Oxford Gray Institute Linac Control Software

1. Introduction

The Oxford Gray Institute Linac is controlled entirely by a host PC through USB connections. There are two USB links, one to the main body of the linac inside the Faraday cage and the other to a beam-switching assembly outside the cage and immediately adjacent to the experimental area.

Both devices are operated by a single control program which runs on the host outside the radiation controlled area. While the link to the beam-switching assembly is straightforward in the sense that conventional copper cabling is used, communications with systems within the Faraday cage (i.e. with the linac control system) use a fibre transmission to preserve the electrical isolation of equipment within the cage. A fibre-optic USB 2.0 extender (Ranger 442, http://www.icron.com, Amplicon part# 96065273, http://www.amplicon.com/) is used in conjunction with 62.5/125 µm multimode fibres terminated in duplex LC connectors.

The software itself is written in ‘C’ using the National Instruments LabWindows CVI software environment.

2. Software tasks

The linac software operates in conjunction with the linac control system hardware and allows the user to change the state of the accelerator system, and to read back the state. The software has the following functionality or attributes:

- A user interface for providing command input and for displaying the current setting or status of the system
- A calibration facility for matching set commands to desired system state and correct display of ‘read-back’ system state
- An event log and status log to provide time-stamped information on past events and system states
- Password security to restrict certain more advanced levels of control such as calibration, electrical interlock overrides and setting of certain critical beam delivery parameters
- Actions to be taken in response to certain events, e.g. the tripping of accelerator vacuum which turns off the gun filament to prevent possible damage
- A control settings file management system which enables fast/convenient machine set-up
- Communication with other programs, via Named Pipes, to enable coordinated irradiation routines
- Specific controlled routines e.g. procedures for energising switching magnets to minimise inrush currents etc.

The fundamental aspect of the linac program is a thread comprising of a ‘while’ loop, which reads the linac state and takes appropriate actions. This is critical for providing real-time information on the system. The loop time is about 40 ms and has the following structure:

- Read raw data from the linac representing logic states and parameter values
- Retrieve, then apply, calibration factors and transforms as required
- Where relevant, apply smoothing algorithms
- Display the information on the graphical user interface panels
- Evaluate significant events and fault conditions
- Act on significant events and fault conditions
- Update the status panel and event log

3. Main user interface panel

The main panel presented to the user at start-up is shown in Figure 1. This panel provides the operator with critical information of the accelerator state and allows access to more detailed operational panels.

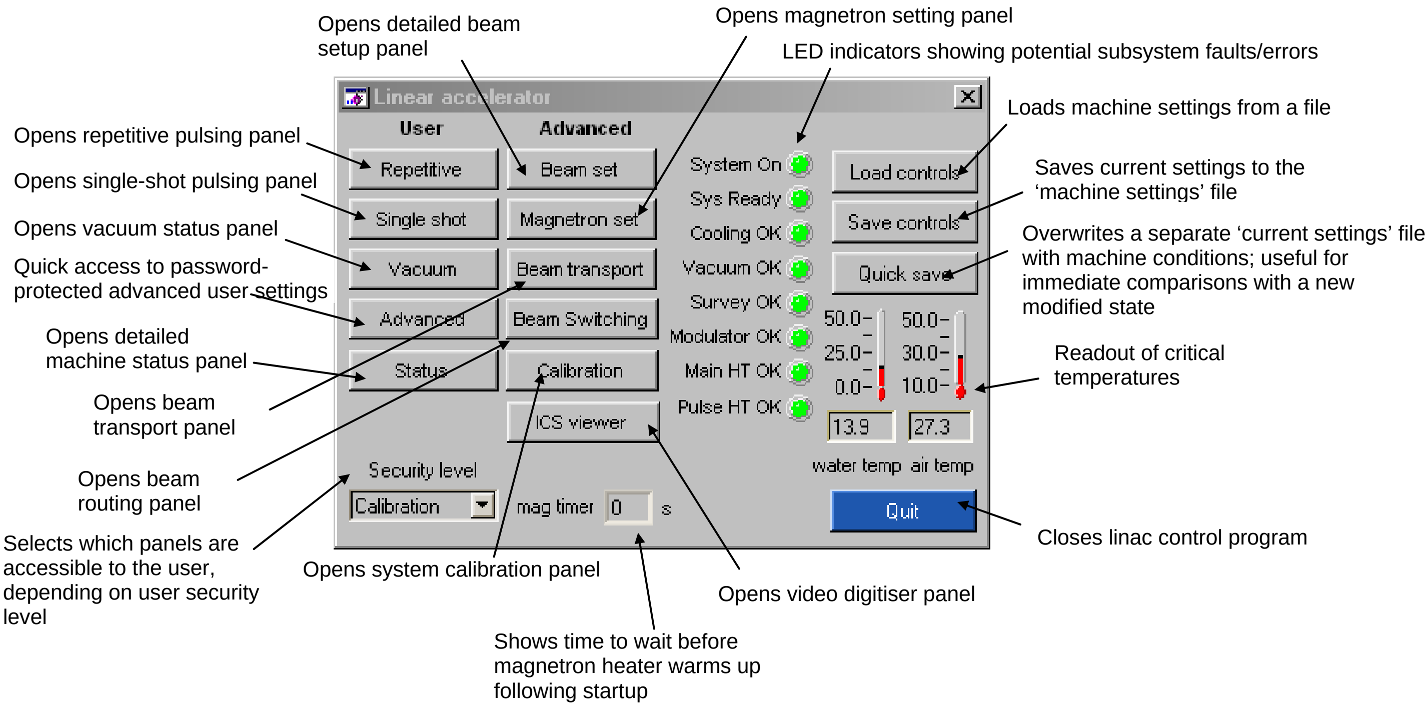

Figure 1: The main user panel shown at linac program start-up

All LED-type indicators reflect real time states, where the relevant signal is read from the accelerator's control system inside the previously mentioned loop (40 ms period) and redisplayed. This is in contrast to a system where the indicator is only updated on a trigger event and which thus might not represent the true current situation.

As well as handling analogue data, the display system also handles logic data (i.e. commands or states which are defined by a single bit, e.g. an on/off switch and a true/false state.

4. Security

Three levels of security are defined. A password is required to access the main panel at which point the operator has access to basic user settings. This gives the operator the ability to initiate irradiations, control gun filament and change pulse widths, and is intended for the experimenter.

Higher levels of security can be requested via a drop down menu on the main panel, which pops up a password entry panel, shown in Figure 2. Advanced user settings give access to more machine parameters and this is intended for the physicist setting up beam conditions. Finally, calibration level security gives access to the calibration panels. Changes on these critically affect the integrity of the linac so warrant the higher security level.

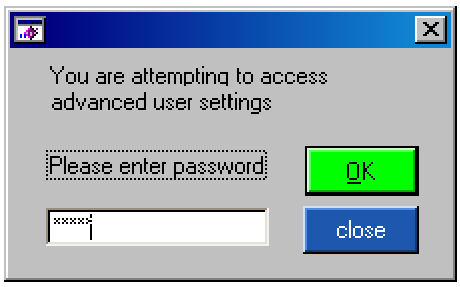

Figure 2: Password security panel

5. Single shot and repetitive pulsing user panels

User-accessible panels to control delivery of single-shot and repetitive pulses are shown in Figure 3. The operation of the controls is self-explanatory.

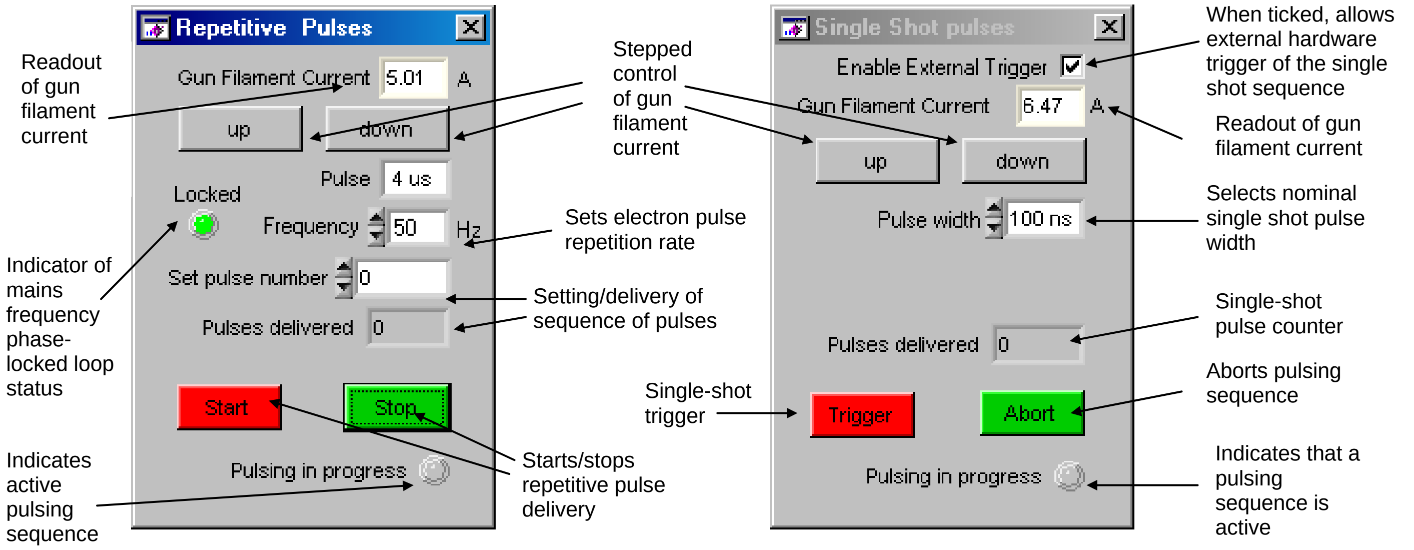

Figure 3: Single shot user panel (right) and repetitive pulsing user panel (left).

6. Beam setup panel

This panel controls all functions associated with the generation of the beam and controls the two machine modulators and is shown in Figure 4.

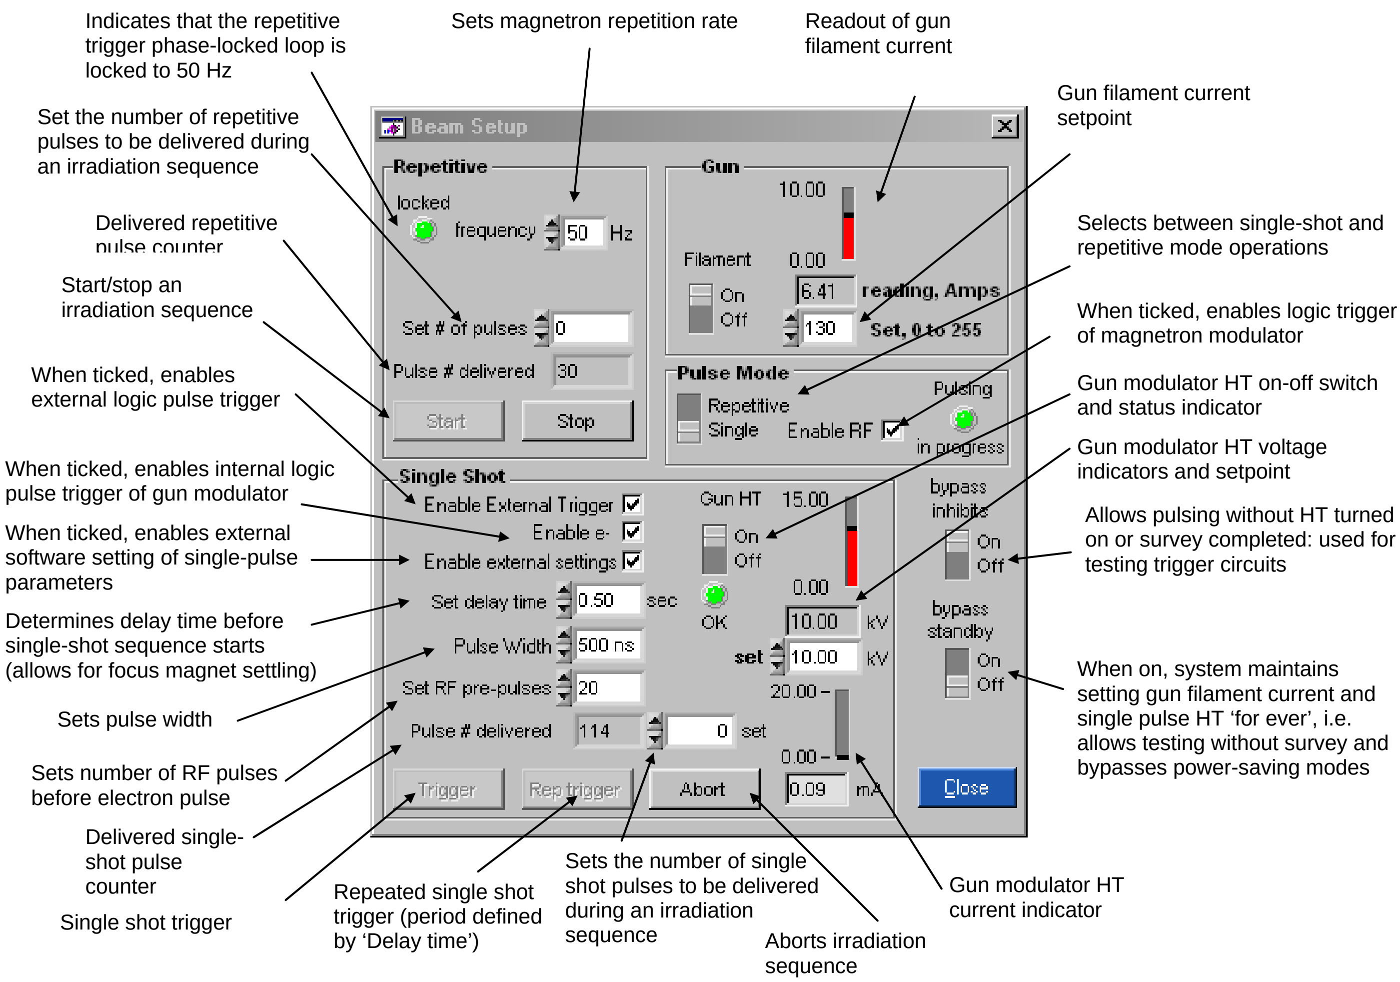

Figure 4: The beam setup panel

The controls associated with the single-pulse modulator (Figure 5) are placed in the bottom left of the panel. This modulator is brought into the gun drive circuit (injection) with the aid of a high voltage relay (energised through the single-repetitive slide switch in Figure 4).

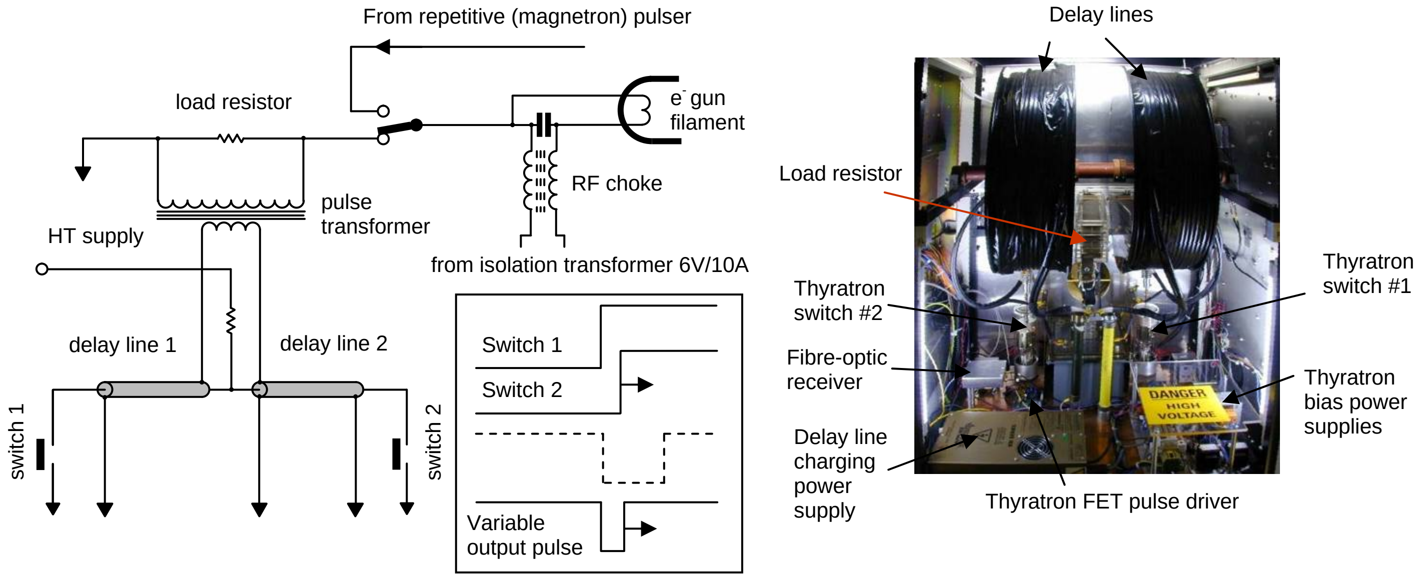

Figure 5: The single-pulse modulator electrical arrangement (left) and practical implementation (right).

In order to prevent possible arcing due to the high pulse voltages present, this relay (Figure 6) is only energised/de-energised when the modulators are not pulsing and the relevant bias voltages are off.

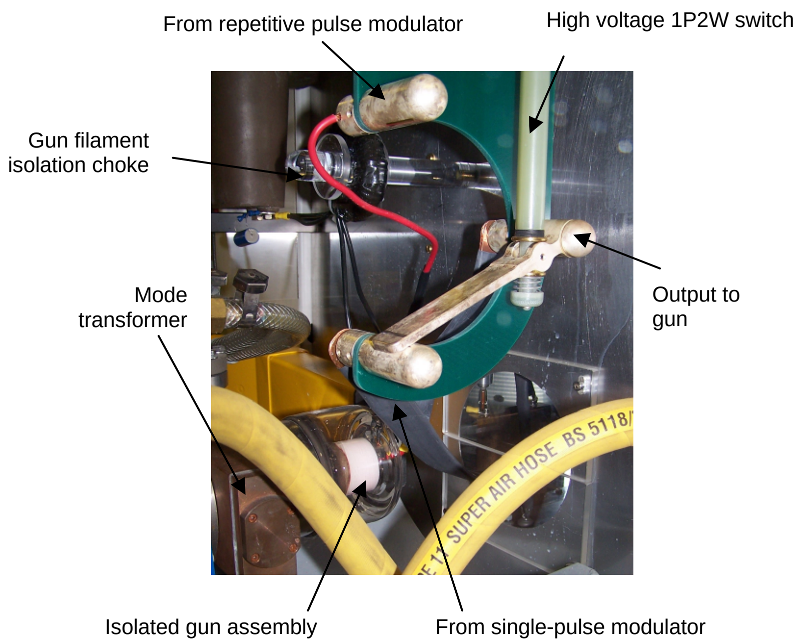

Figure 6: The high voltage changeover relay which selects whether single-shot pulses or repetitive pulses drop the gun potential (to -45 kV) for injecting electrons into the accelerating structure.

The gun filament can float at the pulse injection voltage and is heated by passing current through it provided by a high voltage isolation transformer, delivering up to 10 A at a voltage of 6 V. The current through the filament is not regulated but the isolation transformer ac primary voltage is derived from a motor-driven variac. The primary ac current is proportional to the gun filament current and this is measured with the aid of a current-sensing resistor in the variac output circuit. The voltage developed across this resistor (typ. ~1 V rms) is stepped up to ~30 V rms with a transformer, full-wave rectified and attenuated to provide a dc output voltage approximately proportional to filament current. It is this signal which is displayed in the ‘Gun’ section of the panel shown in Figure 4. The signal is in fact non-linear at very low filament current settings (due to rectifier diode non-linearity below ~1V) but such currents are not useful anyway. Since there is not a linear relationship between variac position and heater current, the filament current is varied by controlling the set-point of the motor servo which sets the variac position, with an 8 bit resolution: the set-point range is thus 0-255.

Returning to the single-pulse modulator controls (bottom left of panel in Figure 4), the parameter that dictates the single pulse injection is the modulator charging power supply, and this may require adjustment in some instances. This is a suitable point to explain the operation of analogue set-points and monitoring. Setting a control typically involves sending values between 0 and 255 to a DAC (digital to analogue converter) in the linac control system, such that the output range of the DAC is matched to the range of the system it is controlling. The operation of the ‘Gun HT’ control is exemplified here. The gun modulator high voltage charging supply covers the range 0 to 15 kV, and so the panel control has (at least) these limits too. The power supply accepts a control signal of 0-10V to generate 0-15 kV. If 10 kV is entered into the edit box the software calculates that this is 2/3 of the full range and sends 2/3 x 255 = 170 to the DAC. The DAC produces a corresponding output voltage according to its output range (2/3 of 10 V) which then acts as the control voltage signal to the supply.

In reality there are a few more steps involved: there are actually two DACs providing coarse and fine voltage contributions; these are followed by an analogue opto-isolator, the output of which is offset and amplified so as to match the relevant control voltage range. Similarly, the signal read back from the ADC (analogue to digital converter) must be corrected before displayed as a real number on the screen, since the monitor output of devices is usually scaled in some specific manner (e..g. the aforementioned gun HT supply provides two monitoring outputs of 0-10 V which represent 0-15 kV output voltage and 0-20 mA output load current).

The pulse width is determined by the time delay between triggering of the switches shown in Figure 5. This time is determined by a delay circuit and a look-up table is used to define the width.

The single pulse modulator delivers its pulse during the RF pulse envelope. However, since magnetrons need time to stabilise (and because the AFC feedback arrangement described later does not have a fast settling time), it is usually necessary to fire the magnetron several times before ‘stable’ operation is obtained. The number of these ‘pre-pulses’ can be set with one of the controls on this panel. In addition, the single pulse trigger is just that: triggering can occur at any time, on user demand. This demand is generally determined by some form of experiment trigger and the experiment is likely to require some form of pre-trigger delay. In addition it is beneficial to maintain the accelerator focus magnets in a standby state, where their current consumption is reduced, and hence the heat output is similarly reduced. In order to ‘wake up’ from this standby state, a fixed delay is introduced. This is delay is set on the ‘set delay’ control, or it may be adjusted by external software, as will be described later.

When the machine is operating in a repetitive mode (or when providing pre-pulses in the single-shot mode) the modulators are locked to the power line frequency (50 Hz) through a phase-locked loop which generates 25, 50, 100, 200 and 300 Hz. This synchronistion signal (20 ms period) is used to define the pre-trigger delay resolution.

7. Magnetron modulator control panel

The magnetron produces high power radiofrequency (3 GHz) pulses of ~4 μs width. Although the magnetron is isolated from the reverse load power with a ferrite isolator, as shown in Figure 7, optimum efficiency (and hence electron beam energy) is obtained when the magnetron is tuned for

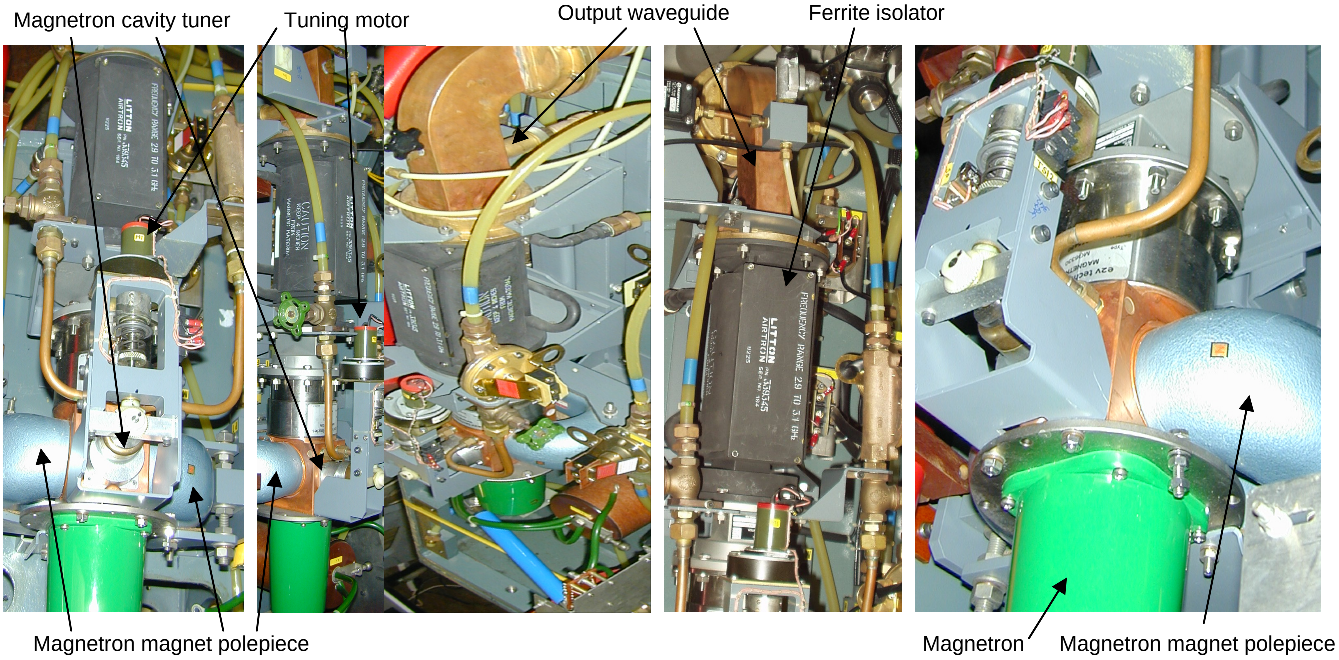

Figure 7: Images of the radiofrequency generation system: the magnetron, its tuning system, the isolator and the output waveguide.

optimal power transfer to the waveguide/accelerating cavity system. This tuning is achieved with the aid of a small DC motor which drives a tuning stub through a high ratio gearbox.

The motor can either be driven manually or by an error signal derived from a pair of detectors (pill-type diodes) loop coupled to two cavities in the output circuit. These cavities are tuned above and below the resonant frequency and the detector outputs are summed in a way to produce a dc signal proportional to frequency difference between actual and resonance frequency, in a manner similar to a Foster Seeley detector. This error signal can then be used as part of an automatic frequency control loop (AFC loop), driving the motor in the appropriate direction so as to minimise the error.

The magnetron system is controlled with an independent software panel, shown in Figure 8. The upper part of the panel deals with the magnetron frequency control and most of the operation of most

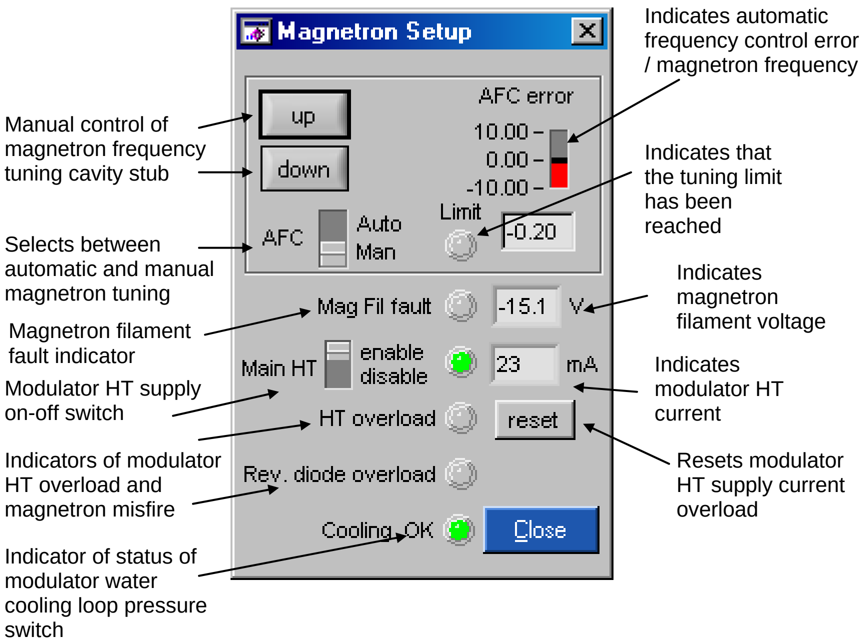

Figure 8: The magnetron modulator control panel.

of the controls is intuitive. An indicator is provided to inform the user that a mechanical limit has been reached by the frequency control motor. In common with most frequency deviation detectors, at large frequency errors, this system's output falls to zero and such a situation would inevitably drive the motor well away from 'correct' tuning.

The lower half of the panel deals with the magnetron and its modulator. An indication of the magnetron filament voltage is provided (this voltage reduces as modulator HT current increases, since the electron cloud in the magnetron can be enough to provide filament heating without additional external power). Provision is made to turn off the main modulator HT and to indicate and recover from the most common modulator faults.

8. Beam transport panel

This panel controls all functions associated with the transport of the beam within the linac body, shown in Figure 9. The focus power supplies are shown in Figure 10. These supplies use old but reliable technology: the output of a low voltage transformer is rectified with a diode/thyristor bridge.

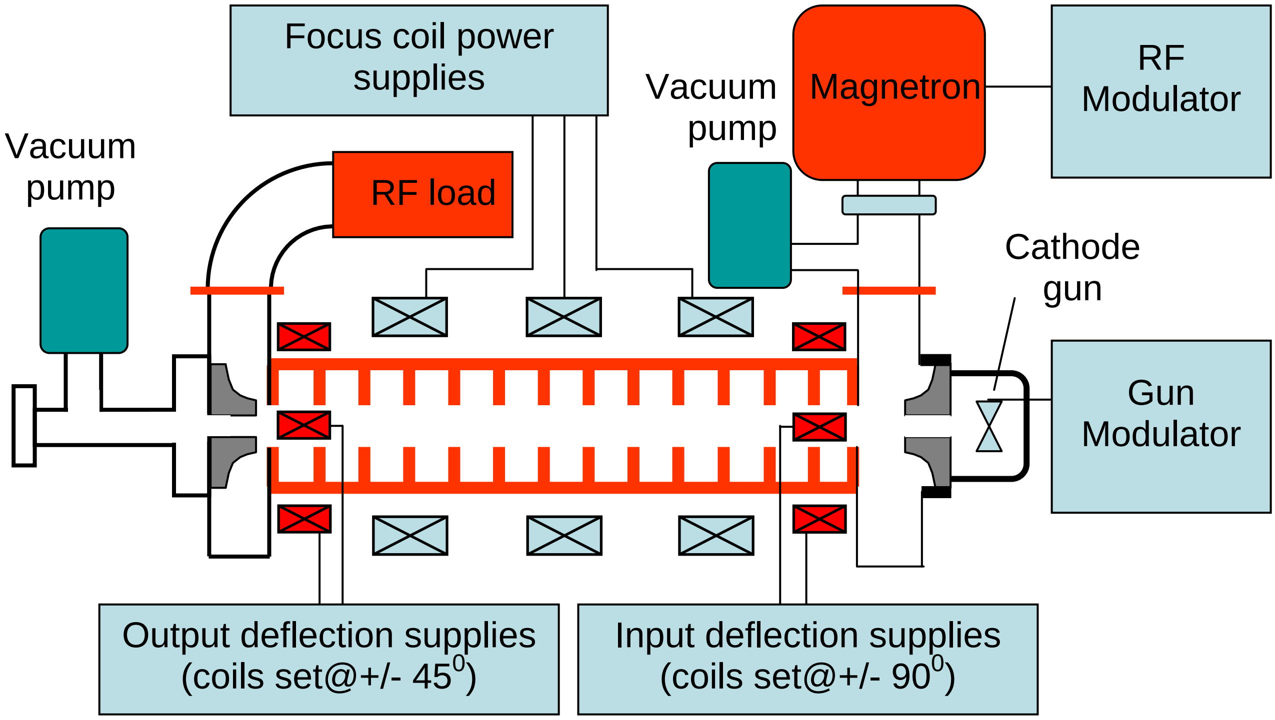

Figure 9: The beam transport arrangement in the linac body.

A high inductance choke in series with the focus magnet smooths the output current. The output current is monitored with a low value current sense resistor and this output is compared with the required set-point to generate an error signal which alters the thyristor drive delay. Three such units are used, driven from multiple digital-to-analogue converters.

The deflection supplies are simple linear units based around an operational amplifier buffered with a

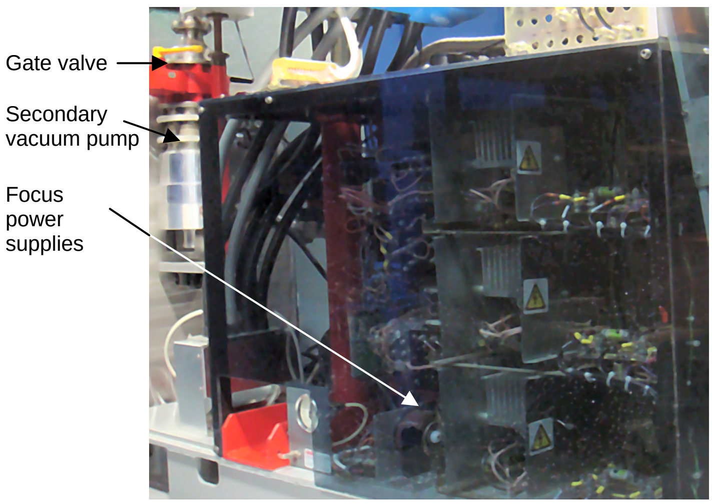

Figure 10: The machine focus power supplies.

power driver, with the output current sensed by a 1 Ω resistor. This time the setpoint input is a bipolar voltage. The software panel is shown in Figure 11 and its operation is self-explanatory.

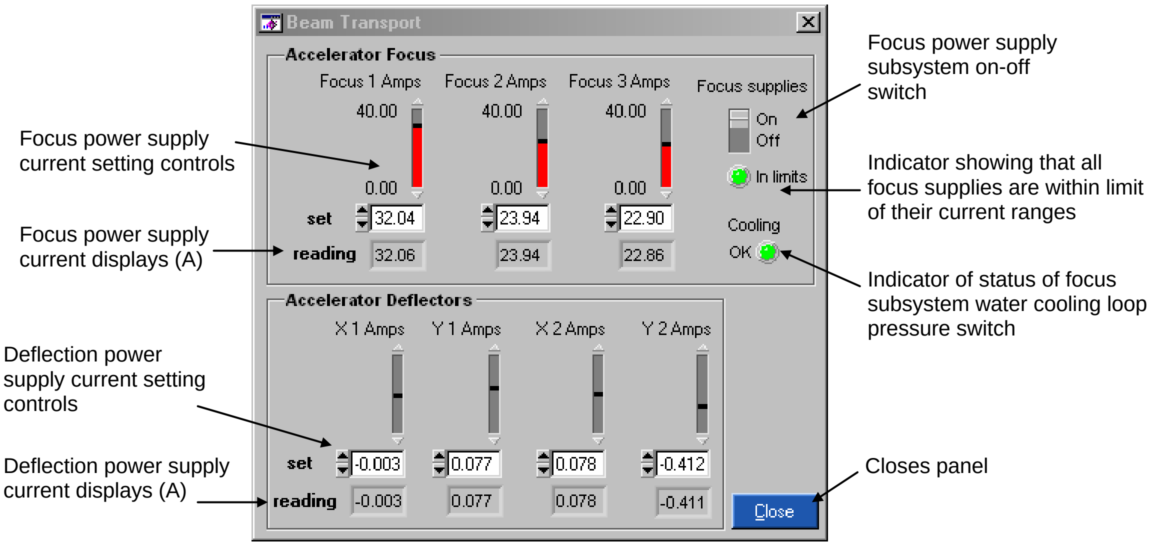

Figure 11: The beam transport panel controlling linac body focus magnet power supplies and two deflector systems, one at the start of the accelerating structure, the other at its end.

9. Vacuum system

The machine's vacuum system is shown in Figure 12 and the relevant software monitor panel is shown in Figure 13. We use two turbomolecular pumps, sharing a single rotary backing pump. One turbo pumps the gun end of the machine, the other the output beam line. The pressures in these two regions are read out by Penning gauges and their readings are displayed on the software panel in Figure 13.

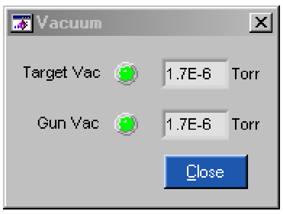

Figure 13: The linac vacuum monitor panel.

Conversion of Penning current to Torr is performed using a simple calibration look-up table.

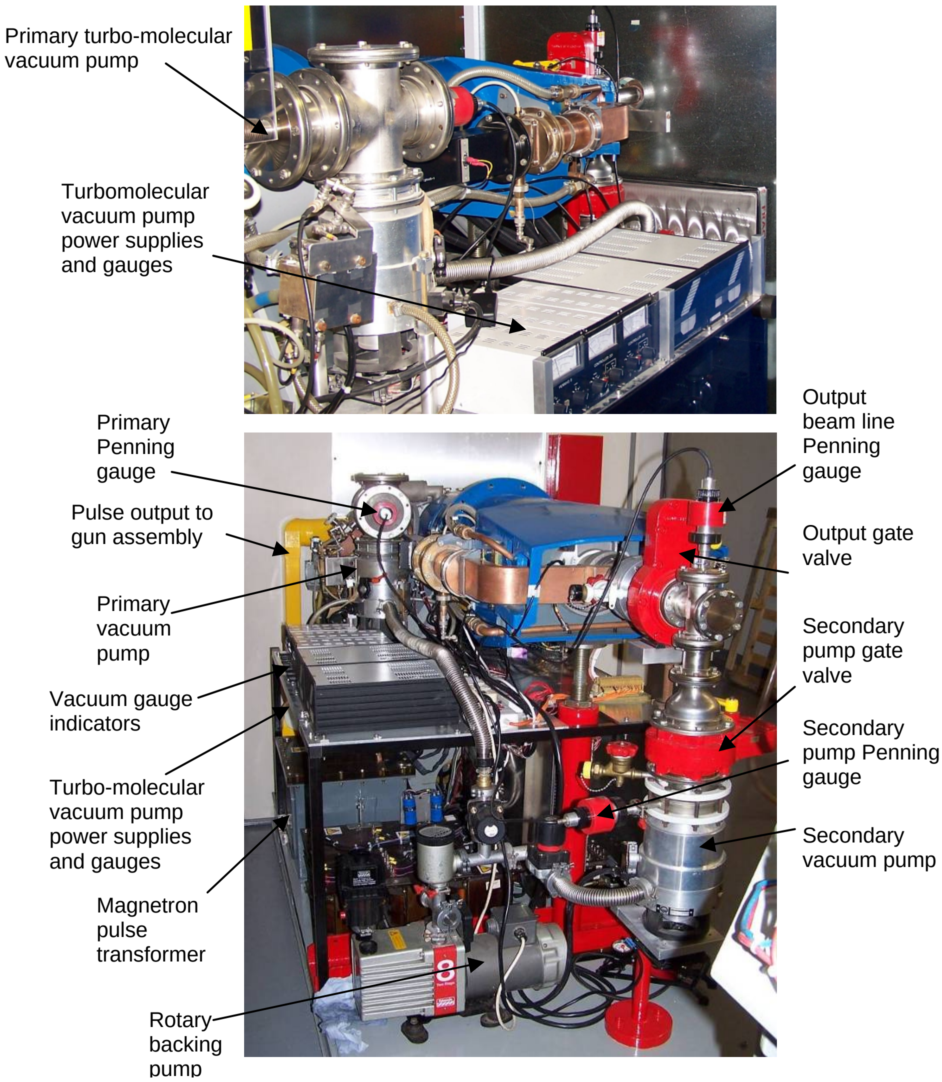

Figure 12: Views of the accelerator's vacuum system.

10. Calibration panels

In order to achieve accuracy in both setting a machine parameter and reading back its value a calibration facility is required. The systems are usually linear, so the corrections to be applied to the software numbers to be sent are a scale factor to control the range of the response and an offset to control the zero point. The same two-point calibration applies to the readout values. The software is designed so all changes are immediate: if a calibration factor is changed then the new setting is sent to the linac, and the readout routine is constantly looking up the calibration factors (as opposed to loading the values once at the start of the program).

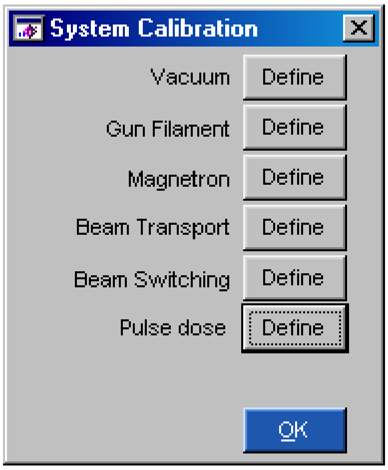

Figure 14: The master system calibration selection panel.

In some cases, the calibration is not linear and either an analytic expression is evaluated or a look-up table is used. Access to machine calibration is only available to high-level users and is reached by selecting the main access calibration access panel, shown in Figure 14. From this panel, the following calibration panels can be accessed:

(1) Vacuum system calibration

This panel, shown in Figure 15, defines a look-up table used to linearise the outputs of the Penning gauges and allows setting of two thresholds: a VacLow threshold, below which no high voltages are allowed to be turned on (risk of sparking) nor can the gun filament current be energised and a Vac Med threshold. This ‘medium’ vacuum region allows energisation of the gun filament and the high voltage supplies but does not permit triggering of the accelerator in either repetitive nor single-shot modes.

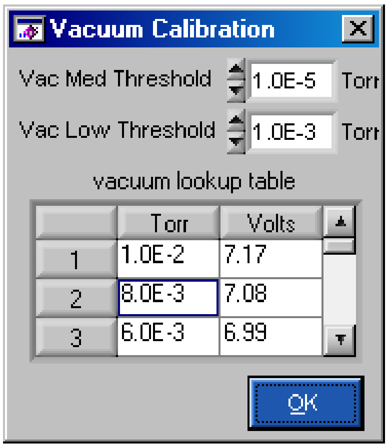

Figure 15: The vacuum system calibration panel.

(2) Gun Filament

Here we use a 3rd order polynomial to linearise the setting and readout of the gun filament power supply (see Figure 5 and associated description). Although this method does provide highly accurate readings at low filament currents, it has been optimised to provide accurate readings over the operating range of filament currents (5-9 A). This arrangement takes into account loading of the filament transformers.

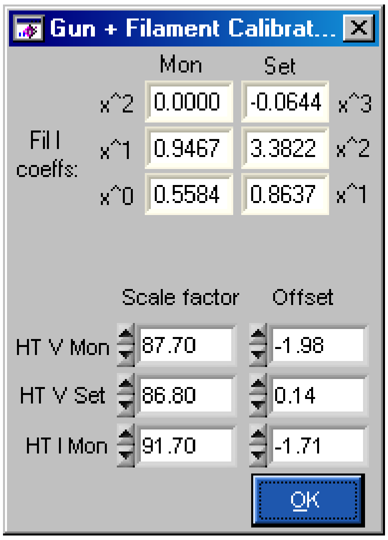

Figure 16: The gun calibration panel.

The single-shot modulator power supply can also be calibrated from this panel, shown in Figure 16. The particular power supply used is a constant voltage power supply, i.e. its output current cannot be set. However, both the output voltage and the output current can be monitored. The derivation of the factors used here is from a range of standard high current and high voltage laboratory meters and their use is only allowed by qualified personnel as floating operation of meters and high voltage probes is required.

(3) Magnetron

The panel shown in Figure 17 is used to calibrate the monitor of the magnetron high voltage supply's output current and the readout of the magnetron filament current. There are no user-defined set-points associated with these supplies. A linear two-point calibration is used here.

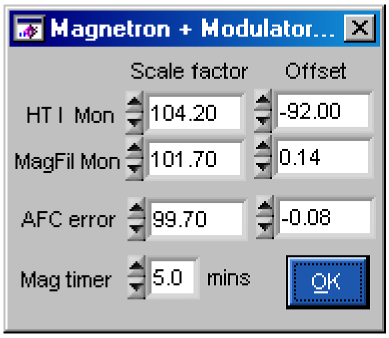

Figure 17: The magnetron modulator calibration panel.

No high voltage must be applied to the magnetron until its filament is fully warmed up and the main modulator thyratron has similarly been warmed up (i.e. that the thyratron hydrogen source has been adequately warmed. The time delay associated with this can be defined on this panel, typically 5 minutes. This delay is backed up with a hardware timer; however, since this is not read out, the software delay is matched to that hardware timer, appropriately disabling a range of controls on the software GUI.

Finally the readout of magnetron tuning (see Figures 7, 8 and associated descriptions) can be calibrated here, using a two-point linear calibration.

(4) Beam transport

The panel shown in Figure 18 is used to calibrate the setpoints and the readouts of the beam transport magnets internal to the machine (as used by Figure 11). Two-point linear calibration is used for the three focus magnets and the two pairs of beam deflectors. The power supplies for the three focus magnets include hardware comparators to indicate whether the magnet currents have been set outside the linear ranges of the power supplies. These comparators are adjusted locally on their respective power supplies and cannot be programmed.

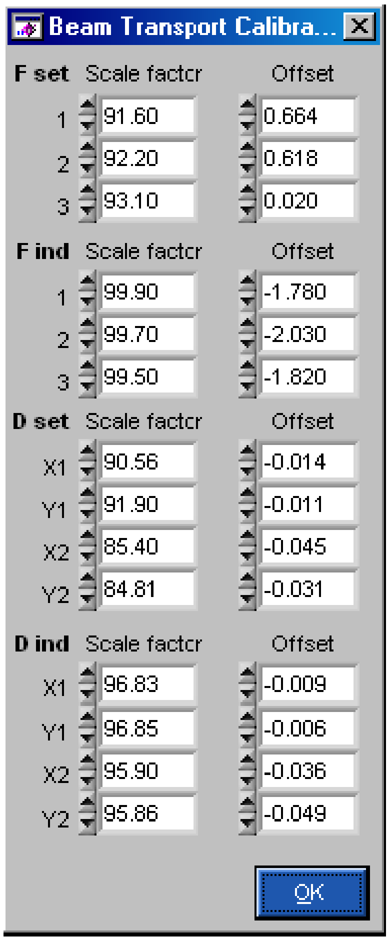

Figure 18: The beam transport calibration panel.

(5) Pulse modulator

The panel shown in Figure 19 is associated with the single shot and the repetitive pulse modulators.

The 'phase' control defines the timing point of the production of the electron pulse with respect to the 50 Hz power line frequency. This allows operation synchronised to the mains cycle, minimising the consequences of any ripple on magnet and other supplies, both on the accelerator and any associated experimental equipment.

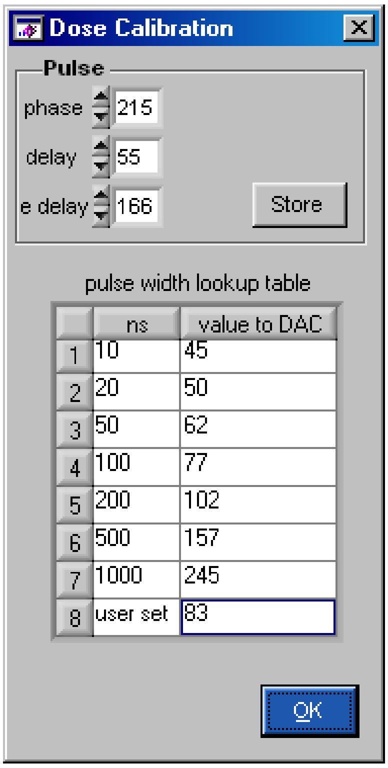

Figure 19: The single-shot modulator calibration panel.

The 'delay' control defines a delay time between the triggering of the accelerator and production of the electron pulse(s) in either repetitive or single-shot modes, while the e delay control defines the timing of the single-shot pulse with respect to accelerator radiofrequency pulse. These values are not expressed in specific units, and are simply 8 bit numbers (0-255) which are interpreted by the accelerator main timing microcomputer. These values can be stored in the microcomputer's memory and form part of the basic timing adjustment of the accelerator.

The single shot pulse widths are similarly defined: a look-up table is used to set the delay time between triggering of the single pulse modulator switches (Figure 5). This delay in turn defines the resulting pulse width.

(6) Beam switching

The beam transport system outside the Faraday cage is calibrated with the aid of panel shown in Figure 20. This defines the current setpoints and current readbacks of the magnets described in section 11.

Since all the power supplies used here are of the constant-current voltage-controlled type, linear two-point calibrations are implemented here.

The hardware system associated with these supplies is distinct from the main accelerator system. Although the software has also been developed as a distinct code, the calibrations have been embedded in the accelerator code.

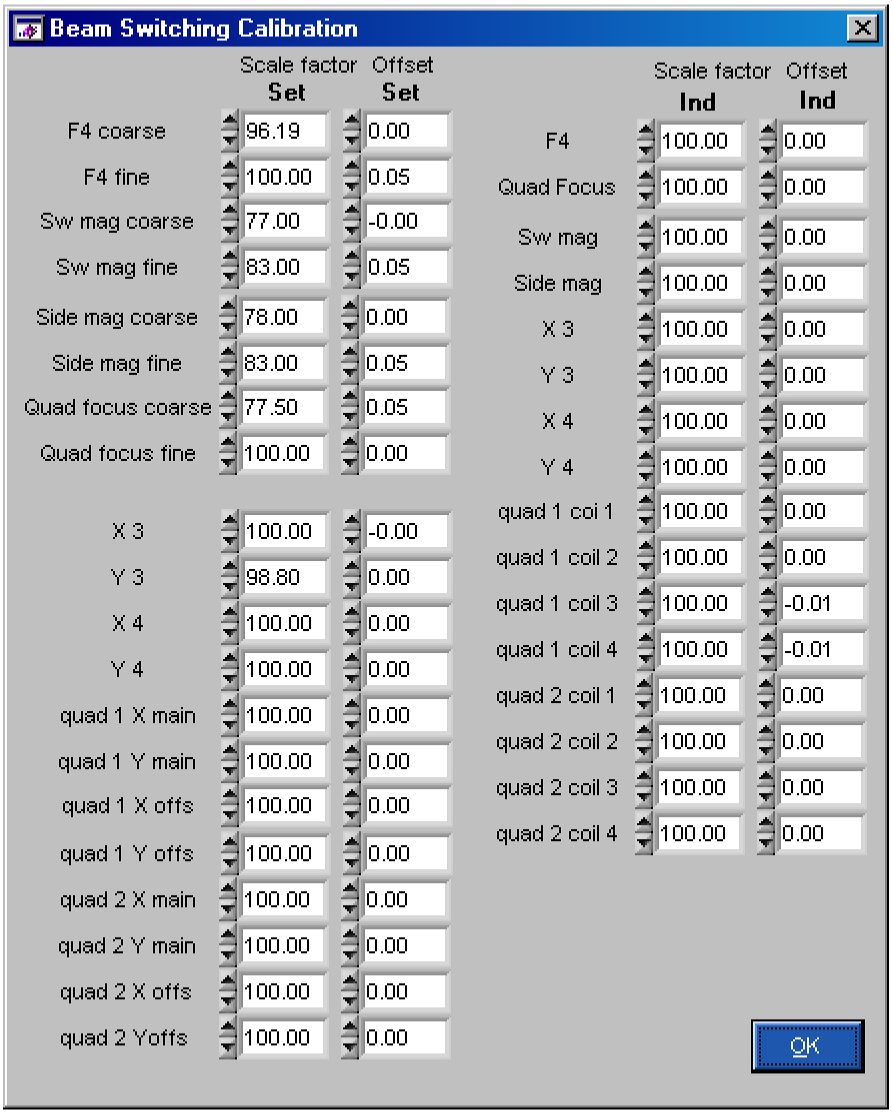

Figure 20: The beam transport calibration panel.

11. Beam switching panel

This panel controls all functions associated with the transport of the beam *outside* the linac body as shown in Figure 21. The associated software panel is shown in Figure 22.

Several magnets are used to focus, switch the direction of the beam and modify the shape of the beam. The magnets requiring higher currents are cooled, where required, by an independent water loop which is not monitored electrically; a water flow indicator is however present near the system, on the wall adjacent to the assembly.

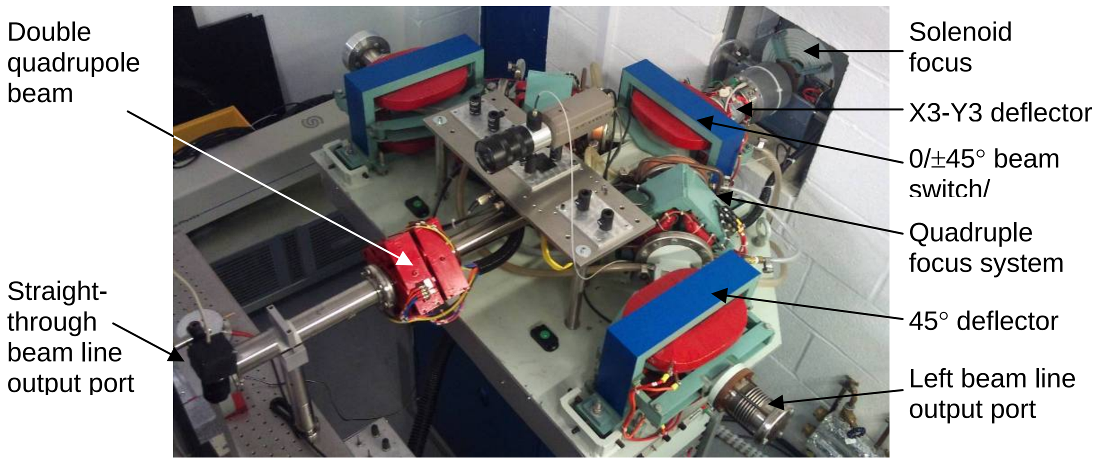

Figure 21: The linac's external beam control optics, used to shape the straight-through beam and to provide two additional orthogonal beam lines.

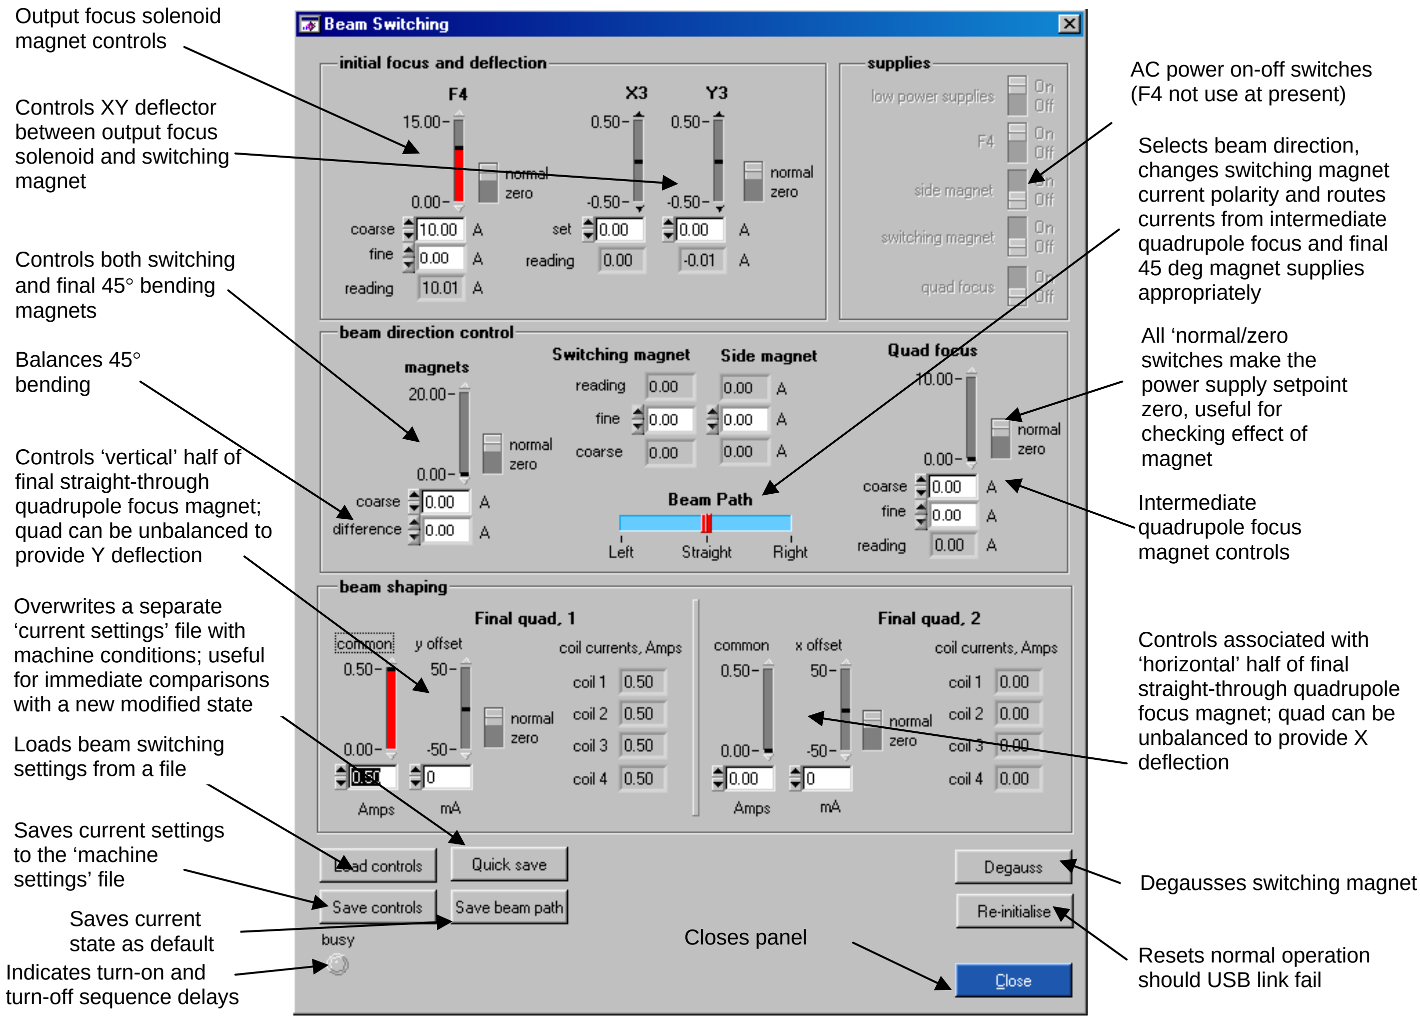

Figure 22: Panel controlling the transport of the beam outside the Faraday cage, allowing routing of the beam to the left or right output ports and control of beam shape in the straight-through beam line.

This software system drives hardware outside the linac Faraday cage through an RJ45 cable extended USB link. This link terminates in an in-house developed digital-to-analogue and analogue-to-digital converter system which interfaces with two groups of power supplies. One of these groups is constructed within the same chassis and is based around low current power supplies, as described in a associated document “Constant current power supply for powering deflection and focus coils”. We don’t show an image of this as it is somewhat boring: it consists of a 3U 19” rack with an on-off switch, not very interesting!

The second group of power supplies delivers much higher currents to drive focus and beam direction switching magnets. The high power supplies used with this group are shown in Figure 23; they are located underneath the beam switching assembly. Analogue voltages are used to interface these, derived form the composite in-house developed controller.

Figure 23: The external beam transport power supplies.

12. Status information

Three methods of conveying information to the operator are used.

One is a scrolling ‘System log’, shown in Figure 24, where significant events are written to a panel along with the time of occurrence. These events represent changes of state, where a software flag is seen to have changed. The software will have two variables representing the last known state of a parameter, which is set immediately after it is read in the main program loop, and the current state. If these are not the same then it indicates that a change event has occurred and an appropriate message is displayed on the panel. The ‘Status’ panel uses the same information but leans toward any system status which would inhibit radiation. It is intended as a quick check when operation is not as expected to be. The edit text box on the status panel is rewritten every program loop and so always shows the current state. For example, the room survey being set appears on the event log but since this is a requirement for radiation it does not show on the status panel. Should the survey be broken however, the event log is updated and this is shown on the status panel too.

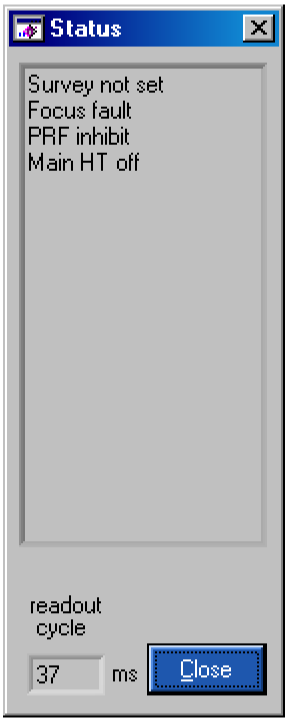

Finally, all machine parameters are written to a ‘log file’ every second. The file is in comma-separated variable (.csv) format, and is backed up every 10 seconds. The idea is that if the file needs to be inspected during linac operation the backed up file is viewed, leaving the original file still accumulating data. When the back-up file is closed by the operator and thus available to the program again, the back-ups continue without any loss of data. While this file is useful to examine previous operating conditions the 1 sec data interval allows detailed inspection of events should a fault occur.

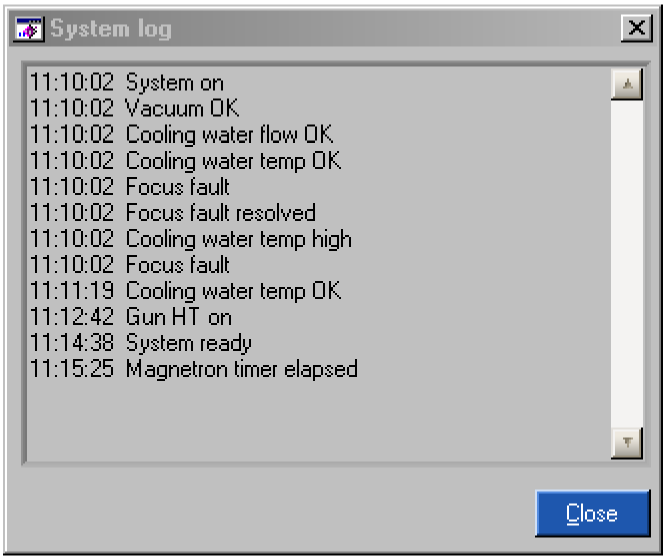

Figure 24: The linac status information panels.

13. Safety interlocks / routines

Interlocks are best implemented in hardware. Software is too prone to change, bugs and unexpected behaviour to be used for critical actions. However, an added level of machine security is useful and has been implemented. As mentioned before, the linac status in general is detected in the main program loop and it is here that certain conditions lead to software actions, for example:

- When a fault occurs while irradiating, send a 'stop pulsing' signal and pop up a warning message
- If a vacuum fault occurs the gun filament is turned off
- If a fault occurs which removes power to the gun supply (such as low water flow) the gun filament current is set to zero
- If the main HT is 'off' during pulsing and automatic magnetron frequency tuning is selected, then this automatic mode is disabled so as to prevent the AFC loop operating from invalid RF cavity signals.

14. Settings file management

If the operator has set up various parameters to achieve a particular beam output, a 'settings' file can be saved so that the machine state can be recalled at a later date. Files are saved as plain text, as shown in Figure 25, to make it straightforward to open outside the main program for either reading or editing purposes. When a previously saved file is loaded into the linac software several checks are needed. The first line of the file indicates what type of file it is so if this is incorrect the file handling routine displays a warning message and exits. This is followed by 'reading in' a series of descriptions followed by the associated setting values. If the expected description is not found, a warning is displayed to indicate that this item has not been set, but the routine carries on with subsequent items.

This feature is of course essential: the settings file should contain information about all of the linac controls, so if anything is missing or incorrectly manually edited, the operator must be alerted. On the other hand we may not want to set every control. For example, a simple switchover from repetitive pulsing to single shot pulsing requires we set the gun selector switch position, turn on the gun HT supply and set it to the required voltage. Rather than do these three steps explicitly in the code we can use a settings file. The technique employed is to create the settings file manually but leave the other settings values blank. This indicates to the program that the file recognises those items as needing to be present, but that the current machine values should be preserved.

The routine for loading files needs to be future-proof and backwards compatible. When searching for a item description it always starts at the beginning of the file, rather than reading items in sequentially. This means that the order of items is not important, particularly relevant if any additions or deletions are made.

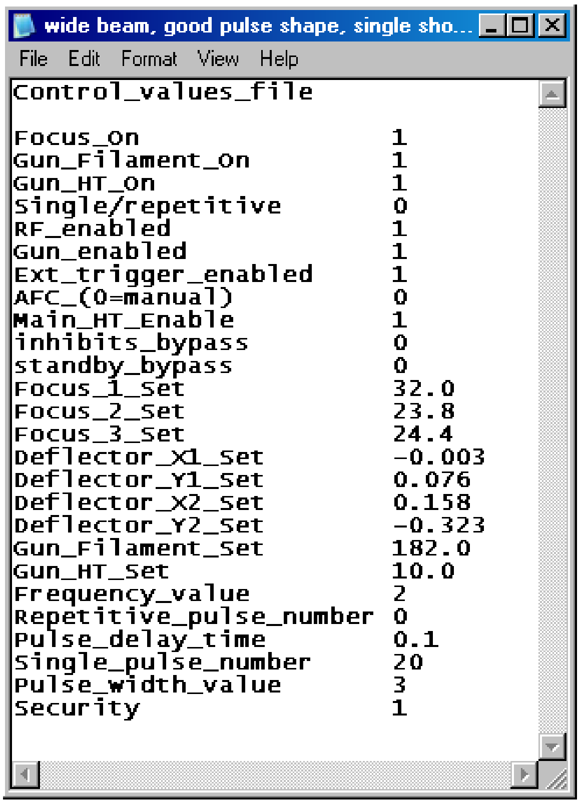


```
Control_values_file

Focus_On                 1
Gun_Filament_On          1
Gun_HT_On                1
Single/repetitive        0
RF_enabled               1
Gun_enabled              1
Ext_trigger_enabled      1
AFC_(0=manual)           0
Main_HT_Enable           1
inhibits_bypass          0
standby_bypass           0
Focus_1_Set              32.0
Focus_2_Set              23.8
Focus_3_Set              24.4
Deflector_X1_Set         -0.003
Deflector_Y1_Set         0.076
Deflector_X2_Set         0.158
Deflector_Y2_Set         -0.323
Gun_Filament_Set         182.0
Gun_HT_Set               10.0
Frequency_value          2
Repetitive_pulse_number  0
Pulse_delay_time         0.1
Single_pulse_number      20
Pulse_width_value        3
Security                 1
```

Figure 25: An example of the linac settings file.

15. Communication with external events

For certain experiments the irradiation has to be synchronized with other events. Also, as a matter of convenience, it is useful to be able to trigger an irradiation from another program on another computer rather than having to return to the main linac program to press a button. Such inter process communication has been set up via the method of 'Named Pipes' described in more detail elsewhere in relation to the robot software.

All pipe code is threaded so as not to interrupt the main program. Three pipes are handled, one receiving messages from the robot, one sending messages to the robot, and one receiving messages from the 'experiment'. The robot interactions consists of the robot triggering an irradiation and the linac triggering a robot sequence when irradiation is complete. The experiment pipe is set up to receive settings such as the pre-RF delay time. This exchange of messages is very flexible and easy to build on, although reprogramming is obviously needed when new functionality is added.

16. Miscellaneous routines

Most of the linac software involves actions which occur immediately in response to a button press or a value change. There are certain routines however which require more careful coding, and these are described below.

- **Setting the gun filament current:** The gun filament lifetime is reduced when subject to sudden increases of current and so a method of ramping up the filament current has been implemented. This is done simply by looking at the target setting and comparing it to the last value sent. If the target is greater than the last value the last value is incremented and updated. If the target setting is e.g. 180 (on a scale of 0 to 255) and the current gun setting stands at 100, a value of 101 is sent. A predetermined 200 ms delay is then started; when it completes, the process repeats, sending the value 102 and so on, until the actual current setting matches the target setting. Since it is possible that the target setting may have changed since the routine began, it must be monitored during each new iteration. Finally, because this loop would block other functions, it takes place in its own thread.
- **Beam switching assembly operations:** The larger power supplies for the switching magnets are controlled in the same way as the gun filament described above, although for different reasons. Due to the high inductance of the coils there is substantial energy stored in the magnetic field when energised. Reduction of the current should happen slowly so that this energy is released slowly and absorbed by the source supply back-emf diodes. Similarly currents are ramped up slowly in order not to 'stress' the constant current power supplies, as they would need to increase their open-circuit voltage and thus inevitably 'clip'. A further consideration involves switching the supplies on and off at power up. This is done sequentially with a delay in between each energisation so that inrush current surges are minimized. During such sequences further changes (say by loading a settings file) are not be allowed and, if attempted, a warning message is displayed.
- **Pulsing timeout:** The gun filament lifetime is reduced when operated at high currents. Furthermore any high voltages on the linac should be avoided when not required. It is for these reasons that a timeout routine was installed, operating as follows: when the 'Survey' is not set and the linac is not pulsing, (i.e. the machine is in an idle state), a timer is started if either the gun filament current is above a given level or the pulsed gun modulator pulse high voltage is above a given level. If both are low then no further action is taken. Now, if the timer is active and both items are reduced sufficiently, the timing is aborted. If a 'room survey' is performed, the timer is also aborted. Similarly, if pulsing occurs the timer is aborted. However, if the timeout period elapses then the gun current and/or pulsed gun modulator high voltage are reduced to lower levels. This action sets a flag in the software. Should a 'room survey' be then performed, indicating that work is about to recommence, the software ensures that the gun high

voltage and the gun filament current are returned to their pre-timeout values. The issues in programming such a routine involve anticipating exceptions to the sequence as described. For example, if the gun filament current has been lowered via the timeout, but then is changed by the operator before surveying, should it be returned to its pre-timeout value or left as it is? The decision made was that the operator has made a deliberate change which should not be overridden.

17. Conclusion

This document provides a brief overview of the software system developed to control the Institute linear accelerator. The in-house developed software was designed to allow simple operation by users of the facility, while allowing trained personnel to set up the accelerator in a safe manner. We have found this system to be reliable over the last 2-3 years and perhaps the simplest way of describing the arrangement is that we have a 'plug and play' USB-controlled accelerator!

This document was prepared by S. Gilchrist and B. Vojnovic during January and February 2012. The assistance of Robert Newman, P.R. Barber and G. Pierce with the development of the software system is gratefully acknowledged. This document should be read in conjunction with other descriptions of our linear accelerator, where details of the hardware are described more thouroughly.

We acknowledge the financial support of Cancer Research UK and the Medical Research Council.