

CANDU Fuel-Management Course

by

B. Rouben

Manager, Reactor Core Physics Branch

Atomic Energy of Canada Ltd.

TABLE OF CONTENTS

1.	Basic Characteristics of the CANDU Lattice	1
1.1	The CANDU Lattice Cell	1
1.2	The Moderator	1
1.3	The Fuel.....	2
1.4	The Coolant	3
1.5	The Pressure-Tube Concept	3
1.6	Fuel Burnup and Irradiation and Effect of Operating Conditions	4
2.	The CANDU Reactivity Devices	6
2.1	Liquid Zone Controllers	7
2.2	Mechanical Control Absorbers.....	7
2.3	Adjuster Rods.....	8
2.4	Moderator Poison	9
2.5	Special Shutdown Systems.....	9
2.5.1	Shutoff Rods (SDS-1).....	10
2.5.2	Liquid-Poison-Injection System (SDS-2)	10
3.	Detector Systems	11
3.1	Zone-Control Detectors	11
3.2	Neutronic Protection System.....	11
3.3	Flux-Mapping System	12
4.	Computational Scheme for CANDU Neutronics	15
4.1	Cell Calculation and the POWDERPUFS-V Lattice Code	15
4.2	Supercell Calculation.....	17
4.3	Finite-Core Calculation and the RFSP Code.....	18
5.	General Discussion on CANDU Fuel Management.....	20
5.1	General Description.....	20
5.2	Overall Objectives	21
5.3	Periods During Operating Life of Reactor	21
5.4	Infinite-Lattice Multiplication Constant	22
6.	Equilibrium-Core Design	24
6.1	Radial Flattening of the Power Distribution.....	24
6.2	Axially Homogeneous Model.....	25
6.3	Time-Average Model	27
7.	Ongoing Reactor Operation with Channel Refuellings.....	32
7.1	The Channel-Power Cycle.....	32
7.2	Channel-Power Peaking Factor	33
7.3	Criteria for Selecting Channels for Refuelling.....	34
7.4	Initial Fuel Load and Transient to Onset of Refuelling.....	35
7.5	Period from Onset of Refuelling up to Equilibrium	36
7.6	Consequences of Fuelling-Machine Unavailability.....	37
7.7	Core-Follow Calculations with RFSP	38
7.7.1	Instantaneous Diffusion Calculations	38
7.7.2	History-Based Methodology for Lattice Properties	39
7.7.3	Modelling of Bulk and Spatial Control.....	39
7.7.3.1	Bulk Control.....	40
7.7.3.2	Spatial Control.....	40
8.	Effects of ¹³⁵ Xe.....	42
8.1	The Xe-I Kinetics	42
8.2	Reactor Startup	43
8.3	Steady-State Xenon Load	43
8.4	Effect of Power Changes on Xenon Concentration.....	44
8.5	Xenon Transient Following a Shutdown.....	44
8.6	Effects of Xenon on Power Distribution	45

9. Summary 47

LIST OF FIGURES

- 1.1 Schematic Face View of CANDU 6 Reactor
- 1.2 Face View of Basic CANDU Lattice Cell
- 1.3 Number of Collisions and Energy Loss per Collision During Moderation
- 1.4 Average Number of Collisions to Thermalize a Fission Neutron
- 1.5 Moderating Ratio of Various Moderators
- 1.6 Various Fuel-Bundle Designs
- 1.7 The CANFLEX Fuel Bundle
- 1.8 Basic Characteristics of CANDU Design
- 1.9 Relationship Between Irradiation and Burnup
- 2.1 Reactivity Worths of CANDU Reactivity Devices
- 2.2 Interstitial Positioning of Reactivity Devices
- 2.3 Plan View of Reactivity-Device Locations
- 2.4 Side-Elevation View of Reactivity-Device Locations
- 2.5 End-Elevation View of Reactivity-Device Locations
- 2.6 Liquid Zone-Control Compartments
- 2.7 Shutoff-Rod Positions
- 2.8 Liquid-Poison-Injection-Nozzle Positions
- 3.1 Positions of Zone-Control Detectors
- 3.2 Ion-Chamber Locations
- 3.3 SDS-1 In-Core-Detector Locations
- 3.4 SDS-2 In-Core-Detector Locations
- 3.5 Triplicated Tripping Logic of CANDU Protection Systems
- 3.6 Positions of In-Core Flux-Mapping Detectors
- 3.7 Harmonic Modes for Flux Mapping
- 4.1 CANDU Supercell
- 4.2a Typical RFSP Reactor Model (Face View)
- 4.2b Typical RFSP Reactor Model (Top View)
- 5.1 The Eight-Bundle-Shift Refuelling Scheme
- 5.2 Infinite-Lattice Multiplication Constant for Standard CANDU-6 Lattice Fuelled with Natural Uranium
- 5.3 Infinite-Lattice Multiplication Constant for Standard CANDU-6 Lattice Fuelled with Depleted Uranium
- 6.1 Schematic of Continuous Bi-Directional Refuelling
- 6.2 Two-Region Model for CANDU 6
- 6.3 Irradiation Values at Beginning and End of Cycle for Eight-Bundle-Shift Refuelling
- 6.4 Multiple-Region Time-Average Model for CANDU 6
- 6.5 Computational Scheme for Time-Average Calculation
- 6.6 Channel-Power Distribution from a CANDU-6 Time-Average Calculation
- 6.7 Channel Dwell Times from the Same CANDU-6 Time-Average Calculation
- 6.8 Neutron Balance in the CANDU-6 Equilibrium Core
- 7.1 Schematic of Maximum Channel Power versus Time
- 7.2 Channels with Depleted Fuel in Initial Core of CANDU 6
- 7.3 Excess Core Reactivity in Initial Period of Reactor Operation
- 7.4 Horizontal Radial Bundle-Power Distribution at 0 FPD
- 7.5 Horizontal Radial Bundle-Power Distribution at 40 FPD
- 7.6 Horizontal Radial Bundle-Power Distribution at 100 FPD
- 7.7 Channel Visit Rate up to 400 FPD
- 7.8 Maximum Channel Power up to 400 FPD
- 7.9 Maximum Bundle Power up to 400 FPD
- 7.10 Average Discharge Burnup as a Function of Time
- 7.11 Average Cumulative Discharge Burnup as a Function of Time
- 7.12 Swing-8 Refuelling Scheme
- 8.1 Xe-I Kinetics

- 8.2 Xenon Reactivity After a Startup
- 8.3 Xenon Reactivity Transients Following Setback to Various Power Levels
- 8.4 Xenon Transient Following Reactor Shutdown

1. Basic Characteristics of the CANDU Lattice

This chapter discusses the basic characteristics of the CANDU lattice. The discussion is largely general but, where necessary, particular reference is made to the CANDU 6 reactor.

1.1 The CANDU Lattice Cell

Figure 1.1 is a schematic face view of the CANDU 6. The CANDU design is modular, with fuel channels set on a square lattice of lattice pitch equal to 28.575 cm. Figure 1.2 shows a face view of the basic lattice cell, dimensions 1 lattice pitch by 1 lattice pitch (28.575 cm x 28.575 cm). In the third direction (perpendicular to the paper) the dimension of the cell is one fuel-bundle length, 49.53 cm. There are twelve fuel bundles in each fuel channel.

The next few sections describe the various components of the lattice cell.

1.2 The Moderator

The CANDU reactor design has opted for heavy water as the moderator, to take advantage of the neutron economy provided by deuterium. This allows the use of natural-uranium fuel and precludes the need for expensive fuel-enrichment technology.

The first desirable property for a moderator is the ability to thermalize neutrons in as few collisions as possible. When the number of collisions required for thermalization is smaller, the average loss of neutron energy per collision is greater, and the probability is enhanced that the neutron will miss the resonance-absorption energy range (or much of it) during moderation (see Figure 1.3). Figure 1.4 shows the average number of collisions needed for various moderators to thermalize a fission neutron.

It is also important that a moderator have a small probability of capturing neutrons (i.e., a small neutron absorption cross section), otherwise there will be a negative impact on neutron economy. In this respect hydrogen is not the best moderator, because it has a relatively high neutron absorption cross section. Heavy hydrogen, or deuterium, on the other hand, has a very low absorption cross section.

A good index of performance for moderators is the moderating ratio, the ratio of the slowing-down power of the material to its neutron absorption cross section:

$$\text{Moderating ratio} = \frac{\text{Slowing-down power}}{\text{Absorption cross section}} = \frac{\xi \Sigma_s}{\Sigma_a}$$

where ξ is the mean logarithmic energy decrement per collision.

The moderating ratio of various moderators is shown in Figure 1.5. It is clear that by far the best moderator for neutron economy is heavy water (D_2O), which is why it was chosen as the moderator for CANDU.

Because of the significant probability of neutron capture by any light water present in the heavy-water moderator, it is crucial that the latter have very high isotopic purity. Reactor-grade moderator must be at least 99.75 % D_2O by weight. Even a reduction of 0.1% in the isotopic purity has a significant effect on the neutron economy of the reactor and on the achievable fuel burnup.

1.3 The Fuel

Natural uranium is used in all currently operating CANDU reactors. This is very convenient for countries which wish not to have to rely on expensive, and most probably foreign, enrichment technology. However, the CANDU design is very flexible and allows the use of advanced fuel cycles, using slightly enriched uranium (SEU), recovered uranium (RU), mixed-oxide fuel (MOX), thorium fuels (Th), and others (DUPIC, actinide burning). These can be introduced into CANDU with few or no hardware changes, when the option becomes attractive.

CANDU fuel is of very simple design. It is manufactured in the form of elements of length ~48 cm. Each element consists of uranium-dioxide pellets encased in a zircaloy sheath. A number of fuel elements are assembled together to form a bundle of length ~50 cm. The elements are held together by bundle end plates. The CANDU fuel bundle contains only 7 different components and is short, easy to handle, and economical.

Various fuel-bundle designs are illustrated in Figure 1.6. Only two bundle types are used in present-generation CANDUs: the 28-element bundle (in Pickering) and the 37-element bundle (in Bruce, Darlington and the CANDU 6). The 28-element bundle has a smaller ratio of sheath mass to fuel mass than the 37-element bundle, which gives the 28-element bundle a reactivity advantage. On the other hand, the 37-element bundle features better thermalhydraulic properties due to the greater fuel subdivision, as the larger number of pins of smaller diameter provide a better heat-removal capability. Thus, the 37-element bundle can operate at a higher power than the 28-element bundle. This tends to further reduce the reactivity of the 37-element bundle, but allows a higher total reactor power for the same mass of fuel, an important economic advantage.

The CANFLEX fuel design (Figure 1.7) has been under development for the last few years as the fuel bundle for the future. In 1998, a demonstration irradiation of 24 CANFLEX fuel bundles has been initiated at the Pt. Lepreau Nuclear Generating Station in New Brunswick, Canada. The CANFLEX bundle has 43 elements, with the outer two rings of elements being of smaller diameter than the inner 7 elements. The CANFLEX bundle features improved thermalhydraulic properties and 20% lower maximum element-power ratings than 37-element fuel, for the same bundle power.

1.4 The Coolant

In all commercial CANDU reactors, heavy water is used as the coolant in the primary heat-transport system, to further improve neutron economy. However, prototype CANDUs have been built using boiling light water or an organic liquid as coolant (Gentilly-1 and WR-1 respectively). The organic coolant, in particular, allows higher temperatures and greater efficiency of conversion of heat to electricity.

1.5 The Pressure-Tube Concept

A major characteristic, selected early in the development of the CANDU reactor, is the pressure-tube design. It is clear that if a liquid is used to remove the large quantity of heat generated inside the reactor, the liquid must be kept at high pressure, otherwise it would boil. The heat-transport-system pressure in CANDU is ~100 atmospheres.

To contain the pressure, the choice is between a pressure-vessel design and a pressure-tube design. In the former (e.g. the PWR) the vessel contains all the fuel and the liquid, which is at once both moderator and coolant. In the pressure-tube design, the moderator and coolant are separate and the coolant flows (at high pressure) through the pressure tubes, which also contain the fuel and in fact comprise the fuel channels. The pressure tubes are made of an alloy of zirconium and 2.5% niobium.

The pressure-tube concept was originally chosen for CANDU because the manufacture of a pressure vessel of the size required for a heavy-water reactor (HWR) would at the time have challenged the capability of Canadian industry. However, the pressure-tube concept has many other advantages in relation to the design and safety of the reactor:

- The rupture of one pressure tube is not as catastrophic as the rupture of an entire pressure vessel. Also, in most cases a pressure tube will leak and give ample warning before rupturing.
- Because the coolant and moderator are physically separated, the moderator can be kept relatively cool. In CANDU the moderator is isolated from the hot pressure tube by a concentric calandria tube made of zircaloy-2. A gas annulus separates the pressure and calandria tubes (see Figure 1.2). Thus the moderator can be kept at about 70 °C and at near atmospheric pressure. Many safety benefits ensue as a consequence. The moderator is a benign, low-pressure and low-temperature environment for interstitial reactivity devices (control rods, etc.); rod-ejection accidents are therefore not a concern. Also, the moderator is a potential ultimate heat sink in case of accident. In addition, a cool moderator further improves neutron economy.
- The pressure-tube concept allows the replacement of fuel in the reactor on power, precluding the need for periodic shutdowns for refuelling. Also,

on-power refuelling means that the excess reactivity in the lattice is never very high, a safety advantage.

These basic characteristics of CANDU are summarized in Figure 1.8.

1.6 Fuel Burnup and Irradiation and Effect of Operating Conditions

Fuel burnup is the amount of energy that is obtained per unit mass of uranium in the fuel. Fuel burnup can be measured in units of MW.h/kg(U) or MW.d/Mg(U). These units are related by the equation:

$$1 \text{ MW.h/kg(U)} = 1 * 1000 / 24 \text{ MW.d/Mg(U)} = 41.67 \text{ MW.d/Mg(U)}$$

In CANDU we also speak of **fuel irradiation**. In the context of the Westcott convention for fluxes and cross sections (see the section below on the lattice code), the irradiation (ω) is the product of Westcott fuel flux by time:

$$\omega = \hat{\phi}t$$

Burnup and irradiation are closely linked. Burnup is a monotonic, nearly linear function of irradiation, as shown in Figure 1.9.

Note that, **for a given type of fuel and reactor**, fuel burnup is essentially the inverse of fuel consumption, i.e., the amount of fuel used to produce a given quantity of energy (electricity), measured for example in units of Mg(U)/GW(e).a. For a given fissile content, a high burnup signifies low fuel consumption, and therefore a small refuelling cost component. High fuel burnup is good, low fuel consumption is good.

However, note that when comparing different fuels or reactors, a higher fuel burnup does not necessarily mean a lower uranium utilization. For instance, the fuel burnup attained in PWRs is much higher than that attained in CANDU, but this is the result of fuel enrichment, not of higher fuel efficiency. In fact, even though PWR fuel burnup may be 3 to 6 times the value in CANDU, the uranium utilization is lower by some 25-28% in CANDU, due to its neutron economy.

A typical fuel burnup attained in the CANDU 6 is 7500 MW.d/Mg(U), or 175-180 MW.h/kg(U). However, the burnup attained depends on the operational parameters of the core.

The burnup is of course influenced by any quantity which affects the core reactivity. Any neutron loss or parasitic absorption which reduces the lattice reactivity will have a negative effect on the attainable fuel burnup. The relationship between reduction in core reactivity and loss of burnup is found to be:

1 milli-k reduction in core reactivity \rightarrow 2.88 MW.h/kg(U) loss in burnup
= 120 MW.d/Mg(U) loss in burnup

Examples of factors which affect the reactivity, and therefore the attainable fuel burnup, are as follows:

- a higher moderator purity increases burnup (the reactivity coefficient of moderator purity = \sim 34 milli-k/atom % purity)
- a higher coolant purity also increases burnup (but much less than a higher moderator purity; the reactivity coefficient of coolant purity = \sim 3 milli-k/atom % purity)
- operating the reactor with moderator poison decreases burnup (the boron reactivity coefficient = \sim 8 milli-k/ppm(B))
- a reflector decreases leakage and increases burnup
- thicker pressure or calandria tubes decrease burnup
- a higher ratio of fuel-sheath mass to fuel mass in a bundle (everything else being equal) decreases burnup
- a lower moderator temperature increases burnup
- flattening the power distribution increases leakage and decreases burnup.

Of course, the lower fuelling cost associated with higher burnup must be weighed against other, possibly opposite, cost components in the total unit energy cost. For instance, a thicker reflector will increase burnup, but beyond a certain point, the cost of the additional reflector may outweigh the benefit in burnup. Similarly, the 37-element bundle is preferred over the 28-element bundle, despite its lower burnup, because it can be operated at a higher power, and it therefore allows a lower reactor capital cost per installed kW. (A bundle with greater fuel subdivision, i.e., a greater number of thin pencils rather than a smaller number of thick pencils, can in general produce higher power for the same pencil rating, measured as kW per m of pencil). The licensed maximum bundle power in the CANDU 6 (using the 37-element fuel bundle) is 935 kW. This corresponds to a fuel rating in the hottest (outermost) element of approximately 60 kW/m. By comparison, the licensed maximum bundle power in the Pickering reactors (using the 28-element fuel bundle) is about 750 kW.

The CANDU 6 has adjuster rods (see next chapter) with a reactivity worth of approximately 15 milli-k. A CANDU reactor which is designed without adjusters has a higher excess reactivity and therefore provides higher burnup. On the other hand, a design with adjusters has the advantages of providing xenon-override capability and the ability to compensate, for a time, for fuelling-machine unavailability.

2. The CANDU Reactivity Devices

Every reactor design must include means of changing and adjusting the system reactivity. These are needed to:

- maintain the reactor critical for normal operation,
- allow power manoeuvres, and
- permit fast reactor shutdown when emergency conditions exist.

In CANDU reactors, the primary long-term method of reactivity control is on-line refuelling, where low-reactivity irradiated fuel is replaced by high-reactivity fresh fuel. On-line refuelling is carried out on a daily or near-daily basis. In addition, however, the CANDU design includes several types of reactivity devices for different kinds of reactivity control. These are described in this chapter.

In the CANDU 6 reactor there are six means of changing the reactivity state of the core besides refuelling. Four of these are used for normal control functions, including controlled shutdown, and two are used by special safety systems for rapid shutdown during accident conditions.

The reactivity devices used for control purposes by the Reactor Regulating System (RRS) in the standard CANDU-6 design are the following:

- 14 liquid-zone-control compartments (H₂O filled)
- 21 adjuster rods
- 4 mechanical control absorbers
- moderator poison.

There are in addition two special shutdown systems (SDS):

- SDS-1, consisting of 28 cadmium shutoff rods which fall into the core from above
- SDS-2, consisting of high-pressure poison injection into the moderator through 6 horizontally oriented nozzles.

Figure 2.1 gives typical reactivity worths and maximum rates of change of reactivity for these devices.

All reactivity devices are located or introduced into guide tubes permanently positioned in the low-pressure moderator environment. These guide tubes are located interstitially between rows of calandria tubes, as shown in Figure 2.2. There exists no mechanism for rapidly ejecting any of these rods, nor can they drop out of the core. This is a distinctive safety feature of the pressure-tube reactor design. The maximum positive reactivity insertion rate achievable by driving all control devices together is about

0.35 milli-k per second, which is well within the design capability of the shutdown systems.

The locations of the reactivity devices are shown schematically in Figures 2.3 (Plan View), 2.4 (Side Elevation) and 2.5 (End Elevation).

The following sections describe the functions of the various types of reactivity devices.

2.1 Liquid Zone Controllers

The purpose of the liquid zone-control (LZC) system is to provide continuous fine control of the reactivity, and hence of the reactor power level. Fine reactivity control is needed because refuelling is not truly continuous, but instead is achieved in small increments (usually eight bundles at one time). Fine reactivity control also compensates for other minor perturbations in parameters, such as temperature changes, which in turn cause small reactivity changes. The liquid zone-control system is also designed to accomplish spatial control of the power distribution, which prevents xenon-induced power oscillations from developing.

The LZC system consists of six vertically oriented units (tubes) running interstitially between the fuel channels from the top to the bottom of the core in the positions shown in Figure 2.6. The two central tubes are divided into three compartments each, and the four outer tubes into two compartments each; the compartments in each unit are separated by appropriately placed bulkheads. There are thus a total of 14 individual zone compartments in the reactor. Variable and controllable amounts of light water (H₂O) are introduced in the compartments, where it serves as a neutron absorber. H₂O is fed to the compartments through small-diameter tubing, and the level of H₂O in each compartment is controlled by varying the relative value of the in-flow and out-flow rates.

The reactor regulating system (RRS) adjusts the H₂O fills in the individual compartments according to the magnitude of the signals from interstitially placed in-core self-powered detectors. The detector systems are described in a later chapter.

2.2 Mechanical Control Absorbers

The zone-control system is normally designed to provide a capability for reactivity control of about ± 3 mk, since this is sufficient to compensate for routine reactivity perturbations due to refuelling, occurring on a semi-continuous basis. For certain, less frequent, events, the reactor regulating system requires a greater reactivity range than the zone-control system can provide. Therefore, two additional reactivity-device systems are provided, to extend the control capability in the positive and negative reactivity directions. These devices are also operated by the reactor regulating system.

The system used to extend the range of control in the negative-reactivity direction is a system of four mechanical control absorbers (MCAs). These are physically the same

as the shutoff rods (see Section 2.5.1 below), but they do not form part of the shutdown system. The control absorbers are normally parked fully outside the core under steady-state reactor operation. They are moved into the core only when circumstances demand a rapid reduction of the reactor power, at a rate or over a range that cannot be accomplished by filling the liquid zone-control system at the maximum possible rate. Modes of control-absorber insertion range from driving the rods in pairs to all four being dropped in by gravity following release of an electromagnetic clutch.

The mechanical-control-absorber system and the zone-control system can be used to reduce power to a very low value without requiring actuation of either of the shutdown systems. The reactivity worth of the MCAs is such that it can compensate for the reactivity increase due to temperature reduction on shutdown.

The positions of the mechanical control absorbers are shown in Figure 2.3.

2.3 Adjuster Rods

The adjuster-rod system extends the range of the reactor regulating system in the positive-reactivity direction beyond that available from the zone-control system. In the CANDU 6, the adjuster-rod system consists of 21 vertical rods, which can be made of stainless steel or cobalt. The reactor is designed to operate with the adjuster rods fully inserted in the core during normal operation. If more positive reactivity is required than the zone-control system can provide, the adjuster rods are withdrawn in groups (banks) as necessary.

There are two circumstances where the reactivity decreases, relative to the normal steady-state-power condition, to a degree that demands withdrawal of some or all of the adjuster rods to permit the continuing operation of the reactor:

- the unavailability of fuelling machines for a period of more than about one week, after which the reactivity decrease due to incremental irradiation of the fuel typically exceeds the range available in the zone-control system, and
- transient increases in the concentration of ^{135}Xe following a reduction of reactor power.

Since the adjuster rods are normally fully inserted in the core, their position in the reactor, and the distribution of absorbing material amongst the rods, are chosen to flatten the power distribution, in conjunction with burnup flattening, to achieve the design power shape.

The positions of the 21 adjuster rods in the CANDU 6 are shown in Figure 2.3. The adjusters are grouped into seven banks, not all composed of an equal number of adjusters. The banks are chosen such that the reactivity worth of any one bank does not exceed the range of the zone-control system. The reactivity worth of the complete

system is about 15 mk. The maximum rate of change of reactivity associated with moving one bank of adjusters is < 0.1 mk per second.

The CANDU-6 adjuster system is nominally designed to have sufficient reactivity to compensate for the increase in ^{135}Xe concentration that occurs within approximately 30 minutes following a reactor shutdown. It also provides capability to operate with fuelling machines unavailable for about a month; however, to operate in steady state with adjuster banks out of core, the power level must be reduced to compensate for the radial power peaking caused by adjuster withdrawal.

Note: Some reactors, such as Bruce A reactors, are designed without an adjuster-rod system. In these reactors, extending the reactivity range in the positive direction can be achieved by routinely operating the reactor with a certain amount of poison in the moderator (see next Section), and removing this poison (in whole or in part) by means of ion-exchange columns when positive reactivity is required.

2.4 Moderator Poison

Moderator poison is used to compensate for excess reactivity:

- in the initial core, when all fuel in the core is fresh, and
- during and following reactor shutdown, when the ^{135}Xe concentration has decayed below normal levels.

Boron is used in the initial core, and gadolinium is used following reactor shutdown. The advantage of using gadolinium after shutdown is that its burnout rate during operation at full power following an extended shutdown period is comparable to the xenon growth rate in terms of reactivity, hence the need to remove poison by ion exchange at a fairly rapid and controlled rate is much less demanding. Poison can be added to the moderator for these purposes either automatically or manually.

It should be noted that the moderator-poison-addition system is completely independent of the very-high-speed liquid-poison injection system which is used as a shutdown system (see Section 2.5.2 below). In the regulating-system function, the poison is inserted into the piping used to circulate the moderator, whereas, in the poison-injection system, the poison is injected through nozzles that are installed horizontally across the core, and a completely independent source of poison is used.

2.5 Special Shutdown Systems

The CANDU 6 reactor is equipped with two physically independent special shutdown systems, SDS-1 and SDS-2. These systems are designed to be both functionally different from each other, and physically separate. These differences are achieved by using vertically oriented mechanical shutoff rods in one system and horizontally oriented liquid-poison-injection nozzles in the second system.

2.5.1 Shutoff Rods (SDS-1)

The shutoff rods are tubes consisting of a cadmium sheet sandwiched between two concentric steel cylinders. The rods are inserted vertically into perforated circular guide tubes which are permanently fixed in the core. The locations of these rods in the CANDU 6 are shown in Figure 2.7. The diameter of the rods is the maximum that can be physically accommodated in the space between the calandria tubes (about 113 mm), when space for the guide tubes and appropriate clearances are considered. The outermost four rods are about 4.4 m long, while the rest are about 5.4 m long. The rods are normally parked fully outside the core and are held in position by an electromagnetic clutch. When a signal for shutdown is received, the clutch releases and the rods fall by gravity into the core, with an initial spring assist.

2.5.2 Liquid-Poison-Injection System (SDS-2)

The alternative way of shutting down the reactor is by high-pressure injection of a solution of gadolinium into the moderator in the calandria. The gadolinium solution is normally held at high pressure in vessels outside of the calandria. Injection is accomplished by opening high-speed valves which are normally closed. When the valves open, the liquid poison is injected into the reactor moderator through six horizontally oriented nozzles that span the core. The nozzles are located in the positions shown in Figure 2.8, and are designed to inject the poison in four different directions in the form of a large number of individual jets. This disperses the poison rapidly throughout a large fraction of the core. The gadolinium solution is held in the retaining pressure vessels at a concentration of typically about 8000 g of gadolinium per Mg of heavy water.

3. Detector Systems

This chapter describes the CANDU detector systems.

3.1 Zone-Control Detectors

To vary the amount of water in the zone-control compartments, the Reactor Regulating System utilizes the readings of detectors associated with the zone controllers. These are fast-response platinum detectors, placed interstitially between fuel channels. There is one detector (plus one spare) for each zone-control compartment. Each detector is located close to the midpoint of the zone-control compartment to which it is associated (see Figure 3.1).

To determine changes required in the water fills of the various compartments, the RRS compares the 14 instantaneous detector readings, ϕ_i , with a set of reference readings, ϕ_i^{ref} , corresponding to the desired power distribution at full power.

In the bulk-control function, the average of the 14 readings ϕ_i is used as the indicator of current power, and the water fills in all compartments are uniformly increased or decreased to move the reactor power down or up to the desired power. Bulk control is exercised automatically by the RRS every half second.

In the spatial-control function, the **relative** values of the ϕ_i are compared to the relative reference values to determine the reactor zones in which the flux is low (i.e., in which power should be raised), and those in which it is high (i.e., in which power should be reduced). The water fills are then moved differentially. In zones where power is to be increased the water level is lowered, and where power is to be decreased the water level is raised. The RRS exercises the spatial-control function automatically every 2 seconds.

Because the zone-control detectors provide essentially “point” readings in the core (the detectors are 3 lattice pitches long but span a very small part of each zone), it is legitimate to ask whether they represent fairly the zones to which they are associated. In order to ensure that the readings used by the RRS do reflect zone-average values, the zone detectors are calibrated every two minutes to zone fluxes obtained by the on-line flux-mapping program (see Section 3.3).

3.2 Neutronic Protection System

CANDU reactors are equipped with protection systems which detect an emergency situation and actuate the safety system(s) discussed in the previous Section. The CANDU-6 neutronic protection systems are described here.

There is a separate neutronic protection system for each of the two shutdown systems. Each protection system is triplicated and consists of out-of-core ion chambers and in-core self-powered detectors. Triplication means that there are three separate

“logic” (or “safety”) channels for each protection system. These channels are labelled D, E, and F for SDS-1 and G, H, and J for SDS-2. In each protection system, it suffices that two of the three logic channels be “tripped” for the corresponding shutdown system to be actuated.

There are three ion chambers in each protection system, one per logic channel. The ion chambers are located at the outside surface of the calandria (see Figure 3.2). Each ion chamber trips its logic channel when the measured rate of change of the logarithm of the flux ϕ , i.e. the quantity $\frac{d \ln \phi}{dt}$ exceeds a pre-determined setpoint (e.g. 10% per second, i.e., 0.10 s^{-1} , for SDS-1 in the CANDU 6).

There are also a number of fast-response (platinum or inconel) in-core detectors in each protection system: 34 for SDS-1, located in vertical assemblies (see Figures 3.3a, 3.3b and 3.3c), and 24 for SDS-2, located in horizontal assemblies (see Figure 3.4). The detectors are distributed among the various logic channels, so that channels D, E and F contain 11 or 12 detectors each, while channels G, H, and J contain eight each. The detectors trip the logic channels on high neutron flux: when the reading of any one detector reaches a pre-determined setpoint, the logic channel to which it is connected is tripped. Because the in-core detectors are designed to protect the reactor against high local flux, the in-core-detector system is sometimes referred to as the regional-overpower-protection (ROP) system.

The setpoints of the in-core detectors are determined by an extensive analysis of hypothetical loss-of-regulation accidents. The analysis involves the calculation of hundreds of different flux shapes which can apply in the reactor. The ROP setpoints are designed to protect against critical values of channel power being reached; the current criterion for critical channel power is fuel dryout. The setpoints must also ensure the efficacy of the shutdown systems in arresting the power pulse which follows a hypothetical loss-of-coolant accident.

In summary, there are two separate ways in which a protection-system logic channel can be tripped:

- on a high rate of log neutron flux at the corresponding ion chamber, and
- on high neutron flux at any one detector belonging to the logic channel.

A shutdown system is actuated whenever two of the three corresponding logic channels are tripped. The triplicated tripping logic described here is shown schematically in Figure 3.5. The triplication assures an extremely high reliability of shutdown-system actuation under accident conditions.

3.3 Flux-Mapping System

The CANDU 6 is provided with a flux-mapping system to synthesize the 3-dimensional flux distribution in the reactor from in-core detector readings. The system

consists of 102 vanadium detectors placed at various positions in the core (see Figure 3.6). Each detector is one lattice pitch long.

The flux-mapping procedure consists of assuming the 3-dimensional flux distribution can be written as a linear combination of a number of basis functions or flux modes, i.e. that the thermal flux at any point r in the core, $\phi(r)$, can be expressed as a linear combination of flux modes $\psi_n(r)$:

$$\phi(r) = \sum_{n=1}^m A_n \psi_n(r) \quad (3.1)$$

where m is the total number of modes used and A_n is the amplitude of the n th mode.

Using this linear expansion, the mode amplitudes A_n are determined by a least-squares fit of the calculated fluxes at the 102 detectors to the measured fluxes. For a detector d at position r_d , the mapped flux is, from Eq. (3.1):

$$\phi(r_d) = \sum_{n=1}^m A_n \psi_n(r_d) \quad (3.2)$$

and this can be compared to the measured flux at the detector, F_d .

The flux-mapping procedure determines the amplitudes A_n by minimizing the sum ε of squares of differences between the mapped and measured fluxes, i.e. minimizing

$$\varepsilon = \sum_{d=1}^{102} w_d (\phi_d - F_d)^2 \quad (3.3)$$

where the w_d are chosen weights.

Once the amplitudes have been evaluated, the flux at any point in the reactor can be calculated very easily from Eq. (3.1). Thus, the 3-dimensional flux and power distributions in the core can be derived. The flux-mapping procedure is very quick.

The flux modes $\psi_n(r)$ used in flux mapping consist in the first instance of a number (~ 15) of pre-calculated harmonics of the neutron diffusion equation. These harmonics represent various possible global perturbations of the flux distribution (see Figure 3.7).

For situations in which the reactor is operated with mechanical control absorbers in-core or adjusters out-of-core, the harmonics are complemented by a number of “device modes” which represent the more localized perturbations due to device movement.

The flux-mapping procedure is carried out automatically in the on-line computer every two minutes. It provides the mapped values of average zonal flux to the regulating system. These zonal fluxes are used to calibrate the zone-control detectors, to ensure that the readings of the zone detectors faithfully represent the overall flux distribution in the reactor.

Flux mapping can also be done “off line”, using recorded flux measurements at the detectors corresponding to any desired time in the reactor history.

4. Computational Scheme for CANDU Neutronics

The computational scheme for CANDU neutronics consists of three stages:

- Cell calculation: to determine lattice properties for basic lattice cells
- “Supercell” calculation: to determine the “incremental” cross sections to be added to the basic-cell properties to account for the effect of reactivity devices
- Finite-core calculation: to solve the neutron-diffusion problem in the reactor core, to calculate the 3-dimensional flux and power distribution.

Computer programs have been developed to perform the calculations corresponding to each stage in the above process. These are now briefly discussed in the following sections.

4.1 Cell Calculation and the POWDERPUFS-V Lattice Code

The cell (or lattice) code which has traditionally been used for CANDU design and analysis is POWDERPUFS-V. This is an empirical-recipe code, based on the results of measurements made on heavy-water-moderated lattices in research reactors ZEEP and ZED-2 at Chalk River Laboratories.

POWDERPUFS-V uses the four-factor formula for the infinite-lattice multiplication constant (k_∞). It also utilizes the Westcott formulation for nuclide cross sections, which is a parametrization in terms of the neutron temperature and the spectral parameter r , applicable to highly thermalized neutron spectra (such as those in the CANDU lattice cell, where over 95% of neutrons in the fuel have a Maxwellian energy distribution). The cross sections are evaluated using the Westcott formula and other simple recipes, with parameter values obtained empirically from experiment. POWDERPUFS-V is applicable to CANDU reactors fuelled with natural uranium, where the amount of plutonium in the fuel is limited by the natural-uranium burnup.

The Westcott convention for calculating the effective cross sections of fuel nuclides is based on assuming that the neutron spectrum can be written as the sum of a Maxwellian function and an epithermal function tending to $1/E$:

$$n(v) = N(1-f)\rho_m(v) + Nf\rho_e(v) \quad (4.1)$$

where $\rho_m(v)$ and $\rho_e(v)$ are the Maxwellian and epithermal normalized density distribution functions, respectively, and

$N =$ total neutron density

$f =$ fraction of the total neutron density in the epithermal spectrum

$$\rho_m(v) = \frac{4}{\sqrt{\pi}} \frac{v^2}{v_T^3} e^{-(v/v_T)^2} \quad (4.2)$$

$$\rho_e(v) = v_T \sqrt{\mu} \frac{\Delta(v)}{v^2} \quad (4.3)$$

and $v_T =$ velocity of a neutron having energy kT

$\Delta(v)$ is an empirical function describing the way the epithermal spectrum (with its $1/E$ “tail”) joins the Maxwellian spectrum. It satisfies $\Delta(v) \rightarrow 0$ for $E < \mu kT$ and $\Delta(v) \rightarrow 1$ for $E > \mu kT$, where μkT represents the lower limit of the $1/E$ spectrum and, by choice of convention, $\mu = 3.681$.

The Westcott flux $\hat{\phi}$ is defined as

$$\hat{\phi} = N v_0 \quad (4.4)$$

where $v_0 = 2200$ m/s

and the effective (Westcott) cross section $\hat{\sigma}$ of a given nuclide is defined so that its product with the Westcott flux gives the reaction rate:

$$\text{Total reaction rate in nuclide} = \hat{\sigma} \hat{\phi} = \hat{\sigma} N v_0 \quad (4.5)$$

(By total reaction rate is meant the reaction rate in the whole spectrum, which includes the Maxwellian and the $1/E$ parts.)

It can be shown that $\hat{\sigma}$ can be written in terms of σ_0 , the 2200-m/s cross section, as follows:

$$\hat{\sigma} = \sigma_0(g + rs) \quad (4.6)$$

where g is the ratio of the reaction rate of the nuclide in a pure Maxwellian spectrum to that of a $1/v$ absorber of the same 2200 m/s cross section (i.e., g is a measure of the ‘non- $1/v$ ’ character of the absorber in a Maxwellian spectrum)

and

r is a measure of the epithermal part (i.e., the ‘hardness’ of the spectrum).

r has a small value in the CANDU lattice: typically, $r \sim 0.05$ for a CANDU lattice fuelled with natural UO_2 . This is what makes the Westcott formulation a good approximation in CANDU reactors.

Both g and s are functions of the neutron temperature T_n , so that Eq. (4.6) is evaluated in fact as

$$\hat{\sigma}(r, T_n) = \sigma_0(g(T_n) + rs(T_n)) \quad (4.7)$$

In POWDERPUFS-V, the factors $g(T_n)$ and $S(T_n)$ are expressed as power series in the neutron temperature. With this database of g and s values for various nuclides, POWDERPUFS-V can calculate reaction rates in the fuel very quickly, using Eq. (4.7). For other materials, constant inputs or simple recipes are used. The methodology requires the evaluation of r and the neutron temperature T_n in an iterative fashion from the lattice parameters via empirical relationships. Fast fission is taken into account in ^{238}U only, and is “lumped” into the thermal-fission cross section; also, up-scattering is ignored. Due to its semi-empirical nature and the simplifying assumptions used, POWDERPUFS-V is very fast-running.

POWDERPUFS-V provides “homogenized-cell” two-energy-group lattice properties for input into finite-core models and calculations. In standard “*fuel-burn*” mode, the lattice properties are provided as functions of fuel irradiation (or burnup) for specified values of lattice conditions, such as fuel, coolant, and moderator temperatures, power level, coolant density, moderator-poison concentration, etc. A “*perturbation*” mode is also provided in POWDERPUFS-V, where lattice properties are evaluated assuming “instantaneous” changes in lattice conditions occurring at various values of fuel irradiation.

POWDERPUFS-V has been used as the lattice code for CANDU reactors for about 30 years, where it has performed very well. One advantage of POWDERPUFS-V is incorporated as a module within the finite-core code RFSP, described in Section 4.3 below.

While POWDERPUFS-V has traditionally been the lattice code for CANDU design and analysis, it will eventually be replaced by a code with a stronger theoretical foundation, a multigroup transport-theory code such as WIMS-AECL. Calculating lattice cross sections with WIMS-AECL is, however, more complex and computationally intensive than with POWDERPUFS-V.

4.2 Supercell Calculation

The effects of reactivity devices on the nuclear properties of the lattice in their vicinity are determined by a supercell calculation, performed with the computer code MULTICELL. A typical supercell is shown in Fig. 4.1. It is essentially a small model volume of the core around a portion of the reactivity device, including a portion of the neighbouring fuel channel (normally oriented perpendicularly to the device). The dimensions of the supercell are typically 1 lattice pitch x 0.5 lattice pitch x 0.5 bundle length. This represents a unit volume over which the effect of the reactivity device is modelled, utilizing the assumption of mirror symmetry about the supercell boundaries.

The calculation provides incremental cross sections, which are to be added to the basic lattice cross sections over “homogenized” supercell volumes along the length of the device.

MULTICELL applies pre-calculated boundary conditions (current-to-flux ratios) on internal surfaces which represent the reactivity device and the fuel (modified to Cartesian geometry). These boundary conditions are calculated using integral transport theory (Kushneriuk's method).

Outside the reactivity-device and fuel regions, MULTICELL uses modified diffusion theory to calculate the 3-dimensional flux distribution. The thermal-neutron flux distribution is assumed flat in this region, based on the contribution to the slowing-down density of the multiple line sources represented by the fuel channels.

Once the flux distribution in the supercell has been calculated, MULTICELL derives the incremental cross sections appropriate for application to the homogenized supercell.

4.3 Finite-Core Calculation and the RFSP Code

Once basic-lattice properties and reactivity-device incremental cross sections are available, the finite-core calculation can proceed. The finite-core computer code RFSP (Reactor Fuelling Simulation Program) is specifically designed for CANDU reactors. It can calculate the steady-state 3-dimensional flux and power distributions in the reactor using two different methods:

- by solving the time-independent finite-difference diffusion equation in two energy groups, and
- by the method of flux mapping (described in Section 3.3), if the readings of the in-core vanadium detectors are available.

The time-independent neutron-diffusion equation solved in RFSP for eigenvalue problems in two energy groups with lattice properties from POWDERPUFS-V is:

$$-\vec{\nabla} \cdot D_1(\vec{r}) \vec{\nabla} \phi_1(\vec{r}) + (\Sigma_{a1}(\vec{r}) + \Sigma_{1 \rightarrow 2}(\vec{r})) \phi_1(\vec{r}) - \frac{\nu \Sigma_f(\vec{r})}{k_{eff}} \phi_2(\vec{r}) = 0 \quad (4.8a)$$

$$-\vec{\nabla} \cdot D_2(\vec{r}) \vec{\nabla} \phi_2(\vec{r}) + \Sigma_{a2}(\vec{r}) \phi_2(\vec{r}) - \Sigma_{1 \rightarrow 2}(\vec{r}) \phi_1(\vec{r}) = 0 \quad (4.8b)$$

In this equation there are no fast-fission or up-scattering terms, consistent with the POWDERPUFS-V methodology, where, in particular, fast fissions are lumped into the thermal-fission term. However, for use with WIMS-AECL lattice properties, RFSP has the capability to solve a true-two-energy-group diffusion equation:

$$-\vec{\nabla} \cdot D_1(\vec{r}) \vec{\nabla} \phi_1(\vec{r}) + (\Sigma_{a1}(\vec{r}) + \Sigma_{1 \rightarrow 2}(\vec{r})) \phi_1(\vec{r}) - \left(\Sigma_{2 \rightarrow 1}(\vec{r}) + \frac{\nu \Sigma_{f2}(\vec{r})}{k_{eff}} \right) \phi_2(\vec{r}) = 0 \quad (4.9a)$$

$$-\vec{\nabla} \cdot D_2(\vec{r}) \vec{\nabla} \phi_2(\vec{r}) + (\Sigma_{a2}(\vec{r}) + \Sigma_{2 \rightarrow 1}(\vec{r})) \phi_2(\vec{r}) - \left(\Sigma_{1 \rightarrow 2}(\vec{r}) + \frac{\nu \Sigma_{f1}(\vec{r})}{k_{eff}} \right) \phi_1(\vec{r}) = 0 \quad (4.9b)$$

A typical reactor model used with RFSP is shown in Fig. 4.2 a and b (face and top views respectively).

Major applications of RFSP are in:

- core-design calculations and analyses, including fuel-management design calculations, and simulations of reactor power histories
- core-follow calculations at CANDU sites, to track the actual reactor operating history, with burnup steps and channel refuellings.

Additional capabilities of the program include, among others:

- the calculation of flux distributions for various reactor configurations
- the simulation of $^{135}\text{Xe}/^{135}\text{I}$ transients
- the capability for simulating (quasi-statically) bulk control and spatial control
- the calculation of harmonic flux shapes for use in flux mapping,
- the calculation of the reactivity increase expected on refuelling of individual fuel channels
- the capability for solving neutron-kinetics problems by the Improved Quasi-Static (IQS) method. RFSP can therefore be used to analyze fast transients, such as those following hypothetical large-loss-of-coolant accidents (LOCA), and to simulate the performance of the shutdown systems.

5. General Discussion on CANDU Fuel Management

5.1 General Description

Refuelling operations in CANDU reactors are carried out with the reactor at power. This feature makes the in-core fuel management substantially different from that in reactors which must be refuelled while shut down.

The CANDU on-power refuelling capability also means that long-term reactivity control can be achieved by an appropriate rate of fuel replacement. Therefore, excess core-reactivity requirements are very small:

- Current CANDU reactors use natural-uranium fuel, and the lattice has much smaller excess reactivity than enriched-fuel lattices
- The CANDU fuel bundle (~50-cm long and containing ~19 kg of uranium) allows adding fuel in small increments
- For continuous or short-term reactivity control, a capability of only a few milli-k is necessary; this is provided in the light-water zone-control compartments
- Other than in the initial core, there are no large batches of fresh fuel, and therefore no need for burnable poison or large amounts of moderator poison to compensate for high excess reactivity; in the initial core, when all fuel is fresh, ~2-3 ppm of moderator boron are required

These factors lead to excellent neutron economy and low fuelling costs. Also, since power production is not interrupted for refuelling, it is not necessary to tailor the refuelling schedule to the utility's system load requirements.

To refuel a channel, a pair of fuelling machines latch onto the ends of the channel. A number of fresh fuel bundles are inserted into the channel by the machine at one end, and an equal number of irradiated fuel bundles are discharged into the machine at the other end of the channel. For symmetry, the refuelling direction is opposite for neighbour channels. In the CANDU-6 reactor, the refuelling direction is the same as that of coolant flow in the channel.

Figure 5.1 illustrates the 8-bundle-shift scheme, where the eight bundles near the outlet end of the channel are discharged, and the four bundles previously nearest the inlet end are shifted nearest to the outlet end. Thus, the four low-power bundles are in-core for two cycles and the high-power bundles are in-core for only one cycle.

Several refuelling operations are normally carried out daily, so that refuelling is almost continuous. CANDU reactors offer extreme flexibility in refuelling schemes:

- The refuelling rate (or frequency) can be different in different regions of the core, and in the limit can in principle vary from channel to channel. By using

different refuelling rates in different regions, the long-term radial power distribution can be shaped and controlled.

- The axial refuelling scheme is not fixed; it can be changed at will. It can be different for different channels. It need not even be the same always for a given channel: it can vary at every visit of the channel. Eight-, 4-, or 10-bundle-shift refuelling schemes have been used.
- A channel can be refuelled without delay if failed fuel exists or is suspected. In such a case, when there is concern that replacing **all** fuel bundles in the channel would drive its power too high, some depleted-uranium bundles can be mixed with standard bundles to limit the power. This is made possible by the subdivision of the fuel in a CANDU channel into short bundles.

5.2 Overall Objectives

The primary objective of fuel management is to determine fuel-loading and fuel-replacement strategies to operate the reactor in a safe and reliable fashion while keeping the total unit energy cost low. Within this context, the specific objectives of CANDU fuel management are as follows:

- The reactor must be kept critical and at full power. On-power fuelling is the primary means of providing reactivity. If the fuelling rate is inadequate, the reactor eventually has to be derated
- The core power distribution must be controlled to satisfy safety and operational limits on fuel power
- The fuel burnup is to be maximized within the operational constraints, to minimize the fuelling cost
- Fuel defects are to be avoided. This minimizes replacement fuel costs and radiological occupational hazards
- The fuel-handling capability must be optimized. This minimizes capital, operating and maintenance costs.

5.3 Periods During Operating Life of Reactor

From the point of view of fuel management, the operating life of a CANDU reactor can be separated into three periods. The first two are short, transitional periods, while the third, the “equilibrium core”, represents about 95% of the lifetime of the reactor.

From First Criticality to Onset of Refuelling

The first period is from first criticality until onset of refuelling. It is of limited duration, about 100 to 150 full-power days (FPD) long. The reactor is initially loaded with natural-uranium fuel everywhere, except for a small number of depleted-fuel bundles at specific core locations, designed to help flatten the power distribution. Consequently, at this time, for the only time in the life of the reactor, there is a fair

amount of excess reactivity. This is compensated by adding boron poison to the moderator.

At about 40-50 FPD of reactor operation, the core reaches its “plutonium peak”, at which time the core reactivity is highest, due to the production of plutonium by neutron capture in ^{238}U , and the as-yet relatively small ^{235}U depletion and fission-product concentration. Following the plutonium peak, the plutonium production can no longer compensate for the buildup of fission products, and the excess core reactivity decreases.

Onset of Refuelling and Transition to Equilibrium Core

When the excess core reactivity has fallen to a small value, refuelling begins in order to maintain the reactor critical. During the transitional period which follows, the reactor gradually approaches the final or “equilibrium” state. The average refuelling rate and in-core burnup are transitional but start to converge towards steady values.

Equilibrium Core

Approximately 400 to 500 FPD after initial start-up, a CANDU reactor has reached a state which may be termed an “equilibrium core”. The overall refuelling rate, the in-core average burnup, and the burnup of the discharged fuel have become essentially steady with time. The global flux and power distributions can be considered as having attained an equilibrium, “time-average”, shape, about which the refuelling of individual channels leads to local “refuelling ripples”. These ripples are due to the various instantaneous values of fuel burnup in the different channels, which are the result at any given instant of the specific sequence of channels refuelled.

Note that with some refuelling operations taking place essentially every day, the equilibrium core contains fuel at a range of burnups, from 0 to some average exit-burnup value.

5.4 Infinite-Lattice Multiplication Constant

The infinite-lattice multiplication constant k_{∞} is a measure of the multiplicative properties of the lattice in the absence of leakage from the lattice cell. The k_{∞} is calculated by a cell code, such as POWDERPUFS-V, and applies to the “ideal” situation of an infinite array of identical cells..

Fig. 5.2 shows the k_{∞} as a function of irradiation for the standard CANDU 6 lattice fuelled with natural uranium. The figure shows that the lattice is ~ 80 milli-k supercritical for fresh fuel (i.e., at zero irradiation). The reactivity **increases** at first with increasing irradiation, reaching a maximum at approximately 0.4-0.5 n/kb, a phenomenon due to the production of plutonium from neutron absorption in ^{238}U . This reactivity maximum is consequently known as the plutonium peak. Beyond the plutonium peak, the reactivity starts to decrease with increasing irradiation, on account of the continuing

depletion of ^{235}U and the increasing fission-product load, and the lattice reaches zero excess reactivity at an irradiation of about 1.6-1.8 n/kb. This marks a natural point at which the fuel can be targeted for removal from the core, since at higher irradiations the lattice becomes increasingly subcritical, i.e., an increasing net absorber of neutrons. **Thus, channels containing fuel approaching or exceeding these irradiation values become good candidates for refuelling.** This very general statement is made more specific in a later section.

It is instructive to examine also the infinite-lattice multiplication constant for the depleted-uranium lattice. This is shown in Fig. 5.3 for depleted uranium with an initial fissile content of 0.52 atom % (as opposed to 0.72 atom % for natural uranium). Note that the plutonium peak is even more pronounced for depleted uranium, a result which is easily explained by the fact that the role of ^{238}U conversion to plutonium is relatively greater when the smaller ^{235}U content. Note also, however, that the depleted-uranium lattice is subcritical at all irradiations, i.e. is always a neutron absorber. This explains the use of depleted fuel to reduce excess reactivity, and also flatten the flux distribution, in the initial core. Depleted fuel is also occasionally used to reduce the power ripple on refuelling.

6. Equilibrium-Core Design

As discussed in Section 5.3 above, the equilibrium core contains fuel with a distribution of burnups, and therefore also of irradiations, ranging from zero to discharge values. Ideally, while the nominal (“average-picture”) equilibrium core could be analyzed by averaging over time simulations of a long period of reactor operation, this is not very practical.

Instead, there are two models, representing different levels of approximation, which have been used for the analysis of the nominal equilibrium core: the axially homogeneous model and the time-average model.

The axially homogeneous model is an older, cruder model which was developed in the early days of CANDU design and which was used up to about 20 years ago. It has since then been essentially abandoned in favour of the more refined time-average model. Although the axially homogeneous model is no longer in routine use, it is still instructive to study it, as it introduces useful quantities and concepts. The two models are described in detail in turn in sections 6.2 and 6.3, following the discussion of power flattening.

6.1 Radial Flattening of the Power Distribution

The flux distribution in the reactor depends on the reactor size and geometry and on the distribution of irradiation. Fuel with a high irradiation has low reactivity, and depresses flux in its vicinity. Similarly, the neutron flux tends (everything else being equal) to be high in regions where the fuel has low irradiation. This fact can be used to “shape” the flux (and power) distributions in the equilibrium core.

Radial flux (and power) flattening can be achieved by **differential fuelling**, i.e. taking the fuel to a higher burnup in inner core regions than in outer core regions. This can be done by judicious adjustment of the relative refuelling rates in the different core regions. In this way the flux and power in the outer region can be increased, resulting in a greater number of channels having power close to the maximum value. Thus, a higher total reactor power can be obtained (for a given number of fuel channels) without exceeding the limit on individual channel power. This reduces the capital cost of the reactor per installed kW.

The radial flattening is quantitatively measured by the radial form factor:

$$\text{Radial form factor} = \frac{\text{Average channel power}}{\text{Maximum channel power}}$$

Radial flattening is further assisted by the use of adjuster rods (described in Section 2.3), whose main purpose is in xenon override. Adjuster rods also provide axial power flattening.

Note that while flattening of the power distribution reduce the reactor capital cost, by reducing the number of channels required to produce a given total power, it does tend to increase the neutron leakage, which is proportional to the flux gradient at the edge of the core. This loss of neutrons does have a consequent increase in fuelling cost.

6.2 Axially Homogeneous Model

The axially homogeneous model is based on:

- the continuous-refuelling approximation, i.e., the approximation that fuel is pushed **continuously** along every channel at a constant rate (which may vary from one burnup region to another),
- the bi-directional feature of refuelling in CANDU, i.e., the fact that neighbouring channels are refuelled in opposite directions, and
- the approximation that the flux, as a function of distance from the refuelling end, is equal in neighbouring channels in the same burnup region.

With these assumptions, it can be shown that the irradiation, averaged over two neighbouring channels, is independent of axial position in the core. The mathematical treatment is simple and is shown in the next few equations, (6.1)-(6.10) .

Consider two neighbouring channels; these are refuelled in opposite directions. Suppose that fuel moves continuously at speed r in these two channels (in opposite directions). See Figure 6.1.

The fuel irradiation ω is defined as the product of Westcott flux in the fuel ($\hat{\phi}(x(t))$) and time. Since the fuel moves through the channels and is exposed to different flux values as it moves, the fuel irradiation in each channel at position $x = x_0$ (measured from one end of the reactor, taken as the origin for the axial co-ordinate) must be calculated as the integral over time of fuel flux as fuel travels through the core to $x = x_0$:

$$\omega_1(x_0) = \int_0^{t_1(x_0)} \hat{\phi}(x(t))dt \tag{6.1}$$

$$\omega_2(x_0) = \int_0^{t_2(x_0)} \hat{\phi}(x(t))dt \tag{6.2}$$

where $t_1(x_0)$ and $t_2(x_0)$ are the times to reach position x_0 :

$$t_1(x_0) = \frac{x_0}{r} \tag{6.3}$$

$$t_2(x_0) = \frac{L - x_0}{r} \tag{6.4}$$

Use Equations (6.3), (6.4) to rewrite ω_1 and ω_2 as integrals over x :

$$\omega_1(x_0) = \int_0^{x_0} \hat{\phi}(x) \frac{dx}{r} = \frac{1}{r} \int_0^{x_0} \hat{\phi}(x) dx \quad (6.5)$$

$$\omega_2(x_0) = \int_L^{x_0} \hat{\phi}(x) \left(-\frac{dx}{r}\right) = \frac{1}{r} \int_{x_0}^L \hat{\phi}(x) dx \quad (6.6)$$

The average fuel irradiation at $x = x_0$ over the two channels is thus:

$$\begin{aligned} \omega_{avge}(x_0) &\equiv \frac{1}{2} [\omega_1(x_0) + \omega_2(x_0)] \\ &= \frac{1}{2r} \left[\int_0^{x_0} \hat{\phi}(x) dx + \int_{x_0}^L \hat{\phi}(x) dx \right] \\ &= \frac{1}{2r} \int_0^L \hat{\phi}(x) dx \end{aligned} \quad (6.7)$$

[The last step requires the assumption that the flux shape (as a function of distance from the inlet end of a channel) is the same in the two neighbouring channels.]

The right-hand side of Eq. (6.7) does not depend on the position x_0 in the channels. This shows that we can treat the average irradiation as uniform along the two channels, equal to one half of the fuel exit irradiation (ω_{exit}) at the channel outlet:

$$\omega_{exit} = \omega_1(L) = \omega_2(0) = \frac{1}{r} \int_0^L \hat{\phi}(x) dx \quad (6.8)$$

Therefore it is appropriate within the assumptions for this “axially-homogeneous” model to use uniform basic-lattice properties along the length of the two channels. Although we could define the average cross sections as the values at half the exit irradiation, a more “correct” treatment defines the average cross section of the fuel as the value which preserves the total reaction rate as the fuel travels through core. Let us label the cross sections in this homogeneous model $\Sigma_i(\text{hom})$, with the subscript i representing the various processes (absorption, production, etc.) The cross section which preserves the reaction rate is

$$\begin{aligned} \Sigma_i(\text{hom}) &\equiv \frac{\text{reaction rate averaged overtime}}{\text{flux averaged over time}} \\ &= \frac{\frac{1}{T} \int_0^T \Sigma_i(\omega(t)) \hat{\phi}(x(t)) dt}{\frac{1}{T} \int_0^T \hat{\phi}(x(t)) dt} \end{aligned} \quad (6.9)$$

where T = transit time of fuel through the core:

$$T = \frac{L}{r}$$

Integrals over time can be replaced by integrals over irradiation ($d\omega = \hat{\phi} dt$), and we get:

$$\Sigma_i (\text{hom}) = \frac{1}{\omega_{\text{exit}}} \int_0^{\omega_{\text{exit}}} \Sigma_i (\omega) d\omega \quad (6.10)$$

where ω_{exit} is specific to the channel pair.

Most often, the axially homogeneous model is used with only two large core regions, the inner core and the outer core, each with a different value of ω_{exit} . Figure 6.2 shows a two-region model for the CANDU 6.

The cross sections defined in Eq. (6.10) are calculated, corresponding to the value of ω_{exit} in each core region, directly in POWDERPUFS-V, in what is called the “reaction-rate-averaged” mode.

As can be seen, the calculational procedure for the axially homogeneous model is simple. It consists of a simple integral, which can be computed within POWDERPUFS-V. The resulting values of lattice cross sections are applied uniformly (both radially and axially) in large regions with a given value of exit irradiation ω_{exit} . However, the continuous-refuelling approximation does not take into account the axial non-uniformity induced by real refuelling schemes. For this reason, the axially homogeneous model has been largely abandoned in favour of the time-average model, described in the next section.

6.3 Time-Average Model

In the time-average model, the lattice cross sections are averaged over the residence (dwell) time of the fuel at each point (fuel-bundle position) in the core. This allows the effect of the actual refuelling scheme used (e.g. 8-bundle shift, 4-bundle shift, etc.) to be captured. Calculations are performed in the *TIME-AVER module of RFSP. The mathematical framework of this module is described in this section, for the specific case of an 8-bundle-shift refuelling scheme as an example.

Time-average nuclear cross sections are defined at each bundle position in core by averaging the lattice cross sections over the irradiation range $[\omega_{in}, \omega_{out}]$ “experienced” over time by fuel at that position, where ω_{in} is the value of fuel irradiation when the fuel enters that position in core and ω_{out} is the fuel irradiation when the fuel leaves that position. As in the axially homogeneous model, the correct definition of cross section is the one which preserves the average reaction rate, but in this case the irradiation range is not 0 to ω_{exit} , but instead ω_{in} to ω_{out} . For example, the time-average thermal neutron absorption cross section at some core position r , $\Sigma_{a2}^{la}(r)$ **Error! Switch argument not specified.**, is

$$\frac{1}{(\omega_{out} - \omega_{in})} \int_{\omega_{in}}^{\omega_{out}} \Sigma_{a2}(\omega) d\omega \text{ Error! Switch argument not specified.} = \Sigma_{a2}^{La}(r) \text{ Error! Switch argument not specified.} \quad (6.11)$$

The basic lattice cross sections inside the integral sign are determined as functions of irradiation using the cell code POWDERPUFS-V, incorporated in its entirety within RFSP as the *POWDERPUF module.

Now, in the time-average model, let ϕ_{jk} be the time-average fuel flux at axial position k in channel j. Here k ranges from 1 to 12 since there are 12 bundles per channel, and j ranges over the channels, e.g. from 1 to 380 in the CANDU 6. The bundle position is labelled jk for short,

Let also T_j be the average time interval between refuellings of channel j (also known as the *dwell time* of channel j).

Then the irradiation increment which the fuel at position jk will experience over its residence time at that position will be

$$\Delta\omega_{jk} = \phi_{jk} T_j \quad (6.12)$$

If the fuel entered position jk with an irradiation $\omega_{in,jk}$, then its exit irradiation from that position, $\omega_{out,jk}$, is given by

$$\begin{aligned} \omega_{out,jk} &= \omega_{in,jk} + \Delta\omega_{jk} \\ &= \omega_{in,jk} + \phi_{jk} T_j \end{aligned} \quad (6.13)$$

When a channel is refuelled with an 8-bundle shift, the first 8 positions in the channel receive fresh fuel and the entrance irradiations for positions 9-12 are simply the exit irradiations from positions 1-4 respectively. Thus we can write in this case (see Figure 6.3):

$$\omega_{in,jk} = 0 \quad k = 1, \dots, 8 \quad (6.14a)$$

$$\omega_{in,jk} = \omega_{out,j(k-8)} \quad k = 9, \dots, 12 \quad (6.14b)$$

More generally, for an N-bundle-shift refuelling scheme, these equations become

$$\omega_{in,jk} = 0 \quad k = 1, \dots, N \quad (6.15a)$$

$$\omega_{in,jk} = \omega_{out,j(k-N)} \quad k = (N+1), \dots, 12 \quad (6.15b)$$

In addition to the refuelling scheme, we have other degrees of freedom in the time-average model. These are the values of exit irradiation $\omega_{exit,j}$ for the various channels j . In principle there are as many degrees of freedom as there are channels. (Of course the values of exit irradiation are not totally free, but are collectively constrained by the requirement to obtain a critical reactor.) The *relative* values of $\omega_{exit,j}$ can be used to “shape” the flux to a desired reference distribution. The exit irradiations are related to the flux in the following way, written here explicitly for the 8-bundle-shift case. Since, in the eight-bundle-shift refuelling scheme, bundles 5 to 12 leave the core at each refuelling, then by definition of exit irradiation

$$\omega_{exit,j} = \frac{1}{8} \sum_{k=5}^{12} \omega_{out,jk} \text{ Error! Switch argument not specified.} \tag{6.16}$$

In view of Equation (6.13) this can be written

$$\begin{aligned} \omega_{exit,j} &= \frac{1}{8} \sum_{k=5}^{12} (\omega_{in,jk} + \phi_{jk} T_j) \text{ Error! Switch argument not specified.} \\ &= \frac{1}{8} \left[\sum_{k=5}^8 (\omega_{in,jk} + \phi_{jk} T_j) + \sum_{k=9}^{12} (\omega_{in,jk} + \phi_{jk} T_j) \right] \text{ Error! Switch argument not} \\ &\hspace{15em} \text{specified.} \end{aligned} \tag{6.17}$$

and in view of Equation (6.15) we can write

$$\omega_{exit,j} = \frac{1}{8} \left[\sum_{k=5}^8 \phi_{jk} T_j + \sum_{k=1}^4 \phi_{jk} T_j + \sum_{k=9}^{12} \phi_{jk} T_j \right] = \frac{T_j}{8} \sum_{k=1}^{12} \phi_{jk} \text{ Error! Switch} \\ \text{argument not specified.} \tag{6.18}$$

It is easy to derive the generalization of this result to an N-bundle-shift refuelling scheme:

$$\omega_{exit,j} = \frac{T_j}{N} \sum_{k=1}^{12} \phi_{jk} \text{ Error! Switch argument not specified.} \tag{6.19}$$

The dwell time T_j therefore satisfies

$$T_j = \frac{N \omega_{exit,j}}{\sum_{k=1}^{12} \phi_{jk}} \text{ Error! Switch argument not specified.} \tag{6.20}$$

We now have all the equations required for the time-average flux distribution to be calculated. These equations are:

- the finite-difference form of the time-independent neutron diffusion equation to solve for the flux distribution,
- Equation (6.20) to compute the dwell time for each channel,
- Equation (6.13) and (6.14) to calculate $\omega_{in,jk}$ and $\omega_{out,jk}$ for each bundle in core,
- Equation (6.11) (and similar equations for the other cross sections) to calculate the time-average lattice properties.

This set of equations must be solved using as input the user-specified target exit irradiations $\omega_{exit,j}$. In order to shape the flux to desired values, and also to take account of the presence of extra “hardware” (device locators, etc., mostly at the bottom of the calandria) which introduces localized absorption, typical time-average RFSP models now use many irradiation regions; see Figure 6.4.

Since consistency must be achieved between the flux, the channel dwell times, the individual-bundle irradiation ranges [ω_{in} , ω_{out}], and the lattice properties, an iterative scheme between the solution of the diffusion equation and the other equations is employed until all quantities converge. Figure 6.5 shows the iterative scheme of calculations.

Typically, several iterations by the user in the values of exit irradiation $\omega_{exit,j}$ in the various regions are also required to obtain a critical reactor and a desired flux shape, often measured by the degree of flattening, or radial form factor.

The time-average model is useful at the design stage, to determine the reference three-dimensional power distribution, the expected refuelling frequency of each channel (or its inverse, the channel dwell time), and the expected value of discharge burnup for the various channels.

Figure 6.6 shows the time-average channel-power distribution obtained with a two-region model for the CANDU 6, where the exit irradiations (approximately 1.8 and 1.6 n/kb in the outer and inner cores respectively) are chosen to produce a radial form factor (ratio of average to maximum channel power) of 0.83. This gives a maximum time-average channel power of about 6.52 MW. This value is substantially lower than the maximum licensed channel power, 7.3 MW. However, remember that there must be sufficient margin above the maximum time-average power to accommodate the refuelling ripple, which could be of the order of 10%. The average in-core fuel irradiation in this model is approximately 1 n/kb.

Figure 6.7 shows the channel dwell times for the same CANDU-6 time-average calculation. It can be seen that the dwell times in the inner core range typically between 150 and 160 full-power days (FPD). In the outer core, the dwell times present a large variation, from about 135 FPD for channels just outside the inner core (where the flux is

still high but the exit irradiation is, by design, lower than in the inner core) to almost 300 FPD for some channels at the outermost periphery of the core.

It is instructive to look at a typical neutron balance in the CANDU-6 equilibrium core. This is displayed in Figure 6.8. This shows that more than 45% of fission neutrons originate from fissions in plutonium. Thus, plutonium contributes approximately half the fission energy produced in a CANDU reactor. (Actually, in fuel near the exit burnup, plutonium contributes about three-quarters of the fission energy generated.) Fast fissions account for 56 fission neutrons out of 1,000. Total neutron leakage is 29 neutrons lost per 1000 neutrons born in fission, representing a 29-milli-k loss (6 milli-k from fast leakage, 23 milli-k from thermal leakage). Resonance absorption in ^{238}U represents a loss of almost 90 milli-k. Parasitic absorption in non-fuel components of the lattice represents a 63-milli-k loss.

7. Ongoing Reactor Operation with Channel Refuellings

After the initial period following first reactor startup, on-power refuelling is the primary means of maintaining a CANDU reactor critical. Thus, a number of channels are refuelled every day, **on the average**. Note that refuelling is not necessarily done **every** calendar day; some stations prefer to concentrate all refuelling operations to 2 or 3 days within each week.

Replacing irradiated fuel with fresh fuel has immediate consequences on the local power distribution and on the subsequent period of operation of the reactor. These must be well understood and are discussed in the following subsections.

7.1 The Channel-Power Cycle

The “refuelling ripple” is the consequence of the daily refuelling of channels and the “irradiation cycle” through which each channel travels. This cycle may be described as follows.

- When a channel is refuelled, its local reactivity is high, and its power will be several percent higher than its time-average power.
- The fresh fuel in the channel then initially goes through its plutonium peak as it picks up irradiation. This means that in fact the local reactivity **increases** for about 40 to 50 FPD, and the power of the channel tends to increase further. The higher local reactivity tends to promote a power increase in the neighbouring channels also.
- Following the plutonium peak, the reactivity of the refuelled channel starts to decrease, and its power drops slowly. Approximately half-way through its dwell time, the power of the channel may be close to the power suggested by the time-average model.
- The reactivity of the channel and its power continue to drop. Eventually, the channel becomes a net “sink” or absorber of neutrons, and nears the time when the channel must be refuelled again. At this time the power of the channel may be 10% or more below its time-average power. When the channel is refuelled, its power may jump by 15 to 20% or even more.

The power of each channel therefore goes through an “oscillation” about the time-average power during every cycle. This cycle repeats every time the channel is refuelled, that is, with a period approximately equal to the dwell time suggested by the time-average model. The cycle length is not **exactly** equal to the dwell time, because channels are not refuelled in a rigorously defined sequence. Instead, as described in the previous section, channels are selected for refuelling based on instantaneous, daily information about the core power and irradiation distributions. In addition, the CANDU fuelling

engineer has much flexibility in deciding how the core should be managed, and in fact can decide to modify the global power distribution by changing the refuelling frequency (dwell time) of various channels.

As individual channels are refuelled and go through their channel-power cycle, the specific sequence of these discrete refuellings results in variability in the instantaneous peak channel and bundle powers in the core. This is illustrated in Figure 7.1, which shows a schematic plot of the maximum channel power versus time and illustrates the difference between maximum time-average channel power, average maximum instantaneous channel power, and absolute maximum channel power.

7.2 Channel-Power Peaking Factor

At any given time, there are several channels in the core which are at or near the maximum power in their cycle. Therefore, the maximum instantaneous channel power is always higher than the maximum time-average channel power, as was evident from the earlier Figure 7.1.

Because many safety analyses are normally carried out in a time-average model, it is very important to quantify how much higher the instantaneous power distribution peaks above the time-average distribution. The Channel-Power Peaking Factor (CPPF) is defined to capture this concept:

$$CPPF = \underset{m}{Max} \left[\frac{CP_{instantaneous}(m)}{CP_{time-average}(m)} \right] \quad (7.1)$$

where m runs over all channels in the core, or at least over all channels except perhaps the channels with the very lowest power, i.e., except the last two outermost rings of channels.

The CPPF value varies from day to day, as the various channels which have fairly recently been refuelled go through their cycle. However, the average CPPF value must obviously depend on the axial refuelling scheme used. The greater the number of bundles replaced at each operation, the greater the reactivity increment, and therefore the greater the refuelling ripple (and therefore the CPPF). When the 8-bundle-shift refuelling is used, a typical value for the CPPF is in the range 1.08-1.10. With a 4-bundle-shift scheme, the typical CPPF is likely to be 1.04-1.05.

The exact value of the CPPF is extremely important because it is used to calibrate the in-core ROP detectors. The hundreds of flux shapes that are used in the ROP safety analysis (to determine the detector positions and setpoints) are all calculated in the time-average model, assuming many different core configurations. But because the real instantaneous channel powers are higher than the time-average powers used in the ROP analysis, channels would reach their “critical channel power” (power at which there is fuel dryout) earlier than in the time-average model. To take this into consideration and

ensure proper safety coverage in the instantaneous power shape, the in-core ROP detectors are calibrated each day to the instantaneous value of CPPF.

In order to maximize the margin to trip, it is obviously important that the CPPF be kept as low as possible. This is why a careful selection of channels to be refuelled needs to be made always. A way in which CPPF can be kept low by design is by using, say, 4-bundle-shift refuelling instead of 8-bundle-shift refuelling, or using a mixed 4- and 8-bundle-shift scheme, where the 4-bundle shifting is done in the inner core (high-power region).

Another way in which poor refuelling strategy could impact on reactor operation is as follows. Concentrated refuelling in the vicinity of an ROP detector will increase its reading, even though this may not increase the CPPF in the core. The high detector reading may lead either to spurious trips or to power deratings (to restore operating margin), both of which lead to loss of power production.

Determining the daily CPPF value, and ensuring detectors are calibrated to the correct value, are on-going duties of the fuelling engineer or reactor physicist at a CANDU nuclear generating station.

7.3 Criteria for Selecting Channels for Refuelling

One of the main functions of the fuel engineer (or site reactor physicist) is to establish a list of channels to be refuelled during the following period (few days) of operation. To achieve this, the current status of the reactor core is determined from computer simulations of reactor operation, the on-line flux mapping system, the ROP and RRS in-core detectors, and zone-control-compartment water fills. The computer simulations of reactor operation provide the instantaneous 3-dimensional flux, power and burnup distributions.

Normally, channel selection will begin with **eliminating** channels which are poor candidates for refuelling. With experience, a fuelling engineer will develop a personal set of rules for eliminating channels. A **typical** (but by no means unique) set of rules may eliminate

- channels with an instantaneous power within 10% of the maximum licensed channel power, as well as their 4 closest neighbours
- channels refuelled recently, say less than 10 FPD prior, as well as their 8 closest neighbours
- channels with a high value of peaking factor (greater than, say, 1.07), as well as their 4 closest neighbours
- channels with low average value of burnup in the bundles which would be discharged (less than, say, 75% of the time-average exit burnup for that channel).

Once channels inappropriate for refuelling have been eliminated, possible lists can start to be developed from the remaining channels. Good combinations of channels for refuelling in the few days to follow will typically contain:

- channels “due to be refuelled”, i.e., channels for which the time interval since the last refuelling is approximately equal to the channel’s dwell time (from the time-average calculation)
- channels with high current value of exit burnup, relative to their time-average exit burnup
- channels with low power, relative to their time-average power
- channels in (relatively) low-power zones (compared to the time-average zone-power distribution)
- channels which, taken together, promote axial, radial and azimuthal symmetry and a power distribution close to the reference power shape
- channels which provide sufficient distance to one another and to recently refuelled channels (to avoid hot spots)
- channels which will result in acceptable values for the individual zone-controller fills (20%-70% range), and
- channels which, together, provide the required reactivity to balance the daily reactivity loss due to burnup (and which will, therefore, tend to leave the zone-controller fills in the desired operational range: average zone fill between 40 and 60%) .

The fuelling engineer will usually have to draw up a list from many options available.

A good way of being confident about a channel selection is to perform a **pre-simulation** of the core following the refuellings. This pre-simulation (especially if it invokes bulk- and spatial-control modelling) will show whether the various power, burnup, and zone-fill criteria are likely to be satisfied, or whether the channel selection should be changed.

7.4 Initial Fuel Load and Transient to Onset of Refuelling

Let us now return to the period which marks the beginning of reactor operation.

In the initial core, fresh fuel is present throughout the core. There is no differential burnup which can assist in flattening the power distribution. Consequently, the power of the central core region would be unacceptably high if no alternate means of flattening the radial power distribution were provided. However, an alternate means **is** readily available: depleted fuel. As we have seen earlier, this depleted fuel is a net absorber of neutrons.

In the initial core of the CANDU 6 (i.e., the initial fuel load), two depleted-fuel bundles (of 0.52 atom percent ^{235}U content) are placed in each of the central 80 fuel channels. This is shown in Figure 7.2. The bundles are located in positions 8 and 9,

where the numbering is from the channel refuelling end. In these axial positions, the depleted-fuel bundles are removed from the core in the first refuelling visit of each of these channels.

Even with some depleted fuel in the core, the fact that all fuel is fresh results in a net excess reactivity in the core. The core reactivity starts at approximately 16 milli-k at full power on FPD 0, and then varies with time as shown in Figure 7.3. Because all the fuel goes through its plutonium peak at about the same time, the excess reactivity initially increases, from about 16 mk to a maximum of about 23 mk between FPD 40 and FPD 50. This excess reactivity is compensated by soluble boron in the moderator. The boron coefficient of reactivity is about -8 milli-k per ppm of boron. Thus the boron concentration (at full power) is initially approximately 2 ppm, rising to about 3 ppm at the plutonium peak. Following the plutonium peak, boron must be removed (by ion exchange) as the excess reactivity drops gradually to zero at about FPD 120.

During this entire first period in the reactor life, refuelling is not necessary since there is already excess reactivity. Actually, refuelling is started about 10 or 20 FPD before the excess reactivity reaches 0, i.e. around FPD 100, because the refuelling rate would be too great if one waited until the last possible moment to start.

Figures 7.4 to 7.6 show representative bundle-power distributions in the horizontal radial direction at FPD 0 (initial core), FPD 40, and FPD 100 respectively. The main feature to be noted is the initial “dishing” of the power distribution in rows containing some channels with depleted fuel. This dishing is quite pronounced at FPD 0, but flattens out with increasing fuel burnup, since the inner core actually goes through its plutonium peak earlier than the outer core, and the depleted-fuel plutonium peak is in fact more pronounced, as pointed out earlier.

Beyond the plutonium peak, the overall power distribution flattens out, and the maximum bundle power drops.

7.5 Period from Onset of Refuelling up to Equilibrium

When refuelling begins, the inner core region has the highest burnup and the lowest power relative to the equilibrium power distribution. Refuelling begins in this region, causing power to rise both because of the addition of fresh fuel and the simultaneous discharge of irradiated and depleted fuel. Only some of the channels in this region can be refuelled, however, otherwise the power would rise excessively.

Refuelling of outer-region channels follows. In this region, channel burnup decreases with increasing distance from the core centre (i.e., decreasing power). Therefore, the refuelling tends to proceed generally from the central core region towards the periphery. However, not all channels at a given radius can be refuelled at the same time. Some channels in each ring are initially bypassed for two reasons: first, it is desirable not to refuel adjacent channels simultaneously, because this would cause a local power peak; and second, it is desirable to have a distribution of burnup in each ring when

equilibrium is reached. Channels missed on the first refuelling of a ring will be refuelled later, until, when the last channels are visited, the burnup in each ring is uniformly distributed between zero and discharge value. Note that this means that the channel with the highest burnup is not always the one which is refuelled.

After refuelling begins, the rate of refuelling rapidly approaches its equilibrium value (approximately 16 bundles per FPD for the CANDU 6). Over short periods, there will be a considerable variation from this average rate. For example, if, for a few days, only the outermost channels were being refuelled, a very high rate (almost 6 channels per full-power day for the CANDU 6) would be required to keep up with the reactivity loss, since these channels would introduce somewhat less reactivity than central channels. It is not possible for the fuelling machines to maintain this rate, and therefore refuelling of outermost channels would have to be intermingled with the refuelling of high-worth, inner-region channels or postponed to some convenient time. Fortunately, this region does not have to be visited very often - about 4 or 5 times per month on the average.

Figure 7.7 shows a plot of refuelling rate vs. time, obtained from a simulation of the CANDU 6 reactor from the onset of refuelling to FPD 400. The corresponding maximum channel powers and bundle powers are shown in Figures 7.8 and 7.9.

Figures 7.10 and 7.11 show the gradual increase in average discharge burnup and in cumulative discharge burnup through this period of operation.

The last figures also show that it is possible to increase the discharge burnup of the initial fuel load, and reduce the channel-visit rate, by adopting the swing-8 refuelling for the **first** visit (only) of each channel. The swing-8 scheme, shown in Figure 7.12, features the reshuffling of fuel (bundles 11 and 12) back into the channel on refuelling, and is used to avoid the very low burnup of the original bundles 11 and 12 if they remained in core for only one cycle.

7.6 Consequences of Fuelling-Machine Unavailability

If refuelling were to stop, core reactivity would continuously decrease. The rate of reactivity decay is about 0.4 mk/FPD in the CANDU-6 core. The reactor regulating system (RRS) would of course attempt to maintain criticality.

The first action that the reactor regulating system (RRS) would take to maintain criticality is to lower the level of water in the liquid zone-control compartments. Eventually, the water will be drained to the lower limit of the control range.

Since the desirable operating range of the zone controllers is between 20% and 70% (i.e., a range of 50%), and since the full reactivity range of the zone controllers (from 100% down to 0%) provides about 7 milli-k of reactivity, the number of days which can be “survived” without refuelling is typically about $3.5 \text{ mk}/(0.4 \text{ mk/FPD})$, i.e., about 8 FPD.

The operator would also ensure that any poison which might exist in the moderator at the time would be removed. Every ppm of boron is worth about 8 mk, however the operating license usually limits the amount of boron in the core in full-power operation to about 0.625 ppm (5 milli-k), so this represents at most about 12 FPD without refuelling.

Continued lack of refuelling would lead to withdrawal of the adjuster rods in their normal sequence. This would permit operation to continue for several weeks. However, as the adjuster rods are withdrawn, the reactor power must be gradually reduced because of changes in the power distribution associated with spatial changes in the distribution of absorption cross section. In effect, withdrawal of the adjusters results in a radially “peaked” power distribution, i.e., higher channel and bundle powers at the center of the core, which forces a power derating in order to remain in compliance with the licensed channel and bundle powers (7.3 MW and 935 kW respectively). The amount of derating necessary increases with the number of adjusters withdrawn.

7.7 Core-Follow Calculations with RFSP

7.7.1 Instantaneous Diffusion Calculations

The main application of RFSP at CANDU sites is in tracking the reactor's operating history. This function is performed with the *SIMULATE module of RFSP.

The core history is tracked by a series of instantaneous snapshots, which can be calculated at any desired frequency. Steps of 2-3 FPD are typically convenient for the site physicist. The code advances the in-core irradiation and burnup distributions at each step, in accordance with the time interval. Individual channel refuellings within a time step are taken into account at the actual time at which they occur.

At each code execution, the zone-control-compartment fills corresponding to the time of the snapshot are input to the code, together with the concentration of moderator poison and any other device movement, so that the instantaneous reactor configuration is captured. As an option, the spatial distribution of ^{135}Xe and its effect on the lattice properties can be modelled in the calculation (see chapter 8); this has an effect on the calculated flux distribution. Bulk and spatial control can also be modelled (see Section 7.7.3).

The presence of in-core detectors in the CANDU 6 allows the validation of the diffusion calculation against actual in-core measurements. The standard deviation of differences between calculated and measured detector fluxes is typically in the range of 2 to 3 %.

The site reactor physicist can also elect to do core tracking using the flux-mapping method in RFSP. In this case, the detector fluxes at the time of the snapshot are input to the code to derive the 3-d mapped flux distribution. This is used to advance the irradiation and burnup distributions from one snapshot to the next. Even in this option,

the diffusion calculation is performed in any case, because results are optimized when the diffusion solution is used as the fundamental mode in the mapping process.

7.7.2 History-Based Methodology for Lattice Properties

Lattice properties are typically calculated by interpolating in irradiation (or, equivalently, in burnup) within “fuel tables” computed by the cell code, assuming core-average values of such parameters as the fuel temperature and the coolant density. In this method, the only independent parameter is the irradiation (or burnup). This conventional methodology is here labelled the “uniform-parameter” method.

However, lattice properties **do** depend on the local values of these parameters, and also on the history of quantities such as the moderator poison concentration.

To take this into account, the “**local-parameter history-based**” methodology has been developed within RFSP for use specifically in core tracking. In this method, fuel tables are **not** employed. Instead, at each core-follow snapshot, an individual POWDERPUFS-V calculation is performed within RFSP for each fuel bundle to update its properties over the incremental burnup step, using *locally appropriate* values of parameters (flux level, fuel temperature, coolant density, whatever parameters the user specifies) for that instant in the core history.

The use of bundle-specific values of parameters other than irradiation gives the method its **local-parameter** label. The **history-based** label originates in the fact that changes in lattice parameters are captured when and only when they actually occur. As a result, the evolution of the nuclear properties of each individual bundle is more properly tracked.

Following the calculation of the lattice properties of bundles, the diffusion equation is solved as usual.

Validation against in-core-detector readings in the CANDU 6, the history-based methodology has resulted in an improvement (reduction) of about 0.2 to 0.5 in the per cent standard deviation of differences between simulation and measurement, compared to the conventional (uniform-parameter) method.

7.7.3 Modelling of Bulk and Spatial Control

In RFSP, the **asymptotic** bulk-control and spatial-control functions of the Reactor Regulating System can be modelled. The label **asymptotic** indicates that the calculation is not in the time domain, but rather that the code attempts to find the long-term **equilibrium** (time-independent) values to which the zone-controller water fills tend.

Bulk and spatial control are modelled by modifying the zone-control-compartment water fills by computed increments every few iterations in the course of the solution of the diffusion equation.

The following subsections describe how the incremental water fills are computed.

7.7.3.1 Bulk Control

Here the zone fills are moved uniformly up or down, depending whether the current value of k_{eff} is higher or lower than the desired value.

That is, the fractional water fill Z_i of each controller i is changed according to

$$Z_i \rightarrow Z_i + \alpha_i (k_{eff} - k_{eff,ref}) \quad i = 1, \dots, N_z$$

where α_i = user-supplied coefficient (usually ~ 140 for the CANDU 6)

7.7.3.2 Spatial Control

In the simplest approximation (selected by setting $ISPCTL = 1$ or 2 in the code), spatial control is modelled by requiring the flux distribution by zone to be proportional to the reference (or target) distribution: in zones where the ratio of flux to target flux is higher (lower) than the average value, the zone fill is increased (decreased). This has the effect of driving the flux distribution towards the target distribution.

That is, the fractional water fill Z_i of each controller i is changed according to

$$Z_i \rightarrow Z_i + \alpha_i \left[\frac{\phi_i}{\phi_{iref}} - \left\langle \frac{\phi}{\phi_{ref}} \right\rangle \right] \quad i = 1, \dots, N_z$$

where :

$$\left\langle \frac{\phi}{\phi_{ref}} \right\rangle = \frac{1}{N} \sum_{i=1}^N \frac{\phi_i}{\phi_{iref}}$$

and :

- ϕ_i = detector (or zone) flux for detector (or zone) i
- ϕ_{iref} = reference detector (or zone) flux for detector (or zone) i
- N_z = total number of zone controllers
- α_i = user-supplied coefficient (usually 2.0)
- N = N_z

The reference or target zone (or detector) flux used in the above equation is normally obtained from a time-average calculation by interpolating in the time-average mesh flux at the detector location or by averaging over the whole zone for the reference zone flux.

In a more sophisticated approximation (selected by $ISPCNTL = 3$), the modelling of spatial control includes “phase-out” factors in the individual “differential lifts” when

the corresponding zone fills reach values that are too high ($> 80\%$) or too low ($< 10\%$). There is also a term which drives each zone fill towards the average zone fill.

8. Effects of ^{135}Xe

The xenon isotope ^{135}Xe plays an important role in any power reactor. It has a very large absorption cross section for thermal neutrons and represents therefore a considerable load on the chain reaction. The ^{135}Xe concentration has an impact on the power distribution, and in turn is affected by the power distribution, by changes in power, and by movements of reactivity devices.

8.1 The Xe-I Kinetics

The $^{135}\text{Xe}/^{135}\text{I}$ kinetics are shown schematically in Figure 8.1.

^{135}Xe is produced to some degree directly in fission, but mostly as the result of the beta decay of its precursor ^{135}I (which has a half-life of 6.585 hours). ^{135}Xe is destroyed in two ways:

- through its own radioactive decay (^{135}Xe has a half-life of 9.169 hours), and by absorption of neutrons to form ^{136}Xe ,
- ^{135}I is a direct product of fission, but can also appear through the radioactive decay chain $^{135}\text{Te} \rightarrow ^{135}\text{Sb} \rightarrow ^{135}\text{I}$. As ^{135}Te and ^{135}Sb have half-lives which are very short (19.0 s and 1.71 s) compared to those of ^{135}I and of ^{135}Xe , it is sufficient to model the decay of ^{135}Te and ^{135}Sb as “instantaneous” and add their yields in fission to that of ^{135}I .

The $^{135}\text{Xe}/^{135}\text{I}$ kinetics in any particular fuel bundle can thus be represented by the following equations:

$$\frac{dI}{dt} = \gamma_i \hat{\Sigma}_f \hat{\phi}_F - \lambda_i I \quad (8.1)$$

$$\frac{dX}{dt} = \gamma_x \hat{\Sigma}_f \hat{\phi}_F + \lambda_i I - \lambda_x X - \hat{\sigma}_x X \hat{\phi}_F \quad (8.2)$$

where

X = average concentration of ^{135}Xe in the bundle in atoms cm^{-3}

I = average concentration of ^{135}I in the bundle in atoms cm^{-3}

γ_x = direct yield of ^{135}Xe per fission (averaged over all fissions)

γ_i = direct yield of ^{135}I in fission, including yields of ^{135}Te and ^{135}Sb (averaged over all fissions)

λ_x = decay constant of ^{135}Xe in s^{-1}

λ_i = decay constant of ^{135}I in s^{-1}

$\hat{\phi}$ = average flux in the fuel in the bundle in $n.cm^{-2}s^{-1}$

$\hat{\Sigma}_f$ = macroscopic fission cross section of the fuel in cm^{-1}

and $\hat{\sigma}_x$ = microscopic ^{135}Xe cross section in cm^2

In the above equations, the term $\gamma_i \hat{\Sigma}_f \hat{\phi}_F$ gives the ^{135}I production rate, while $\lambda_i I$ gives the ^{135}I loss rate (and the production rate of ^{135}Xe by iodine decay). Similarly, the term $\gamma_x \hat{\Sigma}_f \hat{\phi}_F$ gives the production rate of ^{135}Xe due to direct fission, while $\lambda_x X$ gives its decay rate. The term $\hat{\sigma}_x X \hat{\phi}_F$ represents the “destruction” (burnout) rate of ^{135}Xe due to neutron capture. Because of the comparable magnitudes of the various terms, the ^{135}Xe concentration is very sensitive to changes in flux level.

The large absorption cross section of ^{135}Xe plays a significant role in the overall neutron balance in the reactor and thus directly affects the system reactivity, both in steady state and in transients. The $^{135}\text{Xe}/^{135}\text{I}$ kinetics also influences the spatial power distribution in the reactor.

8.2 Reactor Startup

When a reactor is first started, or restarted after a long shutdown, the ^{135}Xe concentration will build up in all fuel bundles according to the equations derived above. In Figure 8.2, the variation of system reactivity as a function of time following startup is given for different final steady-state power levels. It can be seen that it takes ~40 hours for the ^{135}Xe concentration to reach equilibrium.

8.3 Steady-State Xenon Load

At steady state the time derivatives dI/dt and dX/dt are zero. The above equations can then be solved to give the steady state concentrations of ^{135}I and ^{135}Xe (I_{ss} and X_{ss}):

$$I_{ss} = \frac{\gamma_i \hat{\Sigma}_f \hat{\phi}_F}{\lambda_i} \quad (8.3)$$

$$X_{ss} = \frac{(\gamma_i + \gamma_x) \hat{\Sigma}_f \hat{\phi}_F}{\lambda_x + \hat{\sigma}_x \hat{\phi}_F} \quad (8.4)$$

It is obvious from these equations that, as a function of an increasing fuel flux $\hat{\phi}_F$, the steady-state concentration of ^{135}I increases indefinitely, while in contrast that of ^{135}Xe tends to an asymptotic value which will be denoted $X_{ss,\infty}$:

$$X_{ss,\infty} = \frac{(\gamma_i + \gamma_x) \hat{\Sigma}_f}{\hat{\sigma}_x} \quad (8.5)$$

This asymptotic nature of the variation of X_{ss} with $\hat{\phi}_F$ is the reason why ^{135}Xe is termed a “saturating” fission product. (Other saturating fission products are ^{105}Rh , ^{149}Sm , ^{151}Sm , etc.)

The limiting ^{135}Xe absorption rate at very high flux levels leads to a maximum reactivity of ~ -30 mk. In CANDU the equilibrium xenon load is approximately -28 mk. The flux level at full power in CANDU is such that the ^{135}Xe concentration is about 95% saturated, i.e., the average ^{135}Xe concentration is equal to about 95% of the value in an infinite flux.

However, the steady-state ^{135}Xe concentration is not uniform in the core, but varies with flux according to Eq. (8.4). This is discussed in greater detail in Section 8.6 below.

8.4 Effect of Power Changes on Xenon Concentration

Due to the presence of the term $\sigma_x X \hat{\phi}_F$, the variation of the ^{135}Xe concentration with flux is non-linear. The ^{135}Xe reactivity following power (flux) changes will depend on the starting power level, the time at that level, the new power level, and the time spent at the new power level.

Generally speaking, when the power is reduced from a steady level, the ^{135}Xe concentration increases at first. This is due to the fact that ^{135}Xe is still being produced by the decay of ^{135}I , but its burnout rate (by neutron absorption) is decreased because of the reduced neutron flux (reduced power). However, after a certain period (whose length depends on the initial and final power and the rate of power reduction) the ^{135}I decay rate decreases sufficiently (due to the lower fission rate) that the rate of ^{135}Xe production drops below the rate of ^{135}Xe decay (and burnout). At this time, then, the ^{135}Xe concentration reaches a peak value and starts to decrease towards a new (lower) steady-state level.

Conversely, when the power is increased from a steady level, the ^{135}Xe concentration will first decrease and then go through a minimum and start increasing again to a higher steady-state level.

Some typical reactivity variations due to ^{135}Xe following step reductions in power are shown in Figure 8.3. Very similar variations, but in the opposite direction, ensue upon step increases in power. The quantitative effects will be different at different points in the core, due to the initial non-uniform distribution of ^{135}Xe . Thus, for an accurate assessment of xenon transients on the spatial distribution of power, a point-kinetics treatment is generally inadequate and calculations in three dimensions will be required.

8.5 Xenon Transient Following a Shutdown

Following a reactor shutdown, the burnout of ^{135}Xe stops, whereas the production by means of ^{135}I decay continues for several hours. The net result is that there is an initial increase in ^{135}Xe concentration and a decrease in core reactivity. If the reactor is required to be started up shortly after shutdown, extra positive reactivity must be supplied. The ^{135}Xe growth and decay following a shutdown for a typical CANDU is shown in Figure 8.4.

It can be seen from this figure that, at about 10 hours after shutdown, the reactivity worth of ^{135}Xe increases to several times its equilibrium full-power value. At ~35-40 hours the ^{135}Xe has decayed back to its pre-shutdown level. If it were not possible to add positive reactivity during this period, every shutdown would necessarily last some 40 hours, when the reactor again reaches criticality.

To achieve xenon “override” and permit power recovery following a shutdown (or reduction in reactor power), the CANDU-6 reactor has a set of adjuster rods which are in core during normal operation but which can be withdrawn to provide a source of positive reactivity. It is not possible to provide “complete” xenon override capability, as this requires more than 100 mk of positive reactivity. The CANDU-6 adjuster rods provide approximately 15 milli-k of reactivity, which is sufficient for about 30 minutes of xenon override following a shutdown (i.e., a reduction to zero power).

8.6 Effects of Xenon on Power Distribution

Xenon also plays a role in the power distribution. Because the steady-state ^{135}Xe concentration depends on the flux (Eq. 8.2), high-power bundles will have a higher xenon load, and therefore a lower reactivity, than low-power bundles of the same irradiation. The effect of xenon is therefore to **flatten** the power distribution: the reduction in the maximum bundle power due to the local ^{135}Xe concentration can be of the order of 5%, and should be taken into account when accurate results are desired.

The xenon effect also plays a role after refuelling:

- Fresh bundles introduced into the reactor have no xenon at first. As seen in Fig. 8.2, it takes a day or so for the xenon to build in. Thus, the power of these fresh bundles is high at first and subsequently decreases to an equilibrium value, perhaps several percent lower.
- Bundles which are shifted along the refuelled channel experience a change in power. In these bundles, corresponding to the specific power change (decrease *or* increase), there is a xenon transient similar to those illustrated in Fig. 8.3 (or its “opposite”). This effect is not limited to shifted bundles, but occurs in any bundle whose power is affected by the channel refuelling operation, e.g. bundles in neighbouring channels.

Transient-xenon effects due to refuelling may not always be taken into account in routine core-follow calculations, but need to be captured when it is desired to perform the most accurate simulations. The history-based local-parameter methodology in RFSP

permits the simulation of the transient, spatially non-uniform effects of saturating-fission-product kinetics (both Xe-I kinetics and that of other fission-product pairs, such as Sm-Pm) on the lattice properties of every bundle in core individually.

9. Summary

Fuel management in CANDU has both design and operations aspects.

The design component consists of establishing:

- the desired time-average power distribution for the equilibrium core, which will be used as the target power shape by the site fuelling engineer, and
- the configuration of depleted fuel in the initial core.

The design of the time-average distribution is facilitated by the flexibility in selecting region-specific (or, in the limit, channel-specific) target exit-irradiation values and axial refuelling schemes, allowed by the CANDU on-power-refuelling feature.

The operations component is the responsibility of the site fuelling engineer or reactor physicist. It involves:

- core-follow calculations, typically performed 2 or 3 times per week to keep close track of the in-core flux, power, and burnup distributions and of the discharge burnup of individual bundles,
- the selection of channels for refuelling, based on the current core state, power and burnup distributions and zone-control-compartment water fills, and
- the determination of the CPPF (channel-power-peaking factor) value, used as a calibration factor for the ROP detectors.

The job of the site reactor physicist never gets boring. These tasks keep the job interesting and stimulating.