Fuel Management in CANDU

by

B. Rouben
Manager, Reactor Core Physics Branch
AECL

Presented at Chulalongkorn University
Bangkok, Thailand, 1997 December

Abstract

Fuel management in CANDU reactors is discussed. The variation of lattice reactivity with irradiation is shown. The various periods in the operating life of the reactor are discussed. The axially-homogeneous and time-average models for the CANDU equilibrium core are defined. Core-follow simulations are described and illustrated with results for a CANDU 6. Fuel-burnup warranties are described. The effect of Xe-135 on power is described.

Table of Contents

Figures

1. Introduction

1.1 General Discussion of CANDU On-Power Refuelling

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

The capability for on-power refuelling also means that, in CANDU, 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, which has much smaller excess reactivity than enriched fuels;
- 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, it is necessary to provide only a few milli-k (in the light-water zone-control compartments) of capability;
- No burnable poison or large amounts of moderator poison need be provided, since it is not necessary to hold down the excess reactivity of a large batch of fresh fuel; even when all fuel is fresh, in the initial core, only ~2-3 ppm of moderator boron are required;
- If some xenon-override capability is desired, an adjuster-rod system with a reactivity worth of approximately 15 milli-k (as in the CANDU 6) can compensate for about 30 minutes of xenon growth following a reactor shutdown.

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.

Several refuelling operations are normally carried out daily, so that refuelling is almost continuous. By adjusting the refuelling rate in the various core regions, the power distribution in these regions can be effectively controlled over a time scale of the order of weeks. Thus, the fuelling engineer has much flexibility in "shaping" the power distribution.

The fuel is in the form of short bundles (Figure 1.2) contained in a pressure tube or "channel". To refuel a channel, a pair of fuelling machines latch onto each end of the channel. New fuel bundles are inserted at one end, old fuel is pushed along the channel, and spent fuel is discharged into the other machine. Typically, four or eight fresh bundles are added at each fuelling. Figure 1.3 illustrates the 8-bundle-shift scheme, where the eight bundles near the outlet end of the channel are discharged, and the four bundles previously at the inlet end are moved to the outlet end. Thus, the four low-power bundles

are in-core for two cycles (i.e., two visits of the channel) and the high-power bundles are in-core for only one cycle.

CANDU refuelling is bi-directional: alternate channels are refuelled in opposite directions. This promotes axial symmetry in the flux and power distributions.

1.2 Objectives of Fuel Management

The primary objective of fuel management is to determine fuel-loading and fuel-replacement strategies which will result in minimum total unit energy cost while operating the reactor in a safe and reliable fashion. 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.

2. Periods in 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.

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. The excess reactivity is compensated by adding boron poison (~2-3 ppm) to the moderator. Fuel-management calculations are required to assess the effect of the initial fuel loading on the subsequent power operation.

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. When the excess reactivity has fallen to a small value, refuelling begins in order to maintain the reactor critical. During this transitional or "pre-equilibrium" period, the reactor gradually approaches the final or "equilibrium" state.

The equilibrium condition in a CANDU reactor is reached after approximately 400 to 500 FPD. It is characterized by a relatively unchanging core configuration, in which the macroscopic or global power and burnup distributions do not vary significantly with time. The burnup of the discharged fuel, and the fuelling rate of new fuel, become essentially constant. One can think of the flux (and power) distribution in terms of a reference, constant time-average shape about which a "refuelling ripple" is superimposed due to the refuelling of a number of channels each day. Because the equilibrium core continues for the remainder of the reactor's life, most fuel-management studies are performed for this core configuration.

3. Calculational Tools

3.1 Lattice Code

The basic unit of the CANDU reactor is the lattice cell (see Figure 4.1). It consists of a fuel bundle and the associated coolant, contained in concentric pressure and calandria tubes, surrounded by the associated volume of moderator. The lattice pitch in CANDU is 28.575 cm, and the bundle length is 49.53 cm.

The first tool needed for fuel-management calculations is a lattice code. The lattice code computes the nuclear properties of the lattice, assuming an infinite repetition of cells in every direction. The lattice code which has been used for the design of all current CANDU reactors is POWDERPUFS-V, a semi-empirical code, based on research-reactor measurements on heavy-water-moderated lattices. An attractive feature of POWDERPUFS-V is that its mathematical framework is simple, and it executes extremely quickly. Other lattice codes which can be used are the transport codes such as WIMS-AECL.

The lattice properties depend on the physical dimensions and composition of the materials in the cell. In particular, of course, the properties depend on the isotopic composition of the fuel, i.e., the relative amounts of uranium and plutonium isotopes, which will change on irradiation. In fact, for a given lattice and fuel geometry, fuel irradiation is the major determinant of the lattice properties.

Figure 4.2 shows the variation with irradiation of the "infinite-lattice" excess reactivity, $(k_\infty - 1)/ k_\infty$, for both natural-uranium fuel and depleted fuel (depleted to 0.52 atom % ^{235}U). This figure shows one of the characteristics of the irradiation of CANDU fuel: the "plutonium peak".

As natural fuel is irradiated, the reactivity increases at first, 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. The "plutonium peak" (at which the reactivity is highest) occurs at approximately 40-50 FPD of irradiation. Following the plutonium peak, the plutonium production can no longer compensate for the buildup of fission products, and the fuel reactivity decreases.

Depleted fuel also shows a plutonium peak. In fact, the plutonium peak is even more pronounced in depleted fuel, because the effect of plutonium production in ^{238}U is relatively higher in the presence of less ^{235}U. The reactivity of the depleted fuel shown is however negative at all irradiations: that is, this depleted fuel is always a net **sink** (absorber) of neutrons.

3.2 Finite-Core Code

Only crude or conceptual fuel-management calculations can be performed with lattice codes alone. To obtain more accurate results, one needs to perform calculations with finite-core codes.

RFSP (Reactor Fuelling Simulation Program) is the major computer program in use at AECL for the design and analysis of CANDU reactor cores. It is also used for fuel-management calculations at the core-design stage, and for core-follow, either in simulation during design, or to track the core history of an operating reactor, at site.

Some of the features of RFSP are:

- It solves the finite-difference form of the neutron-diffusion equation in 3 dimensions and 2 energy groups;
- It also provides a flux-mapping method for reconstructing the core flux distribution based on the readings of in-core detectors;
- It incorporates the cell code POWDERPUFS-V (developed at AECL for heavy-water, natural-fuel lattices) to calculate the lattice properties, i.e., the nuclear cross sections;
- It can simulate channel refuellings and burnup steps and track the irradiation history of individual bundles;
- It allows variable mesh spacings in the finite-difference model;
- It can emulate the asymptotic bulk-control and spatial-control functions of the reactor regulating system (RRS);
- It can perform spatial-kinetics calculations;
- It is of modular design, with each code functionality programmed into a separate module with its own well-defined input and output. This makes the code flexible and easy to modify. Similarly, the data is organized in blocks according to function and type.

Three-dimensional core models similar to the one illustrated in Figures 4.3a and 4.3b are used in RFSP. Reactivity devices and structural materials are represented by "incremental" cross sections which act over "supercells"; a typical model of a supercell is shown in Figure 4.4. The incremental cross sections, such as $\Delta\Sigma a_1$, $\Delta\Sigma a_2$, $\Delta\nu\Sigma f$, etc., are added to the "bare-lattice" properties of those lattice cells over which the supercells are superimposed. Typical RFSP full-core models have from 5,000 to 30,000 (or more) mesh points, depending on the requirements of the problem.

4. Equilibrium-Core Calculations

The equilibrium core, defined cursorily in Section 2, is attained when the **global** burnup and power distributions have attained a "steady state". While in reality there are always local "ripples" due to the discrete nature of refuelling in CANDU, these are "smoothed out" when thinking of the equilibrium core.

The equilibrium core contains fuel with a distribution of irradiations, ranging from zero to discharge values. Ideally, while the equilibrium core could be analyzed by averaging over time simulations of a long period of reactor operation, this is not very practical.

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

4.1 Axially-Homogeneous Model

This is the simplest model. Chronologically, it was the first to be developed, in the early days of CANDU design. It can still be used for the simplest level of complexity. The axially homogeneous model is based on:

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

With these assumptions, it can be shown that the irradiation, averaged over two neighbouring channels, is independent of axial position in the core. The mathematical treatment is simple and is shown in the next few equations, (4.1)-(4.10) .

Consider two neighbouring channels; these are refuelled in opposite directions. Suppose that fuel moves continuously at speed r in these two channels (in opposite directions). See Figure 4.5.

Fuel irradiation ω is defined as (fuel) flux multiplied by time. Since the fuel moves through the channels and is exposed to different fluxes as it moves, the fuel irradiation in each channel at position $x = x_0$ must be calculated as the integral over time of fuel flux as fuel travels through the core to $x = x_0$:

$$\omega_1(x_0) = \int_0^{t_1(x_0)} \hat{\phi}(x(t))dt \tag{4.1}$$

$$\omega_2(x_0) = \int_0^{t_2(x_0)} \hat{\phi}(x(t))dt \tag{4.2}$$

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

$$t_1(x_0) = \frac{x_0}{r} \tag{4.3}$$

$$t_2(x_0) = \frac{L - x_0}{r} \tag{4.4}$$

Use Equations (4.3), (4.4) to rewrite ω_1 and ω_2 as integrals over x:

$$\omega_1(x_0) = \int_0^{x_0} \hat{\phi}(x)\frac{dx}{r} = \frac{1}{r}\int_0^{x_0} \hat{\phi}(x)\,dx \tag{4.5}$$

$$\omega_2(x_0) = \int_L^{x_0} \hat{\phi}(x)(-\frac{dx}{r}) = \frac{1}{r}\int_{x_0}^{L} \hat{\phi}(x)dx \tag{4.6}$$

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

$$\begin{aligned}\omega_{avge}(x_0) &= \frac{1}{2}\left[\omega_1(x_0) + \omega_2(x_0)\right] \\ &= \frac{1}{2r}\left[\int_0^{x_0} \hat{\phi}(x)\,dx + \int_{x_0}^{L} \hat{\phi}(x)\,dx\right] \\ &= \frac{1}{2r}\int_0^{L} \hat{\phi}(x)\,dx\end{aligned} \tag{4.7}$$

[The last step requires the assumption that the flux shape (from inlet end to outlet end) is the same in the two neighbouring channels .]

The right-hand side of Eq. (4.7) does not depend on the position x_0 in the channels. This shows that we can treat the average irradiation as uniform along the two channels, equal to one half of the fuel exit irradiation (ω_{exit}) at the channel outlet:

$$\omega_{exit} = \omega_1(L) = \omega_2(0) = \frac{1}{r}\int_0^{L} \hat{\phi}(x)\,dx \tag{4.8}$$

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

$$\Sigma_i\,(\text{hom}) = \frac{\text{Reaction rate averaged over time}}{\text{Flux averaged over time}}$$

$$= \frac{\frac{1}{T}\int_0^T \Sigma_i\left(\omega(t)\right)\hat{\phi}\left(x(t)\right)dt}{\frac{1}{T}\int_0^T \hat{\phi}\left(x(t)\right)dt} \qquad (4.9)$$

where T = transit time of fuel through the core:

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

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

$$\Sigma_i\,(\text{hom}) = \frac{1}{\omega_{exit}}\int_0^{\omega_{exit}} \Sigma_i\,(\omega)\,d\omega \qquad (4.10)$$

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

Most often, the axially-homogeneous model is used with only two large core regions, the inner core and the outer core, each with a different value of ω_{exit}. Figure 4.6 shows a two-region model for the CANDU 6. The purpose of defining two separate regions, with different exit burnup, is to use **differential fuelling** to achieve radial flux (and power) flattening. By "driving" the inner core to a higher burnup than the outer core, the radial power distribution can be flattened, and 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 "regional" cross sections of Eq. (4.10) are calculated directly in POWDERPUFS-V, in what is called the "reaction-rate-averaged" mode, as functions of the exit irradiation ω_{exit}.

As can be seen, the calculational procedure for the axially homogeneous model is simple. However, the continuous-refuelling approximation does not take into account the axial effects that any specific refuelling scheme would induce. For this reason, the axially homogeneous model has been largely abandoned, except for the simplest type of analysis, in favour of the time-average model, described next.

4.2 Time-Average Model

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

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

$$\Sigma_{a2}^{t.a.}(r) = \frac{1}{(\omega_{out} - \omega_{in})} \int_{\omega_{in}}^{\omega_{out}} \Sigma_{a2}(\omega)\, d\omega \qquad (4.11)$$

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

Now, in the time-average model, let ϕ_{jk} be the time-average fuel flux at axial position k in channel j, k ranging from 1 to 12 (since there are 12 bundles per channel) and j ranging over the channels, e.g. from 1 to 380 in the CANDU 6.

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 j,k will experience over its residence time at that position will be

$$\Delta\omega_{jk} = \phi_{jk} T_j \qquad (4.12)$$

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

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

$$= \omega_{in,jk} + \phi_{jk} T_j \qquad (4.13)$$

When a channel is refuelled with an 8-bundle shift, the first 8 positions in the channel receive fresh fuel and the entrance irradiations for positions 9-12 are simply the exit irradiations from positions 1-4 respectively. Thus we can write in this case (see Figure 4.7).

$$\omega_{in,jk} = 0 \qquad k = 1,...,8 \qquad (4.14a)$$

$$\omega_{in,jk} = \omega_{out,j(k-8)} \qquad k = 9,...,12 \qquad (4.14b)$$

For an N-bundle-shift refuelling scheme these equations become

$$\omega_{in,jk} = 0 \qquad k = 1,...,N \qquad (4.15a)$$

$$\omega_{in,jk} = \omega_{out,j(k-N)} \qquad k = (N+1),...,12 \qquad (4.15b)$$

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

$$\omega_{exit,j} = \frac{1}{8} \sum_{k=5}^{12} \omega_{out,jk} \qquad (4.16)$$

In view of Equation (4.13) this can be written

$$\omega_{exit,j} = \frac{1}{8} \sum_{k=5}^{12} \left(\omega_{in,jk} + \phi_{jk} T_j \right) = \frac{1}{8} \left[\sum_{k=5}^{8} \left(\omega_{in,jk} + \phi_{jk} T_j \right) + \sum_{k=9}^{12} \left(\omega_{in,jk} + \phi_{jk} T_j \right) \right] \qquad (4.17)$$

and in view of Equation (4.15) we can write

$$\omega_{exit,j} = \frac{1}{8}\left[\sum_{k=5}^{8}\phi_{jk}T^{j} + \sum_{k=1}^{4}\phi_{jk}T_j + \sum_{k=9}^{12}\phi_{jk}T_j\right] = \frac{T_j}{8}\sum_{k=1}^{12}\phi_{jk} \tag{4.18}$$

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

$$\omega_{exit,j} = \frac{T_j}{N}\sum_{k=1}^{12}\phi_{jk} \tag{4.19}$$

The dwell time T_j must therefore satisfy

$$T_j = \frac{N\omega_{exit,j}}{\sum_{k=1}^{12}\phi_{jk}} \tag{4.20}$$

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

- the finite-difference form of the time-independent neutron diffusion equation to solve for the flux distribution,
- Equation (4.20) to compute the dwell time for each channels,
- Equation (4.13) and (4.14) to calculate $\omega_{in,jk}$ and $\omega_{out,jk}$ for each bundle in core,
- Equation (4.11) (and similar equations for the other cross sections) to calculate the time-average lattice properties.

This set of equations must be solved using as input the user-specified target exit irradiations $\omega_{exit,j}$. In order to shape the flux to desired values, and also to take account of the presence of extra "hardware" (device locators, etc., mostly at the bottom of the calandria) which introduces localized absorption, typical time-average RFSP models now use many irradiation regions; see Figure 4.8.

Since consistency must be achieved between the flux, the channel dwell times, the individual-bundle irradiation ranges $[\omega_{in}, \omega_{out}]$, 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 4.9 shows the iterative scheme of calculation.

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

The time-average model is useful at the design stage, to determine the reference 3-dimensional power distribution, the expected refuelling frequency of each channel (or its inverse, the channel dwell time), and the expected value of discharge burnup for the various channels.

Figure 4.10 shows the time-average channel-power distribution obtained with a two-region model for the CANDU 6, where the exit irradiations are chosen to produce a radial form factor (ratio of average to maximum channel power) of 0.83. This corresponds to a maximum time-average channel power about 6.52 MW. This value is substantially lower than the maximum licenses channel power, 7.3 MW. However, remember that there must be sufficient margin above the maximum time-average power to accommodate the refuelling ripple, which could be of the order of 10%.

Figure 4.11 shows the channel dwell times for the same CANDU-6 time-average calculation. It can be seen that the dwell times in the inner core range between 150 and 160 FPD. In the outer core, the dwell times present a large variation, from about 135 FPD for channels just outside the inner core (where the flux is still high but the exit irradiation is, by design, lower than in the inner core) to almost 300 FPD for some channels at the outermost periphery of the core.

It is instructive to look at the neutron balance in the CANDU-6 equilibrium core, displayed in Figure 4.12. This shows that more than 45% of fission neutrons come from fissions in plutonium. Fast fissions account for 56 fission neutrons out of 1,000. Total neutron leakage is 29 neutrons lost per 1000 neutrons born in fission, representing a 29-milli-k loss (6 milli-k from fast leakage, 23 milli-k from thermal leakage). Resonance absorption in ^{238}U represents almost a 90 milli-k loss. Parasitic absorption in non-fuel components of the lattice represents a 63-milli-k loss.

5. Core-Follow Simulations

Core-follow simulations are performed to track the operating history of a reactor. Core-follow simulations, are a sequence of "instantaneous", or "snapshot", calculations in the life of the reactor. The time step ("burn step") between simulations is entirely variable, and is at the user's discretion. Typical burn steps are of the order of 1-3 full-power days (FPD). At each step, the irradiation distribution is incremented from the previous step using the previous flux and irradiation distributions and the length of the burn step (irradiation = flux*time). Individual channel refuellings are simulated within each burn step. Core-follow simulations are carried out with the *SIMULATE module of RFSP.

At the reactor design stage, core-follow is used to simulate the initial transient from startup to equilibrium, to investigate the effect of various refuelling rules, and to obtain accurate estimates of likely maximum powers, discharge burnup, etc.

For operating reactors, core-follow is performed to track the actual history of the core: to calculate and record bundle powers, channel powers, and bundle-irradiation values. At any given time, the current values are used by the fuel engineer at site:

- to select channels for refuelling;
- to ensure that channel and bundle powers are kept within specified limits; and
- for burnup evaluation.

The instantaneous power distribution shows localized channel-power "ripple" effects about the time-average power. These are caused by the discrete refuelling of specific channels on different days in the operating history, since the power distribution in a channel is time-varying. Recently-refuelled channels tend to have the highest powers, and channels near maximum burnup tend to have the lowest powers. When a channel is refuelled, its power increases for several weeks as it proceeds through its plutonium peak, and then tends to decrease until it is refuelled again. The neighbours to channel that is refuelled are also affected, with their power increasing slightly and then decreasing, as the local neutron source increases and decreases. The channel-power ripple depends on the axial refuelling scheme (number of bundles replaced at each channel visit) and on the refuelling frequency.

5.1 Initial Fuel Load and Transient to Onset of Refuelling

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 unless some alternate means of flattening the radial power distribution is provided. This is typically achieved with the use of depleted fuel, of 0.52 atom percent ^{235}U content. 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 are placed in each of the central 80 fuel channels, shown in Figure 5.1. The bundles are located in positions 8 and 9 (bundle positions in the channels being numbered from 1 to 12 in the direction of refuelling). In this axial position, the depleted-fuel bundles are removed from the core in the first refuelling visit of each of these channels.

Even with some depleted fuel in the core, the fact that all fuel is fresh results in a net excess reactivity in the core. The variation of the excess core reactivity with time in the initial period of operation is shown in Figure 5.2. 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 40 and 50 FPD. 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 120 FPD.

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 100 FPD, because the refuelling rate would be too great if one waited until the last possible moment to start.

Figures 5.3, 5.4 and 5.5 show representative bundle-power distributions in the horizontal radial direction at 0 FPD (initial core), 40 FPD, and 100 FPD respectively. The main feature to be noted is the initial "dishing" of the power distribution in rows containing some channels with depleted fuel. This dishing is quite pronounced at 0 FPD, bu flattens out with increasing fuel burnup, since the inner core actually goes through its plutonium peak earlier than the outer core, and the depleted-fuel plutonium peak is in fact more pronounced, as pointed out earlier.

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

5.2 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 center (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 18 bundles per FPD for the CANDU 6). Over short periods, there will be a considerable variation from this average rate. For example, if, for a few days, only the outermost channels were being refuelled, a very high rate (almost 6 channels per day for the CANDU 6) would be required to keep up with the reactivity loss, since these channels have a reactivity worth somewhat smaller than that of central channels. It is not

possible for the fuelling machines to maintain this rate, so that refuelling of outermost-channel would have to be intermingled with refuelling of high-worth, inner-region channels, or postponed to some convenient time. Fortunately, this region does not have to be visited very often - about 4 or 5 times per month on the average.

Figure 5.6 shows a plot of refuelling rate vs. time obtained from a simulation of the CANDU 6 reactor from the onset of refuelling to 400 FPD. The corresponding maximum channel powers and bundle powers are shown in Figures 5.7 and 5.8.

Figures 5.9 and 5.10 show the gradual increase in average discharge burnup and in cumulative discharge burnup through this period of operation.

The last figures also show that it is possible to increase the discharge burnup of the initial fuel load, and reduce the channel-visit rate, by adopting the swing-8 refuelling for the **first** visit (only) of each channel. The swing-8 scheme is shown in Figure 5.11; it features the reshuffling of fuel (bundles 11 and 12) back into the channel on refuelling, and is used to avoid the very low burnup of the original bundles 11 and 12 if they remained in core for only one cycle.

5.3 Equilibrium Fuel Management

Although refuelling rates quickly approach equilibrium values, an equilibrium burnup distribution is usually attained slightly later, at about 500 FPD. Because of this, the fuel discharged from the reactor does not reach its equilibrium burnup value until some time after the equilibrium refuelling rate has been attained. Thereafter, the global core-average characteristics (e.g. average burnup and global power distributions) change very little with time.

Refuelling is the primary method of reactivity addition and maintaining satisfactory power distribution in the core. The rate of refuelling is adjusted to compensate for the reduction in core reactivity caused by burnup. Refuelling rates in different core regions are adjusted to control the power distribution - a lower refuelling rate increases average burnup in a region, and reduces its relative power.

If the fuel were irradiated to a uniform average burnup throughout the core, the channel-power distribution would be radially "peaked", with a relatively small number of channel powers near the maximum value. As explained earlier, to increase the total reactor power with a given number of fuel channels, while respecting channel-power limits, a radially-flattened power distribution is promoted. This reduces the capital cost of the plant per kW installed.

Radial flattening of the power distribution is achieved by differential burnup, in combination with adjuster rods if the latter are provided for xenon override. To attain differential burnup, the core is divided into a number of "burnup regions" radially, with the fuel expressly driven to a different value of burnup in each region.

The on-power-refuelling capability of CANDU gives great flexibility in designing the burnup regions, and therefore the time-average power shape. In principle any number of separate burnup regions may be defined, each with a target average exit burnup. The number and size of burnup regions are at the discretion of the core designer and the fuelling engineer (in a somewhat impractical limit, every individual channel could be defined as a region). The refuelling rates in the various regions can be adjusted to give the desired power distribution.

5.3.1 Axial Refuelling Scheme

During a fuelling operation, new fuel bundles are added to one end of a channel, and old fuel bundles are moved along the channel to new positions. If a small number of bundles are added to the channel at each refuelling, the burnup more closely approaches that for the ideal "continuous" refuelling. Since this would give the most control as far as adding fuel to the core only when necessary, a 1-bundle-shift or 2-bundle-shift refuelling scheme would maximize burnup. However, it is found that only a small burnup loss is suffered when replacing even up to eight bundles per refuelling, while a 10-bundle-shift or 12-bundle-shift scheme produces a significant loss in burnup.

Increasing the number of bundles loaded per refuelling reduces the number of visits of the fuelling machine to the core, which, in turn, reduces operating and maintenance costs for this very important system. Thus, there is an incentive to refuel as many bundles as possible per refuelling, to reduce the load on the fuelling machine.

In order to minimize the load and the maintenance costs on the fuelling machine, the number of refuelling operations should be kept as small as possible. From this point of view, the more bundles loaded per cycle, the fewer visits the fuelling machine has to make. Refuelling schemes where a larger number of bundles is replaced are therefore desirable.

The eight-bundle-shift refuelling scheme is an excellent compromise between maximizing burnup and minimizing channel-visit rate and is the reference fuel-management scheme for the CANDU 6.

Mixed schemes can also be used in principle, e.g., refuelling outer-region channels with an eight-bundle-shift scheme and inner-region channels with a four-bundle-shift scheme.

5.3.2 Channel Selection for Refuelling

During reactor operation, one of the main functions of the fuel engineer is to establish a list of channels to be refuelled during the following period (few days) of operation. To achieve this, the status of the reactor core is determined from computer simulations of reactor operation, the on-line flux mapping system, the ROP and RRS

bundles are added to the channel at each refuelling, the burnup more closely approaches that for the ideal "continuous" refuelling. Since this would give the most control as far as adding fuel to the core only when necessary, a 1-bundle-shift or 2-bundle-shift refuelling scheme would maximize burnup. However, it is found that only a small burnup loss is suffered when replacing even up to eight bundles per refuelling, while a 10-bundle-shift or 12-bundle-shift scheme produces a significant loss in burnup.

Increasing the number of bundles loaded per refuelling reduces the number of visits of the fuelling machine to the core, which, in turn, reduces operating and maintenance costs for this very important system. Thus, there is an incentive to refuel as many bundles as possible per refuelling, to reduce the load on the fuelling machine.

In order to minimize the load and the maintenance costs on the fuelling machine, the number of refuelling operations should be kept as small as possible. From this point of view, the more bundles loaded per cycle, the fewer visits the fuelling machine has to make. Refuelling schemes where a larger number of bundles is replaced are therefore desirable.

The eight-bundle-shift refuelling scheme is an excellent compromise between maximizing burnup and minimizing channel-visit rate and is the reference fuel-management scheme for the CANDU 6.

Mixed schemes can also be used in principle, e.g., refuelling outer-region channels with an eight-bundle-shift scheme and inner-region channels with a four-bundle-shift scheme.

5.3.2 Channel Selection for Refuelling

During reactor operation, one of the main functions of the fuel engineer is to establish a list of channels to be refuelled during the following period (few days) of operation. To achieve this, the 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 produce the neutron flux, and the power and burnup of each fuel bundle and channel as a function of reactor history. The simulations also provide maximum bundle and channel powers in the reactor, so that the compliance with operational limits can be assessed.

The list of good-candidate channels selected for refuelling usually contains:

- 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 burnup from time-average calculation,

- channels with low power, relative time-average power distribution,
- channels in (relatively) low-power zones (compared to 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 result in acceptable values for the individual zone-controller fills (20%-70% range), and
- channels which, together, provide sufficient reactivity gain to leave the zone-controller fills in the desired operational range (average zone fill between 40 and 60%).

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.

The discrete refuelling in CANDU reactors results in variability in the instantaneous maximum channel and bundle powers. This is illustrated in Figure 5.12, which shows a schematic plot of the maximum channel power versus operating time and illustrates the difference between maximum time-average channel power, average maximum instantaneous channel power, and absolute maximum channel power.

5.4 The Channel-Power Cycle

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

- When a channel is refuelled, its local reactivity is high, and its power will be several percent higher than its time-average power.

- The fresh fuel in the channel then initially goes through its plutonium peak as it picks up irradiation. This means that in fact the local reactivity **increases** for about 40 to 50 FPD, and the power of the channel tends to increase further. The higher local reactivity tends to promote a power increase in the neighbouring channels also.

- Following the plutonium peak, the reactivity of the refuelled channel starts to decrease, and its power drops slowly. Approximately half-way through its dwell time, the power of the channel may be close to the power suggested by the time-average model.

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

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

5.5 Channel-Power Peaking Factor

At any given time, there are several channels in the core which are at or near the maximum power in their cycle. Therefore, the maximum instantaneous channel power is always higher than the maximum time-average channel power, as was evident from the earlier Figure 5.12.

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

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

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. This is the reason for the factor D_o in the equation which shows the trip setpoint that is needed at a detector j for that detector to provide protection against a loss of regulation (LPR) from a flux configuration k:

$$TSP_j = D_o \frac{\phi_j(k)}{\phi_o(j)} r_{CPR}(k)$$

In this equation:

- the first factor is the detector calibration factor (CPPF), i.e., essentially the instantaneous detector reading,
- the second factor is the ratio of perturbed to unperturbed flux (due to the perturbation assumed) at the detector, in the time-average model, at 100% FP, and
- the third factor r_{CPR} is the critical power ratio in the perturbed power shape in the time-average model.

The product of the three factors gives the instantaneous detector reading (in the perturbed shape) at the power at which the first channel would reach fuel dryout, that is, it indicates the reading at which the particular detector would have to trip in order to protect against a loss of regulation in the flux shape considered.

The "distance" from actual detector reading to the trip setpoint (i.e., ratio of setpoint to instantaneous reading) is the "margin to trip", or "operating margin". The reactor operator likes to have as great an operating margin as possible because that provides manoeuvering room and reduces the probability of spurious reactor trips.

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.

5.6 Fuel Performance

Another factor considered in selecting the axial refuelling scheme is the fuel performance, specifically the probability of fuel defects. If irradiated fuel bundles are moved along the channel from a low-power position to a high-power position, the sudden boost in power (power ramp) may cause the fuel cladding to rupture. The failure probability normally depends on both the final power and the power boost experienced, and is also a function of irradiation. Such a fuel failure would allow radioactive fission products to enter the PHT system and produce high radiation dose rates near the equipment. Of course, fuel defects are rare in CANDU fuel (and, when they occur, are usually found to be due to some manufacturing defect), and on-power refuelling provides the advantage that defected bundles can be removed promptly from core as soon as they are identified. However, fuel defects are still very undesirable; they are a big nuisance, and, furthermore, they reduce fuel burnup. Refuelling simulations at the design stage provide good statistics on the expected power boost for specific axial refuelling schemes.

The risk of fuel defects results for instance in a 6-bundle-shift refuelling scheme not usually being a good strategy. A 6-bundle shift would push bundles from position 1, the lowest-power position in the channel, to position 7, a very high-power position, and the power boost would therefore be high.

5.7 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. To maintain criticality, the reactor regulating system (RRS) would first tend to drain the liquid zone-control compartments to the lower limit of their control range.

Since the desirable operating range of the zone controllers is at most between 20% and 70%, 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 be 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.

In the unlikely event that the fuelling machine remains unavailable for several weeks, requiring most of the adjuster rods to be withdrawn, a significant change in the fuel burnup distribution, relative to the nominal equilibrium condition, would occur. When the fuelling machine becomes available once again, a greater than normal refuelling rate would probably persist until sufficient reactivity is provided to permit re-insertion of the adjuster rods and hence restoration of full power. The selection of channels for refuelling would be based on the burnup distribution existing when refuelling is resumed.

6. Fuel Burnup and Effect of Operating Conditions

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

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

Burnup is related to irradiation in an almost linear fashion, as shown in Figure 6.1.

Note that, **for a given type of fuel and reactor**, fuel burnup is essentially the inverse of fuel utilization, i.e., the amount of fuel used to produce a given quantity of energy (electricity), measured for example in units of Mg(U)/GW(e).a. High fuel burnup is good, low uranium tilization is good. However, note that when comparing different fuels or reactors, a higher fuel burnup does not necessarily mean a lower uranium utilization. For instance, the fuel burnup attained in PWRs is much higher than that attained in CANDU, but this is the result of fuel enrichment, not of higher fuel efficiency. In fact, even though PWR fuel burnup may be 3 to 6 times that in CANDU, the uranium utilization is lower by some 25-28% in CANDU, due to its neutron economy.

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

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

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

The operating parameters which probably have the greatest influence on the attainable fuel burnup, and their typical reactivity effect, are:

- the moderator purity, ~ 34 milli-k/atom % purity
- the coolant purity, ~ 3 milli-k/atom % purity
- the moderator poison concentration, ~ 8 milli-k/ppm(B)

The core flux shape also affects fuel burnup, mostly through the amount of neutron leakage. Operating the reactor with a flatter flux distribution gives higher neutron leakage, therefore lower reactivity, therefore lower burnup.

The type of fuel used also has an influence on burnup. For instance, 28-element fuel (see Figure 6.2) has a higher ratio of fuel mass to sheath mass (and therefore less parasitic absorption), and therefore provides somewhat higher lattice reactivity. But because the fuel pins in 28-element fuel are thicker and cooling is not as efficient, 28-element fuel cannot be operated at as high a power. From the point of view of fuel burnup, the consequently lower fuel temperature also results in higher reactivity (less resonance absorption because of reduced Doppler broadening of resonances), which further increases attainable fuel burnup. The fuel burnup in Pickering A, which uses 28-element fuel, is about 8,000 MW.d/Mg(U). However, the disadvantage of course is the total reactor power must be lower with 28-element fuel (for the same number of channels): Pickering A full power is 540 MW(e), compared to 680 MW(e) for the CANDU 6.

The CANDU 6 has adjusters with a reactivity worth of approximately 15 milli-k. A CANDU reactor which is designed without adjusters has a higher excess reactivity and therefore provides higher burnup. The Bruce A reactor, without adjusters, can attain burnup values of 9,000 MW.d/Mg(U). Of course, operating without adjusters also has some disadvantages, most notably the lack of xenon-override capability.

7. Fuel-Burnup Warranties

Utilities which buy reactors are interested in attaining as high a fuel burnup as possible. Also, they often ask for warranties on the burnup which is "promised" by the vendor.

There are three distinct types of fuel-burnup warranty which have been offered and negotiated for the CANDU 6 reactors. In each case, the warranted value is related to a set of well-defined reactor operating conditions. A procedure is written, and mutually agreed upon with the client, according to which the warranty demonstration test is to be performed and various corrections are to be calculated to account for differences in actual operating conditions and the conditions stated in the warranty. The procedure specifies the measurements to be performed during the warranty demonstration test period including recording of various operating conditions.

Following is a brief description of the three types of fuel burnup warranty.

7.1 Specific-Fuel-Consumption Warranty

In this type, the warranted value is the amount of fuel consumed per unit of thermal energy produced after equilibrium fuel conditions have been established. The duration of the warranty demonstration test is 150 FPD. During this period, the amount of steam produced at the steam generators is measured along with the number of fuel bundles added to (or discharged from) the reactor core. A record of various operating conditions (parameters) is carefully maintained. These operating conditions would include for instance the moderator isotopic purity, the coolant isotopic purity, the amount of poison in the moderator, etc.

The measured specific fuel consumption is corrected to account for differences in actual operating conditions from conditions specified in the warranty. The corrected specific consumption is compared with the warranted value to demonstrate compliance.

7.2 Initial-Fuel-Load Warranty

In this type, the warranted value is the thermal energy produced by the first 4560 fuel bundles discharged (equal to the number of fuel bundles in the initial load). Since the energy produced by these bundles cannot be measured directly, this type of warranty requires RFSP computer simulations to calculate the energy produced by each individual bundle. A careful record of operating conditions is again maintained and the information is used directly in the computer simulations, or suitable corrections are applied to the results obtained in the simulations. Compliance with the warranty is demonstrated by comparing the corrected value with the warranted value.

7.3 Nuclear Characteristics Warranty

The warranted value in this warranty is the boron concentration required for first criticality. The nuclear characteristics of any nuclear reactor determine the excess reactivity of the core when loaded with unirradiated fuel. In CANDU reactors, this

excess reactivity can be determined by measuring the boron concentration required in the moderator at first criticality. The excess reactivity from that point on depends primarily on the power history and the fuel-management practice. In CANDU reactors, where on-line refuelling is the primary means of reactivity control, the excess reactivity is kept near zero after refuelling starts by replacing fuel at an appropriate rate. Following a transitional period, the refuelling reaches an equilibrium state where the refuelling rate is constant provided the characteristics of the non-fuel core components are unchanged. The refuelling rate is then dependent on the nuclear characteristics of the core only. There is, therefore, a direct relationship between the initial boron concentration in the moderator and the achieved equilibrium refuelling rate (and hence equilibrium fuel burnup) in CANDU reactors.

An obvious advantage of this type of warranty is that it can be discharged immediately after the reactor is commissioned, and detailed measurements over an extended period are not required.

8. Cobalt-60 Production

Cobalt-60 is a radioactive isotope of cobalt. It emits gamma radiation and is used throughout the world in radiation therapy and radiographic equipment.

Cobalt-60 production in the CANDU reactors is achieved by using adjuster rods made of cobalt.

Adjuster rods are normally positioned inside the reactor core and fulfil the following basic functions:

- they provide, on withdrawal, positive reactivity to override the growth of xenon poison immediately after a power reduction or a shutdown, or to compensate for reactivity decay during operation without refuelling.
- they provide the desired flux and power flattening by appropriate distribution of absorption in the various rods

The neutron absorption required in the adjuster design has been provided in the CANDU reactors by using either cobalt or stainless steel. The Picketing 'A' reactors operated with cobalt adjusters for over a decade. Of the four CANDU 6 reactors, Point Lepreau and Wolsong are equipped with stainless-steel adjusters, whereas Cordoba has cobalt adjusters, and G-2 has operated at various times (including currently) with cobalt adjusters.

Cobalt adjusters consist of cobalt "bundles" placed in Zircaloy tubes. These bundles are ~8" long and ~2.5" in diameter. Each bundle has one to six "pencils" equidistantly located on a pitch circle of ~2" diameter. Each pencil in turn is a Zircaloy

tube containing cobalt slugs or pellets. The adjusters are irradiated in the reactor for a period of about 1.5 years before replacement.

Typical Cobalt-60 production in a CANDU reactor is ~ 5.5 MCi per year, with a specific activity of 80 Ci/g after one year of operation at 80% capacity factor.

9. Effects of Xe-135 Poison

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

9.1 The Xe-I Kinetics

The Xe-135/I-135 kinetics are shown pictorially in Figure 9.1.

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

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

The Xe-135/I-135 kinetics in any particular fuel bundle can thus be represented by the following equations:

$$\frac{dI}{dt} = \gamma_i \bar{\Sigma}_f \bar{\phi}_F - \lambda_i I \tag{9.1}$$

$$\frac{dX}{dt} = \gamma_x \bar{\Sigma}_f \bar{\phi}_F + \lambda_i I - \lambda_x X - \bar{\sigma}_X X \bar{\phi}_F$$

where

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

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

γ_x = direct yield of Xe-135 per fission (averaged over all fissions)

γ_I = direct yield of I-135 in fission, including yields of Te-135 and Sb-135 (averaged over all fissions)

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

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

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

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

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

In the above equations, the term $\gamma_i \bar{\Sigma}_f \bar{\phi}_F$ gives the I-135 production rate, while $\lambda_i I$ gives the I-135 loss rate (and the production rate of Xe-135 by iodine decay). Similarly, the term $\gamma_x \bar{\Sigma}_f \bar{\phi}_F$ gives the production rate of Xe-135 due to direct fission, while $\lambda_x X$ gives its decay rate. The term $\bar{\sigma}_x X \bar{\phi}_F$ represents the "destruction" (burnout) rate of Xe-135 due to neutron capture. Because of the comparable magnitudes of the various terms, the Xe-135 concentration is very sensitive to changes in flux level. The large absorption cross section of Xe-135 can significantly affect the value of reactivity of the reactor.

9.2 Reactor Startup

On starting up a reactor Xe-135 will build up according to the equations derived above. In Figure 9.2, the variation of reactivity as a function of time following startup is given for different steady state power levels. It can be seen that it takes ~40 hrs for equilibrium Xe-135 to be established.

9.3 Steady-State Xenon Load

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

$$I_{ss} = \frac{\gamma_i \bar{\Sigma}_f \bar{\phi}_F}{\lambda_i} \qquad (9.2)$$

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

It is obvious from Equation (9.2) that, as a function of an increasing fuel flux $\bar{\phi}_F$, the steady-state I-135 concentration increases indefinitely, while in contrast the steady-state Xe-135 concentration tends to an asymptotic value which will be denoted $X_{ss,\infty}$:

$$X_{ss,\infty} = \frac{(\gamma_i + \gamma_x)\,\bar{\Sigma}_f}{\bar{\sigma}_x} \qquad (9.3)$$

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

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

9.4 Effect of Power Changes on Xenon Concentration

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

Generally speaking, when the power is reduced from a steady level, the Xe-135 concentration increases at first. This is due to the fact that Xe-135 is still being produced by the decay of I-135, but its burnout rate (by neutron absorption) is decreased because of the reduced neutron flux (reduced power). However, after a certain period (whose length depends on the initial and final power and the rate of power reduction) the I-135 decay rate will have reduced sufficiently that the rate of Xe-135 production will have dropped below the rate of Xe-135 decay (and burnout), i.e., the Xe-135 concentration will have begun to decrease towards a new (lower) steady-state level (after having gone through a peak value).

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

Some typical reactivity variations due to Xe-135 following step changes in power are shown in Figure 9.3. It should also be appreciated that Xe-135 is spatially distributed. The absorption effects will therefore be difficult at different points in the core. For these effects a point kinetics treatment is often inadequate.

9.5 Xenon Transient Following a Shutdown

Following a shutdown, the burnout of Xe-135 stops, whereas the production by means of I-135 decay continues for several hours. The net result is that there is an initial increase in Xe-135 concentration and a decrease in core reactivity. If it is required to start up shortly after shutdown, extra positive reactivity must be supplied. The Xe-135 growth and decay following a shutdown for a typical CANDU is shown in Figure 9.4.

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

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

9.6 Effects of Xenon on Power Distribution

Xenon also plays a role in the power distribution. Because the steady-state Xe-135 concentration depends on the flux (Eq. 9.2), high-power bundles will have a higher xenon load, and therefore a lower reactivity, than low-power bundles of the same irradiation. The effect of xenon is therefore to **flatten** the power distribution: the reduction in the maximum bundle power due to the local Xe-135 concentration can be of the order of 5%, and should be taken into account when accurate results are desired.

The xenon effect also plays a role after refuelling. Because there is no xenon in fresh bundles, and it takes a day or so for the xenon to build in, the power of fresh bundles and channels that have just been refuelled is higher than it will be a day or so later. The difference in power could again be of the order of several percent, and needs to be taken into account in the most accurate calculations.

10. Summary

The fuel management in CANDU is a question of daily concern, since the refuelling is on-power. The increased load on the site reactor physicist is compensated by a great flexibility in managing the core and shaping the power distribution. Calculations must be made very frequently, almost every day, to follow the core power history and determine the CPPF value, which is used as calibration factor for the ROP detectors.

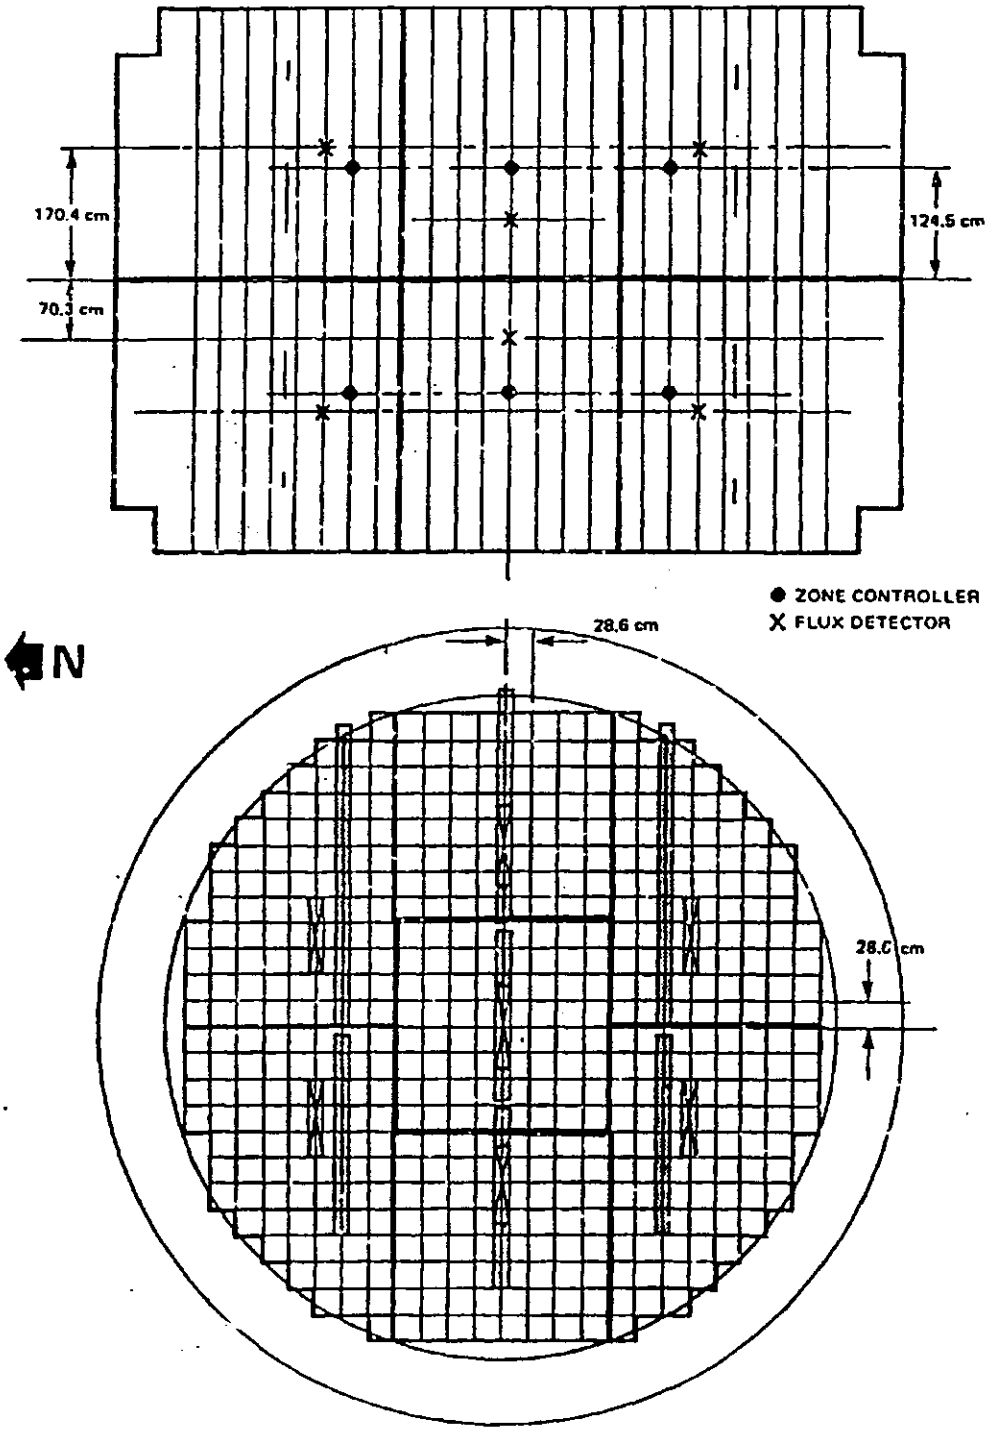

Figure 5.13 Reactor Zones and Controllers

I 1 a

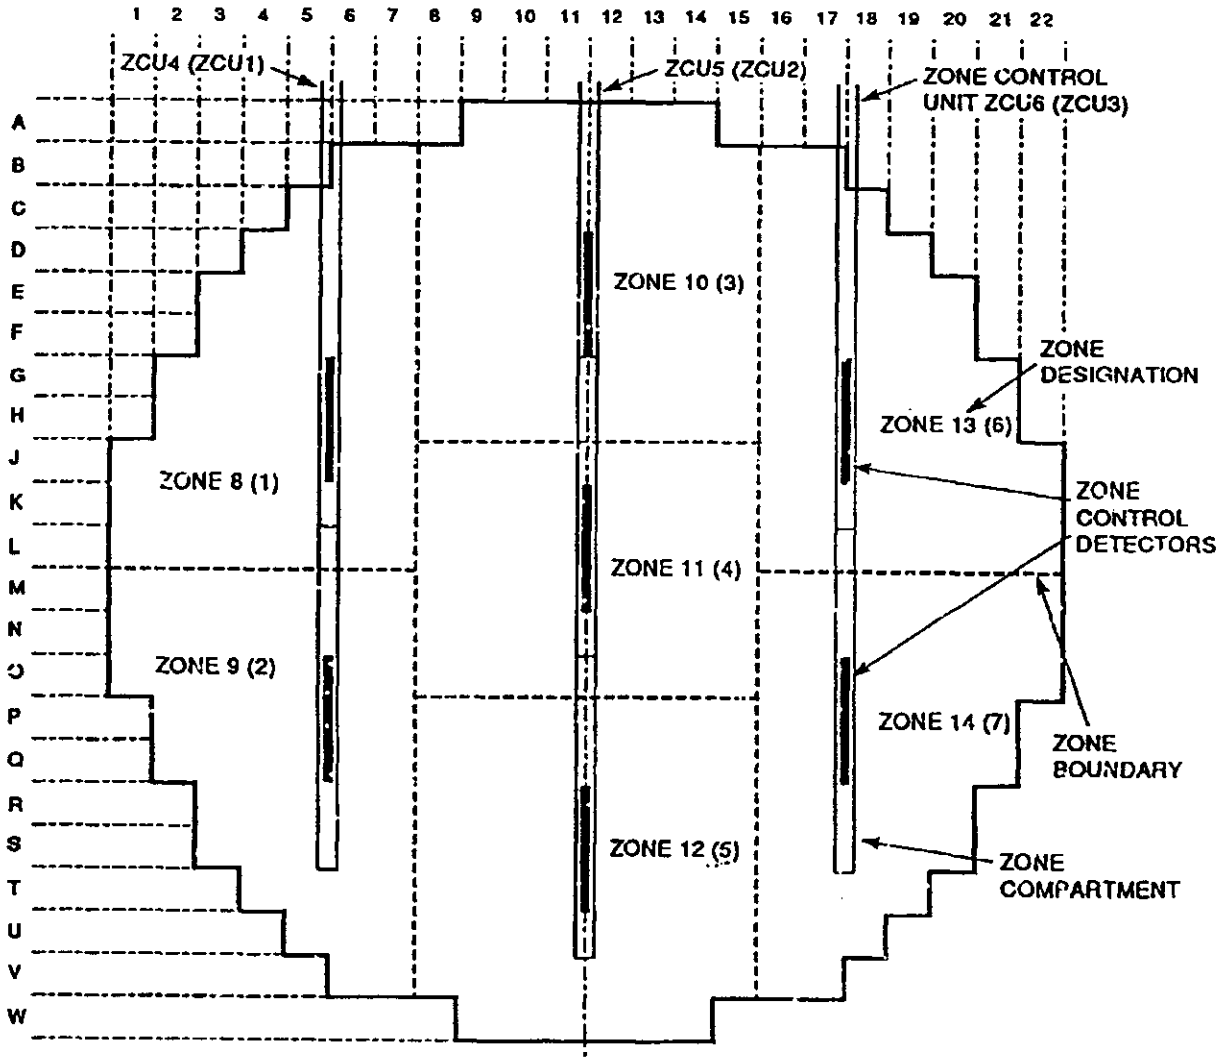

NOTES:

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 5.14 CANDU 6 Zone Locations

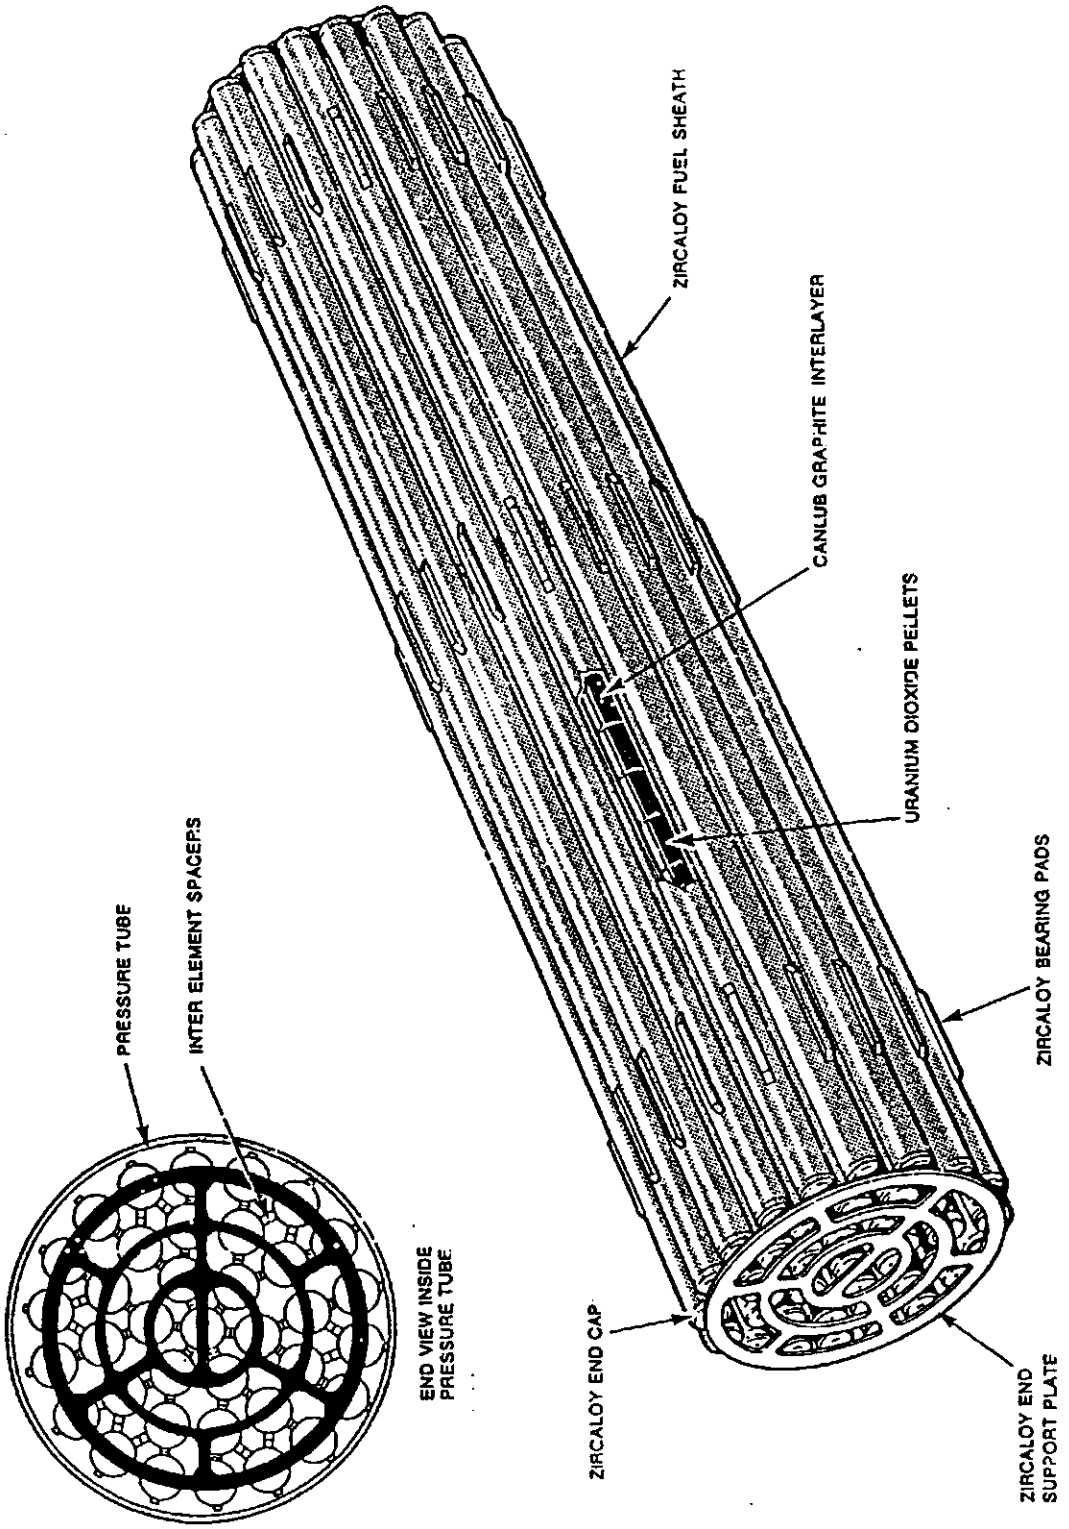

Figure 5.1 37-Element Fuel Bundle in 10 cm Channel

1.2

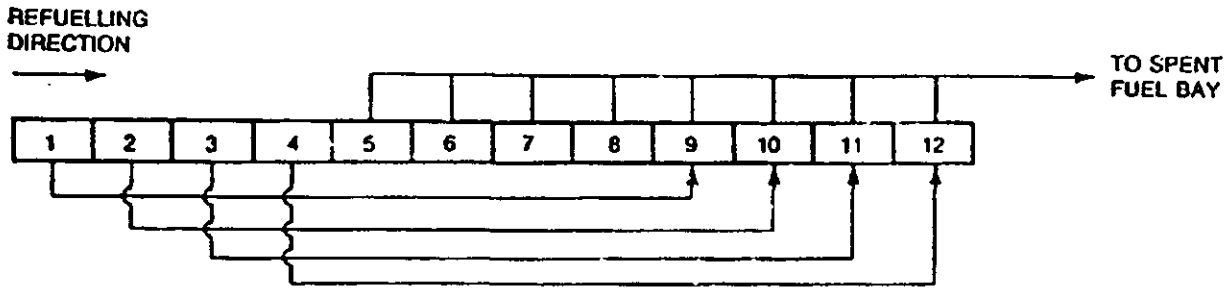

Figure 8.1 Standard 8-Bundle-Shift Refuelling Scheme
1.3

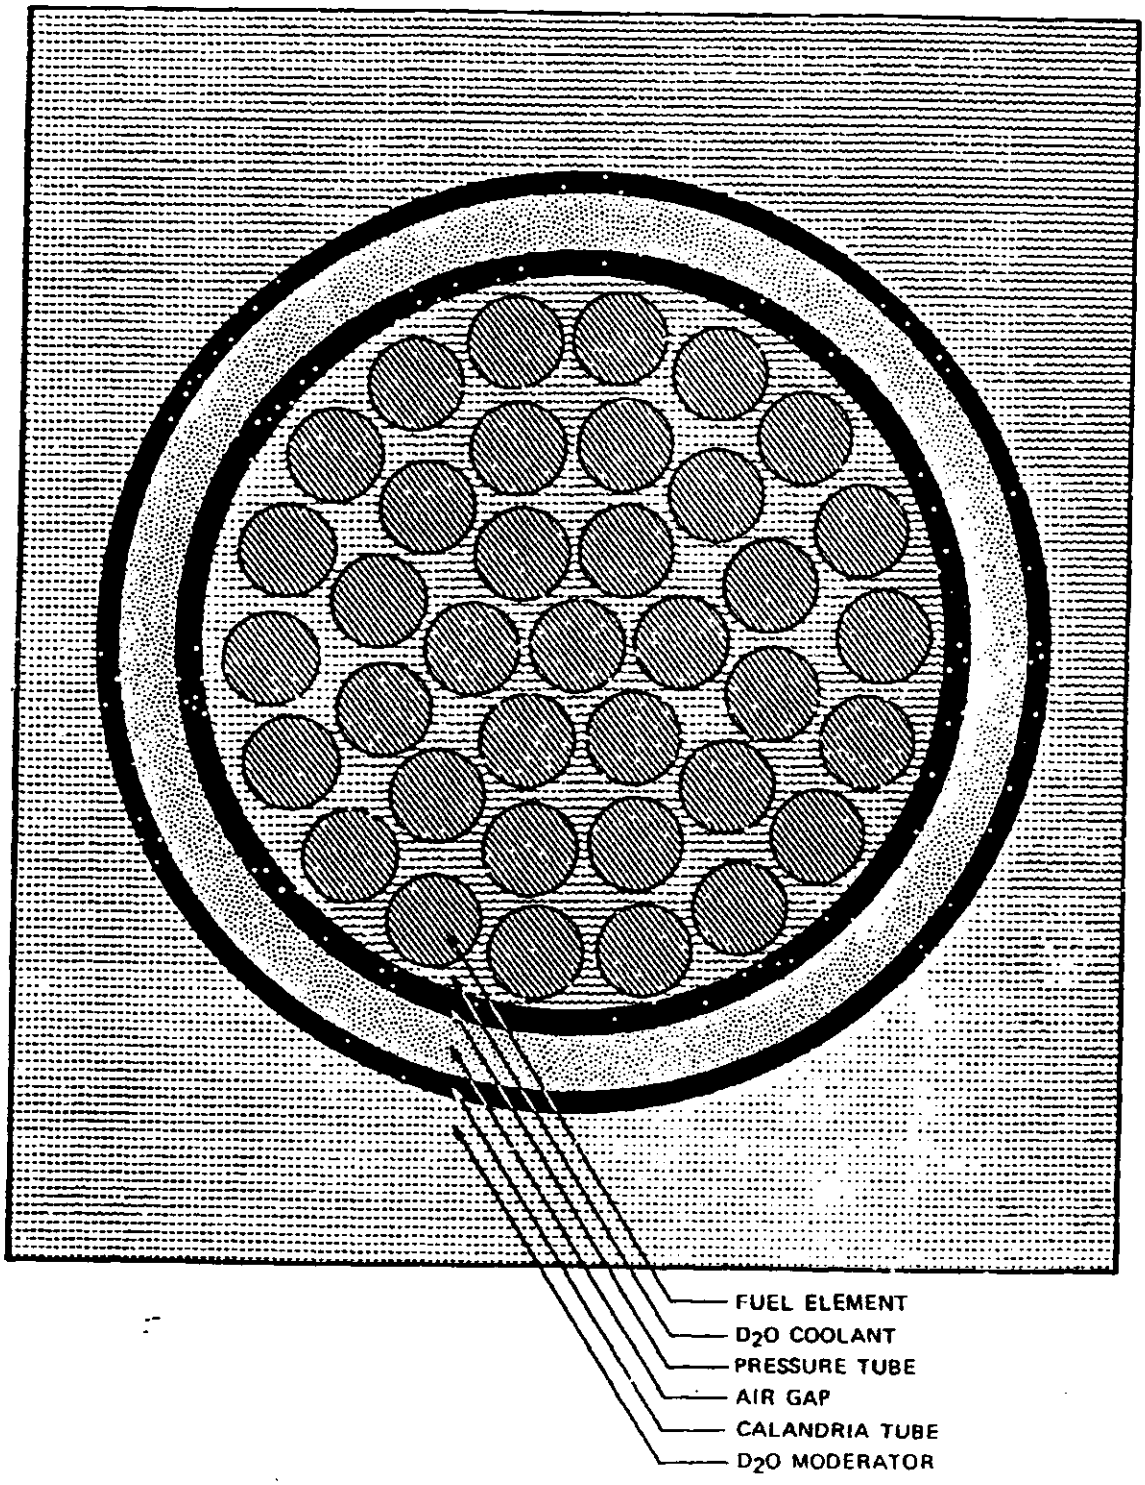

Figure 4.3 Lattice Cell for 37-Element Fuel

4.1

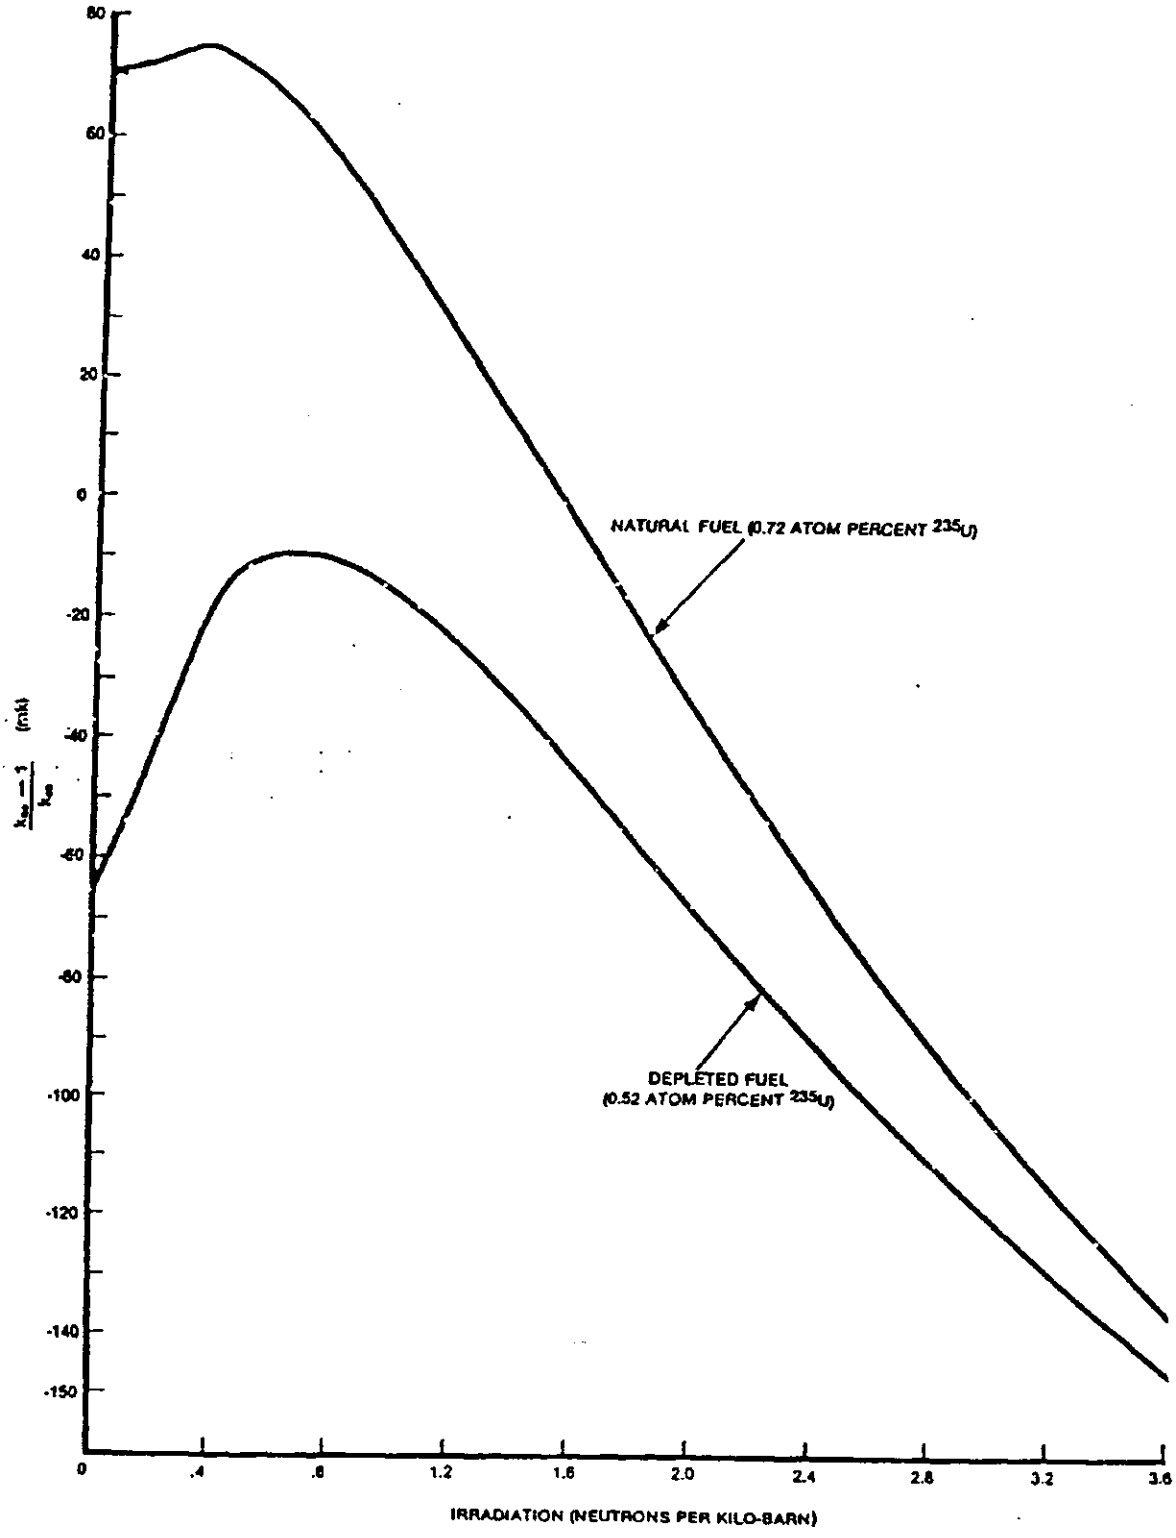

Figure 8.2 Reactivity versus Irradiation of Natural and Depleted 37-Element Fuel

4.2

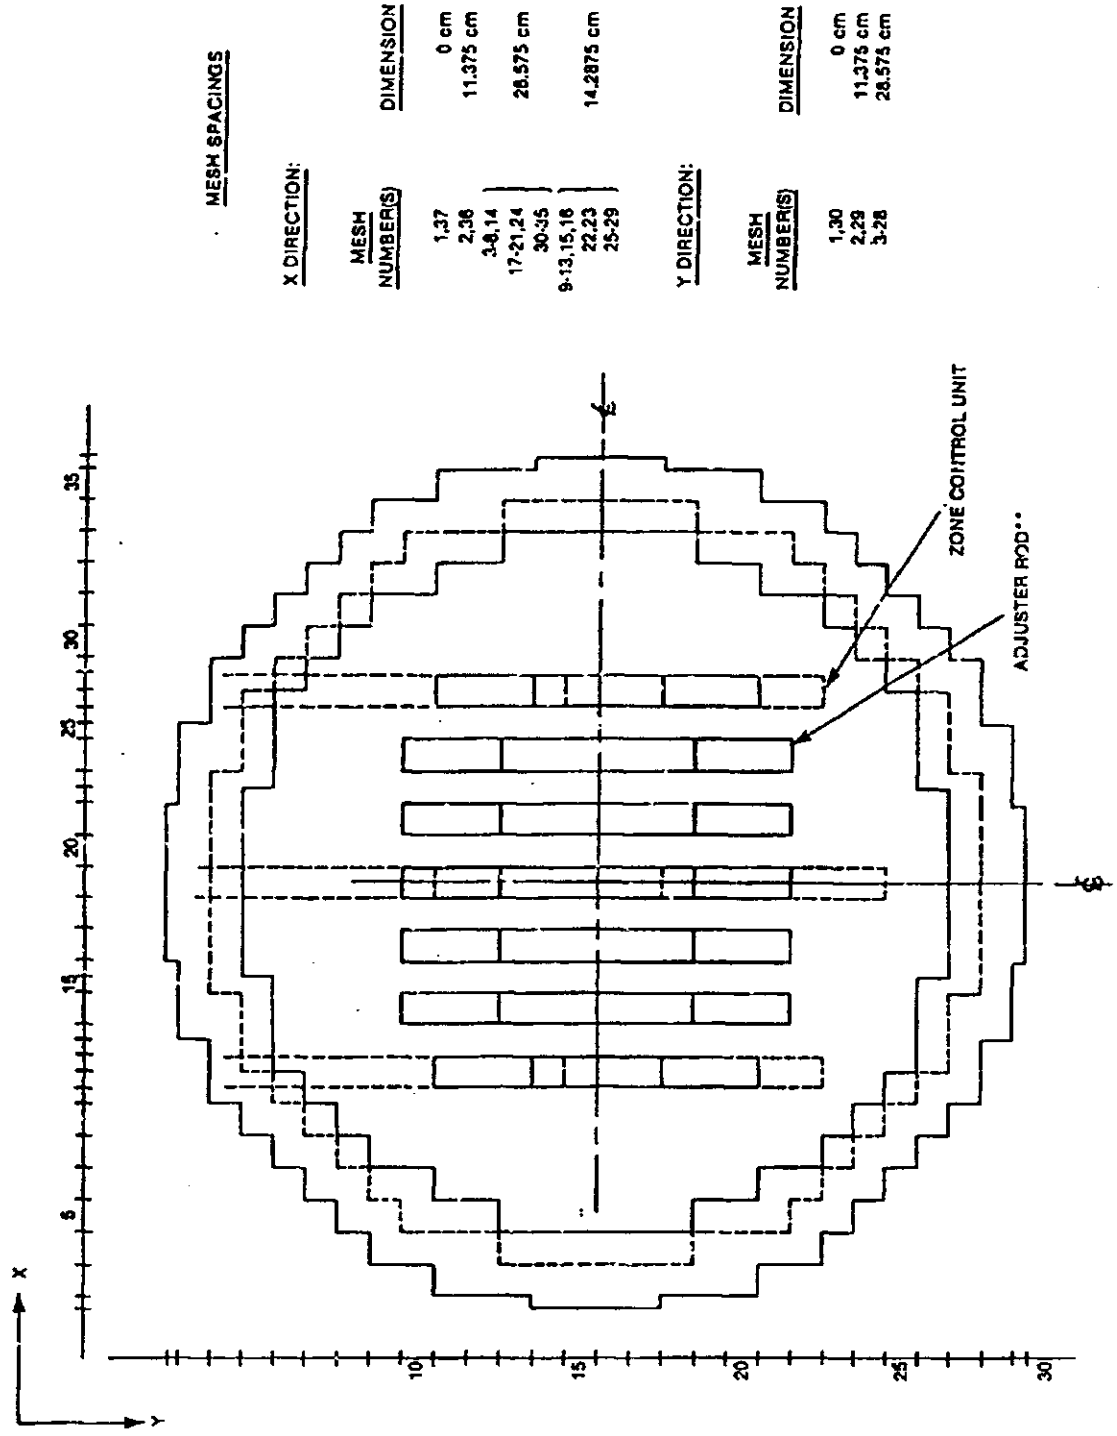

Figure 8.13 (a) CANDU 6 Full Core Model (x-y Direction)

+3

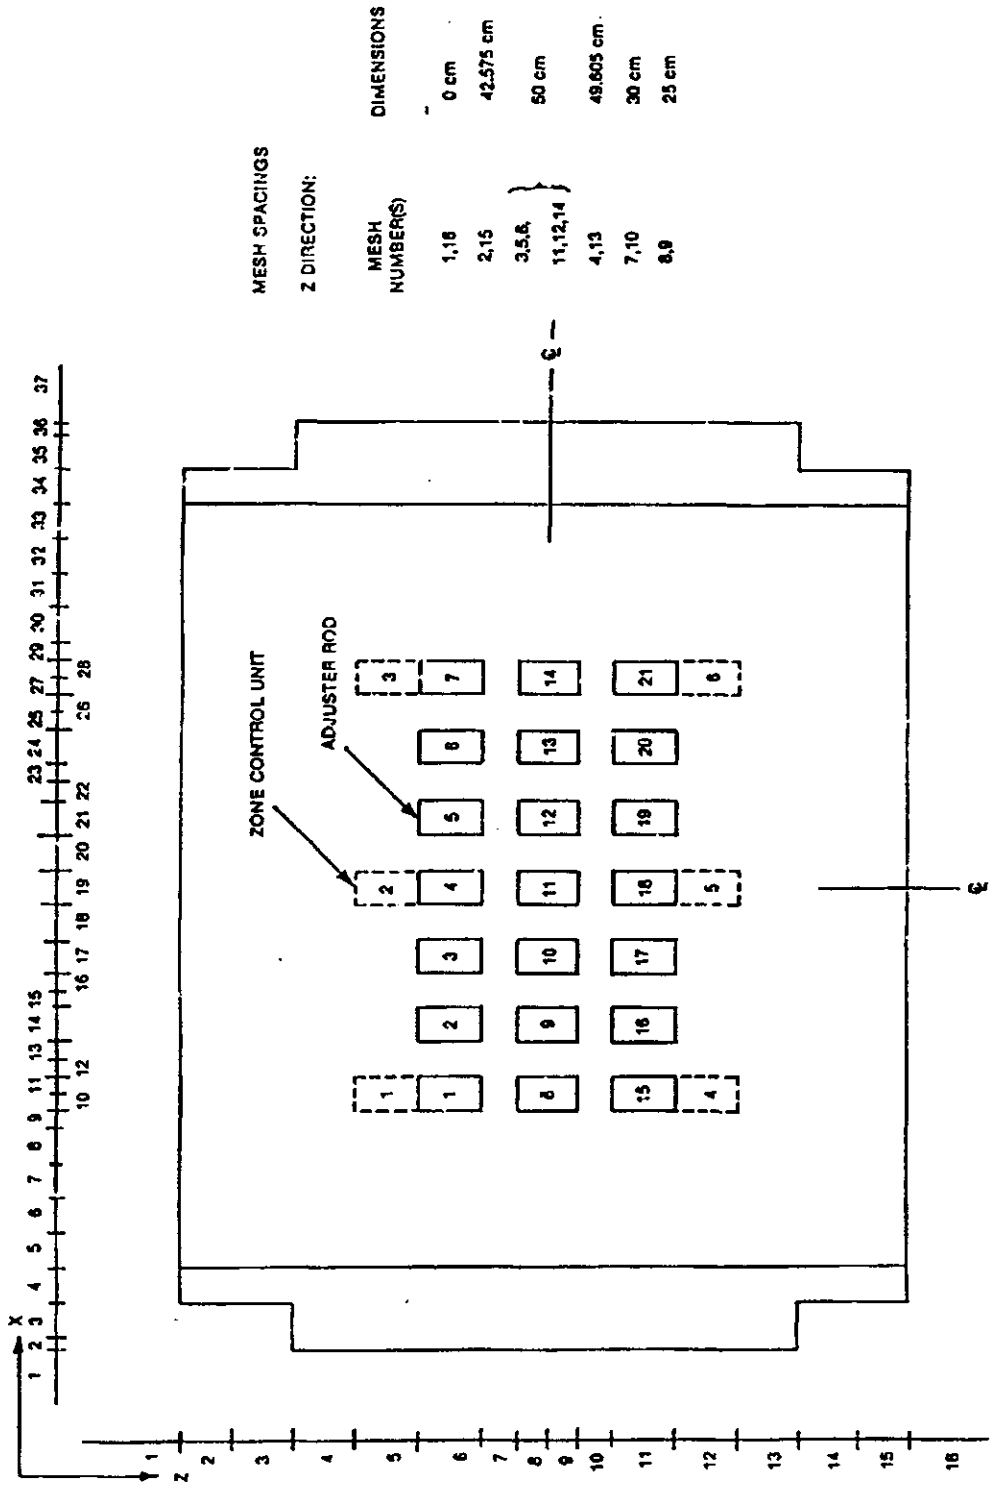

Figure 8.13 (b) CANDU 6 Full Core Model (x-z Direction)

43

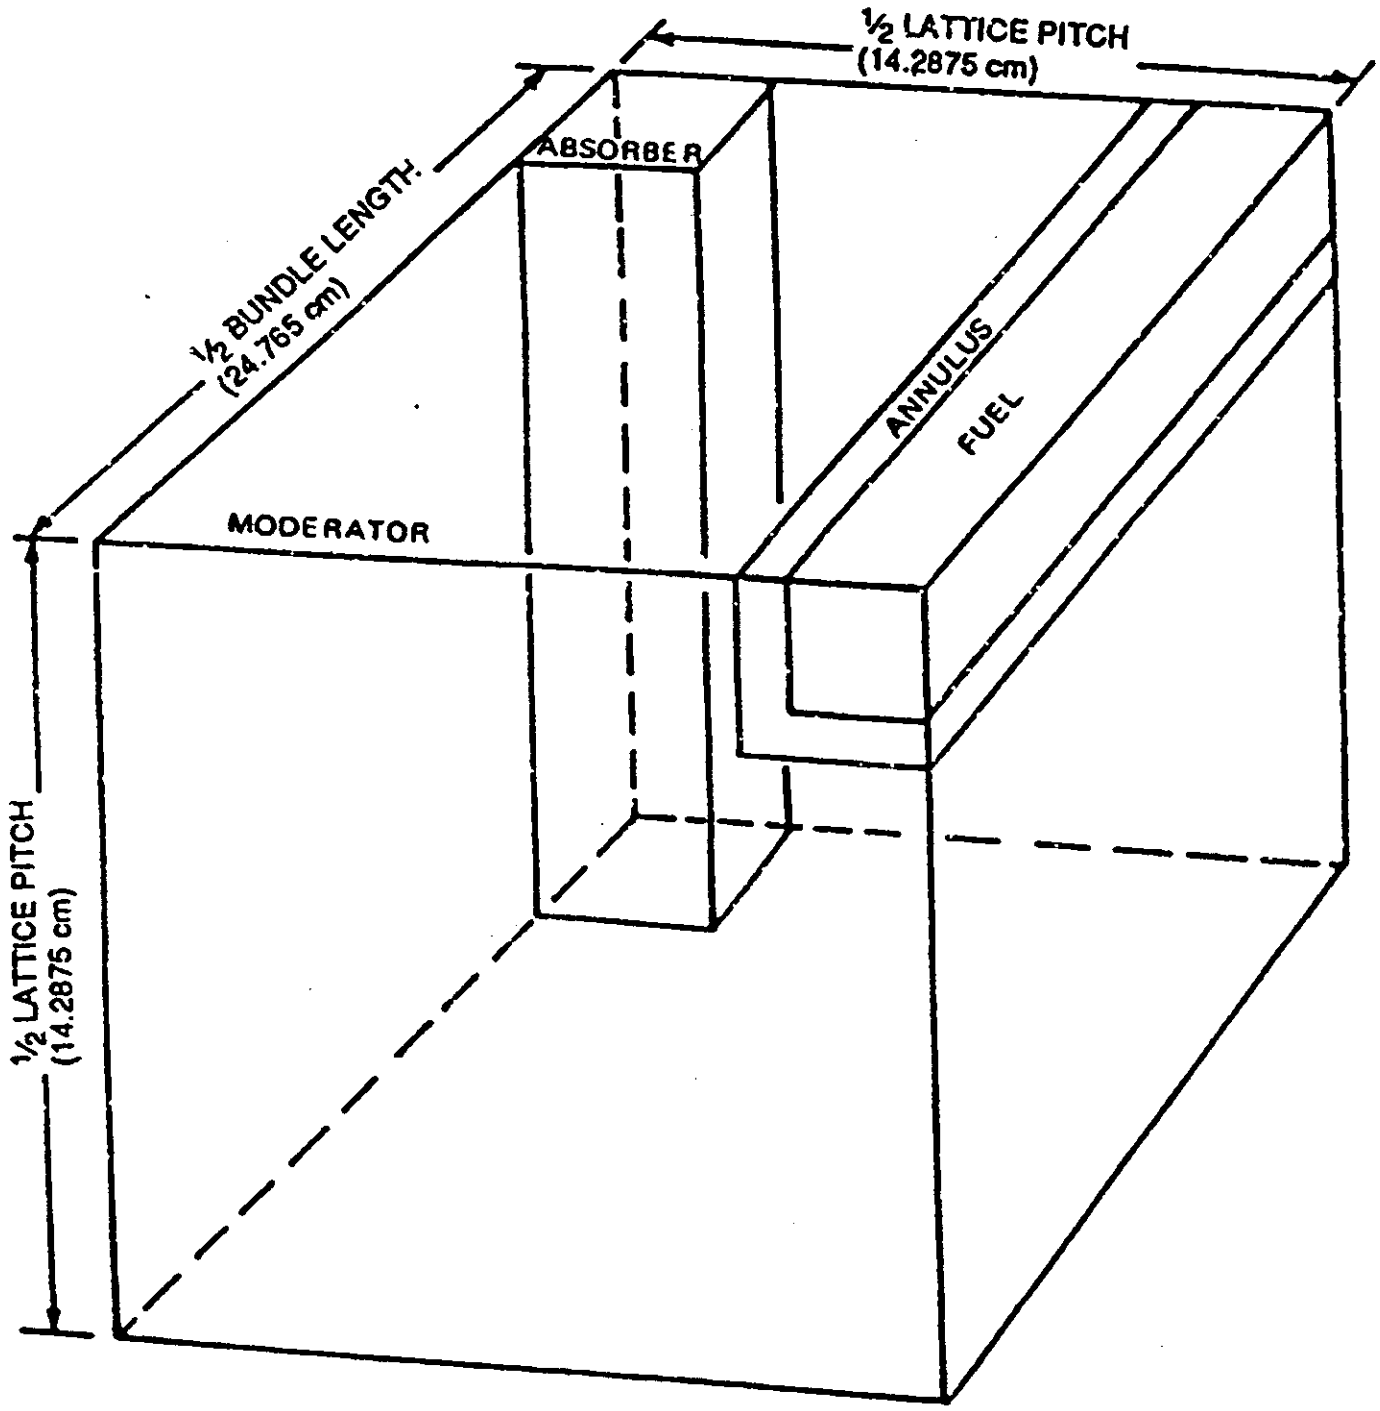

Figure 2.2-6: Typical Supercell Model Used for Multicell Calculation of Device Incremental Cross Sections

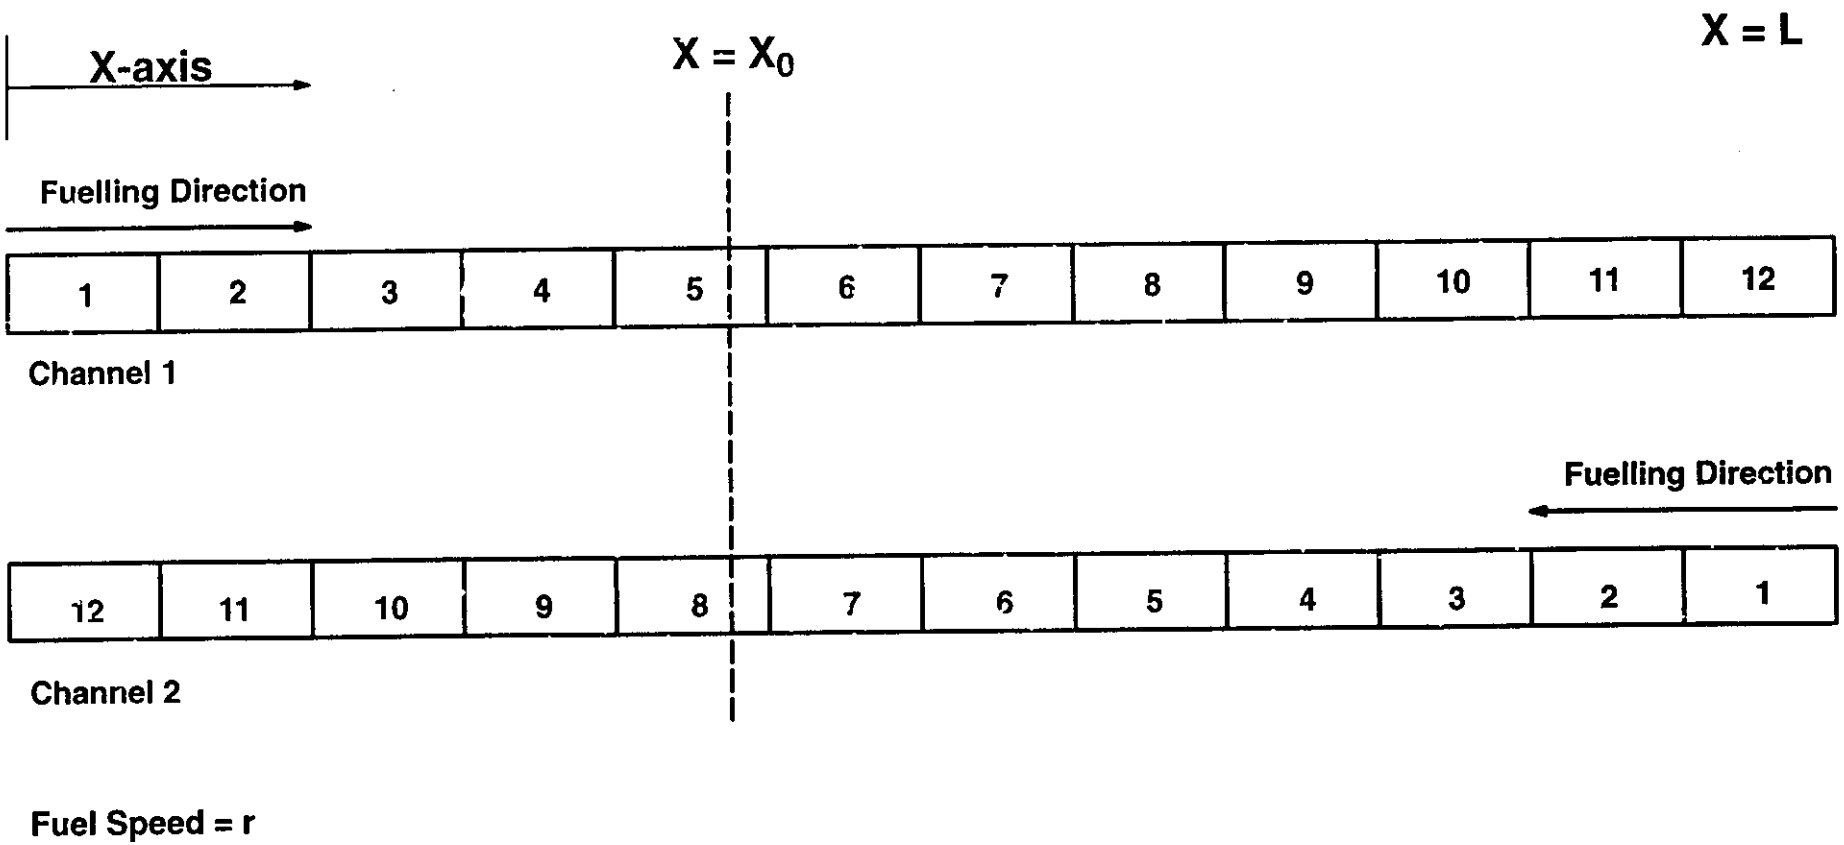

Figure 4.1~~1~~ Bidirectional Fuelling in CANDU in Approximation of Continuous Refuelling ("Homogeneous" Model)

4.5

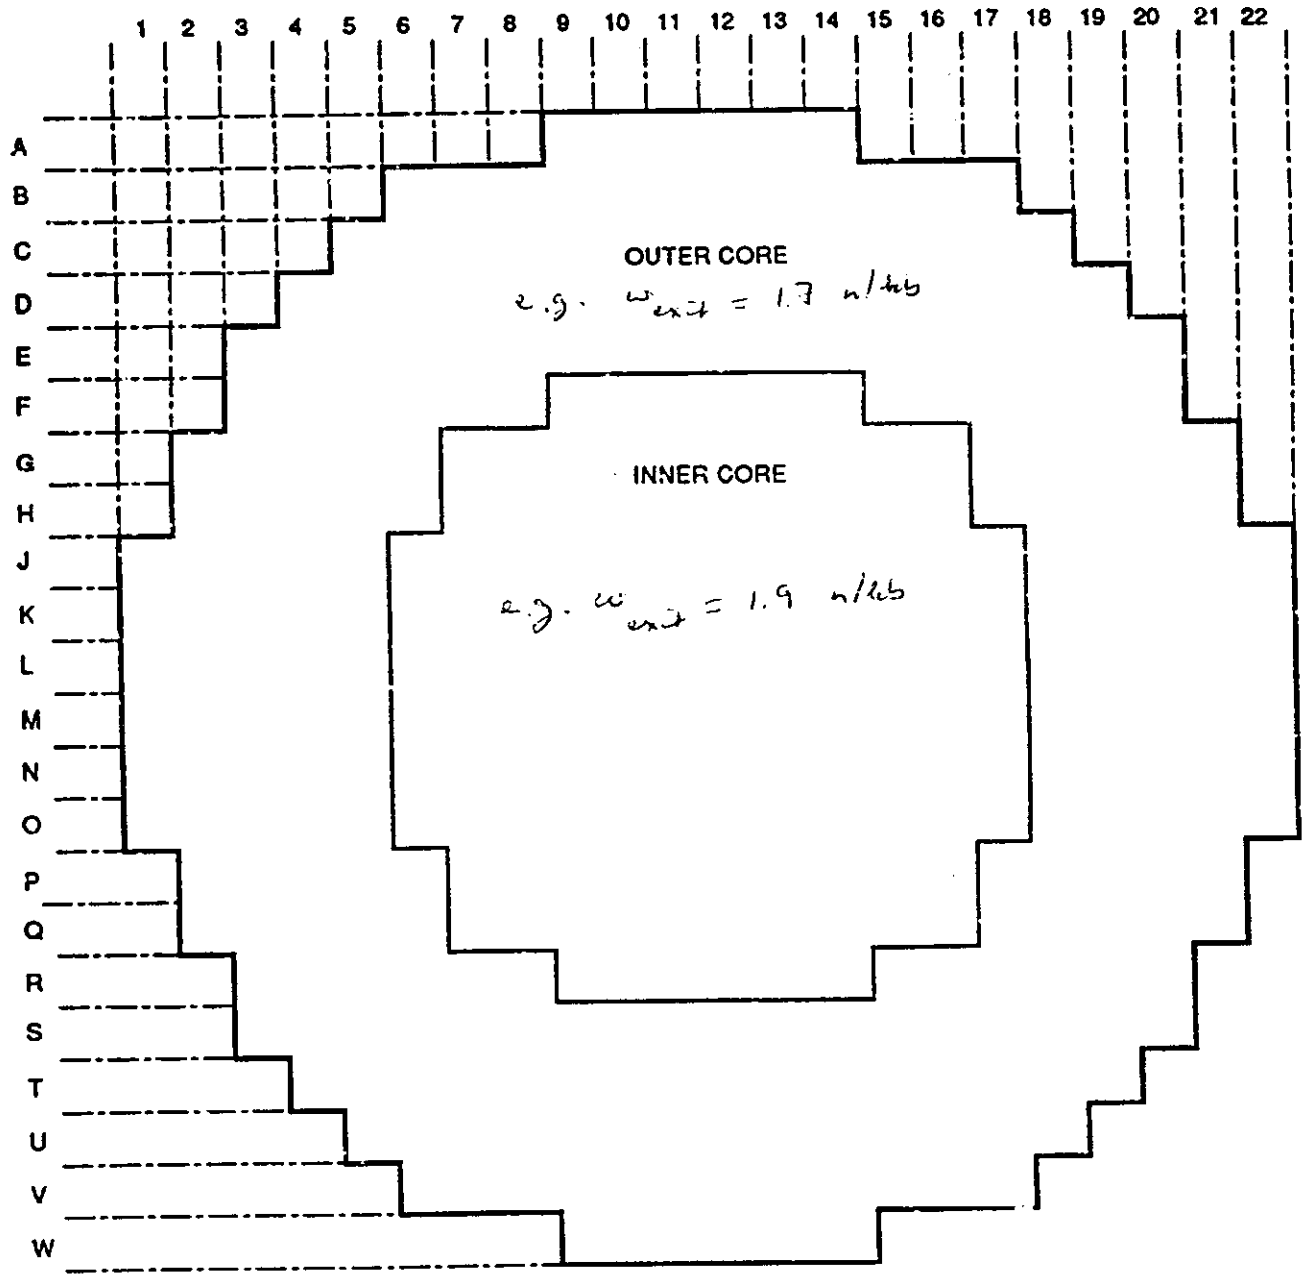

Figure ~~4.1-2~~ 4.6: Two-Region Time-Average Model for CANDU 6

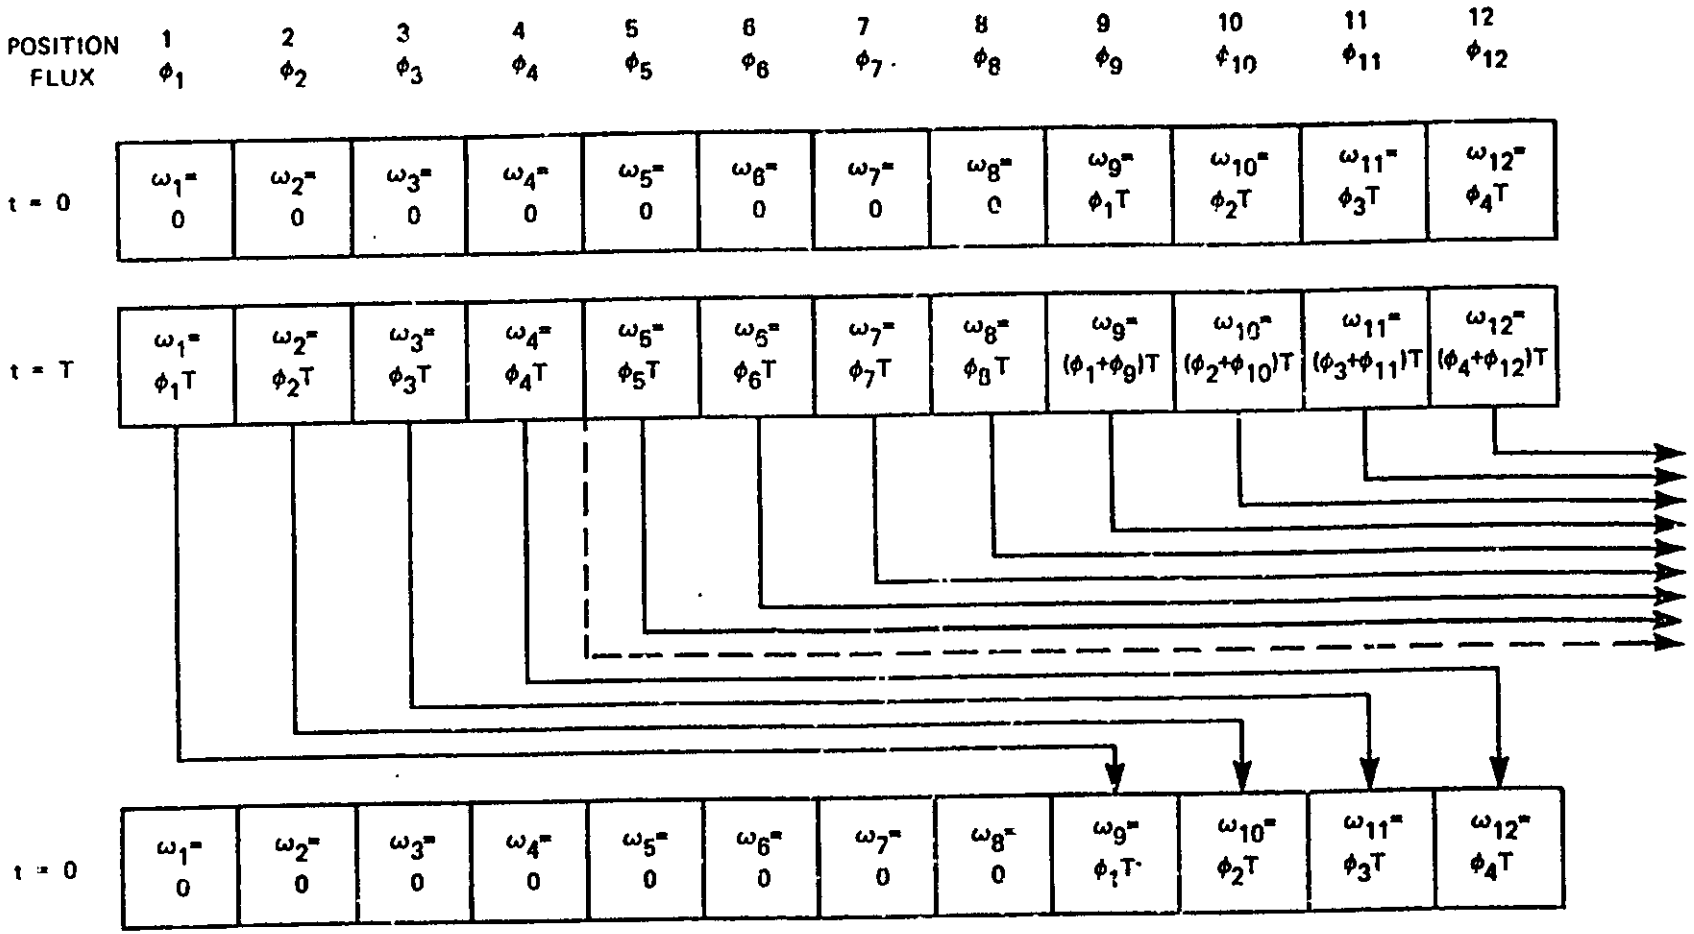

AVERAGE DISCHARGE IRRADIATION

$$= \frac{1}{8} \left\{ \phi_5 T + \phi_6 T + \phi_7 T + \phi_8 T + (\phi_1 + \phi_9)T + (\phi_2 + \phi_{10})T + (\phi_3 + \phi_{11})T + (\phi_4 + \phi_{12})T \right\}$$

$$= \frac{T}{8} \left\{ \phi_1 + \phi_2 + \phi_3 + \phi_4 + \phi_5 + \phi_6 + \phi_7 + \phi_8 + \phi_9 + \phi_{10} + \phi_{10} + \phi_{11} + \phi_{12} \right\}$$

Figure 4.1-3: Eight-Bundle-Shift Channel in Time-Average Mode

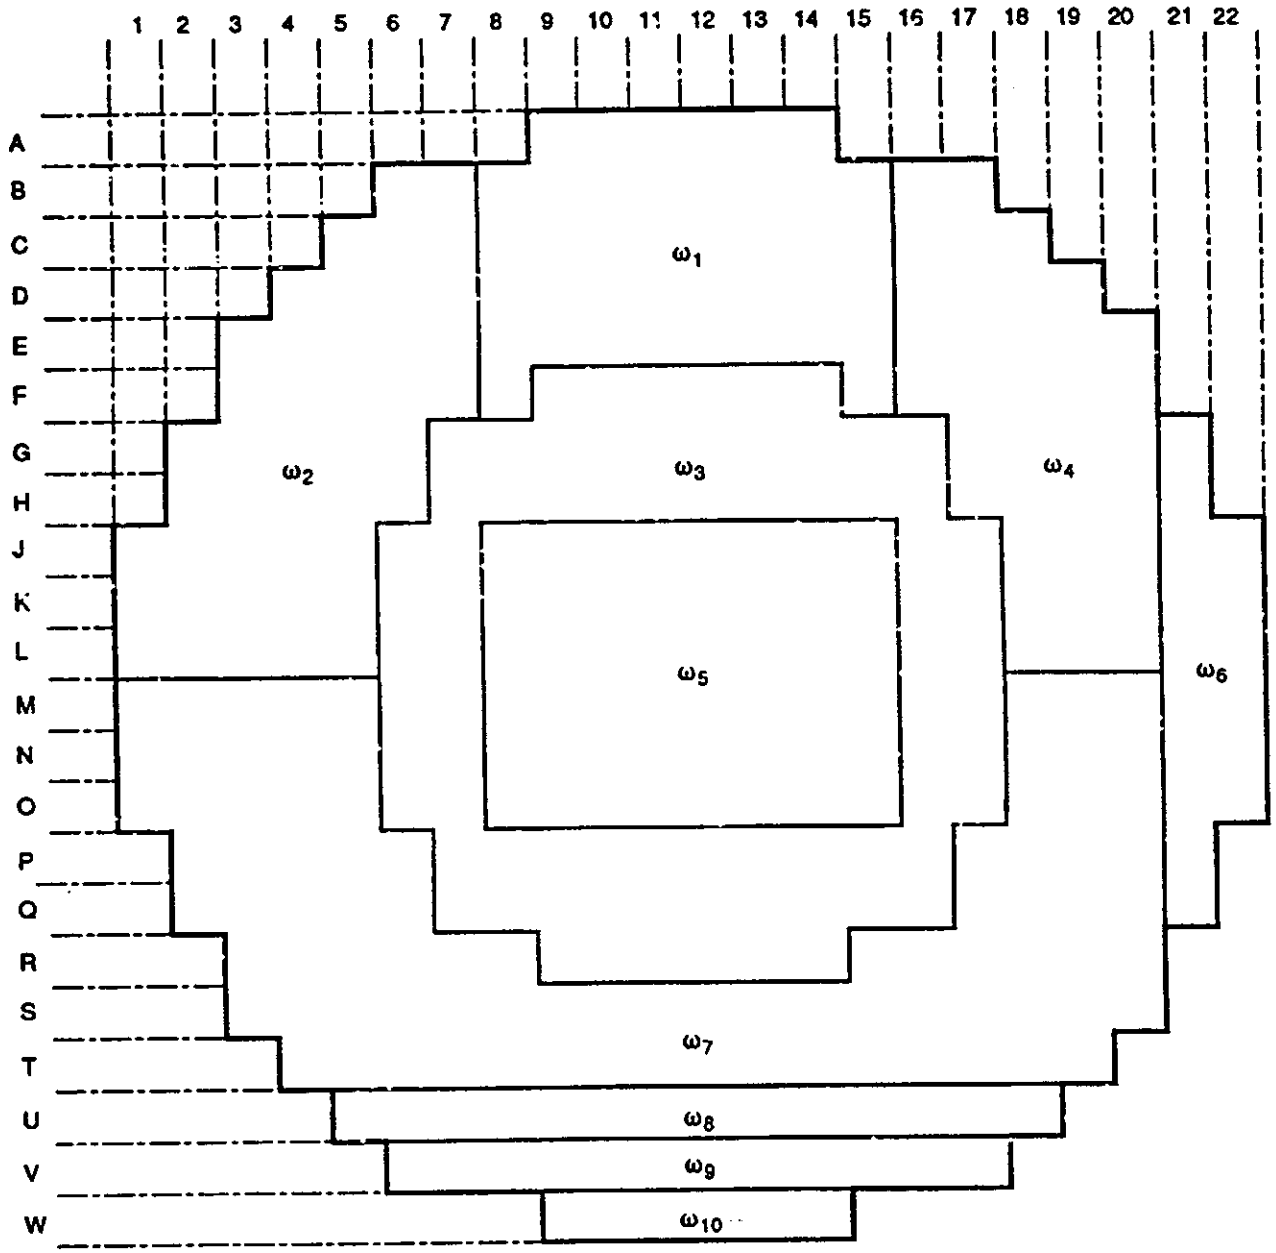

Multi-Region

Figure 4.~~1~~4: Irradiation Zones for a Typical CANDU 6 Time-Average Model

4.8

Fig 4.9

Time–Average Calculation

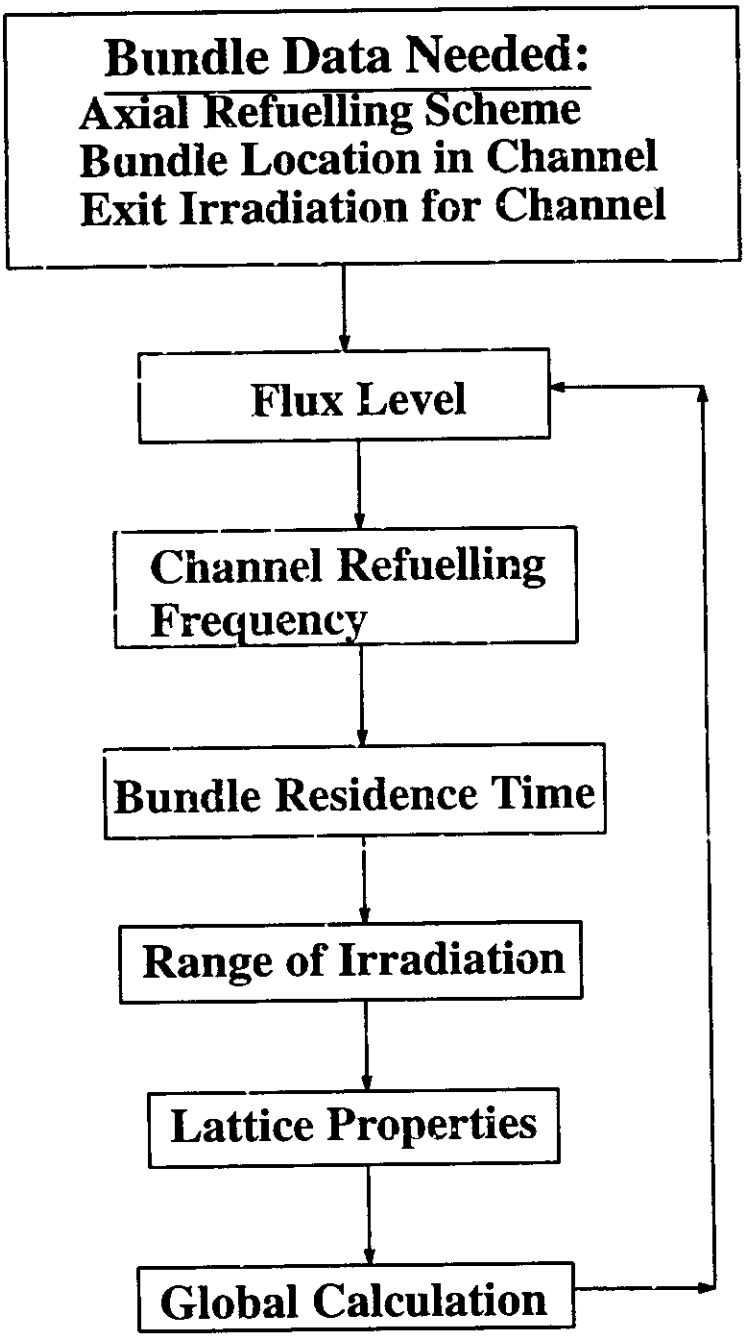

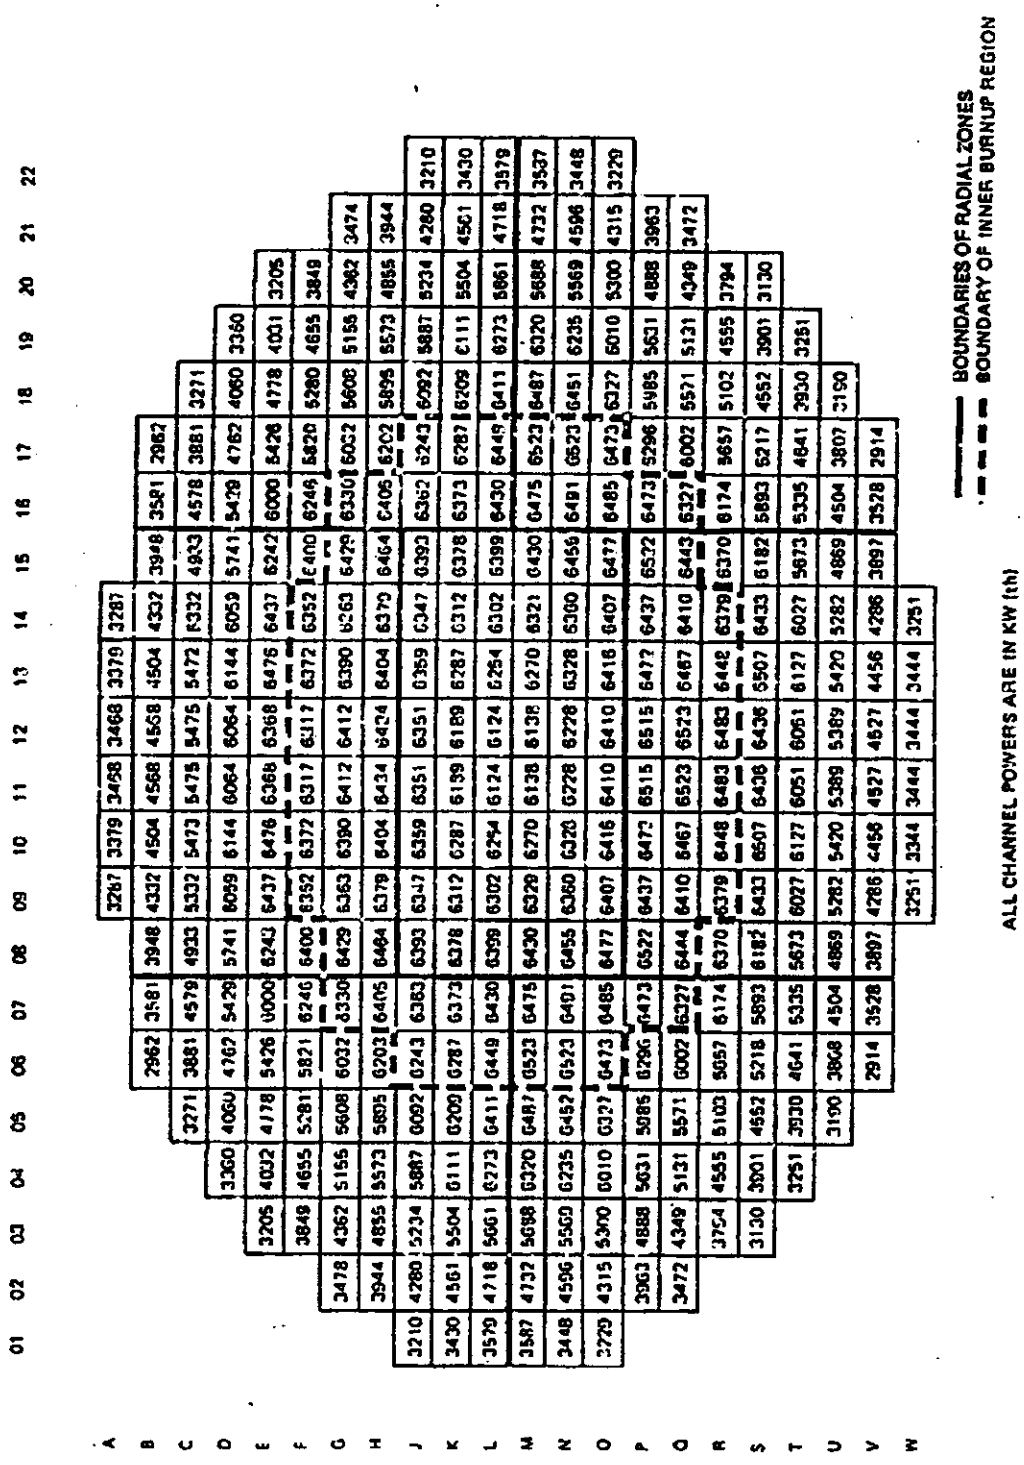

Figure 8.11 Time Average Channel Powers

4.10

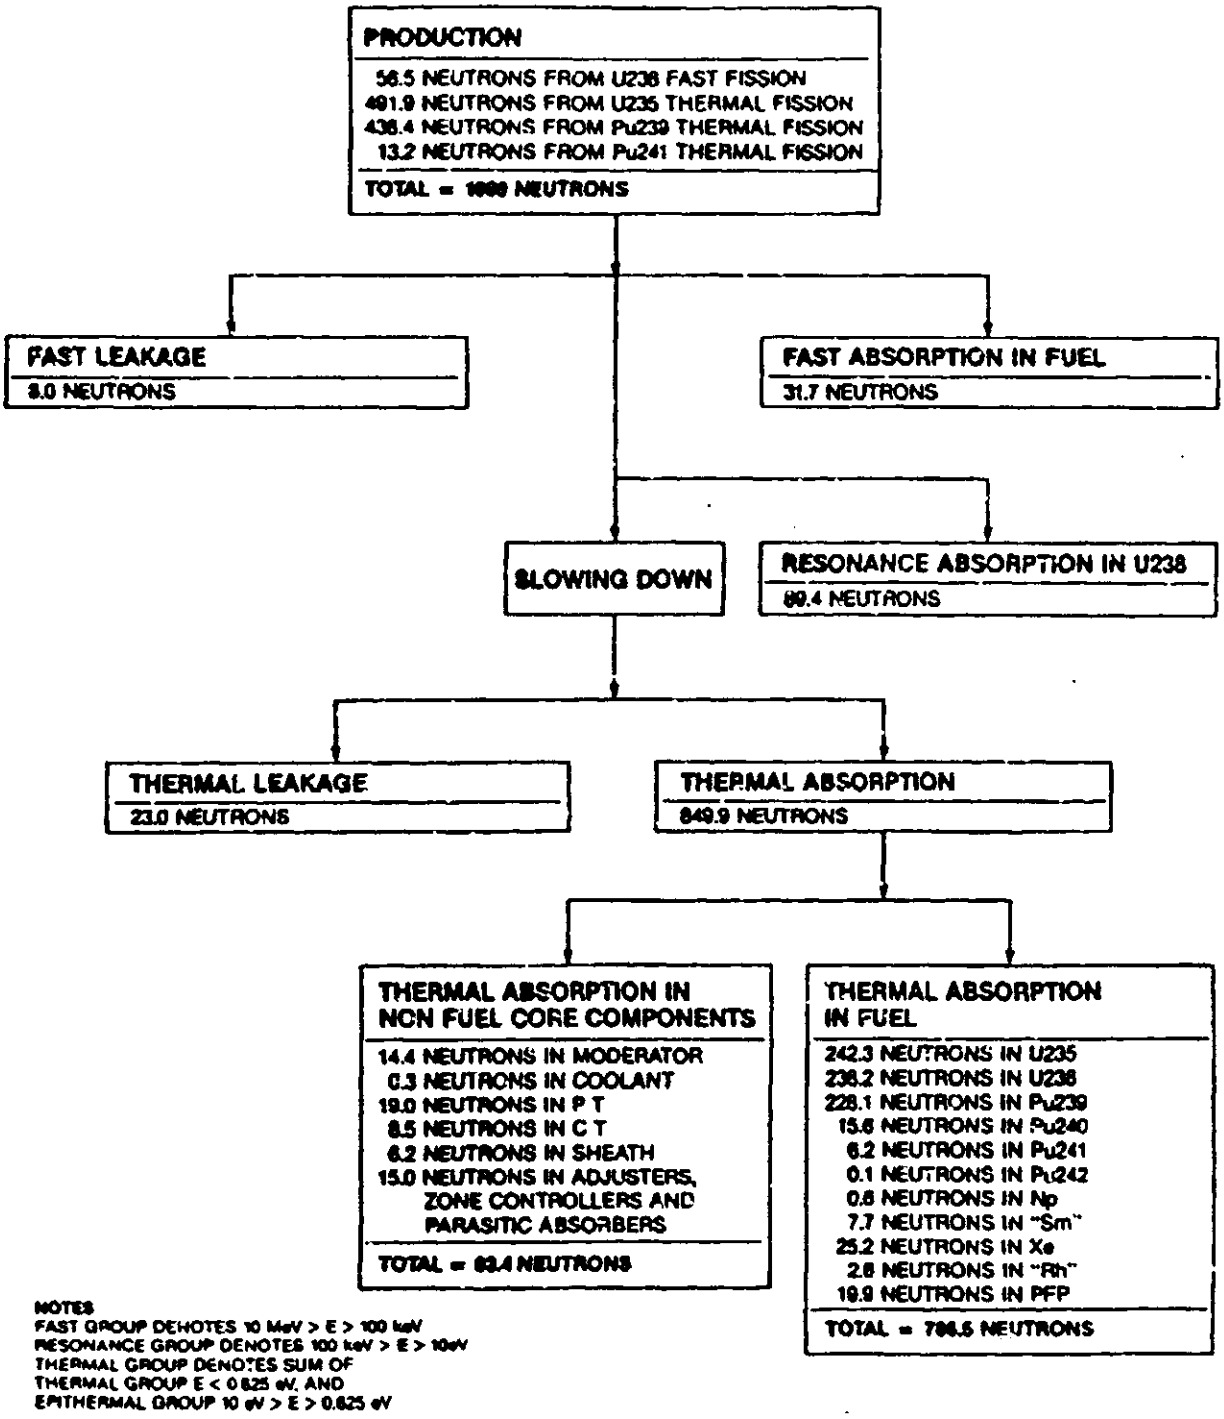

Figure ~~5.6~~ 7.12 Neutron Balance in the CANDU 6 Equilibrium Core

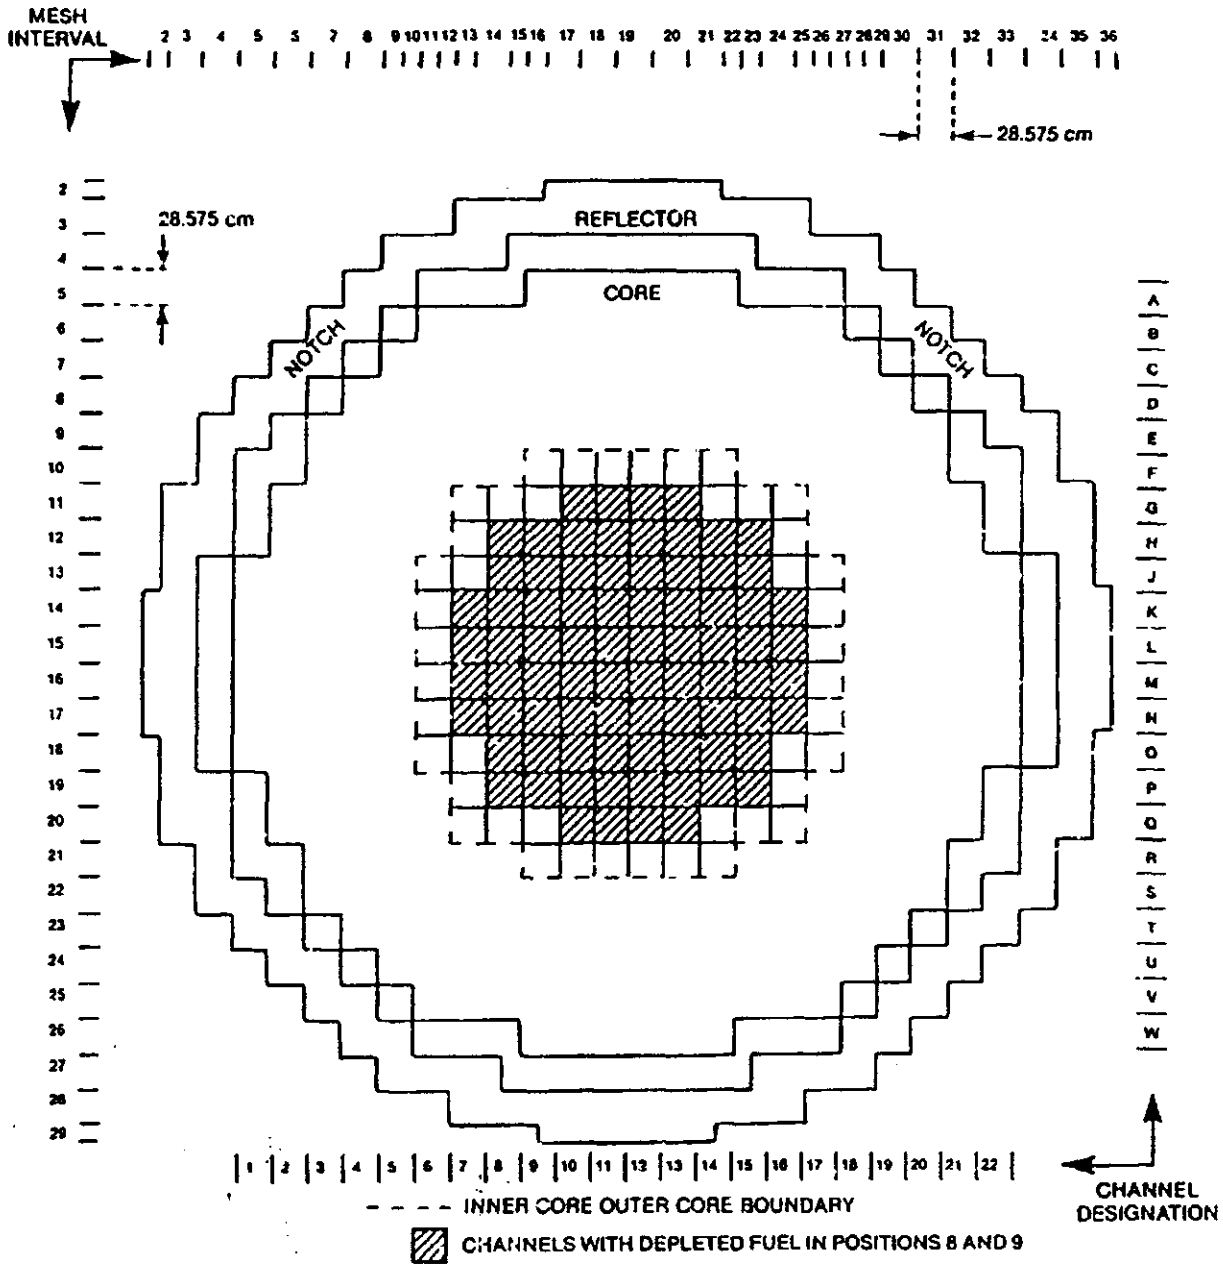

Figure 8.3 600 MW Reactor Model Face View Showing Initial Loading of Depleted Fuel

5.1

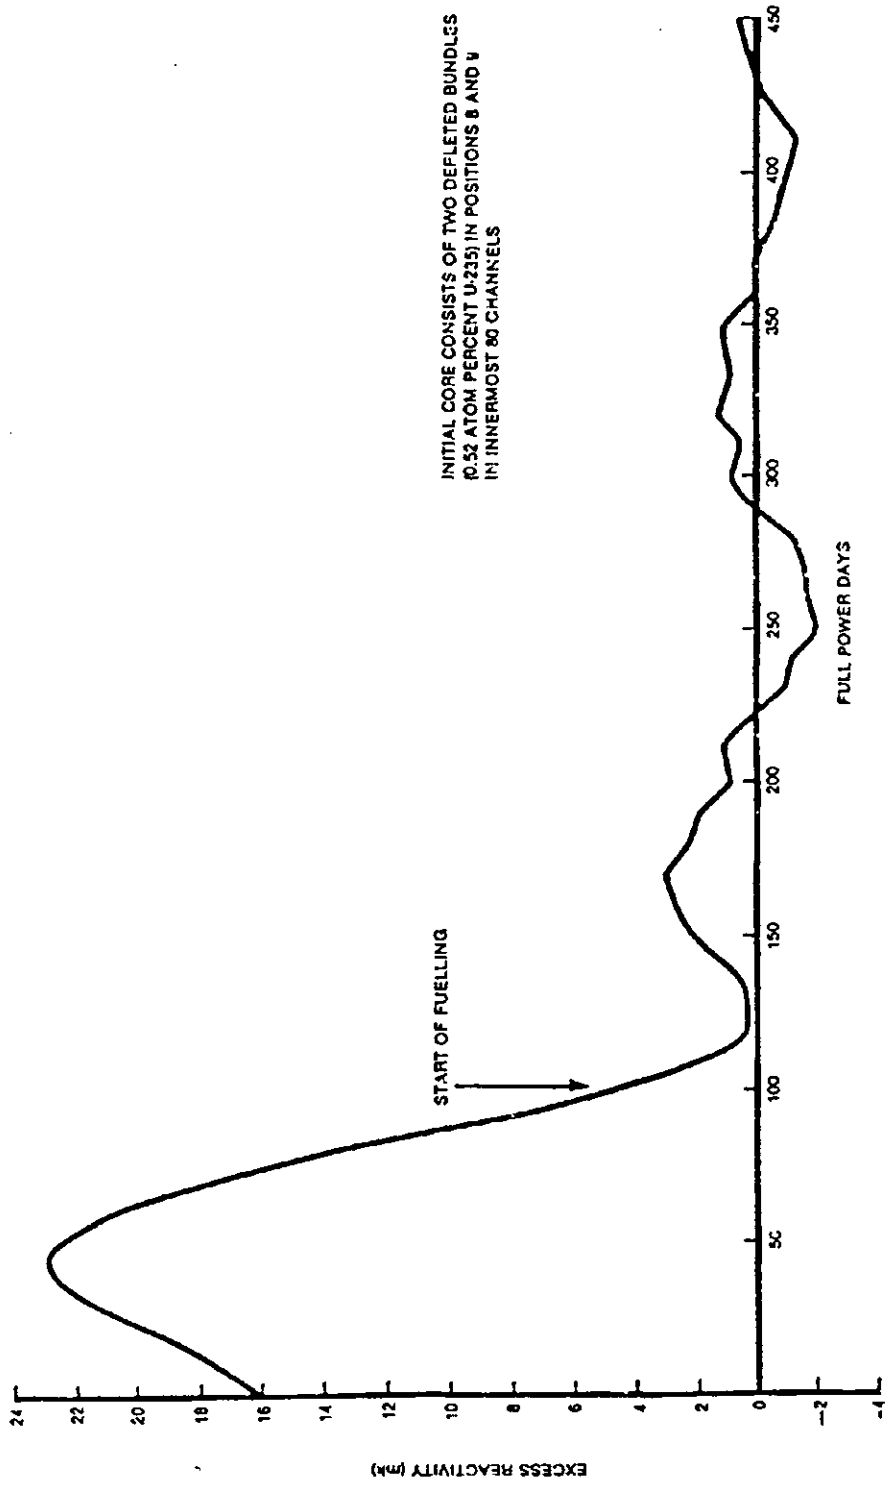

Figure ~~3.4~~ 5.2 Variation of Excess Reactivity During Initial Burnup Period (FMDP Simulations)

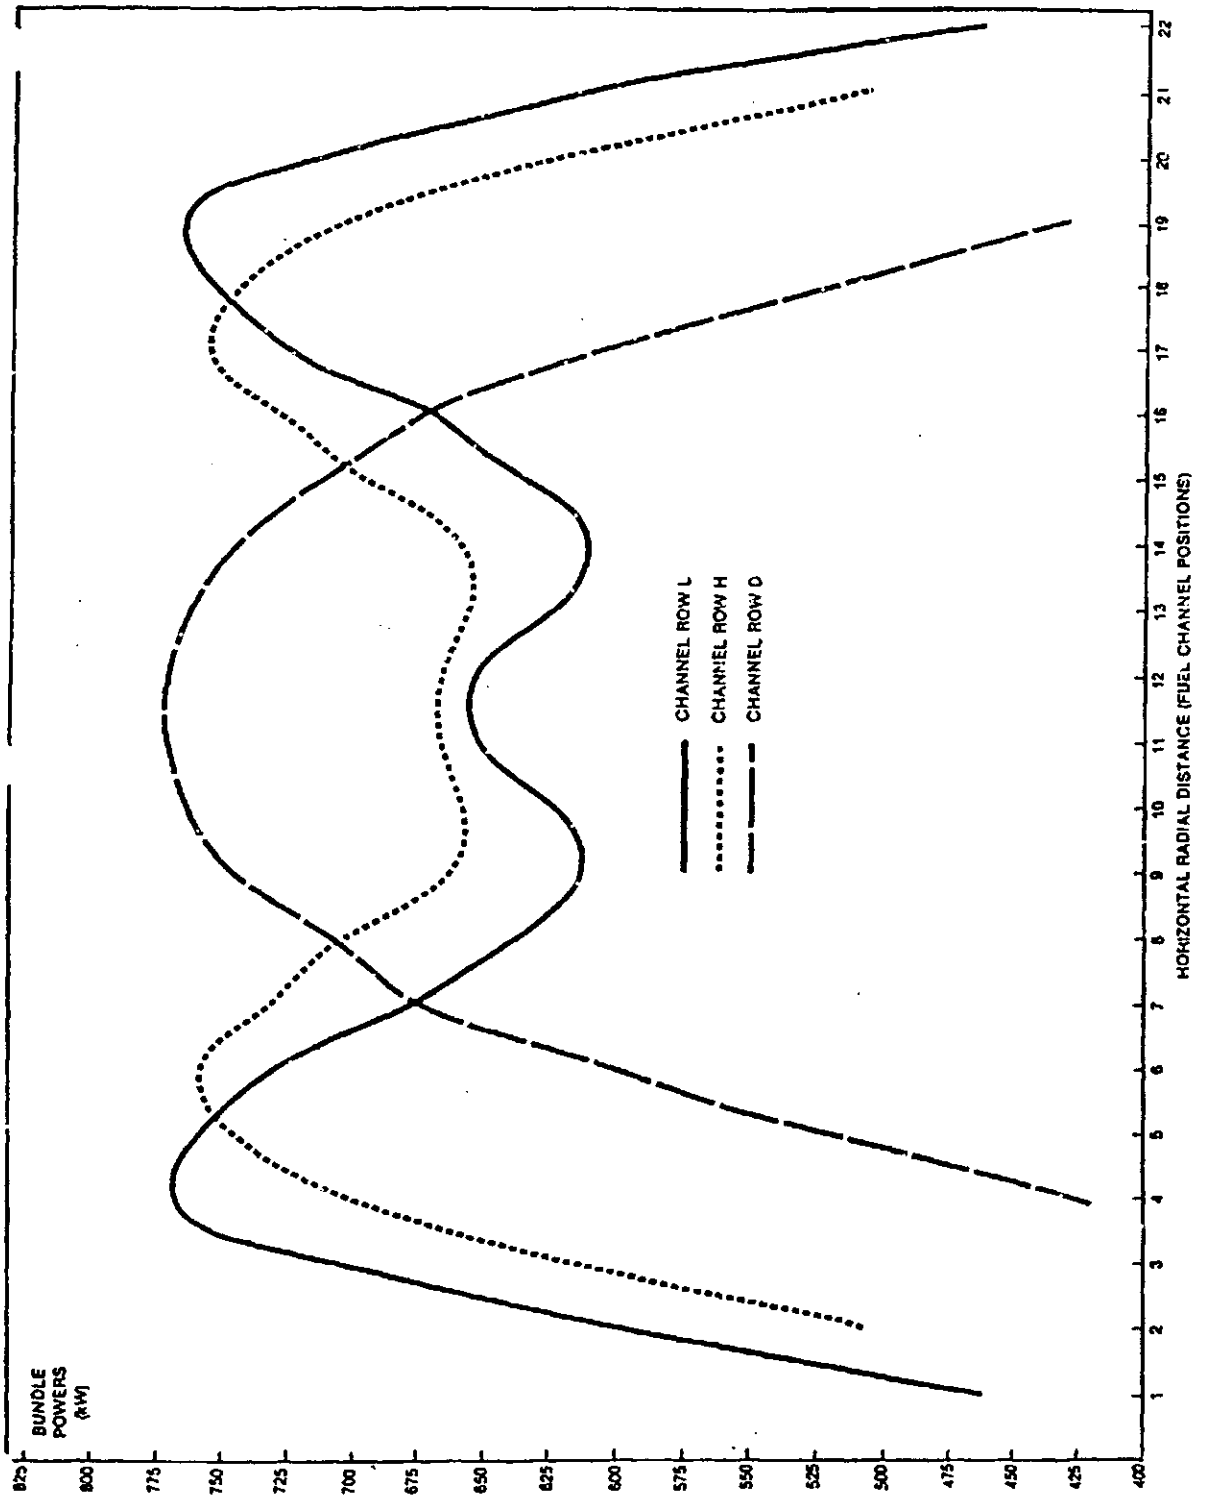

Figure 8.5 Horizontal Radial Bundle Power Distribution at 0 FPD

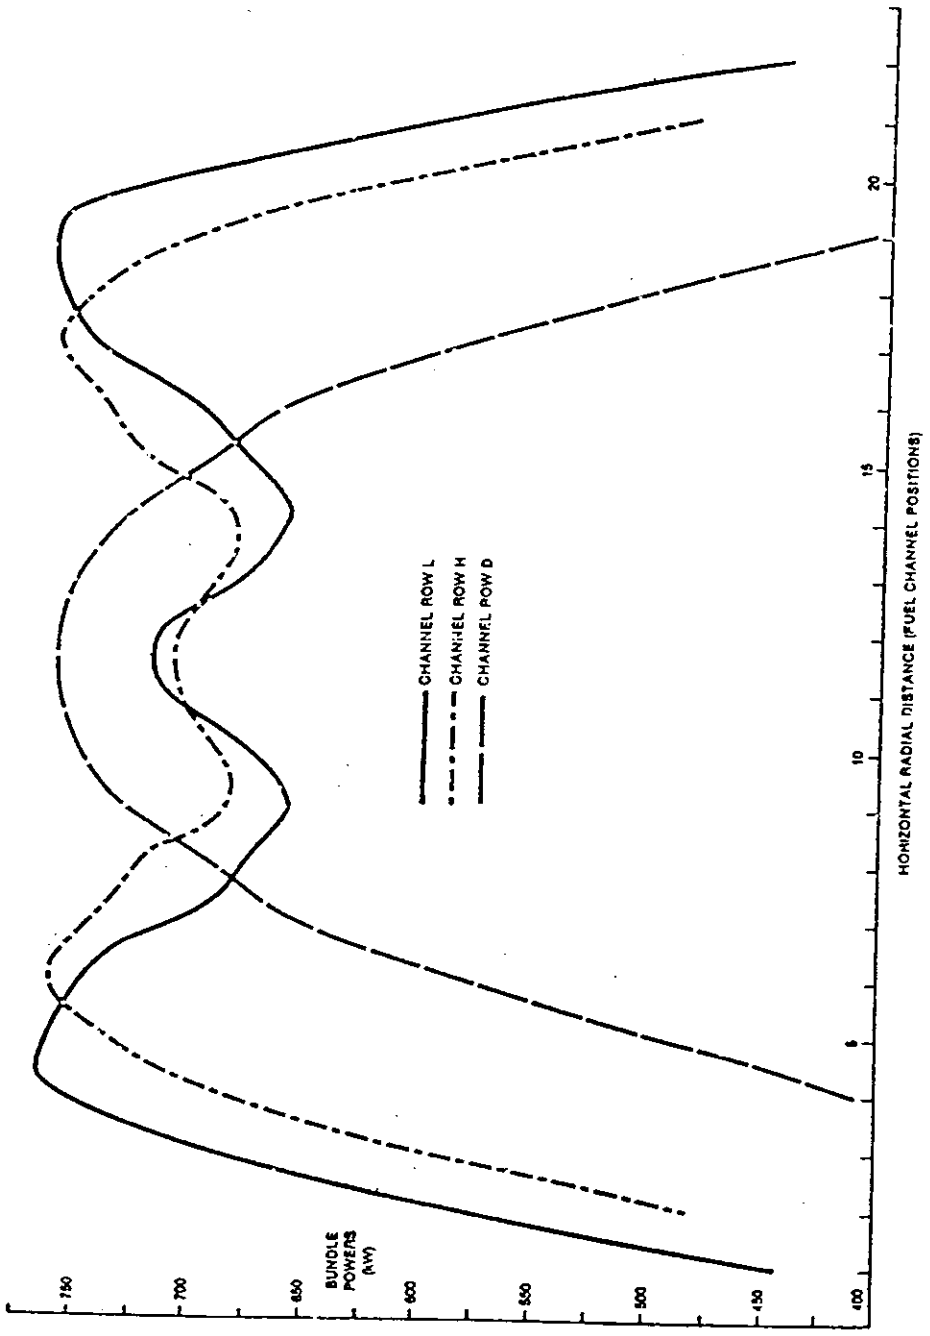

5.4

Figure 8.6 Horizontal Radial Bundle Power Distribution at 40 FPD

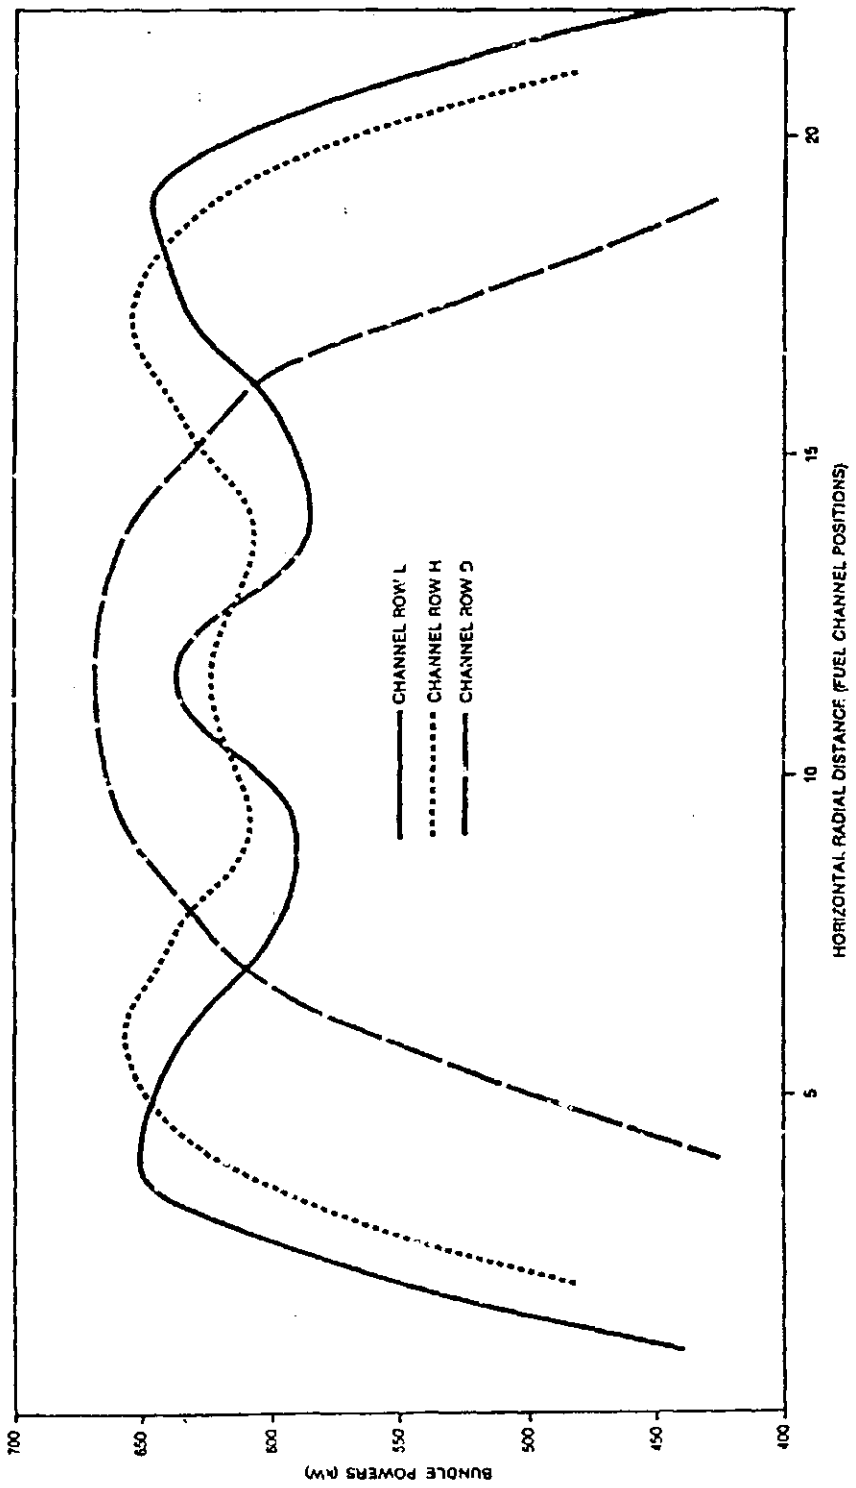

Figure 8.7 Horizontal Radial Bundle Power Distribution at 100 FPD

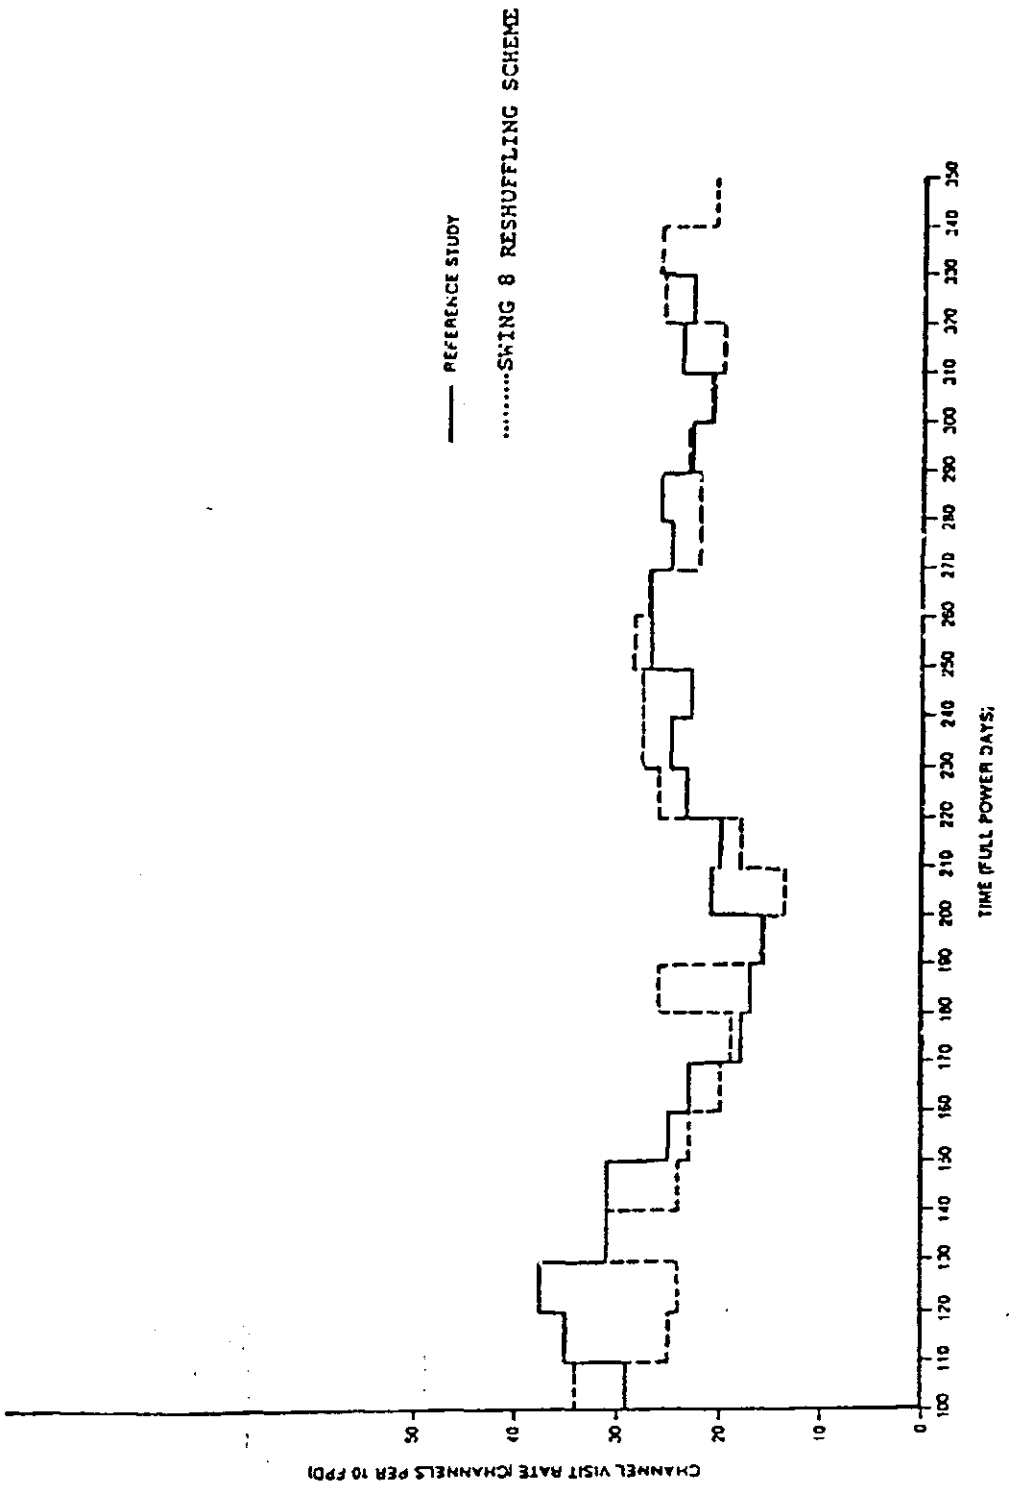

Figure 8.19 Fuelling Machine Visit Rate per 10 FPD Interva

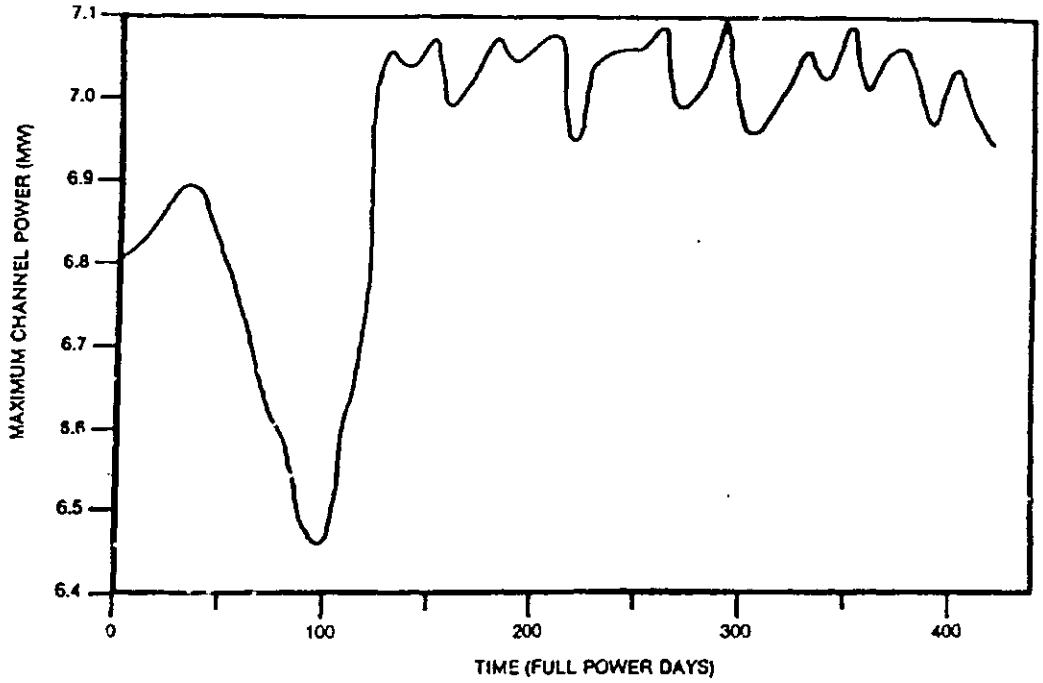

Figure 8.9 Maximum Channel Power

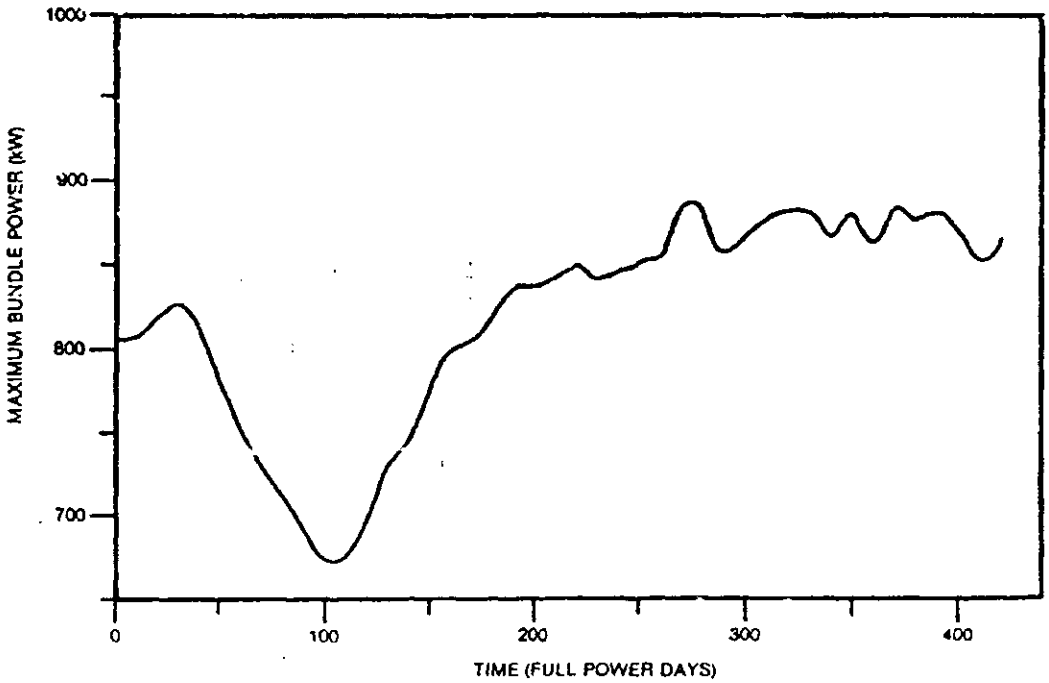

Figure 8.10 Maximum Bundle Power

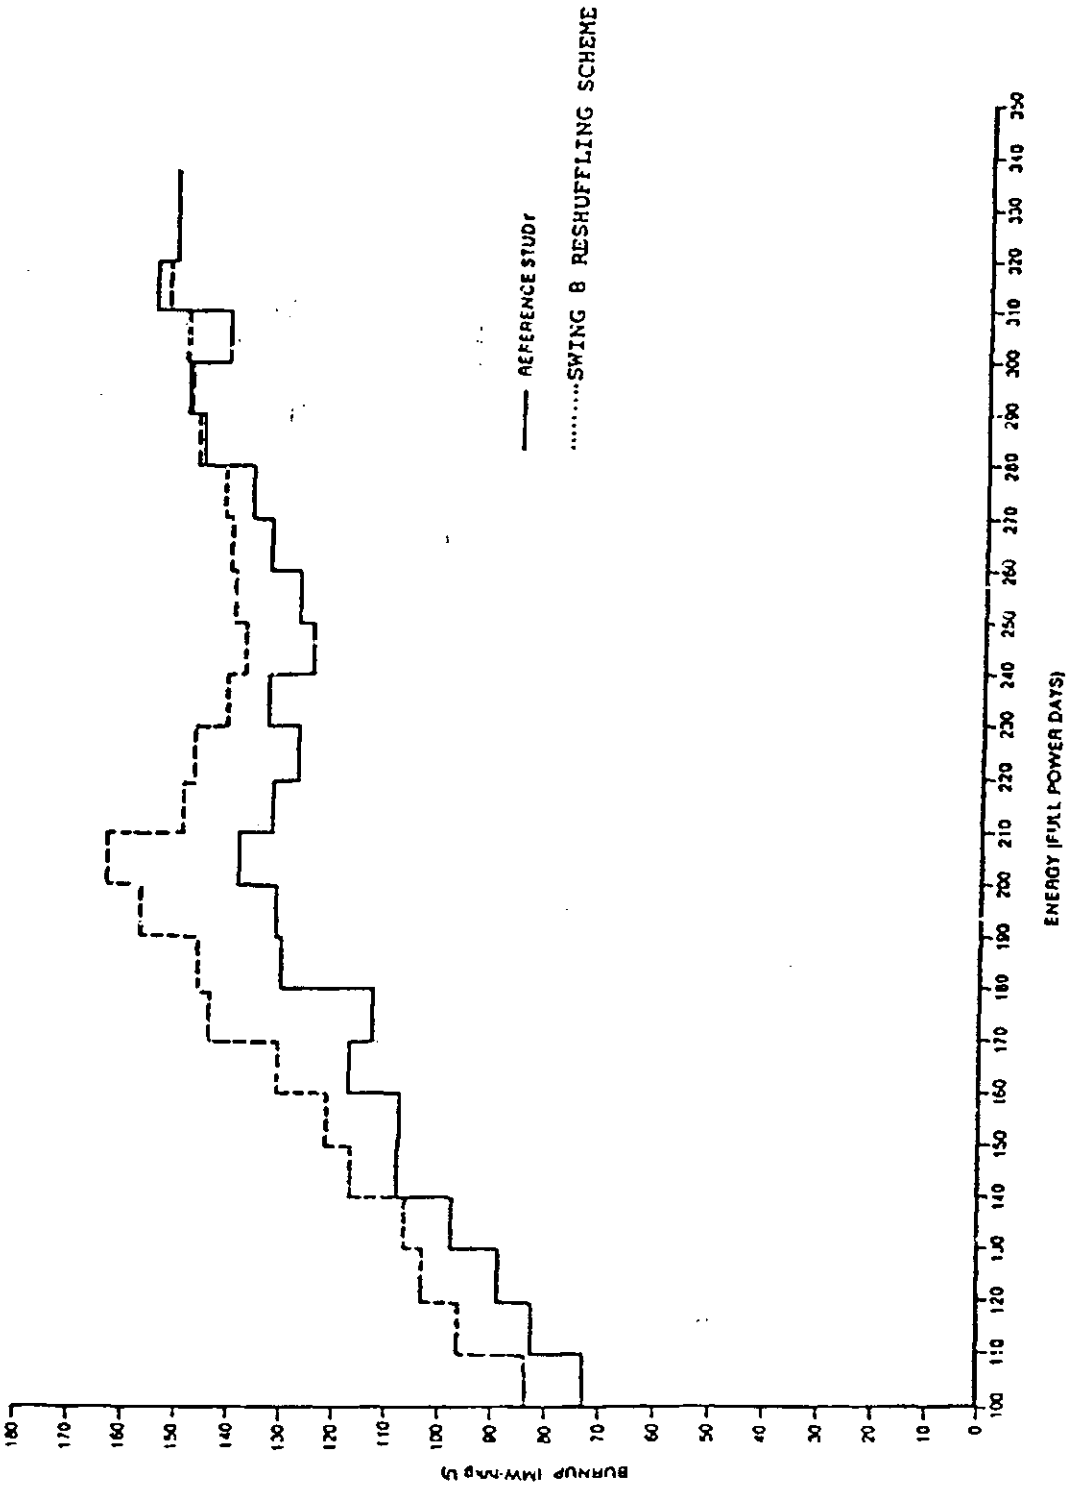

Figure 8.17 Average Discharge Burnup per 10 FPD Interval

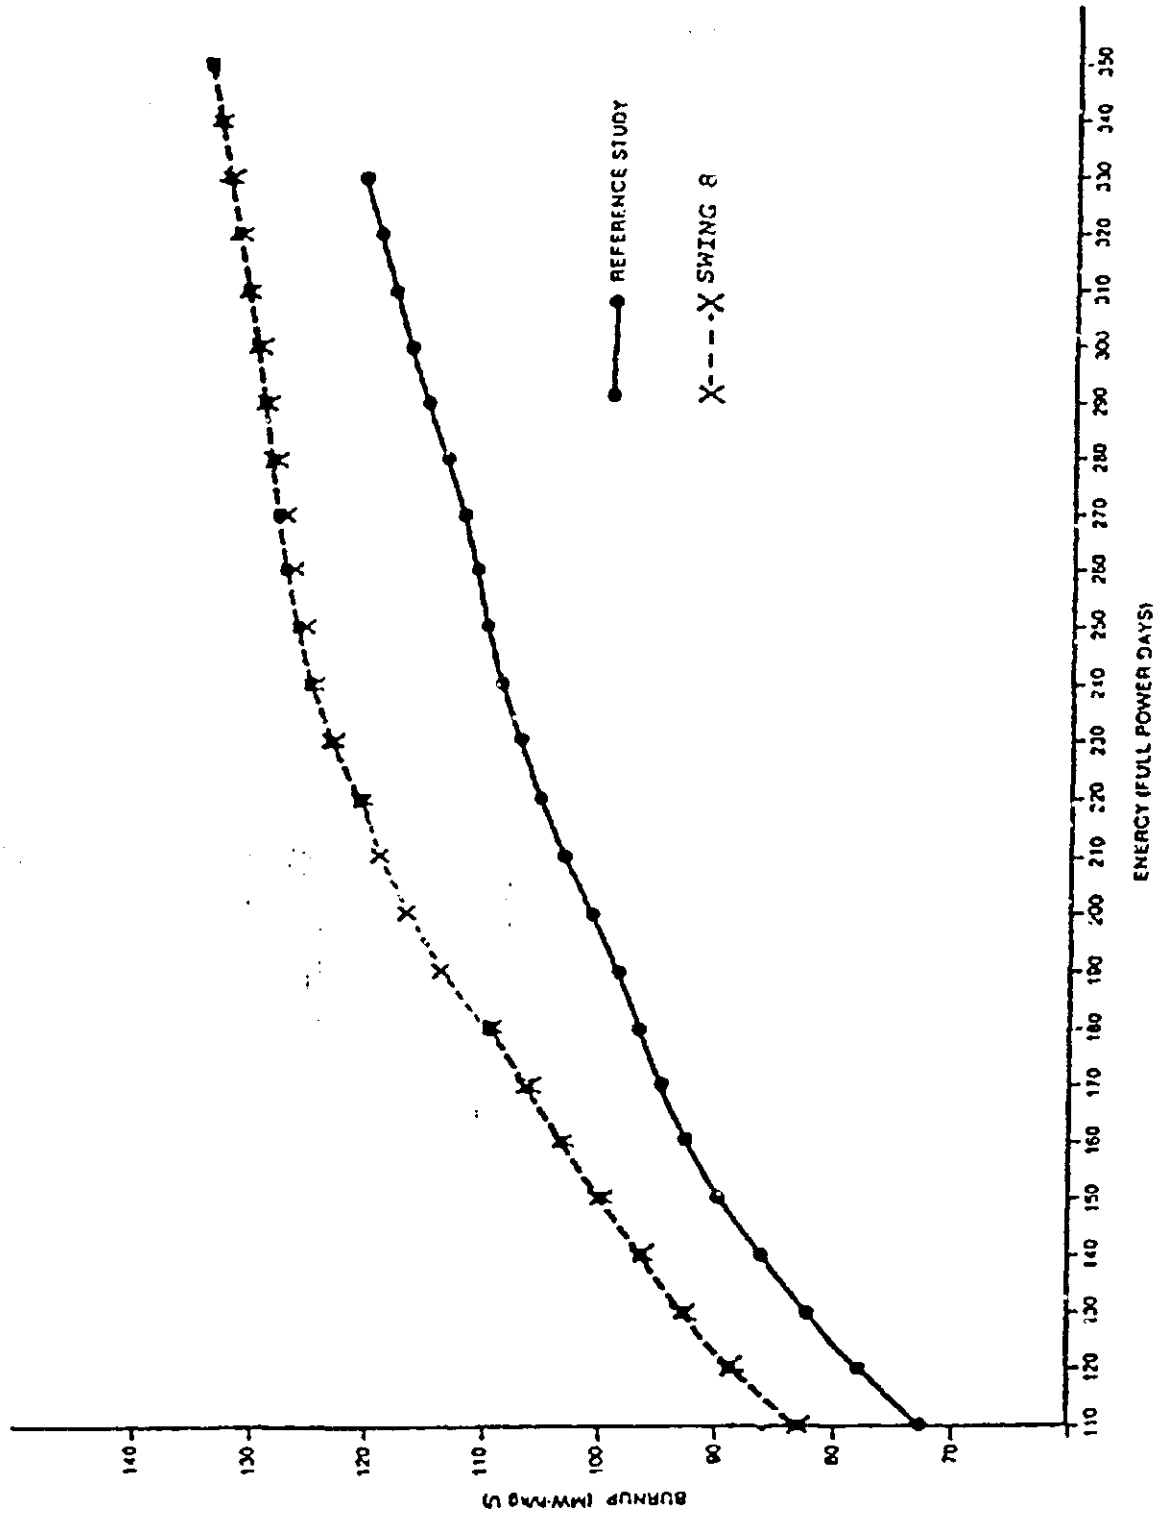

Figure 8.18 Cumulative Discharge Burnup

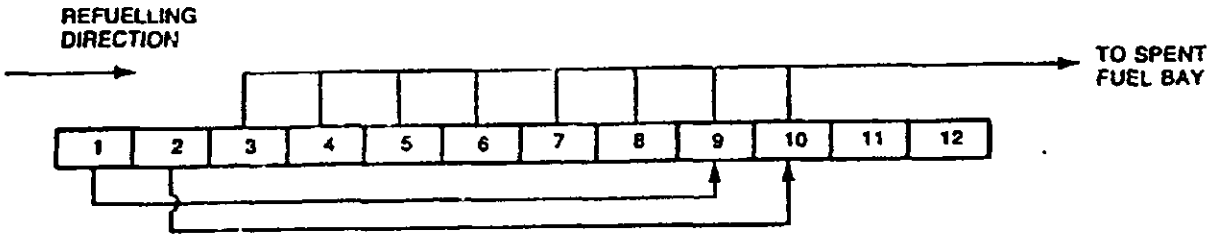

Figure ~~8.16~~ Swing-Eight Refuelling Scheme

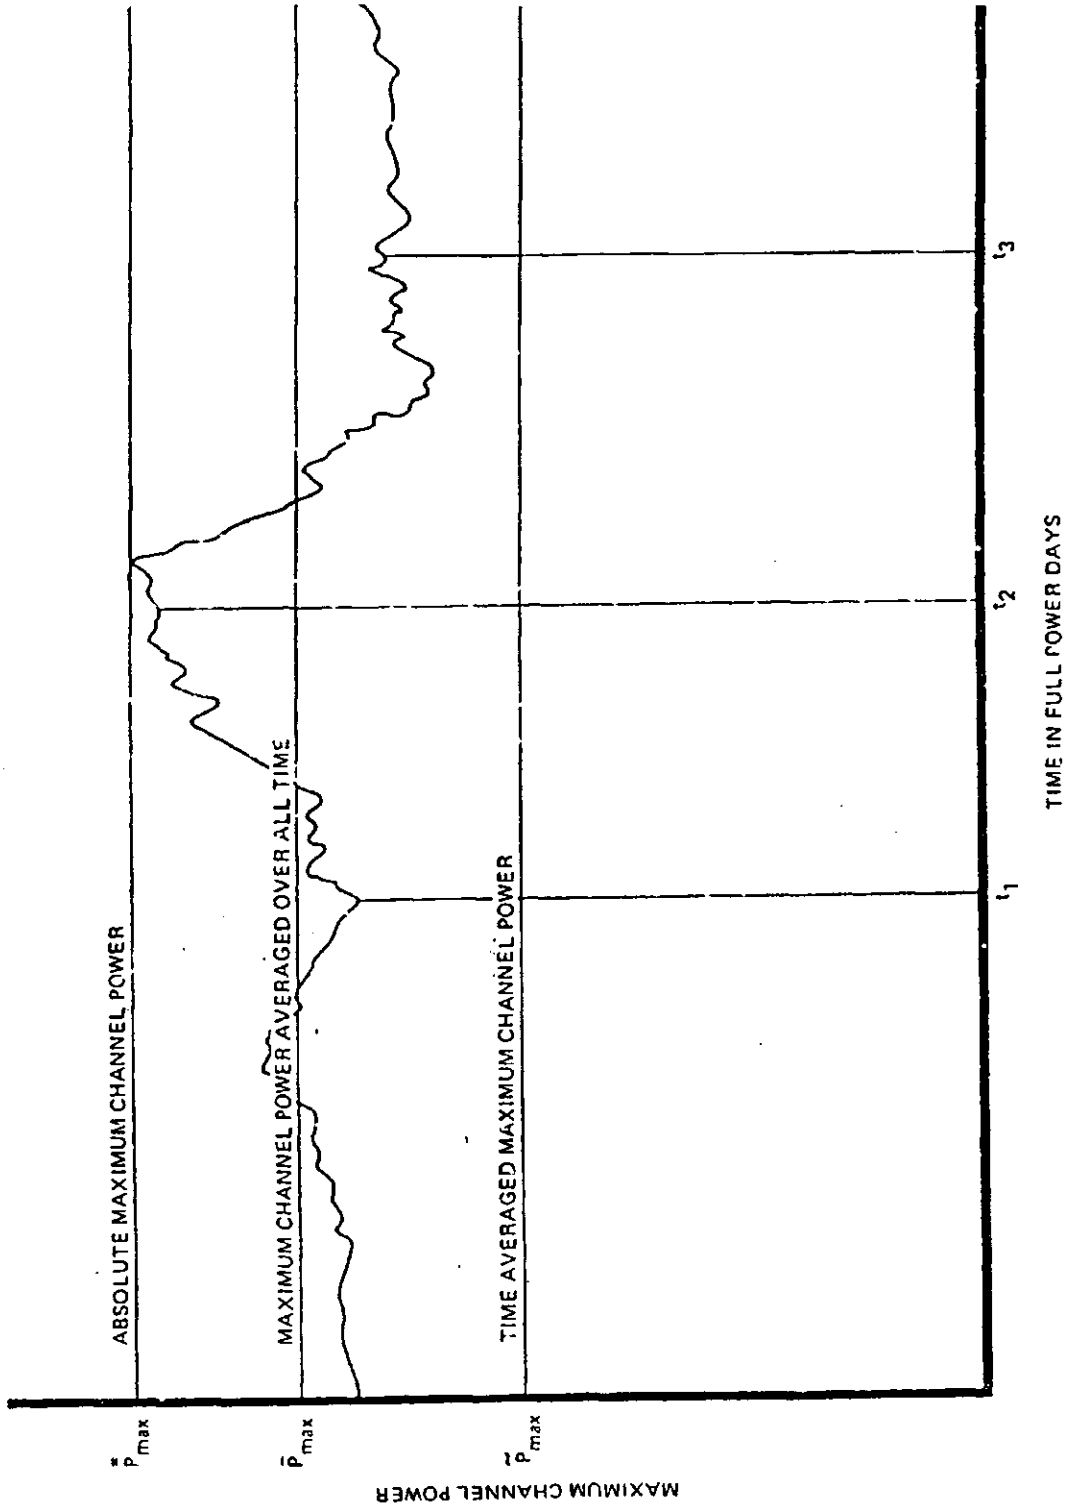

Figure 8.15 Segment of Typical Operation History

5.12

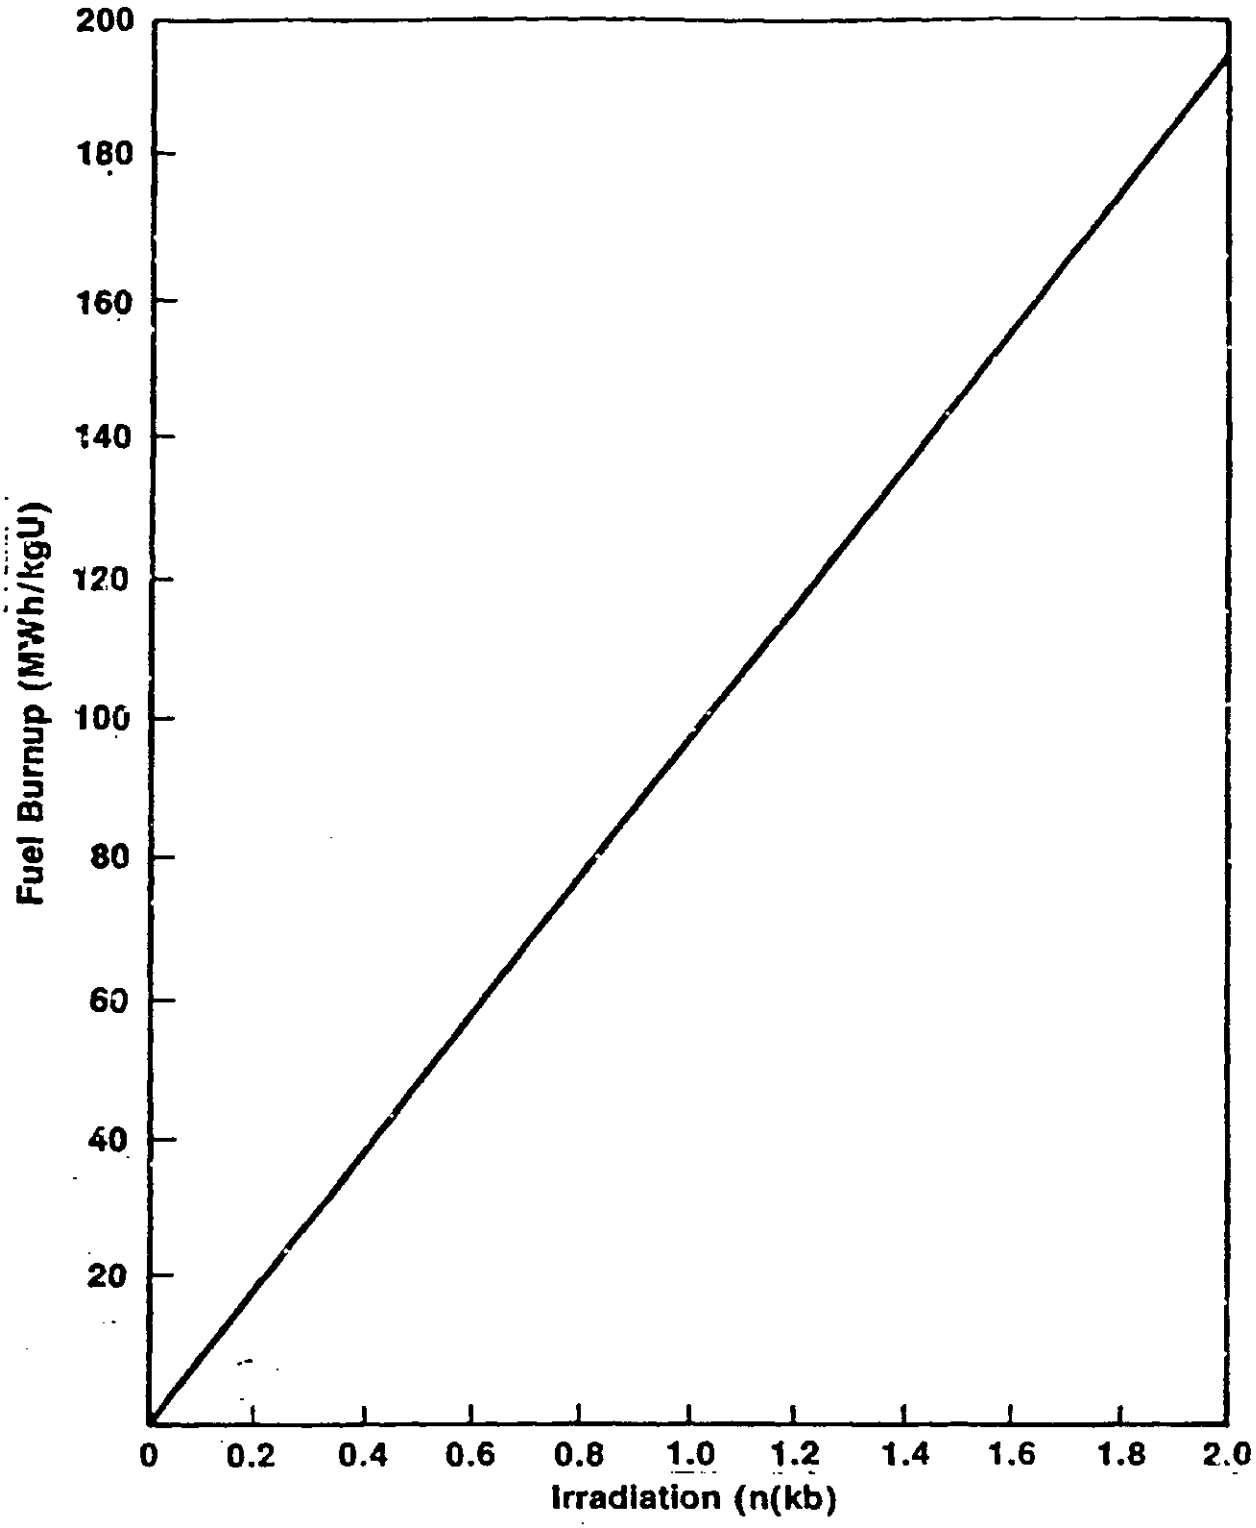

Figure 5.5 Fuel Burnup as a Function of Irradiation

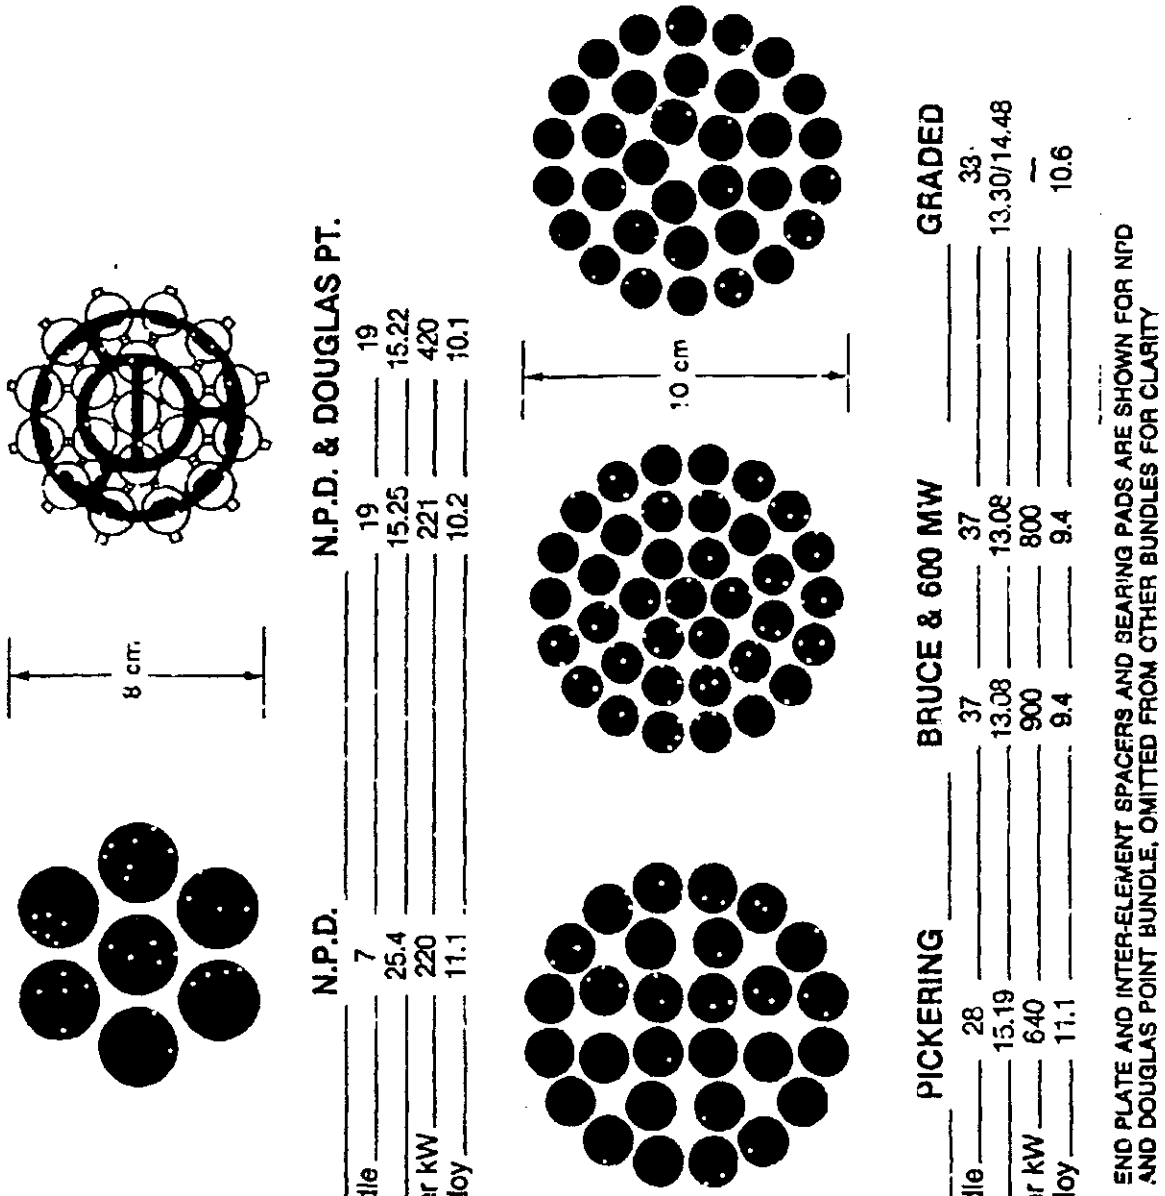

Figure 5.4 (6.1) Basic Data for CANDU Fuel Bundles

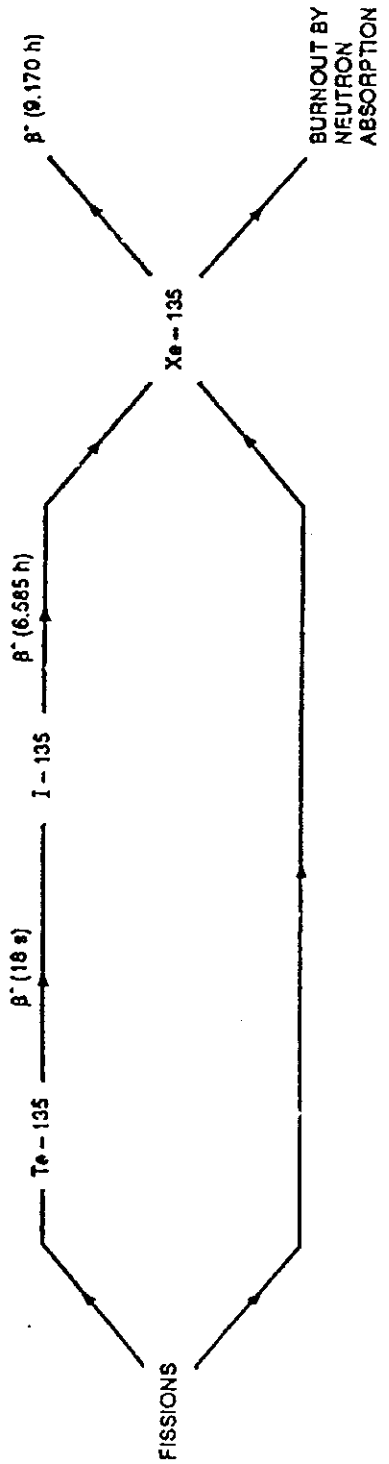

Figure ~~4-14~~ Processes for Xe-135 Creation and Removal

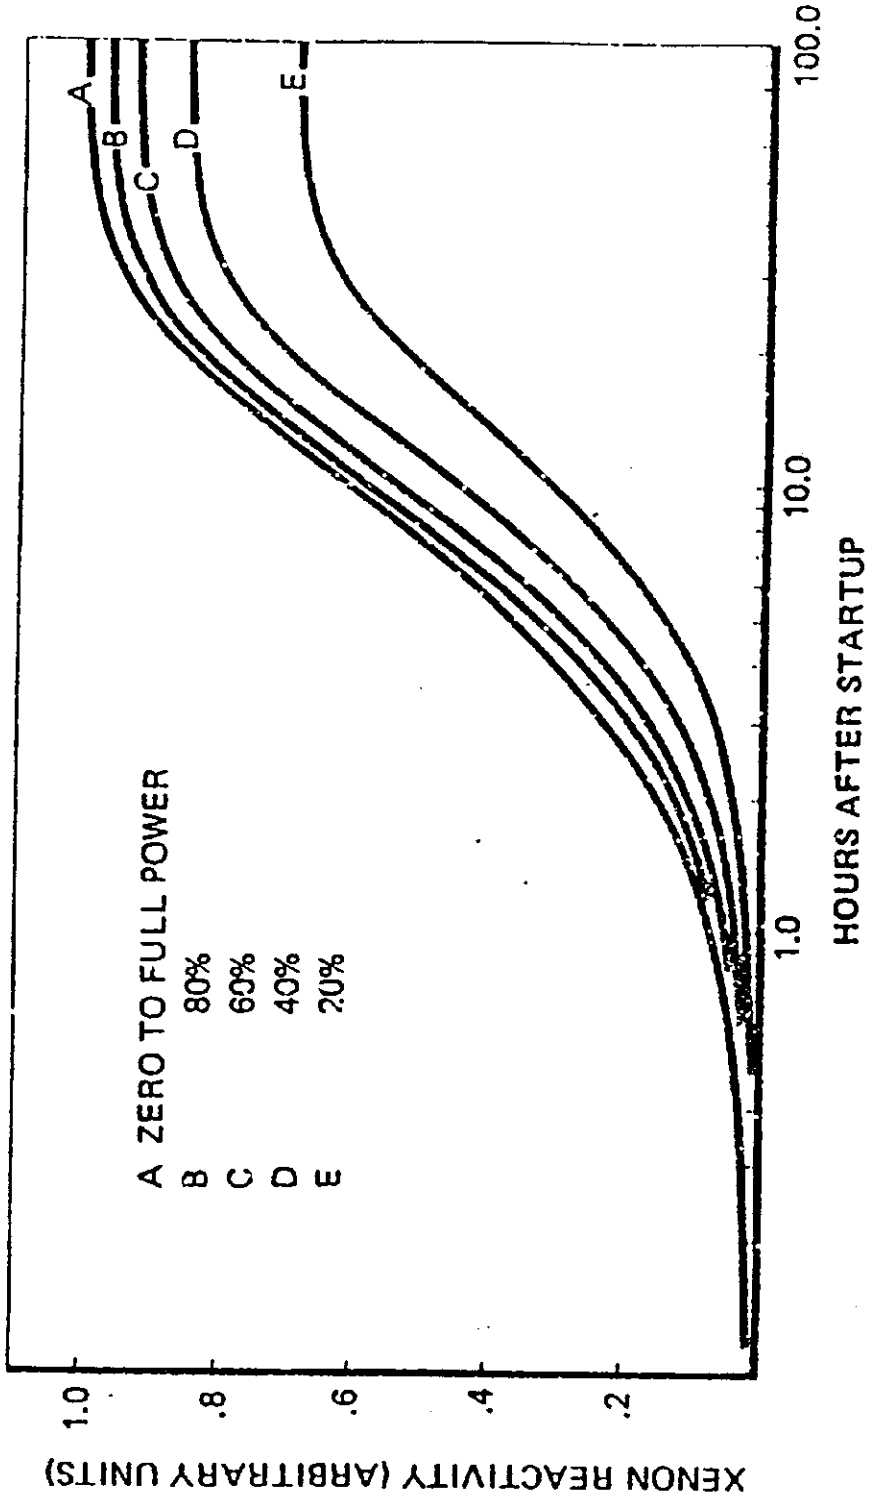

Figure ~~4.15~~ 9.2 Xenon Buildup (Reactor Raised from Zero to Various Power Levels)

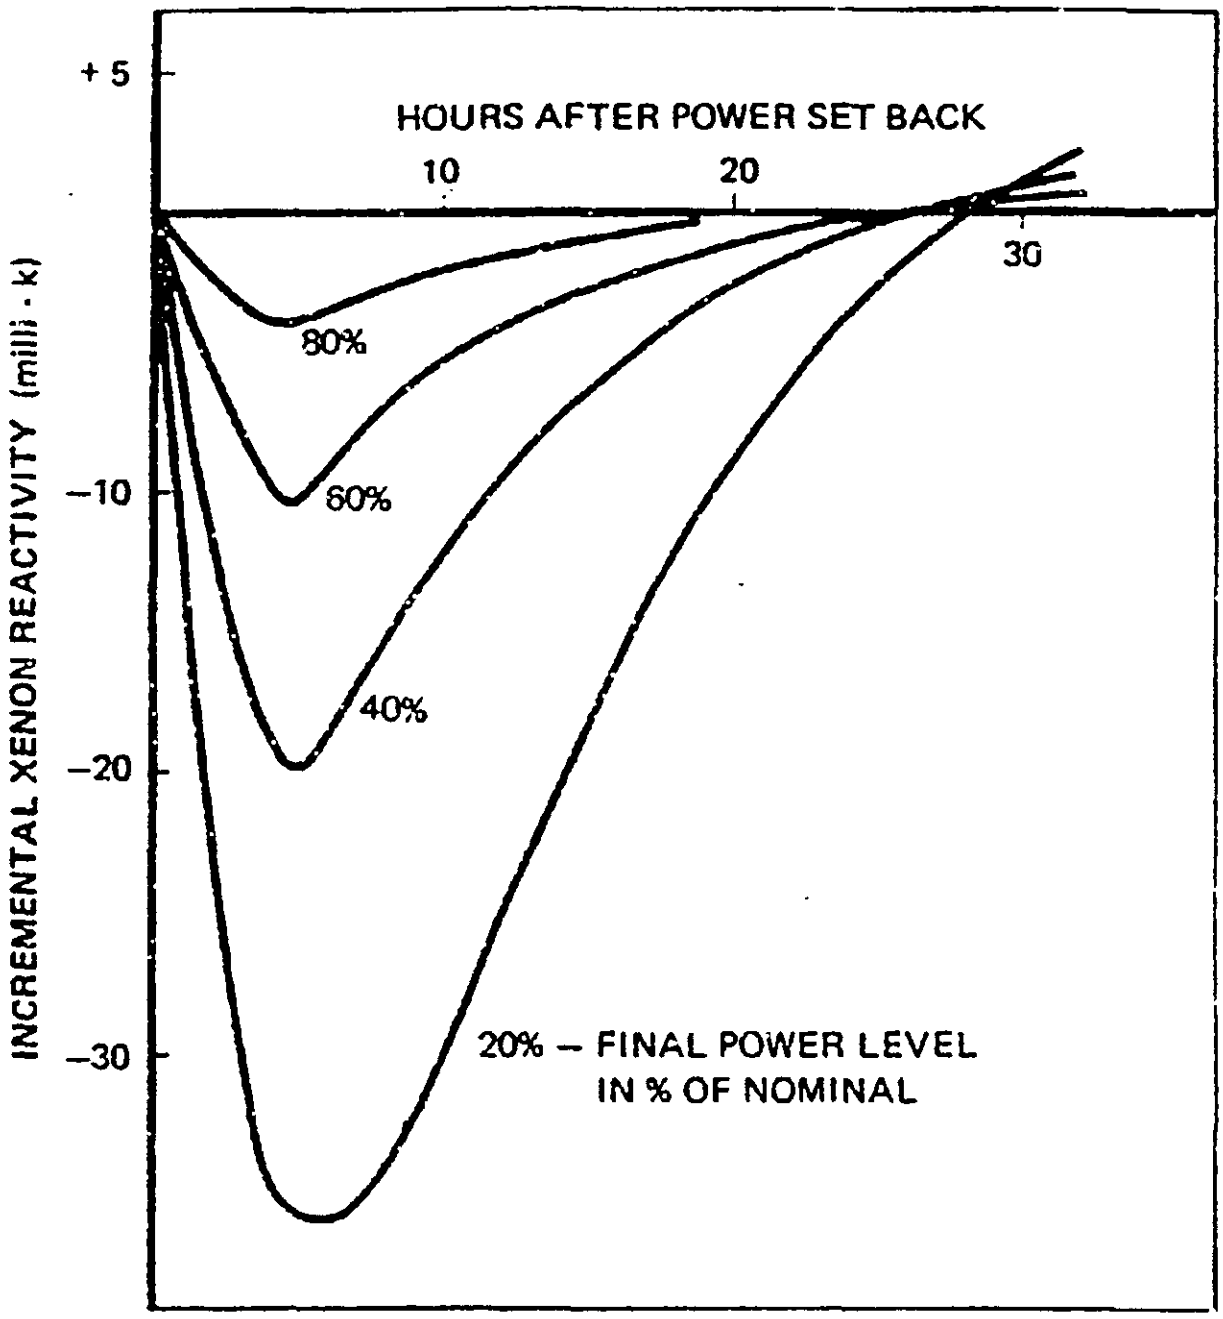

Figure 4.16 Xenon Reactivity Transients Following Power Setbacks from 100%

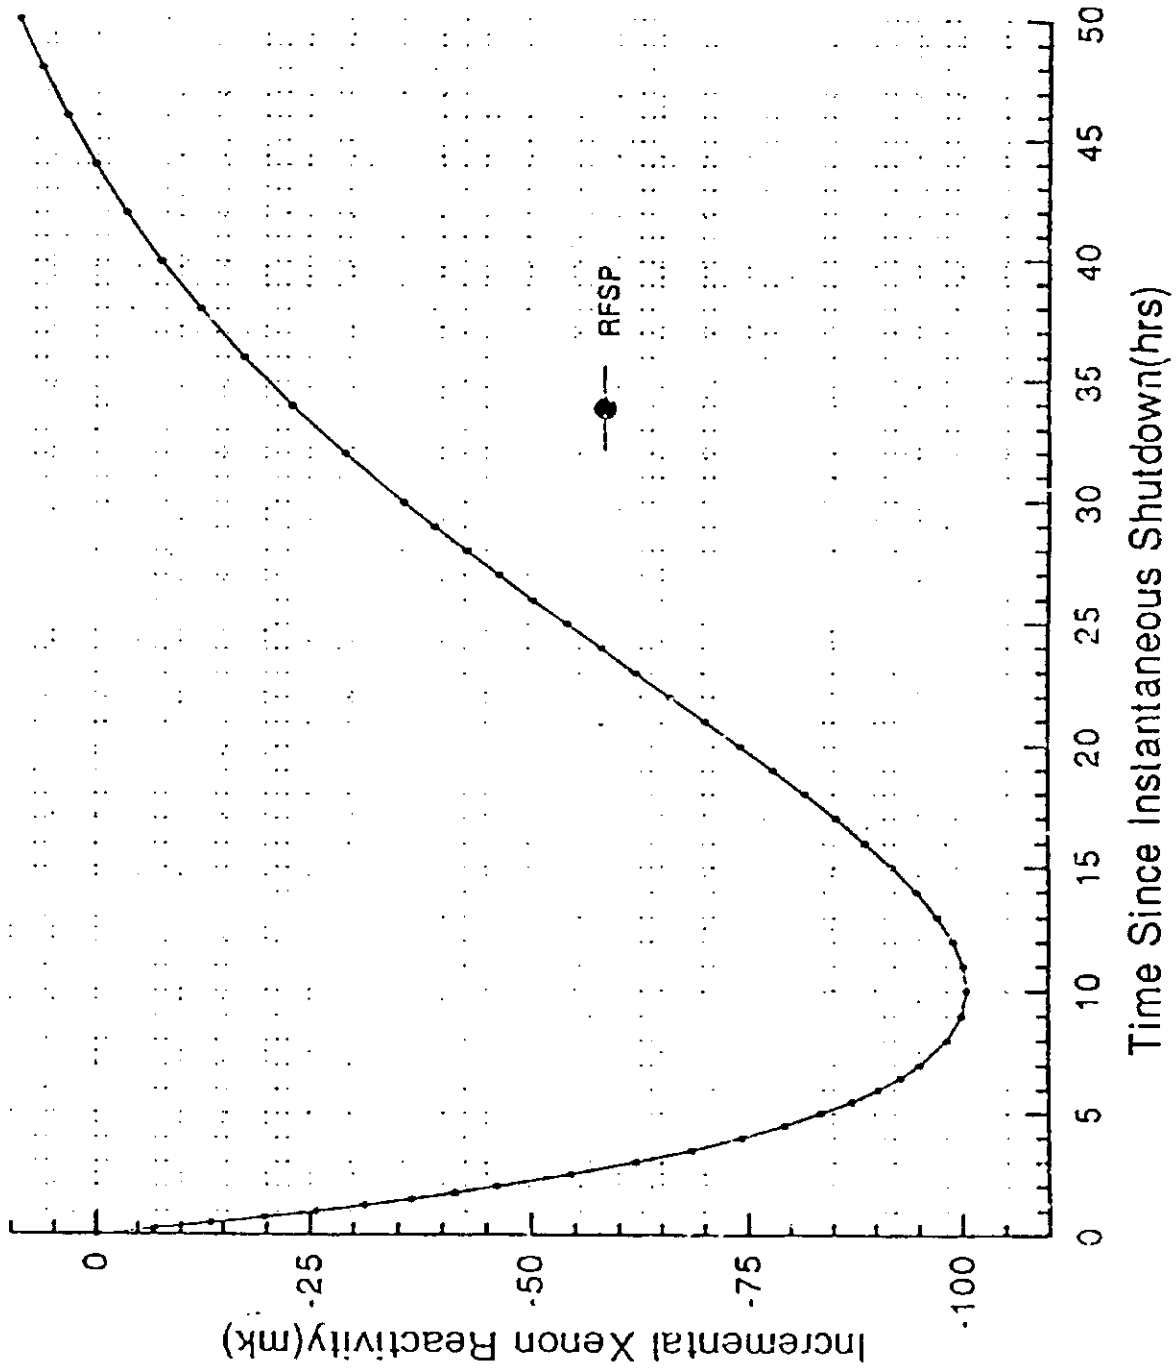

Figure 4.17 Xenon Transient Calculation (Time-Averaged): ~~Setup~~ Shutdown to 0 Power