RF STRUCTURES FOR LINEAR ACCELERATORS*

G.W. Wheeler and S. Giordano
Brookhaven National Laboratory
Upton, L.I., N.Y.

Introduction

This paper will give a brief review of the fundamental requirements for rf accelerating structures and discuss some of the structures which are used or being considered for linear accelerators. No attempt is made to treat the subject in depth.

Linear accelerators are attractive accelerating mechanisms for several reasons. They are capable of accelerating beams of very high peak current and good quality (small momentum spread and small beam cross section and divergence). Injection of the beam into and extraction from the accelerator are accomplished in a most straightforward fashion with high efficiency. For acceleration of electrons to high energies, the linac has the additional virtue of negligible radiation loss from the beam as compared to an electron synchrotron. On the other hand, linacs require very high peak rf power and so are usually pulsed rather than operating continuously. Also, for very high energies, the linac becomes very long.

We may conveniently divide linacs into two classes: those for the acceleration of electrons and those for protons and heavier ions. There are no fundamental differences in the principles of acceleration for the two classes but the particle velocities are greatly different in the energy range of interest. For example, an electron reaches a velocity equal to 99% of the velocity of light ($\beta = v/c = 0.99$) when its kinetic energy is about 3.1 MeV, but a proton does not reach the same velocity until its energy is about 5.7 BeV. Thus, after a short initial section (about 1 meter), the velocity of an electron is essentially constant as it passes through the accelerator. However, the velocity of a proton will still be changing significantly at the output of any proton linac presently under consideration. There are two important results of this difference: the Principle of Phase Stability is of great importance in proton linacs but is not in electron linacs and the accelerating structures for the two classes differ considerably. However, when proton linacs are built for energies above a few BeV, they can be expected to look very much like electron linacs.

At the present time, electron linacs have been operated up to about 1 BeV and proton linacs up to about 70 MeV. An electron linac for about 20 BeV is under construction at Stanford and a proton linac for about 100 MeV is under construction in the USSR. Several proposals have been made to construct proton linacs with energies up to 800 MeV.

In order to demonstrate briefly some of the fundamental principles of linacs,[1] we consider a very simple but practical accelerating structure.

*Work done under the auspices of the U.S. AEC.

Figure 1a shows several right circular cylinders of conducting material (copper) placed end to end. When excited with rf power of the appropriate frequency, each cell will resonate in the TM_{010} mode. The electric field has only an E_z component with a maximum value on the axis and no z dependence. The magnetic field has only an H_θ component which is zero on the axis. Any number of such cells may be placed in line to produce the desired energy gain, each cell receiving rf power from a suitably phased external source. Now we inject a charged particle moving along the axis of the cell. The holes in the end walls of the cell are assumed sufficiently small so that they do not disturb the field distribution. Suppose that the average velocity of the particle while traversing the cell and the length of the cell are such that the particle passes through the cell in exactly one half cycle of the rf. In this arrangement, the phase of the rf field in a cell must differ by 180° from that in adjacent cells.

The peak voltage developed across this cell will be

$$\Delta V_o = \int_{-L/2}^{+L/2} E_o \, dz = E_o L$$

and if the particle crossed the cell at the instant of peak voltage, it would gain an amount of energy, $\Delta W_o = eE_oL$. However, the finite time required to cross the cell will limit the energy gain to a smaller value. Suppose that the particle passes the center of the cell at the time when the electric field is at its maximum value ($\varphi = 0$ and $z = 0$ when $t = 0$ in Fig. 1b, φ is the phase of the particle with respect to the rf at the center of the cell). Now the actual energy gain of the particle will be (neglecting any change in the particle velocity as it crosses the cell, i.e. a relativistic particle)

$$\Delta W_s = e \int_{-L/2}^{+L/2} E_o \cos(\pi z/L) \, dz \, .$$

The ratio of $\Delta W_s/\Delta W_o$ is called the longitudinal transit time factor (T_ℓ) and is always less than unity. Thus the particle energy gain is $\Delta W_s = eE_oT_\ell L$. However, if E_o is increased to E_o' and the particle made to pass the center of the cell at some phase φ_s (Fig. 1b), the same energy gain can be achieved, $\Delta W_s = E_o'T_\ell L \cos\varphi_s$. If a particle has a phase different from φ_s, its energy gain will differ from ΔW_s. For a proton (nonrelativistic), this will result in an incorrect average velocity in the cell and a shift of phase of the particle with respect to the rf wave. If a particle arrives at the center of the cell too early in time ($\varphi = \varphi_1$ in Fig. 1b) it will gain

less energy than ΔW_s and so will slip back in phase, approaching φ_s. Similarly, a late particle ($\varphi = \varphi_2$), will gain more energy than ΔW_s and will catch up to φ_s. The particle with $\varphi = \varphi_s$ and $\Delta W = \Delta W_s$ is known as the synchronous particle. Thus, there is a certain stability range in phase about φ_s, in which there is a restoring force which tends to bunch the particles about φ_s and in which they will oscillate in phase and energy about the synchronous values. It is clear that this phase stability exists only for particles with $\beta < 1$. The existence of phase stability requires that the beam be unstable with respect to radial motion so that auxiliary focusing is required to confine the particle beam along the axis of the accelerator. For an electron with $\beta \cong 1$, the time required to cross the cell is independent of the particle phase and only the energy gain varies. In an electron accelerator, the particles are initially bunched around $\varphi = 0$ in order to achieve the maximum efficiency.

The rf power required to establish the field E_o' in a cell is given as

$$P = \frac{(E_o'L)^2}{ZL}$$

where Z is the normal shunt impedance per unit length for the cell. In accelerator design, it is more convenient to define an effective shunt impedance in terms of the particle energy gain rather than the actual cavity voltage, thus

$$P = \frac{(E_o'TL)^2}{R_s L}$$

where $R_s = ZT^2$.

It should be pointed out that the complete transit time factor (T) is more complex than indicated above where the longitudinal transit time factor (T_ℓ) only was shown for a simple case. Clearly, T and not T_ℓ should be used in the calculation of the actual energy gain. Included in T are the effects of the spatial distribution of the electric fields within the cavity, the time variation of the field due to the finite time that the particles are in the field and the radial variation of the field which is complicated by the hole through which the beam passes. In general, T should be determined from a proper integration of the electric field along the particle trajectory.

The transit time factor for the structure of Fig. 1a can be improved by adding stubs or "drift tubes" as shown in Fig. 2. The drift tubes concentrate the field near the center of the cell and if the total voltage across the cell remains as it was before adding the drift tubes, then $E_oL = V = E_g g$. However, the particle now traverses the region of electric field (E_g) in a time less than one half cycle so that the time average value of the field more nearly approaches its peak value. The presence of the drift tubes, of course, alters the resonant frequency of the cavity.

The major drawback to this accelerating system is the complexity of the rf power system which must feed a large number of cells individually, each with the correct amplitude and phase.

The structure used in this example is designated as a π-mode standing wave structure. The operating mode of a linac is defined as the shift in rf phase between adjacent cells. Most structures can be operated with either a standing or traveling wave configuration depending on which way the ends of the structures are designed, remembering that a standing wave can be considered as two traveling waves moving in opposite directions.

The Drift Tube Accelerator

The structure which has proven most satisfactory for the acceleration of protons at low energies and is used in existing proton linacs is the "drift tube accelerator." This is a 2π-mode standing wave structure in which individual cells are directly coupled by the fields within a large cavity, thus reducing the complexity of the drive and coupling systems as compared with that of Fig. 1. The basic field configuration is that of the TM_{010} mode modified by the presence of the drift tubes. This structure operates in the 2π mode because the fields in the cavity are everywhere in phase. The drift tubes serve to shield the particles from the field during the reverse half cycle as well as improving the transit time factor. Such a structure is shown schematically in Fig. 3.

Most drift tube linacs operate at frequencies near 200 Mc/sec which is a good compromise between several conflicting requirements. It is desirable to use as high a frequency as possible because the shunt impedance for a particular geometry increases with $f^{\frac{1}{2}}$ and the physical dimensions of the structure decrease. On the other hand, the structure cannot be made too small for several reasons. The minimum diameter of the bore hole is set by the size of the beam so that as the drift tube diameter is decreased, the ratio of bore hole diameter to drift tube diameter increases. As this ratio increases, the effective radial variation of the electric field increases also, with the field on the axis being substantially lower than at a radius equal to the bore hole radius. Thus a large radial variation of the transit time factor is introduced which strongly couples the radial and phase motions of the particles. Also it is necessary to incorporate the radial focusing magnets within the drift tubes and the cell length ($\beta_s\lambda$) becomes impractically small at low proton energies (typical injection energies into present linacs are below 1 MeV).

The major problems in the design of the drift tube structure are the determination of the drift tube and cavity geometry so that the structure will resonate at the design frequency and the determination of the electrical properties: the electric field distribution near the axis, the

transit time factor and the shunt impedance. In the early linacs, these quantities were determined by model measurements and approximate analytical treatments which were not very accurate. It should be noted that exact analytical solutions for Maxwell's equations exist for the simple structure of Fig. 1. However, the boundary conditions imposed by the structures of Figures 2 and 3 make exact analytical solutions virtually impossible and approximations poor.

One of the first steps to improve this situation was taken in the design of the 50-MeV injector linac at Brookhaven. A method was developed[2] for computing drift tube shapes by numerical solution of an approximate form of Maxwell's equations. The drift tube shape was calculated without a bore hole for the protons to pass through and the effect of the hole had to be determined from model measurements. Furthermore, the transit time factor was not calculated explicitly by integration of the computed axial electric field but was calculated from the approximate analytical form.

A similar but more accurate approach has since been developed.[3] An unloaded cavity is driven by an oscillating dipole charge on the axis of the cavity. Solutions to Maxwell's equations are obtained numerically to any desired degree of accuracy. With a proper choice of parameters, there will be a node in the electric field on the central plane through the dipole. A metallic surface can be started at this point and drawn perpendicular to the electric field lines until it encloses the dipole, as shown in Fig. 4. The remaining volume of the cavity will resonate at the specified frequency with a known field distribution. The transit time factor and shunt impedance then are obtained by integration of the known fields. Unfortunately, the method does not calculate the effects of the bore holes, which must be added by a perturbation calculation. The resulting drift tube does not have a simple shape which could lead to some additional complexity of fabrication.

Another approach has been taken by a group at MURA.[4] Here a simple cavity and drift tube geometry are specified, including a bore hole. The resonant frequency and field distributions are obtained by a relaxation calculation over a suitable set of meshes conforming to the established boundary conditions. The transit time factor and shunt impedance are then calculated from the known fields. The numerical calculation is extremely complex and has limited accuracy. The frequency is reported to be accurate to 0.1% and the field values to 1%. This accuracy is, however, adequate for most design purposes and this, together with the simple drift tube shape and explicit inclusion of the bore hole effects, makes this program the most suitable at present for practical linac design.

With these techniques, it is possible to search for a drift tube and cavity geometry in each energy range which gives a good shunt impedance while still allowing for reasonable mechanical design. Extensive tabulations of drift tube geometries have been compiled by the MURA group. Typically, values of R_s as high as 50 MΩ/m at 200 Mc/sec can be achieved at low values of β. However, the values of R_s decrease at higher values of β and at $\beta = 0.55$ (about 200 MeV), R_s has dropped below 20 MΩ/m. In principle, the drift tube accelerator will operate at $\beta = 1$, but the decreasing values of R_s make it urgent to look for more efficient structures at the higher energies.

A single drift tube cavity cannot be made in any desired length but is restricted in length to about 20 to 25 free space wavelengths of the operating frequency. This limit arises from the fact that the mode spacing decreases with L^2 resulting in increasing difficulty in maintaining the desired field distribution within the cavity.

The Iris-Loaded Waveguide for Electrons

Most electron linacs operate in the $\pi/2$ or $2\pi/3$ mode with a traveling wave structure of the iris-loaded type at S or L-band frequencies. An electromagnetic wave will propagate in a conventional metallic waveguide with a phase velocity which is greater than the velocity of light. In order for a wave to impart energy to a charged particle, the wave must have a phase velocity equal to the velocity of the particle. There are a number of ways to "load" a waveguide in order to reduce the phase velocity of the wave. The method generally used in electron linacs is shown in Fig. 5. The unit cell is defined between the planes of two adjacent irises. The coupling between cells is predominantly electric with the guide operating in the TM_{01} mode.

The iris-loaded waveguide, operated in this fashion, has proven to be generally the most satisfactory structure and has been extensively studied at Stanford,[5] MIT[6] and Harwell.[7] However, it is worth noting that at least three standing wave electron linacs have been built. The first, at MIT, is an S-band, iris-loaded waveguide operating in the standing wave mode. The other two, at Yale and Los Alamos, are independent cavity standing wave structures similar to that shown in Fig. 1, and operating at 600 Mc/sec and 50 Mc/sec respectively.

The possibility of using superconducting structures for linear accelerators has been investigated at several laboratories with work continuing at Stanford. Although it does not seem possible to have a truly superconducting surface at radio frequencies, it may be possible to increase both the shunt impedance and Q by factors of 10^4 to 10^6 which would reduce the rf excitation power for the cavities to trivial levels. The very high Q will lead to some severe problems associated with the slow build-up times and very tight tuning requirements. Furthermore, the reaction of the beam on the cavity (beam loading) will be particularly severe. If superconducting cavities become practical, it is to be expected that they will find their first applications in rf particle separators (waveguides which are operated in a mode

having transverse electric and magnetic fields) where there will be little beam loading and then in electron linacs and finally, possibly, in proton linacs.

Waveguides for Protons at High Energies

When first looking for a structure better than the drift tube accelerator at high energies, it was natural to start with the structure used in electron linacs. It is possible to design a traveling wave iris-loaded waveguide for particles with $\beta < 1$, but there are several factors which make this unattractive.

For a proton accelerator, the traveling wave configuration is not as good as the standing wave, as can be shown by the following simple argument. The axial electric field distribution for a traveling wave in a section of loaded waveguide can be set up in several ways for a particular value of beam current. Two ways commonly used are (Figures 6a and 6b) the constant attenuation configuration and the constant field strength configuration. If the design current is I_o, then the field distributions (for a constant power input) will look something like the curves shown for I_o and $I \neq I_o$. For a standing wave cavity, the constant field configuration is shown in Fig. 6c for a constant power input and currents I_o and $I \neq I_o$. For electrons, when $I \neq I_o$, the energy gain of the section is different from its design value which results only in a change in the output energy of the accelerator. However, for protons ($\beta < 1$), the incorrect energy gain in the cavity will cause the particles to shift in phase and possibly to be outside of the phase stable area in the following cavity, resulting in the loss of the beam. In order to avoid this situation in a proton linac, it is necessary to change the power input to the cavity so as to maintain the proper energy gain. In the standing wave case (Fig. 6c) the lines for $I \neq I_o$ are straight and parallel to I_o which is not the case for the traveling wave configuration. Thus, it is much easier to maintain the correct energy gain in a standing wave cavity than it is in a traveling wave section. For this reason, all high energy proton linacs are being planned to operate in the standing wave configuration.

The conventional iris-loaded waveguide operating in the standing wave configuration has been examined for the acceleration of protons and found to be rather poor. The structure as used in electron linacs has a reasonably good shunt impedance at $\beta = 1$, but the shunt impedance drops as the guide is more heavily loaded for lower values of β. In the region of $\beta = \frac{1}{2}$, the waveguide is about as bad as the drift tube structure. For this reason the transition from drift tube to waveguide structure is planned near $\beta = \frac{1}{2}$. Some other structures, intermediate between the drift tube and waveguide structures have been examined for use near $\beta = \frac{1}{2}$, but no use of these is planned in this country. However, the crossbar structure[8] is viewed with favor in Europe. The conventional iris-loaded guide also has been found to have a very small bandwidth which is most undesirable as will be shown shortly.

Electron linacs are operated at relatively high frequencies, S and L bands, where the higher frequencies improve the shunt impedance. However, it has been shown[9] from considerations of the particle dynamics in the accelerator, that for a proton linac which has a 200-Mc/sec drift tube accelerator up to about 200 MeV, a frequency of about 800 Mc/sec is the highest which can be used safely for the following waveguide accelerator.

It appears then that a loaded waveguide accelerator for protons should operate in the standing wave configuration at a frequency near 800 Mc/sec. Furthermore, the conventional iris-loaded waveguide does not appear to be suitable. In recent years, a number of studies of suitable structures have been carried out at Brookhaven, Yale University, Los Alamos and the Rutherford Laboratory.

There are several criteria which can be established for a suitable accelerating structure. These criteria often conflict and compromises must necessarily be made. The first and obvious criterion is high shunt impedance. If the shunt impedance is low, the rf power required to excite the structure will be very high, resulting in excessive cost. However, if the power required to excite the structure is equal to or smaller than the power transferred to the beam, still higher shunt impedance is of less importance. The structure must contain bore holes which are large enough to permit passage of the beam without intercepting any of the particles. Unfortunately, as the bore hole diameter is increased, the shunt impedance decreases.

It is desirable to have a structure which is as small as possible but which can be fabricated with reasonable effort, otherwise the cost may be prohibitive. It will be noticed from the relation between the power and energy gain,

$$PL = \frac{(\Delta W)^2}{R_s}$$

that for a given energy gain and R_s, the product of power and length is constant. Thus power can be traded for length to find an optimum combination for lowest cost. In order to keep the multiplicity of components low, it is necessary to make each cavity as long as possible but yet to match the power requirement of the cavity to the capability of the power amplifier. Also the structure must be capable of supporting the peak electric fields which are required.

Finally, it is imperative that the electric field distributions and the phase relationships among the cells in a cavity be relatively insensitive to various perturbing factors. Among the factors are the departure of the resonant frequency of each cell from the overall resonant

frequency of the cavity, temperature effects and beam loading. The sensitivity to these factors can be reduced by increasing the bandwidth of the structure but this usually results in a decrease of the shunt impedance.

The effects of beam loading, which is the reaction of the cavity to an intense beam of particles, may be divided into two parts, the amplitude and phase variation within a single cavity, and the effects of these variations on a multi-cavity machine. There are two effects within a cavity:[10] a resistive component absorbs energy from the cavity, and a reactive component detunes the cavity. The resistive component represents transfer of energy from the power source and cavity to the beam and results in a decrease in the field amplitude in the cavity. The decrease can be kept small in the standing wave cavity by sensing the field amplitude at one point and using this information in a closed loop system to control the power output of the rf amplifier.

The reactive component of beam loading is caused by the fact that the particles are not uniformly distributed in phase but are bunched around a phase $\varphi_s \neq 0$ (Fig. 1b). This effect is generally not present in electron linacs because the beam is bunched around $\varphi = 0$. The reactive effect causes a change in phase between the rf drive power and some fixed point in the cavity and is analogous to driving the cavity off resonance. Estimates[11] show that for a 0.1-ampere beam, the detuning may be of the order of 10 kc at 800 Mc/sec. A change of 10 kc in the resonant frequency should not cause too large a change in the distributed amplitude and phase variations along a single cavity, nevertheless these variations must be taken into account. Measurements[12] made on a short section of a slotted-iris structure show that a considerable reduction of these distributed amplitude and phase variations can be achieved with the use of two radiofrequency power drives in each cavity.

In a multi-cavity linac, it is important to maintain the relative phase between cavities. A change of 10 kc in the resonant frequency will cause a phase change of 10° to 15° for each cavity. However, the phase change will be in the same direction in each cavity. For example, if each cavity were designed so that for a fixed beam current each individual cavity would have a phase change of $+12^\circ \pm\Delta\varphi$, then it would only be necessary to correct the small phase change ($\Delta\varphi$) between each cavity. The overall phase change of $+12^\circ$ can simply be corrected by shifting the initial phase of the bunch by 12° (at the buncher). It should be pointed out that since the individual cavities are operating slightly off resonance, a small increase in rf drive power is required.

The problem of cavity flattening (tuning the cells to achieve the design field distributions) may be divided into two parts: the initial tuning to achieve a flat cavity and the long-time stability. The initial tuning will be done with perturbation tuners. The number and distribution of these tuners along a cavity will depend on the relative sensitivity of a cavity to tuning errors (the sensitivity being inversely proportional to the bandwidth) and the mechanical construction tolerances. The long-time stability will be maintained with water cooling. If a temperature gradient exists along a cavity, it will result in a distributed detuning effect that will cause the field distribution to depart from the design value. Variations of the average temperature will cause the overall resonant frequency of the cavity to change, resulting in relative differences of phase between adjacent cavities. Consideration has been given to the feasibility of controlling the average temperature of the cavity cooling water in a feedback system to adjust the resonant frequency of the individual cavities.

To date, most of the structures that have been considered are symmetrically and periodically loaded, such as the iris-loaded waveguide, cloverleaf, and slotted-iris waveguide, operating in the π-mode standing wave configuration. The advantage of these structures is their high shunt impedance, and the disadvantage is the sensitivity to tuning and beam loading effects. It has been shown, both with model measurement[12] and equivalent circuit calculations,[13,14] that as the frequency spacing between the π mode and adjacent modes increases, a structure becomes less sensitive to tuning and beam loading effects. In a symmetrical structure a greater mode spacing is obtained by making the bandwidth as large as possible.

The bandwidth of a structure is increased by increasing the coupling between the unit cells. For the conventional iris-loaded guide with electric coupling, this can be accomplished by increasing the diameter of the iris hole which results in a rapid decrease in shunt impedance. Tight coupling between the cells can also be achieved by magnetic coupling. In this case, some type of opening in the iris between cells is provided near the outer wall of the guide, and the central bore hole is made only as large as is required to transmit the beam. Furthermore, the bore hole may be fitted with small drift tubes to improve the transit time factor.

Cloverleaf and slotted-iris structures operating in the π-mode standing wave configuration at a frequency of 800 Mc/sec are shown in Fig. 7. The former is a forward wave and the latter a backward wave structure, which does not greatly affect the operation of either structure. It should be pointed out that the discussion which follows applies to both of these structures as well as most loaded waveguides. The π mode is located at the edge of the first TM passband, where the group velocity is zero ($v_g = d\omega/d\beta = 0$). An increase in the frequency separation between the π mode and the adjacent mode is accomplished by increasing the bandwidth, defined as:

$$BW = \left| \frac{\omega_\pi - \omega_o}{(\omega_\pi + \omega_o)/2} \right| .$$

Here ω_π = angular frequency at the π mode and ω_o = angular frequency at the zero mode. For individual cavity lengths of approximately 3 meters, a bandwidth of 8% to 10% is sufficient to reduce cavity flattening and beam loading effects to a reasonable level.

For a symmetrical structure, the $\pi/2$ mode has the greatest mode spacing. The reduction of shunt impedance for the structures of Fig. 7, in this mode, is primarily due to the poor transit time factor.

Figure 8 shows a multiply periodic structure, which at the $\pi/2$ mode has an improved transit time factor as compared to the $\pi/2$ mode of a symmetrical structure. The results of some measurements made on a multiply periodic structure are reported elsewhere.[15] This structure has greater mode spacing and approximately the same shunt impedance as a π-mode cloverleaf or slotted-iris structure.

For the magnetically coupled structures mentioned here, bandwidth of 8 to 10% can be achieved and typical values of the shunt impedance at 800 Mc/sec range from 15 to 30 MΩ/m for β between 0.5 and 0.8. The shunt impedance is affected by the choice of bandwidth and bore hole size. The bandwidth is determined from considerations of minimum cavity length and allowable detuning effects (tuning tolerances and beam loading). Cavity lengths of 3 to 4 meters are presently considered reasonable (although longer cavities are desirable for engineering reasons), with peak power requirements of about 1 MW and electric fields between 1 and 2 MV/m. The electrically coupled, iris-loaded waveguide has about the same value of shunt impedance but the typical bandwidth is only about 1%.

Two new structures that are presently under investigation are shown in Fig. 9. These structures are being considered for operation in the $\pi/2$ mode with the electric field configuration as shown. It should be noted for the structure in Fig. 9, all the cells (including the end cells) are of equal length. The H-type irises shown have, in addition to the center bore hole, four magnetic field coupling holes. The R-type irises consist of a bar and ring (also being considered is a ring with single stem). The field configuration shown in Fig 9 is for a $\pi/2$ mode, and a typical dispersion curve for these structures is shown in Fig. 10. Due to the end termination, the π mode cannot exist in these structures. This work is still very preliminary, but shows some promise. Results of some measurements will be reported elsewhere.[15]

It can be concluded from these studies that practical waveguide structures can be built for the acceleration of protons to high energies. However, while the structures discussed here fill reasonably well the requirements outlined above, they are certainly not as good as one might wish. It is highly desirable that the search continue for other structures which will better meet the requirements at a lower cost.

References

1. Detailed discussions of these principles may be found in the following references: M.S. Livingston and J.P. Blewett, "Particle Accelerators", McGraw-Hill, New York, 1962, and L. Smith, "Encyclopedia of Physics", Vol. 44, p. 341, Springer-Verlag, Berlin, 1959.

2. N.C. Christofilos, "Method of Computation of Drift Tube Shapes", p. 176, Proc. CERN Symposium, 1956. A similar method has been developed at Harwell.

3. R.L. Gluckstern, "Shaped Drift Tubes and Irises for Linear Accelerators", p. 129, Proc. Intern. Conf. on High Energy Accelerators, 1961.

4. D.E. Young, "Design of Drift Tube Linacs", Minutes of Conf. on Linear Accelerators for High Energies at the Brookhaven National Laboratory, August 1962, BNL 6511, p. 76; T.W. Edwards, "Proton Linear Accelerator Cavity Calculations", MURA Internal Report No. 622, June 1961.

5. M. Chodorow et al., Rev. Sci. Instr. 26, 134 (1955).

6. J.C. Slater, Rev. Mod. Phys. 20, 473 (1948).

7. W. Walkinshaw, Proc. Phys. Soc. (London) 61, 246 (1948)

8. A. Carne, "Crossbar and Cloverleaf Structures at Rutherford", Minutes of Conf. on Proton Linear Accelerators at Yale University, October 1963, p. 104.

9. G.W. Wheeler and T.W. Ludlam, "The Effects of Phase and Field Errors on the Longitudinal Motion in a Long Proton Linac", ibid., p. 29.

10. S. Giordano, "Model Measurements and Correction of Beam Loading Effects in Proton Linacs", Minutes 1964 MURA Conf. on Proton Linear Accelerators, p. 252.

11. T. Nishikawa, "An Approach to the Study of Beam Loading for the Linear Accelerator", ibid., p. 214.

12. S. Giordano, "Measurements on Iris-Loaded Waveguides", Minutes of Conf. on Proton Linear Accelerators at Yale University, October 1963, p. 153.

13. R.L. Gluckstern, "Effect of Errors in Repetitive Structures", ibid., p. 190

14. D.E. Nagle and E.A. Knapp, "Steady State Behavior of a Ring or a Chain of Coupled Circuits", ibid., p. 171.

15. S. Giordano, "Some New Radio-Frequency Accelerating Structures", Proc. Particle

Accelerator Conference, March 10-12, 1965, IEEE Transactions on Nuclear Science, to be published.

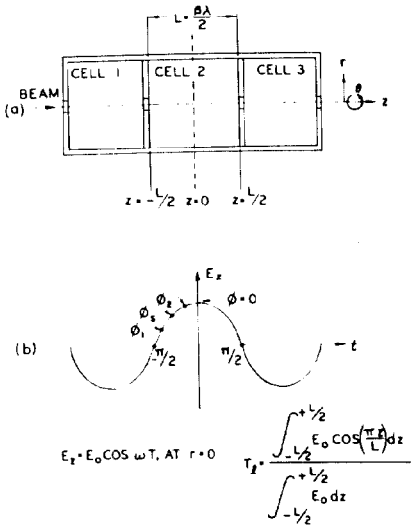

Fig. 1. A simple resonant cavity accelerator.

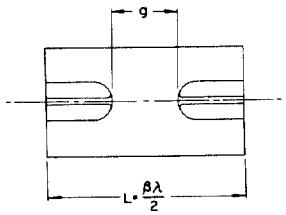

Fig. 2. A drift tube loaded resonant cell.

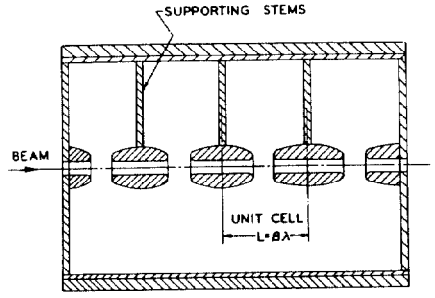

Fig. 3. A drift tube loaded cavity.

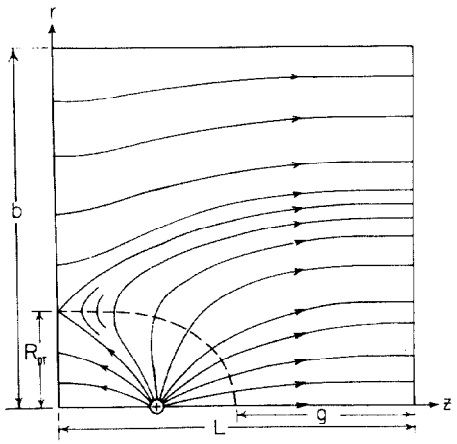

Fig. 4. Half cell for dipole driven cavity.

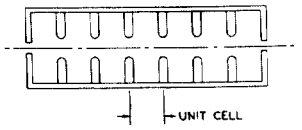

Fig. 5. The conventional iris-loaded waveguide.

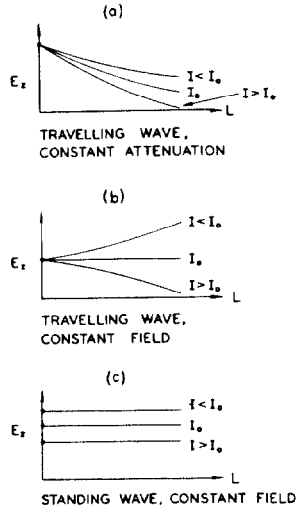

Fig. 6. Effects of beam current on field distribution.

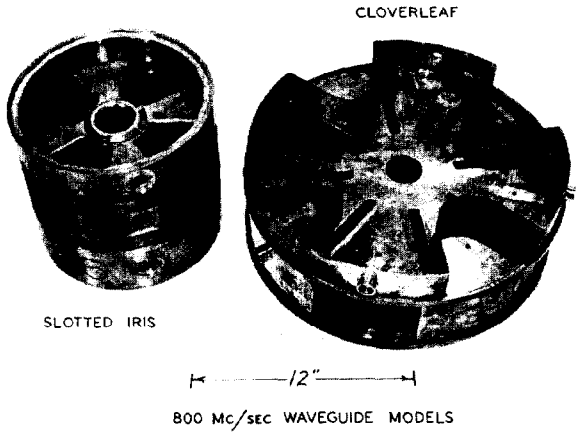

Fig. 7. 800 Mc/sec waveguide models.

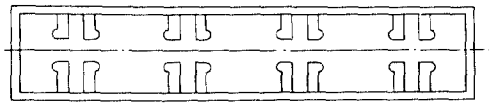

Fig. 8. A doubly periodic waveguide for $\pi/2$ mode operation.

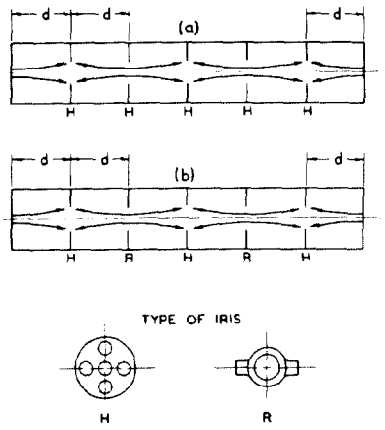

Fig. 9. Two new $\pi/2$ mode structures.

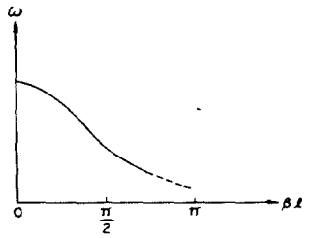

Fig. 10. Typical dispersion curve for structures of Fig. 9.