CLINICAL NEUTRON THERAPY SYSTEM
Control System Specification

PART I:
System Overview
and
Hardware Organization

Jonathan Jacky
Ruedi Risler
Ira Kalet
Peter Wootton

Radiation Oncology Department[1]
University of Washington
Seattle, WA 98195

Technical Report 90-12-01

December, 1990
(Revised March, 1991)

[1]Partially supported by NIH contract no. CM97282 from the National Cancer Institute

Contents

List of Tables

List of Figures

Part I

System Overview and Hardware Organization

Chapter 1

About this document

The Clinical Neutron Therapy System (CNTS) at the University of Washington is a computer controlled cyclotron and neutron therapy treatment facility. It provides particle beams for cancer treatments with fast neutrons, production of medical isotopes, and experiments in medical physics and radiobiology. Its control system includes six processors and handles over one thousand input and output signals. It must meet demanding requirements for availability, equipment protection, and human safety. Devices under computer control include a 900 amp electromagnet and a 39 ton rotating gantry.

This document is the introduction to the functional specification of the CNTS control system. It is intended for an audience including physicists, engineers, computer science researchers and computer programmers. It contains much material describing the the cyclotron and treatment apparatus and explaining how the whole system works. This is necessary to enable the specification to be understood by readers with diverse backgrounds. It is also necessary to explain the rationale for many features, particularly those concerned with safety and equipment protection.

We created an *informal* specification. We describe the control system as thoroughly as is practical using standard technical English, supplemented by tables and diagrams. The only mathematical, or *formal*, contents are some algebraic equations and a few Petri net diagrams (which are explained prior to their use).

We believe that the specification contains sufficient information to permit the control software to be written. It should enable the reader to determine the response of all control system outputs to any combination and sequence of control system inputs, without requiring the reader to refer to the control program code, or perform tests on the actual machine. (Inputs and outputs include those that are are accessible to the users, e.g. video terminal

screens and keyboards and control panel buttons and lamps, as well as those connected directly to the cyclotron and treatment apparatus, e.g. control and status bits, ADC and DAC converters, and communication lines to auxiliary controllers).

The specification does not discuss internals of the control software that are not visible to users of the system, e.g. software design, programming languages, coding standards, etc. Moreover it does not describe the process for developing the software, e.g. integration, verification and validation, configuration control, etc.

1.1 Uses of the specification

The specification introduced by this document has several intended uses:

A guide for designing and coding the control programs The specification is primarily intended to guide software development. We believe that it contains sufficient information to permit the control software to be written. In effect, the specification constitutes an agreement between the developers and the other CNTS facility users concerning how the control system shall behave. Errors and omissions in the specification will result in errors and omissions in the software.

Every decision that affects the way the system appears to its users should be recorded in the specification. In particular, users must *not* have to refer to program design documents or program code in order to determine how the system is supposed to behave.

During design and coding, software developers must only make decisions about how to *implement* the behavior described in the specification; they must not decide what the behavior of the system shall be. If, during design or coding, they discover that ambiguities or omissions exist, or that better ways of performing certain operations exist, they will not make ad-hoc changes that affect how the system behaves. Instead, they will consult with the other facility users to determine how the difficulty should be resolved. The specification will be revised to include the results.

A guide for testing and evaluating the programs The specification is the standard of correctness for the control system. Deviations from the behavior specified there are *errors*. The purpose of testing is to determine whether the system behaves as described in the specification, which should contain sufficient detail to enable test cases to be designed and test results to be predicted. Tests can be specified while the programs are still being developed.

A guide for writing users' manuals To serve its primary purpose, the specification must be long and detailed. Therefore, it is not very well suited to serve as a user's guide or operator's manual. However, it does describe the operation of all the controls in sufficient detail to provide the source material from which those manuals can be derived. Those manuals could be written while the software is still being developed. The specification can also serve as a reference volume to resolve questions not covered in users' manuals.

A statement of safety and performance requirements The control system must perform safely and efficiently. The specification explicitly states the requirements that the system must meet, so that they are not left to default or to the whims of the software developers.

No effort will be made to meet requirements not stated in the specification.

A record of the rationale for design decisions It is not always obvious why a particular method for accomplishing some operation was chosen. The specification explains the rationale for some features.

A repository of engineering data The specification is not intended to be a general guide or comprehensive reference to the CNTS facility. However, material from diverse sources is presented here in a well-organized fashion. Therefore, this may be a useful source for physicists and cyclotron engineers.

A subject of research in computer science and software engineering It is usually considered an essential element of good practice in engineering, including software engineering, to produce a specification for a system before building it. However, it has been found to be quite difficult to produce useful specifications of large software systems. Some practitioners say that developers can only discover what they really need while they are writing the programs, or even later, when they attempt to place the initial version into service. The specification contains our current understanding of everything needed to write the control programs, which we believe is essentially complete. We will test this hypothesis by recording any additional information which we may require in order to produce a useful working system.

A specification that is more precise and compact than this one might be achieved by making greater use of mathematical notations instead of English. Such mathematical specifications are called *formal* specifications. We intend to produce a formal specification of this system, based largely on this specification. A preliminary experiment is reported in [1].

We will record the effort required to produce the present specification, and the additional effort needed to produce the formal specification. We will learn whether the information included here is sufficient to determine a formal specification. We will assess the usefulness of both specifications and compare their relative strengths and weaknesses regarding the uses described in the preceding discussion.

1.2 How to read the specification

The specification is necessarily long and detailed. However, it is arranged so that more general overview precedes detailed explanation, so that readers may start at the beginning and continue until the level of detail exceeds their interests or needs. Readers may consult later sections as needed to answer specific questions.

Part I of the specification, presented in this document, provides an overview of the CNTS facility and its control system. It does not specify the control software in detail; instead it provides context, background, motivation and rationale, and introduces the vocabulary. The following chapters appear in this document:

Chapter 2 describes the cyclotron and treatment apparatus and explains how the whole system works. It emphasizes the controlled equipment itself, rather than the control system. This background material is essential for understanding the rest of the specification. It provides the rationale for many features, particularly those concerned with safety and equipment protection. Moreover, much of the special terminology used in the rest of the specification is introduced here.

Chapter 3 is a "mission statement". It describes the sorts of treatments, experiments and other activities supported by the control system and the roles of the CNTS facility users.

Chapter 4 is an overview of the control system. It describes the hardware components of the control system, how they are connected, and the role of each component.

Chapter 5 describes the design goals of the control system.

Chapters 5 and 6 present requirements for safe and efficient operation. They provide the rationale for many features described subsequently.

Chapter 6 describes our approaches to ensuring safety.

Chapter 7 describes how the control system is protected from inadvertant or malicious

access or alteration. Particular attention is devoted to protection from access through the computer network.

Chapter 8 describes different configurations and future changes of the control system that we anticipate. The design should permit these configurations to be accomplished by adding, removing or substituting hardware and software components, rather than by modifying existing components.

We currently plan that the rest of the specification will be released in two more parts, which will include contributions from additional authors.

Part II will contain detailed specifications of operations which users perform at video terminals and control consoles. The chapters in this part will provide very detailed descriptions of how the control system provides the functions described briefly in Part I, chapters 3 and 4.

Part III will contains detailed specifications of internal operations involving the cyclotron and therapy apparatus itself, which are only indirectly visible to users.

Appendices to each part will contain detailed information which is too voluminous to include in the body of the text, for example a glossary of abbreviations, and the tables of data which largely determine the control system's operating characteristics. Many of these tables reside in a database maintained by the developers and cyclotron engineers. The contents of the database can be revised as needed and current versions of the tables can be examined or printed independently of this document set.

1.3 Other documents

Members of the Radiation Oncology Department at the University of Washington have published several papers describing the CNTS facility [2, 3, 4, 5, 6, 7, 8, 9] and some of its clinical applications [10, 11, 12, 13, 14]

The system as delivered to the University of Washington in 1984 is described in the multi-volume document sets and prints by the manufacturer, Scanditronix AB of Uppsala, Sweden and the major equipment subcontractor, Elven Precision Ltd. of Crawley, England. Particularly relevant to this project is the control system description found in the three separate Scanditronix volumes numbered 7A, 7B and 7C. These describe the PDP-11 based control system which the system described in this document replaces. They also describe a great deal of equipment which will be retained in the new system, in particular the rack of input/output (I/O) interface cards and its contents, its bus, and its analog and digital I/O

cards and their addressing and data representation.

Another source of information is the data and program source files created by Scanditronix which reside on the PDP-11 disks.

Equipment not manufactured or furnished by Scanditronix or Elven is described in the documents from the various manufacturers, e.g. Hewlett-Packard power supplies and the Digital Equipment Corporation computers.

Chapter 2

The Clinical Neutron Therapy System

The Clinical Neutron Therapy System (CNTS) at the University of Washington is a computer controlled cyclotron and neutron therapy treatment facility. It provides proton beams which are used to produce fast neutrons for cancer treatments, and proton as well as other charged particle beams for medical radionuclide production and for experiments in medical physics and radiobiology. The facility has been in operation since 1984.

This chapter describes the cyclotron and treatment apparatus and explains how the whole system works. It emphasizes the controlled equipment itself, rather than the control system. This background material is useful for understanding the rest of the document. It provides the rationale for many features described in following chapters, particularly those concerned with safety and equipment protection. Moreover, much of the special terminology used in the rest of the document is introduced here.

The major components of the CNTS facility are a Scanditronix MC50 cyclotron, an isocentrically mounted neutron beam generator, a fixed horizontal beam neutron generator, an isotope production station, beam transport system, beam diagnostics, hard wired safety system, vacuum and cooling systems and computer control system. A more detailed description of the control system hardware is found in Chapter 4.

The general layout of the facility is diagrammed in Fig. 2.1.

(Insert floor plan diagram that shows both floors)

Figure 2.1: General layout of the facility

2.1 Fast Neutron Production

The Seattle CNTS facility treats cancer patients with fast neutrons. Fast neutrons are neutrons with energies greater than several million electron volt per particle (MeV). Fast neutrons are always produced by a nuclear reaction. At the Seattle facility, a 50.5 MeV proton beam impinges on beryllium. In the resulting reaction the beryllium atoms disintegrate into neutrons, which are electrically neutral, and other fragments, which carry electrical charge (they contain the protons from the beryllium nucleus). The charged fragments have only a short range and are stopped within the target assembly; the neutrons are very penetrating and leave the target. They are then collimated and are used for the irradiation of cancerous tissue.

2.2 Cyclotron

2.2.1 Introduction

The Seattle facility utilizes a cyclotron to produce the proton beam necessary for the production of the fast neutrons. A cyclotron is a cyclic particle accelerator capable of accelerating charged particles to high energies (or velocities). The Seattle machine was built by SCANDITRONIX AB of Uppsala, Sweden. It can accelerate protons to 51 MeV, corresponding to approximately 87 percent of the speed of light. It can also accelerate other charged particles such as deuterons, alpha particles ($^4He^{++}$ ions) and $^3He^{++}$ ions.

Beam intensity is measured in particles per second. Particles from an accelerator are charged; a flow of charged particles is most conveniently expressed as an electric current. For a typical therapy run the proton beam current extracted from the cyclotron and transported to the target in the therapy head is in the order of 60 μA. At 50 MeV particle energy, this corresponds to a beam power of 3 kilowatt.

In a cyclotron the accelerated particles travel in circles (or more accurately in spirals) as opposed to linear accelerators where they travel along a straight line. Linear accelerators are widely used for conventional radiation therapy; in this application they accelerate electrons which are either used directly for therapy or are converted into X-rays. Linear accelerators can also be used for heavy particles like protons or ions. A 50 MeV proton accelerator would be quite long, measuring several tens of meters. A cyclotron, being a circular machine, is more compact. This is an advantage in a hospital based facility, where space is usually restricted. Other cyclic accelerators are betatrons, microtrons, synchrocyclotrons and synchrotrons. Each of these types of accelerators has its special characteristics which

determine its range of applications.

A cyclotron can only accelerate heavy charged particles (and not electrons, which become highly relativistic even at moderate energies).

A cyclotron consists of two major systems: a *magnet* and a *radio frequency* (RF) *system*. The magnet has the task of keeping the particles on circular orbits and to keep them focussed during the acceleration.

The radiofrequency system produces alternating electric fields which are used to give a series of kicks to the particles, accelerating them to higher and higher energies.

In addition there are other important systems:

The *ion source* is a device which produces the particles before they are accelerated by the RF system. The Seattle cyclotron is equipped with an internal ion source, located at the center of the machine. In other cyclotrons the source may also be installed externally and the low energy particles from the source are injected into the accelerator by an injection system.

When the particles reach their final energy they have to be guided out of the cyclotron by the *extraction system*.

The area in which the particle beam circulates is kept evacuated by the *vacuum system*. This is necessary to prevent the particles from being scattered by air molecules, in which case the beam would be dispersed.

The *beam diagnostic system* consists of a set of probes and current pickup electrodes and is used to measure important beam parameters within the cyclotron. It helps the operator to produce the required beam quality without damaging the machine.

The *cooling system* circulates deionized water through all components which need cooling.

Fig. 2.2 shows a general overview of the MC50 cyclotron.

Following is a short description of these subsystems with emphasis on the Scanditronix MC50 cyclotron installed at the CNTS in Seattle.

(Insert cyclotron cutaway view drawing)

Figure 2.2: The SCANDITRONIX MC50 cyclotron and its major components.

2.2.2 Magnet System

In a cyclotron the particles are kept on circular orbits by a magnetic field. The radius of the orbit is determined by balancing the "centrifugal force" acting on the particle with the Lorentz force.[1] This leads to the basic "cyclotron formula":

$$p/q = BR$$

where

p	=	momentum of the particle (m kg / sec)
q	=	charge of the particle (Coulomb)
B	=	magnetic field acting on the particle (Tesla)
R	=	radius of the particle orbit (m)

As the particles are accelerated they gain momentum and the radius of the orbit increases. The particles start near the center of the machine and spiral in increasing orbits until they reach the extraction radius from where they are guided out of the accelerator.

On each orbit they cross electric field gaps where they get accelerating kicks from the RF system. In a cyclotron the RF frequency is fixed. In order for the particles to arrive at the accelerating gaps in synchrony with the RF field, the orbiting time has to remain constant throughout the acceleration. The accelerated particle's apparent mass increases with velocity (a relativistic effect). Therefore, to achieve this isochronous situation, the magnetic field has to increase as a function of radius in order to compensate for the relativistic effect.

The magnetic field is created by an electromagnet (Fig. 2.2). It consists of a 90 ton steel yoke with the main coil, which is powered by a current of up to 900 ampere. This creates an average magnetic field of up to 1.8 Tesla between the 1.40 m diameter pole pieces.

The necessary increase of the field with larger radii is accomplished by proper shaping of the steel of the magnet, together with a set of ten concentric *circular coils* (CC) or *gradient coils* (also called *gradient correction coils* or *circular trim coils*). The coils are numbered from 1 through 10 with coil #1 near the center of the cyclotron and #10 on the outside. Each coil consists physically of two separate coils, one installed on the upper pole piece, the other exactly underneath, installed on the lower pole (Fig. 2.3). The current in these pairs of coils can be adjusted individually to achieve the desired isochronous field shape.

[1] Stated more rigorously, the radius is determined by equating the Lorentz force with the centripetal force required for a circular orbit.

The magnetic field is not uniform in all azimuthal directions. The field strength varies as the particles progress along their circular orbits. This magnetic field "flutter" alternatingly focuses and defocuses the accelerating beam, resulting in an overall focussing effect which prevents the beam from spreading and hitting components inside the cyclotron, thereby getting lost. This *azimuthally varying field* (AVF) was introduced by Thomas [15] and dramatically expanded the energy range of early machines. All modern cyclotrons are therefore AVF machines.

In practice, the azimuthal field shaping is done by steel shims, called *hills*, where the pole gap is reduced and the field strength increased, separated by *valleys*. The SCANDITRONIX MC50 has three spiral shaped hills and three valleys (Figs. 2.2, 2.3, 2.4).

Located in the three valleys are four sets of *harmonic coils* (HC), also called *harmonic correction coils* (HRMCORR), which are designated with letters A – D. Each set consists of three coils in the three valleys of the lower pole and three coils in the valleys of the upper pole piece. All the coils of one set are located at the same radius from the machine center. They are excited by DC currents in such a way that their total contribution to the field is zero. They can then be used to correct the particle paths to help keep the orbits centered within the machine. Otherwise small field deviations can add up over the repeated turns of the particles and the orbits can get uncentered with a resulting poor beam quality.

Harmonic coils D together with circular coil #9 are used to excite coherent radial oscillations of the beam just inside the extraction radius. These oscillations can be adjusted such as to produce increased separation between the last orbit and the particles on the extraction path thereby facilitating extraction.

2.2.3 Radio Frequency (RF) System

The RF system has the task of accelerating the particles. The energy which the particles obtain is supplied through the RF amplifiers which feed the accelerating structures within the cyclotron.

In early cyclotrons there were two accelerating electrodes which were formed like the two halves of a closed pill box, cut in half (Fig. 2.5). Because of their shape they were called *dees* and the name has remained even though the shape of the accelerating structures of modern machines can look quite different. The SCANDITRONIX MC50 has two wedge shaped dees with a tip angle of 90 degrees (see Figs. 2.2 and 2.6).

The particles are accelerated by the electric field at the edge of the dees. The polarity of the field is alternated such that a particle encounters an accelerating field each time it crosses

(Insert drawing of cross-section of cyclotron)

1. Pole piece
2. Hill
3. Harmonic Coils
4. Circular Coils
5. Copper Liner (grounded)
6. Dees
7. Ion Source
8. Ion Source Cathodes
9. Ion Source Chimney
10. Ion Source Window
11. Circulating Beam

Figure 2.3: Schematic cut through the center of the SCANDITRONIX MC50 cyclotron

(Insert drawing of hills, valleys, and harmonic coils)

Figure 2.4: Arrangement of magnetic hills and the four sets of harmonic coils in the magnetic valleys

(Insert drawing of early dees)

Figure 2.5: Dee configuration of early cyclotrons

a dee boundary. Inside the dee (and in the 90 degree dee system also between the dees) there is an electric field free area where the particles just coast.

Each of the dees of the MC50 cyclotron is driven by its own driver/amplifier system. A frequency synthesizer produces the selected RF frequency and feeds it into the two RF modulator units. After appropriate amplification and taking into consideration interlocks which assure safe operating levels, the RF signals get to the two driver amplifiers, which in turn control the grids of the two power tetrodes located at the cyclotron. Inductive loops couple the power from the anodes of the tetrodes to the tuning cavities which together with the dees form quarter wavelenth antennas. The cavities are tuned to the RF frequency by a mechanical moving short using the *coarse tune* servo motor systems. The *fine tune* servo motor system moves a grounded plate facing the dee, thereby changing the capacity and keeping the system in tune. The fine tune system compensates for small deviations of the coarse tune system as well as for other variations like temperature changes.

The total power consumption of the RF system is about 80 kW.

2.2.4 Particle Acceleration

A cyclotron can be operated in several harmonic modes. In the first harmonic mode the frequency of the acceleration voltage equals the revolution frequency of the particles, in the second harmonic mode the RF frequency is twice the revolution frequency and so on. In the Seattle MC50 cyclotron protons are accelerated in the first harmonic mode to final energies of 29 to 51 MeV corresponding to an RF range of 20 to 26 MHz. In the first harmonic mode the dees are operated in "push-pull" mode, indicating that at any given moment one dee has the opposite electric polarity from the other. Fig. 2.6 (with the accompanying text in Table 2.1) illustrates how a particle is accelerated in this mode.

Using the second harmonic mode, deuterons can be accelerated to 14 – 24 MeV, $^3He^{++}$ ions to 21 – 35 MeV and $^4He^{++}$ ions to 28 – 48 MeV.

The description of particle acceleration given in Fig. 2.6 is correct for particles on the main orbit. Real beams always have spreads in angle and position of particles about the main orbit. The special shape of the magnetic field creates focussing forces which draw the particles towards the main orbit. As a result they oscillate around the main orbit both in radial and axial direction. There are also phase focussing forces which keep the particles together longitudinally (or in other words slow particles are accelerated more, such that they catch up, fast particles get accelerated less, with the end effect of keeping the particles together in bunches).

Acceleration of a Particle in a Cyclotron

Acceleration of a particle in first harmonic mode in a cyclotron with 90 degree dees. This is the explanation for Fig. 2.6.

1. The positive particle enters dee #1 and is attracted by the negative potential of $-U_{Peak}/\sqrt{2}$. It increases its energy by $\Delta E = qU_{Peak}/\sqrt{2}$.
2. The particle coasts on a circular orbit (determined by the magnetic field) in the electrically field free region inside dee #1.
3. When the particle arrives at the edge of dee #1 the dee polarity has changed and the particle is accelerated again by the now positive potential $+U_{Peak}/\sqrt{2}$ between dee #1 and the grounded dummy dee.
4. The particle coasts in the field free region between the dees.
5. The particle arrives at dee #2 when the RF voltage of this dee is at $-U_{Peak}/\sqrt{2}$ and gets accelerated again.
6. The particle coasts inside dee #2.
7. The particle gets its 4th accelerating kick of this orbit when it leaves the now positive dee #2.
8. The particle coasts.
9. The cycle starts over again on the next orbit.

The overall energy gain per turn is $\Delta E = 2\sqrt{2}qU_{Peak}$. Note that in this mode of acceleration the phase between the dees is 180 degrees, they operate in push-pull mode. Similar diagrams can be drawn for the higher harmonic modes. In the even modes the dees are operated with 0 degree phase shift in push-push mode to achieve proper polarities for accelerating the circulating particles.

Table 2.1: Acceleration of a particle in a cyclotron

(Insert drawing showing dees and phases)

Figure 2.6: Acceleration of a particle in first harmonic mode in a cyclotron with 90 degree dees

The final energy of the particles is determined by the extraction radius R and the RF frequency f. The velocity of the particles at extraction is $2\pi Rf$. This translates into kinetic particle energy using the relativistic formula

$$T = E_0(\gamma - 1)$$

with

$$\gamma = (1 - v^2/c^2)^{-\frac{1}{2}}$$

where E_0 is the particle rest mass (in energy units)

In order to accelerate a beam to this energy the RF frequency is determined. The magnetic field strength, the shape of the magnetic field and the dee voltage are then adjusted to give isochronous conditions and optimum extraction efficiency.

For a 50 MeV proton beam an RF frequency of 25.9 MHz is used, together with a dee voltage $U_{Peak} = 40kV$. This means the particles make about 440 turns during acceleration.

2.2.5 Ion Source

The Seattle cyclotron is equipped with an internal ion source located at the center of the machine (Figs. 2.2 and 2.3). The ion source delivers the charged particles which are accelerated. The MC50 source is of the PIG type (short for Penning Ion Gauge) and operates in the following way:

Depending on the desired particle type the appropriate gas is delivered to the source via a needle valve (hydrogen for a proton beam, deuterium for deuterons, helium for alpha particles). The gas is then ionized in an electric arc similar to a fluorescent light. The arc discharge burns in a narrow vertical chimney. The discharge is started by applying a negative voltage (1000 to 1500 Volt) to two cathodes at the top and bottom of the chimney (see Fig. 2.3). Electrons are accelerated away from the cathodes towards the grounded wall of the ion source. Because of the magnetic field they cannot immediately reach the wall, instead they travel on spiral paths along the magnetic field lines up and down the chimney. They collide with gas molecules, ionizing them in the process. The ions are then extracted through the side of the chimney through a narrow slit or window (1×6 to 1×10 mm size) by the electric field from the dee tip, whenever the polarity of the dee is negative.

The cathodes in the Seattle source are of the "cold" style, meaning they are heated by the arc current flowing through them and by ion bombardment and not like in a "hot" cathode design where a heated filament is employed.

The power for the arc discharge is supplied by the *arc* power supply. This is a current stabilized supply and the operator can control the particle beam intensity by changing the arc current.

The geometry of the ion source and the surrounding central region of the cyclotron is very crucial for the proper operation of the machine. In order to have optimal conditions both for the first as well as for the higher harmonic modes, the SCANDITRONIX MC50 cyclotron is equipped with a dual ion source. There are two chimneys with the window of the $N = 1$ system facing dee #1 and the $N = 2$ system facing dee #2. All higher harmonic modes are run from the $N = 2$ chimney. The two chimneys have separate gas supply lines and the selection of the appropriate chimney is done via the gas inlet manifold.

2.2.6 Extraction System

One of the most difficult problems in a cyclotron for positive particles is the extraction of a useful beam out of the machine.

In the MC50 cyclotron the magnetic field is shaped (using circular coil #9) to produce what is called the "$\nu = 1$" resonance at a beam radius of 57.0 cm. The expression "$\nu = 1$" describes the case where the wavelength of the radial oscillations of the particles around the main orbit equals the circumference of the orbit. The oscillations then become coherent with the orbiting frequency, resulting in an increase in oscillation amplitude.

This resonance increases the spacing between successive turns of the particles. A narrow copper *septum*, installed in the circulating beam, separates the particles in the last full orbit from the particles which are being guided out of the machine. In order to increase the orbit radius for the particles in the extraction path, an electric field is applied, followed by a region of reduced magnetic field strength.

The mean extraction radius of the MC50 cyclotron is 58.2 cm. This radius together with the RF frequency determines the energy of the particle beam.

The electric field for extraction is created by the negative (up to $-45kV$) *deflector* electrode installed behind the (grounded) septum (see Fig. 2.7). The septum shields the particles on the last orbit from the field, such that it acts only on the particles being extracted. The septum is rather fragile and can easily be burnt by excessive beam losses, which occur when

the beam strikes the septum wall. A thermocouple is used to monitor its temperature, and the beam is turned off if the temperature gets too high.

Following the electrostatic deflector is the *electromagnetic channel* (EMC), consisting of a magnet coil which is powered such as to reduce the magnetic field from the main coil. In the MC50 cyclotron the current running through the main coil is combined with the current from the EMC power supply to feed the EMC coil with a total current of up to 1200 Ampere.

The mechanical position of the deflector/EMC assembly can be adjusted under remote control using three servo motor units (DFLENTR, DFLEXIT, EMCEXIT).

After the EMC two passive mild steel *focussing channels* give the beam the desired focussing properties and an electromagnetic steering magnet (*internal steering magnet*, ISM) helps bring the beam onto the axis of the beamline emerging from the cyclotron.

With a well adjusted set of cyclotron parameters it is possible to extract between 60 and 85 percent of the internal beam and transport it down the beam line. This extraction efficiency varies with the type of particle and the particle energy.

2.2.7 Vacuum System

The area inside the cyclotron where the beam circulates is under vacuum to prevent the beam particles from being scattered by gas molecules. The vacuum is created by two pump groups, PG1 and PG2, located on opposite corners of the machine (Fig. 2.7).

Apart from the control system, the vacuum system together with the associated part of the cooling system is the only part of the facility which is running continuously. It would take too long to obtain good vacuum if the system was shut down overnight.

Vacuum is not only maintained in the cyclotron itself but also in the beam lines. While the beam line pumps are smaller in size, the general arrangement of the pump groups is always the same (Fig. 2.8).

The vacuum pump groups are comprised of a combination of two pumps, a *mechanical pump* (MP) and an oil diffusion pump (DP). These two pumps complement each other. The mechanical pump cannot produce the high vacuum necessary for the operation of the accelerator, but can pump against atmospheric pressure on its exhaust side. The diffusion pump can pump the tank to 10^{-6} mbar but requires a backing pressure of less than 0.5 mbar.

(Insert cyclotron layout print)

Figure 2.7: MC50 cyclotron layout showing extraction elements, beam probes, RF and vacuum system

The mechanical pump is used for the initial pump down of the system to the point where the diffusion pump can take over. This is the purpose of the bypass line. In this mode, the mechanical pump acts as a roughing pump. After the bypass line is closed, the mechanical pump acts as a backing pump for the oil diffusion pump.

Pump group PG2 is not equipped with a bypass line. The roughing of the cyclotron is entirely done by PG1 which has a large mechanical pump.

A water cooled baffle reduces backstreaming of oil vapors from the diffusion pump into the vacuum tank. Medium vacuum and high vacuum gauges monitor the performance of the system.

2.2.8 Beam Diagnostic System

In order to monitor the particle beam inside the cyclotron several beam diagnostic sensors are provided (Fig. 2.7).

The *main probe* is a motor driven rod which can be inserted into the machine and goes in to 17 cm radius. It intercepts the beam, which is stopped on the Tantalum probe tip. The main probe has a dual tip, the first part intercepts a portion of the circulating beam 1 mm wide, the second part stops the rest. The current striking each portion of the tip is measured by a *beam current amplifier* (BCA) and is displayed at the cyclotron control console.

The main probe can for instance be used to check the isochronism of the beam. A proper beam has the same intensity independent of radius, indicating that no beam is lost during acceleration.

The *deflector probe* is similar to the main probe, but has a shorter range. It is used to measure the beam intensity of the last orbit and the intensity of the beam on the extraction path after it has passed the electrostatic deflector and the septum. This probe has a single Tantalum tip.

The beam probes cannot absorb the full beam power. For diagnostic purposes, the beam is run at lower currents to reduce the power.

At the entrance of the EMC, the Focussing Channels and the Internal Steering Magnet electrically isolated graphite pickups are installed. They intercept the outer fringes of the beam, which are not very useful anyway. These pickups are also connected to beam current amplifiers with readouts at the cyclotron control console. The currents are indications of

(Insert diagram of vacuum pump group)

Figure 2.8: Schematic diagram of a vacuum pump group

the beam quality and can help to center the beam on the extraction path.

In order to produce a useful beam from the cyclotron the operator monitors probe currents and losses on the pickups. Additional information is available from diagnostic elements in the beamline outside the cyclotron. The operator adjusts coil currents, the deflector voltage, mechanical deflector positions and RF parameters to optimize the beam intensity and minimize losses in the machine. This activity is called *tuning* the cyclotron and the beam lines. A major function of the control system is to provide the operator with the means to tune the system from the *control desk* (CD), sometimes also called the cyclotron control console.

2.2.9 Cooling System

Cooling water is circulated through all the components which heat up during operation. This includes magnets heated by the current in the coils, the dees and other RF carrying structures, the ion source, where the arc burns and components such as the septum and the probe tips which are exposed to the particle beam. Most of the cooling system only needs to operate when the system is turned on, ready to produce beam. The cooling for the vacuum pumps has to be operational continuously. This part is supplied by a separate loop, the *24-hour loop*, which is running all the time (Fig. 2.9).

The cooling system also cools components along the beam line and in the therapy units. It also removes heat from air conditioning units in the Cyclotron Vault and in the Power Supply Room. The CNTS facility has its own dedicated cooling tower to handle the 500 kW cooling load. The capability of a cooling tower depends on outside weather conditions, in particular on the wet-bulb temperature. In order to adapt to the limitations of the tower, the cooling system has its own (pneumatic) control system which keeps the temperature of the *primary cooling water*, circulating through the cyclotron and beamline equipment, at a constant temperature of 20 degree C in winter and at 25 degree in summer. This is accomplished by modulating the 3-way valve in the *secondary cooling water* system. The temperature of the 24-hour loop for the diffusion pumps is further reduced by a *chiller unit* such that its temperature is about $4 - 5$ degrees cooler than the primary loop.

The primary cooling water is deionized for two reasons. It is necessary for water that is in contact with high-voltage components such as the deflector to be a good insulator. Also, the water must be free of elements such as sodium which would become radioactive. The secondary cooling water is not deionized, but contains glycol to inhibit freezing.

(Insert diagram of cooling system)

Figure 2.9: Cooling system

2.3 Beam Lines

After the beam leaves the cyclotron it is transported to various target stations, depending on its use. This is done by the beam line system (Fig. 2.10). It basically consists of a set of evacuated pipes, of about 60mm diameter, equipped with beam control and monitoring devices.

Following is a short description of the various components installed along the beam lines. Most components are repeated in several locations as similar functions are required along all beam lines.

2.3.1 Vacuum Pump Groups (PG) and Beam Line Valves (BLV)

The vacuum in the beam line system is maintained by six pump groups, separated by *beam line valves*. This ensures good vacuum (a few times 10^{-6} mbar) in all parts of the lines and facilitates maintenance, as the system can get shut down in sections. The *beam exit valve* (BEV) separates the cyclotron from the switching magnet pump group. For a schematic diagram of a group see Fig. 2.8.

2.3.2 Fast Valve (FV)

If an accident occurs in the isotope production line (break of an entry foil of a gas target) a fast valve sensor in this line triggers the fast valve near the exit of the cyclotron. This valve closes within 25 milliseconds and protects the cyclotron from the fast pressure rise and potential debris.

2.3.3 Steering Magnets (STMG)

Steering magnets are used to make small corrections in the direction of the beam, in order to get it centered on the beam line axis. For instance the internal steering magnet (see under extraction system) together with the first steerer (X0) in the beam line is used to line up the beam on the axis of the first beamline section. Steerers are provided for corrections both in the horizontal plane (left/right or x-direction) and in the vertical plane (up/down or y-direction).

(Insert beam line diagram)

Figure 2.10: Layout of the beam line system

2.3.4 Quadrupole Lenses (Q)

Magnetic quadrupole lenses are used to focus the particle beam much like a light beam is focussed by optical lenses. Quadrupoles are used in sets of three called a *triplet* configuration. In the case of quadrupole #1 (Q1) and the gantry quadrupole (QG) all three parts of the lens (LENS 1, LENS 2, and LENS 3) can be controlled separately, in all other cases (Q2A, Q3A, Q2B and Q3B) lens 1 and 3 are electrically run in series and are controlled as one unit (LENS 13).

2.3.5 Stray Beam Detector (SBD)

Stray beam detectors are electrically insulated graphite donuts installed at strategic locations along the beamline. They have several functions. The beam tends to have "tails", particles that travel close to the main portion of the beam but are too far out to be of use. They are intercepted by the stray beam detectors which have center hole diameters chosen to just let the useful beam pass. The beam arriving at the target stations is then free of unusable components. The beam losses on the stray beam detector can be measured and displayed at the control desk. This current is an indication of beam shape and beam position.

The monitoring of the currents from the stray beam detectors is again done in the beam current amplifier units. Apart from amplifying the signals to useful levels and transmitting them to the control desk, the BCA also contains current level monitors which trip if the current on a stray beam detector exceeds a set limit. The BCA then turns off the beam. This protects the beam line equipment from being damaged by a poorly steered or focussed beam.

The stray beam detectors are part of the beam diagnostic system.

2.3.6 Beam Profile Monitor (BPM)

Beam Profile Monitors, sometimes also called Beam Scanners are used to obtain information on the beam shape. They consist of a rotating wire loop which moves through the beam. The beam particles knock out secondary electrons from the wire. This current signal is a measure of the beam intensity at that point. It is amplified and displayed at the control desk on an oscilloscope. The oscilloscope trace then shows the beam intensity as a function of the wire position and a well centered beam without tails appears as a Gaussian curve. The geometry of the wire loop is such that it crosses the beam at different times both in

the horizontal (x) direction and in the vertical (y) direction providing information in both dimensions.

The operator adjusts parameters to obtain the proper beam shape and position at various positions along the beam line while tuning the line. The beam profile monitors are part of the beam diagnostic system.

2.3.7 Faraday Cup (FC)

A Faraday cup is a beam stop which can be inserted into the beam line to intercept the beam. It is used to prevent the beam from continuing down a beam line if that portion of the line is not ready to accept it, and also to measure the intensity of the beam. For instance Faraday Cup 1 (FC1) can be inserted to monitor the beam from the cyclotron in order to tune it, while a patient is being prepared in one of the treatment rooms.

The Faraday cups are built to withstand the full beam power of up to 3.5 Kilowatts. For this purpose they are water cooled. The portion of the cup exposed to the beam is made out of graphite. Graphite can withstand the heat, and it also does not become unduly radioactive from the bombardment. The 50 MeV proton beam penetrates about 13 mm into the graphite. Over this distance the protons get slowed down, primarily by collisions with the electrons of the carbon atoms, until they are stopped. Electrons flow into the cup to balance the protons that are stopped there, giving rise to a current. The Faraday cups are electrically insulated from the rest of the equipment, and the current can be measured via a wire hooked up to a beam current amplifier unit.

The Faraday cups and the beam current amplifiers are part of the beam diagnostic system.

2.3.8 Switching Magnet (SWM, SWTMAGN)

The switching magnet is used to direct the beam to either the *isocentric treatment room* (line A), into the *zero degree line* (at present not installed) or into the *fixed beam treatment room* (line B).

2.3.9 Bending Magnet (BENDMAG)

The bending magnet (also called 48 degree magnet) can divert the beam from the fixed beam line into the *isotope line* (line C) where the production station for radionuclides is

installed.

2.3.10 Nuclear Magnetic Resonance Unit (NMR)

As a result of hysteresis or changes in the thermal environment, it may be necessary to alter the current flowing to a magnet in order to establish or maintain the correct magnetic field which will optimize the beamline tune. This is especially important for the switching magnet and bending magnet, where the fields must undergo large and rapid changes in order to switch the beam into the intended beam line.

Utilizing the interaction between the magnetic field and the nuclear magnetic moment of protons in the sample of a probe, an NMR magnetometer is able to measure the magnetic field in the uniform field region of a magnet and regulate the output of the magnet's power supply to produce the desired field strength.

An NMR unit can be used to regulate the field in either the switching magnet or the C line bending magnet.

2.3.11 Isotope Production Station

The isotope production station at the end of the C line is remote controlled from the isotope laboratory adjacent to the cyclotron facility (not shown in Fig. 2.1).

2.3.12 Beam Plugs (BP)

The beam plugs are removable radiation shields mounted in the shielding walls between the cyclotron vault and the treatment rooms. They shield the treatment rooms from stray radiation from the vault. This allows beam operation in the vault while setting up a patient. The beam plug is a wedge of polyethylene and steel. It can be opened by pulling the wedge up along a sloped rail to let the beam pass underneath. It is not meant to stop the particle beam and is therefore not cooled. The particle beam must be stopped on a Faraday cup when the plug is closed.

2.3.13 Beam Line Elements in the Isocentric Gantry

The beam line extends through the arm of the isocentric gantry to the beryllium target in the head (Fig. 2.11). The bends in the particle path are achieved with the *gantry dipole* (GTYDPOL) which is made up of two parts, the 70 degree magnet and the 160 degree magnet. They are electrically run in series. A correction coil (GTYCORR) on the 70 degree magnet is used to balance the two magnets. A quadrupole triplet (*gantry quadrupole*, QG) focusses the beam and a steerer (XG/YG) is used to center the beam on target. In the XG/YG steerer a 100 Hz AC component (parameter names XWG/YWG) is added to the DC steering current. This causes the beam spot to describe a circle on the target thereby spreading the thermal load and insuring a uniform distribution of the proton beam. This portion of the magnet is called the *wobbler* or *spinner*.

A graphite stray beam detector in the gantry arm (SBD-G), a tantalum stray beam detector in front of the target (SBD-T) and a set of four copper quadrants (UP, DOWN, LEFT, RIGHT), which monitor the beam position at the target entrance, complete the beam line components of the gantry arm.

2.3.14 Beam Line Elements Connected to the Fixed Beam Treatment Unit

The final steering of the beam onto the fixed beam target is done by steering magnet XF/YF immediately after the fixed beam plug. This magnet has the same role as the gantry steering magnet XG/YG in the isocentric gantry and is also excited with a superposition of a DC current for steering and a 100 Hz AC current for wobbling.

2.4 Isocentric Treatment Unit

2.4.1 Introduction

The isocentric treatment unit (Figs. 2.11, 2.12) allows patient treatments with beams coming from different directions. The radiation source, in this case the beryllium target, is mounted on a gantry arm, which can be rotated around the patient. The arm carries the treatment head to which the collimation system is bolted. The arm, head and collimator weigh approximately 11 (metric) tons, the whole rotating assembly, including the counterweights, weighs about 39 tons. The point where the gantry axis and the collimator axis intersect is called the isocenter. This is typically where the tumor within the patient is positioned. The

(Insert isocentric gantry diagram)

Figure 2.11: Isocentric gantry including treatment head

isocenter (as well as the cyclotron midplane and the beamlines) is located 120 cm above the floor. The "Isocenter Distance" or "Source to Axis Distance" (SAD) between isocenter and beryllium target is 150 cm.

During the treatment the patient is lying on the treatment couch, very rarely sitting in a treatment chair. The couch (or chair) is attached to the *patient support assembly* (PSA), which is motorized and which allows the patient to be positioned with an accuracy of 1 to 2 mm.

The isocentric gantry is capable of rotating full 360 degrees. Because of the size of the gantry, a 3 m deep pit is necessary underneath the treatment couch, to allow irradiations from below. For access and safety reasons this pit is covered by a *moving floor*. It consists of two independently movable halves, driven by electric motors, which can be pulled out of the way automatically to let the gantry enter the pit (see Fig. 2.14).

2.4.2 Treatment Head

The treatment head is diagrammed in Fig. 2.11.

Target

The neutron beam is produced by bombarding the beryllium target with 50.5 MeV protons. The Seattle facility uses a "semi-thick" target, in which the protons are slowed down to half their initial energy in the cylindrical beryllium piece and then are stopped in a copper backing. This requires a beryllium thickness of 10.5 mm. The diameter of the beryllium cylinder is 12.7 mm, just slightly larger than the proton beam diameter of 10 mm, defined by the quadrant opening at the target entrance. The copper stop is 3 mm thick and has cooling water passing over its downstream surface.

All protons are stopped in the target assembly. The neutrons, created in nuclear reactions with the beryllium, leave the target primarily in the forward direction. In addition gamma rays are produced. They are also penetrating and are part of the therapy beam. Between 4 and 12 percent of the dose delivered to the patient is from gamma radiation.

The target assembly and other parts of the equipment exposed to the proton or neutron beam become radioactive. The gamma and beta radiation from this activity is present even after the beam has been turned off. This results in radiation exposure to personnel working with the equipment either for therapy or maintenance. The equipment and the building were designed to reduce this exposure and appropriate working procedures have

(Insert isocentric treatment room illustration)

Figure 2.12: Isocentric treatment room showing gantry, patient support assembly and mobile pedestal

been implemented.

X-ray Drawer

Immediately downstream of the target the so-called X-ray drawer is located. It consists of a massive steel block which can slide in a direction perpendicular to the beam axis.

The X-ray drawer can be moved to three positions: SAFE, TREAT and X-RAY. In the SAFE middle position a lead block is moved in front of the target to reduce the amount of gamma radiation coming from the target area when the beam is off. In the TREAT position the *primary collimator* is positioned in front of the target. In the X-RAY position an X-ray tube is positioned on the beam line axis.

The X-ray drawer is moved automatically into the TREAT position at the beginning of each treatment and back to SAFE at the end.

Primary Collimator

The primary collimator is a cone-shaped aperture immediately downstream of the target, which roughly collimates the neutrons emerging from the target in the forward direction. It reduces the intensity of the outer part of the beam which is not used for therapy.

Flattening Filter

Immediately after the primary collimator, and still installed on the X-ray drawer block, the *flattening filter* assembly is located. The flattening filter assembly consists of a set of steel absorbers, which absorb more of the center portion of the beam, thereby flattening the dose distribution across the treated area.

There are three possible filter conditions: no filter, small field filter and large field filter, which are chosen to give optimal conditions for a particular treatment.

X-Ray Tube

An X-ray tube can be positioned downstream of the Beryllium target on the beam line axis. It is used to approximately simulate the neutron beam geometry with X-rays for verification of the treatment setup. A film is exposed through the patient and the treatment field outline can be verified against anatomical landmarks.

The neutron beam cannot be used directly for this purpose because neutrograms do not give enough anatomical detail. There can be discrepancies between the neutron beam and the X-ray picture because the X-ray tube is not located exactly at the target position. This has to be taken into account when interpreting the X-ray film. In critical cases a neutrogram, showing the field outline can be superimposed on an X-ray image, showing the anatomy.

Ion Chamber

A transmission ion chamber follows the x-ray drawer. The neutrons (and gamma rays) pass through this chamber on their way to the patient. The chamber measures the intensity of the beam. It is connected to an electrometer system which integrates the signal. The whole system is called the *dosimetry system*. It is calibrated and is used to determine the dose delivered to the tumor and surrounding tissue. For safety reasons the dosimetry system is split into two completely independent systems with dual chambers, power supplies, electrometers and means to terminate the treatment.

The ion chamber is the last fixed unit in the head. All components downstream of the chamber rotate together with the collimator around the beam axis.

Wedge Filter System

A set of three wedge shaped tungsten blocks is mounted on the Wedge Turret. Inserted into the beam they absorb the neutrons unevenly, creating sloped dose distributions if this is desired. Only one wedge at a time is used and there is a NO WEDGE option. The wedges can be rotated and their orientation with respect to the treatment field can be at any of the four cardinal angles (or in other words the heel of the wedge can be towards the upper, lower, left or right edge of the treatment field).

Beam Defining Lamp

A light bulb is used to project an outline of the treatment field onto the skin of the patient. It is located at the same level in the treatment head as the wedges and is displaced sideways. A mirror brings the light source on the beam axis. The optics are arranged such that the apparent origin of the light coincides with the beryllium target.

Multileaf Collimator

The purpose of the collimation system is to absorb the neutrons not wanted for the irradiation of the tumor. Radiation to the surrounding healthy tissue is eliminated or reduced and complications from the treatment are minimized. The collimator opening is shaped to match the size and form of the *target area* within the patient. In the Seattle multileaf collimator this shaping is achieved by a set of 40 steel leaves which are individually motor driven (Fig. 2.13). The individual leaves are 650 mm long and have some polyethylene inserts as added neutron absorbers. The center leaves are narrower, corresponding to a width of 12.5 mm at the isocenter distance. The outer leaves have a width corresponding to 20 mm. The collimator can be set automatically via the *Leaf Collimator Controller* (LCC) or manually via switches mounted on the unit itself.

Glass Stop

The glass stop, also called the *gamma shutter*, is a 40 mm thick lead glass plate that protects people from residual gamma radiation from the collimator and head. During the treatment it is moved out of the way.

Source to Surface Distance Projector (SSD)

An optical system projects a distance scale at the beam axis. The scale is visible on the skin of the patient and is used for proper positioning.

TV Cameras

Two closed circuit TV cameras are mounted on the outside of the collimator and are used by the radiation technologists to monitor the treatment area during the treatment from

(Insert collimator diagram)

Figure 2.13: Multileaf collimator

their console outside the treatment room.

2.4.3 Patient Support Assembly (PSA)

The treatment couch or Patient Support Assembly (Figs. 2.11, 2.12) has the task of precisely positioning the patient. It is normally equipped with a couch top, but a treatment chair can be fitted if needed.

The PSA is capable of all necessary motions:

- couch HEIGHT
- couch LATERAL translation
- couch LONGITUDINAL translation
- couch ROTATION (around a vertical axis through the isocenter)
- couch TOP rotation (around the support pillar).

The vertical support pillar is also called the *ram*. Apart from the top rotation all motions are motorized and can be controlled either from a local hand pendant or through the *Treatment Motion Controller* (TMC).

2.4.4 Moving Floor

The moving floor covers the 3 m deep pit underneath the gantry when the head is above floor level and it closes gaps around the gantry as far as possible during times when the head is below the treatment couch. The floor is divided into two independent halves constructed of linked wooden slats with carpeting on top. As each half opens up the unused part of the floor is stored in a loop in the pit (Fig. 2.14). In order to allow the patient support assembly to rotate around its vertical axis through the isocenter, the floor is supported on one side of the pit by two rails which can be retracted to let the PSA support ram pass through.

2.4.5 Auxiliary Treatment Room Equipment

The *mobile pedestal* (Fig. 2.12) is a small local control console which can be rolled around in the treatment room and which is used to locally control the equipment. It has two hand

(Insert moving floor diagram)

Figure 2.14: Moving floor

pendants attached to it, one to control the motions of the gantry and PSA, the other for the moving floor.

A large electronic wall display in the treatment room shows the current numerical value (e.g., in cm or degrees) of the gantry and collimator rotations and the PSA motions.

A patient alignment laser system projects a set of lines which intersect at the isocenter and are used for precise positioning.

A room TV camera is used to observe the patient during therapy.

A patient intercom allows the patient to talk to the technologists at the therapy console during the treatment.

2.5 Fixed Beam Treatment Unit

The fixed beam treatment unit does not have a rotating gantry. Instead the head is fitted to the end of a horizontal beam line. The fixed beam head assembly is an exact duplicate of the isocentric unit, providing the same rotational capabilities for wedges and the attached collimator. There is no leaf collimator installed on the fixed beam unit. The collimator consists of a main shielding assembly into which interchangeable cones for different square field sizes can be inserted. As this arrangement is clearly not as versatile as the isocentric gantry with the leaf collimator, the fixed beam unit is only used for experimental physics and radiobiology setups.

2.6 Shielding Doors

There are two concrete shielding doors which control access to the rooms where radiation is present: the vault door and the treatment room door (Fig. 2.1). They are hydraulically driven and run on railroad tracks. No people are inside a room with the beam turned on, with the exception of the patient during therapy. The beam can only be turned on in a room if the door is closed. There is only one therapy door for both treatment rooms because the beam is sent to only one room at a time. A system with doors was chosen over a system with access mazes because of the limited space available.

2.7 Control room

The control room (Fig. 2.1) contains the cyclotron control console (also called the control desk, CD) and two therapy control consoles (also called treatment control desks, TCDA and TCDB), one for each treatment room. Both treatment rooms are reached by passing through the control room.

2.8 Power supply room

A room on the (second) floor above the control room (Fig. 2.1) contains all the power supplies. In addition, this room also contains all the control computers, including the main control computer, the Programmable Logic Controller (PLC) that controls the vacuum and cooling systems and some other devices, the relay rack containing the hardwired safety interlock system (HSIS), the auxiliary control computers provided by Scanditronix including the Leaf Collimator Controller (LCC), Treatment Motion Controller (TMC) and Dose Monitor Controller (DMC), and the control electronics for the Elven equipment. The power supply room also contains the Scanditronix I/O cabinets that contain the input/output cards that interface signals to the main control computer. In addition the floor controller system and the main power switchboards are located in this room.

Chapter 3

Clinical Neutron Therapy System capabilities

This chapter is a "mission statement". It describes the sorts of treatments, experiments and other activities supported by the control system by describing how it assists the various CNTS facility users.

3.1 Cyclotron operations

The cyclotron operator is responsible for bringing up and shutting down the machine, tuning the machine to achieve efficient beam production, switching the beam between the two treatment rooms and the isotope production line, and restoring normal machine operation after shutdowns or interruptions caused by faults. The control system enables the operators to run the cyclotron and associated equipment from the cyclotron console in the control room.

Cyclotron operators turn many subsystems on and off and monitor their status. These subsystems include, among others, the magnet subsystem, the RF subsystem, the ion source, the extraction system, and each beam line.

Cyclotron Operators turn the beam on and off and adjust the beam while observing its characteristics. To obtain satisfactory performance, the operator must adjust, or *tune*, many different analog quantities called *cyclotron control parameters* or just *parameters*. The parameters are physical quantities such as ion source gas flow, deflector voltage, switching

magnet current, etc. The cyclotron designers calculated sets of theoretically ideal parameter values, or *tuning tables*, for each combination of particle and energy. In practice these idealized values are only starting points; variations in beamline characteristics imposed by experimental or treatment configurations and physical factors such as parameter drifts, hysteresis of magnets, variations in beam properties, mechanical misalignments etc. require additional manual tuning. A few parameters are partially controlled by automated servo loops in some operating modes, but the operator's participation is still essential.

Cyclotron operators have at their disposal a database where they can save and automatically restore machine settings, or tunings, consisting of the values of all the adjustable parameters.

The cyclotron provides relatively unattended operation. Once it has been brought up and tuned, the cyclotron operator may press a button that indicates the cyclotron is ready to produce beam. Subsequently, therapy technologists may turn the beam on for therapy operations without requiring the intervention of the cyclotron operator. However, if the beam goes out of tune or any cyclotron-related safety or equipment protection interlocks should trip, the cyclotron operator is required to restore operation.

3.2 Therapy

The CNTS facility provides neutron treatments with all the flexibility of a modern therapy linear accelerator. In addition a multilieaf collimator provides irregular fields. It is possible to perform geometrically complex treatments with large numbers of irregular fields, for example stereotactic treatments for arterio-venous malformations (AVM's).

Treatments are performed by therapy technologists. When each patient arrives for his or her treatment, a technologist uses the terminal at the therapy console to select that patient's stored treatment field specifications from a database. Usually, dosimetrists will have already prepared and stored the treatment field specifications. These specifications describe the treatment machine settings including the prescribed dose. Typically there are several different fields per patient, some or all of which are delivered during each of the patient's multiple visits.

Technologist(s) enter the treatment room with the patient to set up each treatment field: they position the patient on the couch, set the treatment machine rotations (gantry angle, collimator angle, couch turntable angle) to their prescribed values, and translate the couch (in as many directions as needed) to bring the patient into proper position with respect to the machine, as specified in the prescription.

Technologists directly control all external treatment machine motions (e.g., those which pose

potential collision hazards, including all rotations and table motions) by using the control pedestal and hand pendant inside the treatment room. This ensures that technologists are present to observe the motions in progress. The only exception occurs during arc treatments, in which the gantry rotates while the beam is on, controlled from the treatment console in the control room (the gantry and patient are viewed on a TV monitor).

Treatment machine settings which do not present collision hazards, namely wedge selection and positioning, flattening filter selection, and leaf collimator positioning are typically performed automatically (although it is possible to set them manually, that is not the preferred method). This may be done before, during or after manual setup of external motions.

While still inside the room, therapy technologists confirm that the patient is correctly set up with the aid of several devices, including the beam defining lamp which projects the shape of the field onto the patient, the patient alignment laser system which projects lines and crosshairs indicating the position of the machine isocenter onto the patient, and the SSD (source-to-skin distance) projector which indicates the distance from the beam's virtual source to the patient's skin. Technologists may also expose a port film by turning on the X-ray source in the treatment head after the patient is set up (this requires the technologist to leave the treatment room, but does not require closing the treatment room door, so it can be done quickly). When the technologists are satisfied with the setup, they leave the treatment room and close the door. (Another verification technique is to make a neutrogram by exposing a film using the therapeutic neutron beam. The procedure for making a neutrogram is essentially the same as actually treating a field, except the delivered dose is lower).

At the therapy console in the control room, the technologist presses a button to turn on the beam. If the cyclotron is ready to produce beam, and all treatment-related safety interlocks are clear, the beam comes on and remains on until the prescribed dose is delivered. Typically, this takes a few minutes. When the prescribed dose is achieved, the control system automatically turns the beam off. The technologists then open the treatment room door and enter to set up the next field, or to remove the patient and prepare for the next.

The computer control system checks that the settings made by the technologist conform to the prescription for the selected field, and stores a record of each attempt to treat a patient. These are sometimes called check-and-confirm or record-and-verify facilities.

3.3 Clinical physics

Clincal physicists have available at the treatment consoles additional special operations which are not offered to therapy technologists. These include facilities for calibrating the

dose monitors, the leaf collimator and the treatment room equipment motions. There are special facilities that enable clinical physicists to bypass the usual dosimetry, check-and-confirm and record-and-verify mechanisms in order to perform calibrations and other measurements.

A particularly important clinical physics activity is the measurement of the dosimetric characteristics of the radiation fields created by the machine. These measurements include (among others):

- Absorbed dose profiles
- Percent depth dose
- Tissue-phantom ratios
- Photon component
- Slow neutron component

These data are all tables in which the dose is tabulated as a function of position in the absorbing medium, field size and shape, target thickness and bombarding particle (the cyclotron can accelerate other particles besides protons). Typically they are measured by an instrument called a beam scanner or dose plotter, which moves a dosimeter about in a tank of water called a water phantom that acts as an absorbing medium (to simulate a patient). The measurements may require that the beam remain on for up to an hour, without interruption. Other experiments may require operation without the use of the dosimetry system where the beam is turned on and off manually.

For these measurements, it is possible to set field size, collimator angle and other treatment parameters without going through the normal patient treatment procedures. There does not need to be any fictional "patient" (or patients) whose prescription contains the desired fields.

Experimenters can observe and record the equipment operating conditions at all times, including some information that may not be displayed in normal patient treatment.

Once the cyclotron operator tunes the beam and declares it ready for use, it is possible for the experimenter to start and stop irradiation without further assistance unless the beam conditions drift out of the permitted operating windows that are safe for the equipment.

Future versions of the control system will permit machine motions to be controlled from outside the treatment room while the door is closed. This capability is required in order to

make measurements of the dose distribution from conformal therapy fields, which will be obtained as the sum of a series of discrete fixed fields of differing shape and angle. This will only be permitted in experimental modes when there is no one in the treatment room, so it presents no additional hazard to people.

3.4 Diagnostics and maintenance

Cyclotron engineers are provided computer terminal ports in the power supply room and in the control room (from where the terminal can be rolled into either treatment room). These provide many of the same functions also provided to cyclotron operators, therapy technologists and clinical physicists. However, the placement of the terminal ports and a terminal mounted on a cart with a long trailing cable enables the engineers to monitor or operate the system while they are near the equipment under test.

Some functions provided by dedicated buttons on control consoles are also provided at these terminals, for example placing the system in operating modes called Standby 1 and Standby 2. A few special facilities are available only at these diagnostic terminals. These include functions which enable engineers to type control strings to the DMC and other auxiliary controllers (with some syntax checking and error checking).

Diagnostic operations are available while routine operations are underway at the other terminals. In addition, one terminal line is connected to a modem so engineers can dial in through the telephone system and monitor the system from off site.

Cyclotron engineers can create new console terminal displays and alter machine operating characteristics by modifying tables of parameters stored in a database. This does not require additional programming.

3.5 Experiments

It is relatively easy to create modified or augmented versions of the control system to support special measurements and experiments in medical physics and radiobiology.

Chapter 4

Control system overview

The Clinical Neutron Therapy System (CNTS) at the University of Washington is a computer controlled cyclotron and neutron therapy treatment facility. The control system includes several processors and handles over one thousand input and output signals. It must meet demanding requirements for availability, equipment protection, and human safety. Devices under computer control include a 900 amp electromagnet and a 39 ton rotating gantry. This chapter is an overview of the control system. It describes the major components and how they are connected. It summarizes much of the material that is subsequently explained in far greater detail.

4.1 Overall organization

The entire control system can be considered to be divided into a proton beam control system and two treatment control systems. The proton beam control system is controlled by the cycloton operator, and includes all those control system components that are needed to run beam inside the vault (e.g. to tune the beam on FC1 or to produce isotopes), as well as the controls for the beamlines up to the targets. The two treatment control systems are each controlled by therapy operators, and comprise those components needed to perform treatments in each treatment room, including the controls for the gantry and therapy head, and the dosimetry system.

The control system is composed of several loosely-coupled subsystems that contain their own inputs, outputs, processors and memory. There are two reasons for partitioning the system in this way. First, it allows simple but essential functions to be accomplished by the simplest, most reliable, techniques available. For example, most safety and equipment

protection functions are implemented by hard wired relay logic or by a standard commercial programmable logic controller (PLC) programmed in relay logic notation. Those subsystems do not involve complex, error-prone programming language constructs, scheduling issues and so on. The second reason is that it permits staged development: it is possible to modify the control system by changing or replacing one subsystem at a time.

The division into proton beam and treatment control systems is useful for descriptive purposes. However, it must be kept in mind that some processors serve all three of these systems. For example, the same processor drives all video display terminals, including those at the cyclotron operator's console and at both treatment consoles.

4.2 Proton beam control system

The proton beam control system is diagrammed in Fig. 4.1. Isolated in a separate cyclotron vault are the cyclotron itself, beamlines to two treatment rooms, and a third beamline supplying an isotope production station located within the cyclotron vault. Two of the beamlines continue into the two treatment rooms. Also in the vault are the radio-frequency (RF) amplifiers that drive the cyclotron, and various vacuum, cooling, and other subsystems. The power supplies and most other electronics, including all control system processors, are upstairs in the power supply room.

4.2.1 Cyclotron main control computer (CMC) and local area network

Most control system functions are provided by the cyclotron main control computer (CMC). The CMC is a Digital Equipment Corporation (DEC) MicroVAX-II. The CMC is connected by an RS-232 serial line to its own console terminal, which is a DEC LA120 printing terminal and is also located in the power supply room. Other RS-232 lines connect the CMC to the two terminals built into the cyclotron control desk, the terminals at both treatment consoles, and the several auxiliary processors. In addition, there are RS-232 lines to a terminal port in the power supply room and another in the control room. It is possible to connect a VT100-compatible CRT to either port through a long RS-232 cable and run diagnostic and troubleshooting programs anywhere in the power supply room or in either treatment room. Finally, there is a terminal line connected to a modem. This line provides telephone access for remote troubleshooting the cyclotron system. The modem port is similar in function to the diagnostic ports but for security reasons only allows a subset of functions. The diagnostic terminals provide some of the same operations that are available at other terminals and some functions that are not available at any other terminal. All these functions are described in Part II of this manual.

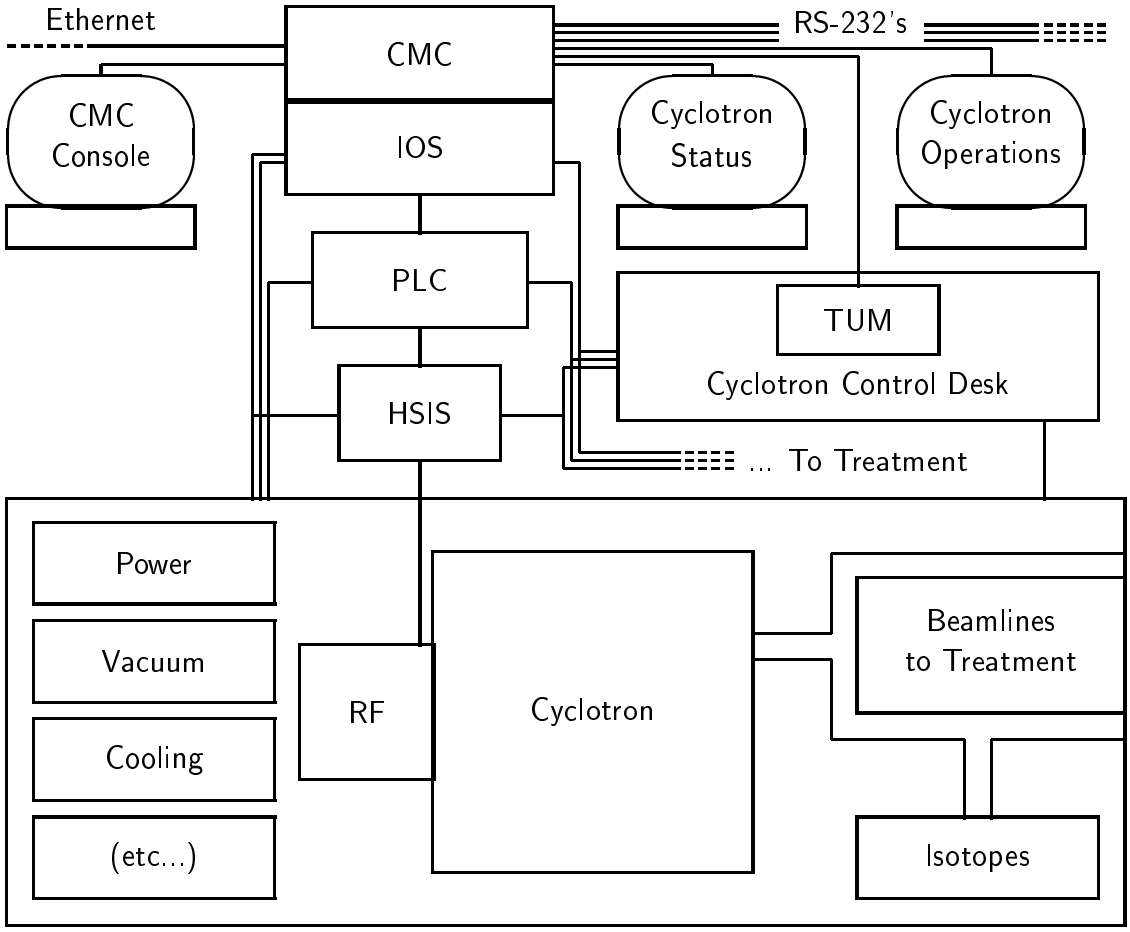

Figure 4.1: Proton beam control system

All terminals are useable only via the control system programs; they are not general purpose terminals. In particular, it is not possible to log on to some other computer from one of them.

The CMC is connected by an Ethernet local area network to a general purpose time sharing VAX computer in the Radiation Oncology Department. In this document, this computer is called the *control system support host*, or more simply the *remote computer*, the *remote host*, or just the *host*. The CMC contains no disk or tape drives. The entire software system is loaded into the CMC memory and remains in memory when running. The software is loaded from a disk drive attached to the control system support host via the Ethernet. The system software used to communicate over the network is the DECNET product from Digital Equipment Corportation.

Patient data and cyclotron system data are also stored on a disk drive on the control system support host. These files can be edited and reports generated by software on the host, independently from the operation of the control system software. Therefore except for writing treatment records and cyclotron operating logs to the host's disk and loading a new set of parameters, no disk drives are necessary to run the control system once it is started up.

Wherever possible, creation, maintenance and use of these files is accomplished by using facilities available host, not the CMC. This helps keep the CMC software simple. Examples of such operations are:

- Edit files of saved machine settings (including deletion of setting records)
- Display/print operator log
- Add/change/delete operator ID's and passwords
- Display/print treatment prescription data and treatment records

It is not necessary to write CMC software to perform these tasks.

When the system is in operation it reads and writes data as necessary to/from the control system support host via the Ethernet using DECNET software. The host and the Ethernet must be operational in order to load the cyclotron control software. Once the system is running, some operations will be possible without network access to the host. The following table shows what will be possible or not possible when the host is inaccessible.

POSSIBLE	NOT POSSIBLE
Produce isotopes	Therapy run on *new* patient
Tuning	Startup of entire system
Save active setting	Store setting in file
Restore saved setting	Restore setting from file
Treatment equipment calibration	Event logging
Operator login/logout	Write messages in log file

The control software must perform safely and gracefully if the network fails while some operation is in progress. Although file operations are not possible when the control system support host is inaccessible, considerable memory for temporary data storage is available in the CMC. This makes it possible to complete treating a patient whose treatment has begun, and to store a record of the treatment and at least some event log messages until contact with the host can be reestablished.

The CMC will not be running a general-purpose operating system such as VMS or UNIX. The CMC software is built as a standalone self-contained system using DEC's VAXELN tool kit. It will not be possible for ordinary users to log in to the CMC. As all files are on the control system support host, access to these files, as well as printing and editing their contents, will be controlled from the host. Chapter 7 describes how the control system is protected from inadvertant or malicious access or alteration. Particular attention is devoted to protection from access through the computer network.

4.2.2 Input/Output System (IOS) and Input/Output (I/O) cards

The CMC is tightly coupled through a direct memory access (DMA) interface and special bus to the cyclotron Input/Output system (IOS), which occupies a large cabinet and provides over one thousand analog and digital input/output (I/O) channels to the cyclotron and other controlled equipment. The IOS interfaces to all signals through four types of I/O cards: digital inputs (DIP), digital outputs (DOP), analog-to-digital converters (ADC), and digital-to-analog converters (DAC). The IOS hardware including the bus and the I/O cards are all proprietary designs from the cyclotron vendor.

4.2.3 Programmable Logic Controller (PLC)

The Programmable Logic Controller (PLC) is a Gould-Modicon 484 controller. It provides largely independent on/off and timing controls for the cyclotron vacuum and cooling systems, so these systems can be brought up in the proper sequence and run with little operator intervention. The PLC also controls some beamline components. For example, the buttons that the cyclotron operator uses to open and close the Faraday cup and beam plugs are connected to the PLC. These inputs from the cyclotron console are merged with others (e.g. to implement interlocks that prevent the cups from moving in some conditions). PLC outputs actually drive the cups and plugs open and closed, as well as the lamps on the cyclotron console that indicate their states.

The buttons that the cyclotron operator and therapy technologists use to actually turn on the beam (by turning on the RF drive amplifiers) are also connected through the PLC.

Some PLC functions can be indirectly observed and controlled by the CMC.

The PLC software is written in relay ladder logic notation; most was provided by the cyclotron vendor. There is a console which permits the PLC software to be modified; some modifications have been designed and installed locally.

4.2.4 Hard-wired Safety Interlock System (HSIS)

The Hard-wired Safety Interlock System (HSIS) is the safety system that determines when the beam may be turned on. It enables the RF drive signals that actually turn the beam on and off (the RF amplifiers generate the electric field that accelerates the particles). The HSIS, using hard-wired relay logic, merges computer-generated commands from the CMC (via the IOS) and the PLC with safety interlocks and safety-related signals under operators' control (doors, keys, emergency OFF buttons, etc.).

The HSIS is a chain of switches which must all be closed in order to enable the RF drive. If any switch in the chain opens, the beam will turn off and cannot be turned on. For example, this is the usual mechanism for turning off the beam at the end of a treatment.

4.2.5 Cyclotron control console and tuning modules (TUM)

The cyclotron control console, or desk, houses two VT100 compatible CRT terminals, the Cyclotron Operations terminal and the Cyclotron Status terminal, which are connected to

the CMC through RS-232 lines. The buttons, switches, dials, lamps, meters and other devices in the console are variously connected to the IOS, PLC, and HSIS; some are connected directly to the controlled equipment.

A VT340 terminal may be used in place of any VT100 (at the cyclotron console, or anywhere else a VT100 is used). The VT340 is a color terminal that is compatible with the VT100. Exactly the same software drives both terminals, so it is not necessary to provide the system with configuration information to tell it which types of terminals are in use. Some displays look better on the VT340.

The console also houses the Tuning Modules (TUM), a set of six dials and displays that can be connected to various devices under software control. Operators use the tuning modules to adjust and observe analog control parameters. TUM contains its own Z-80 processor and programs in ROM and communicates with CMC through an RS-232 serial line. The TUM includes processor, memory and communication cards from Scandia Metric AB of Sweden. Its proprietary components and control programs are by the cyclotron vendor.

4.3 What the cyclotron control programs do

This section briefly describes the functions performed by the cyclotron control programs that run on the CMC. Essentially, the cyclotron control programs mediate between the cyclotron operator and the equipment that actually controls the cyclotron.

The programs identify each user (via a login procedure) to determine that they are authorized to use the cyclotron control console.

The programs allow the operators to control most analog parameters: to turn them on and off, and set and monitor their values. Typical analog control parameters are magnet currents, mechanical positions, voltages etc. The console provides dedicated indicators and controls for only a few essential parameters; it is impractical to place over one hundred parameters within the view or reach of the operator. Instead, they are selected by the operator, singly or in groups, and assigned to the two CRT's or the six tuning modules. The programs rapidly and repeatedly scan these input and output devices to provide operators the impression of continuous control. When the operator turns parameters on or off, the programs change, or ramp, the outputs to their new value as gradually as needed to prevent damage from transients. The programs also scale and transform analog quantities between the engineering units used by the operators (e.g. cubic centimeters per minute, kilovolts, amps) and the binary bit patterns on the ADC and DAC cards. The programs also allow the operator to assign upper and lower boundaries, or windows, of permissible values for each parameter, outside of which a fault is considered to exist; programs automatically turn

off the beam if any parameter value exceeds its window.

In addition to mediating almost all analog controls, the CMC programs also mediate many on/off controls and other digital signals. They encode and decode digital bit patterns to implement commands or obtain status indications that are meaningful to the operators. Many digital signals are mediated only indirectly by the CMC, which is also connected to the PLC and HSIS. This enables the CMC programs to monitor or moderate functions that have been relegated to those simpler processors.

By combining many analog and digital control operations in proper sequence, the CMC programs perform automatic *startup, shutdown, and sequencing.* This enables the operator to conveniently control collections of related components that are usually operated as a group (e.g. all of the power supplies and other components associated with a particular beamline), often by pressing a single button.

None of the other processors besides the CMC has extensive human interface capabilities. The HSIS and PLC are connected to a few dedicated control panel components; none of the processors besides the CMC can drive a video display. Therefore, the CMC handles almost all of the user interface, even for functions that are largely handled by other processors. As a result, many digital signals that are used by the PLC and HSIS are also input to the CMC, so they can be displayed on the CRT displays.

The CMC programs provide access to a database of machine settings, or tunings (which is stored on a disk on the control system support host). Operators may store the current machine setting in the database, or load the machine with a stored setting from the database. In addition, a single setting is saved in volatile memory in the CMC to provide a sort of "undo" capability, and is periodically checkpointed to a disk file to provide some crash recovery.

The programs also perform some *fault detection*, turning off the beam or shutting down subsystems when faults occur, and *interlocking*, preventing certain subsystems from being turned on when it might be unsafe to do so. Patient and operator safety shutdown systems and interlocks are nonprogrammable hard-wired controls (the HSIS and MSIS) but these are backed up, and are displayed to the operators, through the CMC programs. Some internal equipment protection is provided by the PLC, but much is provided solely by the CMC software. The cyclotron delivers up to 3.5 kilowatts in a beam about 5 millimeters in diameter; a poorly tuned beam may strike components within the cyclotron or along the beamline and can burn through tens of thousands of dollars worth of fragile components in seconds.

Finally, the control programs *log* operating data and error reports to disk files on the control system support host for subsequent record keeping, performance analysis, and troubleshoot-

ing.

Primarily for safety and equipment protection reasons some equipment is not controlled by any of the control processors, but are only monitored by them. These systems are:

- Cyclotron vault door
- Water cooling system
- Ion source current
- Harmonic coil azimuths
- Main probe position

Other systems which are not affected by or monitored by the processors are:

- Beam profile monitors
- TV and intercom systems
- Room radiation monitors
- Patient alignment lasers, beam defining lamp and SSD projector

4.4 Treatment control systems

One of the treatment control systems is diagrammed in Fig. 4.2 (this shows the Isocentric Treatment Room, where the beam is steered through a rotating gantry that can move the neutron beam source around the patient, and the neutron beam is shaped by a leaf collimator. The Fixed Beam Treatment Room is similar but there is no rotating gantry and no leaf collimator). The CMC, IOS, PLC and HSIS shown in Fig. 4.1 also serve both treatment control systems; some signals in the treatment consoles and treatment rooms are connected to those components.

Each treatment unit is actually run by a set of controllers: a treatment motion controller (TMC), a dose monitor controller (DMC), and in the isocentric room only, a leaf collimator controller (LCC). Each of these is a semi-autonomous control system with its own sensors, actuators and processor. The CMC sends commands and data and requests data from each

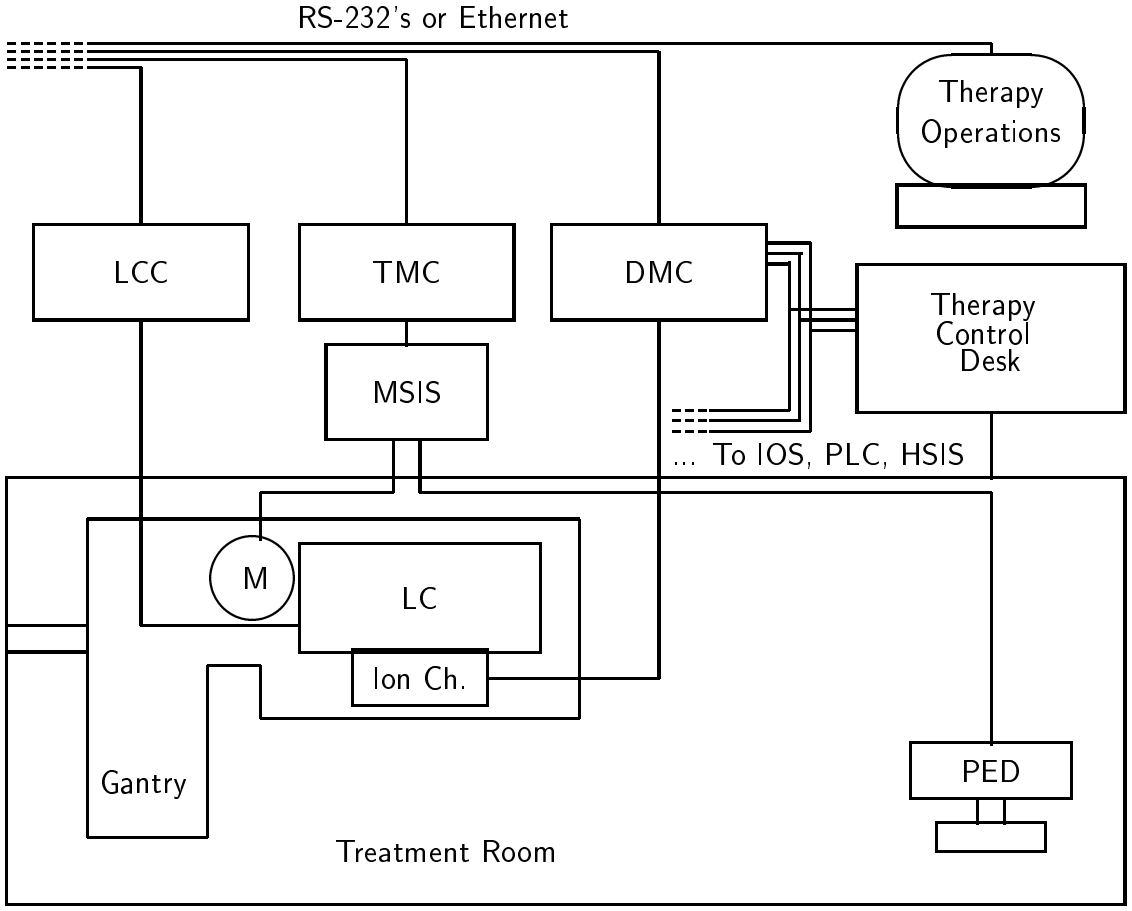

Figure 4.2: Treatment control system, isocentric treatment room

controller by sending strings of ASCII text; the controller carries out commands, reports its status and returns data as ASCII text.

Usually the programs running on the CMC generate and send the control strings without the overt particpation of the operators, but there is a mode of operation (also implemented by software on the CMC) that allows an engineer at a diagnostic terminal to type control strings and observe the status and data that are returned. This mode is used for maintenance and troubleshooting purposes only.

4.4.1 Therapy control console

For each treatment room there is a therapy control console. Each console houses a single VT100-compatible (e.g. VT340) CRT, the Therapy Operations terminal, which is connected to CMC through a serial line. There is a panel of displays driven by the DMC that show the dose rate, accumulated dose, and other quantities. Buttons, lamps, keys and other components on the therapy consoles are connected to the IOS, PLC and HSIS.

4.4.2 Leaf Collimator Controller (LCC)

In the Isocentric Treatment Room, the radiation beam is shaped by 40 independently moveable leaves in a leaf collimator (LC) which are driven by the leaf collimator controller (LCC). Inputs from the CMC to the LCC specify the position that each leaf must achieve to shape the field. The LCC drives the leaf motors, reads the leaf position sensing potentiometers, detects leaf collisions etc.

In the Fixed Beam Room, the beam is shaped instead by a installing a fixed collimator. No LCC is required.

4.4.3 Treatment Motion Controller (TMC), control pedestal and Motion Safety Interlock System (MSIS)

External motions include gantry, collimator, and couch turntable rotation and all couch linear motions. These are controlled either by a manually-operated control pedestal (PED) in the treatment room or under computer control by the treatment motion controller (TMC). Under manual control, the technologist presses a button on the pedestal to indicate which motion is to be set, and then adjusts a variable speed control on the hand pendant while

holding down a "dead-man" enable switch, all while observing the motion and reading the position on the wall display.

For computer controlled external motions, the technologist presses an enable button on the control pedestal (or, for arcs, on the treatment console) while the TMC drives the motion. For computer controlled motions, the CMC loads the TMC with the position and velocity to be achieved and the TMC actually drives the motor and continuously senses the position. The CMC can also command the TMC to report the position and velocity of all external motions. This can always be done, whether the motions are under manual or computer control.

Internal motions include flattening filter selection and wedge selection and orientation. These may be controlled manually from the pedestal or automatically by the TMC. Since there is no collision hazard, it is not necessary to press the pedestal enable button and these may be automatically set while the technologist is outside the room.

The hard-wired Motion Safety Interlock System (MSIS) mediates all potentially hazardous external motions (indicated in Fig. 4.2 by motor M), using relay logic to select either TMC or PED inputs and merge them with signals from safety interlocks, collision detectors, keys, "dead-man" enable switches, etc. The MSIS is a chain of switches, where each switch must be closed to permit motion. If any switch in the interlock chain in the MSIS opens, all motions are disabled.

4.4.4 Dose Monitor Controller (DMC)

The Dose Monitor Controller (DMC) is responsible for measuring the delivered dose and shutting off the beam when the prescribed dose has been delivered. The DMC includes two independent dose terminator (DT) channels. Each DT measures the ionization current created by radiation passing through a different ion chamber. Each DT integrates its current to obtain accumulated dose. The DMC also includes a safety backup timer which can turn off the beam if the dose monitors fail to do so. In addition to turning off the beam, the DMC also mediates some beam steering and dose rate controls.

Because the DMC is normally responsible for turning off the beam and is always responsible for recording the dose that is actually delivered, it is the programmable element in the system which is most critically concerned with patient safety.

The CMC loads the DMC with calibration data, some of which change from day to day (e.g. the conversion constants from ion chamber current to dose) and others which are changed less frequently (e.g. the gain constants in the beam steering and dose rate servos). The

CMC also loads the DMC with the prescribed dose and backup time for each treatment field. The CMC can command the DMC to report the current values of all the calibration constants, and to report the accumulated dose and elapsed time while a treatment is in progress.

The DMC enables each treatment by closing a relay (called the dosimetry relay or DMC relay) in the HSIS after the CMC has loaded the prescribed dose. The DMC ends each normal successful treatment by opening the relay when the prescribed dose has been delivered.

The DMC can interrupt the treatment if it detects some fault conditions, such as an excessive difference between the two DT channels, loss of ion chamber voltage, loss of internal control voltages, etc.

The DMC also drives a panel of displays at the treatment console so the therapy technologist can observe the dose rate, accumulated dose, elapsed time and other pertinent dosimetry information during treatments.

Mechanical dose counters mounted on the DMC chassis record the accumulated dose from each treatment. In the event of a power failure, these counters display the dose that had been accumulated when the power failed.

4.4.5 Moving floor control

There are manual and automatic controls for the moving floor in the isocentric treatment room. Technologists may manually command the floor to open, then move the gantry, and then close the floor. In automatic mode, the floor opens automatically to admit the gantry as it swings under the couch, and then the floor closely tracks the gantry as it moves.

4.5 What the treatment control programs do

This section briefly describes the functions performed by the treatment control programs that run on the CMC. Essentially, the treatment control programs mediate between the facility users and the auxiliary controllers and other treatment equipment. (The users are usually the treatment technologists, but also include the physicists and engineers).

The programs identify each user (via a login procedure) to determine that they are authorized to use the treatment console, and to determine which operations are made available (e.g. some calibration and experiment modes may not be available to all users).

The programs implement different *operating modes*, depending on which beamlines are in use, and whether treatments or physics activities are in progress.

The programs provide access to a database of patients and their treatment field specifications (which is actually stored on a disk on the control system support host). There is also a database of experimental fields for physicists and engineers. Users select field specifications from this database. The programs also provide facilities to create, modify and store field specifications, in case no satisfactory one is found in the database.

The programs may command the LCC and TMC to automatically set up internal motions (those that present no collision hazards). In a future version, the programs may command the TMC to set up external motions semi-automatically, under operator supervision.

The programs poll the LCC and TMC to determine that the treatment setup (field shape and motions) conforms to the prescribed specifications.

The programs load the DMC with calibration data that may change from day to day (e.g. the conversion constants from ion chamber current to dose) or less frequently (e.g the gain constants in the beam steering and dose rate servos).

The programs load the DMC with the prescribed dose and backup time for each treatment.

The programs poll the the DMC to report the current values of all the calibration constants, and to report the accumulated dose and elapsed time while a treatment is in progress.

None of the other processors besides the CMC has extensive human interface capabilities. The DMC drives some dedicated panel displays; the TMC and LCC have none, and none of the processors besides the CMC can drive a video display. Therefore, the CMC handles almost all of the user interface, even for functions that are largely handled by other processors. Its programs display the status of different treatment subsystems. At the operator's command, they display different subsets of the parameters and status data on the CRT screen, performing scaling, transformation and decoding as needed to create a meaningful display.

The programs enable each treatment by closing a relay (called the check and confirm relay or C&C relay) in the HSIS after they have confirmed that the treatment has been set up properly and the prescribed dose has been loaded correctly. The programs can end a treatment by opening the check and confirm relay if a fault is detected (e.g. by polling the LCC, TMC or DMC).

The programs record each attempt to treat a patient (in a file on the control system support host).

Primarily for safety and equipment protection reasons some equipment is not controlled by any of the control processors, but are only monitored by them. These systems are:

- Treatment room door

Other systems which are not affected by or monitored by the processors are:

- TV and intercom systems
- Room radiation monitors
- Patient alignment lasers, beam defining lamp and SSD projector

4.6 Facility history and configuration changes

The facility was installed in 1984, and included a computer control system provided by the cyclotron vendor. This document describes a new, successor control system now being developed by the University; it does *not* describe the system initially provided by the vendor. In the few places in this document where the two systems are compared, it is made clear when the features of the initial system are being described.

To help convey the size and complexity of the system, several characteristics of the control system configuration that was initially installed by the vendor in 1984 are shown in Table 4.1. The programs provided by the vendor comprise over 60,000 lines of FORTRAN and several thousand lines of assembler code.

The new system includes new computer hardware and software and some changes to the peripheral control hardware. The new system is motivated by requirements to make the system easier and quicker to use and easier to maintain. It is also required to accommodate future hardware and software modifications.

The overall hardware organization of the control system, including the set of processors and the allocation of functions to the processors, is substantially the same in the new system as in the initial configuration. However, the hardware used for some of the processors, and the method used to interconnect them, differs.

The software in the new system is so different from the old system that there is no point in attempting a comparison. The new software provides a different user interface and is based on a different design. It does not include any code from the old system.

Name	*Processor*	*Language*	*Program Size*	*Digital In*	*Digital Out*	*Analog In*	*Analog Out*	*Serial I/O*
Main Computer and I/O System (CMC/IOS)	DEC PDP 11/23	RSX-11/M FORTRAN	60,000 lines	568	230	128	77	9
Programmable Controller (PLC)	Gould Modicon 484	ladder logic	112 networks	177	159	—	—	—
Hard-wired Safety Interlock System (HSIS)	relays	hard wiring	31 relays	52	5	—	—	—
Tuning Module Controller (TUM)	Z-80	assembler	8 KB object	16	16	—	—	1
Treatment Motion Controller (TMC)	Z-80	assembler	8 KB object	40	46	7	8	1
Motion Safety Interlock System (MSIS)	relays	hard wiring	70 relays	88	96	—	—	—
Dose Monitor Controller (DMC)	Z-80	assembler	13 KB object	24	49	1	3	1
Leaf Collimator Controller (LCC)	Z-80	assembler	8 KB object	23	81	40	—	1

Table 4.1: Control system processors, inputs and outputs: initial configuration

After we gain experience with the configuration described in this document, we anticipate making additional changes. Some possible future changes are described in chapter 8.

4.6.1 Hardware configuration changes

This section describes the differences between the hardware configuration delivered by the vendor, and the hardware used for the new control system. The control system configuration that was originally installed by the vendor is shown in Table 4.1.

In the initial system the main cyclotron control computer (CMC) was a DEC PDP-11/23. In the new system it is a DEC MicroVAX II.

In the initial system the CMC was connected to the IOS bus through a Z80-based input-output coprocessor. In the new system the CMC is connected to the IOS bus through a DEC DRV11 direct memory access (DMA) interface.

In the initial system the CMC communicated with the PLC through the IOS, using several digital I/O lines. In the new system the CMC communicates with the PLC through an RS-232 serial communications line.

In the initial system there were no diagnostic terminals. The new system adds three diagnostic terminal ports: one in the control room (which can also serve each treatment room), another in the power supply room, and a third connected to a modem.

In the initial system the TMC and LCC were each built from standard microprocessor, memory and communication cards from Scandia Metric AB of Sweden, as well as custom cards built by the cyclotron vendor. Each used a Z-80 microprocessor; programs were in ROM. These control programs were written by the cyclotron vendor. Each processor was connected to the CMC through an RS-232 serial communications line. In the new system, the Z-80 LCC and TMC processors are replaced with a second MicroVAX II processor that implements increased functionality that is upward-compatible with the present LCC and TMC. This MicroVAX is connected to the CMC via the Ethernet, rather than an RS-232 line. It also runs standalone software booted off the control system support host's disk via the Ethernet. The CMC and the auxiliary processors still communicate using text strings; the command sets in the new system are supersets of the initially installed commands.

In the initially installed configuration, the moving floor in the isocentric treatment room was not automatically controlled. In the new system there is an automatic mode, where the floor opens automatically to admit the gantry as it swings under the couch, and then closely tracks the gantry as it moves. This function is provided by a second PLC.

In the new system, some motion control functions that were initially provided by the Elven relay chassis that house the MSIS are allocated to the new PLC that also controls the moving floor.

In the initial configuration, flattening filter selection was not provided at the control pedestal.

Chapter 5

Performance goals and requirements

In this chapter we present the goals and performance requirements that the control system must meet. There are two kinds of requirements, those that pertain to generic or basic operation of the user interface, and those that pertain to the performance of the treatment equipment and the cyclotron machinery itself. Of course we only consider here the aspects of equipment performance that are affected by the control system design.

5.1 Control system design goals

The control system design is motivated by these goals:

- Ease of use.

  The system should be easy to learn and understand and operation should be convenient and quick.

- High reliability and safety.

  Safety requirements should be explicit. The design should be simple and logical, so that safety characteristics can be confirmed and the implementation can be straightforward and relatively bug-free and systematic tests can be designed and performed before the system enters routine service.

- Maintability and adaptability.

It should be easy to change functions and incorporate new ones. It should be possible to perform many modifications *without doing any additional programming.* Most operating parameters should be stored in tables that can be maintained by the engineering staff, not the developers.

Some of the changes that we anticipate are described in chapter 8.

5.2 Terminal response

5.2.1 Typing

Single characters typed in response to prompts should be echoed to the terminal with no perceptible delay. In particular, if the user invokes the terminal's autorepeat function by holding down a key, echoing should stop as soon as the user releases the key. This sets the upper limit on the single-character response time to be the period of the autorepeat function, which is about 30 msec (about 30 characters per second).

5.2.2 Menu selections

Many operations involve the user moving a highlight around an on-screen menu by pressing the terminal arrow keys. It should be possible to move such a highlight rapidly, even if the items occupy a lot of screen space. (All characters whose highlighting turns on or off must be retransmitted to the terminal, so the speed is related to the size of the highlighted items).

Five moves per second is a reasonable target for the largest highlighted items in the system (for example, menu items like saved settings or treatment fields which include a lot of descriptive text). Smaller items (as in the SET PARAMETERS operation) should be proportionately faster.

5.2.3 Updating terminal displays

Much of the screen may need to be updated when the user selects a keypad operation. This should be limited by the terminal line speed, not by program execution speed. For example dense displays like SET PARAMETERS involve rewriting almost every character from lines 3 through 20, or 1440 characters. In addition, there may be many terminal control (escape)

sequences of from 2 to 8 characters. A terminal speed of 9600 baud is about 1000 characters per second, so the display update should be complete within about two seconds after the operator presses the key.

5.3 Equipment performance

The following lists the performance requirements for some high level functions of the system, e.g., treatment operations.

1. Computer system startup and shutdown – Startup should complete within five minutes of the operator typing the "boot" command. Shutdown should be complete within one minute of operator request.

2. Therapy operations – Includes operation on Isocentric Treatment Unit and Fixed Treatment Unit[1]. The main requirement is that the time required to perform a treatment should not be determined by the computer control system; instead it should be determined by the patient, the number and complexity of the fields, and physical limitations of the controlled equipment such as the dose rate and the time required to open and close doors. Staff and equipment should never be required to remain idle, waiting for the computer to complete some operation (such as "check and confirm").

   In addition to this general guidance there are some specific requirements:

   - There should be no built-in upper limits, imposed by the control system, on the number of patients under treatment, or the number of patients that may be treated in a day.
   - The leaf collimator will achieve the required state for any leaf positions, within 15 seconds from operator command. Leaf positions will remain fixed when the beam is on.
   - It will be possible to perform arc treatments in which the gantry rotates while the beam is on, controlled from the treatment console in the control room. During arc therapy collimator positions, beam current and gantry speed will remain constant.
   - The check and confirm system will lock out unsafe treatment setups for patients and treatment setups that do not match a prespecified prescription, unless explicitly overridden.
   - Automated record keeping for all treatments will include all treatment setup parameters, beam on time and monitor units delivered.

[1] Same as isocentric except for fixed collimators, fixed floor position and fixed horizontal beam position

- Couch motion, collimator motion and collimator leaf motion will not be possible when the beam is on.
- Choice of the X-ray unit or the neutron beam for port films will be supported.

3. Beam switching between beam lines should take no longer than a period of 60 seconds.

5.4 General guidelines

The speed of most operations should be limited by the controlled devices, not the computer. For example, power supplies should be ramped at a rate limited by the permitted rate of change of voltage of the supply, not the control system response time.

Operations should be performed concurrently whenever possible. For example, if several power supplies are being ramped, this should happen simultaneously rather than in sequence.

The system should not waste time performing redundant or obviously futile operations. For example if the operator requests a transition from Standby level 1 (SB1) to Standby level 2 (SB2), the system should not spend several minutes attempting to do this, only to announce that some readily apparent condition prevents its completion. Instead, the system should check all relevant conditions at once and immediately report any which would prevent the operation from completing. If some operation is composed of several simpler operations, if possible the system should check whether any of these operations are already complete and should not perform them again if they are.

The system must be robust regarding delays and failed operations. If some operation is delayed or suspended for any reason (e.g. delay or fault in the controlled equipment, software fault in one process), this must not cause appreciable delay in other operations which are not part of the same interdependent sequence *from the operator's point of view.* In particular, if some operation at one terminal is waiting for input from the user or is waiting for some machine operation to complete, this must not delay functionally independent operations in progress within the machine or at other terminals.

At each terminal and console, the system must never wait indefinitely or "hang" if some machine operation fails to complete. Each operation must be associated with a time-out period; if the operation does not complete within its time-out period, the failure is indicated and the operator is provided the options of attempting recovery (perhaps by retrying the operation) or alternatively, of proceeding with something else.

Response times and other performance requirements dictated by safety or equipment pro-

tection considerations are discussed in Part III.

Chapter 6

Safety and equipment protection

The CNTS facility has to ensure proper operation with regard to

- Patient Safety
- Personnel Safety (e.g. technologists, physicists, engineers, etc.)
- Equipment Protection

This chapter describes our approaches to ensuring safety and equipment protection.

The most important potential hazards to people are radiation exposure from the neutron beam, collisions involving the gantry and/or couch, and falls into the pit beneath the moving floor. The most important potential hazard to the cyclotron and beamline components is damage from a poorly steered particle beam. In the rest of this document we emphasize consideration of these hazards. The facility does present other hazards similar to those found in many industrial environments including electrical shock, injury from moving components such as pumps, etc. Protection against these other hazards does not depend mainly on the control system, but is largely provided by other mechanisms, including safety covers, guard rails, shorting devices in high voltage equipment and the like.

We depend on human operators to manually control and observe most potentially hazardous activies. Several safety mechanisms are controlled by the operators. Interlocks implement automatic safety controls. The most critical and essential interlocks, including those devoted to human safety, are implemented in non-programmable hardware. Equipment protection interlocks are implemented in hardware, or on the least complex processor that can provide the needed functions.

The following sections describe the approach in greater detail.

6.1 Human operators and manual controls

Safe operation depends most of all on alert, trained operators who exercise common sense and reasonable caution. There are two main reasons for this. First, most potentially hazardous actions are not performed fully automatically. Instead, they are performed under control of operators who must closely observe operations in progress to ensure they proceed safely. Second, it is often possible to bypass the control system and defeat most of its safety features locally at the controlled equipment; in fact this is necessary to perform some maintenance operations and experimental procedures.

The control system safety features are not intended to limit operator's activities or to prevent them from doing careless things. Instead, they are designed to protect against machine faults or internal machine conditions which the operators cannot directly observe (e.g., analog parameters outside windows), to which they could not respond sufficiently rapidly (e.g. prescribed dose delivered), and to ensure that large collections of interacting components are not used in unsafe combinations (e.g., those involved in beamline selection).

Usually, the inherent physical limitations of the equipment limit the rate at which impending hazards may develop, so it is reasonable to depend on manual controls. For example, the gantry rotates at 3 degrees per second at its fastest, so the nose of the collimator (which could collide with the patient or couch) moves at about 2 cm per second, while the outside of the gantry (which could collide with the floor) moves at about 15 cm per second. The moving floor itself moves up to 30 cm per second. They are not difficult to guide under visual control, and if a poorly planned field or erroneous control signal threatens imminent collision, the therapy technologist has sufficient time to respond.

Similarly, when the proton beam current is at its maximum value of about 70 microamps, the resulting neutron beam dose rate is only about 60 cGy per minute, so it takes about 2 minutes to deliver a typical daily therapeutic dose of 120 cGy. If the control system failed to stop the treatment when the prescribed dose was delivered, the therapy technologist could manually stop the treatment before much damage was done.

Similar physical limitations make the cyclotron operator's work feasible. There are many automatic equipment protection mechanisms that turn off the beam if the system passes outside its normal operating limits, but it is not feasible to automatically protect against all situations which could damage the equipment. Instead, operators use standard procedures to help avoid hazardous modes of operation. For example, there is a standard procedure for tuning the beam and increasing beam current that minimizes the likelihood of damaging

cyclotron and beamline components.

Consistent with this philosophy, much of the engineering effort for the control system is devoted to providing mechanisms that enable the operators to effectively monitor and control safety. Three such mechanisms are the emergency power off switches, lockouts, and displays.

6.1.1 Emergency off switches

The facility is equipped with many emergency off switches so operators can interrupt hazardous conditions even if the control system completely fails (for example, if a runaway gantry threatens imminent collision).

There are two sets of emergency off switches. The first set includes the switches in the cyclotron vault and power supply room, and on the cyclotron control console. These turn off the main power to the facility. However, the power remains on for the control processors, console terminals, intercoms, TV's, room radiation monitors, treatment room and vault doors and all lights.

The second set includes the switches in the treatment rooms and on the therapy control consoles. These remove power from the gantry components and moving floor, so they can be used to stop external motions. They also turn off the beam by removing power from the ion source (this is different from the usual mechanism, which is to open a switch in the HSIS). However, these switches do not turn off the facility's main power.

After an emergency power-off switch has been pressed, it must be reset with a key, and the HSIS must also be reset in order to turn on the beam.

6.1.2 Lockouts

Lockouts are mechanisms that disable actions except while they are explicitly enabled by the operator. The enable switch on the motion controller hand pendant and the console keyswitches are examples. They prevent unintended actions and can also be used to stop dangerous actions. For example, most external therapy motions are disabled unless someone actively exerts pressure on the enable switch (which is colorfully termed the *dead man switch*). This requires the operator to be present where he or she can observe that the motions are safe.

Similarly, the console keyswitches can prevent the beam being turned on while people are

working nearby. Someone who intends to work with potentially hazardous equipment may simply turn the keyswitch off and take the key along with them. Many of the power supplies have LOCAL/REMOTE selector switches. In LOCAL, the supply cannot be activated by the control computer.

6.1.3 Displays

Treatment setup may be partially automated and based on computer generated information (the output of a treatment planning system, for example). It is possible that this information might be incorrect or become corrupted, either due to software errors or to human misunderstandings and mixups. We partially guard against this possibility by displaying automated treatment settings to the technologist on the therapy operations terminal. An important duty of the therapy technologist is to confirm that the information visible on the screen, (and visible by viewing the therapy apparatus itself), conforms to the written prescription and to the identity of the patient.

Console displays usually show two values for each setting: the *set* value which was demanded by the control system and the *read* value actually achieved, which is read back from sensors. It is never necessary to assume that commanded settings have been achieved; instead, the actual read value is available.

6.2 Local safety mechanisms

This and following sections describe built-in safety mechanisms that are not under direct operator control.

Many safety mechanisms are not part of the control system at all. All safety mechanisms which deal with just one specific piece of equipment are provided locally, at the equipment level. These include mechanical features such as safety covers, guard rails, shorting devices in high voltage equipment and the like.

6.3 Interlocks

Interlocks are devices which detect hazardous conditions and disable or interrupt hazardous actions. In contrast to lockouts and displays, which rely upon the active participation of

human operators, interlocks work automatically and operators usually learn of their actions only after they are activated.

Most interlocks are latching – they remain set even when the triggering condition is corrected, and must be explicitly cleared by the operator. This enforces a reasonably systematic recovery procedure, and is also useful for troubleshooting; when the beam suddenly goes off because some transient event sets an interlock, it is often possible to determine which event caused the interruption by noting which interlocks have been set. Exceptions to the latching attribute include safety mechanisms under direct operator control such as doors and keys, and some transient events which are cleared automatically. These transient events include for example, sparks that trigger interlocks in the RF subsystem, which temporarily turn off the RF drive, but for which recovery is automatic.

Information needed to implement interlocks is always derived from sensors. The system does not assume that commanded settings have been achieved; it does not infer what the inputs should be on the basis of past outputs.

Interlocks critical to human safety are implemented by nonprogrammable hardwired devices. Equipment protection interlocks are hardwired or delegated to computers. These guidelines are expressed in greater detail in the following sections.

6.3.1 Local Interlocks

There are sensors that detect unsafe conditions that will directly turn off the particular equipment, including flow and temperature switches, microswitches on access panels, touch sensitive strips on the shielding doors, ground fault protection devices and overcurrent protection. Indication of active safety devices is local but most electrical interlocks are also connected to inputs to the control computer and can be displayed at the operators' consoles.

6.3.2 Hard-wired interlocks

The two most potentially hazardous actions in the facility are turning on the beam and initiating external motions. Both of these are protected by non-programmable hard-wired interlock systems.

All safety related interlocks for the "Beam On" function are achieved through the Hardwired Safety Interlock System (HSIS). This consists of a 24 V DC power source connected through a series of relay and microswitch contactors (the "hardwired safety trace") to sets of two relays which allow beam to be run within the Cyclotron Vault, into the Isocentric Room or

into the Fixed Beam Room. Some of the control system processors control relays in this trace. For example, the CMC can open a relay if analog parameters are not within their windows.

The HSIS is so critical that several design rules are observed:

- All interlocks in the hard wired trace which are related to patient or personnel safety have a back-up, either an independent other interlock or in the case of some relays through a second relay. This means that at least two independent faults must occur to create a potentially dangerous situation.
- Because the safety requirements are different for different beam delivery modes, the selection of these modes is also hardwired (beamline selection). The CMC just monitors this selection.

External motions are similarly protected by a Motion Safety Interlock System comprising the hard-wired relay traces in the Elven relay chassis. The design rule is:

- The controls for the treatment unit external motions require two steps for operation, e.g., a hand pendant or rocker switch simultaneously with a deadman switch, or (when remote) simultaneous activation of an analog motion speed signal and a motion enable relay.

6.3.3 PLC interlocks

Some of the more elaborate equipment safety interlocks are handled by the PLC, including those for the vacuum system and beam line components such as Faraday Cups and Beam Plugs.

6.3.4 CMC interlocks

The CMC programs implement most equipment protection interlocks that require monitoring analog parameters; these are *software* interlocks. In particular, the CMC compares the value of most analog parameters to upper and lower limits, or windows, that are set by the operator. If any parameter exceeds its limits, the CMC turns off and disables the beam by opening a relay in the HSIS (this is accomplished indirectly via the PLC).

The CMC and PLC also monitor and back up the hard-wired safety systems. They monitor the status of the various interlocks through auxiliary contacts on microswitches and relays and check for inconsistencies with the hardwired trace.

6.3.5 Other processors

The HSIS and MSIS work despite the failure of any programmable processor. Where safety is affected if a processor fails, a *watchdog timer* is used to monitor the processor's operation. A watchdog timer is a nonprogrammable hardware circuit external to a processor which can usually detect whether a processor has "crashed" (suffered catastrophic hardware or software failure) or halted. A relay in the HSIS or MSIS opens if a watchdog detects a processor failure.

Chapter 7

Security

The cyclotron main control computer (CMC) is connected by a local area network to a computer called the *control system support host*, or *host*, which houses the control system software and data files. The control system support host serves a department of about 90 people and may also be connected to the hospital and University of Washington wide area networks, which are in turn connected to the Internet, a "network of networks" that is international in scope and includes NSFNET, BITnet, the Defense Data Network and others.

This environment requires that appropriate measures are in place to prevent unauthorized access to the control system while running or to the system software and data files. The system is protected by using a number of standard strategies, described in the following sections.

7.1 Access to the control system support host

7.1.1 Potential hazards

It is important to consider what the consequences could be if someone obtained unauthorized access to the control system support host. It appears that the greatest potential risks involve corruption or loss of files such as treatment specification files, treatment records, the executable software for the CMC, etc. Such losses can occur for reasons other than security breaches (e.g. disk head crashes) , and reasonable precautions (e.g. frequent backups) must be in place to recover from them.

The greatest potential danger to patients would occur if treatment specification files were altered sufficiently to affect the treatment, but not sufficiently corrupted to be unreadable, or to be caught by error-checking code (such as range limits) in the control software. An example of such an error might be deletion of a wedge. In that case we must depend on the therapy technologists to confirm that the information visible on the screen, (and visible by viewing the therapy apparatus itself), conforms to the written prescription for the patient. This is always an important part of the therapy technologist's job, and is necessary to guard against human misunderstandings and software errors as well as sabotage.

7.1.2 File protection

The control system support host is a Digital Equipment Corp. VAX running DEC's VMS operating system, which provides standard file access controls that prevent any user from reading or altering programs or data unless specifically permitted to do so. These same protections apply to network access attempts as well as to access attempts by local users. In addition, access from other systems on the network is controlled by another layer of general access controls on the host.

These controls apply to accidental access attempts as well as deliberate ones.

Files are owned by one User Identification Code (UIC), and the files are protected from access by anyone but the owner (and the system manager) by the VMS file protection mask (S,O:RWE,G,W). Any other user needing access to a file is given explicit access by an ACL (Access Control List, another VMS facility) entry for that file. The owner is responsible for maintaining the ACL entries and deleting them when access is no longer needed.

7.1.3 Network access to the control system support host

Access to the control system support host through the network requires that the user on another network node provide a username (with password) that has been authorized for the host. That username (and its UIC) then determine whether the network user may log in, or has access to any files. So the same protection that is used for host computer users applies to any access from the (world-wide) network.

7.2 Access to the control system computer

It is important to consider what the consequences could be if event someone obtained unauthorized access to the control system main computer (the CMC). In the event that this occurs, the other control processors and especially the hard-wired safety systems are supposed to prevent patient injury or serious equipment damage. The greatest potential hazards would occur if the system were affected sufficiently to alter operations, but not sufficiently to trip interlocks or be caught by error-checking code (such as range limits) in the control software. In such cases we must depend on the cyclotron operators and therapy technologists to observe the system and take corrective action. This is an important part of their duties, and is a necessary precaution against many kinds of problems in addition to security breaches.

7.2.1 Network configuration

The CMC is connected to the Radiation Oncology Department Ethernet, which is connected to the control system support host and also to several other computers located in the Radiation Oncology Department. The department Ethernet is connected to a computer called the *router*. The router is the only connection between the department Ethernet and wide area networks, and makes it possible for department computers to communicate with computers elsewhere. The router runs communication software that complies with the Internet Protocol (IP); IP packets are the only packets that can pass the router. In particular, control system packets are not IP packets and do not pass through the router to the outside world. The only computers that can normally be reached from the outside world through the router are those that have been assigned IP addresses.[1] The control system support host may be assigned an IP address, but the CMC does not have one. This greatly limits the number of systems from which the CMC might be accessed.

7.2.2 Control system software load from a remote host

When the cyclotron control system computer is not running it is not possible to boot it remotely from some other system. It can only be booted by an operator typing a boot command at its own console in the power supply room. This is set by a configuration switch on the control system computer CPU circuit board. Then, the operator must remain present to type responses to several prompts printed by the control programs as they start up.

[1] DECNET packets that are encapsulated within IP packets can also pass the router. These packets are still subject to all DECNET access controls.

When the operator types the boot command at the system console, the control system computer broadcasts a message to all computers on the same Ethernet. In this message, the control computer identifies itself and requests that some other system on the network download software into it. The only computers that can respond are those which have been prepared in advance by specially configuring their own network database. Normally, only the host system that contains the control system software image file will respond, since it is the only one with an entry for the control system computer's Ethernet hardware address. If another computer reachable on the department Ethernet responds to the load request because it contains the control system computer's Ethernet address (by mistake or by malicious intent) the load may proceed from that computer, but the messages on the console will most likely indicate that the wrong software is starting up. The operator may then abort the load by pressing the halt button on the control system computer.

Only computer systems that are connected to the Radiation Oncology Department Ethernet can respond to a downline load request from the cyclotron control system computer. This severely limits the range of candidates for such "false load" responses.

While the control system software is running, it is not possible for a remote host system to reload the software or replace it with a different system software load. The configuration input to the system builder utility used to create the control software specifies that remote triggering is disabled, thus preventing remote reloading of software.

7.2.3 Access to the running system from a remote host

Since there are no file structured devices connected to the CMC, remote file access is not possible. However we need to consider the possibility that a remote system user could attempt to make connections to processes and devices on the CMC via the network.

There is no provision for "login" or creating an interactive process on the control system computer. The development system used to create the control software (DEC's VAXELN product) provides an optional command interpreter called ECL but this is explicitly disabled since it is not included in the running system and each terminal line has use of ECL disabled.

Other kinds of access are controlled by the usual DECNET access control, i.e., in order to connect to an VAXELN object in the cyclotron control system the remote user must provide access control information just as for any other DECNET link request. However, the Authorization Service is not included as part of the control system so no requests can be validated.

7.2.4 Other network problems

It is possible that there could be too much traffic on the department Ethernet when the CMC attempts to communicate with the control system support host (e.g. to load its software, read a treatment specification or write a log file entry). In that case the pending operation will take longer than usual and might even fail (by timing out). The control software must be written to handle such cases gracefully.

It is possible to imagine scenarios involving very resourceful and diabolical saboteurs: one might use special software at a PC or workstation that is connected to the department Ethernet to intercept and read our control system packets and then create counterfeit packets and inject them into the Ethernet; one might install an altered version of our control software (that produces the same startup messages and promps at the system console). Such scenarios presume severe failures of physical security, configuration control, and operational procedures on our part.

If it appears that network traffic or security breaches will create difficulties, many potential problems can be avoided by physically reconfiguring the network. We could connect the cyclotron control processors to their own physically separate Ethernet; this Ethernet would carry only control system packets and nothing else. This net would be connected at the control system support host to a separate Ethernet card, so that no packets pass between this net and other Ethernets.

7.2.5 Local access to the running control system

Cyclotron operators, physicists and treatment technologists run the facility from computer terminals and control consoles. The details are described elsewhere; here we describe the mechanisms that control access to the system.

These access controls are not intended to substitute for physical security, common sense and reasonable caution. It is not necessary to use a console to perform many potentially hazardous activities. Most facility components can be run locally without the assistance of the computer control system. Many safety features can be physically bypassed. The access controls are intended to provide some accountability through improved record keeping, and aid convenience by providing each user with those facilities that they need.

In order to use a computer terminal that provides control functions an operator must type an identification and password which are analogous to the username and password for access to a multiuser computer. Without this, most controls at the corresponding console are disabled. This provides some accountability since a computer log entry is made when each

user logs in or out. In addition, there is a keyswitch at each console. The keyswitch must be in place and turned on to enable the controls and the terminal keyboard (including the login and logout functions). An operator can prevent someone else taking over their console by taking the key along if they have to leave the control room.

There is a file on the control system support host, which contains information about each authorized user. This information includes not only their ID and password, but a list of authorizations that permit that user to perform certain classes of functions. For example, a technologist might have authorization to perform treatment operations but not cyclotron operations. A physicist might be able to perform treatment operations and also run the treatment operations console in experiment mode. A cyclotron operator might be able to perform cyclotron operations but not treatment operations. This is automatically handled by the control system software by looking up the authorization information for the current user at a given console.

Some diagnostic and maintenance functions will not be available to cyclotron operators and therapy technologists during normal operation. Also certain operations are not available to certain terminals. In particular, some operations are not available via the terminal port connected to the modem, since this would allow the system to be manipulated from offsite. The particular allocation of functions to the various terminals is described in forthcoming chapters.

Chapter 8

Subsets and changes

An important goal of this development project is to produce a control system that is easy to modify. This chapter describes some of the modifications that we anticipate. The design should permit these modifications to be accomplished by editing data in tables or by adding, removing or substituting a limited number of hardware or software components, rather than by performing extensive modifications on many different components.

Most of this document has described a single configuration, which we refer to here as the *standard configuration*. Two kinds of variations on the standard configuration are *subsets* and *changes*.

8.1 Subsets

It will be possible to build runnable versions of the control system that contain subsets of the functionality of the standard configuration. We anticipate building several successively larger subsets and running them for testing purposes during the development process.

Most such subset configurations will probably be used only briefly but we anticipate gaining considerable experience with one of them. It comprises the proton beam control system described in section 4.2 and pictured in Fig. 4.1. This subset provides only the functions needed to run the beam inside the cyclotron vault. It is sufficient to produce isotopes but not to perform treatments.

Running beam outside the vault into either treatment room requires that several non-

programmable (hard-wired) interlocks in the HSIS that are associated with treatments (including keys controlled by the therapy technologists) be clear. Therefore, this subset configuration does not have critical human safety implications (other than those involving the HSIS itself). However, it does have equipment protection implications.

8.2 Changes

We anticipate making future changes and additions to the standard configuration. Some likely possibilities are listed here, roughly in order of increasing difficulty.

8.2.1 Provide additional status displays

It may be useful for the operators to have more than one status display available, to help them monitor several subsystems at once.

8.2.2 Reconfigure the IOS

Assignment of signals to particular I/O cards and channels in the IOS may be changed, to accommodate additional controlled devices or improve signal routing. It should be possible to make such reconfigurations rapidly, to accommodate temporary unavailabilty of particular channels due to hardware faults. It should be possible to perform these reconfigurations by changing tables that can be maintained by the engineering staff, not the software developers; it should not be necessary to create new versions of the runnable control software.

8.2.3 Replace Z-80 control processors and RS-232 communication lines

The system as delivered by the vendor included several dedicated control processors based on Z-80 microprocessors connected to the main processor by RS-232 lines. We anticipate eventually replacing many of these processors, and using Ethernet rather than RS-232 for communication. The command set provided by each new processor will be a superset of the commands currently provided. The allocation of functions among the processors may change.

We currently plan that the initial version of the standard configuration will already include some of these replacements (see Chapter 4, section 4.6). Additional replacements may be

made later.

8.2.4 Replace terminals with graphic workstations

A graphic user interface including multiple display windows controlled by "mouse" may be provided by replacing the VT-100 style terminals with graphic workstations or X terminals. These would not be connected by RS-232 serial lines; instead, they would attach to the same Ethernet as used by the other control processors.

8.2.5 Replace the IOS with non-proprietary I/O busses

Some of the devices (e.g. power supplies) that are now controlled through the proprietary IOS may be replaced by other devices that can be operated through industry-standard busses, such as the IEEE-488 bus. The interface to the industry-standard bus would probably be provided by a special interface card that plugs directly into the bus of the CMC, rather than going through the IOS. Most industry standard busses allow many devices to be controlled from a single long cable. This should result in fewer components and simpler wiring than in the standard configuration. The newer bus-controlled devices usually provide some local processing capability, so the control software that runs on the CMC could become simpler and faster.

8.2.6 Modify the fixed-beam treatment apparatus

The treatment head in the fixed beam treatment room might be equipped with a leaf collimator. It might be modified to produce a different type of beam, for example one with a substantial low-energy neutron component to support boron neutron capture therapy (BNCT).

8.2.7 Permit motions to be controlled remotely

The system may be augmented to provide *conformal therapy*. The essential feature of conformal therapy is that the irradiated volume in the patient is made to conform closely to the target shape. In most implementations of conformal therapy, this is accomplished by treating the patient with a large number of complex fields. In order to make this practical,

many successive fields must be set up and administered automatically, without the technologists re-entering the treatment room between fields. This implies that potentially hazardous treatment motions must be controlled from the treatment console, or automatically by the computer, and not from the mobile pedestal and hand pendants.

Furthermore, it may even be desirable to run the beam on the target while the motion is in progress. The standard configuration provides something like this in arc therapy, in which only the gantry may rotate while the room is closed and the beam is on. In conformal therapy, potentially any motion (including leaf motion as well as external motions) might occur while the room is closed and the beam is on.

The capabilities required to support conformal therapy would place great demands upon control system correctness and safety, beyond what is provided by the standard configuration. Therefore, providing such capabilities would require extensive reconsideration which might well conclude that a substantially different design is warranted.

Bibliography

[1] Jonathan Jacky. Formal specifications for a clinical cyclotron control system. In Mark Moriconi, editor, *Proceedings of the ACM SIGSOFT International Workshop on Formal Methods in Software Development*, pages 45–54, Napa, California, USA, May 9–11 1990. (Also in *ACM Software Engineering Notes*, 15(4), Sept. 1990).

[2] Ruedi Risler, Jüri Eenmaa, Jonathan P. Jacky, Ira J. Kalet, Peter Wootton, and S. Lindbaeck. Installation of the cyclotron based clinical neutron therapy system in Seattle. In *Proceedings of the Tenth International Conference on Cyclotrons and their Applications*, pages 428–430, East Lansing, Michigan, May 1984. IEEE.

[3] Ruedi Risler, Stanley Brossard, Jüri Eenmaa, Ira J. Kalet, and Peter Wootton. Two years of operating experience with the Seattle clinical neutron therapy facility. In *Proceedings of the Eleventh International Conference on Cyclotrons and their Applications*, Tokyo, Japan, October 1986.

[4] R. Risler, S. Brossard, J. Eenmaa, J. Jacky, I. Kalet, A. Rask, and P. Wootton. Routine operation of the Seattle cyclotron facility. In *Proceedings of the Twelfth International Conference on Cyclotrons and Their Applications*, 1989.

[5] Peter Wootton and Jüri Eenmaa. Review and assessment of a clinical neutron therapy system: Equipment design and environment. Presented at the EORTC Heavy Particle Therapy Group Meeting and Microsymposium on Dosimetry for Fast Neutron Radiotherapy, May 1986. Douglas Cyclotron Center, Liverpool, United Kingdom.

[6] Kenneth A. Krohn, Jeanne M. Link, and John. H. Courter. Target system for production of PET radionuclides at the University of Washington. In *Proceedings of the 2nd Workshop on Targetry and Target Chemistry*, Heidelberg, FRG, September 1987.

[7] J Eenmaa, I. Kalet, R. Risler, and P. Wootton. The University of Washington clinical neutron therapy facility. *Medical Physics*, 9:620, 1982.

[8] J. Eenmaa and P. Wootton. Comparision of neutron beam treatment characteristics from p(50)Be with d(21)Be and 6 MV photons. presented at the 25th Annual Meeting of the American Association of Physicists in Medicine, New York, August 1983.

[9] J. Eenmaa, I. Kalet, R. Risler, and P. Wootton. Therapy neutron beam characteristics of the University of Washington clinical neutron therapy facility. presented at the 70th Scientific Assembly and Annual Meeting of the Radiological Society of North America, Washington, D.C., November 1984.

[10] George E. Laramore, M. Bauer, Thomas W. Griffin, F. J. Thomas, F. R. Hendrickson, Moshe H. Maor, Brian R. Griffin, J. P. Saxton, and Lawrence W. Davis. Fast neutron and mixed beam radiotherapy for inoperable non small cell carcinoma of the lung: Results of an RTOG randomized study. *American Journal of Clinical Oncology (CCT)*, 9:233–243, 1986.

[11] G. E. Laramore, J. T. Griffith, M. Boespflug, J. G. Pelton, T. W. Griffin, B. R. Griffin, K. J. Russell, W. Koh, R. G. Parker, and L. W. Davis. Fast neutron radiotherapy for sarcomas of soft tissue, bone and cartilage. *American Journal of Clinical Oncology (CCT)*, 12:320–326, 1989.

[12] Thomas. W Griffin, Thomas F. Pajak, George E. Laramore, W. Duncan, M. P. Richter, F. R. Hendrickson, and Moshe H. Maor. Neutron vs photon irradiation of inoperable salivary gland tumors: Results of an RTOG-MRC cooperative randomized study. *International Journal of Radiation Oncology, Biology and Physics*, 15:1085–1090, 1988.

[13] Thomas W. Griffin, John M. Krall, Kenneth J. Russell, Lester J. Peters, Frank J. Thomas, Frank R. Henderson, and George E. Laramore. Fast neutron irradiation of locally advanced prostate cancer. *Seminars in Oncology*, 15:359–365, 1988.

[14] Jüri Eenmaa, Brian R. Griffin, Sharon Warcola, Joseph Eskridge, and H. Richard Winn. A system for neutron beam stereotactic radiosurgery. *International Journal of Radiation Oncology, Biology and Physics*, 1989. (in press).

[15] L. H. Thomas. *Physical Review*, 54, 1938.