CHAPTER 21
CANDU In-Core Fuel-Management

prepared by

Dr. Benjamin Rouben, 12 & 1 Consulting, Adjunct Professor,

McMaster University & University of Ontario Institute of Technology (UOIT)

Summary:

This chapter describes the physical and mathematical models and the computational methods used in the management of fuel in CANDU nuclear reactors. The concepts important to the topic of fuel management are explained: fuel irradiation (fluence), fuel burnup, cross-section averaging. The various levels of physics models used in fuel management are presented in quantitative detail: these are the continuous refuelling model, the time-average model, patterned-age snapshot models, and core-follow models. The typical reactivity curve for the CANDU lattice is presented, and a method to estimate the average attainable fuel burnup from the reactivity curve is explained. The CANDU initial core and the first few months of operation are discussed. Finally, the chapter gives a short qualitative survey of advanced fuel cycles which could be exploited in CANDU reactors.

Table of Contents

List of Figures

List of Tables

1 Introduction

1.1 Overview

CANDU Refuelling is carried out with the reactor at power. This feature makes the in-core fuel management substantially different from that in reactors that must be refuelled during shut-downs. The capability for on-power refuelling means that excess reactivity requirements are at a minimum: only a few milli-k are necessary for continuous and short-term reactivity control. This leads to excellent neutron economy and low fuelling costs.

The primary objective of CANDU in-core 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 in-core fuel management are as follows:

- Adjust the refuelling rate to maintain the reactor critical
- Control the core power to satisfy safety and operational limits on fuel and channel power, thus ensuring that the reactor can be operated at full rating
- Maximize burnup within operational constraints, to minimize fuelling cost
- Avoid fuel defects, to minimize replacement fuel costs and radiological occupational hazards.

To refuel a channel, a pair of fuelling machines latches onto the ends of the channel. A separate chapter in this book describes the fuelling machine and the related fuel handling [Damario 2016]. An even 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 next-neighbour channels. In the CANDU-6 reactor, the refuelling direction is the same as that of coolant flow in the channel.

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;
- A channel can be refuelled without delay if failed fuel exists or is suspected;
- The fuel can be removed completely from a few channels, in order for example to limit the channel axial growth.

This chapter will cover the topics relevant to the management of nuclear fuel in CANDU reactors. Fuel management in CANDU has both design and operations aspects; both are covered in this chapter. The concepts important to in-core fuel management will be explained: fuel irradiation (fluence), fuel burnup, cross-section averaging. The various levels of physics models which are used to carry out fuel-management calculations will be presented: these are the continuous-refuelling model, the time-average model, patterned-age snapshot models, and core-follow

models. The purpose and use of the various calculations will be explained. Additional information can be obtained from [Rouben 2003].

1.2 Learning Outcomes

The goal of this chapter is to help the student understand:

- The basic concepts of fuel irradiation (fluence), fuel burnup, and fuel isotopic changes
- The reactivity curve of the fuel lattice and its importance
- The refuelling process in CANDU
- The significance of the flux and power shapes in reactor
- How to control the flux shape, the means to flatten the flux distribution (adjuster rods, differential fuelling)
- Time-average, snapshot, and core-follow models for CANDU reactors, and their use.

2 Definition of Concepts

Fuel irradiation (also known as fluence) and fuel burnup are separate but related concepts. Both are in a way a measure of the "age" of the fuel in the reactor.

2.1 Fuel Irradiation

Let us use the symbol $\Phi_F(t)$ for the neutron flux in the fuel in a certain fuel bundle, where *t* is the time measured from the moment when the bundle entered the reactor core. The irradiation (or fluence) of the fuel in the bundle, denoted $\omega(t)$, is defined as the time integral of the flux to time t:

$$\omega(t) = \int_0^t \Phi_F(t)\,dt \tag{1}$$

This definition is not quite complete, since the flux is a function of neutron energy *E*. We must remove the vagueness and make the definition complete by specifying which flux is to be used. The usual choice is to select Φ_F as the **thermal** flux in the fuel. Excluding for now the time variable, we can write Φ_F as the integral of the flux as a function of energy:

$$\Phi_F = \int_0^{E_{th}} \Phi_F(E)\,dE \tag{2}$$

where E_{th} is the upper bound of the thermal energy interval. Note: Since the thermal energy interval may be defined differently in different computer codes, the fuel irradiation may vary from code to code, and caution must therefore be exercised when comparing irradiation values using different codes.

Having settled on the definition of fuel flux, let us return now to Eq. (1) and note that fuel irradiation (fluence) is a monotonically increasing function of time, since flux is never negative. The irradiation starts at 0 when the fuel bundle enters the core, and increases up to the time when the fuel bundle exits the core. The value of the irradiation when the bundle exits from the core is then called the **exit** or **discharge** irradiation.

The units of irradiation are units of flux multiplied by units of time. This gives for example $n.cm^{-2}.s^{-1}*s$, i.e., $n.cm^{-2}$. But a much more appropriate (microscopic) unit for area is the barn ($b = 10^{-24}$ cm^2) or kilobarn ($kb = 10^{-21}$ cm^2). Thus, the common unit for irradiation is n/kb.

2.2 Fuel Burnup

The fuel burnup of a given fuel bundle, denoted *B(t)*, is defined as the total **fission energy** (note: not thermal or electric energy) released in the fuel bundle since the moment when it entered the reactor core, divided by the initial mass of heavy element in the bundle. (In the case of the CANDU reactor, where the only initial heavy element in the fuel is uranium, the initial mass of heavy element is the mass of uranium in the bundle.)

Since the energy release depends on the fission rate in the bundle, we can write

$$B(t) = \frac{\int_0^t E_R \Sigma_f(t) \Phi_f(t) dt}{M} \tag{3}$$

where *M* is the mass of uranium in the bundle, Σ_f is the fission cross section in the bundle, and E_R is the fission energy release per fission (~200 MeV).

In contrast to that of fuel irradiation, the definition of fuel burnup does not suffer from vagueness, since energy is a measurable, well-defined quantity. For this reason, if comparisons are to be made between different computer codes, it is much easier, and preferred, to compare on the basis of fuel burnup rather than of irradiation.

Similarly to fuel irradiation, fuel burnup is a monotonically increasing quantity, since the fuel continues to fission and release energy with time. The burnup of the fuel in the fuel bundle starts at 0 when the fuel bundle enters the core, and increases up to the time when the fuel bundle exits the core. The value of the fuel burnup when the bundle exits from the core is then called the **exit** or **discharge** burnup.

The units of fuel burnup are units of energy divided by units of mass. The most commonly used burnup units are megawatt-days per megagram of uranium (MW.d/Mg(U)) or megawatt-hours per kilogram of uranium (MW.h/kg(U)). The typical average exit discharge burnup in existing CANDU reactors varies depending on the design of the core, and may range between 180 and 225 MW.h/kg(U) (or, equivalently, between 7,500 and 9,375 MW.d/MgU).

2.3 Relationship Between Fuel Irradiation and Fuel Burnup

Both fuel irradiation and fuel burnup increase with time and reflect therefore in a way the "age" of the fuel. Of course, fuel "ages" faster in a high neutron flux, since the flux appears in the integrand of both Eqs. (1) and (2).

In a given nuclear lattice, fuel burnup and fuel irradiation have a one-to-one relationship which depends on the fission cross section Σ_f and its variation with time. In the CANDU lattice, there is a relatively small variation in Σ_f (and in E_R) over the typical residence time of a fuel bundle in core, so the relationship between irradiation and burnup is almost linear. For example, the old lattice code POWDERPUFS-V [Rouben 1995] gives the following relationship for fuel burnup as a function of irradiation:

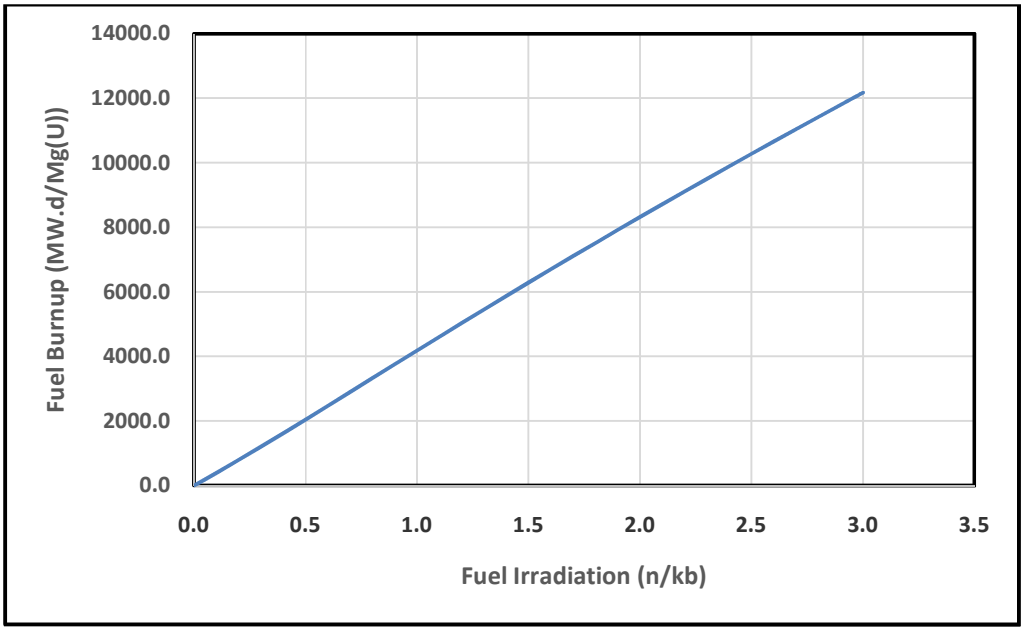

Figure 1 Relationship of Fuel Burnup and Irradiation from POWDERPUFS-V Code

3 CANDU Reactor Physics Computational Scheme

To properly manage the nuclear fuel in the core, one must be able to calculate the neutron flux and power distributions in space and in time. The most accurate computational tool for CANDU neutronics is the Neutron Transport equation (see Section 2.1 of Chapter 4, Reactor Statics). However, solving this equation over the full core on a routine basis demands too much computational effort. Therefore, a 3-stage process involving the Neutron Transport equation and the simpler Neutron Diffusion equation, which involves some approximation (see Section 2.2. of Chapter 4), has been developed for CANDU neutronics. By experience, this 3-stage computational system has been found to give generally very good results. It has been described in Chapter 4 and is reviewed here.

3.1 Lattice Calculation

The first stage involves solving for the flux distribution in the fundamental unit of the CANDU reactor, i.e., the basic lattice cell, which is defined to consist of the fuel, coolant, pressure and calandria tubes, and moderator, but no reactivity devices (see face view in Figure 2; in the third dimension the basic lattice cell is one bundle-length deep).

The lattice calculation is performed using a lattice code. For many years, the lattice calculation was performed with the code POWDERPUFS-V, which was developed specifically for CANDU reactors and which makes use of computational “recipes” to incorporate experimental results obtained with heavy-water lattices. More recently, deterministic multi-group neutron transport codes with a strong mathematical and theoretical basis, such as WIMS-IST [Irish 2002] or DRAGON [Marleau 1999], have been adopted in place of POWDERPUFS-V for lattice calculations. The DRAGON code is currently freely available from École Polytechnique de Montréal.

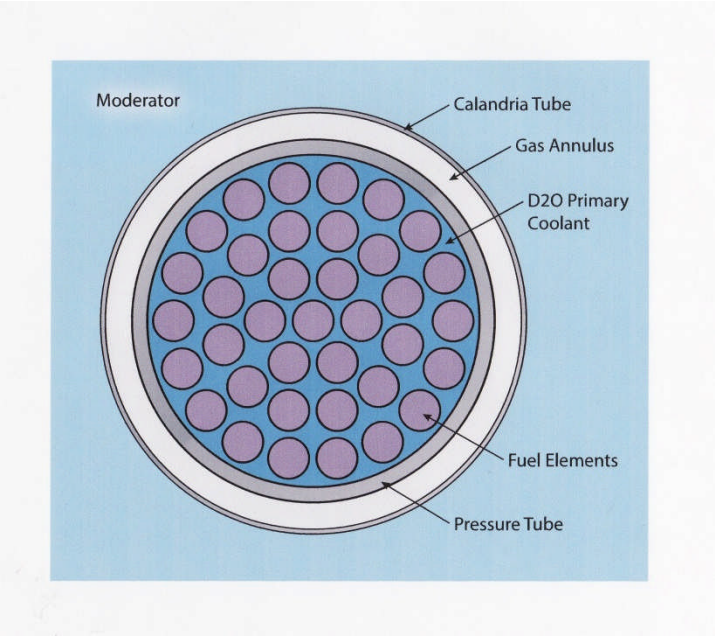

Figure 2 Face View of the CANDU Basic Lattice Cell

The deterministic transport code uses the geometry and properties of materials in the lattice cell to do a very accurate calculation of the flux distribution over the lattice cell. It uses a large or very large number (tens or even hundreds) of neutron energy groups in this calculation, to take account of the neutron moderation accurately. It then averages the nuclear properties (for absorption, moderation, fission, etc.) over the cell according to the calculated reaction rates in each sub-region of the lattice cell to determine the effective neutronic properties of the whole cell. This averaging (homogenization) of the properties of the basic lattice cell dilutes the strong absorption of the fuel with the much weaker absorption in the moderator. The code also "collapses" the cell properties over a very small number of neutron energy groups. Most often in CANDU neutronics, the collapsing is done over 2 energy groups (a thermal group and a slowing-down group). This is sufficient since only about 2% of all fissions are induced by non-thermal neutrons. With these homogenized properties, Fick's Law (see Section 2.2 of Chapter 4) becomes a good approximation over most of the reactor, and the diffusion equation can be used to calculate the neutron flux over the full core.

The lattice code must do one more thing, a "depletion" calculation: Here the code updates the composition of the fuel over a time step, using the various reaction rates that it has determined, and then recalculate the homogenized cell properties. Repeating this over many time steps, the code evolve the lattice-cell neutronic properties to reflect the changes in the nuclide composition of the fuel with time (irradiation/burnup).

The output of the lattice code is then a table of the basic-cell (lattice) nuclear properties as functions of fuel irradiation or burnup. For example, POWDERPUFS-V gives the properties in the 2-energy-group formalism, but in the approximation that the fast-fission and upscattering cross sections are assumed to be 0. A typical fuel table would then appear as in Table 1. Note that the lattice code must also in addition provide the nuclear properties (cross sections) of the

moderator, to be applied outside the fuelled region of the core.

Table 1 Sample Lattice Nuclear Cross Sections versus Fuel Irradiation

Irradiation ω(n/kb)	Fast Diffusion Coefficient D1 (cm)	Thermal Diffusion Coefficient D2 (cm)	Fast Absorption Cross Section Σa1 (cm^-1)	Thermal Absorption Cross Section Σa2 (cm^-1)	Yield Cross Section νΣϕ2 (cm^-1)	Downscattering Cross Section νΣϕ2 (cm^-1)	Ratio of Fuel Flux to Cell Flux
0.00	1.269	0.936	7.68E-04	3.79E-03	4.52E-03	7.40E-03	0.383
0.10	1.269	0.936	7.68E-04	3.82E-03	4.56E-03	7.40E-03	0.380
0.20	1.269	0.936	7.68E-04	3.89E-03	4.66E-03	7.40E-03	0.374
0.30	1.269	0.937	7.67E-04	3.95E-03	4.73E-03	7.40E-03	0.369
0.40	1.269	0.937	7.67E-04	3.99E-03	4.78E-03	7.40E-03	0.365
0.50	1.269	0.937	7.67E-04	4.03E-03	4.81E-03	7.40E-03	0.361
0.60	1.269	0.937	7.66E-04	4.07E-03	4.83E-03	7.40E-03	0.359
0.70	1.269	0.937	7.66E-04	4.10E-03	4.83E-03	7.40E-03	0.356
0.80	1.269	0.937	7.66E-04	4.12E-03	4.83E-03	7.40E-03	0.354
0.90	1.269	0.937	7.66E-04	4.14E-03	4.82E-03	7.40E-03	0.352
1.00	1.269	0.937	7.65E-04	4.15E-03	4.81E-03	7.40E-03	0.351
1.10	1.269	0.937	7.65E-04	4.17E-03	4.78E-03	7.40E-03	0.350
1.20	1.269	0.937	7.65E-04	4.18E-03	4.76E-03	7.40E-03	0.349
1.30	1.269	0.937	7.65E-04	4.19E-03	4.74E-03	7.41E-03	0.348
1.40	1.269	0.936	7.64E-04	4.20E-03	4.71E-03	7.41E-03	0.348
1.50	1.269	0.936	7.64E-04	4.20E-03	4.68E-03	7.41E-03	0.347
1.60	1.269	0.936	7.64E-04	4.21E-03	4.65E-03	7.41E-03	0.347
1.70	1.269	0.936	7.63E-04	4.22E-03	4.62E-03	7.41E-03	0.346
1.80	1.269	0.936	7.63E-04	4.22E-03	4.59E-03	7.41E-03	0.346
1.90	1.269	0.936	7.63E-04	4.22E-03	4.56E-03	7.41E-03	0.346
2.00	1.269	0.936	7.63E-04	4.23E-03	4.53E-03	7.41E-03	0.345
2.10	1.269	0.936	7.62E-04	4.23E-03	4.50E-03	7.41E-03	0.345
2.20	1.269	0.936	7.62E-04	4.23E-03	4.47E-03	7.41E-03	0.345
2.30	1.269	0.936	7.62E-04	4.24E-03	4.45E-03	7.41E-03	0.345
2.40	1.269	0.936	7.62E-04	4.24E-03	4.42E-03	7.41E-03	0.344
2.50	1.269	0.936	7.61E-04	4.24E-03	4.39E-03	7.41E-03	0.344
2.60	1.269	0.936	7.61E-04	4.25E-03	4.37E-03	7.41E-03	0.344
2.70	1.269	0.936	7.61E-04	4.25E-03	4.35E-03	7.41E-03	0.344
2.80	1.269	0.936	7.60E-04	4.25E-03	4.33E-03	7.41E-03	0.343
2.90	1.269	0.936	7.60E-04	4.26E-03	4.30E-03	7.41E-03	0.343
3.00	1.269	0.936	7.60E-04	4.26E-03	4.28E-03	7.41E-03	0.343
3.10	1.269	0.936	7.60E-04	4.26E-03	4.27E-03	7.41E-03	0.343
3.20	1.269	0.936	7.59E-04	4.26E-03	4.25E-03	7.41E-03	0.342
3.30	1.269	0.936	7.59E-04	4.27E-03	4.23E-03	7.41E-03	0.342
3.40	1.269	0.936	7.59E-04	4.27E-03	4.21E-03	7.41E-03	0.342
3.50	1.269	0.936	7.59E-04	4.27E-03	4.20E-03	7.41E-03	0.342
3.60	1.269	0.936	7.58E-04	4.28E-03	4.18E-03	7.41E-03	0.341
3.70	1.269	0.936	7.58E-04	4.28E-03	4.17E-03	7.41E-03	0.341
3.80	1.269	0.936	7.58E-04	4.29E-03	4.16E-03	7.41E-03	0.341
3.90	1.269	0.936	7.58E-04	4.29E-03	4.15E-03	7.41E-03	0.340
4.00	1.269	0.936	7.57E-04	4.29E-03	4.13E-03	7.41E-03	0.340

3.2 Reactivity-Device Calculation

The CANDU basic lattice cell contains, by definition, no reactivity devices. Neither does it contain any structural material other than the fuel channel itself (i.e., the pressure tube and calandria tube). However, there are reactivity devices in appropriate locations in the core. Therefore, the effect of the reactivity devices, their guide tubes and other structural material on the nuclear properties in their vicinity needs to be taken into account. This effect is expressed in the form of homogenized "incremental" cross sections which are added to the (unperturbed) homogenized properties of neighbouring lattice cells in a modelled volume around the device, called a "supercell"; see an example in Figure 3, although the supercell is often defined as having half the width shown, i.e., with only half the width on either side of the device. The incremental cross sections are calculated with a transport code, most often DRAGON; they are identified with the differences in the supercell homogenized properties between the case where the device is inserted into the supercell and the case where it is withdrawn from it. The incremental cross sections are added to the lattice-cell cross sections in the modelled volume around the device's position in the reactor core, when that particular device is inserted.

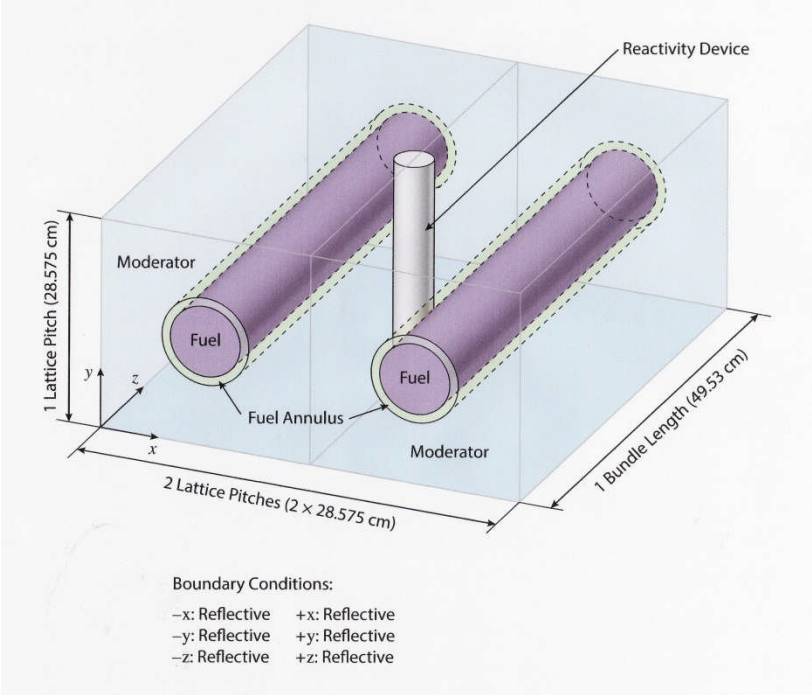

Figure 3 Supercell for Calculation of Device Incremental Cross Sections

3.3 Full-Core Calculation

In the final stage, a full-core reactor model is assembled. As explained above, because homogenized cell and device properties are used, few-group (most often 2-group) diffusion theory can be applied with this model. The full-core model must incorporate, for each cell in the reactor, the homogenized unperturbed-lattice-cell cross sections which apply there, together with the incremental cross sections of nearby reactivity devices in their appropriate locations. The model therefore consists of homogenized basic-lattice-cell properties in the appropriate locations, on which are superimposed reactivity-device properties in the devices' modelled

volumes.

A typical full-core 3-d CANDU reactor model for use with finite-difference diffusion theory is shown in Figure 4 (face view, showing adjuster modelled volumes) and Figure 5 (top view, showing adjuster and liquid-zone-controller modelled volumes).

The full-core flux and power distributions are calculated in CANDU with the finite-difference diffusion-theory code Reactor Fuelling Simulation Program RFSP-IST [Rouben 2002], or for example with DONJON [Hébert 2012]. The DONJON code is currently freely available from École Polytechnique de Montréal.

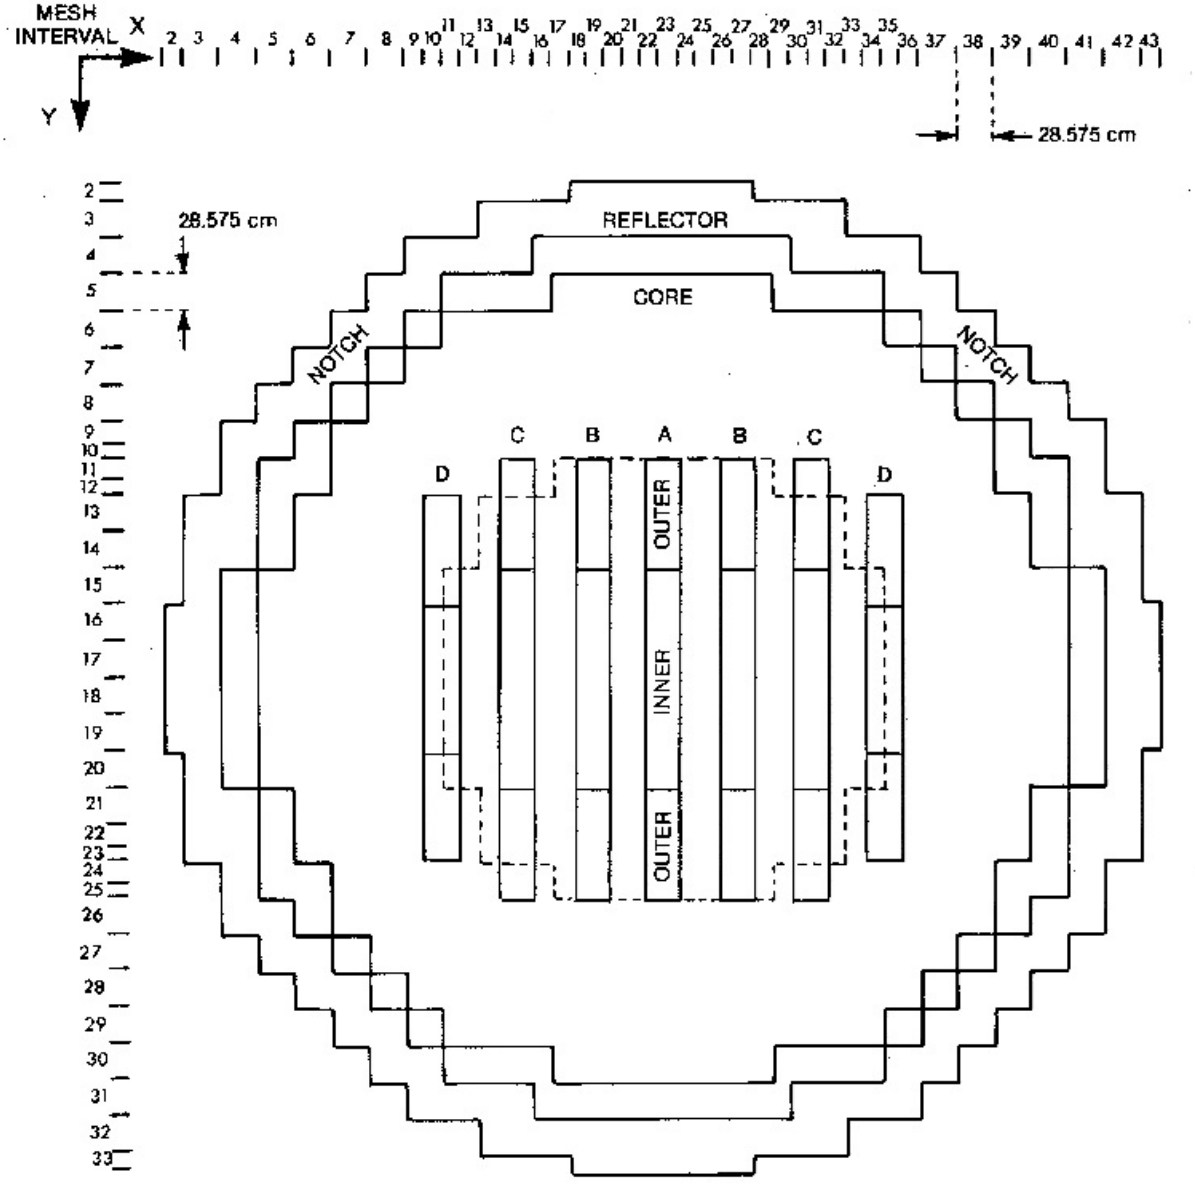

Figure 4 Face View of a Typical CANDU-6 Full Reactor Model

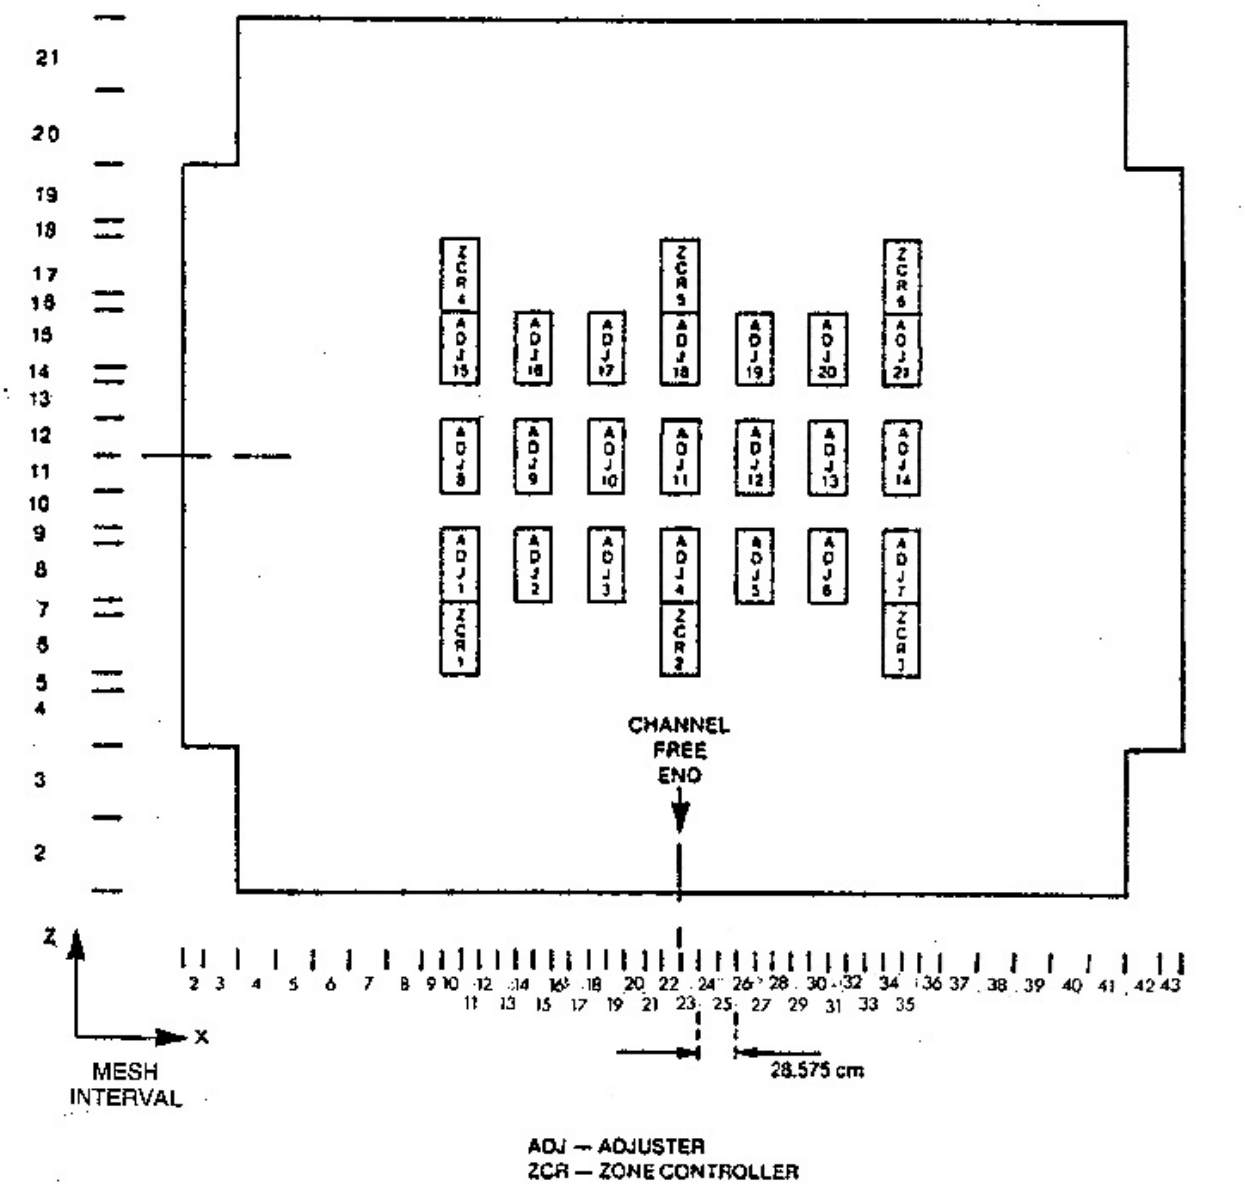

Figure 5 Top View of a Typical CANDU-6 Reactor Physics Model

4 Averaging Nuclear Cross Sections

There are many situations in which one wishes to average nuclear cross sections in space or in time. For example, we have already seen in Section 2.1 that the lattice code averages (homogenizes) the nuclear properties within the basic lattice cell. In the finite reactor also, we may wish to average the properties of a region of the core. How to do this is explained here.

The overarching aim in reactor physics is to calculate reaction rates, because from a knowledge of reaction rates we can derive the values of the important quantities in the system, for example the core reactivity or the power distribution in the core, etc.

A reaction rate *RR* at a point, corresponding to a process with cross section Σ, is calculated as the product of the cross section with the neutron flux:

$$RR = \Sigma\Phi \tag{4}$$

These point reaction rates can then be integrated over space or time. Since the total reaction rate is the quantity of interest, then we must ensure that this is conserved in any averaging.

4.1 Averaging Over Space

Suppose we consider a non-uniform reactor subdivided into a number of regions with different values of the cross section, and we would like to know the "average" cross section for the whole reactor. Let's take a very simple example of a reactor with 2 regions, with uniform cross section value Σ_1 and Σ_2 respectively (see Figure 6).

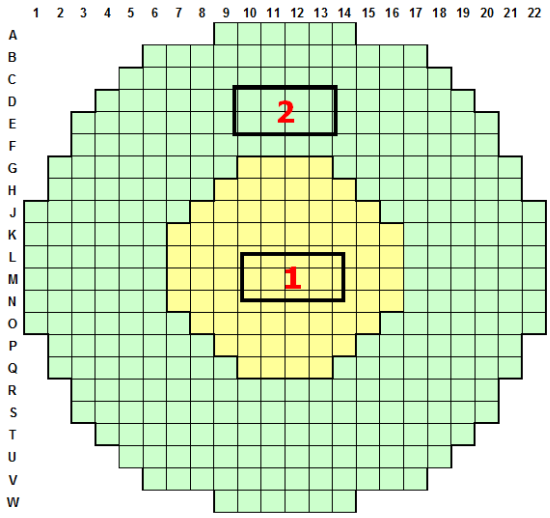

Figure 6 Averaging the properties of a 2-region core

The flux is Φ_1 and Φ_2 in the two regions, and is a function of space in each region. How would we calculate the overall average cross section $\overline{\Sigma}$?

We want to conserve the total reaction rate, which is

$$RR_{total} = \int_{V_1} \Sigma_1\Phi_1 dV + \int_{V_2} \Sigma_2\Phi_2 dV \tag{5}$$

Therefore $\overline{\Sigma}$ should satisfy

$$\int_{V_1} \overline{\Sigma}\Phi_1 dV + \int_{V_2} \overline{\Sigma}\Phi_2 dV = RR_{total}$$

which gives

$$\overline{\Sigma} = \frac{\int_{V_1} \Sigma_1\Phi_1 dV + \int_{V_2} \Sigma_2\Phi_2 dV}{\int_{V_1} \Phi_1 dV + \int_{V_2} \Phi_2 dV} \tag{6}$$

This shows that the proper definition of the average cross section is obtained by averaging the cross section with flux and volume weighting.

If we define also the average flux

$$\bar{\Phi} = \frac{\int_{V_1} \Phi_1 dV + \int_{V_2} \Phi_2 dV}{\int_{V_1} dV + \int_{V_2} dV} = \frac{\int_{V_1} \Phi_1 dV + \int_{V_2} \Phi_2 dV}{V} \tag{7}$$

where $V = V_1 + V_2$ is the total volume, then we can also write

$$\bar{\Sigma} = \frac{\int_{V_1} \Sigma_1 \Phi_1 dV + \int_{V_2} \Sigma_2 \Phi_2 dV}{\bar{\Phi} V} \tag{8}$$

or

$$\bar{\Sigma}\bar{\Phi}V = \int_{V_1} \Sigma_1 \Phi_1 dV + \int_{V_2} \Sigma_2 \Phi_2 dV = RR_{total} \tag{9}$$

4.2 Averaging Over Time

Now let's think about averaging over time at a given point or region (say, for a fuel bundle). Here the properties will vary with time, i.e., with fuel irradiation or burnup. The reaction rate which we must conserve is the time-integrated reaction rate $\int \Sigma(t)\Phi(t)dt$.

The average cross section over a time interval t_1 to t_2 is then defined by

$$\bar{\Sigma} = \frac{\int_{t_1}^{t_2} \Sigma(t)\Phi(t)dt}{\int_{t_1}^{t_2} \Phi(t)dt} \tag{10}$$

Now let's change variables to ω, using $d\omega = \Phi dt$, and we get

$$\bar{\Sigma} = \frac{\int_{\omega_1}^{\omega_2} \Sigma(\omega)d\omega}{\int_{\omega_1}^{\omega_2} d\omega} = \frac{\int_{\omega_1}^{\omega_2} \Sigma(\omega)d\omega}{\omega_2 - \omega_1} \tag{11}$$

i.e., $\bar{\Sigma}$ is basically the average of the cross section over irradiation on the irradiation interval.

5 The Fuel Reactivity Curve

The reactivity of the fuel lattice versus irradiation (or burnup) is provided by the lattice code, together with the lattice properties. For any lattice design, the fuel reactivity curve is a fundamental piece of knowledge for fuel management. A typical reactivity curve for the CANDU lattice is shown below (Figure 7):

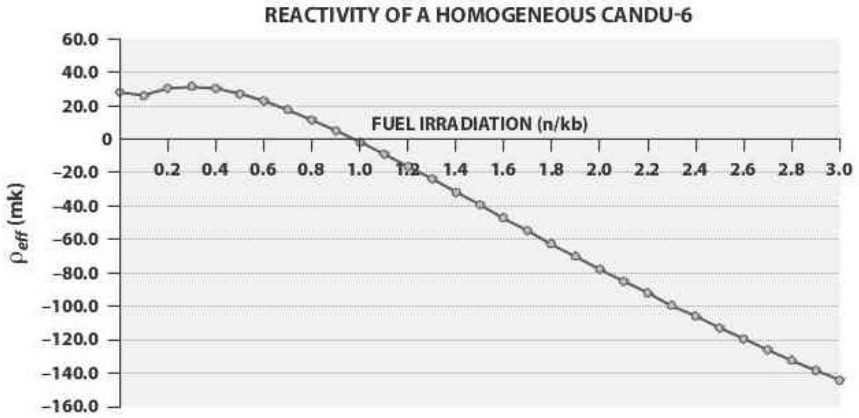

Figure 7 The Reactivity of the CANDU Lattice

The reactivity curve has several interesting features:

- The reactivity at zero irradiation (fresh fuel) is quite high, ~76 mk. It is to be noted that this value is calculated assuming that the saturating fission products (xenon, samarium, etc.) are already at equilibrium, even though the fuel is fresh. To run the infinite lattice in steady state, the positive reactivity would have to be cancelled by adding poison to the moderator, for example ~9.5 ppm of boron in the initial core (assuming a boron reactivity coefficient of -8 mk/ppm).
- The reactivity reduces slowly at first, then increases to a maximum, called the plutonium peak. This behavior is due to the combined effects of the start of U-235 depletion and the start of Pu-239 production. The latter is a slightly better fuel than U-235, and consequently increases reactivity as it is produced, following a short time lag due to the ~2.5-day half-life of Np-239, the immediate parent of Pu-239.
- Following the plutonium peak, the reactivity decreases in a monotonic fashion indefinitely. This is due to the continued net depletion of fissile nuclides and accumulation of fission products. The plutonium peak, however, does not mark the maximum in plutonium concentration. In fact the latter keeps increasing, however it does so more and more slowly, since Pu-239 participates in fission at the same time as it is produced.
- Hypothetically, if the infinite lattice could be run in steady state as the fuel gained irradiation, the boron concentration would have to be adjusted continuously to reflect the positive reactivity.
- In Figure 6, the reactivity reaches a value of zero at an irradiation of ~1.6 n/kb. This would mark the exit irradiation of the infinite lattice. At this point the boron concentration required would of course be 0.

6 Exit-Burnup Estimation from Lattice Calculations

The fuel-reactivity curve above gave us the discharge irradiation attainable for the infinite lattice. Even before doing fuel-management calculations for the full core, we can devise a way to make a fairly good estimate of the average attainable discharge irradiation for the real CANDU reactor, if we have an estimate of the leakage and in-reactor-device worth.

For the CANDU-6 reactor, good estimates of the leakage and of the device worth (adjusters and zone-controller water) are 30 mk and 18 mk respectively. If we subtract a total of 48 mk from the fuel-reactivity curve, we get a curve (Figure 8) which the full reactor would follow if all parts of it advanced in irradiation at the same rate. While this is obviously not exactly the case, we will still get a fairly good approximation of the reactivity curve for the full reactor as per Figure 8, which shows that the CANDU 6 core would initially reach 0 ppm of boron at an irradiation of about 1.0 n/kb (a burnup of ~4,000 MW.d/Mg(U)). This would be the exit irradiation if the CANDU were refuelled in batch mode. [Note that in practical applications, the excess reactivity of the fresh core is reduced with the use of depleted fuel, and the refuelling is started gradually before the excess reactivity of the core reaches decreases to 0.]

However, CANDU can be - and is - refuelled on line. This means that when the moderator poison runs out, daily refuelling can be started, replacing just enough fuel each day to cancel the loss of reactivity due to the increased overall irradiation.

By repeating this day after day, we eventually reach a state where there is approximately an equal amount of fuel at all irradiations, from 0 to a maximum value which determines the attainable discharge irradiation with daily refuelling. This value is obtained by doing a “rolling average” of the reactivity starting from an irradiation of 0, up to the point where the rolling average is 0. It is clear that the 0 rolling average is obtained when the positive area to the left of the point where the reactivity is 0 is exactly cancelled by the negative area beyond that point. By doing the rolling average numerically, we find that the average attainable exit irradiation is ~1.8 n/kb (burnup ~7,500 MW.d/Mg(U)) – see Figure 8. Thus daily refuelling has almost doubled the average exit burnup, compared to batch refuelling!

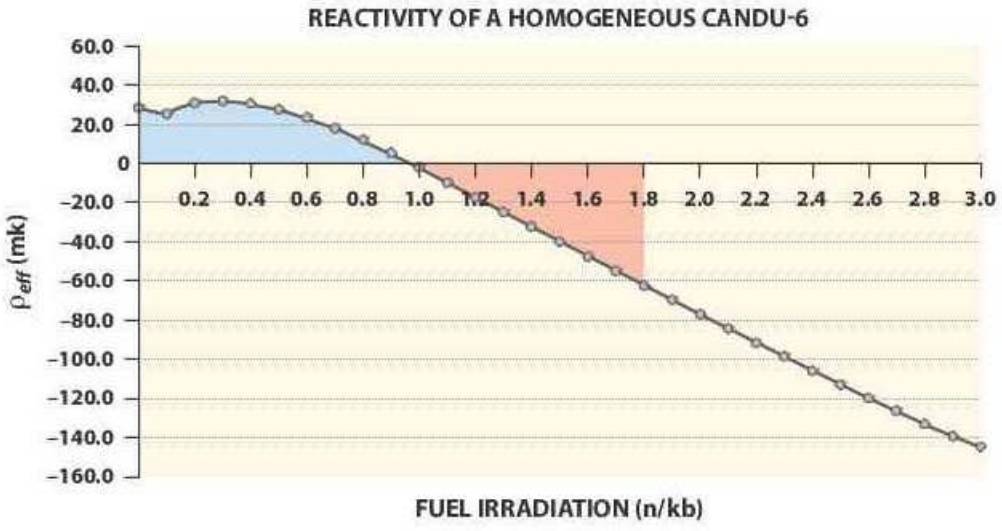

Figure 8 Average Exit Irradiation Attainable with Daily Refuelling

Note of course that all the above analysis, in the previous and current sections, depend on the specific values assumed for various lattice parameters, such as heavy-water purity, fuel and moderator temperatures, etc. A higher value of moderator purity, most importantly, can raise the reactivity curve and consequently increase the attainable exit burnup: a 0.1% increase in moderator purity is estimated to increase the exit burnup by ~360 MW.d/Mg(U). Values around 99.9% for the moderator isotopic purity are common.

7 Flattening the Power Distribution

The reactivity curves in the previous sections provide information on the "local" reactivity in bundles according to their irradiation level. And we have up to now used these figures in the approximation of uniform irradiation throughout the core.

However, of course, fuel in different locations in core actually picks up irradiation at different rates, depending on the local flux. Chapter 4 of this book has shown that the flux distribution in a uniform reactor has a J_0 (Bessel function) shape and a cosine shape in the radial and axial directions respectively. CANDU reactors are not homogeneous, and so the flux distribution will be different. In addition, we may **deliberately** want to change the power distribution.

One reason we may want to change the power distribution is that the J_0 and cosine shapes are "peaked" flux shapes. This means that the power decreases relatively quickly from the maximum value. Now, a peaked flux shape may create a conflict between the reactor power and peak local powers (channel or bundle power). The plant's operating license typically puts limits not only on the total reactor power, but also on the maximum channel and bundle powers. With these restrictions, it may not be possible with a peaked flux shape to generate the reactor power desired without exceeding the licensed maximum channel and bundle powers. For this reason, most power reactors in fact operate with a flattened power distribution, in which the power is redistributed so that the peak is reduced and more channels (radially) and bundles (axially) generate more power. In this way, the desired total power can be generated while

complying with the peak local limits.

In CANDU reactors, power flattening is done in two ways:

1) By placing some absorbing rods in the central area of the core. In CANDU reactors, these are the adjuster rods.
2) By differential fuelling. This means deliberately running the fuel in the central region of the core to a higher discharge irradiation than fuel in the peripheral region. The higher irradiation reduces the local reactivity and therefore the local flux.

Differential fuelling of a core region (it could even be a single channel) in CANDU can also be thought of as being produced by varying the region's refuelling frequency. This can be seen from the following relationship between exit burnup and refuelling rate.

Let P be a region's fission power, B be the average exit burnup in the region, and r_m be the refuelling rate in the region (in mass of uranium per unit time). It is easy then to derive the relationship

$$B = \frac{P}{r_m}, \; also \; written \; as \; r_m = \frac{P}{B} \tag{12}$$

If for example P is given in MW and r_m in MW(U)/d, then the units of B are MW.d/Mg(U). If a CANDU bundle contains a mass M of uranium, then we can write the above relationship in terms of the bundle refuelling rate, r_b:

$$B = \frac{P}{Mr_b} \; or \; r_b = \frac{P}{BM} \tag{13}$$

The most convenient units for r_b are bundles/d, obtained for example from P in MW, B in MW.d/Mg(U), and M in Mg(U)/bundle.

From this analysis we can see that by controlling the refuelling rate one can control the average exit burnup. In fact, in this way we can control the reactivity of the fuel in each region. If we run the fuel in a certain region to a higher burnup, we decrease the region's reactivity, and consequently lower the region's power, as we will see in greater detail further below.

8 Fuel-Management Calculations for Core Design

CANDU reactors are refuelled on-line. Refuelling operations are typically performed every day, or a few times per week. A very small fraction of the fuel in core (typically ~8 fuel bundles, representing something of the order of 0.1-0.2%) is replaced at each refuelling operation. Also, because refuelling is performed on a quasi-continuous basis, the excess reactivity of the core is always very low (except at the very start of the core life, or after a long shutdown).

On account of these features, the overall global power distribution in a CANDU core operated in a well-thought-out manner remains largely constant in time, with relatively small variations ("power ripples") superimposed locally as individual bundles and channels go through their normal burnup cycles from one of their refuellings to the next. Thus, it is meaningful and

useful to want to calculate this long-term average.

In this Section we start to look at fuel-management models which are used at the core-design stage. They represent the reactor in the average over a long period of time, not on any particular day. Such models can then be used to perform diffusion calculations (for example, with full-core codes such as RFSP-IST) to obtain the expected spatial distribution of flux and power in the long term; these distributions will eventually be used as basis for running the reactor over time.

8.1 The Axially-Uniform Model Based on Continuous Refuelling

The first model which we will develop is the Axially-Uniform Model of nuclear cross sections. We start by visualizing the reactor as being continuously refuelled, meaning that fuel is pushed continuously and simultaneously through every fuel channel at a certain speed υ. If channels have length *L*, then the time *T* it takes fuel to traverse the entire channel length is $T = \frac{L}{\upsilon}$.

Let $\Phi(x)$ be the axial flux distribution in a fuel channel, from $x = 0$ to $x = L$ (channel inlet to outlet). Suppose as a first approximation we would like to use the average value $\overline{\Sigma}$ of a basic-lattice cross section $\Sigma(x)$ over the channel [this can be the fission cross section, or absorption cross section, or any other). We can get this average value by applying Eq. (5) to our specific case:

$$\overline{\Sigma} = \frac{\int_0^L \Sigma(x)\Phi(x)\,dx}{\int_0^L \Phi(x)\,dx} \tag{14}$$

Now change variables from x to t by using $x = \upsilon t,\ dx = \upsilon dt$, which gives

$$\overline{\Sigma} = \frac{\int_0^T \Sigma(t)\Phi(t)\,dt}{\int_0^T \Phi(t)\,dt} \tag{15}$$

Now we make a second change of variable, from *t* to fuel irradiation ω:

$$\omega(t) = \int \Phi(t)\,dt,\ \ d\omega = \Phi dt \tag{16}$$

And if we define ω_{exit} as the value of the fuel irradiation at the exit of the fuel channel (but note that we are not limited to using the same value of exit irradiation in all channels), we get

$\therefore$

$$\overline{\Sigma} = \frac{\int_0^{\omega_{exit}} \Sigma(\omega)\,d\omega}{\int_0^{\omega_{exit}} d\omega} = \frac{\int_0^{\omega_{exit}} \Sigma(\omega)\,d\omega}{\omega_{exit}} \tag{17}$$

Thus we would use the average values $\overline{\Sigma}$ calculated as per Eq. (16) (for the various cross-section types) at all bundle positions within each individual fuel channel. Note that actually Eq. (16) is

just another manifestation of Eq. 10). [Once again, note that we can select different values of ω_{exit} in different channels. Another way to visualize this is to realize that we have the option of selecting different speeds υ of the fuel in different channels.]

Once all cross sections are defined for each position in core, and the incremental cross sections of in-core devices are added, then we have a full axially-uniform reactor model for use with the diffusion equation. The solution of that equation will be the core reactivity and the 3-dimensional flux distribution in core, from which all powers, etc., can then be calculated.

Of course this is a simplified model of the properties in the reactor, since in each channel the properties are modelled as uniform and we are losing the axial variation of lattice properties over the channel. This is actually not as bad an approximation as may be thought, because it can be shown that the bi-directional refuelling in CANDU reactors leads to approximately constant lattice properties when these are averaged over 2 neighbouring channels (with the same ω_{exit}). Thus, the axially-uniform model was used extensively in the early stages of CANDU design development and analysis. However, eventually a model which better models the variation of lattice properties in each individual channel was developed. This is the time-average model, described in the next subsection.

8.2 The CANDU Time-Average (Equilibrium-Core) Model

The time-average model [Rouben 2007] is the next, more realistic effort to take into account the spatial distribution of properties over time, arising from different fuel irradiations in different fuel bundles. As opposed to the axially-uniform model, the time-average model does not average lattice properties along a channel. Instead, it takes into account the fact that individual fuel channels are refuelled once in a while, and that fuel bundles remain in a certain location in the channel until the channel is refuelled. It models the effect of the axial refuelling scheme (e.g., 8-bundle-shift (8-bs), 4-bundle-shift (4-bs), etc.) used in each channel on the average lattice properties at each bundle location in the channel.

8.2.1 Degrees of Freedom in CANDU Operation

Let us look at the breadth of freedom in operating a CANDU reactor. The following are "degrees of freedom" available to the operator:

- The fuel type used (e.g., the fuel enrichment or the fuel-bundle geometry, that is 28-element, 37-element, etc.)
- Which channel to refuel on any given day
- Which channel to refuel after the previous one refuelled, and which channel to re-fuel after that one (i.e., ultimately, the sequence of refuellings)
- The "axial" refuelling scheme to be used (i.e., the number of bundles to change at each channel refuelling), e.g., 8-bundle shift, 4-bundle shift, 2-bundle shift, or other, more complicated scheme. In addition, this axial refuelling scheme may vary from core region to core region, or even in principle from channel to channel. [Note that refuelling schemes where bundles are reshuffled from one channel to another are not considered in the current discussion.]
- The refuelling frequency of each channel – the average frequency will depend on the

fuel type(s), on the refuelling scheme, and on the overall core-reactivity decay rate.
- The desired discharge (exit) burnup for the fuel, as an overall average and even by core region
- The power desired for each core zone or region – yes, this is something that can be set at design time (of course the sum of the region powers must be equal to the total core power).

These degrees of freedom are not all independent of one another: refuelling frequencies, discharge burnup, and region powers are all interrelated. Nonetheless, it is clear that there is a very broad range of possibilities available to the operator, even if a single fuel type is used. Therefore, rational and orderly operation of the reactor really requires following guidelines in terms of the average refuelling frequency of each channel.

8.2.2 Definitions

Let us start with some notation and some definitions:

Let *c* be a label for channel number, *c* = 1 to *M*, where *M* is the total number of channels in core

- Let *b* be a label for a bundle-position number within a channel, *b* = 1 to N_c, where N_c is the total number of bundles per channel
- Then we can use *cb* as a label for a specific individual bundle position in core
- Let *x* be a label for a type of nuclear cross section, e.g., absorption, fission, etc.
- Let $\phi_{cb}(t)$ and $\Phi_{cb}(t)$ be respectively the **fuel flux** and the **cell flux** at position cb at some time t. These fluxes are actually multigroup quantities (2-group in the current RFSP-IST). $\phi_{cb}(t)$ and $\Phi_{cb}(t)$ are related by the so-called "F-factor", ratio of flux in fuel to average flux in cell; this is calculated by the lattice code.
- Let $\phi_{tav,cb}$ and $\Phi_{tav,cb}$ be the corresponding **fuel and cell fluxes** in the time-average model. Because the time-average model does not have time as an independent parameter, these fluxes are not functions of time.
- Let D_c be the average time interval between refuellings of channel *c*; this is also known as the **dwell time** for channel *c*. [Note: This is not the same as the average residence time of a bundle in core, since in most refuelling schemes some bundles remain in core for more than 1 cycle.]
- Let $\omega_{entr,cb}$ be the fuel irradiation at the time of entrance of the fuel at position *cb*, and
- Let $\omega_{exit,cb}$ be the fuel irradiation at the time of exit of the fuel from position *cb*.
- Also, let E_c be the total number of bundles which leave channel c on a refuelling operation; these bundles will be labelled β, β = 1 to E_c. For example, in the 8-bs refuelling scheme, we have E_c = 8, and the bundles β are bundles 5 to 12 from the channel inlet end.
- Let $\omega_{exit,c}$ be the **channel exit irradiation** of the fuel from channel *c*, i.e., it is the average of the exit burnup over the bundles β = 1 to E_c which leave the core on a refuelling.
- Let $\Sigma_{x,cb}(\omega)$ be the macroscopic cross section for cell *cb* for reaction type *x*. This is a function of the fuel irradiation ω.

- Then let $\Sigma_{x,cb,tav}$ denote the time-average cell cross section for cell *cb* and reaction type *x*; this will be calculated in the time-average model.

8.2.3 The Selectable Quantities

We shall take the channel exit-irradiation values $\omega_{exit,c}$ values as the essential degrees of freedom, i.e., as the ones most easily selected by the designer. These $\omega_{exit,c}$ values will be dependent on the dwell times D_c, and vice versa. Both of course will depend also on the fuel-flux distribution $\phi_{tav,cb}$, which depends on the core flux distribution $\Phi_{tav,cb}$, which will be determined as the solution of the neutron diffusion equation. The relationship between the fuel and cell fluxes comes from the lattice (or cell) calculation.

Note that in addition the axial refuelling scheme is also a degree of freedom which can be selected a priori by the designer, for each fuel channel (or region) in core.

8.2.4 Relationships Between Bundle and Channel Irradiations and Dwell Times

For positions *cb* in core where a refuelling introduces fresh fuel:

$$\omega_{entr,cb} = 0 \tag{18}$$

For positions *b* to which a bundle moves from a position *b'* in the same channel upon refuelling,

$$\omega_{entr,cb} = \omega_{exit,cb'} \tag{19}$$

Now, since

$$\omega \equiv \mathit{Irradiation} = \mathit{Flux} \times \mathit{Time}(19) \tag{20}$$

then

$$\mathit{Incremental\ irradiation\ accumulated\ by\ fuel\ at\ cb\ between\ refuellings} = \varphi_{tav,cb} \times D_c \tag{21}$$

and we can use this to relate the exit irradiation to the entrance irradiation:

$$\omega_{exit,cb} = \omega_{entr,cb} + \varphi_{tav,cb} \times D_c \tag{22}$$

We can now relate the average channel exit irradiation $\omega_{exit,c}$ to the fuel flux at the various bundle positions in the channel, as follows.

By definition,

$$\omega_{exit,c} = \frac{1}{E_c} \sum_{\beta=1}^{E_c} \omega_{exit,c\beta}$$

(23)

and, using Eq. (22)

$$\omega_{exit,c} = \frac{1}{E_c} \sum_{\beta=1}^{E_c} \left(\omega_{entr,c\beta} + \varphi_{tav,c\beta} \times D_c \right) \tag{24}$$

The bundle positions β are those from which bundles leave the channel on refuelling. Some of these positions may at the same time be positions into which bundles enter as fresh bundles, e.g., bundle positions 5 to 8 for an 8-bundle-shift refuelling scheme. For these bundles, the

term $\omega_{entr,c\beta}$ is zero and does not contribute in Eq. (24).

On the other hand, some β may be positions where bundles come from a previous location in the channel, e.g., bundles 9 to 12 in the channel for the 8-bundle-shifts refuelling scheme. For these bundles, the term $\omega_{entr,c\beta}$ is **not** zero, but can be equated (see Eq. (19)) to $\omega_{exit,c\beta'}$, where β' is the position from which the bundle at β came on the previous refuelling. In this way, it can be seen that the sum of terms in ϕ_{tav} in Eq. (24) now extends over bundles β **and** β'.

In fact, by repeating the exercise for bundles β' and examining from which position **they** came, we see that eventually all the terms $\omega_{entr,cb}$ disappear and are replaced by terms in ϕ_{tav}, with the sum in Eq. (24) covering **all** bundles b in the channel:

$$\omega_{exit,c} = \frac{1}{E_c} \sum_{b=1}^{N_c} \left(\varphi_{tav,cb} \times D_c \right) \tag{25}$$

Since D_c is a channel quantity and is independent of the bundle position *b*, we can take D_c out of the sum. We can also turn the equation around to isolate D_c on one side:

$$D_c = \frac{E_c \times \omega_{exit,c}}{\sum_{b=1}^{N_c} \varphi_{tav,cb}} \tag{26}$$

Eq. (26) links the channel dwell time to the channel average exit irradiation, the refuelling scheme (via E_c), and the total fuel flux in the channel.

8.2.5 Relationships Between Time-Average Cross Sections and Irradiation Values

For any bundle position *cb* in core and any cross section of type x, we can write the time-average cross section during the bundle's residence time at that location using Eq. (11) derived earlier:

i.e.,

$$\Sigma_{x,cb,tav} = \frac{\int_{\omega_{entr,cb}}^{\omega_{exit,cb}} \Sigma_{x,cb}(\omega)\, d\omega}{\omega_{exit,cb} - \omega_{entr,cb}} \tag{27}$$

The integral in Eq. (27) can easily be computed using the cell cross section obtained as a function of irradiation with the cell code.

Using the values selected for the entrance and exit irradiations at each position, the $\Sigma_{x,cb,tav}$ can be computed for all positions in core. These nuclear lattice properties can then be entered into the diffusion core model, and the neutron-diffusion equation can be solved by the usual means to find the time-average flux distribution $\Phi_{tav,cb}$.

8.2.6 Summary of Computational Scheme

The computational scheme for the time-average model is now complete. It consists of the neutron-diffusion equation plus Eqs. (18), (19), (22), (24)-(27).

Remember that:

- The flux distribution $\Phi_{tav,cb}$ depends on the time-average cross sections $\Sigma_{x,cb,tav}$ through the diffusion equation;
- The time-average cross sections $\Sigma_{x,cb,tav}$ depend on the entrance and exit irradiation values via Eq. (27);
- The entrance and exit irradiation values $\omega_{entr,cb}$ and $\omega_{exit,cb}$ depend on the fuel flux in the channel (which will be calculated from the time-average flux $\Phi_{tav,cb}$) and the dwell-time D_c via Eqs. (18), (19) and (22);
- The dwell-time D_c depends on the flux and on the channel-average exit irradiation $\omega_{exit,c}$ via Eq. (26).

It is clear then that this equation set must be solved as a self-consistency problem. An iterative solution scheme is applied until satisfactory self-consistency is attained.

The basic independent data which needs to be input to the time-average model consists of:

a) the axial refuelling scheme for each channel (for instance, 8-bs, 4-bs, 10-bs, etc.). Note that this may vary from channel to channel, i.e., one may choose 8-bs for some channels and 4-bs for others. Typically, however, the axial refuelling scheme must be selected taking into account also non-physics considerations, e.g., fuelling-machine utilization.
b) the channel-average exit irradiations $\omega_{exit,c}$. These are the main degrees of freedom in the problem. While there are in principle as many degrees of freedom as there are channels, the core is usually divided into a small number of relatively large regions, in each of which ω_{exit} is taken as uniform, to reduce the number of degrees of freedom. Figure 9 illustrates a possible time-average model with 8 irradiation regions, for fine control over the target power distribution; note in particular the regions at the bottom and far right-hand side, which represent areas in the calandria with much hardware (to hold device guide tubes in position) and consequently for which different exit irradiations are needed.

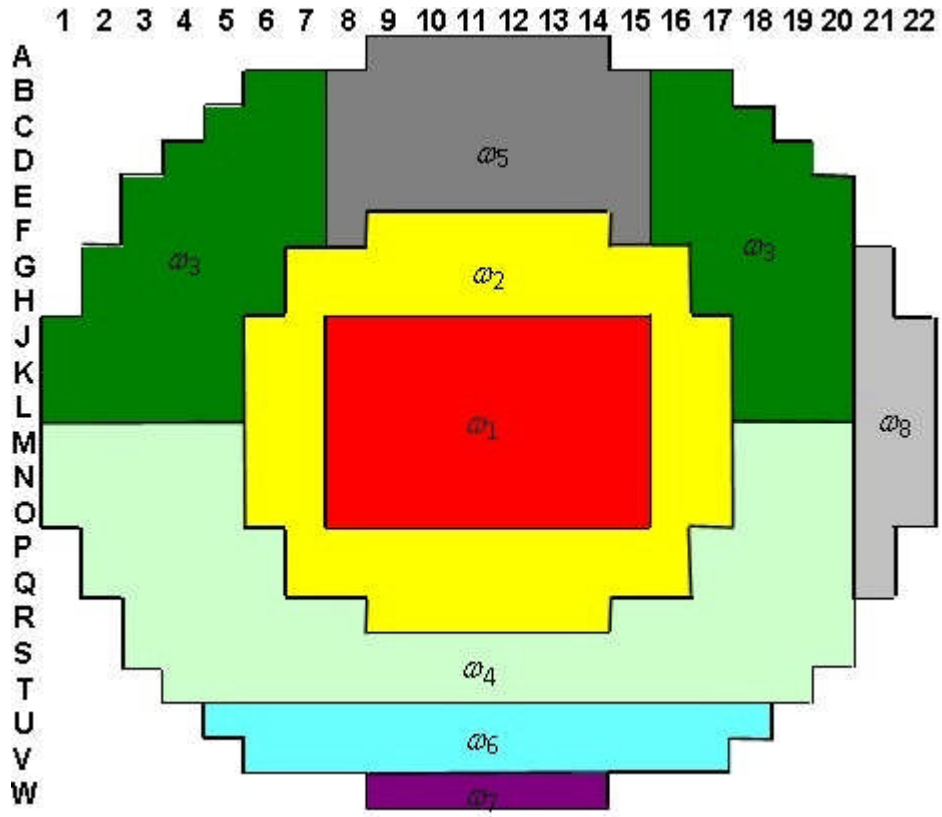

Figure 9 Example of Possible Time-Average Model, with Eight Irradiation Regions

A very important point to note is that there are constraints on the degrees of freedom:

- First, reactor criticality is required. That is, the value obtained for the reactor multiplication constant k_{eff} when solving the diffusion equation must be unity (or close to unity, if we know that the model does not include all material in core, for instance detector material, in which case a small bias is required) – otherwise, the results obtained cannot be expected to represent a realistic time-average picture.
    - Secondly, the absolute and relative values of the channel-average exit irradiations $\omega_{exit,c}$ have a significant impact on the radial power distribution, and in turn on the refuelling frequencies and the attained discharge burnup. Increasing the exit irradiation in a region tends to decrease the local reactivity of the region, leading to a lower power. Decreasing the exit irradiation has the opposite effect.
    - Thus, the channel exit irradiations need to be selected judiciously to achieve both criticality and the desired degree of radial flattening of the power distribution and of consequent burnup and refuelling rates.

The flow chart for the iterative calculation for the time-average flux and power distribution is

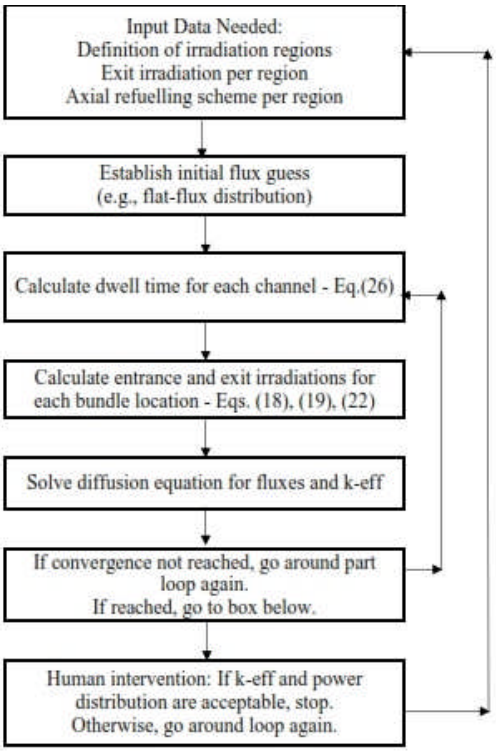

Figure 10 Flow Chart of Time-Average Calculation

shown in Figure 10. First, the user selects the axial refuelling scheme(s) and the core regions with the corresponding values of $\omega_{exit,c}$. Using an initial guess for the flux distribution, the program determines the channel dwell times D_c and the entrance and exit values of irradiation $\omega_{entr,cb}$ and $\omega_{exit,cb}$. It then determines the time-average cross sections $\Sigma_{x,cb,tav}$. The program then proceeds to solve the neutron diffusion equation. The flux distribution thus obtained is used to repeat the cycle, until consistency is attained in all the parameters.

At that point, the user must intervene (outside the code) and examine the results obtained for k_{eff} and for the power distribution. If k_{eff} is not equal (or at least close) to unity, the exit irradiations selected must be adjusted. A standard way of doing this is to increase or decrease all the $\omega_{exit,c}$ values uniformly, depending on whether the k_{eff} is too high or too low. Also, the radial power distribution must be examined (this can be done at this time, or after another cycle of iterations with the new $\omega_{exit,c}$ values).

If the radial power distribution is not satisfactory, e.g., if the degree of radial flattening is not the desired one, or if peak channel powers are too high (or too low), or if the rate of refuelling in certain regions is too high, then the regions and/or exit irradiation values must be "manually" adjusted to move the results in the desired direction. Then another cycle of iterations is started.

This process is repeated until a critical reactor (or the required bias) and a satisfactory power distribution are obtained, at which point the time-average model has been established. In practice, for operating reactors some additional conditions may be imposed when designing the

time-average model, such as matching time-average channel powers determined from relevant operating histories. This is obtained by varying the individual channel exit irradiations; in such instances Figure 9 may be replaced by a map with more irradiation regions, or even by a full map of 380 individual exit irradiations, one for each channel.

8.2.7 Output of Time-Average Calculation

The calculations described above provide or can provide a number of time-average outputs, which are useful as design outputs and also as targets for the reactor operator. These outputs include:

- The reactor multiplication constant (k_{eff})
- The 3-dimensional time-average neutron-flux distribution. From this, the time-average zone fluxes can be obtained, for use as the target distribution by the Reactor Regulating System.
- The channel and bundle powers. Note that these are the time-average values; they do not include the instantaneous ripple corresponding to the stage in the refuelling cycle. The maximum time-average channel power in a CANDU-6 reactor is typically about 6.6-6.7 MW.
- The entrance and exit irradiation values for each bundle position in core.
- The channel dwell times (intervals between refuellings of the individual channels). Figure 11 shows dwell times from a CANDU-6 time-average calculation. One-cycle residence time of bundles at a given position in a CANDU-6 reactor typically range between 160 and 360 full-power days (FPD).
- The refuelling rates in the various irradiation regions.
- The reactivity decay rate in mk/Full-Power Day. This can be determined in the time-average model from the change in reactivity (as obtained from the change in k_{eff}) for an overall change in the exit irradiation in all regions and linking that to the change in residence time corresponding to the change in irradiation. The reactivity-decay rate per FPD is useful in gauging the expected variation in zone-control-compartment fill per FPD if there were no refuelling. [Note that the same quantity can also be calculated by doing 2 core snapshots 1 FPD apart (see next Section) without refuelling any channels.]

See some of these results from a particular time-average calculation for a CANDU-6 reactor in Table 2 below.

Table 2 Example of Summary Results from a CANDU-6 Time-Average Calculation

Total Reactor Power (MW)	2061.4
Total Fission Power (MW)	2158.5
Average Channel Power (kW)	5425
Average Bundle Power (kW)	452
Uranium Mass per Bundle (kg)	19.1
Maximum Channel Power (kW)	6604 (N-6)
Maximum Bundle Power (kW)	805 (O-5, Bundle 7)
k_{eff}	1.00250
Average Exit Burnup	3272 MW.h/Bundle 171.3 MW.h/kg(U) 7139 MW.d/Mg(U)
Avg. Channel Dwell Time (FPD)	192
Feed Rate	1.98 channels/FPD 15.83 bundles/FPD
Reactivity Decay Rate (mk/FPD)	-0.385
Average Zone Fill (%)	50.0

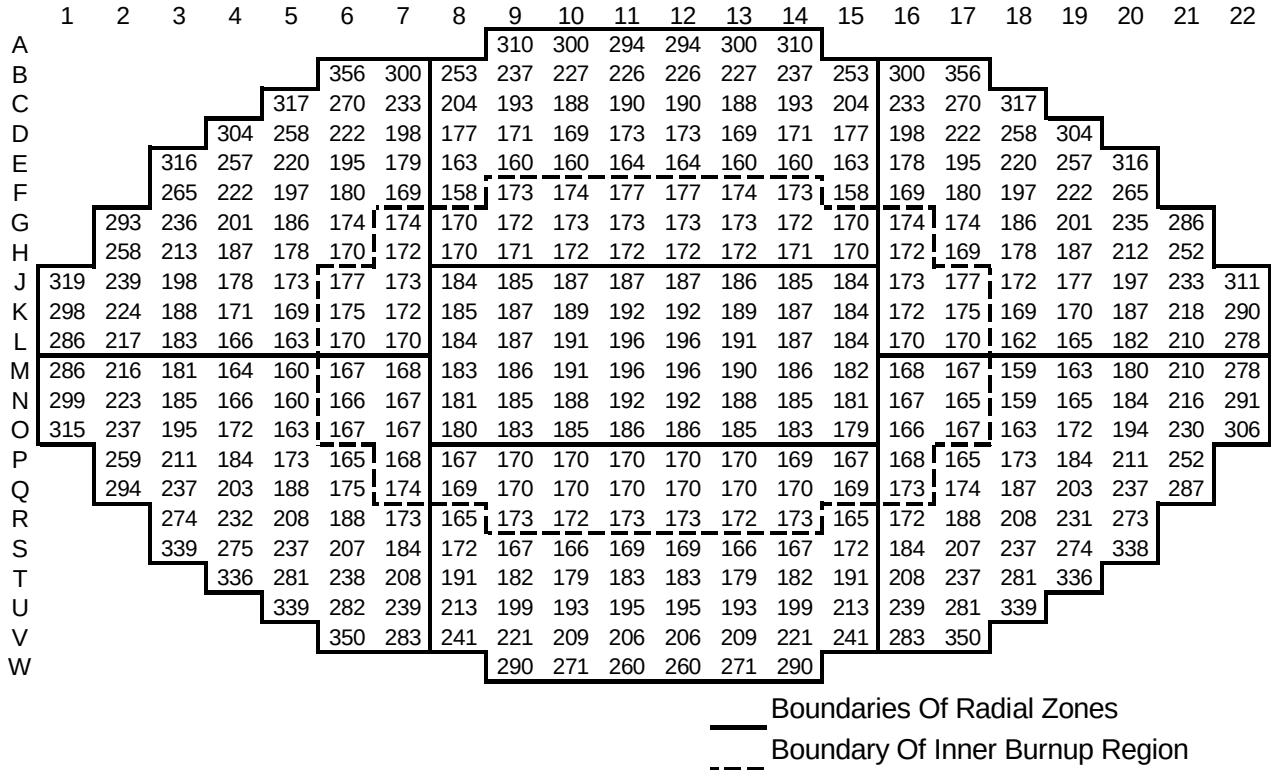

Figure 11 Dwell Times (in FPD) from a Time-Average Calculation for the CANDU 6

9 More Design Calculations: Snapshots Based on the Time-Average Model

In this section we look at another type of design calculations, intended to look ahead at power and overpower results which may be expected during reactor operation.

An important point to remember about the time-average core is that it features no refuelling ripple: all properties are averaged over residence time, no channel is near the beginning or end of its refuelling cycle. However, at the core design stage it is important to be able to collect possible and realistic snapshots of the reactor power distribution during actual reactor operation, without actually following the history of the reactor. We can obtain snapshots of the power distribution if we can define instantaneous distributions of bundle irradiation to calculate instantaneous lattice properties, based on which the diffusion equation can be solved.

On way to obtain instantaneous lattice properties is to create a map of "channel ages". To define channel ages let us first consider a typical "channel-irradiation-and-power cycle".

9.1 Channel-Irradiation-and-Power Cycle

The channel-irradiation-and-power 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. The local reactivity **increases** for ~ 40-50 FPD, and the power of the channel tends to increase further.
- Following the plutonium peak, the reactivity of the refuelled channel starts to decrease as the irradiation increases further, and its power drops slowly.
- The reactivity of the channel and its power then continue to drop. Eventually, the channel becomes a net "sink" (absorber) of neutrons, the time since the channel's last refuelling approaches the channel's "dwell time" D_c (as defined in the time-average model), and the fuel irradiation in the bundles which will be discharged at the channel's next refuelling increases towards the discharge value, and the time comes around when the channel must be again refuelled. 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.

In on-going reactor operation, the power of each channel therefore goes through an oscillation about the time-average power during every cycle. This cycle repeats - not in exactly the same quantitative detail each time- at every refuelling of the channel. This is illustrated in Figure 12.

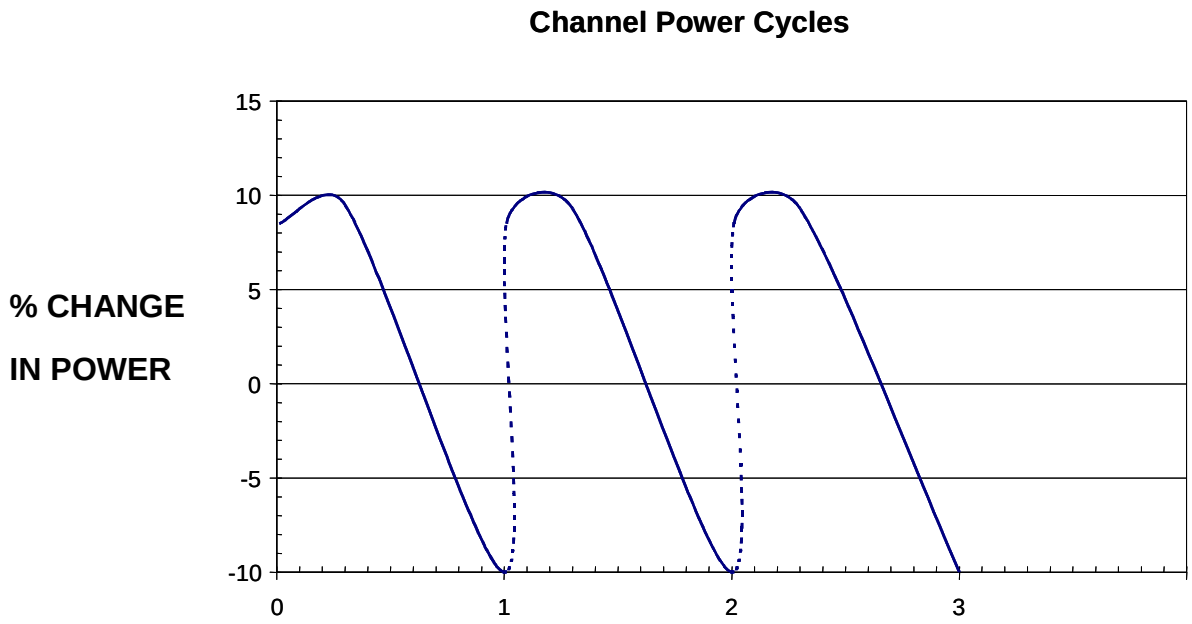

Figure 12 Channel-Power Cycle During On-Going Reactor Operation

9.2 Channel Ages

On any day of operation, we can define the "age" a_c of channel *c* as the instantaneous fraction that it has achieved in its irradiation cycle since its last refuelling. On any given day different channels will of course have different ages: channels which were refuelled recently have a low age (a number close to 0.0), whereas those which were refuelled quite a while ago have a "high" age (a number close to 1.0). The channel ages a_c will be numbers between 0 and 1, although it is possible for a given channel to have an age somewhat higher than 1, if it has not been refuelled for a time longer than its dwell time.

From the meaning of the channel's age, we can define a value for the instantaneous irradiation of all bundles in the channel. For channel *c* and bundles *b*, using the entrance and exit irradiation values for bundle positions (from the time-average model), we can write for bundle's *cb* instantaneous irradiation

$$\omega_{inst,cb} = \omega_{entr,cb} + a_c \times \left(\omega_{exit,cb} - \omega_{entr,cb} \right) \tag{28}$$

The lattice properties for position c_b can then be determined from the instantaneous irradiation value.

9.3 Snapshot with Patterned Ages

To model the core on a given day, we can define a core snapshot by assigning to each channel *c* in core an age value a_c.

We could select random values (between 0 and 1, or slightly higher than 1) for the a_c. However, random selection of ages may easily result in low age values for close-neighbour channels (i.e., clusters of low ages). This would not be very realistic, because low ages in close proximity can

lead to "hot spots" in channel power, and a good operator will consciously avoid refuelling close-neighbour channels very closely in time, so as to avoid low ages close to one another.

It is therefore more desirable to use some "intelligence" in assigning channel ages. One way which has been found to do this is to start by defining an "intelligent pattern" in the order in which channels would be refuelled in a relatively small region of the core, e.g., 6-channels-by-6-channels (see Figure 13). The next step is to subdivide the entire core into 6 x 6 regions and to define a core "refuelling sequence" by selecting channels intelligently in sequence from one region to region to another, making sure that whenever revisiting each 6 x 6 region the order within the region follows the pattern defined for the 6 x 6 square. One obtains (see [Rozon and Shen 2001] a core refuelling sequence such as shown for example for the CANDU 6 in Figure 14, with numbers n_c from 1 to 380 assigned to the channels. The age assigned to channel c can then be defined as

$$a_c = \frac{380.5 - n_c}{380}, \quad n_c = 1,...,380 \tag{29}$$

Once the snapshot model with lattice properties derived using Eqs. (29) and (28) is solved with the diffusion code, useful results are obtained for the maximum channel and bundle powers and, also very importantly, for the Channel Power Peaking Factor (see further below, Section 11.3) and for the overall channel-power, bundle-power, and region-power distributions which can be expected during on-going reactor operation.

From the one snapshot, many others can be obtained by "massaging" the core refuelling sequence, for example by reflecting it horizontally or vertically, or alternatively by redefining each n_c as modulo (n_c + N), where N is a number which can be selected at will. In this manner many sets of results and good statistics can be obtained for maximum channel and bundle powers and the CPPF. This can be very useful in determining rules for selecting channels for refuelling.

22	29	8	17	14	27
7	16	23	34	9	4
30	11	28	19	24	35
3	20	33	2	21	6
12	25	10	15	26	31
5	36	13	32	1	18

Figure 13 Possible Pattern for Order of Refuelling in a 6 x6 Region of Core

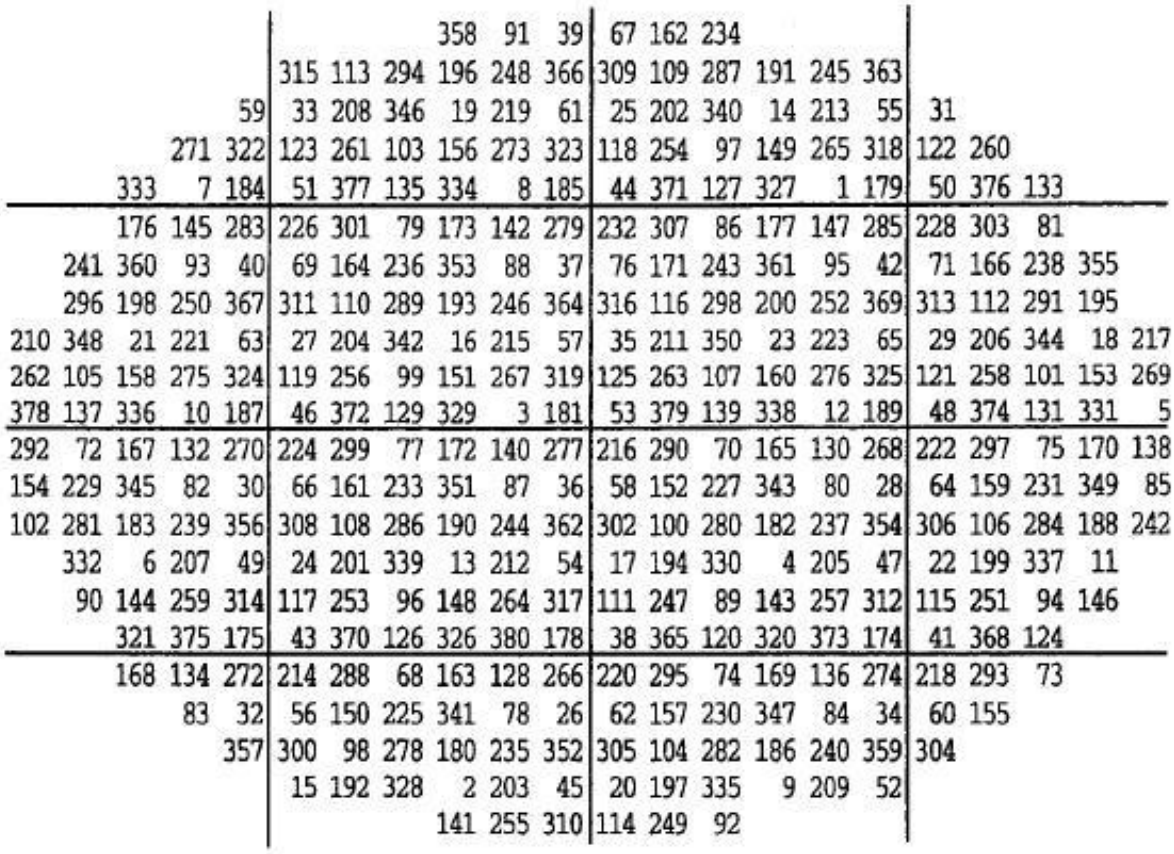

Figure 14 Possible Core Refuelling Sequence for the CANDU 6

10 Design of the CANDU Initial Core

The CANDU initial core is the state of the reactor as it starts life. This is a very specific state, in which the fuel throughout the core is fresh (unirradiated). At zero irradiation, the CANDU lattice has high reactivity. Of course this high reactivity must be cancelled in order to allow for the reactor to be run in steady state.

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.

Two means are available in CANDU to cancel the high reactivity of the initial core and to flatten its power distribution. These are the use of moderator poison (boron, or possibly gadolinium) and the use of depleted fuel (fuel at an enrichment around 0.48-52% of U-235).

Advance calculations are needed to design how exactly to use these two means to render the core critical and to provide appropriate power flattening. Design calculations are done with various concentrations of poison and various numbers and locations of depleted-fuel bundles until a desirable/attractive result is achieved.

The solution is not unique. There are many ways to design the poison concentration and

depleted-fuel assignment for the initial core. However, note:

- using only depleted fuel and no moderator poison would result in having to initiate refuelling very early, the consequence being that too much fuel would be removed with very low burnup – an economic penalty, whereas
- using only boron and no depleted fuel would not achieve sufficient power flattening.

In the first design of the CANDU-6 initial core, the solution which was adopted used two depleted-fuel bundles (of 0.52 atom percent ^{235}U content) in each of the innermost 80 fuel channels, with the depleted-fuel bundles located in positions 8 and 9, numbered 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. The boron concentration in this design turns out to be about 2 ppm at full CANDU-6 power. In recent years some operators have redesigned the initial core loading to factor in the impact of the required initial poison concentration on the safety analysis.

11 Fuel-Management Calculations for Reactor Operation

Once the reactor is operating, it is important to follow the reactor history. This is called a core-follow calculation. It involves calculating updates of the core's power and burnup distributions at specific instants in the reactor's operating history, usually at time intervals of 1 or a few FPD. The objective is to have good knowledge of the state of the core at all times, to be able to confirm compliance with the maximum channel and bundle powers allowed by the operating licence.

Note that the core-follow calculations described here can also be done for a simulated core history, with simulated channel refuellings and device positions. This is to demonstrate how a reactor would/could be operated.

11.1 From Initial Core to Onset of Refuelling

At the very start of reactor operation, the entire initial fuel load goes through the plutonium peak at about the same time (about 40-50 FPD). At this juncture, the core reaches its global plutonium peak and the core reactivity is the highest it will ever be. Following the plutonium peak, the plutonium production can no longer compensate for the depletion of ^{235}U and the build-up of fission products, and the excess core reactivity decreases. Also at this time moderator poison must be removed 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. The core-follow for this period is therefore rather simple. It is a matter only of simulating the core every few days with instantaneous reactivity devices and a poison concentration which yields a critical reactor.

When the excess core reactivity has fallen to a small value, actually, about 10 or 20 FPD before it reaches 0 (i.e., typically around FPD 100), refuelling operations start. It is best not to wait until excess reactivity is exactly 0, because the initial refuelling rate would prove too high.

11.2 After Onset of Refuelling

When on-power refuelling starts, it becomes the primary means of maintaining a CANDU

reactor critical. During the transitional period which follows the onset of refuelling, the reactor gradually approaches the equilibrium core, with the refuelling rate rapidly tending to the time-average value (for example, approximately 15-16 bundles per FPD in the CANDU 6).

When refuelling starts, it is channels in the inner core which tend to be refuelled first, because this region has achieved the highest burnup (on account of its higher flux values). As time goes on, more channels in the peripheral regions of the core will be refuelled.

Approximately 400 to 500 full-power days (FPD) after initial start-up, a CANDU reactor has reached an "equilibrium core" state. The overall refuelling rate, the in-core average burnup, the burnup of the discharged fuel, and the refuelling rates in different core regions have become essentially steady with time. A number of channels are refuelled every day, **on the average**. However, 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.

Doing core-follow calculations at short intervals (1 to a few FPD) then becomes crucial in running the reactor in a steady and effective manner, to comply with licence limits, to ensure that there are no "hot spots" or large power tilts.

Each run of the diffusion code in the core-follow requires inputting the instantaneous 3-dimensional irradiation/burnup distribution. This of course comes from the output of the preceding run. In addition, each run of the core-follow requires modelling all channel refuellings which have occurred since the previous run (and their timing). This allows modelling bundle movements and the entry of fresh-fuel bundles in appropriate locations. The instantaneous positions of reactivity devices are also input. From the power calculated and the time step used, the fuel irradiation in each individual bundle are updated, and the lattice properties updated.

The output of the current run is then obtained and becomes the starting point for the calculation at the next time step. The current 3-dimensional power and fuel-irradiation distributions obtained serve then specifically as the basis on which intelligent selections of channels to refuel are made (see second section below).

In addition, another very significant output of core-follow calculations is the determination of the Channel Power Peaking Factor, which is used in the calibration of in-core protection-system (Regional Overpower/Neutron Overpower) detectors. This is defined in the next Section.

11.3 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 power-and-irradiation cycle. Therefore, the maximum instantaneous channel power is always higher than the maximum channel power in the time-average model.

The Channel-Power Peaking Factor (CPPF) quantifies the degree by which the instantaneous power distribution peaks above the time-average distribution:

$$CPPF = \underset{c}{Max}\left[\frac{Channel-Power_{instantaneous}(c)}{\underset{c}{Channel-Power}_{time-average}(c)}\right] \tag{30}$$

where *c* runs over channels in the core.

The exact value of the CPPF (which varies from day to day) is extremely important because it is used to calibrate the in-core protection (ROP/NOP) detectors. The protection system(s) are designed to actuate a shutdown system in a power increase before any fuel channel risks reaching a "critical channel power" (e.g., a power at which there would be fuel dryout somewhere in the channel). The CPPF has therefore important safety implications.

Hundreds of flux shapes are used in the design of the ROP/NOP detector system to determine detector positions and setpoints. These flux shapes are 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, channels would reach their "critical channel power" earlier than predicted in the time-average model. To take this into consideration and ensure proper safety coverage in the instantaneous power shape, the in-core detectors are calibrated each day to the instantaneous value of CPPF.

On the other hand, a high CPPF value could cut into the reactor's operating margin, i.e., the CPPF has economic implications. It is therefore important to keep the CPPF as low as possible.

The competing safety and economic considerations mean that a careful selection of channels to be refuelled needs to be made always. Determining the daily CPPF value, selecting channels to keep it low, 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.

11.4 Criteria for Selecting Channels for Refuelling

One of the main functions of the fuelling engineer or reactor physicist is to establish a list of channels to be refuelled during the following period (few days) of operation. For instance, for a 5-FPD period, a list of approximately 10 channels must be prepared. 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.

Good combinations of channels selected 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 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%) .

Note that channel selections for refuelling are seldom unique. Many options are 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 the code can predict the resulting zone-control-compartment fills (i.e., the bulk-control and spatial-control responses of the Reactor Regulating System), will show whether the various power, burnup, and zone-fill criteria are likely to be satisfied, or whether a different channel selection should be made. In fact, CANDU operators employ pre-simulations routinely to ensure optimum fuelling selections, including achieving low CPPF values, i.e., good operating margins.

12 Advanced Fuel Cycles

On account of the high neutron economy provided by heavy-water moderation, the CANDU reactor can run on natural-uranium fuel (unlike Light-Water Reactors). But this also means that the CANDU reactor can run on fuels which have higher effective enrichments than natural uranium. This allows a number of Advanced Fuel Cycles to be adopted in CANDU reactors in the future. Many of these cycles illustrate the great synergy possible between CANDU and PWRs, as CANDU can in some way re-use PWR fuel and extend the achievable burnup and reduce waste volumes.

Fuel-management simulations for Advanced Fuel Cycles would proceed in essentially the same manner as for the once-through natural uranium (NU) cycle. Of course the calculation of lattice properties would have to reflect the presence of different nuclides and/or of different nuclide concentrations.

In many cases the Advanced Fuel Cycles can be adopted in CANDU reactors without any changes to the reactor or the control devices. In other cases it may be necessary to adopt different axial refuelling schemes (replacing fewer bundles per visit to a channel). Or it may be appropriate to alter or remove some reactivity devices to optimize reactivity and burnup considerations.

Some examples of Advanced Fuel Cycles in CANDU are:

- The use of Slightly Enriched Uranium (SEU). In this cycle uranium with enrichment in the range 0.9% to, say, 1.2% is used instead of NU. The discharge burnup can be increased significantly, with a consequent significant decrease in the volume of irradiated fuel replaced per year.
- The use of recycled uranium (RU). This uranium would be recovered from the reprocessing of irradiated fuel from CANDU, or more likely, from PWRs. In this case, the presence of certain concentrations of other isotopes of uranium in the fuel has to be recognized in the calculation of lattice properties.
- The MOX (Mixed Oxide) Fuel Cycle, where fuel would be produced from a mixture of uranium and plutonium separated in the reprocessing of PWR fuel. Here again a

range of fissile fractions can be accommodated in the MOX fuel.

- The DUPIC Cycle (Direct Use of PWR Fuel in CANDU). This cycle also uses fuel made from reprocessed PWR fuel, however in this case dry reprocessing is used which does not separate the plutonium from the uranium. A typical fissile fraction in irradiated PWR fuel is about 1.5% (0.9% uranium and 0.6% plutonium). Studies of the DUPIC cycle made from this fuel have shown that the **additional** burnup obtainable in CANDU can be about double the burnup achieved in the original Pressurized Water Reactor.
- The Thorium Cycle. Although Th-232 is not fissile, it is a fertile nuclide: Upon absorption of a neutron and two beta decays it becomes (via Pa-233) the fissile fuel U-233. This cycle draws a lot of interest, as thorium is roughly 3 times as abundant as uranium in the earth's crust, so its use would significantly increase our nuclear natural resources. The thorium cycle requires mixing uranium with a "driver" fuel which provides the neutrons to be absorbed by the thorium. The driver fuel can be uranium or plutonium. In CANDU, the thorium and the driver fuel can be placed in different core regions or different channels, or can even be mixed within the same fuel bundle. The U-233 can be burned in situ or can be separated by reprocessing the fuel after it has been sufficiently irradiated. Much study of the thorium cycle is ongoing currently.

13 Summary

This Chapter has summarized the concepts, models and calculations involved in the management of nuclear fuel in CANDU reactors.

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/NOP detectors.

14 References

[Damario 2016] D. Damario, "The Fuelling Machine", chapter in *The Essential CANDU*, edited by W. Garland, 2016.

[Hébert 2012] A. Débert, D. Sekki and R. Chambon, "A User Guide for DONJON Version4", Technical Report IGE-300, École Polytechnique de Montréal, 2012.

[Irish 2002] J.D. Irish and S.R. Douglas, "Validation of WIMS-IST", in Proceedings of the 23rd Annual Conference of the Canadian Nuclear Society, Toronto, Ontario, Canada, 2002 June.

[Marleau 1999] G. Marleau, "DRAGON, Theory Manual", Technical Report IGE-236, Institut de Génie Nucléaire, École Polytechnique de Montréal, 1999 February.

[Rouben 1995] B. Rouben, "Description of the Lattice Code POWDERPUFS-V", AECL-11357, 1995 October.

[Rouben 2002] B. Rouben, "RFSP-IST, The Industry Standard Tool Computer Program for CANDU Reactor Core Design and Analysis", in Proceedings of PBNC-2002 (13th Pacific Basin Nuclear Conference), Shenzhen, China, 2002 October.

[Rouben 2007] B. Rouben, "Review of the CANDU Time-Average Model and Calculations", in Proceedings of the 28th Annual Conference of the Canadian Nuclear Society, Saint John, New Brunswick, Canada, 2007 June 3-6.

[Rozon and Shen 2001] D. Rozon and W. Shen, "A Parametric Study of the DUPIC Fuel Cycle to Reflect Pressurized Water Reactor Fuel Management Strategy", Nucl. Sci. Eng. 138, 1-25 (2001)

15 Acknowledgements

The following reviewers are gratefully acknowledged for their hard work and excellent comments during the development of this Chapter. Their feedback has much improved it. Of course the responsibility for any errors or omissions lies entirely with the author.

Richard Chambon
Ovidiu Nainer

Thanks are also extended to Bill Garland for expertly editing and assembling the final copy.