Introduction to Reactor Physics

B. Rouben
Manager,
Reactor Core Physics Branch
Atomic Energy of Canada Ltd.

2002 September

Table of Contents

Figures (Included in PowerPoint Presentation)

1. Introduction

This is a **descriptive** course on the physics of CANDU reactors. The objective is to familiarize you with the basic **concepts and ideas, definitions and quantities** which are important in the understanding of reactor physics. The discussion will be at a basic level, as this is an introduction to the subject. There will be very little emphasis placed on complex equations or difficult mathematics. The desired outcome of this course is that you will be able in the future to understand and follow reactor-physics discussions in meetings, reports or presentations.

Nuclear energy is energy from the nuclei of atoms. It is quite natural - it occurs and has occurred in nature in various forms since the beginning of the universe.

Let us start with the law of conservation of energy.

1.1 Law of Conservation of Energy

The total amount of energy in the universe is constant. You cannot create or destroy energy. You can only change it from one form to another, for instance

- potential energy (e.g., of an apple in a tree) to kinetic energy (of the apple when it falls),
- kinetic energy to heat (when the apple hits the ground, its temperature increases slightly).

This seems to run counter to experience! If you burn something, you get energy (heat) without appearing to lose anything.

1.2 Equivalence of Mass and Energy

One of Einstein's greatest discoveries is that the law of conservation of energy must be generalized to include mass as a form of energy!

$$\mathbf{E = mc^2} \qquad \textbf{(1)}$$

1.3 Energy from Mass: Chemical Energy and Nuclear Energy

The energy obtained from burning something, for example coal,

$$\mathbf{C + O_2 \rightarrow CO_2} \qquad (2)$$

comes from a change in mass: the mass of the carbon dioxide molecule is smaller than the sum of the masses of the carbon and oxygen molecules. But the difference is extremely small.

Chemical energy comes from changes in **atoms** and **molecules** (e.g. Eq. 2) - actually their electron clouds. Chemical energy is the true atomic energy!

One unit of energy often used in atomic or molecular physics is the electron-volt (eV), the energy gained by an electron falling through a difference in voltage of 1 Volt. Because the electron is a very small particle, its electric charge is a very small fraction, 10^{-19}, of the macroscopic unit Coulomb. Consequentially, the electron-volt is a small fraction of the macroscopic unit of energy, the Joule (J):

$1\ eV = 1.6 * 10^{-19}\ J.$

However, the eV is a very convenient unit in atomic physics, and as a unit for measuring energies released in chemical reactions.

Nuclear energy comes from changes in the **nuclei** of atoms. Nuclei are made up of nucleons (protons and neutrons), and they are much tinier than atoms: their radii are about **100,000 times smaller** than those of atoms! The amount of energy needed to keep nucleons bound together in such a small space (i.e., the nucleon binging energy) is much larger than the binding energy of electrons within atoms. Consequently, typical changes in energy in nuclear reactions are much larger than those in chemical reactions - typically **hundreds of thousands** or **millions**, of times greater. Therefore, energies released in nuclear reactions are measured in millions of eV (MeV).

A **nuclide** is a nucleus with specific numbers of protons and of neutrons (e.g., the nuclide ^{16}O has 8 protons and 8 neutrons). The number of protons (Z) is the **atomic number**, and characterizes a specific **element** (since Z must also be the number of electrons in the neutral atom, and the number of electrons dictates the unique chemistry of the element). Different nuclides of a **given** element may exist, all with the same Z but with different numbers of neutrons (N); these are different **isotopes** of the element. Protons and neutrons have almost the same mass, and $A \equiv Z+N$ is the **atomic mass**.

The shorthand notation for nuclides is ^{A}X, where X is the chemical symbol of the element with atomic number Z. For instance, uranium (chemical symbol U) has 92 protons, i.e. Z = 92. Natural uranium is made up almost entirely of two isotopes (see Figure "Two Uranium Isotopes"), ^{238}U (with 146 neutrons) and ^{235}U (with 143 neutrons). ^{238}U is by far the most abundant isotope, representing 99.3% of natural uranium; in contrast, ^{235}U makes up only 0.7%. However, ^{235}U is the more important isotope of uranium as far as nuclear reactors are concerned.

1.4 Nuclear Reactions

Nuclear reactions abound in nature:

- Some nuclei transform (decay) by emission of α, β, or γ radiation. An α particle is a nucleus of helium, ^{4}He. A β particle is an electron. And a γ ray is a very energetic X-ray. These types of decay are examples of transmutations of elements (the alchemists' dream!), e.g.
  - α-decay of ^{238}U to thorium:
    $^{238}\text{U} \rightarrow {}^{234}\text{Th} + \alpha$
  - β-decay of molybdenum, produced in a nuclear reactor, to (isomeric) technetium, a radioisotope of great use in diagnostic medicine:
    $^{99}\text{Mo} \rightarrow {}^{99m}\text{Tc} + \beta$
- Nuclei interact in various ways with particles (α, β, γ, or neutrons) which collide with them, e.g.
  - the break-up of deuterium D (i.e., heavy hydrogen ^{2}H) by energetic gamma rays, producing a proton (H) and a **photoneutron:**
    $\text{D} + \gamma \rightarrow \text{H} + \text{n}$
  - the absorption of a neutron by deuterium to yield an even heavier isotope of hydrogen, tritium, T (or ^{3}H)
    $\text{D} + \text{n} \rightarrow \text{T}$
- Nuclei can under some circumstances interact with one another, e.g.
  - the **fusion** reaction between deuterium and tritium, the one of greatest interest at the present time for application in a fusion reactor (but note that the temperature must be millions of degrees!):
    $\text{D} + \text{T} \rightarrow {}^4\text{He} + \text{n}$

1.5 Fission

Nuclear fission is the splitting of a (large) nucleus with the release of energy. The nuclei of some heavy elements - e.g., ^{238}U – do exhibit **spontaneous** fission in nature. However, the rate of such fissions is extremely low. The main decay mode of ^{238}U, however, is α-decay, and the ^{238}U half-life is $4.47*10^9$ years; its half-life for spontaneous fission alone would be $0.8*10^{16}$ years. Spontaneous fission is **not** of much use to us as an energy source!

To apply nuclear energy, we need to have a nuclear reaction which produces energy (by mass conversion), and which can be continuous and controllable.

Neutron-induced fission is such a reaction. Let us look at the neutron-induced fission of ^{235}U (Figure "Neutron-Induced Fission Reaction"):

$$^{235}\text{U} + \text{n} \rightarrow \text{ 2 large pieces (fission products)} + \text{2-3 neutrons} + \beta, \gamma + \text{energy} \qquad (3)$$

1.6 Chain Reaction and Multiplying Medium

Neutron-induced fission [reaction (3)] is the operating principle of fission reactors. It satisfies the criteria for an energy source, since:

- energy is "produced" (actually, liberated), and
- the process is **open to control** by controlling the number of fissions, or, in fact, the number of neutrons**.**

The crucial feature of neutron-induced fission (Eq. 3) is that while a neutron initiates the reaction, some neutrons also emerge. Since the fission reaction "swallowed" one neutron and released 2 or 3 (sometimes more) neutrons, a medium which contains uranium (or other fission fuel) is a **multiplying medium.** Since the **fission neutrons** released can induce more fissions, there is then the possibility of a **chain reaction**, i.e., the potential of a self-perpetuating reaction.

The chain reaction will increase or decrease in intensity, or remain steady, depending on what happens to the neutrons, i.e., on whether the number of neutrons increases, decreases, or remains constant. Under the right conditions, then, a steady and controlled release of energy will be possible.

The medium's **multiplication factor** is the crucial quantity to be considered, which we will do further below.

1.7 Fissionable and Fissile Nuclides

There are only a few nuclides which can fission.

A nuclide which can be induced to fission by an incoming neutron of **any** energy is called **fissile**. There is **only one naturally occurring fissile** nuclide: ^{235}U. Other fissile nuclides are ^{233}U and some isotopes of plutonium (to be discussed below), ^{239}Pu and ^{241}Pu, but none of these occurs in nature to any appreciable extent.

Some nuclides can be induced to fission, but only if the incoming neutron has an energy higher than a certain threshold; these are not fissile nuclides, but they are called **fissionable**. Examples of fissionable, but not fissile, nuclides are ^{238}U and ^{240}Pu. These nuclides are, in a sense, less "easy to work with" as far as the chain reaction is concerned, because, for one thing, the inducing neutron's energy must be maintained high – which is not that easy to do.

1.8 Energy Output

In fission, part of the mass of the original uranium nucleus is converted to energy. The energy produced per single fission is ~ 200 MeV [~ $3.2*10^{-11}$ J]. This is several orders of magnitude greater than the energy produced by combustion. However, it still represents approximately only **0.09%** of the mass energy of the uranium nucleus!

The energy appears mostly (85%) as kinetic energy of the fission fragments (the 2 large pieces), and in small part (15%) as the kinetic energy of the neutrons and the energy in the β and γ particles (rays). **The energy is quickly reduced to heat, which can be used to make steam from water, and generate electricity.** See Figure "Schematic of a CANDU Nuclear Power Plant".

The total power (energy per unit time) generated in a nuclear reactor will depend on the number of fissions which occur per second. The quantities of interest in the nuclear reactor are the **fission power** (total power generated in fission), the **thermal power** (the power removed by the coolant), and the electric power.

In CANDU 5% of the fission power is absorbed in the heavy-water moderator in the calandria, and does not contribute to the thermal power. Thermodynamics limits the fraction of the thermal power which can be turned into electricity to about 30-33%.

As an example, in the CANDU 6:

Fission power = 2156 MW(th)
Thermal Power = 2061 MW(th)
Gross Electric Power = 680 MW(e)

[Exercise: Given that one fission releases 200 MeV, how many fissions occur per second in a CANDU 6 at full power?]

1.9 Products of Fission

The fission products (fission fragments) are nuclides of roughly half the mass of uranium. However, note that they are not always the same in every fission. There are a great number of different fission products, each produced in a certain percentage of the fissions. Most fission-product nuclides are "neutron rich"; they decay typically by β- or γ-disintegration, and are therefore radioactive, with various half-lives. To prevent the release of radioactivity, therefore, the used fuel is safely stored **and contained.**

Some examples of long-lived fission products are:

- ^{85}Kr, half-life 10.4 y
- ^{90}Sr, half-life 28 y

- ^{137}Cs, half-life 30 y
- ^{99}Tc, half-life $2.1*10^5$ y
- ^{129}I, half-life $1.7*10^7$ y

Because many fission products are still decaying long after the originating fission reaction, energy (heat) from this nuclear decay is actually produced in the reactor for many hours, days, even months after the chain reaction is stopped. This **decay heat** is not negligible. When the reactor is in steady operation, decay heat represents about 7% of the total heat generated. Even after reactor shutdown, decay heat must be dissipated safely, otherwise the fuel and reactor core can seriously overheat. Figure "Decay-Heat Curves vs. Time" shows the variation of decay heat with time.

1.10 Plutonium and Other Actinides

As the only naturally occurring fissile nuclide, ^{235}U is the most important nuclide for fission reactors. However, plutonium is also very important, because it is created and burned in the reactor. ^{239}Pu is produced from the absorption of neutrons by ^{238}U:

$$^{238}\text{U} + \text{n} \rightarrow {}^{239}\text{U} \rightarrow {}^{239}\text{Np} + \beta \rightarrow {}^{239}\text{Pu} + 2\,\beta$$

As mentioned earlier, ^{239}Pu itself is fissile. As a result, it participates in subsequent fissions in the fuel, and contributes to the energy production. By further neutron absorption in ^{239}Pu, some ^{240}Pu and ^{241}Pu (the latter also fissile) are also produced. In a CANDU reactor, half the energy produced is from the plutonium created "in situ"!

Because ^{238}U yields fissile material (via neutron absorption), it is said to be **fertile**.

Plutonium is in the **actinide** series of elements and is a **transuranic**, since its atomic number is higher than 92. Other examples of transuranic actinides, also produced in smaller amounts in fission reactors, are curium and americium. [The latter is used in smoke detectors.]

Transuranics tend to have long half-lives, e.g. the half-life of ^{239}Pu is 24,000 years.

1.11 Fast and Thermal Neutrons

Neutrons emerging from fission are typically very energetic: the distribution of fission-neutron energies has its maximum at ~1 MeV (see Figure "Energy Distribution of Fission Neutrons"). From the rest mass of a neutron (940 MeV), one can calculate that this kinetic energy corresponds to a neutron speed of ~13,800 km/s!

But the probability of a neutron inducing a fission reaction in ^{235}U is much larger for slow neutrons than for fast neutrons – see Figure "Schematic View of a Typical Cross Section, Showing Resonances". As a result, to increase the likelihood of perpetuating the fission chain reaction, we want to slow the neutrons down as much as possible. The maximum extent to which neutrons can be slowed down is to energies at which the neutron population is in **thermal equilibrium** with the ambient environment ("thermal" energies). For a temperature of 20^{o} C, thermal energies are of the order of 0.025 eV. A thermal neutron with an energy of 0.025 eV has a speed of 2.2 km/s.

1.12 Moderators

Fast neutrons are slowed to thermal energies by a **moderator**. The moderator most efficient at slowing down neutrons is one whose atoms are of about the same mass as a neutron. Since protons and neutrons have very nearly the same mass, and since the nucleus of a hydrogen atom is a proton, then hydrogen (as, say, in light water) could be assumed to be the best moderator. In fact, pressurized-water reactors (PWR) use ordinary water as the moderator.

However, hydrogen absorbs neutrons easily, and therefore takes neutrons "out of circulation". The fission chain reaction cannot be self-sustaining with a light-water moderator: the fuel must be enriched in the fissile isotope ^{235}U.

The nucleus of heavy hydrogen, or deuterium, consists of one proton and one neutron; the atomic mass is 2. Deuterium (as in heavy water) is also therefore very effective as a moderator. In addition, deuterium does **not** absorb neutrons readily. This is a great advantage as far as perpetuating the chain reaction. We say that the heavy-water moderator promotes excellent **neutron economy**. Heavy water is the moderator used in CANDU reactors.

1.13 Fuel Requirements in Perspective

The energy generated in fission per unit of fuel is immense compared to the energy generated in combustion. One kilogram of uranium in a CANDU reactor produces about 180 MW.h of fission energy, or about 60 MW.h of electricity. If my 4-person household's average electricity use is typical, about 1,000 kW.h per month [12,000 kW.h = 12 MW.h per year], then a single kg of uranium is sufficient for 5 households for one year. That is, a mere **200 g** (< 0.5 lb) of uranium will serve the electricity needs of a household for an entire year. If the electricity were obtained from coal, then the mass of fuel would need to be about 30,000 times as large, i.e. about **6,000 kg** of coal. Consequently, the cost of nuclear electricity is insensitive to fluctuations in the price of uranium.

A CANDU-6 reactor producing about 680 MW of electricity uses approximately 120 tonnes of uranium per year. This corresponds to a volume of uranium of only **12 m^3! This is to be compared to the ~ 4,000,000 tonnes of coal which would be required in a fossil plant.**

The same ratio applies to the amounts of used uranium fuel and ash from the generating stations. The very small mass of used nuclear fuel is safely stored and isolated from the environment. This is not true of the products of combustion (CO_2, SO_2, NO_x, …), which end up in the atmosphere. **Since most coal also contains uranium, a fossil plant may actually release more radioactivity into the environment than a nuclear generating station!**

2. Quantifying the Chain Reaction

2.1 Reactor Multiplication Constant

Several processes compete for neutrons in a nuclear reactor:

- "productive" absorptions, which end in fission
- "non-productive" absorptions (in fuel or in structural material), which do not end in fission
- leakage out of the reactor

The self-sustainability of the chain reaction depends on the relative rates of production and elimination (or loss) of neutrons. It is measured by a quantity called the *effective reactor multiplication constant*, k_{eff}, which can be defined as follows:

$$k_{eff} = \frac{Rate\ of\ production\ of\ neutrons}{Total\ rate\ of\ loss\ of\ neutrons} \tag{2.1}$$

i.e., $$k_{eff} = \frac{Rate\ of\ production\ of\ neutrons}{Rate\ of\ absorption\ of\ neutrons + rate\ of\ leakage} \tag{2.2}$$

There are three possibilities for k_{eff}.

$k_{eff} < 1$: Fewer neutrons are being produced than are eliminated. The chain reaction is not self-sustaining, the number of neutrons decreases with time, and the reactor eventually shuts down. The reactor is subcritical.

$k_{eff} = 1$: Neutrons are being produced at the same rate as they are eliminated. The chain reaction is exactly self-sustaining, the neutron population is stable, and the reactor remains in a steady state. The reactor is critical.

$k_{eff} > 1$: More neutrons are being produced than are eliminated. The chain reaction is more than self-sustaining, and the neutron population and reactor power increase. The reactor is supercritical.

2.2 Critical Mass

The type of fuel in the reactor, and the amount of fuel relative to neutron-absorbing materials, dictate the multiplying properties of the medium within the reactor. But the neutron multiplication also depends on the leakage of neutrons, and consequently on the size of the reactor.

Leakage of neutrons is a phenomenon occurring (obviously) at the surface of the reactor. Therefore, leakage increases as the size (and therefore the surface-to-volume ratio) of the reactor decreases. If the reactor is made smaller while its composition is kept constant, the leakage at some point becomes too large for the chain reactor to remain self-sustaining, i.e., the reactor must have a **minimum** size to work. Below this minimum size or **critical mass,** the leakage is too high and k_{eff} cannot possibly be equal to 1. The critical mass depends on the **shape** of the reactor, the **composition** of the fuel, and the **other materials** in the reactor. The shape with the lowest relative leakage, i.e. the shape for which the critical mass is least, is the shape with the smallest surface-to-volume ratio, i.e. a sphere.

2.3 Reactivity

Reactivity is a quantity closely related to reactor multiplication constant. Reactivity is usually denoted by the symbol ρ, and is defined as

$$\rho \equiv 1 - \frac{1}{k_{eff}} \tag{2.3}$$

$$= 1 - \frac{Neutron\ loss}{Neutron\ production} \tag{2.4}$$

$$= \frac{Neutron\ production - Neutron\ loss}{Neutron\ production} \tag{2.5}$$

$$\equiv Relative\ \textbf{\textit{net}}\ neutron\ production$$

Reactivity is useful because it is of opposite sign for subcritical and supercritical reactor states. Whereas the "central" value for k_{eff} is 1, for reactivity it is 0:

$$k_{eff} < 1 \Rightarrow \rho < 0,$$

$$k_{eff} = 1 \Rightarrow \rho = 0, \tag{2.6}$$

and $$k_{eff} > 1 \Rightarrow \rho > 0.$$

2.4 Control of the Reactivity

To operate a nuclear reactor, we want **most of the time** to keep $k_{eff} = 1$ ($\rho = 0$) so that everything is nice and steady.

When we want to reduce power or shut the reactor down, we need ways to make $k_{eff} < 1$ (i.e., insert negative reactivity). This is usually done by inserting rods or devices made of strong neutron absorbers, such as boron, cadmium, or gadolinium.

When we are starting up the reactor and want to increase power, we need to make k_{eff} slightly greater than 1 (i.e., insert positive reactivity), for a short time. This is usually done by removing a bit of absorption. Note that, in a reactor, we **don't want** to make k_{eff} **much greater than 1**, or greater than 1 for a long time, or the power could increase to high values, potentially with undesirable consequences, e.g. melting of the fuel.

Even when we want to keep $k_{eff} = 1$, we need reactivity devices to counteract perturbations to the multiplication constant. The movement of reactivity devices allows absorption to be added or removed in order to manipulate k_{eff}.

Every nuclear reactor contains regulating and shutdown systems to do all the above jobs.

[A nuclear bomb, in contrast to a reactor, is designed to be **very** supercritical on **fast** neutrons, to generate a huge amount of **uncontrolled** energy in a very short time. **No reactor can explode like a nuclear bomb.**]

2.5 Concept of Nuclear Cross Section

In reactor physics, the rates at which various reactions occur are prime quantities of interest. In particular, it is important to know the rates of reactions that neutrons undergo as they travel in the reactor core. This section and the next two introduce quantities needed to calculate reaction rates.

As neutrons diffuse through the materials of the reactor core, they may enter into a number of reactions with nuclei of various elements: scattering (elastic or inelastic, depending on whether the energy state of the nucleus is changed), absorption, fission, or other reaction.

Imagine a monoenergetic beam of neutrons of speed v, all moving in the same direction, impinging upon a (very thin) slice of area S cm^2 and thickness Δx cm of a target material (see Figure "Neutron Beam Impinging on a Slice of Target Material"). Let the beam intensity be I(x), in units of neutrons.cm^{-2}, where x measures the distance into the material (note that I(x) is attenuated with x, as the beam travels through the material and neutrons undergo reactions).

We may think of the nuclei in the target as presenting a certain effective area to incoming neutrons for a given type of reaction (e.g. scattering, absorption, fission, etc.). This area is called the **microscopic cross section σ**, **the effective area presented to the**

neutron by one nucleus of the material. It is important to note that σ may be totally unrelated to the physical size or radius of the nucleus – it is an effective area which represents the "range" of interaction for the particular nucleus and nuclear reaction with a neutron. In general, σ depends on the type of nucleus, the type of reaction studied, and on the incoming neutron's energy. It has units of energy and can be expressed in cm^2, although a much more appropriate unit is the **barn = 10^{-24} cm^2**, or sometimes the **kilobarn (kb) = 10^{-21} cm^2**.

Now, let us consider how many nuclei of the target material exist in the thin slice. The number of nuclei N per cm^3 of material can be determined from the density of the material and its atomic (or molecular) mass. Then in the thin slice the total number of nuclei of the material must be equal to:

(N per cm^3) times (S cm^2) times (Δx cm)= NSΔx.

So the total effective area presented to each neutron in the beam by all target nuclei in the slice is NSΔx**σ.** This, divided by the surface area S of the slice, gives the probability of the specific reaction occurring in the slice, per neutron in the beam:

Probability of reaction in target material per incident neutron = NSΔx**σ**/S = N**σ**Δx.

Further, we can write this probability as a probability per unit path length of the neutron. This path length traversed by a neutron in the beam is Δx, then:

Probability of reaction in target material
per unit path length traversed per incident neutron = N**σ**Δx/Δx = N**σ**. (2.7)

This quantity is called the **macroscopic cross section Σ** of the material for the nuclear reaction in question; i.e., the relationship between **microscopic** and **macroscopic** cross sections is:

$$\Sigma = N\sigma \tag{2.8}$$

Note: If several different types of nuclei are present in the material, then a number of partial products Nσ for the various nuclide types must be added together to give Σ.

The microscopic cross section **σ** is a basic physical quantity which is determined by experiments of neutron beams of various energies on target materials. Once **σ** is known, then the macroscopic cross section **Σ** can be obtained from Eq. (2.8). **Both σ and Σ depend on the material, the neutron energy or speed, and the type of reaction.**

The macroscopic cross sections for scattering, absorption, fission are denoted Σ_s, Σ_a, and Σ_f respectively. The total cross section Σ_{tot} measures the **total** number of all types of reaction per unit distance:

$$\Sigma_{tot} = \Sigma_s + \Sigma_a \ . \ . \ . \tag{2.9}$$

(Note that the fission cross section is included in the absorption cross section, since it occurs following a neutron absorption.)

Incidentally, another meaning can be attributed to the total macroscopic cross section: it can be shown that it is the reciprocal of the mean free path of neutrons in the material:

$$\lambda = \frac{1}{\Sigma_{tot}} \tag{2.10}$$

For a reaction of type i, and from the definition of the associated macroscopic cross section Σ_i, the total number of reactions i that **1 neutron** is expected to undergo in a projected path length s of travel is given by $\Sigma_i.s$. **In addition, the total number of reactions of a group of neutrons will also be $\Sigma_i.s$ if s is the total cumulative projected path length of all the neutrons.**

2.6 Concept of Neutron Flux

In the previous section, we introduced the concept of probability of reaction per unit path length, and discussed the simple case of a monoenergetic beam of neutrons. The concept of flux generalizes the idea of path length.

Imagine all the neutrons in a unit volume at a given instant. The neutron population density is denoted by n; it can be expressed in units of neutrons/cm^3. In the unit volume, there are many neutrons, travelling at various speeds in various directions. Let us imagine summing all the distances (path lengths) which would be traversed by these neutrons per unit time if they kept travelling at their various speeds. This is the concept of total neutron flux, denoted ϕ (see Figure "Concept of Neutron Flux"). The standard units for ϕ are (neutrons/cm^3).cm/s, i.e., neutrons.cm^{-2}s^{-1}.

In the (hypothetical) case in which all neutrons are travelling at the same speed v, the flux ϕ is the product of the density n of the neutron population and the speed v:

$$\phi(v) = nv \qquad \text{(for monoenergetic neutrons)} \tag{2.11}$$

When there is a distribution of speeds v, the total flux is obtained by integrating over velocity:

$$\phi \equiv \int n(v)v\,dv \tag{2.12}$$

[One may wish to focus on neutrons of a particular speed or range of speeds when calculating a neutron flux. Thus, one may calculate the thermal-neutron flux, the "fast-neutron" flux, etc., by integrating over an appropriate range of speeds.]

In a real reactor, the values of neutron population and flux depend on position in the core. Because flux is an essential ingredient in the computation of reactor rates (see next Section), the calculation of the spatial distribution of flux in the core is an important part of reactor physics (see later Sections). The value of flux at a given point in the core will depend on the distribution of nuclear properties (cross sections) throughout the core, and on the position at which the flux is calculated. The flux drops to zero at, or just beyond, the radial and axial boundaries of the core. For any real reactor, the 3-d flux distribution in the core must be calculated by means of (large) computer codes.

When the neutron flux is not spatially uniform, there is at any point a net leakage of neutrons from regions of high flux to regions of low flux. The net movement of neutrons at any point can be expressed in terms of the neutron current, denoted by the vector quantity $\vec{j}$. Just as the neutron flux depends on the position in the reactor, so does the neutron current. In diffusion theory, it can be shown that the neutron current is proportional to (the negative of) the gradient of the flux:

$$\vec{j} = -D\nabla\phi \tag{2.13}$$

where the proportionality constant D is called the diffusion coefficient. D is a function of the properties of the medium, and also depends on the neutron energy (or speed). It is (in diffusion theory) related to the scattering cross section Σ_s by:

$$D = \frac{1}{3\Sigma_s} \tag{2.14}$$

[Note: in the more exact neutron transport theory, this relation becomes:

$$D = \frac{1}{3\Sigma_{tr}}$$

where Σ_{tr}, called the transport cross section, is a corrected scattering cross section.]

2.7 Calculating Reaction Rates

Putting together the concepts of neutron flux and cross section, one can calculate reaction rates.

The reaction rate for a given process at a given neutron speed v per unit volume per unit time is the product of

- the total path length of neutrons of speed v per volume per time (i.e., flux of neutrons of speed v) and
- the corresponding macroscopic cross section for speed v:

Reaction rate (per unit volume) for neutrons of speed v = $\Sigma(v)\phi(v)$ (2.15)

If there is a distribution of neutron speeds, the reaction rate can be integrated over the distribution, i.e.,

Reaction rate (per unit volume) over all neutrons = $\int \Sigma(v)\phi(v)dv$ (2.16)

2.8 Concept of Irradiation

The irradiation of a material, denoted ω, is a measure of the time spent by the material in a given neutron flux. Mathematically, it is defined as the product of flux by time:

$$\omega = \phi.t \tag{2.17}$$

The units of irradiation are neutrons/cm^2, or more conveniently, neutrons per thousand barns (neutrons per kilobarn). This last unit is written as n/kb.

2.9 Fuel Burnup

Fuel burnup is the (cumulative) quantity of fission energy produced per mass of nuclear fuel during its residence time in the core. The energy considered here is the energy released as a result of fissions, except for the neutrino energy, which cannot be recovered.

The two most commonly used units for fuel burnup are Megawatt-hours per kilogram of uranium, i.e., MW.h/kg(U), and Megawatt-days per Megagram (or Tonne) of uranium, i.e., MW.d/Mg(U).

These units are related by the equation:

1 MW.h/kg(U) = 1,000/24 MW.d/Mg(U) = 41.67 MW.d/Mg(U)

Burnup and irradiation are closely linked. Burnup is a monotonic, nearly linear function of irradiation, as shown in Figure "Relationship Between Irradiation and Burnup". This relationship is a result of the fact that the energy from fission is almost the same whether the fissioning nuclide is one of uranium or of plutonium.

Fuel burnup is an important economic quantity. Fuel burnup is essentially the inverse of fuel consumption, which is the amount of fuel used to produce a given quantity of energy (electricity). Units for fuel consumption are, for example, Mg(U)/GW(e).a. For a given fissile content, a high burnup signifies low fuel consumption, and therefore a small refuelling-cost component.

Note, however, that the true measure of a reactor's efficiency is uranium utilization, the amount of uranium "from the ground" needed to produce a certain amount of energy. Fuel burnup can be high and fuel consumption low because of isotopic enrichment of the fuel, but it is the amount of natural uranium used which is most important from the point of view of resource utilization. Uranium utilization is good (low) in CANDU because of its high 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.

2.10 Delayed Neutrons and Neutron Kinetics

Any imbalance between neutron production and loss causes the neutron population to increase (or decrease) from one generation of neutrons to the next. It is therefore natural to think that the *rate* at which the neutron population (and, consequently, the power) will change will depend on the mean generation time T, the average time interval between successive neutron generations. In a simplistic treatment of kinetics, in fact, the power varies exponentially with reactivity and with time in units of T:

$$P = P_0 \exp\left(\frac{\rho t}{T}\right)$$

However, this simplistic treatment does not account for delayed neutrons. Neutrons produced in fission are either prompt or delayed. The prompt-neutron lifetime Λ (average time interval between birth of a neutron and its absorption in a subsequent fission reaction) in the CANDU lattice is approximately 0.9 millisecond.

If there were no delayed neutrons, the mean generation time would be identical to the prompt-neutron lifetime. In that case, a reactivity of 1 milli-k would lead to a power increase by a factor of 3 in 1 second, a factor of 9 in 2 s, etc. Quite a fast rate of change indeed! (In LWRs, Λ is about 30 times shorter! The rate of change of power would then be 30 times as great for the same reactivity.)

Delayed neutrons, however, reduce the rate of power change considerably. This is so even though the number of delayed neutrons is a very small fraction (~0.6%) of all

neutrons from fission. Most delayed neutrons are produced in the beta decay of fission products, the delayed-neutron precursors. For practical purposes, delayed-neutron precursors from any one fissionable nuclide can be subdivided into six distinct groups, with beta half-lives ranging from 0.2 s to 50 s (see Table 1). In CANDU, there are additional delayed neutrons, produced by the photodisintegration of the deuterium in heavy water. These photoneutrons appear from another eleven distinct precursor groups, with even longer time constants, in the hundreds of seconds to tens of thousands of seconds (see Table 2).

These delayed-neutron and photoneutron time constants are such that, in spite of the small delayed fraction, the "effective" (weighted-average) mean generation time will be much longer than the prompt-neutron lifetime. In fact **in CANDU the mean generation time T is the order of 0.1 second**, about 100 times the ~1-ms prompt-neutron lifetime. In this case, the exponential equation a reactivity of 1 milli-k would then lead to an increase in power of a factor of only about 1.01 per second, compared to 3 per second without delayed neutrons. It is clear that delayed neutrons have a large influence on the evolution of power in transients and facilitate reactor control considerably.

However, it must be noted that the above "exponential" treatment is really inadequate for a proper analysis. The correct treatment involves solving a coupled set of equations for the time-dependent flux distribution and the concentrations of the individual delayed-neutron precursors. This will not be described here, as it is beyond the scope of this short presentation.

3. Basic Characteristics of the CANDU Lattice

This chapter discusses the basic characteristics of the CANDU lattice. The discussion is largely general but, where necessary, particular reference is made to the CANDU 6 reactor.

3.1 The CANDU Lattice Cell

The Figure "Schematic Face View of CANDU 6 Reactor" shows that the CANDU design is modular, with fuel channels set on a square lattice of lattice pitch equal to 28.575 cm. Figure "Face View of Basic CANDU Lattice Cell" shows (not to scale) the basic lattice cell in CANDU, which has dimensions of 1 lattice pitch by 1 lattice pitch (28.575 cm x 28.575 cm) by 1 fuel-bundle length (49.53 cm) – the 3rd dimension is not shown in the figure. There are twelve fuel bundles in each fuel channel.

The next few sections describe the various components of the lattice cell.

3.2 The Moderator

The CANDU reactor design has opted for heavy water as the moderator, to take advantage of the neutron economy provided by deuterium. This allows the use of natural-uranium fuel and precludes the need for expensive fuel-enrichment technology.

The first desirable property for a moderator is the ability to thermalize neutrons in as few collisions as possible. When the number of collisions required for thermalization is smaller, the average loss of neutron energy per collision is greater, and the probability is enhanced that the neutron will miss the resonance-absorption energy range (or much of it) during moderation (see Figure "Number of Collisions and Energy Loss per Collision During Moderation"). Figure "Average Number of Collisions to Thermalize a Fission Neutron" shows the average number of collisions needed for various moderators to thermalize a fission neutron.

Another feature which helps to reduce the neutron capture in fuel resonances is the lumping of the fuel in fuel channels. This enhances the probability of fission neutrons being slowed down in the moderator volume between fuel channels (see Figure "Fission Neutrons Slowed in Moderator Region"), and therefore reducing the probability of neutrons interacting with the fuel when in the resonance-energy range.

It is also important that a moderator have a small probability of capturing neutrons (i.e., a small neutron absorption cross section), otherwise there will be a negative impact on neutron economy. In this respect hydrogen is not the best moderator,

because it has a relatively high neutron absorption cross section. Heavy hydrogen, or deuterium, on the other hand, has a very low absorption cross section.

A good index of performance for moderators is the moderating ratio, the ratio of the slowing-down power of the material to its neutron absorption cross section:

$$Moderating\ ratio = \frac{Slowing-down\ power}{Absorption\ cross\ section} = \frac{\xi\Sigma_s}{\Sigma_a}$$

where ξ is the mean logarithmic energy decrement per collision.

The Figure "Moderating Ratio of Various Moderators" clearly shows that by far the best moderator for neutron economy is heavy water (D_2O), which is why it was chosen as the moderator for CANDU.

Because of the significant probability of neutron capture by any light water present in the heavy-water moderator, it is crucial that the latter have very high isotopic purity. Reactor-grade moderator must be at least 99.75 % D_2O by weight. Even a reduction of 0.1% in the isotopic purity has a significant effect on the neutron economy of the reactor and on the achievable fuel burnup.

3.3 The Fuel

Natural uranium is used in all currently operating CANDU reactors. This is very convenient for countries which wish not to have to rely on expensive, and most probably foreign, enrichment technology. However, the CANDU design is very flexible and allows the use of advanced fuel cycles, using slightly enriched uranium (SEU), recovered uranium (RU), mixed-oxide fuel (MOX), thorium fuels (Th), and others (DUPIC, actinide burning). These can be introduced into CANDU with few or no hardware changes, when the option becomes attractive.

CANDU fuel is of very simple design. It is manufactured in the form of elements of length ~48 cm. Each element consists of uranium-dioxide pellets encased in a zircaloy sheath. A number of fuel elements are assembled together to form a bundle of length ~50 cm. The elements are held together by bundle end plates. The CANDU fuel bundle contains only 7 different components and is short, easy to handle, and economical.

Various fuel-bundle designs are illustrated in Figure "Various Fuel-Bundle Designs – to 37-Element Fuel". Only two bundle types are used in present-generation CANDUs: the 28-element bundle (in Pickering) and the 37-element bundle (in Bruce, Darlington and the CANDU 6). The 28-element bundle has a smaller ratio of sheath mass to fuel mass than the 37-element bundle, which gives the 28-element bundle a reactivity advantage. On the other hand, the 37-element bundle features better thermalhydraulic properties due to the greater fuel subdivision, as the larger number of

pins of smaller diameter provide a better heat-removal capability. Thus, the 37-element bundle can operate at a higher power than the 28-element bundle. This tends to further reduce the reactivity of the 37-element bundle, but allows a higher total reactor power for the same mass of fuel, an important economic advantage.

CANFLEX fuel (Figure "CANFLEX 43-Pin Fuel Bundle)) has been under development for the last few years as the fuel bundle for the future. The CANFLEX bundle has 43 elements, with the outer two rings of elements being of smaller diameter than the inner 7 elements. The CANFLEX bundle features improved thermalhydraulic properties and 20% lower maximum element-power ratings than 37-element fuel, for the same bundle power. In 1998-2000, a demonstration irradiation of 24 CANFLEX fuel bundles was initiated in the Pt. Lepreau CANDU-6 reactor.

3.4 The Coolant

In all commercial CANDU reactors, heavy water is used as the coolant in the primary heat-transport system, to further improve neutron economy. However, prototype CANDUs have been built using boiling light water or an organic liquid as coolant (Gentilly-1 and WR-1 respectively). The organic coolant, in particular, allows higher temperatures and greater efficiency of conversion of heat to electricity.

3.5 The Pressure-Tube Concept

A major characteristic, selected early in the development of the CANDU reactor, is the pressure-tube design. It is clear that if a liquid is used to remove the large quantity of heat generated inside the reactor, the liquid must be kept at high pressure, otherwise it would boil. The heat-transport-system pressure in CANDU is ~100 atmospheres.

To contain the pressure, the choice is between a pressure-vessel design and a pressure-tube design. In the former (e.g. the PWR) the vessel contains all the fuel and the liquid, which is at once both moderator and coolant. In the pressure-tube design, the moderator and coolant are separate and the coolant flows (at high pressure) through the pressure tubes, which also contain the fuel and in fact comprise the fuel channels. The pressure tubes are made of an alloy of zirconium and 2.5% niobium.

The pressure-tube concept was originally chosen for CANDU because the manufacture of a pressure vessel of the size required for a heavy-water reactor (HWR) would at the time have challenged the capability of Canadian industry. However, the pressure-tube concept has many other advantages in relation to the design and safety of the reactor:

- The rupture of one pressure tube is not as catastrophic as the rupture of an entire pressure vessel. Also, in most cases a pressure tube will leak and give ample warning before rupturing.

- Because the coolant and moderator are physically separated, the moderator can be kept relatively cool. In CANDU the moderator is isolated from the hot pressure tube by a concentric calandria tube made of zircaloy-2. A gas annulus separates the pressure and calandria tubes (see Figure 3.2). Thus the moderator can be kept at about 70 °C and at near atmospheric pressure. Many safety benefits ensue as a consequence. The moderator is a benign, low-pressure and low-temperature environment for interstitial reactivity devices (control rods, etc.); rod-ejection accidents are therefore not a concern. Also, the moderator is a potential ultimate heat sink in case of accident. In addition, a cool moderator further improves neutron economy.

- The pressure-tube concept allows the replacement of fuel in the reactor on power, precluding the need for periodic shutdowns for refuelling.

- Also, on-power refuelling means that the excess reactivity in the lattice is never very high, a safety advantage.

Figure "CANDU First Design Principles" summarizes these features.

3.6 Fuel Burnup and Effect of Operating Conditions

The exit fuel burnup attained in the reactor depends on the operational parameters of the core. The burnup is of course influenced by any quantity which affects the core reactivity. Any neutron loss or 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 $\rightarrow$ 2.88 MW.h/kg(U) loss in burnup
= 120 MW.d/Mg(U) loss in burnup

Examples of factors which affect the reactivity, and therefore the attainable fuel burnup, are as follows:

- a higher moderator purity increases burnup (the reactivity coefficient of moderator purity = ~ 34 milli-k/atom % purity)
- a higher coolant purity also increases burnup (but much less than a higher moderator purity; the reactivity coefficient of coolant purity = ~ 3 milli-k/atom % purity)
- operating the reactor with moderator poison decreases burnup (the boron reactivity coefficient = ~ 8 milli-k/ppm(B)

- a reflector decreases leakage and increases burnup
- thicker pressure or calandria tubes decrease burnup
- a higher ratio of fuel-sheath mass to fuel mass in a bundle (everything else being equal) decreases burnup
- a lower moderator temperature increases burnup
- flattening the power distribution increases leakage and decreases burnup.

4. The CANDU Reactivity Devices

Every reactor design must include means of changing and adjusting the system reactivity. These are needed to:

- maintain the reactor critical for normal operation,
- allow power manoeuvres, and
- permit fast reactor shutdown when emergency conditions exist.

In CANDU reactors, the primary long-term method of reactivity control is on-line refuelling, where low-reactivity irradiated fuel is replaced by high-reactivity fresh fuel. On-line refuelling is carried out on a daily or near-daily basis. In addition, however, the CANDU design includes several types of reactivity devices for different kinds of reactivity control. These are described in this chapter.

In the CANDU 6 reactor there are six means of changing the reactivity state of the core besides refuelling. Four of these are used for normal control functions, including controlled shutdown, and two are used by special safety systems for rapid shutdown during accident conditions.

The reactivity devices used for control purposes by the Reactor Regulating System (RRS) in the standard CANDU-6 design are the following:

- 14 liquid-zone-control compartments (H_2O filled)
- 21 adjuster rods
- 4 mechanical control absorbers
- moderator poison.

There are in addition two special shutdown systems (SDS):

SDS-1, consisting of 28 cadmium shutoff rods which fall into the core from above
SDS-2, consisting of high-pressure poison injection into the moderator through 6 horizontally oriented nozzles.

Figure "Reactivity Worths of CANDU Reactivity Devices" gives typical reactivity worths and maximum rates of change of reactivity for these devices.

All reactivity devices are located or introduced into guide tubes permanently positioned in the low-pressure moderator environment. These guide tubes are located interstitially between rows of calandria tubes, as shown in Figure "Interstitial Positioning of Reactivity Devices". There exists no mechanism for rapidly ejecting any of these rods, nor can they drop out of the core. This is a distinctive safety feature of the pressure-tube reactor design. The maximum positive reactivity insertion rate achievable

by driving all control devices together is about 0.35 milli-k per second, which is well within the design capability of the shutdown systems.

The locations of the reactivity devices are shown schematically in Figures "Plan View of Reactivity-Device Locations", "Side-Elevation View of Reactivity-Device Locations", and "End-Elevation View of Reactivity-Device Locations".

The following sections describe the functions of the various types of reactivity devices.

4.1 Liquid Zone Controllers

The purpose of the liquid zone-control (LZC) system is to provide continuous fine control of the reactivity, and hence of the reactor power level. Fine reactivity control is needed because refuelling is not truly continuous, but instead is achieved in small increments (usually eight bundles at one time). Fine reactivity control also compensates for other minor perturbations in parameters, such as temperature changes, which in turn cause small reactivity changes. The liquid zone-control system is also designed to accomplish spatial control of the power distribution, which prevents xenon-induced power oscillations from developing.

The LZC system consists of six vertically oriented units (tubes) running interstitially between the fuel channels from the top to the bottom of the core, as shown in Figure "Liquid Zone-Control Units". The two central tubes are divided into three compartments each, and the four outer tubes into two compartments each; the compartments in each unit are separated by appropriately placed bulkheads - see Figure "Liquid Zone-Control Compartments". There are thus a total of 14 individual zone compartments in the reactor. Variable and controllable amounts of light water (H_2O) are introduced in the compartments, where it serves as a neutron absorber. H_2O is fed to the compartments through small-diameter tubing, and the level of H_2O in each compartment is controlled by varying the relative value of the in-flow and out-flow rates.

The reactor regulating system (RRS) adjusts the H_2O fills in the individual compartments according to the magnitude of the signals from interstitially placed in-core self-powered detectors. The detector systems are described in a later chapter.

4.2 Mechanical Control Absorbers

The zone-control system is normally designed to provide a capability for reactivity control of about ± 3 mk, since this is sufficient to compensate for routine reactivity perturbations due to refuelling, occurring on a semi-continuous basis. For certain, less frequent, events, the reactor regulating system requires a greater reactivity range than the zone-control system can provide. Therefore, two additional reactivity-

device systems are provided, to extend the control capability in the positive and negative reactivity directions. These devices are also operated by the reactor regulating system.

The system used to extend the range of control in the negative-reactivity direction is a system of four mechanical control absorbers (MCAs). These are physically the same as the shutoff rods (see Section 4.5.1 below), but they do not form part of the shutdown system. The control absorbers are normally parked fully outside the core under steady-state reactor operation. They are moved into the core only when circumstances demand a rapid reduction of the reactor power, at a rate or over a range that cannot be accomplished by filling the liquid zone-control system at the maximum possible rate. Modes of control-absorber insertion range from driving the rods in pairs to all four being dropped in by gravity following release of an electromagnetic clutch.

The mechanical-control-absorber system and the zone-control system can be used to reduce power to a very low value without requiring actuation of either of the shutdown systems. The reactivity worth of the MCAs is such that it can compensate for the reactivity increase due to temperature reduction on shutdown.

The positions of the mechanical control absorbers are shown as X's in Figure "X = Mechanical Control Absorbers".

4.3 Adjuster Rods

The adjuster-rod system extends the range of the reactor regulating system in the positive-reactivity direction beyond that available from the zone-control system. In the CANDU 6, the adjuster-rod system consists of 21 vertical rods, which can be made of stainless steel or cobalt. The reactor is designed to operate with the adjuster rods fully inserted in the core during normal operation. If more positive reactivity is required than the zone-control system can provide, the adjuster rods are withdrawn in groups (banks) as necessary.

There are two circumstances where the reactivity decreases, relative to the normal steady-state-power condition, to a degree that demands withdrawal of some or all of the adjuster rods to permit the continuing operation of the reactor:

- the unavailability of fuelling machines for a period of more than about one week, after which the reactivity decrease due to incremental irradiation of the fuel typically exceeds the range available in the zone-control system, and

- transient increases in the concentration of ^{135}Xe following a reduction of reactor power.

Since the adjuster rods are normally fully inserted in the core, their position in the reactor, and the distribution of absorbing material amongst the rods, are chosen to flatten

the power distribution, in conjunction with burnup flattening, to achieve the design power shape.

The positions of the 21 adjuster rods in the CANDU 6 are shown in Figures "Top View Showing Adjuster Positions" and "Face View Showing Adjuster Positions". The adjusters are grouped into seven banks, not all composed of an equal number of adjusters. The banks are chosen such that the reactivity worth of any one bank does not exceed the range of the zone-control system. The reactivity worth of the complete system is about 15 mk. The maximum rate of change of reactivity associated with moving one bank of adjusters is < 0.1 mk per second.

The CANDU-6 adjuster system is nominally designed to have sufficient reactivity to compensate for the increase in ^{135}Xe concentration that occurs within approximately 30 minutes following a reactor shutdown. It also provides capability to operate with fuelling machines unavailable for about a month; however, to operate in steady state with adjuster banks out of core, the power level must be reduced to compensate for the radial power peaking caused by adjuster withdrawal.

Note: Some reactors, such as Bruce A reactors, are designed without an adjuster-rod system. In these reactors, extending the reactivity range in the positive direction can be achieved by routinely operating the reactor with a certain amount of poison in the moderator (see next Section), and removing this poison (in whole or in part) by means of ion-exchange columns when positive reactivity is required.

4.4 Moderator Poison

Moderator poison is used to compensate for excess reactivity:

- in the initial core, when all fuel in the core is fresh, and
- during and following reactor shutdown, when the ^{135}Xe concentration has decayed below normal levels.

Boron is used in the initial core, and gadolinium is used following reactor shutdown. The advantage of using gadolinium after shutdown is that its burnout rate during operation at full power following an extended shutdown period is comparable to the xenon growth rate in terms of reactivity, hence the need to remove poison by ion exchange at a fairly rapid and controlled rate is much less demanding. Poison can be added to the moderator for these purposes either automatically or manually.

It should be noted that the moderator-poison-addition system is completely independent of the very-high-speed liquid-poison injection system which is used as a shutdown system (see Section 4.5.2 below). In the regulating-system function, the poison is inserted into the piping used to circulate the moderator, whereas, in the poison-

injection system, the poison is injected through nozzles that are installed horizontally across the core, and a completely independent source of poison is used.

4.5 Special Shutdown Systems

The CANDU 6 reactor is equipped with two physically independent special shutdown systems, SDS-1 and SDS-2. These systems are designed to be both functionally different from each other, and physically separate. These differences are achieved by using vertically oriented mechanical shutoff rods in one system and horizontally oriented liquid-poison-injection nozzles in the second system.

4.5.1 Shutoff Rods (SDS-1)

The shutoff rods are tubes consisting of a cadmium sheet sandwiched between two concentric steel cylinders. The rods are inserted vertically into perforated circular guide tubes which are permanently fixed in the core. The locations of these rods in the CANDU 6 are shown in Figure "Top View Showing Shutoff-Rod Positions". The diameter of the rods is the maximum that can be physically accommodated in the space between the calandria tubes (about 113 mm), when space for the guide tubes and appropriate clearances are considered. The outermost four rods are about 4.4 m long, while the rest are about 5.4 m long. The rods are normally parked fully outside the core and are held in position by an electromagnetic clutch. When a signal for shutdown is received, the clutch releases and the rods fall by gravity into the core, with an initial spring assist.

4.5.2 Liquid-Poison-Injection System (SDS-2)

The alternative way of shutting down the reactor is by high-pressure injection of a solution of gadolinium into the moderator in the calandria. The gadolinium solution is normally held at high pressure in vessels outside of the calandria. Injection is accomplished by opening high-speed valves which are normally closed. When the valves open, the liquid poison is injected into the reactor moderator through six horizontally oriented nozzles that span the core. The nozzles are located in the positions shown in Figure "Positions of Liquid-Poison-Injections Nozzles", and are designed to inject the poison in four different directions in the form of a large number of individual jets. This disperses the poison rapidly throughout a large fraction of the core. The gadolinium solution is held in the retaining pressure vessels at a concentration of typically about 8000 g of gadolinium per Mg of heavy water.

5. Detector Systems

This chapter describes the CANDU detector systems.

5.1 Zone-Control Detectors

To vary the amount of water in the zone-control compartments, the Reactor Regulating System utilizes the readings of detectors associated with the zone controllers. These are fast-response platinum detectors, placed interstitially between fuel channels. There is one detector (plus one spare) for each zone-control compartment. Each detector is located close to the midpoint of the zone-control compartment to which it is associated (see Figure "Positions of Zone-Control Detectors").

To determine changes required in the water fills of the various compartments, the RRS compares the 14 instantaneous detector readings, φ_i, with a set of reference readings, φ_i^{ref}, corresponding to the desired power distribution at full power.

In the bulk-control function, the average of the 14 readings φ_i is used as the indicator of current power, and the water fills in all compartments are uniformly increased or decreased to move the reactor power down or up to the desired power. Bulk control is exercised automatically by the RRS every half second.

In the spatial-control function, the **relative** values of the φ_i are compared to the relative reference values to determine the reactor zones in which the flux is low (i.e., in which power should be raised), and those in which it is high (i.e., in which power should be reduced). The water fills are then moved differentially. In zones where power is to be increased the water level is lowered, and where power is to be decreased the water level is raised. The RRS exercises the spatial-control function automatically every 2 seconds.

Because the zone-control detectors provide essentially "point" readings in the core (the detectors are 3 lattice pitches long but span a very small part of each zone), it is legitimate to ask whether they represent fairly the zones to which they are associated. In order to ensure that the readings used by the RRS do reflect zone-average values, the zone detectors are calibrated every two minutes to zone fluxes obtained by the on-line flux-mapping program (see Section 5.3).

5.2 Neutronic Protection System

CANDU reactors are equipped with protection systems which detect an emergency situation and actuate the safety system(s) discussed in the previous Section. The CANDU-6 neutronic protection systems are described here.

There is a separate neutronic protection system for each of the two shutdown systems. Each protection system is triplicated and consists of out-of-core ion chambers and in-core self-powered detectors. Triplication means that there are three separate "logic" (or "safety") channels for each protection system. These channels are labelled D, E, and F for SDS-1 and G, H, and J for SDS-2. In each protection system, it suffices that two of the three logic channels be "tripped" for the corresponding shutdown system to be actuated.

There are three ion chambers in each protection system, one per logic channel. The ion chambers are located at the outside surface of the calandria (see Figure "Ion-Chamber Locations"). Each ion chamber trips its logic channel when the measured rate of change of the logarithm of the flux ϕ, i.e. the quantity $\frac{d\ell n\phi}{dt}$ exceeds a pre-determined setpoint (e.g. 10% per second, i.e., 0.10 s^{-1}, for SDS-1 in the CANDU 6).

There are also a number of fast-response (platinum or inconel) in-core detectors in each protection system: 34 for SDS-1, located in vertical assemblies (see some detector positions in Figure "SDS-1 In-Core-Detector Locations"), and 24 for SDS-2, located in horizontal assemblies (see Figure "SDS-2 In-Core-Detector Locations"). The detectors are distributed among the various logic channels, so that channels D, E and F contain 11 or 12 detectors each, while channels G, H, and J contain eight each. The detectors trip the logic channels on high neutron flux: when the reading of any one detector reaches a pre-determined setpoint, the logic channel to which it is connected is tripped. Because the in-core detectors are designed to protect the reactor against high local flux, the in-core-detector system is sometimes referred to as the regional-overpower-protection (ROP) system.

The setpoints of the in-core detectors are determined by an extensive analysis of hypothetical loss-of-regulation accidents. The analysis involves the calculation of hundreds of different flux shapes which can apply in the reactor. The ROP setpoints are designed to protect against critical values of channel power being reached; the current criterion for critical channel power is fuel dryout. The setpoints must also ensure the efficacy of the shutdown systems in arresting the power pulse which follows a hypothetical loss-of-coolant accident.

In summary, there are two separate ways in which a protection-system logic channel can be tripped:

- on a high rate of log neutron flux at the corresponding ion chamber, and
- on high neutron flux at any one detector belonging to the logic channel.

A shutdown system is actuated whenever two of the three corresponding logic channels are tripped. The triplicated tripping logic described here is shown schematically in Figure "Triplicated Tripping Logic for SDS-1". The triplication reduces the chance of a spurious trip and allows the system to be tested on-line. It assures an extremely high reliability of shutdown-system actuation under accident conditions.

5.3 Flux-Mapping System

The CANDU 6 is provided with a flux-mapping system to synthesize the 3-dimensional flux distribution in the reactor from in-core detector readings. The system consists of 102 vanadium detectors placed at various positions in the core (see Figure "Flux-Mapping Detectors"). Each detector is one lattice pitch long.

The flux-mapping procedure consists of assuming the 3-dimensional flux distribution can be written as a linear combination of a number of basis functions or flux modes, i.e. that the thermal flux at any point r in the core, $\phi(r)$, can be expressed as a linear combination of flux modes $\psi_n(r)$:

$$\phi(r) = \sum_{n=1}^{m} A_n \psi_n(r) \tag{5.1}$$

where m is the total number of modes used and A_n is the amplitude of the nth mode.

Using this linear expansion, the mode amplitudes A_n are determined by a least-squares fit of the calculated fluxes at the 102 detectors to the measured fluxes. For a detector d at position r_d, the mapped flux is, from Eq. (5.1):

$$\phi(r_d) = \sum_{n=1}^{m} A_n \psi_n(r_d) \tag{5.2}$$

and this can be compared to the measured flux at the detector, F_d.

The flux-mapping procedure determines the amplitudes A_n by minimizing the sum ε of squares of differences between the mapped and measured fluxes, i.e. minimizing

$$\varepsilon = \sum_{d=1}^{102} w_d (\phi_d - F_d)^2 \tag{5.3}$$

where the w_d are chosen weights.

Once the amplitudes have been evaluated, the flux at any point in the reactor can be calculated very easily from Eq. (5.1). Thus, the 3-dimensional flux and power distributions in the core can be derived. The flux-mapping procedure is very quick.

The flux modes $\psi_n(r)$ used in flux mapping consist in the first instance of a number (~15) of pre-calculated harmonics of the neutron diffusion equation. These harmonics represent various possible global perturbations of the flux distribution (see Figure "Harmonic Modes for Flux Mapping").

For situations in which the reactor is operated with mechanical control absorbers in-core or adjusters out-of-core, the harmonics are complemented by a number of "device modes" which represent the more localized perturbations due to device movement.

The flux-mapping procedure is carried out automatically in the on-line computer every two minutes. It provides the mapped values of average zonal flux to the regulating system. These zonal fluxes are used to calibrate the zone-control detectors, to ensure that the readings of the zone detectors faithfully represent the overall flux distribution in the reactor.

Flux mapping can also be done "off line", using recorded flux measurements at the detectors corresponding to any desired time in the reactor history.

6. Computational Scheme for CANDU Neutronics

The basic aim of reactor-physics calculations is to calculate the neutronics of the core, i.e., the distribution of neutron flux and various reaction rates in the reactor, and, most importantly, the power distribution, as a function of space and time.

The computational scheme for CANDU neutronics consists of three stages. Computer programs have been developed to perform the calculations corresponding to each stage.

6.1 Cell Calculation

The first stage involves calculating lattice properties for basic lattice cells. This is done with lattice (or cell) codes. The previous cell code was POWDERPUFS-V, a semi-empirical code based on the results of experiments in heavy-water moderated reactors at CRL. This is being replaced, especially for safety analysis, by WIMS-IST, a modern code with a firm foundation in neutron transport theory. WIMS-IST is now the Industry Standard Tool for lattice calculations.

In a cell code, a detailed calculation of nuclear properties is made for a basic lattice cell, taking into account the fuel, coolant, pressure and calandria tubes, and moderator. However, the idea is to calculate "homogenized-cell" lattice properties (nuclear cross sections), averaged over the cell. Such homogenized properties, obtained for all cells in the reactor, will be used in the finite-reactor model to calculate the power distribution in the core. Lattice calculations are done assuming "mirror" (reflective) properties at the cell boundaries. This is equivalent to doing calculations for an infinite lattice of identical cells.

Because the neutron spectrum (energy distribution) is very well thermalized in the CANDU heavy-water moderator (~95% of neutrons in the moderator are thermal), two neutron energy groups - thermal and "fast" (or "slowing down") - are certainly sufficient for the homogenized properties.

The 2-group lattice cross sections are:

- Σ_{a1}, neutron absorption cross section in fast group
- Σ_{a2} , neutron absorption cross section in thermal group
- $\Sigma_{1\to2}$, neutron moderation cross section (down-scattering from fast to thermal group)
- $\Sigma_{2\to1}$, neutron up-scattering cross section
- $\nu\Sigma_{f2}$, neutron production cross section in thermal group
- $\nu\Sigma_{f1}$, neutron production cross section in fast group
- Σ_{tr1}, fast-neutron transport cross section

- Σ_{tr2}, thermal-neutron transport cross section

In addition the "H-factors", ratios of bundle power to flux, are also needed. The lattice code must also provide the heavy-water-reflector properties.

[Note that the older cell code, POWDERPUFS-V, makes certain approximations such as neglecting up-scattering and "lumping" fast fissions into the thermal-fission cross section.]

The cell code must be able to derive these lattice properties as a function of fuel irradiation (or burnup) and for different types of conditions in the cell, such as:

- various fuel-bundle designs (28-el, 37-el, CANFLEX, …)
- different fuel, coolant, and/or moderator temperatures
- varying coolant density
- different levels of fuel power
- various concentrations of moderator poison
- etc.

6.2 Supercell Calculation

Various reactivity devices are, or can be, inserted at a number of locations in the CANDU. These devices perturb the basic-lattice properties in their vicinity. The aim of "supercell" calculations is to determine "incremental" cross sections (e.g., incremental absorption cross sections $\Delta\Sigma_{a1}$, $\Delta\Sigma_{a1}$), to be added to the basic-cell properties to account for the effect of the devices. The incremental cross sections are obtained by calculating homogenized properties in the supercell (see Figure "Typical CANDU Supercell Model") with and then without the device, and then subtracting. Supercell calculations used to be performed with MULTICELL, which was based on POWDERPUFS-V lattice properties. This has been supplanted by DRAGON-IST, a 3-d-transport-theory code. DRAGON-IST is now the Industry Standard Tool for reactivity-device calculations.

6.3 Finite-Core Calculation

Putting together the basic-lattice cross sections and reactivity-device incremental cross sections, a model of the whole reactor core is set up and used to calculate the 3-dimensional flux and power distribution. This is called a finite-core calculation.

The Industry Standard Tool for finite-core calculations is RFSP-IST (Reactor Fuelling Simulation Program), specifically designed for CANDU reactors. In its most frequent application, RFSP-IST calculates the steady-state 3-dimensional flux and power distributions in the reactor by solving the time-independent diffusion equation in two energy groups, which is written as the following eigenvalue problem:

$$-\vec{\nabla}\cdot D_1(\vec{r})\,\vec{\nabla}\phi_1(\vec{r}) + (\Sigma_{a1}(\vec{r}) + \Sigma_{1\to 2}(\vec{r}))\,\phi_1(\vec{r}) - \left(\Sigma_{2\to 1}(\vec{r}) + \frac{\nu\Sigma_{f2}(\vec{r})}{k_{eff}}\right)\phi_2(\vec{r}) - \frac{\nu\Sigma_{f1}(\vec{r})}{k_{eff}}\phi_1(\vec{r}) = 0 \qquad (4.9a)$$

$$-\vec{\nabla}\cdot D_2(\vec{r})\,\vec{\nabla}\phi_2(\vec{r}) + (\Sigma_{a2}(\vec{r}) + \Sigma_{2\to 1}(\vec{r}))\,\phi_2(\vec{r}) - \Sigma_{1\to 2}(\vec{r})\,\phi_1(\vec{r}) = 0 \qquad (4.9b)$$

where

- $\vec{r}$ is the position in core,
- the first term in the equation for each energy group calculates the group leakage from one region $\vec{r}$ of the core to another,
- and k_{eff} is the unknown reactor multiplication constant, to be determined at the same time as the flux distribution.

This time-independent equation is an eigenvalue problem because a steady-state (critical) reactor configuration does not have a solution for just **any** given combination of properties. The k_{eff} is an adjustment factor, only certain values of which result in a steady state - the largest is the reactor multiplication constant and is a measure of how far from criticality the given reactor configuration is.

RFSP actually solves the finite-difference form of the diffusion equation. A typical finite-difference reactor model is shown in Figures "Typical RFSP-IST Reactor Model (Face View)" and "Typical RFSP-IST Reactor Model (Top View)".

Major applications of RFSP are in:

- core-design calculations and analyses, including fuel-management design calculations, and simulations of reactor power histories
- core-follow calculations at CANDU sites, to track the actual reactor operating history, with burnup steps and channel refuellings.

Additional capabilities of the program include, among others:

- the calculation of flux distributions for various reactor configurations
- the simulation of $^{135}Xe/^{135}I$ transients
- the capability for simulating (quasi-statically) bulk control and spatial control
- the calculation of harmonic flux shapes for use in flux mapping,
- the calculation of the reactivity increase expected on refuelling of individual fuel channels
- the capability for solving neutron-kinetics problems by the Improved Quasi-Static (IQS) method. RFSP can therefore be used to analyze fast transients, such as those following hypothetical large-loss-of-coolant accidents (LOCA), and to simulate the performance of the shutdown systems.

6.4 CANDU Positive Void Reactivity

Because of the particular relevance of void reactivity in CANDU, this will be described briefly here.

Coolant voiding in CANDU introduces positive reactivity and promotes a power rise. The root cause lies in the fact that CANDU is a pressure-tube reactor, with the coolant separate from the moderator.

In light-water reactors (LWR), one liquid serves as both coolant and moderator, and a loss of coolant is also a loss of moderator, leading to less self-sustainable chain reaction, i.e., a *decrease* in reactivity. In CANDU, however, the loss of coolant does not imply a significant reduction in moderation, but does result in changes in the neutron spectrum (distribution of neutron energy).

These spectrum changes go in the direction of increasing reactivity. A simple explanation for the major causes of this effect is as follows:

- When coolant is present, some of the neutrons from fission are immediately slowed down in the fuel cluster itself; when the coolant is voided, fewer neutrons are slowed down in the fuel cluster, allowing more fast neutrons to induce fast-group fission
- Also, since fewer neutrons are slowed down in the fuel cluster to energies in the resonance range, more neutrons **escape** resonance absorption before entering the moderator
- Most neutrons which re-enter the fuel cluster from the moderator are thermal neutrons. When the hot coolant is present, some of these are promoted to higher energies by collision with coolant, and some are absorbed in fuel resonances. When the coolant is absent, this effect is absent, and there is again increased resonance absorption **escape**.

For irradiated fuel, in which plutonium is present, the change in neutron spectrum gives also a *negative* component in the reactivity change. This is due to a reduction in absorptions in the 0.3-eV **fission resonance of plutonium** (see Figure "Low-Lying (0.3-eV) Fission Resonance in Plutonium-239") and therefore a reduction in fissions from that source. However, the **net** reactivity change on coolant voiding is still positive (but smaller than for fresh fuel).

Full-core void reactivity ranges from 10 to 15 mk or more, depending on the fuel burnup and other parameters. Of course, it is not physically possible to lose all coolant from the core instantaneously. However a Large Loss of Coolant is a hypothetical accident which must be analyzed.

A large loss of coolant is in fact the accident which presents the greatest challenge to CANDU shutdown systems in terms of the rate of positive reactivity insertion. A large LOCA is caused by the rupture of a large pipe such as a Reactor Inlet Header (RIH), Reactor Outlet Header (ROH), or Pump-Suction pipe (see Figure "Examples of Break Locations Giving Rise to a Large LOCA"). In the CANDU 6, such a rupture can inject 4 to 5 mk of positive reactivity in the first second after the break, which is beyond the capability of the Reactor Regulating System to control. Thus a Large LOCA can lead to a sudden power surge (power pulse), which must be terminated by a shutdown system. The manner in which the shutdown systems act (separately) to terminate the power excursion must therefore be carefully studied. This is done with neutron kinetics codes. As described above, RFSP-IST itself has a kinetics capability.

6.5 The Time-Average Model

In the time-average model of the reactor, the lattice cross sections are averaged over the residence (dwell) time of the fuel at each point (fuel-bundle position) in the core. This allows the effect of the actual refuelling scheme used (e.g. 8-bundle shift, 4-bundle shift, etc.) to be captured. Calculations are performed in the *TIME-AVER module of RFSP-IST.

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} is the value of fuel irradiation when the fuel enters that position in core and ω_{out} is the fuel irradiation when the fuel leaves that position. For example, the time-average thermal neutron absorption cross section at some core position r, $\Sigma_{a2}^{t.a.}(r)$, is

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

The time-average calculation is a self-consistency problem, because consistency must be achieved between the flux, the channel dwell times (interval between refuellings), the individual-bundle irradiation ranges $[\omega_{in}, \omega_{out}]$, and the lattice properties. An iterative scheme of solution is employed until all quantities converge.

Typically, in the time-average model, the core is subdivided into many irradiation regions. An average fuel exit irradiation is selected for each region, and the values are designed to achieve an acceptable degree of flattening of the flux shape. Again, the exit irradiation values may have to be determined by several trials. Typical values are shown in Figure "Multiple-Region Time-Average Model for CANDU 6". Figure "Calculational Scheme For Time-Average Calculation" shows the iterative scheme for the time-average calculation.

The time-average model is useful at the design stage, to determine the reference three-dimensional power distribution, the expected refuelling frequency of each channel (or its inverse, the channel dwell time), and the expected value of discharge burnup for the various channels. See typical power results in Figure "Channel-Power Distribution from a CANDU 6 Time-Average Calculation".

Figure "Channel Dwell Times from the Same CANDU 6 Time-Average Calculation" 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 typically between 150 and 160 full-power days (FPD). In the outer core, the dwell times present a large variation, from about 135 FPD for channels just outside the inner core (where the flux is still high but the exit irradiation is, by design, lower than in the inner core) to almost 300 FPD for some channels at the outermost periphery of the core.

It is instructive to look at a typical neutron balance in the CANDU-6 equilibrium core. This is displayed in Figure "Typical Neutron Balance in CANDU 6 (Time-Average Core", which shows that more than 45% of fission neutrons originate from fissions in plutonium. Thus, plutonium contributes approximately half the fission energy produced in a CANDU reactor. (Actually, in fuel near the exit burnup, plutonium contributes about three-quarters of the fission energy generated.) 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 a loss of almost 90 milli-k. Parasitic absorption in non-fuel components of the lattice represents a 63-milli-k loss.

7. Fuel Management

In this chapter we will discuss the effect of fuel burnup (or irradiation) on the lattice properties, and issues relating to fuel management.

7.1 Infinite-Lattice Multiplication Constant

The infinite-lattice multiplication constant k_∞ is a measure of the multiplicative properties of the lattice in the absence of leakage from the lattice cell. The k_∞ is (or can be) calculated from the basic-lattice cross sections provided by the cell code, and applies to the "ideal" situation of an infinite array of identical cells.

Fig. "Infinite-Lattice Multiplication Constant for Standard CANDU-6 Lattice Fuelled with Natural Uranium" shows that the lattice is ~ 80 milli-k supercritical for fresh fuel (i.e., at zero irradiation). The reactivity **increases** at first with increasing irradiation, reaching a maximum at approximately 0.4-0.5 n/kb, a phenomenon due to the production of plutonium from neutron absorption in ^{238}U. This reactivity maximum is consequently known as the plutonium peak. Beyond the plutonium peak, the reactivity starts to decrease with increasing irradiation, on account of the continuing depletion of ^{235}U and the increasing fission-product load, and the lattice reaches zero excess reactivity at an irradiation of about 1.6-1.8 n/kb. This marks a natural point at which the fuel can be targeted for removal from the core, since at higher irradiations the lattice becomes increasingly subcritical, i.e., an increasing net absorber of neutrons. **Thus, channels containing fuel approaching or exceeding these irradiation values become good candidates for refuelling.**

It is instructive to examine also the infinite-lattice multiplication constant for the depleted-uranium lattice. This is shown in Figure "Infinite-Lattice Multiplication Constant for Standard CANDU-6 Lattice Fuelled with Depleted Uranium" (with an initial fissile content of 0.52 atom %, as opposed to 0.72 atom % for natural uranium). Note that the plutonium peak is even more pronounced for depleted uranium, a result which is easily explained by the fact that the role of ^{238}U conversion to plutonium is relatively greater when the ^{235}U content is smaller. Note also, however, that the depleted-uranium lattice is subcritical at all irradiations, i.e. is always a neutron absorber. This explains the use of depleted fuel to reduce excess reactivity, and also flatten the flux distribution, in the initial core. Depleted fuel is also occasionally used to reduce the power ripple on refuelling.

7.2 Radial Flattening of the Power Distribution

The flux distribution in the reactor depends on the reactor size and geometry and on the distribution of irradiation. Fuel with a high irradiation has low reactivity, and

depresses flux in its vicinity. Similarly, the neutron flux tends (everything else being equal) to be high in regions where the fuel has low irradiation. This fact can be used to "shape" the flux (and power) distributions in the equilibrium core.

Radial flux (and power) flattening can be achieved by **differential fuelling**, i.e. taking the fuel to a higher burnup in inner core regions than in outer core regions (cf. the Figure "Multiple-Region Time-Average Model for CANDU 6"). This can be done by judicious adjustment of the relative refuelling rates in the different core regions. In this way the flux and power in the outer region can be increased, resulting in a greater number of channels having power close to the maximum value. Thus, a higher total reactor power can be obtained (for a given number of fuel channels) without exceeding the limit on individual channel power. This reduces the capital cost of the reactor per installed kW.

The radial flattening is quantitatively measured by the radial form factor:

$$\text{Radial form factor} = \frac{\text{Average channel power}}{\text{Maximum channel power}}$$

Radial flattening is further assisted by the use of adjuster rods, whose main purpose is in xenon override. Adjuster rods also provide axial power flattening.

Note that while flattening of the power distribution reduce the reactor capital cost, by reducing the number of channels required to produce a given total power, it does tend to increase the neutron leakage, which is proportional to the flux gradient at the edge of the core. This loss of neutrons does have a consequent increase in fuelling cost.

7.3 The Equilibrium (Time-Average) Core

A consequence of the on-power refuelling in CANDU is that, with refuelling operations taking place essentially every day, the equilibrium core contains fuel at a range of burnups, from 0 to some average exit-burnup value. The **average** in-core irradiation is fairly constant over time, at about half the exit value. The long-term global flux and power distributions in the equilibrium core can be considered as consisting of a constant, "time-average", shape, about which the refuelling of individual channels leads to local "refuelling ripples". These ripples are due to the various instantaneous values of fuel burnup in the different channels, which are the result on any given day of the specific sequence of channels refuelled in the previous days, weeks and months.

Calculations for the time-average model are performed in the *TIME-AVER module of RFSP. The mathematical framework of this module has already been described in Section 6.5. The outputs of the *TIME-AVER module are the time-average flux and power distributions in the core, the individual-channel refuelling rates, and the

lattice properties at each bundle position averaged over the fuel residence time at that location (consistent with the axial refuelling scheme, e.g. 8-bundle shift).

7.4 Ongoing Reactor Operation with Channel Refuellings

After the initial period following first reactor startup, on-power refuelling is the primary means of maintaining a CANDU reactor critical. Thus, a number of channels are refuelled every day, **on the average**. Note that refuelling is not necessarily done **every** calendar day; some stations prefer to concentrate all refuelling operations to 2 or 3 days within each week.

Replacing irradiated fuel with fresh fuel has immediate consequences on the local power distribution and on the subsequent period of operation of the reactor. These are discussed in the following subsections.

7.4.1 The Channel-Power Cycle

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

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

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

- Following the plutonium peak, the reactivity of the refuelled channel starts to decrease, and its power drops slowly. Approximately half-way through the 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.

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. Figure "Schematic of Maximum Channel Power versus Time" illustrates the difference between maximum time-average channel power, average maximum instantaneous channel power, and absolute maximum channel power.

7.4.2 Channel-Power Peaking Factor

At any given time, there are several channels in the core which are at or near the maximum power in their cycle. Therefore, the maximum instantaneous channel power is always higher than the maximum time-average channel power, as was evident from the Figure "Schematic of Maximum Channel Power versus Time".

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:

$$CPPF = \underset{m}{Max}\left[\frac{CP_{ins\,\tan\tan\,eous}(m)}{CP_{time-average}(m)}\right] \qquad (7.1)$$

where m runs over all channels in the core, or at least over all channels except perhaps the channels with the very lowest power, i.e., except the last two outermost rings of channels.

The CPPF value varies from day to day, as the various channels which have fairly recently been refuelled go through their cycle. However, the average CPPF value must obviously depend on the axial refuelling scheme used. The greater the number of bundles replaced at each operation, the greater the reactivity increment, and therefore the greater the refuelling ripple (and therefore the CPPF). When the 8-bundle-shift refuelling is used, a typical value for the CPPF is in the range 1.08-1.10. With a 4-bundle-shift scheme, the typical CPPF is likely to be 1.04-1.05.

The exact value of the CPPF is extremely important because it is used to calibrate the in-core ROP detectors. The hundreds of flux shapes that are used in the ROP safety analysis (to determine the detector positions and setpoints) are all calculated in the time-

average model, assuming many different core configurations. But because the real instantaneous channel powers are higher than the time-average powers used in the ROP analysis, channels would reach their "critical channel power" (power at which there is fuel dryout) earlier than in the time-average model. To take this into consideration and ensure proper safety coverage in the instantaneous power shape, the in-core ROP detectors are calibrated each day to the instantaneous value of CPPF.

In order to maximize the margin to trip, it is obviously important that the CPPF be kept as low as possible. This is why a careful selection of channels to be refuelled needs to be made always. A way in which CPPF can be kept low by design is by using, say, 4-bundle-shift refuelling instead of 8-bundle-shift refuelling, or using a mixed 4- and 8-bundle-shift scheme, where the 4-bundle shifting is done in the inner core (high-power region).

Another way in which poor refuelling strategy could impact on reactor operation is as follows. Concentrated refuelling in the vicinity of an ROP detector will increase its reading, even though this may not increase the CPPF in the core. The high detector reading may lead either to spurious trips or to power deratings (to restore operating margin), both of which lead to loss of power production.

Determining the daily CPPF value, and ensuring detectors are calibrated to the correct value, are on-going duties of the fuelling engineer or reactor physicist at a CANDU nuclear generating station.

7.4.3 Selecting Channels for Refuelling

One of the main functions of the fuel engineer (or site reactor physicist) is to establish a list of channels to be refuelled during the following period (few days) of operation. To achieve this, the current status of the reactor core is determined from computer simulations of reactor operation, the on-line flux mapping system, the ROP and RRS in-core detectors, and zone-control-compartment water fills. The computer simulations of reactor operation provide the instantaneous 3-dimensional flux, power and burnup distributions.

Normally, channel selection will begin with **eliminating** channels which are poor candidates for refuelling. With experience, a fuelling engineer will develop a personal set of rules for eliminating channels. A **typical** (but by no means unique) set of rules may eliminate

- high-power channels (those with current power within 10% of the maximum licensed channel power) and their closest neighbours
- channels refuelled recently, say < 10 FPD prior, and their closest neighbours
- channels with a high value of peaking factor (> , say, 1.07), and their closest neighbours
- low-burnup channels (with exit burnup < , 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 last refuelled approximately one dwell time (from the time-average calculation) prior
- channels with high current 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%).

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.

7.5 Initial Fuel Load and Transient to Onset of Refuelling

Let us now look at the period which marks the beginning of reactor operation.

In the initial core, fresh fuel is present throughout the core. There is no differential burnup which can assist in flattening the power distribution. Consequently, the power of the central core region would be unacceptably high if no alternate means of flattening the radial power distribution were provided. However, an alternate means **is** readily available: depleted fuel. As we have seen earlier, this depleted fuel is a net absorber of neutrons.

In the initial core of the CANDU 6 (i.e., the initial fuel load), two depleted-fuel bundles (of 0.52 atom percent ^{235}U content) are placed in each of the central 80 fuel channels. This is shown in Figure "Channels with Depleted Fuel in Initial Core of CANDU 6". The bundles are located in positions 8 and 9, where the numbering is from the channel refuelling end. In these axial positions, the depleted-fuel bundles are removed from the core in the first refuelling visit of each of these channels.

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 core reactivity starts at approximately 16 milli-k at full power on FPD 0, and then varies with time as shown in Figure "Excess Core Reactivity in Initial Period of Reactor Operation". Because all the fuel goes through its plutonium peak at about the same time, the excess reactivity initially increases, from about 16 mk to a maximum of about 23 mk between FPD 40 and FPD 50 This excess reactivity is compensated by soluble boron in the moderator. The boron coefficient of reactivity is about -8 milli-k per ppm of boron. Thus the boron concentration (at full power) is initially approximately 2 ppm, rising to about 3 ppm at the plutonium peak. Following the plutonium peak, boron must be removed (by ion exchange) as the excess reactivity drops gradually to zero at about FPD 120.

During this entire first period in the reactor life, refuelling is not necessary since there is already excess reactivity. Actually, refuelling is started about 10 or 20 FPD before the excess reactivity reaches 0, i.e. around FPD 100, because the refuelling rate would be too great if one waited until the last possible moment to start.

After refuelling begins, the rate of refuelling rapidly approaches its equilibrium value (approximately 16 bundles per FPD for the CANDU 6). Over short periods, there may however be a considerable variation from this average rate.

7.6 Consequences of Fuelling-Machine Unavailability

If refuelling were to stop, core reactivity would continuously decrease. The rate of reactivity decay is about 0.4 mk/FPD in the CANDU-6 core. The reactor regulating system (RRS) would of course attempt to maintain criticality.

The first action that the reactor regulating system (RRS) would take to maintain criticality is to lower the level of water in the liquid zone-control compartments. Eventually, the water will be drained to the lower limit of the control range.

Since the desirable operating range of the zone controllers is between 20% and 70% (i.e., a range of 50%), and since the full reactivity range of the zone controllers (from 100% down to 0%) provides about 7 milli-k of reactivity, the number of days which can be "survived" without refuelling is typically about 3.5 mk/(0.4 mk/FPD), i.e., about 8 FPD.

The operator would also ensure that any poison which might exist in the moderator at the time would be removed. Every ppm of boron is worth about 8 mk, however the operating license usually limits the amount of boron in the core in full-power operation to about 0.625 ppm (5 milli-k), so this represents at most about 12 FPD without refuelling.

Continued lack of refuelling would lead to withdrawal of the adjuster rods in their normal sequence. This would permit operation to continue for several weeks. However, as the adjuster rods are withdrawn, the reactor power must be gradually reduced because of changes in the power distribution associated with spatial changes in the distribution of absorption cross section. In effect, withdrawal of the adjusters results in a radially "peaked" power distribution, i.e., higher channel and bundle powers at the centre of the core, which forces a power derating in order to remain in compliance with the licensed channel and bundle powers (7.3 MW and 935 kW respectively). The amount of derating necessary increases with the number of adjusters withdrawn.

7.7 Core-Follow Calculations with RFSP-IST

The main application of RFSP-IST at CANDU sites is in tracking the reactor's operating history. This function is performed with the *SIMULATE module of RFSP.

The core history is tracked by a series of instantaneous snapshots, which can be calculated at any desired frequency. Steps of 2-3 FPD are typically convenient for the site physicist. The code advances the in-core irradiation and burnup distributions at each step, in accordance with the time interval. Individual channel refuellings within a time step are taken into account at the actual time at which they occur.

At each code execution, the zone-control-compartment fills corresponding to the time of the snapshot are input to the code, together with the concentration of moderator poison and any other device movement, so that the instantaneous reactor configuration is captured. As an option, the spatial distribution of ^{135}Xe and its effect on the lattice properties can be modelled in the calculation (see chapter 8); this has an effect on the calculated flux distribution. Bulk and spatial control can also be modelled.

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

8. Effects of ^{135}Xe

8.1 The Xe-I Kinetics

The xenon isotope ^{135}Xe plays an 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 ^{135}Xe 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.

Figure "^{135}Xe/^{135}I Kinetics" shows that ^{135}Xe is produced to some degree directly in fission, but mostly as the result of the beta decay of its precursor ^{135}I (which has a half-life of 6.585 hours). ^{135}Xe is destroyed in two ways:

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

The ^{135}Xe/^{135}I kinetics can be represented by two coupled linear differential equations, which we will not write here. However, the consequences of these equations will be described.

The large absorption cross section of ^{135}Xe plays a significant role in the overall neutron balance in the reactor and thus directly affects the system reactivity, both in steady state and in transients. The ^{135}Xe/^{135}I kinetics also influences the spatial power distribution in the reactor.

8.2 Steady-State Xenon Load

The Xe-I kinetics leads to the saturating character of ^{135}Xe: in steady state, the ^{135}I concentration in a fuel bundle increases without limit with the power level, however the ^{135}Xe approaches asymptotically a maximum (saturation) value as the bundle power increases.

The limiting ^{135}Xe absorption rate at very high flux levels leads to a maximum reactivity of ~ -30 mk. In CANDU the equilibrium xenon load is approximately -28 mk. The flux level at full power in CANDU is such that the ^{135}Xe concentration is about 95%

saturated, i.e., the average ^{135}Xe concentration is equal to about 95% of the value in an infinite flux.

8.3 Effects of Xenon on Power Distribution

Xenon plays a role in the 3-d power distribution in the core. Because the steady-state ^{135}Xe concentration depends on the flux, 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 ^{135}Xe concentration can be of the order of 5%, and should be taken into account when accurate results are desired.

8.4 Effect of Power Changes on Xenon Concentration

Due to ^{135}Xe burnout in the neutron flux, the variation of the ^{135}Xe concentration with flux is non-linear. The ^{135}Xe 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 ^{135}Xe concentration increases at first. This is due to the fact that ^{135}Xe is still being produced by the decay of ^{135}I, 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 ^{135}I decay rate decreases sufficiently (due to the lower fission rate) that the rate of ^{135}Xe production drops below the rate of ^{135}Xe decay (and burnout). At this time, then, the ^{135}Xe concentration reaches a peak value and starts to decrease towards a new (lower) steady-state level.

Conversely, when the power is increased from a steady level, the ^{135}Xe 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 ^{135}Xe following step reductions in power are shown in Figure "Xenon Reactivity Transients Following Setback to Various Power Levels". Very similar variations, but in the opposite direction, ensue upon step increases in power. The quantitative effects will be different at different points in the core, due to the initial non-uniform distribution of ^{135}Xe. Thus, for an accurate assessment of xenon transients on the spatial distribution of power, a point-kinetics treatment is generally inadequate and calculations in three dimensions will be required.

8.5 Xenon Transient Following a Shutdown

Following a reactor shutdown, the burnout of ^{135}Xe stops, whereas the production by means of ^{135}I decay continues for several hours. The net result is that there is an initial increase in ^{135}Xe concentration and a decrease in core reactivity. If the reactor is required to be started up shortly after shutdown, extra positive reactivity must be supplied. The ^{135}Xe growth and decay following a shutdown for a typical CANDU is shown in Figure "Xenon Transient Following Reactor Shutdown".

It can be seen from this figure that, at about 10 hours after shutdown, the reactivity worth of ^{135}Xe increases to several times its equilibrium full-power value. At ~35-40 hours the ^{135}Xe 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 some 40 hours, when the reactor again reaches criticality.

To achieve xenon "override" and permit power recovery following a shutdown (or reduction in reactor power), the CANDU-6 reactor has a set of adjuster rods which are in core during normal operation but which can be withdrawn to provide a source of positive reactivity. It is not possible to provide "complete" xenon 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 (i.e., a reduction to zero power).

9. Summary

CANDU reactor physics has both design and operations aspects.

The design component can be (very briefly) summarized as consisting of calculating the reactivity and flux and power distributions for any number of assumed core configurations and scenarios, including the time-average shape and perturbations.

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 design or site reactor physicist is always interesting and stimulating; it never gets boring.