CONTROL SYSTEM PHILOSOPHY FROM THE

NUCLEAR REACTOR FIELD

A. Pearson

Chalk River Nuclear Laboratories

Chalk River, Ontario

Introduction

A Canadian control philosophy began to emerge about 14 years ago when it was first argued that nuclear reactors should be 'automatically' controlled. This beginning is usually overlooked these days when 'control' has become synonymous with 'automatic control', but in the nuclear reactor field the role of the operator has been and still is the subject of considerable debate. On Canadian reactors the designers generally have attempted to minimize the operator's part in closing control loops, leaving him free to keep equipment at its peak and to be a diagnostician when required, for in this role the operator is indispensable. In fact in some cases proper manual controls do not exist and an unserviceable automatic controller forces a reactor shut down.

As a result of this policy our reactors have been highly automated for some time, a situation not shared by other design groups. The differences in approach, however, are due only partly to a differing philosophy since the basic reactor concept is an important factor in the control design.

Having decided that automation was fundamental a concerted effort was made to design very reliable control systems; reliable enough in fact that they might accept some share of the responsibility for safe reactor operation. In the nuclear reactor field a sharp distinction is made between safety and control and the proper mix of safety and availability produces for us a reliable system.

Safety systems guard against abnormal excursions in station parameters and cause irreversible control action in the safe direction whenever an excursion is detected. Redundancy of sensors and shut-off mechanisms is the well established method by which certainty of shutdown is ensured. In earlier systems redundancy was used to such an extent that lack of station availability became an important factor and majority logic became popular to help circumvent the problem.

Control systems on nuclear reactors, in common with control systems anywhere, are supposed to keep station parameters within bounds in the face of operational disturbances. However, since control systems can cause parameters to go up or down generally with equal ease, a faulty control system is considered to be a prime candidate for causing an unsafe excursion. To improve this situation the multi control-channel approach was introduced and majority logic applied here also.

To make these terms clearer Figure 1 shows some simple configurations that could appear in a safety system. In each case an input is connected through a set of switch contacts to an output and safety action occurs whenever the output disappears. Operation is self-explanatory except for the middle one where we note that each switch has two contacts operated together so that opening any two switches removes the output.

If of course one is considering a control system where the requirement is to connect an output rather than disconnect it, then the top and bottom configurations reverse their roles.

Figure 2 shows in a more general way the affect of using various arrangements of redundant components. It is evident that a conflict exists between the desire for a very safe system and for a system with high availability. Similar conflicts are bound to appear in the big accelerator field. Capital investment is high and protective devices must be used to ensure that in the event of misoperation the machine will suffer no serious damage for if it did the accelerator project could end. On the other hand if the approach is too cautious the machine may not be available to the experimenter for long enough periods to be worthwhile.

In this paper we consider various redundancy techniques only as they apply to control systems.

Techniques

It is evident that somewhere in a control system there must be at least one common element. We have only one reactor and, practically, only one element with

which to control its reactivity. The common path however is made as short and as reliable as possible.

Figure 3 shows the basic elements of a multi control-channel system used in several of our reactors. Redundancy begins with the sensors and is maintained until a control signal is derived. The system operates provided that at least two of the channels agree. Differential relays 1, 2 and 3 intercompare the channel outputs and upon detecting a disagreement operate the appropriate contacts to disconnect the straying channel. Whether or not the disconnection is necessary depends on the characteristics of the system. If the fault is not disconnected any error that can appear at the channel output must be limited so that it cannot swamp the efforts of the other two to compensate for it. Also, without rapid disconnection the transient disturbance may be unacceptable.

When only two channels are controlling a disagreement between them leaves us unable to decide which one is correct and generally the plant shuts down.

In cases where the gain in the error amplifier is large the sensor signals must be closely matched otherwise large permanent differences will exist between the control signals. Averaging in earlier stages as in the output is possible but the designer's aim is to avoid such interconnections. Ideally, complete independence is desirable. Without it the elegance of the approach is soon lost and the system failure rate is much higher than that predicted from the chance coincidence of random failures.

Other methods for selecting the correct signal have been used. In Figure 4 a scheme is shown that allows only the median signal to pass. The condition of the diode elements is shown assuming that channel B lies between the other two. The median mode of operation avoids many of the difficulties found in the averaging systems and is now the more common method for bringing together three signal lines. It is an interesting circuit and was devised by F.S. Goulding; implementation by computer program or relay logic is straightforward but the diode arrangement is not at all obvious.

In some of our reactors it is possible to control reactivity by adjustment of the moderator level. This has led to redundancy schemes that leave only the reactor as the single channel component. In Figure 5 the moderator level is held constant by adjusting the three drain valves so that the outflow equals the

inflow. Each control channel operates its own valve and a failure in any line can be compensated by a readjustment of flow in the other two. In this case the magnitude of the fault is fundamentally limited to either a fully open or a fully closed line.

Many other arrangements have been proposed but basic to them all is the need for at least three signals so that a faulty one can be indicated by its disagreement with the other two.

This approach to control system reliability through redundancy has proven itself in power plant service. Faults are rare and consequently statistics are poor but estimates indicate that a partial failure that would permit power levels to just exceed bounds might occur only once in three years and that more severe failures would occur much less frequently. In these circumstances protective systems would shut the plant down. But these too are very reliable so that a chance coincident failure is very remote.

Onto this scene has come the digital computer and from the reliability point of view it faces stiff competition.

Digital Computer Philosophy

Several arguments are put forward for introducing digital computers into the nuclear reactor control field. The arguments are no different from those used in other fields. Compared with alternative techniques computers offer several advantages provided that in the first place the overall system complexity can justify the substantial initial cost.

1. Procurement is faster. Equipment can be specified and purchased before exact requirements are known.

2. More elegant solutions to control problems are possible.

3. The overall control and instrumentation of a plant is unified since in general more time is spent studying dynamics and system interactions.

4. Modifications that are indicated by operating experience can be more readily put into practice through program alterations.

Having sufficiently complex systems to control and with these justifications the digital computer has become well entrenched in our control philosophy and I will outline here the practical steps that have been and are being taken

to bring about the change.

In 1963 we began a computer control experiment on the NRU research reactor at Chalk River. It was aimed at gaining reliability experience on a system having a complexity approaching that of a nuclear power plant, and further, to explore new data storage and handling techniques. Figure 6 shows part of the system. The reactor already had a multi-channel control system regulating the neutron flux level. We added a thermal power control loop to keep the thermal power constant by continuous adjustment of the neutron level set-point.

Since our first concern was reliability we decided on a two stage processor with signal scanning, digitizing, and conditioning being done by a special hardware unit preceding the general purpose computer.

Small computers were not available at the time but the Digital Equipment Corporation evidently had one in the conception stage and proposed it for this preprocessing function as part of the overall system. It scans the analogue signals at high speed, digitizes them, determines if they are within bounds, and selects the median signals from several groups of inputs. The display on the preprocessor provides a direct information read-out to the operator and is useful if the main computer is inoperative.

Information is passed onto the control computer where it is sampled and manipulated to produce a control signal that is returned to the reactor system. If the computer system is not operating, the switch is turned off.

One of the major goals of the NRU computer exercise was to gain reliability experience in as hostile an environment as possible and yet be able to take time to study failures and failure patterns when they did appear. The computer is operated 24 hours a day and is never shut down for routine maintenance.

Several types of failure occur.

1. Catastrophic component failures
2. Faulty peripheral operation
3. Transient failures of two types
   3.1 Those due to noise or marginal operation
   3.2 Those due to the infrequent appearance of a combination of computer commands that cause misoperation

And as a result of 3.,
4. Partial program destruction

5. Complete program destruction
6. Parity errors.

Our experience indicates that we have to live with transient failures, but we do not have to live with the consequences.

Partial program destruction can be very serious. Most of the program will be operating and overall diagnostic checks may not uncover any malfunction. Indeed one gets the impression that some diagnostic programs are designed to prove that the computer is working rather than to expose a fault.

Figure 7 shows a method used to expose faulty subroutines. The various tasks being done by the computer are initiated in one of two ways, either by the appearance of a periodic command or by the random appearance of a signal from some external device. The timed tasks are the most important and are given special attention. Each of these tasks is required to keep track of the number of times it is done. For example, task 2 is done every 1/60 of a second and upon completion a counter is incremented to say that the task actually was completed. The computer then goes on to the next task when the appropriate command appears. One of the jobs carried out during the execution of tasks 4 and 8 is to check all the counters to see that the correct score has been registered. If it has, the counters are reset to zero and the process continues, if not the computer is forced to stop.

Our next concern is with the system performance and a searching overall check is continually made as shown in Figure 8. In the PDP-4 computer a table of numbers is permanently stored. They range uniformly from near zero to a value equivalent to near full scale of the dynamic range of the input variables. One of the numbers is placed in the register of a digital-to-analogue converter and the resulting signal is returned to input of the system and treated as any other parameter. It is converted back into digital form sent to the PDP-4 computer where it is compared (during task 4) with the original number. If it passes all the tests a new number is taken from the table and made ready for the next test when task 4 is again initiated. Any malfunction of the input circuitry or deviation from a linear relationship between input and output is detected and the computer is stopped. If anything halts the computer (including halts produced by subroutine malfunction or parity errors) or in any other way stops the

periodic flow of information around the test loop the 'watchdog' signals and causes the link to the reactor control system to be broken. The set-point that existed prior to the failure is retained external to the computer.

Since many of the computer halts are due to transients it is generally only necessary to reload the program and re-start the system. This is done automatically by taking the program from a magnetic drum. The operation takes about one-half second after which the system is again in operation and switched back onto control.

The above precautions take care of transients satisfactorily, but catastrophic failures remain and occur on the average (taken over four years) about every 2000 hours. This in fact seems to be a feasible figure to specify for computer systems now being considered.

Our first step towards computer control in a nuclear power plant was taken by the designers of the Douglas Point Generating Station. A computer was introduced into the system as shown in Figure 9. About 500 signals are processed, a number approaching the minimum complexity to justify the initial investment. Most of the signals are for data logging and alarm and about 250 of them are associated with the reactor safety system. An output from the computer can cause an automatic reactor shut down but only when the computer signal is in coincidence with another signal that is independent of the computer.

The main control loop is a three channel analogue system with its set-point under computer control. In addition a direct digital control loop is used to control the power distribution in the reactor core. Sensors provide the computer with a temperature profile upon which this control action is based.

The set point control and the power distribution control can be done manually if the computer is inoperative and sufficient sensor data are supplied to the operator for this purpose. Loss of the computer does not mean a station shut down but it does create a greater work load for the operating staff.

Several new power plants are now under construction, four 500 MW(electrical) stations at Pickering near Toronto, and one 250 MW (electrical) station at Gentilly near Three Rivers, Quebec. All of these units are heavily committed to digital computer control with the argument based largely on the comparative

cost of alternate methods. Most of the glamour has gone and there is little room for arguments (so often associated with the installation of computers) that imply some undefined future advantage. Triplication, however, to get adequate reliability, is economically prohibitive and so our now traditional intercomparison schemes are giving way to absolute methods for identifying malfunction.

Figure 10 illustrates in a very diagrammatic way the philosophy behind these new systems. A dual computer arrangement is used with only one computer actually doing the main control functions at any one time. It should now be clear why self-diagnostic techniques have been emphasized for without them the dual computer scheme is unworkable.

As depicted in the figure computer A is controlling and continually sends signals to position several control elements (as many as 14 in the Pickering reactors) so that the correct average power and power distribution is maintained. Computer B is also receiving the necessary sensor information and performing the control calculation. The correct control signals will be appearing at its output but they are not connected. Should the diagnostic program in A indicate faulty operation the switch changes control to B. If B were faulty the plant would be shut down. The design challenge is to make absolutely certain that there is no coupling between the computers that could cause simultaneous malfunction.

Only those tasks that are vital to continuous plant operation are carried out in both computers and a large part of each computer is devoted to different sets of tasks. Loss of these plant functions is permissible for some period.

It will be seen that the number of sensor outputs has risen to nearly 2000. In order to justify the cost one attempts to spread it over as many functions as possible and care must be taken to prevent the computers from becoming a catch-all.

Many more control tasks than the one indicated are done. The control of reactivity is complex and four other systems are involved. Boiler pressure, turbine run-up, and the fuelling machine which is continually loading and unloading fuel, are functions also under computer control.

Concluding Remarks

Whether or not the justifications that have been put forward for using

digital computers as on-line control
elements will be borne out remains to be
seen.

It is my opinion that the computer
has not yet found a decisive role in the
nuclear power field, that is one where
its obvious capabilities could be used
to affect the economics of a power plant.

More elegant solutions to control
problems are not likely to pay off unless
the computer's capability is in some way
factored into the basic design of the
reactor system and I believe this applies
to accelerators and to any other system
that presents a complicated control
problem.

In our analogue world perhaps we are
making too much of direct digital control
especially in view of the rapid advances
being made in linear solid state devices.
A better role for the computer may be
decision making outside the loop.

General References

1. T.J. Thompson and J.G. Beckerley,
   Editors, The Technology of Nuclear
   Reactor Safety, Vol. I, Reactor
   Physics and Control, Chap. 6, Sensing
   and Control Instrumentation, M.I.T.
   Press, 1964.

2. A. Pearson, "The NRU Computer-Control
   Experiment", Symposium on the Use of
   Computers in Analysis of Experimental
   Data and the Control of Nuclear Fa-
   cilities, NBS-Conf-660527, Argonne
   National Laboratory, May 4-6, 1966.

3. E. Siddall and J.E. Smith, "Computer
   Control in the Douglas Point Nuclear
   Power Station", IAEA Symposium on
   Heavy Water Power Reactors, Vienna,
   September 1967.

4. J.E. Smith, "Digital Computer Control
   System Planned for Pickering Nuclear
   Station", Electrical News and Engi-
   neering, pp. 39-41, March 1967.

5. W.R. Whittal and K.G. Bosomworth,
   "Dual Digital Computer Control System
   for the Gentilly Nuclear Power Sta-
   tion", to be presented at the Inter-
   national Federation of Information
   Processing Congress, Edinburgh,
   August 5-10, 1968.

DISCUSSION

(A. Pearson)

PUTNAM, LASL: Could you clarify your remarks about
making changes in the system? We feel the computer
provides a tremendous advantage. If you have the
data and commands available, then you modify your
software to change your operating mode.

PEARSON, AECL: All I am saying is: Nothing in our
experience, and that goes over a good many com-
puters, says it is any easier to do that than to
change the logic of the hardware system. If you
have designed a program that is complicated and
interlaced, as soon as you change one thing, the
chances are something else gets changed. It
really depends upon how much flexibility you have
thought of in the first place. But if you haven't
thought of it in the first place, just as if you
haven't thought of it in any other system in the
first place, the change is just as hard to make.
This is what my experience tells me.

ALLISON, LRL: I would like to back up Tom Putnam's
comment about flexibility in programming. We have
had some experience on the type of diagnostics you
have in accelerator development. It is our experi-
ence that if you are careful with your programming,
and the point was well taken about having to give
thought to it, you can make changes. For example,
when we are in the process of doing a machine ex-
periment, we have frequently been able to change
the program to read instruments, etc. as dictated
by need. This is a difference I think between
something designed primarily for safety and con-
trol on a system of fixed configuration and an
accelerator whose hardware pieces and experimental
aims constantly change.

PEARSON, AECL: I would take my argument into that
field too where I have really had my experience.
Given the same vehicle, that is, a computer which
I can program, and a hardware system where I can
build modules, then a good circuit designer will
produce a hardware configuration as fast as a
programmer can modify the program. If they have
both taken care in the design of their initial
systems to allow for this much flexibility. As
soon as you ask either one to do something that
is not part of the original plan, I still think
the job is about the same for either.

FRANKEL, BNL: I would agree with you sir and say
it depends very much on what computer you are us-
ing in the operating system. If you are using a
relatively small computer, it is much more diffi-
cult to reprogram it than if you have a larger
computer with a very sophisticated operating sys-
tem.

PEARSON, AECL: Yes, and normally the economics
in these systems dictate that we are using the
last "cell space". If we had another four thou-
sand words we would use it. We can't yet afford
the luxury of such a sophisticated operating

system but maybe it has to come, even in these applications.

WATERTON, AECL: I am sorry to say, there is something about our own business I don't know. On your last slide you were mentioning that all five computers were on order, yet you used two computers per station. Is there something odd here?

PEARSON, AECL: There are five stations and ten computers on order. Eight type IBM 1800's for the Pickering reactors and two type SCL 810's for one of the other reactors. People even in the same design office do not agree what computer to buy.
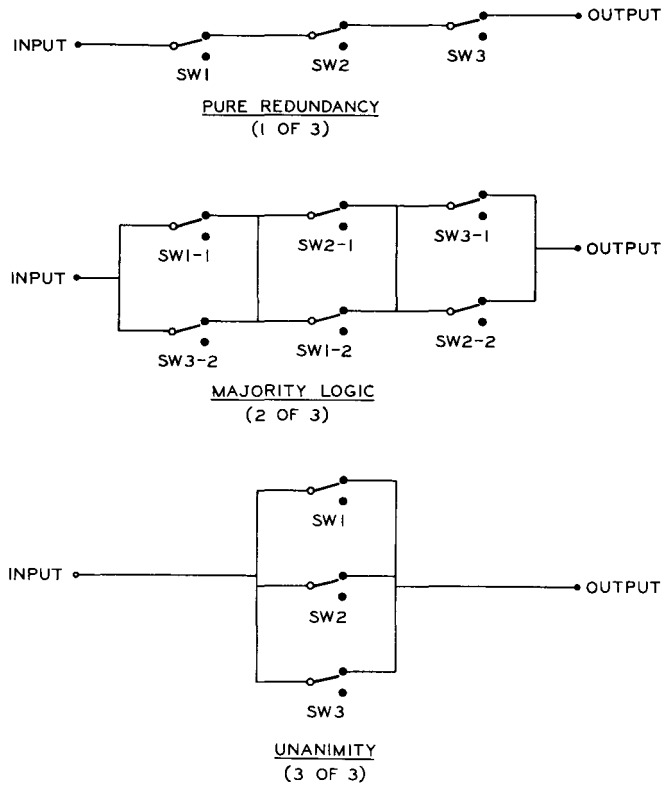
Figure 1

Redundant-Component Configurations
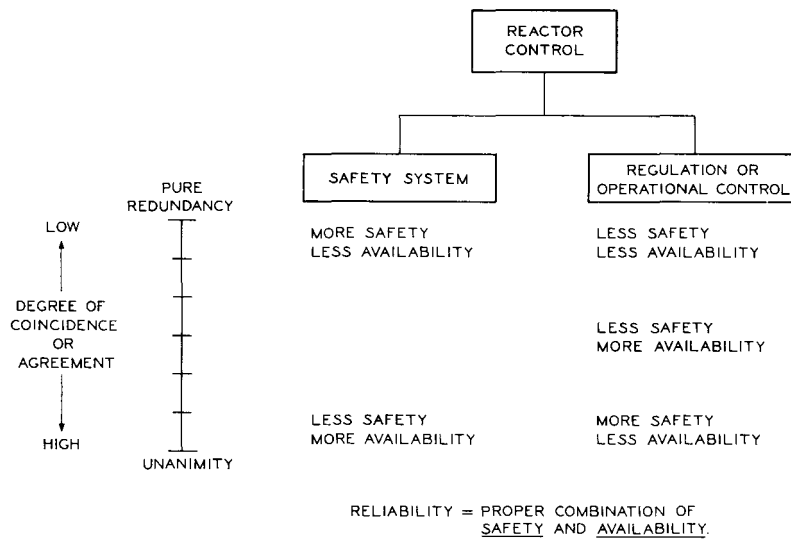
Figure 2

Reactor Control and Safety Concepts
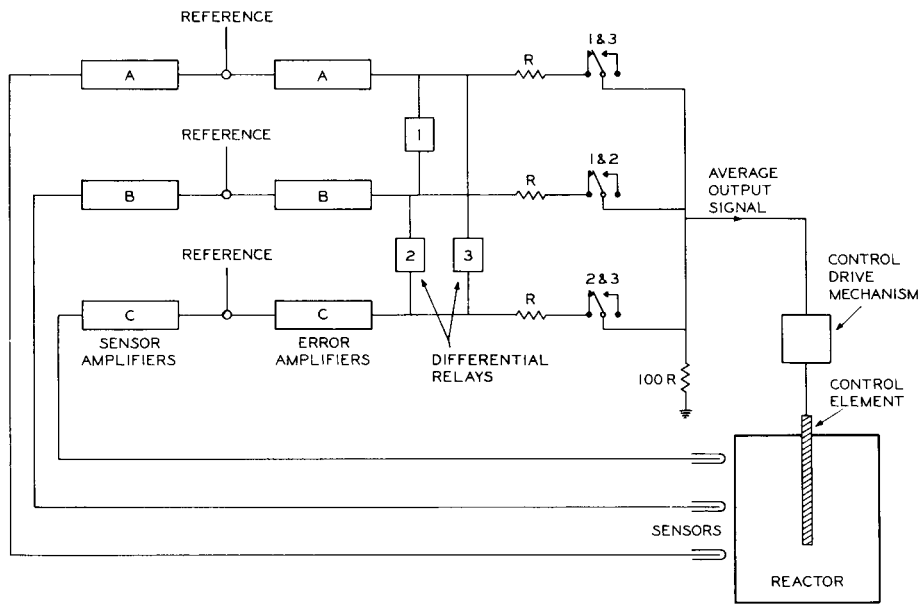
Figure 3

Basic Elements of a Multi-Channel Control System
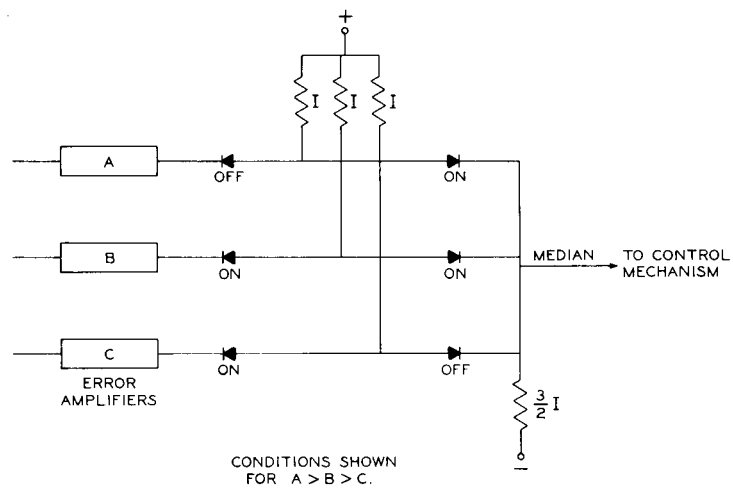
Figure 4

Circuit for Selecting the Median of Three Signals
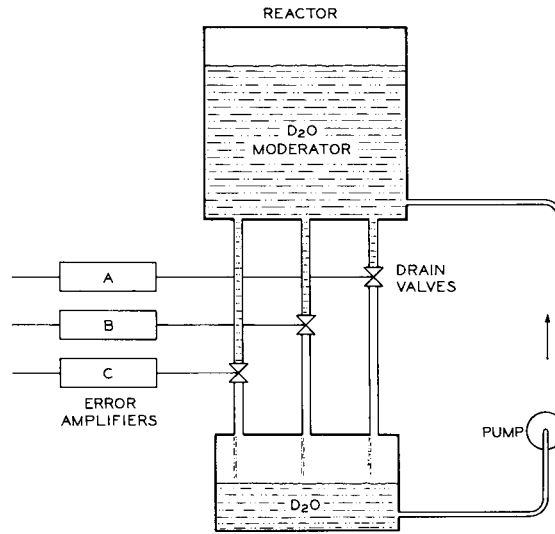
Figure 5

Redundancy Carried into Mechanical Components
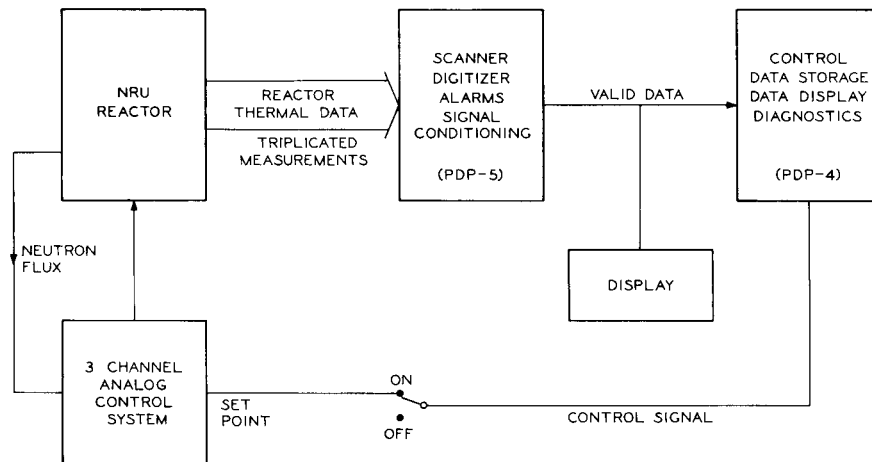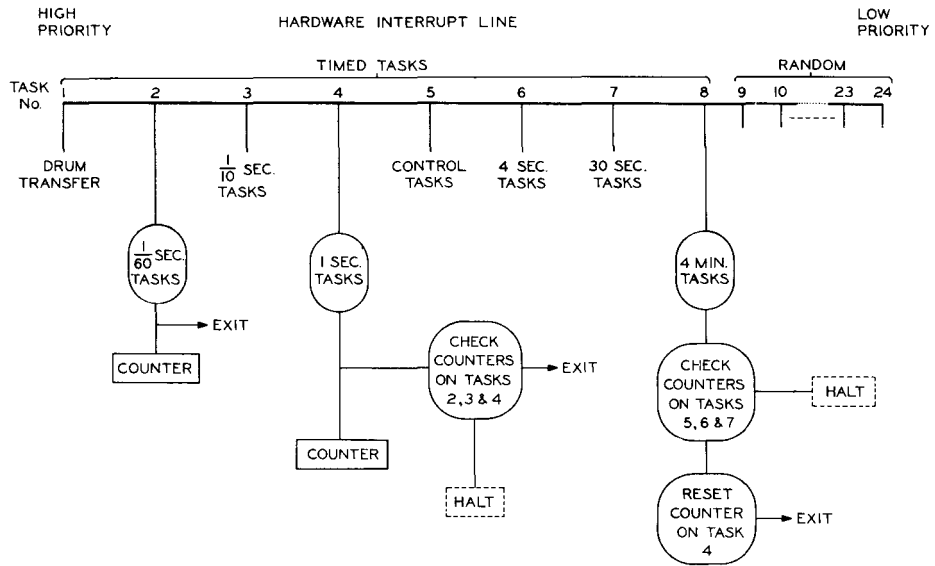

Figure 6

The NRU Control Computer

Figure 7

Checking the Computer Program

Figure 8

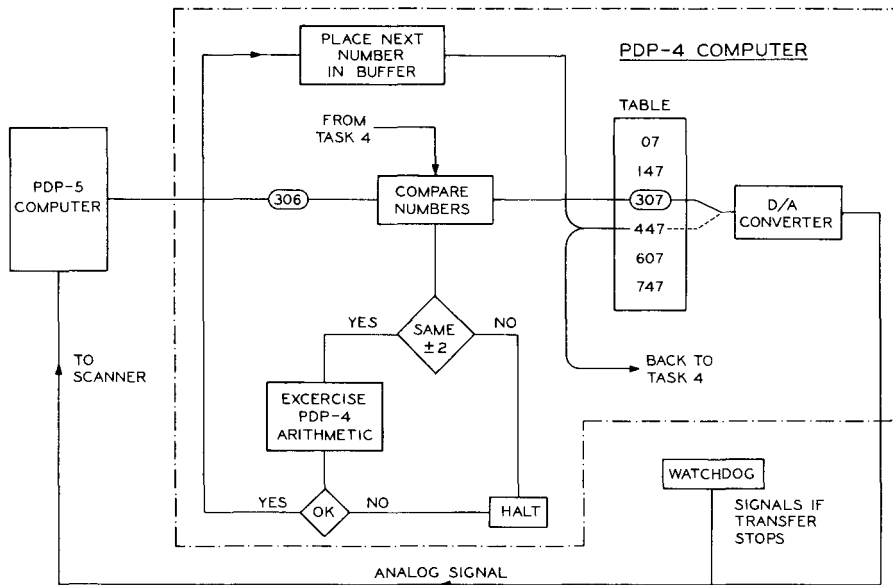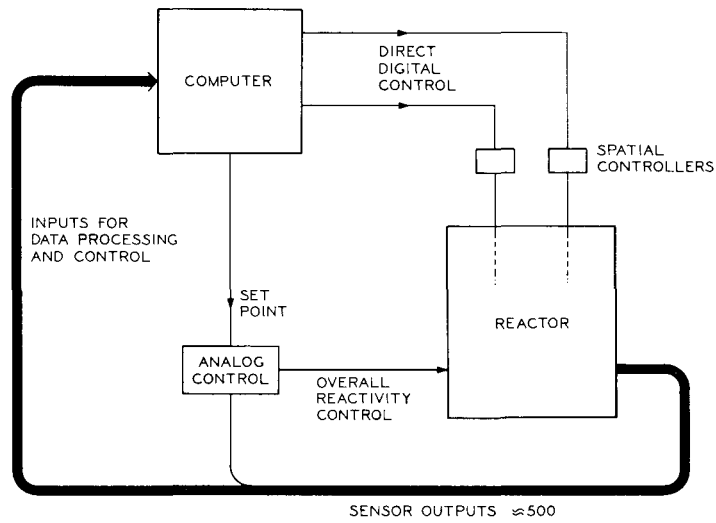Checking the Computer System

Figure 9

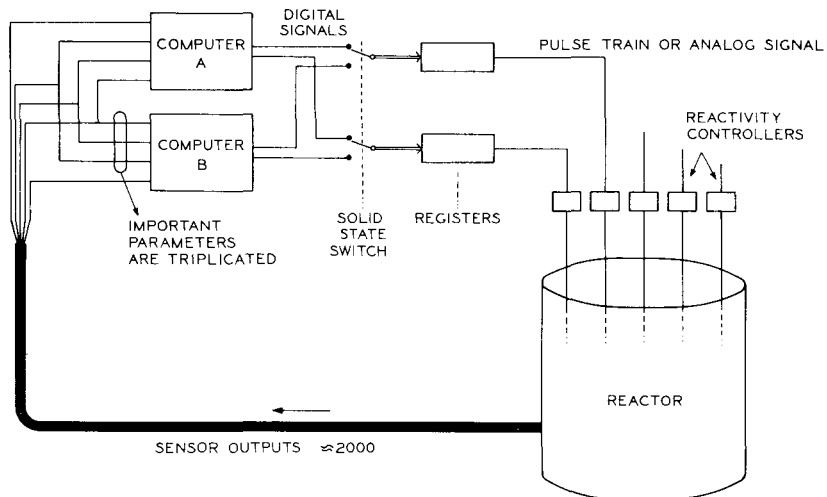Extending the Computer System to Direct Digital Control

Figure 10

A Dual-Computer Concept with Complete Direct Digital Control