Chapter 5 - Safety Systems

Introduction - Special Safety Systems Functions

In previous chapters we have referred to the four safety functions required in a nuclear reactor:

- shut down the reactor
- remove decay heat
- contain any radioactivity
- monitor the state of the plant.

In this chapter we shall describe the major systems that perform these functions. We shall concentrate on CANDU for our examples, although other reactor types have similar systems.

Shutdown Systems

Shutdown is one of the most important safety functions in a reactor because it reduces the amount of energy that has to be removed from the fuel after an accident. It is usually accomplished through rapid insertion of a neutron-absorbing material into the core. Another way is to remove from the core material which is essential to the chain reaction - e.g. the moderator. There are more radical concepts possible in principle, such as removing fuel or changing the core geometry, but they are not in widespread use for fast shutdown.

Before a shutdown system is designed, the requirements should be defined (although again historically, the two went along together). Here are some of the questions that must be asked, and answered:

- how do we get negative reactivity into the core?
- how fast does the system have to act, once it receives a signal?
- how much reactivity *depth* must it have (how many negative milli-k?)
- how reliable must it be?
- what are the acceptance criteria?
- what sort of signals are available and practical to trigger the shutdown system for each accident?
- what sort of environment must the shutdown system be designed to withstand?
- how do we ensure that a fault which could affect the control system or a shutdown system does not affect both? Or both shutdown systems?
- how do we know the systems will work as designed?
- how does the operator know the system has been required, and that it has worked?

We shall cover these topics in turn. Many of the questions are common to other safety systems, so we shall explore them in more detail for the first time here, and just refer to them later on.

Mechanical Design

The most basic part of shutdown design, and the most common mechanism, is inserting a neutron absorber onto the core. Because modern reactors are large, a single absorbing device is generally not sufficient. Almost all reactor types use some form of absorber rod, multiples of which are inserted vertically into the core. Actually most of them are tubes, not rods, but are called rods for historical reasons. In many cases, and especially for non-pressurized reactors such as research reactors, they are inserted from the top, so that gravity can assist them; however in some Boiling Water reactors (BWRs), they are inserted from the bottom. (This has the disadvantage of an unsafe failure mode if the rod should fail mechanically since it can fall out of the core under gravity; such a "rod drop" accident is therefore part of the BWR Design Basis.)

In most reactors in the world, the rods do double duty - being driven in and out of the core for control purposes, and being driven or dropped in rapidly for shutdown purposes. CANDU however separates the control rods from the shutoff rods, as one of the lessons learned from the NRX accident.

Other shutdown mechanisms exist. For example, a second means of shutting down BWRs is by tripping the coolant recirculation pumps - without forced flow, the amount of boiling in the coolant increases and the negative void reactivity shuts down the reactor. If you visit the zero-energy Pool Test Reactor (PTR) reactor at CRL (it is now defuelled - but the pool and the reactor structure are still there), you will be told that the primary means of shutdown was by rods. However if you are observant you will notice a cupboard outside the reactor hall, labelled "Boron for PTR". This was the "ultimate shutdown system" - if for some reason the rods were not able to shut down the reactor, someone - no doubt a graduate student - would grab a bucket of boron from the cupboard and throw it into the pool. A similar "doomsday" shutdown system existed for early gas-cooled reactors in the U.K. - it injected boron *dust* into the core if the primary shutdown system failed. Since the boron dust could never be removed in entirety from the core structure, it would not be used more than once. Later on (removable) boron balls were used, poised above the core.

Paradoxically, a 'doomsday' shutdown system is *not* obviously safe. If its action causes severe economic harm, an operator will be very reluctant to use it, and may even disable it if he fears it will go off spuriously. This is true of all safety systems - if their action wrecks the plant, one risks having them jumpered out (deliberately disabled).

The earliest CANDU shutdown system was moderator dump - large valves at the base of the calandria would open and the moderator would drain out of the calandria vessel under the action of gravity - a bit like the ZED-2 design. The system would be re-poised by closing the valves and pumping the moderator back into the calandria. NPD, Douglas Point and Pickering-A used this system. Pickering A also used a few shutoff rods (recall that designers lost confidence in rods after the NRX accident, and they were not used again until Pickering-A). The moderator dump system was (and is) highly reliable but is somewhat slow compared to shutoff rods, especially for larger cores, and in addition removes a source of water surrounding the fuel channels, which could be used in an emergency (we will discuss this later). A variation of moderator dump is reflector dump - the heavy-water reflector on the Maple reactor series is dumped as part of the shutdown safety system.

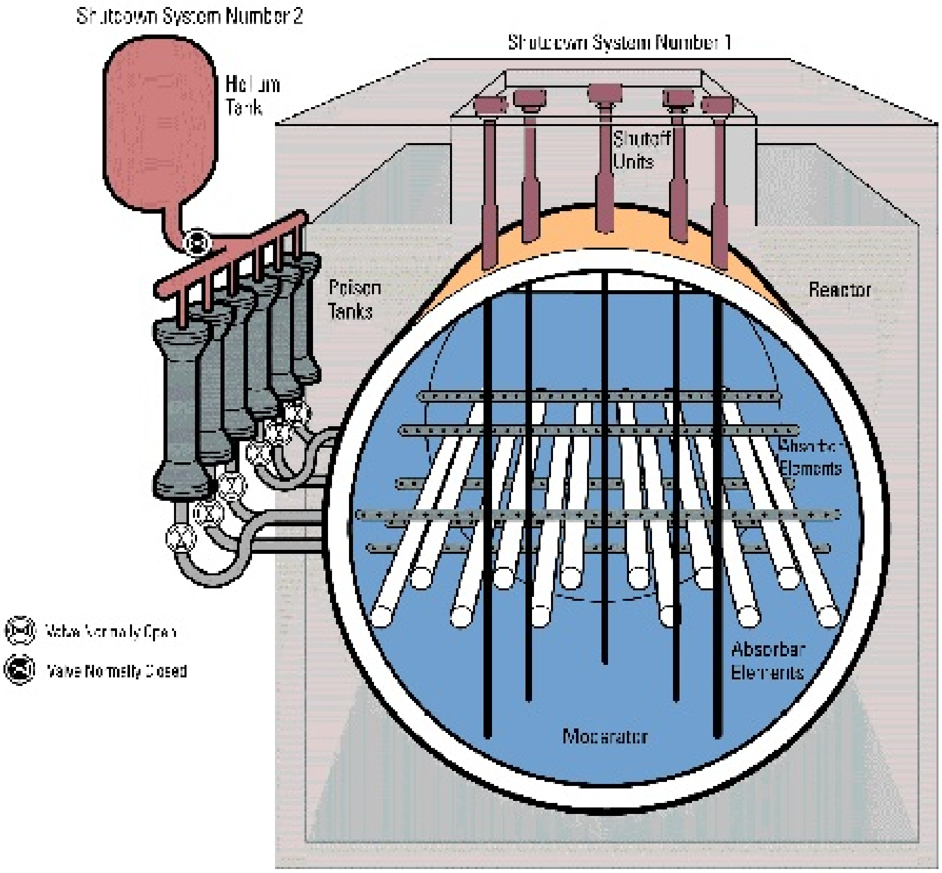

Figure 5-1 - Shutdown Systems #1 and #2

Modern CANDUs have two separate shutdown systems - rods and poison injection. Specifically, Shutdown System #1 for CANDU Classic consists of 26 or 28 shutoff rods, normally suspended above the core, and released on a signal. They act in the moderator, between the rows of pressure tubes, as can be seen in Figure 5-1. They would fall in by gravity, but to give an initial boost to the speed, they are spring loaded, which accelerates them over the first few feet of travel. Mechanically, the rod is suspended on a cable running over a pulley, which is released by a clutch and wound back up by a motor. The rod itself falls into a perforated guide tube within the moderator, whose purpose is to make sure the rod falls straight in and doesn’t tip over or snag.

Shutdown System #2 consists of a liquid neutron absorber (gadolinium nitrate) which injects directly into the moderator water through perforated metal tubes, also as indicated in Figure 5-1. The liquid is accelerated by gas pressure from a common helium tank, pressurizing one poison tank per nozzle, as shown in the figure. The system is actuated by opening the fast acting valves connecting the helium tank to the poison tanks. A disadvantage of this system is that because it injects into the moderator water, it must all be removed chemically (by ion exchange) before the reactor can start up again - a process that takes almost two days. Note that there are no closed valves between the poison tanks and the moderator, for reliability reasons; the consequence however is that poison gradually diffuses toward the core down the pipe, which must therefore be back-flushed from time to time to drive the diffusion front away from the moderator.

Some light-water reactors have a poison injection system which acts as a backup to the combined control/shut-off rods. It injects boron into the reactor coolant system. Typically it is slow and not as effective as the rods, since it needs to flow into the core to work, and the concentration has to be sufficient to counter dilution in the entire coolant system.

Speed

The designer of a shutdown system must know its speed, specifically its required insertion time, particularly the time when it first starts to "bite"[a]. Not surprisingly, the required speed is set by the fastest accident. In CANDU Classic, this is the large loss of coolant accident, which inserts *positive* reactivity at a rate of about 4 mk./sec. The main safety requirement is to prevent melting of the central part of any fuel pin due to the overpower, since significant amounts of molten fuel could risk failing the nearby pressure tube. Induced failure of another pressure boundary component is not acceptable, since if one pressure tube fails, a number of pressure tubes at similar conditions could fail, and one would risk loss of the calandria vessel integrity. It is an easy exercise for you to calculate the allowed energy addition before the fuel melts. It turns out that as long as the **net** positive reactivity is kept below about 6 mk., depending on the design and the assumptions, the energy is not sufficient to melt the centre of the fuel. That then suggests the shutdown system has to start to bite in about a second, and that the initial reactivity insertion rate has to overcome the 6 mk already inserted, plus turn the transient over by 1.5 seconds - in other words, tens of (negative) mk./sec. A subsidiary safety requirement is that the fuel geometry must remain undistorted enough that it can be cooled by ECC later in the accident.

Typically the shutdown system for a large LOCA in CANDU Classic is triggered by either of two very fast signals: high neutron flux, or high rate of rise of neutron flux (log-rate). Trips on low flow in the heat transport system, and high containment pressure, are also triggered but not

[a] First insertion of significant negative reactivity

usually as quickly. One of the perceived disadvantages of CANDU is its positive coolant void coefficient (reactivity and power rise when the coolant voids); the offsetting advantage is that the rise in neutron flux provides very sensitive rapid signals to detect the LOCA.

The Advanced CANDU Reactor (ACR-1000) has a nominal and small negative full-core coolant void reactivity. While the transient void reactivity may become positive for a short while during a large LOCA transient[b], the main signals to trip the shutdown systems are process signals such as low coolant flow and low coolant pressure; clearly the response of the shutdown systems does not have to be quite as fast as for CANDU Classic.

The time at which the shutdown system begins to bite is composed of several components (we'll use the shutoff rods as an example):

Time of "bite" =
+ time of large break[c]
+ time for measured signal to rise to trip set-point
+ response time of detector and amplifiers
+ response time of instrumentation which decides if a signal has passed its set-point
+ response time of trip relay chain
+ time to release clutch holding shutoff rod in place
+ time to accelerate shutoff rod from parked position to ~ the first row of fuel channels

For the poison injection system, the last two items are replaced by:

+ time to open valves connecting helium tank to poison tanks
+ time to accelerate poison through the nozzles and across a couple of lattice pitches within the moderator.

In the end, systems which act within a second are practical but require careful design and maintenance to get them fast enough.

[b] This is due to asymmetric voiding of the two core passes in the broken coolant circuit, so that adjacent channels void at different rates, the so-called "checkerboard void" effect. Explaining why this may cause a positive reactivity transient, especially at the end of pressure-tube life when the tube has crept under stress and irradiation, is well beyond the scope of this course.

[c] Traditionally the time of the accident is taken as t=0. For large LOCA it is assumed the break occurs and grows to full size almost instantaneously, which is a major conservatism. See Ref.[1] for further discussion.

Reactivity Depth

Reactivity depth means the total negative reactivity inserted once the shutdown system has fully operated. For shutoff rods, this occurs when the rods are fully inserted; for poison injection, when the poison is fully injected and mixed with the moderator. The reactivity depth requirement is set, obviously, by the accident which inserts the greatest *positive* reactivity. For CANDU Classic, this is not the large LOCA, but a small LOCA, specifically a pressure-tube break followed by an assumed calandria tube rupture. Why? When a reactor is operating at full power, there is a negative reactivity load due to xenon, a neutron absorber which is formed from the decay of iodine, which in turn is formed from the fission product tellurium, as follows:

$$^{135}\mathrm{Te} \Rightarrow_{<1\mathrm{min}} \Rightarrow {}^{135}\mathrm{I} \Rightarrow_{6.7\mathrm{h}} \Rightarrow {}^{135}\mathrm{Xe} \Rightarrow_{9.2\mathrm{h}} \Rightarrow {}^{135}\mathrm{Cs} \Rightarrow {}^{135}\mathrm{Ba} \text{ (stable)}$$

where the times represent half-lives. The xenon load in a CANDU at full power is about -25 mk and has to be compensated by positive reactivity from the fuel to keep the reactor critical. When the reactor is suddenly shut down, the xenon decays slower than the iodine, so that the absorption due to xenon initially increases to about -40 mk., then decreases as the iodine decays away (Figure 5-2). After a long shutdown, the xenon load is small; so that poison is added to the moderator to compensate for the absent xenon load (since the normal control system does not have the required negative range). As the reactor comes back to power again, the iodine and thus the xenon gradually builds up, and the poison can be chemically removed from the moderator (or sometimes a burnable poison is used - i.e., one which is made into a less absorptive species over time by neutron absorption - typically gadolinium). If a pressure-tube break with assumed failure of the surrounding calandria tube occurs during this startup, then the "poisoned" heavy-water moderator could be replaced by un-poisoned heavy-water coolant. The amount of positive reactivity that could be added in, say, the first 15 minutes is due to:

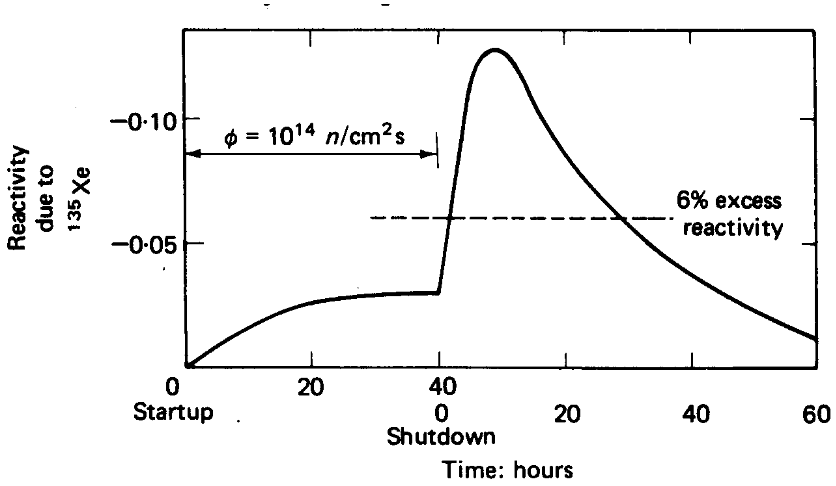

Figure 5-2 - Xenon Reactivity (negative) versus Reactor Power History

- coolant void in the fuel channels, from the fraction of coolant that is lost through the break
- fuel temperature, due to cool-down of the fuel after trip (remember the fuel temperature reactivity feedback is negative from zero to full power)
- ‘clean’ coolant displacing ‘poisoned’ moderator
- increase in moderator temperature (a positive coefficient in CANDU, but small and slow), due to the hot discharging coolant mixing with cold moderator
- decay after shutdown of any xenon that has been formed during power operation

Offsetting this is the negative reactivity due to shutdown, and eventually the very large negative reactivity due to injection of emergency coolant light water. The analysis of the reactivity balance is usually done 15 minutes after the first clear signal of the pressure tube break, when the operator can be assumed to supplement the shutdown reactivity (e.g., by manual poison addition). One then designs the shutdown depth of the shutdown system to achieve a shutdown margin[d] of at least -5 mk. if all the assumptions are obviously conservative, adding more margin if there are uncertainties. Physically one achieves greater shutdown depth by adding more rods, by putting rods preferentially in the high-flux region of the core, etc. There is a practical space limit however, on the number and location of rods. The typical shutdown depth for CANDU is -70 mk. for the shutoff rods and < -200 mk. for poison injection. Shutdown margins for the rods are obviously smaller, typically -5 to -10 mk under the most pessimistic conditions (e.g., assuming the two most effective rods are unavailable).

For the ACR, an in-core break causes an inherent shutdown because of the replacement of the heavy-water moderator by the light-water coolant, which latter is more of a neutron absorber. Shutdown depth is set by the loss of reactivity control accident - i.e. by the amount of positive reactivity that can be added by the control devices.

Unavailability

We have touched on shutdown system unavailability. Recall that this is a demand unavailability and is expressed in dimensionless units, although you will also see it written as hours/year or years/year. The required unavailability is set primarily by the safety goals applied to the reactor by the regulator, and for CANDU is at most 1 failure per 1000 demands (10^{-3} per demand). Experience has shown that the shutdown systems meet about 10^{-4}. Unavailability values much lower than this are theoretically possible, but are usually not credited for a single system, because of:

- the number of tests it would take to establish such a unavailability
- the suspicion that unknown cross-link effects and common-cause failures give a

[d] Shutdown margin = |Total positive reactivity| - |Total negative reactivity|

lower limit to the unavailability of a single system

U.S. reactors typically use 10^{-4} for their shutdown system, based on a prediction done during design plus testing, but they do not test on-line at a sufficient frequency to establish it.

If each shutdown system has an unavailability of 10^{-3} per demand, then one is tempted to say that the unavailability of both systems together is 10^{-6} per demand. This is *only* true if the systems are sufficiently independent and diverse so that they are not subject to common cause failures. We shall return to this later on.

Finally recent CANDUs have used software for both SDS1 and SDS2. Software failure modes can be subtle, so that while testing in operation is important, much more emphasis has to be placed in the design of the software to ensure a rigorous development process, diversity in development tools and platforms, simple programme logic, and strict independence between the software engineers responsible for designing each shutdown system to avoid common cause failures. In CANDU the control system software is completely independent of and separate from the shutdown system software; they run in independent hardware platforms.

Acceptance Criteria

It is not enough to shut the reactor down - one must do so in a timely fashion. We have discussed the two accidents which set the rate and depth of the mechanical and hydraulic design. However, this does not guarantee that for other accidents the system is triggered at the right point - that is a function of the trip signals and trip signal set-points, discussed below. Trip signals are chosen to meet *acceptance criteria*, chosen usually so as not to place an undue burden on the other safety systems - in other words, to prevent or postpone consequences. For example, for accidents which are expected to happen perhaps once or more in the plant lifetime, the shutdown systems should act early enough to prevent fuel failure - thus avoiding a challenge to containment, as well as to the pocketbook of the operating organization. This category would include loss of Class IV power, very small LOCA or a failure in the reactivity control system. For loss of heat sink accidents, the shutdown systems should act soon enough to give the operator lots of time to bring in a backup heat sink - typically 15 to 30 minutes, although in modern designs one aims for at least 8 hours before operator action is required. For accidents such as large LOCA, one should prevent fuel damage early on (to allow time for the containment to isolate) and to ensure that the fuel is not made so brittle from the zircaloy-steam reaction that it cannot be cooled[e], as mentioned earlier. Of course limiting fuel damage is a job shared between shutdown and ECC in a large LOCA - the job of the shutdown system is to deliver the fuel to the ECC in a reasonable

[e] At temperatures reached in large LOCA, the Zircaloy fuel sheath can chemically react with the steam in the channel to produce zirconium oxide, heat and hydrogen.

condition so that ECC can remove decay heat from a known fuel geometry. In any accident, shutdown systems should act early enough so that there is no risk to the pressure boundary (or at least no risk additional to that from the initiating event). Hence preventing fuel melting in CANDU (other than for single channel events) is taken as a conservative surrogate to prevent fuel channel failure.

Signals

For a shutdown system to be effective, it must detect an accident soon enough that the acceptance criteria are met. We can infer some of the commonly-used trip signals from the types of accidents (the list below is exemplary for CANDU Classic and not intended to be complete):

Accident	Symptoms	Typical Trip Signals
Loss of reactor power control	Reactor power rises	**High neutron flux**
	Reactor power rises rapidly	**High log rate of neutron flux**
	Heat transport system pressure rises	**High heat transport system pressure**
Loss of forced circulation	Coolant flow drops	**Low heat transport system flow**
		Low core pressure drop
	Pressure rises	**High heat transport system pressure**
	Reactor power rises[f]	**High neutron flux**
Large loss of coolant	Reactor power rises	**High neutron flux**
	Reactor power rises rapidly	**High log rate of neutron flux**
	Containment pressure rises	**High containment pressure**
	Coolant flow drops	**Low heat transport system flow**
	Coolant pressure drops	**Low heat transport system pressure**
Small loss of coolant	Pressurizer level drops	**Low pressurizer level**
	Coolant flow drops	**Low heat transport system flow**
	Containment pressure rises	**High containment pressure**
	Moderator level rises[g]	**High moderator level**
	Coolant pressure drops	**Low heat transport system pressure**

[f] Why?

[g] For an in-core break only (pressure-tube / calandria-tube rupture)

Accident	Symptoms	Typical Trip Signals
Loss of feedwater	Boiler level drops Feedwater flow drops Heat transport system pressure rises	**Low boiler level** **Low feedwater flow** **High heat transport system pressure**

Canadian practice has been to sense an accident with at least two diverse trip parameters on each shutdown system - typically chosen from the list above. It is not always practical to do so. Moreover Probabilistic Safety Analyses show little benefit of the backup trip. The current trend for new designs in Canada is to de-emphasize backup trips where a clear direct signal exists - as per Ref. [2]. Note that a direct trip parameter is “A value based on direct measurement of a specific challenge to the derived acceptance criteria and, if applicable, a direct measure of the event.” - e.g. reactor power for large LOCA, low flow for loss of Class IV power.

Manual trip is permitted if the time-scales are long - typically fifteen minutes from the first clear signal of the accident in the Main Control Room - and if there is no practical alternative. Even so, it is usually not relied on as a primary trip.

Operating Environment

It seems obvious that if we rely on a safety system to mitigate an accident, it should not be disabled by the accident. Behind this simple statement is a large amount of work to:

- define the conditions which could affect the safety system
- design it to withstand them.

In addition it is not always possible to meet this requirement for a single system - for example a major fire in or near the Main Control Room would require shutdown (because it could affect the control computers) and at the same time *possibly* damage some of the components of Shutdown System #1 so that it would not respond (although it would very likely fail safe).

The shutdown mechanisms in CANDU act mostly within the moderator, which protects them from some of the effects of accidents. But not all. For example, we must design so that:

- there are no high-energy pipes within striking distance of the reactivity mechanisms deck on top of the reactor, where the shutoff rod clutches and pulleys are located
- an in-core break cannot disable shutoff rods to the extent that the system does not meet its reactivity depth requirements (it is not possible to protect **all** the rods from pipe whip and jet forces from an in-core break)
- shutdown system cables and instruments are separated to the extent practical so that a local fire will not incapacitate *both* shutdown systems. This is an example of where absolute protection is not possible - one can only separate cables so far,

and a large fire in an area could disable a shutdown system. The approach then is to ensure that it cannot disable *both* shutdown systems, which is done by placing them in widely separated areas (90° to 180° separation) of the plant. This is an example of the "two-group" philosophy, which we shall cover later.

- the steam, high temperatures, water, and high radiation fields from an accident must not prevent a shutdown system from firing when needed (once it has fired, however, such protection is no longer needed).
- the shaking due to an earthquake must not prevent the shutdown system from actuating, nor slow it down so severely that it cannot meet its acceptance criteria.

Common Cause Failures

We gave an example of where a single cause (fire) could disable more than one system. This is a serious challenge to the protection provided by seemingly independent systems. One must identify all common cause failures and design against them; however to cover off the possibility that we just aren't clever enough to anticipate them all, a *two group* philosophy is followed in CANDU Classic. In summary this philosophy is:

- for each failure, ensure that there are at least two ways of performing the required safety function
- separate these two ways geometrically (so that they are not subject to local damaging hazards such as fire or turbine missiles or aircraft crash)
- use diverse equipment and diverse means of operation
- protect them against the environmental results of the failure
- qualify or protect at least one of the two systems against plant-wide external events such as tornadoes and earthquakes.

Like the fate of all ideals, the practical application involves many tradeoffs: diversity is not always possible; and there is only so much space within which systems can be separated. Thus whereas ideally one would like to route control system cables separate from shutdown system #1 (SDS1) cables *and* shutdown system #2 (SDS2) cables, it would be almost impossible from a plant layout point of view; so one allows *grouping* of control system and SDS1 cables, but they must be *separated* from SDS2. However all 3 logic channels of any given group must be separated, so that a local cable fire cannot generally disable all three channels of any one system. Even this must be qualified - a fire in the SDS1 cabinet in the Main Control Room could disable all 3 SDS1 channels, since they *must* come close together at some point to "vote" on a 2-out-of-3 basis for tripping the reactor. There are further compromises forced by the system operation: a lot of instrumentation *must* be on the reactor, and there the systems come in closer proximity. For example the neutron flux measuring devices of the control system, SDS1 and SDS2 all are in the reactor; the compromise is that the control system and SDS1 devices are allowed to be in close proximity, but SDS2 devices must be on the opposite side of the reactor.

The general rule is that if a compromise in separation must be made, then one must show that one or more of the following applies:

- there is no credible hazard in the area
- another Group 2 system outside the area will mitigate the event
- the system or component is protected by a barrier
- the system or component is fail safe
- the component designed to withstand the hazard

Here is an example of how grouping and separation was implemented on CANDU 6:

Safety Function	**Group 1**	**Group 2**
Shutdown	Reactor Control System Shutdown System 1	Shutdown System 2
Heat Removal From Fuel	Heat Transport System Steam & Feedwater Systems Shutdown Cooling System ECC Moderator	Emergency Water System
Contain Radioactivity	Reactor building air coolers	Containment & containment subsystems
Monitoring & Control	Main Control Centre	Secondary Control Area

The rationale for this particular grouping and separation is as follows:

- two shutdown systems are in separate groups so that a single event cannot prevent shutdown
- ECC and containment are in separate groups so that a single event cannot damage fuel *and* allow radioactivity to escape

The choice is not unique - what would be the advantages and disadvantages of switching ECC and containment?

Most safety systems need services such as air, water and power. These too must be grouped and separated, and here is an example corresponding to the table above:

Safety Support Function	Group 1 Safety Support	Group 2 Safety Support
Electrical Power	Class IV Class III diesels Class II Class I	Emergency Power System (EPS) Diesels Class II Class I
Service Water	Raw Service Water Recirculating Service Water	Emergency Water System
Instrument Air	Instrument Air System	Local Air Tanks

With four safety systems, each of which has three instrumentation channels, grouping and separation of cables is a major challenge. Here is one such implementation, with each letter representing a different instrumentation channel:

System Group	System Name	Channel		
1	Reactor Regulating System	A	B	C
1	Shutdown System 1	D	E	F
1	Emergency Core Cooling System	K	L	M
2	Shutdown System 2	G	H	J
2	Containment System	N	P	Q
1	Emergency Core Cooling System - Seismically Qualified	KK	LL	MM

Figure 5-3 shows how all this comes together in a plant layout. Here is what one attempts to do:

- safety system triplicated instrumentation channels within a Group separated by 1.5 metres
- power supplies split into “ODD” & “EVEN” to serve redundant components within a Group
- “ODD” & “EVEN” cables separated by 1.5 metres
- single channels within the same Group can share common routing (e.g., A, D, K)
- buffering of connections between Main Control Room & SCA
- power cables >600 volts must be 0.45 m. above instrumentation cables

Note that only *one* group (Group 2) is normally required to be seismically qualified - the frequency of severe earthquakes is such that a double line of defence is not needed.

Note that there are other approaches to grouping and separation. LWRs tend to have two to four spatially separated trains, with each train fully qualified for all accidents, and less redundancy within the train, as shown in Figure 5-4. The advantage of a four-train approach is that it allows one train to be taken out for maintenance while the plant is operating. The disadvantage is less diversity of equipment among trains.

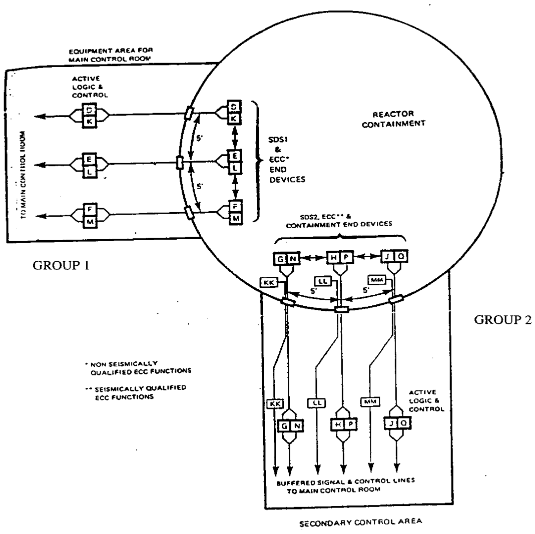
EQUIPMENT AREA FOR MAIN CONTROL ROOM
ACTIVE LOGIC & CONTROL
TO MAIN CONTROL ROOM
D
K
E
L
F
M
5'
5'
SDS1 & ECC* END DEVICES
REACTOR CONTAINMENT
SDS2, ECC** & CONTAINMENT END DEVICES
G
N
H
P
J
Q
KK
LL
MM
GROUP 1
GROUP 2
* NON SEISMICALLY QUALIFIED ECC FUNCTIONS
** SEISMICALLY QUALIFIED ECC FUNCTIONS
ACTIVE LOGIC & CONTROL
BUFFERED SIGNAL & CONTROL LINES TO MAIN CONTROL ROOM
SECONDARY CONTROL AREA

Figure 3 Location and Separation Requirements for Safety Related Systems

Figure 5-3 - Grouping and Separation for Safety-Related Systems

This can be contrasted with the CANDU Classic approach we have just discussed: two diverse separated groups of systems, redundancy within each group, and the level of qualification determined by the safety function of each system, as shown in Figure 5-5.

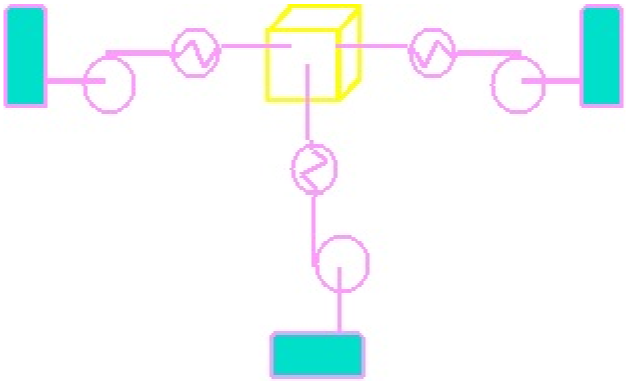

Figure 5-4 - Simplified LWR Three Train Grouping and Separation

Note from the diagrams that LWRs do not normally allow cross-connexion between the trains. CANDU Classic allows cross connexions within the group, but only limited and buffered connexion between Groups. Why allow any at all? For example, the Main

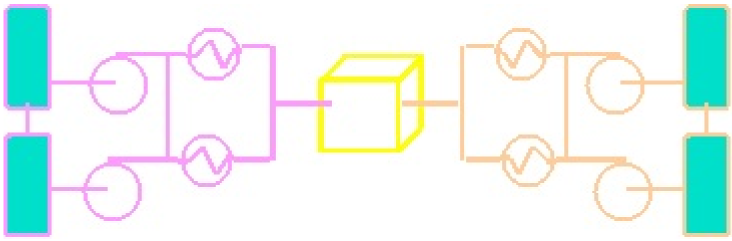

Figure 5-5 - Simplified CANDU Two-Group Grouping and Separation

Control Room (MCR) is a Group 1 area, so ideally should have no Group 2 equipment. But the operator would want to have the capability of firing SDS2 manually from the MCR, or observing its indications - rather than having to walk over to the Secondary Control Area each time. So a buffered connexion is provided - typically the electrical link is broken by an optical link, so that signals can go only one way, and a fault in the Group 1 area cannot propagate to Group 2, and vice versa.

The ACR-1000 design has adopted a four-train approach to ensure maintenance can be done at power with one train (or logic channel) out of service while the remaining trains or channels still meet the safety requirements.

Testing

How do we know the systems will work as designed? For shutdown systems, the answer is fairly easy: they are tested during operation to determine their *reliability* and their *performance*.

Recall that demand availability is supposed to be >0.999, or all but 8 hours per year. The difficulty is that if a system is actually tested[h] at a frequency sufficient to demonstrate this value, there would be severe negative impact on station operation. The decision and action time after a shutdown is only about 20 minutes before xenon build-up overcomes the positive reactivity range of the control system. For SDS1 this is barely enough time to reset the system and restart; for SDS2, a poison-out is inevitable and the plant would be down for 40 hours or more until the xenon has decayed sufficiently. Moreover excessive shutdowns stress the plant (for example each time the reactor is shut down the ‘house load’ - the Class IV electrical power generated by the station itself - is obviously lost, since the turbine is not supplied with steam, and the turbine trips). The problem is addressed by separating the testing of the *logic* from that of the final mechanism; and by testing the mechanism in stages, not all together.

First, the logic. All safety systems in CANDU Classic have three logic channels, as noted, with two-out-of-three being sufficient to initiate the trip or safety system action (Figure 5-6). The requirement that two channels must *both* vote ‘for’ a trip reduces the likelihood of spurious trips due to a single component failure. On the other hand the reactor will still trip if required, even if

[h] That is, if the entire system is tested all together, as it would ‘really’ work

one channel is unavailable (failed unsafe). It is usual when a channel is known to be failed to put it into a safe (tripped) state: this increases somewhat the likelihood of a spurious trip, since only one of the remaining two channels need fire to trip the reactor, but allows continued (safe) operation until the failed channel is repaired. Finally a single channel can be tested without tripping the reactor:

By the same token, one can test hardware devices (as well as logic) without firing the system.
In the ACR-1000, and modern LWR designs, a “two-out-of-four” voting logic is used. There are four instrumentation channels, any two of which can trip the reactor. Again, this allows one channel to be taken out for maintenance and left ‘un-tripped’, reducing the likelihood of spurious trips, but at a cost increase of 33% in instrumentation, maintenance and testing.

The shutoff rods are designed to be partially dropped, individually, in a test. That is, shortly after the clutch has released and the rod begins to fall, the clutch can be re-energized and the rod “caught” before it enters the core. This proves the rod is not stuck (what does it *not* prove?). In practice each rod is partially dropped about once a week.

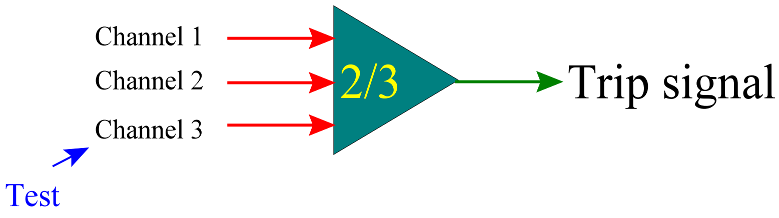

Figure 5-6 - 2 out of 3 trip logic and test

Shutdown system 2 is more complex, as it is not easy to stop an injection. Figure 5-7 shows one way of testing each valve by itself without firing the system, while any two channels which trip *will* fire the system. Testing any single logic channel opens two valves but does not allow (much) poison to leak into the moderator. Thus each component can be tested without activating an injection.

Before a long shutdown, when the poison-out doesn’t matter, SDS2 can be tested “for real” to ensure that some subtle failure has not gone undetected in test.

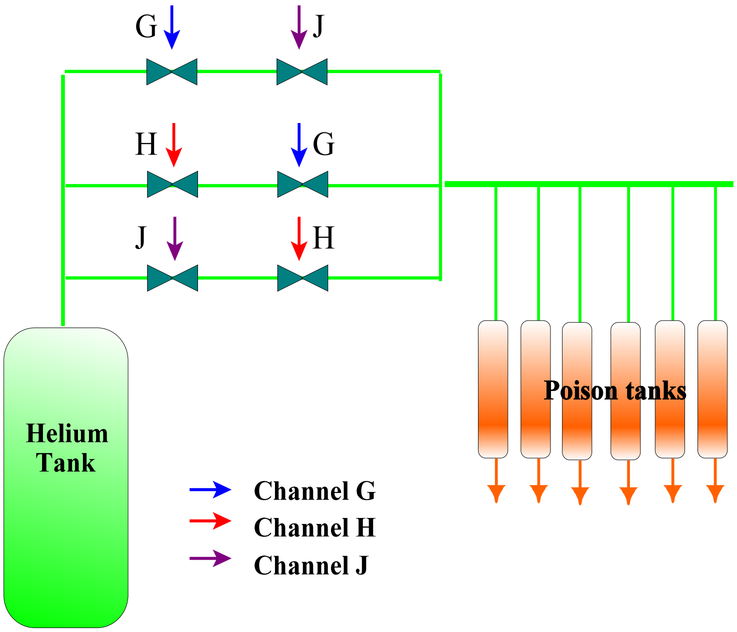

Figure 5-7 - Testing of SDS2 Valves

Performance testing is done usually by measuring the speed of actuation of one or more components. Thus on a partial rod drop, one can determine how long it takes the rod to reach the point when it is caught; or one can test the valve opening times on SDS2. A full test measures the rod drop versus time (time to pass three “gates” on its way down), and of course the actual power rundown is measured during commissioning and on the tests before a long shutdown to verify that the designers got the physics right.

Human Interface

An operator must:

- Be notified that the shutdown system has tripped
- Be able to confirm that it has actuated correctly
- Have procedures to follow in case it has not.

These seem like obvious requirements but there have been a number of accidents where the operator has been misled by his instrumentation - we covered some in the test cases we reviewed. Normally notification consists of an alarm window on the SDS panel, showing (for each channel) that a trip parameter has passed its trip setpoint, and a window showing that the SDS has fired. The latter is *not* sufficient to establish that the system has indeed worked (recall the Three Mile Island case where the valve position indicator in the control room was taken from the control signal which *told* the valve what to do, not what it actually did), and an operator can expect to have available supplementary information such as shutoff rod position; he will also check the neutron power measurement to ensure it is consistent with a shut down reactor. Finally should these latter measurements suggest that the system has not fired, he will have and will follow backup procedures such as: manually trip the shutdown system, manually fire the second shutdown system, drop the mechanical control absorbers (manual stepback) etc.

Our in-depth example of the shutdown systems is echoed in the systems responsible for the other safety functions - heat removal, containment, and monitoring. We will cover these more briefly.

Heat Removal Systems

Some of the specific questions that have to be asked about a heat removal system are as follows:

- how much heat must it remove (full power, decay heat, decay heat after *x* minutes etc.)?
- where is it connected (primary side, secondary side, etc.)?
- how it is initiated?
- what conditions can it operate under (pressure, temperature)?
- what is its reliability?

We will now suggest some answers.

Heat Removal Capability

It seems obvious that at least one system must remove up to the full power generated by the reactor - else how are we to use the power? This is normally done by the main steam and feedwater system. It also seems obvious that one does not need more than one such complete system for safety reasons. However safety does have to cater for the case of a sudden loss of heat removal *from* the secondary side (e.g., turbine trip with loss of condenser vacuum) so that the capability of 100% steam dump to atmosphere is provided. This is accomplished by 16 main steam safety valves (MSSVs) on the main steam lines. Full steam flow is needed only until shortly after the reactor has tripped. The MSSVs have a function for LOCA as well, which we shall cover later.

For economic reasons, another subsystem is provided which can dump up to 60% steam directly from the steam generators to the condenser. This is used for *poison prevent* - that is, if the turbine trips but the operator thinks he can restart it reasonably soon, instead of shutting the reactor down, he can set back the reactor power to the level just sufficient to prevent a poison-out due to xenon buildup. He still has to get rid of this heat, however, so the design allows him to dump the

CANDU ~~Bundle P~~ower after Shutdo

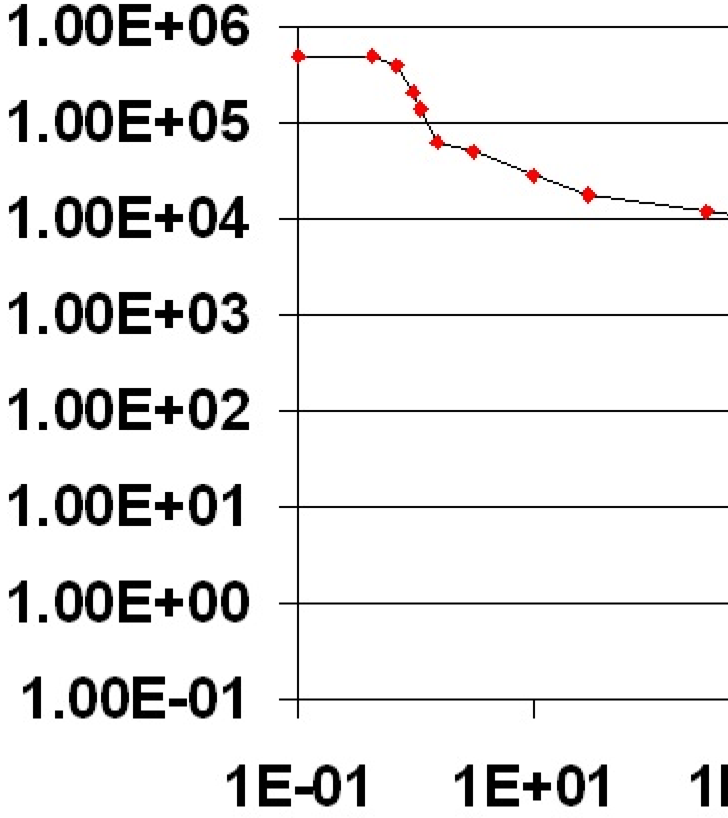

Secondary Side Heat Removal

The secondary cooling system is a good choice if it is available (this depends on the accident, of course), and if the primary cooling system is intact (no large LOCA), since it is normally operating anyway. Thus the usual choice to remove heat is via the steam and feedwater system. The requirements to use the secondary side are:

- the secondary piping should be intact (there are exceptions which we can cover later)
- there should be source of electrical power (or gravity head, or steam pressure) to pump cold water into the steam generators
- there should be a source of water to remove heat from the steam generators

These rather obvious statements give rise to the following design choices:

- if all components and systems on the secondary side are working, the main feedwater pumps provide water to the steam generators; the condenser removes heat from the steam and provides a *continuing* source of cold water to the feedwater pumps. Clearly this relies on the availability of Class IV electrical power (generated by the station itself when it is operating, or from the electrical grid), since the main feedwater pumps are very large.
- if for some reason the main feedwater pumps are not available, the task can be undertaken by one or more auxiliary feedwater pumps, sized to remove decay heat. These feedwater pumps are typically powered by Class III electrical power (generated by station diesels), or directly by a diesel engine, or by a steam turbine connected to the steam generators. It is not necessary to size them for 6% power (the decay power *immediately* after shutdown), as they are not really needed until some of the water already in the steam generators boils away. This takes about half an hour, so that a heat removal capability of the auxiliary feedwater pumps is typically about 4%.
- if the accident involves a loss of Class IV power, then the main feedwater pumps *and* the condenser are unavailable (the latter since the cooling water to the condenser is supplied by large Raw Service Water pumps which are on Class IV power). The auxiliary feedwater pump can still supply water, and the heat is removed by steaming to atmosphere from the steam generators, either via the Atmospheric Steam Discharge Valves (ASDVs), with a capacity of about 10% of full power steam flow, or via the Main Steam Safety Valves (MSSVs), with a capacity of about 115% full power steam flow. This is fine as long as there is water in the feedwater train, but after about an hour or so the water will be used up and the operator will have to establish another heat sink.
- in some CANDUs, the dousing tank located in the top portion of the containment can supply a longer-term source of water by gravity to the steam generators. Conditions for its use are that it cannot have been used up during the accident (e.g., for a pipe break inside containment) and that since it is a low-pressure source of water, the steam generators must be depressurized before it is brought in. The ACR-1000 has an elevated tank at the top of containment (the Reserve Water Tank) for this purpose.

- Ref. [2] has increased the requirements for decay heat removal for new designs, and mandates that they be a safety system - increasing the number of safety systems in ACR-1000, for example, to five. The designation of "safety system" for this so-called Emergency Heat Removal System (EHRS) brings along all the requirements we have been discussing such as reliability, testing, separation etc.
- most recent CANDUs have a seismically qualified source of water (e.g. a large pond) for use after an earthquake. This Emergency Water System (EWS) has its own seismically-qualified power and pumps, and can supply water independently to the steam generators for about 3 days. The water is "dirty" and, barring a real earthquake, would be a system of last resort for the operator.

Primary Side Heat Removal

Since many of the options for secondary side heat removal are only valid for a limited period of time in an accident, most reactors also have a primary-side system to remove decay heat. In CANDU Classic, this is the *Shutdown Cooling System*. It is a closed system connected to the reactor headers (as shown in Figure 5-9) with its own pumps and heat exchangers. It is a high-pressure system. Light Water Reactors have a similar system called a *Residual Heat Removal* (RHR) system, but it is low pressure and requires depressurization of the primary coolant system before it is brought in.

In some CANDUs, there is a connexion from the *Emergency Water System* to the primary cooling system. It requires depressurization of the primary coolant system, and a way for the steam to be removed - e.g., by opening a primary-side relief valve. This is likewise a choice of last resort.

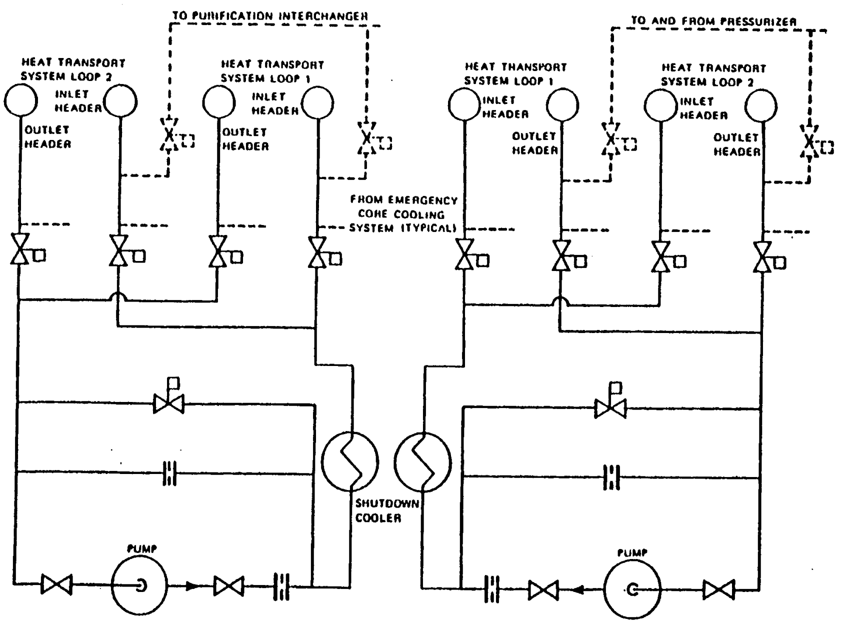
TO PURIFICATION INTERCHANGER
TO AND FROM PRESSURIZER
HEAT TRANSPORT SYSTEM LOOP 2
HEAT TRANSPORT SYSTEM LOOP 1
HEAT TRANSPORT SYSTEM LOOP 1
HEAT TRANSPORT SYSTEM LOOP 2
INLET HEADER
INLET HEADER
INLET HEADER
INLET HEADER
OUTLET HEADER
OUTLET HEADER
OUTLET HEADER
OUTLET HEADER
FROM EMERGENCY CORE COOLING SYSTEM (TYPICAL)
SHUTDOWN COOLER
PUMP
PUMP

Figure 5-9 - Shutdown Cooling System

The *Emergency Core Cooling System* can be viewed as a decay heat removal system for the special case of a break in the primary cooling system piping. We shall cover this system as a special case.

Other Heat Removal Possibilities

We have mentioned that the *moderator* surrounding the fuel channels can be used in a severe accident (LOCA with loss of ECC) to remove decay heat. The heat removal pathway is efficient enough to prevent fuel melting, but will not prevent extensive fuel damage and distortion of the fuel channels (Figure 5-10).

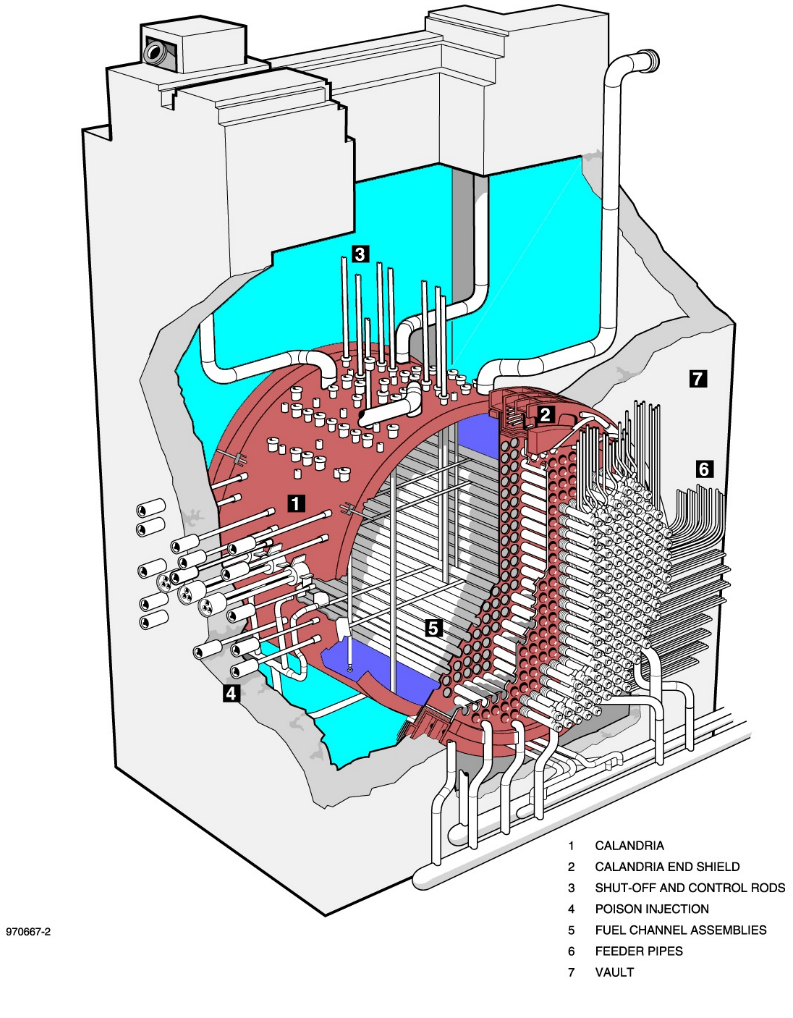

Figure 5-10 - Moderator and Shield Tank Surrounding Core

The *shield tank* surrounding the calandria (see also Figure 5-10) has its own heat removal system (pumps and heat exchangers) and can be used in a severe accident, e.g., if there is a LOCA plus loss of ECC plus loss of moderator heat removal. In CANDU Classic it does not have enough heat removal capability (0.3%) to match that being generated by the fuel; in addition the causes for the failure of ECC and moderator cooling may also have disabled the heat removal capability of the shield tank (e.g., loss of electrical power, loss of service water). The shield tank acts more to delay the progression of core damage than to stop it. However in ACR, the heat removal capacity of the shield tank system has been increased along with the provision for steam relief. Makeup from the elevated Reserve Water Tank has been provided to both the moderator and to the shield tank, and this is sufficient to prevent loss of reactor geometry for about three days. In severe accidents, the more time one has for human action, the better.

The *feeder pipes* connected to each fuel channel give a large surface area which, as long as there is water in the fuel channels, can be used to reject heat to the building atmosphere, if the plant has been shut down for a long time (Figure 5-11), without primary side heat removal or circulation. This means of heat removal was used at Point Lepreau during the long shutdown after wood was left in the boiler.

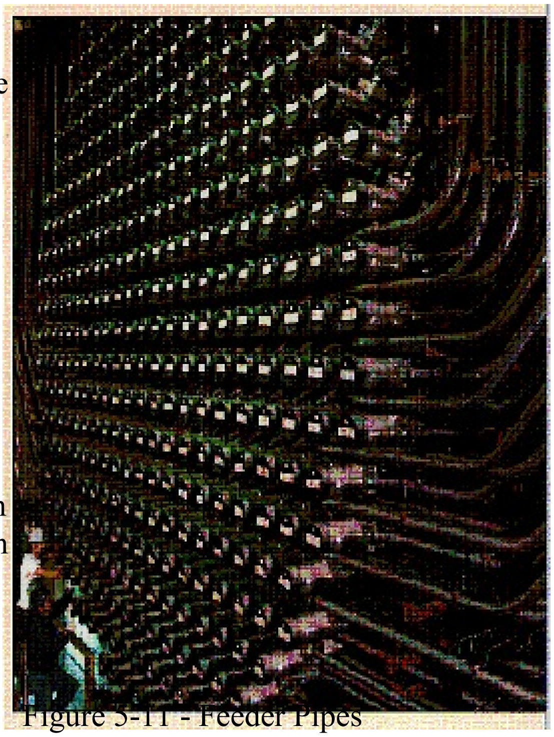

Figure 5-11 - Feeder Pipes

Initiation of Decay Heat Removal

Unlike shutdown systems, which are almost always initiated automatically, decay heat removal systems can be either automatic or manual. The decision is based on *when* they are needed. Typically if they are not needed for 15-30 minutes, they can be manually initiated (e.g., shutdown cooling system, EWS); if they are needed sooner, they must be automatic (e.g., ECC, auxiliary feedwater system).

Operating Pressure

A key design choice is the *operating pressure* of the heat removal system. The table below summarizes the advantages and disadvantages of high vs. low pressure. In many cases the pressure is determined by the nature of the design (e.g.. moderator).

Operating Pressure	Advantages	Disadvantages
High	Can be brought in at any stage of an accident Components tend to be smaller due to more efficient heat removal (larger ΔT) More easily automated	More stringent requirements on code class of piping and components Need to ensure it is tolerant if brought in when system is at low pressure (e.g., risk of pump cavitation)
Low	Can be simpler/cheaper Can be made more passive	Requires prior depressurization of the system (i.e. depends on another system)

Reliability

The mission time of decay heat removal systems can extend to days or (in the case of ECC) to months - the factors of concern are therefore *both* the demand availability (reliability to start) *and* the running reliability once it has started . A typical active decay heat removal system consists of pumps, which require electrical power, and heat exchangers, which require a source of cooling water, which in turn requires electrical power. For systems other than ECC, unavailabilities in the range of 10^{-2} to start, and 10^{-1} over a long mission time are typical[i]; ECC is typically an order of magnitude better. Thus redundancy in decay heat removal systems is necessary.

Some reactor designs have passive systems, where an elevated water tank provides decay heat removal with no pumps required. The connection from the dousing tank to the steam generators in CANDU 6 is of this type. The steam generators must first be depressurized by opening the MSSVs, and power is also required to open the valves to connect the dousing tank to the steam generators, but thereafter no motive power is needed, with water flowing into the steam generators under gravity. Care must be taken to match the water flow to the required heat removal capability, or else the steam generators will repressurize and block further flow; or flood. In designs such as the Westinghouse AP-600/1000 series, decay heat removal from both the primary cooling system and the containment is passive. A passive system is not inherently more or less reliable than an active one, although it is often perceived to be more reliable; nor it is necessarily cheaper or more expensive. As is usual with design, the result depends on the details.

The Table below summarizes in simplified form[j] some of the main characteristics of heat removal systems:

[i] Note the units are different: the starting unavailability is a *demand unavailability* and is dimensionless (think ‘unavailability per attempt’); the mission unavailability is per defined unit of time, e.g., over the defined mission time. Note also that for new designs, the EHRS must meet a safety system demand unavailability of 10^{-3}.

[j] To cover each CANDU type would lead to an overly complex table; thus the reference plant for these statements is a recent CANDU 6. Some comments refer to ACR.

System	Operating Pressure Range	Heat Removal Capacity	Support Systems	Comments
Main steam and feedwater system	Atmospheric to operating (10MPa)	0-115%	Class IV power Water from condenser	Normal power operation Cooldown to 177C after a shutdown
Auxiliary feedwater system	Atmospheric to operating	Decay power	Class III power Water from condenser	Used for loss of Class IV power
Shutdown Cooling system	Operating to atmospheric	Decay power	Class III power (+ Group 2 Emergency Power System on some new designs) Recirculating Cooling Water (RCW)	Used for cooldown from 177C after a shutdown Can be brought in at full system temperature in an emergency
Emergency Water System	Near-atmospheric (up to secondary side operating pressure on some new designs)	Decay power	Group 2 Emergency Power System. Some form of water reservoir (dousing tank, external pond)	Used after an earthquake and as a backup to Group 1 heat removal systems. Requires depressurization of primary or secondary side to be effective. Upgraded to safety system on ACR-1000.

System	Operating Pressure Range	Heat Removal Capacity	Support Systems	Comments
Moderator	Atmospheric (can remove heat when the Heat Transport System is up to about 6 MPa, without pressure-tube failure)	Decay power a few minutes after shutdown (5%)	Class III power RCW	Used in severe accidents where there is no primary-side heat sink e.g., LOCA + LOECC. In ACR-1000, water can be added to the moderator by gravity.
Shield Tank	Atmospheric	0.3% power	Class III power RCW	Will delay progression of core melt due to large amount of water. In ACR-1000, water can be added to the shield tank by gravity.
Feeder pipes	Full pressure to atmospheric	Very low (weeks after shutdown)	Channels should be full of water Heat must be removed from containment	Used at Point Lepreau during the long shutdown, which was required to remove debris from the HTS

ECC

Some of the specific questions that have to be asked about the Emergency Core Cooling System are as follows:

- what are its safety performance requirements?
- where is it connected? (where is the best place to put the water?)
- what is the injection pressure?
- what other functions besides water injection must ECC perform?
- how is it initiated?
- what is its reliability?

We will now suggest some answers.

Performance Requirements

A modern CANDU reactor has between 360 and 480 fuel channels, and therefore 720 and 960 feeder pipes; the ACR-1000 has 520 channels. Such a large amount of small diameter (<3.5") piping would suggest that small breaks are of much interest in CANDU. To date the history has been major failures of two pressure tubes, one in Pickering A and one in Bruce A; leaks in a number of pressure tubes which were not allowed to proceed to failure; feeder pipe leaks but no feeder pipe breaks. The probability of a small pipe break is thus about 10^{-2} per reactor year based on experience[k], which is large enough that it is a concern not only for safety but also for plant investment protection. Large primary pipe breaks have never occurred in a Western nuclear reactor; the probability is between 10^{-4} and 10^{-5} per reactor year. Thus ECC has *safety* requirements for both large and small breaks; and in addition *investment protection* requirements for small breaks only.

It would appear that small breaks would be less of a concern for LWRs. However a small break comes in many guises: a stuck-open relief valve is equivalent to a small break for an LWR (as it is for CANDU); the Three Mile Island partial core melt began as such a small break, exacerbated by the lack of attention paid to small breaks in the design and safety analysis of LWRs prior to the accident.

The safety requirements for large breaks are:

- meet public dose limits
- prevent pressure-tube failure
- ensure the fuel in the fuel channels is coolable

The safety (not regulatory) requirements for small breaks are the same. However for investment protection the requirement is to limit (to the extent practicable) the amount of fuel sheath failure. The hope would be to restart the reactor with the same fuel load after recovery from a small

[k] Once a failure mechanism is identified, it is fixed, so that the frequence of *future* failures may well be much less.

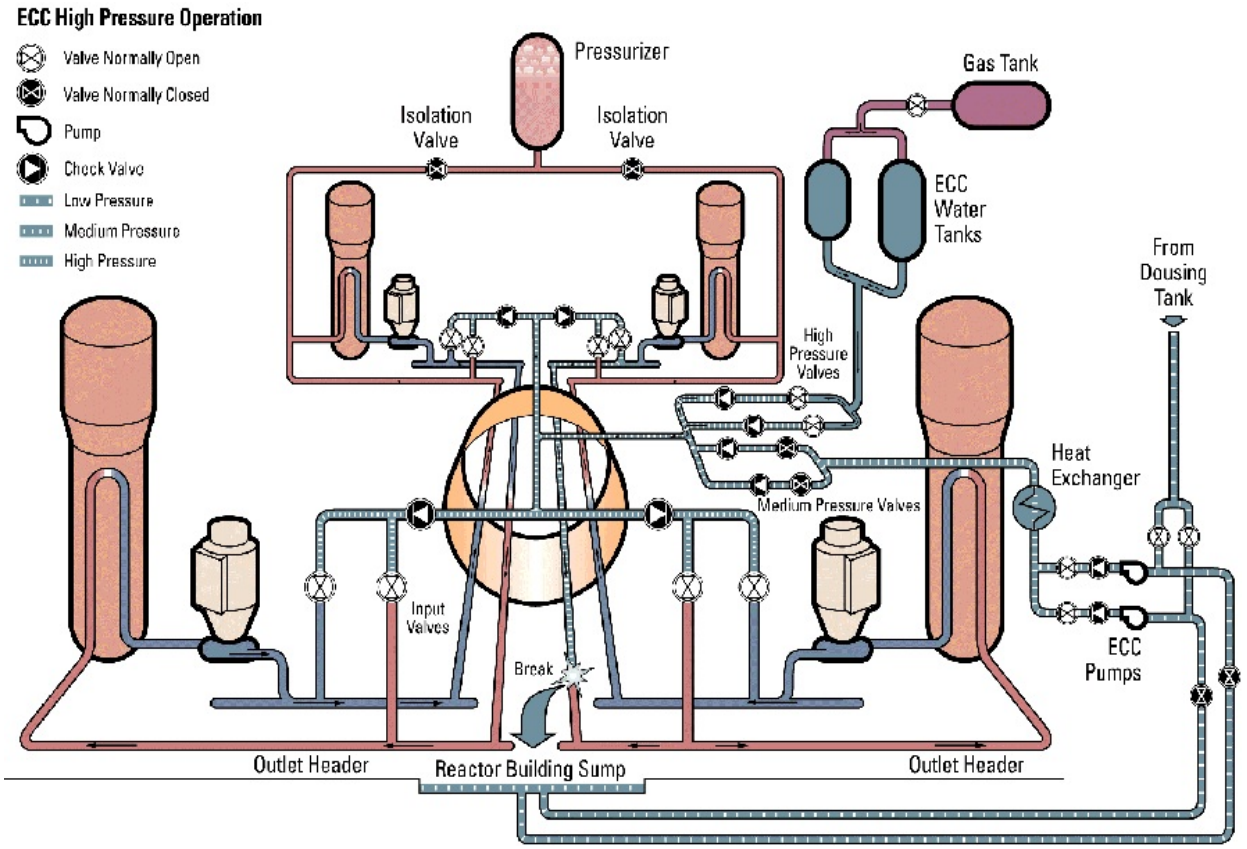

Figure 5-12 - ECC Schematic

break; and at least if defuelling was required, not to have to clean up a highly-contaminated building.

Location of Water Injection

The CANDU ECC injects[1] into the reactor inlet and outlet headers in both heat transport system loops (for reactors which have two loops) - 8 headers in all, in CANDU 6 (Figure 5-12). Since each channel is connected to two headers, and the headers are above the core, this gives a water pathway to every channel in the core. The disadvantages of this scheme are:

- the water injected near or at the break is wasted - i.e., discharges without removing heat from the fuel

[1] The subsystem which performs injection is sometimes referred to as Emergency Coolant Injection, or ECI, in Canada

- the (feeder) pipe to each channel is fairly small-diameter, and contains a lot of stored heat, and the injection water may have to flow countercurrent to the steam exiting the channel as it refills.

Current designs provide sufficient pressure and water volume that wastage of water out of the injection point at or near the break is accounted for.

Early CANDUs such as Douglas Point (and early Indian HWRs) used instrumentation to automatically deduce the end of the reactor where the large break had occurred, by comparing the pressure in the headers. The ECC was then injected into the headers at the opposite end. For small LOCAs, the ECC was injected to all headers. This design does not waste as much ECC water out the break but it is complex and has reliability issues.

Other locations have been tried in different designs. The UK design for a vertical pressure-tube reactor (the Steam Generating Heavy Water Reactor, or SGHWR) used a perforated "sparge tube" (in place of the central fuel element) running down the centre of each channel; the sparge tube sprayed ECC water directly on to the fuel. However analysis of its performance was very difficult, and it could not be proven analytically that each fuel element was cooled sufficiently by the water jets from the central tube. Other prototype pressure-tube reactor designs have had injection into each individual channel through a separate feeder pipe; or have used check valves in each feeder pipe to prevent back-flow. The tradeoff between complexity and reliability is obvious.

Refilling a LWR is in principle much simpler: it is a large pot with the large coolant pipes all located above the core, so all one has to do is pour water into one or more of these pipes (hot or cold leg) and fill the pot nicely from the bottom. The fuel rewets as the water level rises. However the ECC must be borated to avoid a reactivity excursion (why?). There were also some concerns (now solved) about bypass of the ECC water from the ECC inlet location directly to the break, without going through the core.

Light water has been used as the ECC fluid in all modern water-cooled reactors. In CANDU it has the disadvantage (except for ACR) that a spurious injection of ECC will downgrade the heavy-water coolant. Alternative ECC fluids have been studied but not implemented (except see the discussion below re Douglas Point and Pickering-A). Light water injection ensures shutdown in CANDU, of course, without the need for borating.

Injection Pressure and Flowrate

Early CANDUs (Douglas Point, Pickering-A) had a relatively low-pressure heavy-water ECC supplied by the moderator pumps. It avoided heavy-water downgrading for small breaks. However this concept had a number of disadvantages (one can always be clever in retrospect): it cross-linked the moderator and the ECC, so that the moderator could not be relied on as an independent backup to a loss of ECC[m]; and as fuel element powers increased in later designs, the injection pressure was insufficient to provide both economic and safety protection.

The current concept (using CANDU 6 as a model) is to have a three stage ECC:

- A *high pressure* initial injection phase limits fuel overheating for small breaks, and forces rapid refill for large breaks (to limit pressure tube deformation and early fuel damage).

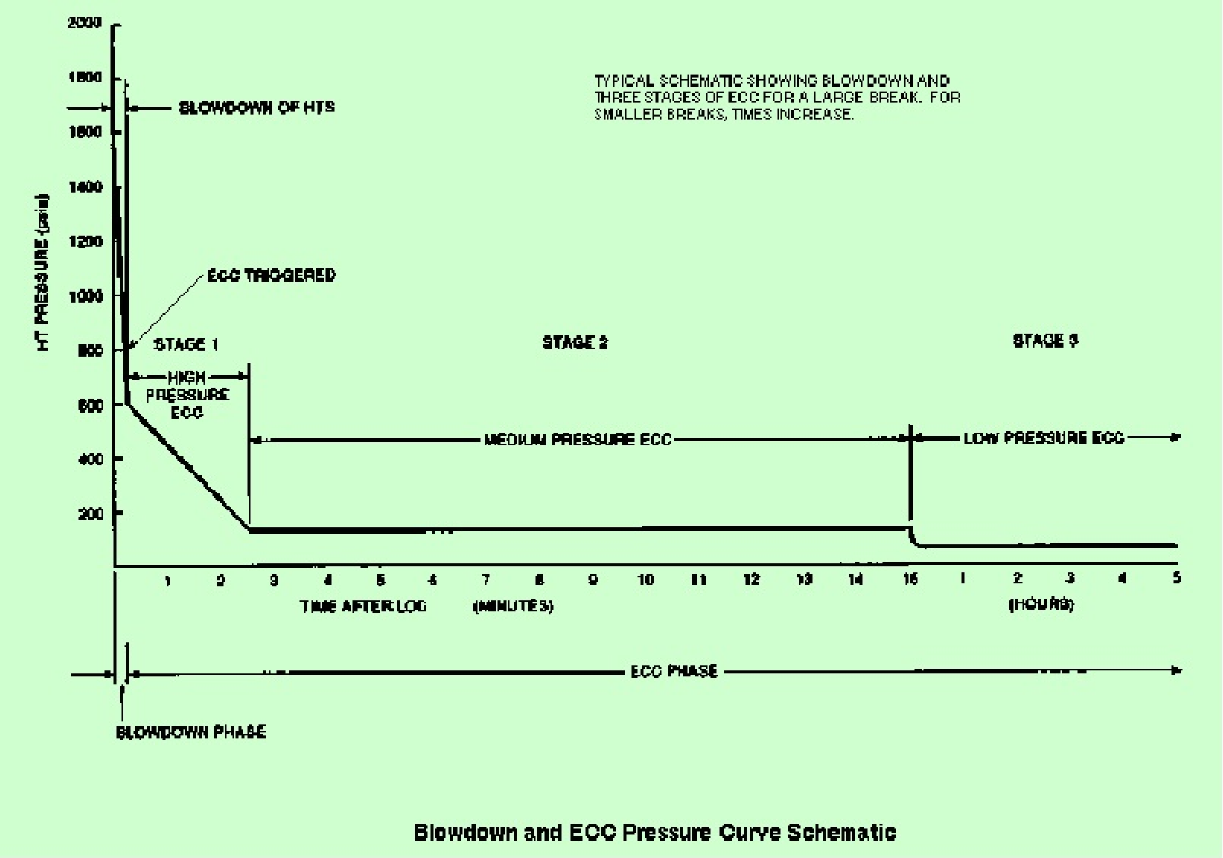

Figure 5-13 - ECC Phases

Typical injection pressure is 5.3 MPa. The pressure creates a good pressure-drop across each channel, allowing water to enter one feeder and steam to exit the other. In CANDU 6, high pressure injection comes from 2 water tanks which are pressurized by gas at the time of a LOCA signal. The tanks have enough water volume to fill one complete heat transport system loop twice over. In some multi-unit CANDU plants, this high pressure

[m] Especially since it might also have been dumped due to shutdown

phase is supplied by pumps. The reason for pumps versus accumulators has a lot to do with the reliability of electrical power. In single unit plants, when the reactor trips, it may cause a collapse of the external electrical grid, especially if the reactor is a major part of the grid. Thus designers cannot count on a reliable source of continuing Class IV power to power ECC pumps; hence the high-pressure phase is through accumulators, which need only control power (local AC power backed up by batteries) to open valves and start to inject water. Class III diesel-generators on the plant site will start automatically on loss of Class IV power, but take a few minutes to run up to speed and synchronize; they are therefore used to power the medium pressure and recovery pumps. The accumulators therefore provide cooling water to the core while the Class III diesel-generators are starting. On multi-unit plants, the reliability of Class IV power tends to be very high (since it can be obtained from the other operating units as well as the grid), so that all parts of ECC can be electrically powered. Internationally, however, it is very unusual to have safety systems powered by Class IV power.

- *Medium pressure* injection takes over when the high pressure water tanks are nearly empty. It uses medium pressure (~1 MPa) pumps and draws cold water from the dousing tank. It pumps this water into the same locations (all headers) as the high pressure ECC. The pumps have a dual electrical supply for reliability: normal diesel-backed Class III power; and seismically qualified Group 2 diesel-generated Emergency Power. This is so they can continue functioning if a LOCA is followed after 24 hours by an earthquake. (CANDU design does not consider a LOCA simultaneous with an earthquake as it is too low in probability.) The medium pressure phase ensures that enough water has collected in the basement of the containment building for the recovery phase to start.
- In *recovery injection*, the medium pressure ECC pumps are switched over to take water from the sump in the basement. They pump this water through dedicated ECC heat exchangers before returning it to the heat transport system.

The high pressure phase lasts for 2.5 minutes *minimum* for the largest break; medium pressure lasts for 12.5 minutes *minimum*; recovery ECC has a mission time of 3 months, after which the moderator can remove heat from the fuel channels without further fuel damage even if ECC is unavailable. Figure 5-13 shows the phases.

In recent CANDU designs (ACR) the medium-pressure phase has been eliminated: after the accumulator phase, the initial water for the recovery phase comes from grade-level tanks inside containment to ensure there is a sufficient head for the ECC recovery pumps.

The flowrates are chosen as expected: the high-pressure phase should refill the core for the largest double-ended pipe break; the medium pressure phase should not lead to uncovering of the fuel; the low pressure phase must keep the fuel covered indefinitely. As with many such performance requirements, deviations may occur and limited fuel uncovering during transitions in cooling mode is both unavoidable and permitted, as long as the heatup is not excessive.

Other Functions

There are two further functions that ECC must perform: loop isolation, and crash cooldown.

Loop isolation simply closes the valves connecting the two heat transport system loops (for those CANDUs which have two loops). In simple terms it limits (most of) the LOCA to one loop. In terms of fuel cooling, the effect is small; indeed it is probably better not to have loop isolation because when the "un-failed loop" is isolated, it gets refilled by ECC quite late - most of the water goes to the failed loop. If the main heat transport pumps are tripped in the interim, the un-failed loop has a period of partial-inventory thermosyphoning, which can be difficult to analyze. However for a LOCA with loss of ECC injection, the loop isolation (which is an independent signal) limits the possible source term of hydrogen. Safety (as usual) is a tradeoff between these two aspects. Loop isolation has been removed from Darlington; and the ACR-1000, while it has two loops which isolate on a LOCA signal, also has pressurized Core Makeup Tanks continually connected to each loop to compensate for limited loss of inventory and shrinkage due to cooldown.

Crash cooldown is more significant. If a break is small, less than say a feeder failure, it is not able to depressurize the heat transport system down to ECC pressure, or to keep it below ECC pressure once injection begins. The problem can be narrowed to smaller breaks, but not solved, by increasing ECC pressure; in any case this rapidly becomes too expensive and too prone to spurious initiation (in CANDU the ECC pressure should be below the lowest pressure to which the heat transport system pressure falls, after a reactor trip without a LOCA). Crash cooldown works by opening all the MSSVs on all the steam lines and blowing down the steam generator secondary side to near atmospheric pressure over about 15 minutes. Since the steam generators are still a heat sink for the primary coolant in a small LOCA, the effect is to force the primary side pressure down over the same timescale. Thus it ensures that ECC is not blocked by the heat transport pressure "hanging up" at secondary side pressure, and that the unfailed loop will be refilled by ECC. Some CANDUs (Darlington) use high-pressure pumps for a small LOCA and are not as dependent on crash cooldown for this purpose.

Note that crash cooldown is not used in PWRs: First, it has a positive reactivity effect due to the negative coolant void coefficient (why?); second the containment philosophy on LWRs isolates (closes) *all* lines penetrating containment, including the steam and feedwater lines, so they are not available for depressurization. We discuss this later in this chapter under 'containment'. The way LWRs achieve primary side depressurization is to open valves on the primary side to discharge coolant inventory, thus converting a small break to a somewhat larger one that depressurizes the system enough for ECC not to be blocked.

The *main heat transport system pumps* play an important role in fuel cooling for small LOCA early in the transient. In Three Mile Island, turning off the main pumps collapsed the water level in the core and lead to fuel overheating (nevertheless the final decision made in the U.S. after TMI was analyzed was to turn off the pumps in a LOCA). For CANDU, the main heat transport pumps are left running during the early phase of a LOCA, since they assist refill; and for small breaks, they provide assurance of fuel cooling. Their electrical supplies and the pumps themselves are therefore environmentally qualified for LOCA. In the longer term, they are tripped, to prevent cavitation once the heat transport system is refilled with cold water. However CANDU safety analysis also predicts the behaviour of LOCA assuming that Class IV power (which powers the pumps) is lost at the time of reactor trip.[n]

Initiation Signals

Given the speed of fuel overheating for a large break (a few seconds), it is obvious that ECC initiation must be automated. Even for small breaks manual initiation is marginal, and the system is likewise automated. The basic signal is low heat transport system pressure (Figure 5-14). By itself this is not unique to a LOCA, so it is *conditioned* (ANDed) by one or more of: high building pressure, sustained low heat transport system pressure, and high moderator level. The last is for an in-core break. A separate signal isolates the loops (in some designs) on low pressure (spurious loop isolation is not a major operational headache). Crash cooldown is shown in Figure 5-14 as being part of the ECC signal;

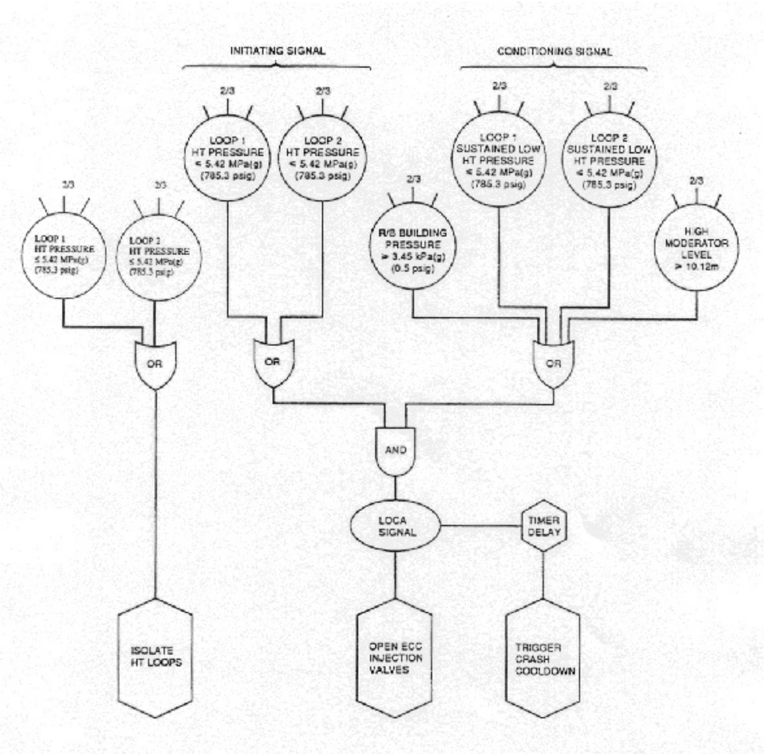
INITIATING SIGNAL
CONDITIONING SIGNAL
2/3
LOOP 1 HT PRESSURE ≤ 5.42 MPa(g) (785.3 psig)
2/3
LOOP 2 HT PRESSURE ≤ 5.42 MPa(g) (785.3 psig)
2/3
LOOP 1 HT PRESSURE ≤ 5.42 MPa(g) (785.3 psig)
2/3
LOOP 2 HT PRESSURE ≤ 5.42 MPa(g) (785.3 psig)
2/3
R/B BUILDING PRESSURE ≥ 3.45 kPa(g) (0.5 psig)
2/3
LOOP 1 SUSTAINED LOW HT PRESSURE ≤ 5.42 MPa(g) (785.3 psig)
2/3
LOOP 2 SUSTAINED LOW HT PRESSURE ≤ 5.42 MPa(g) (785.3 psig)
2/3
HIGH MODERATOR LEVEL ≥ 10.12m
OR
OR
OR
AND
LOCA SIGNAL
TIMER DELAY
ISOLATE HT LOOPS
OPEN ECC INJECTION VALVES
TRIGGER CRASH COOLDOWN

Figure 5-14 - ECC Initiation Signals

[n] This is standard practice in most jurisdictions. The concern is, as noted previously, that a small or unstable electrical grid might collapse if the reactor output to it is suddenly taken off.

however because of its importance[o] recent designs have two independent crash cooldown signals.

Transition to medium pressure ECC is automatic; transition to recovery ECC is automatic in new designs and is being backfit to some existing ones (the time-pressure on the operator was felt to be too severe).

Reliability

As a special safety system, the ECC must meet a demand unavailability of 10^{-3} or less. Because of the large number of valves, this has proven to be a challenging target in the past and is only recently being met consistently. There are no formal requirements for running reliability over the three month mission (after three months, as noted previously, decay heat can be removed directly to the moderator without further fuel damage, even with no water in the channels); a running unreliability of 10^{-2} over the mission time has been used as a target in the PSA, in recognition of the lower consequences of failure and the possibility of repair.

The requirements to have reliable separation between light and heavy water in normal operation, reliable mixing (injection) of light and heavy water during an accident, and the ability to test any ECC valve on power without firing the system, have lead to complex valving arrangements, as shown in Figure 5-15. The picture shows a valve train from the ECC water tanks to one of the headers. Each valve must be testable without firing ECC, and where a valve is required to open for ECC to be effective, it must be duplicated in parallel to ensure sufficient reliability.

In more recent designs (ACR-1000), the number of isolation valves and check valves have been reduced as spurious light-water injection is not a big concern. Indeed ACR-1000 has removed the separate ECC gas tanks, using gas to pressurize the ECC water tanks directly.

Containment & Sub-Systems

The important aspects of containment are as follows:

1. What is the design pressure?
2. What is the leakrate at design pressure?
3. How is pressure controlled? How is heat removed?
4. How is containment isolation ensured?
5. What is the containment reliability?
6. What other functions does containment perform?

[o] Failure of crash cooldown not only blocks ECC for small breaks but leads to pressure-tube deformation at high pressures, which risks channel integrity. This is discussed further in Accident Analysis (Chapter 7).

Design Pressure and Leakrate

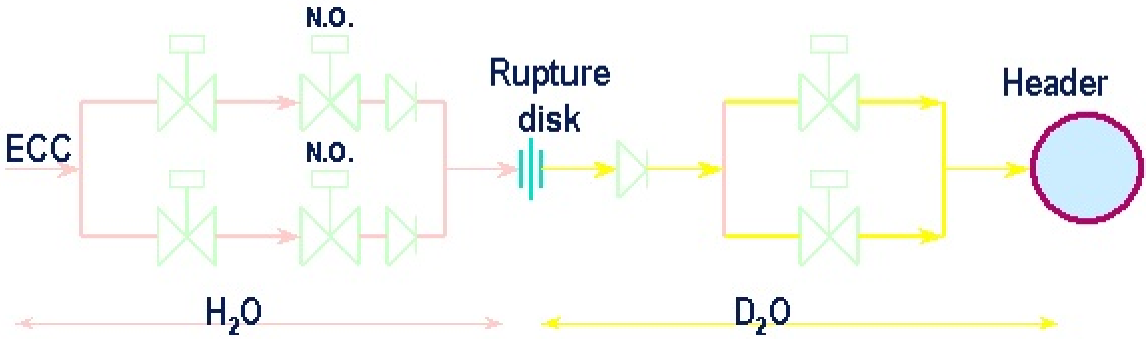

Figure 5-15 - Typical CANDU Classic ECC Valving Arrangement

In its simplest terms, containment is an envelope around those systems containing fission products, and is leak-tight in an accident. Of course any structure will leak in reality, the leakrate increasing with the internal pressure. The *design pressure* is chosen to be greater than the maximum pressure reached in any accident for which a predictable degree of containment leak-tightness is a requirement. To determine the design pressure, all accidents which release significant radioactivity into containment are analyzed, and the peak pressure reached in any one, plus some margin, is chosen to be the design pressure. The important aspect of design pressure is that *the leak-rate at the design pressure is known*. For CANDU 6, for example, the design pressure is 124 kPa(g) and the leakrate at the design pressure is 0.5% of the contained volume per day. Vacuum containments tend to have lower design pressures (because the vacuum building terminates the pressure rise) and (after the initial short-term overpressure) leakrates are *negative* (inward) for several days.

Should the pressure exceed the design pressure, the building will not explode (typical safety margins on massive failure of the building are a factor of about three over the design pressure). However the leakrate will be less easy to predict because the leak area may increase. In particular the leak-tightness of any penetrations and seals above design pressure is dependent on their detailed design and will need to be determined on a case-by-case basis.

Typically the design pressure for CANDU is set by the large LOCA, as this causes both a large short-term pressure and has the potential to release fission products into the containment. The leakage rate at design pressure is confirmed by proof testing before the plant is operational and by periodic testing thereafter.

A large steam-line break within containment in CANDU causes a *higher* peak pressure than a large LOCA[p]. However for existing CANDUs a steam-line break has not been used to set the design pressure, as it is argued that there are no fission products accompanying the break (there could be minor releases of radioactivity if the plant is running with a leaking boiler tube). The requirement for a steam-line break is that there be no structural damage to containment, which could risk damage to other mitigating systems.

The requirement to stay within design pressure for a LOCA applies even assuming failure of an active pressure suppression system (dousing or air coolers, discussed below), consistent with the traditional Canadian approach to licensing.

By contrast, LWRs require that the pressure from both LOCAs and steam/feedwater line breaks stay within containment design pressure. This may be partly because of the potential for a reactivity increase in LWRs for secondary side breaks (in CANDU there is a reactivity decrease) and partly to a view that containment capability should not be damaged after any single accident. International practice adopts the LWR approach - and for new plants in Canada, the requirement (Ref. [2]) is that the containment design pressure is not exceeded for *all* design basis accidents, including main steam line failure.

The calculation of the basis of the design pressure differs between CANDU and LWR practice. In CANDU, the dose used for siting and in accident analysis is calculated using physically-based scientific models: reactor physics, fuel behaviour, heat transport system thermohydraulics, fission product behaviour, containment pressure transient and atmospheric dispersion. The leakage rate at the design pressure is one of the parameters chosen by designers to meet the dose limit. This is different from the approach followed when siting LWRs: the source term and the containment pressure are pre-set (i.e., the dose is calculated assuming the containment is at design pressure for a predetermined period of time and leaking at the design leakrate). The result is that LWRs typically have lower leakage rates (~0.1 - 0.2% / day at design pressure) and (partly because of the requirement on steam line breaks discussed above and partly because of their smaller volume) higher design pressures than existing single-unit CANDUs. Again, the ACR-1000 follows LWR practice on containment leak-rate, with a low leak rate being enabled by the steel liner used in the ACR-1000 containment.

In some severe accidents, the containment pressure can increase to greater than design pressure. It is desirable that the containment leak gradually rather than fail catastrophically in such cases. The Atomic Energy Control Board (AECB - now Canadian Nuclear Safety Commission or CNSC) tested a 1/14 scale-model of a CANDU 6 containment and showed that cracking would

[p] Note that for the Bruce and Darlington containments, the steam lines are entirely outside the containment structure

not occur until the internal pressure was more than twice the design pressure. To fail the building took over 4 times design pressure. However to cause this failure the experimenters had to artificially prevent leakage. In actuality the leakage would increase through the cracks before the pressure reached failure pressure, so it would be impossible for the building to fail suddenly. We shall cover severe accident behaviour in a later chapter (Chapter 7).

Pressure Control and Heat Removal

Without some means of removing heat, the containment pressure in an accident such as a pipe break will rise rapidly as the broken system discharges steam and empties, then more slowly as the decay heat is steamed into containment. It is possible to build a CANDU containment to withstand this; such a "dry" containment would have a high design pressure and/or would be larger. To date CANDU containments have had some means of short-term pressure suppression and/or some means of long term heat removal, to cater for these two phases of an accident.

Canada has used several different methods of pressure suppression:

Ontario Hydro[q] used multi-unit containment, in which parts of the containment envelope are shared among 4 or 8 units. The individual reactor containment buildings are all connected to a common vacuum building kept at very low pressure. Inside the vacuum building is an elevated water tank; when a LOCA occurs, the vacuum valves open (self-actuated on pressure differential), thereby connecting the vacuum building to the reactor building(s); and the contents of the water tank is sprayed over the vacuum building volume. These sprays are also self-actuated on the pressure differential caused by the LOCA. The water sprays condense the steam, and reduce the internal pressure. The containment pressure quickly goes sub-atmospheric, and remains so for several days after an accident, so the leakage is inward, not outward. The vacuum containment concept was developed for two reasons. First, it was thought to be economic on multi-unit plants. Second, because of the near-zero leakage, it is highly effective. It was first proposed by designers on Pickering A, to provide increased safety for a plant that would be located near a major population centre. The *disadvantages* of the vacuum containment concept are: it is more difficult to seismically qualify the long duct connecting the reactor buildings to the vacuum buildings; and if interest rates are high, the economic case is marginal, since one has to build most of the common containment structure before the first reactor can start up.

[q] Now Ontario Power Generation, or OPG

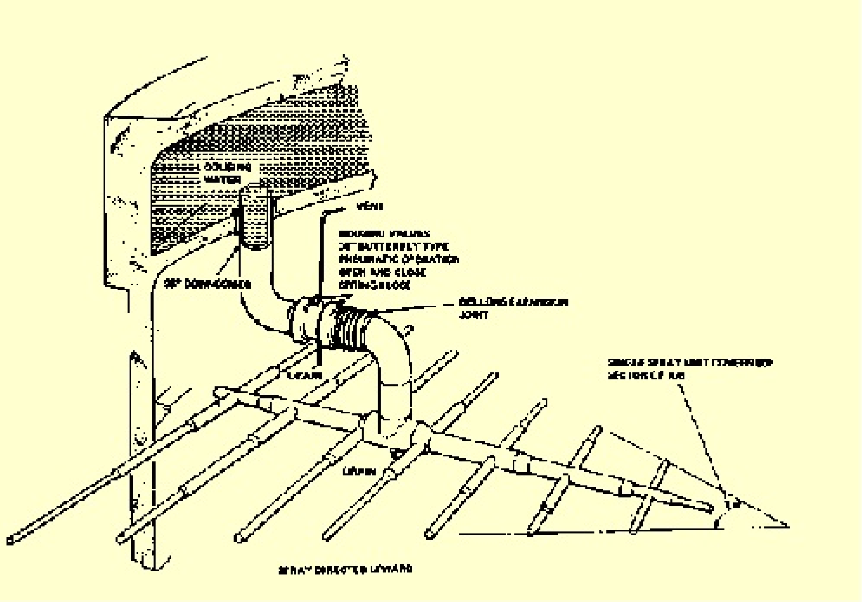

Figure 5-17 - Dousing Spray Arms

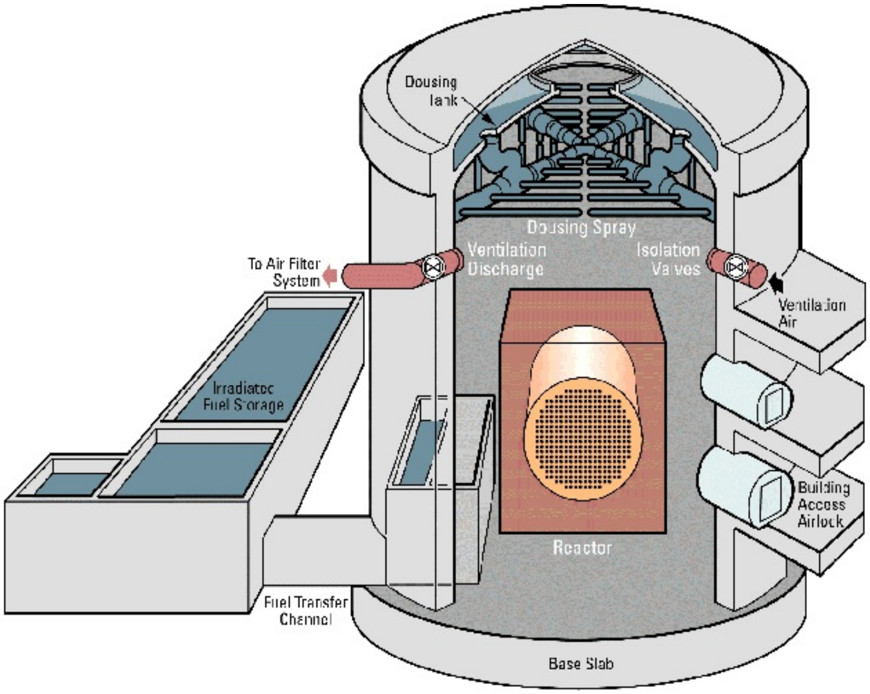

Figure 5-16 - Single-Unit Containment

Single unit CANDU 6s also use pressure-suppression (Fig. 5-16). There is no separate vacuum building, so the sprays and the elevated water tank are located in the actual reactor building. There are six spray arms (Fig. 5-17), each with spray valves. There are two valves in series on each spray arm (to avoid a spurious douse). The valves on three spray arms are pneumatically operated; the valves on the other three are electrically operated, for diversity. The containment building is prestressed, post-tensioned concrete with an epoxy liner for leak-tightness. Dousing starts automatically when the building pressure reaches 14 kPa (g) and is turned off if it falls to 7 kPa (g). This means that it cycles on and off for small breaks. However for large breaks it remains on until the dousing water is used up and is all on the basement floor of the reactor building. Since some of the dousing water is also used for medium pressure ECC, the connections of the spray headers are above the bottom of the dousing tank, to ensure that ECC water is not used up. Dousing is effective in washing soluble fission products out of the containment atmosphere. However it does not control pressure in the long term, since it is used up in the early part of an accident.

The ACR-1000 design is a single-unit containment - i.e. like CANDU 6 but without high-volume dousing sprays. The design pressure is higher and a steel liner is used instead of epoxy for increased leak-tightness at the higher pressure. A low flow spray controls containment pressure in the long term for severe accidents, and has some pressure-suppression capability for design-basis accidents in the shorter term.

Many other containment concepts have been developed. Here are just a few:

1. Many Boiling Water Reactors have a water reservoir circling the base of the building; steam produced in an accident is directed by pipes to this suppression pool and condensed. The "bubbler pond" on Russian RBMK reactors is a variant of this concept.
2. Recent Indian HWRs use a double containment, with a suppression pool to reduce internal pressure. The outer containment prevents leaks from the inner containment from escaping.
3. Some PWRs use ice reservoirs in the building, to which the steam is directed and condensed.
4. Some early reactors (e.g., NPD in Canada) used a pressure-relief containment. The first pulse of steam is released to atmosphere, and then the containment isolates. The idea was that since the fuel ratings were so low, the release of fission products would be delayed relative to the initial release of steam.
5. The Westinghouse AP-1000 containment consists of an inner steel pressure shell surrounded by a concrete outer shell; water flows by gravity over the inner shell to provide passive heat removal via heat conduction through the shell, aided by air convection between the two shells. We discuss this in a later chapter.

In the longer term, heat can be removed by containment air coolers (Figure 5-18). These require both water and electrical power. Alternatively low-flow sprays can also be used, with heat being removed from the spray water by heat exchangers; and the cooled water pumped back up to the spray arms.

In normal operation, containment is *not* a sealed system. Many fluid lines penetrate the building (e.g., steam lines, feedwater lines, service water lines, instrument lines). For CANDUs, the building is normally ventilated for atmospheric temperature control, especially since personnel access to parts of the building is required during operation. All these penetrations are pathways for release of radioactivity if an accident should occur; so on an accident signal, many of them are automatically isolated. There are two dampers in series on each ventilation line which are closed very rapidly on a containment isolation signal (for CANDU 6, high containment pressure

[3.5 kPa (g)] or high radiation in containment). Other lines penetrating containment are also closed on the same signal.

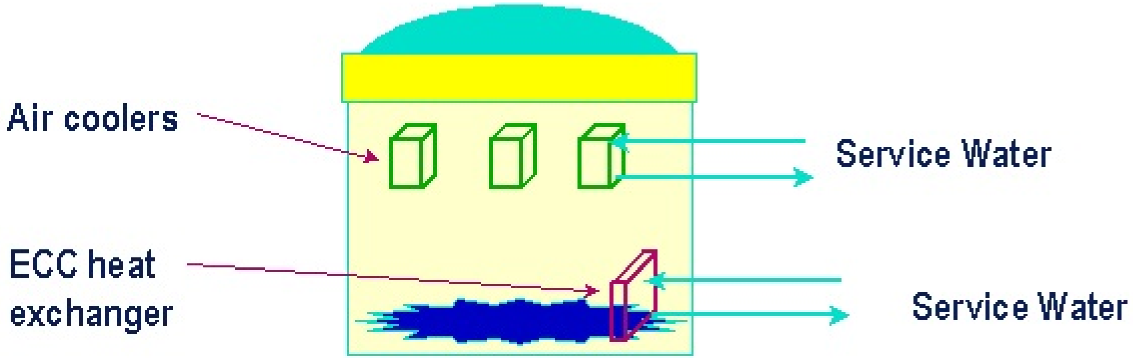

Figure 5-18 - Long-Term Heat Removal

Isolation

Isolation of steam, feedwater, and service water lines on an accident is a controversial issue. They are required to remove heat. If they are isolated, other means of removing heat must be provided. In CANDU, they are *not* isolated immediately, since it is believed that continuing to use a running system is more reliable than stopping it and starting up another one. For LWRs, the opposite approach is taken. For example, in LWRs, the main steam lines are isolated rapidly and automatically as part of the containment isolation. This closes a potential leak path (due to a leaking boiler tube prior to the accident, or a boiler tube failure as the initiating or consequential event) but requires a supplementary heat removal system to be brought in, in short order. In recent CANDUs, a compromise was reached: there are Main Steam Isolating Valves (MSIVs) but they are manually operated and slow - they are used selectively to control leakage in the longer term. Feedwater and service water are not normally isolated.

Reliability

As a special safety system, the containment must meet a demand unavailability of 10^{-3} years/year or less. Containment leak-tightness is tested every few years by pressurizing the building and measuring the leak-rate. This is an invasive and expensive test, and if the leak-rate exceeds the requirement, one must assume there has been a containment impairment for half the interval between tests. Thus on-line leakage monitoring systems are being deployed.

The containment isolation system is a subsystem of the containment special safety system; since the containment as a whole must have an unavailability less than 8 hours per year, the unavailability of the containment isolation subsystem must be less than that. It is tested during operation to prove its unavailability target is not exceeded.

A similar situation holds for the dousing subsystem. With two valves in series in each arm, each can be tested during operation without firing the system. However in two cases (once in Gentilly-2 and once in Point Lepreau), operator errors during test have led to a spurious douse.

Other Functions

Containment also acts as a barrier to protect reactor systems from external events (tornadoes, turbine missiles, aircraft crash and "malevolent events"). These may impose additional design requirements on the structure.

Hydrogen can build up in containment after an accident. After a LOCA, hydrogen is formed slowly by radiolysis of the water circulating through the core. Also a severe accident such as LOCA plus loss of Emergency Core Cooling can produce hydrogen early on due to the chemical reaction between the hot fuel sheaths and the steam in the fuel channels. (It is however less than the amount of hydrogen produced in a core melt in a LWR because the pressure tube conducts heat away from the fuel to the moderator.) The large containment building allows some mixing of the hydrogen due to natural circulation. Air cooler fans provide forced mixing. In addition there are igniters (44 in recent CANDU 6s) placed in various rooms to burn a local hydrogen concentration before it can detonate. Turning on the igniters is done automatically by the containment isolation signal. The requirement is to keep the local concentration of hydrogen below the lower limit for detonation (9%) - controlled slow burning is acceptable.

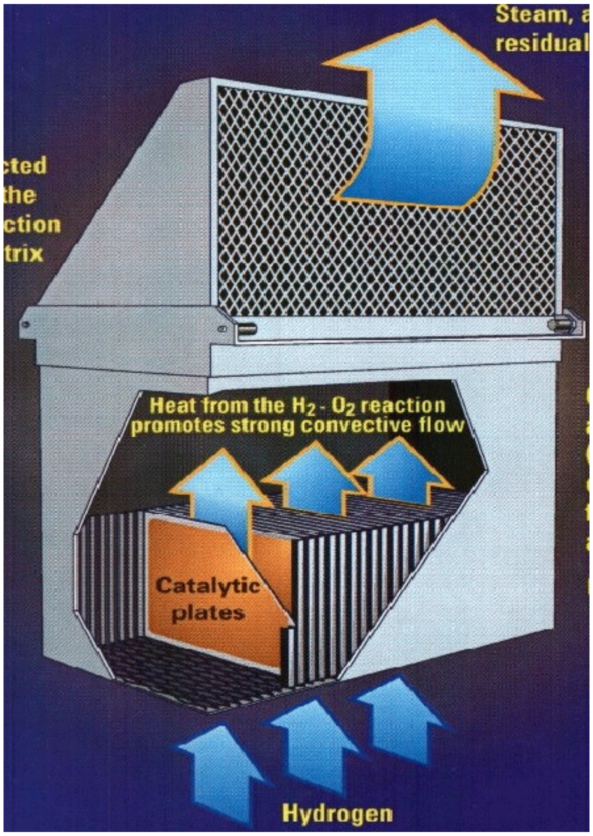
Steam, a
residual
cted
the
ction
trix
Heat from the H_2 - O_2 reaction
promotes strong convective flow
Catalytic
plates
Hydrogen

Figure 19 - Hydrogen Recombiner

Passive autocatalytic recombiners[r] are being developed (and are part fo the ACR-1000 design) for long-term hydrogen control - they have the advantage of not needing electrical power and are always ‘on’ (Figure 5-19).

Monitoring

For most accidents, the plant state is monitored from the main control room (MCR), and the safety functions of shutdown, heat removal and containment can be performed from there. Some accidents however can render the MCR uninhabitable: for example earthquakes[s], fire in the MCR, hostile takeover, aircraft strikes, and high radiation fields. A Secondary Control Area (SCA) is provided for such eventualities: the operators relocate to the SCA and can perform the required safety functions therefrom. In CANDUs, the SCA is required after an earthquake; for ACR-1000, the MCR and essential equipment within it has been qualified for post-seismic accident mitigation.

[r] As the words suggest, they present a catalyst bed to the containment atmosphere, on which the hydrogen recombines with oxygen. The heat of reaction causes a convection flow through the device which helps mix the containment atmosphere.

[s] The MCR is seismically qualified not to collapse and injure operators in an earthquake. However the Group 1 equipment in the MCR is not seismically qualified.

References

1. A.J. Muzumdar and D.A. Meneley, "Large LOCA Margins in CANDU Reactors -- an Overview of the COG Report", Proceedings of CNS 30th Annual Conference, Calgary, AB, 2009 May 31-June 3.
2. “Design of New Nuclear Power Plants”, CNSC report RD-337, November 2008.

Exercises

1. Using the SLOWPOKE 10MW heating reactor which you analyzed in the earlier lecture, describe the possible shutdown system requirements in terms of design, rate, depth, signals, margins, environment, and independence. Again the "right" answer (which of course presumes the designers were "right") is less important than an original well-reasoned case.

2. The ACR design uses light water coolant and heavy water moderator. What are the implications on the ECC design compared to CANDU 6, say? Consider the impact on choice of ECC fluid and design of isolation devices. What simplifications could be made to the overall concept (i.e., can you think of a different concept that would be effective)? Specifically how could the design in Figure 5-15 be simplified?

3. The ACR also uses enriched fuel, in a configuration where the void coefficient is small and negative (about -3 mk. total core void), compared to about +16 mk. for current CANDUs. How would that affect safety system design? Cover each one in turn and give reasons (you will be assessed on your reasoning, so do not just copy what designers have done on ACR). In particular estimate the speed of shutdown required for a large LOCA by assuming the coolant voids in ~2 seconds. (Estimate the amount of energy required to melt the fuel in each case. You will need to estimate the power transient also. You may find it easier to work in units of integrated power - e.g. full-power-seconds. You can assume that the hottest fuel pin has twice the power of a 'core average' fuel pin if you like.

4. Compare two-out-of-three voting logic on a CANDU Classic trip system (Fig. 5-6) with a two-out-of-four voting logic on ACR trip system. In particular, take a trip parameter such as low flow. Assume that each of the three (or four) trip channels has a demand unavailability of 0.01 and a spurious (fail-safe) trip frequency of 1 per year. Assume that a spuriously tripped channel takes an hour to detect and three days to repair. In the CANDU Classic case the spuriously tripped channel must be set in the fail-safe position (tripped) until it has been repaired; but for a four channel system, it may be re-set in the un-tripped position (since the remaining channels still form a 2 out of 3 system). What is the demand unavailability of the loss of flow trip in each case? What is the spurious trip frequency in each case? Name any disadvantages of the two-out-of-four system.

5. What are the options for heat removal from containment after a severe accident (core damage)? What are the pros and cons of each of your options?

6. Propose an alternative ECC fluid for CANDUs which might be better than light water. What are its advantages and disadvantages? How would the ECC design change?

7. What are some of the mechanical failure modes of a CANDU shutoff rod?

8. Calculate the amount of energy that can be added to a fuel pin in a large LOCA, to prevent fuel melting. Use symbols if you don't have the data.

9. Using the ZED-2 zero-power research reactor which you analyzed in the earlier chapter, describe the possible shutdown system (moderator dump) requirements in terms of design, rate, depth, signals, margins, environment, and independence. Again the "right" answer (which of course presumes the designers were "right") is less important than an original well-reasoned case.

10. Calculate the demand unavailability of one ECC train as shown in Figure 5-15. Assume the failure probability of a normally-closed motorized valve to open is 0.01; the failure probability of a normally closed check-valve to open is 0.001; the failure probability of a rupture disk to fail to open when the pressure difference across it reaches the design value is 0.001; and that a normally open valve is left closed in error after maintenance for 1 hour per year on average before being detected and corrected. List the single failure points of the train. How could you test that the check valves were not stuck? (Look it up). What limitations are there to such a test?

11. Look up (e.g. from the USNRC web site) the design of either the EPR or AP-1000 decay heat removal systems. Summarize them and discuss advantages and disadvantages compared to the CANDU decay heat removal systems (choose one CANDU plant for your comparison). If you don't have access to CANDU information, simply compare EPR and AP-1000 decay heat removal systems.