CHAPTER 10

Instrumentation and Control

prepared by
Dr. G. Alan Hepburn
Independent contractor (AECL Retired)

Summary:

This chapter describes the role of instrumentation and control (I&C) in nuclear power plants, using the CANDU 6 design as an example. It is not a text on the general design of instrumentation and control algorithms, a subject which is well covered by many textbooks on the subject. Rather, it describes the architectural design of these systems in nuclear power plants, where the requirements for both safety and production reliability are quite demanding. The manner in which the instrumentation and control components of the various major subsystems co-operate to achieve control of the overall nuclear plant is described. The sensors and actuators which are unique to the nuclear application are also described, and some of the challenges facing designers of a future new build CANDU I&C system are indicated.

Table of Contents

List of Figures

1 Introduction

This chapter describes the role of instrumentation and control (I&C) in nuclear power plants, using the CANDU 6 design as an example. The scope of I&C includes:

- Implementation of control strategies for those control functions which are automated,
- Presentation of information to the operator and receipt of operator inputs for those functions which are under operator control,
- Initiation of reactor shutdown, emergency coolant injection, and containment isolation in the event of failure of the above control functions, and
- Data acquisition.

The automated functions include:

- Automatic control of the reactor, balance of plant, and auxiliary systems;
- Activation of the special safety systems;
- On-power refuelling (CANDU reactors);
- Human/machine interface.

Virtually all the systems in a nuclear power plant contain an instrumentation and control element.

Although some of the material presented is common to other reactor types, the design details described pertain to the CANDU reactor. The CANDU 6 design as implemented at sites in Canada has been chosen as the reference because it is the most widely deployed CANDU design world-wide. While the implementation of I&C in other plants is broadly similar, the differences between the CANDU 6 and other CANDU stations are significant. These differences will be noted in some cases where they are of particular interest to the understanding of I&C in general.

At the time when the CANDU 6 and many of the world's light water reactors were designed, the process I&C subsystems in nuclear plants were implemented using a combination of individual analog control loops and discrete Boolean logic using relay technology. The design of the analog loops is based on classical linear frequency-domain control theory as described in any textbook on the subject.

As a consequence of the large core of the CANDU reactor and of its many fuel channels, reactor inlet and outlet headers, and other components, the CANDU design is very extensively instrumented. As CANDU reactors began to exceed about 250 MWe, the size of the reactor core and the on-power refuelling combined to make manual supervision of the spatial distribution of power in the core more and more unwieldy. The result was a strong motivation to introduce computer control of the reactor and of key process control loops. The resulting control algorithms are quite complex, with many related inputs and outputs. Implementing such a system using conventional equipment would have been impractical. Therefore, a central dual-redundant digital control computer (DCC) system, in which the control logic is defined by software, was introduced in the CANDU design very early on relative to the I&C industry as a whole, to say nothing of the nuclear power industry. All CANDU reactors, starting with Douglas Point, have used computers in their process I&C systems. At least from Pickering on, the use of this technology was a matter of design necessity.

This chapter is intended to provide an introduction to the role of I&C in nuclear power plant applications and is consequently oriented towards issues that are unique to the nuclear context. It is not a text on the basics of instrumentation and control, nor is detailed knowledge of this field necessary to understand the material presented. The details of the control strategies used in each plant system are not covered. The only detailed discussions of these strategies presented are of the subsystems involved in reactor and overall plant control and of the regional overpower trip logic.

For those interested in studying the individual systems in more detail, there is a plethora of material available on the CANTEACH Web site (https://canteach.candu.org). In reviewing this material, the author noted that it contains many references to design features that are unique to individual sites as if they were part of the generic CANDU design. For those who need to know the details of a specific station, the most reliable source is the system design manuals pertaining to that station.

1.1 Overview

The reliability and availability requirements for nuclear power plants tend to differentiate them from many other applications. The implications of these requirements for the major I&C subsystems are discussed from the point of view of both nuclear safety and plant production. References to current standards and regulatory documentation are provided.

The overall architecture of I&C systems is described, using the reference CANDU 6 design as an example. There are close ties here with the architecture of the electrical power systems described in Chapter 11.

The implementation technologies used in nuclear I&C are also described. As an aid to understanding the design decisions evinced by the design, a brief discussion of the technologies available to the designers in the early 1970s is provided.

To give the reader an insight into the more detailed I&C logic functions typical of the CANDU design, the operation of the main I&C systems involved in overall control of the plant (reactor regulating system, boiler pressure control, and unit power regulation) is described in Section 4. The design of the I&C subsystems of the special safety systems is described in Section 5. The tools and techniques used to verify the design are discussed in Section 8.

I&C technology has arguably experienced the most dramatic change of any of the technologies used in power plant design in the years since most of the world's nuclear fleet was constructed. However, at least in the case of CANDU, although a number of I&C subsystems have been replaced by more modern equipment at various stations over the years, the lack of a truly new build design means that the existing I&C design is now very dated. If and when the next new build is undertaken, the designers will face an extremely challenging task in bringing the design concepts and implementation technologies forward into the 21st century.

Several new-build designs have been initiated in the interim, notably for the CANDU 3, CANDU 9, and the Advanced CANDU reactor, but none of these has been carried out to the point where detailed design documentation was produced and implementation planning was documented, let alone where a reactor had been constructed, commissioned, and licensed. Furthermore, I&C technology moves forward perhaps one generation every ten years, so new designs have a rather limited shelf life.

Given the small number of new builds elsewhere in the world, many plants share this tendency to have outdated I&C technology. A discussion of some of the issues facing the designers of an upcoming new build (as opposed to a new replicate design) is included in Section 9, although, to be sure, any such discussion will be overtaken rather rapidly by technological developments.

1.2 Learning Outcomes

The goal of this chapter is for the student to understand:

- The role of I&C systems in the safety and process systems of the CANDU design,
- The high-level safety and production requirements applicable to these systems,
- The role of system architecture in achieving safety and production reliability,
- The technologies used to implement these systems and the constraints present at the time of their design that influenced the design choices made,
- The unique features of the CANDU I&C design and why they were adopted, and
- Some of the major challenges facing the designers of future CANDU I&C systems.

2 Nuclear Safety and Production Requirements for I&C Systems

A frequent reaction of the general population to the idea of nuclear power is, “What happens if it goes out of control”? They immediately have visions of mushroom clouds and grainy black-and-white movies from the Second World War. In fact, a nuclear explosion of this sort is simply not possible with a power reactor. To create a nuclear explosion requires a small mass of highly enriched uranium to be held very firmly together using conventional explosives for the short time it takes for the chain reaction to go massively supercritical and the consequent release of energy to occur.

In a power reactor, maintaining the configuration of the fissile material is also key to maintaining the chain reaction, but the fuel is not nearly as highly enriched (or in the case of CANDU, not enriched at all), and the fissile material must be maintained at a predetermined separation of the order of a few tens of centimetres, with some moderating material, such as graphite, water, or heavy water in the intervening space. If these conditions are not met, the chain reaction will not be sustained.

In any nuclear reactor at constant power, the rate of production of neutrons by fission is exactly matched by the rate of re-absorption in new fissions and by various losses. The neutron multiplication factor $k = 1$. If $k >1$ for some time, power will increase rather quickly to the point that so much heat is produced that the fuel melts and the core disassembles, thereby destroying the geometry required for criticality. But there will not be a nuclear explosion. There may be a steam explosion, as happened at Chernobyl, but that, while undeniably violent, is not in the same league as a nuclear explosion.

Although a nuclear explosion is not possible, in the absence of effective process control and other defense mechanisms, the reactor would be destroyed, and radioactive material could be released to the environment, causing an ecological disaster like that which occurred at Chernobyl. Because there is far more fissile material in a power reactor than in a weapon (tonnes vs. kg), the result would be widespread contamination, which is not an acceptable outcome.

For these reasons, a control system is needed to bring the chain reaction to a useful power level and then to hold it there. For thermodynamic efficiency, the hotter the energy source (fuel elements), the better. However, if dryout is ever allowed to occur, cooling deteriorates quite rapidly, resulting in fuel melting. Therefore, power has to be raised to some point below dryout, then held close to constant by maintaining k = 1. This manoeuvering and control is achieved by adding or removing neutron-absorbing control elements: poison, rods, and in the case of CANDU, light water. Insertion of each reactivity-control mechanism results in a reduction in the neutron multiplication factor. The absorption value of these devices is expressed in milli-k (mk). Insertion of one mk of negative reactivity decreases the neutron multiplication factor by 0.001.

In a light water reactor, fresh fuel requires maximum negative reactivity (rod insertion). As fuel burns up, rods must be gradually withdrawn to sustain the chain reaction. Power distribution within the core is highly predictable because the rods are withdrawn in a predetermined pattern. In CANDU, with continuous refuelling, fuel reactivity can be maintained indefinitely, but local peaks will occur when fresh fuel is inserted. This, together with the larger size of the core, which in the absence of control action can lead to local flux oscillations due to Xe^{135} (a neutron absorber with a fairly short half-life which is produced as a result of the fission process), means that, in a CANDU reactor, flux has to be controlled both as an average value, for the entire core, and spatially, to avoid unwanted flux tilts.

Heat must also be removed from the fuel as it is produced. This is achieved by circulating a heat-transport fluid (e.g., light or heavy water) over the fuel. The energy removed is transferred to the secondary coolant circuit, where it produces steam to drive the turbine. Failure to remove heat can result in damage to the fuel or to heat-transport system components, so control of the heat-transport system is also critical.

The safety role of the process control systems is to keep the various reactor systems operating within predetermined safe limits. Given the potential consequences of process-system failure, however, the defenses against such failure must be extremely robust—much more robust than could credibly be achieved by the process systems themselves. Several additional layers of defense are used in any modern power reactor. The probability that each layer will not be available to accomplish its function is expressed as some unavailability figure, e.g., 10^{-3} years per year. In the case of a defense mechanism, this is equivalent to saying that the mechanism will operate as expected 999 times out of every 1000 challenges.

In any modern reactor design, the additional levels in this defense-in-depth approach are:

- The shutdown systems,
- The emergency core cooling system (ECCS), and
- The containment system.

In CANDU plants, these systems are referred to as the “special safety systems”, while in PWRs, they are referred to as “reactor protection systems” and “engineered safety features”. In the CANDU context, anything else is a “process system”. During plant operation, the special safety systems are poised ready to perform their function when called upon to do so, while the process systems are generally in continuous operation. However, some process-system functions are also normally dormant, but take action to preclude the need for special safety system intervention.

The role of the shutdown system is to stop the chain reaction very rapidly (generally in less than two seconds) if there is an indication that the process parameters are going outside acceptable limits. Shutdown systems achieve their function by rapidly inserting large amounts of negative reactivity. This is typically done in all reactors by inserting neutron-absorbing rods into the core using a highly reliable power source—the force of gravity (or pressurized gas accumulators, in the case of boiling water reactors), augmented initially by springs in the case of CANDU.

The CANDU core has a slightly positive void coefficient. This means that, as power increases through the point that the primary coolant starts to boil, the reactivity would also increase in the absence of any control-system action. In other words, a positive feedback situation exists. On initial consideration, this does not appear to be a desirable characteristic, because an event that results in a power increase is not self-limiting. The Canadian industry's response to this has been to provide a second separate, diverse shutdown system (SDS2). In all designs from Bruce A on, both shutdown systems are equally capable. The second system injects gadolinium nitrate poison into the moderator, using compressed helium gas as the motivating force.

Every effort is made to ensure that the two shutdown systems are independent (different reactivity mechanisms, different design teams, complete electrical and spatial separation of instrumentation and mechanisms, different I&C devices wherever possible and practical), so that the design unavailabilities of 10^{-3} years per year for each system can be multiplied to achieve an overall unavailability of the shutdown function of 10^{-6} years per year.

Proponents of light water reactors tend to make much of the CANDU's positive void coefficient, but the comparison with respect to inherent power dynamics is not all one-sided. Clearly, in addition to the direction that power tends to move, the rate at which it moves is also important. Although a milli-k may not sound like much, the average lifetime of a neutron is less than one millisecond. Considered simplistically, this would mean that an excess reactivity of only 1 mk would result in an approximate tripling of neutron power within one second. Fortunately, this simplistic approach is not applicable in a power reactor core. Due to the effects of a relatively small number of delayed neutrons (about 6%), the increase in power in one second would, for a CANDU, be around 1%. In a light water reactor, it would be around 10%. This is one key difference between the CANDU core and a typical LWR. For a more complete explanation, see [Rouben2002] and Chapter 5, Section 6. Moreover, once power reduction starts, the strong negative power coefficient of a light water reactor tends to resist the desired power reduction, prolonging heat generation in the fuel. In a light water reactor, the shutdown system alone does not insert enough negative reactivity to shut down the reactor. These reactors also rely on an inherent feature of the physics of the light water core—the Doppler resonance phenomenon—to help terminate the reaction [Rouben2008].

Stopping the chain reaction is necessary, but not sufficient to mitigate the effects of a process-system failure. The decay heat, which initially amounts to about 6% of the pre-shutdown value, still has to be removed. In a fossil-fuel plant, cutting off the source of new energy (the fuel feed) instantly removes the heat source. In a power reactor, even after the chain reaction has been terminated, decay heat could cause the fuel to melt, with the resulting undesirable consequences, for many days after the reactor has been shut down, unless continued cooling of the fuel is provided. Usually, the heat-transport system is available to do this, but to guard against any failures in this system, a backup cooling system is required—the emergency core cooling system, whose role is to provide an alternate path for removal of decay heat from the fuel if there is an indication that the process systems responsible for doing this have failed.

The last line of defense is the containment system, whose role is to provide an envelope around the parts of the plant that contain fission products so that this material will not be released to the environment. The containment system also condenses any steam released into this envelope, thus limiting any upward pressure excursion following such a release.

2.1 Requirements for the Special Safety Systems

Oversight of the nuclear industry in Canada is the responsibility of the Canadian Nuclear Safety Commission (CNSC). The forerunner of the CNSC was the Atomic Energy Control Board (AECB). One of the CNSC's responsibilities is to develop the regulations governing the design and operation of nuclear power plants. The regulatory requirements for nuclear power plant safety in Canada are contained in the AECB Regulatory Documents listed in the references.

The reliability requirements for the special safety systems, which have a dominating influence on both the architecture and detailed design of these systems, are:

- SDS1 and 2: All Canadian power reactors are required to have two independent shutdown systems. As stated earlier, this requirement came about as a result of a design solution to an inherent characteristic of the CANDU core. This is now enshrined in a regulatory requirement, AECB regulatory document R-10 [AECB1977], which requires two shutdown systems "unless otherwise approved by the Board".

The overall unavailability for each of these systems is required to be less than
10^{-3} years per year.

- ECCS and containment: The relevant AECB regulatory documents [AECB1991b, 1991c] require that each of these systems have an unavailability of less than 10^{-3} years per year.

These numerical requirements are derived from the need to achieve what is deemed to be an acceptably low release rate of radioactivity to the public following a number of postulated accidents, or design basis events (DBEs). In other words, these are not just arbitrary numbers.

2.2 Safety Requirements for the Process Systems

The process I&C systems also play a role in plant safety. They have to keep the process systems operating within their designed operating envelope. If a convincing case could be made that the process systems would have an extremely high probability of success in doing this, there would be no need for the special safety systems. However, the reality is that, unless the design of the reactor and its associated process systems is inherently fail-safe, independent special safety systems are the only credible way to achieve an acceptable level of safety.

The unavailability targets for the special safety systems are predicated on their not being challenged very frequently by failure of the process systems. For this reason, the process systems do have some safety requirements, depending on the safety-related role of the process system in question. The system singled out for special mention in the case of CANDU reactors is the reactor regulating system, which has its own CSA Standard [CSA2011a]. The requirement for this system, per 4.3.1.1 of that Standard, is that:

> "The design target for failure of reactor power control shall be established by probabilistic safety analysis methods.
> **Notes:**
> In CANDU nuclear power plants, the design target frequency for loss of regulation is historically less than 1 in 100 years."

> A “loss of regulation” (LOR) is defined to be “a failure resulting in an unplanned increase in bulk reactor power”.

Unlike the AECB/CNSC regulatory documents, this CSA Standard has recently been updated to reflect the possibility of non-CANDU designs being licensed in Canada. That “1 in 100 years” number used to be a requirement. Now, whatever design number is used for failures to control reactor power has to be reflected in the overall plant safety analysis.

Because any “unplanned increase in reactor power” should be terminated by the shutdown system, one can conclude that a loss of regulation has by definition occurred if either of the shutdown systems trips, unless it can be shown that the shutdown system itself activated inadvertently. In many cases, it is possible to analyze the event in retrospect and to demonstrate that the shutdown-system intervention was unnecessary, but this analysis still has to be carried out to confirm that the process I&C systems design was not at fault.

In earlier versions of the referenced CSA Standard, which were in force at the time the current Canadian reactors were designed, there was a requirement that “each reactor unit shall be designed and operated such that the combined frequency of all serious process failures does not exceed 1 in 3 years”. This number applied to the entirety of all process systems. In an ideal world, the I&C component of each process system would be allocated a portion of that once in three-year budget, and a reliability analysis would be performed as part of the design to demonstrate that this requirement was met.

Because an LOR is by definition attributable to RRS, it is apparent that the Standard has already taken care of the allocation of the RRS portion of the budget—namely 1 in 100 years. This includes the entire RRS, from sensors to reactivity control mechanisms. Breakdown of this 1 in 100-year number to each RRS subsystem, including the I&C subsystem, is quite correctly left to RRS designers.

Systems design in the 1970s and ‘80s was not pursued in as rigorous a manner as is now considered to be best practice, and the derivation of the safety-related and other reliability requirements for the various process I&C subsystems is typically not fully documented in the design record of CANDU plants. Until recently, design of process systems in the Canadian nuclear industry did not generally follow the detailed process of requirements analysis and allocation which is standard practice in the aerospace and software engineering industries, for example, although these techniques are used in the key areas of plant safety analysis and in the development of safety-critical software.

There are many qualitative requirements included in N290.4, but the reliability requirement quoted above is the only performance parameter stated.

2.3 Production Requirements for the Process Systems

Clearly, for production reasons, quite apart from safety implications, it is desirable that the process systems be sufficiently reliable that they do not result in frequent power reductions or shutdowns. It should be mentioned at this point that the physics of the CANDU reactor are such that, if the reactor is shut down from full power, the buildup of Xe^{135} in the fuel will result in the accumulation of so much negative reactivity that, if the reactor is not restarted within about 30 minutes, it will not be possible to start it up again for about 40 hours. This is known as a “poison-out” and is discussed in more detail in Chapters 5 and 13 (Section 5.3.1). Because many of the control systems must be available if the plant is to run for more than a few seconds, it can be seen that anything longer than a brief outage of one of these key systems will have a severe economic impact. Poisoning out is not an issue for the PWR/BWR.

Production availability is not, of course, a regulatory requirement, and therefore there are no regulatory documents which pertain to this key attribute. It will be seen in subsequent sections that the control systems are designed to fail safe, which will result in a shutdown of the reactor without the need for special safety-system action. It will also be seen that much of the key process control logic is implemented in a pair of dual-redundant digital control computers (DCCs). Although there is no specific requirement for availability of the process systems in general, these systems were designed to be highly reliable and tolerant of individual component failures. The DCCs were designed not to fail in a manner that leads to a poison-out more than once in three years. Note that, while this reads somewhat like the “serious process failure” requirement described in the previous section, the two requirements are quite distinct because, due to the fail-safe design of the process I&C systems, most process I&C outages will not result in a serious process failure as defined in Section 2.2.

3 Overall I&C Architecture

3.1 Architectural Design of the Special Safety Systems I&C

The architectural design of the special safety systems reflects the requirements for low unavailability discussed in Section 2.1. It was stated earlier that the CANDU design includes two independent shutdown systems. Independence is assured by geographical separation of the systems, including their I&C components, and by equipment and design diversity. For example, nucleonic sensors and actuators for SDS1 penetrate the core vertically from above the reactor, on the reactivity mechanisms deck. Nucleonic sensors for SDS2 are located at the side of the reactor, where the poison injection takes place. The actuation technologies are completely diverse, with mechanical neutron-absorbing rods used on SDS1 and injection of a neutron-absorbing solution into the moderator for SDS2.

To maximize physical separation between systems which must be physically independent, all safety-related systems in the CANDU plant are divided into two groups. Group 1 systems include most process systems and SDS1. Most of the other special safety systems, including SDS2, are allocated to Group 2. Within containment, because both groups must, after all, interface with a single reactor, it is possible to separate systems only by adopting approaches such as the horizontal/vertical separation described in the previous paragraph, but outside containment, the two groups are assigned completely separate locations. SDS1 I&C equipment, for example, is located in the control equipment room, adjacent to the process systems I&C equipment, but SDS2 I&C equipment is located in the secondary control area, on the other side of the reactor building. Penetrations through the reactor building wall for the two groups are separated by 90°. For a more complete discussion of grouping, see Chapter 13, Section 5.2.8.

This physical separation guards against relatively localized common-mode events such as mechanical destruction, fires, and flooding which might otherwise disable equipment in both groups simultaneously.

Not all common-mode events can be addressed by geographical separation and technological diversity. The sensors and actuators must be able to perform their functions during the accident conditions against which they are required to provide protection, which in many cases will impact both systems simultaneously (e.g., high radiation fields, temperature, and humidity). The equipment must therefore be qualified to function in the anticipated post-accident environment. Equipment must also be qualified to survive and function during more widespread common-mode events such as electromagnetic disturbances and earthquakes, against which only limited physical protection is possible.

The low unavailability of each special safety system is achieved in part by redundancy of the equipment within each system. To ensure independence of the instrumentation, it is channelized. This involves physical separation of the instrumentation, cable routes, logic equipment, and actuators, and the electrical supplies that power them. Typically, the special systems instrumentation and actuation logic is divided into three or more separate channels. The separation of the electrical supplies closely reflects the separation within the I&C systems, as described in Chapter 11.

Channelization minimizes the probability that many classes of events will disable more than one channel at a time. Because safety-system action requires two of the three channels to call for activation of the safety function, the system will continue to perform its design function even if the third channel has failed in an unsafe manner. The design of the individual channel logic is such that unsafe failure of even a single channel is unlikely.

Channelization also enables the logic in each channel to be tested periodically. A channel under test is placed in a state where it votes for safe action. In calculating the system unavailability, each component failure is assumed to be random and is detected by a test carried out at a specific test interval. On average, then, failures may go undetected for one-half this test interval. Therefore,

$$\text{unavailability} = \text{failure frequency x test interval} / 2.$$

It can be seen that the test frequency of functions that are normally poised but inactive is a key input to the unavailability calculation. Unless such functions are tested periodically, any system which depends on them must be assumed to be unavailable.

The detailed requirements for channelization in CANDU plants are documented in a series of safety design guides which are not in the public domain.

In PWR plants, a single shutdown system is normally used, with four channels of sensor and actuation logic. A reactor trip will result if any two of the four channels call for a trip. For CANDU plants, two three-channel shutdown systems are used. A reactor trip will result if any two of the three channels in a given shutdown system call for a trip. Either shutdown system will trip the reactor if the logic for any trip parameter calls for a trip in at least two channels. This is known as “general coincidence” logic. It is also possible to design a system in which a trip will occur only if the logic for the same parameter in at least two channels calls for a trip. This design is referred to as “local coincidence logic” and is used in some other CANDU plants. It can be seen that, although local coincidence logic is less prone to spurious trips, it requires more inter-channel communication, which compromises channel separation to some extent and adds complexity, particularly when the logic is hard-wired. In practice, the spurious trip rate has been found to be sufficiently small that the additional complexity of local coincidence logic is not warranted.

3.2 Architectural Design of the Process Systems I&C

The redundancy approach used in the special safety systems to minimize their unavailability is also used in the process systems to enhance the availability of key control functions, thus minimizing losses of production. Key control functions typically use triplicated sensors, with voting logic to reject sensors which have failed. Control logic is duplicated or triplicated, and much of the process equipment itself (pumps, valves, etc.) is either triplicated, using three 50%-capacity units, or duplicated using two 100%-capacity units. However, the separation requirements applicable to special safety systems do not apply to the process systems because the latter are not required to operate following an accident.

The architecture of the power supplies for the process I&C reflects that of the I&C equipment.

3.3 Features of the CANDU DCC Design to Enhance Production

As stated in Section 2.3, the economic penalties of even a brief process I&C outage in a heavy water reactor may well extend far beyond the duration of the equipment outage itself, and therefore, for both safety and production reasons, triplicated or duplicated equipment is extensively used.

The DCCs are a prime example of equipment duplication. The DCCs normally act in a master/standby configuration, with only one DCC in control of a specific function at any given time.

The individual control logic subsystems (or “control programs”) running on the DCCs perform their own internal self-checks. Normally, the programs running in the master DCC will be in control of the plant. If a control program in the master DCC determines that it does not have sufficiently reliable input signals to enable it to compute its outputs, that program will stop running, causing control of that specific function to transfer to the standby DCC. Plant control can continue with certain restrictions, with some functions being controlled from the master DCC and some having switched over to the standby machine. This feature is also used to reduce the risk of plant outages when a change is made to the control logic. The new logic can be run in the master DCC only, with the old logic available to take over in the standby DCC at the operator’s discretion if problems are encountered.

3.4 Safety-Related Functions of the Process Systems

The frequency of loss-of-regulation events can be reduced by incorporating logic into the process systems design to detect potential loss of control situations and to take action to reduce reactor power independently of the shutdown systems. In the CANDU design, this is accomplished by fail-safe design of the process I&C logic and by two layers of safety-related logic—the setback and stepback functions—implemented in the control computers. “Failing safe” in this context means failing the output devices in the direction of shutting down the reactor and positioning the process systems for post-shutdown conditions.

3.4.1 Setback function

The setback logic continuously monitors a number of process parameters. Typically, these are:

- High local neutron flux;
- Spatial control outside the normal range of operation;
- Low de-aerator level;

- High steam generator pressure; and
- Upsets in moderator temperature or pressure.

If there is an indication that these parameters are straying outside acceptable limits, a reactor power setback is initiated by gradually ramping back the power set-point generated by the set-point logic (see Section 4.1). In many cases, reducing reactor power will return the process system to its normal operating condition. Once a setback is terminated, operator intervention is required before any further automatic adjustment of reactor power set-point is possible.

3.4.2 Stepback function

Some abnormalities in the plant control systems require more rapid action than is available using a power set-point manoeuvre. The following conditions are monitored:

- Reactor trip;
- Turbine trip;
- Loss of line (grid connection);
- Heat-transport pump trip;
- High heat-transport pump pressure;
- High flux power or high flux rate; and
- Low steam generator level.

If any of these conditions is found to be present, an immediate power reduction, or "stepback" in reactor power, is initiated by releasing the clutches that hold the mechanical control absorbers out of the core. The stepback is normally permitted to continue to zero power, but in the cases of turbine trip or loss of line, it is arrested at an intermediate power level by catching the rods in mid-fall. This partial power reduction permits continued reactor operation at a power level that prevents the growth of xenon in the fuel from shutting down the reactor.

Stepback is the exception to the transfer-of-control approach described in Section 3.3. Both DCCs must initiate a stepback before one will take place. This reduces the risk that the stepback function will act spuriously. If a given DCC is shut down or the stepback function is not running on it, then the remaining machine can initiate a stepback on its own.

Although the stepback initiation logic resides in the DCC, it is independent of the RRS logic and therefore can mitigate some high flux power or flux rate events caused by failure of RRS itself. The stepback function reduces the frequency of demands on the two safety shutdown systems, SDS1 and SDS2.

3.4.3 DCC self-check function

Each DCC incorporates self-checking logic to confirm the availability of certain key functions. For minor losses of capability, such as loss of individual instrument measurements, the system will continue to function normally. However, if a major loss of capability is detected, the DCC in question will be shut down by allowing the computer's watchdog timer to time out. This results in all its outputs being de-energized, which, if the DCC in question is the master, transfers control to the standby machine.

An external watchdog timer, which must be reset every few seconds, is incorporated into each DCC. If this fails to happen, as would be the case, for example, if the DCC software got stuck in a loop, the watchdog timer would time out, and the DCC's outputs would again be de-energized.

The sense of the actuation logic for all systems controlled by the DCC is chosen such that de-energization of the DCC outputs will tend to move the affected process system in a safe direction. If both DCCs fail, for example, the mechanical control absorbers will drop into the core, and the light water zones will fill, resulting in a reactor shutdown.

Although the DCC itself is not seismically qualified, this watchdog timer function is. Hence, it is sometimes referred to as a “seismically qualified stepback”. This ensures that, the DCCs will be reliably disconnected from the process equipment they control should they fail as a result of a seismic event.

4 Overall Plant Control Functionality

The CANDU approach to overall plant control will now be described as an example of the overall plant control required for a nuclear power plant. Note that CANDU reactors are relatively more manoeuvrable than LWRs because they do not have to respect the rather slow manoeuvring rates imposed by a very thick-walled pressure vessel. The manoeuvrability of LWRs also tends to be restricted at the start and end of their fuel cycle.

Although most of the 200-plus systems in the CANDU plant contain an I&C element, three main I&C subsystems are involved in the control of the nuclear steam supply: the reactor regulating system (RRS), boiler pressure control (BPC), and the unit power regulator (UPR). Due to the complex nature of the control logic used, all three of these functions are implemented primarily in the DCCs. The term “boiler”, by the way, was changed at some point to “steam generator”, but the name of the DCC program has stuck.

The plant can be operated in two modes, as specified by the operator, the choice being represented by the position of the A/N (alternate/normal) switch shown in Figure 1:

1. In Normal mode, the operator specifies a plant electrical output in megawatts. In this mode, the UPR logic adjusts the governor valves supplying steam to the turbine to maintain the specified electrical output. The BPC function monitors boiler pressure, and any error will result in changing the requested reactor power set-point which is sent from BPC to RRS.
2. In Alternate mode, the operator specifies the reactor power output, in percentage of full power, and a manoeuvring rate. The alternate mode set-point is then moved to meet this request at the specified rate. RRS sets the reactor power, which determines the energy delivered by the primary heat-transport system to the boilers. BPC then manipulates the turbine governor valves to maintain boiler pressure at a fixed set-point. BPC also has control of the condenser and atmospheric steam discharge valves, which can be opened to receive steam if the turbine load is lost, thus isolating the remainder of the system from abrupt changes in generator load and permitting continued reactor operation without the turbine as a load, thus avoiding a reactor shutdown due to buildup of Xe^{135}.

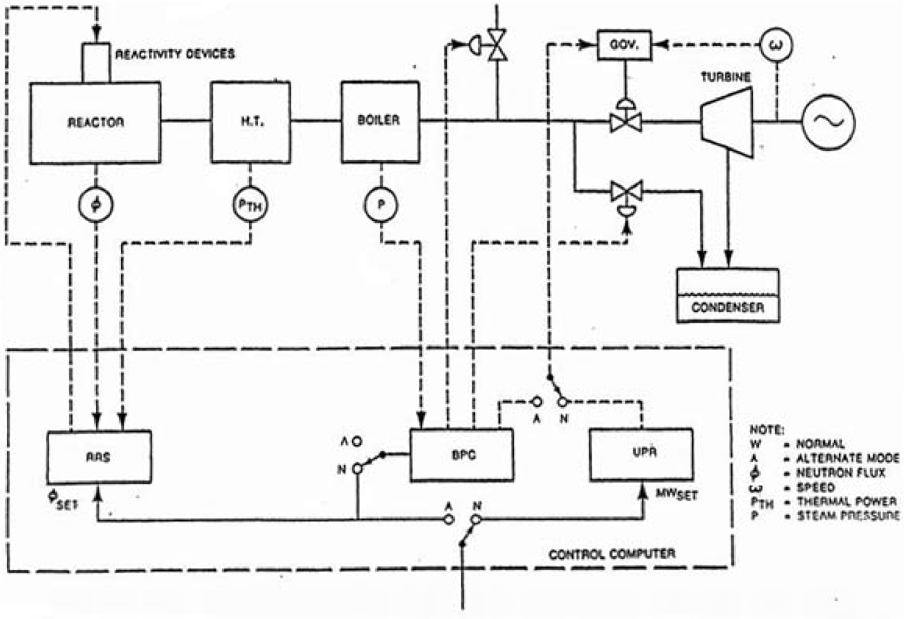

Figure 1 CANDU overall plant control

4.1 RRS control logic

RRS control logic in CANDU accomplishes the following:

1. Because the CANDU core is relatively large and subject to spatial instability due to the dynamic action of Xe^{135}, it has been found necessary to divide the reactor into 14 zones for control purposes (see Figure 2, which shows the arbitrarily labelled “A” and “C” ends of the reactor). There are seven zones at each end of the reactor core.

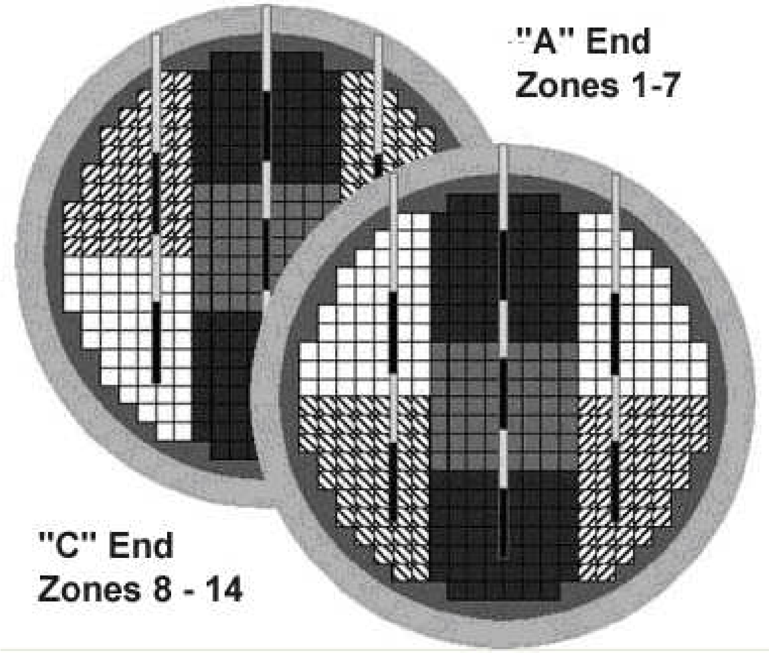

Figure 2 Arrangement of the 14 control zones

The power measurement and calibration logic determines the current reactor power in each of these zones, based on the readings provided by 28 platinum-clad Inconel®

flux detectors (two in each zone) that penetrate the reactor vertically from the reactivity mechanisms deck. The flux power is estimated for each of these 14 zones.

Because the platinum-clad Inconel detectors (see Section 6.1.1) do not provide an absolute value for neutron power, the zone power estimates have to be calibrated against thermal power. This is determined from measurements of temperature rise between the reactor inlet and outlet headers at low power, and from secondary-side measurements of steam flow, feed-water flow, and feedwater temperature at higher power (> 70%), because boiling in the channels renders the primary-side measurements unreliable at these power levels.

2. The appropriate power set-point is selected based on plant mode (Normal/Alternate, see Section 4). This set-point will be overridden if a setback is required (see Section 3.4). In this event, the reactor control mode is automatically switched to alternate mode, and the power set-point is ramped down to a predetermined endpoint. If the operator presses the HOLD POWER button on the RRS panel, automatic adjustment of the power set-point will also be suspended.

3. The demand power logic compares the overall reactor power measured in item (1) above with the required set-point and generates a power error signal E_p, which becomes the basis for the control of the reactivity mechanisms described next.

   Note that internally RRS works with logarithmic power, expressed in decades. This is appropriate because the reactor is a multiplicative device and the reactivity-control mechanisms control the rate of neutron multiplication. The reactivity mechanisms are driven based on E_p, which is also a logarithmic variable, measured in decades. However, in many explanations, and indeed in the displays seen by the operator, the E_p axis is labelled as “Power Error %” (see, for example, Figure 3). What is implied is that the value is a percentage of current power, not of full power.

4. During normal steady-state operation, the reactivity in each of the 14 zones is manipulated by adjusting the level of light water in each of 14 cylindrical compartments, one located in each zone (see Figure 2). In a heavy water reactor, light water is a neutron absorber. Each of the 14 compartments has a fixed outflow and a valve which controls the inflow of light water. These valves are manipulated in response to a combination of the overall power error, E_p, and any deviation of zone flux from the average over all 14 zones. The total worth of all 14 light water zones is about 7 mk. The logic to control overall power is run twice per second. Tilt control is updated every two seconds.

   In a zone where the overall error and the spatial components of the error signal sum to zero, inflow will equal outflow, and the light water level will remain constant at some intermediate level. The level of light water in each zone does not contribute to valve opening. However, as the level approaches either completely full or completely empty, the controlled variable changes from zone power to light water level so that the zones never completely flood or drain. These light water zone controllers are described in Section 6.2.

5. If more negative reactivity is required than the liquid zones are able to provide, as indicated by the zones becoming close to full, then this is accomplished by the me-

chanical control absorber logic, which drives banks of neutron-absorbing rods into the core. These absorber rods are normally located outside the core above the reactor. There are four mechanical control absorbers, with a total worth of 11 mk. The switching logic for them is shown in Figure 3.

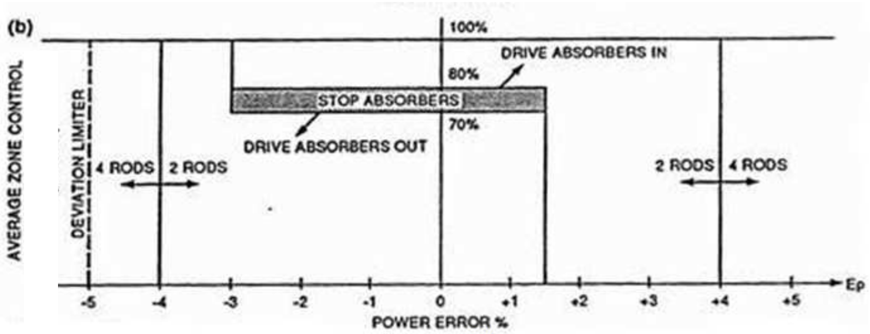

Figure 3 Absorber-rod drive switching

6. Similarly, when there is a lack of positive reactivity, the adjuster control logic can withdraw banks of adjuster rods from the core (see Figure 4). These rods normally reside in the core. There are 21 adjuster rods arranged in seven banks, with a total worth of 15 mk. All adjusters in a given bank are driven simultaneously, but only one bank is driven at a time.

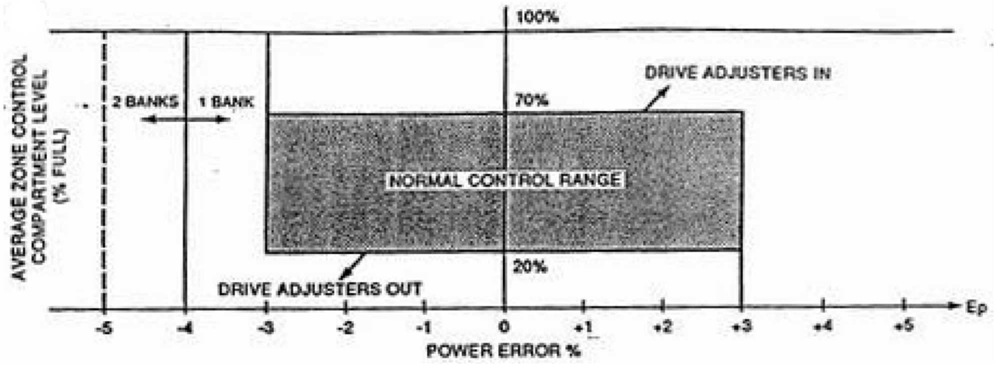

Figure 4 Adjuster-rod drive switching

The speed at which adjuster and absorber rods are driven is determined by a variable-frequency power supply and is a function of power error, as shown in Figure 5.

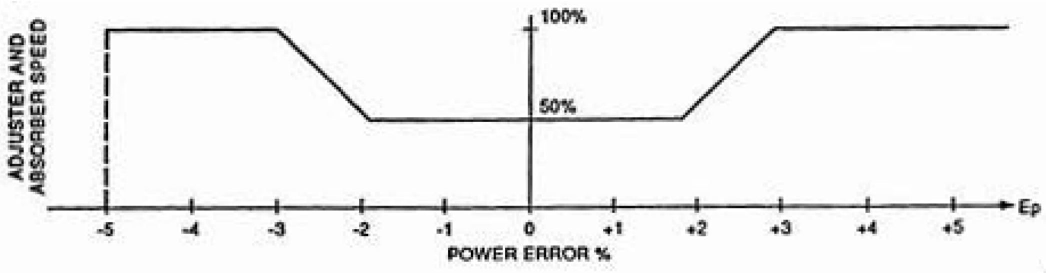

Figure 5 Rod drive speed

Once drive of either adjuster or absorber rods is started, it proceeds to completion to avoid large top-bottom flux tilts.

The absorber rods also incorporate clutches, similar to those used on the SDS1 shutdown rods. When a stepback is required, these clutches are disengaged, allowing the absorbers to fall into the core under gravity. In this way, a much faster power reduction is achieved than if the rods were driven electrically.

In the long term, if the light water zones tend to move significantly away from the 50%-full point, the operator can adjust the available reactivity by adding poison to or removing poison from the moderator.

4.2 BPC Control Logic

During operation at significant power levels, the goal of BPC is to maintain a constant, predetermined boiler pressure. The BPC program also incorporates logic to handle the warming up and cooling down of the plant, but this will not be discussed here in the interests of brevity.

In Normal mode, BPC will pass a power set-point to RRS to maintain constant boiler pressure despite possible variations in the steam taken by the turbine. Hence, the reactor is said to follow the turbine in this mode. The set-point includes proportional and integral terms based on pressure error.

If boiler pressure becomes excessive, BPC can manipulate a combination of condenser steam discharge valves (CSDVs) and atmospheric steam discharge valves (ASDVs) to control boiler pressure. The CSDVs cause steam to bypass the turbine and go straight to the condenser. The ASDVs discharge steam directly to the atmosphere.

In Alternate mode, BPC also controls the steam fed to the turbine by manipulating the turbine governor. In this mode, the reactor power set-point is specified by the operator, and the resulting thermal power generates steam in the boilers. The objective of BPC is to maintain constant boiler pressure. Any deviation in boiler pressure will result in the governor valves being manipulated to maintain boiler pressure, and hence the reactor is said to lead the turbine in this mode. The algorithm in this case uses proportional and derivative terms based on pressure error, plus a feed-forward term based on the rate of change of reactor power.

In addition to manipulating the governor, BPC drives the turbine load limiter, which reads the current load set-point and drives the load limiter to a point 100 MW above this value.

4.3 UPR Control Logic

When the plant is operating in Normal mode, UPR manoeuvres the turbine load towards the target load specified by the operator. It maintains a variable LR (load reference), which is ramped towards the target load at a rate which is also chosen by the operator. The actual turbine load is compared to this load reference, and the turbine speeder and load limiter are adjusted accordingly.

In Alternate mode, UPR has no control function. It fulfils a monitoring role only.

5 Special Safety Systems Functionality

The four CANDU special safety systems will now be described.

5.1 Shutdown Systems 1 and 2

Safety analysis determines a safe operating envelope for the plant. The values of a number of plant parameters are then monitored in each shutdown-system channel to confirm that this envelope has not been exceeded. If any parameter exceeds a predefined trip set-point, then the channel votes for a reactor trip. To guard against modelling errors, the CANDU design uses two diverse parameters in each shutdown system where practical to protect against each postulated initiating event. The case of impracticality has been invoked in the regional overpower protection (ROP) logic, which provides the primary defence against a slow loss-of-regulation accident. The ROP sensors use the same technology in both SDS1 and SDS2, and there is no backup nucleonic parameter for a slow loss of regulation, although the heat-transport high-pressure parameter does provide some protection in this case. The log-rate trip provides some diversity for fast LORs.

Most of the parameters that trip the shutdown systems are sensed using conventional process instrumentation, which provides fairly accurate indications of process values. However, in the regional overpower detection function, the platinum-clad Inconel sensors provide a much less direct indication of the physical parameter of interest: the onset of dryout in each bundle in the core. The logic used in the ROP trip will be described in Section 7, after the limitations of the sensors used have been introduced in Section 6.1.1.

5.2 Emergency Core Cooling System

The purpose of the Emergency Core Cooling System (ECCS) is to provide an alternate means of cooling the fuel when a loss-of-coolant accident (LOCA) has occurred. For a discussion of LOCAs of various degrees of severity, see Chapter 13, Section 5.4.

Emergency cooling occurs in two phases:

1. The ECCS monitors heat-transport system pressure and initiates high-pressure coolant injection whenever this pressure drops below the set-point. Valves are opened, which results in light water being forced into the reactor inlet and outlet headers, using pressurized helium as the motivating force. This is followed by a medium-pressure coolant injection phase in which light water from the dousing tank becomes the source of coolant.
2. When the supply of light water is exhausted, low-pressure recovery is initiated. This involves collecting spilled liquid from the reactor-building sump and recirculating it into the core through the reactor inlet and outlet headers. Both medium pressure and low pressure coolant injection are effected by pumps powered by Class III electrical power.

The sensors and logic used by the ECCS are conventional, but must of course be qualified to function in a LOCA environment.

5.3 Containment System

For a full description of the containment system, see Chapter 13, Section 5.5. The containment itself is the envelope designed to contain any release of radioactive material. In normal operation, it is maintained at sub-atmospheric pressure by vacuum pumps which exhaust to the atmosphere through a filtration system. Following a significant release within the reactor building, which is detected by high pressure, radiation, or both, the containment must be isolated. Any ensuing rise in reactor-building pressure is limited to facilitate the task of keeping leakage within allowed limits. On detection of a release, the containment is isolated by means of dual valves or dampers incorporated in every line or duct that penetrates the containment envelope. The dampers isolate the reactor building ventilation ducts and are pneumatically operated. Examples are the ventilation ducts, the spent fuel port, and the feed-water and steam lines.

The pressure rise is limited by the dousing system. In the upper area of the building, a dousing tank contains water which is released as a spray to condense steam resulting from a heat-transport system or steam-line break. Dousing is initiated by opening a combination of electrically activated and pneumatically actuated valves located beneath the dousing tank. Series valves are used to minimize the probability of inadvertent dosing actuation. Diverse actuation sources enhance the probability that dousing will occur when required. (This applies to single-unit CANDU stations. Multi-unit stations use a vacuum building common to all four units. In the event of a LOCA, the affected containment is connected to this vacuum building, and the steam emanating from the LOCA site is doused there).

A cooling system limits the temperature rise within containment, thus maintaining the integrity of the building in the long term. The cooling system I&C is conventional. A hydrogen ignition system ignites any free hydrogen gas before concentrations can become high enough to be hazardous.

6 I&C Systems Layout and Equipment

In everything that has been said up to this point, the reader may have gained the impression that operation of the nuclear unit is entirely automatic. In the short term (up to about 30 minutes), this is indeed the intent of the design. Experience has shown, however, that there is a role in plant operation for both automated and human control. Humans are not particularly good at performing repetitive, routine tasks without error, and of course most of the control loops involved are simply too fast and have too many interactions with other loops for a human operator to be involved in controlling the loop in real time. However, the human operator has proved to have useful capabilities when things depart from the routine, where perhaps the state of the plant is not completely clear, and when available information is either incomplete or conflicting. He can then assume an overview role and direct the progression of events towards a stable conclusion. To do this, he has to be involved in the operation of the plant in a supervisory role and be presented with available information in a useful way.

The control room is the centralized point for reactor operation. A geographically separate secondary control area is provided to shut down the plant and to provide monitoring of those parameters which confirm safe shutdown following the event, should the main control room become uninhabitable. The Group 2 I&C equipment is located in this area.

The chief operator’s station in the CANDU control room is the operator’s desk, shown in Figure 6. A separate shift supervisor’s office is also provided, but he is not immediately involved in unit operation. The operator’s desk is equipped with video display units (VDUs) driven by the DCCs. The operator can select from a large number of displays which provide information on the status of various systems.

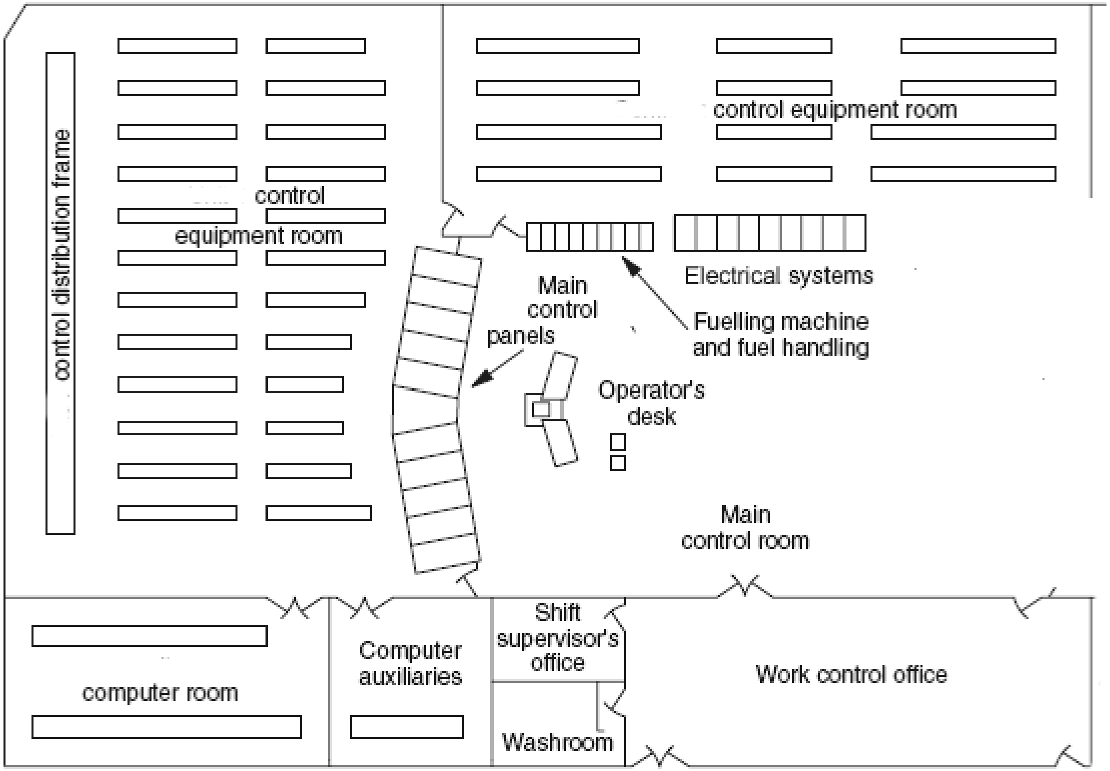

Figure 6 Layout of a CANDU 6 control room

Behind the operators' desk are a number of panels designed for stand-up operation. These panels are organized by system:

1. Containment
2. SDS #2
3. Emergency core cooling
4. SDS #1
5. Moderator
6. Reactor regulation
7. Primary heat-transport system
8. Steam generators
9. Turbine
10. Electrical panels
11. Fuel handling.

Figure 7 shows a view of the CANDU 6 control room, looking from behind the operator's desk. This figure shows a somewhat more modern version of the control room than that depicted in Figure 6. In the centre of the group of stand-up panels are two panels containing large video displays used to display the overall status of the plant. Above these displays are two large monitors which present annunciation information.

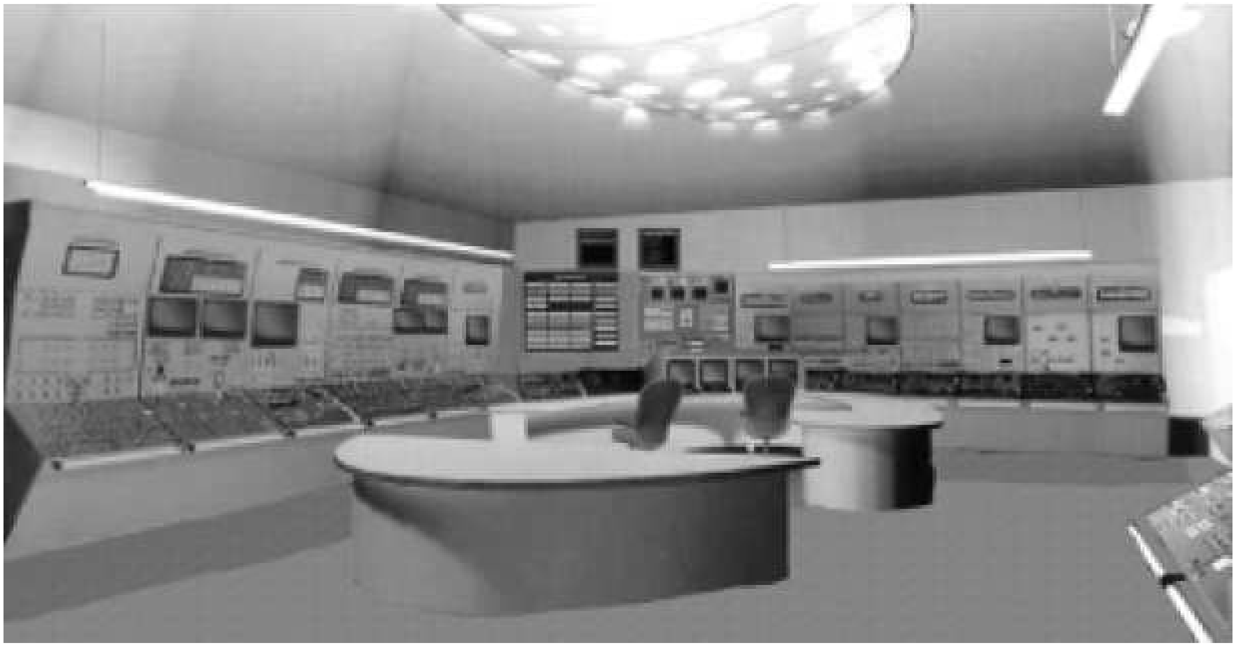

Figure 7 Pictorial view, CANDU 6 control room

At the top of each function panel is a matrix of alarm windows, some of which are driven by the DCCs and some of which are hard-wired to so-called “contact alarm units” in the conventional control logic. These are used to provide system-specific indications of alarm conditions and include essential alarms which will be needed should the control computers fail. The stand-up panels themselves contain a mixture of conventional operating controls (switches, lights, meters, and dedicated PID controllers) and VDUs and keyboards interfaced to the DCCs.

Virtually all CANDU stations have been retrofitted with safety-system monitoring computers to give the operator warning of diminishing margin to trip. These computers monitor buffered signals from the shutdown systems and have a dedicated VDU on the operator’s desk. This enables the operator to take steps to avert a reactor trip. For example, the introduction of new fuel can result in an unanticipated reduction in ROP trip margin during refuelling operations. The trip-monitoring computer alerts the operator, who can then manually reduce the RRS set-point if the margin becomes uncomfortably small.

The CANDU 6 control room uses a “dark panel” approach. When no lights are visible, things are normal.

Behind the control panels is the control equipment room, housing the physical I&C for Group 1 I&C equipment (see Figure 6). Separate equipment rooms are provided for the DCCs.

The detailed control logic for each system is subject to change in the early design stages of a new nuclear plant, and therefore a means has been devised to enable much of the required equipment to be installed before these details are available. Much of the logic equipment, as well as the sensors and actuators, can be installed and wired to the control distribution frame (CDF) before details of the interconnecting logic are known. The interconnection logic is implemented later by wiring discrete connecting wires at the CDF. The CDF is, in effect, a giant junction box about two metres high and several tens of metres long and is a major feature in the CANDU control-equipment room (see Figure 6).

Sensors are typically located close to the parameters being measured. This results in a large number of sensors being located within containment, although the electronics may still be separated from the sensor elements to facilitate maintenance. The electrical signals are taken through the containment wall by penetrations which are designed to preclude leaks from within containment along the cables.

Actuators are likewise close to the process equipment being controlled. They may use electrical power directly or use electrical/pneumatic converters in the case of loads driven by instrument air. Most of the electrical loads are driven by switchgear located in the turbine building.

6.1 Sensors

For the most part, the I&C systems use sensors that are used in other process control applications. However, because the subject of this text is a nuclear reactor, a description of the sensors used for measuring neutron flux is presented below, with particular emphasis on the in-core flux detectors, which are unique to the CANDU reactor. To ensure the necessary independence, the sensors used by each special safety system are dedicated to the system in question and are separate from the sensors, wiring, and associated equipment used by the process systems.

6.1.1 Neutron flux sensors

When starting up and operating a nuclear reactor, it is necessary to measure neutron power over a very wide range, of the order of 10 decades. During start-up, special neutron counters are used to monitor neutron flux as poison is removed from the moderator using ion exchange. After a prolonged shutdown, special fission chambers are also used temporarily. At approximately 10^{-7} x full power, ion chambers located outside the reactor core come on-scale. In the CANDU reactor, control is transferred to the DCCs at this point, and the neutron counters are removed. The ion-chamber readings provide a logarithmic reading of neutron power, which RRS uses for control up to around one decade below full power.

In an ion chamber, two concentric boron-coated cylinders capture thermal neutrons. The resulting alpha emissions cause ionization of hydrogen gas in the space between the cylinders. A high voltage is applied between the cylinders, which captures electrons caused by this ionization. The resulting current (~100 µA at full power) is amplified and provides signals proportional to bulk linear, logarithmic, and log-rate power which are usable over a wide range (10^{-7} to 1.5 x full power).

In the last decade, between 10% and 100% power, the thermal power in the fuel itself is the variable of interest. A prompt measurement of spatial power is required because any control or shutdown action must take place before local fuel damage can occur. Unlike most process variables, there is no single instrument that provides even a close approximation to the power in each channel of the CANDU core. Yet it is crucial to be able to determine this power because the thermal power in each channel determines the point at which dryout will start to occur. Thermal power is closely related to neutron flux, and therefore determining the latter provides a good estimate of the former.

Neutron flux within the core varies continually, even when the reactor is at constant power, due to the rather slow oscillations (in the order of hours) in the distribution of Xe^{135} throughout the core and to local flux peaks resulting from online refuelling. Although it is not possible to control flux down to the location of a few fuel bundles, it is feasible to control flux tilts between the 14 reactor zones. To do this, the flux in each of these zones must be measured locally.

In-core flux detectors are used in the CANDU 6 reactor to measure local flux power for both control and shutdown systems. They resemble coaxial cables with an Inconel sheath and a central core made of either platinum-clad Inconel or vanadium in the sensitive section, as shown in Figure 8.

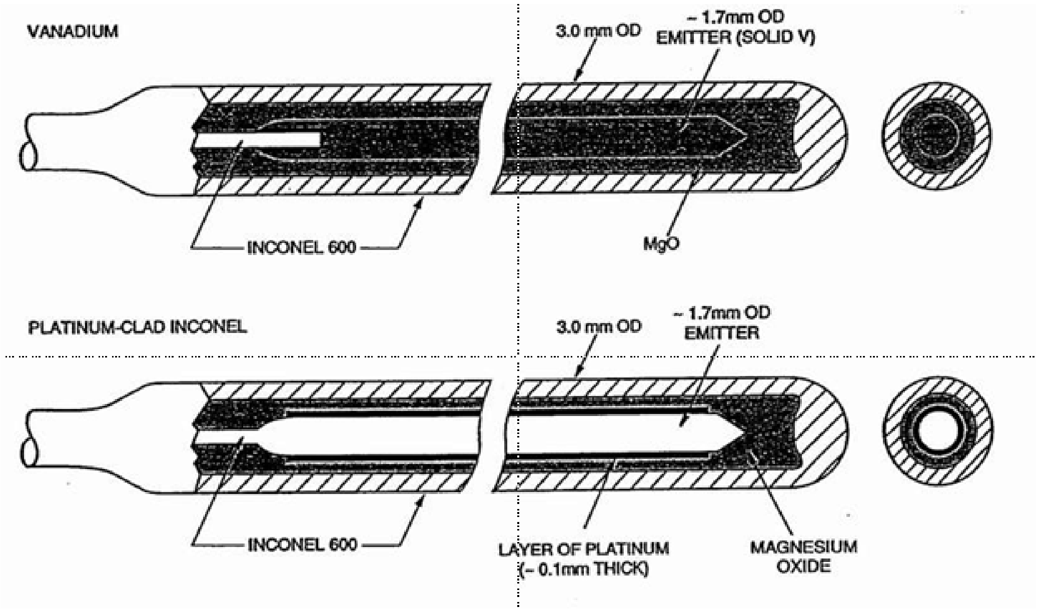

Figure 8 SIR detector geometry.

Vanadium detectors respond primarily to neutrons, but are too slow for either control or shutdown purposes (time constant τ = 335 seconds).

Platinum-clad Inconel detectors provide a signal which is less than fully prompt and varies with age. They respond to both neutrons and γ-rays (proportion approximately 50:50). Because under steady-state conditions γ-flux is proportional to neutron flux, the detector current may be regarded as being proportional to neutron flux. However, although the response of the detector itself to both neutrons and γ-rays is essentially prompt, 30% of the γ-rays which arise indirectly from β-active fission products are delayed, and therefore the overall detector response to a step change in neutron flux is basically about 84% prompt, with the remainder being delayed.

All detectors provide currents of a few microamps and are connected to amplifiers which convert this current to a proportional voltage before passing the signal to the DCCs or to shut-down-system logic. In the case of the SDS detectors, the amplifiers incorporate variable gains, which are adjusted daily to reflect the SDS ROP set-points, and electronic compensation to enable the output signal to compensate for the less than fully prompt detector response (see Section 7).

On Canadian CANDU 6 reactors, all flux detectors are now of the SIR type (straight individually replaceable), incorporating a helium atmosphere. The SIR design facilitates maintenance, and the helium atmosphere has been found to improve reliability.

Neither vanadium- nor platinum-clad detectors provide an absolute response to flux. Their response changes gradually with age and flux exposure in ways that are not yet totally predictable. However, in the short term (e.g., days), their response is repeatable, and therefore if they are calibrated against thermal power, they can provide values which are useful for estimating the thermal power generated in the fuel. This lack of an absolute power measurement contributes significantly to the complexity of both the ROP protection and power control algorithms. A detailed explanation of this is beyond the scope of this text.

The flux detectors used and their locations are as follows:

- 102 vanadium detectors are inserted vertically from the reactivity mechanism deck. These are used by the flux mapping program in the DCCs. The vanadium detectors are also the basis for the setback on high local neutron flux (see Section 3.4.1)
- 28 platinum-clad Inconel detectors (two for each of the 14 zones) are inserted vertically from the reactivity mechanism deck. In each zone pair, one is in channel A and one in channel C. These are used by the RRS program, the flux mapping program, and the stepback program in the DCCs.
- 34 platinum-clad Inconel detectors are inserted vertically from the reactivity mechanism deck. These are used by SDS1. The detectors are divided between channels D, E, and F.
- 24 platinum-clad Inconel detectors are inserted horizontally and are used by SDS2. The detectors are divided between channels G, H, and J. Future builds are expected to require the same number of detectors as for SDS1.

6.1.2 Other sensors

Most process analog signals are transmitted using 4-20 mA current loops (to help minimize the influence of electrical noise) from the point of measurement to the central control area. Temperatures are measured using resistance temperature detectors (RTDs) and to a limited extent thermocouples.

Given the large ground currents to be expected in generating stations, particular care is needed in sensor wiring to ensure that electrical noise and offset voltages do not corrupt these analog signals.

Contact sensing in the CANDU design uses 48VDC logic, although 24VDC is more common in modern process control equipment. This relatively high voltage has been found to minimize problems with degraded field contacts.

Although some I&C systems have been replaced at various sites with more modern digital equipment, there are still very few, if any, instances in CANDU plants of sensors which use digital communication right from the point of measurement. Because the individual stations make their own design decisions about equipment replacement, they would have to be consulted for accurate information in this regard.

6.2 Actuators

Actuators in CANDU are primarily electrically driven by on/off 48VDC outputs from the control system logic. Small loads may be driven directly, while heavy loads are driven by motor control centres which receive their inputs as 48VDC signals. Some larger valves are pneumatically actuated by electrical-to-air converters, and therefore in addition to the electrical supply system, there is an instrument air system with odd and even channels. Because this system is not described elsewhere, an overview of it is included in Section 6.2.1. Air-driven valves can typically move much more rapidly than electrically driven ones. For example, on CANDU, large air-driven valves are used to isolate containment. Air-driven valves typically use local air accumulators to provide a local power source following loss of the main air supply. This power can be used to drive air-driven actuators to a predetermined safe position following such an event.

There are a number of proportional actuators, such as the boiler feed-water control valves and the valves which control the flow of light water to the 14 zone controllers.

The CANDU design predated devices that are commanded directly by digital signals sent over a communications network, and therefore such devices will be found only in areas where the original equipment has been replaced.

Although most actuators used in the CANDU design are conventional, the actuators which control spatial power are unique. The reactor is divided for spatial control purposes into 14 zones (see Section 4.1). Each zone contains a cylindrical compartment with a variable amount of light water (see Figure 9). The level of light water in each compartment is controlled by varying the flow of water into it, using a valve whose position is determined by a 4-20 mA controller (normally an output from the DCC). The light water outflow is essentially constant.

The level of water in each compartment is measured by bubbling helium in at the bottom of the compartment and sensing the difference in pressure between the incoming gas and the pressure in the helium balance header connected to the tops of all 14 compartments (see Figure 7).

This system has the advantage of requiring no active components inside the reactor core. It works well as long as no compartment either floods or drains. However, flooding or draining does occur occasionally, and an alternative design which does not have this disadvantage might be worth investigating.

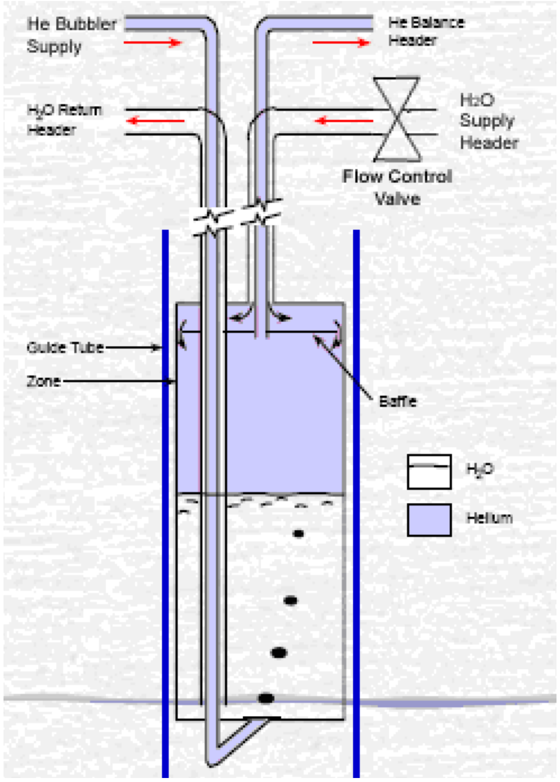

Figure 9 Liquid zone control compartment (simplified).

The absorber rods used for reactor control are very similar to the shut-off rods used on SDS1 (see Section 6.3), except that they are not spring-assisted and use hydraulic damping to limit their rate of movement, allowing them to be caught in mid-fall should a stepback to other than zero power be required,

The adjuster rods have no rapid movement capability. The winches that control their position are driven by electric motors connected to a variable-frequency supply.

6.2.1 Instrument Air System

Pneumatically operated valves are very common in the process control industry in general. These valves include on/off and continuous control valves and are typically commanded by an electrical signal. Compressed air for these valves is supplied by the Instrument Air System, a service system which, in a manner somewhat analogous to the electrical system described in Chapter 11, distributes compressed air to the numerous end-user devices throughout the plant. Like the electrical system, the instrument air system provides channelized sources and distribution paths to ensure that the appropriate degree of redundancy and separation is available. The compressed air provides a powerful force multiplier, and operation of the end devices also tends to be relatively fast. The requirements for both instrument air and electrical power systems are covered in a common document [CSA2011b].

This description will concentrate on the reactor-building instrument air system because it is of greatest interest in terms of nuclear safety and reliability. Within the reactor building, there are about 130 on/off valves and around 70 continuous-action control valves powered by instrument air. Prominent among the former are the valves which initiate liquid poison injection for shutdown system #2, the dousing valves, the containment isolation dampers, and the moderator temperature-control valves, all of which play significant roles in ensuring plant safety. The reactor-building instrument air system is seismically qualified.

Refer to the greatly simplified schematic in Figure 10.

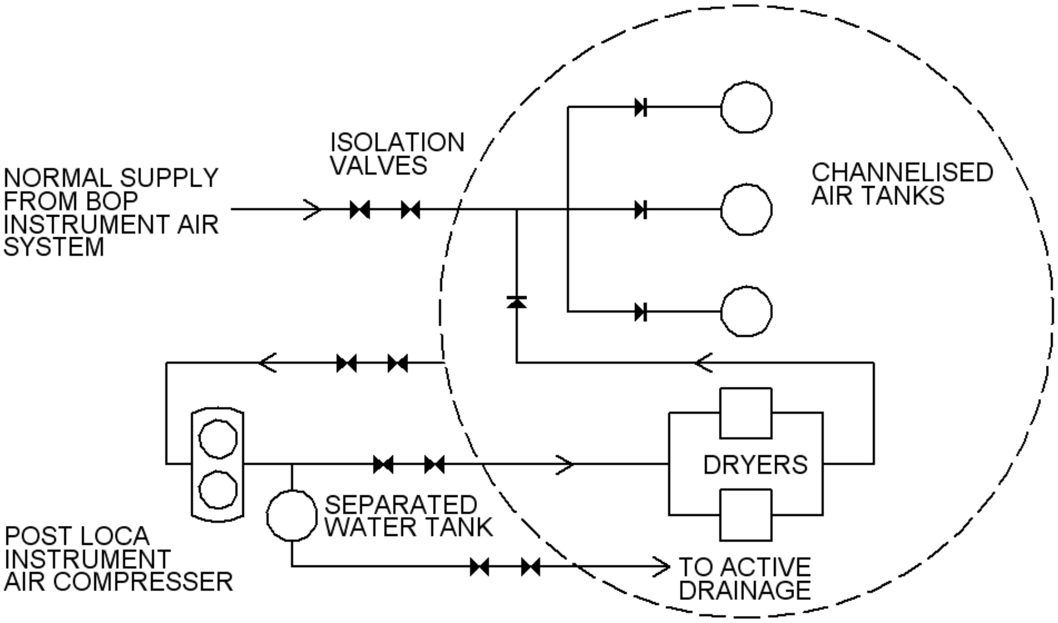

Figure 10 Reactor building instrument air system (simplified).

During normal plant operation, the reactor-building instrument air system receives its supply of compressed air from the same bank of compressors that provide instrument air for the remainder of the plant. These are located in the turbine hall and run off Class 3 electrical power, so that air supply is subject to interruption for a five-minute period in the event of a loss of Class 4 power. The air passes into the reactor building through a single line equipped with two isolation valves in series, which are operable from the control room and provide containment isolation, typically following a LOCA. A number of the pneumatic actuators inside the reactor building are required to operate in the immediate post-LOCA period and must be able to operate during the five-minute interruption caused by loss of Class 4 power. Therefore, a set of three air-storage tanks is provided inside the reactor building to provide the necessary compressed-air reservoir. Note that many of the loads are on/off and have to execute only a single movement in this time frame. The single supply line branches within the reactor building to feed these tanks, each feed being equipped with its own dual isolation valves and a check valve to preclude loss of the tanks' contents in the event of a loss of supply air. Therefore, the three tanks provide a short-term source of triplicated, uninterruptable instrument air within the reactor building. The air, once it has been used, is vented to the reactor-building atmosphere and will eventually be exhausted through the building ventilation system.

Following a LOCA, containment ventilation will be isolated, and continued use of the normal instrument-air source described above would tend to pressurize the reactor building gradually. To address this concern and permit continued long-term use of pneumatic actuators in containment following a LOCA, a separate post-LOCA instrument-air source is provided on recent CANDU 6 plants. This post-LOCA instrument-air (PLIA) system consists of a compressor located outside the reactor building and two redundant dryers located inside the building, together with a number of isolating valves. This equipment is not seismically qualified. The PLIA system feeds the same reactor-building instrument-air tanks described above through appropriate isolating valves. The compressor draws its inlet air supply from within the reactor building, thus avoiding the tendency to pressurize the reactor building. The PLIA compressor will be brought on-line before the valves isolating the supply of normal instrument air to the reactor building are closed. The air handled by the PLIA system is potentially contaminated, and therefore the system components outside containment are located in isolated rooms. Water removed from the flow, which collects in a tank downstream of the compressor, is returned to the reactor building active drainage through an isolated line.

6.3 Actuators for the Special Safety Systems

In the case of SDS #1, 28 cadmium shutoff rods are used, which drop into the reactor core under gravity, assisted initially by springs. They are held out of the core by electromagnetic clutches which are energized by two-out-of-three logic driven by all three channels, as shown in Figure 11. When any two of the three channels vote for a reactor shutdown, the clutches will be de-energized, and a reactor shutdown will take place.

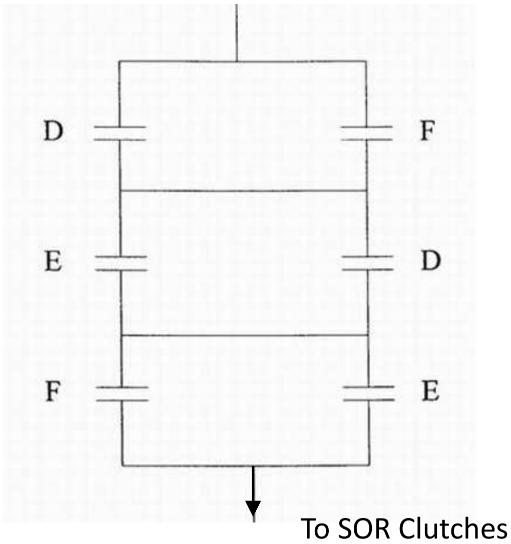

Figure 11 Two-thirds voting logic for SDS1 trip

In the case of SDS2, the contents of six tanks containing gadolinium nitrate/D_2O solution are injected into the moderator using pressurized helium as the motive source. The pneumatically operated valves which connect this pressurized gas to the tanks containing the poison are arranged in a two-out-of-three configuration which is logically similar to that of the electrical contacts shown in Figure 11. Again, a trip in any two of the three SDS2 channels will open the corresponding helium supply valves and trip the reactor.

6.4 Control Logic Technology in the Process I&C Systems

The CANDU design predated the widespread use of digital technology in I&C by perhaps a decade. Nevertheless, the design relies very heavily on digital control technology, which will be described below.

Although the digital control computers (DCCs) are arguably the heart of the CANDU I&C system, they are augmented by a mass of conventional control equipment. For binary control, 48VDC telephone relays were used in the early designs because these were readily available and highly reliable. By the time the detailed design of the CANDU 6 was undertaken, the telephone industry had moved on to solid-state technology, and the telephone relays were replaced by relays designed specifically for the process control industry. By the time Darlington was built, programmable logic controllers (PLCs) were becoming available, and that station used a custom-built PLC to implement much of the Boolean logic.

Much of the closed-loop analog control of the CANDU subsystems is accomplished by general-purpose analog loop controllers for which the parameters are set up as required for the individual application when the plant is commissioned. Because these applications are fairly conventional, they will not be described further here.

6.4.1 Digital I&C technologies in CANDU I&C process systems

The computers used to control the CANDU plant do not use off-the-shelf products designed specifically for digital process control, as is common practice today. At the time of the first CANDU designs, in the mid-1960s to mid-1970s, the number of successful digital control applications world-wide could be counted in no more than double digits, and the available technology was a general-purpose minicomputer with custom-designed process input/output hardware. “Mini” was a relative term. A minicomputer was the size of a small refrigerator and with the associated input/output equipment could well expand to be 10 metres long. These computers were not inexpensive, nor were they noted for their reliability. One did not procure computers in large numbers and deploy them where they were required. Rather, a pair of centrally located computers configured as master and standby was used, and the necessary sensors and actuators were connected to them using discrete pairs of wires (as opposed to communications networks). Networking of computers was unknown. Once the decision was made to use computers, one looked at what other functions, in addition to those that had to be there, could be implemented using this expensive but very flexible resource.

The digital computers on the CANDU 6 implement the following functions:

- Closed-loop control for:
  - the reactor,
  - steam generator level
  - steam generator pressure,
  - turbine loading,
  - heat-transport pressure and inventory, and
  - moderator temperature;
- Control of the fuelling machines, which enable on-power refuelling of the reactor;
- Run-up of the turbine;
- Presentation to the operator of annunciation information for all process systems in the plant;
- Data logging for all process systems in the plant.

A detailed description of the logic implemented in the DCCs is beyond the scope of this text, but as an example, see the description of Overall Plant Control given in Section 4. The architecture of the dual-DCC system is shown in Figure 12.

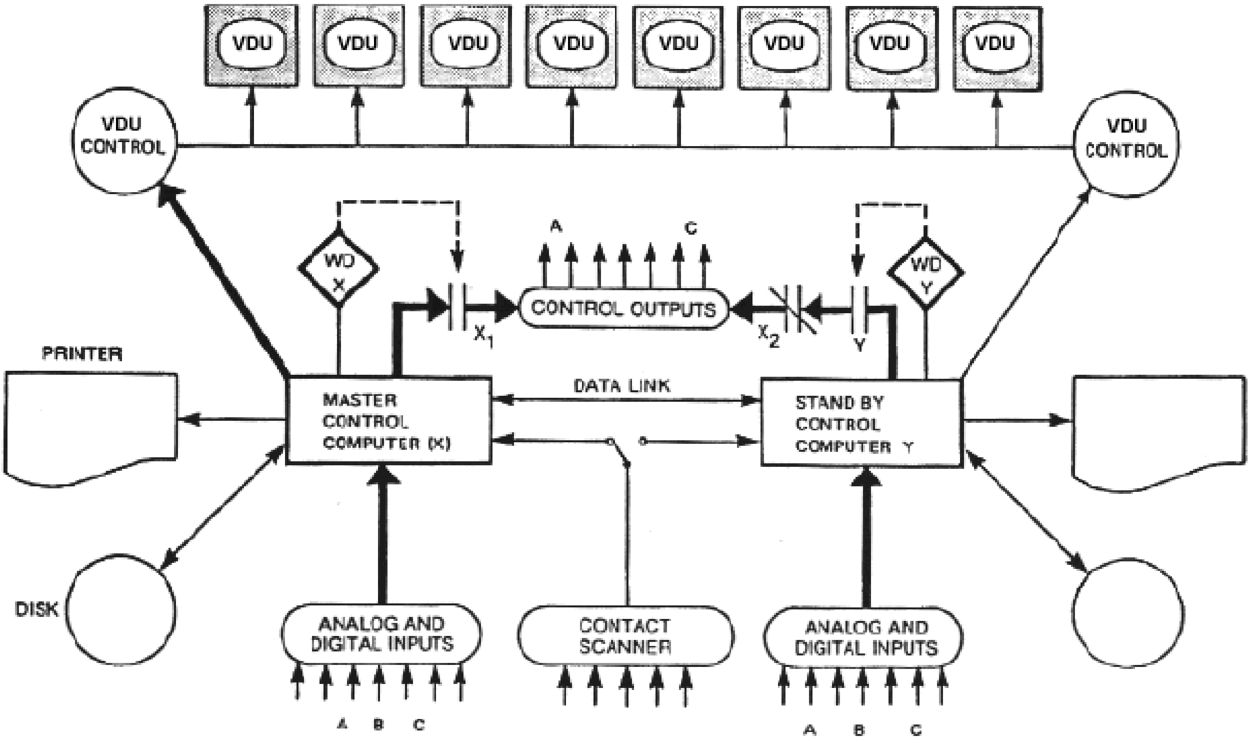

Figure 12 CANDU 6 DCC architecture

All the software was written in assembly language. Although this ensured that the software was as compact and as fast as possible, the result was a program that relied totally on the programmer for its integrity and which was not directly reviewable by other than software cognoscenti.

6.4.1.1 Annunciation and status display

In addition to the inputs needed by the software which controls the plant, several hundred analog and contact inputs are connected to the DCCs to enable alarm information to be presented to the operator and to provide him with status information on the various systems, using the VDU controllers and printers shown in Figure 12.

Alarms are generated within the control logic when abnormal conditions are detected. These conditions are detected either by hardware logic in the control equipment or by tests built into the DCC control software. The alarms are presented to the operator using a combination of illuminated “windows” above the control panels and DCC-driven VDUs in the central area of the stand-up panel. In addition to the alarms generated by the control functions resident in the DCCs, the computers monitor thousands of contact and analog signals, enabling their alarms to be displayed on the VDUs and recorded on a printed log in chronological sequence.

The CANDU alarm annunciation system works well for routine operation and for minor plant disturbances. One can imagine, however, that during major plant disturbances (e.g., loss of Class IV power), the operator is faced with a sudden flood of alarm messages, flashing windows, and horn tones. When Class IV power is lost, many subsystems will no longer be available and will initiate an alarm in consequence, but all that the operator really needs to know is the major event: loss of Class IV power. To mitigate this situation, the CANDU alarm-annunciation scheme uses some very basic prioritization logic. Each alarm is designated as “major” or “minor”. When certain predefined plant upset conditions are in progress, minor alarms are temporarily suppressed on the VDUs. Although this does help, there is no doubt that a more intelligent alarm presentation system would be useful. Such schemes have been contemplated in the past and will probably be incorporated in any new CANDU build, but progress has been slow because the conditions under which it is appropriate to suppress any given alarm typically require very careful examination of the operation of the systems involved, meaning that the construction of the database behind such a system would be very labour-intensive. This would appear to be an ideal application for some kind of expert system. However, there is great reluctance to implement any system whose response to a given plant upset is not completely deterministic. The approach to such intelligent annunciation systems has always been very conservative. One can imagine the repercussions if it were determined in some post-event investigation that suppression of a particular alarm had resulted in incorrect operator action.

6.4.1.2 Fuelling-machine control

The on-line refuelling function is unique to the CANDU design. On the CANDU 6 stations, fuelling-machine control is implemented in the DCCs. On many of the Ontario 4 unit stations, this function has separate, dedicated computer controls provided by the fuelling-machine vendor, but again, minicomputer technology is used. Unlike the closed-loop process-system control discussed earlier, control of the fuelling machines is an example of sequential control. The computer carries out a step in a sequence and then checks that the expected feedback indications have been returned from the plant before moving on to the next step. Each step consists of a single predefined action and a set of post-action checks. A job is made up of a series of steps. Jobs may be executed automatically from start to finish or may progress one step at a time under operator supervision. In any event, if the expected feedbacks do not materialize, execution stops, and the operator is alerted.

The sequential actions required to complete each fuelling activity (i.e., a “job”) are defined in an interpretive language developed specifically for the CANDU fuelling application. Although not as readable as a modern sequential control language, this language does provide the reader with some level of isolation from the underlying assembly-language logic. The operator can allow the job to proceed entirely automatically or can elect to require his permission between each step and the next.

Because fuelling is not required on a continuous basis, the fuelling-machine function is only implemented in a single DCC.

6.4.2 Historical perspective

As an indication of how the technology has changed in 40 years, the CANDU 6 DCCs had 64K bytes of memory and a one-megabyte disc drive. The central processor cycle time was around one microsecond. To maximize performance, the machines were programmed in assembly language. Every byte of software was custom developed. A previously developed operating system was not used, for example. The result is a system which, although very flexible, requires outdated skills to maintain and is also particularly prone to errors in software maintenance. However, all attempts to migrate to more modern equipment have been rejected because of the financial consequences of any delays or errors in implementation of a replacement system. Although aging DCC equipment has been replaced, changes have been made at the individual hardware module level, with custom replacements being engineered to match the fit, form, and function of the old equipment, and particularly to have minimal impact on the existing software. The DCC equipment does not support the standard techniques which have evolved over the years for the development of highly reliable computer systems. Standards such as [CANDU1999] would be applied to the systems and software engineering of any new-build design. The replacement of the DCC platform is undoubtedly one of the greatest challenges facing the designers of a new-build CANDU.

6.5 Logic Technology in the Special Safety Systems

On all CANDUs before the CANDU 6, the algorithms for making the trip decision were implemented entirely in hardware (primarily by analog comparators). This tends to mean that the trip set-points for each parameter are constant values because implementation of set-points which are themselves functions of another plant variable becomes rather complex.

6.5.1 Digital I&C technologies in CANDU 6 shutdown systems

With the need to have two diverse trip parameters for each postulated accident, it became difficult to identify suitable parameters with fixed set-points without imposing a large penalty on reactor power or incurring an increased frequency of unnecessary reactor trips. The decision was made to use set-points for some process trip parameters that were functions of other plant variables, notably reactor power. Means of generating these set-points using analog electronic hardware were examined, but ultimately the decision was made to implement the trip logic for process trips (i.e., trips not based on reactor power measurement) using digital computers.

They were not, however, called computers, but were referred to as “programmable digital comparators” (PDCs).

6.5.2 Digital I&C technologies in Darlington shutdown systems

Given the success of the PDCs on the CANDU 6, the decision was made on Darlington to go further. Not only would both nucleonic and process trips be implemented in computers, but separate computers would be added to support automated testing of the shutdown systems. Darlington would be the first reactor in the world to feature completely digital shutdown-system logic. In recognition of this, the term “PDC” was dropped. At Darlington, the equipment was called “trip computers”. It was recognized that additional steps would have to be taken to ensure high quality in the software engineering which implemented the trip logic. What was not recognized was just how onerous this process would become.

In the ten years that elapsed between CANDU 6 and Darlington, the software industry had matured immensely, and computers were being used to an increasing degree in applications involving public safety. The experience had not been entirely positive. Lives had been lost as a result of software errors. In fact, one of the landmark case studies cited in software engineering texts involved a radiation-therapy machine engineered by AECL which resulted in several deaths before the software design errors were identified and corrected. Interested readers can Google the “Therac-25”.

When the design of the Darlington shutdown-system software was undertaken, it was recognized that there was a problem, but the appropriate systems engineering standards had not yet been written that would establish acceptable ways to address the problem. The Darlington shutdown systems became a stepping stone in the development of these standards, but the program encountered serious delays and cost overruns. A product of this effort was a series of standards written by AECL and Ontario Hydro governing the engineering of software for safety-related applications at various levels of criticality [CANDU1999].

The international safety-related systems industry has since introduced a standard which is applicable not only to software development, but also to systems engineering of safety-related systems using digital technology in general [IEC2001a]. A derivative of this standard applicable specifically to the nuclear industry also exists [IEC2001b], but it has yet to be broadly adopted by the Canadian nuclear industry.

Despite the early experience at Darlington, the end result is a system that is well regarded by operating and regulatory staff.

7 Regional Overpower (ROP) Trip Logic

Description of the ROP logic has been deferred to this point to enable the characteristics of the sensors involved to be discussed before describing how these are accommodated in the design. The following discussion is applicable to both SDS1 and SDS2.

The CANDU 6 ROP systems currently have 34 and 24 flux detectors in shutdown systems 1 and 2 respectively. Hence, there are only 8 to 12 detectors per channel, yet each channel is required to protect against fuel dryout in any one of the 380 channels in the core. As discussed in Section 6.1.1, the detector signals represent only point values in the distribution of flux throughout the core. Moreover, the outputs decline with age and are not completely prompt.

The trip comparator associated with each ROP detector is set to a value determined offline by a computer code "ROVER" to ensure that at least two detectors in each SDS will call for a trip before dryout occurs at any location in the core. The computer code considers a large number of normal and abnormal reactor configurations, but does not account for non-equilibrium fuelling. Typically, with new pressure tubes, this comparator set-point will be around 122%, although it declines as the pressure tubes age. The code is re-run and the comparators are adjusted infrequently (e.g., once every three years) to account for changes that affect thermal hydraulics, such as pressure-tube creep. The comparator setting includes the desired margin to trip.

The RFSP (Reactor Fuelling Simulation Program) code is run frequently (e.g., once every three days). It synthesizes a flux map based on the reactor's power and fuelling history, on reactivity mechanism positions, and on the 102 vanadium detector readings. The power in each channel *j* in the channel-power peaking factor region (all but the outer extremity of the core), CP_j, is then computed and divided by the reference power for that channel, $CP_{ref\ j}$, to yield the normalized channel power:

$$P_{norm\ j} = CP_j/CP_{ref\ j}.$$

The highest $P_{norm\ j}$ becomes the channel-power peaking factor (CPPF, typically around 1.08). Therefore, the CPPF represents the worst-case power peaking factor due to refuelling ripple over all channels. The detector amplifier gains are then adjusted so that they read (current reactor power x CPPF). It can be seen that this approach is quite conservative because it assumes that all detectors will be affected equally by the channel at greatest risk. This purpose of this calibration is to compensate for changes in position of the reactivity mechanisms and for local flux changes due to fuelling. It also incidentally compensates for decreasing detector sensitivity with age. Typically, one or two detectors per channel will be found to be out of limits in any particular adjustment cycle.

Each detector amplifier also includes two filters which compensate for the non-prompt response of the detector itself. These are not normally adjusted. The compensated signals passed to the comparators represent a fairly close approximation to the thermal power in the fuel, which is what causes dryout, and hence is the signal on which the reactor needs to be tripped. The effect of this compensation is shown in Figure 13.

The set-points for both SDS1 and SDS2 are based on the same criteria, including trip margins, and therefore it is a matter of chance which system will trip first. Indeed, in the case of a slow LOR, it is likely that SDS2 will trip first because it has fewer detectors and hence its set-points must be somewhat more conservative. In the past, there was an incentive to arrange things so that SDS1 would trip first because it could be re-poised much more quickly and potentially avoid a reactor poison-out. Nowadays, the analysis associated with restarting the reactor after any trip is sufficiently time-consuming that a poison-out is inevitable in either case.

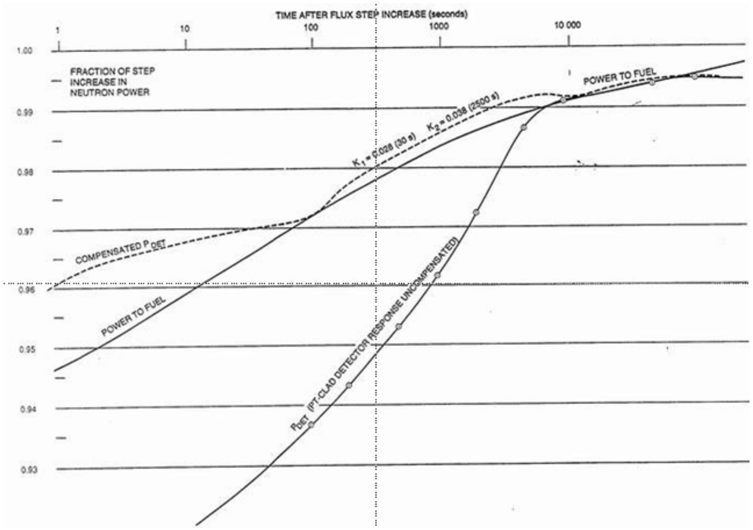

Figure 13 Effect of compensation on Pt-clad detector response

Clearly, there are many uncertainties involved in determining the trip set-points, both in the comparator settings determined by ROVER and in the calibration of the detectors to account for CPPF. Typical uncertainties are:

- simulation uncertainties
- detector uncertainties
- uncertainties specific to a given SDS channel
- common uncertainties
- CCP uncertainties
- refuelling uncertainties
- uncertainties due to the effects of moderator poison.

These uncertainties are combined statistically and incorporated into the comparator settings, the intent being to ensure that at least one detector in each SDS channel will trip in a given shutdown system before dryout occurs, 98% of the time,. As noted in Section 7, two channels must trip before a reactor trip occurs. The CANDU 6 and Ontario reactors use different methods of justifying compliance with this 98% criterion, but achieve similar results.

Uncertainty in any trip set-point results in its being assigned a conservative value. In the case of ROP set-points, this results directly in a loss of plant capacity, and therefore any improvement in understanding the uncertainties is likely to result in higher trip set-points, which will have an immediate economic payoff. Even a 1% increase, for example, results in an increased power capacity of 60,000 MWh per year for a 600-MW station. Depending on the going rate for power, that is worth several millions of dollars every year.

8 Design Verification

All CANDU systems were subject during the design process to a design review carried out by their designers' peers. This review was based on the detailed design documentation produced for the customer. With each succeeding design generation, the documentation became more comprehensive. However, the process-systems review was never based on the detailed comprehensive list of system requirements which is typical of best-practice systems design in a modern engineering context.

In the case of the special safety systems, the situation is better because it was necessary to demonstrate detailed compliance with the safety requirements.

For closed-loop control systems, the dynamic performance of the control algorithms is of major interest because any instability can require close operator attention at best and operator/automatic action up to and including shutting down the plant in the worst case. The design of the CANDU controls was carried out before computers that could simulate plant dynamics in real time were available. The performance of the proposed algorithms was checked using plant models and models of the control algorithms run on a mainframe computer in batch mode. For readers who never knew the days when telephones had wires, a mainframe computer occupied a large room buried in the bowels of the design office, consumed many kilowatts of power, was attended by a staff of dozens of technicians, and executed jobs fed to it in the form of decks consisting of thousands of punched cards. You submitted your job and waited a day or two for the results to be returned. These took the form of endless columns of numbers printed by a line printer on reams of fanfold paper.

The plant models used for such verification were not themselves subject to any formal verification.

In the case of the special safety systems, dynamic performance was of critical interest because any unexpected delays in initiating a trip could lead to fuel damage. In this case, the codes used for verification, such as Cathena, have been subjected to formal validation.

Over the life of the CANDU 6 plants, all sites have added operator training simulators which closely mimic the performance of the subject plant in real time. Although these models have not been formally validated to serve as engineering simulations (i.e., as accurate representations of the process dynamics), it has become common practice to verify any significant control system changes on the operator training simulator before implementation. This not only provides some assurance that the changes to the control algorithms will work, but checks any associated changes to the human/machine interface. The training simulator has been particularly effective in enhancing the level of confidence in changes to the DCC software because the simulators typically emulate the DCC processor, which means that the new DCC software can be loaded and run on the simulator. It should be noted, however, that due to possible lack of fidelity in the plant models used in training simulators, successful operation of the changes on the simulator does not offer a guarantee that the change will be trouble-free on the plant itself.

In the past, no new-build CANDU plant has had a training simulator installed when it was originally commissioned. In today's regulatory environment, availability of an operator training simulator before commissioning has become a requirement. This will lead to a situation where it becomes very important to be able to adapt the operator training simulator rapidly to reflect changes made to the plant design before and during commissioning.

9 Bringing the CANDU I&C Design into the 21st Century

The most advanced CANDU design to be built up to this time is Darlington, which went into service in 1990. Since that time, the instrumentation and control industry has seen a revolution. Key to this change has been the enormous progress in digital electronics and communications. Distributed control system technology is now predominant in I&C system implementation, and standard computer workstations bring almost unlimited computer power to the information management functions supporting plant operation.

During the intervening years, at least three new CANDU designs have been embarked upon, yet none has been developed to the point of implementation. These were the CANDU 3, CANDU 9, and the Advanced CANDU Reactor (ACR). All these featured I&C technology which was current at the time. One might think that, given the advances in computing power, the design of the control systems for these plants would present few problems. Yet this was far from the case.

In a modern distributed control system (DCS), use of which is standard practice in the I&C industry today, the physical inputs are typically digitized close to the field instrumentation. The digital data are then carried over a communications network to one of a number of processing units in which the control algorithms are executed. The control commands are then networked to output stations located near the controlled actuators. The algorithms in the input, processing, and output modules typically run asynchronously, and inter-processor communication is also typically asynchronous. Therefore, the time delays from input to output are no longer deterministic, as is the case for a DCC. If the resulting uncertainties are to be negligible, everything has to run many times faster in a modern system if the same dynamic performance is to be guaranteed in the worst case.

Modern DCS systems have the advantage of being programmed in application-specific high-level languages, frequently using graphical representations of the logic, which make the design much more susceptible to review by outsiders. However, many vendors' products use proprietary languages which are not readily portable to a different vendor's product, should it become necessary to replace the control system at some stage in the life of the plant. An international standard does exist for programmable controller languages [IEC1993], but adoption of this standard has not yet been widespread.

Periodically, the prospect of using some of the more exotic modern I&C technologies such as expert systems, optimal control, neural networks, and fuzzy logic controllers is raised, typically by academics. Many of these techniques have been around since the 1960s, after all. The problem with many of these approaches is that they are typically non-deterministic and therefore difficult to qualify. The regulators quite rightly frown on any design that cannot be guaranteed to produce the exact same response every time it is presented with the same sequence of input conditions.

The old adage, "If it ain't broke, don't fix it", is very good advice. Any change to the existing design must show that it can provide enhanced power output, greater plant availability, and/or lower operating costs if it is to be seriously considered. In this chapter, the author has attempted to point out some features of the existing design that may benefit from re-design in any future CANDU new build.

9.1 Technological Obsolescence

Although Darlington is the last clean-sheet CANDU design developed, there have been several new builds of the CANDU 6, notably in Wolsong, Qinshan, and Cernavoda. Moreover, a number of existing plants have faced the need to replace control equipment as its reliability declined due to age and lack of spare parts.

Perhaps the highest-profile equipment items involved which were not susceptible to plug-in replacement were the shutdown-system computers for the new-build CANDU 6 plants and the DCCs.

The PDCs were replaced with more up-to-date distributed control system hardware. However, due to the relatively small amount of logic involved and the fact that the PDCs are not geographically distributed, the new design was able to avoid reliance on digital communications. Architecturally, then, the new PDCs were little different from the systems they replaced. The functions of the safety-critical software running on the existing equipment were re-engineered for the new platforms. Although this re-engineering was by no means a trivial task, the existence of well-documented requirements and accepted procedures [CANDU1999] for carrying out this work made it fairly routine.

The replacement approach for the DCC hardware took a very different approach. The obvious approach to DCC replacement is to use a modern DCS, as discussed in the previous section. However, DCC replacement has the potential to lead to a very protracted and costly plant outage if anything goes wrong during the replacement or if the software implementation proves to have defects once the plant is brought back online. To minimize this risk, the fundamental requirement of the plant owners has been to replicate the existing control logic in precise detail.

The existing DCCs contain a large body of purpose-designed software, ranging from the operating system to dozens of applications whose details are specific to each CANDU site. The most reliable definition of the functionality of this software is the program listing itself, which is written in a machine-specific assembly language.

As stated in the introduction, the design and implementation of the analog control loops implemented in the DCCs are based on classical frequency-domain techniques. At a detail level, the control logic consists of an interconnected set of gain elements, integrators, and differentiators. However, because these elements are ultimately implemented in assembly language and are specified in the form of difference equations, the basic structure of the control logic is far from obvious to anybody reviewing the detailed design. Translating the existing control logic into the context of a modern DCS platform with a high degree of confidence that the platform change will not introduce latent errors would be a formidable task. Not only is there a risk of errors in the interpretation of the software functionality, but the different architectures of the two platforms make emulation of real-time performance very difficult to guarantee.

Therefore, if a DCS were used to replace these DCCs, the DCC logic would have to be reverse-engineered and a new DCS-based design developed and re-validated from the ground up. No plant owner or CANDU 6 replication project manager has accepted this risk. The alternative was to custom-build replacement computer equipment which would be fit-, form-, and function-compatible with the existing hardware and which would host the existing software with relatively minor changes.

As it happened, the CANDU industry was not the only one to be confronted with this problem, and replacement computer hardware based on more modern components had been developed for both the Pickering A and CANDU 6 DCCs. These “clone” computers served as a starting point for the refurbishment of these units, a process which is still ongoing.

All this means that all existing CANDUs will run with aging I&C technology for the next couple of decades and that the first CANDU new build, should there be one, will be the proving ground for the design, licensing, and implementation of what is now standard practice for the I&C industry, assuming that the “clone” DCC is not retained.

The design life of a CANDU plant is around 40 years. With refurbishment, this could well extend to 70 years or more. Each generation of I&C technology becomes obsolete in about 15 years, particularly if one is speaking about digital equipment. The manufacturers of the equipment typically will not guarantee to support their product beyond about 20 years, though if one happens to pick a particularly successful product line, this date might be extended by another 10 years. Therefore, any nuclear plant will have to face replacement of its digital I&C equipment at least twice in the lifetime of the plant. In the case of the original CANDU DCCs, the original equipment was specifically developed in Canada by a custom systems integration supplier who owned much of the design data and was able to purchase the rest from the original equipment manufacturers. Manufacturers of off-the-shelf I&C products are not interested in this kind of re-engineering business. However, by far the most economic and reliable solution in the future will be to base the I&C on off-the-shelf equipment. Design for replacement has not so far been a requirement on next-generation reactors. This is a challenge that should be accepted. The existing refurbished CANDU plants will probably have to face this issue with their DCCs at least once more during their extended lifetimes because it is hard to imagine the present vendor still providing support in 2040. In future replacements, it is unlikely that development of a replacement platform which can run the existing software will be practical, and therefore the issue of rewriting the software when the equipment is replaced will have to be faced. If a hardware platform which supports the industry-standard languages [IEC1993]is used, this offers some hope that the existing software will be portable. However, the industry in general has shown little inclination to adopt these languages, meaning that there is no guarantee that this approach will yield the desired advantages.

10 Summary of Relationship to Other Chapters

By this point, the reader will be aware that the architectural design of the control systems is heavily dependent on the safety requirements that these systems implement, and on the necessity to avoid unnecessary interruptions of the plant’s *raison d’être* – the production of electrical power. Chapter 13 provides an excellent description of the safety requirements, and should be read in conjunction with the material presented herein.

Chapter 5 provides a discussion of the reactor physics, including the control actuators which the RRS and shutdown systems use to control reactor flux.

I&C systems are only as good as the motive power sources that drive them—in the case of CANDU, primarily Class I and II electrical power and instrument air. The design of the electrical power distribution system is presented in Chapter 11. A brief description of the instrument air system is included in Section 6.2.1 of the present chapter.

11 Problems

1. Starting with an equilibrium reaction, assume that a control rod is moved out of core by a fixed distance. What form will the curve of reactor power vs. time assume?

2. Given the approximate numbers quoted in Section 2 and assuming that the light-water zone controllers are initially half full, in the absence of any other control or shutdown action, what value would the power in a CANDU reactor reach in one second if all the light-water zone controllers suddenly drained? Assume that initial power = 100% and that the total reactivity worth of the zone controllers = 7 mk.

3. From the point of view of reliability, why is it important in a reactor with two shutdown systems that the two systems be independent?

4. In the CANDU design, what is the probability of occurrence of an uncontrolled increase in reactor power which is not stopped by at least one of the shutdown systems?

5. What is the cost, in lost power production, of an outage leading to a 40-hour poison-out of a CANDU reactor? Assume that the electrical output = 700 MW and that the utility is paid 3.5 cents/ kW hour.

6. The two-out-of-four architecture used on many PWRs has one major advantage over two-out-of-three logic. What is it?

7. Some of the factors that can cause a shutdown system to fail to perform its design function are:

   1. Earthquakes
   2. Flooding
   3. Localized physical damage to the system
   4. Electromagnetic effects
   5. Design/analysis shortcomings.

   Steps which are taken to minimize the impact of these factors on the system as a whole include:

   a) Multiple equipment groups (wide physical separation)
   b) Channelization (physical and electrical separation)
   c) Equipment/technology diversity
   d) Equipment qualification
   e) Multiple activation parameters
   f) Fail-safe design (i.e., the safe state is the de-energized state).

   Prepare a table assigning one or more of these lettered steps to the numbered factors above.

8. Prepare a succinct statement which captures the key differences between a setback and a stepback in the CANDU I&C design.

9. Self-check is not listed as one of the stepback parameters. So why is a self-check sometimes referred to as a "seismically qualified stepback"?

10. If it is assumed that every dual DCC failure leads to a poison-out, that DCC failure is required to lead to a poison-out not more than once in three years, and that, on average, it takes 20 minutes to get a DCC back in operation following a failure, what is the design target for mean time between failures of a single DCC?

11. Name a factor that could cause an MTBF calculation based on hardware failure rates alone to be very optimistic.

12. The logic to avoid flooding of the light water zones precludes filling them beyond about 90%. What action will RRS take if the zones fill, but a positive power error persists?

13. What sensors drive the logic providing primary protection against slow LORs in a CANDU? What is the rationale for not providing diverse logic to protect against this event?

14. What is a key difference between the ECCS and containment systems with respect to continued availability of electrical power?

15. What would be some key requirements for the secondary control area?
16. It has been suggested that the control system be made capable of initiating a power setback on low margin to trip. What possible objection could there be to such a proposal?
17. What are the potential effects of a failure of the air-conditioning system which services the room housing the DCCs?
18. Why might the CDF be less relevant in the context of modern I&C technology?
19. List some pros and cons of digitizing instrument readings inside containment.
20. When using thermal measurements to calibrate flux-detector readings, what processing will have to be applied to the flux-detector readings before they are compared to thermal power measurements?
21. Why are ion chambers not used for control in the upper decade (10% to 100%) of reactor power?
22. Why do you think digital sensors have seen limited application in retrofit projects on existing CANDU power plants?
23. What would be the most demanding requirement facing a potential replacement for the level measurement in the light-water zone compartments? Suggest some possible technologies which might be used.
24. What long-term factors might require the operator to adjust bulk reactivity using moderator poison?
25. What characteristics of the CANDU design drove the choice to use digital computer control before it was a widely accepted technology?
26. What are some negative consequences of having the control logic defined in software written in assembly language?
27. What would be a major reason that an annunciation scheme based on expert system technology might be difficult to license?
28. Give an example of a non-proprietary language which would be appropriate to the definition of sequential control logic.
29. List some advantages to making shutdown-system testing an automatic function.
30. List two reasons why software is often considered to be particularly prone to design errors.
31. Which ROP adjustment accounts for the changing shape of flux within the reactor core?
32. Why will the ROP trip comparator set-points need to be adjusted downwards as the pressure tubes age?
33. Give an example of a non-proprietary language which would be appropriate to the definition of (a) analog and (b) Boolean control logic.
34. In what way does the architecture of a distributed control system, as opposed to that of a centralized computer system, affect the execution timing of the control logic?

12 References

[Rouben2002] Rouben, B. *Introduction to Reactor Physics*, Atomic Energy of Canada Ltd., September 2002.

[Rouben2008] Rouben, B., *Reactivity Coefficients*, McMaster University course EP 4D03/6D03, September 2008.

[AECB1977] AECB. *Regulatory Document R-10: The Use of Two Shutdown Systems in Reactors*, AECB, January 11, 1977.

[AECB1991a] AECB. *Regulatory Document R-8, Requirements for Shutdown Systems for CANDU Nuclear Power Plants*, AECB, February 21, 1991.

[AECB1991b] AECB. *Regulatory Document R-7, Requirements for Containment Systems for CANDU Nuclear Power Plants*, AECB, February 21, 1991.

[AECB1991c] AECB. *Regulatory Document R-9, Requirements for Emergency Core Cooling Systems for CANDU Nuclear Power Plants*, AECB, February 21, 1991.

[CSA2011a] CSA. *N290.4-11: Requirements for the Reactor Regulating Systems of Nuclear Power Plants*, Canadian Standards Association, October 2011.

[CSA2011b] CSA *N290.5-06: Requirements for Electrical Power and Instrument Air Systems of CANDU Nuclear Power Plants.*

[CANDU1999] CANDU. *Standard for Software Engineering of Safety Critical Software, CE-10010STD, Revision 2*, CANDU Computer Systems Engineering Centre of Excellence, December 1999.

[IEC2001a] IEC. *IEC 61508, Functional Safety of Electrical/Electronic/Programmable Electronic Safety-Related Systems*, International Electrotechnical Commission, June 2001.

[IEC2001b] IEC. *IEC 61513, Nuclear Power Plants – Instrumentation and Control for Systems Important to Safety – General Requirements for Systems*, International Electrotechnical Commission, March 2001.

[IEC1993] IEC. *IEC 61131-3, Programmable Controllers, Part 3: Programming Languages,* International Electrotechnical Commission, March 1993.

13 Acknowledgements

The following reviewers are gratefully acknowledged for their hard work and excellent comments during the development of this Chapter. Their feedback has much improved it. Of course the responsibility for any errors or omissions lies entirely with the authors.

Werner Fieguth
Dave Fournier
Norm Ichiyen
Alek Josefowicz

Thanks are also extended to Diana Bouchard for expertly editing and assembling the final copy.