

HIGH CURRENT INDUCTION LINACS

William Barletta, Andris Faltens, Enrique Henestroza, and Edward Lee
Accelerator and Fusion Research and Engineering Divisions
Lawrence Berkeley Laboratory
Berkeley, CA 94720

Presented at the
Accelerator Driven Transmutation Technology Conference
Las Vegas, NV
July 25-29, 1994

* This work supported by the Office of Energy Research, Office of of Fusion Energy, U.S.
Department of Energy under Contract No. DE-AC03-76SF00098

MASTER
DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

ep

High Current Induction Linacs

William Barletta, Andris Faltens, Enrique Henestroza, and Edward Lee
*Accelerator and Fusion Research and Engineering Divisions
Lawrence Berkeley Laboratory
Berkeley, CA 94720*

OVERVIEW

Induction linacs are among the most powerful accelerators in existence. They have accelerated electron bunches of several kiloamperes, and are being investigated as drivers for heavy ion driven inertial confinement fusion (HIF), which requires peak beam currents of kiloamperes and average beam powers of some tens of megawatts. The requirement for waste transmutation with an 800 MeV proton or deuteron beam with an average current of 50 mA and an average power of 40 MW lies midway between the electron machines and the heavy ion machines in overall difficulty. Much of the technology and understanding of beam physics carries over from the previous machines to the new requirements. The induction linac allows use of a very large beam aperture, which may turn out to be crucial to reducing beam loss and machine activation from the beam halo. The major issues addressed here are transport of high intensity beams, availability of sources, efficiency of acceleration, and the state of the needed technology for the waste treatment application.

Because of the transformer-like action of an induction core and the accompanying magnetizing current, induction linacs make the most economic sense and have the highest efficiencies with large beam currents. Based on present understanding of beam transport limits, induction core magnetizing current requirements, and pulse modulators, the efficiencies could be very high. The study of beam transport at high intensities has been the major activity of the HIF community. Beam transport and sources are limiting at low energies but are not significant constraints at the higher energies. As will be shown, the proton beams will be space-charge-dominated, for which the emittance has only a minor effect on the overall beam diameter but does determine the density falloff at the beam edge.

BEAM TRANSPORT

The equation of motion for the beam envelope in a focusing channel is

$$\frac{d^2a}{dz^2} = K(z)a + \frac{\epsilon^2}{a^3} + \frac{Q}{a} \quad (1)$$

where a is the beam radius, Ka is the focusing force as a function of the longitudinal coordinate, z , ϵ is the beam emittance, and Q is the perveance, a measure of space charge. By comparing the space charge term with the emittance term, we can solve for the radius at which the two types of forces would be equal. Based on beams for heating of magnetic

fusion energy plasmas, we assume a 100 ampere proton source of 25 cm radius with a temperature of 10 eV, for which, at 1 MeV,

$$a_e = \sqrt{\frac{\epsilon^2}{Q}} \cong 4 \times 10^{-3} \text{ (m)} \quad (2)$$

Beams which are transported with radius small compared to a_e are emittance dominated, and beams which are large compared to a_e are space charge dominated. Beams considered here will be transported at large radius with modest fields, and thus are space charge dominated. There is little incentive to make the aperture smaller by using higher fields, and considerable difficulty in providing strong-enough fields which would make the beams emittance dominated. The solutions to the transport equation for our purposes are

$$I_s = 4 \times 10^5 \left(\frac{q}{A}\right) \beta \gamma (Ba)^2 \quad (3)$$

$$I_q = 8 \times 10^5 (\beta \gamma)^2 (\eta Ba) \quad (4)$$

The subscript s refers to solenoidal focusing, and the subscript q refers to quadrupole focusing with an effective field occupancy η . B is the magnetic field, and β , and γ are the relativistic factors. These results are shown graphically in Figure 1.

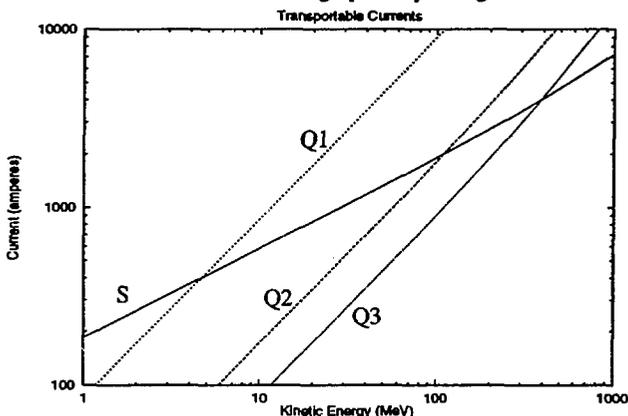


Fig. 1. Transportable proton currents for various solenoidal and quadrupole channels. For solenoids, $B=1T$, and for quadrupoles the field at the edge of the beam is 1T. 'S' denotes a solenoid channel with $a=10\text{cm}$; 'Q1' a quadrupole channel with $a=10\text{cm}$, $\eta=0.5$; Q2 a quadrupole channel with $a=10\text{cm}$ and $\eta=0.1$; and Q3 a quadrupole channel with $a=5\text{cm}$ and $\eta=0.1$.

Transverse focusing gives an upper bound to the beam current. There is a technological limit on attainable magnetic fields of about 10 Tesla and on practical occupancy of about 60% in a FODO lattice. The cost optimum lattices are sparse, with an occupancy close to 10% and fields close to 1 Tesla. Beyond an energy of 100 MeV, focusing ceases to be a major constraint. Longitudinal focusing is mainly a balancing of the bunch field against some externally applied field ramps. It is uneconomical to devote most of the "effort" to just keeping the bunch together, so we choose a bunch shape which requires the focusing field to be small compared to the acceleration field. Longitudinally, the bunch can be thought of as undergoing an oscillation at the reduced plasma frequency where the focusing fields reverse the motion incurred by the beam space charge. Using the well-known approximation that a parabolic bunch maintains its parabolic shape from 1-D space charge expansion and that the fields are linear with z , the longitudinally defocusing space charge fields would be counteracted with occasional ramped focusing fields placed along the accelerator. For a long bunch with rounded ends, a pair of "ears" at the bunch ends, one positive and one negative, is sufficient.

For a space charge dominated beam, transported in a solenoid or quadrupole focused channel, the beam density is nearly uniform within the interior. The density at the edge decays abruptly as $\exp(-z/\lambda_D)$ with the characteristic decay distance,

$$\lambda_D = \frac{v_{th}}{\omega_p} \quad (5)$$

The beam size is 2.6 cm in radius with a decay distance, λ_D , of 2.6 mm for a 1 kA deuteron beam focused with 1T fields at an energy of 400 MeV, which is a midway point in the energy of the accelerator. To this we add an amount corresponding to the beam misalignment oscillation resulting from 500 quadrupoles each with an rms error of 0.1 mm in either plane, $\Delta x = \Delta y = 9$ mm. Another small contribution to the required aperture is the initial mismatch of the beam. As for the other errors originating at low energy, its effect decreases with distance along the machine. There is a small amount of scattering from the background gas which we neglect here. Taking all of these effects together, we expect a beam halo of about one cm. The standard prescription used for HIF for intense beams is that the clearance should be $1.25a + 1$ cm, which corresponds to a clearance of about 2 cm here, which is just adequate. There is an image effect which leads to emittance growth of the beam for closer approaches.

Beam Instabilities

The possibility of transverse and longitudinal instabilities must be considered for the transport and acceleration of any intense beam. Both the transverse and longitudinal instabilities are stabilized by a momentum spread in the beam and driven by resistance in the beam coupling impedances. On both accounts, the waste treatment accelerator is in a favorable situation. Because there is no stringent focusing requirement at the end of the accelerator, there is no significant constraint on the allowable momentum spread. Similarly, there is no need to bend the beams within the accelerator or afterwards. Starting with a monoenergetic beam, which is unstable, it is probably acceptable to let any instability grow and create its own stabilizing momentum

spread. The worst case scenario is stabilized with less than 10% momentum spread, which is transportable by the focusing channel.

The growth rates depend on the beam space charge intensity and the resistive part of the pulser or generator impedance. For the standard PFN and closing switch only type of pulser, the generator impedance is just the characteristic impedance of the PFN, which is roughly the voltage gain per meter divided by the beam current for a matched system. For the capacitor plus controllable switch circuit, the coupling impedance is very low. A pulser with a truly regulated output voltage, of course, would have no resistive impedance.

MACHINE DESIGN

The induction linac accelerates a single bunch of ions, whether deuterons or protons, which contains a fixed total amount of charge. The free variable is the repetition rate. To attain the desired average current of 50 mA, the rep rate could be 50 pps with a bunch charge of 1 mC, or 500 pps with 100 μ C. Given the total charge within the bunch, we may also choose how it is to be accelerated, that is, at some point along the accelerator, we may choose the pulse duration, acceleration rate, and beam radius. Starting with a 100 A proton source, and using simple acceleration scenarios, the 50 pps machine would be 1664 m long and use a maximum induction core buildup of 40 cm, the 500 pps machine would be half as long and require a maximum core buildup of 13 cm, thus there is a major incentive to use the highest rep rate technologically possible. The minimum core inner radius would be about 12 cm at the low energy end, decreasing to about 7 cm at the high energy end. The maximum diameter of these machines would be 1m and 0.5 m respectively. The number of variables is large enough that for the HIF accelerators this selection process has been computerized in the two design programs LIACEP and HILDA. For each set of design choices, these programs calculate the incremental cost of acceleration, and choose the lowest cost design. The entire machine is the sum of these local design optima. In practice, we expect a slow and smooth variation in the character of the machine along its length. The principal beam dynamics activity is to ensure that the bunch is kept together with longitudinally focusing "ears" which overcome the repulsive effects of the space charge, and that a small amount of tilt is provided on the accelerating waveforms if it is desired to lengthen or shorten the bunch length downstream. The major increase in bunch current arises from the increase in velocity, so the bunch length control is a minor adjustment, provided that the bunch is short compared to the machine length.

The results of such a computerized design are shown in Figures 3-7, which display some of the trends along the accelerator. As it is presently set up, this program considers only high field superconducting magnets, which are the correct choice for heavy ions for energy production, and pulsed quadrupoles for machines designed for essentially single pulse operation. As seen in Fig. 1, for the much lighter protons conventional magnets operating at the 1 T level are adequate, and lead to a smaller overall package, but we do not yet have an algorithm for such magnets in the code. Nevertheless, we show the general trends of decreasing bore size and increasing half-periods which result in reduced cost for the available magnet types in the programs. Because of the premium attached to reducing beam loss, the beam aperture was increased to 20 cm for these

cases, that is a clearance of about 10 cm instead of the minimum 2cm, which shows the realizability of the large clearance options.

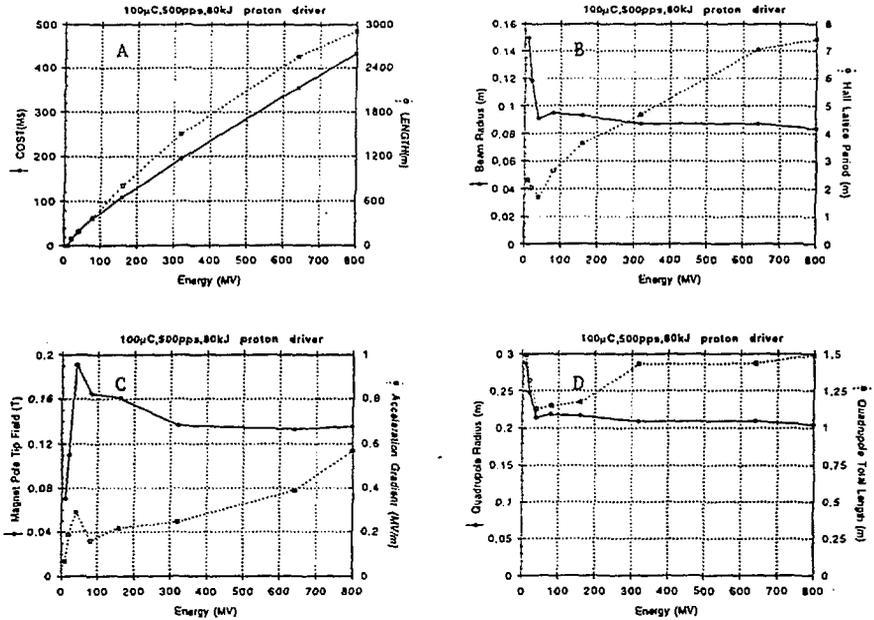


Fig. 2. Example of minimum cost machine designed with an imposed 10 cm clearance between the beam and the vacuum chamber. Panel A shows the cost and length as almost linear functions of energy; panel B shows the decrease of beam radius and increase of half-period with energy; panel C shows the quadrupole field and acceleration rates; and panel D shows the quadrupole radius and length.

The use of magnetic quadrupoles at low energies is inconvenient because of the required short half-periods. A continuous solenoid (or a solenoid with short breaks in it which are required for the acceleration gaps) at low energies can transport a current which is higher than can be transported in a similar strength quadrupole, therefore the first 10 MeV of the machine uses solenoidal transport of a beam of 10cm radius. After that, the beam radius can be gradually decreased. Figure 1 shows the beam transport limits for protons as a function of the particle energy for the continuous solenoid and the alternating gradient quadrupole systems. The corresponding currents for deuteron acceleration are roughly 0.35 times as great in the solenoidal section and 0.5 times as great in the quadrupole transport section. From these curves, it is evident that beam transport ceases to be a major concern past the 100 MeV point. The low-energy limit is not a major hindrance because it is desirable to maintain a low acceleration rate at the beginning of the machine. With a low acceleration rate, relatively less space is required for the acceleration gaps, and more can be devoted to the focusing elements. The low initial acceleration rate is desired for beam transport reasons. It is possible to accelerate at such a high rate that the energy

of the particles passing one location varies by a factor of several, that is, from the front of the bunch to the rear, with the result that a system matched for one energy and current is not well suited for the other. To keep the bunch within reasonable transport limits, we use the criterion that at any one location the energy at the rear of the bunch should be no higher than twice the energy which the front of the bunch had when it passed that location. This corresponds to a doubling of the energy in a distance equal to the bunch length, up to the point where high field or economics limits make this prescription undesirable. Past this point, the acceleration proceeds at a constant rate of 1 MeV/m, which is believed to be the technological limit.

Besides the usual beam transport considerations given above, there are at least two additional ones: whether the energy in the beam pulse may be so high that the machine is damaged by the beam in case of a fault; and whether the beam loss along the machine is sufficiently serious to lead to activation of the accelerator and a subsequent imposed lowering of the operating current. Our work on the HIF beams has shown that an intense beam must be kept within approximately 80% of the aperture, otherwise the nonlinearities associated with the beam image fields lead to emittance growth and beam loss. In addition, we have found that an off center beam is much more susceptible to emittance growth than a centered beam. For these reasons, for the waste disposal application, we have examined cases where the clearance is greatly increased. The induction linac differs from the rf linacs in that there is no significant penalty in the beam dynamics if the bore diameter is made very large. The induction linac could be thought of as a very-low-frequency rf system which is very tightly coupled to the cavity. The electric field distribution across the accelerating gap during the flat portion of the pulse is exactly the same as for a dc gap, and the energy gained by a particle crossing it is just the applied voltage, independent of the transverse dimension of the beam pipe. There actually is a vestige of a low frequency resonance in the induction modules, at a frequency around 1 MHz, but this is greatly swamped by the drive pulse circuitry. That is, the impedance of the induction module is much higher than the impedance of the generator or pulser that is connected across the module, so that the beam interacts with the pulser impedance, and we can neglect the resonant contribution.

Pulsers

The pulser is perhaps the most important component of an induction linac; besides its interaction with the beam it also largely determines the overall machine design. There are two major types of pulse generators in use: the line-type and the controllable switch type. In the line-type modulator a lumped or distributed element pulse forming line is discharged into the load, with the discharge initiated by a closing switch. The usual switch is a thyratron gas tube. The thyratron utilizes a low pressure plasma discharge and remains conducting until the plasma has recombined, which may require some tens of microseconds. The controllable switch type of pulser has the additional capability to be turned off, and in some cases to also regulate the output voltage. The usual controllable switches are hard tubes or transistor arrays. Because of the unfavorable economics or the technical limitations of some of these primary switches, it is also possible to use a magnetic pulse compression line (the Melville line) to increase the output power considerably above that available from the primary switch. Even in the magnetic compression line a single stage is usually limited in its performance by available magnetic materials and insulation, and it has been found preferable to use several stages, each of which handles the entire pulse energy in turn. This adds to the overall cost of the modulator and reduces its effi-

ciency. The ideal switch that can simultaneously handle the pulse energy, regulate the output voltage, present a low generator output impedance, and have a low cost does not exist now. A device which approximates some of these desirable characteristics and which has been declining in cost is the field effect transistor (FET) array. Because of the eventual promise of FETs, we have been making measurements on an FET-driven induction core. Similar and related work is ongoing at LLNL and at FM Technologies.

For the LBL HIF application, the most probable first use of FETs will be in providing small corrections, of the order of 1%, to the main acceleration voltages. The combination of a very high power but economical pulser with a controllable trimmer results in an overall package which has the best features of each constituent. Related developments at LLNL and FM Technologies center on pulsers that are capable at pulsing at several hundred kHz for use in induction accelerator recirculators. The FET pulser at LBL was designed and constructed by American Controls Engineering in San Diego, Ca. The pulser's nominal specifications are an operating voltage of 12 kV and a peak current of 1.5kA. The measured 10-90% risetime and falltime are 150ns and 100ns respectively. The pulser was initially intended to demonstrate burst mode operation at the 100 kHz rate with 1 μ s pulses. It is also well suited to demonstrate the possible performance of induction linacs for other applications, such as the present one. In some respects, the output pulse is similar to that which can be obtained from a thyatron pulser, but using different means.

The circuit used to obtain the pulses shown in Fig.3 is very simple: a large capacitor and the FET array. The capacitor is large enough so that its voltage appears almost constant during the pulse; actually, there is a 1% droop during the pulse due to the charge removed from it. The charge to the capacitor is restored during the interval between pulses, at high efficiency, from a dc power supply which is isolated from the capacitor with a small resistor. The induction cores used in these tests were of several types, principally 50%Ni-50%Fe of 25 μ m thickness and Metglas type 2506 SC of 30 μ m thickness. These cores behave similarly and have essentially square loop character at low frequencies, which allows them to retain a remanent magnetization close to the saturation magnetization value. A dc bias current of 30 amperes was used to reset the cores to the negative remanent value. At constant current near the minimum required to saturate the outermost laminations of the core, the reset would occur in about 50 μ s after the pulse, and of course faster with a higher current. A large choke of 2.5 mHy inductance was used in series with the dc current supply to isolate it from the high voltage pulses applied to the core

During the pulse, the dc level in the choke current changes abruptly by a few amperes. In addition, there is excitation of higher order mode oscillations in the choke. These processes contribute a very small amount to the losses of the system and to the top of the voltage waveform. The spurious oscillations die out in tens of μ s, well before the next pulse. A resistive load was placed across the core to simulate a beam load. The load current could be increased above that shown with a small decrease in the pulse flatness. As shown, the "beam" current corresponds to roughly the middle of a long pulse /low rep rate machine. The nominal efficiency for this point is 50%, with points upstream being slightly lower in efficiency and points downstream being higher in efficiency. Efficiency is the power going into the "beam" divided by the dc power supply input; during the pulse the "beam" current is about twice the core current. The core magnetizing current shown during the pulse is not the minimum possible. Core excitation current can be related to various processes involving domain nucleation and domain wall movements. The

macroscopic result of these is that the drive current can be decreased by using a larger core and reducing the flux swing within it. The drive current is also decreased by avoiding hard saturation in either direction. As shown, about half of the possible flux swing is being utilized. In the design of an induction linac for efficiency instead of for lowest capital cost, it is also advantageous to obtain the required core cross sectional area by using as much axial distance and as little radial distance as possible, to arrive at a smaller core volume. Two good approximations are that the core losses are proportional to the core volume, and that the average drive current is proportional to the average core radius, for a fixed voltage and pulse duration. Improved efficiency comes at the increased cost of the beam transport for the longer accelerator, and the increased cost of the core material for the underutilized flux swing.

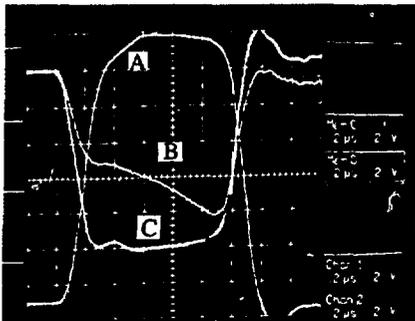


Fig. 3. Waveforms of the FET pulser driving an induction core at 72 pulses per second. A is the load (simulated beam) current (100A/division), B is the core current (80A/division), and C is the output voltage (2kV/division). Horizontal sweep is 200ns/division.

This pulser was operated at 72 pps intermittently for 25 million pulses during exploratory testing and at 100 pps around the clock for a total of 200 million pulses for a life and reliability test of the pulser and core assembly. It was also operated at 200 pps for brief periods. The core was convection cooled with ambient air in these tests. Higher repetition rates require the core package to be re-engineered for better cooling, and such an effort is presently underway.

Sources

High current long pulse sources have already been developed for heating of magnetic fusion energy plasmas. The available species include protons, deuterons, and H^- . The particular sources which are well matched to the front end of an induction linac produced 100 amps of protons at 120 kV and an equivalent deuteron current of 70 amps. The sources were developed in the 1970's and were used reliably for several years in TFTR and Doublet IIIID. These sources used a large plasma chamber and an extraction surface made of sturdy, water cooled grid bars with typical dimension of 0.5 cm diameter with a separation of about 1 cm. The output beams went through an accel-decel region composed of similar grids. The grids were aligned to decrease beam loss. A typical source had dimensions of 12 cmx48 cm. For the induction linac using such a gridded source the geometry would need to be reconfigured toward a more circular

shape, and the emerging ribbon beams would be allowed to merge into one big beam of circular cross section, and accelerated to higher voltages. The emittance increase due to this process is reasonably well understood, and continues to be addressed by the HIF program. The major developmental task for the future is to work on a system which would produce faster risetimes on the injector pulse, either by longitudinal switching of the beam or by transverse deflections.

A considerable amount of hardware remains from the Magnetic Fusion neutral beam program from the previous two decades which could be modified for short pulse operation. It would be relatively inexpensive to look at the high current-short pulse operation while the knowledgeable personnel and hardware are still available.

Needed Developments

In addition to the source development described above, work is needed on low cost induction cores and better switches. The metallic glass induction core material cost has declined dramatically in the past decade, to about \$3/kg. The cost of annealing and insulating the material has not decreased as much, and some development is needed to arrive at a satisfactory solution. In the much more mature silicon steel and nickel iron induction core manufacturing processes the insulation is adequate and takes up negligible space. These are baked-on inorganic coatings which survive the annealing temperatures. For the newer metallic glass induction cores, several types of insulation have been tried, some of which survive the annealing temperatures but have relatively low voltage holding, and others which have very high voltage holding but do not tolerate the annealing temperature. The metallic glasses themselves can not tolerate the temperatures used on the previous silicon steel insulation processes. Several techniques have been suggested but there has not been sufficient investment in the field to follow through.

SUMMARY

The induction linac is used in a few applications which require very high peak currents, such as flash X-ray radiography and heavy-ion fusion. It has been considered for food irradiation, collective acceleration, free-electron laser drivers, and similar applications. The peak currents which have been accelerated have been in the vicinity of 10 kA. The particle used in these applications usually has been the electron. As far as the acceleration process is concerned, it doesn't matter what the particle is. The beam simply is a current source to the acceleration modules. The energy gained by the beam is the result of the drive voltage applied to the module minus the beam generated voltage. Electrically, the induction accelerator functions as a transformer which applies power to a particle beam without the intervening step of going through a radio frequency system. The induction module can be thought of as the equivalent of an rf cavity, but one which is so heavily coupled to the drive circuitry that the cavity is non-resonant. The drive circuitry is some kind of modulator, such as a pulse-forming line or a controllable switch with a capacitor bank. The difference between the two types of accelerators is partly in the power conversion train: the induction linac does not require the step between a modulator and the acceleration cavity, that is, the rf generation; on the other hand, the drive power required by the induction module to establish the acceleration fields is considerable, so the efficiency is high only for high beam currents. While clearly suited for the high peak power applications for which induction

linacs were originally built, they may also be well suited for some high average power applications. This is the subject of this present paper. The high average power is basically a quest for higher average current, since generally the output particle kinetic energy is specified. In some applications it is the upper energy which is limited, and even in these cases it may be desirable to use the highest energies because it is usually easier to increase energy by further acceleration rather than by increased current. The two main ways by which the average current may be increased are by increasing the pulse duration and by increasing the repetition rate. There is a maximum pulse duration which is related to subsequent downstream bunch dynamics, and a maximum rep rate which is related to present day technology. Another, less frequently used technique for increasing the average current is the use of multiple beams: this technique works best where the costs of the induction core and pulsers are dominant and the required core is very large. In this case, the central bore may be enlarged to carry several parallel beams with only a relatively small fractional increase in the core outer diameter and the attendant costs. In some instances, as the core inner diameter is increased to accommodate the multiple beams, the core outer diameter shrinks, because the increased current carrying capacity of the parallel transport channels allows the use of a much shorter pulse duration, which in turn requires the smaller core cross-sectional area. This last option has not been examined for the present application.

The development of a high current proton or deuteron induction linac, after the initial parameter selection for the overall configuration, becomes the development of a few representative modules for several locations along the machine. This is essentially a low cost process which answers most of the technical questions.

ACKNOWLEDGEMENTS

The authors are indebted for the indispensable help provided by Peter Seidl in the preparation of this manuscript and for valuable discussions.