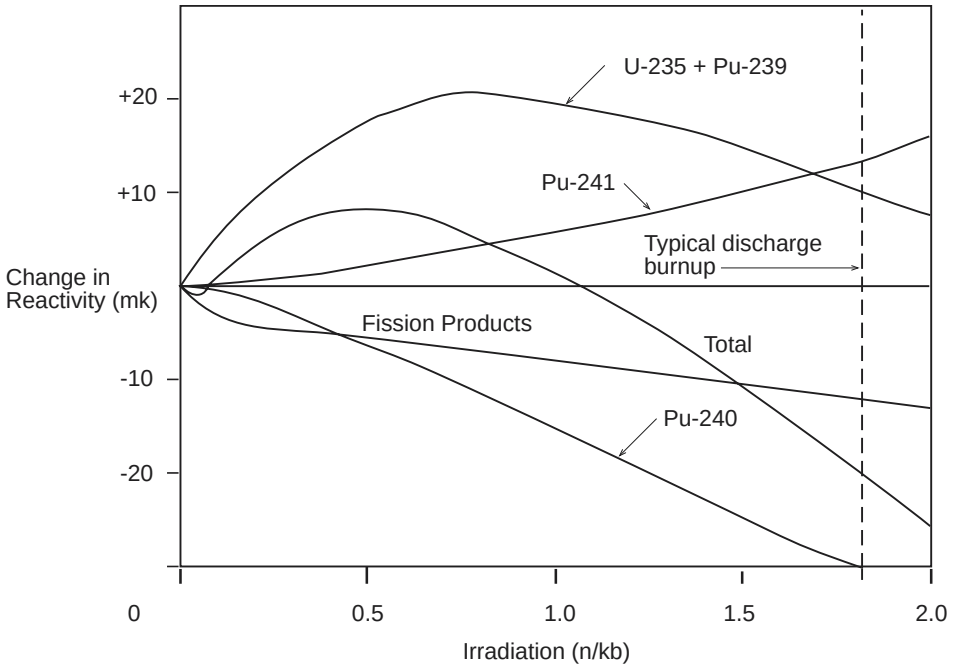

Figure 6.5
Components of Change of Reactivity with Burnup

The net effect of all these contributions is shown by the "total" curve in Figure 6.5. A comparison with Figure 6.1 shows it is essentially the same as the curve for k_∞. The factor that describes the observed changes in k_∞ is the reproduction factor (η). We recall from Section 1.4 that:

$$\eta = \nu \frac{\Sigma_f(\text{fuel})}{\Sigma_a(\text{fuel})}$$

The change in "fuel bundle reactivity" (by which we mean the reactivity of a unit cell containing such a bundle), is determined by how the fission and absorption cross sections in the fuel change due to changes in concentration of fissile isotopes and fission products.

The sharp initial dip in reactivity comes from the accumulation of the saturating fission product Sm-149, which builds to near its equilibrium value over the first week or so of operation. Sm-149, with its high thermal neutron cross-section, increases the absorption cross-section significantly without affecting the fission cross-section. Once this reaches equilibrium, the initial Pu-239 buildup more than compensates for U-235 burnup.

The reactivity increases because the change in concentration of the fissile materials affects the numerator of η more than the denominator. The net positive contribution to fuel bundle reactivity continues until the Pu-239 growth rate falls off as it moves toward equilibrium, and no longer compensates for continuing U-235 burnup. The peak in the total reactivity curve (at an irradiation near 0.5 n/kb) is known as the plutonium peak.

Notes:

The steady buildup of non-saturating fission products and Pu-240 in the fuel makes the decrease in reactivity steeper. Buildup of Pu-241 reduces that rate of reactivity loss somewhat, but is unable to reverse it. The steady buildup of Pu-240/Pu-241 and the non-saturating fission products combine to increase the denominator of η, with only a small increase in the numerator from Pu-241.

At some point, therefore, it becomes necessary to start removing partially burnt-up fuel, replacing it with fresh fuel, well before the fissile content is used up. Section 8 considers the practical aspects of choosing channels for refuelling.

The total reactivity curve crosses the axis (reactivity = 0) near one n/kb. The equilibrium fuelled core contains a mixture of bundles with reactivity (k_∞) values above and below this irradiation. In the equilibrium-fuelled core, reactivity drops steadily with irradiation. The CANDU 600, for example, consumes about 0.4 mk per day (that is, the reactivity worth of the fuel diminishes at that rate). The slope of the curve near k = 1 in Figure 6.1 determines the rate of loss of reactivity, which determines the refuelling rate. To restore 0.4 mk of reactivity each day by refuelling requires replacement of some fuel in each of 2 to 3 fuel channels each day.

The average burnup of fuel discharged from the reactor is larger than one might expect from the graphs, which show k = 1 for an irradiation of 1.1 n/kb. The figures assume a reactor core that is burning up uniformly, while the equilibrium-fuelled reactor contains a mixture of fuel bundles of differing degrees of irradiation. Some of these bundles (the ones nearing discharge) have an irradiation appreciably greater than 1.1 n/kb, but fresh bundles, with irradiation below 1.1 n/kb, compensate for the reactivity "deficit" of the high burnup bundles. The efficiency in using the fissile material in the fuel is generally reported by giving the discharge burnup (sometimes called the exit burnup) which is likely to average 1.8 n/kb or more.

If the refuelling machines are unavailable for some reason, the reactor can continue to operate for only a limited time. For example, if the liquid zone compartments were all 50% full, about 2 mk of excess reactivity would be available. This would allow nearly five days of operation by gradual reduction of liquid zone levels.

REACTOR OPERATIONS AT LOW POWER

7.0 INTRODUCTION

Notes:

This section combines reactor dynamic behaviour, discussed in earlier sections, with reactor operations at low power. We begin with some power definitions and then describe the power rundown of an equilibrium fuelled CANDU reactor that has been operating at steady full power. We will describe the rundown of neutron power in terms of its dependence on prompt neutrons, delayed neutrons, and photoneutrons and compare the neutron power rundown with the much slower decrease in thermal power caused by fission-product decay heat. We then look at the shutdown state and discuss procedures for approaching criticality. Finally, we look at a reactor operated in the low power (almost) critical state and review feedback effects that occur as power rises into the last decade or so of reactor operation.

7.1 THERMAL POWER, NEURON POWER, AND FISSION POWER

The power referred to most frequently in this course is neutron power, which is essentially the fission rate multiplied by the prompt energy release per fission. We cannot measure this directly, but we do monitor the flux and the average neutron flux in the core is proportional to the overall fission rate.

Reactor thermal power (often shortened to thermal power) is the reactor's rate of heat energy production. Thermal power takes account of nuclear decay heating and of conventional heat (pump heat and heat losses to the environment).

Fission power is the name given to the heat generated because of nuclear processes in the fuel. This includes wasted heat, such as that generated in the moderator and shielding. It does not include any conventional heating. The operating license places an upper limit on fission power, which is enforced by regulating the reactor thermal power.

Figure 7.1 illustrates the relationship among these powers for a CANDU 600. The numbers will differ from reactor to reactor, but the per cent values are quite similar.

Notes:

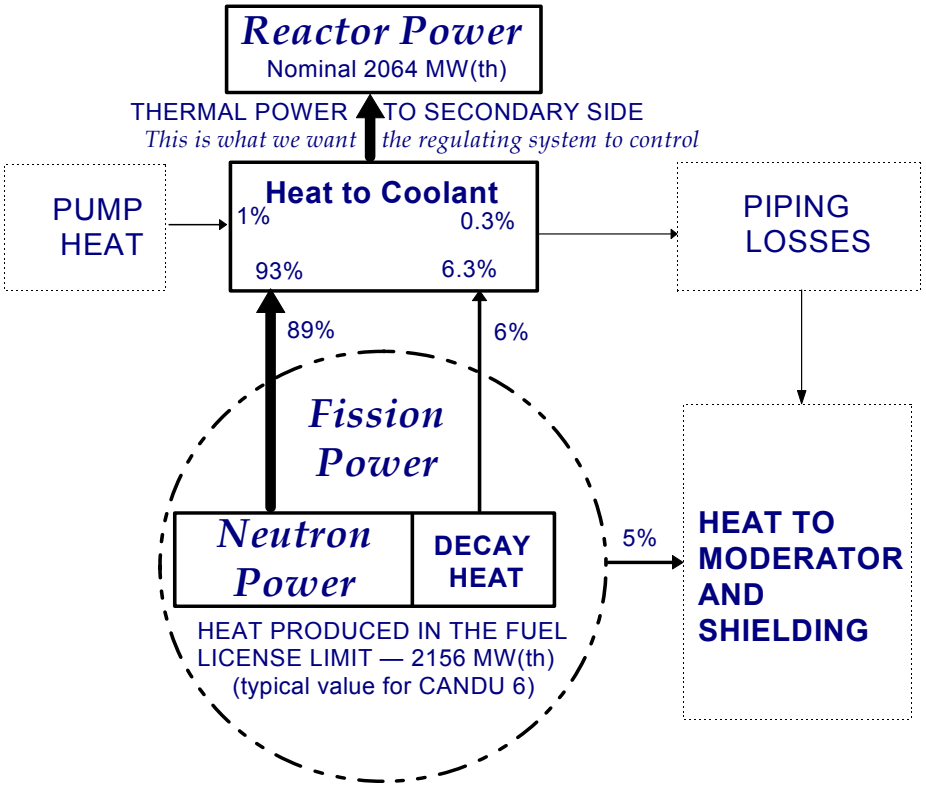

Figure 7.1
Approximate Heat Balance for a CANDU 600 Showing Fission Power, Thermal Power, and Neutron Power

7.2 REACTOR POWER RUNDOWN

The fission reaction is responsible for the power generated in a nuclear reactor, and the fission rate is proportional to the neutron flux, so you might expect that thermal power output is proportional to neutron power. The relationship between flux and thermal power, however, is not linear. For example, if we drop from 100% to 10% neutron power, as indicated by flux-measuring instruments, the thermal power does not fall to 10% of its full power value. Even though the neutron power has decreased by a factor of 10, contributions to thermal power of decay heat, pump heat, and ambient losses do not change, or change slowly. These three sources of non-linearity behave as follows:

The heat produced by the radioactive decay of the fission products in the core (fission-product decay heat) changes slowly following a power change. After sufficiently long reactor operation, 6% to 7% of the steady state heat generated comes from the inventory of fission products. The heat comes from the beta and gamma decay of these products rather than directly from the fission process itself. Thus, if the reactor has operated at 100% power for a long time and is then shut down, even though the fission process stops more or less instantaneously, the thermal output immediately after shutdown is still 7% of its full power value, and will decrease slowly as the fission products decay.

Notes:

Figure 7.2 shows how the reactor thermal power and the neutron power decrease after a shutdown. Note that after a minute or so, the neutron power is contributing only a small fraction of the total power generated in the core.

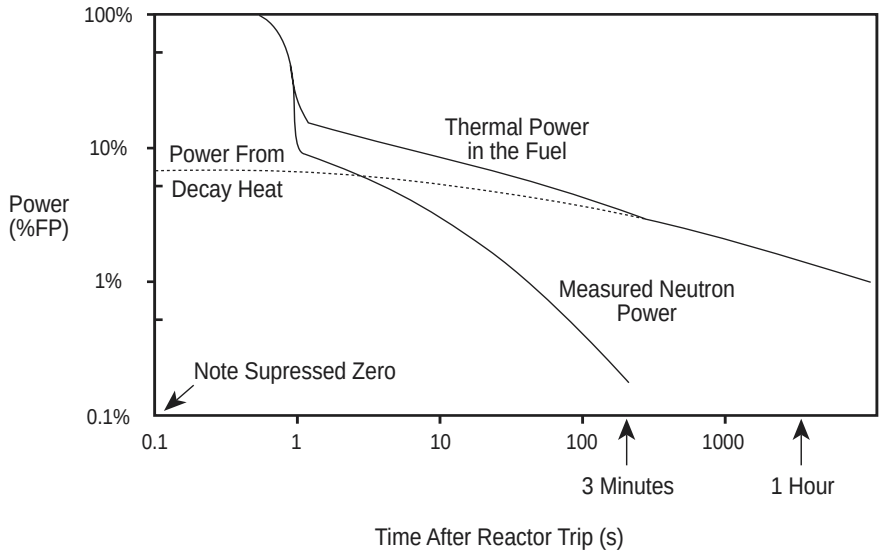

Figure 7.2
Decay of neutron and thermal power after shutdown

The same thing happens, on a less dramatic scale, whenever there is a change from steady operation at one power level to another. When the reactor shifts from 50% to 100% of full power, the ratio of decay heat to neutron power drops to roughly half its previous level, and then builds up slowly as the fission product inventory adjusts to the new power.

Another source of non-linearity is heat generated by fluid friction. About two-thirds of the pressure drop in the heat transport system occurs in the turbulent flow of the coolant channels. This means that about two-thirds of the heat input of the heat transport pumps appears as heat in the coolant channels (about 13 MW(t) at Bruce-A). This input depends exclusively on the coolant flow rate, so it is independent of the reactor power level. This heat, known as pump heat, continues as long as the main pumps are operated.

A third source of non-linearity is the heat lost from the coolant channels, mainly to the moderator (about 4 MW(t) at Bruce-A, for example). The amount of heat lost depends on the temperature difference between the coolant and the moderator, which changes very little over a wide range of reactor power.

Notes:

We will now take a detailed look at the decay of neutron power after a shutdown. The section ends with a quantitative description of the much slower thermal power rundown.

Neutron Power Rundown

To make things concrete, let's assume that an equilibrium-fuelled reactor has been running at full power and is then shut down by a trip which inserts -100 mk of reactivity, giving k = 0.90. Figures 7.3 and 7.4 show the time variation of neutron power following shutdown.

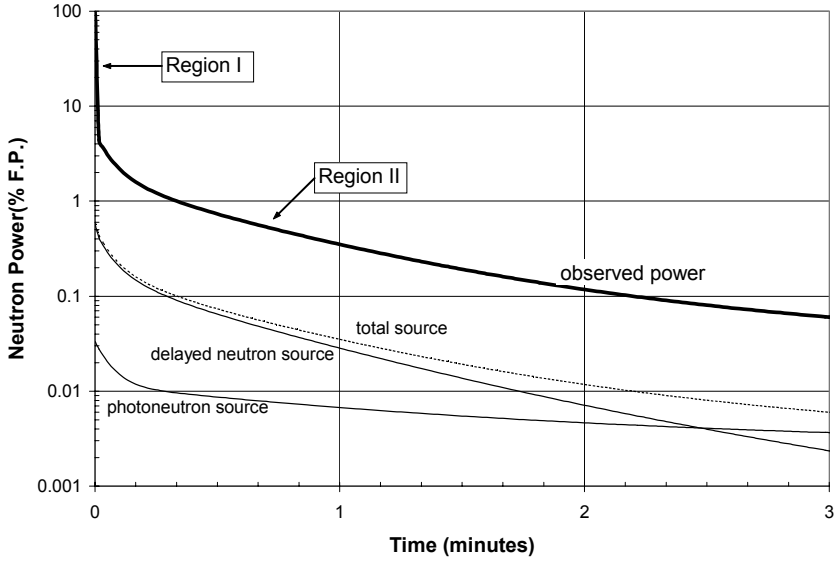

Figure 7.3
Neutron Power after Shutdown (First 3 minutes)

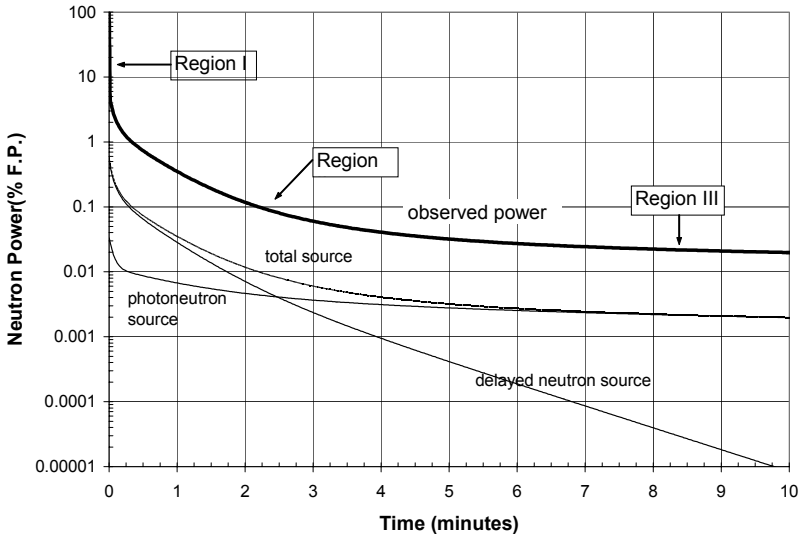

Figure 7.4
Neutron Power after Shutdown (10 minutes)

It is convenient to divide the neutron power rundown curve into three regions and explain them one at a time in the following subsections.

Notes:

Region I – The Prompt Collapse

Notes:

With k = 0.90, the original prompt neutron population will decrease initially by a factor of 0.90 every generation. For example, in 0.5 seconds there is time for 500 prompt neutron lifetimes of 0.001 s each. The neutron power would drop to $(0.90)^{500} \approx 10^{-23}$ of full power, if there were no delayed neutron precursors continuing to produce neutron. This corresponds to the very rapid initial drop shown in Figure 7.3. The collapse of the prompt population in Region I is so fast that, in practical terms, it collapses at the rate at which negative reactivity is inserted into the core.

Region II – Delayed Neutron Hold-up

As we know from Section 3, the neutron population does not continue dropping at this rapid rate, but stabilizes temporarily at a level determined by the subcritical multiplication of the delayed neutrons in the reactor. There are two ways to estimate where the prompt drop will "stop".

In the first, we simply use the prompt jump formula from Section 2. The power immediately after the drop is:

$$P = \frac{\beta}{\beta - \Delta k} P_0$$

which, for an injection of -100 mk into an equilibrium fuelled core ($\beta = 0.005$), gives:

$$P = \frac{0.005}{0.005 + 0.100} P_0 = 0.048\, P_0$$

that is, the insertion of -100 mk of reactivity causes neutron power to fall immediately to about 5% of its original value.

Alternatively, we note that when the reactor is critical at a steady power level, the thermal neutron population originates from prompt and delayed neutrons, with a fraction 0.995 prompt and the remaining 0.005 delayed. The prompt fraction disappears very rapidly on shutdown, before the delayed precursors have had time to decay. The power would drop to 0.005 P_0 (0.5% of full power) without subcritical multiplication. However, the delayed neutron precursors decaying in the now subcritical reactor can be regarded as a source with initial source strength $P_{source} = 0.005\ P_0$.

Because of the subcritical multiplication factor l/(l-k) = 1/0.100 =10, the actual power level immediately after the prompt neutron population disappears is:

$$P = \frac{1}{1-k} \cdot P_{source} = 10 \times (0.005\,P_0) = 0.05\,P_0$$

Notes:

This agrees with the value calculated from the prompt jump formula. Either method is an acceptable approximation.

The "delayed neutron source power" (the power that would be produced by the delayed neutron source alone, in the absence of subcritical multiplication) falls off as shown in Figure 7.4. Initially the drop is rapid as the shorter-lived precursors decay, and then it slows until eventually the decrease is governed by the longest-lived precursor group. The delayed neutron source power is negligible (compared to the photoneutron source) less than ten minutes after shutdown.

Region III – Photoneutron Hold-up

In the critical reactor, the photoneutron source is considerably weaker than the delayed neutron source. However, the half-lives of the fission products whose gamma rays produce the photoneutrons are generally much longer than the half-lives of the fission products that generate the delayed neutrons. After a trip, the photoneutrons become the dominant neutron source in the reactor within two to three minutes (see Figure 7.4).
Table 7.1 lists nine photoneutron groups with the initial strength and half-life of each group. For equilibrium fuel, the photoneutron fraction is about 0.033%.

Group	Fraction of 0.033%	$t_{½}$
9	0.646	2.5 s
8	0.203	41.0 s
7	0.070	2.4 m
6	0.033	7.7 m
5	0.021	27.0 m
4	0.023	1.65 h
3	0.003	4.4 h
2	0.001	53.0 h
1	0.0005	12.8 d

Table 7.1
Photoneutron Data

Figure 7.5 shows the gradual falloff in photoneutron strength after a shutdown. In a week or so, the only significant source strength comes from the group 1 photoneutrons. This photoneutron source has a half-life of 12.8 days and keeps the reactor power measurements on scale for several weeks after shutdown.

Eventually (after a year or more) the photoneutron strength decreases until it is comparable to the spontaneous fission source strength ($\approx 10^{-14}$ F.P.). The photoneutron source strength continues to drop, but the spontaneous fission source is constant, so the power will stop decreasing at a value of about $(1 - k)^{-1} \times 10^{-14}$ F.P.

Notes:

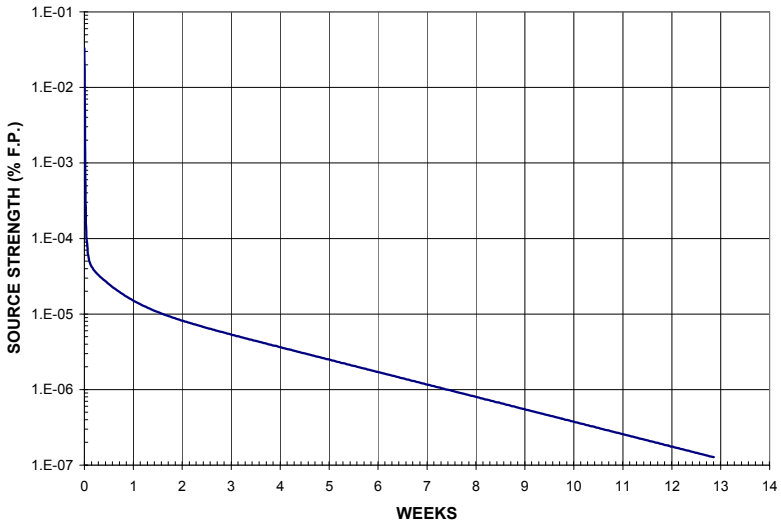

Figure 7.5
Photoneutron Source Strength Following Shutdown

Two points are worth emphasizing about the decay of neutron power:

First, although we have divided the rundown into three distinct regions, in practice the transition between Regions II and III is arbitrary. In Region II, both the delayed neutrons and the photoneutrons contribute to the source, but the balance between the two shifts steadily towards the photoneutrons as time goes on.

Second, throughout the entire rundown, the great majority of neutrons that appear in the reactor are prompt fission neutrons from the fuel. Although, in the subcritical reactor, fission would not take place in the absence of a source of some kind, it is not the source, but the assembly itself that produces most of the neutrons because of its subcritical multiplication. In the case illustrated, where k=0.90, each gamma ray that generates a photoneutron gives rise to an average of 10 neutrons [1/(1-k) = 1/0.1] because of subcritical multiplication.

Thermal Power Rundown

Following shutdown, the thermal power of the reactor decreases much more slowly than the neutron power, mainly because of the decay heat associated with fission products in the fuel. Decay heat produces nearly 7% of the steady state thermal power. Although the fission rate falls off rapidly, decay heat can only fall off at the decay rate of the fission products producing it.

Fission products have half-lives ranging from fractions of a second to thousands of years. The longer-term fall-off in thermal power is therefore a very slow one.

Notes:

Typically, thermal power drops to about 3% full power in about three minutes (which explains, incidentally, why 3% is the capacity chosen for the auxiliary boiler feed-pump), to about half of this (1.5%) in an hour, and to below 1% over a period of eight hours.

As a practical point, we should note another factor that slows the rate of decay of *thermal power* generated in the core. The heat transport pumps generate heat at the rate of nearly 1% of full power. This heat source persists until it is possible to switch to the much smaller shutdown cooling pumps.

7.3 THE SHUTDOWN STATE

During operation, a reactor is always critical or in the deeply subcritical guaranteed shutdown state (GSS), or in active transition between these states. The observed steady-state power in a subcritical reactor is always given by the subcritical multiplication formula:

$$P_{obs} = \left[\frac{1}{1-k}\right] \cdot P_{source}$$

After the reactor is placed in the GSS following extended high power operation, the observed power level changes because both k and P_{source} are changing. Although the poison concentration may not be changed while in the GSS, temperature changes and the decay of fission products both affect core reactivity. Decay of photoneutron precursors gradually changes the neutron source.

Reactivity Changes

A number of inherent reactivity changes in the reactor core follow shutdown from extended high power operation. The effect of the power coefficient is the most immediate. Then, following the xenon transient and decay of xenon, other smaller effects appear. Table 7.2 lists the effects and gives typical CANDU values for each and the time over which the reactivity changes. The range of values in the table is intended to encompass the real variations from reactor to reactor and uncertainties in some of the parameters.

The net effect of these changes is that when the reactor returns to the hot shutdown state ready for restart, the core reactivity is +30 to +32 mk higher than it was before shutdown. Moderator poison keeps the reactor deeply subcritical. On restart, excess poison is removed and the poison concentration at criticality is exactly that needed to offset the increase in reactivity that has occurred during shutdown.

For example, if the core reactivity increases by 31 mk, the poison concentration at critical is equivalent to about 1.1 ppm Gd (1.1 ppm × 28 mk/ppm = 31 mk). The reactivity increase must account for any shutdown fuelling, any change in reactivity device configuration, or any temperature differences as well as the effects in Table 7.2.

Notes:

Off-line calculations to predict these changes in reactivity are unlikely to be more accurate than ± 0.3 mk (1% of 30 mk). In practice, predictions of critical poison concentration are often no better than ± 1 mk, equivalent to ± 15% zone level or ±0.04 ppm Gd. The reason for mentioning this is to point out that it is not possible to know the exact poison concentration at criticality in advance. (The problem of accurate poison concentration measurements compounds the uncertainty.)

	Reactivity Effect (mk)	Time Scale
Temperature Effects		
Full Power Hot to Hot Shutdown	+3	seconds to minutes
Hot Shutdown to Cold Shutdown	-4 to -6	minutes to hours
Fission Product or Precursor Decay		
Xenon Decay	+28	80^{+} hours
Np-239 Decay	+9 to +12	2 weeks
Pm-149 Decay	-8 to -11	2 weeks
Other Saturating FP Decay	-0.5 to -1	5 to 10 days

Table 7.2
Reactivity Changes Following Shutdown

Power Level Changes in a Shutdown Reactor

Once the reactor is cooled, placed in the GSS, and xenon has decayed (after 3 days or so), there are only small reactivity changes in the shutdown core. The observed power then trends downward with the gradual decrease in the source strength.

Given a good estimate of core reactivity, we can determine the source strength from the measured power. Suppose, for example, that the reactor is subcritical by 400 mk (k = 0.6). The subcritical multiplication factor is 2.5 (1/0.4 = 2.5) so the observed power level is two and a half times the source strength. Figure 7.5 shows the neutron source trending downward during the three months following a shutdown. (Source strength estimates beyond about three months have been unreliable.)

7.4 APPROACH TO CRITICAL

Notes:

We will consider two somewhat different approaches to critical from the guaranteed shutdown state (GSS):

- Startup with the reactor regulating system instruments initially off scale low.
- Startup with reactor regulating system instruments on scale.

We do not discuss initial startup or restart after a very long shutdown (~ 1 year or longer) in which spontaneous fission neutrons are the source neutrons. These are "once in a lifetime" events, for which training is given as required.

Reactor power monitoring instruments go off scale low some time after shutdown. The time, typically a few weeks, depends on the particular reactor, its operating history, and the depth of its GSS.

Example:

Suppose your reactor is 250 mk subcritical in the guaranteed shutdown state and your RRS instruments go off-scale low at an observed power level of 4×10^{-6} F.P. The subcritical multiplication factor is 1/0.25 = 4, so the instruments go off scale low when the source power is 10^{-6} F.P. From Figure 7.5, this is about seven weeks after shutdown.

Repeating this calculation with a reactor that is 667 mk subcritical in the GSS, and with RRS instruments that go off-scale at 9×10^{-6} F.P. we find the instruments go off-scale low about three weeks after shutdown.

Any restart from the GSS requires manual operation of the purification system to remove moderator poison. The monitoring instruments differ depending on the length of shutdown. For restarts within a few weeks of shutdown RRS instruments monitor the startup; for restarts after an extended outage, special startup instruments monitor the initial stages of startup, with a transfer to the RRS instruments after they come on scale.

Before looking at these two types of start-up, we will review some principles common to both. First we must discuss what we mean by critical. Then we will look at how the core sensitivity changes in going from the deeply subcritical GSS to the almost critical state.

Criticality

With a neutron source present, there is always a steady state subcritical power level associated with each value of the effective multiplication constant for $k < 1$. For example, if the reactor is 1 mk subcritical, the subcritical multiplication factor is 1000.

You could draw a parallel curve in Figure 7.5 three decades higher than the source strength shown and know that whenever the observed power falls below that curve the reactor is more than 1 mk subcritical (assuming the predicted source strength is correct for the particular reactor).

Notes:

We usually declare the reactor critical when, technically, it is not quite critical. The criterion is whether the regulating system has the capability of taking the reactor slightly supercritical (to manoeuvre power at the demanded rate) without excessively draining the zones. When this condition pertains, the reactor is said to be under direct regulating system control, or operationally critical. Station procedures may differ, but a typical standard is to declare the reactor "critical" (or, under direct regulating system control) when a 10% decrease in zone level (or less) can achieve a subcritical power doubling.

Example:

Recall from Section 3, that power doubles when the subcritical reactivity is halved. If a 10% zone level drop corresponds to adding ~ 0.6 mk, then, by this measure, the reactor is considered critical when $\Delta k = -1.2$ mk ($k = 0.9988$). After the observed power doubling, it is subcritical by 0.6 mk.

Power Increase in the Subcritical Core during Startup

We can estimate the power change that will occur when the reactivity changes from deeply subcritical to almost critical without knowing the source strength. For example, applying the subcritical multiplication formula to the initial and final states, the ratio of observed power for k nearly equal to one and in the GSS is:

$$\frac{P_{obs}("\text{critical}")}{P_{obs}(GSS)} = \frac{1 - k_{GSS}}{1 - k_{"\text{critical}"}}$$

Suppose that the reactor is subcritical by about 600 mk in the GSS. (This takes into account the mk worth of poison added following shutdown, and the underlying changes in inherent core reactivity discussed in Section 7.3.) Assume that the reactor is considered critical when $\Delta k = -0.6$ mk. Then:

$$\frac{P_{obs}("\text{critical}")}{P_{obs}(GSS)} = \frac{0.600}{0.0006} = 1000$$

For this reactor, power will rise by three orders of magnitude during the startup.

Notice that the power level at which the reactor goes "critical" is likely to differ from one startup to the next. It depends on the power level observed prior to startup, which depends on the source strength. Using Figure 7.5 we can see that after a short shutdown (less than one week) the reactor reaches criticality above 10^{-4} F.P. After a long shutdown of, say, twelve or thirteen weeks the power at critical will be a little above 10^{-6} F.P.

The responsiveness of a subcritical core depends on how close the reactor is to critical, not on the power level. A reactivity addition has a much larger effect if the reactor is nearly critical than it does in a deeply subcritical core (see Section 3). The ratio of observed power (count rate for start-up instruments) before and after the addition is:

Notes:

$$\frac{P_{obs}(\text{final})}{P_{obs}(\text{initial})} = \frac{1 - k_i}{1 - k_f}$$

Suppose the reactor is 500 mk subcritical ($1-k_i$) = 0.5 and we add +100 mk of reactivity, so ($1-k_f$) = 0.4. The power would increase by 25% (0.5/0.4 = 1.25). In contrast, adding only +1 mk to a reactor that is initially 2 mk subcritical doubles the power. [In this case, ($1 - k_i$) = 0.002 mk and ($1 - k_f$) = 0.001 mk.]

Clearly, the sensitivity of the core to reactivity insertions changes dramatically as the reactor gets closer to critical. The dynamic response also changes dramatically.

A step increase in reactivity when the reactor is deeply subcritical produces an immediate change in subcritical power. The observed power change keeps up with the reactivity changes during continuous poison removal. When the poison removal stops, the power increase stops.
In contrast, with the reactor almost critical, it may take several minutes for power to reach the new equilibrium power level following each step removal of reactivity. During continuous poison removal, the power rise lags behind the reactivity change. Even if the poison removal stops, power continues to rise for some time. Table 7.3 indicates the stabilization time for different core reactivities.

Core Reactivity (mk)	Stabilization Time [3] ($3 \times |\tau|$)
- 100	35 s
- 10	50 s
- 4	1 ¼ m
- 2	2 m
- 1	3 ½ m
- 0.5	6 m
- 0.25	12 m
- 0.1	30 m

Table 7.3
Stabilization Following a Step Reactivity Insertion (Subcritical Core)

[3] Based on the decay constant for the second longest lived delayed neutron group

Initial Approach to Criticality after an Extended Outage

Notes:

During long shutdowns, where the power drops to a low level, the normal instrument readings are unreliable because background gamma rays contribute to the readings. The installed instruments are considered "off-scale" somewhere between about 10^{-6} and 10^{-7} of full power. The readings below this level are not proportional to the flux. Supplementary "startup instruments" (BF_3 counters—He-3 counters at some stations) are installed before this happens, and the initial stages in the subsequent approach to critical would use these counters.

The initial approach to criticality following removal of the GSS uses the purification system to increase reactivity and the startup instruments to monitor count rate. As there is no automatic regulation, the operating staff must perform the functions of the automatic system. This means observing power level, making sure that changes in power are within the expected range, and adjusting purification flow appropriately. Monitoring usually includes plotting a graph of 1/CR vs. reactivity (or vs. poison concentration), as described near the end of Section 3.

Some stations use a continuous poison pull and plot 1/CR vs. time. The ion-exchange columns have a half time for removal of gadolinium. That is, they will remove half the poison in the moderator in a certain number of hours, half again in the same amount of time, etc. In other words, poison removal slows as poison concentration decreases, so a graph of 1/CR vs. time is not linear.

For either the linear or the non-linear graph, it is important to know the expected trend and to monitor routinely to verify that the actual trend matches the expected one. At some point in this process, the normal instruments come on scale. The operator then uses the RRS ion chambers, the liquid zone control system, and the purification system to take the reactor critical.

Approach to Critical Using RRS

With the normal RRS ion chambers in range, the regulating system responds, but the reactor is not close enough to critical to say the regulating system is in control. The effect of valving in purification with the regulating system holding power is to raise the zone levels with the net reactivity of the core unchanged. With purification valved out, reactivity increases when a request to raise power decreases the zone levels. Procedures differ in detail from station to station, but usually these two processes alternate. This avoids the complication of having simultaneous device operations changing core reactivity.

The following example illustrates the principles common to all procedures. First, we describe the start up; then we explain why it works and what it achieves. We will start our description of the process with the reactor significantly subcritical. Poison removal is stopped. The regulating system is operating and the operator enters power requests directly into RRS (RRS is in Reactor Leading Mode).

The operator notes the power level and request a power double the present level, (that is, 0.3 decades higher). RRS reduces zone levels to increase reactivity. The operator monitors, and when the zones get down to about the 20% level, he stops them from draining by asking the regulating system to hold power (at the present level).

Notes:

Now poison removal is started again *with* RRS holding power. As poison concentration decreases, RRS raises zone levels to compensate, keeping the reactivity of the core constant. The operator monitors, and stops poison removal when the zones are high.

These two steps, RRS response to a request to double the present power with purification stopped, followed by poison removal with RRS holding power, are repeated until RRS achieves a power doubling when the operator requests it. The process is then iterated two (or three) more times, achieving three (or four) definite power doublings in all.

Why does the operator ask for power to double? Why does the operator look for *three* doublings of power? Is the reactor critical?

Let us look at this startup process in more detail, starting with the iteration just before the power doubled the first time. For illustration, we will assume the reactivity worth of the zone control system over the usual operating range (from 20% to 80% full) is about 4 mk. (This corresponds to 6.67 mk from 0% to 100% full).

Suppose the operator drives the zone levels up to a little over 65% by poison removal with RRS holding power and then stops poison removal. Following a request to double the observed power, RRS lowers the zones to eliminate the power error. Power increases but as the zones approach 20%, it becomes apparent that the power is not going to double, so the operator stops the process.

What can we say so far?

First, in dropping from about 65% to 20% the zones insert about 3 mk. (0.45×6.67 mk = 3 mk). Power did not double, so the reactor was initially subcritical by more than 6 mk.

Now the operator removes poison and drives the zones back up to 60%. There is no change in reactivity. Suppose that this time when the doubling request is made, the power doubles when the zones drop to near 20%. The zones have added almost 40% of their full worth, about +2½ mk. Since power doubled, we conclude that the reactor was approximately 5 mk subcritical and is now about 2½ mk subcritical.

Now the operator pulls poison with RRS holding power and drives the zones up to about 55%, stops poison removal and requests a second power doubling.

We know that a power doubling is within the range of the zones and, if the estimate of 2½ mk subcritical is accurate, the next doubling will occur with insertion of +1.25 mk. We expect the zones to drop from about 55% to near 35% (20% of 6.67 mk = 1.3 mk).

Notes:

Again the operator holds power, removes poison till the zones rise to almost 50% and asks for the third power doubling. This time the zones drop a little more than 9% to add a little more than 0.6 mk and double power once again. Average zone level is just over 40% and the reactor is subcritical by about 0.6 mk.

We can see some of the features of this procedure.

- Three doublings allow the operator to fine tune the final zone level, but that is not the main point of requiring three doublings.
- Three doublings on zone level always brings the reactor close to critical.

In our example, the reactor is just close enough to critical to satisfy the requirement for direct regulating system control. If the first doubling in this example had occurred with a zone level drop of, say, 44% (e.g. from 65% to 21%), the second doubling would have required a drop of 22%, and the third a drop of 11%. That would not quite satisfy our criterion for saying RRS is in control (although a 4th doubling would certainly be adequate).

One of the nicest features of this procedure is that the reactor cannot actually go critical! It is a procedure for approaching criticality, cautiously, while ensuring that the reactor necessarily remains subcritical. As long as the reactor is subcritical and the operator requests a power doubling, RRS adds only half of the reactivity required for criticality.

The key to this conservative approach is that when the reactor is nearly critical, power is increased by adjusting power level, not by direct manipulation of reactivity devices. The operator cannot initially know exactly how subcritical the reactor is. Making a positive reactivity insertion into a subcritical reactor by removing a predetermined amount of poison, or by operating purification for a calculated time, could cause the reactor to go critical before expected. This challenges RRS to kick in automatically and take over. If RRS fails, only the Safety Shutdown Systems can prevent a power excursion.

On the other hand, if you monitor reactor power and look for a power doubling the reactor does not actually go critical. The increase can be achieved in a subcritical reactor by adjusting k closer to k = 1.

Restart When the Reactor is Not in the GSS

The reactor is not in the guaranteed shutdown state before restart in two situations. These are poison override (restart within ~ 30 minutes or so of a trip) and recovery immediately following a poison out (typically 35 to 40 hours after a trip).

In each case, the restart uses the regulating system, but the high, rapidly changing xenon concentration adds complications (see Section 5.3.4).

Poison Override

Notes:

Immediate recovery following a trip is possible, if at all, only after an SDS1 trip. Restart is allowed only after re-poising the safety shutdown system, and for SDS2 this may take nearly as long as the poison-out time.

Assuming restart is an option following an SDS1 trip, the shutoff rods must first be withdrawn to re-poise SDS1. The reactor stays subcritical because of the buildup of xenon and because the zones fill and control absorbers drop into the core on a trip. The control absorbers must also be withdrawn (in banks), followed by removal of adjusters, one bank at a time, until there is enough reactivity to overcome the xenon buildup. Suppose criticality is reached (by a decrease in zone level) following removal of the last bank of adjusters. A request to increase power begins the xenon burnout process. Zone levels rise again as xenon load decreases, and the regulating system will request adjuster in-drive (one bank) each time the average zone level reaches 80%.

Recovery from a Poison Out

In restarting after a poison out, the decay of xenon, which is not under operator control, determines the rate of reactivity addition. The operating staff must be ready for criticality, and careful monitoring is required, as predictions of the poison out time may not be accurate.

During a poison outage, as xenon decays in the reactor, the ion chamber signals increase until the reactor power reaches its setpoint. At this point, RRS responds to reactor power measurements. As more xenon decay occurs, the liquid zones start to fill to maintain the reactivity balance. Once the liquid zones reach their control limit and the xenon decay continues, reactor power would increase in the absence of control action. Poison addition and/or adjuster in-drive (if the adjusters are out of core) are required to maintain the liquid zones in control range.

Once the reactor returns to power, the burnout of xenon is rapid and xenon simulation by reactor poison is required (see Section 5). This is straightforward following restart after an SDS2 trip, as the core is heavily poisoned. Poison removal begins after SDS2 is re-poised, but there will we excess poison in the core at criticality. Poison burnout approximately matches xenon buildup so purification can be stopped when power is raised.

One interesting complication is an SDS1 trip immediately following recovery from a poison out. In this case, there is no xenon buildup, and care is required to make sure the reactor does not go critical on rod withdrawal.

7.5 LOW POWER OPERATION FOLLOWING STARTUP

Notes:

Usually following criticality, reactor power is ramped into the high power range to allow electrical power generation. Sometimes, however, extended operation at low power critical may be necessary. We will discuss the characteristics of this operating state before going on to high power operation.

As discussed earlier, the critical reactor at a constant low power is necessarily slightly subcritical because of the photoneutron source. The startup procedure will have verified that the reactor is within something like 10% zone level of critical, ready to be made supercritical and ramped to high power as required. The operator monitors power level, which the regulating system should hold constant, and zone level, which should change very slowly with changes in the photoneutron source strength.

The explanation of gradual zone movements requires us to look more carefully at the photoneutron source.

On return to power, the photoneutrons build back to equilibrium, with the short-lived ones building up first, and the others following along. Following restart, the fission products that produce energetic gamma rays that generate photoneutrons are a mixture of isotopes, some of which are building up from current fissions, while others remain from previous high power operation.

For example, suppose the reactor went critical at 10^{-5} of full power, and power was then ramped to 10^{-3} and held there after a long outage. The source remaining from previous high power operation is small, having decayed during the outage. Short-lived fission products emitting energetic gamma rays that produce photoneutrons are building up, not to the 100% full power level, but perhaps to a level as large as the source remaining from previous operation. As the source strength increases, the regulating system will have to reduce k to hold power constant. The reactor could drift subcritical.

In contrast, the short-lived photoneutron groups do not build to as high a level if the reactor power stays at 10^{-5} of full power. The source will be mainly what remains from previous high power operation. As this source gradually decreases (with a 12.8 day half life) the regulating system will gradually reduce zone levels to keep the reactor at power. With the observed power held constant and the source decreasing, the subcritical multiplication formula shows that k is drifting closer to k = 1.

If the power is now ramped to, lets say, 10^{-3} of full power and held there for a few hours, the short lived photoneutron groups will build up to a level comparable to the left-over source. Now on return to 10^{-5} of full power, the source is much larger, and the subcritical multiplication formula shows that the regulating system must decrease k below its previous value to hold power.

It is quite possible that a reactor that was made critical could, after power manoeuvres (during which it was clearly supercritical), be subcritical on return to its initial power level. The liquid zone control system might not have enough available reactivity to raise power.

Notes:

A Test for Criticality

Operation with the reactor significantly subcritical is not allowed, except for transitions to and from the guaranteed shutdown state. The intent is to avoid long periods of operation in which the regulating system may not be able to effectively control reactor power. (Long periods of time in the GSS are not a problem as there are no equipment failures or failure combinations that could make the core critical.) Startup, the transition between GSS and critical, requires the dedicated attention of control room staff, following approved procedures.

To help avoid situations in which the reactor drifts subcritical, there is a test to decide whether a reactor at low power is close enough to the condition k=1 to be considered "critical", (that is, can the regulating system control power?). The test repeats the final steps of taking the reactor critical. The operator requests a power doubling and if this occurs with only a small change in liquid zone level (for example, less than a 10% decrease, or a Δk of about 0.6 mk), the reactor remains critical. It may be necessary to check this at regular intervals if the reactor is being operated at a very low power level for a long time.

Raising Power

Section 2 describes and explains power increases following small reactivity additions at low power where there are no feedback effects. Power initially rises quickly over less than ½ second (the prompt jump), but only by 1% or so for a small step reactivity insertion. This is followed by a stable rise with a constant rate log of power, that is, a constant fractional power increase (percentage of present power per second = $1/\tau$).

To raise power, the operator requests a power increase by entering a power setpoint in the control computer and selecting an appropriate rate of increase. The formulas in Section 2 model the core response quite well. A reactivity insertion of 0.2 mk or so, achieved by a zone level drop of 3% to 5%, produces a stable rise at a rate near 0.5% P.P./s ($\tau \approx 200$ s). This rate, a faster rate than manoeuvres in the high power range use, will bring power from 10^{-4} of full power to between 1% and 10% F.P. in about ten minutes.

As power rises into the range of 1% to 10%, xenon begins to build up and fuel heating begins to be significant. Xenon buildup is slow; its effect is significant some hours later. Before power reaches 10% of full power, the temperature increases of fuel and coolant will have removed a couple of tenths of a mk of reactivity. Without the active response of the regulating system the power rise would stop.

To continue ramping power, the regulating system will have to continue decreasing zone level to offset the temperature effects. To maintain a rate near the original 0.5% P.P./s, net reactivity (the difference between reactivity inserted by the liquid zones and reactivity removed by increasing temperature) must remain constant at 0.2 mk or so.

Notes:

In the power range from below 10% F.P. to 100% F.P., active changes of reactivity devices are required to change power, or to hold power after a change (as xenon concentration changes). These reactivity insertions are to offset the reactivity feedback effects. The regulating system never needs more than about 0.1 mk net reactivity insertion to raise power at an acceptable rate in the high power region.

The power coefficient results mainly from the fuel temperature coefficient. Its effect is approximately linear, typically about -0.3 mk for each 10 % of full power increase in power.

The equilibrium xenon load builds slowly to a relatively low equilibrium concentration below a few percent of full power, increasing strongly in the range of 10% to 60% full power. Above this range, equilibrium xenon is always within about a mk of its full power equilibrium value (see Section 5).

REACTOR OPERATIONS AT HIGH POWER

8.0 INTRODUCTION

Notes:

For operations producing the rated power output, it is important to limit the power from individual channels or bundles to safe levels. We examine four different ways of flattening the overall flux shape: addition of a reflector, use of adjuster rods, bi-directional refuelling, and differential burnup.

We also describe the role of the liquid zone control system in maintaining a zone-to-zone flux balance, thus limiting high zone flux that could result from flux tilts and oscillations. In addition, we review some practical aspects of fuel management that limit local peaks ("hot spots"). Reactivity device configurations also affect local peaks, so we will look at the reactivity effects of rod drives.

The Regional Overpower System (ROP) protects the reactor from high bulk power and high local power. We will introduce the principles that underlie this system.

One of our main concerns in operating a power reactor is to ensure that fuel bundles are not over-rated, leading to fuel damage. Since the fission rate, and hence power generation, in a bundle is proportional to the thermal neutron flux at the bundle position, we should have some idea of how the flux varies from one part of the reactor to another.

Solving the coupled equations described in Section 1 for a finite**,** homogeneous reactor (one with a uniform mixture of fuel and moderator) leads to smooth flux shapes gently varying from a central peak to zero at the boundary. The CANDU reactor is a heterogeneous system, where fuel is lumped into bundles to reduce the U-238 resonance absorption. Because thermal neutrons are strongly absorbed in concentrated fuel regions, the thermal flux is markedly depressed in the fuel itself relative to the moderator.

The overall flux shape in the CANDU core is generally high in the middle and falls off to the edges, but techniques to be described here extend the high flux over a much larger region of the core, while suppressing the central peak. There are local "peaks and valleys" superimposed on this overall shape.

8.1 FLUX FLATTENING

It is possible to calculate the ratio of the average thermal flux in the core to the maximum thermal flux for the homogeneous core, using the calculated flux variation in a cylindrical reactor. The ratio turns out to be

$$\frac{\phi_{AV}}{\phi_{MAX}} = 0.275$$

This low value for the average to maximum flux raises problems. The total power output of the reactor is proportional to the average flux, so it is advantageous that this be as high as possible. However, a maximum heat rating (to avoid risking damage to the fuel) sets an upper limit on the flux. This limit applies to fuel at the core centre of a homogeneous reactor, where the flux has its maximum value. With ϕ_{MAX} set by safety considerations and ϕ_{AV} equal to only 27.5% of ϕ_{MAX}, the fuel away from the central peak contributes far less than its potential share of the power.

Notes:

The solution is to increase the ratio of ϕ_{AV} to ϕ_{MAX}. Flux flattening is the name given to the improvement of this ratio. In the remainder of this section, we will discuss four ways of flattening the thermal neutron flux in the radial and axial directions. These are:

- addition of a reflector (for radial flattening);
- adjuster rods (for axial and radial flattening);
- bi-directional refuelling (for axial flattening);
- differential burnup (for radial flattening).

The reactor design incorporates the first two of these; fuelling achieves the last two. The net result is an increase of the average flux from 27.5% to about 55% or so of the maximum, approximately doubling the heat output of the reactor without increasing any channel or bundle power above the peak allowed value.

The Reflector

Each CANDU reactor has a radial reflector that consists of a radial extension of the moderator that provides a layer of about 70 cm of heavy water. This layer of heavy water surrounds the core and reflects many leakage neutrons back into the core. This produces two effects: it reduces leakage, and increases the flux near the core boundary (flux flattening).

Figure 8.1 illustrates the effect of the reflector in flattening the thermal flux in the radial direction. The theoretical flux shape for a bare core is labelled “flux without reflector”. The lower flux shape with the reflector in place is normalized to the same total power output as the bare core flux. The higher flux shape with the reflector in place shows the power raised so that the maximum flux in the core is the same as it was for the bare core.

Fast neutrons that are thermalized in the reflector cause the humps there. They "pile up" in the reflector because they are less likely to be absorbed there than in the core.

Notes:

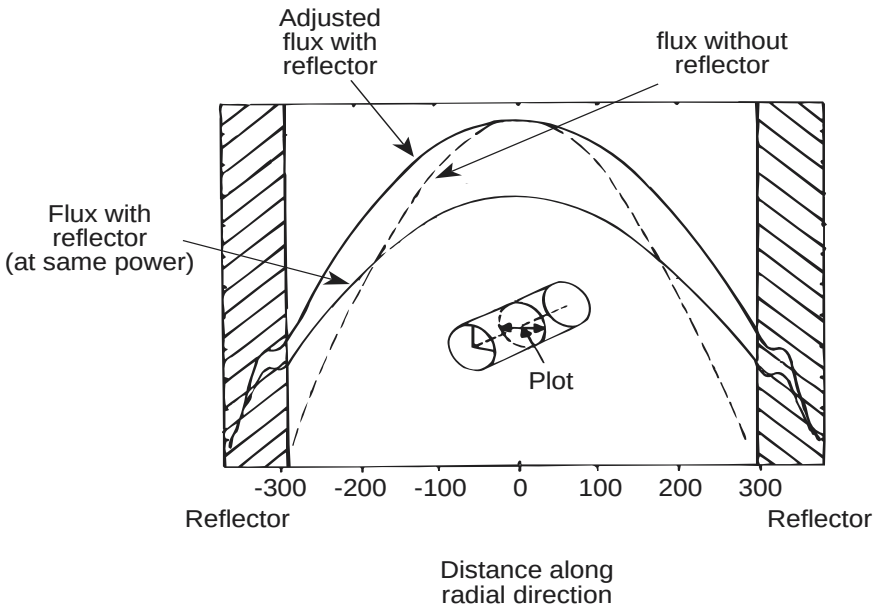

Figure 8.1
Effect on Reflector on Shape of Radial Flux

Adjuster Rods

Adjusters are rods of a neutron-absorbing material normally inserted into the central region of the reactor. They have the dual function of flattening flux (both axially and radially) and of providing excess reactivity when called upon. The name "adjusters" is comes from the flux flattening function—adjusting the flux shape.

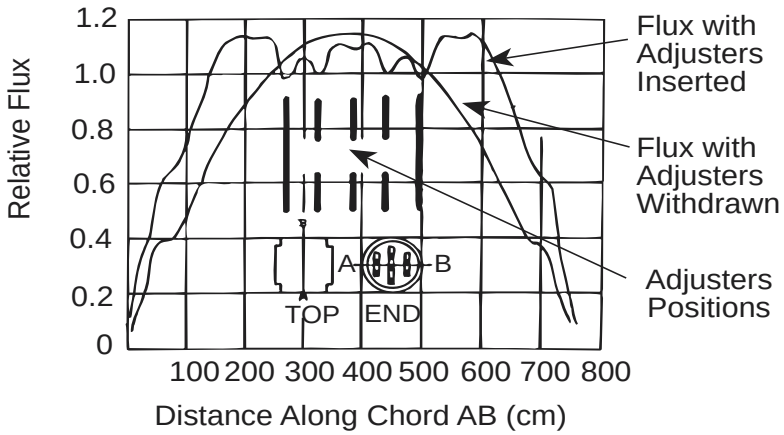

Figure 8.2
Flux Flattening Produced by Adjuster Rods

Figure 8.2 illustrates flux flattening in the radial direction by the adjuster rods. The figure also shows the flux shape with adjusters withdrawn. The curves for the flux with and without adjusters show the same peak flux, since this is what imposes the limit to avoid fuel damage.

As the regulating system withdraws the adjusters from the core, it simultaneously increases liquid zone levels to keep the power output constant. This would preserve the areas under the two curves. To keep the peak flux from increasing, the operator would need to reduce the reactor power setpoint *before* the rods come out of core. This reduces the height of the curve labelled "flux with adjusters inserted". Adjuster withdrawal then produces the more peaked shape shown, with the peak height matching the original peak height. It is clear that a reactor with adjusters inserted produces higher power for the same maximum flux.

Bi-Directional Refuelling

Fuelling a channel involves removing high burnup bundles from a fuel channel (usually four or eight bundles) and replacing them with fresh fuel. Partly burnt-up fuel then occupies one end of the channel, with fresh fuel in the other end. The fission rate in the new fuel is higher so it produces higher thermal flux in that end of the core. If all the refuelling were done from the same end of the reactor, this would eventually create a marked asymmetry in the flux with a high peak near the end with fresh fuel. To avoid this, we refuel adjacent channels in opposite directions, thereby keeping the flux shape essentially symmetrical.

Figure 8.3 shows that the addition of the asymmetric flux shapes produced by channels refuelled in opposite directions produces an axial flux shape that is somewhat flattened. The curve labelled "theoretical shape" in this figure assumes replacement of all fuel bundles during a refuelling operation.

Notes:

This does produce a symmetrical shape, but to get the same total power output, the peak flux would need to be much higher, as shown.

Notes:

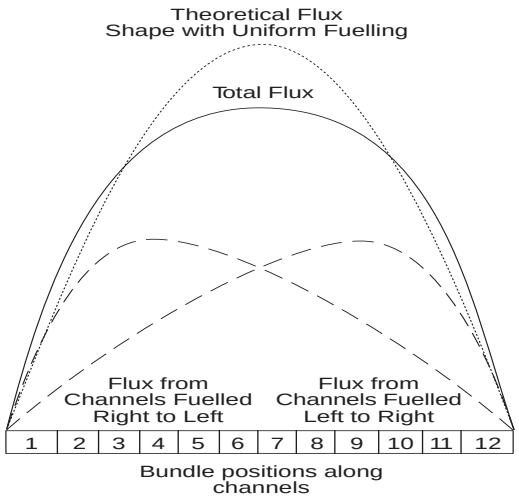

Figure 8.3
Effect of Bi-Directional Refuelling in Flattening Axial Flux Shape

The flattening effect is greater with a four (or two)-out-of-twelve refuelling scheme, but even the eight (or ten)-out-of-twelve replacement produces some flux flattening.

Differential Burnup

Differential burnup is a method of radial flux flattening produced by fuelling strategy. The fuelling engineer allows fuel in the inner zone of Figure 8.4 to reach higher burnup than fuel in the outer zone. To keep the reactor critical, fresher fuel in the outer core supports the high burnup fuel in the inner core. With less fissile material and more fission product absorption in the inner zone, the fission rate distribution is less peaked, so the flux shape is flatter, as shown in Figure 8.4.

Notice that the actual dwell time of fuel in the inner zone is less than in the outer zone. Although the burnup is higher, high burnup occurs more quickly in the higher flux region.

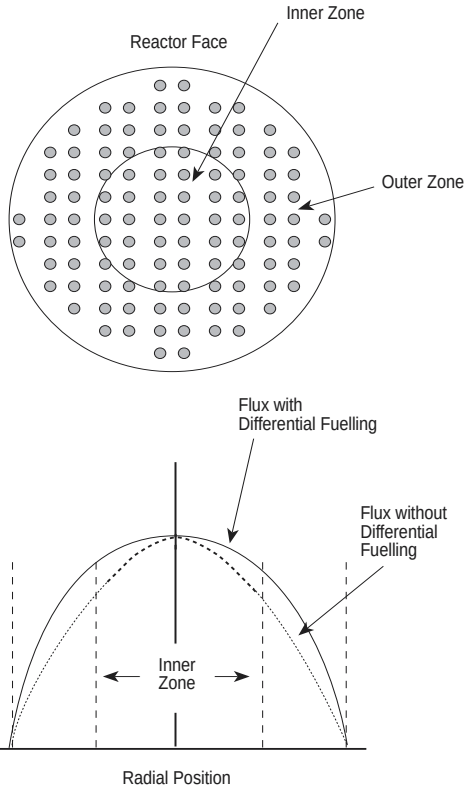

Figure 8 4
Flux Flattening Produced by Differential Fuelling

Notes:

The BNGSA reactors do not have adjusters, which results in some loss of operating flexibility. However, by eliminating the burnup loss caused by adjusters these reactors achieve high fuel burnup. Differential fuelling flattens flux in the radial direction and two bundle shifts on refuelling in the inner zone helps flatten flux axially. The higher parasitic absorption in the inner core has much the same effect as adjuster rods in the inner core.

8.2 FLUX SHAPE DETAILS

The actual flux shape in the reactor core differs somewhat from the overall smooth flux shape just described. Neutron absorption in fuel and in reactivity devices distorts the smooth flux shape. For an equilibrium fuelled CANDU, the overall flux shape is nearly constant during operation, but there are local “bumps” and “hollows” that change strength and location because of changes in xenon concentration, because of fuel burnup and on-power fuelling, and because of changes in the configuration of the reactivity devices.

Local flux peaks must be controlled so that safe operating limits on fuel bundle power and channel power are not exceeded. Peaks can affect fuel power directly or by distorting measurements of power. In routine operation, a good fuelling strategy coupled with effective power regulation limits the peaks.

Anything that could produce an unanticipated local peak, such as a stuck rod, or an uncontrolled xenon oscillation requires attention.

Notes:

Flux Depression around Absorbers

Figure 8.5 shows the thermal neutron flux depression across a CANDU fuel bundle. Notice that the reduced flux extends beyond the boundary of the bundle. The thermal flux just outside the bundle results from the random motions of neutrons moving into the region from all directions.

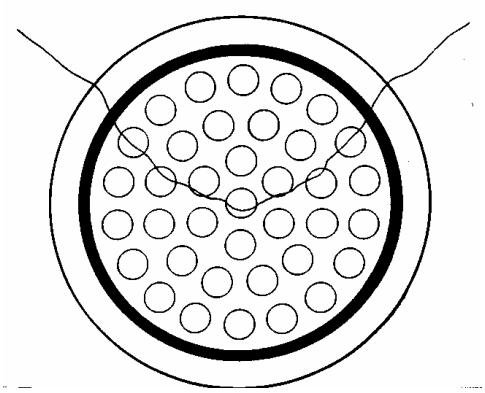

Figure 8.5
Depression of the Thermal Flux in a Fuel Bundle

The presence of a nearby absorber reduces the source of neutrons from that particular direction, so the flux is reduced. This effect is seen for any absorbing material in the core. Not only is the flux reduced in the region where neutrons are absorbed, it is reduced in the surrounding region as well. The flux depression gets less the further you are from the absorber.

Zone to Zone Flux Flattening

The calibrated zone power signals represent the average power output from each zone. The regulating system raises or lowers individual zone levels to keep the average flux in each of the fourteen zones close to the overall average.

The regulating system phases in spatial correction between 15% and 25% zone power, so full spatial control is available when power reaches the threshold for onset of xenon oscillations, about 25% full power. Xenon disturbances below this power do not cause self-sustaining oscillations.

We can estimate the flux level below which xenon disturbances are not a problem. One of the requirements for an oscillation (or transient xenon effect) is a xenon burnout rate significantly greater than the xenon decay rate.

The flux that gives equal xenon burnout and decay rates is

$$\sigma_a^{Xe}\phi N_{Xe} = \lambda_{Xe} N_{Xe} \quad \text{so}$$

$$\phi = \frac{\lambda_{Xe}}{\sigma_a^{Xe}} = \frac{2.12\times10^{-5}\,s^{-1}}{3.5\times10^{-18}\,cm^2} = 6\times10^{12}\,cm^{-2}s^{-1}$$

Notes:

CANDU average fuel flux at full power ranges near 7 to 9 × 10^{13} cm^{-2} s^{-1}, so this is just under 10% of full power.

A local, high power xenon effect not accounted for in Section 5 is the effect of fuelling a high burnup channel. The fresh fuel (with high fissile content and no fission products) causes an immediate increase in flux locally. The high local flux increases burnout in neighbouring channels. Over the next several hours, while xenon is building into the fresh fuel, the high flux continues to burn out xenon in nearby channels. This initially produces a local hot spot. The liquid zones, in responding to an average flux increase in the zone, have some effect in limiting the size of the peak. The hot spot is self-limiting over a number of hours. The buildup of iodine in the neighbouring channels results in increased xenon production and neutron absorption, and so does the buildup of xenon in the fresh fuel.

Usually these fuelling transients cause no particular problem, as long as the problem is not compounded by inappropriate additional fuelling or device movements. At high power, such problems would be detected promptly, and technical advice solicited to correct the problem before additional fuelling takes place, making things worse.

Reactor fuelling should be done with the reactor at high power, so that corrective action can be taken if there are unforeseen reactivity effects such as a hot channel, or flux tilt. These abnormalities cannot be detected with the reactor shut down. A latent flux distortion due to shutdown fuelling activity could go undetected.

Shutdown fuelling can be done on a case by case basis with management approval, but when the reactor is returned to power careful monitoring for flux distortions is required as power is raised.

Local Power Peaks

The liquid zone control system, designed and operated to regulate bulk power, and to adjust zone *average* power, cannot effectively eliminate a *local* "hot spot" that might develop during operation. Such detailed flux shaping is accomplished routinely by selecting appropriate channels for on-power fuelling. The next subsection discusses several considerations taken into account in selecting channels for fuelling, including limiting local power peaks.

To describe the local peaks and valleys in the flux distribution we compare the actual flux distribution with an idealized reference flux shape. The reference flux shape derives from the shape that would exist if day-to-day variations could be averaged out. (This is variously called the time-averaged flux shape or the continuous fuelling model flux shape.) There is no possibility of achieving this theoretical shape, as burnup and fuelling in normal operation produce local fluctuations about the average. The fuelling engineer consults a channel power map based on the reference shape as an ideal "target" for fuelling.

Notes:

Deviations in flux from the reference shape are called fuelling ripple. Fuelling ripple, defined for each channel, is the ratio of the actual (measured) channel power to the reference channel power. (The off-line fuelling code generates a map of reactor fuel channels with ripple values for each.)

For example, if a particular channel has a reference power of 6.5 MW (meaning that over many years of routine, full power operation the analysts expect power from this channel will average 6.5 MW), and the actual power is 6.2 MW before refuelling, the ripple is 6.2/6.5 = 0.95 for that channel. If the channel power reaches a steady 6.9 MW after refuelling, the ripple is then 6.9/6.5 = 1.06. Fuelling ripple changes continually with core conditions.

Given the fuelling ripple for every channel in the core, there will be one highest value of the ripple. This value is the channel power peaking factor (CPPF). (Stations may discount high ripple in the outer one or two rows of channels at the edge of the core when selecting the CPPF channel. Stations may increase the "measured" CPPF to include uncertainties in measurement and analysis.)

All channels and bundles in the core operate below their safe operating limit:

- if bulk power is controlled,
- if the overall flux shape is kept adequately flat, and
- if fuelling keeps local peaks acceptably low.

The regulating system accomplishes the first two of these, provided the operating staff ensures that zone levels operate within their normal operating range, and there are no flux distortions from unusual reactivity device configurations. The safety shutdown systems back up normal regulation. A later section discusses the Regional Overpower Protection System (ROP) - sometimes called the Neutron Overpower Protection System (NOP).

The third item in the list depends on selecting and fuelling the right channels. (Fuelling also affects the second item on the list). The next subsection describes how the fuelling engineer selects channels for fuelling.

Effects of Fuelling

Notes:

On-power fuelling is used for long-term reactivity management and for flux shaping in CANDU reactors. The Fuel Engineer at the station is responsible for ensuring that the optimum fuel cycle is used. A computer program that simulates the core helps select channels for fuelling. Among its many calculations, the simulation determines the axial and radial power distributions, the burnup of each bundle, and the excess reactivity.

A comparison of power output predictions with accurate thermal power measurements (Heat Balance) checks the validity of these calculations. If there are discrepancies, the program is "tuned" by appropriate adjustments to the physics data to obtain satisfactory agreement between the computer code and the measured quantities.

The fuelling engineer reviews the output of the fuelling program and then selects channels for fuelling based on

- safe operation,
- reactivity requirements,
- flux shape control,
- fuel cost, and efficient use of the fuelling equipment.

Regulation and protection of the core in day-to-day normal operation depend crucially on the appropriate choice of channels for fuelling.

On-power fuelling principles and their explanations follow:

- Do not operate with defected fuel elements in the core.

Fuel defects release radioactive fission products to the coolant. This increases chronic dose to station staff. Radioactive releases increase on power ramps. A power change with just a few defective elements in core could result in I-131 in the coolant exceeding allowed limits—limits based on possible accidental release to the public. Defective fuel elements in core also increase the radioactive fission product release from the fuel following a major accident. In addition to these safety issues, contamination of the fuel defect detection and location systems makes them less sensitive to subsequent fuel failures.

- Regularly replace enough fuel to keep the average zone level near midrange.

In a CANDU operated at high power without fuelling, the liquid zones will typically drop about 5% each day to offset the loss of reactivity from burnup. High or low average zone level can make bulk power regulation difficult.

There may not be enough available range to compensate for the xenon and temperature reactivity effects that occur during and following a power manoeuvre.

- Locate the new fuel so that the overall flux shape stays flat and so that individual zone levels do not have to deviate too much from the average to keep it flat.

Notes:

Differential fuelling and fuelling adjacent channels in the opposite direction assist in keeping the overall flux shape flat, as described earlier in this section. In addition, fuel is added preferentially to zones with low zone level to raise them toward the average.

At times, it is not possible to fuel optimally because of technical problems. The liquid zone control system is able to keep axial flux shape flat after fuelling several channels consecutively in the same direction. Similarly, after fuelling several channels consecutively in the same zone, zone to zone flux is equalized by liquid zone level adjustments. In either case, however, individual zone levels will deviate significantly from the average. Very high or very low individual zone levels must not be allowed or spatial regulation is lost. This could lead to a channel or bundle power above the license limits, or an NOP trip or a setback could occur.

- Distribute the fuel to minimize fuelling ripple.

Fuelling near a low burnup channel can result in both channels building to near maximum reactivity (the plutonium peak) at the same time. The computer simulation helps the fuelling engineer anticipate and avoid hot spots, which could necessitate power derating to maintain safe operation. The CPPF could increase, causing reduced margin to trip or, in an extreme case, inadequate trip coverage. Channel or bundle power could increase to the license limit in the hot region.

- Replace fuel with the highest burnup.

It makes economic sense to fuel channels so that the highest burnup fuel is discharged. Typically, this also gives the highest reactivity gain per channel fuelled.

The reactivity gain from fuelling is highest in the central core and falls off to the edges. Typical fuelling strategies insert more bundles per visit to outer core channels than to inner core channels. This results in fewer fuelling machine visits to the outer core, even though the average discharge burnup is lower in the outer core than in the inner core. More frequent visits to the inner core channels may be justified by the smaller fuelling ripple that results from inserting only a few fresh bundles.

Replacing high burnup fuel reduces fuel cost; high reactivity gain per channel fuelled reduces frequency of fuelling operations and thus decreases fuelling cost.

- High burnup bundles should be distributed in the core so they are not exposed to high flux and large power ramps.

Notes:

Old bundles at relatively high power are at risk of failure by stress corrosion cracking if stressed by a large power ramp. Such a ramp could be caused by reactivity device movement or by fuelling. Good overall fuelling strategy, together with correct sequencing of adjuster bank withdrawal and imposition of bulk power limits on adjusted flux shapes have contributed to good CANDU fuel performance. CANDU 37-element fuel bundles have shown no vulnerability to ramp defects, and 28-element fuel bundles only occasionally fail in this way.

Regular fuelling is required to maintain core reactivity. The above principles guide the fuelling strategy, but may at times contradict one another. For example, fuelling to replace the highest burnup fuel may increase ripple; outer core visits are required for overall flux flattening, but do not give as large a reactivity gain as fuelling the inner core. The following is a typical optimum fuelling strategy, in the likely order of priority:

- Remove fuel with known defects at the earliest opportunity.
- Fuel so that the fuelling ripple is controlled
  - Keep the CPPF at an acceptably low value.
  - Ensure adequate operating margin to license limits and trips
- Fuel evenly across the core:
  - Visit the inner and outer core with the right frequency.
  - Alternate the fuelling direction.
  - Avoid fuelling consecutively into the same zone.
- Fuel channels with high discharge burnup and high reactivity gain.

All fuelling operations are documented so the corresponding changes in the reactor core are accounted for in the fuelling simulation program, and so all bundle locations can be traced.

Reactivity Device Movements

A control rod or shutoff rod is typically a cylindrical rod made of some material with a high absorption cross-section for thermal neutrons, for example, cadmium in stainless steel. The worth of a rod is the reactivity change produced when the rod is inserted in the reactor. This depends most obviously on the absorbing material in the rod. It also depends significantly on the position of the rod in the core. The same rod will have a different reactivity worth in different locations.

From time to time individual shutoff rods may become stuck in a reactor at high power, for example on a routine partial rod drop test to verify fast insertion times.

Occasionally, high zone levels cause a bank of control absorbers to drive into the core to reduce zone levels.

Let us consider the worth of a single shutoff rod inserted into the high-flux region at the centre of the core (Figure 8.6). Since the rod is an effective absorber of neutrons, the thermal flux in the region around it goes down significantly. If the reactor continues operating at the same bulk power with the rod inserted, the regulating system must compensate for the decreased flux in the central region by raising the flux in the outer regions, as shown in Figure 8.7. The result is to displace the flux towards the edge of the core, which will lead to a greater thermal neutron leakage.

Notes:

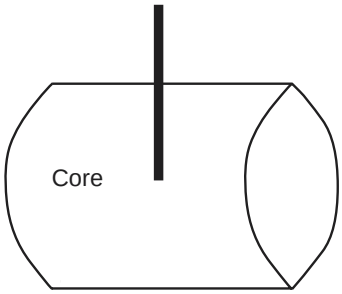

Figure 8.6
Control Road Inserted into Central Core Region

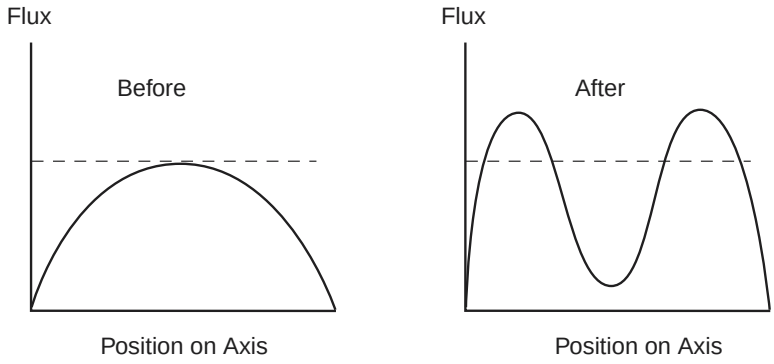

Figure 8.7
Thermal Flux Shape Before and After Insertion of a Rod

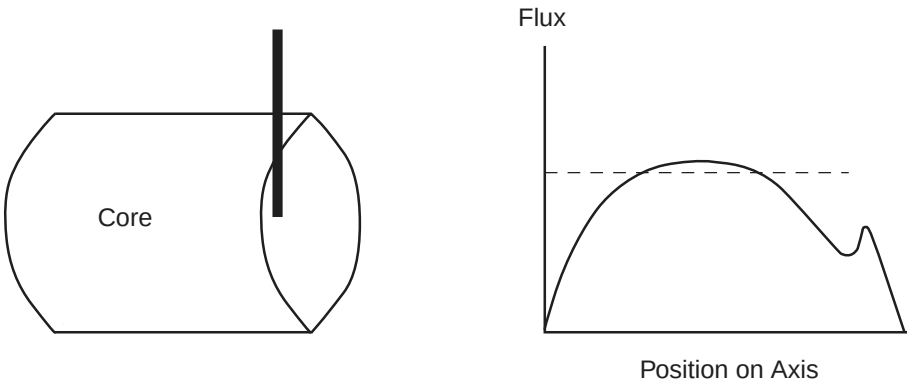

Figure 8.8
Thermal Flux Shape when Rod is Inserted Near the Edge of the Core

Notes:

The worth of a rod therefore results from a combination of two effects:

- increased neutron absorption in the region where it has been inserted, and
- increased leakage caused by the flux distortion produced by the rod.

The size of each effect depends on the placement of the rod in the core. Suppose that the rod insertion is near the edge of the core as shown in Figure 8.8. The absorption effect is obviously less than in the previous case because there are fewer neutrons to be absorbed in this lower flux region. The regulating system therefore increases the flux in the rest of the core by a relatively smaller amount to maintain bulk power. The flux distortion and the increase in neutron leakage will also be less than for the rod insertion into a high-flux region. The overall effect is that the rod worth will be considerably less at the edge of the core than in the central region.

Differential Rod Worth

Let us now examine how the reactivity effect of a shutoff rod or control absorber rod varies as it is gradually inserted into the reactor. We can describe this quantitatively by introducing the differential rod worth, defined as the change in reactivity worth with, say, each extra millimetre of rod added to the core. The total reactivity worth is the worth of the fully inserted rod.

Figure 8.9 shows the variation of the reactivity worth of the rod as a function of its position. The total reactivity of this rod, fully inserted, is about 2 mk. Starting with the rod completely withdrawn, the initial movement produces relatively little effect, since the rod is entering a region where the neutron flux is low. As it enters further into the core, the differential worth (worth per millimetre) steadily rises, reaching a maximum as the leading end reaches the core centre. The slope of the Figure 8.9 curve is the differential worth, which, as you can see, is greatest when the end of the rod reaches the centre line of the core. As the end of the rod

progresses beyond the centre, the differential worth decreases again, reaching a minimum at the fully inserted position.

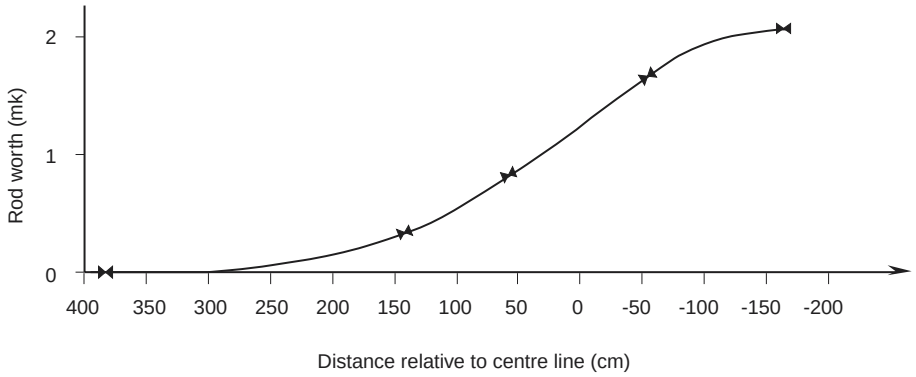

Figure 8.9
Reactivity Worth of Rod as a Function of Position

Notes:

In other words, the rod is most effective when it reaches the high flux central region of the core. This helps explain the need for the spring assisted start that accelerates the shutoff rods when they drop. It also accounts for the great importance placed on testing the shutoff rods to verify the rates of insertion.

In passing, we might note that the partial movement of a rod or a bank of rods into the core would produce the same distortion of the neutron flux shape as described in the previous section, but in the vertical direction. The flux in the upper part of the core will be reduced relative to the flux in the lower part.

Note that the shutoff rods and control absorbers are seldom in the core at high power; they usually drop into core to reduce power rapidly to a low level. The regulating system does not routinely drive control absorbers for reactor regulation. The operator can easily avoid the flux distorting effect of the control absorbers by using the poison addition system to reduce zone levels before rod drives occur. On the rare occasion when the control absorbers drive into core to assist the liquid zone system, the zones drain as a bank of two control absorbers drives in. The way the "bite" of the rod changes with depth has little operating significance; it simply produces a varying rate of zone level decrease.

Rod Shadowing and Anti-Shadowing

The worth of a rod placed in some location depends on whether other rods have already distorted the overall flux shape. To illustrate, suppose that we have inserted a single rod into a high-flux region of the core, as shown in the top portion of Figure 8.10. The regulating system, holding power at setpoint, compensates for the insertion of the rod, distorting the flux shape as shown.

Suppose that we now insert a second identical rod quite close to the first one, as in the middle illustration in Figure 8.10. Since it will be entering a region where the first rod has depressed the flux, it will absorb fewer neutrons than it could have if

the other rod were not there. In addition, the presence of the second rod reduces the flux in the neighbourhood of the first one, thereby reducing its worth compared to the situation where it was the only rod in the core.

Notes:

Thus, if each rod on its own had a worth of, say, x mk, the worth of the two in combination will be less than 2x. The reduction in individual worth of each rod because another is nearby is known as rod shadowing.

Consider now what would have happened if the second rod had been inserted into a region far away from the first one, for example, at the peak of the distorted flux distribution produced by the former, as in the bottom illustration in Figure 8.10. Because it is going into a region where the first rod and the regulating system actions have increased the neutron flux, it absorbs more neutrons than if it were the only rod in the core. Again, the effect is symmetrical and the worth of the first rod increases. The combined worth of the two rods is then greater than 2x. This increase in the worth of each rod due to the presence of the other is known as rod anti-shadowing.

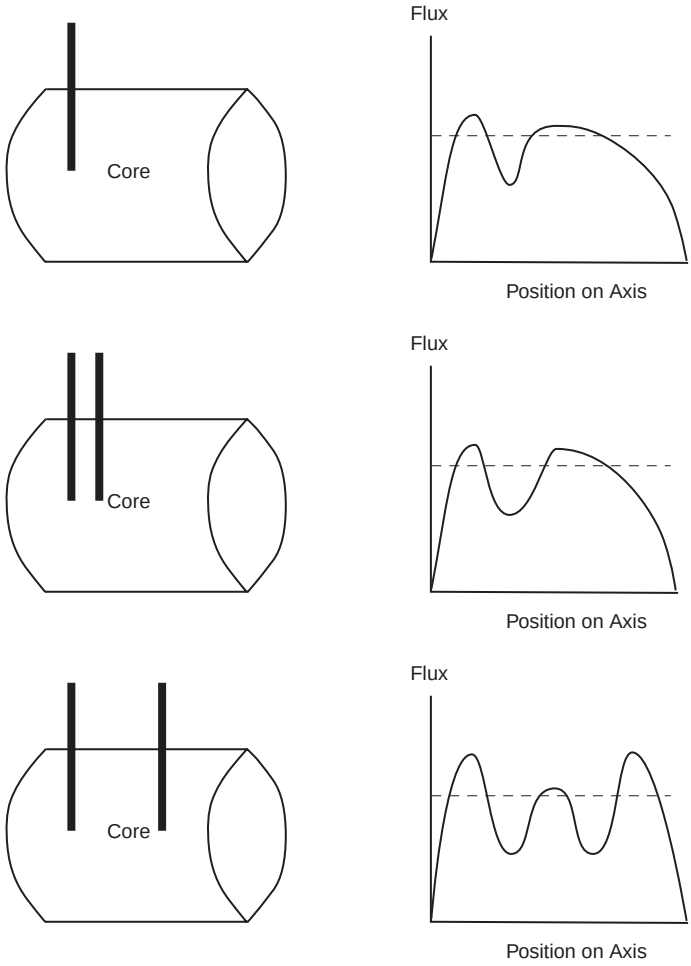

Figure 8.10
Rod Shadowing and Anti-Shadowing

In a CANDU, the effects of rods on one another are most noticeable for the adjuster rods. These are typically arranged in six or seven symmetrical banks of from two to four rods each, with a nominal reactivity worth of a little over 2 mk per bank. The adjusters are made of relatively mild absorbing materials, so they are less flux distorting than the shutoff rods and control absorber rods.

Notes:

Adjuster rods are normally removed from the core (and returned) one bank at a time in an approved sequence. The approved sequence takes account of rod shadowing and anti-shadowing effects so that flux distortions caused by rod interactions do not produce unacceptable flux peaks.

Removal of banks of adjusters in the approved sequence makes the overall flux more peaked—one reason for having the rods in-core is flux flattening. There is an upper limit on bulk reactor power for each allowed configuration of adjuster rods, to limit maximum bundle and channel power.

Equipment problems may on occasion make it impossible to drive the rods in the correct sequence, or may prevent all the rods in a bank from driving. Analysis in advance is required to determine the correct bulk power limit and to make sure local flux peaks are acceptable.

Flux peaking is more severe when adjusters are used as a reactivity shim (for example, if fuelling is not available) than when used for xenon override (for example, following a power decrease). The maximum increase in xenon during the transient is in the highest flux regions of the core, where iodine inventory was highest. This is where the adjusters are located to flatten flux. The increased xenon absorption in these high flux regions helps flatten flux in the region. Operating manuals for some stations may have different bulk power limits on adjuster removal, depending on how the rods are used.

When a bank of adjusters begins to drive, it continues moving until fully out (or until fully in, if it was originally out of core). The regulating system requests adjusters to drive out when there is not enough excess reactivity, that is, when zone levels are low. (RRS requests in-drive on high zone level, but before the level is high enough to call for control absorber in-drive.) The immediate effect of a bank of rods driving out of core is an offsetting zone level increase.

8.3 HIGH POWER PROTECTION

We will consider high power protection under two main headings:

- Regional Overpower Protection (ROP)
- Licence limits on bundle and channel power

Power Rise Limited by NOP/ROP

The ROP systems (one for each shutdown system) provide primary or backup trip coverage for a variety of process failures that cause excessively high power. The system caters particularly to power increases that are too slow to produce a rate log trip.

Notes:

Each ROP system includes three arrays of self-powered in-core detectors to provide three trip channels. For adequate protection, at least one detector in each channel should see any event that produces high bulk power or high local power. Specifically, the trip setpoints ensure that a reactor trip occurs before heat production causes either centerline melting or fuel dryout anywhere in the core.

Centerline melting occurs if the rate of heat production in a fuel element exceeds the capacity of the pellet to transfer heat to the fuel sheath and into the coolant, so that the fuel centerline temperature reaches the melting temperature for UO_2 (about 2750°C). Fuel dryout occurs when the fuel sheath temperature gets so high that excessive boiling of coolant at the fuel surface produces a vapour barrier that limits heat removal from the fuel element.

Dryout results in rapid fuel heating followed by centerline melting. Typically, centerline melting occurs if a bundle power is too high (because of high local flux), and dryout occurs if the channel power is too high (because of high total power output along the channel). Analysis demonstrates protection from both.

The system designers cannot know the initial power in each channel, as these vary from day-to-day because of fuel burnup and refuelling. The designer also does not know what configuration the reactor will be in when a specific event occurs that requires a trip. Finally, the designer must provide coverage for a wide range of initiating events. In the analysis, a trip occurs when any channel power, initially at the reference channel power reaches dryout power (or any bundle reaches centerline melting). Protection for a power rise in a real core (that is, in a reactor with fuelling ripple) occurs because the ROP detector signals are raised (up calibrated) to protect the channel with the highest ripple.

To explain how the trip setpoints are determined and how trip coverage is effected, three items are discussed.

- The design basis set of flux shapes
- Determining the required trips for the reference flux shape
- Taking account of fuelling ripple (CPPF)

The Design Basis Set of Flux Shapes

Trip coverage analysis considers a very large number of events from a variety of initial reactor configurations, at minimum several hundred scenarios. These are events like loss of bulk power control, producing a slow uniform increase in power across the core, or draining of a single zone, causing a bulk power increase with a high local peak superimposed, etc. These and many other upsets are combined

with various initial distortions of the reference flux shape core conditions (for example, reference core at steady high power, reference core with a xenon transient in progress, reference core with a rod stuck in core, etc.) The ROP system must provide overpower protection for a large number of initial flux shapes that could arise due to normal or abnormal movement of reactivity devices, together with changes in xenon concentration. The design basis set of flux shapes is a comprehensive set of flux shapes representing all analyzed upsets.

Notes:

Reference Core Trip Coverage

Each detector must have a trip setpoint such that at least one detector in each array trips before any fuel channel reaches dryout, or before onset of centerline melting, whichever comes first. Conceptually this is quite simple. A computer simulation begins with the reactor in the reference flux shape and simulates each of the analyzed upsets in turn. It allows each design-basis flux shape to increase until the highest channel power reaches the power level corresponding to onset of dryout (or the highest bundle power corresponds to the centerline melting temperature). The computer then notes the thermal flux level at all the in-core detector locations.

As the computer analyzes each of the design basis flux shapes it sorts through the detector readings and selects the least restrictive set of trip setpoints (that is, with maximum margin to trip) that provides the necessary trip coverage.

As more and more restrictive accidents are analyzed, trip setpoints for all detectors move downward, perhaps into the range of 117% to 119%. In the early stages of the design, adjustments to detector locations and their distribution over the arrays improve trip coverage while optimizing the operating margin (margin to trip). The analysts may even make small changes to the reference flux shape (making the reference flux shape slightly different from the time-averaged flux shape). At the end of the process, each ROP system has a set of trip setpoints that provides complete coverage for all incidents in the design basis set**,** provided the reactor fuelling has produced the reference flux shape.

Effect of Fuelling Ripple

The ROP/NOP trip setpoints protect the idealized reference core, as just described. In the real core, fuelling ripple may result in the vulnerable channel (for a particular upset) with a higher power than the reference. The detector expected to provide trip coverage could be reading below the reference value. As power increases uniformly, the vulnerable channel reaches dryout (or a bundle reaches centerline melting) before the detector reaches its trip setpoint. The worst case situation would be if the vulnerable channel was the CPPF channel (the CPPF channel has the highest ripple in the core). This difficulty is solved by recognizing that the actual detector readings don't matter, only the margin to trip matters. The safety system ROP detectors do not need to read 100% with the reactor at full power. Instead, all detectors are adjusted so that they read 100% × CPPF (if the reactor is at 100% of full power). Setting the reading to 100% takes care of the

detectors reading low because of fuelling ripple. Increasing this by the CPPF preserves the margin to trip for the worst case scenario. Trips will occur early if the vulnerable channel is not the CPPF channel.

Notes:

Poor fuelling can produce a high CPPF that results in a small margin to trip. When this happens, the NOP trips can seem very restrictive, with power well below any possibility of dryout or centerline melting. Keep in mind that for some particular slow loss of regulation accident there is a channel that could reach dryout just as the trip setpoint is reached.

If the channel ripple is higher than the CPPF because of some local flux distortion resulting from an unexpected device configuration (perhaps compounded by xenon effects) due, for example, to a shutoff rod stuck partly inserted into the core, trip coverage may be inadequate. The design basis set may include some of these highly unlikely device configurations that give extreme flux shapes. If these configurations require trip settings that are too restrictive in normal operation, they are protected by a hand-switch adjustment that reduces the trip setpoints. (The operator would need to recognize that he is in an abnormal configuration, reduce power appropriately, and adjust the hand-switch).

Inevitably, many possible reactivity device configurations are not in the design basis set. Another hand-switch setting that reduces the trip setpoints even further protects the associated flux shapes, the unanalyzed flux shapes. The analysts check that this trip setpoint provides coverage for a small bounding set of extreme flux shapes, not associated with any particular upset.

We end this section with one final comment about trip coverage. CANDU reactors with 37-element fuel reach dryout conditions before any bundle in the core reaches centerline-melting temperature. The thinner fuel elements (compared to 28-element fuel) provide more total surface area for cooling, and less thickness of uranium dioxide (UO_2 is a poor thermal conductor). CANDU reactors with 28-element fuel are much more likely to reach centerline-melting temperatures before any channel reaches dryout power. The 37-element bundle reactors are described as channel power limited, the 28-element bundle reactors as bundle power limited.

License Limits

An important supplement to the information in the previous section is that the ROP/NOP systems do not prevent the violation of license limits on bundle and channel power. For example, a typical license limit on channel power for 37-element fuel is 7.1 MW. The channel power that will produce dryout, assuming normal coolant conditions, is almost 30% higher than this. Obviously, in this example, a trip intended to prevent dryout may not prevent power increasing beyond the license limit.

The ROP/NOP trip ensures the integrity of the fuel channel by reducing power before there is risk of fuel disassembly that could pose a risk to the channel. An

intact heat transport system guarantees that fission products released to the coolant will not reach the public. The bundle and channel power license limits arise from quite different considerations. Safety analysis demonstrates that radiation releases to the public is within stated limits, and shows adequate primary and backup trip coverage for a whole variety of different upsets. The analysis depends sensitively on the fuel power before the analyzed upset. Adequate trip coverage and protection of the public have not been demonstrated for operation with fuel power that exceeds the values assumed in the analysis; these values establish the license limits.

Notes:

Operation with bundle or channel power above the license limits is operation with excessively high power. Enforcement of the license limits requires the joint efforts of the fuelling engineer, the fuelling crew, and the authorized staff. The fuelling engineer routinely checks the slow evolution of steady state bundle and channel powers with the off-line fuelling code. The code helps select channels for fuelling that provide an adequate margin to the license limits in normal operation.

How does the authorized staff know that a license limit has been violated? There is no direct measurement or panel indication, so the operator does not "know" unless told by the fuelling engineer. However, if the bulk power output is below the license limit and the flux is "flat enough", and the fuelling engineer and fuelling crew have put bundles into the correct locations, then the bundle and channel powers will be within their limits.

This means routine monitoring of device positions and zone levels is part of the enforcement process for license limits on bundle and channel power. Conversely, if the flux shape is not standard, (an unusual tilt—e.g. with zones limiting, or rods stuck in core or some combination of rods giving an off normal configuration) then the operator has little choice but to assume that channel power limits (or bundle power limits) may be exceeded, unless analysis of the particular configuration demonstrates otherwise.