CANDU Fuel Management

by

B. Rouben

Manager, Reactor Core Physics Branch

Atomic Energy of Canada Limited

Presentation at McMaster University
2003 November 10

TABLE OF CONTENTS

1. The CANDU Reactor

Figure 1.1 is a schematic face view of a CANDU reactor, more specifically the CANDU 6. The CANDU design is modular, with fuel channels set on a square lattice of lattice pitch equal to 28.575 cm. Each fuel channel contains 12 fuel bundles; Figure 1.2 shows a 37-element fuel bundle.

The basic building block of the CANDU design is the basic lattice cell, of dimensions 1 lattice pitch by 1 lattice pitch by 1 fuel-bundle length (28.575 cm x 28.575 cm x 49.53 cm). This basic lattice cell, shown in face view in Figure 1.3, has as components:

- a 1-bundle length of pressure tube, with fuel bundle and heavy-water coolant,
- calandria tube concentric with the pressure tube and separated from it by an isolating gas gap, and
- the associated volume of heavy-water moderator.

In the reactor core, there are also reactivity devices (such as zone controllers and adjusters), which are positioned interstitially between fuel channels (See example in Figure 1.4). These devices perturb the nuclear properties of the lattice in their vicinity, and this effect has to be taken into account in the reactor model.

2. Full-Core Neutronics

The computer code which is used for full-core reactor-physics calculations in CANDU is RFSP (Reactor Fuelling Simulation Program). RFSP is a 2-energy-group neutron-diffusion code, and can be used for both time-independent (steady-state) and time-dependent (kinetics) calculations.

The 2-group time-independent neutron-diffusion equation is:

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

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

where

$\nu\Sigma_{f1}$ *and* $\nu\Sigma_{f2}$ are the fast-group and thermal-group neutron-yield cross sections,

Σ_{a1} *and* Σ_{a2} are the fast-group and thermal-group absorption cross sections,

$\Sigma_{1\to 2}$ is the slowing-down (moderation) cross section,

$\Sigma_{2\to 1}$ is the up-scattering cross section,

D_1 *and* D_2 are the fast-group and thermal-group diffusion coefficients, and

k_{eff} is the reactor multiplication constant and is the inverse of the eigenvalue.

RFSP solves the eigenvalue equation above using a reactor model such as the one shown in Figures 2.1a and 2.1b (face and top views respectively), which typically contains several thousand or tens of thousands of mesh points (parallelepipeds).

Generally, each parallelepiped in this model may have different neutronic properties, which are the homogenized properties of the corresponding basic lattice cell, complemented where appropriate by the "incremental" properties of the associated reactivity device.

With "real-life" models, this equation can be solved only by iterative methods (e.g., Successive Over-Relaxation).

3. Fuel Management

3.1 General Considerations

After the initial period following first reactor startup, on-power refuelling is the primary means of maintaining a CANDU reactor critical.

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.

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 is in opposite directions in 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.

On the average, approximately two channels are visited for refuelling (using the 8-bundle-shift scheme) per Full-Power Day (FPD). 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.

3.2 Evolution of Basic Lattice Properties

The homogenized basic-lattice properties which enter into the full-core diffusion equation above are calculated with a lattice code. The lattice (or cell) code which has traditionally been used for CANDU design and analysis is POWDERPUFS-V. This is an empirical-recipe code, based on the results of measurements made on heavy-water-moderated lattices in research reactors ZEEP and ZED-2 at Chalk River Laboratories.

The properties of the basic lattice cell are calculated as if the cell were surrounded by an infinite repetition of itself on all sides in all directions – a uniform "infinite" lattice. In addition to the neutronic cross sections, the lattice code also provides the multiplication constant of this infinite lattice, k_∞. This infinite-lattice multiplication constant k_∞ of course does not include the ~30 milli-k of leakage out of the reactor (since it applies to an infinite lattice), or the ~20 milli-k of in-core reactivity devices (since they are not taken into account in the basic lattice cell). Thus, k_∞ = ~1.050 would correspond approximately to k_{eff} = 1.000, a critical reactor. In addition, in relative terms, k_∞ is a convenient measure of the basic local multiplicative properties of the lattice.

The cell code must be able to do "burnup" or "depletion" calculations, i.e., determine how the composition of the fuel changes as a function of the length of time it resides in the reactor. In addition to their dependence on burnup, the basic lattice properties and k_∞ depend on a multitude of factors: dimensions of the basic cell components, geometry of the fuel bundle, composition and purities of the cell components (especially the purity of the heavy-water moderator and the presence of any neutron "poisons"), densities and temperatures of the cell materials, etc.

A quantity closely related to burnup and often used in place of burnup is the fuel "irradiation" or "exposure":

$$\omega = \phi t,$$

which is a measure of the length of time t the fuel remains in a given flux ϕ. ω is measured in units of "neutrons per kilobarn", n/kb, and its value in POWDERPUFS-V is very closely one hundredth of the fuel burnup measured in MW.h/kg(U).

Figure 3.2 shows k_∞ as a function of irradiation for the standard CANDU 6 lattice fuelled with natural uranium and with typical physical properties. 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. This phenomenon is due to the production of fissile plutonium from neutron absorption in ^{238}U. This reactivity maximum is consequently known as the plutonium peak. Beyond the plutonium peak, k_∞ decreases monotonically with increasing irradiation, on account of the continuing depletion of ^{235}U and the increasing fission-product load.

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).

3.3 Average Exit Irradiation and Burnup

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).

3.4 Snapshot of Reactor Operating History

It is clear that, at any given moment (snapshot) in the operating history of the reactor, the nuclear properties may vary significantly from one point in the core to another, if only because of the different instantaneous burnup of different fuel bundles.

But it is also important to realize that the nuclear properties may vary spatially because of other differences in local core parameters:

- different fuel temperature
- different coolant density and temperature
- different absolute power (and therefore flux) level
- differences in geometry (e.g., different degree of pressure-tube radial creep)
- etc.

In addition, there will be changes in time, on account of advancing burnup, but also because of possible changes in moderator-poison concentration, etc.

And, of course, in the reactor model, the properties at any given mesh point (parallelepiped) - and time - will also differ on account of the presence or absence of specific reactivity devices. RFSP must be able to keep track of all these local differences in lattice properties throughout the core.

3.5 The 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).

A proper set of numerical equations can be written down for the time-average model, but time constraints do not permit writing these here in any detail. Suffice it to say that the time-average equations constitute a problem in self-consistency (see Figure 3.6) between:

- average lattice properties at each location
- 3-d core flux (and, therefore, also power) distribution
- range of irradiations experienced at each location
- dwell (residence) time at each location
- k_{eff}.

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 (which, however, must be in a reasonable range of values if a k_{eff} of unity is to be possible. In particular, 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 3.7.

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 3.8) provide a very useful guideline for the time intervals at which specific channels should be refuelled.

3.6 The Channel-Power Cycle

The time-average model does not feature a "refuelling ripple", because the lattice properties at any point are essentially averages over time. The "refuelling ripple" is a consequence of the daily refuelling of channels and the "irradiation cycle" through which each channel travels, and can be seen only in a snapshot calculation. The "irradiation 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 (time interval between refuellings of the channel), the power of the channel may be close to the power suggested by the time-average model.

- The reactivity of the channel and its power continue to drop. Eventually, the channel becomes a net "sink" or absorber of neutrons, and nears the time when the channel must be refuelled again. At this time the power of the channel may be 10% or more below its time-average power. When the channel is refuelled, its power may jump by 15 to 20% or even more.

The power of each channel therefore goes through an "oscillation" about the time-average power during every cycle. This cycle repeats every time the channel is refuelled, that is, with a period approximately equal to the dwell time suggested by the time-average model. The cycle length is not **exactly** equal to the dwell time, because channels are not refuelled in a rigorously defined sequence. Instead, as described in the previous section, channels are selected for refuelling based on instantaneous, daily information about the core power and irradiation distributions. In addition, the CANDU fuelling engineer has much flexibility in deciding how the core should be managed, and in fact can decide to modify the global power distribution by changing the refuelling frequency (dwell time) of various channels.

As individual channels are refuelled and go through their channel-power cycle, the specific sequence of these discrete refuellings results in variability in the instantaneous peak channel and bundle powers in the core. This is illustrated in Figure 3.9, a schematic plot of the maximum channel power versus time which illustrates the difference between maximum time-average channel power, average maximum instantaneous channel power, and absolute maximum channel power.

The instantaneous "ripple" in a channel's power (say channel m) is defined as

$$Channel - Power\ Ripple\ (m) = \frac{CP_{ins\,\tan\tan\,eous}(m)}{CP_{time-average}(m)}$$

This value will oscillate about unity over the refuelling cycle for channel m.

At any given time, however, 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 3.9.

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:

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

where m runs over all channels in the core, or at least over all safety-significant channels (i.e., the last two outermost rings of channels excepted).

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. If a 4-bundle-shift scheme were used, the typical CPPF would likely be 1.04-1.05.

The exact value of the CPPF is extremely important because it is used to calibrate the in-core ROP (protection) detectors. It determines the operating margin, which is the difference between the CPPF and the trip setpoint of the detectors.

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. 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.

3.7 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 readings of 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 "refuelling ripple" 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** with RFSP. This pre-simulation (especially if it searches for the zone-control-compartment water fills anticipated following the refuellings) will show whether the various power, burnup, and zone-fill criteria are likely to be satisfied, or whether the channel selection should be changed.

3.8 Core-Follow Calculations with RFSP

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.

Then the diffusion equation is solved.

The output of the RFSP calculation for any snapshot consists of the instantaneous 3-d power, irradiation and burnup distributions in the core, as well as the reactor multiplication constant k_{eff}. The power distribution is used to demonstrate compliance with the licensed values of channel and bundle powers. All the results are also extremely useful, if not essential, in assisting the reactor physicist to manage the tasks of fuel management.

4. Summary

Fuel management in CANDU has both design and operations aspects.

The design component consists of establishing:

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

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

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

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

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

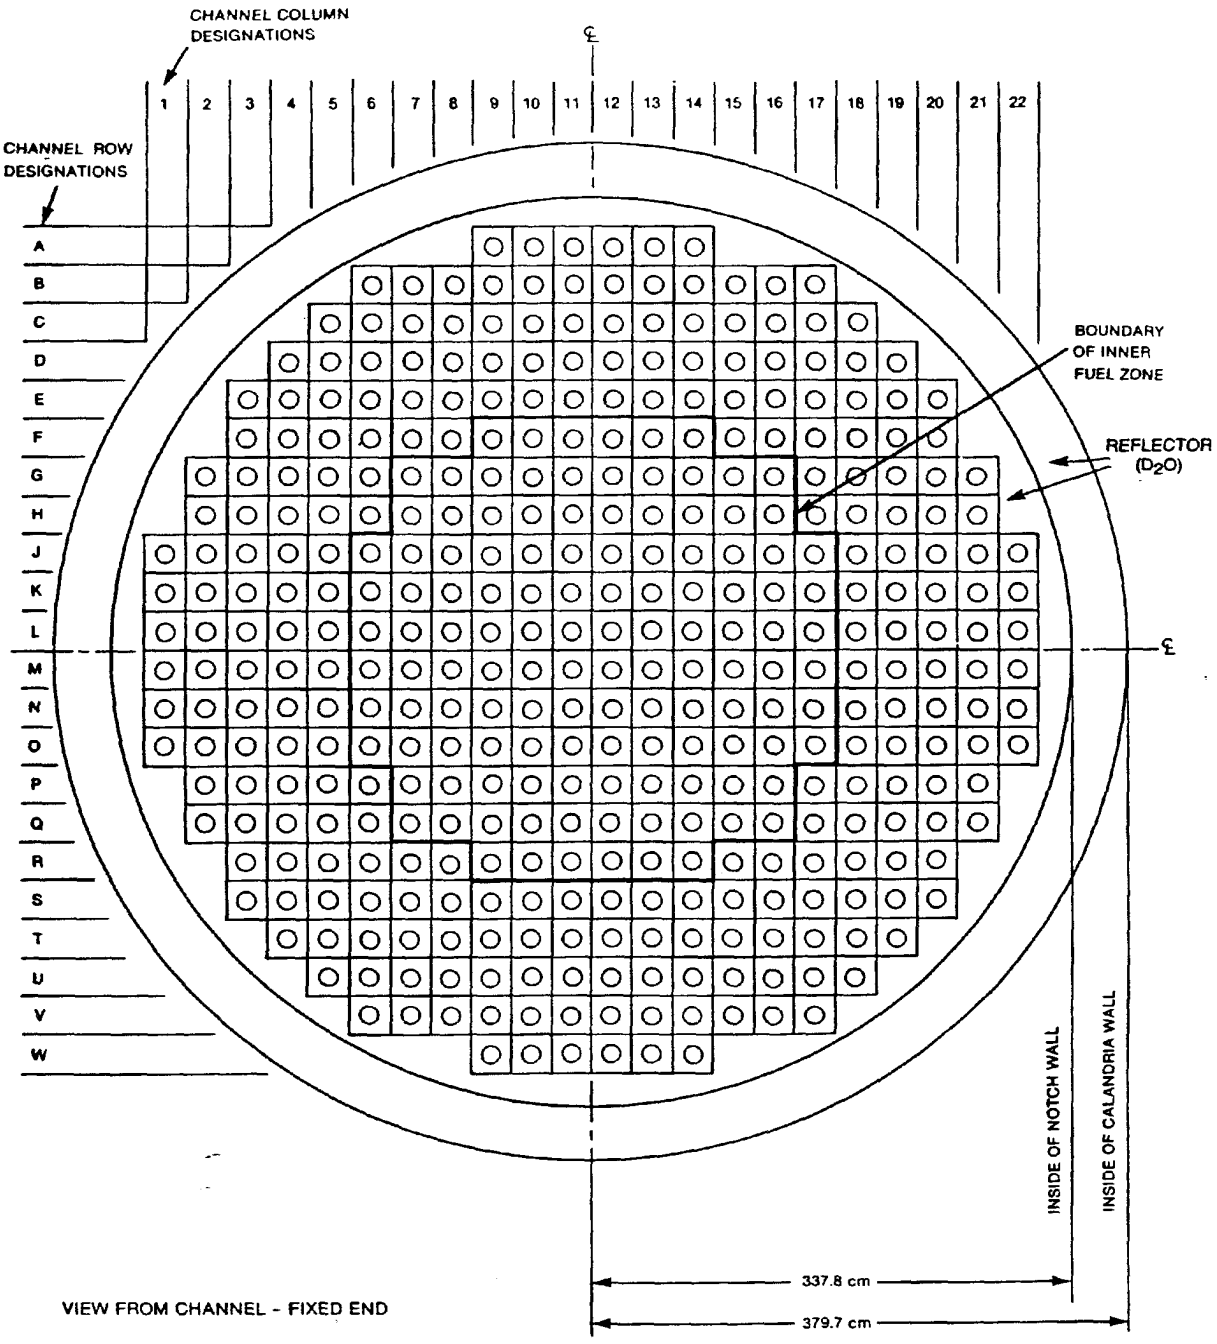

Figure 1.1
Face View of CANDU-6 Reactor Core

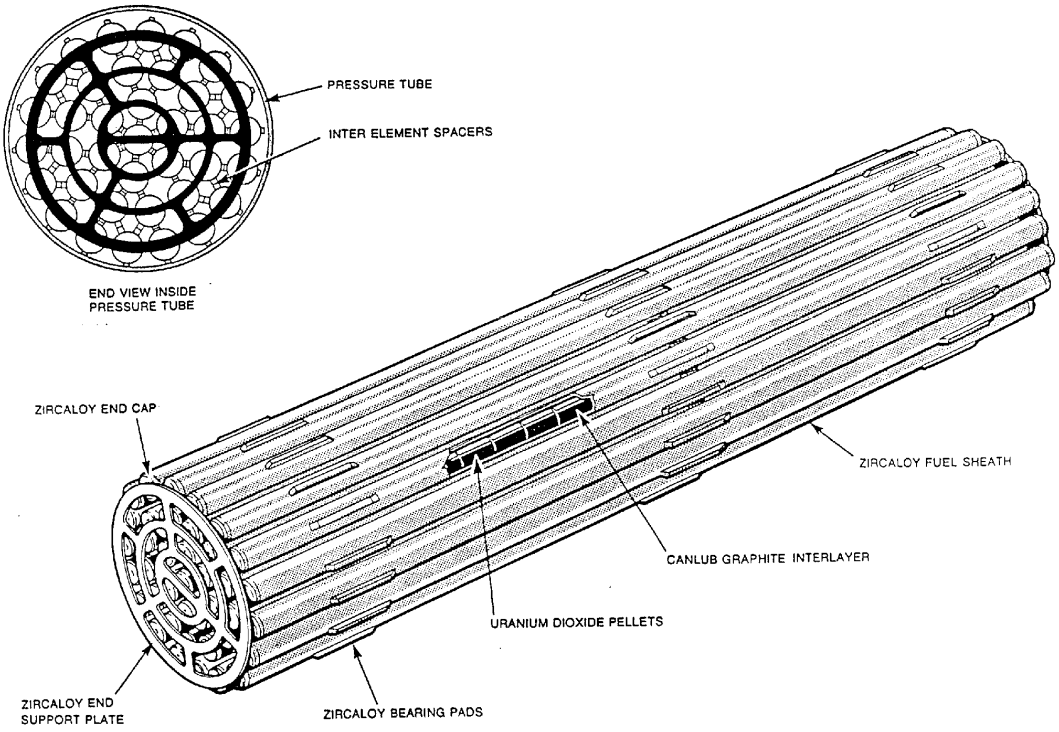

Figure 1.2
CANDU 37-Element Fuel Bundle

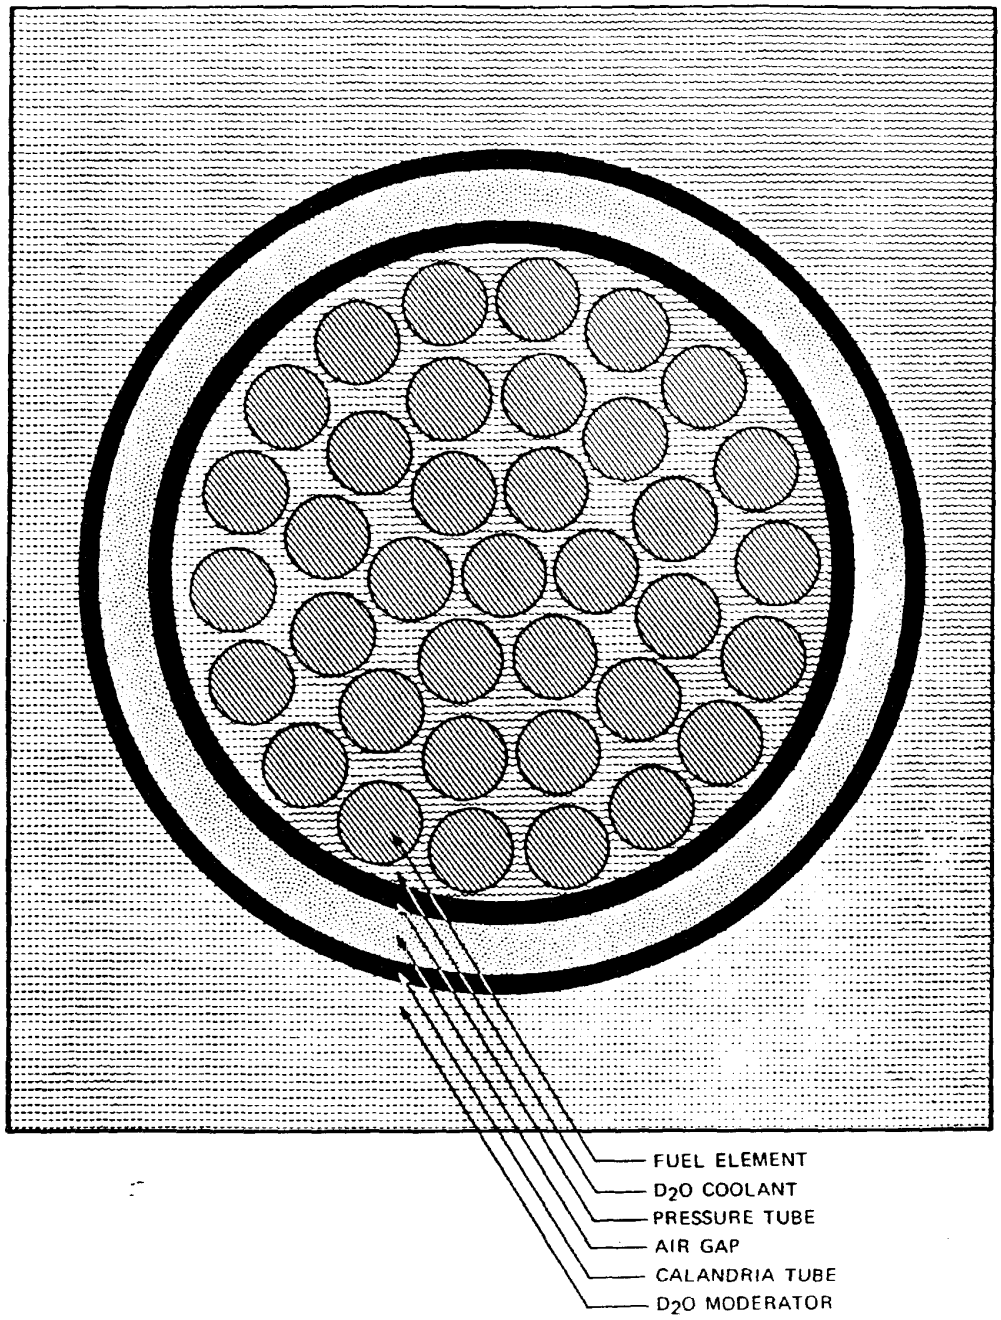

Figure 1.3
CANDU Basic Lattice Cell (Not to Scale)

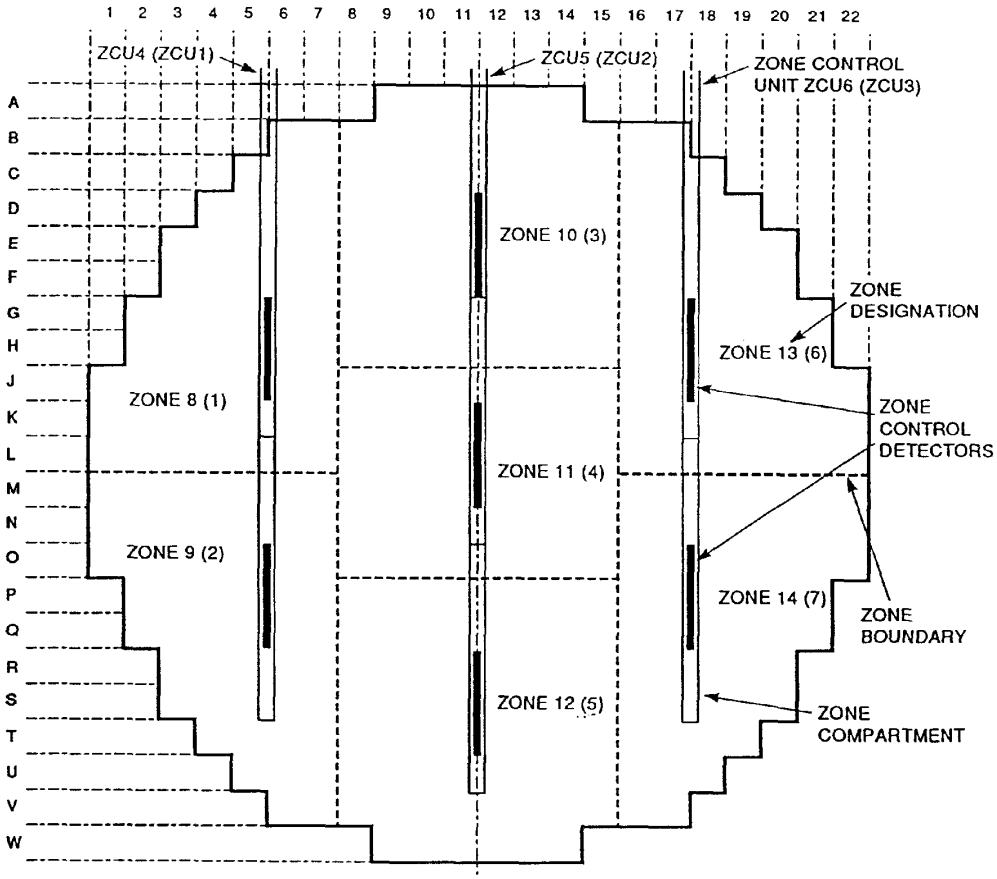

NOTES:

1 FACE VIEW FROM PRESSURIZER END. ZONES AND ZONE CONTROLLERS AT OTHER END (AWAY FROM PRESSURIZER) ARE GIVEN IN BRACKETS.

2 ZONE CONTROL DETECTORS ARE 3 LATTICE PITCHES LONG AND LOCATED IN VFDAs 1, 2, 3 AND 24, 25, 26 AXIALLY OUTBOARD FROM THE ZCUs.

920001

Figure 1.4
Interstitial Position of Liquid Zone Controllers

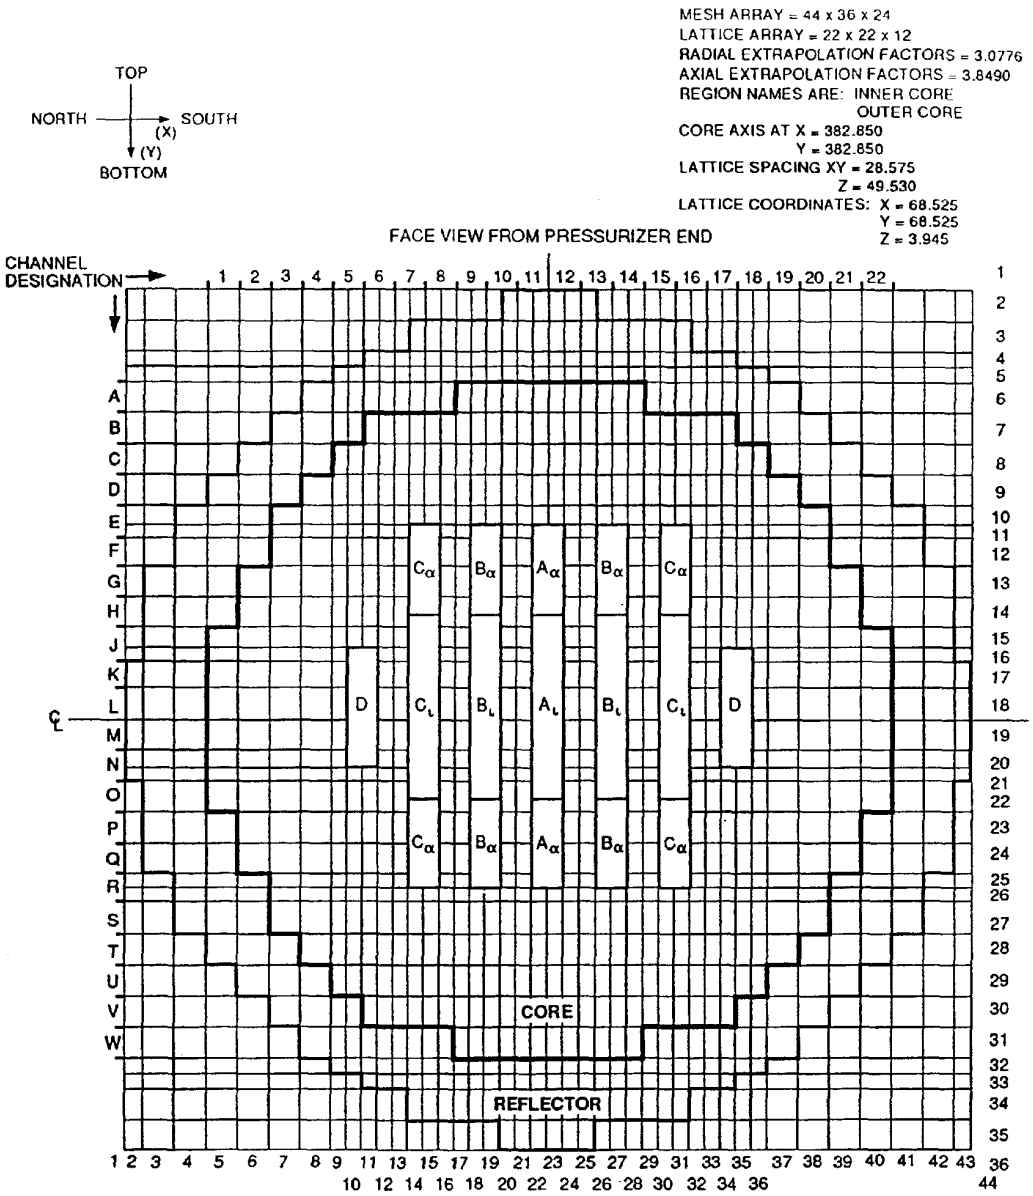

NOTE: ADJUSTERS ARE SHOWN IN THE ASYMMETRIC VERTICAL POSITION IN THE ABOVE DRAWING
ACTUAL VERTICAL POSITION IN MODEL IS DIFFERENT FOR EACH ADJUSTER (SEE TABLE 5.1A)
MESH SPACINGS ARE GIVEN IN TABLE 2.2A
A,B,C AND D ARE ADJUSTER ROD SEGMENTS AND SUBSCRIPTS ι AND α STAND FOR
INNER AND OUTER SEGMENTS RESPECTIVELY

NOT TO SCALE

Figure 2.1a
Typical RFSP Core Model - Face View

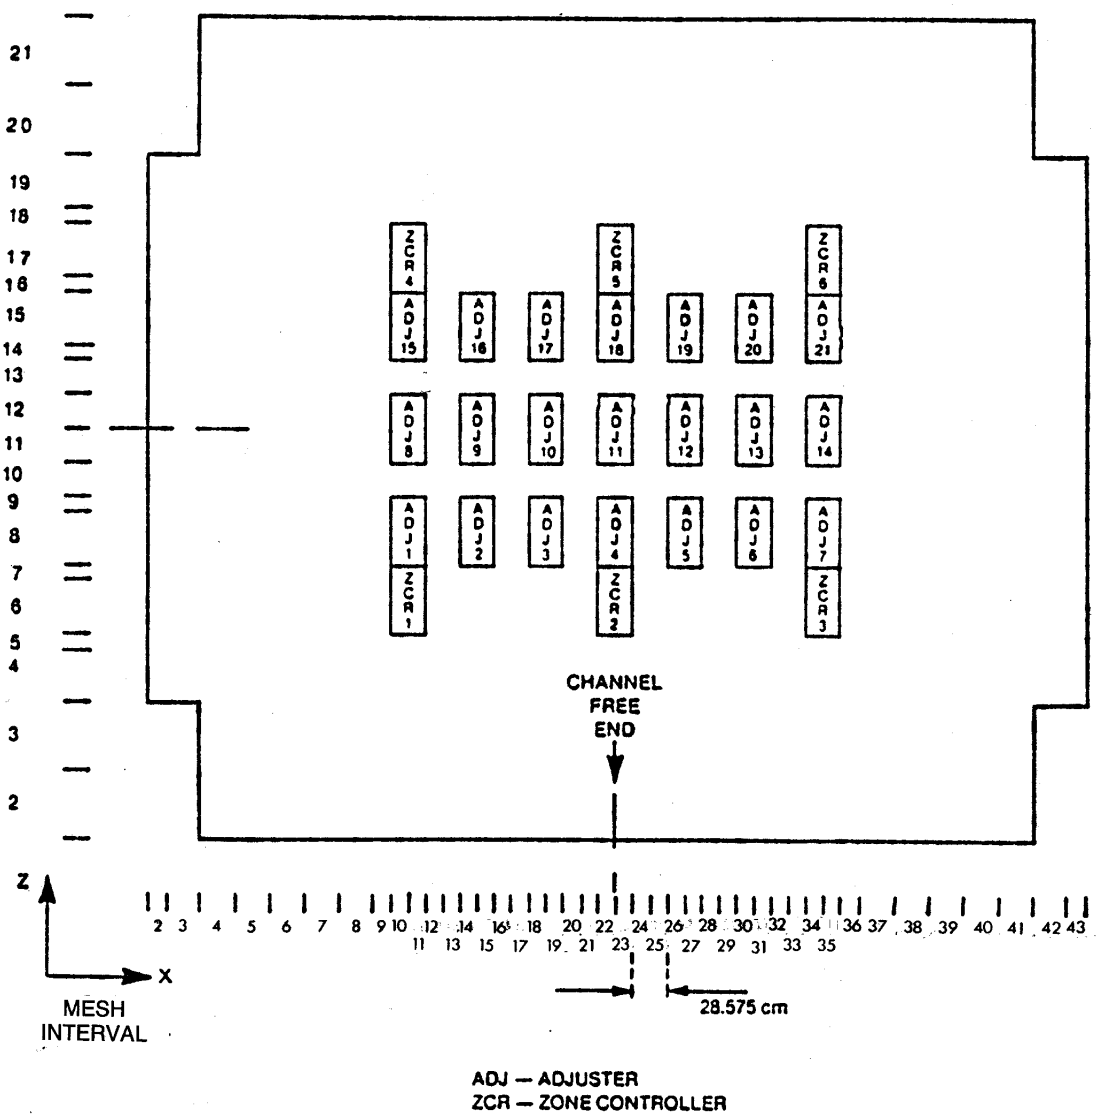

Figure 2.1b
Typical RFSP Core Model – Top View

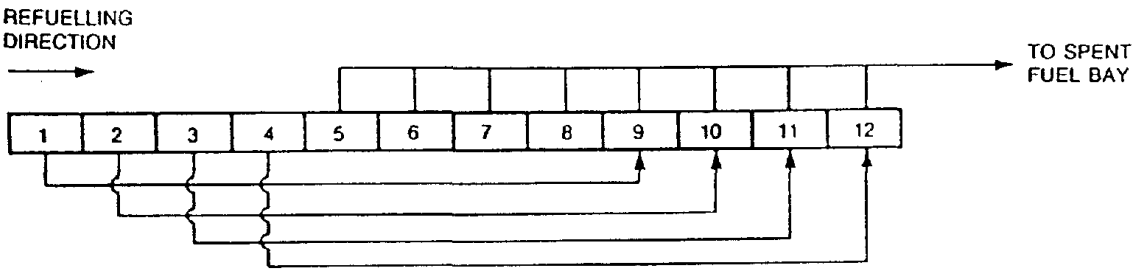

Figure 3.1
8-Bundle-Shift Refuelling Scheme

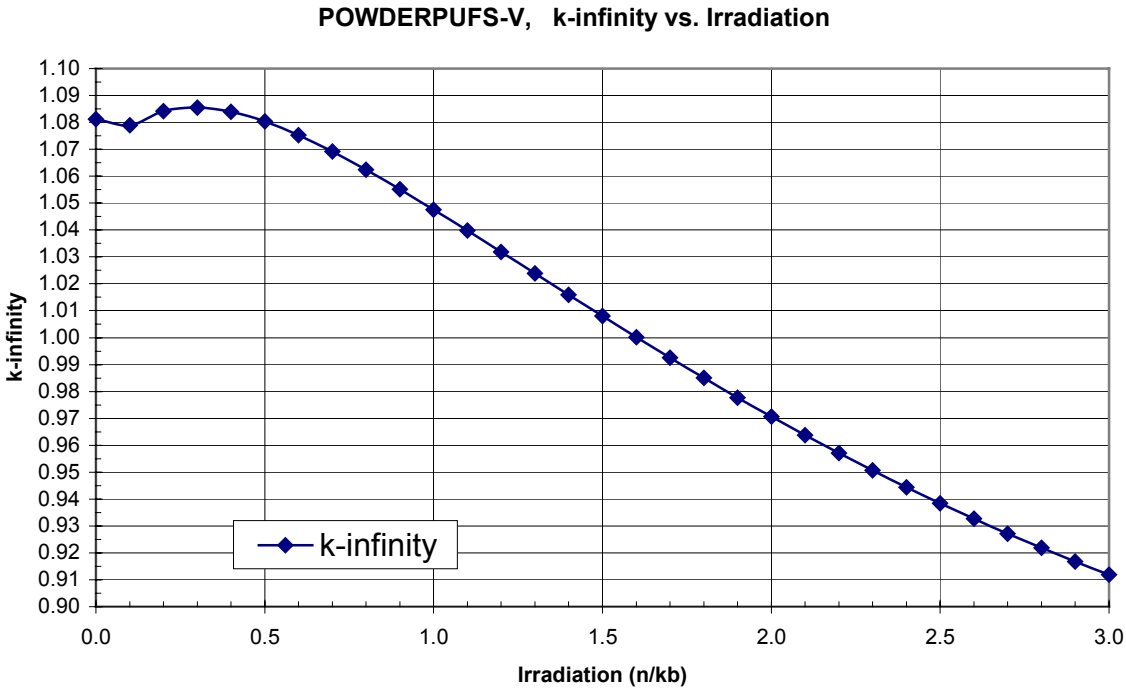

Figure 3.2
k_∞ as a Function of Irradiation

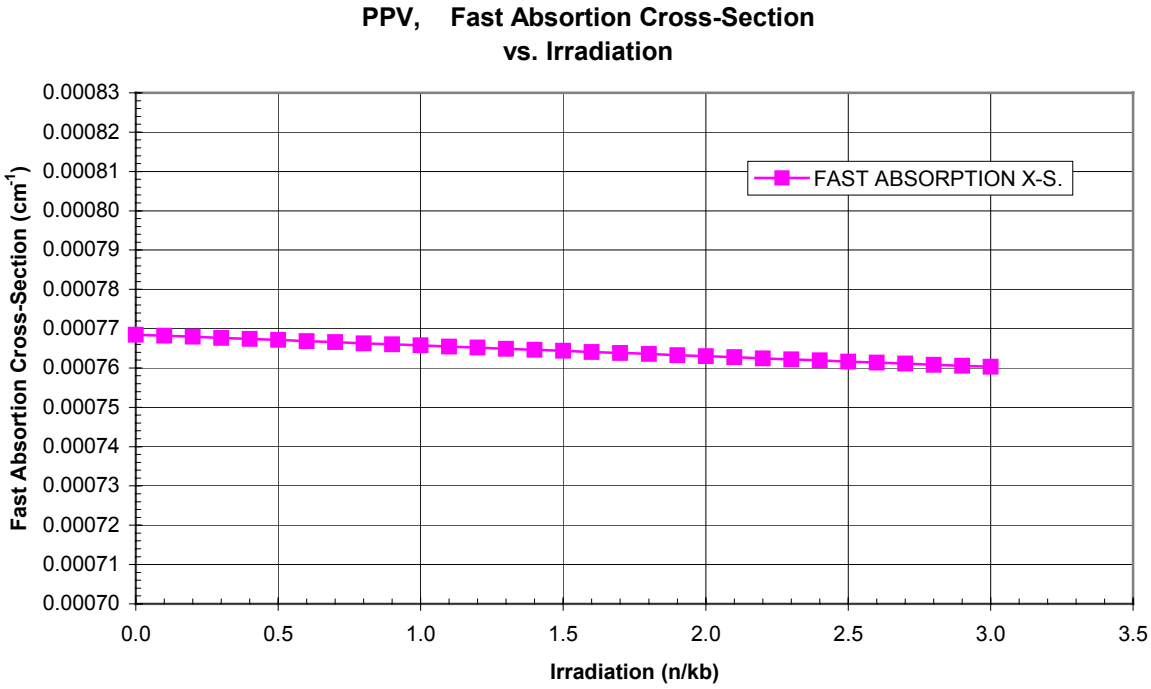

Figure 3.3
Fast-Group Absorption Cross Section (Σ_{a1}) as a Function of Irradiation

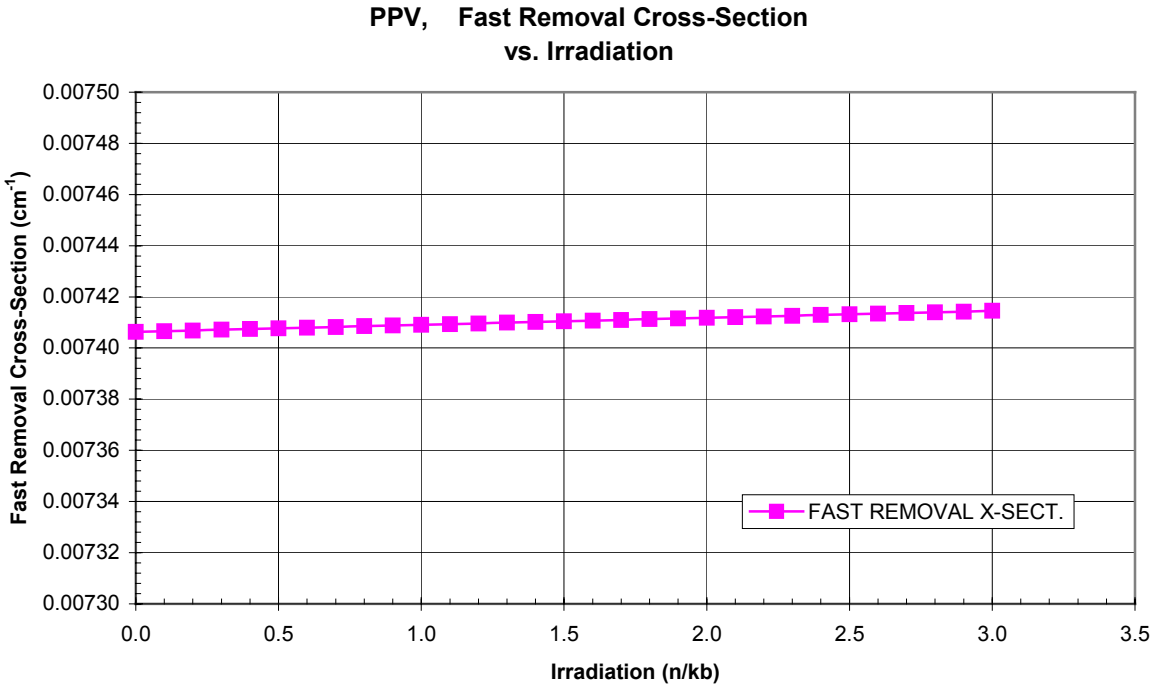

Figure 3.4
Fast-Group Slowing-Down (Moderation) Cross Section (Σ_m) as a Function of Irradiation

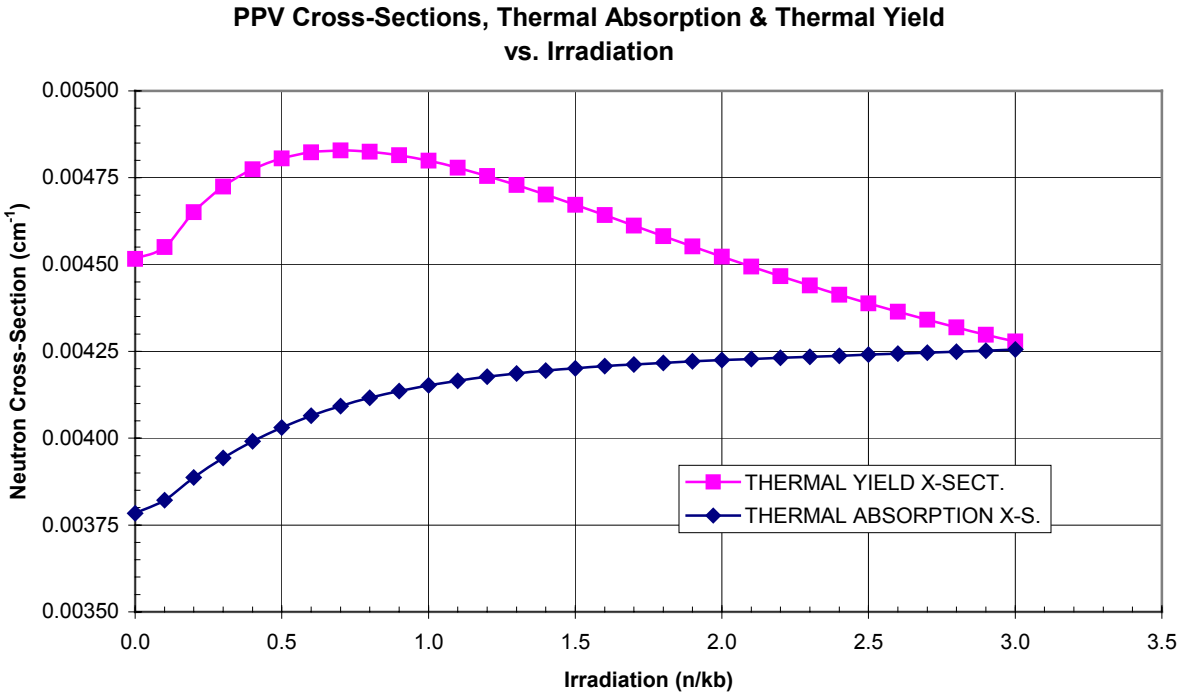

Figure 3.5
Thermal-Group Yield Cross Section ($\nu\Sigma_{f2}$) and
Thermal-Group Absorption Cross Section (Σ_{a2})
As Functions of Irradiation

Time-Average Calculation

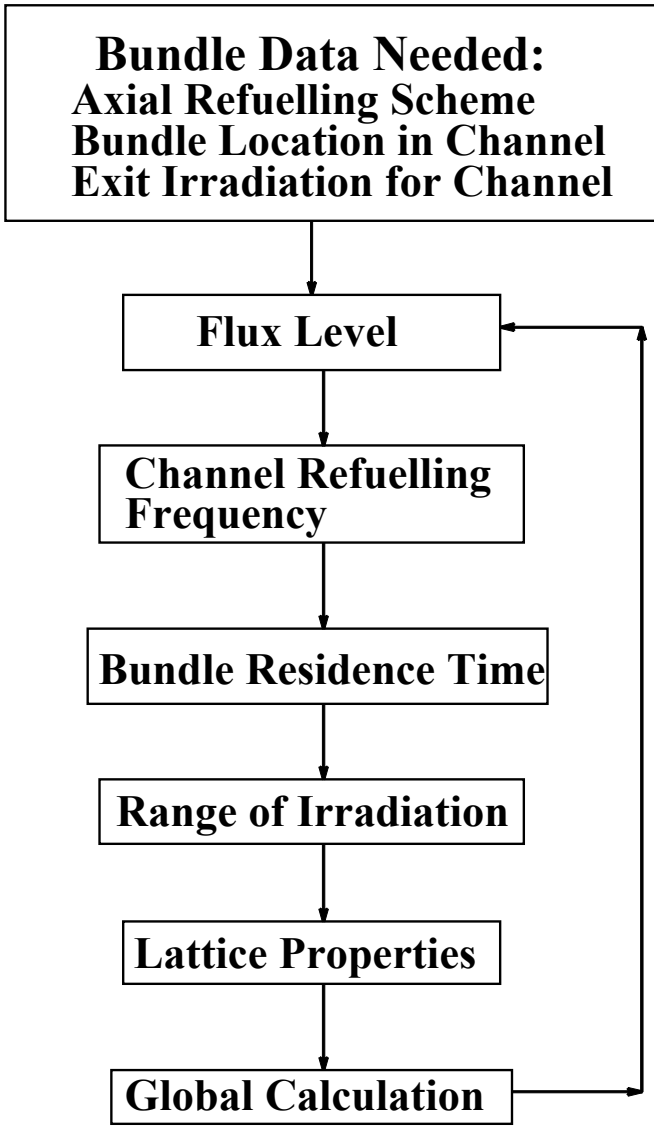

Figure 3.6
Time-Average Self-Consistency Problem

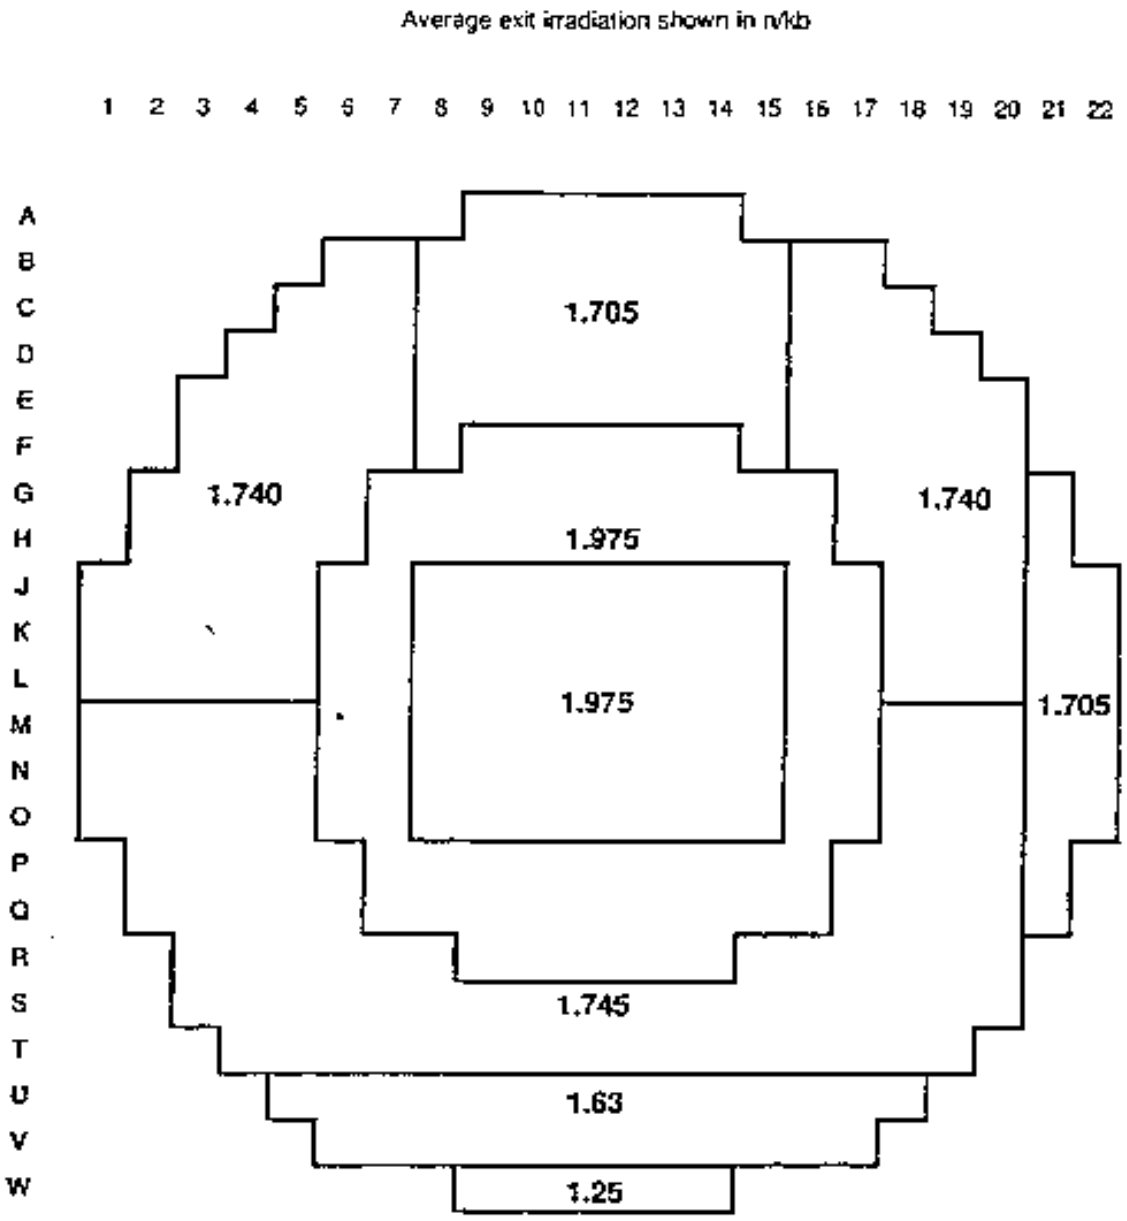

Figure 3.7
Typical Time-Average Model, Showing Irradiation Zones

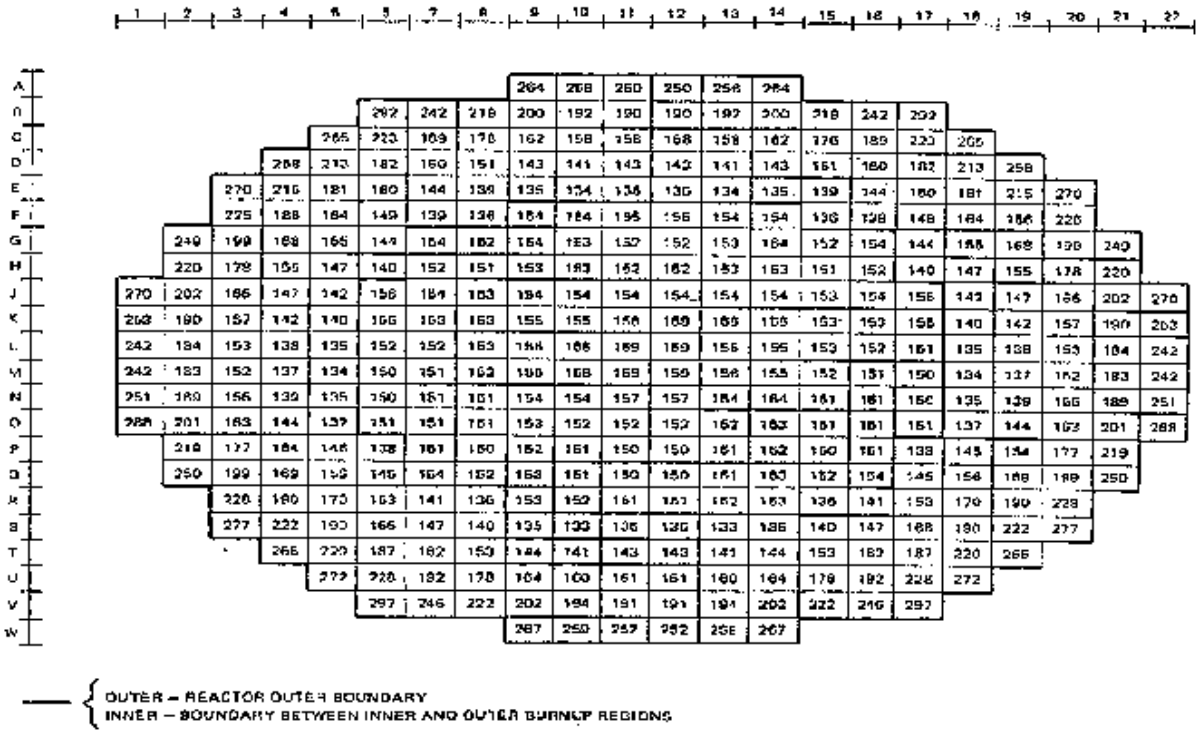

Figure 3.8
Channel Dwell Times from Time-Average Calculation

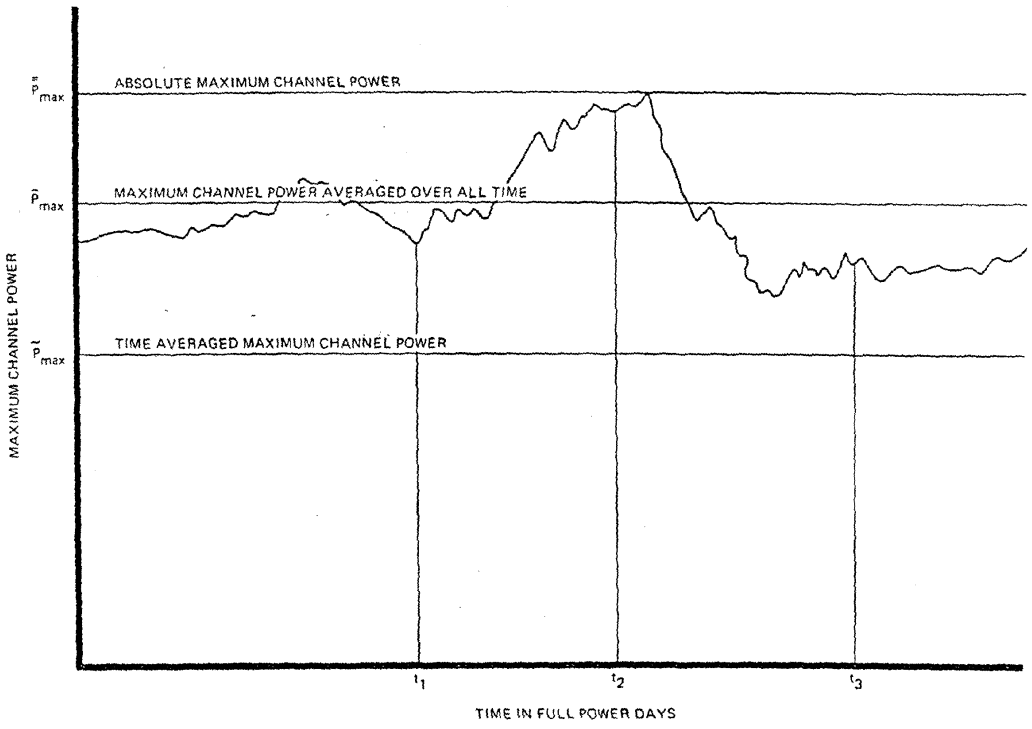

Figure 3.9
Sketch of Maximum Instantaneous and Time-Average Channel Powers