

THE FNAL LINAC

Dan Patterson

March 1986

Table of Contents

	INTRODUCTION.....	vi
1	PREACCELERATOR.....	1-1
1.1	Haefely.....	1-1
1.2	Haefely Operation.....	1-9
1.3	Sources.....	1-12
1.4	Extractor and Magnet.....	1-13
1.5	Column.....	1-17
1.6	Source Operation.....	1-17
2	750 keV TRANSPORT LINE.....	2-1
2.1	750 keV Chopper.....	2-4
2.2	750 keV Transport Line Operation.....	2-6
2.3	Changing Sources.....	2-7
3	TIMING AND PULSE SHIFTER.....	3-1
3.1	Timing.....	3-1
3.2	Enables and Interlocks.....	3-3
3.3	Pulse Shifter.....	3-5
3.4	Beam Inhibit and Enable Links.....	3-6
4	THE ALVAREZ STANDING-WAVE PROTON LINAC.....	4-1
4.1	Cavity Tuning.....	4-1
4.2	Cell Structure and Synchronous Phase.....	4-4
4.3	Design Considerations.....	4-10
5	LINAC RF SYSTEMS.....	5-1
5.1	Low-Level RF.....	5-1
5.2	High-Level RF.....	5-4
5.3	Modulator.....	5-8
5.4	RF System Operation.....	5-15
6	BUNCHER.....	6-1
6.1	Buncher RF System.....	6-1
6.2	Buncher Operation.....	6-3
7	QUADRUPOLES.....	7-1
7.1	Quadrupole Operation.....	7-4
8	200 MeV AREA.....	8-1
8.1	200 MeV Transport Line to Booster.....	8-2
8.2	Momentum and Emittance Analysis Lines.....	8-3
8.3	Linac Steering and Momentum Analysis.....	8-4
9	VACUUM SYSTEMS.....	9-1
9.1	750 keV Line Vacuum.....	9-1
9.2	Cavity Vacuum.....	9-2
9.3	200 MeV Area Vacuum.....	9-3
9.4	NTF Vacuum.....	9-6

10	WATER SYSTEMS.....	10-1
10.1	Preaccelerator Water Systems.....	10-5
10.2	Cavity and RF Water Systems.....	10-5
10.3	Water System Operation.....	10-9
11	NEUTRON THERAPY FACILITY.....	11-1
11.1	Beamline.....	11-1
11.2	Target and Collimator.....	11-2
11.3	Patient Setup.....	11-4
11.4	Control System.....	11-6
11.5	NTF Interlocks.....	11-5
11.6	RF4 Pulse Delay.....	11-10
11.7	NTF Operation.....	11-11
12	CONTROL SYSTEM.....	12-1
12.1	Link.....	12-1
12.2	Secondary Station.....	12-2
12.3	Primary Station.....	12-6
12.4	Operating Modes.....	12-7
12.5	Troubleshooting.....	12-8
12.6	Local Console Operation.....	12-10
13	LINAC SAFETY.....	13-1
13.1	Radiation Hazards.....	13-1
13.2	Linac Safety System.....	13-2
14	LINAC APPLICATIONS PROGRAMS.....	14-1
14.1	L25: Linac RF Status.....	14-2
14.2	L36: Linac Steering.....	14-4
15	LINAC CHRONOLOGY.....	15-1
	GLOSSARY.....	G-1
	REFERENCES.....	R-1
	LIST OF SCHEMATICS.....	S-1

This document assumes that the reader is familiar with the basic terms of electromagnetism and accelerator physics. Terms that are unique to accelerator systems are written in *italics* and may be found in the glossary.

Those not familiar with accelerators are encouraged to read Proton Synchrotron Accelerator Theory by E.J.N. Wilson, which may be found in the back of Basic Booster for Beginners by John Crawford.

Technical documents referenced to in the text are referred to by their DOSYL (Document Systematic Locator) designations, as follows:

OPBULL	Operations Bulletin
TM	Technical Memo
FN	Physics Note
EXP	Experimental Note
HDWREL	Hardware Release
SFTREL	Software Release

Many other useful documents are often not indexed through DOSYL, particularly papers given at conferences and IEEE reports. The reader is encouraged to enter relevant material into the DOSYL system so that it may be more easily referenced. DOSYL is accessed by logging on to the development VAX with username DOSYL; the program will take it from there.

The FNAL Linac




figure 0.1: linac component layout

INTRODUCTION

The FNAL linac (fig. 0.1) produces pulses of H⁻ ions for injection into the booster accelerator. The booster is a rapid-cycling synchrotron with a fixed cycle time of 1/15 second (66 msec) due to the nature of its magnet power supply system, which is designed to resonate at 15 Hz. As an injector for the booster, the linac has the same cycle time of 66 msec.

The linac uses an electrostatic *preaccelerator* (of Cockcroft-Walton design) to produce H⁻ ions at energies up to 750 keV (1 eV = one *electron volt*). These are then accelerated to 200 MeV in a linear series of nine standing-wave *RF cavities* that resonate at 201.24 MHz. *Transport lines* then lead either to the booster injection girder or to beam dumps. The entire linac is about 500 feet long.

The FNAL linac is patterned after, but not identical to, a 200 MeV linac built at Brookhaven National Laboratory beginning in 1964. Construction of the FNAL linac began in May, 1968, and 200 MeV beam was first produced on November 30, 1970.

A second function of the linac is to provide 66 MeV H⁻ ions for the production of neutrons used to treat cancer. The Neutron Therapy Facility (*NTF*) is a program which evaluates the effectiveness of neutron therapy against other methods used in treating certain types of tumors.

1 PREACCELERATOR

The two FNAL preaccelerators each provide up to 750 keV H⁻ ion beams for acceleration in the linac. Each "preacc" consists of an H⁻ ion source inside a *dome* that is kept at a potential of about -750 kV. Negative ions produced at the large negative potential are accelerated as they pass from the source to ground (the wall of the preaccelerator enclosure, or *pit*). Each preaccelerator then has a transport line to the linac, with the last portion common to both systems. At any given time one preaccelerator is on line, while the other is held in reserve or undergoes maintenance. The two sources are called "H-" and "I-" for purely archaic reasons; their construction and operation are virtually identical.

1.1 Haefely

The energy of the ions leaving the preaccelerator comes by virtue of the large potential at which the ions are produced. This potential is provided by a high-voltage multiplier which takes 5 kHz, 75 kVAC and converts it to -750 kVDC. The power supply that does this was manufactured by Emil Haefely CIE AG, Bern, Switzerland; thus the name, *Haefely*. The dome is also built by this company, but "Haefely" generally refers to the power supply itself.

One topic guaranteed to start lengthy arguments and harried references to dusty notebooks is the principle of operation of the Haefely power supplies. Each is simply a five-stage diode voltage multiplier. Figure 1.1 shows two stages of such a multiplier. Two stacks of capacitors are linked by a zigzag array of diodes. The capacitors on the right are "DC"; the charge on them is relatively constant. The capacitors on the left are *coupling capacitors*; they hold a charge plus couple the AC from the source at the bottom up the stack. (Every capacitor on the left side sees the AC voltage from the source.)

An AC transformer at the bottom provides a voltage at point A of $V_0 \sin \omega t$. Point B is at ground. Capacitor A-C couples the voltage at point A to the point C. When C becomes more positive than B, charge flows through diode B-C, charging capacitor A-C to $-V_0$. The voltage at point C is then the sum of the DC and AC components:

$$-V_0 + V_0 \sin \omega t$$




figure 1.1: diode voltage multiplier

On the next half-cycle, the minimum voltage at point C, $-2V_0$, charges capacitor B-D through diode C-D. The voltage at point D is then a constant $-2V_0$ because current cannot flow back through the diode.

The second stage operates in a like manner: the maximum voltage at C is just zero, which causes capacitor C-E to charge to $-2V_0$ (the voltage at point D) through diode D-E. The voltage at point E is then the sum of the DC voltages on capacitors A-C and C-E, plus the AC component (coupled from C to E through capacitor C-E):

$$-V_0 + (-2V_0) + V_0 \sin \omega t = -3V_0 + V_0 \sin \omega t$$

The minimum voltage at point E, $-4V_0$, then charges capacitor D-F through diode E-F to the maximum difference of points D and E, or $-2V_0$ (remember, the voltage at D is a constant $-2V_0$). The voltage at F is a constant $-4V_0$, since, again, current cannot flow back through the diode. Additional stages would likewise add $-2V_0$ per stage.

Figure 1.2 shows a schematic of an actual Haefely. Note that there are two legs with coupling capacitors, connected

The FNAL Linac




figure 1.2: Haefely voltage multiplier

to a common "DC" leg, which is connected to the dome at the top and to ground (through the current-measuring circuitry in the preacc control room) at the bottom. Two coupling legs are used to reduce the ripple at the output of the multiplier. Each coupling leg is driven by its own 75 kVAC transformer; five stages thus give a maximum output of $5(-2V_o) = -10V_o$, or -750 kVDC. Note: all Haefely drawings outside of this document show diode orientation for positive, not negative, high voltage. The diode orientation in figure 1.2 is correct. The actual physical layout of the Haefely power supply is shown in figure 1.3.



The two high-voltage transformers each have two 3200 volt primary windings, a 75 kV secondary, and a 200 volt tertiary winding. The 200 volt winding from one transformer is used for overvoltage protection of the Haefely itself (sec. 1.2). The other is used in one of the voltage regulation feedback loops, which will be discussed later. Both transformers are located in the preaccelerator pit, along with the voltage multiplier and the dome containing the ion source.

The output voltage at the top of the multiplying stack is sent to the dome through a 5 M Ω current-limiting resistor. The dome itself is isolated from ground by a series of *bleeder resistors* with a total resistance of 4250 M Ω . These resistors are contained in one of the five "legs" that support the dome (fig. 1.4). The bleeder resistor for the I- system is also in parallel with a series of capacitors and resistors which are part of the famous *bouncer circuit*. This unused system is designed to reduce the voltage droop caused by ions leaving the dome during acceleration to 750 keV; the only practical difference that the existing circuit components make is that the increased capacitance to ground of the I- system makes the voltage droop during beam time only 3-4 kV for I-, where it is about 7 kV for H-. This droop is not significant for operating purposes, as long as it is consistent.

The potential on the dome is monitored by first dropping the voltage through a *metering resistor* that runs from the side of the dome to the pit wall. This resistor consists of 600 precision resistors in series with capacitors, wired together in five elements of 425 M Ω each for a total resistance of 2125 M Ω . An oil circulating system for cooling this resistor exists, but is not used. The resistor terminates at the pit wall in a small aluminum box containing a voltage divider which drives the analog HV meter on the Haefely control panel as well as the voltage regulation loop electronics.

Regulation of the high voltage on the dome is handled by two feedback loops which act to keep the voltage within a tolerance of 0.25% at -750 kV. A slow feedback loop corrects for fluctuations up to a frequency of about 1 Hz; a fast loop handles frequencies up to about 15 Hz. The slow loop actually compares the measured voltage to the command voltage to produce an error signal. The fast loop looks at the 200 V tertiary winding in one of the HV transformers and

The FNAL Linac



53	2-5256891	BRACES	2
52	2-5255389	BRACES	2
51	3-5254403	BRACES	2
50	2-5256150	CYLINDRICAL ISOLATOR	2
49	2-5259151	CYLINDRICAL ISOLATOR	2
48	0-5255098	STEERING POTENTIOMETER	2
47	1-5256354	BRACE	2
46	0-5256340	POTENTIOMETER	2
45	2-5255391	SPARK-GAP	2
44	3-5241813	BALL	2
43	2-5255369	BRACE	2
42	0-5255502	COLUMN CAPACITOR	2
41	2-5256444	RESISTOR	2
40	0-5256348	SHIELDING HEAD	2
39	0-5256333	HIGH VOLTAGE TABLE	2
38	2-5256122	JUNCTION BOX	2
37	1-5255099	SHIELDING	2
36	0-5255360	INSULATOR	2
35	2-5256122	JUNCTION BOX	2
34	0-5256360	RESISTOR	2
33	2-5255456	FLANGE	2
32	3-660100	COLUMN ISOLATOR	2
31	2-5255238	SUPPORT	2
30	4-5255495	THREADED ROD	2
29	3-5255330	SUPPORT FRAME	2
28	2-5253236	SUPPORT FRAME	2
27	0-5255257	SHIELDING HEAD	2
26	4-5255454	THREADED ROD	2
25	4-5255453	SPACER	2
24	3-5255330	SUPPORT	2
23	2-5255450	SUPPORT	2
22	3-660098	COLUMN ISOLATOR	2
21	2-5255451	FLANGE	2
20	1-5255232	RECTIFIER	2
19	2-5253228	W-HOLDER	2
18	1-5255230	SHIELDS	2
17	1-5255231	SHIELDS	2
16	1-5663417	CAPACITORS	2
15	1-5663417	CAPACITORS	13
14	1-5255230	SHIELD	1
13	1-5255230	SHIELD	2
12	2-5663418	SUPPORT COLUMN	2
11	2-5663418	SUPPORT COLUMN	1
10	3-5255330	SPACER BAR	2
9	2-533445	BALL	3
8	2-5255391	SPARK-GAP	2
7	4-5256123	FOOT	2
6	4-5252799	SHIELD	4
5	2-5254997	TRANSF CONNECT ROD	2
4	0-5256480	GROUND PLATE	1
3	1-5255144	GROUND PLATE	1
2	0-5255983	SHIELD	1
1	4-213434	HIGH VOLTAGE TRANSF	2

figure 1.3: I- Haefely installation

The FNAL Linac

samples successive peaks in the voltage; this is compared with the command voltage to create a second error signal. The two error signals are then summed together with the command voltage to drive the transformer.

The slow loop begins at the voltage divider where the output voltage (about 4.5776 V at -750 kV) is sent via a twinax BNC to the *monitor/control module* in the Haefely control racks. The voltage is sampled and held while a 14-bit A/D digitizes it. The digitized monitor voltage is then fed through 14 *optical isolators* (light links) to isolate the input circuitry from ground and reduce noise due to ground loops. The signal then travels to a latching register which is sampled by the computer to produce the analog readback in the main control room (MCR). A second path leads to the *data selector*, which takes the digitized monitor and command voltages and sends them to the LED displays on the front of the module. The digitized command voltage comes from a 14-bit binary command generator, which is in turn fed by either a computer command or a local knob input.

The digitized monitor and command voltages are also sent via the crate wiring to the adjacent *regulator/oscillator module*. Here the signals are converted back to analog voltages and fed into a differential amplifier to create an error signal, which is then integrated and summed with the output of the fast regulation loop. A FET switch that zeroes the slow loop error signal and disconnects the fast loop input (and thus reduces the HV power supply output to zero) is controlled by the anode power supply on/off status (to protect the power amplifier tubes from being turned on with no drive signal) and by the *reference source* on/off switch (sec. 1.2).

The fast loop begins at the 200 V winding in one of the HV transformers; the 5 kHz sinusoidal output is first passed through a zero-bound circuit which removes the negative portions of the signal. A sample and hold signal, synchronized to a 5 kHz oscillator, samples the peak voltages which are then fed through a low-pass filter to roll off all frequencies above 15 Hz. This signal then represents the variation in the drive voltage of the transformer. It is summed first with the command voltage (from the same A/D as the slow loop) and then with the slow loop error. A zero-bound limiter insures that the resultant drive amplitude lies between zero and some predetermined maximum voltage. An analog multiplier modulates the drive signal to produce a 5 kHz square wave, which then is smoothed by a fourth-order multiple feedback bandpass filter. The resultant sine wave then drives the grid driver amplifier, which in turn modulates the voltage on the grids of two EIMAC 3CW20000A1 *triodes*, which comprise the power amplifier that drives the transformer primary windings. For the I- system, these tubes are located in racks in the

The FNAL Linac





figure 1.4: Haefely monitoring and regulation systems

The FNAL Linac

preaccelerator control room, next to the I- anode supply. The power amplifier and anode supply for the H- system are cleverly hidden beneath the stairs leading to the second floor of the preacc annex.

The HV transformer, voltage multiplier, and dome installation are shown in figure 1.3. The multiplying stack is rather impressive, being just short of 19 feet tall. Diodes are encased by Plexiglas tubes filled with a synthetic silicon fluid; capacitors are covered by *corona rings*. The five legs of the dome consist of G-10 (a glass epoxy insulator) tubes encircled by corona rings at the points where the segments meet. In the H- system, three legs are hollow and the fourth contains the 4250 MΩ bleeder resistor. The I- is the same with the addition of bouncer circuit capacitors in two otherwise hollow legs, linked by horizontal-running resistors. As stated before, the bouncer circuit is not used.



The dome itself contains the power supplies, vacuum pumps, and microcomputer necessary to run the ion source. At a potential of -750 kV, a moment's thought will convince the reader that running power cables to the dome is inadvisable. 208 and 120 VAC for the dome are supplied by a 15 kW generator (more correctly, an alternator) located in the dome itself. The generator is driven by a rotating G-10 shaft that runs up the fifth leg of the dome. The shaft is driven by a 25 kW electric motor located on the floor of the preaccelerator pit. Turning off the motor-generator kills all power to the dome, and even if quickly restored will upset normal source operation for a time. It is generally a bad thing to do.

A Plexiglas tube runs along the top of the metering resistor between the dome and the pit wall. This tube carries fiber-optic cables used for communicating with the microprocessor in the dome and for video signals. A nitrogen gas line supplies pressure to operate a vacuum valve in the dome.

The FNAL Linac

1.2 Haefely Operation

The Haefely HVPS is controllable either through the local control panels or through the computer system. The computer system communicates with the Haefely by way of a shielded, twisted pair of wires (*MIL-STD-1553B* digital bus) that run from the local microcomputer to interface cards mounted in a *Eurobus* crate. This small crate is located in the rack containing the Haefely control relays.





figure 1.5: Haefely control panel

Figure 1.5 is a sketch of the control panels for the H-Haefely. At the upper left is a keyswitch labeled MAINS. Turning this key clockwise removes power from the upper control panel, thereby shutting off the reference source, HV power supply, PA filaments, and motor-generator. The same result can be had by pressing the enormous orange HAEFELY EMERGENCY button on the lower panel. If the source is in

The FNAL Linac

use, this is inadvisable, as it will shut the anode supply off for five minutes, and the source for even longer.

The MOTOR GENERATOR ON/OFF switches control the power to the 25 kW motor that supplies power to the dome. Have a good reason ready before turning this off.

STANDBY ON/OFF controls the PA tubes' filament supply. When turned on, a timer prevents the anode supply from coming on for five minutes. This timer is activated whether or not the filaments are cold. It is not necessary to turn the filaments off unless there is work being done on the tubes themselves.

The BOUNCER switches aren't connected to anything.

HV POWER SUPPLY ON/OFF controls the anode supply for the PA tubes. It must be on before the reference source will turn on (this is done to protect the tubes). READY is the sum of anode power supply and safety system interlocks, which will be described shortly.

REFERENCE SOURCE OFF zeroes the HV transformer drive signal in the monitor/control module. This causes the output of the HV transformer to go to zero, and the dome HV will bleed down through the 4250 M Ω resistor. The anode power supply is cycled on and off in this state; it must be off before an access to the preaccelerator pit is made.

The LOCAL/REMOTE switch disables/enables control through the computer system. All functions are available remotely.

OIL PUMP ON/OFF controls nothing, but it must be on so that other interlocks in the Haefely control circuitry will permit the supply to be turned on.

The following anode power supply interlocks must be made up before HV POWER SUPPLY READY (anode power supply ready) will be lit:

NO COOLING is the sum of water flow and temperature switches on the PA water system, as well as PA cooling fan power. The I- system also has a water pressure switch.

ANODE POWER SUPPLY OVERLOADED is an overvoltage trip of (what else) the anode supply.

BOUNCER OVERLOADED is not connected.

The next three interlocks will trip off the reference source, but leave the HVPS on:

OVERVOLTAGE AC SUPPLY is an overvoltage trip of the HV transformer. This circuit looks at the "other" 200 V transformer winding.

The FNAL Linac

OVERVOLTAGE HVPS is an overvoltage trip of the voltage multiplier stack and dome. It is simply the trip point on the analog voltage meter located in the panel above the upper control panel.

OVERCURRENT HVPS is sensed by a *thyatron* circuit that looks at the current actually being drawn by the voltage multiplier. This circuit is electrically in the path from the base of the DC leg of the multiplier to ground. It is physically located in a box below the metering resistor voltage divider.

The following safety interlocks must be made up before HV POWER SUPPLY READY (anode power supply ready) will be lit:

The DOOR OPEN indications (NORTH and SOUTH for H-, SIDE and PIT for I-) are connected to microswitches on the pit doors. The NORTH DOOR OPEN circuit for the H- system is also connected to a storage room door that is only accessible from inside the pit.

Q4 CENTER OFF (H- only) refers to the center element of the first *quadrupole triplet* in the H- 750 keV transport line. If this element is off, the H- Haefely is turned off to prevent unacceptable levels of X-ray radiation in the H- 750 keV transport line area, which would be caused by high-energy electrons striking the 45° bending magnet (sec. 13.1).

ELEVATOR UP (I- only) indicates that the elevator in the I- pit is not all the way down. The H- elevator is portable and is not interlocked.

GRNDSTICK UP (groundstick up) indicates that the groundstick in the pit is not stored in its proper place. For the I- system this is across the lower pit door entrance. For H- it is across the bottom of the pit stairway.

SAFETY SYS. FAILURE indicates that the linac safety system is down and that the *beam stop* in the downstream 750 keV transport line has not closed (OPBULL 717) (a *critical device failure*). The Haefelys for both H- and I- systems are then automatically shut down to prevent beam from leaving the preaccelerators.

There are also interlocked keys for each system (two for I-, three for H-) that open the appropriate pit doors. These keys must be in their slots and turned clockwise in order for the READY sum to be made up.

The procedure to access the I- (H-) dome (OPBULL 486) is as follows: turn off the reference source and the high voltage power supply. Pull one of the keys on the lower panel (not the key on the upper panel!) and use it to open the lower (north) pit door. Across the doorway (across the

bottom of the stairs) is a groundstick. Use it to ground the HV transformer and the multiplying stack, working from the bottom up. Ground the top of the multiplying stack and the dome. Hang the groundstick from the top of the multiplying stack. A second groundstick may be hung from the dome. Raise the elevator to the up position. Back in the preacc control room, take a second key from the lower panel and the key to the dome, which is hanging on the Haefely control rack. Use the key from the lower panel to open the upper (south) pit door.

To leave the dome, close the dome door and upper (south) pit door. Go back through the lower (north) pit door, and lower the elevator all the way (the H- elevator isn't interlocked, but the door to the storage room is- make sure it is closed). Remove the ground sticks and place the one across the entrance back in its holder, making sure that it makes contact with both microswitches. Close the lower (north) door and return the keys to the lower panel in the Haefely control rack. Turn on the HVPS, and finally, the reference source.

If the Haefely won't maintain full voltage, check the current being drawn by the supply (normally 1-3 mA) and the conductivity in the water system for the water resistors (sec 10.1) If the conductivity is poor (greater than 1.1 $\mu\text{mho/cm}$), the Haefely will draw more current and may trip off before it reaches full voltage.

Of course, it is a good idea to check the pit itself in these situations, to check for major component failures or water leaks.

1.3 Sources

The ion sources are the direct-extraction *magnetron* type, conceived at Novosibirsk around 1972 and developed at Brookhaven and FNAL for use in particle accelerators. The first operational use of the magnetron negative-ion source at FNAL was in 1978. Magnetrons have supplanted the charge-exchange *duoplasmatron* type source which produced protons and is still encountered in older documentation.

The magnetron H- source (OPBULL 600, 953) produces short pulses of negative ions at the linac repetition rate of 15 Hz. These ions are created in a *plasma* formed near a metal surface. A pulsed valve introduces hydrogen gas at low pressure into the volume between two molybdenum electrodes: a matchbox-sized, oval-shaped cathode and a surrounding anode, separated by 1 mm and held in place by a glass ceramic insulator. An external magnet provides a 1-1.5 kG magnetic field parallel to the cathode surface. A low-impedance pulse forming network strikes an 80 μsec , 40 amp arc between the electrodes. Electrons in the arc, spiraling about the magnetic field lines, efficiently ionize the gas to form a dense plasma of H⁺ ions and electrons in the gap.

The FNAL Linac




figure 1.6: magnetron negative-ion source

H^+ ions strike the cathode and occasionally pick up two electrons or "sputter" H^- ions from the surface. H^- ions are repelled from the cathode and charge-exchange with neutral hydrogen atoms at the plasma boundary to produce H^- ions with a smaller energy spread. A pulsed electrostatic extractor then accelerates the negative ions out of the source.

The production of negative ions by the cathode is enhanced by introducing cesium vapor into the source, which coats the electrodes and lowers the surface *work function*. Cesium vapor is supplied by vaporizing solid cesium in an electrically-heated *cesium boiler*. The vapor then travels through a heated tube to the source. Five grams of cesium is enough for about a year of source operation.

Extremely pure hydrogen gas at 10-30 PSIG is supplied by a 30 ft³ gas bottle located inside the dome, which is enough for about six months of source operation. Gas pressure in the source is controlled by a *piezoelectric crystal valve*. When a voltage pulse is applied to the crystal, it bends, admitting gas to the source. A second DC supply on the crystal is used to compensate for temperature changes which affect the operating point of the valve. Source gas pressure is regulated by modulating the width of the voltage pulse that opens the valve. This task is handled by a servo loop residing in computer software that acts to keep the pressure near the set point chosen by the operator.

1.4 Extractor and Magnet

A 1 x 10 mm aperture in the source body allows negative ions and electrons to exit the source when an electrostatic




figure 1.7: FNAL magnetron ion source

extractor, 2 mm below the aperture, is pulsed to about 18 kV (fig. 1.8).

The high density (1 A/cm²) beam exiting the source has a high *divergence* due to *space charge forces* and the defocusing effect of the extractor. To control this, a high-gradient 90° bending and focusing magnet reduces beam divergence from 250 mrad to about 37 mrad. The pole faces of this magnet are slanted to produce a quadrupole field and focus the beam in both planes. Pole-tip extensions on the top straddle the source and provide the magnetic field necessary to sustain the plasma. Extracted electrons are separated from the H⁻ beam by the magnetic field. The pole pieces, cooled to -30° C by a Freon refrigerator located in the dome (fig. 1.9), also capture wayward cesium atoms that are extracted from the source and could cause sparking in the *accelerating column* if not removed. The ion beam leaving the 90° bending magnet is about 50 mA and 1.5 x 2 cm in size.

Sixty percent of the gas which flows through the source is drawn off by a 400 liter/sec *turbomolecular pump* located in the dome. Source pressure is actually measured in the vacuum line to this pump, which may be isolated by a valve, but is normally left open. The valve is operated by nitrogen gas supplied by a tube leading from the outside world. The valve is manually controlled. The hydrogen gas not removed by the turbo flows through the accelerating column and is removed by a 2400 liter/sec *ion pump* (powered by four separate supplies) located outside the preaccelerator pit.

The FNAL Linac




figure 1.8: extractor and magnet

All the hardware necessary to run the source is contained in the dome, which is actually a hollow aluminum cube with rounded corners. Inside are the source assembly, hydrogen gas supply, gas valve pulser, arc supply and pulser, source heater, cesium boiler power supply, cesium tube and valve heaters, extractor power supply and pulser, magnet power supply, magnet pole refrigerator, *ion gauges* for measuring vacuum, *roughing* and turbo vacuum pumps, turbo pump controller, Freon cooling system for turbo pump and magnet coil, generator and breaker panel, oscilloscope and TV cameras, and control microprocessor.

Communication with the control microprocessor in the dome is done by way of fiber optic cables run from the preacc control room (sec. 12.2). The same is true of clock signals (sec. 3.3) and the TV camera video. The control and clock cables are blue, and the video cables are orange or red.

FERRILAB ENGINEERING NOTE	SECTION	PROJECT	SERIAL-CATEGORY	PAGE
	H- FEAACC ION SOURCE ASSEMBLY C.F. 0211-WE-60467			
NAME DRAWINGS/COPIED: PATTERSON		DATE 8-2-82		
SUBJECT H- FEAACC ION SOURCE ASSEMBLY		REVISION DATE		




figure 1.9: ion source installation

1.5 Column

Five centimeters from the exit edge of the bending magnet the beam enters the accelerating column (fig. 1.10). The column consists of seven perforated, disk-shaped titanium electrodes that are designed to guide the ion beam during its acceleration to ground potential. Titanium was chosen by virtue of its high work function, to reduce the number of stray electrons in the column. Ceramic rings insulate the electrodes from the outside world and one another. Two *water resistors* passing from the dome through the outer edges of rings (connected to the electrode disks via rods) to the pit wall maintain roughly equal potential differences between the disks. The distance between the electrodes is small compared to their radius for good field quality on the axis, and to reduce defocusing effects of any charge picked up by the insulating rings. Holes near the outer edge of the electrodes ease the task of vacuum pumping. A set of these electrodes may be seen on the west side of the second floor of the central laboratory.

The column is encased in a spun glass pressure vessel containing sulphur hexafluoride gas at two atmospheres' pressure. This gas is a good insulator and acts to reduce the possibility of the column leads arcing to each other or ground. *Spark gaps* at the outer edge of the electrode rings protect the column by handling most of the current as a result of a spark or bad voltage distribution. Vacuum in the column is maintained by the 2400 liter/sec Ultek ion pump located outside the pit.

1.6 Source Operation

Figure 1.11 is a linac parameter page used to monitor and control operation of the sources. This particular page represents a healthy seven-month old source in the H-preaccelerator; thus all parameter names start with "H". Identical parameters for the I- source start with "I".

The source gas pressure and the source temperature are vital to proper operation. The cathode optimally runs at a temperature of about 400° C, with an cesium coating of 0.6 of a monolayer. If the temperature is too low, cesium tends to condense on the electrodes. Too high a temperature results in an insufficient coating of cesium on the cathode, and in extreme cases may damage the piezoelectric gas valve. Cathode temperature L:HCTEMP is controlled by varying the current in a resistive heater in the source body (controlled manually by a Variac in the dome). The source body temperature itself, L:HSTEMP, is quite a bit cooler than the cathode; the difference is due to heating of the cathode by the arc. The arc pulse width L:HTAW is thus another control of the cathode temperature, although it is not the preferred knob.

Gas pressure affects both the quality of the arc and ion production rates. Too low a pressure results in an unstable arc. Increasing the pressure (by raising the servo loop setpoint L:HPRES) stabilizes the arc, but may reduce beam

The FNAL Linac




figure 1.10: ion source and accelerating column

The FNAL Linac

WED 06-NOV-85 09:57

```

L57 H- SOURCE PARAMETERS.....SET . D/A A/D *COPIES*
<FTP>+ #SS* X-A/D X=TIME Y=MIMKS01 ,M TOR109,M BBS ,
#Save ... ENQ-U I= 0 I=-10 ,-.2 , -10 , -10
< < >+ r_MR AUTO F= 60 F= 80 , .2 , 10 , 10
SURVEY 750 kev timers misc 750 2 .....
-L:HVOLT=.1 H- HIGH VOLTAGE < > -745.1 KV
-L:HARCSV=.1 H- ARC SUPPLY VLTS 277.2 282.4 V
-L:HMAGI=.1 H- MAGNET CURRENT -7.557 7.568 A
-L:HEXTV=.1 H- EXTRACTION VLTS 18.18 18.04 KV
L:HEXTI H- EXTRACTOR CURR -2.734 MA
-L:HGASV=.1 H- GAS PULSED VLTS 141.2 136.1 V
-L:HGASO=.1 H- GAS OFFSET VLTS -.198 -.049 V
-L:HTGW=.01 H- GAS PULSE WIDTH 199 199 US.
-L:HTGDN=.1 H- GAS VALVE DN 700 700 US.
-L:HPRES=.1 H- SRC PRES MAN AUT 19.69 19.58 UT
L:HARCI H- ARC CURRENT -37.6 A
L:HARCV H- ARC VOLTAGE -152.3 V
-L:HTAW H- ARC PULSE WIDTH 83 83 US.
L:HCTEMP H- CATHODE TEMP 410.5 DEG
L:HSTEMP H- SOURCE TEMP 366 DEG
L:HDTEMP H- CS BOILER TEMP 115.4 DEG
-L:HDILV H- CESIUM BOILER -8.887 50 V
L:HVACV H- H VACUUM VALVE 1.020 TOR
L:HVACE H- H VACUUM EXP 6.078
L:HTOR4 H- TOROID #4 #-52.56 MA
-L:HSPTHR H- SRC PR THRESH .198 .198 UT
-L:HSPWIN H- SRC PR WINDOW 20 20 UT
-L:HCPR H- COARSE PR REG 5 5 US
L:HMMAINV H- MAINS VOLTAGE 117.2 V
L:HMMAINI H- MAINS CURRENT 19.57 A
L:HCCHV H- H CHOPPER HV 19.55 KV
L:HCCHI H- H CHOPPER CURR 2.419 MA
L:HTOR4 H- TOROID #4 #-52.21 MA
L:HTOR5 H- TOROID #5 #-53.45 MA
L:HTOR6 H- TOROID #6 #-55.69 MA
L:HTOR7 H- TOROID #7 #-53.63 MA
L:HTOR8 H- TOROID #8 # .256 MA
L:HTOR9 H- TOROID #9 # 1.312 MA
-L:HMIST H- FIRST ELECTRODE 0 27.03 KV

L:IPTEMP PREACC AIR TEMP 75.45 DEG
L:HDTEMP H DOME TEMP # 84.3 DEG
L:IDTEMP I- DOME TEMP # 76.07 DEG
    
```

figure 1.11: ion source parameters

current by stripping the H- ions in the source as the ions collide with gas molecules. High source pressure also increases the load on the ion pump and can induce sparking in the extractor. Nominal source pressure is about 28 μ torr.

Source pressure is regulated by a servo loop (OPBULL 671) which samples the pressure in the vacuum line to the turbo at 15 Hz. Fifteen samples are made, followed by one correction, resulting in an adjustment to the gas pulse width L:HTGW about once a second. If the source pressure is outside the window formed by L:HSPWIN and L:HPRES, no correction will be made and operator intervention may be necessary. To guard against pressure bursts, another algorithm (OPBULL 917) reduces the gas pulse width L:HTGW by a number of μ sec (L:HCPR) every second that the source pressure is more than sum of L:HPRES and L:HSPWIN (the upper limit of the window). The deadband of the loop is set by L:HSPTHR, the source pressure threshold.

During startup or recovery from "loss" of a source, the cathode must be brought to the correct temperature and have a sufficient coating of cesium before reasonable ion

production can be expected. Once the cesium boiler and cesium transfer line heaters have been turned on, the source can be started up after 1-2 hours. In the "old days", it was advised to raise the source pressure, increase the arc pulse length, and raise the magnet current to speed up heating of the cathode; more recent experience has shown that equally good results can be obtained by simply setting all parameters to nominal operating values and waiting. Stubborn sources may benefit by temporarily raising the source pressure to 40 μ torr or so. Arc voltage will initially be high and arc current low as the cathode warms up and acquires cesium. As the arc improves it is possible to slowly decrease the source pressure. Good output current (L:HTOR4 for H-, and L:ITOR1 for I-) should be seen within an hour of good arc operation. In any event, the nominal source values should not be changed without good reason and a note in the logbook.




figure 1.12: ion source scope traces

Figure 1.12 shows a familiar oscilloscope trace that shows source operation at a glance. This trace is sent from a TV camera in the dome via the light link. The upper trace is the sum of the arc voltage and gas valve voltage pulse. (The gas valve voltage pulse occurs much earlier than the arc voltage pulse, so it is necessary to adjust the scope trigger time L:ITSCOP or L:HTSCOP earlier to see it.) The time at which the source parameters are sampled is visible as a glitch on the upper trace. The lower trace is the sum of the extractor voltage (upper dip) and the arc current (lower dip). If the extractor is arcing, the lower trace will jump around considerably. If turning down the source pressure doesn't cure the problem, it may be necessary to reduce the extractor voltage. Be forewarned that the extractor voltage and magnet current both affect the

The FNAL Linac

steering of the beam through the column, as well as the rest of the linac. Do not change nominal values without consulting specialists.

Source lifetimes vary from source to source for not always well understood reasons. The longest-lived source observed at FNAL so far (11/85) is running well at seven months. After a time the output of a source will start to decrease, usually visible as a pronounced slope in the arc current trace. It must then be removed and cleaned of cesium hydride and other sputtered deposits (molybdenum, etc.).

For a temporary shutdown of a source, first shut off the cesium boiler and transfer line heaters, shut off the source body heater (sometimes called the anode heater) and close the cesium supply valve (located in the dome). Wait two hours, then shut off the magnet and arc supplies. Wait another two hours, then turn off the H₂ gas supply.

A permanent source shutdown is much simpler: just turn off everything except the gas supply until the source cools, and then turn off the gas as well.

Spares for the supplies that power the source heaters, magnet, and arc supplies are found in the preacc control room.

The FNAL Linac




figure 2.1: 750 keV transport line layout

2 750 keV TRANSPORT LINE

Each preaccelerator has its own transport line up to a point where the lines merge and continue on to the linac. The I- transport line is the simpler and shorter of the two, as seen in figure 2.1. Once the beam pulse has left the column it passes through a quadrupole triplet Q1, which focuses the beam in both planes. The end elements share a power supply L:IQTM1E, and the center element has its own supply L:IQTM1C. The beam current then is measured by a toroid T1 (L:ITOR1).

The beam then passes through an electrostatic *chopper* which is used to select a short portion of the beam pulse to be sent to the linac. The unwanted beam is deflected into a carbon disk, while the selected beam (the "chop") passes through undeflected. The 750 keV chopper is described in greater detail in sec. 2.1.

The selected beam chop (40 μ S or so for High-Energy Physics (HEP)) then passes through a 90-degree bending magnet which is off for I- operation; the beam travels a straight path through a second quadrupole triplet Q2 (power supplies L:QTM2C and L:QTM2E), the buncher (sec. 6), toroids T2 and T3 (L:IHTOR2 and L:IHTOR3), and a third quadrupole triplet Q3 (power supplies L:QTM3C and L:QTM3C2 for the center element, and L:QTM3E for the end elements).

An *emittance probe* (FN 201, EXP 111) then leads to tank one of the linac. This device (fig. 2.2) measures the angular divergence of the beam by stepping a 0.075 mm slit across the beam and recording the width of the image formed on a pickup plate 10 cm away. The pickup plate is divided into 20 segments on 0.2 mm centers. Each segment thus corresponds to 2 milliradians angular divergence. The stepping motion of the slit and the recording and analyzing of the data are done by computer. This is a destructive measurement and is done only under special conditions.

The H- transport line is a bit more involved. The first element is a quadrupole triplet Q4 which is unique in that it is possible to control all three elements individually (L:HQTM4U, L:HQTM4C, L:HQTM4D), although (4U) usually runs at a fixed value. Toroid T4 (L:HTOR4) measures the beam coming out of the column, and is followed by a vertical dipole trim S1 (L:HVT1). Then come the H- chopper, an emittance probe, toroid T5 (L:HTOR5), horizontal and vertical dipole trims S2 (L:HHT2, L:HVT2), and a single quadrupole Q5 (L:HQ5).

The FNAL Linac




figure 2.2: emittance probe

The next element is a 45° bending magnet which is run in series with the 90° magnet mentioned earlier. These magnets have a bulk supply L:HBENDS and the 90° also has a trim L:TRIM90. A gaussmeter in the aperture of the magnet measures the magnetic field, but is not used for regulation or tuning. L:HBENDS is off for I- operation, and the beam extracted from the H- source at that time is then sent straight through the 45° bending magnet to a carbon block.

Under very special conditions (which have not occurred in the last five years) it may be desirable to run the 45° magnet without the 90°, so as to run H- source beam straight through the 90° magnet. To do this, it is necessary to remove the 90° magnet from the circuit by rotating two rhomboidal plates located on the wall below Q9 in the H- 750 line. In normal operation, both magnets are left in the




figure 2.3: 45°-90° lead switching

circuit, and the supply L:HBENDS is turned on or off depending on whether H- or I- beam is to be used (sec. 2.2).

The FNAL Linac

Next in the line is toroid T6 (HTOR6), a beam profile monitor W1, and a number of single quadrupoles Q6-10. Q6 and Q10 are in series, as are Q7 and Q9. Proper operation of the line requires a specific relation between the current in Q7 and Q9 and the currents in Q8, Q6 and Q10 (see below).

S3 and S4 are dipole trim packages (L:HHT3, L:HVT3, and L:HHT4, L:HVT4).

W2 and W3 are more beam profile monitors. The three monitors W1-3 are manually controlled from the preacc control room and were used for the initial tuneup of the line; they haven't been used for several years.

Toroid T7 (L:HTOR7) is used for rough tunes of the H-transport line. (Older documentation will also show toroids T8 and T9, but these are no longer in place.)

The beam from the H- source is then bent 90° to travel through Q2 and on to the linac, as would the beam from the I- source.

All the quadrupole power supplies are Sorenson SCR20-250s (OPBULL 577) except L:QTM3E. All power supplies for the above quads and dipoles are in the preacc control room, except for the two large EMHP supplies. One of these (below the stairs leading from the upper gallery to the preacc control room) powers the 45°-90° bending magnets, and the other, which powers L:QTM3E, is in the lower linac gallery near the water system for RF station 1. Any of the above supplies will trip off if there is insufficient water flow for their loads.

There are four vacuum valves in the 750 keV transport lines (fig. 2.1 again); two are interlocked and two are controlled manually. The status and control of these valves are handled by the gate valve controller located in the preaccelerator control room (OPBULL 720). The H- and I- valves are interlocked to the vacuum systems; both valves will close if either linac vacuum (monitored by looking at the ion pump power supplies and vacuum in tank one) or H-750 keV transport line vacuum (monitored by a thermocouple in the transport line) goes bad. If the vacuum in either column goes bad (monitored by ion gauges) the valve for that system will close. These valves can also be closed manually from the gate valve controller. If the H- valve closes, the 45°-90° supply L:HBENDS will trip off to keep beam from hitting the valve. The I- valve has a carbon block attached to its upstream face to protect it from beam if it closes.

The H- valve and the Beamline valve should not be closed at the same time. There are no vacuum pumps in the transport line between these valves, and the line cannot be pumped if both valves are closed.

The Beamline valve and the Linac valve are manually controlled from the gate valve controller. The Linac valve is also interlocked to the tank one vacuum pumps and will close if both pumps trip.

Between the quadrupole triplet Q2 and the Linac vacuum valve is the linac safety system *primary critical device*: the beam stop. The beam stop is a block of metal that will

drop into the beam path if the linac safety system drops (to protect personnel), or if the pulse shifter (sec. 3.3) fails to inhibit beam when it should. The block is pneumatically controlled, but should drop in the event of a power failure or loss of gas pressure. If the beam stop is commanded to close and does not, both the H⁻ and I⁻ Haefely power supplies will trip off.

2.1 750 keV Chopper

The amount of beam that is allowed to pass from the source to the linac is determined by the chopper (OPBULL 210), consisting of a pair of conducting plates that straddle the beam path (fig. 2.4). An initial potential difference (a) between the plates causes the beam to be bent off to one side and be lost in a carbon disk. The plates are then brought to the same potential for a controlled period of time (b), which allows the beam to pass through undeflected. The potential difference then returns, and the beam is bent again (c).

The duration of the selected beam pulse (the "chop") varies depending on the use to which the beam pulse will be put: HEP, P-bar (antiproton) production (taking the place of the old cooling ring chop), neutron therapy (NTF, or the old designation CTF), or standby (the do-nothing chop).




figure 2.4: 750 keV chopper operation

The 750 keV chopper power supply (fig. 2.5) consists of two thyratrons connected through a 4 MΩ resistor to a common anode supply that runs at about 28 kV. One thyratron (the *on tube*) has its anode connected through 90 Ω to one of the plates. The other thyratron (the *off tube*) has its anode coupled to the second plate through 90 Ω and a series

capacitor. The second plate is also connected to ground through a 100 kΩ resistor. Initially, then, the first plate is at 28 kV and the second plate is at ground. Incoming beam is deflected by the electric field between the plates. When the "on" tube fires, the tube anode (and thus the first plate) is grounded. With both plates at ground, extracted beam from the column passes through the plates unperturbed. When the "off" tube fires, the series capacitor, which initially had one side at ground and the other at 28 kV, suddenly has the high-potential side clamped to ground. The voltage change (the AC component) is coupled across the capacitor, bringing the second plate to -28 kV; the extracted beam is again deflected into the carbon disk. The second plate then charges back up to ground (relatively slowly) through the 100 kΩ resistor. Likewise, the first plate charges up to 28 kV through the 4 MΩ resistor leading to the anode supply. Note that the 80 μS beam pulse (determined by the arc pulse width) ends long before the plates return to their initial conditions.




figure 2.5: 750 keV chopper circuitry

The chopper on and off times are determined by *predets* in the preacc control room; there are on and off times for each of the four types of chop. Which pair of chop times is actually used for a given pulse is determined by a module in the preacc control room (sec. 3.2). The same times are sent to both H- and I- choppers, and both normally run all the time.

The chopper power supply will trip off if the ion gauge that measures the column vacuum trips off. The power supply will then require a local reset. The H- and I- chopper supplies are located in the upper linac gallery by RF station 1.

2.2 750 keV Transport Line Operation

Tuning the I- 750 keV line is a matter of tuning the horizontal trim L:TRIM90 and the quads for maximum transmission into (L:T01IN, the beam current into tank one) and out of (L:T09OUT, the beam current out of tank nine) the linac. Maximum intensity through both toroids may, with careful tuning, occur at the same time. *Linac steering* (sec. 8.3) should be done after each iteration. Lack of a vertical trim in the I- line may cause problems with the linac steering. The quadrupole currents may have to be "fudged" to improve the steering or beam profiles at 200 MeV. Usually the intensity peaks described while tuning the quads are broad enough to permit altering the magnet currents a bit without adversely affecting the beam intensity. Transmission of the beam through the I- 750 keV line should be almost 100%.

The H- 750 keV line is a bit more involved due to its greater length and number of components. This transport line is designed to be *achromatic*; that is, particles with different momenta that begin on the same trajectory will be on the same trajectory at the end of the line. This implies that tuning is not simply a matter of getting maximum transmission. The quads Q6-10 are an achromatic set and must have a particular current relationship if they are to remain so. A tuning guide for the 750 keV lines was published in OPBULL 628 and a graph showing the correct relationship between the current in Q7 and Q9 (independent variable) and the currents in Q8, Q6 and Q10 (dependent variable) is found there.

To tune the H- line, first find the gross tune of the bends by plotting L:HTOR7 as a function of L:HBENDS. Then fine-tune the field in the 90° magnet by plotting L:T09OUT as a function of L:TRIM90.

The first element of HQTM4, L:HQTM4U, runs flat-out and shouldn't be changed. The other two elements, L:HQTM4C and L:HQTM4D are tuned looking at L:HTOR7 and L:T09OUT, as are L:HQ5 and the pair L:HQ7Q9. The correct values of L:HQ8 and L:HQ6Q10 then depend on the setting of L:HQ7Q9, although some specialists also tune L:HQ6Q10 independently of L:HQ8. The triplets QTM2C,E and QTM3C,E are tuned looking at L:T01IN and L:T09OUT, as with I-.

All the trim dipoles in the H- 750 line are tuned looking at L:HTOR7 and L:T09OUT.

The FNAL Linac

After tuning the trims do the quads again. Several iterations may be necessary to get good intensity and good steering wire profiles at the same time.

2.3 Changing Sources

As sources grow older they diminish in output; source pressure may be raised to maintain nominal arc current, but it is at best a temporary measure, since the extractor may start to spark and the beam current will likely diminish anyway, due to stripping of H⁻ ions by gas molecules. Eventually, the source will slide off the deep end and it will be necessary to switch to the other source (if the other source is already down you have a problem).

To switch from one source to another, first establish that the magnet current, extractor voltage, Haefely voltage, and chopper voltage for the source you are switching to are at nominal values. Check the readbacks of the devices in the new source's transport line against an old listing taken when the line was in use, and make sure that the monitoring for the line is enabled. Then close the gate valve for the dying source (done from the gate valve controller in the preacc control room) and make the appropriate change in the status of L:HBENDS (off for I⁻, on for H⁻). Finally, open the gate valve for the new source and tune up the new transport line.

The FNAL Linac

3 TIMING AND PULSE SHIFTER

Linac operation is synchronous with the booster, which cycles at 15 Hz. The 1.00002 MHz booster clock sent to the preacc control room is phase-locked to the 60 Hz line frequency (60×16667) and is encoded with *phase reversals* to mark the reset events of interest to the booster (see Introduction to the Control System by Dave Johnson for a description of the clock system). The number of reversed-phase cycles of the clock denotes the type of event encoded; one cycle is a "1-gap", four cycles is a "4-gap", and so on. For the booster, a 1-gap is a "nothing" reset. A 4-gap is a *prepulse*, where the pulsed devices in the booster (as well as the 200 MeV chopper and S1) are fired to prepare them for a beam pulse. An 8-gap is a booster beam pulse. As the booster cycles at 15 Hz, one and only one type of phase reversal occurs every fifteenth of a second.

These 15 Hz pulses are used as a reference for all the timers in the linac. By slightly altering the sequence of timers, the linac can be made to run in different ways. The *pulse shifter* is the means by which linac beam is inhibited. Instead of physically blocking the beam pulse from the preaccelerators, it simply delays the timing for the preaccelerators by 1 msec so that the beam enters the linac after the RF pulse is over. The beam is scattered harmlessly in tank 1 and none is transmitted.

3.1 Timing

Refer to figure 3.1. (All modules are located in the preacc control room unless otherwise indicated.) The 1 MHz booster phase-reversal clock is sent from the phase-reversal clock generator in the MAC room (on the floor beneath the MCR) to the *clock module* in the preacc control room. This module has several functions: first, it converts all phase reversals to 1-gaps, effectively producing a 1 MHz clock with 15 Hz phase reversals laid on top. Second, it looks to see whether any phase reversals on the incoming clock are missing. If there are missing phase reversals, the module supplies them with its built-in 15 Hz generator. Finally, if the incoming clock disappears entirely, a built-in 1 MHz generator takes over, while simultaneously sending an error signal to the local microcomputer (sec. 12).

The output of the clock module is daisy-chained to service a number of predets, which produce output pulses delayed with respect to the 1-gaps ("1" on REF TIME thumbwheel on the front panel). The delay of each timer is also settable via a thumbwheel, with minimum resolution of 1 μ sec. The trigger for the timers may also be a second external reference.

The PRIMARY μ P START output of the first pre-det (delay 3000 μ sec) triggers data-collecting by the linac control system (sec. 12.3). The reference time output (coincident with the 1-gap) is the TZERO pulse that informs the rest of the control system that a new cycle is starting.

The second pre-det has outputs QUADS ON (18 μ sec), which triggers all the quadrupole power supplies in the linac cavities, RF ON (1755 μ sec), which is the beginning time for the RF pulse sequences, and TDATA (2002 μ sec), which is a backup trigger for the A/D converters in the linac (sec. 12). The timers TZERO, QUADS ON, and RF ON are sent from the pre-dets to *timing pulse repeaters* by RF system 1, and then by cables down the length of the linac (QUADS ON has more repeaters at systems 6 and 8). Each system picks up the signals from the cables via 5-turn transformers.

The reference output of the second pre-det (again, should be 15 Hz) goes to a *pulse skip detector* which looks for missing pulses that would indicate that the booster phase-reversal clock and clock module have both failed to provide 15 Hz resets to the linac. The pulse skip detector then sends a signal to the local microcomputer, which posts the alarm 15 HZ PULSE ERROR and inhibits beam. This alarm often toggles when the booster gradient magnet power supply (GMPS) cycles on or off.

The third and fourth pre-dets have the on and off times for the 750 keV choppers. The four chops (HEP, NTF, P-bar, and standby) can each have different times. The timers are referenced to an external input, which will be discussed shortly.

After passing through a fifth (spare) pre-det, the 1 MHz clock passes down the linac, as do TZERO, QUADS ON, and RF ON, but without the help of timing pulse repeaters. The 1 MHz clock is used by the microcomputer at each system to run the local timers (sec. 12.2).

3.2 Enables and Interlocks

The linac will not deliver beam unless it is asked to do so. This requires that an *enable* for a particular type of beam pulse exist. Only one type of enable is allowed per pulse. At present there are three types of enables: HEP, for injection into booster (and then on to main ring and the tevatron), P-bar, for injection into booster (and then on to main ring or the AP-4 line), and NTF, for the neutron therapy facility (sec. 11). Whether or not a particular enable eventually generates a beam pulse depends on other interlocks, generated by accelerator hardware and by switches in the MCR.

The *keyswitch module* (also known on old schematics as the pulse shifter controller module) in the MCR is the first module involved in determining whether an HEP enable will exist. An 064 timing module in the MCR generates an HEP ENABLE on tevatron clock resets 13, 15, and 19 (booster resets on the tevatron clock system). This enable will be

The FNAL Linac

defeated if the status of the booster 5° bend switch (read directly from switch itself) is not in the proper position for the current mode of operation (see Basic Booster for Beginners by John Crawford) or if the abort logic/pulse shifter interface (OPBULL 915) indicates that the main ring or tevatron abort loops are down. If any of eight *beam switches* in the MCR are down, no enable will be generated. If the correct conditions exist, an HEP ENABLE will be passed to the to the *prom module* in the preacc control room. As long as the beam switches are up and no abort/ 5° bend fault are present, an HEP REQUEST is sent to the *NTF interlock module* (sec. 11.4, 11.5).

Beam can be produced (given the above conditions) every linac cycle (15 Hz) by pressing the 15 Hz button near the keyswitch module, or commanded by computer via the *Smeds module* next to the keyswitch module. This module is used by the automatic steering program (14.2).

The prom module also receives the NTF ENABLE from the NTF interlock module (NTF enable defers to the HEP enable if both occur simultaneously) and the P-BAR ENABLE, which the abort logic/pulse shifter module generates on tevatron clock events 14 and 17 (more tevatron clock booster resets) if the P-bar abort loop is made up. The first of two proms in the prom module (OPBULL 663) selects the chop width to be used. Note that this prom does not generate the actual widths of the chops. The inputs to the first prom are: H- chopper voltage, 90° magnet status (from a small magnetic switch on the side of the magnet), P-bar enable status, 58° magnet (NTF) status, I- chopper status, NTF enable status, and HEP enable status. The requirements for an HEP chop width are:

<u>H- source</u>	<u>I- source</u>
H- chopper on	I- chopper on
90° bend on	90° bend off
no P-bar enable	no P-bar enable
58° magnet zero	58° magnet zero
no NTF enable	no NTF enable
HEP enable	HEP enable

To get an NTF chop, the requirements are the same except the 58° magnet must be on above some value, there must be an NTF enable, and no HEP enable. A P-bar chop is the same except that there must be a P-bar enable and no HEP or NTF enables. If none of the conditions are met, the choppers default to the standby width. The logic table for the first prom is printed on the front of the module, with the inputs to the first prom as follows:

A0	H- chopper HV
A1	90° bend status
A2	P-bar enable
A3	58° magnet status
A4	I- chopper HV
A5	P-bar enable (same as A2)
A6	NTF enable
A7	HEP enable

The second prom determines whether or not the pulse shifter will shift. Its inputs are: 4-gap disable, linac safety system status, linac microprocessor inhibit, counter disable, linac vacuum valve status, and the chop widths from the first prom. If there is no 4-gap disable (see below), the linac safety system is made up, there is no microprocessor inhibit (see below), and the linac vacuum valves are all open, the pulse will be permitted to occur with the desired chop type.

The linac safety system input to the prom module comes directly from the CARESS racks in the MCR, and it is by means of this input that any safety system in the accelerator complex may turn off linac beam by inhibiting the pulse shifter.

The outputs of the prom module are CHOP TYPE, which is sent to the chop time selector, which takes the times from the third and fourth predets and send them to both 750 keV choppers. Another output, which indicates whether or not an NTF chop is selected, goes to the buncher phase select module (sec. 6) to tell it which buncher phase to select, HEP or NTF. BEAM ENABLE commands the pulse shifter to shift and to permit beam in the linac.

3.3 Pulse Shifter

The pulse shifter module provides the 15 Hz reset pulse for the sources, choppers, and the microprocessor in the preacc control room (secondary #G; sec. 12.2). The reset pulse originates in the same predet that generates the QUADS ON, RF ON, and TDATA (backup) pulses. PULSE SHIFTER IN is delayed 50 μ sec from the real TZERO in order to give the pulse shifter time enough to decide whether or not to shift. This is the source of the apparent 50 μ sec discrepancy between source times and other linac times.

The pulse shifter is a digital delay line (OPBULL 561) of 1 msec duration. The PULSE output of the pulse shifter (the four outputs of the pulse repeater) are the resets for secondary #G, the H- and I- clock generators, and the trigger pulse for the chopper time predets. When PULSE is delayed by 1 msec, the H- and I- source clocks and the secondary #G and chopper timers are delayed by the same amount. The beam is produced normally but doesn't reach the first tank in the linac until about 500 μ sec after the accelerating fields have died away. When PULSE isn't shifted the beam arrives at the linac when the RF pulse is at its maximum, and is accelerated through.

The pulse shifter also counts the number of consecutive beam pulses. If the number exceeds thirteen (the number of booster batches that could be injected into main ring in the days before the tevatron- now it's twelve) the pulse shifter will inhibit itself on a PULSE OVERCOUNT (displayed and reset on the pulse shifter status module in the MCR). If the COUNTER DISABLE input is present, however, this feature

The FNAL Linac

is overridden. COUNTER DISABLE originates in the keyswitch module (if the 15 Hz button is pressed, to allow linac tuning with 15 Hz beam going to the linac beam dump) or the NTF interlock module (if NTF is in operation, to allow uninterrupted production of neutron beam). In either case, the prom module passes it along to the pulse shifter.

If 15 Hz beam is being run to the linac dump, a 'danger exists. If a booster prepulse comes along on the clock (4-gap), the 200 MeV chopper and S1 (sec. 8) will fire, unintentionally sending beam to booster. It is also theoretically possible, with 15 Hz 4-gaps, to run 15 Hz beam to booster, which is somewhat reckless. To keep this from happening, a predet that looks at the "raw" 1 MHz phase-reversal clock produces a pulse whenever a 4-gap comes along. This 1 μ sec pulse is fed to the 4-gap pulse stretcher which makes the pulse long enough for the second prom in the prom module to recognize. This 4-GAP DISABLE pulse will then inhibit the BEAM ENABLE pulse, making beam on 4-gaps impossible.

3.4 Beam Inhibit and Enable Links

Although the computer control system must be relied upon at some level to provide beam-inhibit protection when a critical parameter goes out of tolerance, it was decided to give all the distributed linac microprocessors a direct line to the pulse-shifter circuitry that is independent of control system links. This exists in the form of the beam inhibit line, which originates at the *linac 68k μ P beam inhibit module*, which sends a two-conductor twisted pair all the way down to the far end of the linac gallery (TM 1238). One of these leads is at 12 volts and the other is at ground. All of the distributed microprocessors in the linac control system have a transistor across the leads in open-collector configuration. When the microprocessor wants to inhibit beam, it shorts the two leads together, dragging down the 12 volt line signal and causing the beam inhibit module to send a BEAM INHIBIT signal to the prom module, which in turn removes the BEAM ENABLE input from the pulse shifter. The cause of the inhibit is reported through the control system.

The *H-* and *I-* *beam inhibit light link modules* receive "inhibit beam" commands from the *H-* and *I-* source microprocessors through light links. These modules then clamp the beam inhibit line in the usual way.

In a similar fashion, a BEAM ENABLE from the prom module is also sent to the beam inhibit module. This puts a voltage across another twisted-pair line. With the voltage present, an LED at each microprocessor is lit to tell the μ P that a beam pulse is coming. This is useful for determining the correct baselines for toroid current readbacks.

4 THE ALVAREZ STANDING-WAVE PROTON LINAC

The FNAL linac consists of nine cylindrical, electrically-resonant, water-cooled steel tanks clad inside with OFHC (Oxygen Free, High Conductivity) copper. The tanks are about 16 meters long, except for tank one, which is 7.4 meters long. Each tank is an electrically resonant RF cavity, driven by a 5 MW power amplifier which causes an electromagnetic field to oscillate at a frequency of 201.24 MHz.

4.1 Cavity Fields

The electric field vector points along the long axis of the tank, pointing in the direction of acceleration and then against it, as shown in figure 4.1. This is a simplified example, not taking into account the perturbations of the fields caused by the internal structures of the cavity.




figure 4.1: linac RF cavity fields, TM_{010} mode

The FNAL Linac

The magnitude of the electric field is constant along the length of the tank, except for tank 1. The field in tank 1 is designed to slope from 80% to 120% of average gradient from upstream to downstream to discourage the coupling of beam energy into the transverse directions. The magnetic field component is circumferential in direction and is 90° in phase ahead of the electric field. Since the electromagnetic wave function Ψ is not a product of free charge in the tank, the fields satisfy Laplace's equation $\nabla^2\Psi = 0$. Writing this in cylindrical coordinates and separating variables shows the field components are of the form:

$$R = J_0(kr) \quad k = 2\pi f/c, \quad f = \text{RF frequency}, \quad l = \text{an integer}$$

$$\phi = e^{in\phi} \quad n = \text{an integer}$$

$$Z = \cos(m\pi z/L) \quad m = \text{an integer}, \quad L = \text{length of cavity}$$

TM_{n1m} designates a Transverse Magnetic field with the coefficients n, l, m . In this case, $n = m = 0$ and $l = 1$, so this field is known as the TM_{010} mode.

The radial variation in the electric and magnetic fields are the Bessel functions $J_0(kr)$ and $J_1(kr)$, respectively (fig. 4.2). The first zero of the electric field amplitude ($J_0(2.4048) = 0$) represents the wall of the RF cavity, because the electric field must be zero there.

$$kr = 2.4048 = 2\pi fr/c, \quad f = \text{cavity frequency}, \quad r = \text{cavity radius}$$




figure 4.2: Bessel functions $J_0(x)$ and $J_1(x)$

The FNAL Linac

Thus for a given cavity radius, the TM_{010} mode occurs at the frequency $f = (2.4048)c/2\pi(0.44m) = 260$ MHz. The actual frequency of 201 MHz is lower due to the internal structures of the tank.

Note that the maximum of the magnetic field strength $J_1(kr)$ occurs just before the cavity wall, not at it. All the values for which $J_0(kr) = 0$ are possible wave modes in the cavity; getting the desired mode is a matter of careful design and tuning.

Tuning a cavity to resonate at the correct frequency requires a very precise control of cavity volume. A *bulk tuner* runs along the length of each cavity, resembling a "D" in cross-section (fig. 4.3). The cavity and bulk tuner are both made oversize initially, resulting in an undersized cavity volume; the tuner is then trimmed (and empty space in the cavity increased) until the correct frequency is approximated. Fine control is supplied by a series of copper




figure 4.3: linac cavity cross-section

pistons mounted in the tank wall that can be moved into or out of the tank to vary the volume. These *tuning slugs* are moved during initial set-up of a tank for frequency control and field measurements, and then set in place. One slug in each tank is motorized and is used in a servo loop to control cavity frequency during normal operation (sec. 5.1). Tanks 2-9 each have five fixed and one motorized tuning slug. These slugs are 6" in diameter and have a 5" range of travel. Tank one has twelve fixed slugs, with a thirteenth motorized. This thirteenth slug has a 2" diameter with a 2" range. Tank 1 thus has the smallest tuning range in the linac and is the most vulnerable to variations in the temperature of the cavity's cooling water.

It should be noted that although the active control of the cavity volume is through the tuning slug, this control is effective only if the temperature of the tank, as controlled by the water system, is constant to within a tenth of a degree Fahrenheit or so. If the water temperature takes off, the tuning slug ranges out very quickly and control of cavity volume is lost (sec. 10.3).

4.2 Cell Structure and Synchronous Phase

The interior of each cavity actually consists of a number of resonant cells (fig. 4.4); a cell runs from the




figure 4.4: RF cavity cell structure
 L_n = length of cell n

The FNAL Linac

middle of one *drift tube* to the middle of the next. The relative strengths of the electric fields in adjacent cells is controlled (in tanks 2-9) by the rotation of *post couplers* that stick out from the cavity walls at the cell boundaries. A small tab on the end of the post coupler biases the electric field toward one cell or the other, depending on the amount of rotation of the tab in that direction. (The first cavity in the linac has no post couplers. Cell field levels are determined by the positions of the twelve fixed tuning slugs.)

The length of the cells is such that particles traverse the gap between drift tubes when the electric field vector is pointing in the accelerating direction, and are shielded in the interior of the drift tubes ("drifting" in the absence of electric fields) when the electric field is pointing in the decelerating direction. The particles increase in energy and velocity with every gap crossed (fig. 4.5).




figure 4.5: alternate acceleration and coasting of particles in a linac

The RF period should be equal to the transit time of the particles across the cell. For a particle of charge e in an electric field E_0 ,

$$\text{energy gain} = \Delta W_0 = e \int_{-L/2}^{L/2} E_0 dz = eE_0L \quad L = \text{cell length}$$

If E_0 varies sinusoidally in time, the energy gain becomes:

$$\Delta W_{\text{sin}} = e \int_{-L/2}^{L/2} E_0 \cos(2\pi z/L) dz$$

Of course, this is zero. The trick is that the drift tubes shield the particles at all times except when they traverse the narrow gap g ; this takes a relatively short amount of time with respect to the RF period, so the electric field doesn't vary much. This is equivalent to altering the limits of integration to get a nonzero value for ΔW_{sin} . Define the *longitudinal transit time factor* $T_1 = \Delta W_{\text{sin}}/\Delta W_0$, which is determined by cell geometry ($T_1 < 1$). Then:

$$\Delta W_{\text{sin}} = \Delta W_0 T_1 = eE_0 T_1 L$$

If a particle is made to pass through the center of the cell at some phase ϕ before the maximum electric field, the energy gain becomes:

$$\Delta W_\phi = eE_0 T_1 L \cos\phi$$

There is one ϕ that will make the particle reach the center of the next cell just as the RF has gone through exactly 360° ; this is the *synchronous phase angle* ϕ_s , with energy gain ΔW_s .

If a particle has a phase different from ϕ_s , it will have an energy gain different from ΔW_s ; for a nonrelativistic particle this will result in an incorrect average velocity through the cell with resultant phase shift with respect to the RF when the particle enters the next cell. In the case where ϕ_s is before the RF peak, particles arriving early will gain less energy in the gap and slip back toward ϕ_s . Particles arriving too late will gain more energy in the gap and catch up to ϕ_s . The particle with $\phi = \phi_s$ and $\Delta W = \Delta W_s$ is called a *synchronous particle*. The synchronous phase angle in the FNAL linac is -32° . The stable phase region runs from about $-\phi_s$ to $2\phi_s$ (fig. 4.6), about 105° (TM 279). The synchronous phase angle in tank 1 is a little larger because the gradient is a bit higher.

This stable phase region also defines the length of the *RF bucket*, outside of which particles are rapidly lost.




figure 4.6: stable phase region

A more conventional way to represent RF buckets is to draw them in *phase space*, where the horizontal axis represents particle position (or RF phase) and the vertical axis represents particle momentum (fig. 4.7). The RF bucket is then a curve in phase space. An excellent discussion can be found in Proton Synchrotron Accelerator Theory by E.J.N. Wilson.




figure 4.7: linac RF bucket

If a continuous beam is injected into the linac, only about $105^\circ/360^\circ$, or 25-30%, will be accelerated (sec. 6).

Adding drift tubes to the cells improves the transit time factor by concentrating the field near the center of the cell (fig. 4.8) so that $E_0 L = E_{\text{gap}} g$ ($g = \text{gap length}$). The field has less time to change while the particle crosses the gap.

The total length of the linac, the RF power available, and the maximum surface fields allowed (about 6 MV/m in this case, though average fields are about 1 MV/m) determine the maximum E_0 . The value of ϕ_s may be freely chosen, but it should be large enough to provide a usable region of stable phase space, yet small enough (near the RF peak) to provide




figure 4.8: cell electric field vectors

good acceleration efficiency. Once E_0 and ϕ_s are chosen,

$$dW_s/dz = eE_0 T \cos \phi_s$$

determines β_s (ratio of synchronous particle velocity to the speed of light) and cell length L as functions of distance. There is the "complete" transit time factor, which takes into account the effects of the spatial distribution of E fields in the cavity. For cell n ,

$$L_n = \beta_s \lambda \quad \lambda = \text{free space wavelength of RF (150 cm)}$$

The first cell in the linac is 6.04 cm long; the last is 84.3 cm long, reflecting the change in β_s from injection (0.03) to extraction (0.5). The arrangement of drift tubes in tank 1 is shown in figure 4.9.

The FNAL Linac




figure 4.9: tank 1 drift tubes

The FNAL Linac

It should be noted that this is an important difference between proton and electron linacs. β_S for electrons approaches unity at much lower energies than for protons. Thus $\beta_S = \text{const.}$ for an electron linac. This implies that ϕ_S can be designed to be zero, because it is not necessary to have a region of stable phase; an electron injected at any given phase will stay there. Setting $\phi_S = 0$ for a proton linac would maximize the energy gain for a very small number of particles, but the stable phase region would be so small that overall efficiency would be about zero. Electron linacs are simpler structures and can be made to run up to high energies (example: SLAC). Above a few GeV, proton linacs would more closely resemble electron linacs because the velocity of the protons approaches c , but proton machines of that power would be neither inexpensive nor reliable. The highest-energy proton linac existing at this time (11/85) is the 800 MeV LAMPS facility at Los Alamos, which uses a drift tube-loaded structure at low energies, and a side-coupled structure at higher energies.

4.3 Design Considerations

The RF power required to establish a field E_0 in a cell is:

$$P = (E_0 L)^2 / (ZL)$$

Z = normal *shunt impedance* per unit length. Define the effective shunt impedance $R_S = ZT^2$, where T is the complete transit time factor. Then:

$$P = (E_0 L T)^2 / (R_S L) \propto (\Delta W_S)^2 / (R_S L) \quad R_S \text{ units } M\Omega/m$$

The major problems in the design of the drift tube structure are: 1) the determination of the drift tube and cavity geometry such that the structure will resonate at the design frequency; 2) obtaining the best electric field distribution near the axis (divergence as small as possible); 3) determination of R_S and T . To minimize P we must maximize R_S and T .

It is desirable to use as high an RF frequency as possible, since R_S for a particular geometry varies as $\lambda^{-1/2}$, and the power required to maintain a particular field level varies as λ^{-2} . However, decreasing λ requires that the dimensions of the cell be made smaller; this creates a problem since the borehole diameter b has a minimum size (if it gets too small the beam won't fit through). Also, as the ratio of b to drift tube diameter d increases (fig. 4.4), the effective radial variation of the field near the axis increases, introducing a large radial variation in the transit time factor which couples to radial and phase motions of the particles. Furthermore, it is necessary to incorporate radial focusing magnets (quadrupoles) in the drift tubes which implies a minimum drift-tube size. Finally, since $L = \beta_S \lambda$, there is a practical minimum injection energy for the linac.

The FNAL Linac

Another trade-off involves the ratio g/L . T is optimized by making this small, but making the gap between two drift tubes very small increases the chances of sparking. Optimum values for g/L run from 0.2 to 0.4. Typically, D is fixed, and g/L and d are juggled to obtain a maximum R_s . D is a constant for tanks 2 and 3 and 5-9, but g/L varies from 0.21 in the first cell in tank 1 to 0.47 in the last cell of tank 9.

Cavity Number		1	2	3	4	5	6	7	8	9
Proton energy in	(MeV)	0.75	10.42	37.54	66.2	92.6	116.5	139.0	160.5	181.0
Proton energy out	(MeV)	10.42	37.54	66.18	92.60	116.5	139.0	160.5	181.0	203.3
Cavity length	(m)	7.44	19.02	16.53	16.68	15.58	15.54	15.83	15.88	15.73
Cavity diameter	(cm)	94	90	88	88	84	84	84	84	84
Drift-tube diameter	(cm)	18	16	16	16	16	16	16	16	16
Bore-hole diameter	(cm)	2.0- 2.5	3.0	3.0	3.0	4.0	4.0	4.0	4.0	4.0
Cell length L (First cell)	(cm)	6.04	12.2	41.1	53.3	61.8	68.2	73.3	77.6	81.3
(Last cell)	(cm)	21.8	40.8	53.0	61.5	67.9	73.1	77.4	88.1	84.3
Gap length G (First cell)	(cm)	1.30	4.4	12.2	19.5	22.6	27.1	30.9	34.3	37.3
(Last cell)	(cm)	6.70	12.7	19.3	25.1	26.9	30.8	34.2	37.1	39.7
G/L (First cell)		0.21	0.20	0.30	0.37	0.37	0.40	0.42	0.44	0.46
(Last cell)		0.31	0.31	0.36	0.41	0.40	0.42	0.44	0.46	0.47
Axial transit-time factor (First cell)		0.64	0.86	0.82	0.75	0.73	0.68	0.64	0.61	0.58
(Last cell)		0.81	0.81	0.75	0.69	0.69	0.65	0.61	0.58	0.55
Effective shunt impedance (First cell)	(M Ω /m)	27.0	53.5	44.6	35.0	29.6	24.8	21.5	18.9	16.7
(Last cell)	(M Ω /m)	47.97	44.8	35.2	28.5	25.0	21.7	19	16.8	14.9
Drift space following cavity	(m)	0.22	0.6	0.75	1.0	1.0	1.0	1.0	1.0	--
Number of full drift tubes		55	59	34	28	23	21	20	19	18
Average axial field	(MV/m)	1.60- 2.31	2.0	2.60	2.60	2.56	2.56	2.56	2.56	2.56
Average gap field (First cell)	(MV/m)	7.62	10.0	8.7	7.03	6.9	6.4	6.1	5.8	5.6
(Last cell)	(MV/m)	7.45	6.45	7.2	6.3	6.4	6.1	5.8	5.6	5.4
Peak surface field (First cell)	(MV/m)	9.9	12.6	13.1	12.9	14.0	14.1	14.2	14.3	14.5
(Last cell)	(MV/m)	10.2	9.7	12.9	13.2	14.1	14.2	14.3	14.5	14.9
Cavity excitation power	(MW)	0.61	1.50	2.34	2.44	2.50	2.55	2.65	2.70	2.75
Total power per cavity for 100 mA	(MW)	1.58	4.20	5.20	5.08	4.90	4.79	4.81	4.75	4.68
							Total accumulated length	144.8		
							Total number of unit cells	286.0		
							Total number of full drift tubes	277.0		
							Total cavity excitation power	20.0		
							Total linac rf power for 100 mA	40.0		

figure 4.10: FNAL linac specifications

Since each tank is made up of a series of resonant cells, the decision where to "end" the tanks is somewhat arbitrary. It is necessary for the E fields in the tanks to have a very strict phase relationship to one another, so it would be nice to have all the cells in one tank. This would require an enormous and impractical (and very expensive) RF system. There is also a maximum practical limit for the length of a tank: about 20 free space wavelengths. This is because the separation of electromagnetic standing wave modes decreases as length⁻², and the effects of mechanical imperfections on local frequency errors increases as (length/ λ)², resulting in difficulty maintaining the desired field distribution in the cavity.

The FNAL Linac

Finally, there is a simple trade-off between RF power required and overall length L of the linac. Since

$$PL \propto (\Delta W_S)^2 / R_S = \text{const. (for a given geometry)}$$

one can use less powerful (read: cheaper) RF if the linac is made longer. However, the costs of drift tubes and supporting structures usually make it most practical to obtain the largest reasonably-priced RF systems available. (The RCA 7835 power amplifiers used in the FNAL linac were originally developed for use in radar stations.)

The total cost of a linac is then:

$$\text{cost} \simeq (B/E_0) + (E_0)^2/C + \text{constant}$$

where the first term represents building costs and the second term represents RF system costs.

5 LINAC RF SYSTEMS

Each of the nine accelerating cavities in the linac has its own RF system which drives the cavity fields at the *gradient* (electric field amplitude) and phase necessary for maximum energy gain. A tenth RF station, used for component testing, drives a dummy load.

Every linac cycle (66 msec) the RF systems provide a 400 μ sec pulse for the acceleration of beam. In between RF pulses, the cavity fields collapse. Since the cavities are a highly resonant system ($Q > 60,000$), with a drift tube geometry that assumes a certain energy gain per cell, the degradation of any RF system will strongly affect the transmission and output momentum of the linac; failure of an RF system will make transmission of beam through the linac impossible.

Each RF station consists of a low- and high-level RF system. The low-level system provides a half-watt sine wave signal of the appropriate duration and phase for the high-level system, which amplifies the signal to about 5 MW. This signal is sent via a 9-3/16" coaxial *transmission line* to the cavity, where the magnetic field created by the signal induces surface charges on the drift tubes, which in turn drive the electromagnetic field.

The low-level system (OPBULL 1054) consists of a *master oscillator* and a phase control system, which adjusts RF phase to compensate for: a) cavity tune, b) *beam loading* on the cavity, and c) to maintain the correct phase relationship between adjacent cavities.

The high-level system includes the amplifier chain, associated anode and filament supplies, the *modulator*, the transmission line, and the cavity itself.

5.1 Low-Level RF

A single master oscillator supplies the 201.24 MHz drive signal for all the linac RF stations. This oscillator is located in the NTF control room by system 5. The operating frequency may be read on a *computing counter* there. The oscillator output is fed to a power splitter and runs off in two directions. One output goes to the buncher and systems 1-4. The other output goes to systems 5-10 and the debuncher (see Basic Booster for Beginners by John Crawford for a description of the debuncher). The signal is teed off through an attenuator at each station.

The FNAL Linac

Entering the *frequency control and phase lock module* in the RF system control racks (fig. 5.1) the signal first goes through two phase shifters (SH1 and SH2) that are used in the regulation of intertank phase; the output of these shifters is the desired RF phase for that station. Power splitter PS1 sends the signal to phase shifter SH3, which in turn feeds a mixer, which turns the continuous-wave signal to pulses, as governed by the *waveform generator/sequencer* (not shown). This pulsed RF signal then is amplified by a series of two solid state and three tube amplifiers. The drive signal is then sent via the transmission line to the cavity.

Pickup loops in the cavity measure the magnetic field in the center and at each end of the cavity. The output of the center pickup loop is sent to a power splitter PS2, a phase shifter PS4, and is compared with the "desired" RF phase at the *phase comparator mixer* PC1. The bipolar output of this circuit (*phase lock input*) is designed to be zero when the correct relationship exists between the desired and actual RF phases. The signal is raised by five volts (*phase lock output*) and drives the phase shifter SH3. This counteracts the effects of cavity RF phase shifts caused by beam loading on the cavity (as the beam extracts energy from the cavity), the effects of PA tube and modulator aging, and the vagaries of tuning. This sometimes known as the "fast" feedback loop.

A second loop acts to keep the cavity tuned to the correct frequency. When correctly tuned, the cavity looks like a purely resistive load; the cavity field is in phase with the applied voltage. Any phase difference due to a mistuned cavity can be observed by comparing the RF phase in the transmission line with the RF phase in the cavity. The "slow" loop does just this. A second output of power splitter PS2 goes to phase shifter SH5 and to one input of phase comparator mixer PC2. The other input comes from the *forward power pickup loop* in the transmission line ("forward" is specified because there is also a reverse power loop that measures power reflected from the cavity). The output of PC2 is zero at the desired phase relationship between the inputs, which is set by the frequency adjust potentiometer. This is done after the power amplifier reverse power (transmission line reverse power pickup loop) has been tuned for a minimum.

If the output of PC2 is nonzero, the *frequency control module* will then tell the stepping motor driver to move the tuning slug, which alters the cavity resonant frequency.

The last feedback loop maintains the desired phase relationship between cavities (the distance from one tank to the next may not be an integral number of cell lengths). The low-energy end pickup loop of cavity n is compared with the high-energy end pickup loop of cavity $n-1$. The signal from cavity $n-1$ is first fed through a *mechanical phase shifter* (mounted in a box strung between the

The FNAL Linac




figure 5.1: linac low-level RF system

cavities) which is set so that the output of the phase comparator mixer is zero at the desired phase. The mixer output is filtered, amplified, and sent to the linac computer as *intertank phase*. The computer acts to keep this near zero by controlling a motor-driven potentiometer in the *RF phase adjust module*. The two 180° phase shifters SH1 and SH2 then control the phase at which cavity n tries to run. The greatest phase difference between any two cavities should be less than 2°.

5.2 High-Level RF

The output of the mixer in figure 5.1 ultimately will become the signal that drives the cavity, after being amplified by seven orders of magnitude. In figure 5.2 the signal first passes through a solid-state amplifier with an output of about four watts (sometimes referred to as the *LLPA*, for Low Level Power Amplifier) located in the rear of the RF system control (A5) racks. The signal then travels to the driver (A11) racks, where a second solid-state amplifier (called the first *IPA*, for Intermediate Power Amplifier) boosts the signal to 400 watts (OPBULL 288). This then drives the cathode of the second IPA, an RCA 7651 *tetrode* with an output of 4 kW. Passing through a coupling capacitor, the signal drives the grid of the driver tube, an RCA 4616 *tetrode* with an output of 200 kW. Through another coupling capacitor, the output of the driver passes through a 3-1/8" diameter transmission line to the cathode of the power amplifier. The RCA 7835 *triode* has a peak power output of 5 MW with a duty factor of 0.0075; these powerful tubes cost \$40,000-\$80,000 each. The general layout of RF system components is shown in figure 5.3.

The RCA 7835 operates at a nominal filament current of 6600 amps at a potential of about 5 volts. This is supplied by two double-Y, full-wave bridge rectifiers. Control of the output current is obtained through a brushless motor-driven induction regulator (*inductrol*) connected to the input of the filament power supply. The power supply output current is regulated to about ±0.5% over a 6000-7000 amp operating range.

The *inductrol* varies the input voltage to the filament power supply to control the output current. It thus acts to compensate for variations in line voltage. This is fine, except that the *inductrols* are mechanical devices subject to wear and tear (and smoke and flame); therefore all the *inductrols* are set into the regulate mode only between the main ring reset and the start of main ring ramp. This gating is controlled by the *filament regulator gating module* behind system 6 (OPBULL 121). In its present (3-86) form, it responds only to main ring ramp times. It is destined to be modified for more flexibility to cope with the more elaborate accelerator sequencing that the Tev 1 program has introduced. Binary status is available through the microprocessor for system 6; bit 63C is set to 1 if the *inductrols* are in the regulate mode.

The FNAL Linac

The inductrol and filament power supply are located in the lower linac gallery. Controls for the inductrol voltage regulation loop are on the front panel of the filament supply, but have been disconnected at most stations and are thus not tunable. The filament leads to the PA are water-cooled cables that pass up through the floor and connect to junction blocks at the bottom of the PA. Dropping a wrench across the junction blocks is a bad idea.




figure 5.3: bird's eye view of RF system component layout
 "A" numbers are RF system unit numbers

The capacitor-coupled 5 MW output of the PA drives the 9-3/16" diameter coaxial transmission line (TM 94) which runs to the cavity (fig. 5.4). The transmission line and PA are pressurized with nitrogen at 1.5-2 atmospheres to reduce the chance of sparking. Passing from the rear of the PA, the line runs through the floor to the lower gallery, passes through a $\lambda/2$ range trombone (used to adjust the total length of the line, which must be an integral number of RF wavelengths long), and then runs horizontally through the wall into the linac enclosure.

At this point it is necessary to isolate the nitrogen in the transmission line from the vacuum in the cavity; this is handled by a ceramic annulus, called a *gas barrier*, that fills the area between the conductors of the transmission line but is transparent to electromagnetic fields.




figure 5.4: transmission line layout

Once past the gas barrier, the transmission line runs into the cavity where it terminates in the *drive loop*. The center conductor of the transmission line runs unshielded for six inches, turns two right angles and then is terminated to the outer conductor (fig. 5.5). The drive loop formed by the center conductor does not protrude into the cavity, but is recessed in the wall. The loop transmits the energy to the cavity as the magnetic fields produced by the alternating currents in the center conductor cause charges on the drift tubes to flow, thus creating electric fields.




figure 5.5: RF cavity drive loop

5.3 Modulator

Although the intermediate power amplifiers and driver deliver the RF waveform to the cathode of the power amplifier, the actual pulse duration and power output of the PA are controlled by the voltage on the anode; this is provided by the modulator, the large box to the left of the PA. The linac modulators and PAs were custom-built by the Continental Electronics Manufacturing Company of Dallas, Texas.

The outer envelope of the PA waveform is defined by the modulator waveform (fig. 5.6), which is defined in turn by two modules in the A5 racks. Not only must the pulse be of




figure 5.6: modulator control of RF gradient

correct amplitude and duration, but the sequence of events that turns the RF station on and off must be carefully controlled to prevent damage to the components. For instance, if modulator voltage is applied to the PA's anode when there is no drive on the cathode, the PA could break into oscillation and destroy itself (OPBULL 75). If the modulator pulse starts or stops too suddenly, the resultant abrupt variation of RF drive to the cavity could result in a high voltage standing-wave ratio (VSWR) and arcing across the gas barrier or the PA ceramic (very, very bad).

When an RFON timing pulse is received from the preaccelerator control room (sec. 3.1), the waveform generator/sequencer module (fig. 5.7, 5.8) first looks to see that an interlock enable from the pulse interlock module also exists. If so, after a 5 to 75 μsec delay (MOD START adjustment on the front panel of the waveform generator) a pulse is sent to the driver pulse controller to tell the driver to turn on. At the same time, a gate to the mixer goes high and the permits the drive signal from the low-level system to reach the amplifier chain. The combination of these two should result in 200 kW output of RF from the driver. After another 8.5 μsec, the pulse interlock module looks to see that the driver output is indeed




figure 5.8: modulator pulse timing

above 175 kW. If so, the interlock module sends a waveform enable to the waveform generator; if this signal doesn't appear, the modulator gradient will remain at zero and a driver stop pulse will occur 500 μ sec later.

The modulator pulse begins with a 125 μ sec linear ramp up to the three-volt level. The ramp down portion of the modulator pulse is also 125 μ sec long. The total length of the modulator pulse (200-500 μ sec) is set by the mod stop adjustment.

After a 5-75 μ sec delay (DRIVER STOP adjustment) relative to the end of the modulator pulse, the driver stop pulse turns the driver off. The mixer gate also goes low at this time, removing the low-level signal from the amplifier chain. A second driver off pulse occurs after a fixed 500 μ sec delay relative to the driver start pulse.

There are ten preprogrammed modulator waveforms available, selected through a rotary switch on the front of the waveform generator. A tilt switch affects the flattop portion, allowing it to slope up or down to compensate for any droop in the high-voltage power supply for the modulator, but this has not been needed. There is also a switch that will disable the output of the module to keep the RF system from pulsing.




figure 5.9: modulator pulse-forming circuitry hardware in A5 racks

The *amplitude control module* determines the size of the modulator input pulse and thereby the gradient level in the cavity. It is essentially a variable attenuator. The output of the waveform generator is fed into a digital to analog converter that is controlled by a counter. The counter looks at clockwise (increase D/A output) or counterclockwise (decrease D/A output) pulses that come either from the linac computer or the local knob input. The output of the D/A then is buffered to remove glitches and is sent to the *gradient regulator*. An LED readout on the front of the amplitude control module shows the pulse amplitude in

The FNAL Linac

percent of maximum. A latching pushbutton below the knob zeroes the output of the module.

The gradient regulator loop keeps the cavity gradient within the specified 0.1% tolerance and compensates for beam loading. The gradient regulator module receives the output of the amplitude control module, as well as a gradient signal from the tank and a beam signal from the toroid mounted in the upstream end of the tank. An increase in the gradient signal lowers the output of the regulator while an increase in the beam signal raises it.

A second gradient control loop is handled by the local microprocessor, which samples the A/D channel that represents the cavity gradient once per beam pulse or once every ten seconds in the absence of beam (sec. 14.1, OPBULL 797), and adjusts the setting of the amplitude control module accordingly. This loop is inactive if the gradient's A/D readback falls below 0.8 volts.

The signal then passes from the crate in the A5 racks to the *modulator regulator* module mounted in the front of the modulator itself. The modulator regulator acts to keep the modulator current at the level needed to maintain the desired cavity gradient. It looks at the current flowing from the cap bank to the switch tubes and raises or lowers the input to the modulator to maintain the desired output. A pre-pulse from the waveform generator/sequencer gates the regulator on and off.

The signal is then ready for amplification by the modulator. It is converted to a 200 μ W light pulse and sent through a fiber-optic cable to the first stage of the modulator, which is at a high potential, like the ion source.

Inside the modulator are three nested boxes (fig. 5.10, OPBULL 1056), each the output of one stage of amplification. The three stages are thus shielded from one another, permitting a higher input capacity for each stage. The input waveform from the gradient regulator is sent inside the innermost box; there it is amplified and fed to a *keyer* (a solid-state pulser using a big power transistor) which results in a 1 kV pulse. The output of the keyer is tied to the innermost box, which pulses from -120 V to 1.2 kV.

The first box is tied to the grids of two Maclett 6544 tetrodes. The anodes of these tubes sit at 4.5 kV. The cathode outputs are tied to the second box, which pulses from -2.2 kV to 4 kV.

The second box is tied to the grids of three Westinghouse WL-23646 or ITT F-1123 triodes, known as the *switch tubes*. Their anodes are tied to a 30 μ F capacitor bank in the rear of the modulator which is charged by a 50 kV transformer. The transformer is in turn driven by an SCR-controlled rectifier. Transformer and rectifier are located in the lower gallery.

The capacitor bank stores about 24 kJ of energy at 40 kV. The cathodes of the switch tubes are tied to the third and outermost box, and from there to the anode of the PA. The output of this stage pulses from zero to about 30 kV, turning on the PA to form the RF pulse.

The FNAL Linac




figure 5.10: linac modulator (courtesy Mike Utes)

Around the line from the capacitor bank to the switch tube anodes are three toroids. These are used to measure the current drawn by the modulator.

One toroid is used to sense the modulator current for the gradient regulator circuit. A second toroid is used to provide a signal for the modulator current readback and oscilloscope trace, and to trigger *mod blocks*; the *monitor/inhibitor* module watches the output of the toroid during the pulse. If the current exceeds 400 amps, the pulse is inhibited for the remainder of that cycle by sending a *mod block* signal to the pulse interlock module. Four consecutive *mod blocks* is indicative of a sparking in the PA or the transmission line (OPBULL 692) and causes a *permanent inhibit*. A permanent inhibit both inhibits the pulse and turns off the SCR controller for the HV power supply for the cap bank until a manual reset is done; it

also reduces the setting of the amplitude control module if the system is not reset within 30 seconds. This is done so that when the system turns back on, the reflected power from the cooler and out-of-tune cavity will not be so great.

The third toroid is used to sense a PA crowbar. This requires a current of 600 amps during the pulse or 125 amps between pulses. The *crowbar compare* circuitry, located in a relay rack on the right side of the modulator (behind the PA) receives the crowbar compare level from the waveform generator so that it knows when the pulse occurs. This signal goes high at the start of the modulator pulse, telling the crowbar detector circuit in the modulator to change the crowbar level from 125 to 600 amps. The the end of the pulse, the signal goes low again, and the crowbar level returns to 125 amps.

A crowbar that occurs during the pulse is indicative of a sustained arc in the modulator, usually in the switch tubes. Indicators inside the modulator (left-hand window) will latch when a switch tube arc occurs. The accuracy and usefulness of these indicators has been questioned by some. Exactly how serious a switch tube arc is depends upon the circumstances, and repeated arcs should be brought to the attention of a linac RF specialist. A crowbar that occurs between pulses indicates that the modulator is coming on when there is no drive on the PA cathode, which is potentially lethal to the PA.

A PA crowbar fires the permanent inhibit, and shorts the capacitor bank to ground through a device called an *ignitron*, which is designed to pass large currents in a very short time. Inside the ignitron (fig. 5.11), the cap bank connects to an anode suspended above a pool of liquid mercury, which is at ground potential. The anode and the mercury are in close proximity but do not make electrical contact. Just under the surface of the mercury are two electrodes connected to a thyatron in the crowbar cabinet. A crowbar fires the thyatron, causing the mercury to vaporize. An arc forms in the vapor between the anode and the mercury pool as the current flows to ground. A thirty second timer prevents the pulse from being reset until the ignitron recovers. The ignitron is physically located on the floor in front of the cap bank toward the front of the modulator right next to the toroids. It is warmed by two heat lamps and is thus hard to miss.

After the 30-second timeout an automatic reset occurs; the permanent inhibit has reduced the gradient setting of the amplitude control module to half value, and the computer then runs the gradient slowly back up to 0.8 volts (A/D). The regular gradient regulation loop then takes over if the *auto-gradient* for that station is enabled (OPBULL 797, sec. 14.1). If, after having sent 200 "increase gradient" pulses to the amplitude control module, the RF gradient is still less than 0.2 volts, the computer will stop trying to raise it and operator intervention will be necessary.

The FNAL Linac




figure 5.11: ignitron

5.4 RF System Operation

The cavity gradients and phases have a direct effect on the output momentum of the linac. Nominal values for these are NOT to be changed. Tuning of these parameters consists strictly of keeping them at their nominal values.

The second IPA, driver, and PA tubes are high Q tuned resonators that require occasional tuning, due to tube aging and fluctuations in supply voltages. The forward and reverse power of each stage (the power sent to and reflected from the load) are tuned by adjusting resonant input and output cavities on the tubes. Forward power is tuned to a specific level, while reverse power is tuned for a minimum, matching the impedance of one stage to the next. Reverse power of stage n is tuned by adjusting the input cavity of stage n+1, which affects the forward and reverse power of stage n, which affects the forward power of stage n+1, and so on.

Each tube cavity has two adjustments, tuning and loading. The tuning adjustment actually changes the resonant frequency of the tube cavity, and the loading adjustment changes the coupling between the cavity and its load (the adjacent stage). An overview of the driver racks where most of this tuning is done is shown in figure 5.12.

The FNAL Linac

Tuning an RF system is covered in OPBULL 1050. Some adjustments are relatively insensitive (like driver output tuning) while others are forbidden (PA output tuning). Once all the stages have been tuned, the total phase shift across the amplifier chain will have changed, necessitating an adjustment of the fast feedback loop to zero phase lock input during beam time. The general layout of the RF system control panel is shown in figure 5.13.

Once the procedure in OPBULL 1050 has been followed, connect the PHASE LOCK INPUT signal to the scope. It should be zero during beam time. If not, zero it using a small screwdriver to turn the PHASE REF ADJ pot on the front of the frequency control and phase lock module. PHASE LOCK OUTPUT should then be five volts at beam time.




figure 5.12: All racks (driver PA)

The cavity represents a tuned L-C-R circuit that resonates at 201.24 MHz. The spaces between the drift tubes act like capacitors and the space between the drift tubes and the cavity wall acts like an inductor. When properly

The FNAL Linac

tuned, the capacitive and inductive reactances cancel and the cavity looks like a purely resistive 50Ω load. If not properly tuned at full gradient, substantial power may be reflected back through the transmission line which could damage the PA ceramic and the gas barrier by arcing across them, creating an effective electrical short.

To insure that the cavity is well tuned, the tuning slug tries to keep the PA reverse power at a minimum during beam time (sec. 5.1). An offset may work its way into the feedback loop, though, requiring some manual tweaking: placing the cavity tuner control into manual (fig. 5.13, front of stepping motor driver) run the slug in and out to zero PA reverse power during beam time. Then flip the tuner into auto and watch the reverse power. If all is well it won't move. If it does, manually tune the slug for zero reverse power again and walk around to the back of the module. Using a very small screwdriver adjust the potentiometer on the back of the frequency control and phase lock module until the **FREQ INC** and **FREQ DEC LED's** are both out. Then set the slug control back into auto. The reverse power during beam time should now stay at zero.

Permanent inhibits are indicative of problems in the PA or transmission line. After a permanent inhibit, a local reset should be done and the gradient slowly turned back up while watching for sparking in the PA and transmission line (ears may also be used) and looking for unusual vacuum activity in the form of high current flows in the cavity ion pumps. The operating impedance of the PA should be about 100Ω , which can be calculated from the modulator peak voltage and peak current meters. A calculated value is also displayed on channel **L:MDxLZ**, where **x** is the system number. If the PA impedance is high the filament current may be raised to bring it into line (**OPBULL 692**). The filament current should not exceed 7000 amps without consulting a linac specialist. Note that the current readback calibration is given by a graph posted next to the control switch.

If the RF at a station is off for several minutes, the cavity will cool down and be out of tune. When turning the station back on the gradient should be raised slowly while watching PA reverse power (**OPBULL 74, 596**) and the position of the tuning slug (analog meter on front of the frequency control module).

An occasional PA crowbar from the RF systems is no cause for concern. If one particular station is crowbarring a lot, the usual solution is to turn down the modulator cap bank voltage a couple of kV. This voltage is controlled by a knob on the A5 rack cabinet (fig. 5.13) and read back on the PA **CHARGING CURRENT** meter. Cap bank voltage should not be lowered below 35 kV (49 kV for system 2) because it will get so low that the switch tubes will "max out" during the pulse (remember, the modulator regulator is trying to keep the modulator output the same) and output control will be lost. The modulator voltage waveform will become distorted if this occurs, so watch it while lowering the cap bank voltage.




figure 5.13: A5 racks (RF system control)

The FNAL Linac




figure 5.14: typical RF system waveforms

If problems with the station persist, it may be necessary to turn the system off. The procedure is given in OPBULL 1049, and is briefly given here:

- 1) Turn gradient down (amplitude control module).
- 2) Turn modulator pulse off (waveform generator/sequencer).
- 3) Turn modulator high voltage off.
- 4) Turn EMERGENCY OFF switch to off position. This grounds the modulator's capacitor bank, first through a "soft" ground stick, then a direct short to ground.
- 5) Place PA filament control switch in MANUAL and lower filament current to zero with toggle switch.
- 6) Turn control power off.

Cooling will remain on until timers time out, unless the main breaker is shut off first. Only in an emergency should the main breaker be thrown before cooling shuts off. Note that this procedure is NOT sufficient for safety if the system is to be worked on. Linac specialists will supply further instructions if necessary.

Once repairs have been completed, the procedure for turning back on is:

- 1) Turn control power on.
- 2) Place PA filament control switch in AUTO.
- 3) Wait for small red light next to blue OVERLOAD button to flash. Then press OVERLOAD button. System overload status lights (top of A5 rack) should go out. PA filaments are brought up automatically if the switch is in the AUTO position, and modulator and driver will time out.
- 4) Turn EMERGENCY OFF switch to on position. There are actually five of these switches, and all must be on before modulator high voltage will come on. Switches are located on the A5 control panel, on the front of the modulator, on the left side of the modulator, on the crowbar rack on the right side of the modulator, and on the high-voltage rectifier in the lower linac gallery. (continued on next page)

The FNAL Linac

- 5) Turn modulator high voltage on. Driver high voltage comes on after second IPA screen bias comes on and driver anode supply is ungrounded.
- 6) Make sure gradient is set to zero.
- 7) Turn pulse on, check for nominal driver output.
- 8) Turn gradient up slowly, watching tuner penetration (fig. 5.13) and PA reverse power. If RF has been off for some time, cavity will be cold and out of tune.

This procedure assumes that no difficulties are encountered in reviving the station. A description of the interlock lights is given in OPBULL 1049, but if the problem is not obvious a linac specialist should be consulted.




figure 5.15: modulator exterior

If the MOD READY interlock is lost, it usually means that a circuit breaker in the modulator has blown (fig. 5.15). To access the modulator:

- 1) Turn down the gradient and shut off the pulse and the modulator high voltage.
- 2) Turn the EMERGENCY OFF switches on the control panel and the front of the modulator off.
- 3) Press the door open button and open the front door of the modulator. There is a ground stick mounted just inside the door frame on the right-hand edge.
- 4) While looking away (in case there's an arc) ground out the inner box. Hang the ground stick on the inner box.

(continued on next page)

- 5) Reset the breaker.
- 6) Return the ground stick to its stored position (this will close a microswitch to permit the high voltage to be turned on).
- 7) Close the door, turn the EMERGENCY OFF switches back on, turn the high voltage back on, turn on the pulse, and raise the gradient back to nominal.

Note that although OPBULL 220 gives directions on modulator troubleshooting, any modulator work more involved than resetting tripped breakers should be left to specialists. Modulators are dangerous!




figure 5.16: A3 rack (AC power distribution panel)

If a life-threatening situation should ever arise when working on an RF system, throw the main breaker (fig. 5.16) located in the A3 racks. This breaker is bordered in red and is easy to see. Throwing it will shut down everything in the system immediately. This is a rude shock for many components, so it should be done only in emergencies.

Spare modulator pulse-forming network modules and power supplies are located in cabinets between stations 5 and 6. Other spare items (aside from the ever-ready system 10) are located in a lock-up in the lower linac gallery beneath system 1. A key is available from the key tree in the MCR.

The FNAL Linac

Average values for RF system meter readings are shown here. All RF systems are slightly different, both in tune and temperament, so expect deviations from these numbers.

All racks:

4616 bias	200	VDC
4616 grid	30-50	mA
7651 anode	10-20	mA
7651 anode	3.5	kVDC
4616 filament	0.9	VAC
4616 filament	480	A
7651 bias	15	VDC
7651 screen	.6	kVDC
4616 screen	20-50	mA
4616 screen	1.2	kV (only works at system 1)
4616 anode	17	kVDC
4616 anode	110	mA
4616 filament	0.95	V

A5 racks:

PA high voltage	40	kVDC
PA charging I	0.85	A
modulator peak V	22-25	kVDC
modulator peak I	250	A
PA cathode I	3.1	A
PA filament V	5.2	VDC
PA filament I	6700	A

modulator:

6544 plate V	6.6	VDC
6544 plate I	0.23	A
switch tube bias	2.8	kVDC
switch tube bias	0.3	A

The FNAL Linac




figure 6.1: buncher effect on beam energy and time distribution

The beam pulse coming out of the sources is, depending on the chop selected, anywhere from 10 to 57 μ sec in length. This corresponds to 2,000-11,000 RF cycles. As discussed previously, the synchronous phase angle of the Linac (-32°) determines the stable phase region (105°) of the RF bucket. Capturing 105/360ths of the beam thus yields a capture efficiency of about 35%, which isn't very good.

What is needed is a system that will stuff the particles into 201.24 MHz bunches so that more of them will fit into the RF buckets. This is the function of the buncher. It is a single-gap RF cavity operating at the same frequency as the rest of the linac, but phased so that particles arriving early in the RF cycle see a decelerating voltage and are slowed down, and those arriving later see an accelerating voltage and are speeded up (fig. 6.1).

The ideal buncher for injection would have a single-slope sawtooth waveform, which could bunch beam into arbitrarily small width bunches, and provide 100% capture efficiency (FN 277). This being the real world, sawtooths are difficult to produce at high powers, and so a sinusoidal waveform is used. This reduces the capture efficiency to about 70%. Still, this means the Linac buncher can cram about 240° of beam into a 105° RF bucket.

6.1 Buncher RF System

The buncher RF system, located next to system 1, is similar to the other Linac RF systems, but lacks the driver and power amplifier stages of the amplifier chain—the cavity is driven off the second IPA (RCA 7651). The RF pulse itself is more nearly square, since the much smaller cavity can fill and empty of RF more quickly.

The low-level system consists of some of the same hardware that is found in the other RF systems, as well as some older hardware, which may someday be upgraded. A block diagram is shown in figure 6.2 (compare with figure 5.1). One difference is that there are no intertank phase loops; the buncher phase is set to values commanded by the linac computer.

There are two RF phases associated with the buncher— one used for NTF operation and one used for everything else. The linac computer sends both desired phase settings to the *buncher phase select module*. When the prom module in the preacc control room sends out an NTF chop width, it also sends an NTF ENABLE pulse to the phase select module, which then selects the NTF phase for a time determined by gate pulses from a CAMAC 178 card in the lower linac gallery NTF

racks (sec. 11.4). In the absence of the NTF enable signal it reverts to the HEP phase. The phase select module is the equivalent of the RF phase adjust modules in the other RF systems.




figure 6.2: buncher low-level RF system

Note that in both cases the phase is referenced to tank 1, but the actual number that represents the phase is entirely arbitrary.

Another difference is in the circuitry that forms the pulse (OPBULL 1055). The 3383 sequencer module, when permitted by two interlock modules, generates all the system start and stop pulses, referenced to TZERO. The start and stop times sent to the bias pulser turn the 7651 tube on and off by controlling the grid voltage. Other times are used by the 3384 reference pulse generator to produce a reference

ramp which is shaped and amplified by the 4800 output module. The output module also sums the ramp with the actual cavity field. The output is amplified and sent to the mixer, which both gates and modulates the RF signal passing to the 4-watt amplifier.

Unlike the other RF systems, the amplitude of the buncher RF is varied by changing the size of the input to the amplifier chain, not by changing the cathode-to-anode voltage in the final amplifier.

The copper-plated steel cavity (fig. 6.3), formerly used in the University of Minnesota linac, contains two half drift tubes, drive loop, and slug tuner, all of OFHC copper. It is electrically analogous to a single cell in one of the linac cavities.




figure 6.3: buncher cavity cross-section

6.2 Buncher Operation

By the nature of its operation, the phase of the buncher with respect to tank 1 has a strong effect on linac transmission. Not only can the buncher increase the capture efficiency from 25% to 70%, it can also decrease it to 15%.

The buncher phase is a freely adjustable parameter not fixed in relation to any other system. Usually maximum transmission is desired and the buncher phase is tuned to produce maximum current at the end of tank 9 (L:T09OUT). However this has not always been the case; sometimes it is desired to keep the injector below a certain intensity while main ring or the tevatron is tuned up, but without removing a beam turn from the booster or a booster batch from main ring (usually at the request of experimenters). Then, the buncher is "phased back" to reduce the linac intensity.

It should be noted (OPBULL 1057) that changing the buncher phase changes the position and momentum of the linac beam, and thus the linac steering should be done whenever the buncher phase is changed. Usually the phase is reduced to lower the intensity, but it may also be raised past peak transmission, albeit at the risk of causing strange momentum distributions in the linac (seen as tails on the momentum

The FNAL Linac

wire profile- sec. 8.3). Misphasing the buncher by large amounts is discouraged because it can cause sparking and eventual failure of booster RF stations. Once the buncher is misphased, don't forget that fact. People have tuned booster and main ring for hours on end to increase machine intensity only to discover that the buncher had been misphased earlier and had not been set back.

There are two phase adjustments for the buncher- one for the HEP, standby, and P-bar beam pulses (L:RFBPAH), and one for the NTF beam pulses (L:RFBPAN). These are the D/As sent to the buncher phase select module mentioned earlier. Separate control was given to NTF so that their beam intensity could remain constant in the face of changing requirements for linac/booster intensity. There is no reason to change the buncher phase for NTF pulses unless requested to do so by the NTF staff.

7 QUADRUPOLES

During acceleration of the beam through the linac cavities, the beam tends to blow up due to space charge effects and to a phenomenon known as *RF defocusing*. This occurs because the field lines in the accelerating gap are not parallel, but converge in the first half of the gap and diverge in the second half (refer back to figure 4.8). Since the synchronous phase angle is negative, the fields are increasing during the time of transit across the gap. Thus the diverging electric field has a greater effect than the converging one, and a net defocusing effect results.

To keep the beam size at an acceptable value the beam must be focused in both horizontal and vertical planes during its passage through the cavities. To this end each drift tube in the linac contains a quadrupole magnet, alternating vertically focusing ("D" quad) and horizontally focusing ("F" quad). There is also a quadrupole built into each end of each cavity. Thus a cavity with n drift tubes and $n+1$ cells has $n+2$ quadrupoles. The first quadrupole in tank 1 is vertically focusing.

Quadrupole focusing is critical to the proper operation of the linac, particularly in the first two cavities. If a quadrupole isn't at the correct value, the beam size downstream may become large enough to strike a drift tube and damage it. Therefore every quadrupole is monitored such that it inhibits beam if the current through the magnet is not within a 3 amp tolerance. It is NOT acceptable to run beam through the linac if any quadrupoles are out of tolerance (OPBULL 212).

A cross-section of a typical drift tube and quadrupole is shown in fig. 7.1. Both power leads and cooling water enter and exit through the drift tube stem which protrudes through the top of the linac cavity. The tops of all the drift tube stems are covered by a *stem-box cover* (fig. 7.2) through which power leads and cooling water are fed. Some covers are evacuated to forestall vacuum leaks around the drift tube stems, and in the case of tank 1, to prevent cavity tune shifts due to changes in atmospheric pressure. This is because the drift tubes can, under the influence of air pressure, expand and contract like balloons, changing the gap widths and thus the electrical characteristics of the cavity (sec. 9.2).

The quadrupoles are driven by Acme "Rectifier" supplies (OPBULL 779) located in the upper gallery (some of the supplies for tank 4 are in the lower gallery to make room for NTF equipment upstairs). Each 200 amp supply drives one quad or two quads in series (one F and one D). There are 295 quads and 171 quad supplies in the linac. All the quad supplies are individually controllable. The old *ladder control* scheme, used to conserve D/A channels, is no more, thanks to the upgraded linac control system.





figure 7.1: drift-tube quadrupole

The quadrupole magnets are water-cooled from the same LCW system that controls the temperature of the cavity. Since the available cooling is limited, the power supplies for the quadrupoles are pulsed to reduce the duty factor. Since they are pulsed, they must be synchronized. This is done by the QUADS ON timing pulse from the preacc control room. Every quadrupole power supply is transformer coupled to the cable carrying the pulse.

In each supply, an SCR-controlled bulk supply charges a capacitor bank. The timing pulse initiates a sequence of events resulting in the discharging of the cap bank into the load about 1 msec later. The current waveform resembles a sine wave; a current tolerance of 0.5% at maximum implies a usable pulse length of 100 μ sec. The current pulse flows

The FNAL Linac



QTY	PART No.	DESCRIPTION
1		LINAC TANK
2	ME-4127	D.T. COVER
3	MC-3826 MC-3827	FLANGE SUPPORT
4	CL-4-544	SPHERICAL NUT & WASHER (CARR. LINE)
5	4	#10-24 SOC. HD SCR. S.S. 1/8 LD.
6	MA-3824	CLAMPING RING
7	MA-3810	FLANGE - O RING
8	1	O RING 1/8 X 2 1/2 I.D. FEED
9	1	R.F. CONTACT SPRING RETAINING RING
10	1	#10 SIA. COPPER WIRE 2 1/2 LB.
11	MA-3017	R.F. CONTACT SPRING
12	1	O-RING 1/8 X 1 5/8 I.D. FEED
13	2	1/8 LOCKWASHER S.S.
14	2	1/4-20 HEX HD. NUT S.S.
15	1	MA-3818 CLAMP PLATE
16	2	SP. SW. 85 STRAIGHT TIE CONNECTOR (PARKER) 3/8 S. 22 TURNS
17	1	MA-3825 BUSHING
18	3	1/4-20 SOC. HD. SCR. S.S.
19	1	RETAINING RING (PARKER)
20	1	MA-3821 WATER COLLET
21	3	O RING 1/8 X 1 5/8 I.D. FEED
22	1	D.T. ASSEM.
23	2	1/4-20 HEX HD. NUT S.S.
24	2	WASHER FLAT 1/4 S.S.
25	2	WASHER STAR EXTER. 1/4 S.S.
26	2	YAWG-73 NYLON 1/4 SCREW 11 7/8" SOLID WIRE
27	2	#10-24 HEX. HD. NUT S.S.
28	2	#10 WASHER FLAT S.S.
29	2	WASHER STAR EXTER. #10 S.S.
30	2	YAWG NYLON 7/16 SCREW 11 7/8" SOLID WIRE
31	1	#10-24 STUD 1 1/2 LB. 1/4-20 STUD 1 1/2 LB.
32	1	MA-3964 TERMINAL POST ITEM 1 ITEM 2
33	2	FEED THRU CONAX EXT-125

figure 7.2: stem-box cover

through the load back to the cap bank. Additional SCRs commute after the current peak to reroute the current, charging the cap bank in the original polarity. This way, about half of the energy from the cap bank is saved for the next pulse.

The current waveform of a quad power supply is shown in figure 7.3. The sharp discontinuity is the time of SCR commutation.




figure 7.3: quadrupole current waveform

7.1 Quadrupole Operation

Each supply has eight control cards (OPBULL 779). One of these cards controls the timing of the current pulse relative to the QUADS ON timing pulse. If the pulse comes too early or too late, an overly large maximum current may be required to maintain the desired current at beam time.

When control cards or an entire supply are replaced, the supply should be "timed in" (OPBULL 780) by adjusting the lower potentiometer on the P5 control card while looking at the supply output on an oscilloscope. The output of the supply is sampled from a BNC spigot on the front.

The quadrupole waveform may also be plotted in the control room by plotting the current versus L:TDATA, which controls the time at which the current is sampled. If L:TDATA is changed for such a plot, first disable all the auto-gradient and auto-phase loops in the RF systems (sec. 14) and turn off the beam switch. Then change L:TDATA (expect lots of alarms) for the plot. ALWAYS set L:TDATA back to its nominal value after doing such a plot, and re-enable the auto-gradient and auto-phase loops before turning on beam again.

The FNAL Linac

The quadrupole currents are set and are NOT for tuning by operators. Any change in a quadrupole nominal value or tolerance should be noted in the logbook. Any change to a quadrupole setting should be referred to a linac specialist.

If all the quadrupole power supplies in a cavity trip off, it is a sure sign that the cavity water system has tripped off. If the water system cannot be brought back on quickly, it is necessary to insure that all the quad power supplies are in fact off, since an uncooled quad can burn itself up quickly (sec. 10.3). This is most easily accomplished by turning off the 40-amp wall mounted breakers, which are painted red for easy identification. Panels containing these breakers are located behind systems 1, 4, and in an office between systems 7 and 8.

The FNAL Linac

At the end of the linac, the 200 MeV beam can be steered into either of two beam transport lines (fig. 8.1). One line leads out of the linac enclosure down to the booster injection girder. Details of the construction and operation of this line are given in The Booster Rookie Book, Part II by John Crawford. The other line leads to the system for measuring the momentum spread and emittance of the linac beam. This system will be discussed in more detail in section 8.2.




figure 8.1: 200 MeV area

8.1 200 MeV Transport Line to Booster

After leaving tank 9 the beam passes through a pulsed electrostatic chopper similar to the 750 keV choppers already discussed. The 200 MeV chopper selects beam to be sent to the booster by deflecting it into the field region of a magnetic septum S1 located a few feet downstream. The number of *beam turns* in booster is controlled by the width of this chop. The chopper itself consists of two plates 80 cm long separated by one inch. Both plates are charged to 56 kV before the linac pulse; the chop begins when one plate is grounded in 1 μ sec via a thyatron. At the end of the chop the second plate is grounded in the same manner (fig. 8.2).




figure 8.2: 200 MeV chopper operation

The FNAL Linac

To minimize the chopper voltage needed to shift the beam across the 0.2 inch thick septum plate, the last four quadrupoles in tank 9 are set to focus the beam horizontally to a 1 cm waist at the leading edge of the septum. Dipole trims upstream in the linac are set to position the beam as close as possible to the septum plate. The position and focusing of the beam are monitored by wire scanners W1 and W2 (sec. 8.3).

The magnetic septum S1 consists of an 8-turn water-cooled coil that is pulsed to 1 kA. The resultant 3 kG field bends the beam 9.7° and into the transfer line to the booster.

The chopper and S1 form the first half of an achromatic horizontal bend (OPBULL 397), the second half being provided by horizontal dipole MH1, which bends the beam another 5.9° . Momentum recombination is done with quadrupole Q6. Vertical translation of the beam from the 200 MeV area to the lower elevation of the Booster enclosure is provided by a vertical *dogleg*—two dipoles of equal and opposite bend angles. In this case a 12.8° vertical dipole MV1 bends beam down the *chute* that leads to the booster enclosure, where an identical dipole MV2 renders it horizontal again. This vertical translation is also designed to be achromatic by appropriate settings of quadrupoles Q9-12. A third achromatic set is a horizontal translation and angle change (MH2, S2, and quads Q20-26). Quadrupoles Q7 and Q8 are used for matching between the horizontal and vertical bends.

The achromaticity of the 200 MeV line is dependent on the settings of the quadrupoles, which in turn depend on the line geometry as well as the alignment of the quads with respect to the beam center. These are not variables. Thus quadrupoles in the 200 MeV line, unlike those in the 750 keV line, are NOT intended for casual tuning.

The 200 MeV chopper, S1, major bending dipoles and quadrupoles are monitored not by the linac control system, but by the booster control system.

8.2 Momentum and Emittance Analysis Lines

The beam not selected to go to the booster passes undeflected through the chopper and S1 and on to the momentum/emittance measuring system. Momentum is measured by bending the beam through a 40° *spectrometer* magnet and measuring the beam width downstream with the horizontal wire scanner W5 (fig. 8.3), placed at one *object* of the spectrometer magnet. The spectrometer magnet has a good field width of 10 cm. Quadrupoles Q2 and Q3 are tuned to make the beam nearly fill the horizontal aperture of the magnet; this increases the resolution of the momentum measurement to a theoretical 0.1% by increasing the *dispersion* at W5.

In normal operation the spectrometer magnet runs DC. Any shift of the beam position at W5 is interpreted as a




figure 8.3: linac momentum scan

change in linac momentum (see below).

The spectrometer magnet field is measured by a Hall probe; the output of this probe in volts (seen on the 8-GeV TV system, channel 8) is summed with the input D/A from the computer to produce a drive signal to run the supply. A second, less powerful loop regulates the supply voltage to prevent long-term drifts. Nevertheless, after the spectrometer is turned off for any length of time, it may take about 24 hours to get back to good regulation (the future addition of an NMR probe and another feedback loop may help).

Emittance measurements are usually made only during linac studies. For these measurements, the spectrometer magnet is turned off and the currents in Q2 and Q3 are changed to focus the beam to a waist inside a "straight through" pipe that passes through the spectrometer magnet yoke. Data from simultaneous scans of wires 200-2, 200-3, and 200-4 are used to calculate the beam emittance.

The beam dumps at the end of the momentum and emittance measurement lines are buried in the dirt outside the linac enclosure. Each is a 3 ft diameter steel casting coated by 1.5 ft of concrete. The front face of the steel casings are sloped 3° with respect to the beam axis to distribute the beam energy over a wider area. The casing in the momentum analysis dump is 9 ft long, weighs 15 tons, and is designed to dissipate 3-10 kW. The casing in the emittance measurement dump is 6 ft long and can dissipate 0.3-1 kW.

8.3 Linac Steering and Momentum Analysis

Proper operation of the 200 MeV transport line assumes a specific beam position and angle at the entrance to the

The FNAL Linac




figure 8.4: 200 MeV wire scanner probe

line. To this end scanning wires W1 and W2 have been aligned such that nominal beam position is at the center of the wires. The position is adjusted by two sets of horizontal and vertical dipole trims located between tanks 6-7 and 7-8. Adjusting these trims to center the beam is known as doing the linac steering.

The steering can be done by manually running in the wires, but is usually done by an ACNET applications program on page L36 (sec. 14.2). This program will automatically run the wires through while running 15 Hz beam. The centroids of the beam in both planes are measured (fig. 8.5), and the dipole adjustments that will center the beam are calculated.

Linac steering should be done before any 200 MeV line tuning, before and after any 750 keV line tuning, and once per shift, just for grins.

Momentum analysis is simply a matter of running wire W5 (horizontal) across the beam and observing the beam profile. A change the peak position by 1 cm corresponds to a momentum shift of 1.4 MeV/c (OPBULL 718). The typical momentum spread of the linac is 0.4%; it can be reduced to 0.2%, but doing so degrades the efficiency of the 200 MeV debuncher and hurts booster injection efficiency.

When the beam profile on W5 shows a shift from center, the RF cavity gradients are checked for excessive beam loading. If none is found, and all gradients and intertank phases are at nominal values during beam time, then the linac steering is done. If the beam profile at W5 is still off center, then the field in the spectrometer magnet is changed to center it, and a note is written in the logbook.

The FNAL Linac




figure 8.5: linac beam profiles at steering wires

9 VACUUM SYSTEMS

Maintaining vacuum in the linac transport lines and cavities allows the ion beam to travel through the machine without interference from gas molecules. Vacuum also acts as an electrical insulator, allowing high potentials between objects (such as drift tubes) without arcing. Vacuum in the linac is typically 10^{-7} torr or better. Such low pressures are the domain of the sputter ion pump, which uses electrons to ionize gas molecules. Ions are captured on an anode and complete an electrical circuit. The rate at which ions are hitting the anode is an indication of the absolute gas pressure. Ion gauges, which measure low pressures, work on the same principle. Potentials across the cathodes and anodes in ion pumps and gauges are in the range of several kV.

The ion pumps in the linac are the diode type with titanium anodes. Titanium is good at catching gas molecules, although the anodes wear with time and may eventually have to be replaced, depending on vacuum conditions. Anodes last anywhere from six months to indefinitely.

All the ion pumps in the linac are made of a number of small modules ganged together. The Ultek pumps are made of 25 liter/sec modules and the Varian pumps are made of 30 liter/sec modules.

All the vacuum valves in the linac are electrically controlled and pneumatically operated. Solenoids direct the flow of nitrogen gas that moves the valve actuators. The nitrogen comes from a header that runs the length of the linac; this also supplies gas to pressurize the PAs and transmission lines. Nitrogen is supplied by two LN₂ dewars located at the Central Utility Building (CUB). If the nitrogen supply runs out, all the vacuum valves in the linac will close.

9.1 750 keV Area Vacuum

The vacuum for the 750 keV area is maintained by the Ultek 2400 liter/sec (5000 liter/sec for hydrogen gas) ion pumps at the head of each column, and by the linac cavity pumping systems. The ion pump power supplies are located in relay racks in the preacc control room. Valve interlocks in this system are described in section 2.

Column vacuum is read out in torr on an *ionization gauge controller* mounted in the rack above the ion pump supplies. This meter changes scales (powers of ten) automatically, with an audible "click".

9.2 Cavity Vacuum

Vacuum inside the accelerating cavities is maintained with ion pumps. Each tank has two pump power supplies LEE (low-energy end) and HEE (high-energy end). Each supply powers three Varian 1000 liter/sec ion pumps mounted along the bottom of the tank, which pump directly on the cavity through slots in the cavity floor. The pumps on a common supply are wired in parallel, so the current being drawn by each pump is roughly the supply current divided by three. The ion pump power supplies are located in the racks behind the A5 control console at each station.

Tank 1 is an exception to the above in that it has four ion pumps along the bottom of the cavity instead of six. Each power supply thus feeds only two pumps. Tank 9 has two Varian pumps and two Ultek pumps of 1200 liter/sec capacity, with the other two pump ports blanked off. Two Ultek pumps are available for these ports but are usually held in reserve. The Varian and Ultek pumps have slightly different power supplies.

Each tank has a vacuum valve at each end, except tanks 1 and 2, which are so close together that there isn't room for valves in between- thus tanks 1 and 2 have a common vacuum system.

The upstream vacuum valve for tank 1 is the Linac valve mentioned in the 750 keV section. It is interlocked to the ion pump power supplies in tank 1 such that if both supplies trip off, the Linac valve will close. The HEE valve in tank 2 is likewise interlocked to the pump supplies in tank 2.

In all other tanks, the vacuum valves at each end of a tank will close if both of the pump power supplies for that tank trip off. Each power supply is set to trip if the combined pump current for that supply exceeds 200 mA.

Status of the vacuum valves is monitored by the computer as digital status and the closing of any valve will inhibit beam. Remote control of the vacuum valves is available through the control system (sec. 14.1).

Each ion pump supply has a vacuum monitoring gauge mounted in the rack next to it. This gauge uses the pump supply current to provide its vacuum reading, not an actual gauge (there are none mounted in the cavities). Despite appearances, the gauge does not control the valve. The valve controller (fig. 9.1, OPBULL 416) shows the status of the power supply interlock for each valve, the status of each valve, and provides local control. (Also known as the *gate valve controller*.)

The rear of the crate that holds the vacuum gauge and valve controller also is also connected to the *linlock* vacuum valve status line. This is simply a line coming from a 5-volt power supply in system 9 that is daisy-chained to all the valve controller crates. The line runs to the prom module in the preacc control room, where it forms the LINAC VACUUM VALVE input. When a valve closes, the appropriate valve controller drags this line to ground, inhibiting the BEAM ENABLE output of the prom module to the pulse shifter and thus inhibiting beam.

The FNAL Linac

The ion pump power supply provides the voltage (2-8 kV) necessary to produce the electrons that ionize gas molecules. The meter on the front panel of the supply can read pump current (usually given in milliamps), or the supply voltage. At normal operating pressures (10^{-7} torr) the current and voltage have the usual inverse relation: as the pressure rises the current increases, and the voltage drops. In these ranges the current is the best indicator of the actual vacuum. At relatively high pressures (10^{-3} torr, or below 3 kV), the current stabilizes as the voltage decreases; voltage is then the best indicator of vacuum. There is a function of the meter to show the vacuum value directly but it is of doubtful accuracy.

The most accurate method of determining the cavity vacuum is by reading the pump current, dividing it by the number of pumps it drives, and looking up the equivalent vacuum value on a chart posted on the side of the rack.




figure 9.1: cavity vacuum valve controller

The FNAL Linac

The ion pump power supply voltage is monitored as analog channels L:TKxIPL and L:TKxIPH, where x is the system number; unfortunately, these don't seem to work.

In addition to the main ion pumps for cavity vacuum, tanks 1, 4, and 7 also have small roughing pumps that keep the drift tube stem-box covers under vacuum. These pumps run continuously and are not monitored in any way.

The one-piece stem box cover for tank 1 is kept evacuated to prevent air pressure inside the drift tubes from distorting the drift tube geometry and changing the resonant frequency of the cavity. Although this sounds farfetched, this effect has produced frequency shifts of greater than 10 kHz. This is enough to range out the cavity tuning slug. Therefore, if the tuning slug ranges out and there is no problem with the cavity temperature, the next thing is to peek through the gate to see whether the roughing pump on the stem box cover is still running.

Tanks 4 and 7 have long-standing vacuum leaks around one or more drift tube stems, and the stem-box cover is evacuated with small roughing pumps to help maintain cavity vacuum. If these pumps trip the cavity ion pumps may eventually load down and trip.

The stem-box covers for the other linac tanks can be pumped if necessary but usually remain at atmospheric pressure. The covers for tanks 2-9 are in six sections each, with a vacuum pump port on the top of each section.

The linac tanks are made of three sections welded together end-to-end. The weld joints can develop leaks, so the design of the tank includes *belly pans*; these are small volumes that may be pumped on to keep the area outside of the weld under vacuum. Tank 3 has a leaky weld, so another roughing pump is devoted to pumping out one of the belly pans on this tank.

After major work on a cavity is completed, some provision must exist for pumping a cavity back down to the point where the ion pumps can handle the load. This is done with large pump stations consisting of a vaned roughing pump and a two-stage Roots blower. These are connected to the tank for initial pumpdown and may be removed for normal operation. Two of these stations are mobile and can be moved up and down the linac to the required location. A third is permanently affixed to tank 2. These stations are locally controlled and are not monitored.

9.3 200 MeV Area Vacuum

A map of the 200 MeV vacuum system is shown in fig. 9.2. Vacuum is maintained by ion pumps located underneath the chopper, near Q6, and in the line leading to beam dump #2. More pumps exist downstream of this area in the booster enclosure, but that is beyond the scope of this document. The ion pump near Q6 is separated from the beam pipe by a vacuum valve, but the other two are connected directly to the pipe. In addition to the ion pumps, *turbo carts* can be

The FNAL Linac

valved in to assist in pumpdown. Turbo carts are roughing pump/turbo pump combinations that are manually valved in and out, as well as manually operated.

Vacuum isolation valves provide the ability to work on a part of the 200 MeV transport lines without having to let the entire vacuum system up to air. These valves (VL-T9, VL-S1, VL-02, VL-DMP, and VI-Q9) can be remotely controlled from a rack located in the booster gallery just downstream of linac system 9.




figure 9.2: 200 MeV vacuum layout

Vacuum is monitored in the 200 MeV area by a number of pig gauges that operate on the same principle as ion pumps. The readbacks for gauges PG-S1, PG-BD01, PG-BD02, PG-Q8, and PG-Q13 are in the same rack as the isolation valve controllers. These provide vacuum readings to control the interlocked vacuum isolation valves. When the pressure, as measured by the gauge, increases from the normal 10^{-6} torr to about 5×10^{-4} torr, the valves connected to that gauge will close (OPBULL 29, 513). Valves VL-T9 and VL-02 are connected to gauge PG-S1. Gauge PG-Q8 is connected to valves VL-S1 and VI-Q9. VI-Q9 is also connected to gauge PG-Q13. The status of all the isolation valves, turbo cart valves, and pig gauge status (OK/BAD) is monitored by the booster control system.

In the event of a power outage, or any case where the vacuum is too poor to be able to use the ion pumps, it is necessary to use the turbo carts (OPBULL 267). The procedure is as follows:

The FNAL Linac

- 1) If beampipe is up to air, open valve between beampipe and turbo cart first. If vacuum is only moderately bad, leave it closed. Start roughing pump.
- 2) When vacuum gets down to 30" (as indicated on cart) start turbo pump. Wait five minutes.
- 3) Open valve between turbo cart and beampipe.
- 4) When vacuum gets down to 5×10^{-4} , start ion pumps, turn on pig gauges.
- 5) When ion pump voltage gets above 4.5-5 kV, valve out turbo cart. Turn turbo off, wait for it to spool down, turn off roughing pump.

9.4 NTF Vacuum

The NTF (sec. 11) vacuum system is quite simple. There is a single vacuum valve in the beamline just after the 58° bending magnet. The vacuum is continuous all the way to the beryllium target. The valve is interlocked to an ion gauge in the beamline, but not to linac vacuum. The gauge readout and valve controller are found behind the A5 racks for RF system 4.

There are no pumps in the NTF beamline itself, but a *foreline* runs upstairs to a roughing and turbo pump station behind the rack that contains the valve controller. The turbo pump runs continuously.

10 WATER SYSTEMS

Conventional high-power electrical devices, like those found in the linac, produce a good deal of waste heat. To keep components at operating temperatures many linac devices are cooled using water that has been specially treated to reduce the number of free ions and thereby lower the conductivity. The *LCW* (Low Conductivity Water) systems are treated, temperature-regulated, and pressurized at the Central Utility Building.

Two separate *LCW* systems service the linac: the 95° system and the 55° system. Cooling for NTF power supplies, magnets, and target, and all the 200 MeV devices is provided by the 95° system. Devices cooled by this system connect to the supply and return headers through hoses or copper tubing, and thus rely on the proper supply and return pressures, as well as temperature, to provide sufficient cooling. This system typically runs at about 85° F, with supply and return pressures of 110 and 20 psi, respectively.

The preaccelerators, 750 keV line, RF stations, tank quadrupoles and cavity walls are cooled indirectly by the 55° system. For these systems a number of closed-loop pumping stations exchange heat with the 55° system to regulate system temperatures, but do not draw water from the 55° system continuously.

Three types of modular *LCW* pumping stations are used in the linac. Two preaccelerator pumping stations of the first type service the column water resistors, Haefely power amplifiers and 750 keV transport line elements. Each linac cavity has one of each of the other two types of water station: the RF water systems provide cooling for the components of the RF systems, and the cavity water systems provide cooling and precise temperature control of the cavity walls and drift tubes.

The 55° *LCW* system typically runs at 52° ±4° F. The supply and return lines from CUB enter the lower linac gallery between systems 5 and 6. The remote readback for the temperature is behind system 8, and the pressures are monitored behind system 5.

Maintaining proper temperature in the 55° system is vital to linac operation, as are supply and return pressures (150 and 50 psi, respectively), and, to a lesser extent, conductivity (1.1 μmho/cm or less).

All the linac cooling systems are closed systems that exchange heat with the 55° system to regulate their temperature, and top off their volume with water taken from

The FNAL Linac

the 55° system when necessary. Each has its own heat exchanger, temperature regulator, pump, deionizer bottle, and reservoir. Temperature regulation is achieved by varying the 55° water flow through the heat exchanger to keep the return water temperature constant. In order to operate, the systems must have primary cooling water present in the heat exchangers and sufficient reserve levels in the reservoirs. All three types of systems are located in the lower linac gallery.




figure 10.1: H- 750 keV line water system

The FNAL Linac




figure 10.2: I- 750 keV line water system

10.1 Preaccelerator Water Systems

Two cooling systems service the preaccelerators. The distribution of the load is a bit uneven due to the different ages of the preaccelerator systems (I- predates H- by a few years). One system cools the I- 750 keV line, the I- Haefely power amplifier, and the water resistors for both H- and I- systems. This system is in the lower linac gallery just outside the door to the preaccelerator control room (fig. 10.3). The other system cools the H- 750 keV line and the H- Haefely power amplifier. It is located by the water system for linac RF station 1 and is easily distinguished by its blue front panel. Both systems have local control of pump motors, and local readouts of pressures, temperature, and conductivity.

All of the elements in the 750 keV transport lines are water-cooled, with flow switches to trip off their power supplies if there is insufficient water flow. The status of the flow switches, temperature-sensing *klixons*, and other interlocks are shown at the top of the relay racks that contain the power supplies, located in the preacc control room. The flow switches for the I- 750 keV line are mounted in a box on the wall to the left of the upper door to the I-pit. Valves for the individual elements are located beneath the floor plates there. The flow switches for the H- 750 keV line are mounted in a box on the wall at the west end of the catwalk over the 750 keV line; valves are mounted on the wall beneath the box. Schematics of the H- and I- 750 keV line water flows are shown in figures 10.1 and 10.2.

10.2 Cavity and RF Water Systems

The cavity and RF water systems for a given station have a common control panel with two sets of gauges. System pressures and temperatures can be read and the pumps controlled from here. Note that the large indicator light at the top is lit if that system's pump is off, not on. The RF water system plumbing is immediately behind the control panel, and the cavity water system is toward the rear wall of the gallery. The RF system conductivity meter and a duplicate set of controls for the RF system pump are located in the A7 racks in the upper gallery.

The RF water cooling systems (fig. 10.4) each remove about 150 kW of waste heat from the RF components. Water temperature is regulated at about 80°F. If the temperature exceeds 120°, the system automatically shuts off.

A master control station in the A7 racks regulates, monitors, and interlocks the flow of water to the various components of the RF system (fig. 10.4). Conductivity is monitored and displayed by a meter here. Flow control is accomplished by using variable area valves for each circuit. The supply valve for each circuit is on the front of the

The FNAL Linac




figure 10.4: A7 rack (RF water system control panel)

cabinet, and the return valves are in the rear; some circuits have more than one return line. If a system is to be isolated, close the supply valve first, then the return (get this backwards and you're in for a shower). Each individual circuit may be purged of water with nitrogen gas.

All water-cooled RF components (fig. 10.5) are interlocked to their power source by flow and thermal switches in the return lines from each circuit. The interlock cards are located behind a door at the top of the panel. Spare cards are located there also, as well as the cage in the lower linac gallery by system 1.

The power removed from each cavity by the cooling system (fig. 10.6) is determined by the combination of RF losses and quadrupole power required by each cavity, which range from 10 kW for tank 1 to 25 kW for tank 9. To maintain a stable cavity resonant frequency (at least within the range where the tuning slug can maintain control), the supply temperatures for these systems is maintained to $\pm 0.1^\circ\text{F}$ of a preset temperature between 72° and 78° F by a Honeywell temperature controller located behind the A5 racks. The loop set point, present value, and recorded values are presented there. The closed servo loop regulates the flow of cold water to the primary side of the heat exchanger.

The FNAL Linac




figure 10.5: RF water system




figure 10.6: cavity water system

The FNAL Linac

10.3 Water System Operation

The fact that the linac is very sensitive to variations in water system temperature cannot be overstressed. Tank 1, with its small tuning slug, is particularly sensitive to temperature fluctuations. In fact, temperature problems are usually heralded by the tuning slug in tank 1 hitting the stops with resultant alarm: SYSTEM 1 FREQ CCW LIMIT (although this may have another, interesting cause: sec. 9.2).

Sometimes this can be remedied by altering the set point on the Honeywell temperature controller (0.1° or less, please) up or down depending on which way the slug is ranged out (warmer to lower the frequency of the tank, colder to raise the frequency). This is done by opening the front panel of the controller and changing the position of the green arrow with the thumbwheel. The knob on the front of the controller is a direct manual control of the temperature-regulating valve and will only cause wild oscillations if turned. Don't try it.

More often, however, the problem is with the temperature of the 55° system itself. The 55° system water temperature L:LCWT is set to alarm below 48° or above 56° F. 55° system temperature alarms should be investigated by CUB personnel or the duty mechanic at once. If the temperature continues to rise above the set limit:

- 1) Shut off the beam.
- 2) Run down gradients, turn off pulse and modulator HV and run PA filaments down to zero at all stations.

If the temperature continues to rise, call the linac system specialists. In extreme cases:

- 4) Shut off PA filament power, control power, and all other RF power supplies at all stations.
- 5) Shut off all quadrupole power supplies (this is done most efficiently by turning off all the wall-mounted breakers located behind systems 1, 4, and in an office between systems 7 and 8).
- 6) As a last resort turn off the pumps for the RF and cavity water systems.

The above procedure is also used in the event of hose leaks in the modulator or the downstairs water system. A cavity water leak is cause for shutting down the system and calling linac system experts immediately!

When the system shuts down (or is shut down) it must be manually reset with the switch at the top of the A7 rack. If the system is reluctant, check: thermal overload switch, circuit breakers for motor contactors, and the system instrumentation breaker.

The FNAL Linac

Occasionally it is necessary to make up volume in one of the water systems by filling the reservoir from the 55° LCW system. The procedure is as follows:

- 1) Set Auto/Manual switch on control panel to AUTO.
- 2) Open valve labeled 55° HEADER.
- 3) Water level on reservoir sight glass should rise and stabilize at correct level when automatic valve closes. Close 55° header valve, and return system to MANUAL.
- 4) If level in reservoir falls, close 55° header valve. Locate flow gauge next to valve labeled VALVE #2. Slowly close valve #2 until indicator in flow gauge drops 1/8". This reduces the system pressure to less than that of the 55° header so that the system may draw water from the header.
- 5) Open 55° header valve and fill reservoir to correct level.
- 6) Close 55° header valve, open valve #2, and return system to MANUAL.

11 NEUTRON THERAPY FACILITY

The Neutron Therapy Facility (NTF, or the old designation CTF, for Cancer Therapy Facility) uses high energy neutrons for the treatment of malignant tumors. Neutrons are produced by bombarding a beryllium target with 66 MeV H⁻ ions and exposing a patient to the collimated neutron beam under tightly controlled conditions. The first facility was started up in July 1975. Patient treatment with a rebuilt and more sophisticated setup began in September 1976 and continues today.

The operation of the linac as a booster injector takes about one second during the accelerator cycle. In the days before P-bar, the linac was left idle much of the time and much beam could be devoted to NTF operation. More recently, rapid cycling of the booster as an accumulator/debuncher injector has created a greater demand for linac beam.

11.1 Beamline

The linac gallery floor plan determined the location of NTF, which required the use of an existing freight elevator at the junction of tanks 4 and 5. The beam energy is 92 MeV at this point but the strongest conventional pulsed magnet that will fit in the 1-meter space between the tanks and still bend beam through the required angle can only bend beam of 65-70 MeV. The solution is to accelerate NTF beam pulses through the first three linac cavities and then permit the ions to drift through tank 4 without additional acceleration. This is accomplished by delaying the timing of the RFON pulse for system 4 by 1 msec so that the RF pulse does not occur until after the beam has passed. The tank quadrupoles pulse at the regular time.

The ion beam is extracted by a pulsed C-type magnet L:C58DEG which bends the beam through 58° to miss tank 5. The current in this magnet is ramped down when not sending beam to NTF to permit normal linac operation, taking about 0.2 second to change states. A second set of coils in the same magnet driven by the supply L:C58VER bucks out any residual fields when the main supply is ramped down (OPBULL 400). A Hall probe in the magnet aperture measures the magnetic field and may be read out through the control system. A 32° H-magnet, L:C32DEG, completes the 90° bend to send the beam through a transport line that passes through the wall of the linac enclosure. This supply runs DC.

The transport line contains seven quadrupoles and two horizontal trims, as well as a toroid for monitoring beam intensity. The transmission efficiency of the line is typically 95%. Although the line is not achromatic, it is

The FNAL Linac




figure 11.1: NTF beamline

designed to have zero dispersion at the target (TM 556). In this way the effects of variations in beam momentum and focusing are minimized. Total length of the NTF transport line is about 20'; power supplies are located in the lower linac gallery below RF system 4.

11.2 Target and Collimator

At the end of the transport line, the beam is focused to a 6 mm diameter spot, limited by the gradients available in the beamline quadrupoles. A tantalum collimator (fig. 11.2) with a 15.9 mm bore prevents off-axis ions from striking the target. The target itself is a beryllium disk (TM 928) 25.4 mm in diameter and 22.1 mm thick. Beryllium is used due to its low atomic number (increases neutron production) and good mechanical properties (easier to work with than lithium).




figure 11.2: target and mounting block

The target removes 19 Mev from the incident ions, with the residual energy being spent in a gold disk that backs the beryllium block (TM 834) and improves thermal conductivity to the heat sink. The target assembly is housed in a water-cooled aluminum holder, which is electrically grounded. Experience has shown that targets last at least five years before they must be replaced due to cumulative radiation damage.

The neutrons pass through a *primary collimator* (fig. 11.3) in the form of a steel cone 12.7 cm long. At the downstream end of this collimator two *air ionization chambers* are used to monitor the neutron flux. These chambers consist of aluminum plates (alternating high-voltage and signal planes) separated by air. As the neutrons ionize the air, charge accumulates on the signal planes. When corrected for chamber temperature and barometric pressure (the interior of the chambers are open to the atmosphere) the charge on the plates is proportional to the neutron intensity. During treatment the ratio of the output charges between the two chambers and the ratio of the beam current (measured by the toroid in the transport line) to the output charge of the first chamber are monitored. If the ratios fall outside a narrow range, beam is inhibited.

The FNAL Linac




figure 11.3: NTF collimators

The output of the chambers is calibrated to the output of a tissue-equivalent ionization chamber (TM 834) that is itself calibrated using a cesium-137 source in a fixture of known geometry. These two calibrations are performed at the beginning of every treatment day, and the output of the cesium-137 source is measured once a week, using a chamber calibrated by the National Bureau of Standards on a yearly basis (TM 1018).

Downstream of the primary collimator, the 94 cm long *main collimator* shapes the neutron beam to the desired size. A steel cylinder is fitted on the inside with a series of concentric *Benelex* liners (similar to G-10), which surround the collimator. Collimators are formed of a mixture of 1/8" polyethylene pellets, cement, and water. The Benelex liners can be removed to permit collimators with various size cross-sections to be used. Neutron beam cross sections, as measured in the plane of the area to be treated, can be up to 30 x 30 cm². The entire main collimator can be rotated about its long axis if a particular orientation is required. Teflon wedges may be attached to the downstream end of the main collimator to act as neutron attenuators or "prisms" to produce neutron flux gradients across the width of the subject.

11.3 Patient Setup

The treatment room is built on a freight elevator (fig. 11.4). Entrance, simulation, and patient positioning take

place on the upper level. For actual treatment, the room travels to the lower level where the neutron beam enters from the wall-mounted collimator. A geometric reference point at each level (*isocenter*) used in patient setup and alignment is defined by two orthogonal pairs of HeNe laser beams. The geometric relation of an X-ray source to the upper-level isocenter (or X-isocenter) is the same as the relation of the neutron source (assumed to be in the center of the target) to the lower-level isocenter (N-isocenter) with a 180° difference in orientation (fig. 11.4). The patient is usually positioned such that the tumor to be treated lies at the isocenter.




figure 11.4: NTF treatment room geometry

The patient is immobilized in the restraining chair usually by means of a custom-made cast. The areas where the neutron beam will enter the patient ("portals") are determined by diagnostic radiographs (X-ray pictures), using the known geometry of the X-ray source and the X-isocenter. The positions of the restraining chair (adjustable in rotation, elevation, and two orthogonal directions) are adjusted to produce the desired field cross-section at the *isoplane*, the plane containing the isocenter and perpendicular to the neutron beam axis. In most cases, the transition from portal to portal is done by rotation only. The restraining chair coordinates at each portal are recorded so they may be easily reproduced.

Reference points are drawn on the patient's skin (or cast) using the alignment laser beams in the upper level. The patient is then lowered to the neutron level, and the chair rotated through 180° and translated to the N-isocenter. The alignment laser beams at this level then duplicate the geometry of the lasers at the upper level, and

the distance from the N-isocenter to the neutron target is the same as the distance from the X-isocenter to the X-ray source (190 cm). The lasers thus permit quick confirmation of proper alignment of the patient, in addition to the known coordinates of the restraining chair. The position can be further verified by taking neutron radiographs and comparing them to the original diagnostic radiographs (TM 834). As a final test, a small *field lamp* illuminates the patient through the main collimator by means of an aluminum mirror placed between the primary and main collimators at a 45° angle to the beam axis.

Once the correct collimator is in place, the air ionization chambers calibrated, the patient properly positioned, and the linac and NTF transport line are ready (see below), treatment may begin.

Individual neutron therapy programs average about four weeks in duration, with two to three treatments per week.

11.4 Control System

The treatment apparatus is controlled by *medical* and *beamline microcomputers* (fig. 11.5). The beamline microcomputer monitors transport line dipole magnet currents, beam current, interlock status, commands the phase shifter for linac RF system 4, and handles communication with linac secondary #C. It also measures barometric pressure and integrated ionization chamber voltages which are then used to calculate dosages (expressed in *monitor units*), given parameters loaded from the medical microcomputer. The beamline microcomputer monitors neutron dosages during patient treatment, automatically terminating exposure when the desired dosage is reached. The beamline microcomputer, beamline power supplies, measuring and integrating electronics, A/D converters, timers, and interlock systems are located in the lower linac gallery near the water system for RF station 5.

The medical microcomputer in the NTF control room interfaces with the radiation therapy technologists via a console, and performs calculations for cesium-137 source and air ionization chamber calibration. Auxiliary hardwired *scalers* and timers in the same racks will inhibit beam if the beamline microcomputer should fail to end the exposure when the correct number of monitoring units have accumulated.

Any unusual variation of 32° or 58° magnet current or ionization chamber response during treatment will cause the beamline microcomputer to inhibit beam via the NTF interlock system (sec. 11.5). Quadrupole currents in the transport line are monitored by secondary #C and will also inhibit beam via the beam inhibit line (sec. 3.4) if they wander out of tolerance. Once all the necessary conditions for a beam permit have been satisfied, actual exposure begins when two beam request switches on the *NTF beam control module* are pushed simultaneously.

The FNAL Linac




figure 11.5: NTF control system

The FNAL Linac

11.5 NTF Interlocks

The following information on the NTF interlock system was the best available at the time of this writing (3-86). This system is due for an upgrade in the near future; hopefully it will be properly documented.

When a beam request and appropriate set of conditions exist at the prom module, the pulse shifter will permit beam in the linac. An elaborate system of interlocks (TM 613) allows an NTF beam request to occur only under the proper conditions.

The task of collecting interlock data is allocated to the *NTF interlock module* (fig. 3.1), also known as the *Gannon box*. This box receives all kinds of hardwired information (see below) as well as information from the beamline microcomputer. This information is displayed on the front of the box and at two remote *NTF status modules*, one in the NTF control room and one in the MCR. Be forewarned that all three of these boxes are labeled differently- an interlock's name may change from box to box. The interlocks listed here follow the names on the module in the MCR.

There are three levels of interlocks in the NTF system, each a superset of the previous level. The ANDed sum of the first level of interlocks is a *ramp enable*, which will permit the 58° bend magnet to ramp. All the inputs to the first two levels follow two separate and redundant electrical paths, labeled "A" and "B". Thus there is an "A ramp enable" and a "B ramp enable" that look at the same interlocks. Both are required to permit the 58° supply to ramp.

Among the second-level interlocks is one called CURRENTS NOMINAL. This is an indication that the 58° and 32° supplies are on at the proper level, which means that the A and B ramp enables must exist. These enables are thus inputs to the second level of interlocks, which are ANDed together like the first level in two redundant electrical paths to produce *system ready* indications, "A system ready" and "B system ready". These are then ANDed together to form an input to the third level of interlocks.

If all three levels are made up, NTF will receive beam if the linac is physically capable of producing it.

The the status module in MCR relay rack MCRR #1 shows the status of the NTF interlocks. An LED next to each input will be out if that input is "bad", inhibiting NTF beam. Starting with the bottom (first level) interlocks (all inputs hardwired unless otherwise indicated):

A RAMP ENABLE	"A" system first-level sum.
B RAMP ENABLE	"B" system first-level sum.

TIMED ON	Gate, defining the times when the 58° may ramp, exists. These times are generated by a CAMAC 178 card in the rack above the Gannon box, and are referenced to tevatron clock events 1A and 1B. This card also
----------	---

The FNAL Linac

generates the gate for switching the buncher RF phase for the NTF pulse (sec. 6.1). Also see description of "NTFP" bit on help pages for the Time Line Generator, ACNET page D69.

- LINAC SECURE Linac safety system status, hardwired from CARESS module (sec. 13.2, fig. 5.1) in MCR.
- NTF SECURE NTF safety system status.
- MCR RAMP SWITCH Toggle switch to right of LED. If off, will inhibit 58° from ramping.
- ~~MED. μ P1 ENABLE~~ Also SCALER ENABLE. Scaler module in NTF control room is on, providing backup dose limiting in case beamline μ P fails to end dose at correct level.
- Beamline μ P*
- MED. μ P2 ENABLE Sum of 8 bits from beamline μ P (also DOSE OK, BEAMLINE μ P ENABLE, or CB2). Includes inhibits for missing beam pulses, bad ion chamber ratios, treatment limit reached, and bad μ P power supplies.
- MED. ~~μ P1 OK~~ *15 Hz* Also TO OK. Ion chamber integrator clamp pulse (1947 μ sec.) appearing. From predet above Gannon box, referenced to TZERO.
- MED. ~~μ P2 OK~~ *+5V* +5 volt power supply in Gannon box.

The second-level interlocks are as follows:

- A SYSTEM READY "A" system second-level sum.
- B SYSTEM READY "B" system second-level sum. Two unlabeled red buttons on the front of the Gannon box will remove either the "A" or "B" system ready to test the interlock chain. This inhibits NTF beam, of course.
- LINAC COMPUTER OK 15 Hz interrupt from secondary #C's CPU (not timer board- sec. 12.2).
- CURRENTS NOMINAL 58° and 32° magnet currents within limits for NTF beam. (A,B ramp enable required.)
- NTF VAC. VALVE Vacuum valve downstream of 58° magnet open.
- COMPUTER INHIBIT Also LINAC MONITOR. Status of linac beam inhibit line (sec. 3.4).
- TARGET H₂O OK Target cooling water flow OK.
- RF4 INHIBIT Linac RF station 4 pulse properly delayed by 1 msec to allow beam to coast through tank 4. From NTF interface module at system 4.

The FNAL Linac

58° CURRENT Also 58° MAGNET INHIBIT. Status of external interlock chassis for 58° magnet. From rack above Gannon box.

This section also contains a switch labeled CONTINUOUS BEAM. This enables the 58° magnet to stay on DC, ignoring HEP enables from the keyswitch module in the MCR (sec. 3.2), thereby inhibiting HEP beam.

The third-level interlocks are as follows:

TREATMENT ENABLED Sum of all interlocks. Linac will send beam to NTF when this is lit, if beam is possible.

MEDICAL OR TUNEUP KEY Key must be turned in either tuneup keyswitch (on status module in MCR) or medical keyswitch in NTF control room. Key issued by MCR crew chief.

MCR BEAM SWITCH Toggle switch to right of LED. Inhibits NTF beam, but not 58° magnet ramping.

MEDICAL BEAM REQUEST Request from pushbuttons on NTF beam control module to initiate beam.

NTF PULSE RATE Beam enable from beamline μ P.

SYSTEM READY Sum of first and second-level interlocks. This is the signal passed through the tuneup and medical keyswitches.

At the bottom of the NTF status module are two keyswitches. The MCR keyswitch must be in the ON position for NTF beam to occur. The tuneup keyswitch allows beam when the key is turned and the tuneup beam switch above it is pressed. This is used when tuning up the transport line. The same key is turned in the medical keyswitch in the beam control module during patient treatment.

11.6 RF4 Pulse Delay

When the 58° magnet is enabled, the beamline μ P sends a signal to linac RF system 4 to delay the RF pulse. The delay system consists of two modules in the racks behind the A5 racks for system 4. At 1700 μ sec into the cycle (55 μ sec before the normal RFON pulse) the first module looks to see whether an NTF enable exists at the input from the beamline microcomputer. If it does, the NTF enable output to the second module is latched, to prevent the system from "changing its mind" during the beam pulse. When the second module receives the NTF enable, it delays the incoming RFON timing pulse from the preacc control room by 1 msec before passing it on to the low-level system. When NTF beam is enabled, the first module also sends a delayed trigger

signal to the A/D converters that are responsible for RF system readbacks, so that they sample at the correct time with respect to the RF pulse.

To determine whether the RF for system 4 is actually being shifted, *NTF interface modules* at systems 3 and 4 look to see when the RFON pulse occurs at their respective systems. The module at system 3 sends its pulse to the module in system 4. If the two pulses differ in time by 1 msec, a line from the module in system 4 to the Gannon box goes high, permitting NTF beam. If the signals don't differ by 1 msec, the line goes low, inhibiting NTF beam by removing the RF4 INHIBIT input from interlock chain. This removes the NTF enable output from the NTF status module in the MCR. The interface modules are located in the crates containing the modulator pulse-forming circuitry at stations 3 and 4.

Another NTF enable signal is sent to the low-level system for the buncher (via the prom module in the preacc control room) to select the RF phase used for NTF beam pulses. The duration of the shift in phase is controlled by the CAMAC 178 card in the rack above the Gannon box, which sends its times to the buncher phase select module at the buncher RF station (sec. 6.1).

11.7 NTF Operation

The responsibility for operating the Neutron Therapy Facility beamline lies with the staff of NTF. The responsibilities of the operations group as regards NTF are to keep the preaccelerator, 750 keV transport line, buncher, RF systems 1-3, and tank quadrupoles in systems 1-4 in operation (OPBULL 923, 1047). During actual patient treatment time, these items have top priority.

Daily operation of NTF begins with the issuing of three keys from the MCR. The NTF beam enable key (brown tag), the NTF reset and enable key (red tag), and the air door override key (white tag). The beam enable key is turned in the MCR keyswitch on the NTF status module and is left there for the duration of NTF operation. The reset and enable key is placed in the tuneup keyswitch during tuneup of the transport line, and in the medical keyswitch in the beam control module in the NTF control room during patient treatment. The air door override key allows the large pneumatically-controlled shielding door on the upper level to remain partially open during treatment. This is because shielding requirements are met by closing a smaller and more easily opened door on the upper level. The air door override key fits in the NTF radiation safety system module in the NTF control room.

Since NTF does not depend on RF system 4 or systems 5-9 to operate, a beam inhibit from one of these areas may be bypassed in order to run NTF beam. However, HEP beam should NEVER be run under these conditions. A prerequisite for bypassing an inhibit for the benefit of NTF is the removal of the HEP enable key from the abort logic/pulse shifter

The FNAL Linac

interface, which will disallow HEP and P-bar beam enables. This key is then locked up in the Crew Chief's cabinet until the inhibit monitoring is restored.

12 CONTROL SYSTEM

The text of this section borrows heavily from TM 1238, written by Mike Shea.

The linac control system (fig. 12.1) consists of seventeen *secondary* microcomputers that converse on a fiber-optic *serial data link*, and a *primary* microcomputer which controls the link and communicates with outside *host* computers. The secondary stations are physically distributed throughout the linac (fig. 12.2); each major equipment area has its own dedicated microcomputer. The primary station is located by RF system 6.




figure 12.1: linac control system

12.1 Link

The serial communication link is a single fiber-optic cable that carries *Manchester* encoded data signals using

IBM's *SDLC* (Synchronous Data Link Control) protocol. Manchester encoding consists of phase reversals of a periodic signal that operates at a specified frequency, in this case 1 MHz. The link starts and ends at the primary control station, carrying signals in one direction around in a loop (fig. 12.2). Each secondary station has a *fiber optic link repeater*, which receives and decodes signals from the upstream link, then encodes and transmits data to the downstream link. The order of the secondary stations on the link is:

G H I B 1 2 3 4 C 5 6 A 7 8 9 E D

where G = ground station (H-, I- Haefely control, 750 keV lines)
 H = H- source
 I = I- source
 B = buncher
 1-4 = RF stations 1-4
 C = NTF (formerly CTF)
 5,6 = RF stations 5,6
 A = RF station 10 (test station)
 7-9 = RF stations 7-9
 E = 200 MeV area
 D = debuncher




figure 12.2: *SDLC link routing*

The fiber-optic system was chosen for its immunity from noise (important for operating around RF systems) and the simplicity of communicating with the H and I secondaries located in the -750 kV domes.

12.2 Secondary Station

Each secondary (fig. 12.3) is a complete miniature control system consisting of a Motorola MC68000-based single-board computer, 32K of magnetic (nonvolatile) RAM, two nine-byte binary I/O boards, a nine-channel timer board,




figure 12.3: secondary microprocessor architecture

and a communication board used for talking to the link, all housed in a *Multibus* bin (fig. 12.4).

The CPU board is a FNAL design that runs at 8 MHz. Secondary programs are contained in onboard ROM and are the same for all secondaries; thus all secondary CPU boards are interchangeable.

The magnetic-core memory board (CMOS RAM in some secondaries) contains fixed data tables and local database information in the first 12K bytes. This data is specific to a given station, so memory boards are not directly interchangeable. The remaining 20K are used for device readings, setting, nominals, tolerances, and SDLC communications.

One of the binary I/O boards (the *blue dot board*) is used for normal binary interface and the other (the *red dot board*) is connected to *OPT022* devices providing one byte each of optical input and output, and to external D/A and A/D converters (the number of these converters varying from station to station).

The normal binary interface provides status and control of relays, stepping motors, and vacuum valves. A typical RF station requires seven bytes of binary I/O. Secondary B has a second blue dot board used to drive the stepping motors for the emittance probes at the downstream end of the 750 keV transport line.




figure 12.4: secondary Multibus crate

The OPTO22 optically-coupled binary interface is used for signals that require isolation from local grounds or need to be amplified for special purposes, in particular the local connections to the hardwired beam enable line, and the beam inhibit line, which allows any secondary to inhibit beam without relying on the SDLC link (sec. 3.4).

External 16-channel, 12-bit, sample-and-hold A/D converters are used to digitize analog signals for monitoring. Each secondary has one to three of these A/D converters; usually all the A/D's in the linac are triggered simultaneously by the TDATA timing pulse. The analog input range for these A/D's is -10 to 10 volts. Secondaries that control RF systems (B, 1-9, and D) also have one A/D chassis with an input range of -2.5 to 2.5 volts for 4x resolution of RF signals. These A/D's are identified by red labels on the front panels.

All secondaries except A, D, and E also have external 16-channel, 10-bit D/A converters with 0-10 volt outputs. These are used to control source parameters, linac quadrupoles, and 750 keV line devices.

The D/A and A/D chassis are located in the rack below the multibus crate, except the H and I secondaries, which have D/A and A/D boards mounted in the Multibus crate, which is itself inside a Faraday cage to shield against electrical noise.

The FNAL Linac

Each secondary has a pre-det timer board with a 2 MHz clock (phase locked to the hardwired 1 MHz booster clock coming from the pre-acc control room) and nine timers that may be independently set in 0.5 μ sec increments to maximum delays of 16384 μ sec. Three of these timers have divide-by-eight circuits and may be set to maximum delays of 131.072 msec with 4 μ sec resolution. All the timers are usually referenced to the linac reset TZERO, but may also be triggered by each other.

The timer board in secondary #B (buncher) has a special function in that it produces the TDATA timing pulse used to trigger the A/D units in the linac. This is done so that the value of L:TDATA may be changed through the control system. The other TDATA pulse, produced by a pre-det in the pre-acc control room, is a backup in case the timer board in secondary #B fails. Switching between the two is a matter of changing a cable on a patch panel next to the Multibus crate for secondary #B. The signal is sent down the linac via a timing pulse repeater just as QUADS ON, RF ON, and TZERO are.

A watchdog timer with a delay of 150 msec is reset whenever a special register on the board is addressed. The normal operation of the CPU will cause this to occur every cycle (66 msec) as TZERO is received; if the register is not addressed and the watchdog times out, a system reset is generated, which should suffice to restore the secondary to periodic operation. If TZERO disappears, the secondary will reset itself every 150 msec until the reset returns.

The communication card serves two primary functions: it handles the serial I/O and video RAM for a local console (HDWREL 3), and interfaces the secondary to the local fiber optic link repeater chassis. The link repeater (fig. 67) is housed in a separate chassis with its own power supply; the secondary can be taken offline (bypassed) without affecting link data traffic. A switch on the front of the link repeater chassis will put it into bypass; it may also be put into bypass by flipping a dip switch on that secondary's annunciator board (the aluminum panel in the Multibus crate that shows the various interrupt levels), or if power to the secondary is lost, or if the cable between the link repeater and communication card is unplugged. The death of the link repeater has no effect on the secondary aside from isolating it from the rest of the control system.

The link repeater is also responsible for decoding the 1 MHz clock encoded on the link. This clock is not the same as the 1 MHz booster clock and is used for data transmission only. If the link should fail upstream of a given secondary, a clock built into that secondary's repeater will be enabled and used for transmitting data and a beacon message with the address of the secondary will be sent out on the link, for the benefit of the primary.

Each secondary operates synchronously with the 15 Hz repetition rate of the linac. Every 66 msec, the secondary reads all the analog and binary parameters in its area, updates lists of data for requestors that need data from its

area, scans the devices in its area to update the alarm conditions, adjusts parameters that are under software loop control, executes the application program (if any) selected by the operator of the local console, updates the local console display, makes any requested settings to local or remote parameters, and handles all communication with the SDLC link. (Whew!)

12.3 Primary Station

The primary hardware (fig. 12.5) is a bit different from the secondaries'. In addition to the standard core memory board, there is a 128K RAM memory board used to buffer data coming from and going to the hosts. There are two MC68000 CPU boards, one to drive the SDLC link and the other to drive the local console. Each CPU has a communication board; the board for the link driver CPU connects to the fiber optic link repeater chassis but not to the console, and the board for the console driver CPU connects to the console but not the link repeater. The CPU and communication boards are not interchangeable with their counterparts in the secondaries.




figure 12.5: primary Multibus crate

The FNAL Linac

The console CPU board also has an RS-232 port that sends alarm messages to a TEC terminal in the same rack. The video output of this terminal is buffered and drives the scrolling linac alarm screen in the MCR.

The primary has two sets of console hardware; one is by the primary station itself, and the other is in the MCR. Only one can be driven at a time, selected by switching BNC cables on a patch panel below the primary.

A single red dot binary I/O board in the primary is used to inhibit beam by pulling the hardware loop.

The *Ethernet controller board* handles communication on the link to the host computers. Ethernet is a half-duplex, 10 megabit/sec protocol, which is an international standard for intercomputer communications. The Ethernet system can in theory support up to 100 nodes, but in this application has only the primary and two hosts.

The fiber optic link repeater for the primary is also unique, because the primary communicates on the link in full duplex (simultaneous transmission and reception) rather than the loop mode used in the secondaries. This link repeater also provides the 1 MHz clock normally used for link communications.

The primary station initiates all traffic on the link--secondaries do not speak unless spoken to. At the direction of a host computer the primary will relay settings and requests for data to the secondaries, collect and organize the returned data, and return it to appropriate host. The primary polls each secondary every two seconds to see whether it is still online, and will create an alarm if any are not. The primary also looks for beacon messages from secondaries that would indicate a link failure. Note that an offline secondary will not generate a beacon.

12.4 Operating Modes

There are three modes of secondary operation. In the *local mode* the secondary controls the devices in its area without any help from any other computers. The local data base contains the names, titles, conversion factors, settings, nominals, tolerances, and control characteristics of all parameters in the local system. All the devices in that secondary's area may be controlled and monitored using the local console. The console also has a rudimentary plotting capability (sec. 12.6).

In the *global mode*, available when the SDLC link is up, any secondary can be used to monitor and control any device controlled by any other secondary, i.e., the entire linac (excluding some of the 200 MeV area). Aside from typing in different device mnemonics, the operation of the local console is the same as in the local mode. Communication through the SDLC link implies that the primary station and all link repeaters are functioning normally, otherwise link traffic stops. There is no direct communication between secondaries-- all requests and responses go through the primary.

The FNAL Linac

Communication between the secondaries and a host computer via the primary comprises the *gateway mode*, which is the most common mode of operation. One host is itself an MC68000-based station used for generating daily operational summaries used by the linac and controls support groups. It is located next to secondary 3. The other host is the linac *front end* PDP11/34 found in the cross gallery computer room. The front end interfaces the linac-dedicated control system to the Accelerator Control NETWORK (ACNET) which provides monitoring and control of any and all linac parameters from ACNET consoles in the MCR and (shudder) elsewhere. The accelerator control system (see Introduction to the Control System by Dave Johnson) also provides data logging and alarm monitoring services. The linac front end issues data requests and commands to the primary station in order to satisfy the needs of the control system.

All three secondary modes are supported simultaneously; local control does not preclude remote control and vice versa.

12.5 Troubleshooting

Every CPU board is connected to an annunciator panel that is studded with LEDs. Normally most of these flash at 15 Hz. If no activity is seen, it may be necessary to RESET the CPU with the button so marked. If more than two bus errors occur in a row, the CPU will halt and try to restart itself. If more than fifteen halts occur in an hour, the CPU will halt and require a manual RESET. If there is no obvious problem (like a power failure) and the CPU fails to show activity after a RESET, call the controls group for help.

Constant 15 Hz activity should also be observed on the link repeater chassis, the A/D chassis, and the local console.

If a secondary has gone offline the primary will show an alarm to that effect and the link repeater (fig. 12.6) for that secondary will show a red BYPASS LED lit, with a yellow LED indicating the reason. TX CLK ENABLE is an indication that the local link repeater clock is enabled because the link has failed upstream. A secondary in this condition will be sending a beacon message with its signature to the primary; the primary is the only station with TX CLK ENABLE normally lit.




figure 12.6: fiber-optic link repeater

The FNAL Linac

A small video screen above the primary will show the status of all the secondaries at a glance (fig. 12.7). Included is the number of data lists being requested by the hosts and other secondaries, the secondaries receiving data list requests, the secondaries transmitting data lists, the secondaries responding to the test poll, and the length of time required to collect all the data requested on that cycle.

If there is trouble with the primary or a secondary, the first thing to try is a RESET (located on the bottom right of the Multibus chassis). This starts the CPU from the first step of the program, as would cycling the power. Since analog and digital settings are stored in nonvolatile RAM, they are unaffected by crashes and power outages. The INT (interrupt) button merely causes the CPU to pause and then resume from where it left off, so it is of little value in a failure situation.

The secondaries in the high-voltage domes can be reset through the control system. To do this, type in channels G64 (H) or G65 (I) and interrupt under the RESET field. All other secondaries and the primary must be reset locally.




figure 12.7: primary link status display

If a secondary goes offline no data are available from it and local investigation is necessary. The offline condition can result from the secondary being powered down, halted, bypassed, reset, or caused by a faulty communication board or SDLC link problems.

The linac control system is nothing if not modular. Secondary components (except the link repeater) can be taken from secondary A (and secondary A subsequently bypassed). The binary I/O cards, CPU, and communication card should operate in any other secondary, as well as A/D and D/A chassis. The core memory board may be "known good" but it will have to be downloaded from the VAX when it is moved to

another station. Another source for secondary spares is from the test setup in the controls lab near RF station 2. A third source is a cabinet near the PA for RF station 2.

Primary spares are available only from the controls lab next to RF station 2. Remember the primary's link repeater is also unique. A good spare Ethernet board may be taken from the 68000 host near RF system 3.

IMPORTANT! Multibus uses printed circuit edge connectors, so the bin must be powered down to remove or insert cards. Also make a concerted effort to put the new board in the same slot as the old one (extra points for right-side-up).

External A/D and D/A chassis are easy to change, the only catch is that there is a chassis address switch on the back that must be set to match the address of the unit being replaced. Of course, the -10 to 10 volt input A/D's should not be mixed up with the -2.5 to 2.5 volt input A/D's.

All local consoles are identical and may be interchanged freely, but since all parameters are available from any console it is not necessary to change them just to establish control over a particular device.

A spare TEC terminal for the scrolling linac alarms monitor is located in the rack next to the active terminal.

12.6 Local Console Operation

The motion of the cursor on the screen is controlled by holding in either the up-down or left-right buttons and turning the knob (fig. 12.8). The same knob is used to change the value of a device when the cursor is in the number field. The INCREASE and DECREASE buttons may also be used, although they may be too fast for some applications, like running up RF gradients. A keyboard is mounted in a drawer beneath the screen, with two more interrupt buttons and a HOME key that sends the cursor to the upper-left hand corner.




figure 12.8: local console

In this context, to "enter" a character or string of characters means to type them in on the keyboard and hit the keyboard interrupt.

The FNAL Linac

Pages are selected as on an ACNET console: either move the cursor to the desired page title and interrupt, or enter the page number directly.

Up to fourteen channels of linac data can be displayed on the local parameter page, gathered from anywhere in the network. The mode switches select one particular aspect of the devices:

A/D	readbacks in thirteen-pulse averages
D/A	current settings
NOM	nominal value of channels in alarm scan
TOL	tolerance of channels in alarm scan
SET	D/A when line or page was initialized

In all cases the number field will be in reverse video if the channel is out of tolerance, blinking reverse video if it is inhibiting beam.

The unit switches select:

ENG	engineering units (amps, degrees, volts, percent)
VOLT	measured voltage from the A/D converter
HEX	raw data in hexadecimal (base 16)

Devices may be changed by knobbing or by entering a value in the number field, as with ACNET consoles. Pulses may be sent to stepping motors by holding down the "S" key.

The parameter pages are nominally set up with standard lists of local devices. A line may be blanked by interrupting in the first column in the line. Entering two digits of the desired channel number and interrupting will call up the desired channel from the local system. Three digits will retrieve a channel from anywhere in the system, a complete device mnemonic does the same.

To blank a page from any line to the bottom, blank the line, type in four spaces, and interrupt.

To get sequential channels from an area, enter the first channel on a line, blank the line below it, and interrupt on the blank line. Holding the interrupt button down will cause sequential channels to appear on the remaining blank lines.

To get the same channel number from sequential areas, enter the first channel on a line, blank the line below it, type the next desired system number in the first column in the blank line, and interrupt. Holding the interrupt button down will cause channels from sequential systems to appear on the remaining blank lines.

To start a time plot, enter a "T" on line 13 (hit HOME and four up arrows to get to this line). Enter the y-axis device on line 14, and the desired scales on line 13 as indicated. The plot may be restarted by interrupting under the prompt.

To do a parameter plot, enter a "P" on line 13. Enter the y-axis device on line 14, the x-axis device on line 15, and the scales on line 13 as indicated.

In both cases the plot may be terminated by interrupting on the first column in line 13.

The FNAL Linac

In addition to standard parameter pages, there are analog and binary descriptor pages, which show information about a channel not always displayed on a parameter page. The analog descriptor page shows the channel number and conversion factors for the channel in question, and also shows the setting of the alarm monitoring bits "A", "I", and "2".

If the "A" bit is set, the channel is included in the alarm scan and will cause an alarm if the readback goes out of tolerance. If the "I" bit is set the channel will inhibit beam if it goes out of tolerance and the "A" bit is set. If the "2" bit is set, two consecutive out-of-tolerance readbacks are required before an alarm or inhibit will occur. This is useful for filtering out nuisance alarms caused by a channel wandering out of tolerance for a single cycle.

The binary descriptor page shows the setting of the alarm monitoring bits "A", "I", "N", and "2". Setting the "N" bit toggles the nominal condition of the channel (0 or 1) no matter how the other bits are set.

Personnel hazards associated with the linac fall into two categories: electrical and radiation. Electrical hazards, in the form of RF systems and high-voltage power supplies, are not particular to the linac. Basic electrical safety is a topic too general to be discussed here. Radiation hazards posed by the linac are somewhat unique due to the construction of the linac enclosure and the nature of linac operation. A complete discussion of radiation hazards and appropriate safety procedures are found in the Accelerator Division Safety Manual, available in the MCR.

13.1 Radiation Hazards

Remanent radioactivity, caused by accelerated protons colliding with the linac structure, is negligible below energies of 10 MeV. There are thus no significant hazards from remanent radioactivity in the 750 keV line area (EXP 90), which may be accessed when the linac is in operation.

The most active area in terms of remanent radioactivity is the 200 MeV area. Highest rates occur near equipment that intersects the ion beam. The bending magnets and wire scanners are good examples of this. The NTF beamline is another "hot spot".

The walls of the linac enclosure are a concrete-soil-concrete sandwich ranging from 3' thick at the low-energy end to 12' thick in the 200 MeV area. In addition, soil is piled on top of the linac enclosure to form a shielding *berm* several feet thick. During normal operation, rates outside the linac enclosure are so small as to be unmeasurable.

An exception to the above rule is the area downstream of the NTF treatment area, in the lower linac gallery. Dose rates as high as 50-100 mR/hr have been measured there during NTF patient treatment (OPBULL 425). A "chipmunk" ionization chamber is permanently installed at this location to monitor neutrons and γ -rays, and warns personnel if rates exceed modest limits.

A second radiation hazard comes not from activation of materials by high-energy protons but from X-rays produced by the RF gradients in the linac cavities. *Bremsstrahlung* radiation is caused when high-energy electrons- like those formed by Haefely power supplies and RF cavities- are decelerated by the atomic nuclei in solid objects. RF cavities are particularly adept at producing this radiation, with dose rates as high as 2 R/hr measured at one foot from the cavity. Since most X-rays are formed at 70% of nominal gradient and above, rates in the vicinity of the cavities

are easily averted by turning the cavity gradients down to 50% of their nominal values and disabling the PA CROBAR bit on ACNET page L25 (sec. 14.1). This is standard procedure if any work is to be done near a linac cavity.

RF system modulators, too, are a source of X-rays. With proper shielding, normal rates around an operating modulator are less than 0.5 mR/hr. Since much personnel activity occurs in these areas, the modulator exteriors are surveyed every couple of months.

13.2 Linac Safety System

The status of the linac radiation safety system (*RSS*) is generated and displayed in the CARESS racks in the main control room. The ten controlled-access keys that open the gates to the linac enclosure are also kept there. Removing any one of these keys will inhibit the critical devices for the linac, preventing beam from being accelerated. A critical device failure for any other accelerator enclosure will likewise inhibit the critical devices for the linac (OPBULL 314).

The critical devices for the linac are the beam stop (primary) and pulse shifter (secondary) (OPBULL 694, 721). Their status is displayed on the linac RSS panel in the CARESS racks. There is no electrical safety system as such, and the RF systems are not interlocked. For this reason it is necessary to manually reduce the cavity gradients to 0.5 volts or less (OPBULL 404) when working in the vicinity of a linac cavity. Flashing beacons in the tunnel above each tank warn personnel in the tunnel when a controlled-access key is pulled and the gradient in the tanks is above 0.65 volts (OPBULL 282).

Lack of an electrical safety system also implies that no electrical devices (magnets, kickers, septa) are automatically turned off for an access (two exceptions- see below). It is the responsibility of personnel making the access not only to remain aware of the electrical hazards in the enclosure, but also to turn off and lock off any device upon which they plan to work. The 58° and 32° supplies for the NTF beamline will trip off when the safety system drops, but the above rule still applies.

The linac enclosure has four entrances: two at the upstream end of tank one, one between tanks 4 and 5, and one in the 200 MeV area. There are also 24 *penetrations* running from the lower linac gallery into the linac enclosure. The penetrations are 30" in diameter and carry the 9-3/16" transmission lines, power and signal cables, and water pipes. Some are shielded with lead (depending on the local neutron flux) and all are covered with metal plates and padlocked.

If it is necessary to access a penetration, a single "penetration key" is available in the main control room. This key is interlocked to the linac safety system in a unique way: if the key is removed, the linac safety system

The FNAL Linac

drops, requiring a search-and-secure (OPBULL 396), beginning with a special interlock in the lower linac gallery by system 9 (OPBULL 1048). This is a reminder to check all the penetration covers (also in the lower gallery) to make sure they are locked.

If the penetration key has not been removed, the "special" interlock will not drop, and the search-and-secure may begin in the usual manner at the box in the upper linac gallery by system 9.




figure 13.1: linac interlock boxes

The FNAL Linac

14 LINAC APPLICATIONS PROGRAMS

Most linac tuning is done through ACNET, using consoles in the MCR. An index page (fig. 14.1) lists all the applications programs available. Aside from parameter pages, two programs of the most interest to the operator are the RF system digital status page on L25, and the 200-MeV steering page on L36.

A display of the linac toroid outputs (showing the beam currents all along the linac) may be had by going to a parameter page, toggling the plot package options to INJ, and interrupting under *LIN_TOR (fig. 14.2) in the right-hand corner.

TUE 11-FEB-66 10:48

L	PARAMETERS	LINAC/PREACC INDEX PAGE	*PAGES *COPIES*
		RF STATIONS	I- SOURCE
2	RF SYSTEMS	25 LINAC RF STATUS	48 I- SOURCE PARAMS
3	RF SURVEY	26	49
4	PARAMETER	27	50
5	CTF PARAMETERS	28	51
6		29	52
7		30	53
8		31	54
9		32	55
	DIAGNOSTICS	33	H- SOURCE
11	LINAC ANALOG DATA	STEERING / QUADS	57 H- SOURCE PARAMS
12	LINAC BINARY DATA	35	58
13	LINAC MEMORY DATA	36 LINAC STEERING	59
14	LINAC I/O TEST	37	60
15	LINAC ANALOG PREV	38	61
16	LINAC CORE MEMORY	39	62
17	LINAC DESCRIPTORS	40	63
18	LINAC TANK QUADS	41	64
19		EMITTANCE MSMT	MISCELLANEOUS
20		42 200-MEV WIRE SCANS	66 CTF MISCELLANEOUS
21	PREACCELERATOR	44 WIRE CHOMP	67 E-Z WRITER
22	GENERAL STATUS	45 EMITPLOT	68
23		46 EMITTANCE (LOW EN)	69 TOROID PLOT

figure 14.1: linac index page

Besides the 750 keV and 200 MeV line toroids, there are toroids mounted at the downstream end of tanks 2-9. A toroid plot is useful in situations where the linac beam disappears in that it shows where beam was last seen. Analog readbacks are also available on L3, subpage TOROIDS.

Loss monitors are also available in tanks 4-9 and the 200 MeV area to show how much beam is being lost in their areas. The downstream end of tanks 4-9 each have a loss monitor near the center axis of the tank and one on top of the tank. Tank 9 loss monitors are useful when tuning the 750 keV lines, as they are more sensitive than toroids to variations in beam intensity (although their output may vary for reasons unrelated to linac transmission). 200 MeV loss monitors are useful when tuning the straight-ahead dump line

The FNAL Linac

for emittance measurements. The readouts for these monitors are available through hardwired spigots in a patch panel near ACNET console #7, and on page L3, subpage LOSSES.




figure 14.2: toroid plot

14.1 L25: Linac RF Status

Digital status bits from the linac RF systems are displayed here (fig. 14.3). "Good" bits are represented as green dots, while "bad" or "naughty" bits are red numbers, corresponding to the system number. Some bits are controllable through the computer system; these have yellow dashes after their names, and will be discussed here.

- Comp Enbl Computer enable. Setting this bit turns the modulator pulse off and inhibits the auto-reset function.
- Reset Rdy Reset ready. Interrupting on this bit resets the RF station, the same as the blue OVERLOAD button on the A5 console.
- System HV This turns the modulator high-voltage on, the same as the big red button on the A5 console.
- VVL Open Remote control of low- (VVL) and high- (VVH) energy end vacuum valves.
- VVL Close
- VVH Open
- VVH Close

The FNAL Linac

WED 15-JAN-86 12:51

```

L25 Linac RF Status  Data
Test Poll.123456789ABCDE.GHI.....V      List Err: 0000      3

Pulse Off 123456789.  Ipr An OV .....      Beam Inhib.....
Comp Enbl-.....      Dr Scr OV .....
Dr RF Lo .2.....      Dr An OV .....      Auto-Enbl:
Perm Inh .....      Dr Crowbar .....      Dr Reset -.....
PA Crowbar .....      Freq CH .....      Perm Inh -.....
Reset Rdy-.....      Freq CCU .....      QPS Reset-.....
PA Water .....      VVL Open -.345....      PA Crowbar-123.....
Mod Water .....      VVL Close-.5.....      Auto-Grad-123456789.
Dr Water .....      VVH Open -.23....9.      Auto-Phas-1.....
Dr Rdy .....      VVH Close-.2.....9.
PA Fil Lo .....      VV Local ..34....9.
Mod Rdy .....
HVPS Rdy .....
System HV.....      Rad System .....

PA Fil DC .....      Dt Water .....
PA DC OV .....      QPS Intlk-123.....
PA Cth OV .....      LEE IPPS ..3.....
PA HVXr .....      HEE IPPS ..3.....
    
```

figure 14.3: L25: linac RF status

- QPS Intlk Quadrupole power supply interlock. Interrupting under this bit sends a reset to all the quad power supplies in that system. This will not reset a blown breaker.

- Dr Reset Driver reset. Setting this bit inhibits the automatic reset of a driver overload.

- Perm Inh Permanent inhibit. The permanent inhibit function includes an automatic reduction of cavity gradient to half value if a reset is not performed within 30 seconds. This is done to keep the RF from being turned on too fast after the cavity has had a chance to cool down. Setting this bit inhibits the gradient reduction function.

- QPS Reset Quadrupole power supply reset. Setting this bit inhibits the automatic reset of quad power supplies in that system.

- PA Crowbar Setting this bit inhibits the automatic reset of a PA crowbar, which includes running the gradient back up to 0.8 volts after the system has reset. For this reason this bit should be set when accesses to the linac enclosure are made.

- Auto-Grad Automatic gradient regulation. Setting this bit inhibits computer regulation of gradient levels above 0.8 volts.

- Auto-Phas Automatic intertank phase control. Setting this bit inhibits computer control of RF system phase used to regulate intertank phases to minimum levels. Always disabled for system 1.

The FNAL Linac

An important note: the Auto-Grad and Auto-Phas controls should be inhibited whenever the timer L:TDATA is changed from its nominal value of 2000 (or so) μ sec. The software that regulates the gradients and phases samples these values at L:TDATA time, and chaos will result if L:TDATA is not occurring during the RF peak. NEVER change the nominal setting of L:TDATA.

14.2 L36: Linac Steering

This page (fig. 14.4) provides automatic measurement of beam positions at the beginning of the 200 MeV transport line, as well as calculating corrections to the steering trim magnets (sec. 8.3) that should center the beam.

Interrupting under *Find Beam Positions automatically runs in wires 200-1 and 200-2 (wire scanners 1 and 2 in fig. 8.1). The computer initiates 15 Hz beam to make the scan (via the Smeds module). At the end of the scan, interrupting under *CALC and *SEND will calculate and send the corrections to the trim dipoles. A second scan is then done by interrupting under *Find Beam Positions again.

```
WED 15-JAN-86 12:52
L36  LINAC STEERING  I- 200 MEV  *
Wire 200-1 Position          CM
Wire 200-2 Position          CM

*Find Beam Positions  HORZ  VERT
WIRE 200-1           -.003  -.012  CM
WIRE 200-2           -.012  -.003  CM

*CALC HORZ  *SEND  D/A  A/D
Horz Trim Tank 6-7  /           AMP
Horz Trim Tank 7-8           AMP

*CALC VERT  *SEND  D/A  A/D
Vert Trim Tank 6-7           AMP
Vert Trim Tank 7-8           AMP

*Display Previous Trim Values

Steering Last Checked 01/08/86 at 0539
```

figure 14.4: L36: linac steering program

It should be noted that the trim supplies are so small that D/A values of 10 or greater will probably not have the desired effect on the beam positions. If the program cannot center beam without turning a trim on that hard, the 750 keV line should probably be retuned.

Steering should be checked before and after any 750 keV line tuning, before any 200 MeV line or booster tuning, before any momentum scans are done, and once per shift to forestall the natural tendency for the positions to drift.

15 LINAC CHRONOLOGY

The sequence of events that occur in an HEP beam pulse are listed here. All times are referenced to a phase reversal on the booster phase reversal clock, which occurs about 2 msec before the guide field minimum in booster. All times are given in microseconds, and may represent approximate values.

- 0 TZERO system reference timer, timing board reference for all secondaries except G, H, I
- 10 spectrometer A/D hold trigger
- 18 QUADS ON quadrupole power supply trigger
- 50 pulse shifter input. Up to this time the pulse shifter has been deciding whether or not to shift. For a beam pulse, the TZERO-equivalent reset is sent to the G, H, I secondaries at this time.
- 750 source gas valve opens
- 860 source gas valve closes
- 1000 most linac scope triggers
- 1100 quadrupole power supplies start to ramp
- 1700 RF4 NTF enable detection trigger- if an NTF enable does not exist at this time the pulse for RF station 4 will not be allowed to shift
- 1755 RF ON- RF system reset
- 1795 (variable) modulator start pulse, driver start pulse, mixer gate opens to allow RF to amplifier chain
- 1800 referenced to above event- modulator ramp starts, crowbar compare level goes high
- 1910 source extractor on
- 1920 Haefely voltage sampled
- 1928 referenced to mod ramp start- full RF gradient reached

The FNAL Linac

1935 source arc on- beam in preaccelerator column

1983 750 keV chopper on tube fires- beam in linac

2000 TDATA hold trigger for all linac A/D units; A/Ds for sources and two A/Ds at secondary G track the shifted pulse- for a beam pulse they coincide with unshifted A/Ds.

2013 750 keV chopper off tube fires- linac beam ends

2020 source arc off- beam pulse ends

2025 source extractor off

2050 (variable) quadrupole power supplies commute

2078 (variable) RF gradients start to ramp down

2200 referenced to above event- modulator pulse ends, crowbar compare level goes low

2235 (variable) driver stop pulse, mixer gate closes

2295 referenced to mod start pulse- second driver stop pulse

3000 μ P reset for all secondaries and primary- primary counters start, secondaries collect data

3100 quadrupole currents reach zero

8120 index poll time- primary requests data indexes from secondaries

13240 answer poll time- primary requests data from secondaries

16000-21000 secondaries finish sending data to primary

33720 command poll time- primary requests commands from secondaries

38840 test poll time- primary polls secondaries to see if they answer

54200 alarms poll time- primary requests alarm status from secondaries

66000 new cycle begins

GLOSSARY

- 3383 sequencer** One of the modules in the buncher low-level RF system, which produces the times necessary for defining the pulse and running the system. Sends pulses to 3384 reference pulse generator and the bias pulser for the 7651.
- 3384 reference pulse generator** One of the modules in the buncher low-level RF system, which defines the RF pulse envelope according to times sent from the 3383 sequencer.
- 4-gap pulse stretcher** A module that stretches 4-gap pulses so that the prom module will recognize them and inhibit beam on 4-gaps.
- 4616** Linac RF driver PA tube, with an output of 200 kW. Cathode driver for the 7835 PA tube. Manufactured by RCA.
- 4800 output module** Last module in the buncher low-level RF system before the mixer and amplifier chain. Driven by 3384 reference pulse generator and feedback from cavity gradient.
- 7651** Linac IPA2 tube, with an output of 2 kW. Drives 4616 driver tube. Manufactured by RCA.
- 7835** Linac PA tube, with an output of 4 MW. Drives linac RF cavity via 9" coaxial transmission line. Manufactured by RCA.
- abort logic/pulse shifter interface** Produces status of main ring and tevatron abort loops. Input to keyswitch module.
- accelerating column** Set of seven titanium electrodes (eight gaps) arranged in Pierce geometry to accelerate ions to 750 keV. Situated between -750 kV dome and pit wall.
- achromatic** Quality of a transport line or optical system where particle momentum has no effect on its trajectory through the system.

The FNAL Linac

ACNET Acronym for Accelerator Control NETwork. A system of computers that monitors and controls the accelerator complex. Interfaced to users through consoles in the MCR and elsewhere.

air ionization chamber Devices used by NTF to monitor neutron flux during patient treatment.

amplitude control module Linac low-level RF system component that controls the amplitude of the RF gradient by varying the size of the modulator input pulse.

annunciator board Status panel in linac primary and secondary microprocessors showing status and interrupt levels.

auto-gradient Feature where the linac RF gradients are controlled by computer through the amplitude control module. Normally engaged.

beacon Message on serial data link flashed by a secondary microprocessor when the link repeater upstream of it fails.

beam inhibit light link module Module through which the H and I secondary microprocessors can pull the beam inhibit line.

beam loading Phenomenon whereby beam being accelerated by an RF cavity changes the gradient and phase of the RF in the cavity.

beam stop Linac primary critical device in the 750 keV line that blocks the beam path to prohibit beam in linac. Controlled by CARESS and the pulse shifter.

beam switch Toggle switch that will inhibit linac beam on HEP pulses. Eight of these are found in the MCR.

beam turns The width of the beam pulse going to booster divided by the revolution period of the booster at injection.

beamline microcomputer An NTF computer that monitors beamline devices, dose rates, and communicates with the NTF medical microcomputer as well as the local linac secondary.

belly pan A design feature of the linac tanks that allows leaks in the tank welds to be put under vacuum.

Benelex A substance similar to G-10 used to hold NTF neutron collimators in position.

The FNAL Linac

- berm** Earth shielding over the top of a radiation enclosure.
- bias pulser** A module that drives the grid of the 7651 tube in the buncher RF system.
- bleeder resistors** In the linac application, the resistors between the high-voltage dome and ground that bleed down the charge on the dome when the power supply is turned off. 4250 M Ω .
- blue-dot board** Binary interface board in a linac secondary microcomputer.
- bouncer circuit** Haefely-designed system used to boost dome voltage to compensate for charge leaving the dome during the beam pulse. Not installed at FNAL.
- bremstrahlung** X-ray radiation caused by deflection of high-energy electrons by atomic nuclei.
- bulk tuner** A long copper bar, D-shaped in cross section, that runs the length of a linac RF cavity and roughly determines the correct cavity volume.
- CARESS** Acronym for Central Accelerator Radiation and Electrical Safety System. Hardware dedicated to monitoring and controlling the accelerator safety systems and critical devices. Independent of any other control system. Located in and below the MCR.
- centroid** Technically, the center of mass. Used here to describe the center of a beam profile.
- cesium boiler** Electrically-heated crucible in the dome used to vaporize cesium used to coat the source cathode.
- chop time selector** Module in the preacc control room that selects times generated by predets and sends them on to the 750 keV choppers.
- chopper** Electrostatic device that selectively deflects a portion of the beam pulse to control intensity in the linac or beam turns in booster. Two are found in the 750 keV lines, and one in the 200 MeV area.
- chute** Passageway leading from the 200 MeV area to the booster radiation enclosure. The 200 MeV transport line to the booster travels through this.
- clock generator** Modules in the preacc control room (one for H-, one for I-) that take the 1 MHz clock and convert it to light pulses to send to the domes via fiber optic cables.

The FNAL Linac

- clock module** Module in the preacc control room that converts all the phase reversals on the booster clock to 1-gaps, as well as providing backup 15 Hz and 1MHz clocks for the linac in case the booster clock fails.
- Cockcroft-Walton accelerator** A high-voltage DC accelerator, especially for the acceleration of protons. The DC voltage is produced from a circuit of rectifiers and capacitors to which a low AC voltage is applied.
- computing counter** A device that measures pulse frequencies with great accuracy. Used for the display of the linac master oscillator frequency.
- corona rings** Curved metal surfaces shielding the outside of high-voltage devices to forestall the ionization of air by electric discharges. Used in the HV sources for the preaccelerators.
- coupling capacitor** A capacitor used to couple AC voltages to parts of an electrical circuit while blocking DC voltages.
- critical device failure** A safety system condition where the radiation safety system (RSS) for an area drops before the critical device for that area has been turned off to remove beam from the area. Results in a pulse shifter inhibit and must be reset at the CARESS racks.
- crowbar compare circuit** Modulator and PA protection circuit that shuts down the RF system if the current from the modulator to the PA anode exceeds certain limits.
- data selector** Part of the monitor/control module for the Haefely power supply which selects either the voltage readback or the command voltage for display on the front of the module.
- dispersion** Quality of a beam transport system at a given point that defines the variation of the transverse position of the beam with variations in beam momentum. Usually expressed in meters.
- divergence** The angle that the trajectory of each particle makes with the beam axis. Accelerator systems always try to reduce beam divergence.
- dogleg** A shifting of the beam axis by a pair of dipole magnets, such that the initial and final beam axes are parallel.
- dome** A large rounded box, maintained at a high potential, that houses the ion source and supporting equipment in a preaccelerator.

The FNAL Linac

- drift tube** A long copper torus, containing a quadrupole magnet, through which beam passes during acceleration in a linac RF cavity. The beam is shielded from the electric field in the cavity while in the drift tube, and is accelerated by the field while passing between drift tubes.
- drive loop** Termination of the coaxial transmission line leading from the PA to the RF cavity, which couples the energy in the transmission line to the magnetic fields in the cavity.
- driver** See 4616.
- duoplasmatron** A type of ion-producing source that develops protons by extracting positive ions from an arc struck in hydrogen gas. The I- source used to be a duoplasmatron; both sources are now magnetrons.
- electron-volt** Unit of energy equal to the energy gained by an electron as it falls through a potential of one volt.
- emittance probe** Device used to measure the size of the phase space ellipse occupied by the beam, done by measuring the divergence of the beam at a number of points across the beam axis.
- enable** A request for beam sent to the prom module in the preacc control room. Three types of enables exist: HEP, NTF, and P-bar.
- Ethernet controller board** Board found in the linac primary microcomputer that handles communication with the Ethernet link.
- Eurobus crate** One of two small crates in the preacc control room containing cards used to interface the local microcomputer to the Haefely power supply controls.
- extractor** An electrostatic device used to extract ions from the magnetron ion source, consisting of a pair of plates mounted beneath the source connected to a pulsed power supply.
- fiber-optic link repeater** Chassis associated with primary and all secondary microcomputers that maintain continuity of the serial data link and control the microcomputers' interfacing with it.
- filament regulator gating module** Module behind system 6 that tells the inductrols to regulate or not regulate, depending on timing pulses associated with the main ring ramp. Due for modifications (3-86).

The FNAL Linac

- field lamp** Light source between the NTF primary and main collimators designed to duplicate the profile of the neutron beam with any given collimator insert. Used in patient set-up.
- foreline** A vacuum line leading from a vacuum pump to the device or devices to be pumped.
- forward power** The power transmitted from one stage of an RF system to the next, which may not represent the full power of the first stage due to power reflected from the second (reverse power).
- frequency control and phase lock module** Module in the low-level RF system that provides RF input to the amplifier chain at the correct phase.
- frequency control module** Module in the low-level RF system that commands the tuning slug controller in response to error signals that represent the difference between the transmission line RF phase and the cavity RF phase.
- front end** A PDP11/34 computer that interfaces with a specific accelerator system, responsible for collecting data required by ACNET. The linac front end communicates with the primary microcomputer via an Ethernet link.
- G-10** A green glass-epoxy electrical insulator material used throughout the laboratory.
- Gannon box** Another name for the NTF interlock module, built by Jeff Gannon.
- gas barrier** A ceramic insulator in the 9-3/16" transmission line that isolates RF cavity vacuum from the nitrogen that pressurizes the transmission line.
- gate valve controller** A module in the preacc control room which shows the status of the 750 keV line vacuum valves and permits local control of those valves. Also the module at each cavity that monitors the ion pump power supply status and controls the vacuum valves at each end of the cavity.
- gateway mode** One of three modes of linac control system operation where data requests and commands from the front end are serviced by the secondary microcomputers.
- global mode** One of three modes of linac control system operation where data requests and commands from secondary microcomputers are serviced by other secondaries.

The FNAL Linac

- gradient** A measure of the electric field amplitude in an RF cavity. Usually expressed in percent of nominal.
- gradient regulator** Module in the modulator pulse-forming circuitry that acts to keep RF cavity gradient at the desired level.
- Haefely** Trade name for a Swiss company that manufactures high-voltage equipment. In this application, the high-voltage power supply that maintains the preaccelerator dome at -750 kV.
- host** One of two computers that communicate with the primary microcomputer on the linac control system Ethernet link. One host is the front end, the other is a 68000-based computer used for generating status lists for the linac support group.
- ignitron** Device found in modulators used to dump the capacitor bank voltage in the event of a PA crowbar. An ignitron passes electrical current to a pool of liquid mercury at ground potential.
- inductrol** Voltage regulator for the PA filament supply that compensates for variations in line voltage due to other large loads at the laboratory.
- intertank phase** A signal that represents the difference in phase of two adjacent RF cavities in the linac.
- ion gauge** A vacuum-measuring gauge that works by ionizing gas molecules with electrons and measuring the amount of ion current drawn to an anode.
- ion pump** A type of vacuum pump that works by ionizing gas molecules with electrons and catching the ions on a titanium anode. Used at pressures of 10^{-4} torr or less.
- ionization gauge controller** One of two ion gauges in the preacc control room that monitor column vacuum and control the vacuum valves leading to the columns, as well as being interlocked inputs to the 750 keV chopper supply controllers.
- IPA** Acronym for Intermediate Power Amplifier. An intermediate level of amplification in the high-level RF system. The 400 W solid-state amplifier is the first IPA, and the 7651 tube is the second IPA.
- isocenter** One of two points in space defined by alignment fixtures in the NTF set-up and treatment rooms. Patient set-up is usually done on the premise that the area to be treated will lie at the isocenter.

isoplane An imaginary plane perpendicular to the neutron beam axis and containing the isocenter. Used in NTF patient setup.

keyer A solid-state device providing the first stage of amplification in a linac modulator. The light-pipe input to the modulator is amplified to about 1 kV by the keyer.

keyswitch module Also known as the pulse shifter controller module. This module produces an HEP enable to be sent to the prom module. Keyswitches on the front of the module selectively bypass beam switches in the MCR, as well as other inputs to the module.

klixon A temperature-sensitive electrical switch used in interlock circuits of power supplies and magnets that opens when a certain temperature is exceeded.

ladder control Obsolete method of controlling quadrupole power supplies whereby the settings of a few supplies determined the settings for a number of others. No longer used.

LCW Acronym for Low Conductivity Water. Water that has had the free ions removed in order to increase the resistivity of the water to 9 M Ω /cm or greater. Found in the 55° and 95° systems used to cool linac components.

linac 68k μ P beam inhibit module Module in the preacc control room that sends the beam/no beam status to the linac secondaries, and causes a beam inhibit if one of the secondaries or the primary pulls the beam inhibit line.

linac steering Procedure whereby beam is aligned at the head of the 200 MeV transport line by steering dipoles between tanks 6-7 and 7-8. Can be performed automatically by ACNET page L36.

linlock line Vacuum valve status input to the prom module and pulse shifter, powered by a small 5 volt supply at RF station 9. Vacuum valve controllers at each station short out the line whenever a valve closes, inhibiting beam.

LLPA Acronym for Low Level Power Amplifier. 5-watt solid-state amplifier in the rear of the A5 racks that is the first step in the high-level RF system.

local mode One of three secondary modes where the secondary microcomputer satisfies data requests and commands from the local console.

The FNAL Linac

- longitudinal transit time factor** Ratio of the energy gain of particle traversing a cell with a constant electric field to the energy gain of a particle traversing a cell with a sinusoidally-varying electric field. Simple, no?
- magnetron** Ion source that produces negative ions by extracting them from a plasma formed by an electric arc formed in hydrogen gas. Used in H- and I- sources at FNAL.
- main collimator** NTF collimator assembly following the primary collimator and containing inserts to shape the neutron beam to the desired size.
- Manchester (encoding)** System of data encoding on a bipolar clock used in serial data link communications in the linac control system.
- master oscillator** Oscillator that provides the RF signal for all linac RF systems. Typically runs at 201.24 MHz.
- mechanical phase shifter** Device strung between RF cavities with inputs from each cavity, which produces the intertank phase signal used for phase regulation.
- medical microcomputer** NTF computer that drives the NTF console and interfaces with the beamline microcomputer.
- metering resistor** Precision resistor-capacitor network used for measuring the voltage on the preaccelerator dome. 2125 M Ω .
- MIL-STD-1553B** Digital multiplex data bus standard used for interface between secondary #G and Eurobus crate cards that control the Haefely power supplies. Originally developed by the U.S. military.
- mixer** Final component in low-level RF system that passes the RF signal to the amplifier chain when gate signal arrives from waveform generator/sequencer.
- mod block** Premature termination of a modulator pulse triggered by a modulator current of greater than 400 amps during the pulse.
- modulator** High-voltage device connected to the anode of the PA that controls the RF output to the cavity. The modulator consists of three nested amplifiers with a final output of 30 kV at 350 A.
- modulator regulator** Module that regulates modulator current to maintain constant RF cavity gradient.

The FNAL Linac

- monitor/control module** Module in the preacc control room that monitors voltage on the preaccelerator dome, digitizes it, and sends it along with the command voltage to the regulator/oscillator module to complete the voltage regulation feedback loop.
- monitor/inhibitor module** Module that oversees a number of RF system interlocks and inhibits the modulator pulse if conditions warrant.
- monitor unit** Artificially-derived unit used by NTF to define neutron dose rates.
- Multibus** IEEE standard format for microcomputer systems hardware. The crates that house the primary and secondary microcomputers conform to this standard.
- NTF** Acronym for Neutron Therapy Facility, a medical facility investigating the treatment of malignant tumors with neutrons. Neutrons for the facility are generated by steering 66 MeV ions from the linac into a beryllium target.
- NTF beam control module** The module in the NTF control room that actually initiates beam when two switches on the front panel are pressed simultaneously.
- NTF interface module** Modules located at RF systems 3 and 4 that sense the time of the RFON pulse. The module at system 4 will send an RF4 INHIBIT to the NTF interlock module if the RFON pulse at system 4 does not shift at the proper time.
- NTF interlock module** Also known as the Gannon box. This is the locus of all the NTF interlocks. The logic for producing an NTF enable is contained in this unit, which then sends the enable to the prom module. The interlock module also drives the status modules in the NTF control room and the MCR.
- NTF status module** Module that shows the status of the NTF interlocks, generated by the NTF interlock module. The module in the MCR also has keyswitches and toggle switches that are themselves part of the interlock chain. The module in the NTF control room has status only.
- object** Optical term referring to one focus of a dipole magnet. Parallel beams entering a dipole will cross at the object.
- off tube** One of two thyratrons that control the operation of a 750 keV chopper. After the on tube fires, the off tube fires to ground one side of a series capacitor and draw one of the chopper plates from ground to a negative potential.

The FNAL Linac

on tube One of two thyratrons that control the operation of a 750 keV chopper. The on tube fires to bring one of the chopper plates from a high potential to ground. The off tube fires at a later time.

optical isolator An element in an electrical circuit that converts a signal to a light pulse at one point in the circuit and then back again in another in order to isolate the grounds of each part of the circuit from each other.

OPT022 Device which converts a binary signal to a light pulse. Found in the red dot binary I/O boards in the linac secondaries.

PA crowbar RF system trip that occurs when the modulator current is greater than 600 amps during the RF pulse or 125 amps between pulses. A PA crowbar fires the permanent inhibit, shorts the capacitor bank to ground through the ignitron, and resets automatically after 30 seconds.

penetration One of 24 short tunnels leading from the lower linac gallery into the linac enclosure. The transmission lines, water lines, electrical and signal cables pass through the penetrations.

permanent inhibit RF system trip caused by four successive mod blocks in a row. A permanent inhibit shuts off the modulator pulse and shuts down the high-voltage power supply for the capacitor bank.

phase comparator mixer An element of the low-level RF systems that looks at the relative phases of two RF signals and produces a signal proportional to the difference in phase.

phase lock input Output of a phase comparator mixer that looks at the desired RF phase and the cavity RF phase. This signal is raised by 5 volts to produce phase lock output.

phase lock output Signal proportional to the difference between the desired RF phase and the cavity RF phase, used to drive a phase shifter that compensates for shifts in cavity RF phase due to beam loading.

phase reversal A method of encoding data on a clock where the shifting from one level to the other is shifted in phase for a number of cycles. The number of cycles for which this reversed phase exists constitutes the data being transmitted.

phase space A six-dimensional space consisting of position and momentum in each of three orthogonal directions.

The FNAL Linac

- pickup loop** A small (1/8" square) loop at the end of a coaxial cable that is driven by the RF cavity magnetic field to produce a signal proportional to the electric field in the cavity.
- piezoelectric crystal** A symmetric crystal that bends when an electric potential is applied to it (or vice versa). Used in the ion sources to regulate the gas flow into the source. A piezoelectric crystal is also the heart of a crystal oscillator, such as the master oscillator.
- pig gauge** Penning Ion Gauge. A type of ion gauge used to measure vacuum and control valves in the 200 MeV area.
- pit** A large temperature- and humidity-controlled room containing the Haefely high-voltage transformers, voltage multiplier, preaccelerator dome, column, and motor-generator for a preaccelerator.
- plasma** A high-temperature mixture of ionized atoms and electrons.
- post coupler** A copper stem with a tab on the end that is used to control the relative field levels in two adjacent cells in a linac RF cavity.
- preaccelerator** Generally refers to all the hardware associated with the production of 750 keV ions, including the source itself and the Haefely power supply. The general layout is of the Cockcroft-Walton type.
- predet** A standard module that produces timing pulses referenced to specific phase reversals on an incoming clock. The delays of the output pulses are set by thumbwheels on the front of the module and may be as short as 1 μ sec.
- prepulse** A 4-gap phase reversal on the booster clock that commands the 200 MeV, booster, and 8 GeV pulsed devices to fire in preparation for a beam pulse. Linac beam is not allowed on prepulses.
- primary collimator** The first collimator after the beryllium target used to produce neutrons for NTF. The primary collimator is made of steel and is of fixed geometry.
- primary critical device** The principle device that determines whether or not beam will enter a certain enclosure. Critical devices are controlled by the safety system, as well as through the control system. The primary critical device for the linac is the beam stop in the 750 keV line. The secondary critical device for the linac is the pulse shifter.

The FNAL Linac

primary microprocessor Choke point of the linac control system that controls messages on the serial data link, and interfaces with the host computers via the Ethernet link.

prom module Logic module in the preacc control room that looks at the various beam enable inputs and decides which, if any, will produce a beam pulse. The prom module also selects the chop width to be sent to the choppers.

pulse interlock module Component in the modulator pulse-forming circuitry that will inhibit the modulator pulse if conditions warrant.

pulse repeater Module immediately after the pulse shifter in the preacc control room that sends the pulse to the chopper predets, the H- and I- clock generators, two of the A/D units for secondary #G, and secondary #G itself.

pulse shifter Module in the preacc control room that will delay the reset pulse to the sources by 1 msec if no beam is desired in the linac. Controlled by the prom module, safety system status, and vacuum valve status.

pulse shifter status module Module in the control room that shows the pulse shifter status and can reset a pulse overcount if one occurs.

pulse skip detector Module in the preacc control room that senses when one or more of the 1-gap phase reversals from the clock module is missing (should occur at 15 Hz) and sends a message to secondary #G so that it can post an alarm.

quadrupole triplet A series of three quadrupoles of alternating polarity that produce a net focusing of the beam in both planes. Typically the two outer elements of a triplet are wired in series.

ramp enable A permit sent to the NTF 58° magnet power supply telling it that it may ramp. The ramp enable comes from the NTF interlock module.

red-dot board One type of binary I/O board found in a linac secondary connected to OPT022 devices to provide optical input and output to the secondary.

reference source Part of the monitor/control module in the Haefely control system that produces a command voltage for the power supply in response to computer commands or local input.

The FNAL Linac

regulator/oscillator module Module in the Haefely control system that receives the digitized command and monitor voltages from the monitor/control module and uses them to generate a drive signal for the power amplifier. The regulator/oscillator also receives an input that represents the high-voltage transformer output and uses it for regulation.

reverse power The RF power from an amplifier that is reflected back by the load. The power not reflected back is the forward power. The sum of the two represent the total power of the amplifier. In RF station tuning, the reverse power between stages is always tuned for a minimum.

RF bucket That area in RF phase where particles oscillate about the synchronous phase angle. Particles in an RF bucket will normally remain in the bucket. Particles outside the bucket will not be accelerated.

RF cavity An electrically-resonant standing-wave cavity designed to impart energy to particles as they pass through a gap or number of gaps in the cavity by virtue of the electric field gradient across the gap(s).

RF defocusing Phenomenon caused by the curved fields in the gap between drift tubes and the changing electric field strength that results in radial defocusing of the beam as it is accelerated between drift tubes.

RF phase adjust module Module in the low-level RF system that adjusts the phase of the input RF from the master oscillator under the direction of the local secondary microcomputer. This adjustment controls the intertank phase.

roughing pump A rotating-vane pump used in the initial stage of vacuum pumpdown, good down to about 10^{-3} torr.

RSS Acronym for Radiation Safety System. Logic system for a particular enclosure that controls the primary and secondary critical devices for that area. Inputs to the RSS include the status of the gates leading to the enclosure, the emergency switches in the enclosure, and whether or not the enclosure has been secured.

scaler Hardwired device in the NTF control room that acts to terminate patient exposure in case the beamline μP fails to do so at the proper point. The scaler contains counters and timers not associated with the control system and inputs directly to the NTF interlock module.

SCR Acronym for Silicon Controlled Rectifier. Also called a thyristor. A semiconductor rectifier whose forward anode-cathode current is controlled by a signal applied to a third electrode. Widely used in power supplies throughout the laboratory.

SDLC Acronym for Synchronous Data Link Control. A data link communication protocol developed by IBM and used to control the serial data link (or SDLC link) that connects the primary and secondary microcomputers in the linac control system.

secondary microprocessor Also secondary microcomputer. One of seventeen computers distributed throughout the linac that control and monitor the functions of the linac. The secondaries converse on a serial data link under the control of the primary microcomputer.

septum A magnetic or electrostatic device used to deflect charged particles along one of two paths. Typically, a solid metal sheet or plane of wires separates a region with and electric or magnetic field from a region of no field. Beam entering the first region is deflected while beam entering the second region is not.

serial data link Also SDLC link. A fiber-optic data link that connects the primary and secondary microcomputers in the linac control system. Data is transmitted on a 1 MHz Manchester-encoded clock under SDLC protocol.

shunt impedance Technically, an impedance in parallel with an electrical circuit. In this application, the ratio of the square of the electric field strength on the axis of an RF cavity to the power dissipated per meter of length.

Smeds module Module in the MCR built by Jim Smedinghoff (Smeds) that allows computer enabling of 15 Hz beam for use by the linac steering program (L36).

space charge forces Divergent forces on a charged particle beam caused by nonzero net charge density, i.e.: mutual repulsion.

spark gap Electrodes found on the Haefely voltage multiplier and the accelerating column that will arc in the event of a voltage imbalance and thus prevent an arc from occurring in expensive/inaccessible components.

spectrometer A device for producing a spectrum. The 40° bending magnet in the momentum analysis line comprises a magnetic spectrometer. Particles with different momenta will follow different paths through the magnet and will appear in different positions at the scanning wire.

The FNAL Linac

- stem-box cover** A cover over the tops of the drift-tube stems that protrude from the top of a linac RF cavity, which may be evacuated to prevent vacuum leaks around the stems.
- synchronous particle** A particle that is at the synchronous phase angle with respect to the accelerating RF.
- synchronous phase angle** That RF phase angle at which a particle will receive just enough energy in each cell to reach the next cell just as the RF goes through 360° .
- system ready** Sum of the second-level interlocks in the NTF interlock module. Two such sums ("A" and "B") are required to be present in order to make up the next level of interlocks.
- tetrode** A thermionic emission tube with four electrodes: cathode, anode, grid, and screen.
- thyatron** A gas-filled triode in which the voltage on the grid can trigger ionization of the gas in the tube. Once the gas is ionized, current flows from cathode to anode until the potential across the two falls below a certain level. Thyatrons are used as high-voltage relays in the chopper power supplies and in the RF modulators to trigger the ignitrons.
- transmission line** An electric line uniform in series resistance, series inductance, shunt inductance, and shunt capacitance. The transmission line from a PA to an RF cavity is coaxial in construction.
- transport line** A system of bending and focusing magnets used to transport beam from one area to another.
- triode** A thermionic emission tube with three electrodes: cathode, anode, and grid.
- trombone** A U-shaped device for adjusting the length of coaxial cables or transmission lines.
- tuning slug** Copper cylinders mounted in the wall of the RF cavities that may be moved in and out to adjust cavity tune. Specifically, one motorized slug in each tank under the control of the low-level RF system.
- turbo cart** Vacuum pumping stations found in the 200 MeV area, consisting of a combination of a roughing pump and a turbo (turbomolecular) pump. These stations are connected to the 200 MeV line vacuum through manually-operated valves.

The FNAL Linac

turbo pump Turbomolecular vacuum pump. A pump containing spinning turbine disks used in the second stage of vacuum pumpdown, good down to about 10^{-9} torr. Ion pumps are used at operational pressures for the sake of economy and practicality.

water resistor Plastic tubes running from the preaccelerator pit wall along the column and to the dome, used to control the potential drop among the column electrodes. The resistance of the tubes are controlled by LCW flowing through them. This resistance also has an effect on the current drawn by the Haefely power supply.

waveform generator/sequencer Module in the modulator pulse-forming circuitry that generates RF system times in response to the RFON timing pulse. This module also forms the basic modulator waveform which is then passed to the amplitude control module.

work function The difference in energy of the Fermi level of a solid and the energy of free space outside the solid. In this application, the amount of energy required to liberate an electron from a metal surface.

REFERENCES

The list that follows by no means covers all the material written about the FNAL linac; much more can be discovered in conference papers, IEEE reports, journals, and books. Some references are not entirely up-to-date, as improvements to the linac are constantly being made (in particular the control system was changed in 1983). It takes a seasoned eye to separate the wheat from the chaff in these cases. Furthermore, some documents (noted at the end of this section) are totally obsolete and are best avoided entirely.

GENERAL INFORMATION

OPBULL 1046 L. Allen, Your Basic Injector Gotchas
Compilation of many little-known facts about
the linac, oriented toward operations.

PREACCELERATOR

OPBULL 486 D. Kindelberger, Entering the Preaccelerator
Dome Ignore steps E and F.

OPBULL 600 J. Crisp, (no title) FNAL magnetron ion
source. Parameters have changed with the
introduction of the grooved cathode.

OPBULL 671 C. Schmidt, H- Source Pressure Servo System
Control system channel numbers obsolete.

OPBULL 717 R. Hren, Addition to Preaccelerator High
Voltage Interlock System

OPBULL 917 C. Schmidt, Correcting for Source Pressure
Bursts Control system channel numbers obs.

OPBULL 953 C. Schmidt, The H- Preaccelerator and Source
Changes Introduction of grooved cathode.

C. Curtis Ion Source and Injector Development
1976 Proton Linear Accelerator Conference.
Comparison of several types of ion sources.

The FNAL Linac

- C. Schmidt A 50-mA Negative Hydrogen Ion Source, 1979
Particle Accelerator Conference, San
Francisco. FNAL magnetron development.
- C. Curtis, et al. Linac H- Beam Operation and Uses at
Fermilab IEEE Transactions on Nuclear
Science, 1979.
- D. Howard Description of Solid State Monitor/Regulator
for Haefely 750 keV P.S., 1980.

750 keV TRANSPORT LINE

- OPBULL 210 C. Curtis, 750 keV Chopper Operation
- OPBULL 577 F. Cilyo, Preacc Sorenson Power Supplies
Troubleshooting Flow charts and everything.
- OPBULL 628 L. Allen, Tuning 750 the keV Transport Line
Old control system nomenclature. Try to get
one with the chart on the back.
- OPBULL 720 J. Wendt, Preaccelerator Vacuum Interlock
- FN 201 C. Curtis, et al. The Operation of the
First Section of the NAL Linear Accelerator
Old control system nomenclature, but good
reading.
- EXP 111 S. Ohnuma, The Beam Emittance... Discussion
of beam emittance throughout the
accelerator.

TIMING AND PULSE SHIFTER

- OPBULL 561 L. Wahl, Pulse Shifter General Information
Some good information but largely obsolete.
- OPBULL 663 L. Allen, New H- Chopper Controls
Prom module logic; some new inputs.
- OPBULL 915 M. Utes, Abort Logic/Pulse Shifter Interface
This system is rapidly changing. Beware.
- OPBULL 1058 L. Allen, Linac 15 Hz Running and P-bar Beam
Switches, Generation of P-bar enables. This
is also changing.
- OPBULL 1059 L.Allen, NTF Switch vs TLG, Description of
NTFP bit on ACNET page D69.

THE ALVAREZ STANDING-WAVE PROTON LINAC

- TM 279 S. Snowdon, Modified Phase Acceptance in Linacs Calculation of stable phase region.
- TM 611 S. Ohnuma, A Simple Program to Calculate the Alvarez-Type Linac Cavity
Simple for *him*, maybe.
- L. Smith Linear Accelerators, Instrumentelle Hilfsmittel der Kernphysik I, 1959.
Don't worry, it's in English. (Well, most of it anyway....)
- Livingston & Blewett Particle Accelerators, chapter 10
(John Crawford has a copy.)
- G. Wheeler RF Structures for Linear Accelerators, IEEE Transactions on Nuclear Science, 1965.
Discussion of drift-tube accelerators.
- G. Wheeler Progress in Proton Linear Accelerators, 1969
Comparison of BNL, NAL, LASL linacs.
- P. Grand The Physical Design of a 200-MeV Linac Facility, IEEE Transactions on Nuclear Science, 1967. Brookhaven linac.
- The following six papers are from the 1970 Proton Linear Accelerator Conference at FNAL:
- M. Palmer The Alignment of Tanks and Drift Tubes for the NAL 200-MeV Linac Using a Laser Light Beam
- D. Young, et al. Construction and Initial Performance of the NAL 200-MeV Linear Accelerator
Written before 200 MeV beam was produced.
- C. Owen, et al. Experience with Post Coupler Stabilized Structures in the NAL Linac
Good discussion of post couplers.
- M. Trump Bead Perturbation Measurement
Method by which cavity fields are measured.
- J. Aggus Linac Tank Installation and Tune-Up for 200-MeV Linac Description of fabrication of tanks for the Brookhaven linac.
- K. Halbach Effect of Drift Tube Tolerances on the Electric Field Distribution Along the Length of an Alvarez Cavity Masochists only.

The FNAL Linac

LINAC RF SYSTEMS

- OPBULL 74 D. Mendenhall, Linac RF Systems The virtues of turning on RF slowly.
- OPBULL 75 C. Owen, Linac Driver Minimum Forward Power Interlock
- OPBULL 121 C. Owen, (no title) PA filament regulator system. Old control system nomenclature.
- OPBULL 288 D. Mendenhall, 400 Watt Solid-State Power Amps First IPA.
- OPBULL 596 J. Lackey, Permanent Inhibits and Long Linac Downtimes, or How to Save a Gas Barrier More virtues of turning on RF slowly.
- OPBULL 692 L. Allen, PA Crowbar vs Permanent Inhibit
- OPBULL 797 L. Allen, Remote Control of Linac Gradients Ignore the second paragraph; the modules are different now.
- OPBULL 1049 R. Gerig, Linac Turn On: Turn Off Procedure Includes description of status lights.
- OPBULL 1050 J. Bauman, Tuning a Linac RF Station An ageless classic.
- OPBULL 1054 L. Allen, Linac Low-Level RF
- OPBULL 1056 L. Allen, Linac Modulators Y631s replaced by keyer; three toroids instead of two.
- TM 94 R. Featherstone, Coaxial Transmission Line for the NAL 200 MeV Linac Tests of several types of transmission lines.
- L. Wahl Waveform Generator/Sequencer Module, Amplitude Control Module, and Gradient Regulator, none yet published.
- A. Donaldson Modulator Reliability and Replacing Tetrodes With MOSFETS 1981 Proton Linear Accelerator Conference. Modulator keyer. (Haven't read this one, myself.)
- C. Curtis, et al. Operating Experience with the NAL 200-MeV Linac 1972 Proton Linear Accelerator Conference. Emittance measurements and RF system reliability.
- J. Leiss Beam Loading in Linear Accelerators, IEEE Transactions on Nuclear Science, 1965.

The FNAL Linac

- M. Hoover Grid-Cotrolled Power Tubes in Particle-Accelerator Applications, IEEE Transactions on Nuclear Science, 1965.
- D. Mendenhall The National Accelerator Laboratory Linear Accelerator Radio-Frequency System, 1970 Proton Linear Accelerator Conference.

BUNCHER

- OPBULL 1055 L. Allen, Linac Pulse Forming Network Old low-level RF system, still partially applicable to the buncher.
- OPBULL 1057 L. Allen, The Buncher and Beam Proper buncher tuning methods.
- FN 277 L. Teng, High-Efficiency Buncher System for Linac Actually a scheme for two parallel bunchers.

QUADRUPOLES

- OPBULL 212 C. Owen, Linac Operational Procedures Proper monitoring of linac quadrupoles.
- OPBULL 779 R. Benson, The Linac Quadrupole Power Supply System
- OPBULL 780 R. Benson, Timing In a Linac Quadrupole
- R. Goodwin, et al. Pulsed Power Supplies for the NAL Linear Accelerator Quadrupole Magnets, 1970 Proton Linear Accelerator Conference

200 MeV AREA

- OPBULL 397 E. Gray, 200-MeV Transport Line Setup No more horizontal stacking at booster injection since negative-ion sources used.
- OPBULL 718 R. Webber, Checking the Linac Beam Momentum Old control system nomenclature.
- FN 227 S. Ohnuma, Momentum Spread of the 200 MeV Linac Beam at NAL Changing the momentum spread of the linac by varying the RF phase between the last two cavities.
- TM 443 C. Curtis, Variability of NAL Linac Energy Changing output momentum of the linac by varying RF phases.

The FNAL Linac

- E. Hubbard, et al. System for Transport and Analysis of 200 MeV Linac Beam, 1970 Proton Linear Accelerator Conference. Good description of 200 MeV line setup and diagnostics, but does not reflect experience with beam.
- R. Goodwin, et al. Beam Diagnostics for the 200 MeV NAL Linac, 1970 Proton Linear Accelerator Conf. Descriptions of emittance probes and wire scanners.

VACUUM SYSTEMS

- OPBULL 29 G. Ortlieb, Vacuum Interlock Between Linac and Booster Tank 9 vacuum valve interlock to 200 MeV area Penning gauge.
- OPBULL 267 T. Larson, Power Failure in 200 MeV Vacuum System, Turbo Pumps Starting turbo carts.
- OPBULL 416 J. Erickson, (no title) Vacuum valve controller interlocks.
- OPBULL 513 W. Andrews, 200 MeV Vacuum System

WATER SYSTEMS

- W. Andrews Linac Water System General operation.
- F. Kzirch The Cooling System for the NAL 200-MeV Linac 1970 Proton Linear Accelerator Conference.

NEUTRON THERAPY FACILITY

- OPBULL 400 D. Rohde, Cancer Therapy Supply C58 58° magnet vernier. Old control system nomenclature.
- OPBULL 923 L. Allen, Suspending Linac Parameters Suspending inhibits on systems 4-9 without burning up the linac.
- OPBULL 1047 L. Allen, Linac Monitoring The same as OPBULL 923.
- TM 556 D. Johnson, Fermilab Medical Facility: Beam Transport Line NTF beam transport line.
- TM 613 J. Gannon, CTF Beam Interlocks Wildly obsolete for the most part, some good info.

The FNAL Linac

- TM 834 Awschalom & Rosenberg, Neutron Beam Calibration and Treatment Planning
Long but very informative description of NTF facility; patient setup procedure.
- TM 928 Awschalom & Rosenberg, Conceptual Design of Beryllium Targets for the Generation of Neutron Beams for Radiation Therapy by the (p,n) Reaction Discussion of target design. Also one of the longest titles I've come across.
- TM 993 Awschalom, et al., The Fermilab Neutron Therapy Facility: Treatment Planning for Neutron and Mixed Beams Patient setup.
- TM 1018 Awschalom, et al., High Precision in Dose Delivery: Routine Use of a Microcomputer Calibration procedures.
- FN 314 L. Cohen, et al., Second Annual Progress Report: Fermilab Cancer Therapy Facility, Batavia, Illinois Some medical results of neutron therapy; not written for the layman.
- FN 342 L. Cohen, et al., Fast Neutron Radiation Therapy Mechanisms of tissue response to neutron beams.
- M. Awschalom, et al., The Fermilab Neutron Radiotherapy Facility, IEEE Transactions on Nuclear Science, 1977.

CONTROL SYSTEM

- TM 1238 M. Shea, The Fermilab 200 MeV Linac Control System Hardware Long and thorough discussion of the new (1983) linac control system. Recommended reading.
- HDWREL 3 R. Florian, et al., A Serial Console for Linac Secondaries Description of local console hardware. Good if you want to build your own.
- SFTREL 105 R. Goodwin, Linac Analog Data Page Description of ACNET page L11, with that plot package no one can stand. Good insights into control system capabilities, however.

The FNAL Linac

Goodwin & Shea The Distributed Control System for the Fermilab 200 Mev Linac, 1984? Overview.

LINAC SAFETY

- OPBULL 282 J. Erickson, (no title) Radiation warning lamps in linac.
- OPBULL 314 R. Parry, Critical Device Failure Mode Module Control of beam stop and pulse shifter.
- OPBULL 396 R. Parry, Linac Radiation Safety System #3 no longer applicable. Search and secure procedure. Only OPBULL containing the word "ameliorated".
- OPBULL 399 R. Parry, CTF Radiation Safety System Not operations group responsibility.
- OPBULL 404 G. Urban, Linac Accesses Turning down RF gradients for accesses to linac.
- OPBULL 425 S. Baker, Cancer Therapy Facility Radiation Safety Radiation levels in lower linac gallery during patient treatment.
- OPBULL 694 R. Parry, Everything You Wanted to Know, and Need to Know, About the Linac 750 keV Beam Stop and Gate Valve Beam stop interlocks. Pressure and vacuum valve interlocks are no more.
- OPBULL 721 R. Parry, Changes and/or Additions to the Safety System New beam stop interlocks.
- OPBULL 1048 R. Parry, Linac Penetration Interlock
- EXP 90 C. Curtis, Accelerator Experiment: Preaccelerator Radiation Survey Radiation levels outside of accelerating column.
- M. Awschalom Radiation Shielding and Safety Considerations for the NAL 200 MeV Linac 1970 Proton Linear Accelerator Conference. Calculations of radiation levels in and about the linac, done before the facility was finished. Walls not solid concrete.

The FNAL Linac

The following documents are obsolete and are not suitable for informative purposes:

OPBULL 276
OPBULL 324
OPBULL 525
OPBULL 527
OPBULL 643
OPBULL 646
OPBULL 688
OPBULL 782
OPBULL 805
TM 309
SFTREL 65
SFTREL 67
SFTREL 99

The FNAL Linac

The FNAL Linac

LIST OF SCHEMATICS

All numbers are FNAL drawing numbers unless otherwise noted. Copies of all microfilmed FNAL drawings are available from the duplicating department in the basement of the central laboratory.

PREACCELERATOR

0211-ME-3288	Haefely power supply and preaccelerator dome
0211-ME-60008	Haefely power supply and preaccelerator dome
0212-EB-60602	Haefely power amplifier
0212-EB-60603	voltage divider (Haefely monitor circuit)
0212-ED-60599	monitor/control module
0212-ED-60600	regulator/oscillator module
0211-ME-60469	source assembly
0211.MD-60266	accelerating column electrode

Haefely 0-525-65 Haefely bouncer circuit

750 keV TRANSPORT LINE

0210-EC-60481	chopper power supply
0211.ME-60519	chopper electrodes

TIMING AND PULSE SHIFTER

0212-EE-60192	keyswitch module
0212-ED-60201	pulse shifter
0212-ED-60215	pulse shifter doghouse wiring
0212-EC-60307	pulse shifter buffer
0214-ED-60493	prom module, first prom
0214-ED-60494	prom module, second prom

THE ALVAREZ STANDING-WAVE PROTON LINAC

0271.ME-3141	tank 1 external view
0272.ME-3885	post coupler
0272.MC-4541	post coupler
0272.MD-4581	tuning slug

The FNAL Linac

LINAC RF SYSTEMS

0230.ME-4103 RF system overhead view
0230.ME-4095 9-3/16" transmission line plan view
0275.ME-4905 gas barrier and RF drive loop
0272-MD-4027 RF drive loop
0275.MC-3882 RF pickup loop

QUADRUPOLES

0201.MC-3173 drift tube assembly
0272.MC-3018 drift tube #1 (end wall of tank)
0201.ME-3174 drift tube alignment fixture
0272.MD-4410 stem-box cover

200 MeV AREA

0310-ME-19427,28,29,30 200 Mev line components (5 sheets)
0310-MC-19634 component listing for above
031--ME-19644 200 MeV area

WATER SYSTEMS

0230-MD-4058 preaccelerator water system
0230.MC-4626 preaccelerator water flow diagram
0230.MD-4132 RF water system temperature control
0261-MD-3918 identical to 0230.MD-4132 (had to happen)
0230.MC-3874 RF water system control, A7 rack
0230-MD-4627 RF water flow diagram
0261-MD-3919 cavity water system temperature control
0230-MC-4626 cavity water flow diagram

Fermilab

TM-1238
0802.000

THE FERMILAB 200 MeV LINAC CONTROL SYSTEM HARDWARE

M. F. Shea

January 1984

11

CHAPTER 1 INTRODUCTION

CHAPTER 2 ARCHITECTURE

2.1	INTRODUCTION	2-1
2.2	THE PRIMARY STATION	2-1
2.3	THE SECONDARY STATION	2-1
2.4	OPERATING MODES	2-2
2.4.1	Local Mode	2-2
2.4.2	Global Mode	2-3
2.4.3	Gateway Mode	2-3

CHAPTER 3 OVERVIEW OF HARDWARE COMPONENTS

3.1	INTRODUCTION	3-1
3.2	THE MICROCOMPUTER	3-2
3.2.1	The CPU Board	3-2
3.2.2	Non-volatile RAM Board	3-3
3.3	INTERNAL I/O BOARDS	3-3
3.3.1	Communication Card	3-3
3.3.2	Binary I/O	3-5
3.3.3	Predet Timers	3-6
3.3.4	Multibus Analog I/O	3-6
3.4	EXTERNAL I/O COMPONENTS	3-7
3.4.1	The Local Console	3-7
3.4.2	Fiber Optic Link Repeater Chassis	3-8
3.4.3	OPTO22 Binary I/O	3-9
3.4.4	Sample-and-Hold / A-D Converters	3-9
3.4.5	External D-A Converters	3-10
3.4.6	The MIL-1553 Interface	3-11

CHAPTER 4 TIMING SYSTEM FOR THE LINAC

4.1	INTRODUCTION	4-1
4.2	MANUAL PREDET TIMERS	4-1
4.3	VARIABLE TIMING TRIGGERS	4-1
4.3.1	TDATA	4-2
4.3.2	High Voltage Dome Timing	4-2
4.3.3	Local Secondary Timers	4-3
4.4	NTF TIMING	4-3
4.5	CLOCK BACKUP MODULE	4-3

CHAPTER 5 NEUTRON THERAPY FACILITY

5.1	INTRODUCTION	5-1
5.2	THE NTF 6800 SYSTEMS	5-1
5.3	THE NTF SECONDARY	5-2
5.4	NTF TIMING	5-2

CHAPTER 6 OPERATIONAL CONSIDERATIONS

6.1	INTRODUCTION	6-1
6.2	STATUS OF THE PRIMARY	6-1
6.3	SUPERFICIAL VIEW OF A SECONDARY	6-2
6.4	WHAT DOES A RESET DO?	6-3
6.5	A VIEW FROM THE ALARMS SCREEN	6-4
6.5.1	The Alarm Messages	6-4
6.5.1.1	The OFFLINE Message	6-4
6.5.1.2	The RESET Message	6-5
6.6	THE TDATA CONNECTION	6-5
6.7	ETHERNET	6-5
6.8	REPLACEABLE PARTS AND SPARES	6-6
6.8.1	Linac Spares	6-6
6.8.2	Changing A-D And D-A Chassis	6-7
6.8.3	Local Console Replacement	6-7
6.8.4	Replacing Multibus Cards	6-7
6.8.5	Primary Station Precautions	6-8
6.8.6	Alarm Terminal	6-8
6.8.7	Ethernet Controller	6-8
6.9	ANOMALIES AND LITTLE KNOWN FACTS	6-9
6.9.1	MADC Module	6-9
6.9.2	Ethernet Controller	6-9

CHAPTER 7 SYSTEM DETAILS

7.1	THE MULTIBUS BIN	7-1
7.2	LOCATION OF SECONDARY STATIONS	7-2
7.3	PRIMARY/SECONDARY INTERRUPT LEVELS	7-
7.4	BEAM INHIBIT/ BEAM ENABLED	7-3
7.5	MEMORY MAPS FOR CPU BOARD	7-4
7.6	BINARY I/O BOARD CONFIGURATION	7-4
7.7	PIA ASSIGNMENTS	7-4
7.8	OPTO22 ASSIGNMENTS	7-5
7.9	MISCELLANEOUS ADDRESS ASSIGNMENT	7-5
7.9.1	Communication Board Addresses	7-5
7.9.2	Timer Board Addresses	7-7

APPENDIX A THE MC68000-BASED SINGLE BOARD COMPUTER

APPENDIX B THE LINAC COMMUNICATION CONTROLLER

B.1	GENERAL INFORMATION	B-1
B.2	DESIGN CONSIDERATIONS	B-1
B.2.1	The Bus Interface	B-1
B.2.2	Console Support	B-1
B.2.3	Video Generator	B-2
B.2.4	Console Serial I/O	B-2
B.2.5	SDLC/DMA	B-2
B.2.6	Byte-Serial DMA I/O	B-3
B.3	DESIGN DETAILS	B-3
B.3.1	Timing And Clocks	B-

B. 3. 2	Address Decoding	B-4
B. 3. 3	I/O Connectors	B-5
B. 3. 4	Miscellaneous Registers	B-5
B. 4	DMA OPERATION	B-6

APPENDIX C BINARY I/O CARD

APPENDIX D LOCAL CONSOLE FOR LINAC CONTROL STATIONS

APPENDIX E FIBER OPTIC LINK REPEATER CHASSIS

E. 1	INTRODUCTION	E-1
E. 2	DESCRIPTION OF THE REPEATER CHASSIS	E-1
E. 2. 1	Receiver Section	E-1
E. 2. 2	The Logic Section	E-2
E. 2. 2. 1	BYPASOUT	E-2
E. 2. 2. 2	TXCLK	E-2
E. 2. 2. 3	TXDATA	E-3
E. 2. 2. 4	RXENV (Receive Data Envelope)	E-3
E. 2. 3	Transmitter Section	E-3
E. 3	FRONT PANEL	E-3
E. 4	REAR PANEL	E-4
E. 5	HP FIBER OPTIC RECEIVERS	E-4

APPENDIX F 16-CHANNEL S/H-AD CHASSIS

F. 1	INTRODUCTION	F-1
F. 1. 1	Description Of The Chassis	F-1
F. 2	JUMPER OPTIONS	F-2
F. 3	FRONT PANEL	F-3
F. 4	REAR PANEL	F-3

APPENDIX G 16 CHANNEL DA CHASSIS

G. 1	INTRODUCTION	G-1
G. 2	GENERAL DESCRIPTION	G-1
G. 3	FRONT PANEL OF THE D-A CHASSIS	G-2
G. 4	REAR PANEL	G-2

APPENDIX H MULTIBUS 9-CHANNEL PREDT BOARD

H. 1	INTRODUCTION	H-1
H. 2	DESCRIPTION OF TIMER BOARD	H-1
H. 3	WATCHDOG TIMER	H-2
H. 4	PREDT TIMERS	H-2
H. 5	TIME BASE INPUT	H-2
H. 6	TRIGGER INPUT	H-3
H. 7	INITIALIZE	H-3
H. 8	INTERRUPTS	H-3

CHAPTER 1

INTRODUCTION

This report is a description of the present Linac distributed control system that replaces the original Xerox computer and interface electronics with a network of 68000-based stations. In addition to replacing the obsolete Xerox equipment, goals set for the new system were to retain the fast response and interactive nature of the original system, to improve reliability, to ease maintenance, and to provide 15 Hz monitoring of all Linac parameters.

Our previous experience with microcomputer installations showed that small, stand-alone control systems are rather straightforward to implement and have been proven to be reliable in operation, even in the severe environment of the 750-keV preaccelerator. The overall design of the Linac system incorporates the concept of many relatively small, stand-alone control systems networked together using an intercomputer communication network. Each station retains its local control system character but takes advantage of the network to allow an operator to interact with the entire Linac from any local console. At the same time, a link to the central computer system allows Host computers to also access parameters in the Linac.

The individual local control stations are called Secondary stations and the system that controls the communication links to Hosts and Secondaries is known as the Primary. An overall view of the organization of the Linac control system is given in Figure 1.0.




FIGURE 1.0 Organization of the Linac Control System

CHAPTER 2 ARCHITECTURE

2.1 INTRODUCTION

The Linac control system consists of seventeen stand-alone microcomputers interconnected using IBM's SDLC (Synchronous Data Link Control), a serial communication link protocol. Control stations are physically located so that each major equipment area has its own dedicated control station as shown in Figure 2.1. The communication protocol requires a Primary station to control the data transfers to the Secondary stations.

2.2 THE PRIMARY STATION

The Primary serves as the master station as defined by the SDLC Loop communication protocol and it also provides the connection between external Host computers and the Linac system as a whole. In its role as the SDLC primary, it initiates all message traffic on the link-- Secondaries may not initiate transmissions. On behalf of an external Host, the Primary will relay commands and settings to Secondaries, collect, sort and organize requested data and return it to the appropriate Host. The Primary operates synchronously with the Linac 15 Hz repetition rate.

2.3 THE SECONDARY STATION

Each Secondary station is, in effect, a complete, miniature control system. As shown in Figure 2.3.1, it contains all the essential elements of a real-time control system: a computer, analog input and output, binary input and output, and a complete console to allow operator interaction with the system. Figure 2.3.2 is a photograph of an installed Secondary showing the Multibus bin, the link repeater, four A-D and two D-A chassis.




Figure 2.1 Location of Linac Secondaries




FIGURE 2.3.1 Block Diagram of Linac Secondary




FIGURE 2.3.2 PHOTOGRAPH OF SECONDARY STATION

The Secondary operates synchronously with the 15 Hz repetition rate of the Linac. Every 66 ms, the computer reads all the analog and binary parameters in its area, updates lists of data for requestors that need selected data from its area, performs the monitoring function to look for changes in alarm conditions, adjusts parameters that are under software closed loop control, executes the application program currently selected to run on behalf of the operator at the local console, updates the local console display as required, makes settings to local or remote parameters as requested, and handles all message-oriented communications that occur on the SDLC Loop. This sequence begins about beam time when the 15 Hz interrupt arrives.

All Secondaries receive their interrupt at the same time, so the activity described above is carried out in all Secondaries in parallel.

2.4 OPERATING MODES

Three distinct modes of operation are possible; the Local mode, the Global mode, and the Gateway mode. All Modes are supported simultaneously.

2.4.1 Local Mode

In the local mode, a Secondary can be used to control all the equipment for its own area. This is a stand-alone mode which implies that no other computers or communication links are required to operate. A local control console is an integral part of each Secondary, and from this console, the user can call up groups of parameters on a small video screen, observe the present values, and make changes to the analog and digital settings for the equipment in that area. An installed local console is shown in Figure 2.4.1. A local data base stored in non-volatile memory contains the names, titles, calibration constants, settings, nominal and tolerance values, and control characteristics of each analog and digital parameter in the local system.

Software in the local stations (Secondaries) consists of 1) a system part that performs the communications, interrupt handling, data acquisition, parameter monitoring, closed loop control, and console support; and 2) a selection of application programs that provide the operator's interface with the system. System programs are written in 68000 assembly language; the application programs are written in Pascal.




FIGURE 2.4.1 PHOTOGRAPH OF A LOCAL CONTROL CONSOLE

2.4.2 Global Mode

The global mode includes all the capability of the local mode along with the additional features of the Linac system that are available when the SDLC link is in operation. The net effect of the global mode is to allow an operator at any console to control and monitor any device in the entire Linac. This added capability is transparent to the user -- the same actions performed at a console will cause readings to be returned from (or settings to be made to) parameters in either local or remote locations. In this mode, ion source parameters may be adjusted by an operator using the console at, for example, Tank 3. From a single console, a technician can check the RF system parameters for each station without walking down the entire Linac gallery.

2.4.3 Gateway Mode

The Gateway Mode refers to the operation of the Linac system as a whole by way of the communication link that exists between the Primary and external Host computers. One Host, a PDP-11/34, serves as the Linac Front End computer for the central control system. A 68000-based Host, located in the Linac gallery, provides automatic, daily operational summaries for Linac and Controls group support personnel. Communication between the Hosts and the Primary is by way of Ethernet, a 10 megabit serial protocol that is now an international Local Area Network standard.

Using the appropriate messages, a Host can request and receive Linac data, and make settings to Linac devices. When a data request is received from a Host, the Ethernet address of the Host is remembered by the Primary so that the answer messages may be returned to the appropriate Ethernet Host address. Gateway accessing of the Linac system does not interfere with Local or Global Mode activity.

CHAPTER 3

OVERVIEW OF HARDWARE COMPONENTS

3.1 INTRODUCTION

The Linac stations consist of a microcomputer based in Multibus/IEEE-796 hardware, connected to a number of I/O cards and external chassis. The bins and power supplies are standard 12 slot Intel ICS-80 products and the board complement for a typical Secondary is:

- 1- Motorola MC68000-based CPU board
- 1- Communication board
- 1- 32K byte core memory
- 1- 9-channel predet timer
- 2- Binary I/O card

Slot assignments for a typical Secondary are as follows:

- Slot 1- 9-Channel Predet Timer
- Slot 2- Binary I/O (for normal binary control and sense)
- Slot 3- Binary I/O (A/D, D/A, and OPT022)
- Slot 4-
- Slot 5-
- Slot 6-
- Slot 7-
- Slot 8-
- Slot 9- 32K Core Memory
- Slot 10-(not available- used by 32K Core Board)
- Slot 11-Communication Controller
- Slot 12-68000 CPU Card

Slots 4-8 are used for additional binary and analog I/O as required by individual Secondaries. Most of the AD/DA converters and the SDLC link repeater chassis are external to the Multibus bin. Each of the hardware components will be briefly described below. More complete descriptions for the boards are given in the Appendices.

3.2 THE MICROCOMPUTER

The CPU card and core memory card form the microcomputer that operates the Secondary station.

3.2.1 The CPU Board

Each Secondary is controlled by a 68000-based single board computer located in slot 12, the rightmost slot of the ICS-80 chassis. This computer is a Fermilab design because it was needed before commercial boards became available. General features of the board are:

- 8MHz 68000 CPU
- 8 pairs of 24 pin memory sockets
- 20-bit addressing
- 2 RS-232 Serial I/O ports
- 4 Bytes of parallel I/O
- 3 16-bit timers
- Multibus/IEEE-796 compatible
- Programmable memory addressing

The CPU card is currently configured to use 6 pairs of onboard memory sockets for program storage and two pairs for onboard scratchpad RAM. All Secondary programs reside in onboard ROM so the computer does not need to use the backplane for accessing instructions, but only for data and I/O transfers. No programs are downloaded and executed. The ROM sockets are all configured for 8K byte parts (96KB total) and the RAM sockets are for 2K byte parts (8K bytes total). Note that all Secondary CPU cards are identical -- including software. A CPU card from any Secondary can replace the CPU card in any other Secondary. Differences between systems are accommodated in data tables in the offboard non-volatile RAM board.

All the onboard I/O facilities are brought out to the annunciator panel to the right of the CPU board. Two bytes of input and two bytes of output are used for option switch inputs and LED indicators, respectively. The serial I/O ports are terminated in standard RS232 25 socket "D" connectors.

A small RESET button is provided on the CPU board, but the large RESET switch on the ICS-80 chassis may also be used.

Timers on the CPU card are not brought off the board. These timers are used for internal housekeeping functions by the system.

A complete description of the CPU card is given in Appendix A.

3.2.2 Non-volatile RAM Board

A 32K magnetic core memory board is used in most of the Secondaries (a CMOS RAM card may also be used). This is a commercially available card (Micro-Memory Model 8086) that occupies two slots in the Multibus bin. All fixed data specific to a given station resides in this memory, and therefore these cards are not directly interchangeable between Secondaries. The contents must be downloaded from the VAX. A single toggle switch, accessible from the front, activates the memory protect feature that disallows writing data into a preselected portion of the memory. Currently, the first 12K bytes are write-protected.

The 32K core board is the only Multibus memory in a Secondary station. It contains fixed data tables and the local data base information in the first 12K bytes. The remaining space is used for readings, settings, nominal and tolerance values and dynamically allocated communication buffers. SDLC communication messages are transmitted from and received into this memory under DMA control.

It is the non-volatile property of this memory board that allows a Secondary to operate stand-alone and to come up after a power outage in an operating state without the need to download information from another computer. During initialization following RESET, the most recent setting values for external parameters are sent to the hardware.

3.3 INTERNAL I/O BOARDS

The space available in the Multibus chassis is not sufficient to house all the I/O requirements for most Secondary stations. Communications cards, binary I/O and predet timers are Multibus cards and are included in the chassis. Analog I/O is placed in external chassis controlled by a Multibus binary I/O board.

3.3.1 Communication Card

A single Multibus card provides the communication I/O for each Secondary station. This card controls the console serial I/O, houses the video RAM for the console display and interfaces the Secondary to the Link Repeater chassis.

The local console is driven by asynchronous serial protocol like an RS232 terminal using a Motorola MC6850 ACIA. Each fifteenth of a second, two bytes of data are transmitted to the terminal to light the appropriate indicator lamps, and four bytes are read to acquire the pushbutton switch status, the keyboard data, and the shaft encoder knob reading. These serial signals are interfaced to the ACIA using opto-isolators to minimize ground loops between the computer and the remote terminal.

A composite video signal generated on the communication board is sent to the 5 inch monitor in the console, and an independent video output is available on a Lemo connector to use for an auxiliary display. The 32 character by 16 line display is generated using an MC6847 video display controller circuit. The display memory is in the memory space of the 68000 processor so data may be updated by simply storing ASCII data in RAM. An entire screen of information may be written in less than one vertical scan time. For flicker free operation, the program changes data only during the horizontal flyback time. Four characters are updated at the end of each 64 microsecond horizontal scan time. (~0.5 Megabaud) Limited graphics capability is supported by the 6847 controller and this allows some parameter and time plotting to be done from the local console.

The majority of space on the communications board is dedicated to the SDLC/DMA circuitry. A Motorola 6854 chip is the protocol interface to the serial link. It connects to the serial link chassis by a 26-conductor ribbon cable that carries the serial transmit/receive clocks and data along with several additional control signals intended to bypass the Secondary for several failure modes.

One megahertz serial data is received by the 6854 and automatically transferred to a 16-byte FIFO memory from where it is normally placed in memory under DMA control. The purpose of the FIFO is to allow some buffering time to prevent receiver overrun in case the DMA controller cannot immediately access the Multibus. Data to be transmitted is handled in a similar way and is stored in a transmit FIFO by the DMA controller and automatically transferred to the 6854 to be sent out serially on the SDLC link.

Interrupts generated by the SDLC and DMA controller chips are connected to Interrupt Priority Level 6 on the 68000 processor. This is the highest active level so that link interrupts may be processed quickly.

A more detailed description of the communication board is given in Appendix B.

3.3.2 Binary I/O

The binary I/O provided by the Xerox control system was organized as full bytes of output and full bytes of input. Each byte was cabled to the end rack in the control area and connected to the field wiring using AMP terminal blocks. To replace the Xerox equipment, a simple Multibus binary board was designed that provided for nine bytes of binary I/O. These signals are tied directly to the original Xerox cables. A typical RF station requires only seven bytes of binary I/O.

The binary card appears as nine bytes of memory in the I/O space of the 68000, so the binary data is handled with normal MOVE instructions. The card is constructed so each byte may be latched, and the output of the latch goes to both the 74XX output buffers and to the input of tristate buffers that allow the computer to read the stored data. To configure the byte for input, the latch is removed, and the output buffers are replaced by pull-up resistors. The tristate buffer then reads the state of the input lines. The binary I/O board receives only a power on Reset. Normal Resets do not change the stored values being output to the hardware. Although commercial Multibus boards are available with nine bytes of binary I/O, these designs use LSI interface chips that output both states during initialization. This is unacceptable for a system that is to control hardware (like vacuum valves)

To accommodate the need for stepping motor pulses (short ~20 microsec) and longer (fraction of a sec) pulses, the Secondary software forms the pulse by driving the selected bit to the active state and later returning it to the inactive state.

High power binary output and optically isolated binary inputs are achieved using OPTO22 devices. Each station has one byte of OPTO22 output and one byte of OPTO22 input.

The nine bytes of binary signals are grouped into 3 connectors each containing three bytes of signals. The 50 conductor ribbon cables interface the 24 bits to external equipment. Alternate conductors are grounded. Fused +5V power may be optionally brought off the board to power external devices (such as OPTO22 boards).

Secondaries normally have two binary I/O boards; one for normal binary interface and a second that connects to the OPTO22, the external A/D chassis and the external D/A chassis. Details of the binary I/O board are in Appendix C.

3.3.3 Predet Timers

Each Secondary has a Multibus compatible predet timer board that can output nine independent timing pulses. Each pulse is usually referenced to TZERO, the beginning of the Booster cycle, although jumpers on the board allow one timer to be triggered by the delayed output of the previous (or any other) timer. A 1MHz signal derived from the Booster clock is doubled by an onboard phase locked loop to form the 2MHz time base used by all timing channels.

Each of the 9 timers (known as timer 0, 1...8) can have a resolution of 0.5 microsec and a maximum delay setting of 16384 microsec. Timers 2, 5, and 8 have a divide-by-eight prescaler that can be enabled to allow delays that cover the full 66 ms at a reduced resolution of 4 microsec.

These predet timers use Motorola 6840 chips. The output pulse width is set to 1 microsecond, and it is buffered off the board with a transformer coupled driver that can drive a terminated 50 ohm cable.

A watchdog timer included on the board is activated during the initialization sequence by the 68000. This timer is simply a retriggerable one-shot that is set for about 150 ms. During the execution of a task triggered by the 15 Hz interrupt routine, the 68000 addresses a register location on the timer board causing the watchdog timer to be retriggered. Failure to address that particular register will allow the one-shot to time out causing the watchdog timer circuit to generate a system RESET.

In normal operation, the watchdog timer should never time out. Failure modes that can cause a watchdog timer reset are: loss of 15 Hz clock, a bus error that causes the processor to halt, or any occurrence that keeps the processor at an interrupt level so that it does not perform its normal tasks. Note that tasks are executed at the lowest level.

3.3.4 Multibus Analog I/O

Most of the analog I/O for a Secondary station is external to the Multibus bin. The exceptions are: areas that require bipolar D/A converters and the A/D and D/A converters in the H and I high voltage domes. Multibus compatible analog interface is used in the preaccelerators to allow these cards to be located in the shielded enclosure with the microcomputer. Using a large NEMA enclosure provides enough shielding to allow the system to survive arc-down of the HV terminal without interruption.

The D/A board is a commercially available 8-channel, 12-bit, bipolar unit (Datel Model ST800 DA). Internal D/As are used in systems 7, E, G, H, and I. The A/D board is a single ended 32-channel, 12-bit, +10V to -10V MADC (Datel Model ST800 S32). A/D converters in all other systems are external to the Multibus bin.

3.4 EXTERNAL I/O COMPONENTS

A typical Linac station supports 64 analog input channels and 16 D/A output channels. The present limit for a Secondary is 128 analog channels. For several reasons it is desirable to keep the A/D channels external to the Multibus bin: 1. Concentrating the analog cabling to Multibus cards is difficult, 2. Some Secondaries have insufficient space in the bin, 3. An external chassis is an electrically quieter environment for analog signals, 4. For trouble shooting and maintenance it is easier to swap out an external chassis than to replace an internal Multibus card. In any case, the amount of data transfers to I/O devices is much less than, for example, to memory cards, so a simpler, lower speed interconnecting scheme is adequate. Other I/O equipment is located external to the Multibus bin for reasons of physical placement and independent power.

3.4.1 The Local Console

A local console is included as an integral part of each Secondary station. It is normally within 20 meters of the Multibus bin, but longer distances are possible. Three coaxial cables interconnect the console and the Multibus communication board that drives it-- a transmit, a receive, and a composite video cable. The serial TX/RX cables carry 4800 baud data, and the composite video signal drives the 5 inch TV monitor in the console.

General features of the console are: a 16X32 alphanumeric display, a shaft encoder knob, a 64 key keyboard and 13 illuminated momentary pushbuttons. The lights are controlled entirely by the computer-- only the lights that are appropriate for the currently active application program will be operative. A key operated switch can disallow making settings from the local console although parameters may still be displayed normally. Details of the console hardware are given in Appendix D.

3.4.2 Fiber Optic Link Repeater Chassis

Data transmitted around the SDLC loop is carried on fiber optic medium chosen for its noise immunity and ease of incorporating stations H and I (the High Voltage domes) into the Linac system.

The fiber optic loop begins at the Primary located in Area 6, connects to G, the preaccelerator ground station and then goes to the H and I dome systems before going to B, the buncher system, and on down the rest of the Linac stations, to the 200 MeV area, and to "D" the Debuncher and back to the Primary. At each station, the data on the link is received and retransmitted by an external Fiber Optic Link Repeater chassis shown in Figure 3.4.2. Clock and data signals are encoded as a single Manchester signal so that only one fiber is needed. At each station the light signal is received, decoded, and sent to the communications card. Transmit data from the communication card is sent back to the Repeater chassis to be encoded and sent to the next station.

Repeater functions are separately powered chassis so that Secondary stations can be powered down without interrupting the link traffic. Provision is made to Bypass the local station in order to maintain the integrity of the SDLC loop, even if a problem exists in individual Secondaries. The Repeater may be switched into Bypass manually, under program control, power loss in the Secondary station, or if the 26-conductor interconnecting cable is unplugged.

In the Bypass mode, the Repeater will decode and re-encode the data it receives and transmit it to the next station. Although the decoded data is still sent to the communication board, the 68000 will not act upon any received message if the station is off-loop.

Each repeater recovers the encoded clock from the received data, and this clock is used for transmitting. Another clock is built into each Repeater for use when a "beacon" message is transmitted. According to SDLC Loop protocol, a station should receive flag characters when no messages are present. If a station does not receive flags, it must transmit a special Beacon message that includes its own address to aid in locating the link problem. The Secondary will enable its own local clock to use if the upstream link has failed. In normal operation, only the Primary uses its local clock. All other stations recover the clock from the incoming signal.

Details of the Repeater chassis are in Appendix E.




FIGURE 3.4.2 PHOTOGRAPH OF REPEATER CHASSIS


3.4.3 OPTO22 Binary I/O

Some of the binary I/O signals require higher drive outputs or isolation from local grounds. A 16-channel OPTO22 mounting rack is located at each Secondary--- one byte for input and one byte for output. Typical signals interfaced to the OPTO22 devices are: beam inhibit, RF reset, RF enable, and quad reset outputs and beam cycle input bit.


3.4.4 Sample-and-Hold / A-D Converters

Nearly all Linac signals are pulsed and need to be sample-and-held before digitizing. A combined 16-channel S-H / A-D was developed for use throughout the Linac. Signals are input on two 8-contact Burndy coax connectors to individual S-H amplifiers, and the held outputs are input to a 16-channel, 12 bit MADC module. Each Secondary (except H and I) has 3 to 5 such chassis. A single trigger causes all signals to be held simultaneously, and all chassis are normally triggered at one time using the Linac TDATA pulse.

The S-H/A-D chassis are driven by a single 50-conductor ribbon cable using one of the three-byte interface connectors on the binary I/O card. The three bytes are used as shown below. Figure 3.4.4 is a photograph of the A-D chassis.



(A) A-D DATA BYTES



(B) A-D COMMAND BYTE




FIGURE 3.4.4 PHOTOGRAPH OF 16 CHANNEL S/H A/D CONVERTER




FIGURE 3.4.5 PHOTOGRAPH OF 16 CHANNEL D/A CONVERTER

Note that data are returned left justified in a 16-bit two's-complement word and therefore higher or lower resolution A-D converters could be used for more or less precise measurements. In the data returned by the twelve-bit digitizers, the 4-bit channel number is returned in the four least significant bits of the 16-bit word.

The least significant bit of the command byte is a software-generated "digitize" strobe, and the remaining bits select both the channel number within the chassis and the chassis number. This allows the addressing of eight 16-channel digitizers or 128 analog channels. The chassis address switch is a DIP switch accessed through the rear panel of the individual A-D chassis.

Using the addressing arrangement described, all the external A-D for a Secondary station are interfaced through a single daisy-chained 50-conductor ribbon cable that plugs into any of the 3-byte binary I/O connectors.

The normal MADC module (Analogic Model 6812) is set for +10V to -10V analog input range. In Secondary stations that control RF systems, the first A-D is configured for +2.5V to -2.5V full scale range in order to improve the resolution for the RF power level measurements. These signals are typically in the range of zero to one volt. A red label reading "FULL SCALE +-2.5 VOLTS" is located in the upper right of the front panel on these units. Details of the A-D chassis are given in Appendix F.

3.4.5 External D-A Converters

External 16-channel D-A converter chassis are located in most Secondary stations. They are used to set the reference level for the 171 Linac quadrupole power supplies and the low energy beam transport magnet power supplies. Analog output of these D-A chassis are 0-10 volts.

Sixteen integrated circuit D-A converters are housed in a chassis. Each D-A chip has onboard latches and outputs current to a buffer amplifier. The Secondary processor controls the D-A chassis in a manner similar to the A-D addressing discussed in Section 3.4.4, and the interface to the D-A chassis is the same, except that all three bytes of data are output; two for data and one for command. A photograph of the D-A chassis is shown in Figure 3.4.5. Details of the D-A chassis are given in Appendix G.

3.4.6 The MIL-1553 Interface

Because of the noisy electrical environment, a serial data transmission system has been implemented in the Preaccelerator Ground station area to interface the two Haafely high voltage generators. This serial protocol is military standard MIL-1553, a 1 MHz multiplexed bus that operates in a half duplex command/response mode. The bus itself is a shielded, twisted pair, multidrop arrangement that is transformer coupled, at each drop and at each station, to provide good electrical isolation between stations. Because it is a popular standard, at least in military circles, commercially available LSI chips are available to support the data transmission protocol.

On a MIL-1553 bus, each data transfer is initiated by the controller and received by a Remote Terminal (RT). The command is followed immediately by data if the RT is to receive information. After a delay of 4-12 microseconds, the RT sends a status word back to the controller. For a transmit command the controller sends only the command and the RT returns a status word followed by requested data. An attractive feature of the MIL-1553 standard is that the RT can be simple hardware; it may be, but is not necessarily, intelligent. The maximum total length of a MIL-1553 bus is determined by specified voltage limits at the transmitter and receiver, and by the required response time of the Remote Terminal. In practice, a MIL-1553 bus can reach anyplace on a big airplane; a few hundred feet.

All words in the MIL-1553 messages are 16-bit of data preceded by a 3-bit sync character and followed by a parity bit. The MIL-1553 system can be understood by examining the command word in Figure 3.3.6. The command word consists of a 5-bit RT address, a 5-bit subaddress, a 5-bit word count, and a Transmit/Receive bit (T/R). The single word command can then address one of thirty-two RT's (actually only thirty-one are possible because RT=31 is used for broadcast commands). Within a selected RT, the controller can access one of thirty subaddresses and command it to receive or transmit up to 32 sixteen-bit words of data.

The command word establishes the general size of a MIL-1553 system; thirty-one RTs with thirty subaddresses each. Note that although each subaddress can transmit or receive up to 32 words, these words are only accessed sequentially and cannot be addressed individually or randomly. The characteristic of the MIL-1553 standard that the RT does not need to be a computer means that it can be as simple as hardware that strobes data into a latch, or enables status input data to be transmitted on the bus. It is the ease of connecting several devices in an area with a single twisted pair cable, that makes MIL-1553 an attractive standard for use in a control system.

The MIL-1553 specification defines the cabling, signal levels, termination and data protocols of this standard. It does not define the internal organization of the RT hardware beyond the connection to the bus; that is left to the system designers. The implementation chosen for the preaccelerator area uses Eurocard hardware interconnected with a simple bussed backplane. The RT electronics is contained on a single sized Eurocard (100 by 160 mm). This card provides the interface between the twisted pair MIL-1553 bus and the parallel backplane. Subaddresses within the RT are decoded on other single sized Eurocard function modules.

To control the Haefely high voltage set, only four interface cards are needed; a relay driver output card, two binary input cards and a stepping motor pulse serializer card. The Haefely is operated by actuating relays that parallel the local control panel push buttons and the digital status is read from contact closures of relays driven by the local control panel. An existing Haefely analog control module provides a 14-bit digital measurement of the high voltage terminal potential and this value is read into a binary input module. The reference voltage that determines the setting of the high voltage is generated by a D-A that is interfaced to appear as a stepping motor. That is, the computer can only send incremental up/down pulses to cause relative adjustments to the present value. This control is provided by a Eurocard module that receives a 16-bit word by a MIL-1553 command and counts the value to zero by counting up or down depending on the sign of the 16-bit word. Up/down pulses are output as CCW/CW pulse trains to the analog control module.

Physically the control electronics for each Haefely consists of a 5.25 inch by 19 inch rack-mounted Eurocard chassis containing a modular power supply, the MIL-1553 RT module and the four interface cards described above. Data that is accessed by a MIL-1553 connection has been made one of the data types recognized by a Secondary station. As with other allowed data types, this data is known to the system by entries in the data access table that resides in non-volatile core memory.

CHAPTER 4

TIMING SYSTEM FOR THE LINAC

4.1 INTRODUCTION

All Linac timing is derived from the Booster phase reversal clock. This bipolar clock is driven from the MAC room to relay rack LUO-RR1-3 in the preaccelerator area where it is received into a clock backup module and passed on for use by all Linac timers. From the clock backup module the clock signal goes to the H- and I- domes via fiber optic light pipes, to manual predet timers in the preaccelerator area and to a cable that supplies the Booster clock to all the Linac Secondary stations. A block diagram of the timing system is shown in Figure 4.1.

4.2 MANUAL PREDET TIMERS

Timing requirements for the Linac are rather modest. A fixed trigger is required for the pulsed quadrupole power supplies, the RF systems, and the 15-Hz interrupt for the Primary station. These fixed timing triggers are generated by a manual predet in the preaccelerator area. Settings of these triggers are seldom changed. In the upgrade of the Linac system, only the source of the triggers was changed--- the buffering and distribution of trigger signals along the Linac was left unchanged.

4.3 VARIABLE TIMING TRIGGERS

In the original Linac control system, the only computer controlled variable timer channel was TDATA, a pulse used to trigger all the sample-and-hold amplifiers in the Linac. When the H- preaccelerator was installed, both the dome and H- ground station used local variable predets. In the present system, stations G, H, and I are operated using variable timers that are delayed from the Pulse Shifter output.




Figure 4.1 Block Diagram of Linac Timing System

4.3.1 TDATA

The variable TDATA timer is located in Secondary B, the Buncher system, because it is the first system not timed by the Pulse Shifter. The output of the TDATA timer goes to a high level pulse repeater that drives TDATA the length of the Linac. At each Secondary, transformer coupled taps from the TDATA line passively pick off the signal for use as the sample and hold trigger.

Note that if TDATA is mistimed, nearly every monitored parameter in the entire Linac system will be reported as out-of-tolerance. As an aid to diagnosing this problem, a panel is installed in LE1-RR2-10 to allow the TDATA pulse repeater input to be connected either to the normal variable TDATA or to a test pulse T2000. This pulse is the output from a manual predet in the preaccelerator area set to 2000 microseconds from TZERO, the same as the nominal setting for variable TDATA. Selecting T2000 as the input to the TDATA pulse repeater will allow normal operation of the Linac in the event of a malfunction in the TDATA timer.

4.3.2 High Voltage Dome Timing

In normal operation, the ion source in the high voltage dome produces beam each cycle, even though Linac beam was not requested. To avoid accelerating beam that is not needed, the timing of the preaccelerator/ground-station systems is changed to cause beam to occur after the end of the Linac RF pulse. A Beam-Enable module determines that beam is being requested and other conditions are met (vacuum OK, safety system permit, ...). The enable signal is input to the pulse shifter that outputs a pulse at TZERO time for a beam cycle, or at TZERO+ims for a non-beam cycle. TZERO occurs at phase reversal time of the Booster clock, so Linac timers triggered by TZERO will be consistent with Booster time delay values. (Note that in the old Linac control system, TZERO was a pulse that preceded beam by about 10 microseconds. There is no equivalent timing pulse in the present system.)

Because of the operation of the pulse shifter, the trigger for the preaccelerator-related systems jumps from TZERO to a later time when beam is not being requested. The pulse shifter output is the start time for the preacc Secondaries, so these three systems appear to have beam continuously. The pulse shifter output is used as the 'TZERO' pulse in stations G, H, and I, and all time delays in these systems are computer controlled relative to it.

4.3.3 Local Secondary Timers

Each Secondary is supplied with a programmable 9-channel timer board. One of these timers is used for the Secondary's own 15 Hz interrupt that determines when the processor reads its analog and digital data. Other channels, accessible from the back panel of the Multibus bin, are for local and future use.

4.4 NTF TIMING

A manual predet timer, triggered at TZERO time, is installed in the Neutron Therapy Facility. It provides the timing used in the 6800-based system that collects and accumulates the neutron dose data.

4.5 CLOCK BACKUP MODULE

The loss of 15Hz triggers is harmful for some Linac systems. Therefore, a clock backup module was developed to insure a continuous source of timing triggers in the absence of the Booster clock.

This module receives the Booster clock, decodes the phase reversals, regenerates the clock signal and outputs it as a normal bipolar clock. All phase reversals are transformed to single phase reversals so that manual predet timers can trigger continuously at 15 Hz by selecting a "1" on the digiswitch.

This module includes a 1 MHz crystal oscillator and a 15 Hz generator locked to the 60 Hz line. If either the phase reversals or the 1 MHz are lost, the backup module will continue to output a clock using its own internal generators. These operate independently so that if only the phase reversals are lost, the output clock will be the Booster 1MHz with the locally generated 15 Hz phase reversals.

It is intended that an additional clock backup module be placed in each high voltage dome to insure continuous triggers for the ion source in the event the clock from the ground station is lost.

This module also outputs a TTL-compatible 1 microsecond pulse at 15 Hz and a trigger that occurs for a three phase reversal event on beam cycles. For use in the domes, the backup module should encode its phase reversals at pulse-shifter-out time. The pulse shifter trigger is sent to the dome as missing pulses on the clock light pipe. Provision is made to input to the backup module, an external

trigger to be encoded as a phase reversal on the backup module's output. This feature will be used in the dome systems only, where the pulse shifter signal is input to make it the reference time for dome timers.

When the input clock to a backup module fails, the next output trigger will be delayed by at most one fifteenth of a second. No two phase reversals are allowed to occur closer than three line cycles (about 50ms). This eliminates the possibility of two closely spaced TZERO pulses when switching from internal to external clock.

CHAPTER 5

NEUTRON THERAPY FACILITY

5.1 INTRODUCTION

Because the Neutron therapy facility is a user of Linac beam, its controls are integrated into the Linac control system. The original 6800-based beam line microcomputer was left intact and interfaced to a normal Secondary station (Sec C). This Secondary receives data from, and makes settings to, both the 6800 system and its own analog and binary interface equipment.

5.2 THE NTF 6800 SYSTEMS

In 1975, a 6800-based system was put into operation to allow local control of the NTF facility. This system consists of two microcomputers; the Medical Control Room microcomputer and a Beam Line microcomputer interconnected by a byte serial FIFO link. The Beam Line system reads and integrates ion chamber data to provide dose accumulations for patient treatment and radiobiology studies. Because of the critical nature of the function it provides, considerable effort has gone into testing and verifying the operation of this system to insure its readings and calculations accurately record the actual dose the patient receives. We therefore decided to leave the 6800 system intact and simply connect it to Sec C in the same way it was interfaced to the Xerox 530. The NTF system is then essentially unchanged, except that its data is available to the new control system.

Originally, the Beam Line system had a second FIFO port to connect with the X530 and the Medical Control Room. The connection to the Medical Control Room is also the same, so the operation of the facility by the medical technicians is unchanged.

5.3 THE NTF SECONDARY

Secondary C, the NTF station operates as a normal Secondary. It contains all the functions and controls all the devices that pertain to the operation of NTF. Data acquisition is a combination of readings from Sec C's own analog and binary interface, plus data received from the Beam Line 6800 system. The choice of which system reads a given channel was made by simply accepting from the 6800, any data it is reading and interfacing all other channels directly to Sec C. In this way, there is exactly one reading of each parameter, and Sec C has a copy of the values being used internally by the 6800. In all other respects, Sec C is a normal station, so NTF data is available to any Linac Host or Secondary.

5.4 NTF TIMING

To provide a hardware source of interrupts for the 6800 Beam Line system, a manual predet timer is installed near the 6800. This timer is driven by a Booster clock from the MAC room and uses a TZERO trigger as its external reference. Two output pulses from this timer are used to provide the two interrupt triggers needed by the Beam Line system.

CHAPTER 6

OPERATIONAL CONSIDERATIONS

6.1 INTRODUCTION

A cursory check of the operation of the Linac system is available on the top line of most Linac programs on the control room consoles. There the operator will see a green character to indicate the station number of each secondary that responded to a test poll issued by the Primary. There are 30 possible secondaries but only 18 currently exist. They are numbered 1...9, A, B, C, D, E, G, H, I, V corresponding to Linac tanks 1 through 9, the test station (system 10), buncher, NTF(CTF), debuncher, 200MEV, preaccelerator ground station, H minus dome, I minus dome, and the primary console, respectively. A system that fails to respond will be replaced by a dot on these displays. If only dots appear on the top line when Linac application programs are called up, then the Linac Front End PDP-11 is down or the Linac primary is down or the SDLC Link is not working.

6.2 STATUS OF THE PRIMARY

The Primary system is located in the controls area of Linac system 6. It is a dual - 68000 processor system; one processor is a dedicated link driver and the other drives both the Ethernet link to the Front End and its own console.

A small video screen above the primary is controlled by the Link Driver. Figure 6.2 is a copy of this display. Line 1 is the title, date and time. Line 4 is again the display of Secondaries that responded to the test poll. Line 3 shows the Secondary stations that are currently returning data to requestors. The requestor may be a Host like the Front End or a Secondary. If a Secondary is a requester receiving data, its position on line 3 will be displayed as inverse video (non blinking). On line 2, the numbers following the H: and L: are the numbers of active lists of data that are being collected by the Primary on behalf of the Host and the Linac stations. The time given on Line 2 shows when the Primary has finished collecting

Number of active Host(H:) and Local(L:) lists

Time within this cycle that all answers are returned

NBS time MM/DD HHMM:SS

T: Number of online Secondaries (Test response OK)

C: Number of commands transmitted this cycle

A: Number of alarms received this cycle

Number of lists returned by each Secondary

Stations responding to test poll

Diagnostic information output after a fault

Includes: Fault type

Date/Time

Link activity

Processor status

Register contents

```

LINAC/PRIMARY LINK 11/03 0925:17
H: 4 L: 7 22MS T: 17 C: A:
... 1... 112.1.1... 1.2...
. 123456789ABCDE.GHI...
BE D0=000003FF A0=00000002
FC=32C5 00000000 A1=00040000
AA=00040000 00000008 A2=00000678
IR=32C1 1013FFFF A3=00000E00
SR=2704 0000FFFF A4=00042000
PC=0002214C 00000008 A5=00000780
US=00000500 0004F6C0 A7=00000234
LINAC PRIMARY LINK 10/14 0517:37
H: 3 L:
.....
. 123456789AB.DE.GHI...

```

Figure 6.2 Status Display for Primary Station

data and has sent the last list of data to a requestor. The number displayed measures milliseconds from 15 Hz interrupt time. Currently, the Primary begins collecting data at ten ms into the cycle. The C: and A: give a count of the number of command and alarm messages that were serviced during the present Linac cycle.

If the Linac SDLC link is broken, then the next operating secondary will send a special "Beacon" message to the Primary. This feature of the SDLC Loop protocol is intended to help locate a link problem. Indication of a beacon message is shown by a station number on line 3 that is blinking between normal and inverse video. The problem then exists between that station and the next station upstream. The order of stations is G H I B 1 2 3 4 C 5 6 A 7 8 9 E D so a beacon from station 5 would indicate a problem with station C. A Secondary station that is offline or even powered down will not cause a beacon message because the Link repeater will bypass the station and maintain link continuity.

Inside the door covering the Multibus chassis, interrupt and task activity can be observed in the LEDs of the annunciator panel. If no activity is seen, the Primary is down. This is unusual because both the software and the hardware try to restart the system under normal failure modes. Note that these systems will not be found in mode where the processor is randomly executing meaningless data; the 68000 will typically enter its halt state if it receives two bus errors or address errors in a row. If the system tries to restart itself after an error condition and finds that it has had fifteen such conditions within the past hour, it will halt. A manual reset will allow the processor to start, but how long it runs will depend upon when it receives its fifteenth abort within a one hour period. Typical error or abort conditions include bus errors, address errors, unimplemented interrupts, illegal instructions and trap errors.

If the primary is found to be halted, the operator may try to RESET it, but if it does not start a more severe problem exists and someone should be called. There may be some helpful clues on the Primary console or on the Link Driver TV display.

6.3 SUPERFICIAL VIEW OF A SECONDARY

A Secondary station contains a single 68000 processor. As with the Primary, the interrupt and task level activity can be observed from the front of the Multibus chassis. No activity indicates the processor is halted and only a reset will restart the processor.

When a Secondary is operating, activity will be seen on the LEDs of the communicator panel, and the S-H/A-D chassis below the Multibus bin. The LEDs will exhibit a 15 Hz repetitive pattern that is nearly identical in all stations. The green LED on the front of the local console should also flash at 15 Hz.

A Secondary that is suspected of malfunctioning may be reset to see if the problem disappears. Some of the LSI chips in the system are software configurable and on rare occasions the internal configuration may be changed by some severe external perturbation - like a lightning strike.

When checking on a Secondary, the status of its Link Repeater chassis should also be checked to insure that the station has not been bypassed. A bypass condition is indicated by the large red "BYPASS" LED on the Link Repeater front panel, and the cause of the Bypass condition is shown by small yellow LEDs. The secondary's yellow "TX CLK ENABLE" LED should be off. (Only the Primary TX CLK ENABLE LED will be on for normal operations). Constant or 15 Hz activity should be observed on the green indicators of the Link Repeater chassis. A Secondary's TX CLK ENABLE indicates the local station is sending a Beacon message because it is not receiving data from upstream.

If local parameters are available on the local console, but no remote data is displayed, the problem is related to communications. If some, but not all, remote data is available, the problem is with the stations that do not respond - the link and local station are probably OK.

6.4 WHAT DOES A RESET DO?

A reset will restart the 68000 from the very beginning of the program. There is no attempt to perform a "warm" Reset - all resets, whether generated by software or hardware, manual or automatic, perform the same restart of the system. During initialization all the analog and digital settings are sent to the hardware, but in general no effect will be observed in the external equipment because the settings are the same as those stored in the hardware. When a Secondary is powered up it retransmits the analog and digital settings from non-volatile RAM where they are stored.

A Reset will also reinitialize all the programmable LSI controllers in the system. In general, if a system is suspected of malfunctioning it may be reset to see if the apparent problem disappears.

6.5 A VIEW FROM THE ALARMS SCREEN

Alarms and Limit monitoring of Linac analog and digital parameters has been operational since the system was installed. Each Secondary provides the monitoring function for its own parameters and the necessary alarm messages are collected by the Primary stations. Each channel can be monitored to inhibit Linac beam for out-of-tolerance parameters. Because each station has only its own parameters to monitor, only a few milliseconds (4 ms typically) are spent to perform this task. All channels are monitored every Linac cycle and beam is inhibited on the next cycle if necessary.

6.5.1 The Alarm Messages

The format of analog and binary alarm messages that appear on the alarm screen contains the date and time, cycle number, and name or descriptor.

Usually the alarm messages relate to Linac devices, but self checking features of the system cause alarm messages that indicate problems with the systems themselves. Of particular interest are the "SEC #n OFFLINE" and the "SEC #n RESET".

6.5.1.1 The OFFLINE Message -

The SECONDARY #n OFFLINE message is generated by the Primary whenever a Secondary fails to respond to a test poll. One of the possible 32 Secondaries is sent a test poll each fifteenth of a second, so every station will be tested every two seconds. A binary status bit is set or cleared by the Primary to record the response received from each station. These status bits are then monitored with the 2x filter to inhibit beam for improper response received from the existing stations.

A Secondary can be "OFFLINE" for many reasons; powered down, halted, manually bypassed, automatically bypassed, being reset, faulty communication board, and SDLC link problems can all cause OFFLINE messages. In general, it's a bad omen and bears watching.

6.5.1.2 The RESET Message -

The SEC #n RESET message is a comment that is generated by the Secondary to announce that it has just been Reset and reported by the Primary along with Analog and Binary alarms. During the reset of a Secondary, that Secondary may not receive a test poll and therefore an OFFLINE condition may not be detected. However, a Secondary that has been reset, for whatever reason, will voluntarily send the RESET message. These messages are rare and indicate some traumatic experience at the Secondary. They have been observed occasionally during heavy crowbar activity of an RF system or severe arcing of the Preaccelerator. The RESET comment may, but need not, be accompanied by an OFFLINE message. The time associated with the RESET message is the time last known to the system before it RESET; hence it is the time the system went down. At each even minute, the time of all systems is updated with the value received from the NBS clock in the Main Control Room.

6.6 THE TDATA CONNECTION

Most Linac Parameters are sampled and held at a variable time, TDATA, that normally occurs at 2000 us into the cycle. It is generated by a variable timing channel in system B, buffered by a pulse repeater, driven down the entire length of the Linac, and transformer coupled at each station to trigger the sample and hold amplifiers. When TDATA fails for two successive Linac cycles, several hundred messages will be generated and the alarms screen will begin to scroll indicating all manner of bad news. Such a flood of messages usually suggests a problem with TDATA. One simple check can be made to determine if the variable timer has failed. At relay rack LE1-RR2-10 a panel has been installed to allow the TDATA pulse repeater to get its input from the variable TDATA or from a manual predet pulse called T2000. As its name suggests, T2000 arrives at 2000 us and the use of this pulse could allow continued operation of the Linac if the variable TDATA fails. Selecting between the two sources is done by simply moving a BNC cable from TDATA to T2000.

6.7 ETHERNET

The Linac Primary to PDP-11 Front End Host connection is made using an Ethernet link. Besides being a fast, widely accepted international standard for intercomputer communication, Ethernet has the prominent feature of being supported by commercial suppliers who make interface cards for both Unibus and Multibus. Ethernet is a half duplex, serial, 10 megabit protocol. By way of the Ethernet

interface, requests for data from Host computers are sent to the Linac Primary and answers are collected and returned to the host by the primary at the specified repetition rate. For repetitive data, the host only needs to ask once and data will be returned till the request is cancelled. A typical parameter page on a console will cause a short (~200 us) burst of data to be returned at 15 Hz.

Because an Ethernet cable can support many (up to 100) nodes, multiple Host Computers can be connected simultaneously. The Hosts "learn" the (6 byte) address of the Primary from a broadcast transmission that is sent by the Primary at 17 sec intervals. The Primary remembers the address of the Host that was included in the request for data so the answers may be returned to the proper address.

In the Linac system there are only two Hosts - the Front End PDP-11 and a 68000-based Host located next to Sec 3 in relay rack LE3-RR3-1.

The Ethernet connection has been very trouble-free particularly since the new generation controller card was installed in the PDP-11. However, if trouble with the Ethernet link is suspected, a quick check of the 68000-based Host can determine whether the Link is working or not. This Host operates in the same way as a normal secondary, except that its only connection with the system is to the Primary via the Ethernet cable. If the 68000 Host can get data from the Primary, the Ethernet link works. In that case, the suspected problem must be related to the Front End.

6.8 REPLACEABLE PARTS AND SPARES

Considerable effort has gone into making the Linac system very modular and easy to service. The Link Repeaters, S-H/A-D converters and D-A converters are housed in external chassis.

6.8.1 Linac Spares

Hot spares for the Linac system are available in System A, the Linac test station. This system is normally powered on with a full complement of cards and chassis. Binary I/O cards, CPU and communication cards from this system should operate in any station. Spare ADs and DAs are also maintained in this area. Note that the contents of the core memory cards are peculiar to each station because this memory contains the data base. Although System A is a source for a "known good" core memory, the contents would have to be downloaded from the VAX if the core memory is moved to another station.

Another viable, but less certain source of warm spares is the small Primary-and-2-Secondaries test setup in the lab at Linac station 2. This setup is maintained to test cards offline. It is also a source for spares for the Primary station.

A third source for spares is in a cabinet near 7835 for Linac system 2. An attempt is being made to outfit this cabinet with spares that have seen several months of continuous service online in the system.

6.8.2 Changing A-D And D-A Chassis

If an A-D or D-A chassis has been determined to be faulty it may be replaced with a spare. Two Burndy 8-conductor coax connectors, a ribbon cable, one BNC cable and the AC cord are unplugged and the chassis is free to be swapped with the spare. The chassis address switch setting of the spare should be set to match the unit being replaced. Two types of A-D chassis exist, +2.5V F.S. and +10V F.S. The 2.5V units are identified by an obvious red label in the upper right corner of the front panel. Only one of the 2.5V units is used per system and they are used only in systems 1, 2, 3, 4, 5, 6, 7, 8, 9, B and D.

D-A chassis are nearly identical to replace. They also have a chassis address switch that must match the online unit being replaced.

6.8.3 Local Console Replacement

Local consoles are connected to the system by three BNC coax cables; transmit, receive and composite video. All consoles are identical so any malfunctioning console may be replaced with a spare. Since the architecture of the Linac system allows any parameter to be accessed from any terminal, the consoles at neighboring stations can be used, in many cases, to accomplish local control.

6.8.4 Replacing Multibus Cards

Only normal care must be exercised to replace Linac system Multibus cards. Because Multibus uses printed circuit edge connectors, the bin must be powered down before replacing cards. Some cards access a reset line on the upper, P2, connector, so care should be exercised to locate the replacement card in the same slot as the original.

Two types of binary I/O cards exist although they differ only in the arrangement of input and output bytes. The two types are identified by red or blue dots on the card ejectors. When replacing binary boards, the original must be replaced with the same type.

6.8.5 Primary Station Precautions

The Linac Primary is a dual processor system and contains the Link Driver and the Primary Console CPU cards in slots 12 and 11 respectively. They each have communication cards in slots 10 and 9. The Link Driver connects to a fiber optic link repeater chassis but has no console; the Primary Console CPU has a local console connection but does not attach to a link repeater.

When replacing CPU and communication cards in the Primary station, spares should be taken only from the test station Primary in the lab. Interrupts and software for the Primary differ from a normal secondary. The Primary in an SDLC system operates in full duplex rather than the loop mode used in Secondaries, so the Fiber Optic Link Repeater chassis is special for the Primary. Because all messages originate in the Primary, the link repeater chassis is the source of the 1 MHz clock used for the entire link. Its "TX CLK ENABLE" LED is always on (the secondary TX CLK ENABLE LEDs are off unless the station is transmitting a beacon message).

The primary contains one binary I/O card it uses to inhibit beam. This is a standard (red dot) binary I/O card.

6.8.6 Alarm Terminal

For reporting alarms, the primary console outputs the messages using ACIA-2, an RS232 port on the Primary console CPU. These messages go to a TEK terminal and the video is buffered to the Main Control Room. In case of failure, a second TEK terminal is available. Both are located in the relay rack with the Primary station (LE6-RR2-1).

6.8.7 Ethernet Controller

The connection to the Linac Front End PDP-11 is made using Ethernet. The Ethernet controller board is located in slot 2 of the Primary Multibus bin and connection to the transceiver cable is by way of a 16 conductor ribbon cable. The 68000-based Host in Area 3 has an identical controller that is available for use as a hot spare.

6.9 ANOMALIES AND LITTLE KNOWN FACTS

As with any system, there are a few curiosities in the operation of some of the equipment.

6.9.1 MADC Module

On rare occasion, an MADC module used in the A-D chassis reports a full scale value for all of its sixteen analog channels. Although Analogic denies that such a thing can happen, we have observed the effect a few times. It happens only when the chassis is being powered up and the cure is to power cycle the chassis. In normal operation the A-Ds are never turned off so the effect will only happen when recovering from a power outage.

6.9.2 Ethernet Controller

In the past, the Multibus Ethernet controller has been observed in a state where it fails to transmit. Status of the board (address V:1E001) indicates the transmit buffer is being used to transmit. The proper status is \$CC; the anomalous value is \$EC. This fault is cleared by unplugging and replugging the 16-conductor ribbon cable on the card edge. Again, the manufacturer claims never to have seen this effect. Power cycling the primary will also clear the fault, but a reset of the primary does not effect it.

CHAPTER 7

SYSTEM DETAILS

This section contains details of the implementation and is included here for completeness.

7.1 THE MULTIBUS BIN

The Multibus chassis used in all Linac stations is a standard Intel ICS-80. The 4-wide card cages are not used, but have been replaced by an Electronics Solutions 12-slot card cage and backplane. This was done to eliminate the segmented backplane that comes from stacking three of the 4-slot backplanes; Electronic Solutions 12-slot backplanes are continuous. Also, a 12-slot backplane in an ICS-80 chassis gives a board spacing of 0.75 inches, instead of the normal 0.6 inches, to allow better cooling. The ICS-80 chassis itself was chosen because of the vertical format and front access to the Multibus cards. Vertical card orientation improves cooling air flow and reduces board warpage.

The ICS-80 chassis has a key-operated power switch and two push buttons labeled "RESET" and "INTR". The signal generated by the INTR pushbutton connects to the INT7 (Pin P1-36) of the Multibus backplane. It causes a Level 7 interrupt on the CPU board, which is the highest priority, non-maskable interrupt. The Level 7 Interrupt service routine causes a restart of the program.

The RESET pushbutton connects to pin 38 of the P2 connector -- a pin designated as external RESET by Intel. (This pin has since been reserved by the IEEE-796 committee, and subsequently redefined by Intel as part of the iLBX bus.) In the Linac system, pin 38 of P2 connects to the slots occupied by CPU cards, timer boards and core memory boards. Both the core memory and the timer boards are capable of initiating a system reset.

7.2 LOCATION OF SECONDARY STATIONS

Secondary stations are located close to the equipment they control, usually in the same bay of relay racks with the end rack containing the AMP terminal strip. The location of each secondary is listed below.

Sec. #	Location
1	LE1-RR4-2
2	LE2-RR4-2
3	LE3-RR3-2
4	LE4-RR3-2
5	LE5-RR3-2
6	LE6-RR3-2
7	LE7-RR2-2
8	LE8-RR2-2
9	LE9-RR2-2
A	System 10 Area
B	LE1-RR2-9
C	LU5 Area
D	Booster Gallery
E	LG1 Area
G	H- Control Racks
H	H- HV Dome
I	I- HV Dome
Pri.	LE6-RR3-1
Host	LE3-RR3-1

Note that the Secondary numbers are mnemonically chosen and are different from the installed order on the loop. The order of Secondaries around the loop is:

G H I B 1 2 3 4 C 5 6 A 7 8 9 E D

This order on the loop is necessary to identify the location of a problem when a "Beacon" message is received by the Primary.

7.3 PRIMARY/SECONDARY INTERRUPT LEVELS

Interrupt levels for the Primary Console, Link Driver, and Secondary Stations are listed IN Table 7.3. The INTn values are the Multibus backplane interrupt lines. Unless otherwise noted, the INTn line used is the same as the interrupt level. Unused interrupt levels are connected to the INTO line on the backplane. The timer interrupt lines are from the MC6840 that is on the CPU board. PIA chips are also on the CPU card, but these interrupt levels are input using the Lemo connector on the display panel. The front panel board includes circuitry to condition external strobes before inputting to the control line inputs of the PIA.

LEVEL	SECONDARY	PRI. CONSOLE	LINK DRIVER
7	ICS-80 INT	ICS-80 INT	ICS-80 INT
6	ADLC/DMA		LINK RECEIVE
5			
4	SER. CONSOLE	SER. CONSOLE	SER. CONSOLE (INT3 Line)
3		ETHERNET (INT2 LINE)	
2	CONSOLE/MOTOR	CONSOLE/MOTOR	CYCLE EVENT
1	15 HZ	15 HZ	15 HZ

TABLE 7.3 Interrupt Assignments for Linac Systems

7.4 BEAM INHIBIT/ BEAM ENABLED

Some Linac parameters are monitored to turn off the beam if out-of-tolerance conditions are found. The mechanism for inhibiting beam should not depend upon the link or any other Secondary, so a hardwired "open collector" type connection is made between the pulse shifter electronics and the Secondary stations. A three pair cable is daisy-chain connected to each Secondary. One pair is dedicated to beam inhibit; one for informing each Secondary that the beam is enabled; one is spare.

The Inhibit signal is powered in the Preaccelerator area and connected to an output OPT022 at each Secondary. Any Secondary can then short the pair, resulting in an inhibit signal for the pulse shifter.

The Beam Enabled signal is powered in the Preaccelerator area and detected by an OPT022 input module at each Secondary. The beam enabled signal is used by the Secondary to differentiate between beam and no-beam cycles.

7.5 MEMORY MAPS FOR CPU BOARD

The memory map shown in Figure 7.5 is the map used for all Secondary stations and for the Primary console. The organization of memory is determined by the programmable logic devices on the CPU card. Eight memory socket pairs on the CPU card are mapped for 2 Kbyte RAM chips in socket pairs 0, 1, 2, and 3 and for 8 Kbyte PROMS in socket pairs 1, 2, . . . , and 7. Secondaries now have 8 Kbytes of scratch RAM in pairs 0 and 1, leaving space for 96 Kbytes of PROM space (64 Kbytes are currently installed). The 32 Kbyte core memory board at \$10000 is the only offboard RAM used in Secondary stations. The Primary system has a 128 Kbyte RAM board located at \$40000.

With reference to Figure 7.5, the video RAM block at \$4000 is an offboard memory area. The video RAM and peripheral registers on the communication board occupy \$400 memory locations. All systems have communication card at \$4000 and the Primary has a second communication card at \$4400 for use by the Primary Link Driver.

Onboard I/O is at \$5000. Because a 68000 has no I/O instructions, a separate block of memory is decoded and assigned for use as Multibus I/O. This block, from \$6000 to \$6FFF, transfers data using the Multibus IDRC and IOWC strobes on the backplane. To the 68000, I/O is just a block of memory.

7.6 BINARY I/O BOARD CONFIGURATION

All binary I/O is supplied by two nine-byte boards that are the same except for the arrangement of buffers and resistor terminations on the interface lines. Figure 7.6 shows the configuration of buffers, latches and resistors for the two board types. The Buncher system (B) has an additional (blue dot) board installed to operate the stepping motors associated with the emittance probes.

7.7 PIA ASSIGNMENTS

The parallel I/O facilities of the 68000 CPU board are used for the front panel. Two PIAs provide four bytes of binary I/O and eight control lines. The A side of each PIA is used for binary input to sense an eight pole DIP switch. The B sides are configured for outputs, buffered, and used to drive LED indicators and test points that show the interrupt and task activity of the 68000. Only four of the eight control lines are used. Interrupt signals can be input on Lemo connectors on the front panel. These drive the CA1 inputs of the two PIAs. The CA2 outputs are

0000	MEM-0
1000	MEM-1
2000	MEM-2
3000	MEM-3
4000	Video RAM
5000	Onboard I/O
6000	OffBoard I/O
7000	
10000	
	32K Core Memory
18000	
20000	
	MEM-7
24000	
	MEM-6
28000	
	MEM-5
2C000	
	MEM-4
30000	
	MEM-3
34000	
	MEM-2
38000	
	MEM-1
3C000	
40000	
	Offboard Memory
80000	

FIGURE 7.5 Memory Map for Linac CPU Boards




Figure 7.6 Buffer Arrangement for Binary I/O Cards

DATE: 8/03/83

P1A 1 A-SIDE: 68000 SYSTEM STATION ADDRESS

ADDRESS: DATA REGISTER 5041 ; CONTROL REGISTER 5045
CONFIGURATION: DATA DIR REG. 00, CONTROL REG. 04

CONTROL BITS: CA1- NOT USED
CA2- NOT USED

DATA: PA 7- SPARE
PA 6- SPARE
PA 5- SPARE
PA 4- STATION ADDRESS BIT 4
PA 3- STATION ADDRESS BIT 3
PA 2- STATION ADDRESS BIT 2
PA 1- STATION ADDRESS BIT 1
PA 0- STATION ADDRESS BIT 0

P1A 1 B-SIDE: 68000 SYSTEM INTERRUPT LEVEL INDICATORS

ADDRESS: DATA REGISTER 5043 ; CONTROL REGISTER 5047
CONFIGURATION: DATA DIR REG. FF, CONTROL REG. 04

CONTROL BITS: CB1- NOT USED
CB2- NOT USED

DATA: PB 7- SPARE
PB 6- 6854/6844 SDLC/DMA INTERRUPT ACTIVITY
PB 5- SPARE
PB 4- 6850 SERIAL CONSOLE RECEIVE
PB 3- SPARE
PB 2- 6840 TIMER (MOTOR INT, CONSOLE INT)
PB 1- 6821 P1A-0 15 HZ INTERRUPT
PB 0-

TABLE 7.7a Bit Assignments for P1A1

DATE: 8/03/83

PIA 2 A-SIDE: 68000 SYSTEM MODE SWITCHES

ADDRESS: DATA REGISTER 5040 ; CONTROL REGISTER 5044
CONFIGURATION: DATA DIR REG. 00, CONTROL REG. 07

CONTROL BITS: CA1- 15 HZ INTERRUPT
CA2- NOT USED

DATA: PA 7- SBUG/SYSTEM SELECT ON RESET
PA 6- INHIBIT AUTO SETTING RESTORE
PA 5-
PA 4-
PA 3- MESSAGES TO ACIA2
PA 2- MESSAGES TO CONSOLE
PA 1- ALARMS INHIBIT
PA 0- COMPUTER SDLC LINK BYPASS

PIA 2 B-SIDE: 68000 SYSTEM TASK INDICATORS

ADDRESS: DATA REGISTER 5042 ; CONTROL REGISTER 5046
CONFIGURATION: DATA DIR REG. FF ; CONTROL REG. 04

CONTROL BITS: CB1- NOT USED
CB2- NOT USED

DATA: PB 7- MESSAGE MONITOR
PB 6- DATA LIST UPDATING
PB 5- DATE AND TIME
PB 4- SMALL MEMORY DUMP
PB 3- APPLICATION PROGRAM
PB 2- CONSOLE PROCESSING AND KEYBOARD INTERRUPT
PB 1- ALARMS AND CLOSED-LOOPS
PB 0- PROCESS LINK COMMANDS

TABLE 7.7b Bit Assignments for PIA2

DATE: 8/03/83

OUTPUT BYTE ADDRESS: \$6020

OPTO22 TERMINAL # SIGNAL,RTN	BINARY I/O CONNECTOR PIN #	BIT #	FUNCTION
1,2	47	0	BEAM INHIBIT
3,4	45	1	COMPUTER ENABLE
5,6	43	2	
7,8	41	3	
9,10	39	4	QUAD RESET
11,12	37	5	SYSTEM RESET
13,14	35	6	HV OFF
15,16	33	7	HV ON

INPUT BYTE ADDRESS: \$6022

OPTO22 TERMINAL # SIGNAL,RTN	BINARY I/O CONNECTOR PIN #	BIT #	FUNCTION
17,18	31	7	BEAM ENABLED
19,20	29	6	SPARE DAISY CHAIN INPUT
21,22	27	5	SPARE
23,24	25	4	SPARE
25,26	23	0	SPARE
27,28	21	1	SPARE
29,30	19	2	SPARE
31,32	17	3	SPARE

- NOTES: 1. ALL EVEN CONNECTOR PINS ARE COMMON
2. CONNECTOR PIN #49 IS FUSED +5 VOLTS
3. ALL DC INPUT SIGNALS ARE ASSUMED POSITIVE

TABLE 7.8 Bit Assignments for OPTO22 I/O

buffered to Lemo connectors on the front panel. The CB1 and CB2 signals are not used. PIA configuration data and bit assignments are given in Table 7.7.

7.8 OPTO22 ASSIGNMENTS

A 16-channel OPTO22 mounting rack is installed in each Secondary. Eight bits of input and eight bits of output are available. For completeness, the pin assignments and OPTO22 terminal numbers are documented in Table 7.8 below along with the current bit assignments for a typical station that includes an RF system.

7.9 MISCELLANEOUS ADDRESS ASSIGNMENT

This section lists the addresses that have been assigned to various devices in the Linac system. They are included here for reference.

7.9.1 Communication Board Addresses

The base address of the Communication card is \$4000. Listed in Table 7.9.1 are the base addresses of the LSI chips along with the address of various registers on the board. Bit assignments for the registers are included. Registers internal to the LSI chips are addressed as described in the Motorola data sheets with the base address added as an offset.

Addr.		Function
\$4000	R/W	Video RAM
\$4200	R/W	Tx,Rx FIFO
\$4220	R	Misc. Status
\$4240	R/W	Parallel Tx,Rx FIFO
\$4260	R	FIFO Status
\$4260	W	FIFO Clear
\$4300	R/W	ACIA Base Address
\$4340	R/W	DMA Base Address
\$4380	R/W	SDLC Base Address

TABLE 7.9.1a Communication Board Addresses

Bit No.	Function
7	
6	
5	
4	
3	Clear SDLC Rx FIFO
2	Clear SDLC Tx FIFO
1	Clear Parallel Rx FIFO
0	Clear Parallel Tx FIFO

FIFO Clear Register (\$4260 Write)

Bit No.	Function	
7	IR SDLC Rx FIFO	
6	OR SDLC Rx FIFO	
5	OR SDLC Tx FIFO	
4	IR SDLC Tx FIFO	
3	IR Parallel Rx FIFO	IR=Input Ready
2	OR Parallel Rx FIFO	OR=Output Ready
1	OR Parallel Tx FIFO	
0	IR Parallel Tx FIFO	

FIFO Status Register (\$4260 Read)

Bit No.	Function
7	TV Ready
6	FLAG DETECT Missing
5	Link Bypass
4	
3	
2	
1	
0	

Miscellaneous Status Register (\$4220 Read)

TABLE 7.9.1b Bit Assignments for Comm. Controller Registers

7.9.2 Timer Board Addresses

The timer board contains three Motorola 3-channel timer circuits and the system watchdog timer. A WRITE operation to the watchdog timer address triggers a one-shot; the data is ignored. Addresses for this board are listed below.

Addr.	Function
\$6100	TIMER CHIP 0 BASE ADDRESS (CHAN 0, 1, 2)
\$6108	Timer Chip 1 Base Address (Chan 3, 4, 5)
\$6110	Timer Chip 2 Base Address (Chan 6, 7, 8)
\$6118	Watchdog Timer Trigger

APPENDIX A

THE MC68000-BASED SINGLE BOARD COMPUTER

The MC68000 CPU board used in the Linac Control System is described in Fermilab FN-330.

APPENDIX B

THE LINAC COMMUNICATION CONTROLLER

B.1 GENERAL INFORMATION

A multifunction Multibus compatible communications controller has been developed for use in the Linac control system. Figure B.1 is a block diagram of this controller. Included on the board are:

- Synchronous Data Link Controller (SDLC)
- Direct Memory Access controller (DMA)
- Byte Serial I/O port
- Asynchronous serial port
- Video RAM display generator

In operation the Communication controller connects the local Linac stations to the SDLC link repeater chassis, and supports the local console.

B.2 DESIGN CONSIDERATIONS

B.2.1 The Bus Interface

The controller board is designed to be compatible with the Multibus/IEEE-P796 with Compliance characteristics Master D8 M20 and Slave D8 M20 VO. That is, it operates either as a Slave or Master (when performing DMA), decodes/drives 20 address lines, uses only an 8-bit data path and is the source of non-vectored interrupts. Because an 8-bit data path is used, the board can be operated by either 8 or 16-bit processors. The schematic diagram is given in Figure B.2.1.




Figure B.1 Block Diagram of Communication Controller

B.2.2 Console Support

The Communicationn board includes circuitry to support the local consoles of the Linac Secondary stations. The console provides an alphanumeric video display, keyboard, shaft encoder knob and a selection of lighted pushbuttons. Appendix D describes this console.

B.2.3 Video Generator

The video generator is built around an LSI video display driver, the MC6847. Although designed for video games, the circuit has a 16 line by 32 character alphanumeric mode that is used for this application. All of its features are preprogrammed and the character generator, shift register, and address counters are all internal so that only a clock (3.5795 MHz) and 512 bytes of display RAM are needed for the chip to output the composite video signal. For simplicity, the processor accesses the display RAM during horizontal retrace time to avoid flicker on the screen.

B.2.4 Console Serial I/O

A Motorola MC6850 ACIA is used for the serial connection to the console. Because the console operates at a slow (4800) baud rate the 6850 is allowed to interrupt the processor to signal that data from the console is available. Opto-insolators are used in the transmit and receive lines to decouple this board from the remote console. Current for these are driven from the console chassis.

B.2.5 SDLC/DMA

Data received and transmitted on the 1 MHz link are transferred to and from memory under DMA control. This data rate corresponds to one byte per 8 us for one channel (Tx or Rx). Operation of the Primary station requires simultaneous operation of the transmit and receive channels, giving a data rate of 4 us per byte. To allow this data rate without placing severe restrictions on all Multibus cycle times, some on-board buffering was needed. A 16 byte FIFO is provided in each DMA channel and transfers by the DMA controller are made between these FIFO buffers and offboard memory. The DMA gains access to the bus by normal multi-master bus arbitration logic on the board. In this way the synchronous SDLC chip transfers of data to the FIFOs are decoupled from the bus cycles and the 1 MHz data rate can be supported without providing large RAM buffer space on

this communications board.

To allow normal communication with the SDLC chip itself, access is also provided from Multibus to the chip directly. Contention for access by the FIFO and Multibus is resolved by logic contained in PAL devices on the board.

Off board connection to the SDLC link data and clock signals are interfaced through a 26-pin card edge connector. The link repeater chassis is described separately in Appendix E

B.2.6 Byte-Serial DMA I/O

Because the DMA controller has four channels, a byte serial I/O path is included on the board. This capability is not required at each Linac station, but may be used for other inter-computer connections. Two 16-byte FIFO buffers are included; one for the input and one for the output. These buffers can be accessed either by the processor directly or by the DMA controller. Handshake signals are available, along with the data, on the 26-pin card edge connector.

B.3 DESIGN DETAILS

Each LSI chip comes with its own set of peculiarities and surprises - only some of which are detailed in the data sheets. The LSI components on the communication controller are from the Motorola 68XX family. Much of this choice was dependent upon prior experience with these chips and the fact that Motorola designs seem to have fewer surprises than designs from some other vendors.

B.3.1 Timing And Clocks

One outcome of using the Motorola parts is the necessity of interfacing synchronous chips with the asynchronous Multibus. An Enable clock must be provided for each chip and each access from Multibus must be synchronized to this E clock. On a read cycle the 68XX data is valid on the trailing edge of E. This data is latched so that it remains valid until the asynchronous deassertion of the MRDC command signal. For a Multibus write cycle the XACK signal is asserted by the communication board after the trailing edge of the E clock causing the processor to maintain valid data on the bus during one full period of E.

A common E clock is used for the 6844 (DMA Controller) and the 6850 ACIA (Asynchronous Communication Interface Adapter). It is derived from the MC6875 clock generator "Memory clock" output. When a DMA cycle is in progress, the bus arbitration circuit asserts the MR (Memory Ready) signal that causes the E pulse to be stretched until the XACK time, the completion of the data transfer. This stretch is essential for the 6844 and tolerable for the 6850. However, the E clock for the 6854 cannot be stretched. In normal operation the SDLC receive clock shifts link data into the 6854, but it is the E clock that transfers the newly assembled byte from the input shift register to the next location in the 6854's internal 3 byte FIFO buffer. This internal use of the E clock requires the E clock frequency to be 1.5 times the receive clock frequency and disallows stretching.

To accommodate these constraints, the "2xFo" output of the MC 6875 clock is used to generate the E clock for the SDLC. For use in the Linac control system an 8 MHz crystal on the 6875 generates a 2 MHz E clock for the 6844 and 6850, and a 4 MHz output from the 2xFo output. This 4 MHz signal is divided by 2 to obtain the 2 MHz E clock for the 6854. Jumpers on the board allow the use of a 4MHz crystal to generate a 1 MHz E clock for the DMA and ACIA and a 2 MHz clock for the 6854. The 2 MHz option is chosen to reduce the Multibus duty cycle during periods of DMA activity. Using a 2 MHz E clock, two channels of DMA servicing the simultaneous 1 MHz SDLC transmit and receive consume less than 50% of the Multibus time. The 2 MHz option does require the use of the faster "B" version of the 68XX parts.

B.3.2 Address Decoding

As viewed by a Multibus master, the communication controller appears as a 1 K block of memory. Half of this space is the 512 byte display RAM for the video generator; the other functions of the board occupy the second 512 bytes. The memory map is shown in Table 7.9.1.

The location of the 1 K block of memory within the 20-bit address field is determined by switch settings on the board. The next level of decoding makes use of a Programmable Array Logic device, a PAL 10L8. A diagram and the PAL equations are given in Fig. 3. In addition to simple address decoding, the PAL synchronizes the 6844 and 6850 chip-selects to the E clock. T68 is a signal generated when a 68XX cycle is to begin, and it is latched to form TCY68 at the trailing edge of E. This technique provides a PAL input, synchronized to E, that is one E cycle long. This TCY68 term is included in equations that provide chip selects to the 6850, 6844 and the 74245/74373 buffers and latches involved in 68XX cycles.

Outputs CSENDR and CSENWR are chip selects for a 74139 that selects the asynchronous registers on the board; TX and RX FIFO's, status registers and the FIFO clear register.

Multibus accesses to the 6854 are indicated by the ADDR54 output, but these accesses must be arbitrated with the autonomous transfers between the 6854 and the TX,RX FIFO's. A 12L6 PAL shown in Fig. 4. performs this arbitration.

B.3.3 I/O Connectors

These card edge I/O connectors are used for the serial port, the parallel port and the console I/O. Also a Lemo coaxial connector is available to drive a second TV monitor with the information displayed on the console.

Connector J1 carries the I/O for the SDLC link. In addition to the TX and RX clock and data lines, this connector carries the Flag Detect signal and Secondary power and ground to the link repeater chassis. of J1 is given in Table 1.

J2 is the interface to the byte serial section of the board.

J3 connects the console serial Transmit and receive, and the composite video signal to the console chassis.

B.3.4 Miscellaneous Registers

Three miscellaneous registers are available on the board. The bit definitions of these registers are given in Table 7.9.1. Writing a "1" to a bit in the FIFO clear register clears the associated FIFO. The value of the bit locations of the FIFO status register reflects the state of the individual input ready and output ready signals of the FIFOs. A "1" means ready.

In the miscellaneous status register, the MSB is a signal derived from the horizontal sync output of the 6847. This bit allows the software to store data in the Video RAM only during the horizontal retrace time.

Bit 6 of Misc. Status is driven by a 100 ms one-shot that is retriggered by flag detect pulses from the 6854. A "0" indicates no flags were detected for 100 ms.

Link Bypass is an input from the Link Repeater chassis that allows the processor to determine if the station is being bypassed. A "0" means bypassed.

B.4 DMA OPERATION

The 6844 DMA controller is operated in the four channel mode, transferring data between the various FIFOs and memory. During the DMA memory accesses, the upper address lines are asserted to reflect the setting of the 4-bit dip switch S3. This is necessary because the 6844 drives only 16 address lines. Because the upper four address lines are set by switches, all DMA transfers are constrained to lie within a single 64K block of address space.

For DMA cycles, shift-in/shift-out pulses for the selected FIFO are generated by IC61, a PAL10LB. Normal Multibus access is provided so the FIFOs can be written and read by the processor.

An SDLC controller chip must be told when the last byte of a message is transferred to it, so the 16 bit frame check sequence and terminating flag can be appended to the bit stream. In the 6854 this is accomplished by storing the last byte in the Frame Terminate register rather than the Frame Continue register used for all other bytes of the message. The last byte flag is stored in the FIFO by the DMA controller. This flag then progresses through the FIFO along with the 1st byte of data and when this byte is transferred, the last byte flag selects the frame terminate register by asserting RSD on the 6854. By this means, the transmitted frames are terminated without processor intervention. For sending short messages under processor control, a Multibus access to the address of a TX FIFO+1 will set the frame terminate flag with the stored byte of data. A short frame can then be sent without DMA control.

APPENDIX C
BINARY I/O CARD

The Binary I/O card is described in Controls Hardware
Release No.4.

APPENDIX D

LOCAL CONSOLE FOR LINAC CONTROL STATIONS

The local console is described in Controls Hardware
Release No. 3.




Figure E.2 Diagram of Fiber Optic Link Repeater Chassis

APPENDIX E

FIBER OPTIC LINK REPEATER CHASSIS

E. 1 INTRODUCTION

The Link repeater chassis associated with each Linac Secondary station allows each secondary to receive data from and transmit data to the Primary station. The repeater is built as a separate chassis so that an individual station can be powered down without interfering with the rest of the communication loop.

E. 2 DESCRIPTION OF THE REPEATER CHASSIS

The repeater contains a receiver section, a logic section, and a transmitter section. The circuit diagram is given in Figure E. 2.

E. 2.1 Receiver Section

Optical data is input to the Motorola MFOD404 photodetector module that is mounted in the chassis connector. This module contains the photodiode and an integral amplifier that outputs a differential signal to the MFOC600 receiver chip. Optical power input to the MFOD404 is converted to an electrical signal with a sensitivity of 30/mv/uw. This signal is converted to TTL by the MFOC600.

The incoming signal uses self clocking Manchester encoding. Clock and data information are separated and input to the logic section. A phase-locked loop (U7) serves to eliminate accumulated jitter in the detected clock signal.

E. 2. 2 The Logic Section

Most of the conventional logic for the repeater chassis is contained in U11, a 12L6 PAL (Programmable Logic Array). The PAL generates four output signals:

- BYPASOUT - Causes this station to be bypassed.
- TXCLK - Selects whether the received clock or the local clock is used to transmit.
- TXDATA - Selects either the incoming data or the output data from this station to transmit.
- RXENV - An envelope of the messages received on the link.

E. 2. 2. 1 BYPASOUT -

A Secondary station may be bypassed to allow the rest of the link to operate while an individual system is offline. Four conditions will cause a bypass; 1) Loop Online Control requests the repeater chassis to bypass. This LOC line is operated by the 68000 and the MC6854 in a way that the station can go online and offline without interfering with messages on the link. The MC6854 transitions the LOC bit between messages. 2) A manual Bypass switch will force a bypass but the bypass occurs at the instant the switch is thrown. An error will be generated if a message is in progress. 3,4) If the secondary station power (5v) is turned off or if the ribbon cable connecting the link chassis and communication card is disconnected, the link will cause the station to be bypassed. An error may be generated.

When bypass is in effect, the decoded receive clock and data are output to the encoder circuitry and the fiber optic transmitter.

E. 2. 2. 2 TXCLK -

TXCLK is the clock signal used by the transmitter section to encode the transmit data. Under both normal operating and bypass conditions TXCLK is the same as the RXCLK. Under some failure conditions, a local oscillator is switched in and used as the TXCLK. The local clock is enabled by the RTS output from the MC6854, a pin operated by the 68000. The local clock frequency is 1.0MHz.

E. 2. 2. 3 TXDATA -

TXDATA, the serial bit stream of data input to the transmit encoder, is usually the output data from the MC6854 SDLC chip. During BYPAS decoded receive data is sent directly to the transmitter by way of the PAL.

E. 2. 2. 4 RXENV (Receive Data Envelope) -

In order to observe the incoming message traffic receive data envelope detector circuit is included in the link repeater chassis. This circuit determines if the incoming data is different from SDLC flags (01111110) and if a flag was detected within the past 40 ms. The first condition indicates that data is present while the second condition requires that the link is working.

E. 2. 3 Transmitter Section

The PAL outputs TXCLK and TXDATA are encoded into a self clocking Manchester signal using an exclusive OR and a dual latch (431 and 434). The combined signal is input to a 75450 driver that powers the fiberoptic transmitter, a Motorola MFOE106 Infra-Red emitter. A 500 ohm variable series resistor adjusts the amount of drive in order to control the optical output power. This power is normally set to about 10 db mw of optical power arriving at the next receiving station.

E. 3 FRONT PANEL

The front panel of the fiber optic repeater chassis shown in Fig. 3.4.2 has a combination of LEDs and test points to aid in determining the status and operation of the link.

Each of the four Bypass conditions described in 2.2.1 is given a small yellow LED and the Bypass Status is shown by a large red LED.

Along the left side of the front panel is a group of LEDs and test points. Those are described below. All LEDs are tied to +5V and pulled low in the active state.

TX CLK ENABLE A yellow LED that is on when the local 1 MHz crystal oscillator is enabled by the processor. It is activated by the M6854 RTS pin. The test point is the low active side of the LED.

FLAG DET A green LED driven by a 10 us one shot triggered by the MC6854 Flag Detect output bit. For continuous flags, the LED is on DC. The test point is the buffered Flag Detect pulse from the MC6854.

RX CLK The decoded receive clock output by the phase locked loop.

RX DATA Receive Data Envelope. Test point is the LED signal.

TX CLK Transmit clock used by encoding circuit. The test point is the LED signal.

TX DATA Output of a one-shot triggered by the transmitted data. The one-shot causes the LED to extinguish if constant high or low data are output by the PAL. The test point is directly the buffered transmit data output from the PAL.

The only front panel control is the locking toggle switch for manual bypass. Operating this switch will force the link into its bypass mode asynchronously. An error may be caused by the action. The bypass switch is debounced by logic on the main circuit board.

E. 4 REAR PANEL

The rear panel contains connectors for the transmit and receive fiber optic cable and a 26 pin ribbon cable that connects the link chassis to the communication board of the local station.

SMA type fiber optic connectors are used (Amphonal 905 series). Cable ends are assembled by Belden using 200 micron glass fiber optic cable (Beldon 220001).

Only one control is located on the rear panel - the 500 ohm variable resistor used to adjust the fiber optic output power level. A test point is available to examine the output of the MFOD404 DETECTOR.

E. 5 HP FIBER OPTIC RECEIVERS

Later models of the Link Repeater chassis may use HP - fiber optic receivers. These devices output a TTL compatible signal that is input directly to the Manchester decoder circuitry, so the MFOD 600 chip is not used. In this case the test point is the output of the HP device.

APPENDIX F

16-CHANNEL S/H-AD CHASSIS

F.1 INTRODUCTION

The majority of Linac analog signals are pulsed. Examples include beam current, RF parameters and drift tube quadrupole currents. Because of the pulsed nature of these signals, they must be sample and held before being digitized for input to the computer.


A small modular 16 channel chassis is used to house both that sample-and-hold circuit and the digitizer for a group of analog signals. Each station contains several of these chassis. The circuit diagram for this chassis is shown in Fig. F.1.

F.1.1 Description Of The Chassis

The SH/AD chassis receives analog input signals on two 8-pin coaxial connectors. Each channel is routed to a Harris 2425 sample-and-hold amplifier and the outputs of these amplifiers are input to a 16 channel multiplexed analog-to-digital converter (Analogic model 6812). The digitizer output is buffered onto a 50 conductor ribbon cable bus by way of a card edge connector at the rear of the chassis.

The buffers are tri-state devices so that all chassis may be daisy-chain connected. Channel and chassis address data, a digitize strobe and output data are all carried on this bus. The entire group of chassis for a local station are driven by a three byte I/O connector included as part of a 9-byte binary I/O board in the local station.

In the three byte cable, two bytes are input and carry the twos complement left justified data from the digitizer. The third byte is output data to the chassis and the bit usage is defined as shown below:



The digitize strobe is a 2 us high active strobe generated by the computer under program control. The chassis and channel number define a unique analog input channel and cause data from that channel to be output to the daisy-chained bus. When a digitize strobe occurs, the channel selected by the chassis/channel number will be digitized.

All digitizer chassis are set up to read +10V F.S. signals except chassis AD-0 is jumpered for +2.5V F.S. in order to obtain two more bits of resolution for RF power readings. These signals seldom exceed one volt.

F.2 JUMPER OPTIONS

Jumper connections for both the 10V and the 2.5V full scale ranges are provided. The jumper connections for the +10V version are indicated on the schematic drawing by dashed lines.

As installed, all of the Harris 2425 sample and hold circuits are operated as single-ended gain of one followers. The hold pulse is generated by a 40 ms one-shot triggered by sample and hold trigger input on the J3 BNC input connector on the rear panel.

Jumper J1 is removed so that the A.D digitizes the selected (1 of 16) channel each time a digitize strobe is generated. That is, the chassis number is not examined to determine if the selected channel number is located in a given chassis.

When operated in this mode, the computer can issue a single digitize strobe and then collect the result from each chassis in turn. This speeds up the reading process because the computer does not need to wait for the conversion time for each channel.

F.3 FRONT PANEL

There are no operator controls on front panel. Indicator LEDs are provided to show End-of-Convert, chassis select and +5, +15 power, convert strobe and the internal Sample-and-Hold output of 6812 digitizer module. The strobe test point will show the time at which the computer is gathering data and the S-H OUT test points shows the analog value of the selected channel during the digitize time. This test point is useful to determine if the digitizer is functioning properly.

F.4 REAR PANEL

The rear panel contains interface connectors J1 and J2 for the analog inputs, J3 for the S-H Trigger and the daisy-chained ribbon connector. The three-bit chassis number is set by way of a DIP switch accessible through the rear panel. Chassis numbers are normally chosen to begin at zero and count up for as many chassis as are installed.

APPENDIX G

16 CHANNEL DA CHASSIS

G.1 INTRODUCTION

A 16-channel D-A chassis was developed for supplying the reference voltage for the Linac drift tube quadrupoles. An external chassis is used for several reasons. Commercially available Multibus D-A boards are available with only eight output channels. Some Linac stations would require five such boards consuming nearly half of the Multibus bin space. External, independently powered chassis allow the Multibus bin to be powered down without interrupting the reference voltages. Replacement of a faulty D-A can be accomplished without powering the computer off.

The circuit diagram is given in Figure G.1.

G.2 GENERAL DESCRIPTION

The 16-channel D-A is housed in a three and one-half inch chassis. Interface to the computer is by way of a 50-conductor ribbon cable that carries three bytes of output data. Two bytes are the D-A setting and the third byte carries the unique D-A address and a software generated store strobe. The chassis address is set by a dip switch accessible through the rear panel. Data for the D-A value is sent from the computer as a twos complement number that is left justified in the two bytes of data. The D-A circuit provides a positive only output, so the sign bit is unused. Because each chassis compares its address to the address on the control byte, the entire 50-conductor ribbon cable can daisy-chain to as many as seven chassis.

Individual DA circuits are National DAC1006 monolithic D-A converters followed by an op-amp used as a current to voltage converter and an output buffer. A zero and full scale potentiometer is included on each channel, but these are not accessible.

G.3 FRONT PANEL OF THE D-A CHASSIS

The front panel contains only LED indicators and a single test point connected to the strobe signal. Yellow LED indicators are provided to monitor the +5V and the +15V power supply outputs. A green LED is driven by the strobe pulse when the chassis is selected. An additional red LED is connected to the chassis select comparator. In normal operation the chassis select time is very short so the LED is not visible. The computer selects chassis number seven after a DA setting is made in order to deselect the chassis. This prevents noise spikes from storing erroneous data.

G.4 REAR PANEL

The rear panel contains two 8-pin Burndy coaxial connectors that carry the 16 DA output signal, a cutout for the 50-conductor ribbon cable, one access hole for setting the chassis number and the AC power connection.

APPENDIX H

MULTIBUS 9-CHANNEL PREDET BOARD

H.1 INTRODUCTION

A nine-channel timer board has been designed to provide local variable timing for Linac secondary stations. Timer channels are integrated into the Secondary's software to make them appear as a controllable analog parameter. The timers are triggered by TZERO and use the Booster clock as a time base.

H.2 DESCRIPTION OF TIMER BOARD

The Linac timer board is designed as a Multibus slave board that is in the I/O space of the Linac Secondary. Each of three 3-channel timers use 8 bytes of I/O address; the board occupies 32 bytes of I/O space. Register assignments are given in the Motorola data sheets offset by the addresses listed Section 7.9.2.

This board was designed specifically as a timing generator rather than making it a general purpose counter board. Default traces on the printed circuit board are arranged so that each timer is independent, all timers start by an external trigger and the time base is a 2 MHz signal phase locked to the Booster clock.

The trigger and clock inputs and the nine output pulses are available on a 26 pin card edge connector on the user side of the board. Because the timer circuitry does not fill the board, a large prototype area is included to allow custom curcuitry to be added. A 50-pin card edge connector is available to access this space.

H.3 WATCHDOG TIMER

A somewhat unrelated feature of the board is the system watchdog timer. This timer consists of a 220 ms retriggerable one shot that triggers a 100 us one-shot that provides a system reset on P2 pin 38. After system initialization the 68000 performs a write to address \$6118 to trigger the 220 ms one-shot. This is done each Linac cycle by a task that is triggered during the 15 Hz interrupt. The timer is retriggered each cycle but for any reason if the processor does not trigger the one-shot for three cycles, the watchdog timer will time out and trigger the second one-shot that generates a system reset. One and only one reset pulse is generated, but under normal circumstances this reset will restart the system. This feature is enabled by a toggle switch on the front edge of the board. A cut trace option will disconnect the signal from P2.

H.4 PREDET TIMERS

Three Motorola MC6840 chips provide the counters used in the predet channels. Delay times are loaded into these chips during system initialization. The timers are configured to count the E clock and a trigger input normally starts all timers simultaneously. Once a timer is triggered it counts for a time determined by the number in the preload register, and then the output transitions from high to low. This trailing edge is used to trigger a one shot that outputs a 1 us pulse which is buffered with a 75452 and transformer coupled off the board. A power on inhibit circuit disallows any output pulses during a power-up sequence.

Although the board is designed to make each timer independent, cut trace options and jumpers points are provided to allow one timer to be triggered by the timeout of another timer. These jumpers are grouped near each 6840 and are organized to allow one timer output to trigger either of the other timers within the chip. This is useful to generate pulse-on/pulse-width functions.

H.5 TIME BASE INPUT

The timer board is designed to expect a 1 MHz input time signal that is normally the Booster clock. This signal is input to a NE564 phase locked loop with a divide by two in the feedback loop to generate a 2 MHz clock at the phase locked loop output. If the input 1 MHz signal is removed, the phase locked loop continues to operate, but at a slightly different frequency. Only a few tens of millivolts

of 1 MHz signal is required for the lock condition. Typically 200 mv of signal is provided to each Secondary from a common line that carries the Booster clock down the length of the Linac.

H.6 TRIGGER INPUT

The input trigger signal is expected to be a positive 0.5 us pulse. It is transformer coupled and terminated in 50 ohms.

H.7 INITIALIZE

The INIT signal on the Multibus backplane is buffered and connected to the reset pin of the MC6840 chips. All timers are disabled by this pulse and the internal preload registers will be set to FFFF. The INIT signal resets the watchdog timer so that no reset pulse will be generated.

H.8 INTERRUPTS

The IRQ interrupt outputs of the three MC6840 chips are each buffered and available to connect to the INTx lines on the Multibus backplane. No interrupts are connected by traces on the board.