Session 17
CANDU Full-Core Calculations and Fuel Management

by
B. Rouben

In this Session we will look at some of the typical full-core reactor-physics calculations done for CANDU, and we will then focus on CANDU fuel management.

1. CANDU Full-Core Reactor-Physics Calculations

In a previous session we discussed the 3-stage process for CANDU reactor-physics calculations. We then focused on lattice-physics calculations and did a number of in-class and assignment problems with the lattice code POWDERPUFS-V.

In this Session we will discuss briefly some of the kinds of full-core calculations done for CANDU. We will then focus on fuel management in CANDU. The discussion is largely general but, where necessary, particular reference is made to the CANDU 6 reactor.

2. Finite-Core Calculations and the RFSP Code

Once basic-lattice properties and reactivity-device incremental cross sections are available, the finite-core calculation can proceed. We will discuss here the finite-core computer code RFSP (Reactor Fuelling Simulation Program), now an Industry Standard Tool (RFSP-IST), which has been specifically designed for CANDU reactors. There are other full-core reactor-physics codes, e.g., SORO at OPG, which differ from RFSP in specific functionalities and features, but the general idea is similar.

RFSP 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, for certain CANDU reactor models (e.g., the CANDU-6) which have in-core vanadium detectors whose readings can be used to reconstruct the full 3-d-core flux distribution. We will not cover this method here.

2.1 The Time-Independent Diffusion Equation in RFSP

The time-independent neutron-diffusion equation solved in RFSP for eigenvalue problems is the finite-difference version of the neutron-diffusion equation in two energy groups (2 = thermal group, 1 = “fast” or “slowing-down” group):

$$-\vec{\nabla} \cdot D_1(\vec{r})\,\vec{\nabla}\phi_1(\vec{r}) + (\Sigma_{a1}(\vec{r}) + \Sigma_{1\to 2}(\vec{r}))\,\phi_1(\vec{r}) - \left(\Sigma_{2\to 1}(\vec{r}) + \frac{\nu\Sigma_{f2}(\vec{r})}{k_{eff}}\right)\phi_2(\vec{r}) = 0 \qquad (2.1a)$$

$$-\vec{\nabla} \cdot D_2(\vec{r})\,\vec{\nabla}\phi_2(\vec{r}) + \left(\Sigma_{a2}(\vec{r}) + \Sigma_{2\to 1}(\vec{r})\right)\phi_2(\vec{r}) - \left(\Sigma_{1\to 2}(\vec{r}) + \frac{\nu\Sigma_{f1}(\vec{r})}{k_{eff}}\right)\phi_1(\vec{r}) = 0 \qquad (2.1b)$$

Here the Σ_{a1}, etc., are the "homogenized" basic-lattice cross sections provided by the lattice code, to which the reactivity-device incremental cross sections have been added at the appropriate mesh points. And k_{eff}, which divides the fission cross sections, is the reactor multiplication constant (or inverse of the eigenvalue) which makes possible a solution for the homogeneous diffusion equation: because the equation is homogeneous, a solution does not exist for just any value of k_{eff}.

The "system reactivity" is defined as $\rho \equiv 1 - \frac{1}{k_{eff}}$. This normally has almost the same value as (k_{eff} - 1), except if k_{eff} is far from 1. ρ is usually the parameter taken as the quantitative measure of departure from criticality for the reactor with the exact set of cross-section values specified in the equations (i.e., in the model). $\rho = 0$ for an exactly critical reactor, > 0 for a supercritical reactor, and < 0 for a subcritical reactor.

2.2 Typical RFSP Reactor Model

A typical reactor model used with RFSP is shown in Figures 2.1 and 2.2 (Face and Top views respectively). The total number of mesh points in typical full-core models runs to 10,000, sometimes perhaps to 30,000.

Any RFSP calculation must start with a determination of the specific lattice properties at each fuel-bundle position and of the reflector properties, and the superposition of device incremental cross sections. Thus RFSP must at all times manipulate a very large database of basic model data.

The diffusion equation can then be solved only by iterative methods (e.g., successive over-relaxation in RFSP). The iterative solution, which commonly takes several hundred iterations to converge, provides both the full 3-d flux distribution in the reactor and the k_{eff} value.

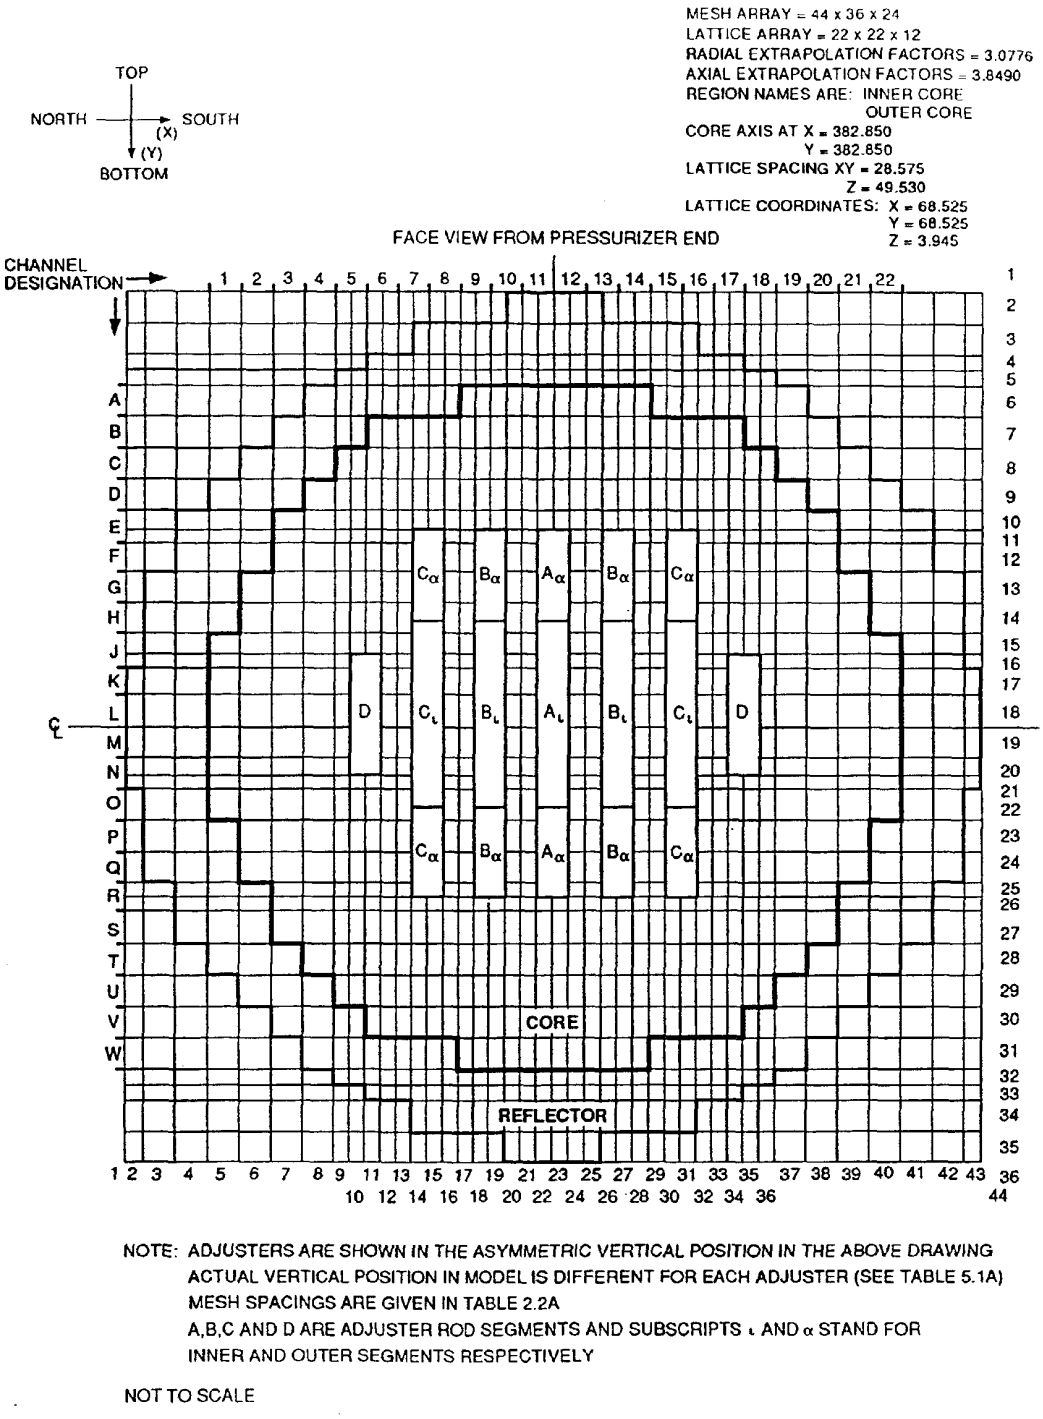

Figure 2.1 - Face View of Typical RFSP Reactor Model

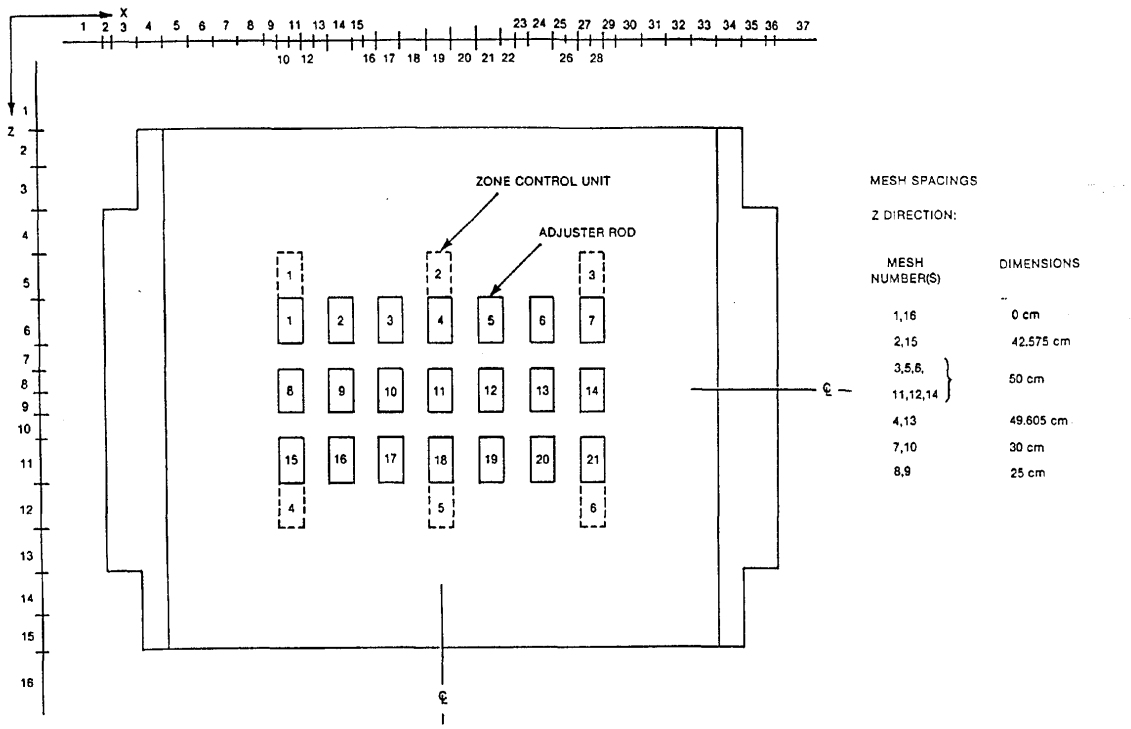

Figure 2.2 - Top View of Typical RFSP Reactor Model

2.3 General Applications of RFSP

Major applications of RFSP are in:

- core-design calculations and analyses, including the calculation of flux distributions for various reactor (or reactivity-device) configurations
- fuel-management design calculations, including the design of the time-average reactor "picture" or model
- core-follow of "virtual" reactor power histories
- simulations of isolated, "random" snapshots of a reactor history
- core-follow calculations at CANDU sites, to track the actual reactor operating history, with all burnup steps and channel refuellings modelled
- calculations in support of safety analysis; these can be time-independent calculations, or time-dependent (neutron-kinetics) calculations, such as those required for loss-of-coolant accidents [Note: in neutron-kinetics problems the time-dependent diffusion equation is the one which must solved; this has terms with time derivatives and explicit terms for delayed neutrons, and is not an eigenvalue problem – more in a different session].

Additional capabilities of the program include, among others:

- the simulation of the 3-d steady-sate distribution of ^{135}Xe in the core, and of ^{135}Xe /^{135}I transients in the case of power variations

- the capability for simulating (quasi-statically) bulk control and spatial control, i.e., the prediction of the values at which zone-compartment water fills will stabilize corresponding to a given reactor configuration
- the calculation of the reactivity increase expected on refuelling of individual fuel channels
- the calculation of harmonic flux shapes for use in flux mapping.

2.4 Inter-Code Communication

It is clear that RFSP communicates with the other codes in the physics suite, i.e., the lattice code and the reactivity-device code, since it needs to determine the properties at each model mesh point.

It is usual to think of this as a 1-way communication **from** the lattice and device codes **to** the full-core code, but it is not always so. For instance, lattice properties of a fuel bundle **do** depend on the absolute value of the bundle flux or power: therefore, in principle, once the full-core calculation is performed and the 3-d power distribution is obtained, the power values should be fed back to the lattice code to recalculate each bundle's nuclear cross sections. Such iteration **back from** the full-core code **to** the lattice code is indeed carried out when self-consistency is desired in the physics calculation.

But there is more than that. Inter-code communication is not limited to the physics suite. Other self-consistencies may be important.

For example, lattice properties depend on the coolant density (shades of the void reactivity!). Therefore, where there is boiling in a channel, the lattice properties should account for the reduced value of coolant density in specific bundles. Thus, the 3-d distribution of coolant density in the core, calculated by a thermalhydraulics code, should in principle be fed to the lattice calculations. But wait. The coolant density distribution in a channel may (will!) depend on the bundle-power distribution in the channel. Therefore, the thermalhydraulics calculation cannot be done in isolation: an "outer" iteration between RFSP and the thermalhydraulics code is required for self-consistency. This is in fact done for some applications, e.g., "ROP" (Regional Overpower Protection) calculations, where "Critical Channel Powers" (important in safety analysis) are to be determined for a large number of reactor configurations.

There may be any number of other instances when codes must communicate. For example, if moderator temperature and therefore density vary over the core, this needs to be communicated to the physics codes. Also, where "fuel-channel" codes are used to predict distributed values of radial creep in the core, this information could/should be used in the physics computations, etc.

2.5 Some Types of Calculations with RFSP

Examples of calculations which are performed with RFSP:

2.5.1 Changes in Reactivity-Device Configurations

In this type of calculation, RFSP is used to calculate the reactor multiplication constant and the core flux distribution for reference and perturbed reactivity-device configurations, for example:

- adjusters in core and then adjusters withdrawn from core, or
- shutoff rods in parked position (out of core), then fully inserted or half inserted, or
- zone controller water fills at 20% and then at 80%.

From such runs the "reactivity worth" of a device or a set of devices is obtained, by subtraction of the reactivity values corresponding to the device withdrawn and in-core (or vice versa).

In addition, the perturbation to the flux distribution (and, by extension, to the core power distribution) by the movement of the device is a very important quantitative output of such calculations. It is very important to know whether the movement of the device(s) will cause a flux "tilt", by how much, and in which direction.

2.5.2 Reactivity Coefficients

In this type of calculation, RFSP is used to calculate the derivative of system reactivity with respect to an overall parameter, such as moderator temperature, fuel temperature, moderator-poison concentration: $\left(\frac{d\rho}{dT_M}, \frac{d\rho}{dT_f}, \frac{d\rho}{dp\,oison} \right)$ respectively.

This is the reactivity coefficient as it <u>should</u> be calculated, i.e., as a full-core property. In some of our exercises we have approximated this coefficient by using k_{eff} instead of ρ and by using PPV results (admittedly, with the reactor-rate option) as surrogates for full-core results.

2.5.3 Changes in Fuel Configuration

In this type of calculation, a change in reactor configuration is imagined. For example, depleted-uranium fuel may be placed in some locations in the core. Then the consequent changes in reactivity and in flux and power distribution is determined. Again, it is important to know the size of these changes, to ensure as much as possible that the Reactor Regulating System (RRS) and protective systems will function as predicted.

3. General Discussion on CANDU Fuel Management

3.1 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 fuel management 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 3.1 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.

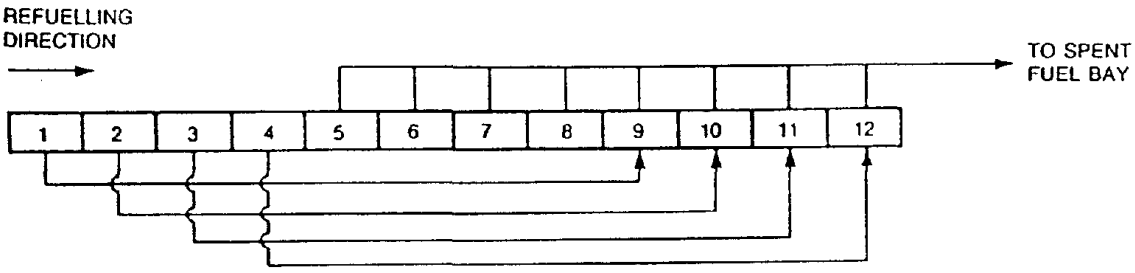

Figure 3.1
8-Bundle-Shift Refuelling Scheme

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.

3.2 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.

3.3 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). 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.

3.4 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. 3.2 shows the k_∞ as a function of irradiation for the standard CANDU 6 lattice fuelled with natural uranium.

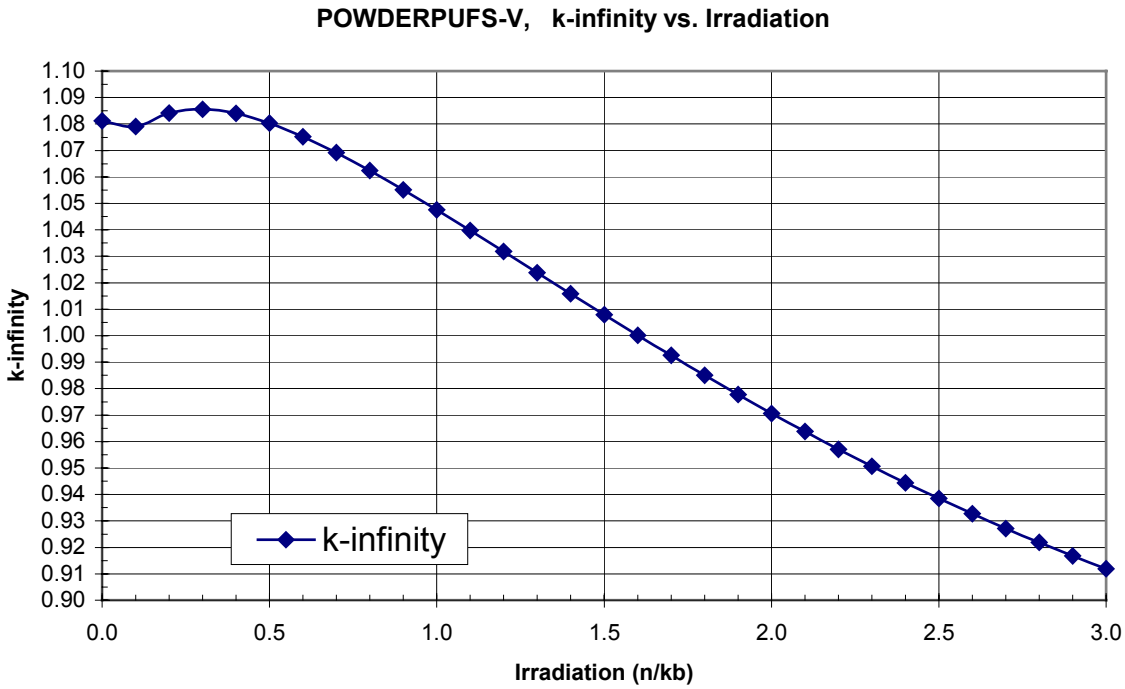

Figure 3.2
k_∞ as a Function of Irradiation

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.

The infinite-lattice multiplication constant k_∞ reaches the "magic" value of 1.050 (corresponding roughly to a critical reactor) at an irradiation of approximately 0.9-1.0 n/kb. Beyond this value of irradiation, the fuel becomes a net absorber of neutrons. It can still remain in the reactor, since there is other, younger fuel in the core, on account of on-power refuelling.

Thus, in fact, fuel need not be discharged from the reactor until it reaches an irradiation such that the k_∞ averaged from 0 to that irradiation is about 1.050. This average "discharge" irradiation is typically 1.7-1.8 n/kb in the CANDU 6. The corresponding average discharge or "exit" burnup is about 170-180 MW.h/kg(U).

Figures 3.3-3.5 show the basic-lattice cross sections, which "make up" k_∞, as a function of irradiation. It is clear that only two cross sections show significant variation with irradiation: the thermal fission cross section Σ_{f2}(on account of the depletion of fissile material) and the thermal absorption cross section Σ_{a2} (on account of the accumulation of fission products).

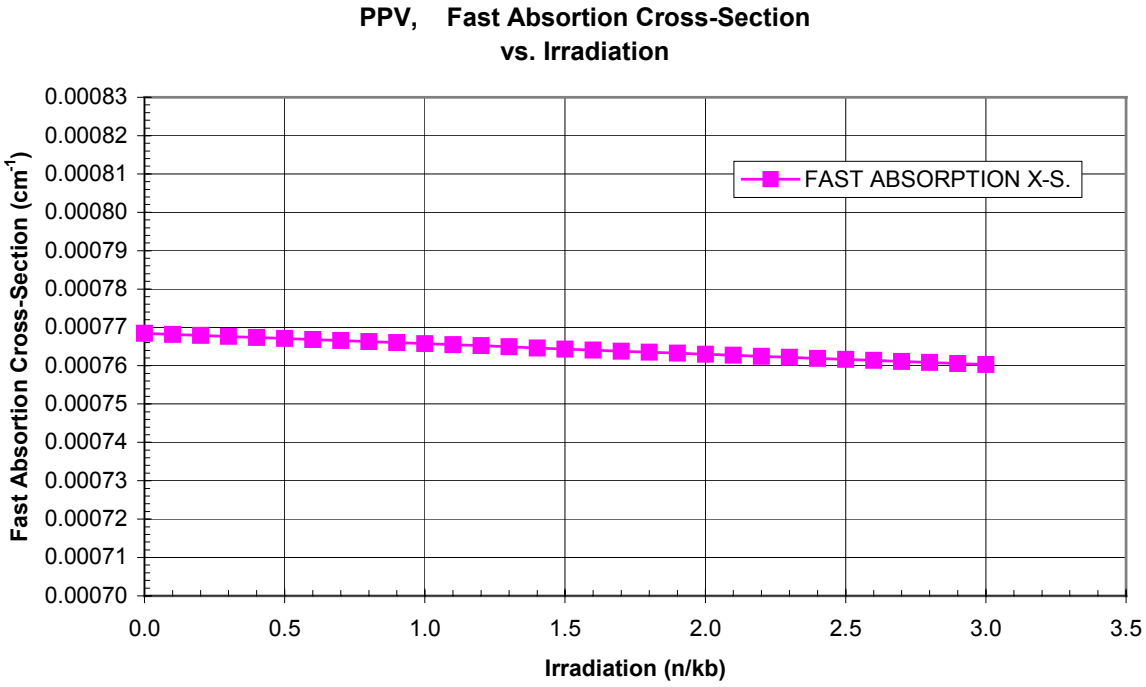

Figure 3.3
Fast-Group Absorption Cross Section (Σ_{a1}) as a Function of Irradiation

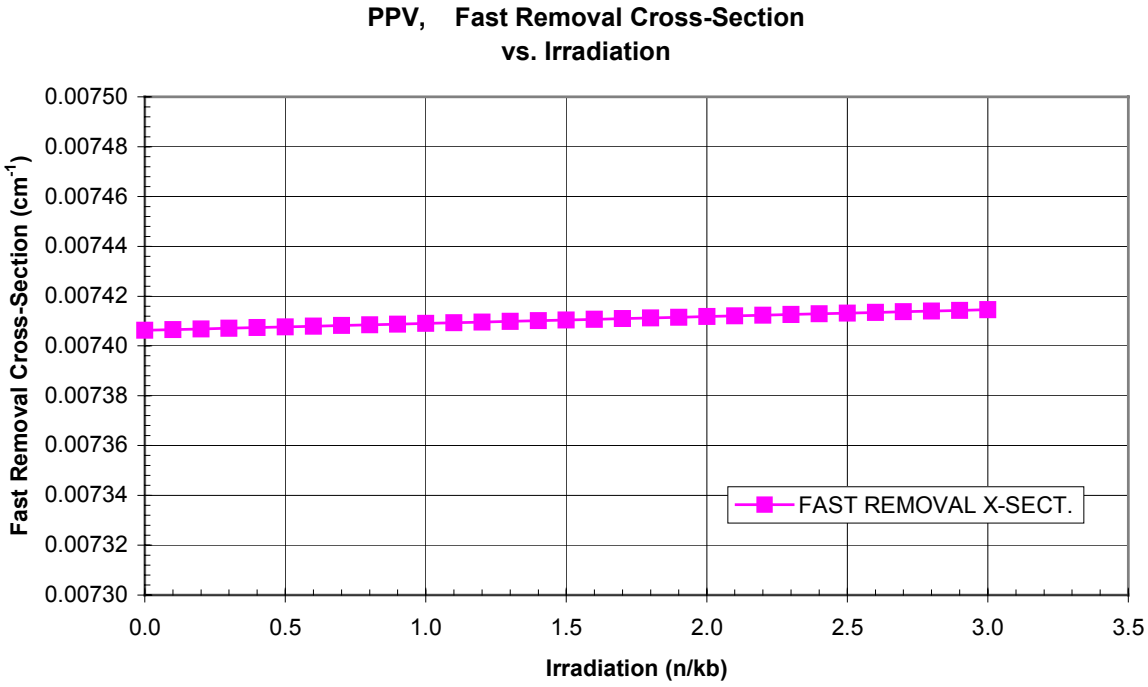

Figure 3.4
Fast-Group Slowing-Down (Moderation) Cross Section (Σ_m) as a Function of Irradiation

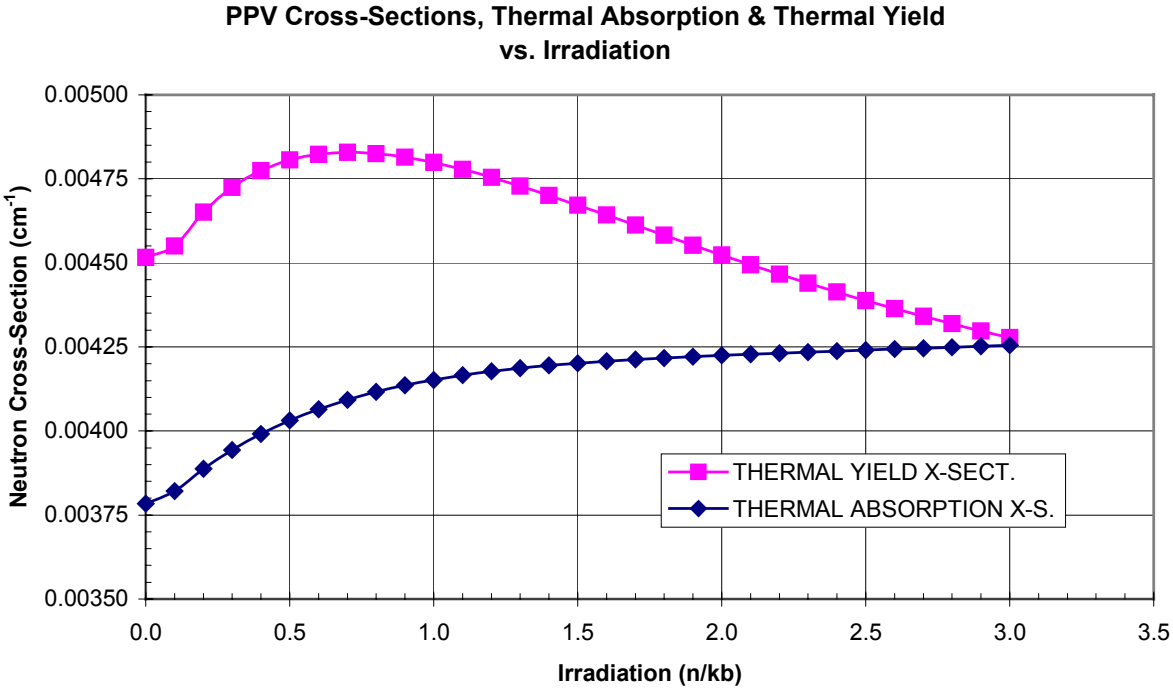

Figure 3.5
Thermal-Group Yield Cross Section ($\nu\Sigma_{f2}$) & Thermal-Group Absorption Cross Section (Σ_{a2}) As Functions of Irradiation

It is instructive to examine also the infinite-lattice multiplication constant for the depleted-uranium lattice. The infinite-lattice reactivity is shown in Figure 3.6 for depleted uranium with an initial fissile content of 0.52 atom % (as opposed to 0.72 atom % for natural uranium).

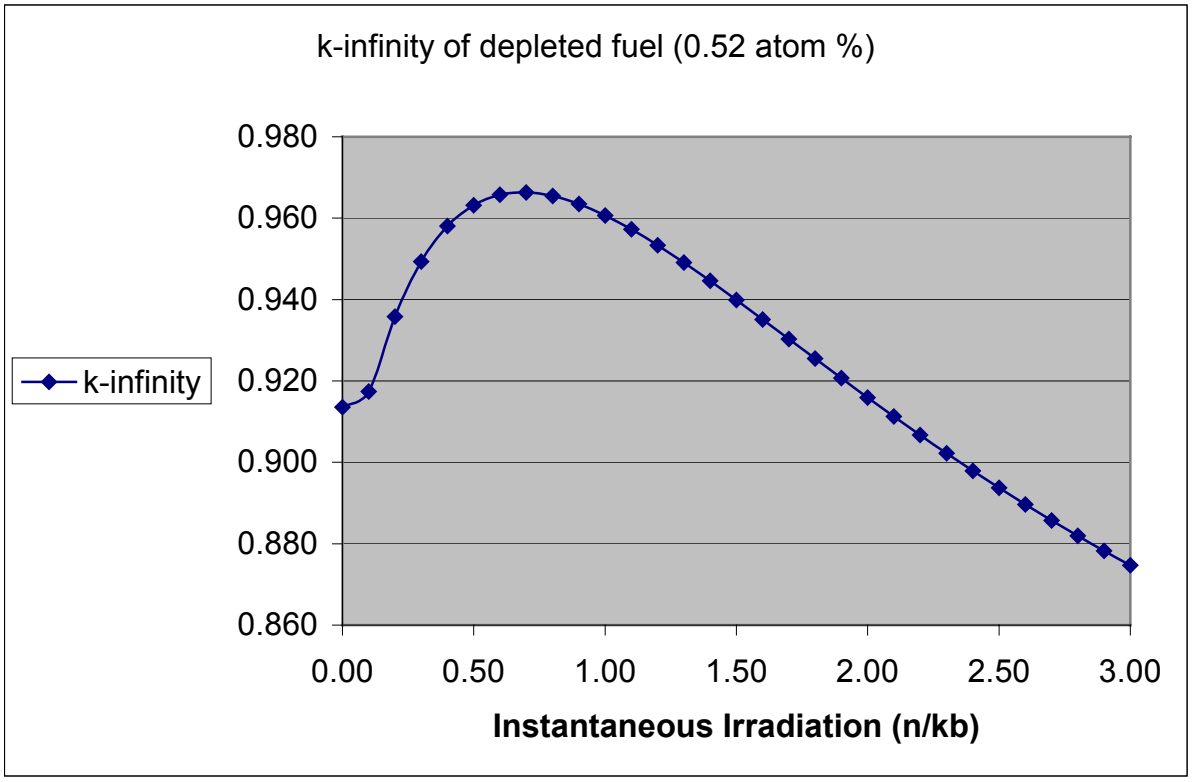

Figure 3.6 - Infinite-Lattice Excess Reactivity of CANDU for Depleted-Uranium Fuel

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. it 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.

3.5 Isotopic Densities

The lattice code solves the nuclide-chain equations as a function of burnup; i.e., it calculates the depletion or build-up of the various nuclides in the fuel as irradiation or burnup proceeds. It is instructive to see how the "isotopic densities" of the various nuclides evolve. This is shown in Figure 3.7.

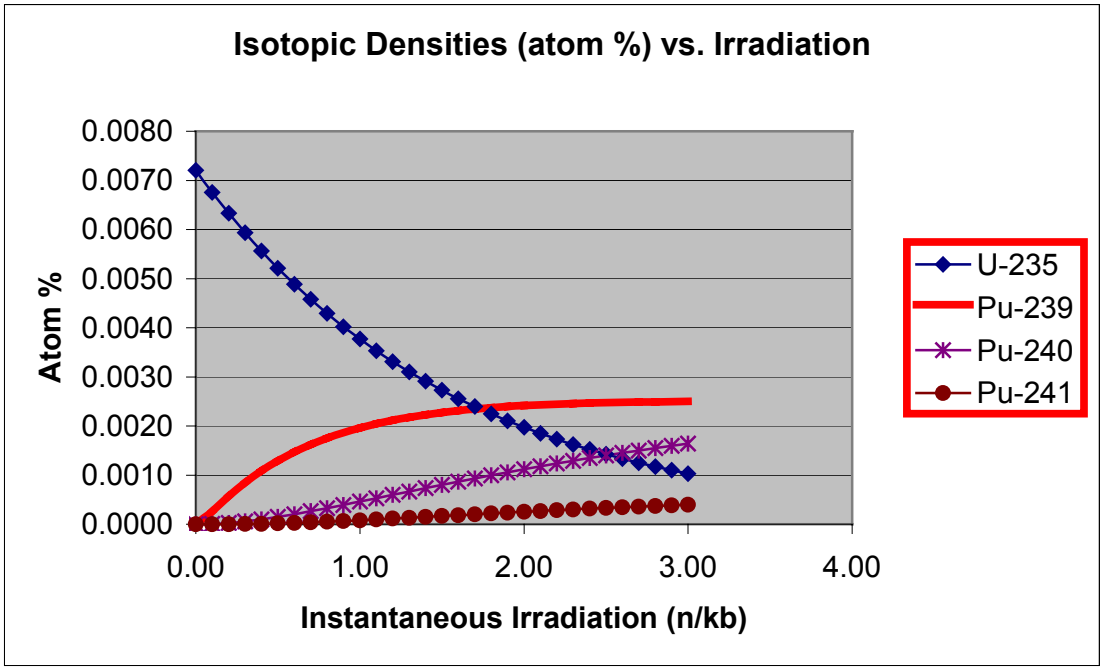

Figure 3.7 Evolution of Isotopic Densities of Fuel Nuclides vs. Irradiation

The figure shows how U-235 depletes and how Pu-239 builds up to a near asymptotic value as the burnup increases. Pu-240 and Pu-241 also increase. At the average exit irradiation (1.8 n/kb, burnup ~ 180 MW.h/kg(U)), the nuclide densities of U-235 and Pu-239 are about equal, at ~0.0023. The total fissile content of the fuel (including Pu-241) at this exit irradiation is about 0.55 atom %.

4. Equilibrium-Core Design

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 doing time averages of quantities over simulations of a long period of reactor operation, this is not very practical.

Instead, a "time-average" model is developed to analyze the nominal equilibrium core to a good approximation. This model will be described in detail following the discussion of power flattening.

4.1 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, whose main purpose is in xenon override. Adjuster rods also provide axial power flattening.

Note that while flattening of the power distribution reduces 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.

4.2 Time-Average Model

This is a model in which the lattice properties at each bundle location are not determined from an instantaneous or snapshot burnup. Instead, the properties at each location are averaged over the range of irradiations experienced by the fuel during the interval of time it resides at that location between refuellings (called the "dwell" time for that location).

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 – see Figure 4.1.

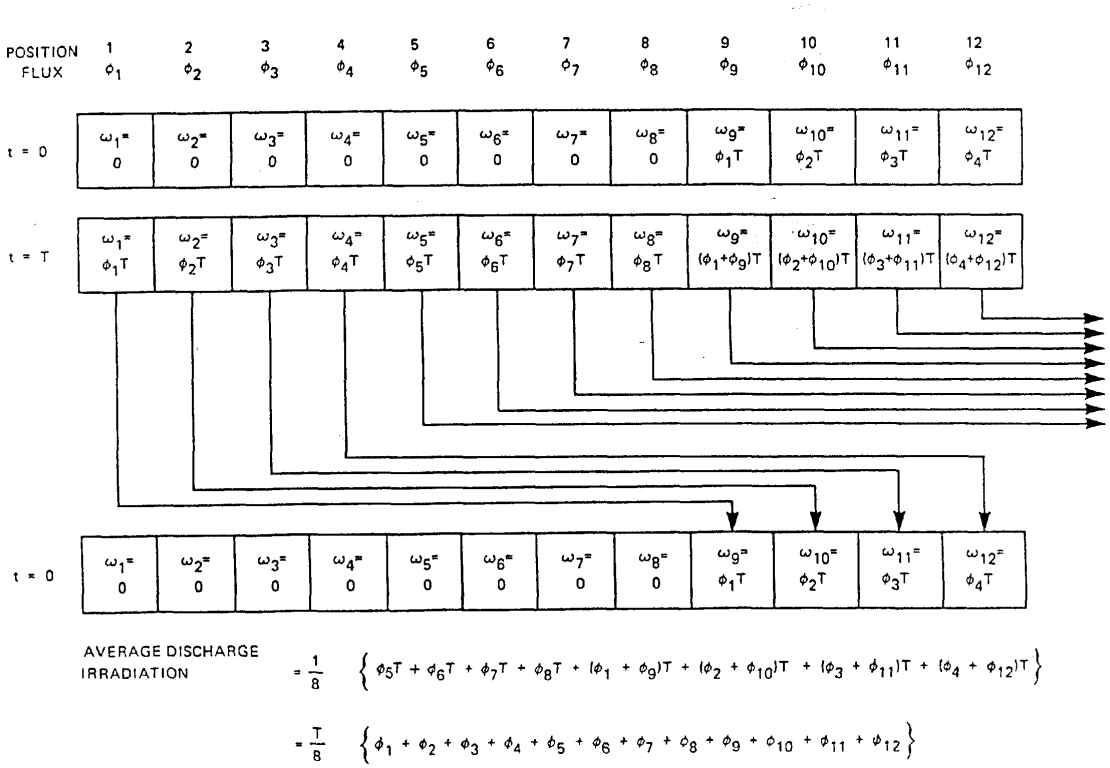

Figure 4.1 – Bundle Irradiations During a Refuelling Cycle for an 8-Bundle Shift

Time-average nuclear cross sections are defined at each bundle position in core by averaging the lattice cross sections over the irradiation range [ω_{in}, ω_{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. The irradiation range the fuel experiences is then ω_{in} to ω_{out}.

The time-average thermal neutron absorption cross section (for example) at some core position r, $\sum_{a2}^{t.a.}(r)$, is then

$$\Sigma_{a2}^{t.a.}(r) = \frac{1}{(\omega_{out} - \omega_{in})} \int_{\omega_{in}}^{\omega_{out}} \Sigma_{a2}(\omega)\, d\omega \qquad (4.1)$$

where the basic lattice cross sections inside the integral sign are determined as functions of irradiation using the lattice code (e.g., POWDERPUFS-V).

Now, 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 (in most CANDU reactors), and j ranges over the channels, e.g. from 1 to 380 in the CANDU 6. The bundle position is then 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

$$\Delta\omega_{jk} = \phi_{jk}.T_j \qquad (4.2)$$

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

$$\omega_{out,jk} = \omega_{in,jk} + \Delta\omega_{jk}$$

$$= \omega_{in,jk} + \phi_{jk}.T_j \qquad (4.3)$$

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):

$$\omega_{in,jk} = 0 \qquad k = 1,...,8 \qquad (4.4a)$$

$$\omega_{in,jk} = \omega_{out,j(k-8)} \qquad k = 9,...,12 \qquad (4.4b)$$

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

$$\omega_{in,jk} = 0 \qquad k = 1,...,N \qquad (4.5a)$$

$$\omega_{in,jk} = \omega_{out,j(k-N)} \qquad k = (N+1),...,12 \qquad (4.5b)$$

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,

which we can select to our liking. 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 relationship between exit irradiations and fluxes can be derived as follows (done 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

$$\omega_{exit,j} = \frac{1}{8}\sum_{k=5}^{12} \omega_{\mathrm{out,jk}} \tag{4.6}$$

In view of Equation (4.3) this can be written

$$\omega_{exit,j} = \frac{1}{8}\sum_{k=5}^{12}\left(\omega_{\mathrm{in,jk}} + \phi_{\mathrm{jk}}\,\mathrm{T_j}\right) = \frac{1}{8}\left[\sum_{k=5}^{8}\left(\omega_{\mathrm{in,jk}} + \phi_{\mathrm{jk}}\,\mathrm{T_j}\right) + \sum_{k=9}^{12}\left(\omega_{\mathrm{in,jk}} + \phi_{\mathrm{jk}}\,\mathrm{T_j}\right)\right] \tag{4.7}$$

and in view of Equation (4.5) we can write

$$\omega_{exit,j} = \frac{1}{8}\left[\sum_{k=5}^{8}\phi_{\mathrm{jk}}\,\mathrm{T_j} + \sum_{k=1}^{4}\phi_{\mathrm{jk}}\,\mathrm{T_j} + \sum_{k=9}^{12}\phi_{\mathrm{jk}}\,\mathrm{T_j}\right] = \frac{\mathrm{T_j}}{8}\sum_{k=1}^{12}\phi_{\mathrm{jk}} \tag{4.8}$$

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

$$\omega_{exit,j} = \frac{\mathrm{T_j}}{\mathrm{N}}\sum_{k=1}^{12}\phi_{\mathrm{jk}} \tag{4.9}$$

The dwell time T_j therefore satisfies

$$T_j = \frac{\mathrm{N}\,\omega_{\mathrm{exit,j}}}{\sum_{k=1}^{12}\phi_{\mathrm{jk}}} \tag{4.10}$$

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 (4.10) to compute the dwell time for each channel,
- Equations (4.3) and (4.4) to calculate $\omega_{in,jk}$ and $\omega_{out,jk}$ for each bundle in core,

- Equation (4.1) (and similar equations for the other cross sections) to calculate the time-average lattice properties.

The degrees of freedom in this time-average problem, which can be chosen by the reactor designer, are the channel-specific axial refuelling scheme - e.g. 8-bundle-shift, 4-bundle-shift, etc…(need not be the same for all channels) - and the channel-specific exit irradiations. The latter, however, must be in a reasonable range of values if a k_{eff} of unity is to be possible. And the variation of the exit irradiations over the core (e.g., inner core vs. peripheral region) will determine the degree of radial flattening of the power distribution. A typical subdivision of the core in irradiation regions for purposes of a time-average calculation is illustrated in Figure 4.2.

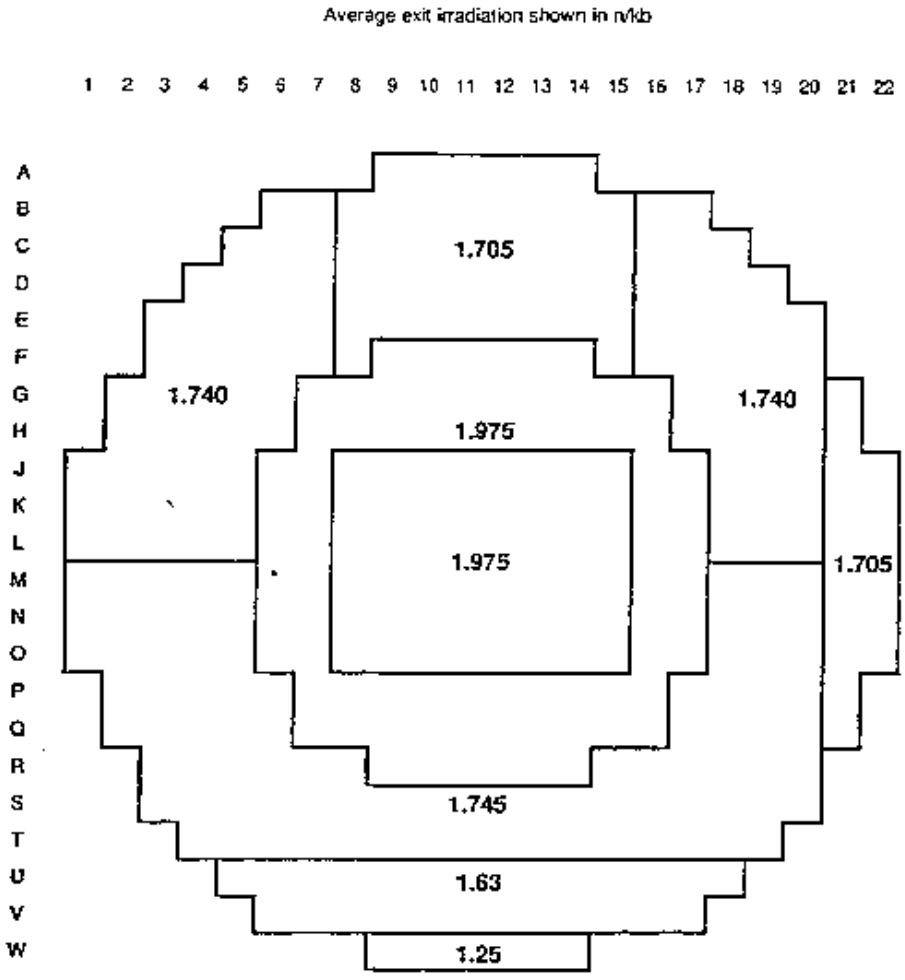

Figure 4.2 Typical Time-Average Model, Showing Irradiation Zones

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 4.3 shows the iterative scheme of calculations.

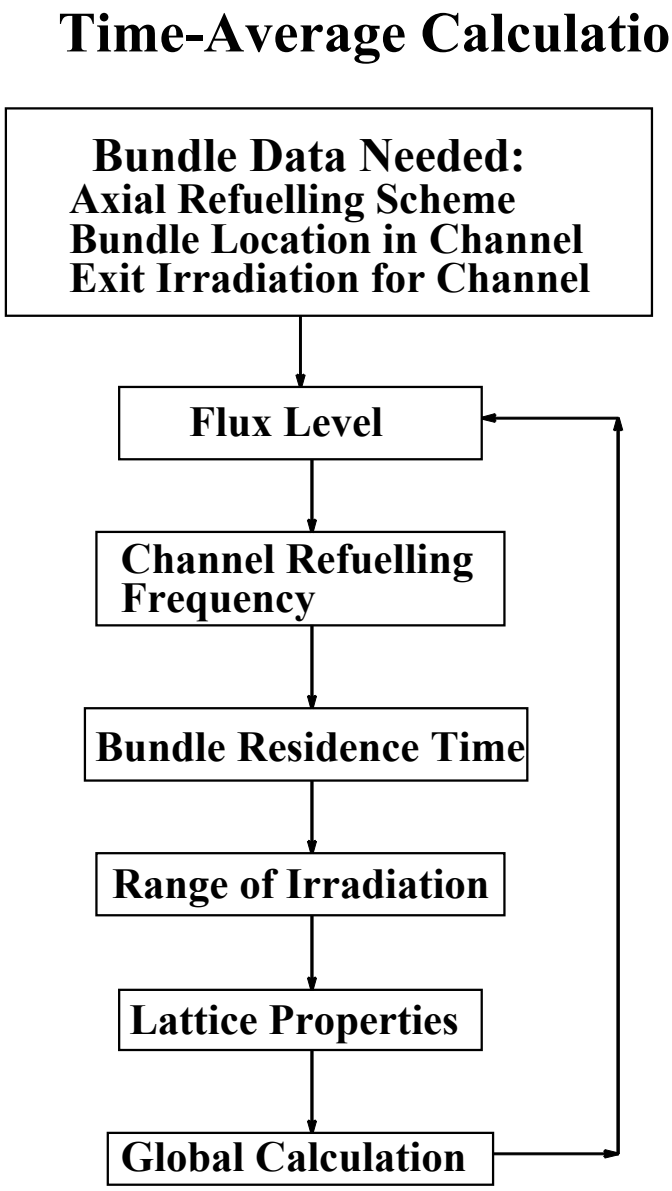

Figure 4.3 Time-Average Self-Consistency Problem

Typically, several "outer" iterations by the human user in selecting values of exit irradiation $\omega_{exit,j}$ in the various regions are required to attain a final solution giving a critical reactor and a desired flux shape (the flux shape is 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.

Once the time-average calculation is completed (with a reasonable result), the time-average power distribution becomes the target power distribution for fuel management. Also, the channel dwell times from the time-average calculation (see Figure 4.4) provide a very useful guideline for the time intervals at which specific channels should be refuelled.

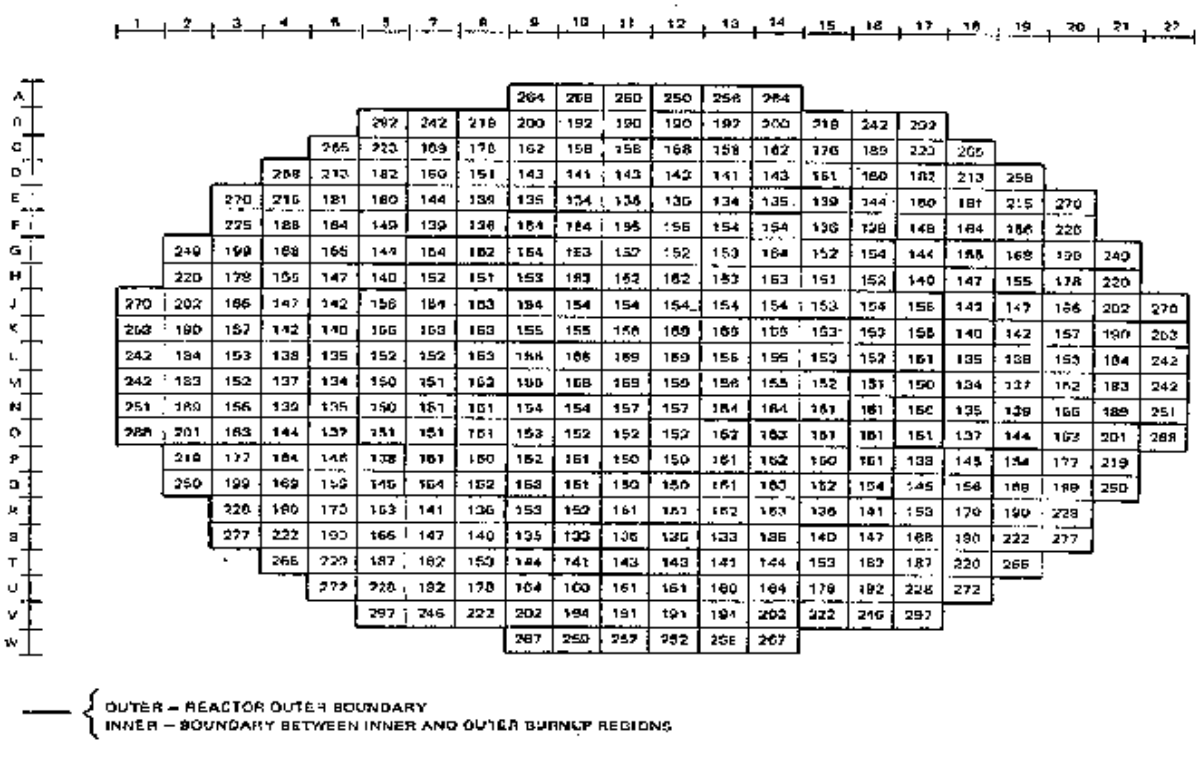

Figure 4.4 Channel Dwell Times (in Full-Power Days) from Time-Average Calculation

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.

5. 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.

5.1 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 time-average model does not exhibit refuelling ripple, because there is no "new" or "old" fuel in this model: the properties at all points are averages over the range of irradiation experienced.

In an operating reactor, with channel refuellings, the power of each channel 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 5.1, which is 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.

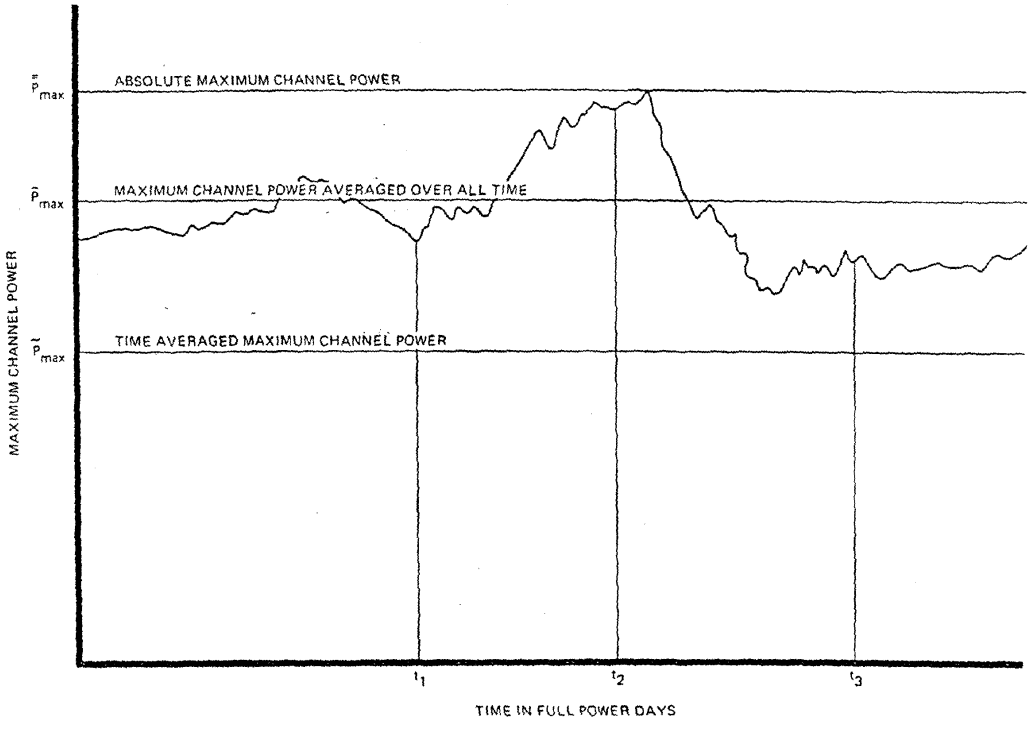

Figure 5.1 Sketch of Maximum Instantaneous and Time-Average Channel Powers

5.2 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 Figure 5.1.

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:

$$CPPF = \underset{m}{Max}\left[\frac{CP_{ins\,\tan\tan\,eous}(m)}{CP_{time-average}(m)}\right] \tag{5.1}$$

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.

5.3 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.

Figure 5.2 shows, as an example, a selection of about 60 channels selected to be refuelled within successive ~3-day intervals in a 30-day period in the life of a CANDU 6 reactor. This selection exhibits good symmetry left-right, up-down, and in-out, and a good distribution of channels in the various zones of the core (delineated by line boundaries).

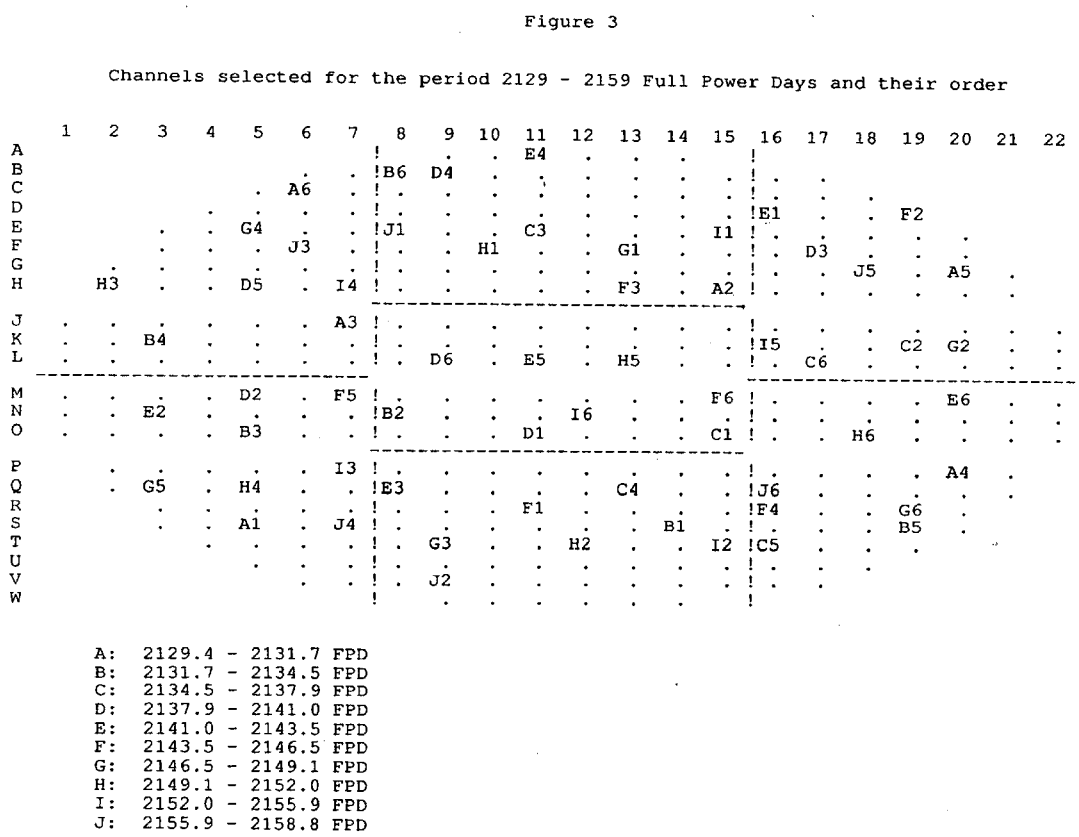

Figure 5.2 Selection of Channels For Refuelling in 3-Day Intervals

5.4 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.

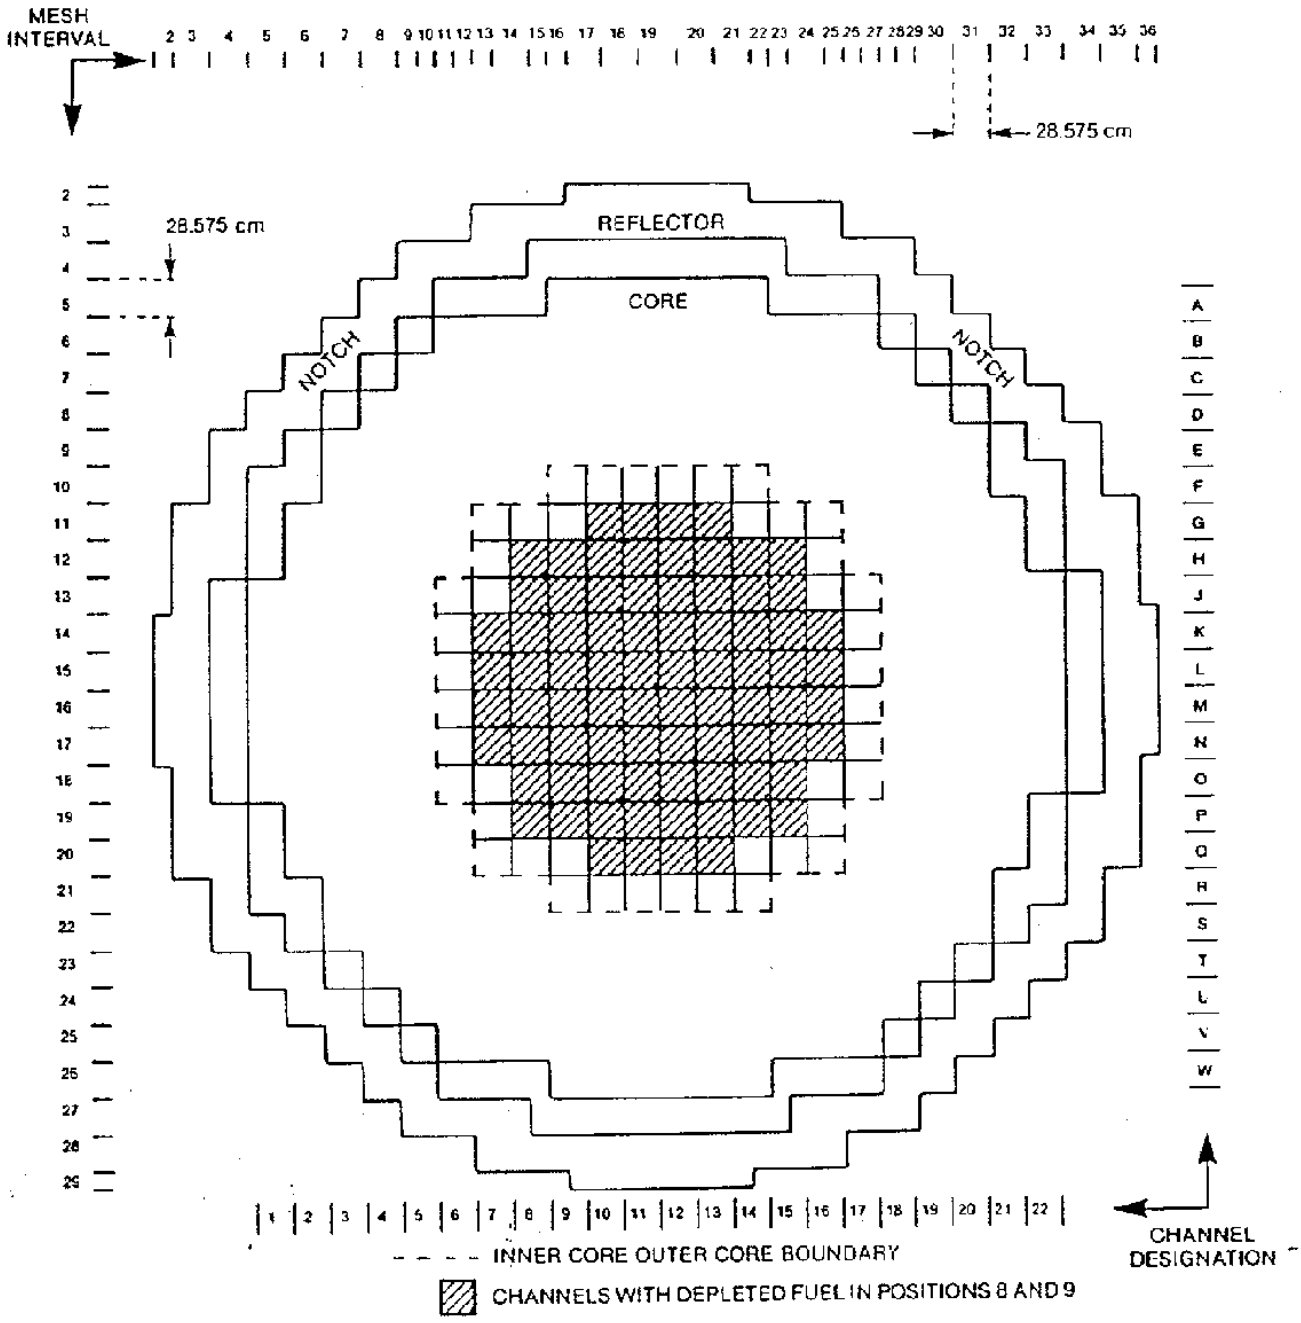

Figure 5.3 Channels with Depleted Fuel in Initial Core of CANDU 6

Even with some depleted fuel in the core, the fact that all fuel is fresh results in a net excess reactivity in the core – see Figure 5.4.

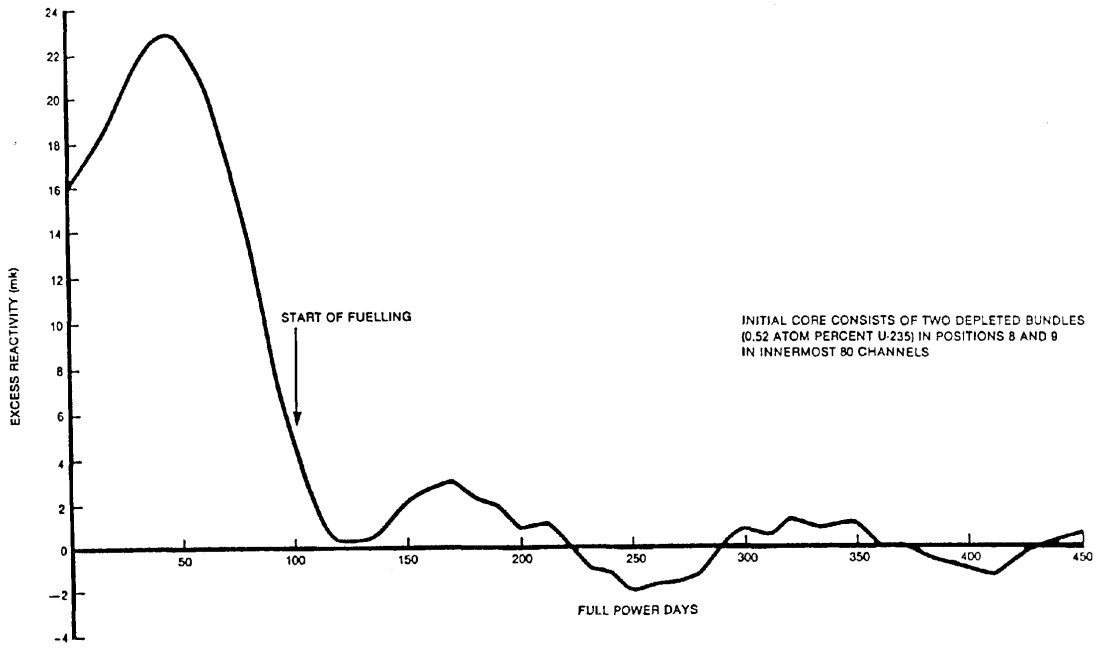

Figure 5.4 Core Excess Reactivity as a Function of Time from First Full Power

The core reactivity starts at approximately 16 milli-k at full power on FPD 0. 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.

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

5.5 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 can be a considerable variation from this average rate.

5.6 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 mk/(0.4 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.

5.7 Core-Follow Calculations with RFSP

5.7.1 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; this has an effect on the calculated flux distribution. Bulk and spatial control can also be modelled (see Section 5.7.2).

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 %.

5.7.2 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.

5.7.2.1 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 $k_{eff\text{-}ref}$.

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

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

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

5.7.2.2 Spatial Control

In the simplest approximation, 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

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

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.

6. Summary

There is a very wide variety of reactor-physics calculations. They touch all aspects of reactor design, analysis, and operation. The physics calculations require an entire suite of specialized computer codes, ending with the full-core code. The physics codes exchange data among themselves and with codes of other disciplines, especially thermalhydraulics. Many analyses in fact require self-consistency between all components of the calculation.

One specific application of reactor-physics codes is in fuel management. 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.