

CANDU Fuel Management

by

Shao-hong Zhang⁺ and B. Rouben^{*}

⁺Xi'an Jiaotong University

^{*}Atomic Energy of Canada, Limited

Abstract

CANDU fuel management is described, with specific application to CANDU 6 reactors, of the Qinshan CANDU type. 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 while shut down. The primary objective of CANDU 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. At the design stage, fuel management involves the design of the time-average flux/power distribution. Operations-related responsibilities of the site fuelling engineer include selecting channels for refuelling, tracking the reactor power history, and ensuring licensed maximum powers are respected.

1. Introduction

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 while shut down.

To refuel a channel, a pair of fuelling machines are latched 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.

For example, in the 8-bundle-shift fuelling scheme, 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. In this scheme, the four low-power bundles are in-core for two cycles, while the high-power bundles are in-core for only one cycle.

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

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

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

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

Fuel management in CANDU has both design and operations aspects. Both types of activities will be discussed here.

2. Overall Objectives of Fuel Management

The primary objective of CANDU 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:

- Adjust the refuelling rate to maintain the reactor critical
- Control the core power to satisfy safety and operational limits on fuel 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.

3. Design Aspects of Fuel Management

3.1 Equilibrium Core

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, 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. A number of refuelling operations are carried out essentially every day. As a consequence, the equilibrium core contains fuel at a range of burnups, from 0 to exit-burnup values. The refuelling of individual channels generates local “refuelling ripples” about the time-average shape. These ripples are due to the various instantaneous values of fuel burnup in the different channels, due at any given instant to the specific sequence of channels refuelled.

The main design activity relating to fuel management is establishing the time-average power distribution in the equilibrium core. During actual reactor operation, this time-average distribution is used as the power target by the fuelling engineer at site. The next section describes the calculational design of the time-average core.

3.2 The Time-Average Model

Calculations for the time-average model are performed in the *TIME-AVER module of the fuel-management code RFSP (Reactor Fuelling Simulation Program)^{[1],[2]}. In this model, the lattice cross sections are averaged over the residence (dwell) time of the fuel at each point (fuel-bundle position) in the core. This allows the effect of the actual refuelling scheme used to be captured. The mathematical framework is described for the specific case of an 8-bundle-shift refuelling scheme as an example.

The concept of fuel irradiation needs to be discussed, as it is used in this model. Irradiation ϕ is defined as the product of Westcott flux^[1] in the fuel (which we label as ϕ) and time t : $\phi = \phi t$. Irradiation is measured in units of neutrons/kilobarn (n/kb).

Time-average nuclear cross sections are defined at each bundle position in core by averaging the lattice cross sections over the irradiation range $[\phi_{in}, \phi_{out}]$ “experienced” over time by fuel at that position, where ϕ_{in} is the value of fuel irradiation when the fuel enters that position in core and ϕ_{out} is the fuel irradiation when the fuel leaves that position. The cross section defined is the one that preserves the average reaction rate over the irradiation range ϕ_{in} to ϕ_{out} . For example, the time-average thermal neutron absorption cross section at some core position r , $\sigma_{a2}^{t.a.}(r)$, is

$$\sigma_{a2}^{t.a.}(r) = \frac{1}{(\phi_{out} - \phi_{in})} \int_{\phi_{in}}^{\phi_{out}} \sigma_{a2}(\phi) d\phi \quad (3.1)$$

The basic lattice cross sections inside the integral sign are determined as functions of irradiation using the cell code POWDERPUFS-V^{[1],[3]}, incorporated in its entirety within RFSP as the *POWDERPUF module.

Let ϕ_{jk} be the time-average fuel flux at axial position k in channel j (k ranges from 1 to 12 since there are 12 bundles per channel, and j is a channel index, 1 to 380 in the CANDU 6. The bundle position is labelled jk for short. Let also T_j be the average time interval between refuellings of channel j (also known as the *dwell time* of channel j).

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

$$\Delta\phi_{jk} = \Delta_j \cdot T_j \quad (3.2)$$

If the fuel entered position jk with an irradiation $\phi_{in,jk}$, then its exit irradiation from that position, $\phi_{out,jk}$, will be given by initial value plus increment, i.e.

$$\phi_{out,jk} = \phi_{in,jk} + \Delta\phi_{jk} = \phi_{in,jk} + \Delta_j \cdot T_j \quad (3.3)$$

When a channel is refuelled with an 8-bundle shift, the first 8 positions in the channel receive fresh fuel and the entrance irradiances for positions 9-12 are simply the exit irradiances from positions 1-4 respectively. Thus we can write (see Figure 3.1):

$$\phi_{in,jk} = 0 \quad k = 1, \dots, 8 \quad (3.4a)$$

$$\phi_{in,jk} = \phi_{out,j(k-8)} \quad k = 9, \dots, 12 \quad (3.4b)$$

In addition to the refuelling scheme, there are other degrees of freedom in the time-average model. These are the values of exit irradiation $\phi_{exit,j}$ for the various channels j , which are the target values of average irradiation for fuel as it exits from these channels. 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 $\phi_{exit,j}$ can be used to “shape” the flux to a desired reference distribution. The exit irradiances are related to the flux in the following way. Since, in the 8-bundle-shift refuelling scheme, bundles 5 to 12 leave the core at each refuelling, then by definition of exit irradiation

$$\phi_{exit,j} = \frac{1}{8} \sum_{k=5}^{12} \phi_{out,jk} \quad (3.5)$$

In view of Eq. (3.3) this can be written

$$\phi_{exit,j} = \frac{1}{8} \sum_{k=5}^{12} (\phi_{in,jk} + \Delta_j T_j) = \frac{1}{8} \sum_{k=5}^8 (\phi_{in,jk} + \Delta_j T_j) + \sum_{k=9}^{12} (\phi_{in,jk} + \Delta_j T_j) \quad (3.6)$$

and in view of Eqs.(3.4b) and (3.3) we can write

$$\phi_{exit,j} = \frac{1}{8} \sum_{k=5}^8 \Delta_j T_j + \sum_{k=1}^4 \Delta_j T_j + \sum_{k=9}^{12} \Delta_j T_j = \frac{T_j}{8} \sum_{k=1}^{12} \Delta_j \quad (3.7)$$

The dwell time T_j therefore satisfies

$$T_j = \frac{8\phi_{exit,j}}{12 \sum_{k=1} \phi_{jk}} \quad (3.8)$$

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 (3.8) to compute the dwell time T_j for each channel,
- Equation (3.3) and (3.4) to calculate $\phi_{in,jk}$ and $\phi_{out,jk}$ for each bundle in core,
- Equation (3.1) (and similar equations for the other cross sections) to calculate the time-average lattice properties.

This set of equations must be solved using as input the user-specified target exit irradiations $\phi_{exit,j}$. In order to shape the flux to desired values, and also to take account of the presence of extra “hardware” (device locators, etc., mostly at the bottom of the calandria) that introduces localized absorption, time-average RFSP models may use many irradiation regions; see Figure 3.2.

Since consistency must be achieved between the flux, the channel dwell times, the individual-bundle irradiation ranges $[\phi_{in}, \phi_{out}]$, and the lattice properties, an iterative scheme between the solution of the diffusion equation and of the other equations is employed until all these quantities converge. Typically, the user will also need to iterate on the values of exit irradiation $\phi_{exit,j}$ in the various regions to obtain a critical reactor and the desired flux shape.

The time-average model is useful at the reactor design stage. It determines the reference (target) 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. Dwell times range typically between 150 and 160 FPD in the inner core. In the outer core, they present a large variation, from ~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.

The calculation also provides an estimate of the decay rate of core reactivity if refuelling were not performed; this is typically 0.4 milli-k/FPD in the CANDU 6.

By judicious choice of the relative exit irradiation values, the CANDU designer can achieve essentially any desired global power distribution in the time-average model. The

most common parameter used to characterize the power distribution is the degree of radial flattening; this is measured by the radial form factor, defined as the ratio of average to maximum channel power. By *differential fuelling*, i.e. taking the inner core region to a higher value of exit irradiation (or burnup) than peripheral core regions, the degree of radial flattening can be controlled. For example, for the CANDU 6, a value typical chosen for the radial form factor is 0.83. Since the total power is 2061.4 MWt from 380 channels, this value of radial form factor is equivalent to a maximum time-average channel power of 6.54 MW.

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. The refuelling ripple is superimposed on time-average powers to give instantaneous powers. A typical refuelling ripple of 8-10 % in the CANDU 6 provides margin to the maximum licensed channel power, currently 7.3 MW.

3.3 Depleted Fuel in Initial Core

Another design aspect of fuel management is the use of depleted fuel in the CANDU initial core. This is the only time that fresh fuel is present throughout the core. There is no differential burnup that 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. At the design stage, the number and location of depleted-fuel bundles are selected to provide appropriate power flattening and a smooth transition to the equilibrium core as burnup proceeds, channels start to be refuelled, and depleted-fuel bundles exit from the core.

In the standard CANDU-6 initial core, two depleted-fuel bundles (of 0.52 atom percent ^{235}U content) are placed in each of the innermost 80 fuel channels. 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.

4. Operational Aspects of Fuel Management

4.1 From Initial Core to Onset of Refuelling

Even with some depleted fuel in the initial core, the core reactivity at full power starts at approximately 16 milli-k on FPD 0. The entire initial fuel load goes through the plutonium peak at about the same time: at about 40-50 FPD, the core reaches its global “plutonium peak”, at which time the core reactivity is highest (~23 milli-k), 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.

The initial-core excess reactivity is compensated by soluble boron in the moderator. The boron concentration at full power is ~2 ppm at FPD 0, rising to ~3 ppm at the plutonium peak. Following the plutonium peak, boron 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. 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 is started; it is best not to wait until excess reactivity is exactly 0, because the initial refuelling rate would prove too high.

After the initial period of reactor operation, on-power refuelling is 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 (approximately 16 bundles per FPD in the CANDU 6).

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.

4.2 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. The local reactivity **increases** for ~ 40-50 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 then continue to drop. Eventually, the channel becomes a net “sink” (absorber) of neutrons, and the time approaches 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.

The power of each channel therefore goes through an “oscillation” about the time-average power during every cycle. This cycle repeats every time the channel is refuelled, that is, with a period approximately equal to the dwell time suggested by the time-average model. The cycle length is not **exactly** equal to the dwell time, because channels are not refuelled in a rigorously defined sequence. Instead, 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.

4.3 Channel-Power Peaking Factor

At any given time there are several channels in the core that 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.

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

$$CPPF = \underset{m}{Max} \frac{CP_{instantaneous}(m)}{CP_{time\text{-}average}(m)} \quad (4.1)$$

where m 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 ROP (safety) detectors. The hundreds of flux shapes that are used in the ROP analysis (to determine the detector positions and setpoints) 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” (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 operating margin, it is important to keep the CPPF as low as possible. This is why 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.

4.4 Criteria for Selecting Channels for Refuelling

One of the main functions of the fuelling engineer 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’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%) .

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 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.

4.5 Core-Follow Calculations with RFSP

4.5.1 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. Asymptotic bulk and spatial control can also be modelled.

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

The site reactor physicist can also elect to do core tracking using the flux- mapping method in RFSP. In this case, the detector fluxes at the time of the snapshot are input to the code to derive the 3-d mapped flux distribution. This is used to advance the irradiation and burnup distributions from one snapshot to the next. Even in this option, the diffusion calculation is performed in any case, because results are optimized when the diffusion solution is used as the fundamental mode in the mapping process.

4.5.2 History-Based Methodology for Lattice Properties

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

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

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

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

5. Summary

Fuel management in CANDU has both design and operations aspects.

The design component consists of establishing:

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

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

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

- Core-follow calculations, typically performed 2 or 3 times per week to keep close track of the in-core flux, power, and burnup distributions and of the discharge burnup of individual bundles,

- The selection of channels for refuelling, based on the current core state, power and burnup distributions and zone-control-compartment water fills, and
- The determination of the CPPF (channel-power-peaking factor) value, used as a calibration factor for the ROP detectors.

The site reactor physicist's tasks keep the job interesting and stimulating.

6. References

1. H.C. Chow, "Physics Codes and Methods Used at AECL for CANDU Reactor Design and Analysis", companion paper to present paper, CJNPE, 1999...
2. B. Rouben, "Overview of Current RFSP-Code Capabilities for CANDU Core Analysis", AECL Report, AECL-11407, 1996 January.
3. B. Rouben, "Description of the Lattice Code POWDERPUFS-V", AECL Report, AECL-11357, 1995 October.

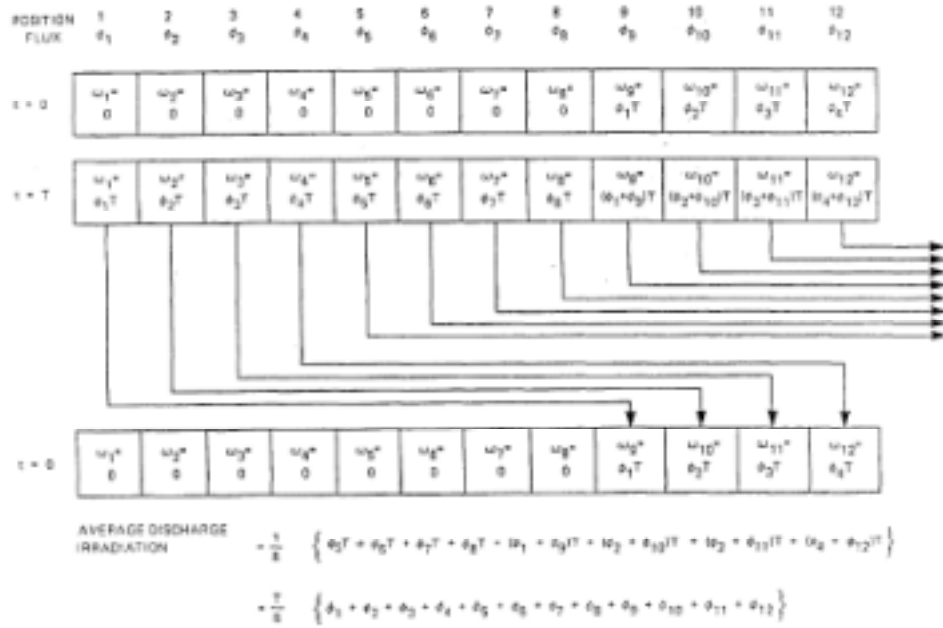


Fig. 3.1 Beginning and End of Cycle for 8-Bundle-Shift Refuelling




Fig. 3.2 Multiple-Burnup-Region Model for CANDU 6