THE NOVICE'S GUIDE TO ELECTRON LINEAR ACCELERATORS

Introduction

In a linear accelerator electrons are accelerated by the action of radio-frequency electromagnetic waves. Relatively low energy electrons are injected into an accelerating structure and gain energy as they travel down the structure. In most electron linear accelerators, very high frequency waves (usually of wavelength of around 10 cm, i.e. around 3 GHz) are used and these are made to propagate down the accelerating structure in the same direction as the electrons. For reasons that will hopefully become clear, the accelerating structure is complex and its function is to impart a significant proportion of the radio-frequency power to the electrons. This structure is basically a waveguide, usually of length of a metre or more, with corrugations on its inside. These corrugations, or iris diaphragms, cause the waves to travel at a velocity determined by the iris and waveguide dimensions. An example of such a corrugated waveguide is shown in Figure 1.

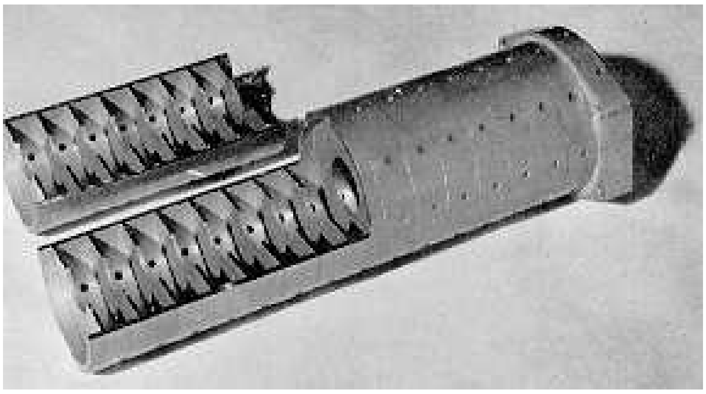

Figure 1: Cutaway view typical accelerating waveguide.

A very rough analogy of the operation of linear accelerators is to envisage electrons as being surfers carried along by a sea wave. Just as in the case of a surfer, electrons move at approximately the same velocity as the wave. Continuing the analogy, it is difficult for surfers to catch the right sea wave and in just the same way electrons have to be injected at the right energy to be carried along by the wave. Similarly, it takes a fair amount of power in the wave to perform the acceleration of the rider and in linacs such high powers are usually generated by magnetrons or by klystrons; the mode of operation of these devices is described later. It is not certain if the analogy could be continued when it comes to surfing acrobatics, unless it is to imagine the behaviour of poorly focused electron beams in the accelerator!

The electric field component in the RF wave acts on the injected electrons, initially by forming them into a bunch, then accelerating them down the corrugated waveguide. This process is illustrated in Figure 2. Each cycle of the RF electric field increases the energy of the particles so that when they emerge at the output end of the waveguide, the effect is the same as if they were accelerated by a static electric field. It is important to realise that the velocity of electrons is not constant along the waveguide. Typical injection energies of around 40 kV result in electrons travelling with a velocity of around 40% of that of light, while their velocity when fully accelerated differs by less than 0.1% of that of light and the design of the iris diaphragms is thus quite tricky. In the same way, the initial mass of the electrons is increased by over twenty-fold in the acceleration process. The more distant sections of the waveguide thus increase the relativistic mass of the electrons by a large factor, while their velocity remains almost unchanged.

In practice the RF power levels needed to perform the process described above are such that this action cannot be sustained continuously and almost all linacs operate in a pulsed repetitive mode. The envelope of the RF system's output is a square (-ish) pulse of duration of a few microseconds and the pulse repetition rate is of the order of tens to hundreds per second. During this pulse RF power levels of several megawatts are needed and typical average powers of a few kilowatts are thus necessary. In most conventional linacs, electrons are injected synchronously with the RF pulse and thus typical accelerated electron pulse widths of a few microseconds are produced. There are a number of ways that the pulse width may be reduced and some of these are described later.

Electrons are commonly generated by thermionic emission from cathodes either directly heated or indirectly heated (oxide cathodes). In the case of magnetron driven accelerators, it is thus common to apply the same voltage pulse to excite the magnetron and the electron injector. This pulse amplitude is usually 20-50 kV and the pulse is generated using a step-up autotransformer driven from a thyratron-discharged delay line or pulse-forming network. In fact most RF generation systems are direct descendants of RF systems developed just before and during WW2 for radar purposes. The requirements there are similar, namely generation of RF pulses of duration of a few microseconds at peak powers in the megawatt range and most of the techniques used in linacs owe much to this wartime work.

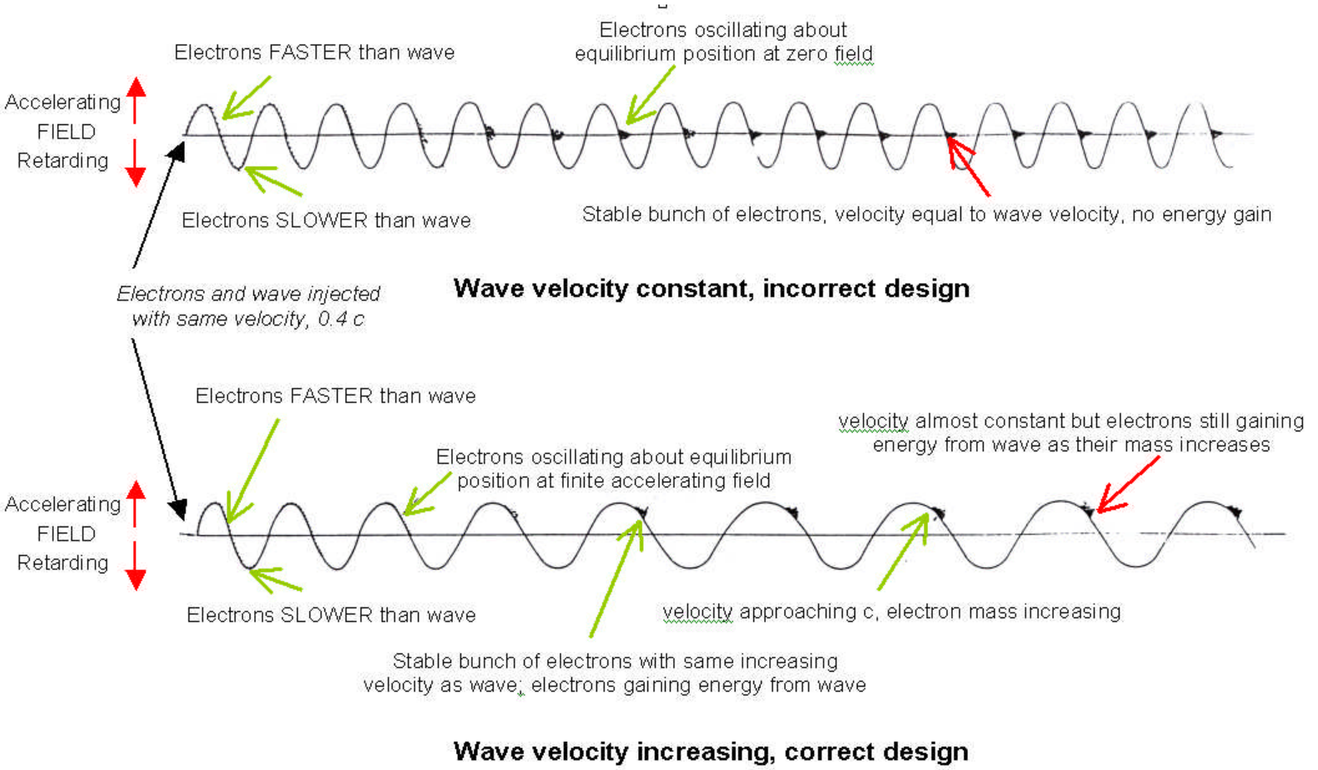

Figure 2: Operation of a correctly and incorrectly designed linac, showing electron bunches 'riding' the wave.

A basic circuit of a linear accelerator pulsing system is shown in Figure 3 and consists of a high voltage supply (E), which is used to charge the pulse-forming network (PFN). This is a lumped component delay line (a conventional transmission line would be impractical due to the length required for microsecond-wide pulses). A high current switch, a thyratron (TH), when triggered, discharges the PFN into the primary of a trifilar-wound pulse transformer (T). The operation of the thyratron is described below. Two AC (50 Hz) voltage sources (V1 & V2) are applied at the 'bottom' end of the pulse transformer and are used to power two filaments, one in the magnetron (M), the other in the accelerator's electron gun. Both filaments are pulsed negatively by the pulse transformer.

Radio-frequency power is led out of the magnetron using conventional rectangular waveguides and often a circulator or isolator is used to prevent reflected power from adversely affecting the operation of the magnetron. A vacuum window which is transparent to RF (usually mica) is placed in the path of this waveguide as a barrier to the vacuum inside the accelerating structure. RF power enters the corrugated guide through a waveguide transformer which both matches the guide to the magnetron and changes the wave propagation mode. Now that both RF and electrons are launched, acceleration can take place provided that the operating frequency is correct for the design of the accelerating waveguide. A similar RF/electron beam decoupling arrangement is used at the output end of the accelerator and any remaining RF power is led out of the machine, again through a window, into a conventional waveguide and then to a matched load where it is absorbed. Sometimes the RF from the output end is guided back to the input transformer where it is recombined at the appropriate phase with the input wave, reinforcing the latter and thus increasing efficiency. The high-energy electrons exit the accelerator through a thin (a few microns thick) metallic window. One often describes the output as being a pulse though this is not true. At the beginning it was stated that the RF field first bunches the electrons; they come out of the accelerator still bunched, i.e. they follow the RF cycles. At an operating frequency of 3 GHz the period is around 330 ps and a series of very narrow pulses, separated by 330 ps is actually what is delivered during the acceleration time of a few microseconds. One important consideration, essential to achieve satisfactory acceleration has been left out in the above description. Electron beams diverge due to the repulsive force between similarly charged particles. One or more magnetic lenses are used to confine the electron beam along its path through the structure.

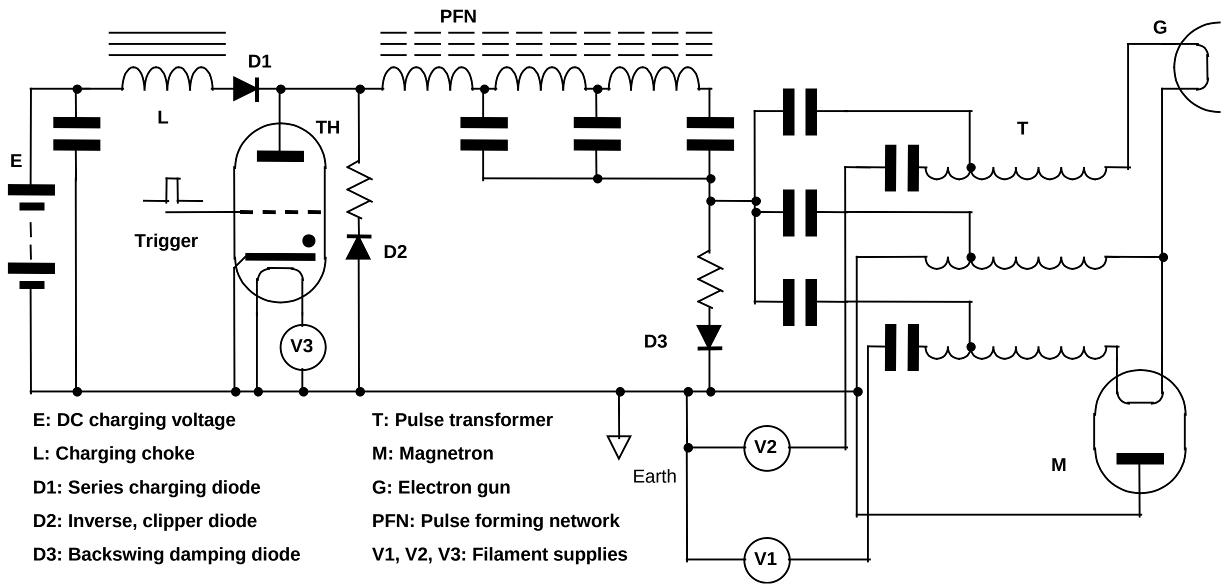

Figure 3: Basic circuit used to trigger the magnetron, M, by charging the pulse forming network, PFN, by E and discharging it through thyratron TH.

These are usually electromagnets operating at electrical power levels of hundreds of watts to a few kilowatts and are wound as solenoids and placed along the length of the accelerator. In practice other magnetic elements (deflection coils, beam shaping elements) are also used to thread the beam through the various restrictions and to make up for any misalignments (physical and magnetic) that are inevitably present in a practical machine.

The power devices that make it all work

- **The thyratron**

A thyratron is a fast controllable electronic switch. Thyratrons can be made to withstand many tens of thousands of volts when they are in the ‘off’ state and can conduct many hundreds of amperes when they are turned on. They are similar in construction to one of the earliest electronic devices developed: the triode valve, but differ from triodes in that they are filled with a gas, most often hydrogen. In a vacuum triode, the current through the valve is controlled by the potential on a control grid which repels electrons on their way from a cathode to an anode. In a thyratron, the action of the control grid on the current through the valve is quite different. The potential on the grid is normally negative and that on the anode is positive, relative to the cathode. This grid potential is sufficiently negative to prevent electrons emitted by the cathode from reaching the anode and the thyratron is in the ‘off’ state. If the grid potential is now made to be zero or positive, electrons can leave the cathode and current starts to flow through the tube. These electrons now form positive ions by collision with the gas in the space between the grid and the anode, provided that the anode voltage is greater than the ionisation potential of the gas filling the tube. A glow discharge now forms and the anode current increases rapidly to the full emission current of the cathode. The thyratron is now ‘on’. If the grid potential is made negative again, the positive ions drift towards it to form a cloud of positive space charge which shields it from the rest of the ions. In other words, the grid no longer controls the tube current and the tube stays in the ‘on’ state until the anode potential falls below the ionisation potential of the gas. At that point, the tube reverts to the ‘off’ state and no current flows. This switching action can be made to be very fast (much less than a microsecond) and the thyratron is thus an ideal device to use when large currents from the pulse-forming network need to be switched into the pulse transformer and thence into the magnetron.

- **The magnetron**

A magnetron is an oscillator which can produce very high powers of very high frequency radio waves, or microwaves. Their uses are more widespread than is realised: there's one in every microwave oven! Schematic views are shown in Figure 4. Oscillations build up from noise contained initially in a thermionic-emission current from a heated cathode Eventually electronic interaction takes place between a circulating current and a 'pi-mode' radio-frequency standing wave on a slow-wave circuit. During the operation of the magnetron, this interaction takes place due to a circulating hub of space charge, supplied primarily by secondary electrons from the cathode. This is shown as the shaded outline in figure 4 (a). The electron emission is such as to produce initially circular particle trajectories due to the presence of an axial magnetic field provided by a permanent magnet. The internal structure of magnetrons incorporates a number of cavities which determine the oscillation frequency. Although they are essentially devices that operate at a single frequency, they can be mechanically tuned, sometimes by as much as 10% or more.

The details of their operation are a bit tricky to understand and for a full explanation the reader is referred to the texts in the references. A simplified explanation is given here which assumes that oscillations have started, i.e. that the puzzle whether chickens came before eggs or vice-versa has been satisfactorily resolved in what is essentially a circular (!) argument. The combined action of crossed electric and magnetic fields (supplied by the supply voltage and by the permanent magnet respectively) causes the cathode-emitted electrons to spiral outward towards the anode in a cycloidal path. If oscillations have started, by the time the electrons have reached the surrounding cavities, the electric field is the resultant of the supplied pulsed DC field and the RF electric field across the cavity gaps. The so called 'pi mode' of operation means that, at any one time, the phase of the RF field will be reversed at alternate cavities i.e. there will be a 180 degree phase shift between adjacent cavity gaps. At each cusp of the cycloidal electron path, their kinetic energy is zero and all the energy imparted to the electron by the pulsed DC field has been transferred to the RF field, in other words the efficiency is high. When tuned correctly, the spiralling electrons continue on to the anode because the phase of the RF field is reversed across the cavities. Proper adjustment of the electric and magnetic fields causes the electrons to progress laterally the width of one cavity each half-cycle of the RF field and oscillations will thus be sustained in the cavities. At the end of this process high levels of RF power are generated within the magnetron and this may be extracted through a slot located in the sidewall of one of the cavities at a high current point. The slot couples to the output waveguide.

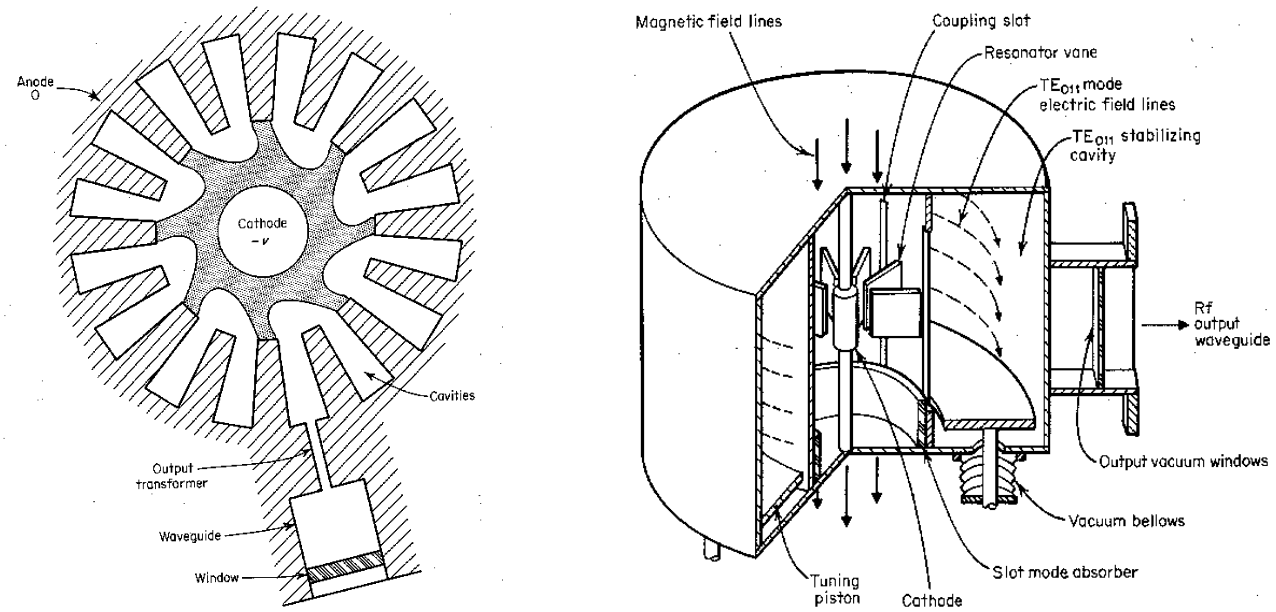

Figure 4: Typical conventional (left) and coaxial (right) magnetron structures. In both cases RF power is extracted through a window in the waveguide. A series of cavities around an axial cathode, surrounded by an anode are responsible for the generation of microwaves.

- **The klystron**

Klystrons are used as radio frequency sources when higher frequencies and powers are required in linacs capable of generating higher energy particles. Klystrons can be configured to act as amplifiers (as shown in Figure 5) or as oscillators, where RF oscillations are sustained by internal feedback mechanisms within the structure. Their construction is in many ways similar to that of a linear accelerator itself in that they rely on a linear electron beam accelerated towards a collector, travelling past a series of cavities and being constrained by a series of focusing coils.

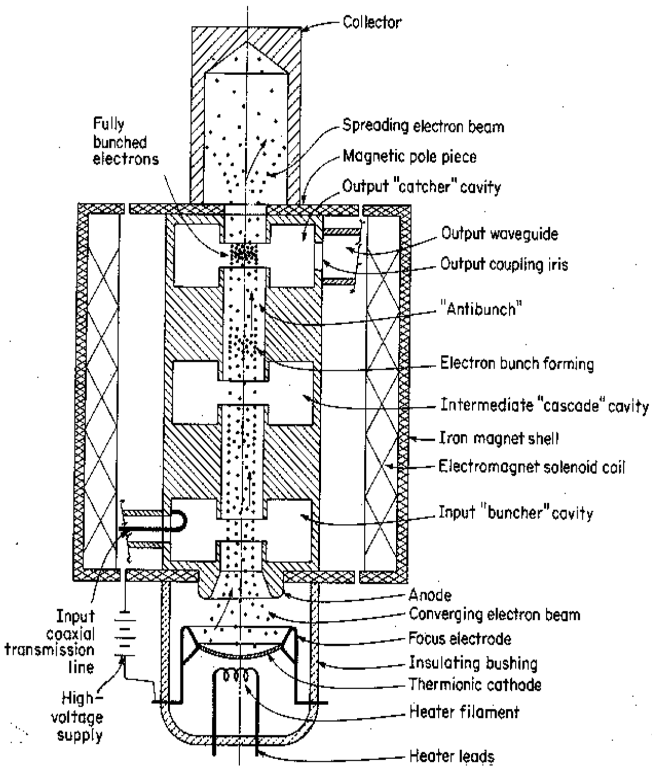

Figure 5: Cross-section through a klystron tube.

However the similarity ends there. Once again a detailed description of the mode of operation is not presented here. The electrons that leave the gun of a klystron are uniform in density, but the co-operative action of the drift tubes and cavities imparts a high degree of density modulation to appear in the beam at the output gap. This action is called bunching and is akin to that which takes place in the accelerating structure of a linac. An input coaxial line excites the input cavity and a sinusoidally varying electric field thus appears across the input gap. This electric field causes the electron velocity to vary in a similar fashion, except for electrons which pass through the gap when the RF voltage is zero; these pass into the drift tube at constant velocity. The result of this velocity modulation is that electrons tend to bunch progressively around the DC electrons which are travelling at a velocity determined by the cathode-collector voltage. Subsequent cavities reinforce this action and ultimately large amounts of power may be extracted from the output cavity. Klystrons can generate many tens of megawatts at microwave frequencies.

The circulator or isolator

When high powers of RF are generated using the devices described above, it is desirable to provide optimum matching of the load to the generator in order to ensure maximum power transfer. In practice this is not always possible and standing waves, as a direct result of reflections within the system, are quite likely to occur. This situation is undesirable as far as the generator is concerned because significant power levels may be fed back to the generator causing excessive dissipation or stresses which would hasten its destruction or reduce its lifetime. A device called an isolator is thus often inserted between the generator and the load, in our case the linac itself.

Isolators are based on devices called circulators which behave like roundabouts for RF with three exits (or 'ports'). Power entering one of the ports (#1 e.g. from the generator) exits the adjacent port (#2). Power entering that adjacent port #2 (e.g. reflected power from the load) must go onto the next port (#3). Here a matched load is placed which absorbs the reflection. The whole arrangement behaves as a one way valve for RF. Isolators are passive devices relying on Faraday rotation in a magnetically biased ferrite.

Shortening the electron output pulse

As described above, linear accelerators produce an output pulse, or rather a series of bunched fine-structure pulses, during the time that the radio-frequency power is present. If output pulses shorter than this are needed, several techniques may be employed. Shortening the RF pulse itself is not possible since a short RF pulse would imply a system with a broad resonance, or low ‘Q’ value, i.e. the RF frequency would not be stable. Systems with a sharp resonance necessarily take a long time to get to resonance (think of a tuning fork which takes a long time to emit the maximum sound level or the pendulum of a mechanical clock).

One simple way to shorten the output pulse is to prevent injection of electrons for most of the time that the RF is present, and to supply them to the accelerator for a short time only. The simplest way to achieve this is to separate the gun filament from the magnetron filament in the circuit of figure 3. This requires an independent pulsing circuit with its own thyratron and pulse transformer, triggered some time after the RF is triggered, for a time determined by a shorter pulse-forming network. The gun filament is thus pulsed during the RF pulse and electrons enter the accelerator only during this short time. Practical limitations in the design of the pulse transformer needed to produce the tens of kilovolts required to inject electrons are such that pulses shorter than about 100 ns are very difficult to achieve.

Another analogous method which is somewhat faster is to use electrostatic deflection to steer the beam away from an input slit on the input of the accelerator. At some point during the RF pulse, the electron beam is swept past the slit to a new deflected position on the other side of the slit. The output pulse is now dependant on the length of time that the beam spends in front of the slit.

An intrinsically faster method is to place a grid between the filament and the accelerator filament. This is biased negatively with respect to the filament so that electrons are repelled back to the filament. When the pulse is needed the grid is biased positively and electrons are allowed to enter the accelerator. Because the grid can be placed in close proximity to the cathode (filament) it can exert significant influence on the beam with very low biasing voltages and consequently low voltage swings are needed to turn on the beam (a few hundred volts). Pulses down to a nanosecond or so may be generated using this technique. One limitation is now the cathode itself which needs to be indirectly heated (oxide type of cathode). This requires a very good vacuum to prevent poisoning of the cathode surface and output current are generally limited to few amperes. A variant of this method is to select or pick off just one of the fine structure pulses discussed at the beginning. In a 3 GHz machine these pulses occur every 330 ps and the control grid thus needs to be pulsed for only 300 ps or so in order to generate a much shorter pulse of a few picoseconds. Synchronisation problems can be quite severe however and time jitter must be reduced as much as possible.

All these methods essentially ‘waste’ the electron beam during the time when it is not wanted. There are, however, ways of inserting separate bunching systems, too complex to describe here in detail, which compress the charge available in a long pulse into a much shorter output pulse. Such systems can thus produce very high peak beam currents.

Conclusion

Linear accelerators are very versatile generators of high energy particles. They can be designed to cover energies from 1 MeV or so all the way to several hundred MeV and can produce very short pulses either repetitively or as single, triggered pulses. Average beam powers of several kilowatts are quite feasible and indeed often essential in cases such as clinical therapy machines. They require relatively little maintenance (often limited to periodic changes of filament) and many have been in continuous service for tens of years.

Although easily pronounced, the term ‘linac’ now commonly used, was felt a not to do justice to the elegance of the arrangement; there was an attempt in the early fifties to have the term ‘orthotron’ adopted (ortho: straight). This was to impart a degree of majesty to these machines and to give an idea of their straight-line mode of operation; sadly it was not to be and the term linac stuck. It is hoped that some of the mysteries surrounding linear accelerators have been clarified. The greatest mystery, however, has not been tackled, namely why the word linac is pronounced as if it had two n's in the middle.

Finally some pictures of a real linac, one that came into service at the Gray Laboratory and which prompted the writing of this document by B Vojnovic in late 1995.

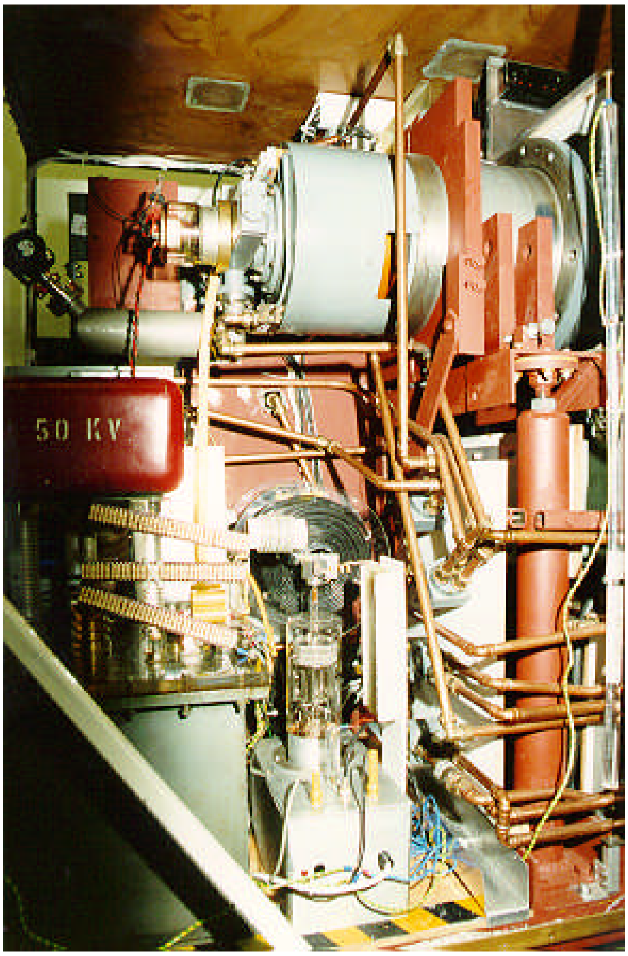

The gun end of the Gray Lab 5 MeV linac. The pulse thyratron can be seen at the bottom of the picture, with the gun at the top left and the focus coils at top right.

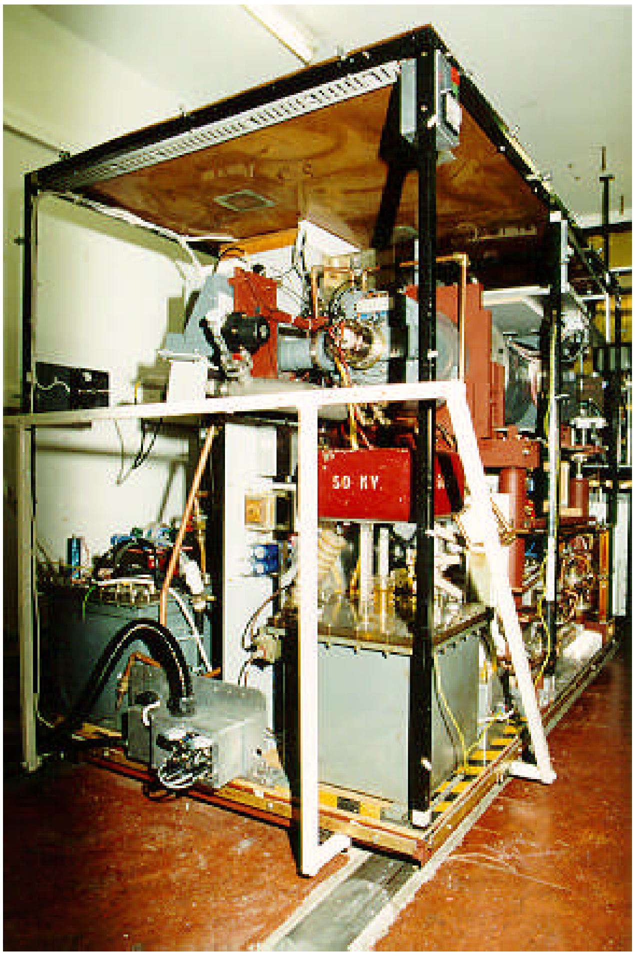

View of the accelerator from the gun end, injection pulse transformer in the foreground, magnetron pulse transformer at the bottom left.

References:

Power microwave devices:
Electronic Engineer's Handbook, edited bt Fink and Christiansen, 2nd Edition, McGraw Hill,1982 edition, ISBN 0-07-020981-2.
Electricity and Magnetism, by Bleaney and Bleaney, University of Oxford Press, 1957.
Electronic Designer's Handbook,edited by LJ Giacoletto, 2nd Edition, McGraw Hill Book Co, 1977 edition, ISBN 0-07-023149-4
Waveguides, cavities etc:
Microwaves, by AJ Baden Fuller, 2nd Edition, Pergammon Press, 1979, ISBN 0-08-024227-8
Transmission Lines for communications by CW Davidson, Macmillan Press 1985 ISBN 0-333-32738-1
Linear Accelerators:
Nuclear Physics by AES Green, McGraw Hill Book Co,1955
Introduction to Atomic and Nuclear Physics by H Semat, 3rd Edition, Chapman and Hall

