LIGHT WATER REACTOR CONTROL SYSTEMS

The control and operation of nuclear reactors that use uranium for fuel and light water for both moderator and coolant is discussed in this article. The starting point is the physics of the neutron life cycle, which determines the dynamic behavior of a nuclear reactor. Reactor startup, low-power operation while critical, and operation in the presence of feedback effects are then enumerated. Finally, specific aspects of the operation of both pressurized and boiling water reactors (*PWR*s and *BWR*s) are reviewed. The material in the first two sections of this article is extendible to reactors that utilize other fuels, moderators, and coolants.

NEUTRON LIFE CYCLE

An isotope is defined as being fissile if its nucleus will split in two or, in other words, undergo nuclear fission, if struck by a neutron that has an extremely small kinetic energy, typically 0.025 eV, which corresponds to a speed of 2200 m/s. There are only four fissile isotopes (^{233}U, ^{235}U, ^{239}Pu, and ^{241}Pu). Of these, only ^{235}U, which constitutes 0.71% of all uranium atoms, is naturally occurring. Because light water absorbs neutrons, the fuel utilized in light water reactors (*LWR*s) must be enriched in the ^{235}U isotope in order to permit the establishment of a self-sustaining neutron chain reaction. If the LWR is large and therefore has a small surface-to-volume ratio, few neutrons will diffuse out of the fueled region, or *core.* The neutron leakage is said to be small and a low degree of enrichment is possible. This is the case with PWRs and BWRs, which have fuel enriched to 2% to 3%. In contrast, if the LWR is small and has a large surface-to-volume ratio, neutron leakage will be large and high enrichments will be needed. This is the case for research and test reactors, which may have enrichments of 20% or more.

Figures 1 and 2 show the cross sections of ^{235}U and ^{238}U for neutrons over the energy range 1 eV to 10,000 eV. Cross sections are traditionally given in *barns,* with 1 barn equaling 10^{-24} cm^2. Cross sections may be defined for any type of interaction (scattering, absorption, fission, total) between a neutron and the nucleus in question. The cross section specifies the probability that the interaction will occur. For the ^{235}U fission cross section, the important features are that the probability of an interaction with neutrons at high energy (keV to MeV range) is quite small, a few barns, and that the probability of an interaction at low energy ($\approx$0.025 eV) is high. For the ^{238}U total cross section, the important features are that some neutron absorption will occur at high energies and that there are several sharply defined resonances of large magnitude, with the first at 6.67 eV. The first of these features may result in fission if the incident neutron has an energy in excess of 1.5 MeV. The second results in neutron capture.<figureAnchor figures="W5206-fig-0001 W5206-fig-0002"/>

The fissioning of ^{235}U produces an average of 2.5 neutrons. These are emitted over a certain energy distribution, the most probable energy being slightly below 1 MeV and the maximum at about 10 MeV. Neutrons in this energy range are classified as *fast* because of their high kinetic energies. The probability of these neutrons interacting with ^{235}U to cause fission is remote, given the small ^{235}U cross section in this energy range. Examination of Fig. 1 suggests that a self-sustaining neutron chain reaction is possible if the fast neutrons are slowed down, or *thermalized,* in order to take advantage of ^{235}U's large fission cross section for low-energy, or *thermal,* neutrons. The slowing-down process is called *neutron moderation.* Moderation is optimized by causing the fast neutrons to collide with something of similar mass. A head-on collision between two objects of the same mass will result in a complete transfer of energy. Thus, for nuclear reactors, the moderators of choice are hydrogen-rich substances such as light water. It is important to recognize that the slowing down occurs in discrete steps. Neutrons do not lose energy continuously. Rather they lose it every time they undergo a collision.

The physics of the neutron life cycle that is needed to support a self-sustaining chain reaction can now be understood. This cycle is shown in Fig. 3. Assume that a certain number (n) of fast neutrons have been produced by the thermal fission of ^{235}U. What can happen to these neutrons? The fuel consists of the isotopes ^{235}U and ^{238}U. So, to understand the fate of the neutrons, consider what occurs as the fast neutrons move from high to low energy while subject to the interaction probabilities shown in Figs. 1 and 2. The possibilities are:

1. Some fast neutrons may be absorbed by ^{238}U and cause fission. This process is called *fast fission* and is quantitatively represented by the fast fission factor ε which is the ratio of the total number of neutrons produced from fast and thermal fission to the number produced from thermal fission alone. Values of ε depend on the enrichment, with ε approaching unity (from above) for fully enriched fuel.
2. Fast neutrons may also escape from the reactor core. This is referred to as *fast-neutron leakage.* Such neutrons are lost from the life cycle and will not cause a fission reaction. The quantity L_f is defined as the ratio of the total number of fast neutrons escaping leakage to the total number produced from fast and thermal fission. L_f is often called the non-leakage probability and therefore the fraction that does leak out is given by $(1 - L_f)$.
3. The remaining neutrons collide with moderator nuclei and, as a result, lose energy in a discontinuous manner. If, after a collision, a neutron's energy corresponds to one of the ^{238}U resonances, it will be absorbed in ^{238}U. Such interactions do not continue the neutron chain reaction, because only neutrons with kinetic energies in excess of 1.5 MeV can cause ^{238}U to fission. Hence, during the slowing-down process, any neutrons that happen to have energies corresponding to a ^{238}U resonance are, like those that leaked out of the reactor core, lost from the life cycle. The quantity p is defined as the resonance escape probability. It is the ratio of the total number of thermalized neutrons to the total number of fast neutrons that escaped leakage. The quantity p is a function of the enrichment, and it approaches unity (from below)

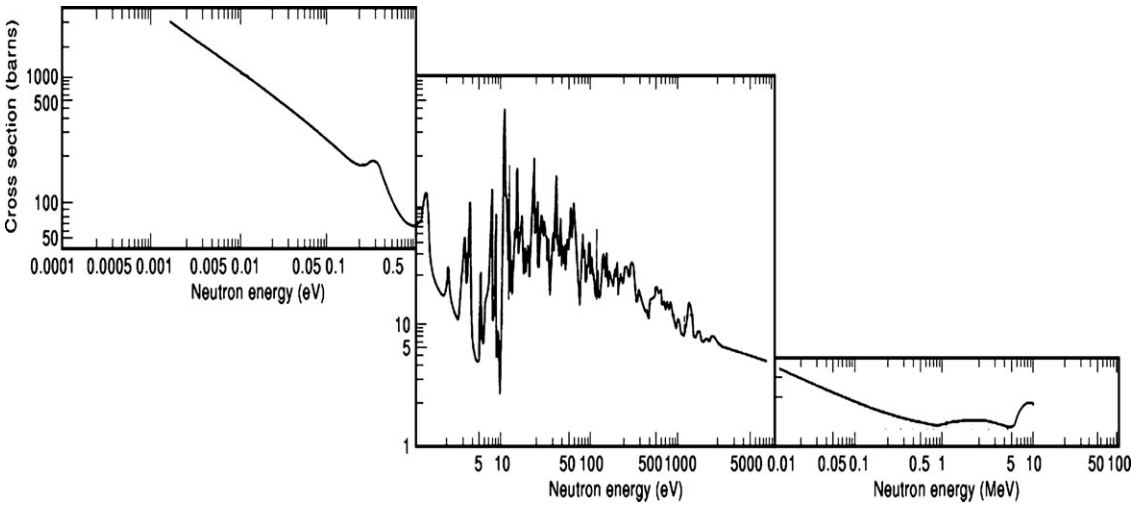

Figure 1. Microscopic cross section of ^{235}U. Source: BNL-325.

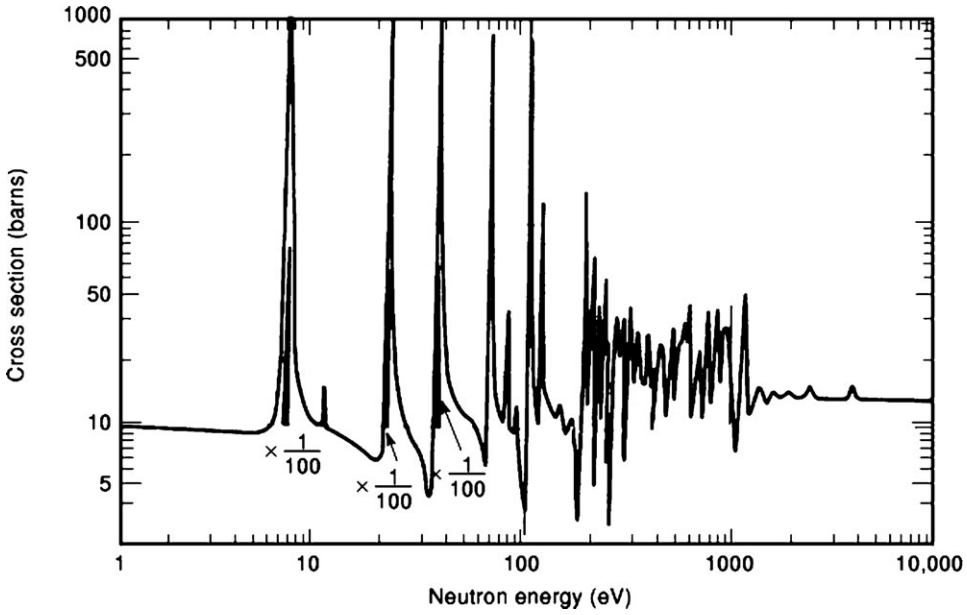

Figure 2. Microscopic cross section of ^{238}U. Source: BNL-325.

for fully enriched fuel.

4. The next fate that might befall a neutron is that it leaks out of the core while at thermal energies. This is referred to as thermal neutron leakage. The quantity L_t is defined as the ratio of the total number of thermal neutrons escaping leakage to the total number of thermalized neutrons.
5. The neutrons are now at thermal energies. Some will be absorbed in the ^{235}U. Others will be absorbed in ^{238}U, core structural materials, control devices, or even the coolant. The thermal utilization factor, f, is defined as the ratio of the thermal neutrons absorbed in the fuel to the total number of thermal neutrons escaping leakage. This parameter f is the quantity that reactor operators alter when they manipulate control devices.
6. Of those neutrons absorbed in ^{235}U, only about 80% will cause fission. The rest result in a transformation to ^{236}U. The ratio of the fission and the total absorption cross sections gives the fraction that contribute to the neutron life cycle. For every fission, a certain number of neutrons are produced. This quantity is denoted by the symbol ν. The thermal reproductive factor $\nu\Sigma_f/\Sigma_a$ is defined where Σ_f and Σ_a are the macroscopic cross sections for fission and absorption respectively. (*Note:* The cross sections shown in Figs. 1 and 2 are the microscopic ones, denoted by the symbol σ. Macroscopic cross sections are the product σN, where N is the number density (nuclei/cm^3) of the material, ^{235}U in this instance.) The thermal reproductive factor η is defined as the ratio of the number of fast neutrons produced from thermal fission to the number of thermal neutrons absorbed in fuel.

The product of the above six quantities is the core multiplication factor K. Thus,

$$K = \epsilon L_f p L_t f \eta \quad (1)$$

Equation (1) is commonly referred to as the *six-factor formula.* If the core is surrounded by a reflector of such size that the probability of non-leakage becomes unity ($L_f = L_t = 1.0$), then Eq. (1) reduces to:

$$K_\infty = \epsilon p f \eta \quad (2)$$

which is called the *four-factor formula.* The symbol K_∞ implies an infinite reflector. The quantity K may be interpreted physically as

$$K = \frac{\text{neutrons produced from fission}}{\text{neutrons absorbed and neutrons lost to leakage}} \quad (3)$$

It is useful, although admittedly artificial, to consider neutrons as moving through the life cycle shown in Fig. 3 in successive groups, or *generations.* In that case, the quantity K may also be written as

$$K = \frac{\text{neutrons in present generation}}{\text{neutrons in preceding generation}} \quad (4)$$

This interpretation follows because neutrons in the present generation were, aside from some small source contributions that are discussed below, all produced from fission, and those in the preceding generation were either absorbed (including absorption in fuel) or lost to leakage.

If a reactor's core multiplication factor is exactly unity, then the reactor is said to be *critical.* That is, the number of neutrons produced from fission is exactly balanced by those that are absorbed or lost to leakage. A reactor can be critical at any power level, because the number of neutrons in the life cycle does not matter to the achievement of criticality as long as a balance exists between production and removal. However, a reactor's power level is proportional to the number of neutrons in the life cycle. More neutrons imply more fissions and hence a greater energy release. Thus, the reactor with the greater number of neutrons in its life cycle will be at the higher power level.

REACTOR OPERATION

If a reactor's core multiplication factor is known, the rate of change of its neutron population can be determined. For subcritical behavior such as occurs during a startup, it is useful to work with the parameter K directly. However, once critical, it is preferable to define a quantity called the *reactivity* ρ. The definition is

$$\rho \equiv \frac{K - 1}{K} \quad (5)$$

Hence, ρ may be thought of as the fractional change in the neutron population per generation. A reactivity of zero implies that the reactor is exactly critical. If the reactivity is less than unity, the reactor is subcritical. If it exceeds unity, the reactor is supercritical.

Reactivity is a global property of a nuclear reactor. Nevertheless, it is common practice to associate a *reactivity worth* with each core component and feedback effect as if these components and/or effects existed independently. For example, movement of a control device will alter the core multiplication factor and hence insert a certain amount of reactivity. The actual magnitude of the change is usually expressed as a function of the distance over which the device is moved. In reality, it is a function of both that distance and the current core configuration. If the latter were changed, a different reactivity worth might be observed for the same movement of the device.

The reactivity associated with a control device may be tabulated in terms of either a differential or an integral value. The former is the change in reactivity per unit distance of travel at a given position. It has units of millibeta per meter. The latter is the total change in reactivity generated by moving a device from one position to another. Its units are millibeta.

Operation of a reactor may be subdivided into three distinct regimes. These are subcritical, critical but below the *point-of-adding-heat,* and critical at appreciable power. Each regime is associated with a different dynamic behavior.

The reactor dynamics of the first regime are characterized by the subcritical multiplication process. The neutron population grows in response to the slow removal of the control mechanisms until a self-sustaining neutron chain reaction is achieved.

The second regime is that in which the power is raised from whatever level existed when criticality was attained to the point-of-adding-heat, which is usually 1% to 3% of full power. The point-of-adding-heat is defined as the power level above which a change in power is reflected as a change in temperature. At lower powers, the plant heat capacity is such that power changes do not affect temperature. Reactor dynamics are described by the point kinetics equations. The distinguishing feature of operation in this regime is the absence of reactivity feedback effects associated with changes in temperature, void fraction, or fission product poisons such as xenon.

The third regime, operation at power, is also describable in terms of the point kinetic equations provided that allowance is made for the many feedback effects.

Reactor Startup and Subcritical Operation

A prerequisite for a safe reactor startup is that there are sufficient neutrons present so that the nuclear instrumentation is operable. In a shut-down reactor, the ratio of photons (gamma rays) to neutrons may be as great as 100:1. Thus, it is important that the instruments be capable of distinguishing neutrons from photons and that the neutron population be sufficient to generate on-scale signals. Possible neutron sources in a shut-down reactor include spontaneous fission, photoneutron reactions, and installed sources.

Spontaneous fission is important as a source of neutrons in both PWRs and BWRs because of the large amount of ^{238}U present. The magnitude of this source decreases as enrichment rises or as reactors become smaller, because in both instances there is less ^{238}U.

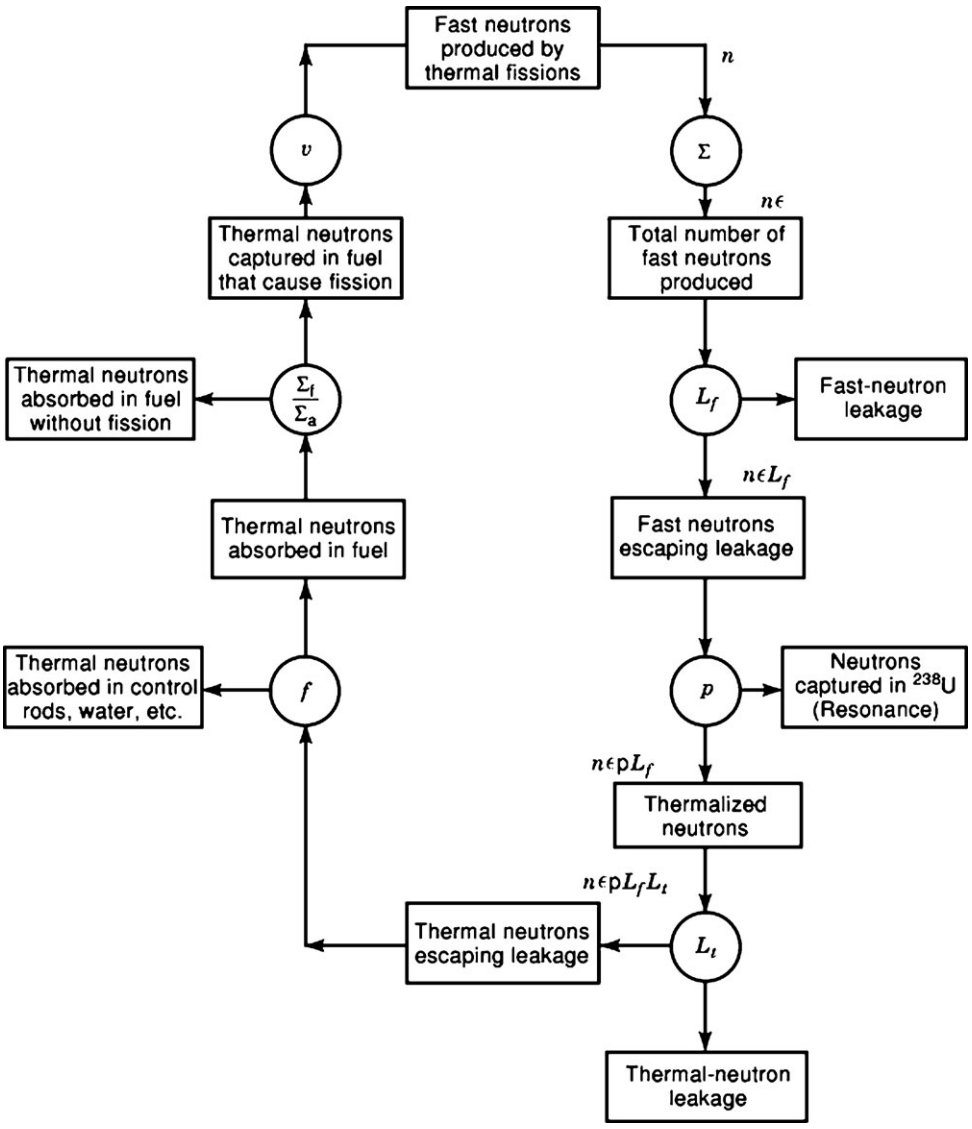

Figure 3. Neutron life cycle.

Photoneutrons are produced when gamma rays that are emitted by certain fission products interact with a deuterium (heavy hydrogen) nucleus to yield a neutron and proton: <chemStruc id="W5206-cstr-0001" graphicIsNumbered="no" numbered="no">

$$^{0}_{0}\gamma + {}^{2}_{1}\mathrm{H} \rightarrow {}^{1}_{0}\mathrm{n} + {}^{1}_{1}\mathrm{H}$$

The photoneutron reaction is important in LWRs because deuterium oxide, which is heavy water (D_2O), is present in ordinary water (H_2O). The gamma rays needed to initiate the reaction require an energy of 2.2 MeV. Fission products that produce such gamma rays decay within a few months of a reactor shutdown. Thus, for a photoneutron source to be of use, the reactor must have been recently operated at full power for an extended time.

There are several possible types of installed sources. One of the most common is plutonium–beryllium, or PuBe. The reaction is <chemStruc id="W5206-cstr-0002" graphicIsNumbered="no" numbered="no">

$$^{239}_{94}\mathrm{Pu} \rightarrow {}^{4}_{2}\mathrm{He} + {}^{235}_{92}\mathrm{U}$$

$$^{4}_{2}\mathrm{He} + {}^{9}_{4}\mathrm{Be} \rightarrow {}^{12}_{6}\mathrm{C} + {}^{1}_{0}\mathrm{n}$$

A PuBe source consists of a mixture of plutonium and beryllium powder that is doubly encased in stainless steel. Heat transfer from these sources is therefore poor, and they must be removed from a reactor prior to the production of any appreciable power. [It is important to recognize that removal of such a source from a critical reactor will cause a power increase because moderator (light water) replaces the space previously occupied by the steel casing. This is further explained below under feedback effects.] Sources of the type described by the chemical equations above, can also be manufactured using certain other alpha emitters such as polonium.

Another type of installed source is antimony–beryllium. The reaction is <chemStruc id="W5206-cstr-0003" graphicIsNumbered="no" numbered="no">

$$^{1}_{0}n + {}^{123}_{51}Sb \rightarrow {}^{124}_{51}Sb$$

$$^{124}_{51}Sb \rightarrow {}^{0}_{0}\gamma$$

$$^{0}_{0}\gamma + {}^{9}_{4}Be \rightarrow {}^{1}_{0}n + 2({}^{4}_{2}He)$$

$$^{4}_{2}He + {}^{9}_{4}Be \rightarrow {}^{1}_{0}n + {}^{12}_{6}C$$

Radioactive antimony (^{124}Sb) is a prerequisite for the operation of this source.

Both spontaneous fission and the photoneutron reaction provide a homogeneous source of neutrons. No consideration need be given to source–detector geometry. However, if an installed source is used, then its placement becomes an issue. Installed sources should be positioned at the center of the core, with neutron detectors located symmetrically on the core perimeter, one in each quadrant.

A shut-down reactor that contains fissile material will have a core multiplication factor that is greater than zero but less than one. The actual neutron population in this reactor will be greater than the source population alone. This is evident from the neutron life cycle as shown in Fig. 3. Suppose a source emits S_0 neutrons every generation. Once these S_0 neutrons have completed their life cycle sequence, they will have contributed KS_0 neutrons towards the next generation. The total neutron population after one generation will therefore be $S_0 + KS_0$. Similarly, two generations later there will be $S_0 + KS_0 + K^2S_0$ neutrons, the contribution from the original S_0 neutrons now being the term K^2S_0. The original S_0 neutrons contribute to each successive generation until K^n approaches zero, where n is the number of generations. The existence of a neutron population in excess of the source level in a subcritical fissile medium is referred to as *subcritical multiplication.* After n generations, that population will be

$$\text{total neutrons} = S_0 + KS_0 + K^2S_0 + K^3S_0 + K^4S_0 + \cdots + K^{n-1}S_0 + K^nS_0 \quad (6)$$

If the core multiplication factor K is less than unity, then the series $1 + K + K^2 + \cdots K^n$ will converge to $1/(1 - K)$ as K^n approaches 0. Thus, the total neutron population is $S_0/(1 - K)$. This relation may be used to calculate the equilibrium neutron level in a subcritical reactor. To do this, recognize that the count rate (CR) obtained over a designated time interval is proportional to the neutron population. Thus,

$$CR = \frac{S_0}{1 - K} \quad (7)$$

Several caveats apply. First, this formula for subcritical multiplication does not allow calculation of the time required to attain criticality, because time does not appear in it. Second, as the core multiplication factor K approaches unity, the number of generations and hence the time required for the neutron population to stabilize get longer and longer. This is one reason why it is important to conduct a reactor startup slowly. If the startup were to be done rapidly, there would not be sufficient time for the neutron population to attain its equilibrium value. Third, the equilibrium neutron level in a subcritical reactor is proportional to the neutron source strength. This is why it is important to have neutron count rates in excess of some minimum prior to initiating a startup. Fourth, the formula is only valid while subcritical. This restriction applies because Eq. (6) will converge to a limit only if K is less than unity. Fifth, K is not zero in a shutdown reactor, but typically ranges between 0.90 and 0.95.

Equation (7) is not very useful, because a method to measure K does not exist. However, both the count rate and the source level are measurable. The latter is simply the count rate with the reactor in a shut-down state. Thus, it is of benefit to rearrange Eq. (7) to

$$1 - K = CR_0/CR \quad (8)$$

where CR_0 is the initial count rate. This is the equation of a straight line where K is the independent variable and CR_0/CR is the dependent variable. Often, the latter term is written as S_0/CR. It can be used to predict the attainment of criticality provided that S_0 remains constant. For example, suppose an installed source is at the core center and fuel assemblies are being loaded. Neutron counts are recorded after every fifth assembly is installed. A plot of S_0/CR versus the number of assemblies can be used to estimate the fuel loading required to achieve criticality. This is illustrated in Fig. 4. The broken line is an extrapolation of the measured data. It shows that criticality will be achieved at about 60 assemblies. Such graphs are referred to as $1/M$ *plots,* where M stands for multiplication. Such plots can also be used to predict criticality during a reactor startup. In that case, the horizontal axis will be the position of the control devices.

Actual $1/M$ plots may not be linear because of several factors. First, poor counting statistics may affect the initial portion of the plot. Second, the source–detector geometry may not result in a uniform distribution of neutrons throughout the core region. Third, the effect of the independent variable (number of assemblies, control device position, etc.) may not be linear. Fourth, the source level may not be constant. For example, if spent fuel is replaced with fresh, then the photoneutron contribution to the source will diminish because the fission products that produce the 2.2 MeV photons are being removed. Care must therefore be exercised in the use of these plots.

If a reactor is to be loaded with fresh fuel, then two $1/M$ plots will be made: the first with all control devices fully inserted, and the second with certain devices removed from the core. The difference in the projections of the two plots represents the minimum amount, or margin, by which the reactor can be shut down. Fuel will be loaded only until the projection of the second plot is attained.

Low-Power Operation While Critical

The starting point for the derivation of the time-dependent relations needed to describe neutron dynamics once criticality has been achieved is the fission process. The fission of a ^{235}U nucleus normally yields two fission fragments, an average of 2.5 neutrons, and an assortment of beta particles, gamma rays, and neutrinos. The neutrons that are produced directly from the fission event are referred to as *prompt,* because they appear almost instantly. Most of the neutrons produced in a reactor are prompt. However, certain fission fragments, which are called *precursors,* undergo a beta decay to a daughter nuclide that then emits a neu-

Assemblies loaded	Source	Count Rate	S_0/CR
0	S_0	S_0	1.00
5	S_0	$1.09S_0$	0.92
10	S_0	$1.20S_0$	0.83
15	S_0	$1.33S_0$	0.75
20	S_0	$1.50S_0$	0.67
25	S_0	$1.71S_0$	0.58
30	S_0	$2.00S_0$	0.50

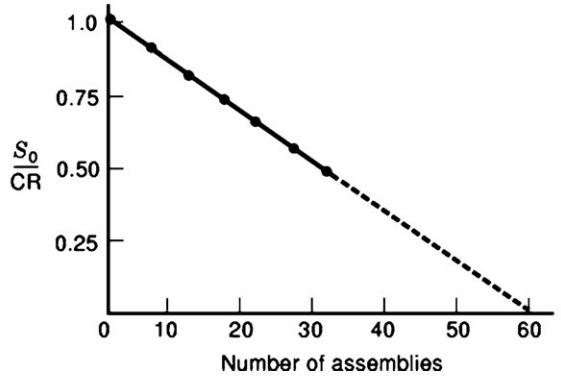

Figure 4. Example of 1/M plot.

tron. Neutrons produced in this manner are referred to as *delayed*. The delay is the time that must elapse for the precursor to undergo its beta decay. Delayed neutrons constitute an extremely small fraction of a reactor's total neutron population. The fraction of delayed neutrons that exists at fast energies is denoted by the symbol β.

Figure 5 illustrates the fission process. There are three parallel paths whereby neutrons may be produced. These paths or mechanisms yield prompt neutrons, delayed neutrons, and photoneutrons respectively. Delayed neutrons and photoneutrons are born at energies that are less than those of their prompt counterparts. Nevertheless, all three types of neutrons are fast when born, and all three types require moderation in order to continue the neutron chain reaction. However, because they are born at somewhat lower energies, the delayed neutrons and photoneutrons are less likely to be lost from the core because of fast leakage than are the prompt neutrons. Hence, the fraction of these neutrons increases during the slowing-down process. To summarize: In absolute numbers, all three neutron populations decrease as they move through the neutron life cycle. However, the loss of the prompt neutrons is greatest because they are born at the highest energies. So the fraction of the delayed neutrons and photoneutrons increases. The fraction of these neutrons at thermal energies is denoted by the symbol $\mathring{\beta}$, which is commonly called the *effective* delayed neutron fraction. Its value for LWRs is typically 0.0065.

The effective delayed-neutron fraction $\mathring{\beta}$ is a very small number. Nevertheless, delayed neutrons are crucial to the safe operation of a nuclear reactor. The reason is that they lengthen the average neutron lifetime so that control of the neutron chain reaction is possible. The time required for a prompt neutron to be born, thermalize, and cause a fission is about 1×10^{-4} s in an LWR. The corresponding average life time for a delayed neutron is about 12.2 s. Denote these times as t_p and t_d. Then the average lifetime of a neutron is $(1 - \mathring{\beta})t_p + (\mathring{\beta})t_d$, or about 0.79 s. This is considerably longer than the 10^{-4} s prompt-neutron lifetime.

Reactivity was defined earlier as $(K - 1)/K$, which is the fractional change in the neutron population per generation. It is common practice to quantify reactivity by reference to the delayed-neutron fraction. Thus, even though reactivity is a fraction and therefore dimensionless, several systems of units have been developed. The most common are $\Delta K/K$ and the beta. The conversion for an LWR is that 1 beta of reactivity equals 0.0065 $\Delta K/K$. There are 1000 millibeta (mbeta) in 1 beta. Reactivity values are also sometimes quoted in *dollars* and *cents,* with 1 beta being equal to a reactivity of one dollar or 100 cents.

The amount of positive reactivity present in a nuclear reactor should always be limited to some small percentage of the delayed neutron fraction. In this way, the delayed neutrons will be the rate-determining factor in any transient, and their 12.2 s lifetime will govern the reactor's dynamics. The rationale for this approach is illustrated in Table 1. Shown are three cases, all with the initial condition that the reactor is critical with a population of 10,000 neutrons. For the first case, no change is made. One generation later there are 9935 prompt and 65 delayed neutrons. Criticality cannot be maintained without the delayed neutrons. Hence they are the rate-determining step. For the second case, 0.500 beta of reactivity is added. This corresponds to 0.00325 $\Delta K/K$, or 33 additional neutrons in the first generation. Thus, after one generation, there will be a total of 10,033 neutrons, of which 9968 will be prompt and 65 will be delayed. The reactor power, which is proportional to the number of neutrons in the life cycle, is rising. However, the delayed neutrons are still controlling, because it takes 10,000 neutrons to remain critical and there are only 9968 prompt ones. For the third case, 1.5 beta of reactivity is added. This corresponds to 0.00975 $\Delta K/K$ or 98 neutrons in the first generation. Thus, after one generation

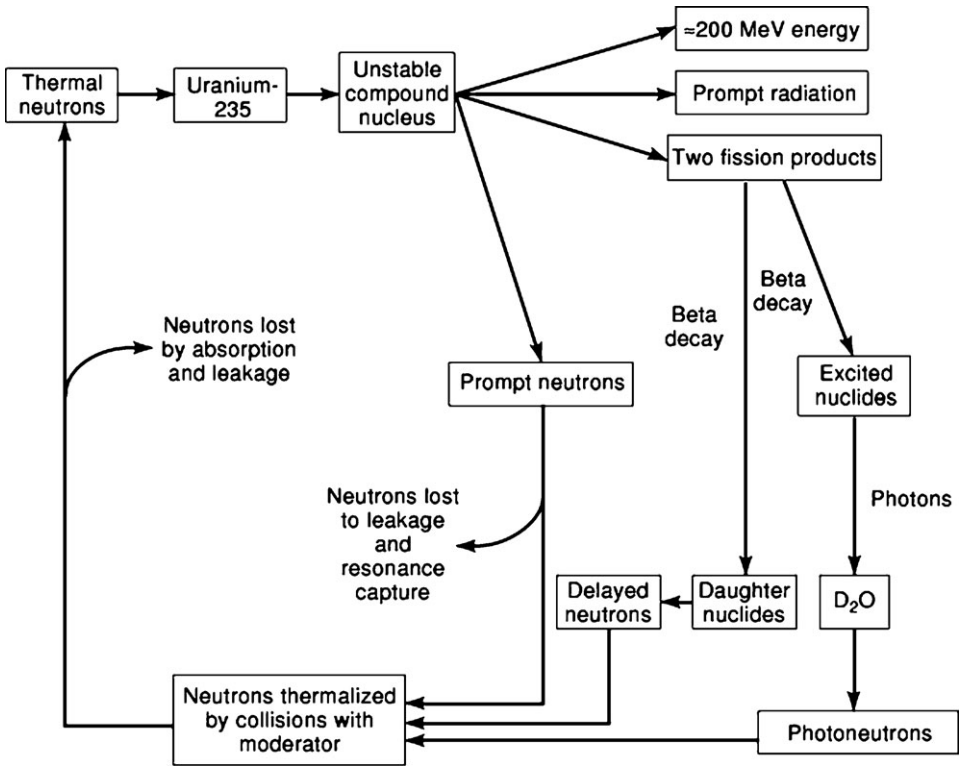

Figure 5. Fission process.

there are 10,032 prompt and 66 delayed neutrons. There are more than enough prompt neutrons to maintain criticality. Thus, the prompt neutrons with their 10^{-4} s lifetime are the controlling factor.

If an amount of reactivity equal to delayed neutron fraction is inserted, the reactor is said to be in a *prompt critical* condition. This terminology should not be construed to mean that reactivity insertions of less than that value are safe and those greater than it are unsafe. There is no sharp division between safe and unsafe. Rather, as reactivity is increased, there is a continuous decrease in the importance of the delayed neutrons to the reactor's dynamic response. For that reason, transients in LWRs are normally performed with small reactivity additions such as 100 mbeta.

Reactivity is not directly measurable, and hence most reactor operating procedures do not refer to it. Instead, they specify a limiting rate of power rise, commonly called a *reactor period.* The reactor period is denoted by the Greek letter τ and is defined as the power level divided by its rate of change:

$$\tau(t) = \frac{P(t)}{dP(t)/dt} \tag{9}$$

where $P(t)$ is the reactor power. Thus, a period of infinity (zero rate of change) corresponds to the critical condition. If the period is constant, the integration of Eq. (9) gives the relation

$$P(t) = P_0 e^{t/\tau} \tag{10}$$

where P_0 is the initial power. A reactivity addition of 100 mbeta (0.00065 $\Delta K/K$) would create a period of about 100 s. For such a period, Eq. (10) shows that an increase in reactor power by an order of magnitude (from 10% to 100%, for example) would require 230 s. Such a rate of rise is quite manageable.

In order to understand the time-dependent behavior of a nuclear reactor, equations are needed that describe the response of the prompt and delayed neutron populations to changes in reactivity. This problem is mathematically complex, because the neutron population in a reactor is a function of both space (position in the core) and time. The spatial dependence is significant if the dimensions of the core exceed the distance that a neutron will travel while slowing down and diffusing. For many practical situations, it is acceptable to assume that the spatial and temporal behavior are separable. The result is the space-independent equations of reactor kinetics. These are often called the *point kinetics* equations. They are

$$\frac{dT(t)}{dt} = \frac{\rho(t) - \overline{\beta}}{l^*} T(t) + \sum_i^N \lambda_i C_i(t) \tag{11}$$

$$\frac{dC_i(t)}{dt} = \frac{\overline{\beta}_i T(t)}{l^*} - \lambda_i C_i(t) \qquad \text{for} \quad i = 1, \ldots, N \tag{12}$$

where $T(t)$ amplitude function
$\rho(t)$ net reactivity
β effective delayed neutron fraction
l^* prompt neutron lifetime
λ_i decay constant of the ith precursor group
C_i concentration of the ith precursor group
N number of delayed-neutron precursor groups

Table 1. Rationale for Limiting Reactivity

Initial Condition	Initial Population	Reactivity Addition	One Generation Later	Condition
Critical	10,000 neutrons 9,935 prompt 65 delayed	0% $\Delta K/K$ (0.0 beta, 0 neutrons)	10,000 9,935 65	Steady state; not critical on prompt neutrons alone
Critical	10,000 neutrons 9,935 prompt 65 delayed	0.325% $\Delta K/K$ (0.5 beta, 33 neutrons)	10,033 9,968 65	Supercritical; not critical on prompt neutrons alone
Critical	10,000 neutrons 9,935 prompt 65 delayed	0.98% $\Delta K/K$ (1.5 beta, 98 neutrons)	10,098 10,032 66	Power runaway; critical on prompt neutrons alone

The reader is referred to one of the classic texts on reactor physics for the derivation of these equations (1, 2). The amplitude function $T(t)$ is a weighted integral of all neutrons present in the reactor. It is common practice to equate $T(t)$ to the reactor power, denoted here as $P(t)$. However, this simplification ignores an important restriction on the validity of the point kinetics equations. Namely, the derivation requires that the shape (as opposed to the amplitude) of the neutron flux remain constant during a transient. This is in turn implies that significant physical movement of a reactor's control device would invalidate the point kinetics approach to the analysis of a transient. However, because the allowed magnitude of reactivity insertions is already limited for reasons of safety, this additional restriction often has little practical effect in the analysis of operational transients on LWRs.

Table 2 gives the physical meaning of the quantities that appear in the point kinetics equations. The first kinetics equation describes the behavior of the neutrons. It states that the rate of change of the total neutron population equals the sum of the rates of change of the prompt and delayed neutrons. The second kinetics equation describes the behavior of the precursors. It says that the rate of change of the precursors is the difference between their production and loss, the latter being described by radioactive decay. Twenty fission fragments may undergo a beta decay to produce a daughter nuclide that then emits a delayed neutron. Each has its own unique half-life. However, some of the half-lives are sufficiently similar so that the precursors can be represented as six effective groups. Thus, the value of N is usually six.

Table 2. Physical Interpretation of Parameters in Point Kinetics Equations

Quantity	Physical Meaning
$\rho(t)$	Fractional change in the total neutron population per generation
β	Fraction of neutrons that are delayed
$\rho(t) - \bar{\beta}$	Fractional change in the prompt-neutron population per generation
$1/\ell^*$	Number of neutron generations per unit time
$T(t)$	Amplitude function (equivalent to total neutron population)
$\{[\rho(t) - \bar{\beta}]/\ell^*\}T(t)$	Change in prompt-neutron population per unit time
$\lambda_i C_i$	Rate of decay of delayed-neutron precursors. This equals the rate of appearance of the delayed neutrons
$\bar{\beta}_i T(t)/\ell^*$	Rate of production of delayed-neutron precursors per unit time

The next step in the analysis of reactor transients is to relate the reactor period to the reactivity. This has traditionally been achieved through use of the Inhour equation, which is valid only for step insertions of reactivity and then only after sufficient time has elapsed since the insertion to achieve asymptotic conditions (1, 2). A more general approach is to combine the point kinetics equations through a process of differentiation and substitution to obtain the dynamic period equation (3). This relation is rigorous, subject only to the aforementioned limitations on Eqs. (11) and (12). The derivation is done in terms of the quantity $\omega(t)$ which is the reciprocal of the instantaneous reactor period. Thus, $\tau(t) = 1/\omega(t)$, where

$$\dot{T}(t) \equiv \omega(t)T(t) \tag{13}$$

This definition is substituted into Eq. (11).

The next step in the derivation is to define the standard effective multigroup decay parameter:

$$\lambda_e(t) \equiv \frac{\sum \lambda_i C_i(t)}{\sum C_i(t)} \tag{14}$$

This parameter is time-dependent because the relative concentrations of the delayed-neutron precursor groups change as the power is raised and lowered. Specifically, as the power is increased, short-lived precursors dominate and the value of λ_e increases. As the power decreases, the reverse occurs. Application of this definition to the neutron kinetics equations yields

$$\omega(t)T(t) = \frac{\rho(t) - \bar{\beta}}{l^*} T(t) + \lambda_e(t) \sum C_i(t) \tag{15}$$

The next and crucial step in the derivation is to differentiate Eq. (15), the modified version of the neutron kinetics equation. So doing, and using the definition of the instantaneous period from Eq. (13) to eliminate the derivative of the amplitude function, we obtain

$$\dot{\omega}(t)T(t) + [\omega(t)]^2 T(t) + \frac{\bar{\beta} - \rho(t)}{l^*} \omega(t)T(t) = \frac{\dot{\rho}(t)T(t)}{l^*} + \dot{\lambda}_e(t) \sum C_i(t) + \lambda_e(t) \sum \dot{C}_i(t) \tag{16}$$

A series of substitutions are now made. First, the rate of change of the precursor concentrations is eliminated by substitution of Eq. (12):

$$\dot{\omega}(t)T(t) + [\omega(t)]^2 T(t) + \frac{\bar{\beta} - \rho(t)}{l^*} \omega(t)T(t) = \frac{\dot{\rho}(t)}{l^*} T(t) + \dot{\lambda}_e(t) \sum C_i(t) + \frac{\lambda_e(t)\bar{\beta}T(t)}{l^*} - \lambda_e(t) \sum \lambda_i C_i(t) \tag{17}$$

Next, the quantity $\Sigma C_i(t)$ is eliminated by use of Eq. (14), the definition of the effective multigroup decay parameter. The quantity $\Sigma\lambda_i C_i(t)$ is then eliminated by substitution of Eq. (11) after modification by Eq. (13). The result is

$$\begin{aligned}\dot{\omega}(t)T(t) + [\omega(t)]^2 T(t) + \frac{\overline{\beta} - \rho(t)}{l^*}\omega(t)T(t) \\ = \frac{\dot{\rho}(t)}{l^*}T(t) + \frac{\dot{\lambda}_e(t)}{\lambda_e(t)}\left(\omega(t)T(t) + \frac{\overline{\beta} - \rho(t)}{l^*}T(t)\right) \\ + \frac{\lambda_e(t)\overline{\beta}T(t)}{l^*} - \lambda_e(t)\left(\omega(t)T(t) + \frac{\overline{\beta} - \rho(t)}{l^*}T(t)\right)\end{aligned} \tag{18}$$

Division by the amplitude function $T(t)$ and consolidation of terms yields

$$\begin{aligned}\omega(t)\left(\frac{\dot{\omega}(t)}{\omega(t)} + \omega(t) + \frac{\overline{\beta} - \rho(t)}{l^*} + \lambda_e(t) - \frac{\dot{\lambda}_e(t)}{\lambda_e(t)}\right) \\ = \frac{\dot{\rho}(t)}{l^*} + \frac{\lambda_e(t)\rho(t)}{l^*} + \frac{\dot{\lambda}_e(t)}{\lambda_e(t)}\left(\frac{\overline{\beta} - \rho(t)}{l^*}\right)\end{aligned} \tag{19}$$

Solving for $\omega(t)$ yields

$$\omega(t) = \frac{\dot{\rho}(t) + \lambda_e(t)\rho(t) + \dfrac{\dot{\lambda}_e(t)}{\lambda_e(t)}[\overline{\beta} - \rho(t)]}{\overline{\beta} - \rho(t) + l^*\left(\dfrac{\dot{\omega}(t)}{\omega(t)} + \omega(t) + \lambda_e(t) - \dfrac{\dot{\lambda}_e(t)}{\lambda_e(t)}\right)} \tag{20}$$

The expression for the instantaneous reactor period is therefore

$$\tau(t) = \frac{\overline{\beta} - \rho(t) + l^*\left(\dfrac{\dot{\omega}(t)}{\omega(t)} + \omega(t) + \lambda_e(t) - \dfrac{\dot{\lambda}_e(t)}{\lambda_e(t)}\right)}{\dot{\rho}(t) + \lambda_e(t)\rho(t) + \dfrac{\dot{\lambda}_e(t)}{\lambda_e(t)}[\overline{\beta} - \rho(t)]} \tag{21}$$

This relation is the standard dynamic period equation. It is a rigorously derived exact relation. Given that the prompt neutron lifetime is extremely small, terms containing l^* can often be deleted and Eq. (21) reduces to

$$\tau(t) \approx \frac{\overline{\beta} - \rho(t)}{\dot{\rho}(t) + \lambda_e(t)\rho(t) + \dfrac{\dot{\lambda}_e(t)}{\lambda_e(t)}[\overline{\beta} - \rho(t)]} \tag{22}$$

This result is a general relation that accurately predicts the instantaneous reactor period associated with any reactivity pattern provided that the prompt critical value is not approached. It can generally be considered valid provided that the reactor period is longer than 10 s.

The instantaneous reactor period may also be expressed in terms of the alternate dynamic period equation. The two relations are mathematically equivalent. The alternate form avoids the presence of the term $(\beta_e^2/\lambda_e)(\beta - \rho)$, which is difficult to evaluate numerically. The derivation is similar except that the differentiation is performed before substitution of the effective decay parameter and that parameter is defined differently as

$$\lambda'_e(t) \equiv \frac{\sum \lambda_i^2 C_i(t)}{\sum \lambda_i C_i(t)} \tag{23}$$

The alternate form of the dynamic period equation is

$$\tau(t) = \frac{[\overline{\beta} - \rho(t)] + l^*\left(\dfrac{\dot{\omega}(t)}{\omega(t)} + \omega(t) + \lambda'_e(t)\right)}{\dot{\rho}(t) + \lambda'_e(t)\rho(t) + \sum \overline{\beta}_i[\lambda_i - \lambda'_e(t)]} \tag{24}$$

Further information on both equations is given in Ref. 4. It has been shown that the Inhour equation is a special case (asymptotic conditions following a step reactivity insertion) of the dynamic period equation (5).

Examination of either form of the dynamic period equation shows that the instantaneous reactor period is a function of the rate of change of reactivity, the reactivity, and the rate of redistribution of the delayed-neutron precursors within the defined groups. Furthermore, review of the derivation of this equation shows that the rate of change of reactivity is proportional to the prompt-neutron population, while the terms that contain the reactivity and the rate of redistribution of the precursors are related to the delayed-neutron population. Those realizations provide certain physical insights relative to reactor control and operation:

1. The period observed at any given moment in a reactor will depend on both the distance that a control device has been moved beyond the critical position and the rate at which that device is being moved. The former corresponds to the reactivity, and the latter to the rate of change of reactivity.
2. Changes in the velocity of a control device will have an immediate effect on the period because such changes alter the prompt-neutron population.
3. A high-speed insertion of the control devices over a short distance can be used to reduce the reactor power rapidly. Such power reductions, which are referred to as *cutbacks,* are used on some LWRs as an alternative to an abrupt reactor shutdown, or *scram.* Advantages of this alternative are that the reactor is subject to less thermal cycling and the time to restore full operation is reduced.
4. Reactivity cannot be changed on demand. Rather, a control device's position or a soluble poison's concentration has to be adjusted first. This takes time.
5. The dynamic response of a reactor is determined by that of its prompt- and delayed-neutron populations. Prompt neutrons appear essentially simultaneously with the fission event and are therefore a function of the current power level. Delayed neutrons appear some time after the fission event and are therefore a function of the power history. This dependence on the power history means that delayed neutrons will not be in equilibrium with the observed power during a transient. Hence, upon attaining a desired power level, the delayed neutrons will continue to rise, and an overshoot will occur unless the controller is capable of reducing the prompt-neutron population at a rate sufficient to offset the still-rising delayed-neutron population.

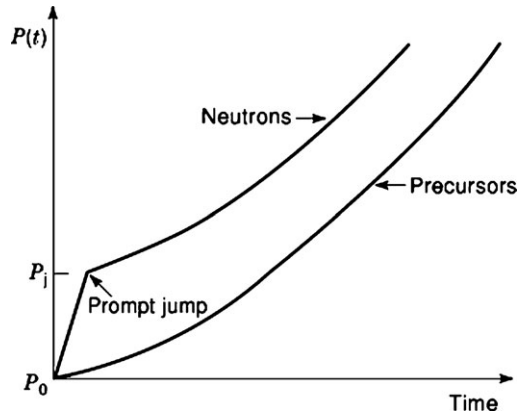

Figure 6. Response to a step change in reactivity.

These insights are best illustrated through analysis of a reactor's response to step and ramp reactivity changes. The point kinetics equations can be solved analytically for step changes in reactivity. Figure 6 shows the behavior of both the neutrons and precursors following a step insertion such that $\rho < \beta$. The initial response is a rapid increase in the prompt-neutron population. This is called the *prompt jump,* and it represents the start of a nuclear runaway. However, the runaway cannot continue, because the reactivity is less than the delayed-neutron fraction. The prompt jump occurs almost instantly. (*Note:* The time scale on the figure is distorted.) This rapid rise is then followed by a more gradual one that corresponds to the growth of the delayed neutrons. Once asymptotic conditions are established, and if no feedback effects exist, the rise becomes exponential.

The power after the prompt jump is given by

$$P_j = P_0 \left(\frac{\overline{\beta}}{\overline{\beta} - \rho} \right) \tag{25}$$

where P_0 is the power level before the step insertion. Note that the effect of a step insertion depends on the initial power level. Suppose a reactivity of 0.2β is inserted as a step, and also assume the initial power to be 10% of allowed. The final power is 12.5% of allowed, a minor change. But what if the initial power had been 90% of allowed? The final power would be 112.5% of allowed, a potentially serious problem.

The power behavior following the prompt jump may be approximated as

$$P(t) = P_j e^{t/\tau} \tag{26}$$

$$\tau \approx \frac{\overline{\beta} - \rho}{\lambda_e \rho} \tag{27}$$

where ρ is zero, λ_e can be assumed to be 0.1 s^{-1}, and ρ is the step reactivity insertion. Equation (27) is only valid for asymptotic conditions.

Reactivity would never be deliberately added in a stepwise manner to an LWR. In contrast, ramp reactivity insertions are commonplace. These occur whenever control devices are moved or when the concentration of a soluble poison is adjusted. Numerical analysis is required to determine the effect of a ramp insertion accurately. The material presented here is qualitative and is shown with the objective of illustrating important aspects of a reactor's dynamic behavior.

Assume that reactivity is inserted at the rate of 10 mbeta/s for 10 s, held constant at 100 mbeta for 40 s, and then removed at the rate of −5 mbeta/s. The reactor is exactly critical at the start of the insertion. It is useful to calculate the period immediately before and after each change in the reactivity insertion rate. Thus, calculations are done at $t = 0^-, 0^+, 10^-, 10^+, 40^-, 40^+, 60^-$, and 60^+ s. In addition, a calculation is done at 50 s. The relation used is

$$\tau(t) \approx \frac{\overline{\beta} - \rho(t)}{\dot{\rho}(t) + \lambda_e(t)\rho(t)} \tag{28}$$

This relation is, as noted earlier, only approximate. Further, it is assumed that $\lambda_e(t)$ is equal to 0.1 s^{-1}. The problem is solved with reactivity in millibetas. Thus, β equals 1000.

At $t = 0^-$, both the reactivity and its rate of change are zero. So the period is infinite. At $t = 0^+$, the reactivity is still zero. But the rate of change of reactivity is now a positive 10 mbeta/s. Hence, $\tau = (1000 - 0)/[10 + (0.1)(0)] = 100$ s. So the mere act of initiating a reactivity insertion has immediately placed the reactor on a positive period of 100 s. At $t = 10^-$, there is 100 mbeta of reactivity present, and reactivity is still being added at the rate of 10 mbeta/s. The period has gone from 100 s to 45 s. The power is rising at an ever increasing rate. At $t = 10^+$, rod withdrawal stops and the rate of reactivity change become zero. This causes the period to lengthen to 90 s. The power is still rising but at a slower rate. At $t = 40^-$, conditions are the same as at $t = 10^+$. This means that the period was constant for $10 < t < 40$. So, during that segment of the transient, the power rose on a pure exponential. At $t = 40^+$, the reactivity is still 100 mbeta, but the rate of change of reactivity is now −5 mbeta/s. Hence, $\tau = (1000 - 100)/[-5 + (0.1)(100)] = 180$ s. The reactor power is still rising.

This is a very important observation. The fact that the control devices are being inserted does *not* necessarily mean that the reactor power is decreasing. It may still be rising, but at an ever decreasing rate. This behavior is the result of the delayed-neutron precursors that have not yet achieved their equilibrium value for the current reactor power level.

At $t = 50$, another interesting observation is made. The rate of change of reactivity is still −5 mbeta/s. The reactivity present is now 50 mbeta. Thus, $\tau = (1000 - 50)/[-5 + (0.1)(50)]$. The period is infinite. The power increase has been halted even though there is still positive reactivity present. A decreasing prompt-neutron population is offsetting the still rising contribution of the delayed neutrons so as to establish a point of unstable equilibrium. This is sometimes called the "point of power turning." At $t = 60^-$, the period is −200 s and the power is actually decreasing. At $t = 60^+$, the period is infinite and the reactor is again exactly critical. However, it is at a higher power level than it was at the outset of the transient.

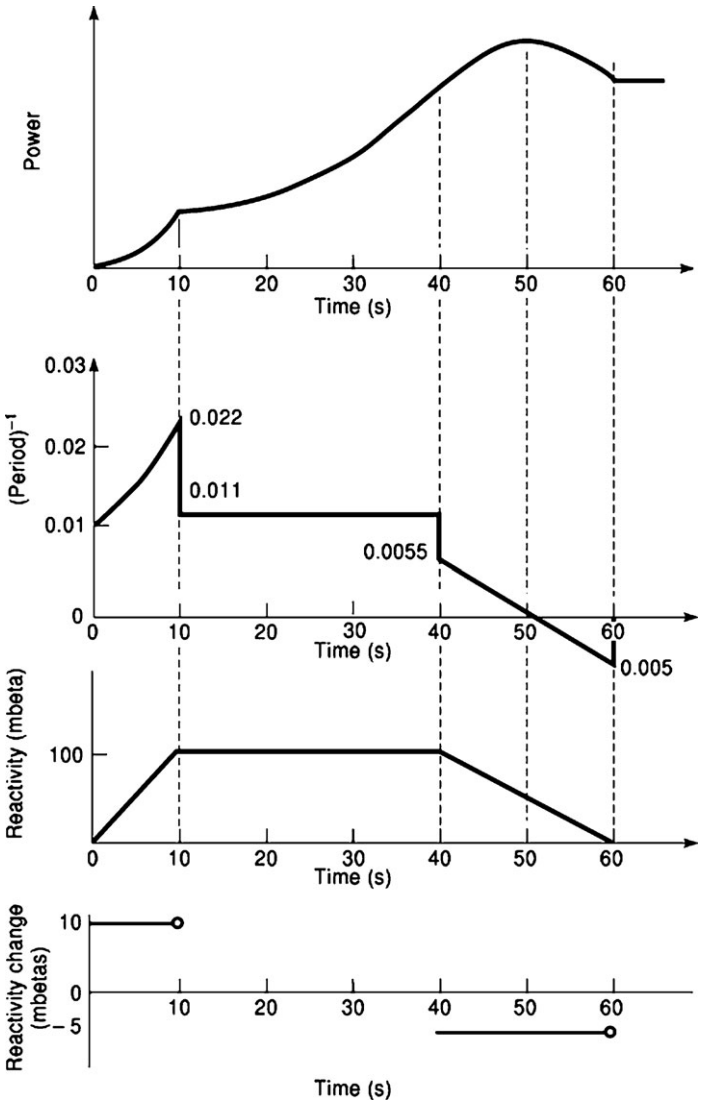

Figure 7. Response to ramp reactivity insertion.

Table 3. Qualitative Analysis of Ramp Reactivity Addition

Time	$\dot{\rho}$	ρ	τ	Comment
0−	0	0	∞	Reactor critical
0+	10	0	100	Period stepped to a positive value
10	10	100	45	Power rising faster than exponential
10+	0	100	90	Period stepped to a longer value
40	0	100	90	Power rising on pure exponential
40+	−5	100	180	Power rising but at an ever slower rate
50	−5	50	∞	Unstable equilibrium; period infinite
60−	−5	0	−200	Power decreasing
60+	0	0	∞	Reactor critical at higher power

Table 3 summarizes the calculation, and Fig. 7 shows the results. The reciprocal of the period is plotted because of the difficulty in displaying an infinite value. Observe that whenever the rate of change of reactivity is altered, there is a step change in the period and a change of slope in the power profile. This calculation was idealized. In the actual case, the effect of the different half-lives of each precursor group would make the features shown in the figure less distinct.

A digital control algorithm that utilizes the dynamic period equation as a model of a reactor's dynamics and thereby permits a desired power trajectory to be tracked is described by Bernard and Lanning (6). This reference also enumerates some of the safety issues associated with the use of digital controllers for nuclear reactors.

Operation in the Presence of Feedback Effects

The equations that were developed to describe a reactor's response while critical at low power are also applicable to high-power operation. The only difference is that the reactivity now becomes a function of the reactor power because of the presence of various power-dependent feedback effects. Thus, the point kinetics equations become nonlinear and can only be solved via numerical methods. Several of the more important feedback effects are described here.

Moderator Temperature. Negative reactivity feedback associated with the temperature of the moderator is one of the most important passive safety features in nuclear re-

actors. The physical basis for the effect in LWRs is that the coolant also functions as the moderator and as the temperature of the coolant rises, it becomes less dense. This in turn results in both less moderation and more leakage, because there are fewer thermalizing collisions between the fast neutrons and hydrogen nuclei. Leakage is a minor effect in a large reactor, so the decrease in moderation is usually the dominant factor in a large LWR. The neutron spectrum becomes *hardened,* that is, it is shifted to higher energies. As a result, negative reactivity is generated. This effect makes LWRs inherently self-regulating. The time scale for the effect is the time required for coolant to circulate through the primary loop. The specific sequence is evident from Fig. 8, which is a schematic of a PWR.

Suppose that the demand on the turbine increases. The following sequence then occurs:

1. The turbine first-stage steam pressure decreases.
2. The steam flow from the steam generator increases. This causes the steam generator pressure P_{sg} to drop.
3. The steam generator is a saturated system. So its temperature T_{sg} also drops.
4. The decrease in steam generator temperature causes a decrease in the cold leg temperature T_{CL} of the primary coolant. The cold leg is the piping through which the coolant flows on exiting the steam generator.
5. Cooler primary coolant enters the reactor core. This denser coolant increases neutron moderation.
6. The reactor power increases, and so does the temperature of the hot leg (T_{HL}). The hot leg is the piping through which coolant flows on exiting the core.
7. Hotter primary coolant reaches the steam generator. The steam generator temperature and pressure rise, and the steam supply equals the demand.

The above sequence may be summarized as follows:

$$\text{demand} \uparrow P_{sg} \downarrow T_{sg} \downarrow T_{CL} \downarrow \rho \uparrow \text{power} \uparrow T_{HL} \uparrow T_{sg} \uparrow P_{sg} \uparrow$$

The final result is that the reactor power has increased to equal the demand. Also, the difference between the hot and cold leg temperatures has increased.

Negative moderator temperature coefficients of reactivity are achieved by deliberately designing a reactor so that it is undermoderated. In the case of an LWR, this means that insufficient light water is present in the core to cause all of the fast neutrons produced from fission to slow down completely. The wisdom of this approach to reactor design is evident if the opposite is considered. If a core is overmoderated, then all neutrons attain thermal energies. If the reactor power becomes excessive, the coolant heats up and becomes less dense. So some moderator is lost from the core. However, this loss of moderator has no effect on neutron thermalization, because there was already more than enough moderator present. So the reactivity is unchanged, and the power excursion continues.

Failure to design a reactor so that it is undermoderated may even result in the existence of a positive moderator temperature coefficient of reactivity. Materials selected to be moderators are characterized by a high energy loss per neutron collision, a large cross section for scattering, and a small cross section for absorption. If sufficient moderator is present so that all neutrons are fully thermalized, then the absorption properties of the moderator, even though small, will affect the reactor's dynamics. Specifically, a decrease in the moderator's density will result in there being both fewer neutron scatters and fewer neutron absorptions. If the core was initially overmoderated, the loss of the scattering interactions will have no effect. However, the loss of some neutron absorption will generate positive reactivity. The power increase that caused the initial loss of moderator density will therefore accelerate. One of the design flaws in the Soviet style RMBK reactors was the existence of a positive coolant temperature coefficient of reactivity over a portion of the allowed temperature range. This was one of the contributing factors to the Chernobyl accident. Most reactors are required by national law to be undermoderated. (The RMBK reactors were not.)

It should be recognized that there are certain accident scenarios for which negative temperature coefficients of reactivity make the situation worse. These include steam line breaks and control rod drops. In both cases, the reactor will cool off if no corrective action is taken, and the negative coefficient will cause a positive reactivity insertion. Protection against these accidents is often provided by quick-closing shutoff valves on the steam lines and a requirement to scram the reactor on a dropped rod.

Void Coefficient. Negative void coefficients of reactivity are analogous to the negative moderator temperature coefficient. The formation of a void results in a decrease in the amount of moderator and hence the generation of negative reactivity. This effect is important in BWRs, which operate with a significant vapor fraction. In fact, BWR power can be controlled by adjusting the recirculation flow, which in turn controls the rate at which voids (steam bubbles) are generated in the core. The time scale for the void effect in a BWR is that for the coolant to flow through the recirculation loop.

Fuel Coefficient. If a reactor's fuel heats up, it will expand and become less dense. The result is the generation of negative reactivity, because there will be fewer fissions and more leakage. The fuel coefficient of reactivity is a minor effect in LWRs, because the uranium dioxide fuel has a very small coefficient of thermal expansion. It is, however, a significant effect in some other reactor types, particularly ones with metallic cores.

Doppler Coefficient. The Doppler effect is another means whereby a change in fuel temperature may alter reactivity. Unlike the fuel coefficient, it is a very important passive safety feature in LWRs. LWR fuel consists of both ^{235}U and ^{238}U. As shown in Fig. 2, one of the distinguishing features of the ^{238}U total cross section is the presence of six narrow resonances, with the first occurring at 6.67 eV. Neutrons with energies that correspond to one of these resonances will be absorbed in reactions that do not lead to fission. (The absorption of a neutron by ^{238}U will, after two successive beta decays, yield ^{239}Pu, which,

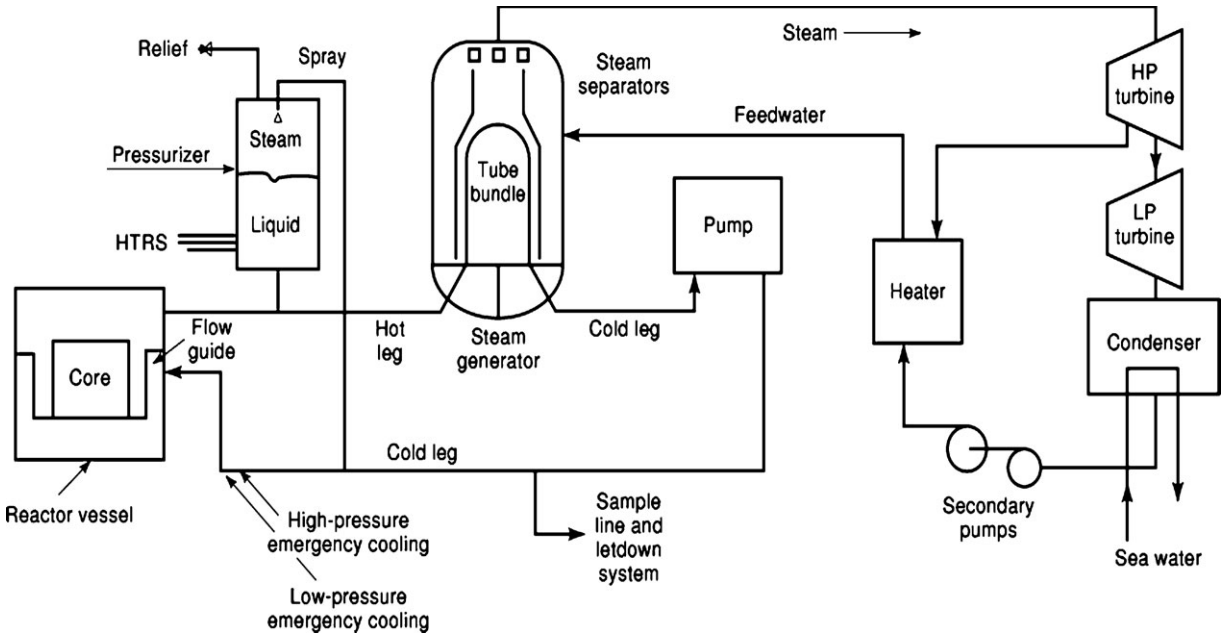

Figure 8. Schematic of a pressurized water reactor.

if it then absorbs a neutron, may fission. The half-lives for the two beta decays are 23.5 min and 56 h, respectively.) The basis of the Doppler effect is that nuclei vibrate, at a frequency that increases with temperature. That is, the nuclei undergo thermal motion. The effect of the motion is to broaden the energy range over which the resonance is effective. This is illustrated in Fig. 9. If the ^{238}U nuclei were at so low a temperature as to be almost at rest, the only neutrons absorbed would be those with energies of E_0. Now suppose that a neutron with an energy that is slightly below the lower limit of the resonance strikes a vibrating ^{238}U nucleus while that nucleus is moving towards the neutron. The collision's effect will be the same as if the neutron had slightly more energy and had hit an almost stationary ^{238}U nucleus. Hence, the neutron will be absorbed. So, as temperature increases, the effective width of the resonance increases.

So far there is no net increase in absorption, because for every neutron that is now absorbed and that would not have been so previously, there is one that would have been absorbed and that now escapes. The increase in absorption occurs because the resonance has been broadened and because the fuel is heterogeneous, or *lumped.* Neutrons that scatter off ^{238}U only lose a small amount of energy, because ^{238}U is so massive. Moreover, because the fuel and moderator are spatially distinct, these neutrons may very well remain in the vicinity of the ^{238}U and interact with it again. If the neutron's energy relative to the vibrating ^{238}U neutrons is in the vicinity of the resonance, these repeated interactions make it likely that an absorption will ultimately occur.

The Doppler effect occurs on a very short time scale because there is no need for any heat transfer to occur. Fission energy is deposited directly in the fuel. Hence, the fuel temperature is immediately increased in the event of a power excursion. This makes the effect of more value to reactor safety than either the temperature or the void coefficient of reactivity. However, unlike those two effects, the Doppler coefficient is of little use in routine regulation of a reactor, because the temperature change needed in the fuel for the effect to occur is quite large.

Xenon. Xenon-135 has a cross section for thermal neutron absorption of 2.7×10^6 barns. That of ^{235}U is, by way of comparison, only 577 barns. Nuclides with exceptionally high absorption cross sections are called *poisons.* Unfortunately for reactor operation, ^{135}Xe is a by-product of the fission process and, as a result, affects a reactor's dynamic behavior. The time scale is on the order of hours. Figure 10 shows that ^{135}Xe is produced both directly from fission and indirectly from the decay of iodine-135, which is in turn produced from the decay of the fission product tellurium-135. Production from iodine is the dominant of the two processes. (*Note:* ^{135}Te is so short-lived that it is often assumed that ^{135}I is produced from fission at the ^{135}Te yield.) ^{135}Xe is removed by decay to cesium-135 and by neutron absorption (burnup) to ^{136}Xe. The dominant means of removal of ^{135}Xe is by burnup. ^{135}Xe affects all facets of a reactor's operation, including startup, power maneuvers, shutdowns, and restarts.

Assume that an LWR is initially xenon-free. On startup, the ^{135}Xe concentration begins to rise, because it is produced directly from fission and, once an inventory of ^{135}I is established, from that source as well. An equilibrium is eventually established between the production and loss mechanisms. This occurs in about 40 h, with the equilibrium value being given by

$$X_{eq} = \frac{\gamma \Sigma_f \varphi}{\sigma \varphi + \lambda} \tag{29}$$

where X_{eq} equilibrium ^{135}Xe concentration

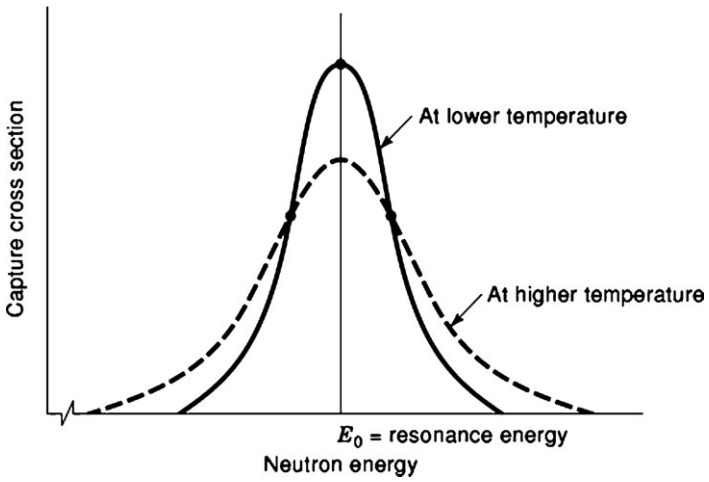

Figure 9. Effect of Doppler broadening on a resonance.

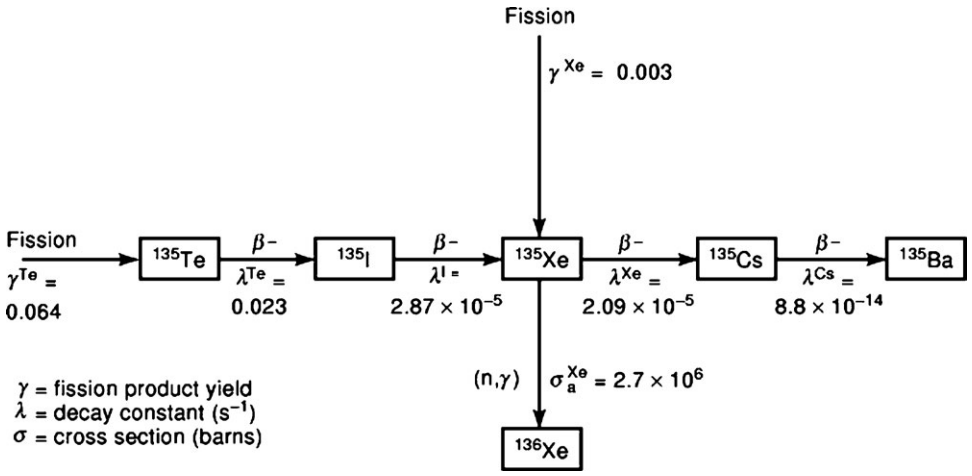

Figure 10. Xenon-135 production and removal mechanisms.

γ fission product yield for ^{135}Xe
Σ_f macroscopic fission cross section
ϕ neutron flux
σ microscopic absorption cross section for ^{135}Xe
λ decay constant for ^{135}Xe

For low neutron fluxes, which correspond to low power levels $\sigma\phi < \lambda$, Eq. (30) indicates that the ^{135}Xe concentration is proportional to the neutron flux. For high power levels $\sigma\phi > \lambda$ and the equilibrium concentration of ^{135}Xe approaches a limiting value. LWR neutron fluxes approach but do not attain the magnitude needed for the ^{135}Xe concentration to reach its limit. Thus, in an LWR, the ^{135}Xe concentration at 50% of full power will be about 70% of the ^{135}Xe value that exists at full power. The principal detriment associated with equilibrium ^{135}Xe is that additional fuel must be loaded in order to provide enough reactivity to offset that associated with the ^{135}Xe. This increases the size of the reactor and requires the presence of additional control devices.

Reactivity transients occur as the result of the changing ^{135}Xe concentrations that follow both power maneuvers and shutdowns. Assume that an LWR is shut down following extended full-power operation. This eliminates the major means of removal of ^{135}Xe which is burnup. However, the major means of production, which is decay from ^{135}I, continues. Production exceeds removal, so the ^{135}Xe concentration rises. This rise continues for about 11 h, at which time the iodine supply is sufficiently depleted so that ^{135}Xe removal by decay to ^{135}Cs becomes dominant. At this time, the ^{135}Xe concentration has peaked. Peak xenon causes operational problems in that, if an unplanned shutdown occurs late in the fuel cycle of an LWR, there may not be enough fuel present to permit restart until the ^{135}Xe peak has decayed. Such reactors are referred to as being *xenon-precluded.*

Further xenon-related reactivity transients occur when a reactor is restarted. The major means of removal (burnup) is instantly resumed. However, the major source of supply (iodine-135) is below its equilibrium value. The concentration of ^{135}Xe therefore decreases to a level below its equilibrium value. While this occurs, the reactor's control devices must be used to insert negative reactivity to offset the positive reactivity being generated by the decrease of ^{135}Xe. This is a counterintuitive maneuver in that control devices are normally withdrawn slowly to offset fuel burnup. Figure 11 shows the ^{135}Xe profile for a power history in which the reactor is started up, operated at full power for 40 h, shut down for 11 h, and then restarted to 50% of

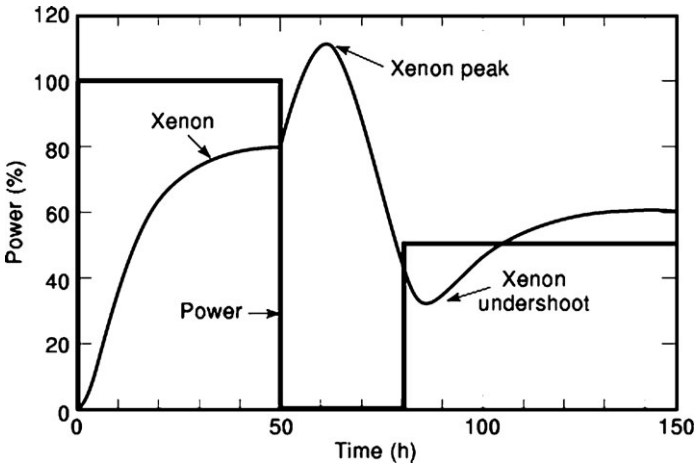

Figure 11. Effect of power history on xenon concentration.

full power. ^{135}Xe first rises to its equilibrium value, then peaks on shutdown, drops below equilibrium on restart, and eventually settles out at about 70% of its full-power value.

Varying ^{135}Xe concentrations can cause spatial oscillations of the neutron flux in large LWRs. Consider an LWR as comprising several sectors, each of which behaves individually subject to the constraint that the total power for the reactor must be kept constant. Suppose the power level rises slightly in one sector. In order for the total power to remain constant, that in the other sectors must be decreased. ^{135}Xe therefore starts to peak in those sectors. At the same time, it is burning out a little faster in the sector that is now at high power. So the change in ^{135}Xe reinforces the power imbalance, and the condition worsens. A flux tilt and possible power peaking problems may result. Balance is eventually restored, because more ^{135}I is generated in the high-power sector and less in the others. The tilt then reverses. Damping can be achieved through appropriate manipulation of the control devices or through the design of negative moderator temperature/void feedback mechanisms that are of sufficient magnitude to override the tilt.

A final operational problem associated with xenon is that it may alter the shape of a reactor's neutron flux distribution. Suppose an LWR is at full power. The shape of the radial neutron flux will be a sinusoid with the peak in the core center. The xenon distribution will have a similar shape. On shutdown, the xenon peaks with the peak being greatest in the core center. So, on restart, the neutron flux will be depressed in the core center and augmented on the core periphery. Control devices on the core perimeter will therefore have greater reactivity worths. The signal seen by nuclear instruments may be similarly affected. That is, those on the core perimeter may read higher than normal and those in the center less for the same actual power production in the core as a whole.

Samarium. ^{149}Sm is another poison. It has only one source of supply, a fission product decay chain, and one means of removal, burnup. The production sequence is <chemStruc id="W5206-cstr-0004" graphicIsNumbered="no" numbered="no">

$$\text{fission} \rightarrow {}^{149}_{60}\text{Nd} \xrightarrow[2\,\text{h}]{\beta^-} {}^{149}_{61}\text{Pm} \xrightarrow[54\,\text{h}]{\beta^-} {}^{149}_{62}\text{Sm}$$

The fission product yield of ^{149}Nd is 0.113. The absorption cross-section of ^{149}Sm for thermal neutrons is 40,800 barns. ^{149}Sm attains an equilibrium concentration about 300 h after startup. This value is independent of the neutron flux and hence reactor power level. On shutdown of the reactor, the ^{149}Sm peaks because production continues and there is no removal. The ^{149}Sm remains at its peak value until the reactor is restarted. It then asymptotically approaches its original equilibrium value.

^{149}Sm does not affect a reactor's dynamics to the extent that ^{135}Xe does, because it has only one means of production and removal. Fuel manufacturers often load ^{149}Sm, which is stable, into the fuel so that its concentration and hence reactivity effect will not be a variable during the initial startup of the core.

LIGHT WATER REACTOR OPERATION

Thus far, the emphasis in this article has been on the attainment of criticality and the dynamic behavior of a reactor once critical. However, the thermal–hydraulic aspects of the reactor also require consideration. The need is for there to be integrated control of pressure and temperature while also observing heatup limits during both reactor startup and while maneuvering at power. The reader may find it useful to consult a reference on reactor design. That by Rahn et al. (7) is suggested.

Pressurized Water Reactor Operation

A distinguishing feature of a PWR is that both the initial pressurization of the plant and the control of pressure during operation are achieved independently of the production of energy from the reactor core. In addition, heatup of the coolant to the normal operating temperature can be done without the use of fission reactor power. As a result, two options exist for the startup of a PWR. The reactor can be taken critical while at low temperature and pressure, and fission energy then be used for plant heatup; or the reactor can be left subcritical until the normal operating pressure and temperature have been established. Both approaches have been used. At present, the latter is favored because it limits the consequences of certain accident sce-

narios. These include a continuous rod withdrawal with the reactor initially at low power, and a loss of pressure. The impact of the first scenario will be less severe under the second startup approach because the magnitude of the moderator temperature coefficient will be greatest with the plant at normal operating temperature and hence negative reactivity feedback will be generated to offset the effect of the rod withdrawal. Also, under the second approach, the safety system's pressure sensors will be operable and will automatically shut the reactor down if pressure is lost. Under the other approach to startup (critical at low pressure and temperature), the pressure sensors would, by definition, be bypassed.

Pressure Control. Plant pressure in a PWR is produced by the pressurizer, which is a tank that is connected to the hot leg and physically located above the reactor vessel. As shown in Fig. 8, the tank is equipped with electric heaters and a spray valve. The heaters are used to keep the water in the pressurizer at saturated conditions. That is, the liquid in the pressurizer is kept boiling so that a liquid–vapor interface exists. In contrast, coolant in the reactor vessel and the primary piping is kept subcooled with no bulk boiling allowed. The volume occupied by vapor in the pressurizer is called the *bubble.* It is this vapor pressure that is transmitted throughout the primary system. In addition to providing the source of pressure for a PWR's primary side, the pressurizer dampens pressure oscillations. For example, if the power is increased, the coolant exiting the core heats up and expands. This creates a surge of water that flows into the pressurizer and compresses the bubble. The pressure rises. This in turn activates the spray valve, which causes some of the vapor to condense, thereby restoring pressure to normal. The reverse occurs on a downpower, only the heaters energize to raise the pressure.

If a PWR is shut down and completely cooled down for maintenance, the pressurizer will be completely filled with water. This minimizes corrosion. Pressure control of the primary system is then achieved by coordinating flow into the piping from the charging system and flow out of the piping through the letdown system. When the pressurizer is completely filled with liquid, it is referred to as being *solid.* The process of reestablishing a vapor–liquid interface is referred to as *drawing a bubble.*

Temperature Control. Under full-power operating conditions, the temperature is controlled by balancing the heat produced from the core with the energy removed by the turbine. If the temperature drops out of a narrow band centered on a designated setpoint, control rods are withdrawn to increase the power. Likewise, if temperature should rise, the rods are inserted. Plant heatup could, as noted earlier, also be done by using the energy produced from fission. The reactor would be critical at a few percent of full power with the energy used to offset plant heat capacity, so that a specified heatup rate is achieved. However, the preferred approach to plant heatup is to operate the primary coolant pumps with the reactor shut down. Friction losses provide a source of energy that raises the plant temperature to its normal operating value in about 6 h.

Reactivity Control. Five mechanisms are available for reactivity control. Of these, the first three are within the purview of the reactor operator. The latter two are design features. The five mechanisms are:

1. Movable Control Rods
   a. Full-Length Shutdown Rods These are normally fully withdrawn at startup and kept withdrawn during operation. Their function is to protect the core against a sudden positive reactivity insertion.
   b. Full-Length Control Rods These are used to create subcritical multiplication and to achieve criticality. They are also used to compensate for the reactivity associated with temperature changes when at power, to maneuver the reactor at up to 5%/min, and to compensate for reactivity changes associated with changes of reactor power.
2. Part-Length Rods These are used for power shaping. They are not inserted on a trip.
3. Soluble Poison (Boric Acid) This is used to control slow, long-term changes of reactivity, including fuel depletion, long-lived fission product buildup, and reactivity effects associated with plant heatup and xenon.
4. Negative Moderator Temperature Coefficient As described earlier, this is a design feature that promotes passive safety under normal conditions and makes reactors self-regulating.
5. Burnable Poison The presence of the soluble poison reduces the magnitude of the negative temperature coefficient and might even result in a positive and hence destabilizing coefficient. In order to reduce reliance on the soluble poison, burnable ones are incorporated in the fuel. These are materials that burn out faster than the fuel and hence allow additional ^{235}U to be loaded without increasing the core size.

PWR Startup

Certain limiting conditions such as observance of a maximum linear heat generation rate, the avoidance of interactions between the UO_2 fuel pellets and the surrounding cladding, and the need to balance power production in the upper and lower portions of the core apply under all operating conditions. During startup, operation is further restricted by the need to observe (1) maximum pressure to avoid brittle fracture of the core vessel, (2) minimum temperature for criticality safety, and (3) minimum pressure for coolant pump operation.

Maximum Pressure to Avoid Brittle Fracture. Reactor pressure vessels are made of carbon steel, which has a body-centered cubic lattice structure and is therefore subject to brittle fracture. That is, should the vessel temperature drop below a certain value and the vessel be subject to a high pressure, the vessel could fail in a brittle or catastrophic manner rather than in a ductile one. The temperature at which the failure mode shifts from ductile to

brittle is called the *nil ductility temperature* or *NDT*. It is a function of the reactor vessel material and the integrated neutron exposure of the vessel's inner wall. The NDT rises as the fluence seen by the vessel increases.

The maximum allowed pressure of the primary system is determined by the ductility of the pressure vessel steel. As temperature decreases, so does the allowed pressure, because the vessel steel is more susceptible to failure at low temperature. Stresses are different for heatup and cooldown, and hence there are two separate curves of pressure versus temperature. The curve for cooldown is more limiting, because under that condition the thermal, pressure, and accumulated fatigue stresses are all tensile for the inner wall of the pressure vessel.

Minimum Temperature for Criticality Safety. This restriction reflects the need to ensure that the moderator temperature coefficient of reactivity is negative prior to the attainment of criticality. Also, if the coefficient is positive, then the reactor must be kept subcritical by an amount equal to or greater than the reactivity that would be inserted should there be a depressurization. (*Note:* There are some exceptions to these rules. The technical specifications of some plants do permit them to be taken critical in the presence of a positive moderator coefficient provided that the magnitude of the coefficient is small. All plants must have a negative moderator temperature coefficient at their normal operating temperature.)

The moderator temperature coefficient of reactivity may not be sufficiently negative at temperatures below a certain minimum, because the coefficient of thermal expansion of water is small at low temperature and because of the presence of the soluble poison.

Minimum Pressure for Coolant Pump Operation. The net positive suction head (*NPSH*) is defined as the pressure at the pump suction position less the saturation pressure that corresponds to the temperature of the fluid being pumped. If it is sufficiently positive, a centrifugal pump can operate. Otherwise, water will flash to steam in the eye of the pump and cause the pump either to cavitate or to become gas-bound. Hence, there is a certain minimum pressure that must be maintained if the primary coolant pumps are being operated. This minimum pressure increases as plant temperature rises, because the saturation pressure rises with temperature.

Figure 12 shows the pressure–temperature limits for a PWR. There are two curves showing the maximum pressure to avoid brittle fracture, one for heatup and one for cooldown. The primary system pressure must be kept below the appropriate curve at all times. The brittle fracture curves will shift to the right and become more restrictive as the pressure vessel incurs additional neutron damage.

There is one curve for the minimum temperature for critical operation. If the reactor is critical, then the system temperature must be maintained to the right of the curve so that the moderator temperature coefficient of reactivity will be sufficiently negative. As a result, the consequences of both depressurization and continuous-rod-withdrawal accidents will be less severe.

There is one curve showing the minimum pressure at which the reactor coolant pumps may be operated. The plant pressure should be kept above this curve.

Also shown in Fig. 12 is a typical trajectory for a plant startup in which the plant is taken from a cold shutdown condition with a solid pressurizer to a hot operating condition. Table 4 lists the principal steps in the startup of a PWR.

Table 4. PWR Startup Sequence

1. Perform precritical checklists.
2. Adjust steam generator water chemistry.
3. Coordinate pressure–temperature control of primary via charging/letdown systems.
4. Energize pressurizer heaters and open spray valve to heat up pressurizer.
5. Charge nitrogen gas to pressurizer to create a gas-filled volume that will dampen pressure pulses.
6. Start primary coolant pumps if above NPSH curve. Raise primary temperature using pump heat until above minimum temperature for the attainment of criticality.
7. Adjust chemistry of primary coolant.
8. Reduce concentration of soluble boron so criticality can be attained.
9. Draw a steam bubble in the pressurizer.
10. Drain steam generators so level is within operating range.
11. Recirculate condenser hot well to eliminate stratification.
12. Use friction losses from the primary coolant pumps to heat the primary at the specified rate, usually 50°F/h (28 K/h). Drain coolant to compensate for density changes.
13. Take reactor critical by first withdrawing the full-length shutdown rods and then the full-length control rods. The latter are used to create subcritical multiplication.
14. Position control rods for a 0.5 decade per minute startup rate. Take critical data at 10^{-8} A. Raise power to point-of-adding-heat (1% to 3%).
15. Admit steam to the secondary by opening the bypasses on the main steam isolation valves. Commence warmup of turbine.
16. Raise power to 10% and open main steam stops.
17. Place steam generators and secondary feed pumps in automatic.
18. Adjust control rods and soluble boron to compensate for buildup of short-lived fission product poisons such as xenon.
19. Adjust control rods and soluble boron to compensate for fuel depletion. The full-length control rods are left partially inserted to provide a means to alter reactor power in response to changes of load. Use part-length rods to correct for power distribution anomalies.

BWR Operation

Figure 13 is a schematic diagram of a boiling water reactor. Comparison of this diagram with that of a PWR shows several differences that affect reactor operation. First, BWRs are direct cycle plants. The steam that drives the turbine is produced by boiling action within the core. There is no use of a primary loop that transfers heat from the fuel to a secondary system via steam generators. Hence, BWRs can operate at lower pressures than PWRs while still achieving the same thermal efficiency. Second, coolant flow through the core is a mixture of both natural and forced circulation. The density difference between the liquid–vapor mixture in the core and the liquid in the downcomer (the region on the core perimeter where the jet pumps are located) causes flow. This flow can be enhanced by the recircula-

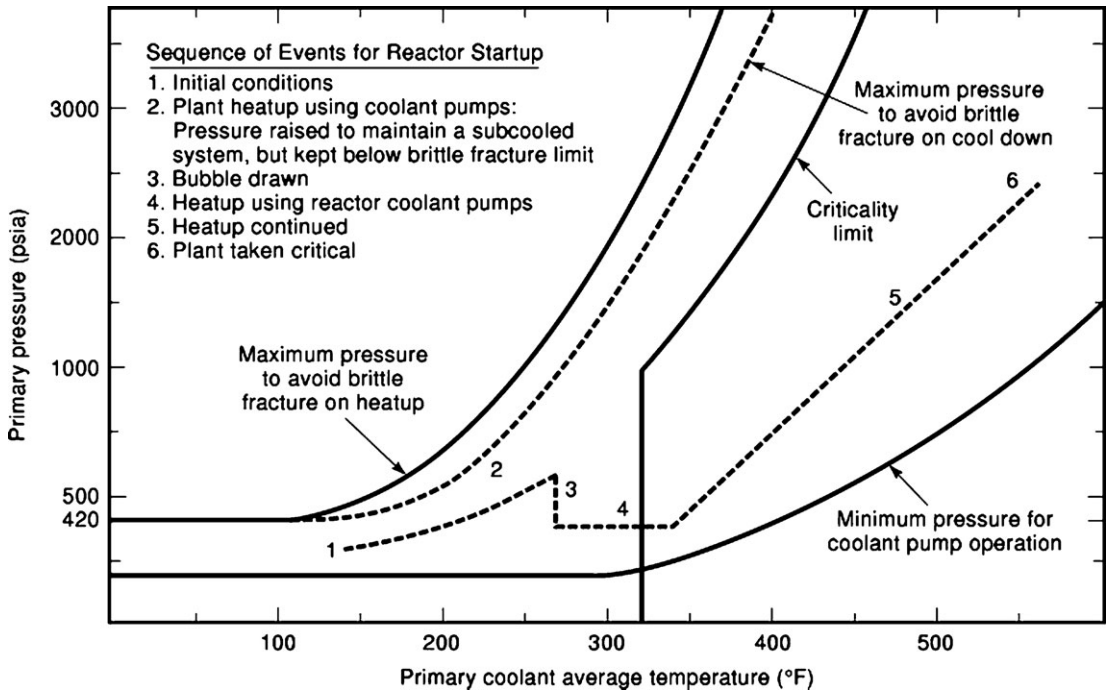

Figure 12. Pressure–temperature limits for a PWR.

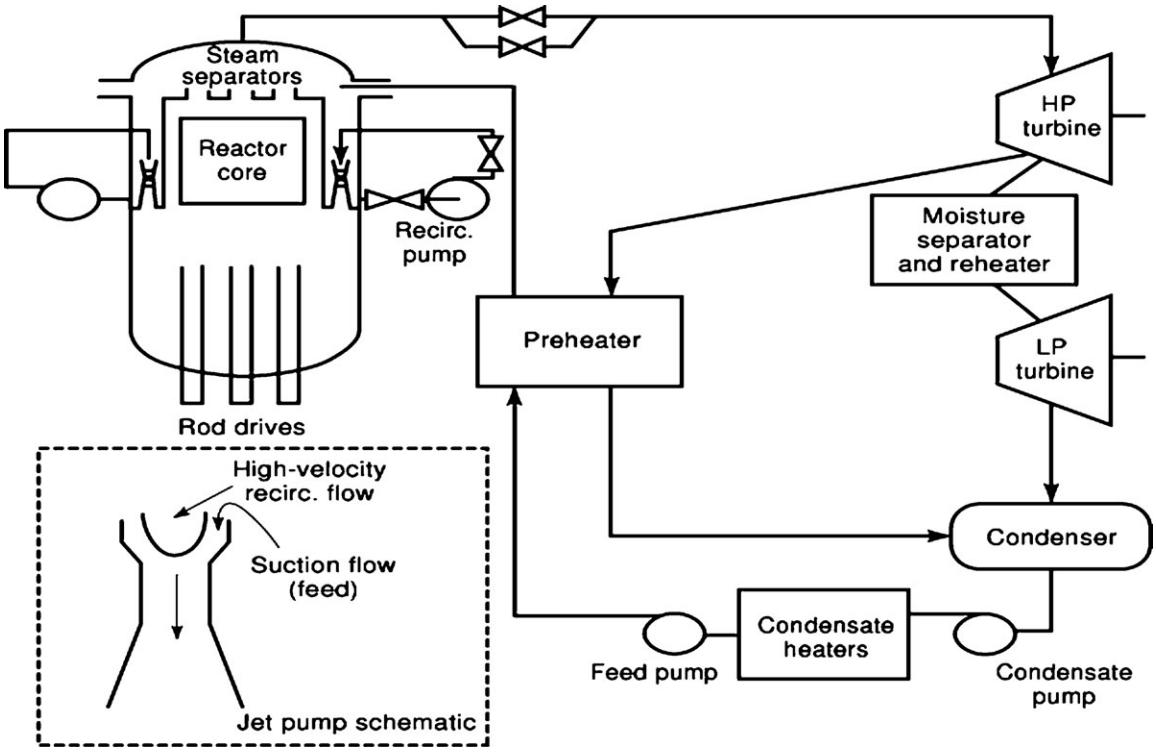

Figure 13. Schematic of a boiling water reactor.

tion pumps, which drive water through the jet pumps. A region of low pressure therefore exists at the jet pump suction, so that feedwater combines with liquid from the steam separators and flows through the core. A third difference between BWRs and PWRs is that the source of pressure in a BWR is the boiling that occurs in the core. Thus, the option does not exist in a BWR to achieve normal operating temperature and pressure and then take the plant critical.

There are two mechanisms available for reactivity control in a BWR:

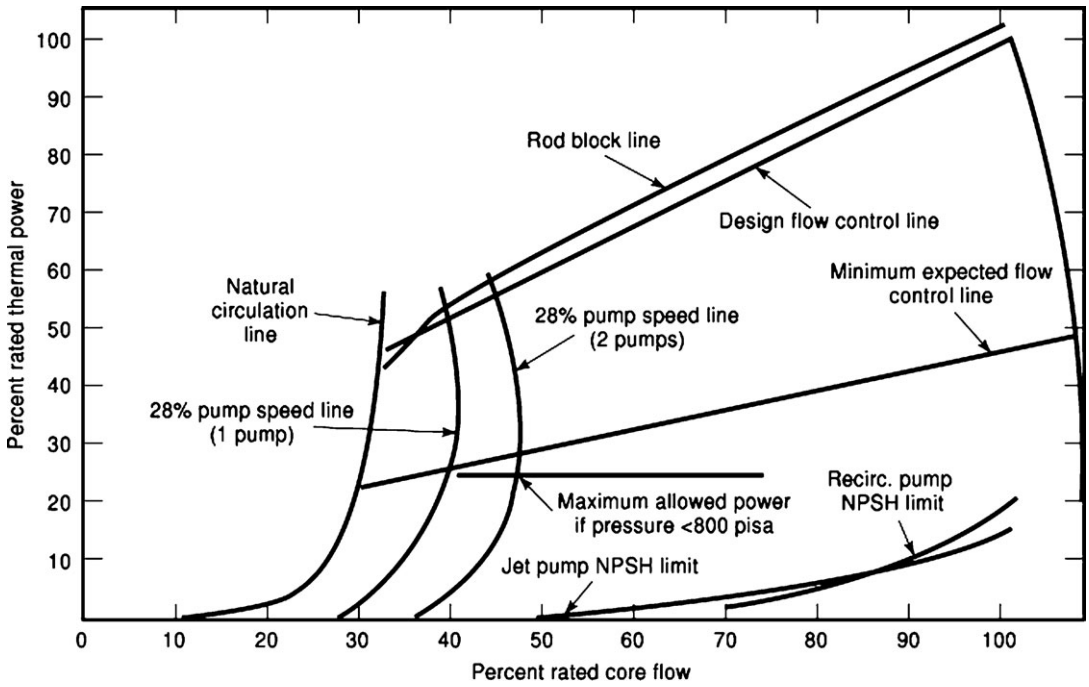

Figure 14. BWR thermal power and core flow restrictions.

1. Movable Control Rods One cruciform-shaped rod is present for every four fuel assemblies. This large number of rods provides a means to ensure that each set of fuel assemblies is properly operated. The boiling process could result in some assemblies producing too much power and others too little. In a PWR, the rods are driven from above the core. This is not possible in a BWR, because of the steam separators. So in a BWR the rods, which contain boron carbide, are driven upwards from the bottom of the core. Power changes in excess of 25% are accomplished with these rods.
2. Negative Moderator and Void Reactivity Coefficients The negative moderator and void reactivity feedback coefficients are within the control of the operator. They are used to accomplish power changes of as much as 25% of rated. An increase in the recirculating pump speed increases the flow rate through the core, which in turn increases heat transfer and decreases the vapor fraction. Therefore, the average density of the liquid–vapor mixture that is flowing through the core increases, and so does the neutron moderation. The net effect is that positive reactivity has been added to the core. The reactor power rises until increased vapor formation and higher fuel temperatures reverse the process. The reactor then settles out, so that it is critical but at a higher power level. A decrease in the recirculation flow will have the opposite effect.

BWR Startup

The major steps in the startup of a BWR are:

1. The recirculation pumps are started and set at the desired speed, 28% of rated for the plant shown in Fig. 14.
2. The reactor is made critical by the withdrawal of the control rods.
3. The control rods are further withdrawn, thereby placing the reactor on a positive period. The power is raised to above the point-of-adding-heat, and the core vessel is pressurized. In accordance with the safety limit specifications, power may not be increased above 25% until the plant pressure exceeds 800 psia.
4. Once the pressure rises above 800 psia, the reactor power is raised to 55% of rated. A check is then made of the estimated critical position and the rod pattern.
5. The recirculation pump speed is then increased so that total core flow rises to 100%. As it does, the core power also rises because of the moderator-temperature–void reactivity coefficient. Power is then leveled at the desired value.

This procedure takes approximately 18 h to complete. Not mentioned in the above summary are many other significant factors such as heatup limitations associated with thermal stress.

Figure 14 shows the relations between thermal power and core flow that must be observed for a boiling water reactor. The initial conditions for startup are 0% power and 36% flow. The percentage of flow through the core exceeds the recirculation pump speed (28%) because of the effect of the jet pumps. Once critical, the plant is coordinated so that the power and flow move along the 28% pump speed line. The recirculation pump speed is kept constant while moving along this line. The reason for the increase in flow as

the power rises is natural circulation, which is also termed thermally induced flow. The line marked "maximum allowed power" must not be crossed until the plant pressure is above 800 psia. The power is then raised to 55% of rated. The corresponding core flow is 45%. This point is the intersection of the 28% pump speed line and the design flow control line. The control rods and the recirculating pump speed are then adjusted as needed to move the plant's power and flow along the design flow control line until conditions of 100% power and 100% flow are attained.

Also shown in the figure are net positive suction head limits for the jet and recirculation pumps and the rod block line. The latter represents the combinations of core flow and power for which control rod withdrawal is prohibited automatically.

Certain combinations of core flow and thermal power have been associated with oscillatory behavior in BWRs. Accordingly, the operating curves shown in Fig. 14 have been modified to preclude or at least restrict operation in the knee of the curve. Analysis of the factors involved in this instability is given by Tong and Weisman (8).

BIBLIOGRAPHY

1. A. F. Henry *Nuclear Reactor Analysis*, Cambridge, MA: MIT Press, 1975.
2. J. J. Duderstadt L. J. Hamilton *Nuclear Reactor Analysis*, New York: Wiley, 1976.
3. J. A. Bernard A. F. Henry D. D. Lanning Application of the "reactivity constraint approach" to automatic reactor control, *Nucl. Sci. Eng.*, **98** (1), 87–95, 1988.
4. J. A. Bernard Formulation and experimental evaluation of closed-form control laws for the rapid maneuvering of reactor neutronic power, Report MITNRL-030, Massachusetts Inst. Technol., 1989.
5. J. A. Bernard and L. W. Hu," Dynamic Period Equation: Derivation, Relation to Inhour Equation, and Precursor Estimation," *IEEE Transactions on Nuclear Science*, Vol.**46**, No. 3,Part I, June 1999, pp 433–437.
6. J. A. Bernard D. D. Lanning Considerations in the design and implementation of control laws for the digital operation of research reactors, *Nucl. Sci. Eng.*, **110** (1), 425–444, 1992.
7. F. J. Rahn *et al. A Guide to Nuclear Power Technology*, Malabar, FL: Keiger, 1992.
8. L. S. Tong J. Weisman *Thermal Analysis of Pressurized Water Reactors*, La Grange Park, IL: American Nuclear Soc., 1996.

JOHN A. BERNARD
Massachusetts Institute of
Technology, Cambridge, MA